

Kai-Min Chung
Yu Sasaki (Eds.)

LNCS 15492

Advances in Cryptology – ASIACRYPT 2024

30th International Conference on the Theory
and Application of Cryptology and Information Security
Kolkata, India, December 9–13, 2024
Proceedings, Part IX

9
Part IX



 Springer

Lecture Notes in Computer Science

15492

Founding Editors

Gerhard Goos
Juris Hartmanis

Editorial Board Members

Elisa Bertino, *Purdue University, West Lafayette, IN, USA*

Wen Gao, *Peking University, Beijing, China*

Bernhard Steffen , *TU Dortmund University, Dortmund, Germany*

Moti Yung , *Columbia University, New York, NY, USA*

The series Lecture Notes in Computer Science (LNCS), including its subseries Lecture Notes in Artificial Intelligence (LNAI) and Lecture Notes in Bioinformatics (LNBI), has established itself as a medium for the publication of new developments in computer science and information technology research, teaching, and education.


LNCS enjoys close cooperation with the computer science R & D community, the series counts many renowned academics among its volume editors and paper authors, and collaborates with prestigious societies. Its mission is to serve this international community by providing an invaluable service, mainly focused on the publication of conference and workshop proceedings and postproceedings. LNCS commenced publication in 1973.


Kai-Min Chung · Yu Sasaki
Editors

Advances in Cryptology – ASIACRYPT 2024

30th International Conference on the Theory
and Application of Cryptology and Information Security
Kolkata, India, December 9–13, 2024
Proceedings, Part IX

Editors

Kai-Min Chung 
Academia Sinica
Taipei, Taiwan

Yu Sasaki 
NTT Social Informatics Laboratories
Tokyo, Japan

ISSN 0302-9743

ISSN 1611-3349 (electronic)

Lecture Notes in Computer Science

ISBN 978-981-96-0946-8

ISBN 978-981-96-0947-5 (eBook)

<https://doi.org/10.1007/978-981-96-0947-5>

© International Association for Cryptologic Research 2025

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Singapore Pte Ltd.
The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

If disposing of this product, please recycle the paper.

Preface

The 30th Annual International Conference on the Theory and Application of Cryptology and Information Security (Asiacrypt 2024) was held in Kolkata, India, on December 9–13, 2024. The conference covered all technical aspects of cryptology and was sponsored by the International Association for Cryptologic Research (IACR).

We received a record 433 paper submissions for Asiacrypt from around the world. The Program Committee (PC) selected 127 papers for publication in the proceedings of the conference. As in the previous year, the Asiacrypt 2024 program had three tracks.

The two program chairs are greatly indebted to the six area chairs for their great contributions throughout the paper selection process. The area chairs were Siyao Guo for Information-Theoretic and Complexity-Theoretic Cryptography, Bo-Yin Yang for Efficient and Secure Implementations, Goichiro Hanaoka for Public-Key Cryptography Algorithms and Protocols, Arpita Patra for Multi-Party Computation and Zero-Knowledge, Prabhanjan Ananth for Public-Key Primitives with Advanced Functionalities, and Tetsu Iwata for Symmetric-Key Cryptography. The area chairs helped suggest candidates to form a strong program committee, foster and moderate discussions together with the PC members assigned as paper discussion leads to form consensus, suggest decisions on submissions in their areas, and nominate outstanding PC members. We are sincerely grateful for the invaluable contributions of the area chairs.

To review and evaluate the submissions, while keeping the load per PC member manageable, we selected the PC members consisting of 105 leading experts from all over the world, in all six topic areas of cryptology, and we also had approximately 468 external reviewers, whose input was critical to the selection of papers. The review process was conducted using double-blind peer review. The conference operated a two-round review system with a rebuttal phase. This year, we continued the interactive rebuttal from Asiacrypt 2023. After the reviews and first-round discussions, PC members and area chairs selected 264 submissions to proceed to the second round. The remaining 169 papers were rejected, including two desk-rejects. Then, the authors were invited to participate in a two-step interactive rebuttal phase, where the authors needed to submit a rebuttal in five days and then interact with the reviewers to address questions and concerns the following week. We believe the interactive form of the rebuttal encouraged discussions between the authors and the reviewers to clarify the concerns and contributions of the submissions and improved the review process. Then, after several weeks of second-round discussions, the committee selected the final 127 papers to appear in these proceedings. This year, we received seven resubmissions from the revise-and-resubmit experiment from Crypto 2024, of which five were accepted. The nine volumes of the conference proceedings contain the revised versions of the 127 papers that were selected. The final revised versions of papers were not reviewed again and the authors are responsible for their contents.

The PC nominated and voted for three papers to receive the Best Paper Awards. The Best Paper Awards went to Mariya Georgieva Belorgey, Sergiu Carпов, Nicolas Gama,

Sandra Guasch and Dimitar Jetchev for their paper “Revisiting Key Decomposition Techniques for FHE: Simpler, Faster and More Generic”, Xiaoyang Dong, Yingxin Li, Fukang Liu, Siwei Sun and Gaoli Wang for their paper “The First Practical Collision for 31-Step SHA-256”, and Valerio Cini and Hoeteck Wee for their paper “Unbounded ABE for Circuits from LWE, Revisited”. The authors of those three papers were invited to submit extended versions of their papers to the Journal of Cryptology.

The program of Asiacrypt 2024 also featured the 2024 IACR Distinguished Lecture, delivered by Paul Kocher, as well as an invited talk by Dakshita Khurana. Following Eurocrypt 2024, we selected seven PC members for the Distinguished PC Members Awards, nominated by the area chairs and program chairs. The Distinguished PC Members Awards went to Sherman S. M. Chow, Elizabeth Crites, Matthias J. Kannwischer, Mustafa Khairallah, Ruben Niederhagen, Maciej Obremski and Keita Xagawa.

Following Crypto 2024, Asiacrypt 2024 included an artifact evaluation process for the first time. Authors of accepted papers were invited to submit associated artifacts, such as software or datasets, for archiving alongside their papers; 14 artifacts were submitted. Rei Ueno was the Artifact Chair and led an artifact evaluation committee of 10 members listed below. In the interactive review process between authors and reviewers, the goal was not just to evaluate artifacts but also to improve them. Artifacts that passed successfully through the artifact review process were publicly archived by the IACR at <https://artifacts.iacr.org/>.

Numerous people contributed to the success of Asiacrypt 2024. We would like to thank all the authors, including those whose submissions were not accepted, for submitting their research results to the conference. We are very grateful to the area chairs, PC members, and external reviewers for contributing their knowledge and expertise, and for the tremendous amount of work that was done with reading papers and contributing to the discussions. We are greatly indebted to Bimal Kumar Roy, the General Chairs, for their efforts in organizing the event, to Kevin McCurley and Kay McKelly for their help with the website and review system, and to Jih-Wei Shih for the assistance with the use of the review system. We thank the Asiacrypt 2024 advisory committee members Bart Preneel, Huaxiong Wang, Bo-Yin Yang, Goichiro Hanaoka, Jian Guo, Ron Steinfeld, and Michel Abdalla for their valuable suggestions. We are also grateful for the helpful advice and organizational material provided to us by Crypto 2024 PC co-chairs Leonid Reyzin and Douglas Stebila, Eurocrypt 2024 PC co-chairs Marc Joye and Gregor Leander, and TCC 2023 chair Hoeteck Wee. We also thank the team at Springer for handling the publication of these conference proceedings.

December 2024

Kai-Min Chung
Yu Sasaki

Pedro Branco	Max Planck Institute for Security and Privacy, Germany
Gaëtan Cassiers	UCLouvain, Belgium
Céline Chevalier	CRED, Université Paris-Panthéon-Assas, and DIENS, France
Avik Chakraborti	Institute for Advancing Intelligence TCG CREST, India
Nishanth Chandran	Microsoft Research India, India
Jie Chen	East China Normal University, China
Yu Long Chen	KU Leuven and National Institute of Standards and Technology, Belgium
Mahdi Cheraghchi	University of Michigan, USA
Nai-Hui Chia	Rice University, USA
Wonseok Choi	Purdue University, USA
Tung Chou	Academia Sinica, Taiwan
Arka Rai Choudhuri	NTT Research, USA
Sherman S. M. Chow	Chinese University of Hong Kong, China
Chitchanok Chuengsatiansup	University of Melbourne, Australia
Michele Ciampi	University of Edinburgh, UK
Valerio Cini	NTT Research, USA
Elizabeth Crites	Web3 Foundation, Switzerland
Nico Döttling	CISPA Helmholtz Center, Germany
Avijit Dutta	Institute for Advancing Intelligence TCG CREST, India
Daniel Escudero	JP Morgan AlgoCRYPT CoE and JP Morgan AI Research, USA
Thomas Espitau	PQShield, France
Jun Furukawa	NEC Corporation, Japan
Rosario Gennaro	CUNY, USA
Junqing Gong	East China Normal University, China
Rishab Goyal	University of Wisconsin-Madison, USA
Julia Hesse	IBM Research Europe, Switzerland
Akinori Hosoyamada	NTT Social Informatics Laboratories, Japan
Michael Hutter	PQShield, Austria
Takanori Isobe	University of Hyogo, Japan
Joseph Jaeger	Georgia Institute of Technology, USA
Matthias J. Kannwischer	Chelpis Quantum Corp, Taiwan
Bhavana Kanukurthi	Indian Institute of Science, India
Shuichi Katsumata	PQShield and AIST, Japan
Jonathan Katz	Google and University of Maryland, USA
Mustafa Khairallah	Lund University, Sweden
Fuyuki Kitagawa	NTT Social Informatics Laboratories, Japan

Karen Klein	ETH Zurich, Switzerland
Mukul Kulkarni	Technology Innovation Institute, United Arab Emirates
Po-Chun Kuo	WisdomRoot Tech, Taiwan
Jooyoung Lee	KAIST, South Korea
Wei-Kai Lin	University of Virginia, USA
Feng-Hao Liu	Washington State University, USA
Jiahui Liu	Massachusetts Institute of Technology, USA
Qipeng Liu	UC San Diego, USA
Shengli Liu	Shanghai Jiao Tong University, China
Chen-Da Liu-Zhang	Lucerne University of Applied Sciences and Arts and Web3 Foundation, Switzerland
Yun Lu	University of Victoria, Canada
Ji Luo	University of Washington, USA
Silvia Mella	Radboud University, Netherlands
Peihan Miao	Brown University, USA
Daniele Micciancio	UCSD, USA
Yusuke Naito	Mitsubishi Electric Corporation, Japan
Khoa Nguyen	University of Wollongong, Australia
Ruben Niederhagen	Academia Sinica, Taiwan and University of Southern Denmark, Denmark
Maciej Obremski	National University of Singapore, Singapore
Miyako Ohkubo	NICT, Japan
Eran Omri	Ariel University, Israel
Jiaxin Pan	University of Kassel, Germany
Anat Paskin-Cherniavsky	Ariel University, Israel
Goutam Paul	Indian Statistical Institute, India
Chris Peikert	University of Michigan, USA
Christophe Petit	University of Birmingham and Université libre de Bruxelles, Belgium
Rachel Player	Royal Holloway University of London, UK
Thomas Prest	PQShield, France
Shahram Rasoolzadeh	Ruhr University Bochum, Germany
Alexander Russell	University of Connecticut, USA
Santanu Sarkar	IIT Madras, India
Sven Schäge	Eindhoven University of Technology, Netherlands
Gregor Seiler	IBM Research Europe, Switzerland
Sruthi Sekar	Indian Institute of Technology, India
Yaobin Shen	Xiamen University, China
Danping Shi	Institute of Information Engineering, Chinese Academy of Sciences, China
Yifan Song	Tsinghua University, China

Katerina Sotiraki	Yale University, USA
Akshayaram Srinivasan	University of Toronto, Canada
Marc Stöttinger	Hochschule RheinMain, Germany
Akira Takahashi	J.P. Morgan AI Research and AlgoCRYPT CoE, USA
Qiang Tang	University of Sydney, Australia
Aishwarya Thiruvengadam	IIT Madras, India
Emmanuel Thomé	Inria Nancy, France
Junichi Tomida	NTT Social Informatics Laboratories, Japan
Monika Trimoska	Eindhoven University of Technology, Netherlands
Huaxiong Wang	Nanyang Technological University, Singapore
Meiqin Wang	Shandong University, China
Qingju Wang	Telecom Paris, Institut Polytechnique de Paris, France
David Wu	UT Austin, USA
Keita Xagawa	Technology Innovation Institute, United Arab Emirates
Chaoping Xing	Shanghai Jiaotong University, China
Shiyuan Xu	University of Hong Kong, China
Anshu Yadav	IST, Austria
Shota Yamada	AIST, Japan
Yu Yu	Shanghai Jiao Tong University, China
Mark Zhandry	NTT Research, USA
Hong-Sheng Zhou	Virginia Commonwealth University, USA

Additional Reviewers

Hugo Aaronson	Jiawei Bao
Damiano Abram	Jyotirmoy Basak
Hamza Abusalah	Nirupam Basak
Abtin Afshar	Gabrielle Beck
Siddharth Agarwal	Hugo Beguinet
Navid Alapati	Amit Behera
Miguel Ambrona	Mihir Bellare
Parisa Amiri Eliasi	Tamar Ben David
Ravi Anand	Aner Moshe Ben Efraim
Saikrishna Badrinarayanan	Fabrice Benhamouda
Chen Bai	Tyler Besselman
David Balbás	Tim Beyne
Brieuc Balon	Rishabh Bhadauria
Gustavo Banegas	Divyanshu Bhardwaj
Laasya Bangalore	Shivam Bhasin

Amit Singh Bhati
Loïc Bidoux
Alexander Bienstock
Jan Bobolz
Alexandra Boldyreva
Maxime Bombar
Nicolas Bon
Carl Bootland
Jonathan Bootle
Giacomo Borin
Cecilia Boschini
Jean-Philippe Bossuat
Mariana Botelho da Gama
Christina Boura
Pierre Briaud
Jeffrey Burdges
Fabio Campos
Yibo Cao
Pedro Capitão
Ignacio Cascudo
David Cash
Wouter Castryck
Anirban Chakrabartha
Debasmita Chakraborty
Suvradip Chakraborty
Kanad Chakravarti
Ayantika Chatterjee
Rohit Chatterjee
Jorge Chavez-Saab
Binyi Chen
Bohang Chen
Long Chen
Mingjie Chen
Shiyao Chen
Xue Chen
Yu-Chi Chen
Chen-Mou Cheng
Jiaqi Cheng
Ashish Choudhury
Miranda Christ
Qiaohan Chu
Eldon Chung
Hao Chung
Léo Colisson
Daniel Collins
Jolijn Cottaar
Murilo Coutinho
Eric Crockett
Bibhas Chandra Das
Nayana Das
Pratish Datta
Alex Davidson
Hannah Davis
Leo de Castro
Luca De Feo
Thomas Decru
Giovanni Deligios
Ning Ding
Fangqi Dong
Minxin Du
Qiuyan Du
Jesko Dujmovic
Moumita Dutta
Pranjal Dutta
Duyen
Marius Eggert
Solane El Hirsch
Andre Esser
Hülya Evkan
Sebastian Faller
Yanchen Fan
Niklas Fassbender
Hanwen Feng
Xiutao Feng
Dario Fiore
Scott Fluhrer
Danilo Francati
Shiuan Fu
Georg Fuchsbauer
Shang Gao
Rachit Garg
Gayathri Garimella
Pierrick Gaudry
François Gérard
Paul Gerhart
Riddhi Ghosal
Shibam Ghosh
Ashrujit Ghoshal
Shane Gibbons
Valerie Gilchrist

Xinxin Gong	Yuval Ishai
Lorenzo Grassi	Ryoma Ito
Scott Griffy	Amit Jana
Chaowen Guan	Ashwin Jha
Aurore Guillevic	Xiaoyu Ji
Sam Gunn	Yanxue Jia
Felix Günther	Mingming Jiang
Kanav Gupta	Lin Jiao
Shreyas Gupta	Haoxiang Jin
Kamil Doruk Gur	Zhengzhong Jin
Jincheol Ha	Chris Jones
Hossein Hadipour	Eliran Kachlon
Tovoheray Hajatiana Randrianarisoa	Giannis Kaklamanis
Shai Halevi	Chethan Kamath
Shuai Han	Soumya Kanti Saha
Tobias Handirk	Sabyasachi Karati
Yonglin Hao	Harish Karthikeyan
Zihan Hao	Andes Y. L. Kei
Keisuke Hara	Jean Kieffer
Keitaro Hashimoto	Jiseung Kim
Aditya Hegde	Seongkwang Kim
Andreas Hellenbrand	Sebastian Kolby
Paul Hermouet	Sreehari Kollath
Minki Hhan	Dimitris Kolonelos
Hilder Lima	Venkata Koppula
Taiga Hiroka	Abhiram Kothapalli
Ryo Hiromasa	Stanislav Kruglik
Viet Tung Hoang	Anup Kumar Kundu
Charlotte Hoffmann	Péter Kutas
Clément Hoffmann	Norman Lahr
Man Hou Hong	Qiqi Lai
Wei-Chih Hong	Yi-Fu Lai
Alexander Hoover	Abel Laval
Fumitaka Hoshino	Guirec Lebrun
Patrick Hough	Byeonghak Lee
Yao-Ching Hsieh	Changmin Lee
Chengcong Hu	Hyung Tae Lee
David Hu	Joohee Lee
Kai Hu	Keewoo Lee
Zihan Hu	Yeongmin Lee
Hai Huang	Yongwoo Lee
Mi-Ying Huang	Andrea Lesavourey
Yu-Hsuan Huang	Baiyu Li
Zhicong Huang	Jiangtao Li
Shih-Han Hung	Jianqiang Li

Junru Li
Liran Li
Minzhang Li
Shun Li
Songsong Li
Weihan Li
Wenzhong Li
Yamin Li
Yanan Li
Yu Li
Yun Li
Zeyong Li
Zhe Li
Chuanwei Lin
Fuchun Lin
Yao-Ting Lin
Yunhao Ling
Eik List
Fengrun Liu
Fukang Liu
Hanlin Liu
Hongqing Liu
Rui Liu
Tianren Liu
Xiang Liu
Xiangyu Liu
Zeyu Liu
Paul Lou
George Lu
Zhenghao Lu
Ting-Gian Lua
You Lyu
Jack P. K. Ma
Yiping Ma
Varun Madathil
Lorenzo Magliocco
Avishek Majumder
Nikolaos Makriyannis
Varun Maram
Chloe Martindale
Elisaweta Masserova
Jake Massimo
Loïc Masure
Takahiro Matsuda
Christian Matt
Subhra Mazumdar
Nikolas Melissaris
Michael Meyer
Ankit Kumar Misra
Anuja Modi
Deep Inder Mohan
Charles Momin
Johannes Mono
Hart Montgomery
Ethan Mook
Thorben Moos
Tomoyuki Morimae
Hiraku Morita
Tomoki Moriya
Aditya Morolia
Christian Mouchet
Nicky Mouha
Tamer Mour
Changrui Mu
Arindam Mukherjee
Pratyay Mukherjee
Anne Müller
Alice Murphy
Shyam Murthy
Kohei Nakagawa
Barak Nehoran
Patrick Neumann
Lucien K. L. Ng
Duy Nguyen
Ky Nguyen
Olga Nissenbaum
Anca Nitulescu
Julian Nowakowski
Frederique Oggier
Jean-Baptiste Orfila
Emmanuela Orsini
Tapas Pal
Ying-yu Pan
Roberto Parisella
Aditi Partap
Alain Passelègue
Alice Pellet-Mary
Zachary Pepin
Octavio Perez Kempner
Edoardo Perichetti

Léo Perrin
Naty Peter
Richard Petri
Rafael del Pino
Federico Pintore
Erik Pohle
Simon Pohmann
Guru Vamsi Policharla
Daniel Pollman
Yuriy Polyakov
Alexander Poremba
Eamonn Postlethwaite
Sihang Pu
Luowen Qian
Tian Qiu
Rajeev Raghunath
Srinivasan Raghuraman
Mostafizar Rahman
Mahesh Rajasree
Somindu Chaya Ramanna
Simon Rastikian
Anik Raychaudhuri
Martin Rehberg
Michael Reichle
Krijn Reijnders
Doreen Riepel
Guilherme Rito
Matthieu Rivain
Bhaskar Roberts
Marc Roeschlin
Michael Rosenberg
Paul Rösler
Arnab Roy
Lawrence Roy
Luigi Russo
Keegan Ryan
Markku-Juhani Saarinen
Éric Sageloli
Dhiman Saha
Sayandeep Saha
Yusuke Sakai
Kosei Sakamoto
Subhabrata Samajder
Simona Samardjiska
Maria Corte-Real Santos
Sina Schaeffler
André Schrottenloher
Jacob Schuldt
Mark Schultz
Mahdi Sedaghat
Jae Hong Seo
Yannick Seurin
Aein Shahmirzadi
Girisha Shankar
Yixin Shen
Rentaro Shiba
Ardeshir Shojaeinasab
Jun Jie Sim
Mark Simkin
Jaspal Singh
Benjamin Smith
Yongha Son
Fang Song
Yongsoo Song
Pratik Soni
Pierre-Jean Spaenlehauer
Matthias Johann Steiner
Lukas Stennes
Roy Stracovsky
Takeshi Sugawara
Adam Suhl
Siwei Sun
Elias Suvanto
Koutarou Suzuki
Erkan Tairi
Atsushi Takayasu
Kaoru Takemure
Abdullah Talayhan
Quan Quan Tan
Gang Tang
Khai Hanh Tang
Tianxin Tang
Yi Tang
Stefano Tessaro
Sri AravindaKrishnan Thyagarajan
Yan Bo Ti
Jean-Pierre Tillich
Toi Tomita
Aleksei Udoenko
Arunachalaeswaran V.

Aron van Baarsen	Kyosuke Yamashita
Wessel van Woerden	Jiayun Yan
Michiel Verbauwhede	Yingfei Yan
Corentin Verhamme	Qianqian Yang
Quoc-Huy Vu	Rupeng Yang
Benedikt Wagner	Xinrui Yang
Julian Wälde	Yibin Yang
Hendrik Waldner	Zhaomin Yang
Judy Walker	Yizhou Yao
Alexandre Wallet	Kevin Yeo
Han Wang	Eylon Yogev
Haoyang Wang	Yusuke Yoshida
Jiabo Wang	Aaram Yun
Jiafan Wang	Gabriel Zaid
Liping Wang	Riccardo Zanotto
Mingyuan Wang	Shang Zehua
Peng Wang	Hadas Zeilberger
Weihao Wang	Runzhi Zeng
Yunhao Wang	Bin Zhang
Zhedong Wang	Cong Zhang
Yohei Watanabe	Liu Zhang
Chenkai Weng	Tianwei Zhang
Andreas Weninger	Tianyu Zhang
Stella Wohnig	Xiangyang Zhang
Harry W. H. Wong	Yijian Zhang
Ivy K. Y. Woo	Yinuo Zhang
Tiger Wu	Yuxin Zhang
Yu Xia	Chang-an Zhao
Zejun Xiang	Tianyu Zhao
Yuting Xiao	Yu Zhou
Ning Xie	Yunxiao Zhou
Zhiye Xie	Zhelei Zhou
Lei Xu	Zibo Zhou
Yanhong Xu	Chenzhi Zhu
Haiyang Xue	Ziqi Zhu
Aayush Yadav	Cong Zuo
Saikumar Yadugiri	

Artifact Chair

Rei Ueno

Kyoto University, Japan

Artifact Evaluation Committee

Julien Béguinot	LTCI, Télécom Paris, Institut Polytechnique de Paris, France
Aron Gohr	Independent Researcher
Hosein Hadipour	Graz University of Technology, Austria
Akira Ito	NTT Social Informatics Laboratories, Japan
Haruto Kimura	University of Melbourne, Australia and Waseda University, Japan
Kotaro Matsuoka	Kyoto University, Japan
Florian Mendel	Infineon Technologies, Germany
Hiraku Morita	Aarhus University, University of Copenhagen, Denmark
Prasanna Ravi	Nanyang Technological University, Singapore
Élise Tasso	Tohoku University, Japan

Contents – Part IX

Quantum Cryptography

Quantum Unpredictability	3
<i>Tomoyuki Morimae, Shogo Yamada, and Takashi Yamakawa</i>	
Quantum Money from Class Group Actions on Elliptic Curves	33
<i>Hart Montgomery and Shahed Sharif</i>	
Generalized Hybrid Search with Applications to Blockchains and Hash Function Security	65
<i>Alexandru Cojocaru, Juan Garay, and Fang Song</i>	
Unclonable Non-interactive Zero-Knowledge	94
<i>Ruta Jawale and Dakshita Khurana</i>	
Unclonable Secret Sharing	129
<i>Prabhanjan Ananth, Vipul Goyal, Jiahui Liu, and Qipeng Liu</i>	
Improved Quantum Lifting by Coherent Measure-and-Reprogram	158
<i>Alexandru Cojocaru, Juan Garay, Qipeng Liu, and Fang Song</i>	
Adaptive Hardcore Bit and Quantum Key Leasing over Classical Channel from LWE with Polynomial Modulus	185
<i>Duong Hieu Phan, Weiqiang Wen, Xingyu Yan, and Jinwei Zheng</i>	
Mind the Bad Norms: Revisiting Compressed Oracle-Based Quantum Indistinguishability Proofs	215
<i>Ritam Bhaumik, Benoît Cogliati, Jordan Ethan, and Ashwin Jha</i>	

Symmetric-key Cryptography

Toward Full n -bit Security and Nonce Misuse Resistance of Block Cipher-Based MACs	251
<i>Wonseok Choi, Jooyoung Lee, and Yeongmin Lee</i>	
General Practical Cryptanalysis of the Sum of Round-Reduced Block Ciphers and ZIP-AES	280
<i>Antonio Flórez-Gutiérrez, Lorenzo Grassi, Gregor Leander, Ferdinand Sibleyras, and Yosuke Todo</i>	

Robust AE With Committing Security 312
Viet Tung Hoang and Sanketh Menda

Author Index 343

Quantum Cryptography



Quantum Unpredictability

Tomoyuki Morimae¹(✉), Shogo Yamada¹, and Takashi Yamakawa^{1,2,3}

¹ Yukawa Institute for Theoretical Physics, Kyoto University, Kyoto, Japan
tomoyuki.morimae@yukawa.kyoto-u.ac.jp

² NTT Social Informatics Laboratories, Tokyo, Japan

³ NTT Research Center for Theoretical Quantum Information, Atsugi, Japan

Abstract. Unpredictable functions (UPFs) play essential roles in classical cryptography, including message authentication codes (MACs) and digital signatures. In this paper, we introduce a quantum analog of UPFs, which we call unpredictable state generators (UPSGs). UPSGs are implied by pseudorandom function-like states generators (PRFSs), which are a quantum analog of pseudorandom functions (PRFs), and therefore UPSGs could exist even if one-way functions do not exist, similar to other recently introduced primitives like pseudorandom state generators (PRSGs), one-way state generators (OWSGs), and EFIs. In classical cryptography, UPFs are equivalent to PRFs, but in the quantum case, the equivalence is not clear, and UPSGs could be weaker than PRFSs. Despite this, we demonstrate that all known applications of PRFSs are also achievable with UPSGs. They include IND-CPA-secure secret-key encryption and EUF-CMA-secure MACs with unclonable tags. Our findings suggest that, for many applications, quantum unpredictability, rather than quantum pseudorandomness, is sufficient.

Keywords: Unpredictability · Secret-key encryption · Message authentication codes with unclonable tags

1 Introduction

1.1 Background

Pseudorandom functions (PRFs), first formalized by Goldreich, Goldwasser and Micali in 1984 [15], are one of the most fundamental primitives in classical cryptography. A PRF is an efficiently-computable keyed function that is computationally indistinguishable from a random function for any polynomial-time adversary that can query the function. PRFs have many important applications in cryptography, and in particular, they are essential building blocks of EUF-CMA-secure message authentication codes (MACs) and IND-CPA-secure secret-key encryption (SKE).

Naor and Reingold [29] introduced a related primitive so-called unpredictable functions (UPFs). Like PRFs, a UPF is an efficiently-computable keyed function, but the crucial difference is that the goal of the adversary is not to distinguish it

from the random function, but to predict the output corresponding to an input that was not queried before. More precisely, let $f := \{f_k\}_k$ be an efficiently-computable keyed function. Then f is a UPF if it satisfies the following property, which is called unpredictability:

$$\Pr[y = f_k(x) : k \leftarrow \{0, 1\}^\lambda, (x, y) \leftarrow \mathcal{A}^{f_k(\cdot)}] \leq \text{negl}(\lambda) \quad (1)$$

for any polynomial-time adversary \mathcal{A} , where x was not queried by \mathcal{A} . It is easy to see that PRFs imply UPFs. The other direction is not straightforward, but Naor and Reingold showed that UPFs imply PRFs [29], and therefore PRFs and UPFs are actually equivalent.

What happens if we consider quantum versions of PRFs and UPFs? Recently, quantum analogs of elementary primitives, including one-way functions (OWFs), pseudorandom generators (PRGs), and PRFs, have been extensively studied [5–8, 10, 18, 25, 27, 28, 32]. For example, pseudorandom states generators (PRSGs) introduced by Ji, Liu, and Song [18] are a quantum analog of PRGs. One-way states generators (OWSGs) introduced by Morimae and Yamakawa [28] are a quantum analog of OWFs. EFIs introduced by Brakerski, Canetti, and Qian [10] are a quantum analog of EFID [14].¹ There are mainly two reasons why studying such new quantum elementary primitives are important. First, they could be weaker than (quantumly-secure) OWFs [22, 23], which are the most fundamental assumption in classical cryptography. More precisely, even if $\mathbf{BQP} = \mathbf{QMA}$ or $\mathbf{P} = \mathbf{NP}$ and therefore OWFs do not exist, these new primitives could exist (relative to oracles). Second, despite that, they have many useful applications, such as private-key quantum money, SKE, non-interactive commitments, digital signatures, and multiparty computations, etc. These facts suggest that these primitives will play the role of the most fundamental assumptions in quantum cryptography, similar to OWFs in classical cryptography.

Quantum versions of PRFs were already studied. There are two quantum analogs of PRFs. One is pseudorandom unitary operators (PRUs) that were introduced by Ji, Liu, and Song [18].² It is a set $\{U_k\}_k$ of efficiently implementable unitary operators that are computationally indistinguishable from Haar random unitary operators. The other quantum analog of PRFs is pseudorandom function-like states (generators) (PRFSs) that were introduced by Ananth, Qian and Yuen [7]. A PRFS is a QPT algorithm that, on input a secret key k and a classical bit string x , outputs a quantum state $\phi_k(x)$. The security roughly means that no QPT adversary can tell whether it is querying to the

¹ An EFID is a pair of two efficiently samplable classical distributions that are statistically far but computationally indistinguishable. An EFI is its quantum analog: a pair of two efficiently generatable quantum states that are statistically far but computationally indistinguishable.

² Weaker variants, so-called pseudorandom states scramblers [24] and pseudorandom isometries [4] were recently introduced. They are shown to be constructed from OWFs.

PRFS oracle or to the oracle that returns Haar random states.³ EUF-CMA-secure MACs (with quantum tags) and IND-CPA-secure SKE (with quantum ciphertexts) can be constructed from PRFSs [7].

On the other hand, no quantum analog of UPFs was explored before. Is it equivalent to a quantum analog of PRFs, such as PRUs or PRFSs? Does it imply EUF-CMA-secure MACs and IND-CPA-secure SKE like PRFSs and PRUs? Can we gain any meaningful insight for quantum cryptography by studying it?

1.2 Our Results

The goal of the present paper is to initiate the study of a quantum version of UPFs which we call unpredictable state generators (UPSGs). We define UPSGs and construct several cryptographic applications from UPSGs. UPSGs are implied by PRFSs, and therefore UPSGs could exist even if OWFs do not exist, similar to PRSGs, OWSGs, and EFIs. As we will explain later, the equivalence between PRFSs and UPSGs are not clear, and UPSGs could be weaker than PRFSs. Despite this, we show that all known applications of PRFSs are also achievable with UPSGs.⁴ This finding provides us with an insightful observation: *For many applications, quantum unpredictability, rather than quantum pseudorandomness, is sufficient.* Relations among our results and known results are summarized in Fig. 1.

Defining UPSGs. Our first contribution is to define UPSGs. A UPSG is a QPT algorithm Eval that, on input a secret key k and a classical bit string x , outputs a quantum state $\phi_k(x)$. Intuitively, the security (unpredictability) is as follows: no QPT adversary, which can query the oracle $\text{Eval}(k, \cdot)$, can output (x^*, ρ) such that x^* was not queried and ρ is close to $\phi_k(x^*)$.⁵

In the classical case, PRFs and UPFs are equivalent [29]. What happens in the quantum case? In fact, we can show that PRFSs imply UPSGs. However, the other direction is not clear. In the classical case, the construction of PRFs from UPFs is done by using the Goldreich-Levin [16, 29]: if $f_k(\cdot)$ is a UPF, $g_{k,r}(x) := f_k(x) \cdot r$ is a PRF with the key (k, r) , where $x \cdot y$ is the inner product between bit strings x and y . However, we cannot directly apply that idea to UPSGs: In particular, what is $\phi_k(x) \cdot r$?

In summary, a quantum analog of UPFs, UPSGs, are implied by PRFSs, which especially means that UPSGs could also exist even if OWFs do not exist.

³ If the query x was not queried before, the oracle samples a new Haar random state ψ_x and outputs it. If the query x was done before, the oracle outputs the same ψ_x that was sampled before.

⁴ Strictly speaking, MACs with unclonable tags that are realized with PRFSs satisfy the security against QPT adversaries that query the oracle *quantumly*, but those realized with UPSGs satisfy that only for the classical oracle query.

⁵ We could consider classical query or quantum query. In the latter case, it is not clear what we mean by “not queried”. One possible formalization, which we actually adopt, is to define that a bit string x was not queried if the weight of $|x\rangle$ is zero for all quantum queries. For more precise statements, see Sect. 3.1.

However, the equivalence is not clear, and UPSGs could be weaker than PRFSs. Then, a natural question is the following: Do UPSGs have useful applications like PRFSs?

IND-CPA-Secure SKE. Our second contribution is to construct IND-CPA-secure SKE (with quantum ciphertexts) from UPSGs. In the classical case, unpredictability implies pseudorandomness [29], which implies encryption. However, in the quantum case, as we have explained before, we do not know how to convert unpredictability to pseudorandomness, and therefore it is not self-evident whether SKE can be constructed from UPSGs. Despite this, we show that it is actually possible:

Theorem 1.1. *If UPSGs exist, then IND-CPA-secure SKE exist.*

IND-CPA-secure SKE can be constructed from PRFSs [7]. Theorem 1.1 shows that such SKE can be constructed from a possibly weaker primitive, UPSGs.

MACs with Unclonable Tags. Our third contribution is to define and construct EUF-CMA-secure MACs with unclonable tags from UPSGs.⁶ The unclonability of tags roughly means that no QPT adversary can, given t -copies of a quantum tag, output a large (possibly entangled) quantum state that contains at least $t+1$ valid tag states. MACs with unclonable tags are useful in practical applications. For example, consider the following attack (which is known as the *replay attack* in the classical cryptography): Alice sends the message “transfer \$100 to Bob” with a MAC tag to a bank. Malicious Bob can steal the pair of the message and the tag, and sends it ten times to the bank so that he can get \$1000. In the classical cryptography, the standard EUF-CMA security of MACs cannot avoid such an attack, and some higher-level treatments are necessary. For example, common techniques are using counters or time-stamps, but they require the time synchronization among users.

If tags are unclonable, we can avoid such a replay attack. Actually, it is easy to see that UPSGs imply EUF-CMA-secure MACs with quantum tags. (We have only to take $\phi_k(x)$ as the tag of the message x .) However, the mere fact that tags are quantum does not automatically imply the unclonability of tags. Moreover, it is not self-evident whether the quantum unpredictability implies unclonability. (Quantum pseudorandomness implies unclonability [18], but it is not clear whether a possibly weaker notion of quantum unpredictability also implies unclonability.) Despite that, we show that MACs with unclonable tags can be constructed from UPSGs.

Theorem 1.2. *If UPSGs exist, then EUF-CMA-secure MACs with unclonable tags exist.*

⁶ We will see that the unclonability of tags automatically implies EUF-CMA security, and therefore we have only to focus on the unclonability of tags.

EUF-CMA-secure MACs with unclonable tags can be constructed from PRFSs [7].⁷ Theorem 1.2 shows that EUF-CMA-secure MACs with unclonable tags can be constructed from a possibly weaker primitive, UPSGs.⁸

Private-Key Quantum Money. The definition of MACs with unclonable tags straightforwardly implies that of private-key quantum money schemes in [18]. We therefore have the following as a corollary of Theorem 1.2. (For the definition of private-key quantum money schemes and a proof of Corollary 1.1, see the full version.)

Corollary 1.1. *If UPSGs exist, then private-key quantum money schemes exist.*

OWSGs and EFIs. IND-CPA-secure SKE implies one-time-secure SKE, and one-time-secure SKE implies OWSGs and EFIs [27]. We therefore have the following as a corollary of Theorem 1.1.

Corollary 1.2. *If UPSGs exist, then OWSGs and EFIs exist.*

However, thus obtained OWSGs are mixed OWSGs (i.e., the ones with mixed states outputs), because ciphertexts of the SKE from UPSGs are mixed states. We can actually directly show that UPSGs imply pure OWSGs, which means UPSGs are broken if $\mathbf{PP} = \mathbf{BQP}$ [12]:

Theorem 1.3. *If UPSGs exist, then pure OWSGs exist and $\mathbf{PP} \neq \mathbf{BQP}$.*

1.3 Technical Overview

IND-CPA-Secure SKE from UPSGs. Let us first recall a construction of IND-CPA-secure SKE from UPFs in classical cryptography. In the classical case, we first use the Goldreich-Levin [16] to construct PRFs from UPFs: Let $f_k(\cdot)$ be a UPF. Then $g_{k,r}(x) := f_k(x) \cdot r$ is a PRF with the key (k, r) [29]. With a PRF $F_k(\cdot)$, an IND-CPA-secure SKE scheme can be constructed as follows: The secret key is the key of the PRF. The ciphertext of a message m is $\text{ct} = (r, F_k(r) \oplus m)$ with a random bit string r .

However, a similar strategy does not work in the quantum case. In particular, we do not know how to convert UPSGs to PRFSs: what is $\phi_k(x) \cdot r$!?

Our idea is to use the duality between the swapping and the distinction [1, 17, 21]. The duality intuitively means that distinguishing two orthogonal states $|\psi\rangle$ and $|\phi\rangle$ is as hard as swapping $|\psi\rangle + |\phi\rangle$ and $|\psi\rangle - |\phi\rangle$ with each other. Our ciphertext for a single bit message $b \in \{0, 1\}$ is, then, $\text{ct}_b := (x, y, |ct_{x,y}^b\rangle)$, where

⁷ [7] only showed that PRFSs imply EUF-CMA-secure MACs with quantum tags, but we can easily show that tags are actually unclonable because their tags are pseudorandom.

⁸ Strictly speaking, there is a difference: MACs with unclonable tags that are realized with PRFSs satisfy the security against QPT adversaries that query the oracle *quantumly*, but those realized with UPSGs satisfy only the security against the classical query.

$|ct_{x,y}^b\rangle := |0\|x\rangle|\phi_k(0\|x)\rangle + (-1)^b|1\|y\rangle|\phi_k(1\|y)\rangle$, and x and y are random bit strings. Here, $|\phi_k(0\|x)\rangle$ and $|\phi_k(1\|y)\rangle$ are outputs of UPSGs on inputs $0\|x$ and $1\|y$, respectively. The secret key of our SKE scheme is the key k of the UPSGs. If a QPT adversary can distinguish ct_0 and ct_1 , then due to the duality, we can construct another QPT adversary that can convert $|\phi_k(0\|x)\rangle$ to $|\phi_k(1\|y)\rangle$. However, it contradicts the unpredictability of the UPSGs.

This argument seems to work. There is, however, one subtle issue here. The adversary of the IND-CPA security can query the encryption oracle, but in general we do not know whether the duality works if the distinguisher queries to an oracle, because the swapping unitary is constructed from the distinguishing unitary and its inverse.

We can solve the issue by observing that the oracle query by the adversary can actually be removed. Because the oracle is an encryption algorithm for single-bit messages and because the adversary queries to the oracle only polynomially many times, we can remove the oracle by giving sufficiently many outputs of the oracle to the adversary in advance as an auxiliary input. The duality in [17] takes into account of the auxiliary inputs to the adversary, and therefore now we can use the duality.

MACs with Unclonable Tags from UPSGs. It is straightforward to see that UPSGs imply EUF-CMA-secure MACs with quantum tags, because we have only to take the output $\phi_k(x)$ of the UPSG on input x as the tag corresponding to the message x . However, the mere fact that the tags are quantum does not automatically mean that they are unclonable. PRFSs also imply EUF-CMA-secure MACs with quantum tags, and in that case, the unclonability of tags is straightforward, because quantum pseudorandomness implies unclonability [18]. However, in the case of UPSGs, it is not clear whether the quantum unpredictability is also sufficient for unclonability.

Our idea to construct unclonable tags is to use the unclonability of random BB84 states. (In other words, to use Wiesner money [31].) Assume that a UPSG exists. Then, there exists an EUF-CMA-secure MAC. (Actually, in the following argument, any EUF-CMA-secure MACs even with classical tags are fine.) Let τ_m be a tag corresponding to a message m . Then, if we set $\tau'_m := \tau_m \otimes |x\rangle\langle x|_\theta$ as a new tag, it becomes unclonable. Here, x, θ are random bit strings, $|x\rangle_\theta := \bigotimes_i H^{\theta^i} |x^i\rangle$, H is the Hadamard gate, and x^i and θ^i are i th bit of x and θ , respectively.

However, the verifier who wants to verify the tag cannot verify τ'_m , because the verifier does not know x and θ . Let us therefore modify our tag as $\tau''_m := (x, \theta, \tau_m \otimes |x\rangle\langle x|_\theta)$. Now, this can be verified by doing the projection onto $|x\rangle_\theta$, but the unclonability is no longer satisfied because x and θ are open.

To solve the issue, we introduce IND-CPA-secure SKE. Fortunately, as we show in this paper, IND-CPA-secure SKE exists if UPSGs exist. Let us modify our tag as $\tau'''_m := \text{Enc}(\text{sk}, (x, \theta)) \otimes \tau_m \otimes |x\rangle\langle x|_\theta$, where Enc is the encryption algorithm of the SKE scheme. Now it is unclonable due to the security of the SKE scheme, but it is no longer authenticated: $\text{Enc}(\text{sk}, (x, \theta)) \otimes \tau_m \otimes |x\rangle\langle x|_\theta$ could be replaced with $\text{Enc}(\text{sk}, (x', \theta')) \otimes \tau_m \otimes |x'\rangle\langle x'|_{\theta'}$ with another x' and θ' chosen by the adversary, because encryption does not necessarily mean authentication. The adversary who knows x' and θ' can of course make many copies of the tag.

The problem is finally solved by considering the following tag: $\tau_m''' := \text{Enc}(\text{sk}, \tau_{m||x||\theta} \otimes |x, \theta\rangle\langle x, \theta|) \otimes |x\rangle\langle x|_{\theta}$, where $\tau_{m||x||\theta}$ is the tag corresponding to the message $m||x||\theta$.

1.4 Open Problems

To conclude Introduction, let us provide some interesting open problems.

1. Do UPSGs imply PRFSs? Or can we separate them?




Fig. 1. Relation among primitives. The red color arrows represent our results. A dotted arrow from primitive A to primitive B represents that primitive A with pure outputs implies primitive B. (Color figure online)

2. Is there any application that is possible with PRFSs, but not with UPSGs? So far, all known applications of PRFSs are achievable with UPSGs.
3. We show that EUF-CMA-secure MACs are possible with UPSGs. How about EUF-CMA-secure digital signatures? Can we realize them with UPSGs? So far, we do not know how to realize them even with PRUs.⁹
4. Do OWSGs imply UPSGs? It is neither known whether PRSGs imply PRFSs.

2 Preliminaries

2.1 Basic Notations

We use the standard notations of quantum computing and cryptography. For a bit string x , x^i denotes the i th bit of x . For two bit strings x and y , $x||y$ means the concatenation of them. We use λ as the security parameter. $[n]$ means the set $\{1, 2, \dots, n\}$. For any set S , $x \leftarrow S$ means that an element x is sampled uniformly at random from the set S . We write negl to mean a negligible function and poly to mean a polynomial. PPT stands for (classical) probabilistic polynomial-time and QPT stands for quantum polynomial-time. For an algorithm A , $y \leftarrow A(x)$ means that the algorithm A outputs y on input x .

For simplicity, we sometimes omit the normalization factor of a quantum state. (For example, we write $\frac{1}{\sqrt{2}}(|x_0\rangle + |x_1\rangle)$ just as $|x_0\rangle + |x_1\rangle$.) $I := |0\rangle\langle 0| + |1\rangle\langle 1|$ is the two-dimensional identity operator. For the notational simplicity, we sometimes write $I^{\otimes n}$ just as I when the dimension is clear from the context. We use X , Y and Z as Pauli operators. For a bit string x , $X^x := \bigotimes_i X^{x^i}$. We use Y^y and Z^z similarly. For two density matrices ρ and σ , the trace distance is defined as $\text{TD}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1 = \frac{1}{2} \text{Tr} \left[\sqrt{(\rho - \sigma)^2} \right]$, where $\|\cdot\|_1$ is the trace norm.

2.2 Lemmas

We use the following lemma by Hhan, Morimae and Yamakawa [17] (based on [1]).

Lemma 2.1 (Duality Between Swapping and Distinction [17], Theorem 5.1). *Let $|\psi\rangle$ and $|\phi\rangle$ be orthogonal n -qubit states. Assume that a QPT algorithm \mathcal{A} with some m -qubit advice state $|\tau\rangle$ can distinguish $|\psi\rangle$ and $|\phi\rangle$ with advantage Δ . Then, there exists a polynomial-time implementable unitary V over $(n+m)$ -qubit states such that*

$$\frac{|\langle \alpha | \langle \tau | V | \beta \rangle | \tau \rangle + \langle \beta | \langle \tau | V | \alpha \rangle | \tau \rangle|}{2} = \Delta, \quad (2)$$

where $|\alpha\rangle := \frac{|\psi\rangle + |\phi\rangle}{\sqrt{2}}$ and $|\beta\rangle := \frac{|\psi\rangle - |\phi\rangle}{\sqrt{2}}$.

⁹ Recently, [13] showed an oracle separation between PRUs and EUF-CMA-secure digital signatures with classical signatures.

We also use the security of Wiesner money [26, 31].

Lemma 2.2 (Security of Wiesner Money [26]). *Let us consider the following security game:*

1. *The challenger \mathcal{C} chooses $x, \theta \leftarrow \{0, 1\}^\lambda$ and sends $|x\rangle_\theta$ to the adversary \mathcal{A} . Here, $|x\rangle_\theta := \bigotimes_{i \in [\lambda]} H^{\theta^i} |x^i\rangle$.*
2. *\mathcal{A} sends a 2λ -qubit state ρ to \mathcal{C} .*
3. *\mathcal{C} projects ρ onto $|x\rangle_\theta^{\otimes 2}$. If the projection is successful, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .*

For any unbounded adversary \mathcal{A} , $\Pr[\top \leftarrow \mathcal{C}] \leq \text{negl}(\lambda)$.

2.3 Cryptographic Primitives

The following is the standard definition of IND-CPA-secure SKE schemes for classical messages. However, in this paper, we consider general cases where ciphertxts can be quantum states.

Definition 2.1 (IND-CPA-Secure SKE for Classical Messages). *An IND-CPA-secure secret-key encryption (SKE) scheme for classical messages is a set of algorithms (KeyGen, Enc, Dec) such that*

- $\text{KeyGen}(1^\lambda) \rightarrow \text{sk}$: *It is a QPT algorithm that, on input the security parameter λ , outputs a classical secret key sk .*
- $\text{Enc}(\text{sk}, m) \rightarrow \text{ct}$: *It is a QPT algorithm that, on input sk and a classical bit string (plaintext) m , outputs a ciphertext ct , which can be a quantum state in general.*
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$: *It is a QPT algorithm that, on input sk and ct , outputs m .*

We require the following two properties.

Correctness: For any bit string m ,

$$\Pr[m \leftarrow \text{Dec}(\text{sk}, \text{ct}) : \text{sk} \leftarrow \text{KeyGen}(1^\lambda), \text{ct} \leftarrow \text{Enc}(\text{sk}, m)] \geq 1 - \text{negl}(\lambda). \quad (3)$$

IND-CPA Security (Against Classical Query): For any QPT adversary \mathcal{A} ,

$$\Pr \left[b = b' : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ (m_0, m_1, \text{st}) \leftarrow \mathcal{A}^{\text{Enc}(\text{sk}, \cdot)} \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{Enc}(\text{sk}, m_b) \\ b' \leftarrow \mathcal{A}^{\text{Enc}(\text{sk}, \cdot)}(\text{st}, \text{ct}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda), \quad (4)$$

where \mathcal{A} can only classically query $\text{Enc}(\text{sk}, \cdot)$.

We also need IND-CPA-secure SKE for quantum messages.

Definition 2.2 (IND-CPA-Secure SKE for Quantum Messages [3,11]).

An IND-CPA-secure secret-key encryption (SKE) scheme for quantum messages is a set of algorithms (KeyGen, Enc, Dec) such that

- KeyGen(1^λ) \rightarrow sk : It is a QPT algorithm that, on input the security parameter λ , outputs a classical secret key sk.
- Enc(sk, ρ) \rightarrow ct : It is a QPT algorithm that, on input sk and a quantum state ρ on the register \mathbf{M} , outputs a quantum state ct on the register \mathbf{C} .
- Dec(sk, ct) \rightarrow ρ : It is a QPT algorithm that, on input sk and a state ct on the register \mathbf{C} , outputs a state ρ on the register \mathbf{M} .

We require the following two properties.

Correctness:

$$\mathbb{E}_{\text{sk} \leftarrow \text{KeyGen}(1^\lambda)} \|\text{Dec}(\text{sk}, \cdot) \circ \text{Enc}(\text{sk}, \cdot) - \text{id}\|_\diamond \leq \text{negl}(\lambda), \quad (5)$$

where id is the identity map, Enc(sk, \cdot) is a CPTP map¹⁰ that runs the encryption algorithm Enc with sk on the plaintext state, Dec(sk, \cdot) is a CPTP map that runs the decryption algorithm Dec with sk on the ciphertext state, and Dec(sk, \cdot) \circ Enc(sk, \cdot) is the composition of Dec(sk, \cdot) and Enc(sk, \cdot). Here $\|\mathcal{F} - \mathcal{E}\|_\diamond := \max_\rho \|(\mathcal{F} \otimes \text{id})(\rho) - (\mathcal{E} \otimes \text{id})(\rho)\|_1$ is the diamond norm between two CPTP maps \mathcal{F} and \mathcal{E} acting on n qubits [30], where the max is taken over all $2n$ -qubit states ρ .

IND-CPA Security: Let us consider the following security game:

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} can query the oracle Enc(sk, \cdot). (This means that \mathcal{A} can apply the CPTP map Enc(sk, \cdot) on the register \mathbf{M} of any \mathcal{A} 's state $\rho_{\mathbf{M}, \mathbf{Z}}$ over the registers \mathbf{Z} and \mathbf{M} , and get another state $\rho'_{\mathbf{Z}, \mathbf{C}}$ over the registers \mathbf{Z} and \mathbf{C} .)
3. \mathcal{A} sends two registers \mathbf{M}_0 and \mathbf{M}_1 to \mathcal{C} .
4. \mathcal{C} chooses $b \leftarrow \{0, 1\}$ and applies the CPTP map Enc(sk, \cdot) on \mathbf{M}_b . \mathcal{C} then sends the output to \mathcal{A} .
5. \mathcal{A} can query the oracle Enc(sk, \cdot).
6. \mathcal{A} sends $b' \in \{0, 1\}$ to \mathcal{C} .
7. If $b = b'$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

For any QPT adversary \mathcal{A} , $\Pr[\top \leftarrow \mathcal{C}] \leq \frac{1}{2} + \text{negl}(\lambda)$.

The following lemma is essentially shown in [11]. We give its proof in the full version.

¹⁰ In this paper, we sometimes use the same notation Enc for an algorithm and a CPTP map, but we believe there is no confusion.

Lemma 2.3 (IND-CPA security for classical messages implies that for quantum messages [11]). *If IND-CPA-secure SKE schemes for classical messages that are secure against QPT adversaries that query the encryption oracle classically exist, then IND-CPA-secure SKE schemes for quantum messages exist.*

The following lemma can be shown with the standard hybrid argument [11].

Lemma 2.4 (IND-CPA-multi security [11]). *Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA-secure SKE scheme for quantum messages. Let t be a polynomial. Let us consider the security game that is the same as that of Definition 2.2 except for the following two modifications.*

- In step 3, \mathcal{A} sends two registers \mathbf{M}'_0 and \mathbf{M}'_1 to \mathcal{C} . Here, \mathbf{M}'_0 consists of t registers $\{\mathbf{M}'_0^i\}_{i \in [t]}$, and \mathbf{M}'_1 consists of t registers $\{\mathbf{M}'_1^i\}_{i \in [t]}$. For each $i \in [t]$ and $b \in \{0, 1\}$, $|\mathbf{M}'_b^i| = |\mathbf{M}_b|$, where $|\mathbf{A}|$ is the size (i.e., the number of qubits) of the register \mathbf{A} .
- In step 4, \mathcal{C} chooses $b \leftarrow \{0, 1\}$ and applies the CPTP map $\text{Enc}(\text{sk}, \cdot)$ on each \mathbf{M}'_b^i for $i \in [t]$. \mathcal{C} then sends the all outputs to \mathcal{A} .

Then, in this modified game, $\Pr[\top \leftarrow \mathcal{C}] \leq \frac{1}{2} + \text{negl}(\lambda)$ for any QPT adversary \mathcal{A} and any polynomial t .

Definition 2.3 (One-way States Generators (OWSGs) [27,28]). *A one-way states generator (OWSG) is a set of algorithms $(\text{KeyGen}, \text{StateGen}, \text{Ver})$ such that*

- $\text{KeyGen}(1^\lambda) \rightarrow k$: It is a QPT algorithm that, on input the security parameter λ , outputs a classical key k .
- $\text{StateGen}(k) \rightarrow \phi_k$: It is a QPT algorithm that, on input k , outputs a quantum state ϕ_k .
- $\text{Ver}(k', \phi_k) \rightarrow \top/\perp$: It is a QPT algorithm that, on input ϕ_k and a bit string k' , outputs \top or \perp .

We require the following correctness and security.

Correctness:

$$\Pr[\top \leftarrow \text{Ver}(k, \phi_k) : k \leftarrow \text{KeyGen}(1^\lambda), \phi_k \leftarrow \text{StateGen}(k)] \geq 1 - \text{negl}(\lambda). \quad (6)$$

Security: For any QPT adversary \mathcal{A} and any polynomial t ,

$$\Pr[\top \leftarrow \text{Ver}(k', \phi_k) : k \leftarrow \text{KeyGen}(1^\lambda), \phi_k^{\otimes t} \leftarrow \text{StateGen}(k)^{\otimes t}, k' \leftarrow \mathcal{A}(1^\lambda, \phi_k^{\otimes t})] \leq \text{negl}(\lambda). \quad (7)$$

Here, $\phi_k^{\otimes t} \leftarrow \text{StateGen}(k)^{\otimes t}$ means that the StateGen algorithm is run t times.

3 Unpredictable State Generators

3.1 Definition

In this subsection, we define UPSGs. The syntax is given as follows.

Definition 3.1 (Unpredictable States Generators (UPSGs)). *An unpredictable states generator is a set $(\text{KeyGen}, \text{Eval})$ of QPT algorithms such that*

- $\text{KeyGen}(1^\lambda) \rightarrow k$: *It is a QPT algorithm that, on input the security parameter λ , outputs a classical key k .*
- $\text{Eval}(k, x) \rightarrow (x, \phi_k(x))$: *It is a QPT algorithm that, on input k and a bit string x , outputs x and a quantum state $\phi_k(x)$.*

In general, $\phi_k(x)$ could be mixed states, but in this paper, we restrict them to pure states.

The security, which we call unpredictability, roughly means that no QPT adversary (who can *quantumly* query to $\text{Eval}(k, \cdot)$) can output (x^*, ρ) such that x^* was *not queried* and ρ is close to $|\phi_k(x^*)\rangle$. In order to formally define it, we have to clarify what we mean by “quantumly query” and “not queried before”.

Quantum Query. We assume that $|\phi_k(x)\rangle \leftarrow \text{Eval}(k, x)$ is the following QPT algorithm: on input k and x , it applies a unitary U_k on $|x\rangle_{\mathbf{X}}|0\dots 0\rangle_{\mathbf{Y}, \mathbf{Z}}$ to generate $|x\rangle_{\mathbf{X}}|\phi_k(x)\rangle_{\mathbf{Y}}|\text{junk}_k\rangle_{\mathbf{Z}}$ and outputs the \mathbf{X} and \mathbf{Y} registers. Note that it is not the most general case. First, as we have mentioned, we assume that the output $|\phi_k(x)\rangle$ is pure. Second, in general, the junk state $|\text{junk}_k\rangle$ could depend on x , but we here assume that it depends only k . These two restrictions seem to be necessary to well define the quantum query.

With such Eval , the quantum query to the oracle $\text{Eval}(k, \cdot)$ means the following:

1. A state $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |\xi_x\rangle$ is input to the oracle, where $\{\alpha_x\}_x$ are any complex coefficients and $\{|\xi_x\rangle\}_x$ are any states.
2. The oracle adds the ancilla state $|0\dots 0\rangle_{\mathbf{Y}, \mathbf{Z}}$ and applies U_k on the registers $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ of $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |0\dots 0\rangle_{\mathbf{Y}, \mathbf{Z}} |\xi_x\rangle$ to generate $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |\phi_k(x)\rangle_{\mathbf{Y}} |\text{junk}_k\rangle_{\mathbf{Z}} |\xi_x\rangle$.
3. The oracle removes the junk register \mathbf{Z} and outputs the state $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |\phi_k(x)\rangle_{\mathbf{Y}} |\xi_x\rangle$.

Remark 3.1. If the output states are not pure, we cannot define such a quantum query. In that case, we can consider a classical query as in [5]. However, in this work, we focus on the case when the output states are pure and on the quantum query.

Not Queried. We define the word “not queried” as follows. Assume that \mathcal{A} queries the oracle q times. For each $i \in [q]$, let $|\psi_i\rangle$ be the entire \mathcal{A} ’s state immediately before its i th query to the oracle. (Without loss of generality, we can assume that \mathcal{A} postpones all measurements to the last step, and then \mathcal{A} ’s entire state is always pure.) We say that x^* is not queried if $\langle \psi_i | (|x^*\rangle\langle x^*|_{\mathbf{X}} \otimes I) | \psi_i \rangle = 0$ for all $i \in [q]$. Here, for each $i \in [q]$, $|\psi_i\rangle = \sum_x c_x |x\rangle_{\mathbf{X}} \otimes |\eta_x\rangle$.

Now we define the unpredictability.

Definition 3.2 (Unpredictability). *Let us consider the following security game:*

1. The challenger \mathcal{C} runs $k \leftarrow \text{KeyGen}(1^\lambda)$.
2. The adversary $\mathcal{A}^{\text{Eval}(k, \cdot)}(1^\lambda)$ outputs a bit string x^* and a quantum state ρ , and sends them to \mathcal{C} . Here, \mathcal{A} can make quantum queries to $\text{Eval}(k, \cdot)$. x^* should not be queried by \mathcal{A} .
3. \mathcal{C} projects ρ onto $|\phi_k(x^*)\rangle$. If the projection is successful, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

For any QPT adversary \mathcal{A} , $\Pr[\top \leftarrow \mathcal{C}] \leq \text{negl}(\lambda)$.

Remark 3.2. Note that the projection of ρ onto $|\phi_k(x^*)\rangle$ can be done as follows:

1. Prepare $|x^*\rangle\langle x^*| \otimes \rho \otimes |\text{junk}_k\rangle\langle \text{junk}_k|$.
2. Apply U_k^\dagger on $|x^*\rangle\langle x^*| \otimes \rho \otimes |\text{junk}_k\rangle\langle \text{junk}_k|$.
3. Measure all qubits in the computational basis. If the result is $x^*||0\dots 0$, the projection is successful. Otherwise, the projection is failed.

Remark 3.3. It is easy to see that UPSGs with $O(\log \lambda)$ -qubit output do not exist.¹¹

Remark 3.4. In [9], they define a security of digital signatures against quantum adversaries. Their security definition is as follows: any QPT quantum adversary, who queries the signing oracle t times, cannot output $t+1$ valid message-signature pairs. We could define a quantum version of unpredictability based on their security definition, but exploring this possibility is beyond the scope of the present paper. At least, their definition seems to be incomparable to Definition 3.2. In particular, we do not know how to construct IND-CPA-secure SKE from their definition, because we do not know how to use the duality in that case.

3.2 Relation to PRFSs

In this section, we recall the definition of PRFSs and construct UPSGs from PRFSs.

Definition 3.3 (Pseudorandom Function-Like States (PRFSs) [5, 7]). *A pseudorandom function-like state (PRFS) (generator) is a set of algorithms (KeyGen, Eval) such that*

¹¹ The adversary has only to output $0\dots 0$ and maximally-mixed state.

- $\text{KeyGen}(1^\lambda) \rightarrow k$: It is a QPT algorithm that, on input the security parameter λ , outputs a classical secret key k .
- $\text{Eval}(k, x) \rightarrow |\phi_k(x)\rangle$: It is a QPT algorithm that on input k and a bit string x , outputs a quantum state $|\phi_k(x)\rangle$.

We require the following security. For any QPT adversary \mathcal{A} ,

$$|\Pr[1 \leftarrow \mathcal{A}^{\text{Eval}(k, \cdot)}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Haar}}}(1^\lambda)]| \leq \text{negl}(\lambda). \quad (8)$$

Here, $\mathcal{A}^{\text{Eval}(k, \cdot)}$ means that \mathcal{A} can quantumly query the oracle $\text{Eval}(k, \cdot)$ in the sense of Sect. 3.1¹². $\mathcal{A}^{\mathcal{O}_{\text{Haar}}}$ means that \mathcal{A} can quantumly query the oracle $\mathcal{O}_{\text{Haar}}$ in the following sense.

1. A state $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |\xi_x\rangle$ is input to the oracle, where $\{\alpha_x\}_x$ are any complex coefficients and $\{|\xi_x\rangle\}_x$ are any states.
2. The oracle returns $\sum_x \alpha_x |x\rangle_{\mathbf{X}} |\psi_x\rangle_{\mathbf{Y}} |\xi_x\rangle$, where $|\psi_x\rangle$ is a Haar random state.

Theorem 3.1. *If PRFSs exist then UPSGs exist.*

Proof of Theorem 3.1. Let $(\text{KeyGen}, \text{Eval})$ be a PRFS. We show that it is a UPSG. Assume that it does not satisfy the unpredictability. Then, there exist a polynomial p and a QPT adversary \mathcal{A} that can quantumly query $\text{Eval}(k, \cdot)$ such that

$$\sum_k \Pr[k \leftarrow \text{KeyGen}(1^\lambda)] \sum_{x^*} \langle x^* | \langle \phi_k(x^*) | \mathcal{A}^{\text{Eval}(k, \cdot)}(1^\lambda) | x^* \rangle | \phi_k(x^*) \rangle \geq \frac{1}{p(\lambda)} \quad (9)$$

for infinitely many $\lambda \in \mathbb{N}$. Here, $\mathcal{A}^{(\cdot)}(1^\lambda)$ denotes the state of $\mathcal{A}^{(\cdot)}$ before the measurement. Then, the following QPT adversary \mathcal{B} breaks the security of PRFS.

1. The challenger \mathcal{C}' of the PRFS chooses $b \leftarrow \{0, 1\}$.
2. Run \mathcal{A} on input 1^λ . When \mathcal{A} queries the oracle, \mathcal{B} simulates it by querying \mathcal{B} 's oracle (that is $\text{Eval}(k, \cdot)$ if $b = 0$ and $\mathcal{O}_{\text{Haar}}$ if $b = 1$).
3. \mathcal{B} measures the first register of $\mathcal{A}^{(\cdot)}(1^\lambda)$ to get x^* . Query x^* to \mathcal{B} 's oracle to get $|\xi\rangle$, which is $|\xi\rangle = |\phi_k(x^*)\rangle$ if $b = 0$ and a Haar random state $|\psi_{x^*}\rangle$ if $b = 1$.
4. \mathcal{B} does the swap test between the second register of $\mathcal{A}^{(\cdot)}(1^\lambda)$ and $|\xi\rangle$. If the swap test succeeds, \mathcal{B} outputs 1. Otherwise, \mathcal{B} outputs 0.

If $b = 0$,

$$\Pr[1 \leftarrow \mathcal{B}] = \frac{1}{2} + \frac{1}{2} \sum_k \Pr[k \leftarrow \text{KeyGen}(1^\lambda)] \sum_{x^*} \langle x^* | \langle \phi_k(x^*) | \mathcal{A}^{\text{Eval}(k, \cdot)}(1^\lambda) | x^* \rangle | \phi_k(x^*) \rangle \quad (10)$$

$$\geq \frac{1}{2} + \frac{1}{2p(\lambda)} \quad (11)$$

¹² In [5, 7], they do not explicitly consider the junk state $|\text{junk}_k\rangle$. Here, we assume that $|\text{junk}_k\rangle$ is independent of x similarly to the case of UPSGs.

for infinitely many λ . Here we have used Eq. (9). On the other hand, if $b = 1$,

$$\Pr[1 \leftarrow \mathcal{B}] = \frac{1}{2} + \frac{1}{2} \sum_k \Pr[k \leftarrow \text{KeyGen}(1^\lambda)] \sum_{x^*} \mathbb{E}_{|\psi\rangle \leftarrow \mu} \langle x^* | \langle \psi | \mathcal{A}^{\text{OHar}}(1^\lambda) | x^* \rangle | \psi \rangle \quad (12)$$

$$\leq \frac{1}{2} + \text{negl}(\lambda), \quad (13)$$

where μ denotes the Haar measure and we have used $\mathbb{E}_{|\psi\rangle \leftarrow \mu} \langle \psi | \sigma | \psi \rangle \leq \text{negl}(\lambda)$ for any state σ . Therefore, \mathcal{B} breaks the security of the PRFS. \square

3.3 Pure OWSGs from UPSGs

In this section, we show that UPSGs imply OWSGs with pure output states.

Theorem 3.2. *If UPSGs exist, then pure OWSGs exist.*

Proof. Let $(\text{UPSG.KeyGen}, \text{UPSG.Eval})$ be a UPSG. From it, we construct a pure OWSG $(\text{KeyGen}, \text{StateGen})$ as follows.

- $\text{KeyGen}(1^\lambda) \rightarrow k' : \text{Run } k \leftarrow \text{UPSG.KeyGen}(1^\lambda)$. Choose $x_i \leftarrow \{0, 1\}^\ell$ for $i \in [n]$. Here, $n := |k| + \lambda$. Output $k' := (k, x_1, \dots, x_n)$.
- $\text{StateGen}(k') \rightarrow |\psi_{k'}\rangle : \text{Parse } k' = (k, x_1, \dots, x_n)$. Run $|\phi_k(x_i)\rangle \leftarrow \text{UPSG.Eval}(k, x_i)$ for $i \in [n]$. Output $\psi_{k'} := (\bigotimes_{i=1}^n |\phi_k(x_i)\rangle) \otimes (\bigotimes_{i=1}^n |x_i\rangle)$.

For the sake of contradiction, assume that this construction is not secure. This means that there exist polynomials p and t , and a QPT adversary \mathcal{A} such that

$$\frac{1}{p(\lambda)} \leq \sum_k \Pr[k] \frac{1}{2^{n\ell}} \sum_{x_1, \dots, x_n} \sum_{s, x'_1, \dots, x'_n} \Pr'[s, x'_1, \dots, x'_n | k, x_1, \dots, x_n] \prod_{i \in [n]} |\langle \phi_k(x_i) | \phi_s(x'_i) \rangle|^2 \delta_{x_i, x'_i} \quad (14)$$

$$= \sum_k \Pr[k] \frac{1}{2^{n\ell}} \sum_{x_1, \dots, x_n} \sum_s \Pr'[s, x_1, \dots, x_n | k, x_1, \dots, x_n] \prod_{i \in [n]} |\langle \phi_k(x_i) | \phi_s(x_i) \rangle|^2 \quad (15)$$

for infinitely many λ . Here, $\Pr[k] := \Pr[k \leftarrow \text{UPSG.KeyGen}(1^\lambda)]$ and

$$\Pr'[s, x'_1, \dots, x'_n | k, x_1, \dots, x_n] := \Pr[(s, x'_1, \dots, x'_n) \leftarrow \mathcal{A}(1^\lambda, ((\bigotimes_i |\phi_k(x_i)\rangle) \otimes (\bigotimes_i |x_i\rangle))^{\otimes t})]. \quad (16)$$

From the \mathcal{A} , we construct a QPT adversary \mathcal{B} that breaks the security of the UPSG as follows.

1. Sample $x_1, \dots, x_n \leftarrow \{0, 1\}^\ell$ and $x^* \leftarrow \{0, 1\}^\ell$.
2. For each $i \in [n]$, query x_i to the oracle $\text{UPSG.Eval}(k, \cdot)$ t times to get $|\phi_k(x_i)\rangle^{\otimes t}$.
3. Run $(s, x'_1, \dots, x'_n) \leftarrow \mathcal{A}(1^\lambda, (\bigotimes_{i=1}^n |\phi_k(x_i)\rangle) \otimes (\bigotimes_{i=1}^n |x_i\rangle))^{\otimes t}$. If $x'_i \neq x_i$ for at least one $i \in [n]$, abort.
4. Run $|\phi_k(x^*)\rangle \leftarrow \text{UPSG.Eval}(s, x^*)$. Output $(x^*, |\phi_k(x^*)\rangle)$.

By the standard average argument, we can show that the probability that \mathcal{B} wins is non-negligible for infinitely many λ . For its proof, see the full version. \square

4 IND-CPA-Secure SKE from UPSGs

In this section, we construct IND-CPA secure SKE from UPSGs.

Theorem 4.1. *If UPSGs exist, then IND-CPA-secure SKE schemes for classical messages secure against classically querying QPT adversaries exist.*

Remark 4.1. From Lemma 2.3, IND-CPA-secure SKE schemes for classical messages secure against classically querying QPT adversaries imply IND-CPA-secure SKE schemes for quantum messages. Therefore, the above theorem also shows the existence of such SKE schemes if UPSGs exist.

Proof of Theorem 4.1. It suffices to construct an IND-CPA-secure SKE scheme for single-bit messages because, from it, we can construct an IND-CPA-secure SKE scheme for multi-bit messages by parallel repetition.¹³ Let $(\text{UPSG.KeyGen}, \text{UPSG.Eval})$ be a UPSG. As is explained in Sect. 3.1, we assume that UPSG.Eval is the following algorithm: on input k and $x \in \{0, 1\}^\ell$, it applies a unitary U_k on $|x\rangle_{\mathbf{X}}|0\dots 0\rangle_{\mathbf{Y}, \mathbf{Z}}$ to generate $|x\rangle_{\mathbf{X}}|\phi_k(x)\rangle_{\mathbf{Y}}|\text{junk}_k\rangle_{\mathbf{Z}}$, and outputs the registers \mathbf{X} and \mathbf{Y} . From $(\text{UPSG.KeyGen}, \text{UPSG.Eval})$, we construct an IND-CPA-secure SKE scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ for single-bit messages as follows.

- $\text{KeyGen}(1^\lambda) \rightarrow \text{sk}$: Run $k \leftarrow \text{UPSG.KeyGen}(1^\lambda)$ and output $\text{sk} := k$.
- $\text{Enc}(\text{sk}, b) \rightarrow \text{ct}$: Parse $\text{sk} = k$. Choose $x, y \leftarrow \{0, 1\}^\ell$. Generate

$$|\text{ct}_{x,y}^b\rangle_{\mathbf{X}, \mathbf{Y}} := \frac{|0\rangle_{\mathbf{X}}|\phi_k(0|x)\rangle_{\mathbf{Y}} + (-1)^b|1\rangle_{\mathbf{X}}|\phi_k(1|y)\rangle_{\mathbf{Y}}}{\sqrt{2}} \quad (17)$$

and output $\text{ct} := (x, y, |\text{ct}_{x,y}^b\rangle)$. Here, $|\text{ct}_{x,y}^b\rangle$ is generated as follows:

1. Prepare $|0\rangle_{\mathbf{X}}|\phi_k(0|x)\rangle_{\mathbf{Y}, \mathbf{Z}} + (-1)^b|1\rangle_{\mathbf{X}}|\phi_k(1|y)\rangle_{\mathbf{Y}, \mathbf{Z}}$.
2. Apply U_k on the registers \mathbf{X} , \mathbf{Y} , and \mathbf{Z} to generate

$$|0\rangle_{\mathbf{X}}|\phi_k(0|x)\rangle_{\mathbf{Y}}|\text{junk}_k\rangle_{\mathbf{Z}} + (-1)^b|1\rangle_{\mathbf{X}}|\phi_k(1|y)\rangle_{\mathbf{Y}}|\text{junk}_k\rangle_{\mathbf{Z}}. \quad (18)$$

3. Remove the register \mathbf{Z} .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow b'$: Parse $\text{sk} = k$ and $\text{ct} = (x, y, \rho_{\mathbf{X}, \mathbf{Y}})$. Run the following algorithm.
 1. Prepare $\rho_{\mathbf{X}, \mathbf{Y}} \otimes |\text{junk}_k\rangle\langle \text{junk}_k|_{\mathbf{Z}}$.
 2. Apply U_k^\dagger on $\rho_{\mathbf{X}, \mathbf{Y}} \otimes |\text{junk}_k\rangle\langle \text{junk}_k|_{\mathbf{Z}}$.
 3. Apply $|0\rangle\langle 0| \otimes X^x + |1\rangle\langle 1| \otimes X^y$ on the register \mathbf{X} .
 4. Measure the first qubit of the register \mathbf{X} in the Hadamard basis to get $b' \in \{0, 1\}$. Output b' .

Correctness is clear. To show the security, we define Hybrid 0, which is the original security game of the IND-CPA-secure SKE scheme between the challenger \mathcal{C} and the QPT adversary \mathcal{A} , as follows.

¹³ See [19].

Hybrid 0

1. The challenger \mathcal{C} runs $k \leftarrow \text{UPSG.KeyGen}(1^\lambda)$.
2. \mathcal{C} chooses $b \leftarrow \{0, 1\}$ and $x, y \leftarrow \{0, 1\}^\ell$. \mathcal{C} generates $|\text{ct}_{x,y}^b\rangle$ by running $\text{UPSG.Eval}(k, \cdot)$ coherently. Here,

$$|\text{ct}_{x,y}^b\rangle = \frac{|0\rangle\langle x| \phi_k(0|x) + (-1)^b |1\rangle\langle y| \phi_k(1|y)}{\sqrt{2}}. \quad (19)$$

3. \mathcal{C} sends $\text{ct} := (x, y, |\text{ct}_{x,y}^b\rangle)$ to the adversary \mathcal{A} .
4. \mathcal{A} can classically query to the oracle \mathcal{O}_k , where \mathcal{O}_k works as follows:
 - (a) On input $c \in \{0, 1\}$, it chooses $x', y' \leftarrow \{0, 1\}^\ell$ and generates $|\text{ct}_{x',y'}^c\rangle$.
 - (b) It outputs $(x', y', |\text{ct}_{x',y'}^c\rangle)$. \mathcal{A} sends $b' \in \{0, 1\}$ to \mathcal{C} .
5. If $b = b'$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

For the sake of contradiction, assume that our construction is not IND-CPA secure. This means that there exist a polynomial p and a QPT adversary \mathcal{A} such that

$$\Pr[\top \leftarrow \text{Hybrid 0}] \geq \frac{1}{2} + \frac{1}{p(\lambda)} \quad (20)$$

for infinitely-many $\lambda \in \mathbb{N}$.

Our goal is to construct a QPT adversary \mathcal{B} that breaks the unpredictability of the UPSG. For that goal, we use the duality between swapping and distinction [17]. However, we cannot directly use it here, because our \mathcal{A} queries to the encryption oracle \mathcal{O}_k , but the distinguisher in Lemma 2.1 does not access any oracle. To solve the issue, we have to remove the oracle \mathcal{O}_k from Hybrid 0. Fortunately, \mathcal{O}_k is an encryption oracle for single-bit messages, and \mathcal{A} makes classical queries only polynomial times. Therefore, we can give \mathcal{A} enough number of outputs of \mathcal{O}_k in advance as auxiliary inputs, and \mathcal{A} can use these states instead of the outputs of \mathcal{O}_k . In this way, we can remove the oracle \mathcal{O}_k . We formalize this as Hybrid 1.¹⁴ It is clear that $\Pr[\top \leftarrow \text{Hybrid 1}] = \Pr[\top \leftarrow \text{Hybrid 0}]$.

Hybrid 1

1. The challenger \mathcal{C} runs $k \leftarrow \text{UPSG.KeyGen}(1^\lambda)$.
2. \mathcal{C} chooses $b \leftarrow \{0, 1\}$ and $x, y \leftarrow \{0, 1\}^\ell$. \mathcal{C} generates $|\text{ct}_{x,y}^b\rangle$ by running $\text{UPSG.Eval}(k, \cdot)$ coherently. Here,

$$|\text{ct}_{x,y}^b\rangle = \frac{|0\rangle\langle x| \phi_k(0|x) + (-1)^b |1\rangle\langle y| \phi_k(1|y)}{\sqrt{2}}. \quad (21)$$

3. \mathcal{C} sends $\text{ct} := (x, y, |\text{ct}_{x,y}^b\rangle)$ to the adversary \mathcal{A} .
4. ~~\mathcal{A} can classically query to the oracle \mathcal{O}_k , where \mathcal{O}_k works as follows:~~

¹⁴ In Hybrid 1, text struck through with red is the step in the previous hybrid, and the red text is the new step in the current hybrid.

- (a) ~~On input $c \in \{0, 1\}$, it chooses $x', y' \leftarrow \{0, 1\}^\ell$ and generates $|\text{ct}_{x', y'}^c\rangle$.~~
 (b) ~~It outputs $(x', y', |\text{ct}_{x', y'}^c\rangle)$.~~

\mathcal{A} receives $|\tau\rangle := \bigotimes_{i \in [t], c \in \{0, 1\}} |x_c^i\rangle |y_c^i\rangle |\text{ct}_{x_c^i, y_c^i}^c\rangle$ as an auxiliary input, where t is the maximum number of \mathcal{A} 's queries to \mathcal{O}_k in the step 4 of Hybrid 0, and $x_c^i, y_c^i \leftarrow \{0, 1\}^\ell$ for each $i \in [t]$ and $c \in \{0, 1\}$. When \mathcal{A} queries $c_i \in \{0, 1\}$ to \mathcal{O}_k in its i th query, it does not query to \mathcal{O}_k . Instead, it uses $|x_c^i\rangle |y_c^i\rangle |\text{ct}_{x_c^i, y_c^i}^c\rangle$ as the output of \mathcal{O}_k . \mathcal{A} sends $b' \in \{0, 1\}$ to \mathcal{C} .

5. If $b = b'$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

Let $\mathbf{w} := \{x_c^i, y_c^i\}_{i \in [t], c \in \{0, 1\}}$, where $x_c^i \in \{0, 1\}^\ell$ and $y_c^i \in \{0, 1\}^\ell$ for each $i \in [t]$ and $c \in \{0, 1\}$. Let $\Pr[\top \leftarrow \text{Hybrid 1} | k, x, y, \mathbf{w}]$ be the conditional probability that \mathcal{C} outputs \top given $k \leftarrow \text{UPSG.KeyGen}(1^\lambda)$ and x, y, \mathbf{w} are chosen in Hybrid 1. We define a ‘‘good’’ set of (k, x, y, \mathbf{w}) as follows:

$$G := \left\{ (k, x, y, \mathbf{w}) : \Pr[\top \leftarrow \text{Hybrid 1} | k, x, y, \mathbf{w}] \geq \frac{1}{2} + \frac{1}{2p(\lambda)} \wedge x \notin \mathbf{w} \wedge y \notin \mathbf{w} \wedge x \neq y \right\}. \quad (22)$$

Let $\Pr[k, x, y, \mathbf{w}]$ be the probability that k, x, y and \mathbf{w} are chosen in Hybrid 1. Then, we can show the following lemma by the standard average argument. For its proof, see the full version.

Lemma 4.1. $\sum_{(k, x, y, \mathbf{w}) \in G} \Pr[k, x, y, \mathbf{w}] \geq \frac{1}{4p(\lambda)}$ for infinitely many $\lambda \in \mathbb{N}$.

Let us fix (k, x, y, \mathbf{w}) . Moreover, assume that $(k, x, y, \mathbf{w}) \in G$. Then from Eq. (22), \mathcal{A} of Hybrid 1 can distinguish $|\text{ct}_{x, y}^0\rangle$ and $|\text{ct}_{x, y}^1\rangle$ with an advantage greater than $\frac{1}{2p}$ using the auxiliary input $|\tau\rangle$. By using Lemma 2.1, we can construct a polynomial-time implementable unitary V^{15} such that

$$\frac{1}{2p(\lambda)} \leq \frac{|\langle 0 \| x \rangle \langle \phi_k(0 \| x) | \langle \tau | V | 1 \| y \rangle \langle \phi_k(1 \| y) \rangle | \tau \rangle + \langle 1 \| y \rangle \langle \phi_k(1 \| y) | \langle \tau | V | 0 \| x \rangle \langle \phi_k(0 \| x) \rangle | \tau \rangle|}{2} \quad (23)$$

$$\leq \max\{|\langle 0 \| x \rangle \langle \phi_k(0 \| x) | \langle \tau | V | 1 \| y \rangle \langle \phi_k(1 \| y) \rangle | \tau \rangle|, |\langle 1 \| y \rangle \langle \phi_k(1 \| y) | \langle \tau | V | 0 \| x \rangle \langle \phi_k(0 \| x) \rangle | \tau \rangle|\} \quad (24)$$

$$\leq \max\{\|\langle \phi_k(0 \| x) |_{\mathbf{Y}} (V | 1 \| y)_{\mathbf{X}} \langle \phi_k(1 \| y) \rangle_{\mathbf{Y}} | \tau \rangle_{\mathbf{Z}}\|, \quad (25)$$

$$\|\langle \phi_k(1 \| y) |_{\mathbf{Y}} (V | 0 \| x)_{\mathbf{X}} \langle \phi_k(0 \| x) \rangle_{\mathbf{Y}} | \tau \rangle_{\mathbf{Z}}\|\}. \quad (26)$$

From this V , we construct the QPT adversary \mathcal{B} that breaks the security of the UPSG as follows:

1. Choose $b \leftarrow \{0, 1\}$.
2. Choose $x, y \leftarrow \{0, 1\}^\ell$. If $x = y$, output \perp and abort. Choose $x_c^i \leftarrow \{0, 1\}^\ell$ and $y_c^i \leftarrow \{0, 1\}^\ell$ for each $i \in [t]$ and $c \in \{0, 1\}$. Set $\mathbf{w} := \{x_c^i, y_c^i\}_{i \in [t], c \in \{0, 1\}}$. If $x \in \mathbf{w}$ or $y \in \mathbf{w}$, output \perp and abort.

¹⁵ Note that this V is independent of (k, x, y, \mathbf{w}) since, in the proof of Lemma 2.1, we use \mathcal{A} only as a black-box. For details, see [17].

3. If $b = 0$, get $|0\|x\rangle|\phi_k(0\|x)\rangle$ by querying $0\|x$ to $\text{UPSG.Eval}(k, \cdot)$. If $b = 1$, get $|1\|y\rangle|\phi_k(1\|y)\rangle$ by querying $1\|y$ to $\text{UPSG.Eval}(k, \cdot)$.
4. For each $i \in [t]$ and $c \in \{0, 1\}$, generate $|\text{ct}_{x_c^i, y_c^i}^c\rangle$ by making the coherent query $|0\|x_c^i\rangle + (-1)^c|1\|y_c^i\rangle$ to $\text{UPSG.Eval}(k, \cdot)$. Set $|\tau\rangle := \bigotimes_{i \in [t], c \in \{0, 1\}} |x_c^i\rangle |y_c^i\rangle |\text{ct}_{x_c^i, y_c^i}^c\rangle$.
5. If $b = 0$, apply the unitary V on $|0\|x\rangle|\phi_k(0\|x)\rangle|\tau\rangle$ and output the second register and $1\|y$. If $b = 1$, apply the unitary U on $|1\|y\rangle|\phi_k(1\|y)\rangle|\tau\rangle$ and output the second register and $0\|x$.

Since \mathcal{B} does not abort if $(k, x, y, \mathbf{w}) \in G$, the probability that the adversary \mathcal{B} wins is

$$\Pr[\mathcal{B} \text{ wins}] \geq \sum_{(k, x, y, \mathbf{w}) \in G} \frac{\Pr[k, x, y, \mathbf{w}]}{2} \left(\|\langle \phi_k(0\|x) |_{\mathbf{Y}} (V|1\|y)_{\mathbf{X}} | \phi_k(1\|y) \rangle_{\mathbf{Y}} |\tau\rangle_{\mathbf{Z}} \|^2 \right) \quad (27)$$

$$+ \|\langle \phi_k(1\|y) |_{\mathbf{Y}} (V|0\|x)_{\mathbf{X}} | \phi_k(0\|x) \rangle_{\mathbf{Y}} |\tau\rangle_{\mathbf{Z}} \|^2 \quad (28)$$

$$\geq \sum_{(k, x, y, \mathbf{w}) \in G} \frac{\Pr[k, x, y, \mathbf{w}]}{2} \max \left\{ \|\langle \phi_k(0\|x) |_{\mathbf{Y}} (V|1\|y)_{\mathbf{X}} | \phi_k(1\|y) \rangle_{\mathbf{Y}} |\tau\rangle_{\mathbf{Z}} \|^2 \right. \quad (29)$$

$$\left. \|\langle \phi_k(1\|y) |_{\mathbf{Y}} (V|0\|x)_{\mathbf{X}} | \phi_k(0\|x) \rangle_{\mathbf{Y}} |\tau\rangle_{\mathbf{Z}} \|^2 \right\} \quad (30)$$

$$\geq \sum_{(k, x, y, \mathbf{w}) \in G} \frac{\Pr[k, x, y, \mathbf{w}]}{2} \frac{1}{4p(\lambda)^2} \geq \frac{1}{32p(\lambda)^3} \quad (31)$$

for infinitely many λ , where we have used Eq. (26) in Eq. (30), and Lemma 4.1 in Eq. (31). This shows that \mathcal{B} breaks the security of the UPSG. Hence we have shown the theorem. \square

5 MACs with Unclonable Tags

In this section, we define MACs with unclonable tags and construct it from UPSGs.

5.1 Definition

First, we give the definition of the standard EUF-CMA-secure MACs. However, in this paper, we consider more general case where the tags could be quantum states. MACs with classical tags can be considered as a special case where the tags are computational-basis states.

Definition 5.1 (EUF-CMA-Secure MACs). *An EUF-CMA-secure MAC is a set $(\text{KeyGen}, \text{Tag}, \text{Ver})$ of QPT algorithms such that*

- $\text{KeyGen}(1^\lambda) \rightarrow \text{sigk}$: *It is a QPT algorithm that, on input the security parameter λ , outputs a classical key sigk .*
- $\text{Tag}(\text{sigk}, m) \rightarrow \tau$: *It is a QPT algorithm that, on input sigk and a classical message m , outputs an n -qubit quantum state τ .*
- $\text{Ver}(\text{sigk}, m, \rho) \rightarrow \top/\perp$: *It is a QPT algorithm that, on input sigk , m , and a quantum state ρ , outputs \top/\perp .*

We require the following two properties.

Correctness: For any m ,

$$\Pr \left[\top \leftarrow \text{Ver}(\text{sigk}, m, \tau) : \begin{array}{l} \text{sigk} \leftarrow \text{KeyGen}(1^\lambda) \\ \tau \leftarrow \text{Tag}(\text{sigk}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda). \quad (32)$$

EUF-CMA Security: For any QPT adversary \mathcal{A} ,

$$\Pr \left[\top \leftarrow \text{Ver}(\text{sigk}, m^*, \rho) : \begin{array}{l} \text{sigk} \leftarrow \text{KeyGen}(1^\lambda) \\ (m^*, \rho) \leftarrow \mathcal{A}^{\text{Tag}(\text{sigk}, \cdot)}(1^\lambda) \end{array} \right] \leq \text{negl}(\lambda), \quad (33)$$

where \mathcal{A} queries the oracle only classically, and \mathcal{A} is not allowed to query m^* .

The following corollary is straightforward from the definition of UPSGs.

Corollary 5.1. *If UPSGs exist, then EUF-CMA-secure MACs exist.*

Proof of Corollary 5.1. Let $(\text{KeyGen}', \text{Eval}')$ be a UPSG. We construct EUF-CMA-secure MAC $(\text{KeyGen}, \text{Tag}, \text{Ver})$ as follows:

- $\text{KeyGen}(1^\lambda) \rightarrow \text{sigk}$: Run $k \leftarrow \text{KeyGen}'(1^\lambda)$ and output it as sigk .
- $\text{Tag}(\text{sigk}, m) \rightarrow \tau$: Parse $\text{sigk} = k$. Run $|\phi_k(m)\rangle \leftarrow \text{Eval}'(k, m)$ and output it as τ .
- $\text{Ver}(\text{sigk}, m, \rho) \rightarrow \top/\perp$: Parse $\text{sigk} = k$. Project ρ onto $|\phi_k(m)\rangle\langle\phi_k(m)|$. If the projection is successful, output \top . Otherwise, output \perp .

The correctness is clear. The EUF-CMA-security follows from the unpredictability of UPSG. \square

Next, we define MACs with unclonable tags.

Definition 5.2 (MACs with Unclonable Tags). *Let $(\text{KeyGen}, \text{Tag}, \text{Ver})$ be an EUF-CMA-secure MAC. If $(\text{KeyGen}, \text{Tag}, \text{Ver})$ satisfies the following property (which we call unclonability), we call it MAC with unclonable tags: For any QPT adversary \mathcal{A} and any polynomials t and ℓ ,*

$$\Pr \left[\text{Count}(\text{sigk}, m^*, \xi) \geq t + 1 : \begin{array}{l} \text{sigk} \leftarrow \text{KeyGen}(1^\lambda) \\ (m^*, \text{st}) \leftarrow \mathcal{A}^{\text{Tag}(\text{sigk}, \cdot)}(1^\lambda) \\ \tau^{\otimes t} \leftarrow \text{Tag}(\text{sigk}, m^*)^{\otimes t} \\ \xi \leftarrow \mathcal{A}^{\text{Tag}(\text{sigk}, \cdot)}(\tau^{\otimes t}, \text{st}) \end{array} \right] \leq \text{negl}(\lambda), \quad (34)$$

where \mathcal{A} queries the oracle only classically, and \mathcal{A} is not allowed to query m^* . $\tau^{\otimes t} \leftarrow \text{Tag}(\text{sigk}, m^*)^{\otimes t}$ means that Tag algorithm is run t times and t copies of τ are generated. ξ is a quantum state on ℓ registers, $\mathbf{R}_1, \dots, \mathbf{R}_\ell$, each of which is of n qubits. Here, $\text{Count}(\text{sigk}, m^*, \xi)$ is the following QPT algorithm: for each $j \in [\ell]$, it takes the state on \mathbf{R}_j as input, and runs $\text{Ver}(\text{sigk}, m^*, \cdot)$ to get \top or \perp . Then, it outputs the total number of \top .

Remark 5.1. EUF-CMA security is automatically implied by the unclonability, Eq. (34).¹⁶

¹⁶ The proof is easy. Let \mathcal{A} be a QPT adversary that breaks the EUF-CMA security, which outputs (m^*, ρ) . Then the QPT adversary \mathcal{B} that breaks the unclonability is constructed as follows: it first simulates \mathcal{A} to get (m^*, ρ) . It then sends m^* to the challenger to get its tag τ . It finally sends τ and ρ to the challenger, both of which are accepted as valid tags.

5.2 Construction from UPSGs

In this subsection, we construct MACs with unclonable tags from EUF-CMA-secure MACs and IND-CPA-secure SKE schemes.

Theorem 5.1. *If EUF-CMA-secure MACs (secure against classically querying QPT adversaries) and IND-CPA-secure SKE schemes for classical messages (secure against classically querying QPT adversaries) exist, then MACs with unclonable tags exist.*

Because EUF-CMA-secure MACs (secure against classically querying QPT adversaries) can be constructed from UPSG (Corollary 5.1), and IND-CPA-secure SKE schemes for classical messages (secure against classically querying QPT adversaries) can be constructed from UPSGs (Theorem 4.1), we have the following corollary:

Corollary 5.2. *If UPSGs exist, then MACs with unclonable tags exist.*

Proof of Theorem 5.1. Let $(\text{MAC.KeyGen}, \text{MAC.Tag}, \text{MAC.Ver})$ be an EUF-CMA-secure MAC secure against classically querying QPT adversaries and $(\text{SKE.KeyGen}, \text{SKE.Enc}, \text{SKE.Dec})$ be an IND-CPA-secure SKE scheme for quantum messages. (From Lemma 2.3, such SKE schemes exist if SKE schemes for classical messages secure against classically querying QPT adversaries exist.) We construct a MAC with unclonable tags $(\text{KeyGen}, \text{Tag}, \text{Ver})$ as follows:

- $\text{KeyGen}(1^\lambda) \rightarrow \text{sigk}'$: Run $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$. Output $\text{sigk}' := (\text{sk}, \text{sigk})$.
- $\text{Tag}(\text{sigk}', m) \rightarrow \tau'$: Parse $\text{sigk}' = (\text{sk}, \text{sigk})$. It does the following:
 1. Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$. Here, $|x\rangle_\theta := \bigotimes_{i \in [\lambda]} H^{\theta^i} |x^i\rangle$, where H is the Hadamard gate, and x^i and θ^i denote the i 'th bit of x and θ , respectively.
 2. Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$.
 3. Run $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta \rangle \otimes \tau)$.
 Output $\tau' := |x\rangle_\theta \langle x | \theta \rangle \otimes \text{ct}$.
- $\text{Ver}(\text{sigk}', m, \rho) \rightarrow \top / \perp$: Parse $\text{sigk}' = (\text{sk}, \text{sigk})$. Let ρ be a state on two registers \mathbf{A} and \mathbf{C} . (If ρ is honestly generated, $\rho_{\mathbf{A}, \mathbf{C}} = (|x\rangle_\theta \langle x | \theta \rangle_{\mathbf{A}} \otimes \text{ct}_{\mathbf{C}})$.) It does the following:
 1. Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C} to get another state $\rho'_{\mathbf{A}, \mathbf{M}}$ on the registers \mathbf{A} and \mathbf{M} .
 2. Measure the first 2λ qubits of \mathbf{M} in the computational basis to get the result $x' \| \theta'$.
 3. Run $\text{MAC.Ver}(\text{sigk}, m \| x' \| \theta', \cdot)$ on the remaining qubits of the register \mathbf{M} to get $v \in \{\top, \perp\}$. Project the register \mathbf{A} onto $|x'\rangle_{\theta'}$. If the projection is successful and $v = \top$, output \top . Otherwise, output \perp .

The correctness is clear. Since the unclonability implies EUF-CMA security, it suffices to show our construction satisfies the unclonability. Let t and ℓ be polynomials. We define the Hybrid 0 as follows, which is the original security game of unclonability between the challenger \mathcal{C} and QPT adversary \mathcal{A} .

Hybrid 0

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:
 - (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
 - (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta \rangle \otimes \tau)$.
 - (c) Output $|x\rangle \langle x | \theta \rangle \otimes \text{ct}$.
3. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i | \theta_i \rangle \otimes \tau_i)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i | \theta_i \rangle \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the following.
 - (a) Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j .
 - (b) Measure the first 2λ qubits of the register \mathbf{M}_j in the computational basis to get the result $x'_j \| \theta'_j$.
 - (c) Run $\text{MAC.Ver}(\text{sigk}, m^* \| x'_j \| \theta'_j, \cdot)$ on the remaining qubits of the register \mathbf{M}_j to get $v_j \in \{\top, \perp\}$.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful and $v_j = \top$, set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^\ell w_j \geq t + 1$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

To show the theorem, let us assume that there exists a QPT adversary \mathcal{A} such that $\Pr[\top \leftarrow \text{Hybrid 0}] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many λ . Our goal is to construct an adversary that breaks the security of the Wiesner money scheme from \mathcal{A} . To demonstrate that, we define some hybrids.¹⁷ To construct an adversary against the Wiesner money, we want to make sure that two copies of $|x\rangle_\theta$ are generated when \mathcal{C} outputs \top . The next Hybrid 1 ensures such a situation, and the hop from Hybrid 0 to 1 can be done by invoking the EUF-CMA security of the MAC.¹⁸

Hybrid 1

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:

¹⁷ In each hybrid, text struck through with red is the step in the previous hybrid, and the red text is the new step in the current hybrid.

¹⁸ This is actually a well-known technique to construct a full money from a mini-scheme [2].

- (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
- (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta \rangle \otimes \tau)$.
- (c) Output $|x\rangle \langle x |_\theta \otimes \text{ct}$.
3. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i | \theta_i \rangle \otimes \tau_i)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i |_{\theta_i} \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j .
 - (b) Measure the first 2λ qubits of the register \mathbf{M}_j in the computational basis to get the result $x'_j \| \theta'_j$.
 - (c) Run $\text{MAC.Ver}(\text{sigk}, m^* \| x'_j \| \theta'_j, \cdot)$ on the remaining qubits of the register \mathbf{M}_j to get $v_j \in \{\top, \perp\}$.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful and $v_j = \top$, set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^\ell w_j \geq t + 1$ and the event E does not occur, then \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp . Here E is the event defined as follows:
 - Event E : there exists $j \in [\ell]$ such that $(x'_j, \theta'_j) \notin \{(x_i, \theta_i)\}_{i \in [t]}$ and $w_j = 1$.

Lemma 5.1. $\Pr[\top \leftarrow \text{Hybrid } 0] \leq \Pr[\top \leftarrow \text{Hybrid } 1] + \text{negl}(\lambda)$.

Proof of Lemma 5.1. We can show

$$\Pr[E] \leq \text{negl}(\lambda) \tag{35}$$

whose proof is given later. If $\Pr[E] \leq \text{negl}(\lambda)$,

$$\Pr[\top \leftarrow \text{Hybrid } 0] = \Pr[\top \leftarrow \text{Hybrid } 0 \wedge E] + \Pr[\top \leftarrow \text{Hybrid } 0 \wedge \bar{E}] \tag{36}$$

$$\leq \text{negl}(\lambda) + \Pr[\top \leftarrow \text{Hybrid } 1], \tag{37}$$

which shows the lemma.

Let us show Eq. (35). Assume that $\Pr[E] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many $\lambda \in \mathbb{N}$. Then the following QPT adversary \mathcal{B} breaks the EUF-CMA security of the MAC:

1. The adversary \mathcal{B} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$.
2. \mathcal{B} simulates the interaction between \mathcal{C} and \mathcal{A} in Hybrid 1 by querying to $\text{MAC.Tag}(\text{sigk}, \cdot)$ up to the step 5. Then, \mathcal{B} gets a classical message m^* and a state ξ on the registers $\mathbf{R}_1, \dots, \mathbf{R}_\ell$, where m^* is a challenge message that \mathcal{A} sends to \mathcal{C} in the step 2. Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
3. For each $j \in [\ell]$, \mathcal{B} does the following:
 - (a) Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j .

- (b) Measure the first 2λ qubits of the register \mathbf{M}_j in the computational basis to get the result $x'_j \parallel \theta'_j$.
4. \mathcal{B} chooses $j^* \leftarrow [\ell]$. If $(x'_{j^*}, \theta'_{j^*}) \in \{(x_i, \theta_i)\}_{i \in [t]}$, \mathcal{B} aborts. Otherwise, \mathcal{B} outputs $m^* \parallel x'_{j^*} \parallel \theta'_{j^*}$ and the all qubits of the register \mathbf{M}_{j^*} except for the first 2λ -qubits.

It is clear that \mathcal{B} does not query $m^* \parallel x'_{j^*} \parallel \theta'_{j^*}$. Let $\Pr[\mathcal{B} \text{ wins}]$ be the probability that \mathcal{B} wins the above security game of EUF-CMA security. Then, we have $\Pr[\mathcal{B} \text{ wins}] \geq \frac{1}{\ell} \Pr[E]$. Therefore, \mathcal{B} breaks the EUF-CMA security if $\Pr[E] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many $\lambda \in \mathbb{N}$. This means $\Pr[E] \leq \text{negl}(\lambda)$. \square

If $\Pr[\top \leftarrow \text{Hybrid 1}] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many λ , at least two copies of $|x\rangle_\theta$ for some x and θ should be generated due to the pigeonhole principle. In the following Hybrid 2, we randomly guess the indexes of such states. Then we have the following lemma.

Lemma 5.2. $\Pr[\top \leftarrow \text{Hybrid 2}] \geq \frac{1}{\ell} \Pr[\top \leftarrow \text{Hybrid 1}]$.

Hybrid 2

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:
 - (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
 - (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \parallel x \parallel \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x \parallel \theta| \otimes \tau)$.
 - (c) Output $|x\rangle \langle x|_\theta \otimes \text{ct}$.
3. \mathcal{C} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \parallel x_i \parallel \theta_i)$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i \parallel \theta_i| \otimes \tau_i)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i|_{\theta_i} \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j .
 - (b) Measure the first 2λ qubits of the register \mathbf{M}_j in the computational basis to get the result $x'_j \parallel \theta'_j$.
 - (c) Run $\text{MAC.Ver}(\text{sigk}, m^* \parallel x'_j \parallel \theta'_j, \cdot)$ on the remaining qubits of the register \mathbf{M}_j to get $v_j \in \{\top, \perp\}$.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful and $v_j = \top$ and $(x'_j, \theta'_j) = (x_{i^*}, \theta_{i^*})$, set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^{\ell} w_j \geq t + 1$ and the event E does not occur, If $\sum_{j=1}^{\ell} w_j \geq 2$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp . ~~Here E is the event defined as follows:~~

- ~~Event E :~~
~~there exists $j \in [\ell]$ such that $(x'_j, \theta'_j) \notin \{(x_i, \theta_i)\}_{i \in [t]}$ and $w_j = 1$.~~

Let us define Hybrid 3 as follows. The following lemma is straightforward.

Lemma 5.3. $\Pr[\top \leftarrow \text{Hybrid 3}] \geq \Pr[\top \leftarrow \text{Hybrid 2}]$.

Hybrid 3

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:
 - (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
 - (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta \rangle \otimes \tau)$.
 - (c) Output $|x\rangle \langle x | \otimes \text{ct}$.
3. \mathcal{C} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i | \theta_i \rangle \otimes \tau_i)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i | \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Run $\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j .
 - (b) ~~Measure the first 2λ qubits of the register \mathbf{M}_j in the computational basis to get the result $x'_j \| \theta'_j$. Set $(x'_j, \theta'_j) := (x_{i^*}, \theta_{i^*})$.~~
 - (c) Run $\text{MAC.Ver}(\text{sigk}, m^* \| x'_j \| \theta'_j, \cdot)$ on the remaining qubits of the register \mathbf{M}_j to get $v_j \in \{\top, \perp\}$.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful and $v_j = \top$ ~~and $(x'_j, \theta'_j) = (x_{i^*}, \theta_{i^*})$~~ , set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^\ell w_j \geq 2$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

Now in Hybrid 3 two copies of $|x_{i^*}\rangle_{\theta_{i^*}}$ are generated. In order to use it to break the security of the Wiesner money scheme, we have to remove the classical description of BB84 states “hidden” in the ciphertexts. If we introduce Hybrid 4 as follows, the following lemma is straightforward.

Lemma 5.4 $\Pr[\top \leftarrow \text{Hybrid 4}] \geq \Pr[\top \leftarrow \text{Hybrid 3}]$.

Hybrid 4

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:

- (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
- (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta | \otimes \tau)$.
- (c) Output $|x\rangle \langle x |_\theta \otimes \text{ct}$.
3. \mathcal{C} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i | \theta_i | \otimes \tau_i)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i |_{\theta_i} \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Run ~~$\text{SKE.Dec}(\text{sk}, \cdot)$ on the register \mathbf{C}_j to get the state on the registers \mathbf{A}_j and \mathbf{M}_j . Does nothing in this step.~~
 - (b) Set $(x'_j, \theta'_j) := (x_{i^*}, \theta_{i^*})$.
 - (c) Run ~~$\text{MAC.Ver}(\text{sigk}, m^* \| x'_j \| \theta'_j, \cdot)$ on the remaining qubits of the register \mathbf{M}_j to get $v_j \in \{\top, \perp\}$. Does nothing in this step.~~
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful and ~~$v_j = \top$~~ , set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^\ell w_j \geq 2$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

Now, we are ready to remove the information about the BB84 state from ct_i by invoking IND-CPA security. We formalize it as Hybrid 5.

Hybrid 5

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:
 - (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
 - (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m \| x \| \theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\rangle_\theta \langle x | \theta | \otimes \tau)$.
 - (c) Output $|x\rangle \langle x |_\theta \otimes \text{ct}$.
3. \mathcal{C} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.
 - (b) Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$ and ~~$\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |x_i\rangle_{\theta_i} \langle x_i | \theta_i | \otimes \tau_i)$~~ . Run $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |0\dots 0\rangle \langle 0\dots 0 |)$.
4. \mathcal{C} sends $\{|x_i\rangle \langle x_i |_{\theta_i} \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Does nothing in this step.
 - (b) Set $(x'_j, \theta'_j) := (x_{i^*}, \theta_{i^*})$.
 - (c) Does nothing in this step.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.

- (e) If the projection is successful, set $w_j := 1$. Otherwise, set $w_j := 0$.
7. If $\sum_{j=1}^{\ell} w_j \geq 2$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

Lemma 5.5 $|\Pr[\top \leftarrow \text{Hybrid 4}] - \Pr[\top \leftarrow \text{Hybrid 5}]| \leq \text{negl}(\lambda)$.

Proof of Lemma 5.5. Note that the difference between Hybrid 4 and Hybrid 5 lies only in the step 3b.

Let us consider the following security game of IND-CPA security between a challenger \mathcal{C}' and a QPT adversary \mathcal{B} :

1. The challenger \mathcal{C}' runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$.
2. \mathcal{B} runs $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
3. \mathcal{B} simulates \mathcal{A} in Hybrid 4 by querying to $\text{SKE.Enc}(\text{sk}, \cdot)$ up to the step 2. \mathcal{B} gets m^* , where m^* is the challenge message that \mathcal{A} sends to \mathcal{C} in the step 2.
4. \mathcal{B} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{B} does the following: \mathcal{B} chooses $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and prepares the state $\eta_i^0 := |x_i\|\theta_i\rangle\langle x_i\|\theta_i| \otimes \tau_i$ by running $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^*\|x_i\|\theta_i)$. \mathcal{B} also prepares the state $\eta_i^1 := |0\dots 0\rangle\langle 0\dots 0|$.
5. \mathcal{B} sends the states $\bigotimes_{i=1}^t \eta_i^0$ and $\bigotimes_{i=1}^t \eta_i^1$ to \mathcal{C}' .
6. \mathcal{C}' chooses $b \leftarrow \{0, 1\}$ and gets $\bigotimes_{i \in [t]} \text{ct}_i$ by running $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, \eta_i^0)$ if $b = 0$ and $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, \eta_i^1)$ if $b = 1$ for each $i \in [t]$. \mathcal{C}' sends $\bigotimes_{i \in [t]} \text{ct}_i$ to \mathcal{B} .
7. \mathcal{B} generates $\bigotimes_{i \in [t]} |x_i\rangle_{\theta_i}$. \mathcal{B} simulates the interaction between \mathcal{C} and \mathcal{A} from the step 4 of Hybrid 4 to the last step by using $\bigotimes_{i \in [t]} \text{ct}_i \otimes \bigotimes_{i \in [t]} |x_i\rangle_{\theta_i}$ and querying to $\text{SKE.Enc}(\text{sk}, \cdot)$. If \mathcal{C} outputs \top , \mathcal{B} sends $b' := 0$ to \mathcal{C}' . Otherwise, \mathcal{B} sends $b' := 1$ to \mathcal{C}' .
8. \mathcal{C}' outputs \top if $b = b'$. Otherwise, \mathcal{C}' outputs \perp .

Let $\Pr[b' \leftarrow \mathcal{B}|b \leftarrow \mathcal{C}']$ be the probability that \mathcal{B} sends $b' \in \{0, 1\}$ to \mathcal{C}' when \mathcal{C}' chooses $b \in \{0, 1\}$. It is clear that $\Pr[0 \leftarrow \mathcal{B}|0 \leftarrow \mathcal{C}'] = \Pr[\top \leftarrow \text{Hybrid 4}]$ and $\Pr[0 \leftarrow \mathcal{B}|1 \leftarrow \mathcal{C}'] = \Pr[\top \leftarrow \text{Hybrid 5}]$. Therefore, if $|\Pr[\top \leftarrow \text{Hybrid 4}] - \Pr[\top \leftarrow \text{Hybrid 5}]| \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many $\lambda \in \mathbb{N}$, \mathcal{B} breaks the IND-CPA security. \square

Let us define Hybrid 6 as follows. The following lemma is straightforward.

Lemma 5.6. $\Pr[\top \leftarrow \text{Hybrid 5}] = \Pr[\top \leftarrow \text{Hybrid 6}]$.

Hybrid 6

1. The challenger \mathcal{C} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
2. The adversary \mathcal{A} sends m^* to \mathcal{C} , where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here, $\mathcal{O}_{\text{sk}, \text{sigk}}$ takes a bit string m as input and works as follows:
 - (a) Choose $x, \theta \leftarrow \{0, 1\}^\lambda$ and generate $|x\rangle_\theta$.
 - (b) Run $\tau \leftarrow \text{MAC.Tag}(\text{sigk}, m\|x\|\theta)$ and $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, |x\|\theta\rangle\langle x\|\theta| \otimes \tau)$.
 - (c) Output $|x\rangle_\theta\langle x\|\theta| \otimes \text{ct}$.
3. \mathcal{C} chooses $i^* \leftarrow [t]$. For each $i \in [t]$, \mathcal{C} does the following.
 - (a) Choose $x_i, \theta_i \leftarrow \{0, 1\}^\lambda$ and generate $|x_i\rangle_{\theta_i}$.

- (b) ~~Run $\tau_i \leftarrow \text{MAC.Tag}(\text{sigk}, m^* \| x_i \| \theta_i)$. Run $\text{ct}_i \leftarrow \text{SKE.Enc}(\text{sk}, |0\dots 0\rangle\langle 0\dots 0|)$.~~
4. \mathcal{C} sends $\{|x_i\rangle\langle x_i|_{\theta_i} \otimes \text{ct}_i\}_{i \in [t]}$ to \mathcal{A} .
 5. \mathcal{A} sends $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$, where \mathcal{A} can make classical queries to the oracle $\mathcal{O}_{\text{sk}, \text{sigk}}$ and does not query m^* . Here \mathbf{R}_j has two registers \mathbf{A}_j and \mathbf{C}_j for each $j \in [\ell]$.
 6. For each $j \in [\ell]$, \mathcal{C} does the followings.
 - (a) Does nothing in this step.
 - (b) Set $(x'_j, \theta'_j) := (x_{i^*}, \theta_{i^*})$.
 - (c) Does nothing in this step.
 - (d) Project the register \mathbf{A}_j onto $|x'_j\rangle_{\theta'_j}$.
 - (e) If the projection is successful, set $w_j := 1$. Otherwise, set $w_j := 0$.
 7. If $\sum_{j=1}^{\ell} w_j \geq 2$, \mathcal{C} outputs \top . Otherwise, \mathcal{C} outputs \perp .

Finally, we construct an adversary that breaks the security of the Wiesner money scheme from \mathcal{A} of Hybrid 6, which concludes our proof of the theorem.

Lemma 5.7. $\Pr[\top \leftarrow \text{Hybrid 6}] \leq \text{negl}(\lambda)$.

Proof of Lemma 5.7. Let us assume that there exist polynomials t, ℓ and a QPT \mathcal{A} adversary such that $\Pr[\top \leftarrow \text{Hybrid 6}] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many $\lambda \in \mathbb{N}$. From this \mathcal{A} , we can construct a QPT adversary \mathcal{B} that breaks the security of the Wiesner money scheme as follows:

1. The challenger \mathcal{C}' chooses $x, \theta \leftarrow \{0, 1\}^\lambda$ and sends $|x\rangle_\theta$ to \mathcal{B} .
2. \mathcal{B} runs $\text{sk} \leftarrow \text{SKE.KeyGen}(1^\lambda)$ and $\text{sigk} \leftarrow \text{MAC.KeyGen}(1^\lambda)$.
3. \mathcal{B} simulates the interaction between the challenger and \mathcal{A} in Hybrid 6, where, in the step 3, \mathcal{B} chooses $i^* \leftarrow [t]$ and replaces $|x_{i^*}\rangle_{\theta_{i^*}}$ with $|x\rangle_\theta$. Then, \mathcal{B} gets $\xi_{\mathbf{R}_1, \dots, \mathbf{R}_\ell}$ from \mathcal{A} . \mathcal{B} chooses $j_0, j_1 \leftarrow [\ell]$ and outputs the register \mathbf{A}_{j_0} and \mathbf{A}_{j_1} .

The probability that \mathcal{B} wins is

$$\Pr[\mathcal{B} \text{ wins}] \geq \binom{\ell}{2}^{-1} \Pr[\top \leftarrow \text{Hybrid 6}] \geq \frac{2}{\ell(\ell-1)} \frac{1}{\text{poly}(\lambda)}. \quad (38)$$

However, this contradicts the security of the Wiesner money scheme, Lemma 2.2. Therefore, $\Pr[\top \leftarrow \text{Hybrid 5}] \leq \text{negl}(\lambda)$. \square

By combining Lemmata 5.1 to 5.7, we have $\Pr[\top \leftarrow \text{Hybrid 0}] \leq \text{negl}(\lambda)$, but it contradicts the assumption that $\Pr[\top \leftarrow \text{Hybrid 0}] \geq \frac{1}{\text{poly}(\lambda)}$ for infinitely many λ . Therefore we have $\Pr[\top \leftarrow \text{Hybrid 0}] \leq \text{negl}(\lambda)$. \square

Acknowledgments. TM is supported by JST CREST JPMJCR23I3, JST Moonshot R&D JPMJMS2061-5-1-1, JST FOREST, MEXT QLEAP, the Grant-in Aid for Transformative Research Areas (A) 21H05183, and the Grant-in-Aid for Scientific Research (A) No. 22H00522.

References

1. Aaronson, S., Atia, Y., Susskind, L.: On the hardness of detecting macroscopic superpositions. *Electron. Colloquium Comput. Complex.* p. 146 (2020)
2. Aaronson, S., Christiano, P.: Quantum money from hidden subspaces. In: Karloff, H.J., Pitassi, T. (eds.) 44th ACM STOC. pp. 41–60. ACM Press (May 2012). <https://doi.org/10.1145/2213977.2213983>
3. Alagic, G., Broadbent, A., Fefferman, B., Gagliardoni, T., Schaffner, C., St. Jules, M.: Computational security of quantum encryption. In: *Information Theoretic Security: 9th International Conference, ICITS 2016, Tacoma, WA, USA, August 9–12, 2016, Revised Selected Papers 9*. pp. 47–71. Springer (2016)
4. Ananth, P., Gulati, A., Kaleoglu, F., Lin, Y.T.: Pseudorandom isometries. arXiv preprint [arXiv:2311.02901](https://arxiv.org/abs/2311.02901) (2023)
5. Ananth, P., Gulati, A., Qian, L., Yuen, H.: Pseudorandom (function-like) quantum state generators: New definitions and applications. In: Kiltz, E., Vaikuntanathan, V. (eds.) TCC 2022, Part I. LNCS, vol. 13747, pp. 237–265. Springer, Heidelberg (Nov 2022). https://doi.org/10.1007/978-3-031-22318-1_9
6. Ananth, P., Lin, Y.T., Yuen, H.: Pseudorandom strings from pseudorandom quantum states. *Cryptology ePrint Archive*, Paper 2023/904 (2023), <https://eprint.iacr.org/2023/904>, <https://eprint.iacr.org/2023/904>
7. Ananth, P., Qian, L., Yuen, H.: Cryptography from pseudorandom quantum states. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part I. LNCS, vol. 13507, pp. 208–236. Springer, Heidelberg (Aug 2022). https://doi.org/10.1007/978-3-031-15802-5_8
8. Behera, A., Brakerski, Z., Sattath, O., Shmueli, O.: Pseudorandomness with proof of destruction and applications. *Cryptology ePrint Archive*, Paper 2023/543 (2023), <https://eprint.iacr.org/2023/543>, <https://eprint.iacr.org/2023/543>
9. Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 361–379. Springer, Heidelberg (Aug 2013). https://doi.org/10.1007/978-3-642-40084-1_21
10. Brakerski, Z., Canetti, R., Qian, L.: On the computational hardness needed for quantum cryptography. *ITCS 2023: 14th Innovations in Theoretical Computer Science* (2023)
11. Broadbent, A., Jeffery, S.: Quantum homomorphic encryption for circuits of low t-gate complexity. In: *Annual Cryptology Conference*. pp. 609–629. Springer (2015)
12. Cavalari, B., Goldin, E., Gray, M., Hall, P., Liu, Y., Pelecinos, A.: On the computational hardness of quantum one-wayness. arXiv preprint [arXiv:2312.08363](https://arxiv.org/abs/2312.08363) (2023)
13. Coladangelo, A., Mutreja, S.: On black-box separations of quantum digital signatures from pseudorandom states. arXiv preprint [arXiv:2402.08194](https://arxiv.org/abs/2402.08194) (2024)
14. Goldreich, O.: A note on computational indistinguishability. *Information Processing Letters* 34.6 (1990), pp.277–281. (1990). [https://doi.org/10.1016/0020-0190\(90\)90010-U](https://doi.org/10.1016/0020-0190(90)90010-U)
15. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. *Journal of the ACM* **33**(4), 792–807 (1986)
16. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: 21st ACM STOC. pp. 25–32. ACM Press (May 1989). <https://doi.org/10.1145/73007.73010>
17. Hhan, M., Morimae, T., Yamakawa, T.: From the hardness of detecting superpositions to cryptography: Quantum public key encryption and commitments. In:

- Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part I. LNCS, vol. 14004, pp. 639–667. Springer, Heidelberg (Apr 2023). https://doi.org/10.1007/978-3-031-30545-0_22
18. Ji, Z., Liu, Y.K., Song, F.: Pseudorandom quantum states. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 126–152. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96878-0_5
 19. Katz, J., Lindell, Y.: Introduction to Modern Cryptography. Chapman and Hall/CRC Press (2007)
 20. Khurana, D., Tomer, K.: Commitments from quantum one-wayness. Cryptology ePrint Archive, Paper 2023/1620 (2023), <https://eprint.iacr.org/2023/1620>, <https://eprint.iacr.org/2023/1620>
 21. Kitagawa, F., Morimae, T., Nishimaki, R., Yamakawa, T.: Quantum public-key encryption with tamper-resilient public keys from one-way functions. arXiv preprint [arXiv:2304.01800](https://arxiv.org/abs/2304.01800) (2023)
 22. Kretschmer, W.: Quantum pseudorandomness and classical complexity. TQC 2021 (2021). <https://doi.org/10.4230/LIPICS.TQC.2021.2>
 23. Kretschmer, W., Qian, L., Sinha, M., Tal, A.: Quantum cryptography in algorithmica. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing. pp. 1589–1602 (2023)
 24. Lu, C., Qin, M., Song, F., Yao, P., Zhao, M.: Quantum pseudorandom scramblers. arXiv preprint [arXiv:2309.08941](https://arxiv.org/abs/2309.08941) (2023)
 25. Metger, T., Poremba, A., Sinha, M., Yuen, H.: Simple constructions of linear-depth t-designs and pseudorandom unitaries. arXiv preprint [arXiv:2404.12647](https://arxiv.org/abs/2404.12647) (2024)
 26. Molina, A., Vidick, T., Watrous, J.: Optimal counterfeiting attacks and generalizations for wiesner’s quantum money. In: Conference on Quantum Computation, Communication, and Cryptography. pp. 45–64. Springer (2012)
 27. Morimae, T., Yamakawa, T.: One-wayness in quantum cryptography. Cryptology ePrint Archive, Paper 2022/1336 (2022), <https://eprint.iacr.org/2022/1336>, <https://eprint.iacr.org/2022/1336>
 28. Morimae, T., Yamakawa, T.: Quantum commitments and signatures without one-way functions. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part I. LNCS, vol. 13507, pp. 269–295. Springer, Heidelberg (Aug 2022). https://doi.org/10.1007/978-3-031-15802-5_10
 29. Naor, M., Reingold, O.: From unpredictability to indistinguishability: A simple construction of pseudo-random functions from MACs (extended abstract). In: Krawczyk, H. (ed.) CRYPTO’98. LNCS, vol. 1462, pp. 267–282. Springer, Heidelberg (Aug 1998). <https://doi.org/10.1007/BFb0055734>
 30. Watrous, J.: The theory of quantum information. Cambridge university press (2018)
 31. Wiesner, S.: Conjugate coding. SIGACT News **15**(1), 78–88 (1983)
 32. Yan, J.: General properties of quantum bit commitments (extended abstract). In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part IV. LNCS, vol. 13794, pp. 628–657. Springer, Heidelberg (Dec 2022). https://doi.org/10.1007/978-3-031-22972-5_22



Quantum Money from Class Group Actions on Elliptic Curves

Hart Montgomery¹ and Shahed Sharif²

¹ Linux Foundation, San Francisco, USA
hart.montgomery@gmail.com

² CSU San Marcos, San Marcos, USA

Abstract. We construct a quantum money/quantum lightning scheme from class group actions on elliptic curves over \mathbb{F}_p . Our scheme, which is based on the invariant money construction of Liu-Montgomery-Zhandry (Eurocrypt'23), is simple to describe. We believe it to be the most instantiable and well-defined quantum money construction known so far. The security of our quantum lightning construction is exactly equivalent to the (conjectured) hardness of constructing two uniform superpositions over elliptic curves in an isogeny class which is acted on simply transitively by an exponentially large ideal class group.

However, we needed to advance the state of the art of isogenies in order to achieve our scheme. In particular, we show:

- An efficient (quantum) algorithm for sampling a uniform superposition over a cryptographically large isogeny class.
- A method for specifying polynomially many generators for the class group so that polynomial-sized products yield an exponential-sized subset of class group, modulo a seemingly very modest assumption.

Achieving these results also requires us to advance the state of the art of the (pure) mathematics of elliptic curves, and we are optimistic that the mathematical tools we developed in this paper can be used to advance isogeny-based cryptography in other ways.

Keywords: Elliptic Curve Isogenies · Quantum Money

1 Introduction

Quantum money is a way of implementing digital money where “banknotes” that represent the money are quantum states. The idea for quantum money was first sketched out by Wiesner [Wie83], and since then quantum money has captivated the quantum computing research community. In this work, we focus on *publicly verifiable* quantum money [Aar09], which means that any observer without privileged information can verify the correctness of the banknotes, and quantum lightning [Zha19], which guarantees that even the mint cannot cheat by producing duplicate banknotes.

Unfortunately, constructing publicly verifiable quantum money has proven to be rather elusive. Farhi, Gosset, Hassidim, Lutomirski, Nagaj, and Shor showed

that, even with some natural modifications, Wiesner’s quantum money scheme cannot be used to directly build a publicly verifiable scheme [FGH+10]. The first candidates for truly publicly verifiable quantum money were given by Aaronson [Aar09] and Aaronson and Christiano [AC12], and gave publicly verifiable quantum money constructions relative to quantum and classical oracles, respectively. Unfortunately, the proposed instantiations of oracles in both constructions were later broken [LAF+10] [CPDDF+19], casting doubt on the possibility that such oracles could be securely implemented in the real world. Zhandry’s concrete construction of quantum lightning [Zha19] was also broken by Roberts [Rob21]. More recently, the lattice-based construction of Khesin, Lu, and Shor [KLS22] was broken by Liu, Montgomery, and Zhandry [LMZ23].

On the other hand, there are a handful of candidates have been proposed that have not been broken, including constructions from knots [FGH+12] and quaternion algebras [Kan18, KSS21]. In addition, Zhandry [Zha19], as suggested by [BDS16], showed how to build publicly verifiable quantum money from quantum-secure indistinguishability obfuscation (iO). Unfortunately, none of these assumptions have received much cryptanalytic attention at all, and all known candidates of post-quantum iO [GGH15, BGMZ18, BDGM20, WW21] do not have strong connections to well-studied cryptographic assumptions.

Liu, Montgomery, and Zhandry recently showed a generic, oracle construction for quantum money that they called *invariant money* [LMZ23]. They showed a number of possible constructions, but all of these were either oracle constructions or schemes that were not known to be efficiently instantiable. In particular, they showed an uninstantiable construction from class group actions on elliptic curves. In particular, they mention how invariant money could potentially be instantiated from isogenies. However, they also comment, “We do not know if it is even possible to instantiate such a scheme, as it would likely require new ideas in isogeny-based cryptography.”

Very recently, Zhandry [Zha24] showed how to build a construction of quantum money from regular group actions that was loosely based on the generic construction of [LMZ23]. While Zhandry’s construction was the first instantiable construction from group actions, it unfortunately resorts to nonstandard assumptions that seem significantly stronger than the assumptions used in the invariant money construction would be for an instantiable construction.

1.1 Our Contributions

In this work, we prove new results in isogeny-based cryptography necessary to build quantum money from the invariant money framework of [LMZ23]. We use these results to build a new quantum money scheme based on class group actions on elliptic curves, sharing some features with a proposal in [LMZ23]. Rather than come up with a new verification algorithm like that of Zhandry [Zha24], we advance the state of the art in the mathematics of elliptic curves, which enables us to build a very simple construction with simple (although still non-standard) assumptions.

Our quantum money scheme is conceptually very simple (modulo the new mathematical techniques required for its construction) and has an extremely easy to state security assumption. For instance, for quantum lightning, the scheme is exactly as secure as the problem of sampling two uniform superpositions over all elliptic curves over \mathbb{F}_p with N points, for some N corresponding to elliptic curves with large class group. We also show that this new assumption is hard assuming that it is hard to compute isogenies between random isogenous elliptic curves (which is a standard assumption in cryptography at this point) and a quantum knowledge assumption similar to the one proposed recently by Zhandry [Zha24].

Our quantum money/lightning scheme is the first fully instantiable construction with a simple, offline (albeit quantum) security assumption. We also believe it is the simplest scheme proposed so far, and the one that rests on assumptions that are closest to traditional cryptographic security assumptions.

To enable such a simple scheme with a very nice security assumption, we advance the state of the art in elliptic curve mathematics and algorithms. In particular, for the minting algorithm to be efficient, it is necessary that a random elliptic curve mod p has non-negligible probability of having large associated class group. Previous work has either relied on heuristic assumptions, or worked exclusively with supersingular elliptic curves. In our work, we give new explicit lower bounds on the number of elliptic curves whose endomorphism ring has large discriminant, which together with a result of Tatzawa show that at least 14% of elliptic curves mod p have endomorphism ring with exponentially large class group. We also, to our knowledge, for the first time apply a heuristic model due to Erdős-Rényi on encodings of class group elements.

Heuristics on sizes of class groups and efficiency of encodings for class groups are ubiquitous in isogeny-based cryptography. Consequently our formalization of these heuristics should find other uses in cryptography and mathematics outside of quantum money.

1.2 Other Related Work

Quantum money has been a key primitive in quantum computing, especially as quantum money and quantum lightning are closely tied to numerous other areas in quantum computing. For instance, the first message in the quantum key distribution protocol of Bennet and Brassard [BB87] is just a banknote in Wiesner’s quantum money scheme. Recent works on copy protection [ALL+21, CLLZ21] [LLQZ22] require at a minimum a computational assumption that implies quantum money.¹

In the isogeny realm, our work makes precise estimates used for parameter selection in CRS-style cryptosystems. Previous work (for example [Cou06, RS06, DKS18]) relies completely on heuristics that imply a randomly chosen elliptic curve mod p will have large associated class group. Those heuristics are made

¹ This holds true even for certain weaker versions such as copy *detection*, also known as infinite term secure software leasing.

precise here. We also show that, under a very plausible heuristic assumption, with overwhelming probability every element of the class group has a compact encoding.

1.3 Outline

The rest of this paper proceeds as follows. In Sect. 2, we provide a brief technical overview of our construction and new isogeny results. We hope this enables the reader to understand our ideas at a high level. Then, in Sect. 3, we provide preliminary material. In Sect. 4, we build the mathematical tools we need for our constructions. We note this section utilizes math that is outside the scope of knowledge of most cryptographers, but we do our best to make it as accessible as possible. We then present our full algorithm for sampling a superposition of elliptic curves in Sect. 5 and our algorithm for verifying these superpositions in Sect. 6. We finally present our full quantum money construction in Sect. 7.

Unfortunately, due to space constraints, we must defer some content to the full version of the paper. In particular, we defer the full proof of our superposition verification algorithm and proofs of security to the full version, available on eprint.

2 Technical Overview

In this section we explain our main results and contributions at a high level. We begin by giving an overview of the new mathematical tools that are needed for its construction. We then outline our new quantum money scheme.

2.1 New Techniques and Facts on Elliptic Curves

In [LMZ23], it is stated that “generating superpositions over X , where X is the set of all elliptic curves with some (even polynomially likely) property seems difficult. In fact, we do not even know how to generate a uniform superposition over all elliptic curves efficiently.” Our minting algorithm solves this problem by first efficiently generating a uniform superposition over all elliptic curves, by encoding elliptic curves as pairs (j, b) , where j is the j -invariant and b is *twisting data*. We then show how to efficiently generate the superposition over a random exponentially large isogeny class. Note that we believe generating a superposition over a *specified* exponentially large isogeny class is difficult, and indeed the security of our scheme depends on this being the case.

Passing from the uniform superposition over all elliptic curves over a particular finite field to a superposition over a large isogeny class is accomplished by an explicit estimate on sizes of class groups associated to elliptic curves. The previous literature either restricts to isogeny classes of supersingular elliptic curves (which constitute a negligible fraction of all elliptic curves mod p), or relies on heuristics stating that the class group associated to a random elliptic curve mod p has size $O(\sqrt{p})$. We make these heuristics precise by focusing on

elliptic curves whose Frobenius discriminant (that is, the discriminant of the characteristic polynomial of Frobenius) is both square-free and $\geq 3p$. We show that for $p > 2^{63}$, at least 14% of elliptic curves mod p have such a Frobenius discriminant. We then use Tatzuza's effective version of Siegel's Theorem [Tat51] to obtain the following:

Corollary 4.14. *Let $p > 2^{63}$ be prime. Given an ordinary elliptic curve E/\mathbb{F}_p , let \mathcal{O} be $\text{End}(E)$. If an elliptic curve E is drawn from either the distribution \mathcal{D}_p or the distribution \mathcal{U}_p , then with probability at least 14%, the class group of \mathcal{O} and the isogeny class of E both have size at least*

$$0.089 \frac{\sqrt{p}}{\log p},$$

Here, \mathcal{U}_p is the uniform distribution on elliptic curves, and \mathcal{D}_p is a related distribution defined in Sect. 3.3.

In addition to [Tat51], the proof of Corollary 4.14 relies on statistical analysis of the trace of Frobenius due to Murty-Prabhu [MP19b], and statistical analysis of the values of $4p - t^2$, $1 \leq t < \sqrt{p}$, using methods of Friedlander-Iwaniec [FI10].

Finally, we give a new treatment of computations in class groups. Typically, one cannot directly compute ideals coming from random ideal classes, since these ideals can have exponentially large generators. Instead, one attempts to express ideal classes as products of prime ideals with small norm. By treating small norm prime ideals as random elements in the class group, we use results of Erdős-Rényi to show that with all but negligible probability, every ideal class can be represented by a product of distinct prime ideals of small norm; see Sect. 4.4 and Sect. 4.5.

2.2 Quantum Money from Class Group Actions on Elliptic Curves

As we have alluded before, our scheme generally falls into the invariant framework of [LMZ23]. However, other than the obvious choice of using isogeny classes (determined by the number of points on elliptic curves ($\#E(\mathbb{F}_p)$)) as the invariant, essentially every other choice we make deviates from the suggestions of [LMZ23] for implementing invariant money using class group actions on elliptic curves, and, as we have previously mentioned, our scheme accomplishes or circumvents some tasks they find difficult or impossible.

Recall that, informally speaking, a (public key) quantum money scheme is a tuple of algorithms (Gen, Ver) where Gen creates a quantum money state $|\psi\rangle$ which we refer to as a banknote and a serial number σ , and Ver takes as input a banknote $|\psi\rangle$ and a serial number σ and outputs 0 or 1, depending on whether or not the quantum money state is valid.

We say that a money scheme satisfies *quantum money unforgeability* if, given a random valid banknote and serial number pair $(|\psi\rangle, \sigma)$ it is hard to generate *two* banknotes $|\psi'\rangle, |\psi''\rangle$ that both verify with serial number σ . A scheme constitutes secure quantum lightning if it is hard for an adversary to find two states $|\psi'\rangle, |\psi''\rangle$ that both verify for *any* serial number σ . We formally define quantum money

and quantum lightning in Sect. 3 and, for familiar readers, note that we use the “mini-scheme” definition from [AC12].

Our Construction. For an elliptic curve E , let t be the trace of Frobenius acting on E , and define the *Frobenius discriminant* to be $\Delta_{\text{Fr}}(E) = 4p - t^2$. (This is the negative of the usual definition.) Note that Δ in the elliptic curve literature designates the discriminant of E itself; that is, if E is given by $y^2 = x^3 + ax + b$, the discriminant is $4a^3 + 27b^2$. We use Δ_{Fr} exclusively for the Frobenius discriminant.

Suppose we let $\mathcal{E}_{\text{sf},3p}$ represent the set of isomorphism classes of elliptic curves E over \mathbb{F}_p for prime p with $\Delta_{\text{Fr}}(E) \geq 3p$ and square-free, and let \mathcal{I}_N denote the set of elliptic curves over \mathbb{F}_p with N points. Recall that two elliptic curves over \mathbb{F}_p are isogenous if and only if they have the same number of points.

Our construction at a high level works as follows:

Money States: a valid quantum money state is a uniform superposition over a single isogeny class in $\mathcal{E}_{\text{sf},3p}$.

Gen: To generate a money state, at a high level we do the following:

1. Sample a uniform superposition of elliptic curves in $\mathcal{E}_{\text{sf},3p}$, getting a state $|\psi\rangle = \sum |E\rangle$.
2. In superposition, compute, in an adjacent register, the number of points in $|E\rangle$ using Schoof’s algorithm.
3. Measure this new adjacent register and denote its value as N .
4. Output the tuple $(|\psi\rangle, N)$ as the money state. Note that $|\psi\rangle$ will have been altered by the measurement.

We note that sampling such a superposition as we do in step (1) was previously unknown and listed as an open problem in [LMZ23], and it requires new results in number theory to solve; namely, we require precise lower bounds on the proportion of elliptic curves with $\Delta_{\text{Fr}}(E)$ large and square-free. This follows from our mathematical work that we explained earlier in this overview.

Ver: Let $|\mathcal{I}_N\rangle := \frac{1}{\sqrt{\#\mathcal{I}_N}} \sum_{E \in \mathcal{I}_N} |E\rangle$. In other words, $|\mathcal{I}_N\rangle$ is the state corresponding to a uniform superposition over all elliptic curves with N points. Our goal, similar to [LMZ23], is to compute an approximation of the projection-valued measure $V_N = |\mathcal{I}_N\rangle\langle\mathcal{I}_N|$.

To do this, we take a similar approach as both [LMZ23] and [FGH+12]. We simulate a (invertible) random walk in superposition over the isogeny class group and continually check to see if the state has changed using projection-valued measures. Intuitively, correct money states will not change when we compute invertible maps on the state since these maps will map uniform superpositions to uniform superpositions. On the other hand, most incorrect money states will have changed substantially by such a walk. For instance, a classical state consisting of a single elliptic curve will likely be completely different after a random walk and thus fail verification.

While our verification algorithm closely resembles that of the invariant money in [LMZ23] at a high level, we emphasize that we need substantially new techniques to make it work. Most notably, to our knowledge the fact that a random elliptic curve has large associated class group with non-negligible probability has never been formally shown.

Security. The security of our construction is based on a simple assumption:

Definition 2.1. *The Elliptic Curve Superposition Collision Problem ($ECSCP_p$):* The prime p is fixed. Create two (possibly entangled) quantum states that are each negligibly far from the superposition of all elliptic curves over \mathbb{F}_p in some isogeny class with Frobenius discriminant $\Delta_{\text{Fr}} \geq 3p$ and square-free.

We justify below why mathematicians believe this to be a hard problem. An astute reader will note that this security assumption is tied very closely to our scheme itself, and a reduction is simple. This is true—see the full version of our paper for the formal reduction—but we believe this to be a positive facet of our construction, in the same vein as the fact that the security of ElGamal encryption trivially following from the DDH assumption is a positive of that scheme, too.

[LMZ23] and [Zha24] argue that their constructions are secure using quantum assumptions of knowledge, and we can actually argue that not just our construction but the security assumption of our construction that we informally mentioned above is secure using assumptions of knowledge. As in [LMZ23] and [Zha24], we use two problems to prove the security of the $ECSCP_p$. First, we assume that it is hard to compute isogenies between random isogenous elliptic curves (the group action discrete log problem [ADMP20] over isogeny class groups). Second, we make a quantum assumption of knowledge: we assume that if there exists an adversary that breaks our main assumption (generating two superpositions) with non-negligible probability, we assume that there also exists some adversary that breaks our main assumption with similar probability from whose state “paths” (isogenies) between elliptic curves that it uses can be extracted. This is a relatively new type of assumption and we defer to [Zha24] for an extensive discussion on this sort of quantum knowledge assumption.

Construction Rationale. An adversary can efficiently fabricate money states if they can solve the group action discrete log problem for ideal classes acting on elliptic curves; fortunately, this problem is believed to be hard. But the hardness requires that the ideal class group is large. By Tatuzawa’s estimate [Tat51], with overwhelming probability it is sufficient that $\Delta_{\text{Fr}}(E)$ be large for any elliptic curve E in the isogeny class. Thus we choose only isogeny classes for which $\Delta_{\text{Fr}} \geq 3p$. (Note that as E varies over elliptic curves mod p , by the Hasse-Weil bounds $\Delta_{\text{Fr}}(E)$ varies between 0 and $4p$.) Finally, it may happen that the isogeny class is a disjoint union of sets, each acted upon by different a class group corresponding to a different choice of endomorphism ring inside a given imaginary quadratic field (for instance, elliptic curves with endomorphism ring isomorphic to $\mathbb{Z}[\sqrt{-3}]$ along with elliptic curves with endomorphism ring

isomorphic to $\mathbb{Z}[\frac{1+\sqrt{-3}}{2}]$). We introduce a square-freeness condition to guarantee that the entire isogeny class forms a single homogeneous space.

Comparison to Zhandry’s Scheme [Zha24]. Zhandry’s recent quantum money construction is also loosely based on the invariant money construction in [LMZ23]. However, his construction turns out to be very different from ours. Zhandry makes a very nice observation that the Fourier domain can be useful for creating verifiable quantum states over group actions and builds a quantum lightning scheme based on this observation. On the other hand, our construction is more closely tied in nature to the invariant money scheme itself, although we need to solve or formalize several open problems in the mathematics of elliptic curves in order to make the scheme work.

We also note that Zhandry’s assumptions are more complicated and seemingly stronger than ours: he proves his scheme secure in two different ways: one way uses what he calls the “D2X” assumption, which is an interactive assumption that requires a quantum oracle, and the other uses a quantum knowledge assumption. On the other hand, we can prove our scheme secure using a simple, noninteractive (albeit quantum) assumption, and we can prove this assumption secure using a simpler knowledge assumption as compared to Zhandry.

3 Definitions

In this section we provide basic definitions. For our quantum notations and definitions, we mostly borrow from and mimic [LMZ23].

Basic Cryptographic Notation. When we say a function is *negligible*, we mean that it is (asymptotically) smaller than $\frac{1}{f(\lambda)}$ for any polynomial f and security parameter λ . When we say a function is *non-negligible*, we mean that there is some polynomial function f for which the function grows (asymptotically) faster than $\frac{1}{f(\lambda)}$ for security parameter λ .

3.1 Quantum Specifics

We attempt to avoid any complicated quantum specifics. For general background and notation on quantum computing, we highly recommend [NC10].

Notation. Following [LMZ23], for quantum notation, we denote $|\cdot\rangle$ as the notation for a pure state and $|\cdot\rangle\langle\cdot|$ for its density matrix. ρ denotes a general mixed state. We let “ \dagger ” denote the conjugate transpose.

Definition 3.1. *A projection-valued measure on a Hilbert space \mathcal{H} is a set of outcomes $i \in M$ and, for each outcome, a positive semi-definite matrices \mathbf{P}_i , such that:*

1. $\sum_M \mathbf{P}_i = \mathbf{I}$,
2. each \mathbf{P}_i is Hermitian,

3. $\mathbf{P}_i^2 = \mathbf{P}_i$ and $\mathbf{P}_i = \mathbf{P}_i^\dagger$, and
4. for $i \neq j \in M$, $\mathbf{P}_i \mathbf{P}_j = 0$.

We say that the probability of obtaining output i on a pure state $|\psi\rangle$ is just $\langle\psi|\mathbf{P}_i|\psi\rangle$, with the measured state collapsing to $\frac{\mathbf{P}_i|\psi\rangle}{\sqrt{\langle\psi|\mathbf{P}_i|\psi\rangle}}$; the analogous result for a mixed state is defined in the natural way using the trace.

For simplicity, we will abuse notation and refer to a single PVM operator \mathbf{P}_i as a measurement, and the complementary operator $\mathbf{I} - \mathbf{P}_i$ will be implicit. In this case, a “successful” measurement on a state $|\psi\rangle$ will result in the output $\frac{\mathbf{P}_i|\psi\rangle}{\sqrt{\langle\psi|\mathbf{P}_i|\psi\rangle}}$; we will typically be less concerned about the output of “failed” measurements.

3.2 Quantum Money and Quantum Lightning

Here, we define public key quantum money and quantum lightning. We use the definitions of [LMZ23] verbatim. As they do, following Aaronson and Christiano [AC12], we will only consider so-called “mini-schemes”, where there is only a single banknote.

Both quantum money and quantum lightning share the same syntax and correctness requirements. There are two quantum polynomial-time algorithms Gen, Ver such that:

- $\text{Gen}(1^\lambda)$ samples a classical serial number σ and a quantum state $|\psi\rangle$.
- $\text{Ver}(\sigma, |\psi\rangle)$ outputs a bit 0 or 1, and, if the output bit is 1, also outputs a state $|\psi'\rangle$.

Definition 3.2. *We say that our quantum money scheme is correct if there exists a negligible function negl such that, for any polynomially sized integer i , we have $\Pr[\text{Ver}^i(\text{Gen}(1^\lambda))] \geq 1 - \text{negl}(\lambda)$. In other words, a correctly generated money state can be verified any polynomial number of times and verification will still pass.*

Where public key quantum money and quantum lightning differ is in security. The differences are analogous to the differences between one-way functions and collision resistance.

Definition 3.3 (Quantum Money Unforgeability). *(Gen, Ver) is secure public key quantum money if, for all quantum polynomial-time A , there exists a negligible negl such that A wins the following game with probability at most negl :*

- The challenger runs $(\sigma, |\psi\rangle) \leftarrow \text{Gen}(1^\lambda)$, and gives $\sigma, |\psi\rangle$ to A .
- A produces a potentially entangled joint state $\rho_{1,2}$ over two quantum registers. Let ρ_1, ρ_2 be the states of the two registers. A sends $\rho_{1,2}$ to the challenger.
- The challenger runs $b_1 \leftarrow \text{Ver}(\sigma, \rho_1)$ and $b_2 \leftarrow \text{Ver}(\sigma, \rho_2)$. A wins if $b_1 = b_2 = 1$.

Definition 3.4 (Quantum Lightning Unforgeability). (Gen, Ver) is secure quantum lightning if, for all quantum polynomial-time A , there exists a negligible negl such that A wins the following game with probability at most negl :

- A , on input 1^λ , produces and sends to the challenger σ and $\rho_{1,2}$, where $\rho_{1,2}$ is a potentially entangled joint state over two quantum registers.
- The challenger runs $b_1 \leftarrow \text{Ver}(\sigma, \rho_1)$ and $b_2 \leftarrow \text{Ver}(\sigma, \rho_2)$. A wins if $b_1 = b_2 = 1$.

In summary, the difference between quantum lightning and quantum money is therefore that in quantum lightning, unclonability holds, even for adversarially constructed states.

3.3 Elliptic Curves

Consider a large prime p . We represent isomorphism classes of elliptic curves over \mathbb{F}_p by pairs $(j, b) \in \mathbb{F}_p \times \mathbb{Z}$, where $b \in \{0, 1\}$ except in the following cases:

- If $j \equiv 1728 \pmod{p}$ and $p \equiv 1 \pmod{4}$, then $0 \leq b \leq 3$.
- If $j \equiv 0 \pmod{p}$ and $p \equiv 1 \pmod{3}$, then $0 \leq b \leq 5$.

Fix $\alpha \in \mathbb{F}_p$ a quadratic nonresidue. When $j \neq 0, 1728$, define the elliptic curve associated to the pair (j, b) to be given by

$$y^2 = x^3 + \frac{3j\alpha^{2b}}{j-1728}x + \frac{2j\alpha^{3b}}{j-1728}.$$

If $j = 1728$, define (j, b) to be given by

$$y^2 = x^3 + \alpha^b x.$$

If $j = 0$ and $p \equiv 1 \pmod{3}$, we require that α be both a quadratic and a *cubic* nonresidue. Then (j, b) is given by

$$y^2 = x^3 + \alpha^b.$$

There is a bijection between pairs (j, b) and \mathbb{F}_p -isomorphism classes of elliptic curves; see [Sil09, Cor. X.5.4.1]. In the literature, j is known as the *j -invariant* of the elliptic curve, and b enumerates *twists* of elliptic curves with given j .

Definition 3.5. Let \mathcal{U}_p denote the uniform probability distribution on isomorphism classes of elliptic curves over \mathbb{F}_p , given by uniformly randomly choosing a pair (j, b) as above.

There are many statistical results in the literature which refer to a different distribution \mathcal{D}_p , defined as follows.

Definition 3.6. We say $(a, b) \in \mathbb{F}_p^2$ is a Weierstrass pair if $4a^3 + 27b^2 \neq 0$. We associate to the Weierstrass pair the elliptic curve given by $y^2 = x^3 + ax + b$. The Weierstrass distribution \mathcal{D}_p is the probability distribution on isomorphism classes of elliptic curves over \mathbb{F}_p induced by the uniform distribution on Weierstrass pairs (a, b) .

Note that there are exactly p pairs (a, b) for which $4a^3 + 27b^2 = 0$: when $-3a$ is a quadratic residue, there are 2 values for b , plus the pair $(0, 0)$ yields $2 \cdot \frac{p-1}{2} + 1 = p$. Thus there are $p^2 - p$ pairs (a, b) yielding elliptic curves.

Looking ahead, in Lemma 4.11 we show that the distance between \mathcal{U}_p and \mathcal{D}_p is at most $2/p$, and hence is negligible.

Definition 3.7. *For an elliptic curve E/\mathbb{F}_p , let $t = t(E) = p + 1 - \#E(\mathbb{F}_p)$ be the trace of Frobenius on E , let $\frac{t}{2\sqrt{p}}$ be the normalized trace, let $\Delta_{\text{Fr}} := \Delta_{\text{Fr}}(E)$ be $4p - t^2$ the Frobenius discriminant, let $K = K(E) = \mathbb{Q}(\sqrt{-\Delta_{\text{Fr}}})$, and let $D = D(E)$ be the discriminant of K . Let $\mathcal{E}_{\text{sf}, 3p}$ be the set of elliptic curves E over \mathbb{F}_p such that $\Delta_{\text{Fr}}(E) \geq 3p$ and $\Delta_{\text{Fr}}(E)$ is square-free.*

The trace t is an \mathbb{F}_p -isogeny invariant of E , and hence an invariant of the isomorphism class of E . Thus the derived invariants Δ_{Fr} and D depend only on the isomorphism class of E , not on the equation used to define E . Thus by abuse of notation, we write $E \in \mathcal{E}_{\text{sf}, 3p}$ to mean that the isomorphism class of E is in $\mathcal{E}_{\text{sf}, 3p}$.

Definition 3.8. *Given an imaginary quadratic field K with discriminant D , let \mathcal{O}_K be the ring of integers of K , and define the Bach generating set B_K to be the set of ideal classes of unramified primes \mathfrak{l} of \mathcal{O}_K with $N(\mathfrak{l}) < 6(\log D)^2$.*

Bach [Bac90, p. 376] showed that, assuming GRH, B_K generates $\text{Cl}(\mathcal{O}_K)$.

4 Isogeny Building Blocks

In this section we build the isogeny-related tools we will need for our quantum money construction. Some are already known, some have been folklore (but, to our knowledge, never formalized), and some are new. For example, the prior isogeny literature uses the heuristic that the size of the class group is proportional to \sqrt{D} , where D is the discriminant of $\text{End}(E) \subset K$; in fact, the heuristic is sometimes made precise by noting that the constant of proportionality is given by a certain infinite series depending on D . While this suffices if the elliptic curve E is fixed, it is not strong enough to give bounds across large sets of elliptic curves where D varies. In Sect. 4.1 we give, to our knowledge, the first precise lower bounds for the number of elliptic curves mod p having large class groups—see Corollary 4.14. As a consequence for our protocol, we will obtain that minting is efficient, and that forgery attacks which rely on solving the isogeny problem are hard.

Ideal classes are typically encoded as products of prime ideals with small norm. After giving an algorithm for enumerating these prime ideals, we give a new analysis for the effectiveness of these encodings in Sect. 4.4 and Sect. 4.5, specifically by proving an eigenvalue bound for the adjacency matrix of the associated Cayley graph in Proposition 4.22. Our quantum money verification algorithm utilizes the class group action on elliptic curves, and the eigenvalue bound guarantees that verification is efficient.

The results in this section often use number theory that is new to the cryptography literature; while this enables proofs of new results, there is no easy way to present this material to readers who do not have the requisite background in mathematics. However, we have done our best to make the mathematics readable to cryptographers and have also summarized the necessary number theory. We explain the relevant number theory in more details in the full version of this paper.

4.1 Probability of Large Isogeny Class

The goal of this section is to prove Corollary 4.14, which shows that for large p , a random elliptic curve belongs to an exponentially large isogeny class with probability $\geq 14\%$.

Our estimate will be accomplished in three steps:

- In Theorem 4.1, we compute the probability that a random elliptic curve E has Frobenius discriminant in a specified range.
- We determine how likely a random number of the form $4p-t^2$ is square-free in Theorem 4.9, and combine this with Theorem 4.1 to obtain, in Corollary 4.12, that a random elliptic curve is in $\mathcal{E}_{\text{sf},3p}$ with probability $\geq 14\%$.
- In Corollary 4.14, we use a result of Tatzawa to show that for elliptic curves in $\mathcal{E}_{\text{sf},3p}$, the size of the associated class group is at least $0.089 \frac{\sqrt{p}}{\log p}$.

In the next result, we use the Weierstrass representation of elliptic curves, so a pair $(a, b) \in \mathbb{F}_p \times \mathbb{F}_p$ corresponds to the elliptic curve $y^2 = x^3 + ax + b$.

Theorem 4.1. *Let $I \subset [-1, 1]$ be an interval, and let $\mu_{ST}(I) = \frac{2}{\pi} \int_I \sqrt{1-X^2} dX$. Let $N_I(p)$ be the number of elliptic curves over \mathbb{F}_p , encoded as pairs (a, b) with $4a^3 + 27b^2 \neq 0$, with normalized Frobenius trace $\frac{t}{\sqrt{p}} \in I$. Then*

$$\left| \frac{N_I(p)}{p^2} - \mu_{ST}(I) \right| \leq \frac{8}{3p^{1/4}} + \frac{4}{3p^{1/2}} + \frac{4}{p^2} + \frac{4}{p^{9/4}} + \frac{4 \log p}{p^{5/2}}.$$

The measure μ_{ST} is known as the *Sato-Tate measure*.

Proof. Birch [Bir68] showed that

$$\lim_{p \rightarrow \infty} \left| \frac{N_I(p)}{p^2} - \mu_{ST}(I) \right| = 0.$$

An asymptotic error bound for $|N_I(p)/p^2 - \mu_{ST}(I)|$ is given in [MP19b]; we follow that proof to compute the explicit constant. Our notation follows that of [MP19b], with one exception. Suppose that $I = [x_0, x_1] \subset [-1, 1]$, and define $J = [\alpha, \beta] \subset [0, 1]$ where $\cos \pi\alpha = x_1$, $\cos \pi\beta = x_0$. Let $\mu'_{ST}(J) = 2 \int_J \sin^2 \pi\theta d\theta$. Murty-Prabhu [MP19b] work with the measure μ'_{ST} in place of μ_{ST} , but the substitution $x = \cos \pi\theta$ translates between the two.

Let $m \geq 2$ be an integer. For the elliptic curve $y^2 = x^3 + ax + b$, define $\theta_{a,b}$ as the angle in $[0, \pi]$ for which $\cos \theta_{a,b}$ is its normalized trace. If $m = 2k$ is even, we have

$$\sum_{\substack{a,b \in \mathbb{F}_p \\ 4a^3 + 27b^2 \neq 0}} \frac{\sin 2k\theta_{a,b}}{\sin \theta_{a,b}} = 0 \tag{1}$$

as follows. Fix $c \in \mathbb{F}_p$ a quadratic nonresidue. Consider the elliptic curve corresponding to (a, b) , having Frobenius trace t . Then the elliptic curve corresponding to (c^2a, c^3b) has Frobenius trace $-t$. Thus $\theta_{ca,cb} = \pi - \theta_{a,b}$, and so the set of θ values appearing in the sum are symmetric with respect to $\theta \mapsto \pi - \theta$. But $\sin(\pi - \theta) = \sin \theta$, while $\sin 2k(\pi - \theta) = -\sin 2k\theta$. Therefore, $\sum \frac{\sin 2k\theta}{\sin \theta} = 0$.

Now suppose m is odd. By [DS05, Thm. 3.5.2], the dimension of the space of weight $m + 1$ cusp forms is $\leq \lfloor \frac{m+1}{12} \rfloor$. Let $T_{m+1}(p)$ be the p -Hecke operator acting on the latter space of cusp forms. By [Del74, Thm. 8.2], the eigenvalues of $T_{m+1}(p)$ all have magnitude $p^{m/2}$. Therefore the trace of $T_{m+1}(p)$ is bounded by $\frac{1}{12}(m+1)p^{m/2}$. Substituting into the Eichler-Selberg trace formula appearing on [MP19b, p. 31] and combining with [MP19b, eq. (3.3)], we obtain

$$\left| \sum_{\substack{a,b \in \mathbb{F}_p \\ 4a^3 + 27b^2 \neq 0}} \frac{\sin m\theta_{a,b}}{\sin \theta_{a,b}} \right| \leq \frac{1}{6}(m+1)p^{3/2} + 2p^{-\frac{m-3}{2}}. \tag{2}$$

Let $M \geq 2$ be an integer. Let $\widehat{S}_{J',M}^\pm$ denote the Fourier transform of the M th Beurling-Selberg polynomial for the interval $J' = \frac{1}{2}J = [\frac{\alpha}{2}, \frac{\beta}{2}]$; see [MP19b, §2.3].² From [MP19b, p. 30, (c)], we have

$$\widehat{S}_{J',M}^\pm(0) = \frac{\beta - \alpha}{2} \pm \frac{1}{M+1} \tag{3}$$

and, for $0 < m \leq M$,

$$\left| \widehat{S}_{J',M}^\pm(m) + \widehat{S}_{J',M}^\pm(-m) - \frac{\sin \pi m\beta - \sin \pi m\alpha}{\pi m} \right| \leq \frac{2}{M+1}. \tag{4}$$

Since $\frac{2}{M+1} \leq \frac{2}{m}$ and $|\frac{\sin \pi m\beta - \sin \pi m\alpha}{\pi m}| \leq \frac{2}{m}$,

$$|\widehat{S}_{J',M}^\pm(m) + \widehat{S}_{J',M}^\pm(-m)| \leq \frac{4}{m}. \tag{5}$$

An elementary calculation shows that

$$\mu_{ST}(I) = (\beta - \alpha) - \frac{\sin 2\pi\beta - \sin 2\pi\alpha}{2\pi}. \tag{6}$$

² Note that there are several mistakes in the calculation of [MP19b], including an incorrect choice for J' . See the preprint version [MP19a] for a more accurate treatment.

From (3), (4) with $m = 2$, and (6),

$$\left| 2\widehat{S}_{J',M}^{\pm}(0) - \widehat{S}_{J',M}^{\pm}(2) - \widehat{S}_{J',M}^{\pm}(-2) - \mu_{ST}(I) \right| \leq \frac{4}{M+1}. \quad (7)$$

From (1),

$$\left| (\widehat{S}_{J',M}^{\pm}(1) + \widehat{S}_{J',M}^{\pm}(-1)) \sum_{4a^3+27b^2 \neq 0} \frac{\sin 2\theta_{a,b}}{\sin \theta_{a,b}} \right| = 0. \quad (8)$$

From (2) and (5),

$$\left| (\widehat{S}_{J',M}^{\pm}(2) + \widehat{S}_{J',M}^{\pm}(-2)) \sum_{4a^3+27b^2 \neq 0} \frac{\sin 3\theta_{a,b}}{\sin \theta_{a,b}} \right| \leq \frac{4}{3}p^{3/2} + 4 \quad (9)$$

and

$$\begin{aligned} & \left| (\widehat{S}_{J',M}^{\pm}(m) + \widehat{S}_{J',M}^{\pm}(-m)) \sum_{4a^3+27b^2 \neq 0} \left[\frac{\sin(m+1)\theta_{a,b}}{\sin \theta_{a,b}} - \frac{\sin(m-1)\theta_{a,b}}{\sin \theta_{a,b}} \right] \right| \\ & \leq \frac{4(m+2)}{3m}p^{3/2} + \frac{16}{m}p^{-\frac{m-2}{2}}. \end{aligned} \quad (10)$$

From (10), we obtain

$$\begin{aligned} & \sum_{3 \leq m \leq M} \left| (\widehat{S}_{J',M}^{\pm}(m) + \widehat{S}_{J',M}^{\pm}(-m)) \sum_{4a^3+27b^2 \neq 0} \left[\frac{\sin(m+1)\theta_{a,b}}{\sin \theta_{a,b}} - \frac{\sin(m-1)\theta_{a,b}}{\sin \theta_{a,b}} \right] \right| \\ & \leq \frac{8}{3}Mp^{3/2} + 16p^{-\frac{1}{2}} \log M. \end{aligned} \quad (11)$$

Let

$$\begin{aligned} B^{\pm} = & p^2(2\widehat{S}_{J',M}^{\pm}(0) - \widehat{S}_{J',M}^{\pm}(2) - \widehat{S}_{J',M}^{\pm}(-2)) \\ & + \sum_{m=1}^2 \left((\widehat{S}_{J',M}^{\pm}(m) + \widehat{S}_{J',M}^{\pm}(-m)) \sum_{4a^3+27b^2 \neq 0} \frac{\sin(m+1)\theta_{a,b}}{\sin \theta_{a,b}} \right) \\ & + \sum_{m=3}^M \left((\widehat{S}_{J',M}^{\pm}(m) + \widehat{S}_{J',M}^{\pm}(-m)) \sum_{4a^3+27b^2 \neq 0} \frac{\sin(m+1)\theta_{a,b}}{\sin \theta_{a,b}} - \frac{\sin(m-1)\theta_{a,b}}{\sin \theta_{a,b}} \right). \end{aligned}$$

From [MP19a, eq. (6)], we have

$$B^- \leq N_I(p) \leq B^+. \quad (12)$$

Letting $M = \lfloor p^{1/4} \rfloor$, and putting together (7)–(12) we obtain the claim. \square

Next, we determine what portion of elliptic curves with $|t| < \sqrt{p}$ satisfy that $4p - t^2$ is square-free. The following argument through Corollary 4.10 is modeled after the proof of [FI10, Theorem 2.1].

Lemma 4.2 ([You91]). *Let $\gamma \approx 0.577$ be Euler’s constant. Then for $x \geq 2$ an integer,*

$$\sum_{n=1}^x \frac{1}{n} \leq \log x + \gamma + \frac{1}{2x}$$

For $n \in \mathbb{N}$, let $\tau(n)$ be the number of positive divisors of n .

Lemma 4.3. *For $x \geq 2$,*

$$\sum_{d \leq x} \tau(d) \leq x \log x + (2\gamma - 1)x + 4\sqrt{x}.$$

Proof. Let $\{x\}$ denote $x - [x]$. According to the proof of [BKZ18, Theorem 2], we have

$$\sum_{d \leq x} \tau(d) = 2 \sum_{d \leq \sqrt{x}} \left\lfloor \frac{x}{d} \right\rfloor - [\sqrt{x}]^2.$$

Continuing with that proof, but keeping track of the error terms, we get

$$\begin{aligned} \sum_{d \leq x} \tau(d) &\leq 2 \sum_{d \leq \sqrt{x}} \left(\frac{x}{d} + 1 \right) - (\sqrt{x} - \{\sqrt{x}\})^2 \\ &\leq 2x \sum_{d \leq \sqrt{x}} \frac{1}{d} + \sqrt{x} - x + 2\sqrt{x}. \end{aligned}$$

Simplifying and applying Lemma 4.2, the claim follows. □

Let $\rho(n)$ be the number of solutions to $t^2 \equiv 4p \pmod{n}$ with $t \in \mathbb{Z}/n\mathbb{Z}$.

Lemma 4.4. *For $d \geq 1$, $\rho(d^2) \leq 4\tau(d)$.*

Proof. Suppose the prime factorization of d is $q_1^{e_1} \cdots q_s^{e_s}$. By Sun Tzu’s Theorem (Chinese Remainder Theorem), $\rho(d^2) = \prod_i \rho(q_i^{2e_i})$. Observe that $4p$ is not a square mod p^2 , and hence $\rho(d^2) = 0$ if $q_i = p$ for any i . Assume $q_i \neq p$ for all i . If q_i is odd, then $\rho(q_i^{2e_i}) = \rho(q_i)$ by Hensel’s Lemma, and $\rho(q_i) \leq 2$. If $q_i = 2$, then $\rho(q_i^{2e_i}) \leq 8$. Finally, observe that $2^s \leq \tau(d)$, with equality precisely when $e_i = 1$ for all i . The claim follows. □

Lemma 4.5. *If $2 \leq Y < p$, then*

$$\sum_{d=1}^Y \rho(d^2) \leq 4Y \log Y + (8\gamma - 4)Y + 16\sqrt{Y}.$$

Proof. Combine Lemmas 4.3 and 4.4. □

Let $A(n)$ be the number of solutions to $t^2 \equiv 4p \pmod{n}$ with $1 \leq t \leq \sqrt{p}$.

Lemma 4.6. For $d \geq 1$, $|A(d^2) - \rho(d^2) \frac{\sqrt{p}}{d^2}| \leq \rho(d^2)$.

Proof. Let $M = \lfloor \frac{\sqrt{p}}{d^2} \rfloor$. In each of the intervals $[1, d^2], [d^2 + 1, 2d^2], \dots, [(M - 1)d^2 + 1, Md^2]$, there are exactly $\rho(d^2)$ contributions to $A(d^2)$. In the interval $[Md^2 + 1, \frac{\sqrt{p}}{d^2}]$, there are at most $\rho(d^2)$ additional contributions. \square

Lemma 4.7. If $2 \leq Y < \sqrt{p}$, then $\sum_{Y < d < \sqrt{p}} A(d^2) \leq \frac{36p}{Y^2}$.

Proof. For positive integers κ, Y , let

$$\begin{aligned} N(\kappa) &= \#\{(t, d) \in \mathbb{Z}^2 : d > Y, 1 \leq t \leq \sqrt{p}, 4p - t^2 = \kappa d^2\} \\ &= \#\{(t, d) \in \mathbb{Z}^2 : d > Y, 1 \leq t \leq \sqrt{p}, \text{Nm}_{\mathbb{Q}(\sqrt{-\kappa})/\mathbb{Q}}(t + d\sqrt{-\kappa}) = 4p\}. \end{aligned}$$

If (t, d) is a pair counted by $N(\kappa)$, then we have $\kappa = \frac{4p - t^2}{d^2}$, and hence $N(\kappa) = 0$ for $\kappa > \frac{4p}{Y^2}$. We have $\sum_{Y < d < \sqrt{p}} A(d^2) \leq \sum_{1 \leq \kappa \leq \frac{4p}{Y^2}} N(\kappa)$. We claim that $N(\kappa) \leq 9$,

from which the lemma follows. To see this, suppose that (2) and (p) each split as the product of two principal ideals in $\mathbb{Q}(\sqrt{-\kappa})$; say $(2) = \mathfrak{l}_1 \cdot \bar{\mathfrak{l}}_1$ and $(p) = \mathfrak{l}_2 \cdot \bar{\mathfrak{l}}_2$. Then (t, d) is counted by $N(\kappa)$ if and only if $t + d\sqrt{-\kappa}$ is a generator for an ideal of norm $4p$, which must lie in the list $\mathfrak{l}_1^2 \mathfrak{l}_2, (2)\mathfrak{l}_2, \bar{\mathfrak{l}}_1^2 \mathfrak{l}_2, \bar{\mathfrak{l}}_1 \bar{\mathfrak{l}}_2, (2)\bar{\mathfrak{l}}_2, \mathfrak{l}_1^2 \bar{\mathfrak{l}}_2$. Conjugation changes the sign of d , and since we only count $d > 0$, we need only consider half of the above ideals. The number of generators for each ideal is at most the size of the units in the ring of integers of $\mathbb{Q}(\sqrt{-\kappa})$, which is at most 6 (occurring when $\kappa = 3$). But multiplication by -1 only changes the sign of the pair (t, d) , and so there are at most 3 generators per ideal which contribute to the count of $N(\kappa)$. Therefore $N(\kappa) \leq 9$.

Finally, if the splitting behavior of (2), (p) differs from our assumption above, then $N(\kappa)$ can only shrink. \square

Let $\mu(n)$ denote the Möbius function.

Lemma 4.8. If $2 \leq Y < \sqrt{p}$ is an integer, then

$$\left| \sum_{d > Y} \mu(d) \frac{\rho(d^2)}{d^2} \right| \leq 12Y^{-0.7}.$$

Proof. We have

$$\left| \sum_{d > Y} \mu(d) \frac{\rho(d^2)}{d^2} \right| \leq \sum_{d=Y+1}^{\infty} \frac{4\tau(d)}{d^2}.$$

The first displayed inequality in the proof of Lemma 5 from [NR83] yields $\tau(d) \leq 5d^{0.3}$. The result follows from the fact that $\sum_{d=Y+1}^{\infty} \frac{20}{d^{1.7}} \leq \int_Y^{\infty} 20x^{-1.7} dx$. \square

Theorem 4.9. *Let $C = \prod_{q \text{ prime}} \left(1 - \frac{2}{q^2}\right)$ and let p be an odd prime. Let n_p be the number of values of t for which $1 \leq t < \sqrt{p}$ and $4p - t^2$ is square-free. Then*

$$n_p \geq Cp^{\frac{1}{2}} - \frac{4}{3}p^{\frac{1}{3}} \log p - (8\gamma + 32)p^{\frac{1}{3}} - 12p^{\frac{4}{15}} - 16p^{\frac{1}{6}}.$$

Proof. By inclusion-exclusion, $n_p = \sum_{d=1}^{\infty} \mu(d)A(d^2)$. Let $Y = \sqrt[3]{p}$. Using Lemma 4.6, we have

$$\begin{aligned} \sum_{d=1}^{\infty} \mu(d)A(d^2) &= \sum_{d \leq Y} \mu(d)A(d^2) + \sum_{d > Y} \mu(d)A(d^2) \\ &\geq \sqrt{p} \sum_{d \leq Y} \mu(d) \frac{\rho(d^2)}{d^2} - \sum_{d \leq Y} \rho(d^2) - \sum_{d > Y} A(d^2) \end{aligned}$$

Additionally,

$$\sqrt{p} \sum_{d \leq Y} \mu(d) \frac{\rho(d^2)}{d^2} \geq \sqrt{p} \sum_{d=1}^{\infty} \mu(d) \frac{\rho(d^2)}{d^2} - \sqrt{p} \left| \sum_{d > Y} \mu(d) \frac{\rho(d^2)}{d^2} \right|.$$

Combining Lemmas 4.5, 4.7, and 4.8 with the fact that

$$\sum_{d=1}^{\infty} \mu(d) \frac{\rho(d^2)}{d^2} = \prod_{q \text{ prime}} \left(1 - \frac{\rho(q^2)}{q^2}\right) \geq C,$$

we obtain the result. □

Corollary 4.10. *If $p > 2^{63}$ is prime, then the probability that $4p - t^2$ is square-free, where t is randomly chosen in $1 \leq t < \sqrt{p}$, is at least 25%.*

Proof. We wish to bound n_p/\sqrt{p} . The Feller-Tornier constant [OEI24, Seq. A065493] is

$$\frac{1}{2} + \frac{1}{2} \prod_{q \text{ prime}} \left(1 - \frac{2}{q^2}\right) > .66,$$

from which it follows that $C > 0.32$. Thus by Theorem 4.9, $\frac{n_p}{\sqrt{p}} \geq 0.32 - \epsilon(p)$, where

$$\epsilon(p) = \frac{4}{3}p^{-\frac{1}{6}} \log p + (8\gamma + 32)p^{-\frac{1}{6}} + 12p^{-\frac{7}{30}} + 16p^{-\frac{1}{3}}.$$

Since $p > 2^{63}$, $\epsilon(p) \leq \epsilon(2^{63}) < 0.07$, the claim follows.

Recall from Definitions 3.5 and 3.6 the probability distributions \mathcal{U}_p and \mathcal{D}_p on the set of isomorphism classes of elliptic curves over \mathbb{F}_p .

Lemma 4.11. *For any prime $p \geq 5$, the ℓ_2 distance between the distributions \mathcal{D}_p and \mathcal{U}_p is $\leq \frac{2}{p}$.*

Proof. In the distribution \mathcal{D}_p , we represent elliptic curves by the $p^2 - p$ pairs $(a, b) \in \mathbb{F}_p \times \mathbb{F}_p$ satisfying $4a^3 + 27b^2 \neq 0$. As observed in [Bir68, §1], for each isomorphism class of elliptic curves, there are $\frac{(p-1)}{2}$ pairs (a, b) giving rise to it, except for the cases of $y^2 = x^3 + ax$ and $y^2 = x^3 + b$. For $y^2 = x^3 + ax$, there are $(p-1)/4$ pairs $(a, 0)$ if $p \equiv 1 \pmod{4}$ for each of 4 isomorphism classes of elliptic curves, and $(p-1)/2$ pairs for each of 2 isomorphism class if $p \equiv 3 \pmod{4}$. For $y^2 = x^3 + 1$, there are $(p-1)/6$ pairs if $p \equiv 1 \pmod{3}$ yielding 6 isomorphism classes of elliptic curves, and $(p-1)/2$ pairs if $p \equiv 2 \pmod{3}$ for 2 isomorphism classes of elliptic curves. Then \mathcal{D}_p chooses isomorphism classes with $a, b \neq 0$ with probability $\frac{1}{2p}$; curves with $b = 0$ with probability between $\frac{1}{2p}$ and $\frac{1}{4p}$; and curves with $a = 0$ with probability between $\frac{1}{2p}$ and $\frac{1}{6p}$. The largest discrepancy from the uniform distribution \mathcal{U}_p occurs when $p \equiv 1 \pmod{12}$. In this case, for \mathcal{D}_p there are $2p - 2$ isomorphism classes of elliptic curves with $a, b \neq 0$, 4 with $b = 0$, and 6 with $a = 0$, while \mathcal{U}_p is uniform across all $2p + 8$ isomorphism classes. A routine calculation now yields the result. \square

Corollary 4.12. *Let $p > 2^{63}$ be prime. The probability that an elliptic curve drawn from the distribution \mathcal{D}_p lies in $\mathcal{E}_{\text{sf}, 3p}$ is at least 14%. The same holds if we draw from the distribution \mathcal{U}_p .*

Proof. We first choose an elliptic curve according to \mathcal{D}_p . Note that $|\frac{t}{2\sqrt{p}}| < \frac{1}{2}$ implies that $\Delta_{\text{Fr}}(E) \geq 3p$.

Next, we cannot directly use Corollary 4.10 since the distribution of t values is not uniform. But since the density function $\frac{d\mu_{\text{sr}}}{dX} = \frac{2}{\pi}\sqrt{1-X^2}$ is decreasing as a function of $|X| = |\frac{t}{2\sqrt{p}}|$, the probability that $\Delta_{\text{Fr}}(E)$ is square-free, given that $|t| < \sqrt{p}$, is lowest when all of the square-free values occur for $|t|$ as large as possible. (In fact, this will never occur, since even t values result in $4 \mid \Delta_{\text{Fr}}(E)$; but the probability we obtain with this assumption will in any case be a lower bound for the true probability.) Via Corollary 4.10, we therefore assume that the elliptic curves with square-free t and $4p - t^2 \geq 3p$ occur when $0.75\sqrt{p} < |t| < \sqrt{p}$. Applying Theorem 4.1, we see that at least 14% of all elliptic curves mod p have trace t in this range.

The analogous results for \mathcal{U}_p follow from applying Lemma 4.11. \square

Theorem 4.13. *Suppose $0 < \epsilon < \frac{1}{2}$ and K is an imaginary quadratic field. Let D be the absolute value of the discriminant of K , and h the class number of K . If $D > \max(e^{1/\epsilon}, e^{11.2})$, then $h > \frac{0.655\epsilon}{\pi} D^{\frac{1}{2}-\epsilon}$ except for at most one choice of K .*

Proof. See the remarks immediately following Theorem 2 in [Tat51]. The idea is that according to the Dirichlet class number formula, $h = \frac{1}{\pi}\sqrt{D} \cdot L_D(1)$, where L_D is the L -function associated to K (this is a power series related to the Riemann zeta function which encodes number theoretic information about K). Tatzawa in [Tat51] provides lower bounds for $L_D(1)$. \square

For somewhat better bounds, see [Hof80, Theorem 1].

Corollary 4.14. *Let $p > 2^{63}$ be prime. Given an ordinary elliptic curve E/\mathbb{F}_p , let \mathcal{O} be $\text{End}(E)$. If an elliptic curve E is drawn from either the distribution \mathcal{D}_p or the distribution \mathcal{U}_p , then with probability at least 14%, $E \in \mathcal{E}_{\text{sf},3p}$ and both the class group of \mathcal{O} and the isogeny class of E have size at least $0.089 \frac{\sqrt{p}}{\log p}$.*

Proof. By Corollary 4.12, there is a probability of at least 14% such that $E \in \mathcal{E}_{\text{sf},3p}$. For $E \in \mathcal{E}_{\text{sf},3p}$, consider Fr the Frobenius endomorphism and $K = \mathbb{Q}(\text{Fr})$. Since $\Delta_{\text{Fr}}(E)$ is square-free, $\mathbb{Z}[\text{Fr}] = \text{End}(E) \cong \mathcal{O}_K$, and this must hold for any elliptic curve isogenous to E . Therefore $\Delta_{\text{Fr}}(E) = D(E)$ and the isogeny class of E is acted upon simply transitively by the class group, and so has the same size as the class group.

For the size claim, choose $\epsilon = \frac{1}{\ln p}$ in Theorem 4.13. Note that the one possible exceptional K in that Theorem is subsumed by the round-off error in the probability calculation of Corollary 4.12. \square

4.2 The SEA Isogeny Algorithm

We recall the complex multiplication theory of elliptic curves. Given an imaginary quadratic number field K and an order $\mathcal{O} \subset K$, we say that E has complex multiplication by \mathcal{O} if $\text{End}(E) \cong \mathcal{O}$, or if \mathcal{O} is isomorphic to a subring of $\text{End}(E)$, and there is no larger order of K isomorphic to a subring of $\text{End}(E)$; this second case is only necessary for supersingular elliptic curves.

Identify \mathcal{O} with (the corresponding subring of) $\text{End}(E)$. Given an integral ideal I of \mathcal{O} , define $E[I] = \{\cap \ker(\alpha) : \alpha \in I\}$; it is a subgroup of E of order the ideal norm $N(I)$. Write φ_I for the canonical isogeny $\varphi_I : E \rightarrow E_I := E/E[I]$. Observe that $\deg(\varphi_I) = N(I)$. The isomorphism class of E_I depends only on the ideal class of I . We let $M(p)$ be the complexity of one arithmetic operation in \mathbb{F}_p . Then we have

Theorem 4.15. *Let E be an elliptic curve over \mathbb{F}_p with complex multiplication by \mathcal{O} , and let $\mathfrak{l} \subset \mathcal{O}$ be a prime ideal of norm ℓ , where ℓ is a rational prime. Then there is a classical algorithm which computes the isogeny $\varphi_{\mathfrak{l}}$ in time $O(\ell M(p) \log \ell \log \log \ell \log p)$.*

Proof. See [DKS18, p. 12], specifically “Elkies steps” (Algorithms 3 and 4). See the full version of this paper for more details. \square

As we’ll see in Sect. 4.3, we will be concerned with the case where $\ell \leq 6 \log^2(4p)$.

4.3 Building the Generating Set of Isogenies

Below we will give an algorithm to list prime ideal classes in B_K . We will omit prime ideals of the form $\ell \mathcal{O}_K$ for ℓ a rational prime, since such so-called *inert* primes are principal and hence yield the trivial class. (Additionally, the ideal norm of $\ell \mathcal{O}_K$ is ℓ^2 , so the inert primes will quickly exceed the Bach bound.) The prime ℓ is inert if and only if $-D$ is a quadratic nonresidue mod ℓ . If ℓ is not

inert and does not divide D , then ℓ factors as the product of a prime ideal, say \mathfrak{l} , and its conjugate.

We now give an algorithm to generate a list of prime ideals \mathfrak{l}_i representing B_K .

Algorithm 4.1: Bach Generating Set Algorithm

Input: K an imaginary quadratic field with discriminant D

Output: (\mathfrak{l}_i) , list of primes in \mathcal{O}_K

- 1 Initialize $i = 1$ and $\ell = 2$.
 - 2 Check if ℓ is inert by determining if the Legendre symbol $(\frac{-D}{\ell}) = -1$; if yes, go to step 5.
 - 3 By enumeration, find the smallest positive x satisfying $x^2 \equiv -D \pmod{\ell}$.
Let $\mathfrak{l} = (\ell, x + \sqrt{-D})$ and $\mathfrak{l}' = \bar{\mathfrak{l}} = (\ell, x - \sqrt{-D})$.
 - 4 Check if the class of \mathfrak{l} has already been generated by determining if $\mathfrak{l} \cdot \mathfrak{l}_j$ is a principal ideal for any \mathfrak{l}_j with $j < i$. If yes, go to the next step.
Otherwise, set $\mathfrak{l}_i = \mathfrak{l}$. Check if \mathfrak{l}^2 is principal, and if not, then also set $\mathfrak{l}_{i+1} = \mathfrak{l}'$. Increment i by 1 or 2 accordingly.
 - 5 Increment ℓ to the next larger rational prime. If $\ell > 6(\log D)^2$, then output the list (\mathfrak{l}_i) and terminate.
-

We remark on accomplishing each of these steps. Step 2 is clear. For step 3, such an x with $1 \leq x \leq \frac{\ell}{2}$ is guaranteed to exist, as the Legendre symbol is $+1$ and solutions come in additive inverse pairs. For step 4, take the pairwise products of the generators of $\mathfrak{l} \cdot \mathfrak{l}_j$, and let $A \subset \mathbb{C}$ be the lattice generated by these products (where we view $\mathbb{Q}(\sqrt{-D}) \subset \mathbb{C}$ using either field embedding). Then apply Lagrange-Gauss reduction to the lattice; $\mathfrak{l} \cdot \mathfrak{l}_j$ is principal if and only if $\lambda_1(A) = \ell \cdot \ell_j$.

Proposition 4.16. *If \tilde{B}_K is the output of Algorithm 4.1, then the map $\tilde{B}_K \rightarrow B_K$ given by $\mathfrak{l} \mapsto [\mathfrak{l}]$ is a bijection. Furthermore, if $[\mathfrak{l}] \in B_K$, then $[\mathfrak{l}]^{-1} \in B_K$.*

Proof. The definition of B_K immediately yields that the map is well-defined and surjective. Let $\mathfrak{l} \in \tilde{B}_K$ of norm ℓ . If \mathfrak{l}^2 is principal, then $[\mathfrak{l}]$ is its own inverse. Otherwise in step 4, we also have $\bar{\mathfrak{l}} \in \tilde{B}_K$ and since

$$\mathfrak{l} \cdot \bar{\mathfrak{l}} = \ell \mathcal{O}_K$$

is principal, we get $[\mathfrak{l}]^{-1} = [\bar{\mathfrak{l}}] \in \tilde{B}_K$. This proves the second claim.

For the first claim, it suffices to show that the map is injective. Step 4 of the algorithm shows that $\forall \mathfrak{l}, \mathfrak{l}' \in \tilde{B}_K$, $\mathfrak{l} \cdot \mathfrak{l}'$ is not principal. Since \tilde{B}_K is closed under inversion, this implies that if $\mathfrak{l}_i, \mathfrak{l}_j \in \tilde{B}_K$ with $i \neq j$, then $\mathfrak{l}_i \cdot \mathfrak{l}_j^{-1}$ is not principal; in other words, $[\mathfrak{l}_i] \cdot [\mathfrak{l}_j]^{-1} \neq [(1)]$, and hence $[\mathfrak{l}_i] \neq [\mathfrak{l}_j]$. \square

Proposition 4.17. *Algorithm 4.1 requires $O((\log D)^7)$ bit operations and uses $O((\log D)^2 \log \log D)$ bits of memory.*

Note that $\log D = O(\log p)$.

Proof. Each of the first three steps take $O((\log D)^2)$. The length of the generators for the lattice Λ in step 4 is $O(D)$, and hence Lagrange-Gauss take $O((\log D)^3)$. There are $O((\log D)^2)$ pairs to check in one invocation of step 4, and the algorithm repeats $O((\log D)^2)$ times, whence the time estimate.

Each prime is recorded as a pair of generators whose coefficients are of size $\log \ell = O(\log \log D)$. Since there are $O((\log D)^2)$ primes, the space complexity follows. \square

4.4 The Distribution of Class Group Generators

Let G be a finite abelian group. We say that a sequence $h_1, \dots, h_t \in G$ is *weakly Erdős-Rényi* if, for every $g \in G$, $\exists e_1, \dots, e_t \in \{0, 1\}$ such that $g = h_1^{e_1} h_2^{e_2} \dots h_t^{e_t}$.

Definition 4.18. *We define the Bach-Erdős-Rényi game as follows. Given a parameter λ , an adversary wins if it can find an imaginary quadratic field K with discriminant $D > 2^\lambda$ such that the classes $B_K \subset \text{Cl}(\mathcal{O}_K)$ are not weakly Erdős-Rényi.*

Assumption 4.19 For every quantum polynomial time adversary, the probability of winning the Bach-Erdős-Rényi game is a negligible function of λ .

Why should we believe this assumption? First, consider the following theorem.

Theorem 4.20 ([ER65], Theorem 2). *If G is a finite abelian group, $\delta > 0$, and*

$$t \geq \log(\#G) + \log \log(\#G) - 2 \log \delta + 5,$$

then a randomly chosen sequence $h_1, \dots, h_t \in G$ is weakly Erdős-Rényi with probability at least $1 - \delta$.

Taking $G = \text{Cl}(\mathcal{O}_K)$ and $\delta = \frac{1}{D}$, the assumption holds as long as B_K acts like a random set of elements from G . In the full version of this paper, we give heuristic evidence that B_K acts “sufficiently randomly.”

4.5 Eigenvalue Bounds

A *vertex-transitive graph* Γ is one for which for every pair of vertices v, w , there is an automorphism of the graph γ such that $\gamma(v) = w$.

Proposition 4.21 (Lemma 6.1, [Bab91]). *Let Γ be a vertex-transitive graph of degree d and diameter δ . Then the second largest eigenvalue of the adjacency matrix of Γ is $\leq d - \frac{1}{16.5\delta^2}$.*

Fix an isogeny class $\mathcal{I}_N \subset \mathcal{E}_{\text{sf},3p}$. Let X be the graph with vertex set \mathcal{I}_N and for which E, E' are adjacent if and only if $\exists [1] \in B_K$ such that $[1] * E = E'$.

Proposition 4.22. *Suppose B_K is weakly Erdős-Rényi for $\text{Cl}(\mathcal{O}_K)$, and let $r = \#B_K$. Then the second largest eigenvalue μ_2 of the adjacency matrix of X satisfies $\mu_2 \leq r - \frac{1}{16.5r^2}$.*

Proof. Let $E_1, E_2 \in \mathcal{I}_N$. Since $\Delta_{\text{Fr}}(E_1), \Delta_{\text{Fr}}(E_2)$ are square-free, we must have $\Delta_{\text{Fr}}(E_1) = D(E_1) = \Delta_{\text{Fr}}(E_2)$. Therefore $\exists c \in \text{Cl}(\mathcal{O}_K)$ such that $c * E = E'$. Since B_K is weakly-Erdős-Rényi, $\exists [l_1], \dots, [l_t] \in B_K$ such that $c = \prod [l_i]$. The map $E \mapsto (\prod [l_i]) * E$ yields an automorphism of X which sends E to E' , and hence X is vertex-transitive.

Observe that X is a regular graph with degree equal to r . If B_K is weakly Erdős-Rényi, then the diameter of X is bounded above by r . The result now follows from the previous proposition. \square

5 Sampling a Superposition of Elliptic Curves over \mathbb{F}_p

Suppose p is a large prime. Recall that we represent elliptic curves over \mathbb{F}_p by a pair (j, b) where $j \in \mathbb{F}_p$ and b is twisting data. In this subsection, we show how to sample a uniform superposition over elliptic curves over \mathbb{F}_p . To our knowledge, this is not known for *supersingular* elliptic curves [MMP22], and most “natural” ways of generating random elliptic curves run into the index erasure problem [AMRR11] when used to try to generate a superposition of elliptic curves.

Algorithm 5.1: Algorithm ECSupGen

Input: p a prime

Output: $|E\rangle$ a quantum state

Let \mathcal{S} be a register that can store a pair (j, b) , where $j \in \mathbb{F}_p$ and $0 \leq b \leq 5$.

Generate a uniform superposition $|\psi\rangle \in \mathcal{S}$ over all pairs (j, b) , where

- If $j \not\equiv 0, 1728 \pmod{p}$, then $b = 0$ or 1 .
 - If $j \equiv 1728$ and $p \equiv 1 \pmod{4}$, then $0 \leq b \leq 3$. If $p \equiv 3 \pmod{4}$, then $b = 0$ or 1 .
 - If $j \equiv 0$ and $p \equiv 1 \pmod{3}$, then $0 \leq b \leq 5$. If $p \equiv 2 \pmod{3}$, then $b = 0$ or 1 .
-

Proposition 5.1. *Let $|\psi\rangle$ be the output of Algorithm 5.1. Then $|\psi\rangle$ is a uniform superposition over all isomorphism classes of elliptic curves over \mathbb{F}_p . The algorithm takes time $O(\log p)$.*

Proof. The first claim is immediate from the discussion in Sect. 3.3. The complexity estimate comes from generating the superposition over all j , which dominates the conditional superposition over b -values. \square

We remark that encoding a superposition using the Weierstrass encoding (a, b) (corresponding to the elliptic curve $y^2 = x^3 + ax + b$) is also possible. However,

the Weierstrass encoding requires about twice as many qubits. Additionally, since there are $O(p)$ pairs (a, b) corresponding to the same isomorphism class, the group action will not be as nicely behaved; for instance, the set of elliptic curves is the disjoint union of $O(p)$ orbits under the class group action.

6 Verifying a Superposition of Elliptic Curves

In this section, we give an algorithm that verifies that a quantum state is negligibly close to a uniform superposition of elliptic curves over \mathbb{F}_p with a given number of points. Our algorithm is based on the verification algorithm of [LMZ23] (which is in turn an abstraction of the verification procedure of [FGH+12]); at a high level, it is the same as that of [LMZ23], although we have made a number of changes that are specific to our scheme.

6.1 Overview of Verification Algorithm

Let \mathcal{I}_N denote the set of elliptic curves over \mathbb{F}_p with N points, and let $|\mathcal{I}_N\rangle := \frac{1}{\sqrt{\#\mathcal{I}_N}} \sum_{E \in \mathcal{I}_N} |E\rangle$. In other words, $|\mathcal{I}_N\rangle$ is a uniform superposition over all elliptic curves with N points.

Our goal, similar to [LMZ23], is to compute an approximation of the projection-valued measure $V_N = |\mathcal{I}_N\rangle\langle\mathcal{I}_N|$. Unlike [LMZ23], for us there is only a single orbit in \mathcal{I}_N (because the class group acts transitively on the isogeny class), so we can simplify our corresponding PVM relative to theirs.

Note that if we started with a uniform superposition $|\mathcal{I}_N\rangle$, then $V_N|\mathcal{I}_N\rangle = |\mathcal{I}_N\rangle$ immediately. If instead we compute $V_N|\psi\rangle$ for some superposition $|\psi\rangle$ that does not put much weight on $|\mathcal{I}_N\rangle$, we know that V_N is likely to reject. We emphasize that such a projection does not disturb a “correct” state $|\mathcal{I}_N\rangle$. We will show that our algorithm closely mimics the behavior of the PVM V_N .

At a rough level, our algorithm works as follows: we first check to make sure that we are given a state that contains a representation of a (possible) superposition of elliptic curves with N points. Then, as in [LMZ23], we mimic taking a random walk (in superposition) over all elliptic curves with N points. If we ensure that our “steps” in the random walk are invertible (i.e., the mapping is one-to-one), then we have the following nice property: if we start with $|\mathcal{I}_N\rangle$, then taking a step of our walk brings us to $|\mathcal{I}_N\rangle$, which is where we started. If, on the other hand, we started with, say, a single elliptic curve E , then taking a random walk would likely leave us with a different elliptic curve, which is a totally different state. As in previous work [FGH+12, LMZ23], we build this intuition into a full verification algorithm.

6.2 Verification Algorithm Definitions

Definition 6.1. *The Isogeny Computation σ_i . Fix $E_0 \in \mathcal{E}_{\text{sf}, 3p}$, and let \mathcal{I}_N be its isogeny class. Let $\mathcal{V} = \mathbb{C}^{\mathcal{I}_N}$; that is, the complex vector space with orthonormal basis given by $|E\rangle$ for $E \in \mathcal{I}_N$. Let $K = \text{End}(E_0) \otimes \mathbb{Q}$. Let ι_1, \dots, ι_r denote prime*

ideals which form a system of representatives for the classes in B_K . For each $i \in [1, r]$ we let $\sigma_i : \mathcal{V} \rightarrow \mathcal{V}$ denote the unitary given by $\sigma_i|E\rangle = |i * E\rangle$.

Note that by giving only p and $N := \#E_0(\mathbb{F}_p)$, \mathcal{I}_N is determined, while K, B_K can be efficiently computed.

Recall that if $[l] \in B_K$, then $[l]^{-1} \in B_K$ as well. Therefore the σ_i come in pairs which are inverses of each other.

Definition 6.2. The State $|\mathbf{1}_n\rangle$. If n is an integer, define the state $|\mathbf{1}_n\rangle := \frac{1}{\sqrt{n}} \sum_{i=1}^n |i\rangle$ and if $n' > n$, define $|\mathbf{1}_{n,n'}\rangle := \frac{1}{\sqrt{n'-n+1}} \sum_{i=n}^{n'} |i\rangle$.

Definition 6.3. For integers n and k , let $\mathbf{P}_{n,k} := |\mathbf{1}_n\rangle\langle\mathbf{1}_n| \otimes \mathbf{I}_k$, where \mathbf{I}_k denotes the identity matrix acting on k qubits.

We view $\mathbf{P}_{n,k}$ as a projection-valued measure with outputs 0 and 1. By Definition 3.1, the probability of obtaining output 1 when we apply the operator $\mathbf{P}_{n,k}$ to a pure state $|\psi\rangle$ is $\langle\psi|\mathbf{P}_{n,k}|\psi\rangle$, with the measured state collapsing to $\frac{\mathbf{P}_{n,k}|\psi\rangle}{\sqrt{\langle\psi|\mathbf{P}_{n,k}|\psi\rangle}}$.

Lemma 6.4. We have the following:

1. Each $\mathbf{P}_{n,k}$ is a positive semi-definite Hermitian matrix.
2. $\mathbf{P}_{n,k}^2 = \mathbf{P}_{n,k}$.
3. $\mathbf{P}_{n,k} = \mathbf{P}_{n,k}^\dagger$, where “ \dagger ” denotes conjugate transpose.

Definition 6.5. Let $r = \#B_K$, $k = \#\mathcal{I}_N$, and $\mathcal{W} = \mathbb{C}^r \otimes \mathcal{V}$. Let $\mathbf{U} : \mathcal{W} \rightarrow \mathcal{W}$ be the unitary given by

$$\mathbf{U} := \sum_{i=1}^r |i\rangle\langle i| \otimes \sigma_i + \sum_{i=r+1}^{2r} |i\rangle\langle i| \otimes \mathbf{I}_k.$$

Equivalently, for $1 \leq i \leq r$ and $E \in \mathcal{I}_N$,

$$\begin{aligned} \mathbf{U}(|i, E\rangle) &= |i, i * E\rangle \text{ and} \\ \mathbf{U}(|i+r, E\rangle) &= |i+r, E\rangle. \end{aligned}$$

Recall that instead of \mathcal{V} , we work in the larger vector space of all pairs (j, b) . But if $E \notin \mathcal{I}_N$ (which can be efficiently determined via Schoof’s algorithm), then we may define σ_i to act trivially on $|E\rangle$. However, valid bank notes lie in \mathcal{V} , which can be checked efficiently. Thus the action of σ_i on $|E\rangle$, $E \notin \mathcal{I}_N$, will not be relevant.

6.3 Verification Algorithm

Algorithm 6.1: Algorithm ECSupVer

Input: a prime p , integers N and τ , and a quantum state $|\psi\rangle$ stored in a register \mathcal{S}

Output: a bit 0 or 1. If it returns 1, then ECSupVer alters $|\psi\rangle$ to a state $|\psi'\rangle$ which it then outputs.

- 1 Check that $|\psi\rangle$ is properly formatted as a superposition over pairs $(j, b) \in \mathbb{F}_p \times \{0, \dots, 5\}$ with b following the restrictions of Algorithm 5.1, and that $4p - (p + 1 - N)^2 \geq 3p$ and is square-free. If not, output 0.
- 2 Use Schoof's algorithm to compute the number of points in the elliptic curve representation of $|\psi\rangle$ in a new register.
- 3 Measure the value in the new register. If it is *not* N , output 0 and terminate. From N , compute K, B_K with Algorithm 4.1, and \mathbf{U} as in Definition 6.5. Then discard this register.
- 4 Let $r = \#B_K$. Using a new register, create the state $|\varphi\rangle := \mathbf{1}_{2r} \otimes |\psi\rangle \in \mathcal{W} := \mathbb{C}^{2r} \otimes \mathcal{V}$.
- 5 Repeat the following τ times:
 1. Apply the unitary \mathbf{U} to $|\varphi\rangle$.
 2. Apply the projection-valued measurement corresponding to $\mathbf{P}_{2r,k}$ to the resulting state. If the measurement fails (i.e., we do not get a state lying in the set $\mathbf{1}_{2r} \otimes \mathcal{V}$) output 0 and terminate.

Discard the first register and output 1 as well as the resulting state.

We will say that ECSupVer “accepts” if it returns 1 and a state.

6.4 Verification Algorithm Efficiency

We next prove that our verification algorithm is efficient. We do this with the following lemma.

Lemma 6.6. *On input a prime p , integers N and τ , and a quantum state $|\psi\rangle$, the algorithm ECSupVer runs in time*

$$\max(O(\log^8 p), O(\tau(\log^3 p)(\log \log^2 p)(\log \log \log^2 p))).$$

Proof. Note that Schoof's algorithm used in step 2 takes time $O(\log^8 p)$, which dominates the running time of the algorithm before the loop in step 4.

Step 4 is dominated by the cost of step (a), which is the application of the unitary \mathbf{U} which performs the isogeny computation. If we let $M(p)$ be the complexity of one arithmetic operation in \mathbb{F}_p , then we know from Theorem 4.15 that there is a classical algorithm which computes a degree ℓ isogeny in time $O(\ell M(p) \log \ell \log \log \ell \log p)$. Since, in our case, $\ell = O(\log^2 p)$ and $M(p) = O(\log^3 p)$, we know that each iteration of step 4 is upper-bounded by a function which is $O((\log^5 p)(\log \log^2 p)(\log \log \log^2 p))$. The claim follows. \square

6.5 Proof of Verification Algorithm

We now argue that our verification algorithm accepts with all but negligible probability.

Definition 6.7. *The Matrix \mathbf{M} .* Define a unitary operator $\mathbf{M} : \mathcal{V} \rightarrow \mathcal{V}$ by $\mathbf{M} := \frac{1}{r} \sum_{i=1}^r \sigma_i$.

Equivalently, for $E \in \mathcal{I}_N$, $\mathbf{M}|E\rangle = \frac{1}{r} \sum_{i=1}^r |t_i * E\rangle$.

The operator \mathbf{M} is analogous to a similar operator in [LMZ23], and we can borrow from, and make more precise, their analysis and explanation.

Lemma 6.8. *The eigenvalues of \mathbf{M} are real. The largest eigenvalue of \mathbf{M} is 1; the corresponding eigenvector is precisely $|\mathcal{I}_N\rangle = \sum_{E \in \mathcal{I}_N} |E\rangle$. Furthermore, if B_K is weakly Erdős-Rényi, then the second largest eigenvalue λ_2 for M is at most $1 - \frac{1}{16.5r^3}$.*

See the full version of our paper for the proof, which relies on Proposition 4.22.

Let $|\psi_1\rangle, \dots, |\psi_k\rangle$ be an eigenbasis for \mathbf{M} , where we let $|\psi_1\rangle = |\mathcal{I}_N\rangle$. For each j , let a_j denote the eigenvalue corresponding to the eigenstate $|\psi_j\rangle$; if B_K is weakly Erdős-Rényi, then $a_j \leq 1 - \frac{1}{16.5r^3}$ for $j \geq 2$.

The next result shows that Algorithm ECSupVer approximately implements the PVM V_j . Theorem 6.9 implies that “proper” money states (where $\alpha_1 = 1$) are always accepted, and “bad” money states (where α_1 is negligible) are not accepted with noticeable probability. Moreover, if we start with a “good enough” money state, we will still have one post-verification.

Theorem 6.9. *Let p be a prime. Let N be a positive integer for which $D := 4p - (p + 1 - N)^2$ is square-free and larger than $3p$. Let $K = \mathbb{Q}(\sqrt{-D})$, B_K the Bach generating set for K , and suppose B_K is weakly Erdős-Rényi. Let $r = \#B_K$, and suppose $\tau = 33r^3\lambda$.*

1. If $|\psi\rangle = \sum_{j=1}^k \alpha_j |\psi_j\rangle \in \mathcal{V}$ is an arbitrary state, then the probability that ECSupVer accepts on input $|\psi\rangle$ is at least $|\alpha_1|^2$ and at most $|\alpha_1|^2 + 2^{-\lambda}$.
2. If ECSupVer accepts on some input state $|\psi\rangle = \sum_{j=1}^k \alpha_j |\psi_j\rangle$, and $|\psi'\rangle = \sum_{j=1}^k \alpha'_j |\psi_j\rangle$ is the corresponding output of ECSupVer, then $|\alpha'_1| \geq \frac{|\alpha_1|}{\sqrt{|\alpha_1|^2 + 2^{-\lambda}}}$.

Please see the full version of the paper for the proof of this theorem.

7 The Protocol

Our elliptic curve-related security parameter is a large prime p . We can derive this from the “true” security parameter λ , where $\lambda \approx \log p$. In general, the choice of p will be derived from security of elliptic curve isogeny-related problems.

Minting. The minting algorithm $\text{Gen}(1^\lambda)$ takes as input parameter p (determined implicitly by the security parameter λ) and proceeds as follows.

Algorithm 7.1: Minting Algorithm *Mint*

Input: A prime $p \in \mathbb{Z}$

Output: $|\psi\rangle, \sigma \in \mathbb{Z}$

- 1 Let \mathcal{S} be a quantum register that is capable of holding a representation of an elliptic curve. In \mathcal{S} , construct a superposition $\sum|E\rangle$ over all elliptic curves over \mathbb{F}_p using the Algorithm 5.1.
 - 2 Use Schoof's algorithm to compute the number of points in $|E\rangle$ in superposition, and store the result in a new register, yielding the state $\sum|E\rangle|\#E(\mathbb{F}_p)\rangle$
 - 3 In superposition, compute $\Delta_{\text{Fr}}(E)$, then set a third register to be 1 if $\Delta_{\text{Fr}}(E)$ is square-free and $\Delta_{\text{Fr}}(E) > 3p$, and 0 otherwise. Measure this last register; if the result is 0, start over at step 1.
 - 4 Measure the 2nd register (containing $|\#E(\mathbb{F}_p)\rangle$), and output the resulting state, which we refer to as $|\psi\rangle$ and the measured value σ . The state $|\psi\rangle$ is the bank note and σ is the serial number.
-

Observe that the state is a superposition over elliptic curves in a specific isogeny class. Note that $\text{Gen}(1^\lambda)$ outputs tuples of the form $(|\psi\rangle, \sigma)$ as desired.

Verification. The verification algorithm $\text{Ver}(|\psi\rangle, \sigma)$ does the following. Recall that $|\psi\rangle$ is (supposed to be) a superposition of elliptic curves, and σ is supposed to be the number of points in each of the elliptic curves in superposition.

Algorithm 7.2: Verification Algorithm *Ver*

Input: $|\psi\rangle, \sigma \in \mathbb{Z}$

Output: $\{0, \perp\}$ or $\{1, |\psi'\rangle\}$

- 1 Run $\text{ECSupVer}(|\psi\rangle, \sigma)$ and receive an output tuple $(|\psi'\rangle \in \mathcal{S}, b \in \{0, 1\})$.
 - 2 **if** $b = 0$ **then return** 0 and \perp and discard $|\psi'\rangle$
 - 3 **else return** 1 and $|\psi'\rangle$
-

We have deferred a considerable amount of complexity to the actual description of ECSupVer here which is located in Sect. 6.

7.1 Correctness of the Scheme

We next argue that our construction is correct and efficient. We note that this follows almost immediately from the analysis of our ECSupGen and ECSupVer algorithms, but we will present formal arguments here regardless. We start by arguing that, with all but negligible probability, our Gen and Ver algorithms are efficient.

Proposition 7.1. *Let $p > 2^{63}$. For some parameter λ , the minting algorithm Gen on inputs p and λ runs in time $O(3\lambda \log^8 p)$ with probability $1 - 2^{-\lambda}$.*

Proof. Schoof's algorithm for point counting takes time $O(\log^8 p)$, and this dominates asymptotically the cost of minting. From Corollary 4.12, we know that the probability of failure in step 3 is at most 86%. Thus, the probability that the algorithm has not terminated after 3λ iterations of step 3 is less than $2^{-\lambda}$. \square

We make one further (but important) note on the minting algorithm.

Proposition 7.2. *Let $p > 2^{63}$ be prime. Let \mathcal{O} be $\text{End}(E)$ for all E output by the minting algorithm. Then the class group of \mathcal{O} and the isogeny class of E both have size at least $0.089 \frac{\sqrt{p}}{\log p}$.*

Proof. This follows immediately from Corollary 4.14. \square

Proposition 7.3. *The verification algorithm Ver on inputs $|\psi\rangle$ and σ runs in time*

$$\max(O(\log^8 p), O(\tau(\log^5 p)(\log \log^2 p)(\log \log \log^2 p))).$$

where we set $\tau = 33r^3\lambda$ for $r = \#B_K$.

Proof. This follows immediately from Lemma 6.6. \square

Proposition 7.4. *Our quantum money/lightning protocol (Gen, Ver) is correct. More precisely, as required by Definition 3.2, for any polynomially sized (in λ) integer i , we have*

$$\text{Ver}^k(\text{Gen}(1^\lambda)) = (|\psi'\rangle, 1) \tag{13}$$

with probability at least $1 - i2^{-\lambda+1}$ for all $k \leq i$ for some state $|\psi'\rangle$.

Proof. This follows from our earlier results evaluating ECSupGen and ECSupVer . More precisely, from Lemma 4.11, we know that, for any prime $p \geq 5$, the algorithm ECSupGen outputs a state $|\psi\rangle$ that has distance $\leq \frac{2}{p}$ from the uniform superposition of all elliptic curves over \mathbb{F}_p . Therefore, if we write $|\psi\rangle$ in terms of an eigenbasis, the weight α_1 of the eigenstate $|\psi_1\rangle$ must be at least $\alpha_1 \geq \sqrt{1 - \frac{2}{p}}$.

From Theorem 6.9 we know that $\text{Ver}(|\psi\rangle, j)$ accepts on an input with probability at least $|\alpha_1|^2$. Thus, Ver accepts on $\text{Gen}(1^\lambda)$ with probability at least $1 - \frac{2}{p}$.

Also from Theorem 6.9, we know that the output state of ECSupVer (and thus Ver) on some input $|\psi\rangle$ can be written in eigenbasis form where the state $|\psi_1\rangle$ has weight at least $\sqrt{\frac{|\alpha_1|^2}{|\alpha_1|^2 + 2^{-\lambda}}}$. Note that, for $\frac{2^{-\lambda}}{|\alpha_1|^2} \geq 2$, we have

$$\sqrt{\frac{|\alpha_1|^2}{|\alpha_1|^2 + 2^{-\lambda}}} \geq \sqrt{1 - \frac{2^{-\lambda}}{|\alpha_1|^2}} \tag{14}$$

Therefore, in the output state $|\psi''\rangle = \text{Ver}(\text{Gen}(1^\lambda))$, we have that $|\psi''\rangle$ contains the eigenstate $|\psi''_1\rangle$ with weight $\alpha'_1 \geq \sqrt{1 - \frac{2^{-\lambda}}{1 - \frac{2}{p}}}$. Note that this is exceptionally close to 1.

In particular, as long as $|\alpha_1|^2 \geq \frac{1}{2}$ for any $|\psi\rangle$ that is input to `ECSupVer`, we know that the output will be at most a distance of $2^{-\lambda+1}$ from $|\psi_1\rangle$. Through a simple inductive argument, we can see that this will always be the case. Therefore `ECSupVer` and, correspondingly, `Ver` will accept on an input from `Gen` with probability at least $1 - 2^{-\lambda+1}$ in each verification. The final result follows from a simple union bound. \square

Acknowledgements. The authors would like to thank Alice Silverberg, Alina Bucur, and Mark Zhandry for helpful conversations. Thanks also to Neha Prabhu for assistance with the proof of Theorem 4.1, and to Henryk Iwaniec for assistance with the proofs of Lemmas 4.3, 4.4, and 4.5.

References

- [Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity, CCC '09*, pages 229–242, Washington, DC, USA, 2009. IEEE Computer Society.
- [AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 41–60. ACM Press, May 2012.
- [ADMP20] Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In Shihō Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 411–439. Springer, Heidelberg, December 2020.
- [ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 526–555, Virtual Event, August 2021. Springer, Heidelberg.
- [AMRR11] Andris Ambainis, Loïck Magnin, Martin Roetteler, and Jérémie Roland. Symmetry-assisted adversaries for quantum state generation. In *2011 IEEE 26th Annual Conference on Computational Complexity*, pages 167–177. IEEE, 2011.
- [Bab91] László Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. In *23rd ACM STOC*, pages 164–174. ACM Press, May 1991.
- [Bac90] Eric Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.
- [BB87] Charles H. Bennett and Gilles Brassard. Quantum public key distribution reinvented. *SIGACT News*, 18(4):51–53, July 1987.
- [BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for iO: Circular-secure LWE suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. <https://eprint.iacr.org/2020/1024>.
- [BDS16] Shalev Ben-David and Or Sattath. Quantum tokens for digital signatures, 2016. <https://arxiv.org/abs/1609.09047>.
- [BGMZ18] James Bartusek, Jiaxin Guan, Fermi Ma, and Mark Zhandry. Return of GGH15: Provable security against zeroizing attacks. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 544–574. Springer, Heidelberg, November 2018.

- [Bir68] B. J. Birch. How the number of points of an elliptic curve over a fixed prime field varies. *Journal of the London Mathematical Society*, s1-43(1):57–60, 01 1968.
- [BKZ18] Bruce C. Berndt, Sun Kim, and Alexandru Zaharescu. The circle problem of Gauss and the divisor problem of Dirichlet–still unsolved. *The American Mathematical Monthly*, 125(2):99–114, 2018.
- [CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 556–584, Virtual Event, August 2021. Springer, Heidelberg.
- [Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291, 2006. <https://eprint.iacr.org/2006/291>.
- [CPDDF+19] Marta Conde Pena, Raul Durán Díaz, Jean-Charles Faugère, Luis Hernández Encinas, and Ludovic Perret. Non-quantum cryptanalysis of the noisy version of Aaronson–Christiano’s quantum money scheme. *IET Information Security*, 13(4):362–366, 2019.
- [Del74] Pierre Deligne. La conjecture de Weil : I. *Publications Mathématiques de l’IHÉS*, 43:273–307, 1974.
- [DKS18] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 365–394. Springer, Heidelberg, December 2018.
- [DS05] F. Diamond and J. Shurman. *A First Course in Modular Forms*. Graduate Texts in Mathematics. Springer, 2005.
- [ER65] P. Erdős and A. Rényi. Probabilistic methods in group theory. *Journal d’Analyse Mathématique*, 14(1):127–138, Dec 1965.
- [FGH+10] Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, Daniel Nagaj, and Peter Shor. Quantum state restoration and single-copy tomography for ground states of hamiltonians. *Physical review letters*, 105(19):190503, 2010.
- [FGH+12] Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, and Peter W. Shor. Quantum money from knots. In Shafi Goldwasser, editor, *ITCS 2012*, pages 276–289. ACM, January 2012.
- [FI10] J. B. Friedlander and H. Iwaniec. Square-free values of quadratic polynomials. *Proc. Edinb. Math. Soc. (2)*, 53(2):385–392, 2010.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Heidelberg, March 2015.
- [Hof80] Jeffrey Hoffstein. On the Siegel-Tatuzawa theorem. *Acta Arithmetica*, 38:167–174, 1980.
- [Kan18] Daniel M. Kane. Quantum money from modular forms, 2018. <https://arxiv.org/abs/1809.05925>.
- [KLS22] Andrey Boris Khesin, Jonathan Z Lu, and Peter W Shor. Publicly verifiable quantum money from random lattices, 2022. <https://arxiv.org/abs/2207.13135v2>.
- [KSS21] Daniel M. Kane, Shahed Sharif, and Alice Silverberg. Quantum money from quaternion algebras. Cryptology ePrint Archive, Report 2021/1294, 2021. <https://eprint.iacr.org/2021/1294>.

- [LAF+10] Andrew Lutomirski, Scott Aaronson, Edward Farhi, David Gosset, Jonathan A. Kelner, Avinatan Hassidim, and Peter W. Shor. Breaking and making quantum money: Toward a new quantum cryptographic protocol. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 20–31. Tsinghua University Press, January 2010.
- [LLQZ22] Jiahui Liu, Qipeng Liu, Luowen Qian, and Mark Zhandry. Collusion resistant copy-protection for watermarkable functionalities. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 294–323. Springer, Heidelberg, November 2022.
- [LMZ23] Jiahui Liu, Hart Montgomery, and Mark Zhandry. Another round of breaking and making quantum money: How to not build it from lattices, and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 611–638. Springer, Heidelberg, April 2023.
- [LO77] J. C. Lagarias and A. M. Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pages 409–464. Academic Press, London-New York, 1977.
- [MMP22] Marzio Mula, Nadir Murru, and Federico Pintore. Random sampling of supersingular elliptic curves. Cryptology ePrint Archive, Report 2022/528, 2022. <https://eprint.iacr.org/2022/528>.
- [Mon94] H.L. Montgomery. *Ten Lectures on the Interface between Analytic Number Theory and Harmonic Analysis*. Conference board of the mathematical sciences regional conference series in mathematics. Conference Board of the Mathematical Sciences, 1994.
- [MP19a] M. Ram Murty and Neha Prabhu. The error term in the Sato-Tate theorem of Birch, 2019. <https://arxiv.org/abs/1906.03534>.
- [MP19b] M. Ram Murty and Neha Prabhu. The error term in the Sato-Tate theorem of Birch. *Bulletin of the Australian Mathematical Society*, 100(1):27–33, 2019.
- [NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [NR83] J. L. Nicolas and G. Robin. Majorations explicites pour le nombre de diviseurs de n . *Canadian Mathematical Bulletin*, 26(4):485–492, 1983.
- [OEI24] OEIS Foundation Inc. The On-Line Encyclopedia of Integer Sequences, 2024. Published electronically at <http://oeis.org>.
- [Rob21] Bhaskar Roberts. Security analysis of quantum lightning. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 562–567. Springer, Heidelberg, October 2021.
- [RS06] Alexander Rostovtsev and Anton Stolbunov. Public-Key Cryptosystem Based On Isogenies. Cryptology ePrint Archive, Report 2006/145, 2006. <https://eprint.iacr.org/2006/145>.
- [Sch95] René Schoof. Counting points on elliptic curves over finite fields. *Journal de théorie des nombres de Bordeaux*, 7(1):219–254, 1995.
- [Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.
- [Tat51] Tikao Tatzuzawa. On a theorem of Siegel. *Japanese journal of mathematics: transactions and abstracts*, 21:163–178, 1951.

- [Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983.
- [WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Heidelberg, October 2021.
- [You91] Robert M. Young. 75.9 Euler’s constant. *The Mathematical Gazette*, 75(472):187–190, 1991.
- [Zha19] Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 408–438. Springer, Heidelberg, May 2019.
- [Zha24] Mark Zhandry. Quantum money from abelian group actions. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2024.



Generalized Hybrid Search with Applications to Blockchains and Hash Function Security

Alexandru Cojocaru¹(✉), Juan Garay², and Fang Song³

¹ School of Informatics, University of Edinburgh, Edinburgh, UK
a.cojocaru@ed.ac.uk

² Department of Computer Science and Engineering, Texas A&M University,
College Station, USA
garay@tamu.edu

³ Department of Computer Science, Portland State University, Portland, USA
fang.song@pdx.edu

Abstract. In this work we first examine the hardness of solving various search problems by hybrid quantum-classical strategies, namely, by algorithms that have both quantum and classical capabilities. We then construct a hybrid quantum-classical search algorithm and analyze its success probability.

Regarding the former, for search problems that are allowed to have multiple solutions and in which the input is sampled according to arbitrary distributions, we establish their hybrid quantum-classical query complexities—i.e., given a fixed number of classical and quantum queries, determine what is the probability of solving the search task. At a technical level, our results generalize the framework for hybrid quantum-classical search algorithms recently proposed by Rosmanis [Ros22]. Namely, for an *arbitrary* distribution D on Boolean functions, the probability that an algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds in finding a preimage of 1 for a function sampled from D is at most $\nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2$, where ν_D captures the average (over D) fraction of preimages of 1.

Regarding our second contribution, we design a hybrid algorithm which first spends all of its classical queries and in the second stage runs a “modified Grover” in which the initial state depends on the target distribution D . We then show how to analyze its success probability for arbitrary target distributions and, importantly, its optimality for the uniform and the Bernoulli distribution cases.

As applications of our hardness results, we first revisit and generalize the formal security treatment of the Bitcoin protocol called the *Bitcoin backbone* [Eurocrypt 2015], to a setting where the adversary has both quantum and classical capabilities, presenting a new *hybrid honest majority* condition necessary for the protocol to properly operate. Secondly, we re-examine the generic security of hash functions [PKC 2016] against quantum-classical hybrid adversaries.

The full version of the paper can be found at [CGS23].

1 Introduction

The query model is an elegant abstraction and is widely adopted in cryptography. A notable example is the random oracle (RO) model [BR93], where a hash function f is modeled as a random black-box function, and all parties including the adversary can evaluate it only by issuing a query x and receiving $f(x)$ in response. Numerous cryptosystems have been designed and analyzed in the RO model—e.g., [BR94, BR96, Sho01, FOPS04, FO13].

The advent of quantum computing brings about a new query model, where *superposition* queries to the hash function f in the form of $\sum_{x,y} \alpha_{x,y} |x\rangle |y\rangle \mapsto \sum_{x,y} \alpha_{x,y} |x\rangle |y \oplus f(x)\rangle$ are permitted, which equips quantum adversaries with new capabilities. Indeed, some classically secure digital signature and public-key encryption schemes are broken in the *quantum* random oracle (QRO) model, where a quantum adversary is able to make such superposition queries to f [YZ21]. As such, a significant amount of effort has been devoted to address such quantum-query adversaries (cf. [BDF+11, ES15, Unr15, HHK17, AHU19, DFMS19, CMS19, ES20, DFMS22]), often resulting in considerable efficiency overhead, such as more complex constructions or larger key sizes, in order to maintain security.

However alarming this threat is, it does not come for free, as it requires running a large-scale quantum computer coherently for an extended amount of time, while in the near-to-intermediate term the available quantum devices are likely to be computationally restricted as well as expensive [Pre18]. This reality inspires a *hybrid* query model, where the computational entity (the adversary) is granted a quota of both classical and quantum queries, resulting in a model which subsumes the classical and quantum query models as special cases. Thus, establishing a trade-off between classical and quantum queries allows giving a more accurate estimation of security and hence optimized parameter choices for cryptosystems depending on what resources are likely to be available to near-term quantum adversaries.

Recently, Rosmanis studied the basic unstructured search problem in the hybrid query model [Ros22], where given oracle function $f : X \rightarrow \{0, 1\}$, one wants to find a “marked” input, i.e., x with $f(x) = 1$. This search problem and many variants, such as multiple or randomly chosen marked inputs, are well understood when all queries are quantum [Gro96, BBBV97, Zal99, DH09, Zha19], and where Grover’s quantum algorithm gives a quadratic speedup over classical algorithms, which is also proven to be optimal [BBBV97]. To reiterate, Rosmanis’s work proves the hardness of searching in the domain of a function with a *unique* marked input x^* in the hybrid query model. Specifically, any quantum algorithm with τ_c classical queries and τ_q quantum queries succeeds in finding x^* with probability at most $\frac{1}{|X|} \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2$. This hardness bound is also shown in [HLS22], by a new recording technique tailored to the hybrid query model.

1.1 Our Contributions and Technical Overview

Bounding the Hardness of Hybrid Search

In this work, we consider an arbitrary distribution D on the function family $\mathcal{F} = \{f : X \rightarrow \{0, 1\}\}$, and prove a precise upper bound on the probability of finding a preimage x with $f(x) = 1$ when $f \leftarrow D$, for any algorithm \mathcal{A} spending τ_c classical and τ_q quantum queries. Specifically, we show that:

$$\Pr_{f \leftarrow D} [f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2,$$

where $\nu_D \stackrel{\text{def}}{=} \sup_{\varphi: \|\varphi\| \leq 1} \left(\mathbb{E}_{f \leftarrow D} \left\| \left(\sum_{x: f(x)=1} |x\rangle \langle x| \right) \varphi \right\|^2 \right)$ captures the *average* fraction of preimages of 1 and is solely determined by the distribution D .

Our generalized bound then allows us to derive hardness bounds for specific relevant distributions. “All” we need to do is to analyze ν_D , and this usually can be done by simple combinatorial arguments. For example, let D be the uniform distribution over functions with exactly one marked input. Then we can observe that $\nu_D = \Pr_{f \leftarrow D} [f(x) = 1] = 1/|X|$ for an arbitrary x , which reclaims the result by Rosmanis [Ros22]. The hardness of searching given a function with $w > 1$ marked items can be similarly derived.

We further demonstrate our result on another distribution D_η , where each input is marked according to a Bernoulli trial. Namely, for every $x \in X$, we set $f(x) = 1$ with probability η *independently*. By determining ν_D in this case, we derive the hardness of search when the function is drawn from D_η . This search problem under D_η , which we call *Bernoulli Search*, is particularly useful in several cryptographic applications. Firstly, we can prove generic security bounds for hash function properties, such as preimage-resistance, second-preimage resistance and their multi-target extensions, against hybrid quantum-classical adversaries. This follows by first adapting the reductions in [HRS16], where the hash properties are connected to the *Bernoulli Search* problem in the fully quantum query setting, and then plugging in our hybrid hardness bound of *Bernoulli Search*. In another application, *Bernoulli Search* was shown to dictate the security of proofs of work (PoWs) and security properties of Bitcoin-like blockchains in the RO model (with fully quantum queries) [CGK+23]. This allows us to identify a new *honest-majority* condition under which the security of the PoW-based Bitcoin blockchain holds against hybrid adversaries equipped with both classical and quantum queries.

At a technical level, the proof of our hardness bound follows the overall strategy of [Ros22]. As in the standard optimality proof of Grover’s algorithm [BBV97], one would consider running an adversary’s algorithm with respect to the input function $f \leftarrow D$ or a constant-0 function. Then one argues that each query diverges the states in these two cases, which is called a *progress measure*, by a small amount. On the other hand, in order to find a marked input in f , the final states need to differ significantly. Therefore, sufficiently many queries are necessary for the cumulative progress to grow adequately.

Now, when classical queries are mixed up with quantum queries, the quantum states would collapse after each classical query, and it becomes unclear how to measure the progress. To address this, Rosmanis considers instead an intermediate oracle named *pseudo-classical*. Namely, consider a quantum query with the output register initialized to $|0\rangle$: $\sum_x \alpha_x |x\rangle |0\rangle \mapsto \sum_x \alpha_x |x\rangle |f(x)\rangle$. We can then view a classical query as the result of measuring the input register that collapses to x and receiving $f(x)$, whereas a pseudo-classical oracle measures the output register, resulting in one of two possible outcomes: $\sum_{x:f(x)=0} \alpha_x |x\rangle |0\rangle$ (denoted as the *0-outcome branch*) or $\sum_{x:f(x)=1} \alpha_x |x\rangle |1\rangle$ (denoted as the *1-outcome branch*). With this change, one instead tracks the progress between: (i) the 0-outcome branch in case of $f \leftarrow D$, and (ii) the state in case of the constant-0 function (which always stays in the 0-outcome branch). The algorithm fails if its state stays in the 0-outcome branch and is close to the state in the constant-0 case. A key ingredient in our proof is to deliberately separate the evolution of various objects on an *individual* function and which *characteristics* of the distribution D influence the evolution and in what way. This enables us to obtain a clean and concise lower bound for the generalized hybrid search problem.

Hybrid Search Algorithms: Design and Analysis. In the second part of our work we focus on constructing a hybrid search algorithm for an arbitrary distribution D and show that in several interesting cases (e.g., Bernoulli) the algorithm is optimal. Inspired by our hardness analysis, our algorithm proceeds in a two-stage fashion:

- The first stage is purely *classical*. We query the τ_c inputs that are the most likely to be assigned the value 1 under D . More precisely, for any x in the input domain, let the function $\omega(x) = \sum_f D(f) \cdot f(x)$, which can be viewed as the (unnormalized) probability that $f(x) = 1$ with f drawn from D . Let S be the set of inputs whose $\omega(x)$ values are the τ_c -highest (ties are broken arbitrarily). Then the algorithm queries all the points $x \in S$. If none of them give a solution, we move on to the second stage.
- The second stage is fully *quantum*. We run a modified Grover algorithm \mathcal{A} which is tailored to the prior knowledge on the distribution D . Instead of starting from an equal superposition of all points in the search space as in the standard Grover search algorithm, we construct an initial state in which the amplitude of each point is proportional to $\omega(x)$. Then, for each of the τ_q quantum queries, two reflection operators are applied to rotate the initial state towards a target state encoding the solutions. We give a comprehensive analysis and derive a precise lower bound for the success probability of \mathcal{A} on the distribution D , which amounts to $\tau_q^2 \cdot \frac{\sum_x \omega^2(x)}{\sum_x \omega(x)}$. In other words, for the algorithm in the second stage, we define an induced distribution \tilde{D} by restricting and (re-normalizing) D to functions f satisfying $f(x) = 0$ for all $x \in S$. We then invoke \mathcal{A} on \tilde{D} in a modular way.

Note that the hybrid algorithm needs to compute the values $\omega(x)$ from the description of the target distribution D , and during the quantum procedure, the

algorithm will implement a unitary dependent on the $\omega(x)$ values, hence the algorithm does not need to be time efficient.

We can show that the success probability of the hybrid algorithm is at least the average of the success probabilities of the classical stage and of the quantum stage. In some special cases, such as the Bernoulli distribution, both the classical probability (i.e., at least one success in τ_c Bernoulli trials) and the weights $w(x)$ (hence the quantum success probability) are easy to derive. We can show that the hybrid algorithm gives matching lower bounds to the hardness bounds proven in the first part of our work.

Discussion and Directions for Future Work. We believe that the hybrid query model is both of theoretical and practical importance. Since near-term quantum computers are limited and expensive, it is to the interest of a party to supplement it with massive classical computational power. This also reflects the fact that those parties who have early access to quantum computers (e.g., large tech companies and government agencies) largely coincide with those who are capable of employing classical clusters and supercomputers. Next, we discuss some future directions.

One immediate question is to study other problems in the hybrid query model. The work of [HLS22] proves the hardness of the collision problem by their generalized recording technique in the hybrid query model. It would be useful to further develop techniques and establish more query complexity results.

Our applications to hash functions and Bitcoin-like blockchains can be seen as analyzing cryptographic constructions in the QRO model against hybrid adversaries. Many block ciphers rely on a different model, known as the *ideal cipher* model. As a simple example, the Even-Mansour cipher encrypts a message m by $E_k : m \mapsto \sigma(k \oplus m) \oplus k$, where σ is a random permutation given as an oracle and k is the secret key. As it turns out, this classically secure cipher is completely broken when quantum queries are allowed to both E_k and σ [KM10]. Since the secret key k is managed by honest users, it is debatable whether superposition access to E_k is realistic, and there has been progress in re-establishing the cipher's security under a partially quantum adversary with quantum access to σ but classical access to E_k [JST21, ABKM22]. The hybrid query model we consider in this work suggests further relaxing the queries to σ to be a hybrid of classical and quantum ones, and it would be valuable to re-examine the security of such schemes in the ideal cipher model.

Querying an oracle also occurs more broadly in many other cryptographic scenarios. Security definitions often give some algorithm as an oracle to the adversary, such as an encryption oracle in the chosen-plaintext-attack (CPA) game, and a signing oracle in formalizing the unforgeability of digital signatures. There has been a considerable effort of settling appropriate definitions and constructions (e.g., quantum-accessible pseudorandom functions, encryption and signatures) when quantum adversaries are granted superposition queries to these oracles (cf. [BZ13, Zha15, AMRS20, Zha21, CEV23]). Extending such efforts to the hybrid-adversary landscape would offer fine-grained security assessments of post-

quantum cryptosystems. Finally, in the context of complexity theory, the study of hybrid algorithms is further motivated by related models focusing on the interplay between classical computation and near-future quantum devices [CCHL22], and between circuit depth and quantum queries [SZ19, CM20, CCL23].

Organization of the Paper. The rest of the paper is organized as follows. The generalized search problem we are considering, which we call *Distributional Search*, is stated in Sect. 2, together with its hybrid quantum-classical hardness; two case studies: Multi-Uniform Search and Bernoulli Search; as well as our proposed hybrid search algorithm. Detailed proofs and analyses of our main results above are presented in Sect. 3—hardness in Sect. 3.1 and the quantum algorithm analysis in Sect. 3.2, respectively. Due to space constraints, the applications of Bernoulli Search, as well as some of the proofs are presented in the full version of the paper [CGS23].

2 Problem Definition(s) and Main Results

2.1 The Distributional Search Problem

The underlying problem we consider is the search for a preimage of 1 of an arbitrarily distributed black-box boolean function.

Distributional Search Problem (Dist-Search)

Let D be an arbitrary distribution supported on the function family $\mathcal{F} = \{f : X \rightarrow \{0, 1\}\}$.

Given: Black-box access to a function f drawn from distribution D .

Goal: Find x such that $f(x) = 1$ if there exists such an x .

It is not surprising that the problem’s hardness is crucially influenced by the number of solutions *on average* under D ; however, what is interesting about our study is that we can show a clean quantitative relation.

Let $f : X \rightarrow \{0, 1\}$ be an arbitrary function. We define the projector on the space spanned by the preimages of 1 as: $\pi_f \stackrel{\text{def}}{=} \sum_{x:f(x)=1} |x\rangle\langle x|$.

Denote by $\pi_f^\perp \stackrel{\text{def}}{=} \mathbb{1} - \pi_f$, and let D be a distribution on \mathcal{F} . We define the value that captures the *average* fraction of preimages of 1 as:

Definition 1 (ν_D). *The average fraction of solutions in \mathcal{F} is defined as:*

$$\nu_D \stackrel{\text{def}}{=} \sup_{\varphi: \|\varphi\| \leq 1} \left(\mathbb{E}_{f \leftarrow D} \|\pi_f \varphi\|^2 \right), \quad (1)$$

where $\|\varphi\|$ denotes the Euclidean norm of the quantum state φ .

Characterization of ν_D . To better understand the ν_D value, we now derive an alternative characterization. For simplicity, assume without loss of generality

that the domain of our target functions is $X = [m] \stackrel{\text{def}}{=} \{1, \dots, m\}$, for some positive integer m . We will write down the truth table to represent each $f : [m] \rightarrow \{0, 1\}$ as a bitstring $x \in \{0, 1\}^m$ and denote by x_i the i -th bit of x .

In this way, D becomes a distribution on $\{0, 1\}^m$, and we write $d_x \stackrel{\text{def}}{=} D(x)$ as the probability of sampling x from the distribution D . Then, from Definition 1, we can rewrite ν_D as:

$$\nu_D = \sup_{\varphi} \left(\mathbb{E}_{x \leftarrow D} \|\pi_x \varphi\|^2 \right), \quad \text{where } \pi_x \stackrel{\text{def}}{=} \sum_{i: x_i=1} |i\rangle \langle i|.$$

Let $\varphi := \sum_{i=1}^m \alpha_i |i\rangle$, with $\|\alpha\| \leq 1$. We have:

$$\begin{aligned} \nu_D &= \sup_{\alpha: \|\alpha\| \leq 1} \mathbb{E}_{x \leftarrow D} \sum_{i: x_i=1} \alpha_i^2 \\ &= \sup_{\alpha: \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \sum_{x \in \{0,1\}^m} d_x \cdot x_i \\ &= \sup_{\alpha: \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_i, \end{aligned}$$

where, for each $i \in [m]$, we define $\omega_i \stackrel{\text{def}}{=} \sum_{x \in \{0,1\}^m} d_x \cdot x_i$. In other words, ω_i captures the likelihood that x_i is assigned value 1 under D . Then it becomes clear that the supremum is achieved by a vector α having 0 entries except taking 1 on i^* where ω_{i^*} is maximized: $\sup_{\alpha: \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_i \leq \sup_{\alpha: \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_{i^*} = \omega_{i^*} \sup_{\alpha: \|\alpha\| \leq 1} \alpha_i^2 \leq \omega_{i^*}$. Therefore,

$$\nu_D = \omega_{i^*} = \max_{i \in [m]} \omega_i. \quad (2)$$

We also note that for any $i \in [m]$, ω_i/ω , where $\omega \stackrel{\text{def}}{=} \sum_i \omega_i$, can be viewed as the probability¹ that $x_i = 1$, when x is sampled according to D .

2.2 Hardness of Dist-Search

Next, we turn to establishing the following bound for the success probability of solving Dist-Search, which constitutes one of our main results:

Theorem 1 (Hardness of Dist-Search – fixed query order). *For any algorithm \mathcal{A} making up to τ_c classical queries and τ_q quantum queries (with a fixed order of the queries independent of f), \mathcal{A} solves the Dist-Search problem with probability:*

$$\text{Succ}_{\mathcal{A}, D} = \Pr_{f \leftarrow D} [f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2.$$

¹ We remark that normalization is required as ω might not be 1; for example, in the case of the Bernoulli distribution, $\omega = m\eta$. Hence, normalization is needed so as to view ω_i/ω as a probability distribution.

The proof can be found in Sect. 3.1. By relying on the result by Don *et al.* [DFH22], the above hardness result can be directly extended to any general hybrid algorithm in which the order of the classical and quantum queries can be adaptive (and can depend on the underlying oracle), at the cost of only a constant factor; i.e. increasing the number of classical and quantum queries by a factor of 2:

Theorem 2 (Hardness of Dist-Search). *For any algorithm \mathcal{A} making τ_c classical queries and τ_q quantum queries, \mathcal{A} solves the Dist-Search problem with probability:*

$$\text{Succ}_{\mathcal{A},D} := \Pr_{f \leftarrow D} [f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2.$$

As this bound for general adversaries is directly derived from the hardness of hybrid algorithms with a fixed query order, in the sequel we will only focus on proving Theorem 1.

2.3 Case Studies

In this section, we will apply our hardness result to two common function distributions. As a common ingredient, it will be helpful to consider the following indicator random variable:

$$\mathbb{1}_x^f \stackrel{\text{def}}{=} \begin{cases} 1 & , \text{ if } f(x) = 1 ; \\ 0 & , \text{ if } f(x) = 0 , \end{cases}$$

for all $f \in \mathcal{F}$ and $x \in X$. Then, for a distribution D :

$$\mathbb{E}_{f \leftarrow D}(\mathbb{1}_x^f) = \Pr_{f \leftarrow D} [f(x) = 1].$$

2.3.1 Multi-uniform Search The first interesting case is a general Grover-type search. We consider a distribution D_w which is *uniform* over functions that map exactly w inputs to 1. In other words, drawing $f \leftarrow D_w$ is equivalent to sampling a subset $S \subseteq X$ with $|S| = w$ uniformly at random and set $f(x) = 1$ if and only if $x \in S$. We consider the resulting multi-uniform search problem:

Multi-Uniform Search

Given: $f \leftarrow D_w$, which maps a uniform size- w subset to 1.

Goal: Find x such that $f(x) = 1$.

Theorem 3. *For any adversary \mathcal{A} making τ_c classical queries and τ_q quantum queries,*

$$\text{Succ}_{\mathcal{A},D_w} \leq \frac{w}{M} \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2,$$

where $M = |X|$ is the domain size.

Proof. We just need to show that $\nu_D = \sup_{\varphi: \|\varphi\| \leq 1} \mathbb{E}_{f \leftarrow D_w} (\|\pi_f \varphi\|^2) \leq \frac{w}{M}$ in this case. Consider an arbitrary unit vector $\varphi = \sum_x \alpha_x |x\rangle$ with $\sum_x |\alpha_x|^2 = 1$.

$$\begin{aligned} \mathbb{E}_{f \leftarrow D_w} (\|\pi_f \varphi\|^2) &= \mathbb{E}_{f \leftarrow D_w} \left(\left| \sum_x \alpha_x \mathbb{1}_x^f |x\rangle \right|^2 \right) \\ &= \sum_x |\alpha_x|^2 \cdot \mathbb{E}_{f \leftarrow D_w} (\mathbb{1}_x^f) \\ &= \sum_x |\alpha_x|^2 \cdot \Pr_{f \leftarrow D_w} [f(x) = 1] = \frac{w}{M}. \end{aligned}$$

Alternatively, using the characterization of ν_D (Eq. 2), we can derive this result, by first noticing that for the multi-uniform distribution we have:

$$d_x = \begin{cases} \frac{1}{\binom{M}{w}}, & \text{if } \text{hw}(x) = w; \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Then, by relying on the characterization of ν_D , we can directly conclude that:

$$\begin{aligned} \nu_D &= \max_{i \in [M]} \omega_i = \max_{i \in [M]} \sum_{x \in \{0,1\}^M} d_x \cdot x_i \\ &= \frac{1}{\binom{M}{w}} \cdot \sum_{x: \text{hw}(x)=w} x_i \\ &= \frac{1}{\binom{M}{w}} \cdot \binom{M-1}{w-1} = \frac{w}{M}. \end{aligned} \quad (4)$$

□

Next, we note two special scenarios. When $w = 1$, our result reproduces Rosmanis's result [Ros22], and when $\tau_c = 0$, it reproduces the fully quantum query complexity of Grover search with multiple marked items (cf. [BBBV97, Zai99]).

2.3.2 Bernoulli Search The second interesting case we consider is what we call a Bernoulli distribution D_η on \mathcal{F} , as specified below:

Bernoulli Search

Given: $f \leftarrow D_\eta$ drawn via the following sampling procedure.

For each $x \in X$, *independently* set:

$$f(x) = \begin{cases} 1, & \text{with probability } \eta; \\ 0, & \text{otherwise.} \end{cases}$$

Goal: Find x such that $f(x) = 1$.

Theorem 4. For any adversary \mathcal{A} making up to τ_c classical queries and τ_q quantum queries,

$$\text{Succ}_{\mathcal{A}, D_\eta} \leq \eta \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2.$$

Proof. Consider an arbitrary unit vector $\varphi = \sum_x \alpha_x |x\rangle$ with $\sum_x |\alpha_x|^2 = 1$. Again, we just need to show that $\mathbb{E}_{f \leftarrow D_\eta} (\|\pi_f \varphi\|^2) \leq \eta$. Similarly as before,

$$\mathbb{E}_{f \leftarrow D_\eta} (\|\pi_f \varphi\|^2) = \sum_x |\alpha_x|^2 \cdot \Pr_{f \leftarrow D_\eta} [f(x) = 1] = \eta.$$

Alternatively, using the characterization of ν_D (Eq. 2), we can derive this result directly by noting that every position is marked independently with probability η . Hence $\nu_D = \max_i w_i = \eta$. \square

Note that when $\tau_c = 0$, this bound reproduces the complexity of Bernoulli Search using fully quantum queries (cf. [HRS16, ARU14]).

2.4 Designing Hybrid Search Algorithms

In the remaining of this section we propose a hybrid algorithm for the Dist-Search problem, analyze its success probability and show that in several relevant cases, the algorithm is optimal, hence leading to tight query complexity in the hybrid search model.

As a first step, we next describe a quantum search algorithm that, by adapting Grover's algorithm, takes into account a given distribution D .

2.4.1 Quantum Search Algorithm on D A main distinction from standard Grover is that the amplitudes in our initial state are proportional to the weights ω_i (capturing the likelihood that x_i is a solution under D), rather than a uniform superposition.

Quantum Search Algorithm \mathcal{A} for an Arbitrary Distribution D

Given: $x \in \{0, 1\}^m$ drawn from D .

Goal: Find $i \in [m]$ such that $x_i = 1$ making τ_q quantum queries to x .

Initialization: \mathcal{A} constructs a unitary U_D such that

$$|\phi_0\rangle \stackrel{\text{def}}{=} U_D |0\rangle = \frac{1}{\sqrt{\omega}} \sum_i \sqrt{\omega_i} |i\rangle.$$

Modified Grover iteration: Repeatedly apply $G := R_0 R_x$, where

$$R_0 \stackrel{\text{def}}{=} -(\mathbb{1} - 2|\phi_0\rangle\langle\phi_0|),$$

$$R_x \stackrel{\text{def}}{=} \sum_i (-1)^{x_i} |i\rangle\langle i|.$$

Output: Measure the state in the computational basis and output the measurement outcome i .

Note that once U_D is available, $R_0 = -U_D(\mathbf{1} - |0\rangle\langle 0|)U_D^\dagger$ can be readily implemented, and one application of R_x can be realized by *one* query to x .

For any fixed x , we let ε_x denote the probability that \mathcal{A} finds a solution (i.e., some i with $x_i = 1$); thus, $\varepsilon = \mathbb{E}_{x \leftarrow D}(\varepsilon_x)$ represents the success probability of \mathcal{A} averaged over the distribution D . Next, we turn to lower-bounding this success probability; the proof is deferred to Sect. 3.2.

Theorem 5. *Algorithm \mathcal{A} with τ_q quantum queries finds an i with $x_i = 1$ with probability:*

$$\varepsilon \geq \tau_q^2 \cdot \frac{\sum_i \omega_i^2}{\omega}.$$

2.4.2 A Hybrid Algorithm for Distributional Search

We are now ready to describe a hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries. The basic idea is as follows: Given distribution D , let $S = \{i_1, \dots, i_{\tau_c}\} \subseteq [m]$ be the set of indices with the τ_c largest values of ω_i . (In case of ties, we break them arbitrarily.) Our algorithm will first issue the τ_c classical queries on S to verify whether there exists an index $i \in S$ such that $x_i = 1$; if not, it will run the quantum search algorithm \mathcal{A} from before, but on the reduced search space $[m] - S$.

In order to run the quantum algorithm in a modular fashion, we define an induced distribution \tilde{D} on $\{0, 1\}^{m-\tau_c}$. We will denote by x_T the substring of x of size $|T|$ obtained from concatenating the bits x_i for all $i \in T$, and by \tilde{S} the set defined as $\tilde{S} \stackrel{\text{def}}{=} [m] - S$.

To define \tilde{D} , we first define $d \stackrel{\text{def}}{=} \sum_{x \in \{0,1\}^m: x_S=0} d_x$. Then for each $\mathbf{x} \in \{0, 1\}^{m-\tau_c}$, we define $\tilde{d}_{\mathbf{x}} \stackrel{\text{def}}{=} \frac{d_{\tilde{x}}}{d}$, where \tilde{x} is the unique string with $x_S = 0$ and $x_{\tilde{S}} = \mathbf{x}$. Note that there is a fixed mapping that matches every index $\mathbf{i} \in \tilde{S}$ with an index $i \in [m]$ such that $\mathbf{x}_{\mathbf{i}} = 1$ if and only if $x_i = 1$. We assume that this mapping is performed implicitly whenever necessary. Therefore, for every $\mathbf{i} \in \tilde{S}$, we can write the weight under \tilde{D} as:

$$\tilde{\omega}_{\mathbf{i}} = \sum_{\mathbf{x} \in \{0,1\}^{m-\tau_c}} \tilde{d}_{\mathbf{x}} \cdot \mathbf{x}_{\mathbf{i}} = \frac{\sum_{x: x_S=0} d_x \cdot x_{\mathbf{i}}}{\sum_{x: x_S=0} d_x}.$$

Our hybrid algorithm can now be described as follows.

Hybrid Search Algorithm \mathcal{A}_h for an Arbitrary Distribution D

Given: $x \in \{0, 1\}^m$ drawn from D .

Goal: Find $i \in [m]$ such that $x_i = 1$ by making τ_c classical queries τ_q and quantum queries to x .

Classical Stage. \mathcal{A} makes classical queries for each $i \in S$, where S , defined as above, consists of the indices with the τ_c largest ω_i . If some $x_i = 1$, output i and exit; otherwise, continue.

Quantum Stage. Run the quantum algorithm \mathcal{A} on induced distribution \tilde{D} .

The algorithm's success probability can be split into analyzing the classical and quantum stages separately, as we show below. First, we define the following binary random variables:

- $Z_c^x = 1$ if and only if $x_i = 1$ for some $i \in S$ (i.e., the classical stage succeeds);
- $Z_q^x = 1$ if and only if the quantum stage is successful.

Lemma 1. *For any distribution D , the probability that hybrid algorithm \mathcal{A}_h succeeds is:*

$$\Pr[\text{Hybrid Success}] \geq \frac{1}{2} (\mathbb{E}_{x \leftarrow D}(Z_c^x) + \mathbb{E}_{x \leftarrow D}(Z_q^x)) .$$

Proof. The algorithm fails if both classical and quantum stages fail. Hence the failure probability is

$$\mathbb{E}_{x \leftarrow D}((1 - Z_c^x)(1 - Z_q^x)) = 1 - \mathbb{E}_{x \leftarrow D}(Z_c^x) - \mathbb{E}_{x \leftarrow D}(Z_q^x) + \mathbb{E}_{x \leftarrow D}(Z_c^x \cdot Z_q^x) .$$

Then, by using the Cauchy-Schwartz inequality (Lemma 5), and as Z_c^x and Z_q^x are both binary variables, we have

$$\begin{aligned} \mathbb{E}_{x \leftarrow D}(Z_c^x \cdot Z_q^x) &\leq \sqrt{\mathbb{E}_{x \leftarrow D}(Z_c^x) \cdot \mathbb{E}_{x \leftarrow D}(Z_q^x)} \\ &\leq \frac{1}{2}(\mathbb{E}_{x \leftarrow D}(Z_c^x) + \mathbb{E}_{x \leftarrow D}(Z_q^x)) . \end{aligned}$$

We can then conclude that the algorithm's success probability is

$$\Pr[\text{Hybrid Success}] = 1 - \mathbb{E}_{x \leftarrow D}((1 - Z_c^x)(1 - Z_q^x)) \geq \frac{1}{2} (\mathbb{E}_{x \leftarrow D}(Z_c^x) + \mathbb{E}_{x \leftarrow D}(Z_q^x)) .$$

□

Applying Theorem 5, we can immediately give an expression for the quantum success probability. Namely:

$$\mathbb{E}_{x \leftarrow D}(Z_q^x) \geq \tau_q^2 \frac{\sum_{i \in \bar{S}} \tilde{\omega}_i^2}{\sum_{i \in \bar{S}} \tilde{\omega}_i} .$$

2.4.3 Success Probability for Special Distributions

We now show that for some special cases the hybrid algorithm above is optimal. We note that in these cases, the quantum stage actually coincides with the standard Grover search, and thus the quantum success probability can be obtained by the known result. Our analysis can be viewed as an alternative approach following the general result expressed by Theorem 5.

When $x \leftarrow D$ assigns a single i with $x_i = 1$ uniformly at random, \tilde{D} can be seen as the same distribution but restricting to x with $x_S = 0$. For all $\mathbf{i} \in \tilde{S}$, we have $\tilde{\omega}_{\mathbf{i}} = \frac{1}{m-c}$, and hence:

$$\mathbb{E}_{x \leftarrow D}(Z_q^x) = \tau_q^2 \cdot \frac{\sum_{\mathbf{i} \in \tilde{S}} \tilde{\omega}_{\mathbf{i}}^2}{\sum_{\mathbf{i} \in \tilde{S}} \tilde{\omega}_{\mathbf{i}}} = \tau_q^2 \frac{1}{m-c}.$$

It is also easy to observe that $\mathbb{E}_{x \leftarrow D}(Z_c^x) = \tau_c \frac{1}{m}$.

Lemma 2 (Uniform Search Hybrid Success and Optimality). *When D is the uniform distribution, our hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds with probability at least*

$$\Pr[\text{Hybrid Success}] \geq \frac{1}{2} \left(\frac{\tau_c}{m} + \frac{\tau_q^2}{m - \tau_c} \right).$$

Except for constant factors and lower-order terms, this matches the hardness bound shown in Theorem 3, and hence the hybrid query complexity for the uniform distribution is $\Theta\left(\frac{1}{m}(\tau_c + \tau_q^2)\right)$.

Similarly, we can obtain a tight bound for the Bernoulli distribution, by the observation that \tilde{D} in this case is just another Bernoulli distribution with the same η . Hence,

$$\mathbb{E}_{x \leftarrow D}(Z_q^x) = \eta \cdot \tau_q^2.$$

On the other hand,

$$\mathbb{E}_{x \leftarrow D}(Z_c^x) = 1 - (1 - \eta)^{\tau_c} \geq \frac{1}{2} \eta \cdot \tau_c.$$

Lemma 3 (Bernoulli Search Hybrid Success and Optimality). *When D is the Bernoulli distribution, our hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds with probability at least*

$$\Pr[\text{Hybrid Success}] \geq \frac{1}{2} \eta \left(\frac{1}{2} \tau_c + \tau_q^2 \right).$$

Again, except for constant factors and lower-order terms, this matches the hardness bound shown in Theorem 4, and hence the hybrid query complexity for the Bernoulli distribution is $\Theta(\eta(\tau_c + \tau_q^2))$.

3 Proofs of the Main Results

3.1 Hardness of Dist-Search

In this section we will prove the main hardness result stated in Theorem 1. For convenience, we restate it again here:

Theorem 1 (Hardness of Dist-Search – fixed query order). *For any algorithm \mathcal{A} making up to τ_c classical queries and τ_q quantum queries (with a fixed order of the queries independent of f), it holds that \mathcal{A} solves the Dist-Search problem with probability:*

$$\text{Succ}_{\mathcal{A},D} := \Pr_{f \leftarrow D} [f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2.$$

3.1.1 Preliminaries and Overview

We first formally describe an oracle function for the case of quantum and pseudo-classical queries.

Definition 2 (Query Operators). *We define the following operators, describing the actions of quantum and pseudo-classical oracles for a hybrid algorithm given a boolean function f .*

– A pseudo-classical oracle is described by

$$P_{f,b} \stackrel{\text{def}}{=} \sum_{x:f(x)=b} |x\rangle\langle x| \otimes \mathbb{1} \otimes |b\rangle$$

– A quantum oracle is described by

$$Q_f \stackrel{\text{def}}{=} \sum_{x,b} |x\rangle\langle x| \otimes \mathbb{1} \otimes |b \oplus f(x)\rangle\langle b|$$

We denote $\Pi_f \stackrel{\text{def}}{=} \pi_f \otimes \mathbb{1}$ ($\mathbb{1}$ operates on the output and ancilla registers) and $\Pi_f^\perp \stackrel{\text{def}}{=} \mathbb{1} - \Pi_f$ ($\mathbb{1}$ operates on the entire system). Then on a pseudo-classical query, the two operators $P_{f,0} = \Pi_f^\perp \otimes |0\rangle$ and $P_{f,1} = \Pi_f \otimes |1\rangle$ correspond to the two possible measurement outcomes. It is more convenient to answer quantum queries by the corresponding phase oracle:

$$Q_f \stackrel{\text{def}}{=} \mathbb{1} - 2\Pi_f.$$

This can be seen as setting the output register of the standard oracle in $|-\rangle$, and as a result, a quantum query flips the signs of the 1-preimages.

When running a hybrid query algorithm with f , we will keep track of the (sub-normalized) pure state $\psi_f^{(t)}$, which denotes the state of the algorithm on input f after t queries in the situation where every pseudo-classical query measures 0 (we will call this the 0-branch of \mathcal{A}^f). Namely, consider

an arbitrary algorithm with at most τ queries (τ_q quantum and τ_c pseudo-classical) specified by a sequence of unitary operators² ($U^{(0)}, U^{(1)}, \dots, U^{(\tau)}$). Let $T_c = \{t : t\text{-th query is pseudo-classical}\}$ and $T_q = \{t : t\text{-th query is quantum}\}$. Then $\psi_f^{(t)}$ is defined recursively by

$$\psi_f^{(t)} \stackrel{\text{def}}{=} \begin{cases} U^{(t)} P_{f,0} \psi_f^{(t-1)}, & \text{if } t \in T_c; \\ U^{(t)} Q_f \psi_f^{(t-1)} & \text{if } t \in T_q. \end{cases} \quad (5)$$

From this definition, the projection of $\psi_f^{(t)}$ under Π_f^\perp characterizes the event that an algorithm fails to find a 1-preimage.

Lemma 4. *For any algorithm \mathcal{A} , the failure probability of finding a 1-preimage of f after t queries is*

$$\delta_f^{(t)} = \Pr[f(x) \neq 1 : x \leftarrow \mathcal{A}^f] \geq \left\| \Pi_f^\perp \psi_f^{(t)} \right\|^2.$$

Hence, the failure probability with respect to distribution D satisfies

$$\delta_D^{(t)} = \mathbb{E}_{f \leftarrow D} \delta_f^{(t)} \geq \mathbb{E}_{f \leftarrow D} \left\| \Pi_f^\perp \psi_f^{(t)} \right\|^2.$$

Thus, our goal becomes lower-bounding $\left\| \Pi_f^\perp \psi_f^{(t)} \right\|$. To do this, we consider running the same algorithm, but with a null function:

$$f_\emptyset : x \mapsto 0, \forall x \in X.$$

In this case, a quantum query is equivalent to applying identity (denoted $Q_\emptyset \stackrel{\text{def}}{=} \mathbb{1}$), and a pseudo-classical query does not tamper the input state either, but just appends $|0\rangle$. To be precise, we define

$$P_{\emptyset,0} \stackrel{\text{def}}{=} \mathbb{1} \otimes |0\rangle,$$

and at each step $t \geq 0$, the state of the algorithm denoted by $\phi^{(t)}$ can be described as:

$$\phi^{(t)} = \begin{cases} U^{(t)} P_{\emptyset,0} \phi^{(t-1)}, & \text{if } t \in T_c; \\ U^{(t)} \phi^{(t-1)} & \text{if } t \in T_q. \end{cases}$$

Without loss of generality we assume initially $\psi_f^{(0)} = \phi^{(0)} = |0\rangle$, and hence $\left\| \Pi_f^\perp \psi_f^{(0)} \right\| = \left\| \Pi_f^\perp \phi^{(0)} \right\| = 1$. In order to succeed, algorithm \mathcal{A}^f needs to move $\psi_f^{(t)}$ away from the kernel of Π_f^\perp or reduce its norm. This motivates defining the progress measures below.

² Dimensions may grow depending on the arrangement of the pseudo-classical queries.

Table 1. Summary of variables and quantities used in our Dist-Search analysis.

π_f	$\sum_{x:f(x)=1} x\rangle \langle x $
Π_f	$\pi_f \otimes \mathbb{1}$ ($\mathbb{1}$ on ancilla registers)
δ_f	$\Pr[f(x) \neq 1 : x \leftarrow \mathcal{A}^f]$ (Failure probability with fixed f)
δ_D	$\mathbb{E}_D(\delta_f)$ (Failure probability with $f \leftarrow D$)
$\phi^{(0)} = \psi^{(0)}$	Initial state
$\phi^{(t)}$	State after t -th query in $\mathcal{A}^{f_\emptyset}$
$\psi_f^{(t)}$	State on the 0-branch after t -th query in \mathcal{A}^f
Q_f	$\mathbb{1} - 2\Pi_f$ (quantum oracle of f)
Q_\emptyset	$\mathbb{1}$ (quantum oracle of f_\emptyset)
$P_{f,0}$	$\Pi_f^\perp \otimes 0\rangle$ (pseudo-classical oracle of f)
$P_{f,1}$	$\Pi_f \otimes 1\rangle$ (pseudo-classical oracle of f)
$P_{\emptyset,0}$	$\mathbb{1} \otimes 0\rangle$ (pseudo-classical oracle of f_\emptyset)
$\gamma_f^{(t)}$	$\ \Pi_f \phi^{(t)}\ ^2$
$\gamma^{(t)}$	$\mathbb{E}_D(\gamma_f^{(t)})$

Definition 3 (Progress Measures). For any function f and $t \geq 0$, define

$$A_f^{(t)} \stackrel{\text{def}}{=} \left| \langle \phi^{(t)}, \psi_f^{(t)} \rangle \right|^2, \quad B_f^{(t)} \stackrel{\text{def}}{=} \left\| \psi_f^{(t)} \right\|^2 - \left| \langle \phi^{(t)}, \psi_f^{(t)} \rangle \right|^2.$$

Given a distribution D on \mathcal{F} , define the expected progress measures by

$$A_D^{(t)} \stackrel{\text{def}}{=} \mathbb{E}_{f \leftarrow D} \left(A_f^{(t)} \right), \quad B_D^{(t)} \stackrel{\text{def}}{=} \mathbb{E}_{f \leftarrow D} \left(B_f^{(t)} \right).$$

Notice that:

$$A_f^{(t)} + B_f^{(t)} = \left\| \psi_f^{(t)} \right\|^2, \quad A_f^{(0)} = 1, \quad B_f^{(0)} = 0.$$

We will show that $A_D^{(t)} - B_D^{(t)}$ essentially lower bounds the failure probability $\delta_D^{(t)}$ (Lemma 8). Hence, an algorithm's objective would be to *reduce* $A_D^{(t)}$ and *increase* $B_D^{(t)}$. However, we can limit how much change can occur after τ queries (Proposition 1). This is by carefully analyzing the effect of each quantum or pseudo-classical query (Lemmas 10 and 11). Roughly speaking,

- A quantum query reduces $A_D^{(t)}$ by at most $4\sqrt{\nu_D \cdot B_D^{(t)}}$ and increases $B_D^{(t)}$ by the same amount (as a quantum query does not affect $\left\| \psi_f^{(t)} \right\|^2$), and
- A pseudo-classical query increases $B_D^{(t)}$ by at most ν_D , while a part $z^{(t)}$ of $B_D^{(t)}$ can also be spent to decrease $A_D^{(t)}$ by $\sqrt{\nu_D \cdot z^{(t)}}$ (Table 1).

3.1.2 Proof of Theorem 1 First off, we state the Cauchy-Schwarz inequality for random variables and derive a corollary that is useful in several places.

Lemma 5 (Cauchy-Schwarz). *For any random variables X, Y , it holds that: $|\mathbb{E}(XY)|^2 \leq \mathbb{E}(X^2) \cdot \mathbb{E}(Y^2)$.*

Corollary 1. *Let Z be a discrete random variable, and $g(Z)$ and $h(Z)$ be two non-negative functions. Then it holds that: $\mathbb{E}_Z(\sqrt{g(Z) \cdot h(Z)}) \leq \sqrt{\mathbb{E}_Z(g(Z)) \cdot \mathbb{E}_Z(h(Z))}$.*

It will be helpful to consider a two-dimensional plane in our analysis, which we now define explicitly.

Definition 4 (Useful 2-D Plane). *For $t \geq 0$, let*

$$\phi_f^{(t)} \stackrel{\text{def}}{=} \frac{\Pi_f \phi^{(t)}}{\|\Pi_f \phi^{(t)}\|} = \Pi_f \phi^{(t)} / \sqrt{\gamma_f^{(t)}}, \quad \phi_f^{(t)\perp} \stackrel{\text{def}}{=} \frac{\Pi_f^\perp \phi^{(t)}}{\|\Pi_f^\perp \phi^{(t)}\|} = \Pi_f^\perp \phi^{(t)} / \sqrt{1 - \gamma_f^{(t)}}$$

be the normalized vectors resulting of projecting $\phi^{(t)}$ on the orthogonal subspaces spanned by 1 and 0 preimages of f , respectively, and let $\Phi^{(t)}$ be the 2-dimensional plane spanned by $\{\phi_f^{(t)}, \phi_f^{(t)\perp}\}$. Then $\phi^{(t)\perp}$ is identified as the normalized state perpendicular to $\phi^{(t)}$ in $\Phi^{(t)}$, i.e.,

$$\phi^{(t)\perp} \stackrel{\text{def}}{=} \phi_f^{(t)} \sqrt{1 - \gamma_f^{(t)}} - \phi_f^{(t)\perp} \sqrt{\gamma_f^{(t)}}.$$

It is useful to decompose $\psi_f^{(t)}$ with respect to $\Phi^{(t)}$:

Lemma 6 (Decomposition of $\psi_f^{(t)}$ wrt $\Phi^{(t)}$). *Let a and b be projecting $\psi_f^{(t)}$ on the plane $\Phi^{(t)}$ and then decomposing it under basis $\{\phi^{(t)}, \phi^{(t)\perp}\}$, and let c be the remaining component of $\psi_f^{(t)}$ orthogonal to $\Phi^{(t)}$, i.e., $c \perp \Phi^{(t)}$. Then $\psi_f^{(t)}$ can be expressed as $\psi_f^{(t)} = a + b + c$ with*

$$a = \phi^{(t)} \sqrt{A_f^{(t)}}, \quad b = \omega \sqrt{B_f^{(t)} - \|c\|^2} \cdot \phi^{(t)\perp},$$

where ω is a complex phase ($|\omega| = 1$) of the vector $\psi_f^{(t)} - \langle \psi_f^{(t)}, \phi_f^{(t)} \rangle \cdot \phi^{(t)} - c$. Thus,

$$\Pi_f^\perp \psi_f^{(t)} = \phi_f^{(t)\perp} \left(\sqrt{1 - \gamma_f^{(t)}} \sqrt{A_f^{(t)}} - \sqrt{\gamma_f^{(t)}} \cdot \omega \sqrt{B_f^{(t)} - \|c\|^2} \right) + c_f^\perp,$$

with $c_f^\perp := \Pi_f^\perp c$.

Intuitively, for the next result, the goal is to relate the failure probability with the progress measures A and B . To do so, we will first relate the failure probability with the norm of the non-solution component. By decomposing this norm in terms of the two progress measures A and B and an orthogonal component which can be removed, we can determine a lower bound on the failure probability as a function of the two progress measure after each performed query.

Lemma 7. For any fixed f and $t \geq 0$,

$$\delta_f^{(t)} \geq A_f^{(t)} - \gamma_f^{(t)} - 2\sqrt{\gamma_f^{(t)} \cdot B_f^{(t)}}.$$

Proof. For convenience, we omit writing the superscript (t) in this proof. We first show that $\|\pi_f^\perp \psi_f\| \geq \sqrt{(1 - \gamma_f)A_f} - \sqrt{\gamma_f B_f}$. By Lemma 6, we have that

$$\Pi_f^\perp \psi_f = \phi_f^\perp \left(\sqrt{1 - \gamma_f} \sqrt{A_f} - \sqrt{\gamma_f} \cdot \omega \sqrt{B_f - \|c\|^2} \right) + c_f^\perp,$$

with $c_f^\perp := \pi_f^\perp c$. Since $c \perp \Phi$, it follows that

$$\langle \phi_f^\perp, c_f^\perp \rangle = \langle \phi_f^\perp, \Pi_f^\perp c \rangle = \langle \Pi_f^\perp \phi_f^\perp, c \rangle = \langle \phi_f^\perp, c \rangle = 0.$$

We can then obtain:

$$\|\Pi_f^\perp \psi_f\| = \left| \sqrt{1 - \gamma_f} \cdot \sqrt{A_f} - \sqrt{\gamma_f} \cdot \omega \sqrt{B_f - \|c\|^2} \right| + \|c_f^\perp\|$$

Hence by choosing $c = 0, \omega = 1$, we get: $\|\Pi_f^\perp \psi_f\| \geq \sqrt{(1 - \gamma_f)A_f} - \sqrt{\gamma_f B_f}$. Therefore we can lower bound the failure probability:

$$\begin{aligned} \delta_f &\geq \|\pi_f^\perp \psi_f\|^2 \geq (1 - \gamma_f)A_f - 2\sqrt{(1 - \gamma_f)\gamma_f B_f} \\ &\geq A_f - \gamma_f - 2\sqrt{\gamma_f B_f} \quad (A_f, \gamma_f \leq 1) \end{aligned}$$

□

Taking the expectation over D , we can express the failure probability with respect to the distribution.

Lemma 8. For any distribution D and $t \geq 0$,

$$\delta_D^{(t)} \geq A^{(t)} - \gamma^{(t)} - 2\sqrt{\gamma^{(t)} \cdot B^{(t)}}.$$

Proof.

$$\begin{aligned} \delta_D^{(t)} &= \mathbb{E}_{f \leftarrow D}(\delta_f^{(t)}) \\ &\geq \mathbb{E}_D(A_f^{(t)}) - \mathbb{E}_D(\gamma_f^{(t)}) - 2\mathbb{E}_D\left(\sqrt{\gamma_f^{(t)} \cdot B_f^{(t)}}\right) \quad (\text{Linearity of expectation}) \\ &\geq A^{(t)} - \gamma^{(t)} - 2\sqrt{\mathbb{E}_D(\gamma_f^{(t)}) \cdot \mathbb{E}_D(B_f^{(t)})} \quad (\text{Corollary 1}) \\ &= A^{(t)} - \gamma^{(t)} - 2\sqrt{\gamma^{(t)} \cdot B^{(t)}} \end{aligned}$$

□

We can also relate $\gamma^{(t)}$ to the value ν_D determined by the distribution D :

Lemma 9. For any $t \geq 0$ and any distribution D , we have: $\gamma^{(t)} \leq \nu_D$.

Proof.

$$\gamma^{(t)} := \mathbb{E}_{f \leftarrow D} \left(\left\| \Pi_f \phi^{(t)} \right\|^2 \right) = \mathbb{E}_{f \leftarrow D} \left(\left\| (\pi_f \otimes \mathbf{1}) \phi^{(t)} \right\|^2 \right)$$

We write $\phi^{(t)} = \sum_i \alpha_i |u_i\rangle \otimes |v_i\rangle$ under the Schmidt decomposition, where $\alpha_i \geq 0$ such that $\sum_i \alpha_i^2 = 1$ are the Schmidt coefficients, and $\{|u_i\rangle\}$ are orthonormal states on the system of the input register and $\{|v_i\rangle\}$ are orthonormal states on the system of output and ancilla registers. Then we can rewrite $\gamma^{(t)}$ as:

$$\begin{aligned} \gamma^{(t)} &:= \mathbb{E}_{f \leftarrow D} \left(\left\| (\pi_f \otimes \mathbf{1}) \phi^{(t)} \right\|^2 \right) = \mathbb{E}_{f \leftarrow D} \left(\left\| (\pi_f \otimes \mathbf{1}) \left(\sum_i \alpha_i |u_i\rangle \otimes |v_i\rangle \right) \right\|^2 \right) \\ &= \mathbb{E}_{f \leftarrow D} \left(\left\| \sum_i \alpha_i (\pi_f |u_i\rangle) \otimes |v_i\rangle \right\|^2 \right) \\ &= \mathbb{E}_{f \leftarrow D} \left(\sum_i \alpha_i^2 \left\| (\pi_f |u_i\rangle) \otimes |v_i\rangle \right\|^2 \right) \quad (|v_i\rangle \text{ are orthogonal}) \\ &= \mathbb{E}_{f \leftarrow D} \left(\sum_i \alpha_i^2 \left\| \pi_f |u_i\rangle \right\|^2 \cdot \left\| |v_i\rangle \right\|^2 \right) \quad (\|a \otimes b\| = \|a\| \cdot \|b\|) \\ &= \mathbb{E}_{f \leftarrow D} \left(\sum_i \alpha_i^2 \left\| \pi_f |u_i\rangle \right\|^2 \right) = \sum_i \alpha_i^2 \cdot \mathbb{E}_{f \leftarrow D} \left(\left\| \pi_f |u_i\rangle \right\|^2 \right) \\ &\leq \sum_i \alpha_i^2 \nu_D \quad (\text{definition of } \nu_D) \\ &= \nu_D \sum_i \alpha_i^2 = \nu_D \end{aligned}$$

Proposition 1 (Bounding Progress Measures). After $\tau = \tau_c + \tau_q$ queries,

$$A^{(\tau)} \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2, \quad B^{(\tau)} \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2.$$

Proving Proposition 1 is the most involved step technically speaking. We present the details separately in Sect. 3.1.3 and here we apply it to prove Theorem 1.

Proof of Theorem 1. Assuming the bounds above on the two progress measures, we obtain that:

$$\begin{aligned} \delta^{(\tau)} &\geq 1 - 4\gamma^{(\tau)} \cdot (\sqrt{\tau_c} + \tau_q)^2 - \gamma^{(\tau)} - 2\gamma^{(t)} \cdot (\sqrt{\tau_c} + 2\tau_q) \quad (\text{Proposition 1}) \\ &= 1 - \gamma^{(\tau)} \cdot (4(\sqrt{\tau_c} + \tau_q) + 2\sqrt{\tau_c} + 4\tau_q + 1) \\ &\geq 1 - \gamma^{(\tau)} \cdot (2(\sqrt{\tau_c} + \tau_q) + 1)^2 \quad (\tau_c \geq 0) \\ &\geq 1 - \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 \quad (\gamma^{(\tau)} \leq \nu_D \text{ Lemma 9}) \end{aligned}$$

Therefore,

$$\text{Succ}_{\mathcal{A},D} \leq 1 - \delta^{(\tau)} \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2.$$

□

3.1.3 Bounding the Progress Measures (Proposition 1)

We repeat the proposition statement for convenience here:

Proposition 1 (Bounding the Progress Measures). *After $\tau = \tau_c + \tau_q$ queries,*

$$A^{(\tau)} \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2, \quad B^{(\tau)} \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2.$$

Firstly, we will consider a fixed function f , and bound how much each query can possibly reduce $A_f^{(t)}$ and increase $B_f^{(t)}$.

Lemma 10 (Progress Measures for a Fixed Function). *For every t the progress measures after the $t+1$ -th query satisfy the following recurrent relations:*

- *If the $t+1$ -th query is pseudo-classical, then there exists a sequence $\left(z_f^{(t)}\right)_{t \geq 0}$, satisfying $0 \leq z_f^t \leq B_f^{(t)}$, such that:*

$$\begin{aligned} A_f^{(t+1)} &\geq A_f^{(t)} - 2\gamma_f^{(t)} - 2 \cdot \sqrt{z_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \\ B_f^{(t+1)} &\leq B_f^{(t)} + \gamma_f^{(t)} - z_f^{(t)} \end{aligned} \tag{6}$$

- *If the $t+1$ -th query is quantum, then:*

$$\begin{aligned} A_f^{(t+1)} &\geq A_f^{(t)} - 4\gamma_f^{(t)} - 4 \cdot \sqrt{B_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \\ B_f^{(t+1)} &\leq B_f^{(t)} + 4\gamma_f^{(t)} + 4 \cdot \sqrt{B_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \end{aligned} \tag{7}$$

Proof. The proof can be found in the full version of the paper [CGS23]. □

Lemma 11 (Progress Measures for Dist-Search). *For every t , the progress measures after the $t+1$ -th query satisfy the following recurrent relations:*

- *If the $t+1$ -th query is pseudo-classical, there exists $z_t \in [0, B^{(t)}]$ such that:*

$$\begin{aligned} A^{(t+1)} &\geq A^{(t)} - 2\nu_D - 2\sqrt{\nu_D} \cdot \sqrt{z_t} \\ B^{(t+1)} &\leq B^{(t)} - z_t + \nu_D \end{aligned} \tag{8}$$

- *If the $t+1$ -th query is quantum, then we have:*

$$\begin{aligned} A^{(t+1)} &\geq A^{(t)} - 4 \cdot \nu_D - 4 \cdot \sqrt{\nu_D} \cdot \sqrt{B^{(t)}} \\ B^{(t+1)} &\leq B^{(t)} + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{B^{(t)}} \end{aligned} \tag{9}$$

Proof. Letting $z_t \stackrel{\text{def}}{=} \mathbb{E}_{f \leftarrow D}(z_f^t)$, we can observe that $z_t \in [0, B^{(t)}]$. Taking expectations over D , and applying Corollary 1 ($\mathbb{E}(\sqrt{g(Z)} \cdot h(Z)) \leq \sqrt{\mathbb{E}(g(Z)) \cdot \mathbb{E}(h(Z))}$) and Lemma 9 ($\gamma^{(t)} \leq \nu_D$), the relations for $A^{(t)}$ and $B^{(t)}$ follow. \square

Next, since we intend to lower bound $A^{(\tau)}$ and upper bound $B^{(\tau)}$, we can change the inequalities to equalities and analyze instead the new sequences (a_t, b_t) defined below. It is clear that $A^{(\tau)} \geq a_\tau$ and $B^{(\tau)} \leq b_\tau$.

Definition 5 (Sequences $(a_t)_{t \geq 0}, (b_t)_{t \geq 0}$). We define the following sequences based on the evolution of the progress measures A and B :

$$\begin{aligned} a_0 &\stackrel{\text{def}}{=} A^{(0)} = 1 \quad ; \quad b_0 \stackrel{\text{def}}{=} B^{(0)} = 0 \\ a_{t+1} &\stackrel{\text{def}}{=} \begin{cases} a_t - 2 \cdot \nu_D - 2 \cdot \sqrt{\nu_D} \cdot \sqrt{z_t}, & \text{if } t+1 \in T_c \\ a_t - 4 \cdot \nu_D - 4 \cdot \sqrt{\nu_D} \cdot \sqrt{b_t}, & \text{if } t+1 \in T_q \end{cases} \\ b_{t+1} &\stackrel{\text{def}}{=} \begin{cases} b_t + \nu_D - z_t, & \text{if } t+1 \in T_c \\ b_t + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{b_t}, & \text{if } t+1 \in T_q \end{cases} \end{aligned}$$

where $(z_t)_{t \geq 1}$ is the sequence defined in the proof of Lemma 11, which satisfies $0 \leq z_t \leq B^{(t)}$ for any t .

Lemma 12 (Bounding a_τ and b_τ).

$$a_\tau \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2, \quad b_\tau \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2. \quad (10)$$

Proof. The proof consists of four steps.

(1) First we show that $b_\tau \leq (\sqrt{\tau_c} + 2\tau_q)^2 \cdot \nu_D$.

To get an upper bound for each term of this sequence, we can let $z_t = 0$ and instead consider the sequence:

$$d_{t+1} \stackrel{\text{def}}{=} \begin{cases} d_t + \nu_D, & \text{if } t+1 \in T_c \\ d_t + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t}, & \text{if } t+1 \in T_q \end{cases}$$

As a result we have: $b_t \leq d_t$ for any $t \in [\tau]$.

Our task is to bound the last term d_τ in the sequence. Every hybrid strategy A that uses τ_c classical queries and τ_q quantum queries can be expressed by $A = [x_1, \dots, x_\tau]$, where if $x_i = 0$ (resp. $x_i = 1$) indicates that the i -th query of A is classical (resp. quantum), and there are exactly τ_c values of 0 and τ_q values of 1. Therefore, the sequence $(d_t)_t$ parameterized by the strategy A , denoted as $(d_t^A)_t$, can be re-written as:

$$d_{t+1}^A \stackrel{\text{def}}{=} \begin{cases} d_t^A + \nu_D, & \text{if } x_{t+1} = 0 \\ d_t^A + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t^A}, & \text{if } x_{t+1} = 1 \end{cases} \quad (11)$$

Our task then becomes determining the strategy A^* which achieves the maximum $d_\tau^{A^*}$. We claim that

$$A^* \stackrel{\text{def}}{=} [0, \dots, 0, 1, \dots, 1],$$

namely the strategy of making all classical queries upfront is optimal. This follows from a greedy argument.

Consider two arbitrary strategies $A = [x_1, \dots, x_i, x_{i+1}, \dots, x_\tau]$ and $B = [y_1, \dots, y_i, y_{i+1}, \dots, y_\tau]$ which only differ in the i and $i+1$ -th queries. Namely, $x_i = 0, x_{i+1} = 1$ and $y_i = 1, y_{i+1} = 0$ and $x_j = y_j$ for $j \in \{1, \dots, \tau\} - \{i, i+1\}$. We next show that $d_\tau^A > d_\tau^B$. As $x_1 = y_1, \dots, x_{i-1} = y_{i-1}$, this implies directly that $d_{i-1}^A = d_{i-1}^B$. Then for the i -th and $i+1$ terms of the two sequences we have:

$$\begin{aligned} d_i^A &= d_{i-1}^A + \nu_D & ; & \quad d_{i+1}^A = d_{i-1}^A + 5\nu_D + 4\sqrt{\nu_D} \sqrt{d_{i-1}^A + \nu_D} \\ d_i^B &= d_{i-1}^B + 4\nu_D + 4\sqrt{\nu_D} \sqrt{d_{i-1}^B} & ; & \quad d_{i+1}^B = d_{i-1}^B + 5\nu_D + 4\sqrt{\nu_D} \sqrt{d_{i-1}^B} \end{aligned}$$

Then, as $d_{i-1}^A = d_{i-1}^B$ it is clear that $d_{i+1}^A > d_{i+1}^B$. As $x_j = y_j$ for all $i+2 \leq j \leq \tau$, this also implies that $d_\tau^A > d_\tau^B$.

Denote the following swap operation on strategies. Given as input a strategy $A = [x_1, \dots, x_i, x_{i+1}, \dots, x_\tau]$ the function swap_i outputs a strategy A' :

$$\text{swap}_i(A) = A' \text{ where } A' = [x_1, \dots, x_{i+1}, x_i, \dots, x_\tau]$$

Our previous argument implies that for a strategy A such that $x_i = 0$ and $x_{i+1} = 1$, we have: $d_\tau^A > d_\tau^{\text{swap}_i(A)}$. Therefore, we can see that any strategy $A = [x_1, \dots, x_\tau]$ can be obtained from a sequence of applications of swap_i on A^* .

$$A^* \stackrel{\text{def}}{=} [0, \dots, 0, 1, \dots, 1] \xrightarrow{\text{swap}_{i_1}} \dots \xrightarrow{\text{swap}_{i_k}} A \text{ for some indices } i_1, \dots, i_k.$$

It hence follows that $d_\tau^{A^*} \geq d_\tau^A$, i.e., A^* is the optimal strategy.

Now, let us compute the last term of the optimal strategy, i.e.: $d_\tau^{A^*}$. We can rewrite the sequence d_t as:

$$d_{t+1}^{A^*} = \begin{cases} d_t^{A^*} + \nu_D, & \text{if } 0 \leq t < \tau_c \\ d_t^{A^*} + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t^{A^*}} = \left(\sqrt{d_t^{A^*}} + 2\sqrt{\nu_D} \right)^2, & \text{if } \tau_c \leq t < \tau \end{cases}$$

As $d_0^{A^*} = 0$, it is clear that we have: $d_{\tau_c}^{A^*} = \tau_c \cdot \nu_D$. For $\tau_c \leq t \leq \tau$, we will prove by induction that:

$$d_t^{A^*} = (\sqrt{\tau_c} + 2(t - \tau_c))^2 \cdot \nu_D$$

For the base case $t = \tau_c$, we already showed that $d_{\tau_c}^{A^*} = \tau_c \cdot \nu_D$. For the inductive step, we have that:

$$d_{t+1}^{A^*} = \left(\sqrt{(\sqrt{\tau_c} + 2(t - \tau_c))^2 \cdot \nu_D} + 2\sqrt{\nu_D} \right)^2 = (\sqrt{\tau_c} + 2(t - \tau_c + 1)) \cdot \nu_D$$

which concludes the inductive proof. Hence, by putting things together:

$$b_\tau \leq d_\tau \leq d_\tau^{A^*} = (\sqrt{\tau_c} + 2\tau_q)^2 \cdot \nu_D \quad (12)$$

(2) Secondly, we show that $\sum_{t \in T_q} \sqrt{b_{t-1}} \leq \sqrt{\nu_D} \cdot \tau_q (\sqrt{\tau_c} + \tau_q - 1)$.

As for b_τ , to get an upper bound we let $z_t = 0$ and use the sequence $(d_t^A)_t$. From the definition of the sequence (Eq. 11), it is clear that $(d_t^A)_t$ is a strictly increasing sequence for any strategy A . This also implies that for any strategy A we have:

$$\sum_{t \in T_q} \sqrt{d_{t-1}^A} \leq \sum_{\tau_c \leq t \leq \tau} \sqrt{d_t^A}$$

In other words, $\sum_{t \in T_q} \sqrt{d_{t-1}^A}$ is maximized when the strategy performs first all τ_c classical queries and then the τ_q quantum queries. Hence, the maximum is achieved for the strategy described above by the sequence $(d_t^{A^*})_t$.

Using the previous result in Eq. 12:

$$\sum_{\tau_c \leq t \leq \tau} d_t^{A^*} = \nu_D \cdot \sum_{\tau_c \leq t \leq \tau} (\sqrt{\tau_c} + 2(t - \tau_c))^2$$

This gives us:

$$\begin{aligned} \sum_{t \in T_q} \sqrt{b_{t-1}} &\leq \sum_{\tau_c \leq t \leq \tau} \sqrt{d_t^{A^*}} = \sqrt{\nu_D} \sum_{\tau_c \leq t \leq \tau} \sqrt{\tau_c} + 2(t - \tau_c) \\ &\leq \sqrt{\nu_D} \left(\tau_q (\sqrt{\tau_c} - 2\tau_c) + 2 \sum_{\tau_c \leq t \leq \tau} t \right) \\ &= \sqrt{\nu_D} \tau_q (\sqrt{\tau_c} + \tau_q - 1) \end{aligned}$$

(3) Thirdly, we show that $\sum_{t \in T_c} \sqrt{z_{t-1}} \leq \sqrt{\nu_D} \cdot (\tau_c + 2\sqrt{\tau_c} \tau_q)$.

By definition of the sequence z_t (Definition 5), we know that for $t \in T_c$:

$$\sum_{t \in T_c} z_{t-1} = \nu_D \cdot \tau_c + \sum_{t \in T_c} (b_{t-1} - b_t)$$

Thus it suffices to derive an upper bound on $\sum_{t \in T_c} (b_{t-1} - b_t)$. We rewrite b_τ as:

$$b_\tau = b_0 + \sum_{t=1}^{\tau} (b_t - b_{t-1}) = \sum_{b_t \geq b_{t-1}} (b_t - b_{t-1}) + \sum_{b_t < b_{t-1}} (b_t - b_{t-1})$$

As a result, we have that:

$$\sum_{t \in T_c \wedge b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{b_t < b_{t-1}} (b_{t-1} - b_t) = \sum_{b_t \geq b_{t-1}} (b_t - b_{t-1}) - b_\tau$$

In other words we also have:

$$\sum_{t \in T_c \wedge b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t \in T_c \wedge b_t \geq b_{t-1}} (b_t - b_{t-1}) + \sum_{t \in T_q \wedge b_t \geq b_{t-1}} (b_t - b_{t-1})$$

For $t \in T_q$, from sequence definition (Definition 5), we have that $b_t > b_{t-1}$ and hence:

$$\sum_{t \in T_c \wedge b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t \in T_c \wedge b_t \geq b_{t-1}} (b_t - b_{t-1}) + 4\tau_q \cdot \nu_D + 4\sqrt{\nu_D} \sum_{t \in T_q} \sqrt{b_{t-1}}$$

By applying step (2), we get:

$$\sum_{t \in T_c \wedge b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t \in T_c \wedge b_t \geq b_{t-1}} (b_t - b_{t-1}) + 4\nu_D \tau_q + 4\nu_D \tau_q (\sqrt{\tau_c} + \tau_q - 1)$$

By subtracting the first sum from the right hand side we get:

$$\sum_{t \in T_c} z_{t-1} = \nu_D \cdot \tau_c + \sum_{t \in T_c} (b_{t-1} - b_t) < \nu_D \cdot (\tau_c + 4\tau_q^2 + 4\tau_q \sqrt{\tau_c})$$

Finally, by using the Cauchy-Schwarz inequality:

$$\sum_{t \in T_c} \sqrt{z_{t-1}} \leq \sqrt{\nu_D \cdot (\tau_c + 4\tau_q^2 + 4\tau_q \sqrt{\tau_c})} \cdot \sqrt{\tau_c} \leq \sqrt{\nu_D} \cdot (\tau_c + 2\tau_q \sqrt{\tau_c})$$

(4) In the final step, we show that $a_\tau \geq 1 - 4\nu_D(\sqrt{\tau_c} + \tau_q)^2$.

From the definition of a_t (Definition 5):

$$\begin{aligned} a_\tau &= a_0 + \sum_{t=1}^{\tau} (a_t - a_{t-1}) \\ &= 1 - \sum_{t \in T_c} (2\nu_D + 2\sqrt{\nu_D} \cdot \sqrt{z_{t-1}}) - \sum_{t \in T_q} (4\nu_D + 4\sqrt{\nu_D} \cdot \sqrt{b_{t-1}}) \\ &= 1 - 2\tau_c \nu_D - 4\tau_q \nu_D - 2\sqrt{\nu_D} \sum_{t \in T_c} \sqrt{z_{t-1}} - 4\sqrt{\nu_D} \sum_{t \in T_q} \sqrt{b_{t-1}} \end{aligned}$$

Using the bounds derived in steps (2) and (3), we get :

$$\begin{aligned} a_\tau &\geq 1 - 2\tau_c \nu_D - 4\tau_q \nu_D - 2\nu_D \cdot (\tau_c + 2\sqrt{\tau_c} \tau_q) - 4\nu_D \cdot \tau_q (\sqrt{\tau_c} + \tau_q - 1) \\ &= 1 - 4\nu_D (\sqrt{\tau_c} + \tau_q)^2 \end{aligned}$$

□

3.2 Quantum Algorithm Analysis

In this section we will prove the success probability of our proposed quantum algorithm described in Sect. 2.4.1.

Theorem 5. *Algorithm \mathcal{A} with τ_q quantum queries finds an i with $x_i = 1$ with probability:*

$$\varepsilon \geq \tau_q^2 \cdot \frac{\sum_i \omega_i^2}{\omega}.$$

Proof. We adapt the geometric analysis of standard Grover's algorithm to analyze \mathcal{A} . First for any x , define two states below:

$$|A_x\rangle := \frac{1}{\sqrt{\alpha_x}} \sum_{i:x_i=1} \sqrt{\omega_i} |i\rangle, \quad |B_x\rangle := \frac{1}{\sqrt{\beta_x}} \sum_{i:x_i=0} \sqrt{\omega_i} |i\rangle,$$

with normalization factors

$$\alpha_x := \sum_{i:x_i=1} \omega_i = \sum_i \omega_i x_i, \quad \text{and} \quad \beta_x := \sum_{i:x_i=0} \omega_i = \sum_i \omega_i (1 - x_i).$$

We will focus on the two dimensional plane spanned by $|A_x\rangle$ and $|B_x\rangle$. Observe that $|\phi_0\rangle$ belongs to this plane, and can be decomposed under the basis $\{|A_x\rangle, |B_x\rangle\}$:

$$\begin{aligned} |\phi_0\rangle &:= \sin\theta |A_x\rangle + \cos\theta |B_x\rangle, \text{ where:} \\ \sin^2\theta &= |\langle\phi_0|A_x\rangle|^2 = \frac{1}{\omega \cdot \alpha_x} \left(\sum_i \omega_i x_i\right)^2 = \frac{\alpha_x}{\omega}. \end{aligned}$$

We then show that on the two dimensional plane, R_0 is a reflection about $|\phi_0\rangle$ and R_x is a reflection $|B_x\rangle$. We introduce a state $|\phi_0^\perp\rangle$ on the plane orthogonal to $|\phi_0\rangle$, which can be written as

$$|\phi_0^\perp\rangle = \cos\theta |A_x\rangle - \sin\theta |B_x\rangle.$$

Clearly $\{|\phi_0\rangle, |\phi_0^\perp\rangle\}$ forms another basis on the plane, under which we can express $|A_x\rangle$ and $|B_x\rangle$ as below.

$$|A_x\rangle = \sin\theta |\phi_0\rangle + \cos\theta |\phi_0^\perp\rangle, \quad |B_x\rangle = \cos\theta |\phi_0\rangle - \sin\theta |\phi_0^\perp\rangle.$$

It then becomes easy to verify that

$$R_0 |A_x\rangle = \sin\theta |\phi_0\rangle - \cos\theta |\phi_0^\perp\rangle, \quad R_0 |B_x\rangle = \cos\theta |\phi_0\rangle + \sin\theta |\phi_0^\perp\rangle.$$

Hence R_0 reflects about $|\phi_0\rangle$. Similarly, R_x reflects about $|B_x\rangle$ as shown below.

$$R_x |\phi_0\rangle = -\sin\theta |A_x\rangle + \cos\theta |B_x\rangle, \quad R_x |\phi_0^\perp\rangle = -\sin\theta |\phi_0\rangle - \cos\theta |\phi_0^\perp\rangle.$$

As a consequence, $G = R_0 R_x$ composes two reflections and effectively amounts to an rotation of 2θ . Therefore, after τ_q iterations, the state becomes

$$|\phi_{\tau_q}\rangle := \sin((2\tau_q + 1)\theta) |A_x\rangle + \cos((2\tau_q + 1)\theta) |B_x\rangle.$$

This is illustrated in Fig. 1.

When measuring $|\phi_{\tau_q}\rangle$, an outcome i with $x_i = 1$ occurs with probability

$$\varepsilon_x = \sin^2((2\tau_q + 1)\theta) \geq \left(\frac{2\tau_q + 1}{2}\theta\right)^2 \geq \tau_q^2 \sin^2\theta = \tau_q^2 \frac{\alpha_x}{\omega}.$$

Thus:

$$\varepsilon = \mathbb{E}_{x \leftarrow D} \varepsilon_x \geq \tau_q^2 \frac{\mathbb{E}_x \alpha_x}{\omega} = \tau_q^2 \frac{\sum_i \omega_i \sum_x d_x x_i}{\omega} = \tau_q^2 \frac{\sum_i \omega_i^2}{\omega}.$$

□




Fig. 1. Illustration of the evolution in the two-dimensional plane.

Optimality for Permutation-Invariant Distributions. Consider a special family of distributions, where ω_i are identical for all $i \in [m]$ implying that every i is mapped to 1 with equal probability. We call such a distribution D *permutation invariant*, and in this case our quantum algorithm \mathcal{A} becomes identical to the standard Grover’s algorithm. It also follows immediately Eq. (2) that for any $i, \omega_i = \nu_D$. Therefore we obtain that

$$\frac{\sum_i \omega_i^2}{\omega} = \frac{\sum_i \omega_i^2}{\sum_i \omega_i} = \frac{m\nu_D^2}{m\nu_D} = \nu_D.$$

As a result, quantum algorithm \mathcal{A} succeeds with probability $\Omega(\tau_q^2 \nu_D)$ in the case of permutation-invariant distribution, which is in turn *optimal* by our hardness bound (Theorem 1). This also reproves the tight quantum query complexity for multi-uniform search and Bernoulli search. We summarize it below.

Corollary 2. *For a permutation-invariant distribution D , the quantum algorithm \mathcal{A} coincides with the standard Grover’s algorithm, and it succeeds with probability $\Omega(\tau_q^2 \cdot \nu_D)$ with τ_q quantum queries which is tight.*

In particular, multi-uniform search and Bernoulli search have tight quantum query complexity $\Theta(\tau_q^2 \frac{w}{m})$ and $\Theta(\tau_q^2 \eta)$ for quantum algorithms with τ_q queries.

Acknowledgements. J.G. was partially supported by NSF SaTC grants no. 2001082 and 2055694. F.S. was partially supported by NSF grant no. 1942706 (CAREER). J.G. and F.S. were also partially support by Sony by means of the Sony Research Award Program. A.C. acknowledges support from the National Science Foundation grant CCF-1813814, from the AFOSR under Award Number FA9550-20-1-0108 and the support of the Quantum Advantage Pathfinder (QAP) project, with grant reference EP/X026167/1 and the UK Engineering and Physical Sciences Research Council.

References

[ABKM22] Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the even-mansour cipher. In *Advances in Cryptology – EUROCRYPT 2022*, pages 458–487. Springer, 2022.

[AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In *Advances in Cryptology – CRYPTO 2019*, pages 269–295. Springer, 2019.



- [AMRS20] Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song. Quantum-secure message authentication via blind-unforgeability. In *Advances in Cryptology – EUROCRYPT 2020*. Springer, 2020.
- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 474–483. IEEE, 2014.
- [BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997.
- [BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology – ASIACRYPT 2011*, pages 41–69. Springer, 2011.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and Communications Security*, pages 62–73, 1993.
- [BR94] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption. In *Advances in Cryptology–EUROCRYPT 1994*, pages 92–111. Springer, 1994.
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures-how to sign with rsa and rabin. In *Advances in Cryptology–Eurocrypt 1996*, pages 399–416. Springer, 1996.
- [BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Advances in Cryptology – CRYPTO 2013*, pages 361–379. Springer, 2013.
- [CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. The complexity of nisq, 2022.
- [CCL23] Nai-Hui Chia, Kai-Min Chung, and Ching-Yi Lai. On the need for large quantum depth. *J. ACM*, 70(1), jan 2023.
- [CEV23] Céline Chevalier, Ehsan Ebrahimi, and Quoc-Huy Vu. On security notions for encryption in a quantum world. In *Progress in Cryptology – INDOCRYPT 2022*, pages 592–613. Springer, 2023.
- [CGK+23] Alexandru Cojocaru, Juan Garay, Aggelos Kiayias, Fang Song, and Petros Wallden. Quantum Multi-Solution Bernoulli Search with Applications to Bitcoin’s Post-Quantum Security. *Quantum*, 7:944, 2023.
- [CGS23] Alexandru Cojocaru, Juan Garay, and Fang Song. Generalized hybrid search and applications. Cryptology ePrint Archive, Paper 2023/798, 2023.
- [CM20] Matthew Coudron and Sanketh Menda. Computations with greater quantum depth are strictly more powerful (relative to an oracle). In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, page 889-901, New York, NY, USA, 2020. Association for Computing Machinery.
- [CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In *17th International Theory of Cryptography Conference – TCC 2019*, pages 1–29. Springer, 2019.
- [DFH22] Jelle Don, Serge Fehr, and Yu-Hsuan Huang. Adaptive versus static multi-oracle algorithms, and quantum security of a split-key prf. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 33–51, Cham, 2022. Springer Nature Switzerland.
- [DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *Advances in Cryptology – CRYPTO 2019*, pages 356–383. Springer, 2019.

- [DFMS22] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In *Advances in Cryptology – EUROCRYPT 2022*, pages 677–706. Springer, 2022.
- [DH09] Cătălin Dohotaru and Peter Høyer. Exact quantum lower bound for grover’s problem. *Quantum Information & Computation*, 9(5):533–540, 2009.
- [ES15] Edward Eaton and Fang Song. Making Existential-unforgeable Signatures Strongly Unforgeable in the Quantum Random-oracle Model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography – TQC 2015*, volume 44 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 147–162. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.
- [ES20] Edward Eaton and Fang Song. A note on the instantiability of the quantum random oracle. In *International Conference on Post-Quantum Cryptography*, pages 503–523. Springer, 2020.
- [FO13] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, 2013. Preliminary version in CRYPTO 1999.
- [FOPS04] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the rsa assumption. *Journal of Cryptology*, 17(2):81–104, 2004. Preliminary version in CRYPTO 2001.
- [Gro96] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996.
- [HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the fujisaki-okamoto transformation. In *15th International Theory of Cryptography Conference – TCC 2017*, pages 341–371. Springer, 2017.
- [HLS22] Yassine Hamoudi, Qipeng Liu, and Makrand Sinha. Quantum-classical tradeoffs in the random oracle model. *CoRR*, abs/2211.12954, 2022.
- [HRS16] Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In *19th IACR International Conference on Public-Key Cryptography — PKC 2016*, pages 387–416. Springer, 2016.
- [JST21] Joseph Jaeger, Fang Song, and Stefano Tessaro. Quantum key-length extension. In *19th International Theory of Cryptography Conference – TCC 2021*, pages 209–239. Springer, 2021.
- [KM10] Hidenori Kuwakado and Masakatu Morii. Quantum distinguisher between the 3-round feistel cipher and the random permutation. In *2010 IEEE International Symposium on Information Theory*, pages 2682–2685. IEEE, 2010.
- [Pre18] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [Ros22] Ansis Rosmanis. Hybrid quantum-classical search algorithms. arXiv preprint [arXiv:2202.11443](https://arxiv.org/abs/2202.11443), 2022.
- [Sho01] Victor Shoup. OAEP reconsidered. In *Advances in Cryptology–CRYPTO 2001*, pages 239–259. Springer, 2001.
- [SZ19] Xiaoming Sun and Yufan Zheng. Hybrid decision trees: Longer quantum time is strictly more powerful, 2019.
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Advances in Cryptology – EUROCRYPT 2015*, pages 755–784. Springer, 2015.

- [YZ21] Takashi Yamakawa and Mark Zhandry. Classical vs quantum random oracles. In *Advances in Cryptology – EUROCRYPT 2021*, pages 568–597. Springer, 2021.
- [Zal99] Christof Zalka. Grover’s quantum searching algorithm is optimal. *Physical Review A*, 60(4):2746, 1999.
- [Zha15] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. *International Journal of Quantum Information*, 13(04):1550014, 2015. Preliminary version in IACR CRYPTO 2012.
- [Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In *Advances in Cryptology – CRYPTO 2019*, pages 239–268. Springer, 2019.
- [Zha21] Mark Zhandry. How to construct quantum random functions. *Journal of the ACM (JACM)*, 68(5):1–43, 2021. Preliminary version in FOCS 2012.



Unclonable Non-interactive Zero-Knowledge

Ruta Jawale^(✉)  and Dakshita Khurana 

University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
jawale2@illinois.edu

Abstract. A non-interactive ZK (NIZK) proof enables verification of NP statements without revealing secrets about them. However, an adversary that obtains a NIZK proof may be able to clone this proof and distribute arbitrarily many copies of it to various entities: this is inevitable for any proof that takes the form of a classical string. In this paper, we ask whether it is possible to rely on quantum information in order to build NIZK proof systems that are impossible to clone.

We define and construct *unclonable non-interactive zero-knowledge arguments (of knowledge)* for NP, addressing a question first posed by Aaronson (CCC 2009). Besides satisfying the zero-knowledge and argument of knowledge properties, these proofs additionally satisfy unclonability. Very roughly, this ensures that no adversary can split an honestly generated proof of membership of an instance x in an NP language \mathcal{L} and distribute copies to multiple entities that all obtain accepting proofs of membership of x in \mathcal{L} . Our result has applications to *unclonable signatures of knowledge*, which we define and construct in this work; these *non-interactively* prevent replay attacks.

Keywords: Unclonable · Zero-Knowledge · Quantum Money

1 Introduction

Zero-knowledge (ZK) [27] proofs allow a prover to convince a verifier about the truth of an (NP) statement, without revealing secrets about it. These are among the most widely used cryptographic primitives, with a rich history of study.

Enhancing Zero-Knowledge. ZK proofs for NP are typically defined via the simulation paradigm. A simulator is a polynomial-time algorithm that mimics the interaction of an adversarial verifier with an honest prover, given only the statement, i.e., $x \in \mathcal{L}$, for an instance x of an NP language \mathcal{L} . A protocol satisfies zero-knowledge if it admits a simulator that generates a view for the verifier, which is indistinguishable from the real view generated by an honest prover. This captures the intuition that any information obtained by a verifier upon observing an honestly generated proof, could have been generated by the verifier “on its own” by running the simulator.

Despite being widely useful and popular, there are desirable properties of proof systems that (standard) simulation-based security does not capture. For example, consider (distributions over) instances x of an NP language \mathcal{L} where it is hard to find an NP witness w corresponding to a given instance x . In an “ideal” world, given just the description of one such NP statement $x \in \mathcal{L}$, it is difficult for an adversary to find an NP witness w , and therefore to output *any* proofs of membership of $x \in \mathcal{L}$. And yet, upon obtaining a *single proof* of membership of $x \in \mathcal{L}$, it may suddenly become feasible for an adversary to make many copies of this proof, thereby generating *several* correct proofs of membership of $x \in \mathcal{L}$.

Unfortunately, this attack is inevitable for classical non-interactive proofs: given any proof string, an adversary can always make multiple copies of it. And yet, there is hope to prevent such an attack quantumly, by relying on the *no-cloning* principle.

Indeed, a recent series of exciting works have combined cryptography with the no-cloning principle to develop quantum money [2, 24, 34, 48, 49], quantum tokens for digital signatures [16], quantum copy-protection [1, 3, 8, 23], unclonable encryption [6, 7, 19, 28, 39], unclonable decryption [26], one-out-of-many unclonable security [35], and more. In this work, we combine zero-knowledge and unclonability to address a question first posed by Aaronson [1]:

*Can we construct unclonable quantum proofs?
How do these proofs relate to quantum money or copy-protection?*

1.1 Our Results

We define and construct unclonable non-interactive zero-knowledge argument of knowledge (NIZKAoK). We obtain a construction in the common reference string (CRS) model, as well as one in the quantum(-accessible) random oracle model (QROM). The CRS model allows a trusted third-party to set up a structured string that is provided to both the prover and verifier. On the other hand, the QROM allows both parties quantum access to a truly random function \mathcal{O} .

In what follows, we describe our contributions in more detail.

Definitional Contributions. Before discussing how we formalize the concept of unclonability for NIZKs, it will be helpful to define hard distributions over NP instance-witness pairs.

Hard Distributions over Instance-Witness Pairs. Informally, an efficiently samplable distribution over instance-witness pairs of a language \mathcal{L} is a “hard” distribution if given an instance sampled randomly from this distribution, it is hard to find a witness. Then, unclonable security requires that no adversary given an instance x sampled randomly from the distribution, together with an honestly generated proof, can output *two accepting proofs* of membership of $x \in \mathcal{L}$.

More specifically, a hard distribution $(\mathcal{X}, \mathcal{W})$ over $R_{\mathcal{L}}$ satisfies the following: for any polynomial-sized (quantum) circuit family $\{C_{\lambda}\}_{\lambda \in \mathbb{N}}$,

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_{\lambda}, \mathcal{W}_{\lambda})} [C_{\lambda}(x) \in R_{\mathcal{L}}(x)] \leq \text{negl}(\lambda).$$

For the sake of simplifying our subsequent discussions and definitions, let us fix a NP language \mathcal{L} with corresponding relation \mathcal{R} . Let $(\mathcal{X}, \mathcal{W})$ be some hard distribution over \mathcal{R} .

A Weaker Definition: Unclonable Security. For NIZKs satisfying standard completeness, soundness and ZK, we define a simple, natural variant of unclonable security as follows. Informally, a proof system satisfies unclonable security if, given an honest proof for an instance and witness pair (x, w) sampled from a hard distribution $(\mathcal{X}, \mathcal{W})$, no adversary can produce two proofs that verify with respect to x except with negligible probability.

Definition 1. (*Unclonable Security of NIZK*). *A NIZK proof (Setup, Prove, Verify) satisfies unclonable security if for every language \mathcal{L} and every hard distribution $(\mathcal{X}, \mathcal{W})$ over $R_{\mathcal{L}}$, for every poly-sized quantum circuit family $\{C_{\lambda}\}_{\lambda \in \mathbb{N}}$,*

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_{\lambda}, \mathcal{W}_{\lambda})} \left[\begin{array}{l} \text{Verify}(\text{crs}, x, \pi_1) = 1 \\ \text{and } \text{Verify}(\text{crs}, x, \pi_2) = 1 \end{array} \middle| \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^{\lambda}) \\ \pi \leftarrow \text{Prove}(\text{crs}, x, w) \\ \pi_1, \pi_2 \leftarrow C_{\lambda}(x, \pi) \end{array} \right] \leq \text{negl}(\lambda).$$

In the definition above, we aim to capture the intuition that one of the two proofs output by the adversary can be the honest proof they received, but the adversary cannot output any other correct proof for the same statement. Of course, such a proof is easy to generate if the adversary is able to find the witness w for x , which is exactly why we require hardness of the distribution $(\mathcal{X}, \mathcal{W})$ to make the definition non-trivial.

We also remark that unclonable security of proofs *necessitates* that the proof π keep hidden any witnesses w certifying membership of x in \mathcal{L} , as otherwise an adversary can always clone the proof π by generating (from scratch) another proof for x given the witness w .

A Stronger Definition: Unclonable Extractability. We can further strengthen the definition above to require that any adversary generating two (or more) accepting proofs of membership of $x \in \mathcal{L}$ given a single proof, must have generated one of the two proofs “from scratch” and must therefore “know” a valid witness w for x . This will remove the need to refer to hard languages.

In more detail, we will say that a proof system satisfies *unclonable extractability* if, from any adversary \mathcal{A} that on input a single proof of membership of $x \in \mathcal{L}$ outputs two proofs for x , then we can extract a valid witness w from \mathcal{A} for at least one of these statements with high probability. Our (still, simplified) definition of unclonable extractability is as follows.

Definition 2 (Unclonable Extractability). *A proof (Setup, Prove, Verify) satisfies unclonable security there exists a QPT extractor \mathcal{E} which is an oracle-aided circuit such that for every language \mathcal{L} with corresponding relation $\mathcal{R}_{\mathcal{L}}$ and for every non-uniform polynomial-time quantum adversary \mathcal{A} , for every instance-witness pair $(x, w) \in \mathcal{R}_{\mathcal{L}}$ and $\lambda = \lambda(|x|)$, such that there is a polynomial $p(\cdot)$*

satisfying:

$$\Pr \left[\text{Verify}(\text{crs}, x, \pi_1) = 1 \wedge \text{Verify}(\text{crs}, x, \pi_2) = 1 \left| \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \text{Prove}(\text{crs}, x, w) \\ \pi_1, \pi_2 \leftarrow \mathcal{A}_\lambda(\text{crs}, x, \pi, z) \end{array} \right. \right] \geq \frac{1}{p(\lambda)},$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr[(x, w_{\mathcal{A}}) \in \mathcal{R}_{\mathcal{L}} | w_{\mathcal{A}} \leftarrow \mathcal{E}^{\mathcal{A}}(x)] \geq \frac{1}{q(\lambda)}.$$

In fact, in the technical sections, we further generalize this definition to consider a setting where the adversary obtains an even larger number (say $k - 1$) input proofs on instances x_1, \dots, x_{k-1} , and outputs k or more proofs. Then we require the extraction of an NP witness corresponding to any proofs that attempt to “clone” honestly generated proofs (i.e. the adversary outputs two or more proofs w.r.t. the same instance $x_i \in \{x_1, \dots, x_{k-1}\}$). All our theorem statements hold w.r.t. this general definition. Finally, we also consider definitions and constructions in the quantum-accessible random oracle model (QROM); these are natural generalizations of the definitions above, so we do not discuss them here.

We also show that the latter definition of unclonable extractability implies the former, i.e. unclonable security. Informally, this follows because the extractor guaranteed by the definition of extractability is able to obtain a witness w for x from any adversary, which contradicts hardness of the distribution $(\mathcal{X}, \mathcal{W})$. We refer the reader to the full version [33] for a formal proof of this claim.

Moreover, we can generically boost the unclonable-extractor’s success probability from $1/q(\lambda)$ to $1 - \text{negl}(\lambda)$ with respect to a security parameter λ . For details, see Sect. 4.2 and Sect. 5.2.

Realizations of Unclonable NIZK, and Relationship with Quantum Money. We obtain realizations of unclonable NIZKs in both the common reference string (CRS) and the quantum random oracle (QRO) models, assuming public-key quantum money mini-scheme and other (post-quantum) standard assumptions. We summarize these results below.

Theorem 1 (Informal). *Assuming public-key quantum money mini-scheme, public-key encryption, perfectly binding and computationally hiding commitments, and adaptively sound NIZK arguments for NP, there exists an unclonable-extractable NIZK argument of knowledge scheme in the CRS model.*

Adaptively sound NIZK arguments for NP exist assuming the polynomial quantum hardness of LWE [40].

Theorem 2 (Informal). *Assuming public-key quantum money mini-scheme and honest verifier zero-knowledge arguments of knowledge sigma protocols for NP, there exists an unclonable-extractable NIZK argument of knowledge scheme in the QROM.*

Is Quantum Money Necessary for Unclonable NIZKs? Our work builds unclonable NIZKs for NP by relying on any (public-key) quantum money scheme (mini-scheme), in conjunction with other assumptions such as NIZKs for NP. Since constructions of public-key quantum money mini-scheme are only known based on post-quantum indistinguishability obfuscation [2, 50], it is natural to wonder whether the reliance on quantum money is inherent. We show that this is indeed the case, by proving that unclonable NIZKs in fact imply public-key quantum money mini-scheme.

Theorem 3 (Informal). *Unclonable NIZK arguments for NP imply public-key quantum money mini-scheme.*

Applications Unclonable Signatures of Knowledge. A (classical) signature scheme asserts that a message m has been signed on behalf of a public key pk . However, in order for this signature to be authenticated, the public key pk must be proven trustworthy through a certification chain rooted at a trusted public key PK . However, as [21] argue, this reveals too much information; it should be sufficient for the recipient to only know that *there exists* a public key pk with a chain of trust from PK . To solve this problem, [21] propose *signatures of knowledge* which allow a signer to sign *on behalf of an instance x of an NP-hard language* without revealing its corresponding witness w . Such signatures provide an anonymity guarantee by hiding the pk of the sender.

While this is ideal for many applications, anonymity presents the following downside: a receiver cannot determine whether they were the intended recipient of this signature. In particular, anonymous signatures are more susceptible to *replay attacks*. Replay attacks are a form of passive attack whereby an adversary observes a signature and retains a copy. The adversary then leverages this signature, either at a later point in time or to a different party, to impersonate the original signer. The privacy and financial consequences of replay attacks are steep. They can lead to data breach attacks which cost millions of dollars annually and world-wide [32].

In this work, we construct a signature of knowledge scheme which is the first *non-interactive* signature in the CRS model that is *naturally secure against replay attacks*. Non-interactive, replay attack secure signatures have seen a lot of recent interest including a line of works in the bounded quantum storage model [11] and the quantum random oracle model [10]. Our construction is in the CRS model and relies on the quantum average-case hardness of NP problems, plausible cryptographic assumptions, and the axioms of quantum mechanics. We accomplish this by defining *unclonable signatures of knowledge*: if an adversary, given a signature of a message m with respect to an instance x , can produce two signatures for m which verify with respect to the same instance x , then our extractor is able to extract a witness for x .

Theorem 4 (Informal). *Assuming public-key quantum money mini-scheme, public-key encryption, perfectly binding and computationally hiding commitments, and simulation-sound NIZK arguments for NP, there exists an unclonable-extractable signature of knowledge in the CRS model.*

Our construction involves showing that an existing compiler can be augmented using unclonable NIZKs to construct unclonable signatures of knowledge. The authors of [21] construct signatures of knowledge from CPA secure dense cryptosystems [44, 45] and simulation-sound NIZKs for NP [42, 43]. Signatures of knowledge are signature schemes in the CRS model for which we associate an instance x in a language \mathcal{L} . This signature is simulatable, so there exists a simulator which can create valid signatures without knowledge of a witness for x . Additionally, the signature is extractable which means there is an extractor which is given a trapdoor for the CRS and a signature, and is able to produce a witness for x . We show that, by switching the simulation-sound NIZKs for unclonable simulation-extractable NIZKs (and slightly modifying the compiler), we can construct unclonable signatures of knowledge.

Relationship with Revocation. A recent exciting line of work obtains *certified deletion* for time-lock puzzles [46], non-local games [25], information-theoretic proofs of deletion with partial security [22], encryption schemes [13, 18], device-independent security of one-time pad encryption with certified deletion [36], public-key encryption with certified deletion [30], commitments and zero-knowledge with certified everlasting hiding [31], and fully-homomorphic encryption with certified deletion [9, 12–14, 41]. While certified everlasting deletion of secrets has been explored in the context of *interactive* zero-knowledge proofs [31], there are no existing proposals for *non-interactive* ZK satisfying variants of certified deletion. Our work provides a pathway to building such proofs.

In this work, we construct a quantum *revocable/unclonable anonymous credentials* protocol in which the issuer of credentials uses a pseudonym to anonymize themselves, receivers of credentials do not require any trusted setup, and the issuer has the ability to remove access from other users. Our work follows a line of work on (classical) revocation for anonymous credentials schemes using NIZK [4, 15, 20].

In particular, our construction involves noting that NIZK proof systems that are unclonable can also be viewed as supporting a form of certified deletion/revocation, where in order to delete, an adversary must simply return the entire proof. In other words, the (quantum) certificate of deletion is the proof itself, and this certificate can be verified by running the NIZK verification procedure on the proof. The unclonability guarantee implies that an adversary cannot keep with itself or later have the ability to generate *another proof* for the same instance x . In the other direction, in order to offer certifiable deletion, a NIZK must necessarily be unclonable. To see why, note that if there was an adversary who could clone the NIZK, we could use this adversary to obtain two copies, and provably delete one of them. Even though the challenger for the certifiable deletion game would be convinced that its proof was deleted, we would still be left with another correct proof.

1.2 Related Works

This work was built upon the foundations of and novel concepts introduced by prior literature. We will briefly touch upon some notable such results in this section.

Unclonable Encryptions. *Unclonable encryption* [6, 7, 19, 28, 39] imagines an interaction between three parties in which one party receives a quantum ciphertext and splits this ciphertext in some manner between the two remaining parties. At some later point, the key of the encryption scheme is revealed, yet both parties should not be able to simultaneously recover the underlying message. While our proof systems share the ideology of unclonability, we do not have a similar game-based definition of security. This is mainly due to proof systems offering more structure which can take advantage of to express unclonability in terms of simulators and extractors.

Signature Tokens. Prior work [17] defines and constructs *signature tokens* which are signatures which involve a quantum signing token which can only be used once before it becomes inert. The setting they consider is where a client wishes to delegate the signing process to a server, but does not wish the server to be able to sign more than one message. They rely on quantum money [2] and the no-cloning principle to ensure the signature can only be computed once. For our unclonable signatures of knowledge result, we focus on the setting where a client wishes to authenticate themselves to a server and wants to prevent an adversary from simultaneously, or later, masquerading as them.

One-shot Signatures. The authors of [5] introduce the notion of *one-shot signatures* which extend the concept of signature tokens to a scenario where the client and server only exchange classical information to create a one-use quantum signature token. They show that these signatures can be plausibly constructed in the CRS model from post-quantum indistinguishability obfuscation. Unless additional measures for security, which we discussed in our applications section, are employed, classical communication can be easily copied and replayed at a later point. In contrast, we prevent an adversary from simultaneously, or later, authenticating with the client’s identity.

Post-quantum Fiat-Shamir. Our QROM results are heavily inspired by the recent post-quantum Fiat-Shamir result [37] which proves the post-quantum security of NIZKs in the compressed quantum(-accessible) random oracle model (compressed QROM). These classical NIZKs are the result of applying Fiat-Shamir to post-quantum sigma protocols which are HVZKAoKs. We further extend, and crucially rely upon, their novel proof techniques to prove extractability (for AoK) and programmability (for ZK) to achieve extractability and programmability for some protocols which output quantum proofs.

1.3 Concurrent Works

Unclonable Commitments and Proofs. A recent, concurrent work [29] defines and constructs unclonable commitments and interactive unclonable proofs. They additionally construct commitments in the QROM that are unclonable with respect to any verification procedures, and they show that it is impossible to have (interactive) proofs with the same properties. The authors also observe a similar relationship between non-interactive unclonable proofs and public-key quantum money via unclonable commitments. They also briefly mention a connection between unclonable commitments and unclonable credentials.

In contrast, we define unclonable-extractable proofs which we construct in the *non-interactive* setting in *both* the crs model and the QROM. We also show a relationship between non-interactive unclonable-extractable proofs and quantum money in *both* the crs model and the QROM. Our work also *formalizes* the relationship between unclonable-extractable proofs and unclonable *anonymous* credentials.

2 Technical Overview

In this section, we give a high-level overview of our construction and the techniques underlying our main results.

2.1 Unclonable Extractable NIZKs in the CRS Model

Our construction assumes the existence of public-key encryption, classical bit commitments where honestly generated commitment strings are perfectly binding, along with

- *Public-key quantum money mini-scheme* (which is known assuming post-quantum $i\mathcal{O}$ and injective OWFs [50]). At a high level, public-key quantum money mini-scheme consists of two algorithms: Gen and Ver . Gen on input a security parameter, outputs a (possibly mixed-state) quantum banknote $\rho_{\mathbb{S}}$ along with a classical serial number s . Ver is public, takes a quantum money banknote, and outputs either a classical serial number s , or \perp indicating that its input is an invalid banknote. The security guarantee is that no efficient adversary given an honest banknote $\rho_{\mathbb{S}}$ can output two notes $\rho_{\mathbb{S},0}$ and $\rho_{\mathbb{S},1}$ that both pass the verification and have serial numbers equal to that of $\rho_{\mathbb{S}}$.
- *Post-quantum NIZKs for NP*, which are known assuming the post-quantum hardness of LWE. These satisfy (besides completeness) (1) soundness, i.e., no efficient prover can generate accepting proofs for false NP statements, and (2) zero-knowledge, i.e., the verifier obtains no information from an honestly generated proof beyond what it could have generated on *its own* given the NP statement itself.

Construction. Given these primitives, the algorithms (Setup, Prove, Verify) of the unclonable extractable NIZK are as follows.

SETUP(1^λ): The setup algorithm samples a public key pk of a public-key encryption, the common reference string crs of a classical (post-quantum) NIZK for NP, along with a perfectly binding, computationally hiding classical commitment to 0^λ with uniform randomness t , i.e. $c = \text{Com}(0^\lambda; t)$. It outputs $(\text{pk}, \text{crs}, c)$.

PROVE: Given the CRS $(\text{pk}, \text{crs}, c)$, instance x and witness w , output $(\rho_{\mathbb{S}}, s, ct, \pi)$ where

- The state $\rho_{\mathbb{S}} \leftarrow \text{Gen}$ is generated as a quantum banknote with associated serial number s .
- The ciphertext $ct = \text{Enc}_{\text{pk}}(w; u)$ is an encryption of the witness w with randomness u .
- The proof string π is a (post-quantum) NIZK for the following statement using witness (w, u) :

EITHER $(\exists w, u : ct = \text{Enc}_{\text{pk}}(w; u) \wedge R_L(x, w) = 1)$ OR $(\exists r : c = \text{Com}(s; r))$,

where we recall that pk and c were a part of the CRS output by the Setup algorithm.

VERIFY: Given CRS $(\text{pk}, \text{crs}, c)$, instance x and proof $(\rho_{\mathbb{S}}, s, ct, \pi)$, check that (1) $\text{Ver}(\rho_{\mathbb{S}})$ outputs s and (2) π is an accepting NIZK argument of the statement above.

Analysis. Completeness, soundness/argument of knowledge and ZK for this construction follow relatively easily, so we focus on unclonable extractability in this overview. Recall that unclonable extractability requires that no adversary, given an honestly generated proof for $x \in \mathcal{L}$, can split this into *two accepting proofs* for $x \in \mathcal{L}$ (as long as it is hard to find a witness for x). Towards a contradiction, suppose an adversary splits a proof into 2 accepting proofs $(\rho_{\mathbb{S},0}, s_1, ct_1, \pi_1)$, $(\rho_{\mathbb{S},1}, s_2, ct_2, \pi_2)$. Then,

- If $s_1 = s_2 = s$, the adversary given one bank note with serial number s generated two valid banknotes $\rho_{\mathbb{S},0}$ and $\rho_{\mathbb{S},1}$ that both have the same serial number s . This contradicts the security of quantum money.
- Otherwise, there is a $b \in \{1, 2\}$ such that $s_b \neq s$. Then, consider an indistinguishable hybrid where the adversary obtains a simulated proof generated *without witness w* as follows: (1) sample quantum banknote $\rho_{\mathbb{S}}$ with serial number s , (2) sample public key pk along with secret key sk , (3) generate $c = \text{Com}(s; t)$, $ct = \text{Enc}_{\text{pk}}(0; u)$, (4) generate proof π using witness t (since $c = \text{Com}(s; t)$) instead of using witness w . Send common reference string $(\text{pk}, \text{crs}, c)$ and proof $(\rho_{\mathbb{S}}, s, ct, \pi)$ to the adversary. Now, the proof that the adversary generates with $s_b \neq s$ *must* contain $ct_b = \text{Enc}_{\text{pk}}(w; u)$, since c being generated as a commitment to $s \neq s_b$ along with the perfect binding property implies that $(\nexists r : c = \text{Com}(s_b; r))$. That is, given instance x , the adversary can be used to compute a witness w for x by decrypting ciphertext ct_b , thereby contradicting the hardness of the distribution.

Our technical construction in Sect. 4.4, while conceptually the same, is formalized slightly differently. It uses NIZKs with an enhanced simulation-extraction property, which can be generically constructed from NIZK (see Sect. 4.1). Having constructed unclonable extractable arguments in the CRS model, in the next section, we analyze a construction of unclonable extractable arguments in the QROM.

2.2 Unclonable Extractable NIZK in the QROM

We now turn our attention to the QRO setting in which we demonstrate a protocol which is provably unclonable. Our construction assumes the existence of public-key quantum money mini-scheme and a *post-quantum sigma protocol for NP*. A sigma protocol (P, V) is an interactive three-message honest-verifier protocol: the prover sends a commitment message, the verifier sends a uniformly random challenge, and the prover replies by opening its commitment at the locations specified by the random challenge.

Construction. The algorithms $(\text{PROVE}, \text{VERIFY})$ of the unclonable extractable NIZK in the QROM are as follows.

PROVE: Given an instance x and witness w , output $(\rho_s, s, \alpha, \beta, \gamma)$ where

- The quantum banknote ρ_s is generated alongside associated serial number s .
- P is run to compute the sigma protocol’s commitment message as α given (x, w) as input.
- The random oracle is queried on input (α, s, x) in order to obtain a challenge β .
- P is run, given as input (x, w, α, β) and its previous internal state, to compute the sigma protocol’s commitment openings as γ .

VERIFY: Given instance x and proof $(\rho_s, s, \alpha, \beta, \gamma)$, check that (1) the quantum money verifier accepts (ρ_s, s) , (2) the random oracle on input (α, s, x) outputs β , and (3) V accepts the transcript (α, β, γ) with respect to x .

Analysis. Since the completeness, argument of knowledge and zero-knowledge properties are easy to show, we focus on unclonable extractability. Suppose an adversary was able to provide two accepting proofs $\pi_1 = (\rho_{s,0}, s_1, \alpha_1, \beta_1, \gamma_1)$ and $\pi_2 = (\rho_{s,1}, s_2, \alpha_2, \beta_2, \gamma_2)$ for an instance x for which it received an honestly generated proof $\pi = (\rho_s, s, \alpha, \beta, \gamma)$. Then,

- Suppose $s_1 = s_2 = s$. In this case, the adversary given one bank note with serial number s generated two valid banknotes $\rho_{s,0}$ and $\rho_{s,1}$ that both have the same serial number s . This contradicts the security of quantum money.
- Otherwise, there is a $b \in [1, 2]$ such that $s_b \neq s$. By the zero-knowledge property of the underlying HVZK sigma protocol, this event also occurs when the proof π that the adversary is given is replaced with a simulated proof. Specifically, we build a reduction that locally programs the random oracle

at location (α, s, x) in order to generate a simulated proof for the adversary. Since the adversary's own proof for $s_b \neq s$ is generated by making a distinct query $(\alpha_b, s_b, x) \neq (\alpha, s, x)$, the programming on (α, s, x) does not affect the knowledge extractor for the adversary's proof, which simply rewinds the (quantum) random oracle to extract a witness for x , following [37]. This allows us to obtain a contradiction, showing that our protocol must be unclonable.

2.3 Unclonable NIZKs Imply Quantum Money Mini-Scheme

Finally, we discuss why unclonable NIZKs satisfying even the weaker definition of unclonable security (i.e., w.r.t. hard distributions) imply public-key quantum money mini-scheme. Given an unclonable NIZK, we build a public-key quantum money mini-scheme as follows.

Construction. Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \text{NP}$. Let $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ be an unclonable NIZK protocol for \mathcal{L} .

$\underline{\text{GEN}}(1^\lambda)$: Sample $(x, w) \leftarrow (\mathcal{X}, \mathcal{W})$, $\text{crs} \leftarrow \text{Setup}(1^\lambda, x)$, and an unclonable NIZK proof π as $\text{Prove}(\text{crs}, x, w)$. Output a (possibly mixed-state) quantum banknote $\rho_{\mathbb{S}} = \pi$, and associated serial number $s = (\text{crs}, x)$.

$\underline{\text{VER}}(\rho_{\mathbb{S}}, s)$: Given a (possibly mixed-state) quantum banknote $\rho_{\mathbb{S}}$ and a classical serial number s as input, parse $\rho_{\mathbb{S}} = \pi$ and $s = (\text{crs}, x)$, and output the result of $\text{Verify}(\text{crs}, x, \pi)$.

Analysis. The correctness of the quantum money scheme follows from the completeness of the unclonable NIZK Π . We will now argue that this quantum money scheme is unforgeable. Suppose an adversary \mathcal{A} given a quantum banknote and classical serial number $(\rho_{\mathbb{S}}, s)$ was able to output two banknotes $(\rho_{\mathbb{S},0}, \rho_{\mathbb{S},1})$ both of which are accepted with respect to s . We can use \mathcal{A} to define a reduction to the uncloneability of our NIZK Π as follows:

- The NIZK uncloneability challenger outputs a hard instance-witness pair (x, w) , a common reference string crs , and an unclonable NIZK π to the reduction.
- The reduction outputs a banknote $(\rho_{\mathbb{S}}, s)$ to the adversary, where $\rho_{\mathbb{S}} = \pi$ and $s = (\text{crs}, x)$. It receives two quantum banknotes $(\rho_{\mathbb{S},0}, \rho_{\mathbb{S},1})$ from \mathcal{A} , and finally outputs two proofs (π_0, π_1) where $\pi_0 = \rho_{\mathbb{S},0}$ and $\pi_1 = \rho_{\mathbb{S},1}$.

If \mathcal{A} succeeds in breaking unforgeability, then the quantum money verifier accepts both banknotes $(\rho_{\mathbb{S},0} = \pi_0, \rho_{\mathbb{S},1} = \pi_1)$, with respect to the same serial number $s = (\text{crs}, x)$. By syntax of the verification algorithm, this essentially means that both *proofs* (π_0, π_1) are accepting proofs for membership of the same instance $x \in \mathcal{L}$, w.r.t. crs , leading to a break in the unclonability of NIZK.

2.4 Unclonable Signatures of Knowledge

Informally, a signature of knowledge has the following property: if an adversary, given a signature of a message m with respect to an instance x , can produce

two signatures for m which verify with respect to the same instance x , then the adversary *must know* (and our extractor will be able to extract) a witness for x .

We obtain unclonable signatures of knowledge assuming the existence of an unclonable extractable *simulation-extractable* NIZK for NP. Simulation-extractability states that an adversary which is provided any number of simulated proofs for instance and witness pairs of their choosing, cannot produce an accepting proof π for an instance x which they have not queried before and where extraction fails to find an accepting witness w . Our unclonable extractable NIZK for NP in the CRS model can, with some extra work, be upgraded to simulation-extractable.

We informally describe the construction of signatures of knowledge from such a NIZK below.

Construction. Let $(\text{Setup}, \text{P}, \text{V})$ be non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for NP. Let \mathcal{R} be the NP relation corresponding to \mathcal{L} .

SETUP: The setup algorithm samples a common reference string crs of an unclonable-extractable simulation-extractable NIZK for NP. It outputs crs .

SIGN: Given the CRS crs , instance x , witness w , and message m , output signature π where

- The proof string π is an unclonable-extractable simulation-extractable NIZK with tag m using witness w of the following statement:

$$(\exists w : (x, w) \in \mathcal{R}).$$

VERIFY: Given CRS crs , instance x , message m , and signature π , check that π is an accepting NIZK proof with tag m of the statement above.

Analysis. The simulatability (extractability) property follows from the zero-knowledge (resp. simulation-extractability) properties of the NIZK. Suppose an adversary \mathcal{A} given a signature σ was able to forge two signatures $\sigma_1 = \pi_1$ and $\sigma_2 = \pi_2$, and, yet, our extractor was to fail to extract a witness w from \mathcal{A} . Then,

- Either both proofs π_1 and π_2 are accepting proofs for membership of the same instance w.r.t. crs . However, this contradicts the unclonability of the NIZK.
- Otherwise there exists a proof π_i (where $i \in \{1, 2\}$) for an instance which \mathcal{A} has not previously seen a proof for. We can switch to a hybrid where our signatures contain simulated proofs for the NIZK. But now, we have that the verifier accepts a proof for an instance which \mathcal{A} has not seen a simulated proof for and, yet, we cannot extract a witness from \mathcal{A} . This contradicts the simulation extractability of the NIZK.

Roadmap. In Sect. 4, we define and construct unclonable NIZKs in the CRS model, and in Sect. 5, in the QROM. Along the way, we also show that unclonable NIZKs imply quantum money (in the CRS and QRO model respectively). Later, we show how to define and construct unclonable signatures of knowledge from unclonable NIZKs in the CRS model.

3 Preliminaries

We defer definitions to the full version [33]; below we recall some useful theorems.

3.1 Post-quantum Commitments and Encryption

Theorem 5 (Post-quantum Commitment). [38] *Assuming the polynomial quantum hardness of LWE , there exists a non-interactive commitment with perfect binding and computational hiding.*

3.2 NIZKs in the CRS Model

Theorem 6 (Post-quantum NIZK Argument for NP in the CRS Model). [40] *Assuming the polynomial quantum hardness of LWE , there exists a non-interactive adaptively computationally sound, adaptively computationally zero-knowledge argument for NP in the common reference string model.*

Theorem 7 (Simulation Sound Compiler). [43] *Given one-way functions and a single-theorem NIZK proof system for NP, then there exists a non-interactive simulation sound, adaptively multi-theorem computationally zero-knowledge proof for NP in the common reference string model.*

Corollary 1 (Post-quantum Simulation Sound NIZK for NP). *Assuming the polynomial quantum hardness of LWE , there exists a post-quantum non-interactive simulation sound, adaptively multi-theorem computationally zero-knowledge proof for NP in the common reference string model.*

Proof. This follows from Theorem 6 and Theorem 7.

3.3 NIZKs in the QRO Model

Theorem 8 (NIZKAoK in QROM [37, 47]). *Let Π be a post-quantum sigma protocol. The Fiat-Shamir heuristic applied to Π yields a classical post-quantum NIZKAoK in the QROM.*

3.4 Quantum Money

Theorem 9 (Quantum Money from Subspace Hiding Obfuscation [2, 50]). *If injective one-way functions and post-quantum iO exist, then public-key quantum money exists.*

4 Unclonable Non-interactive Zero-Knowledge in the CRS Model

4.1 Simulation-Extractable NIZK

We defer the definition, and proofs to the full version [33]; below we state our results.

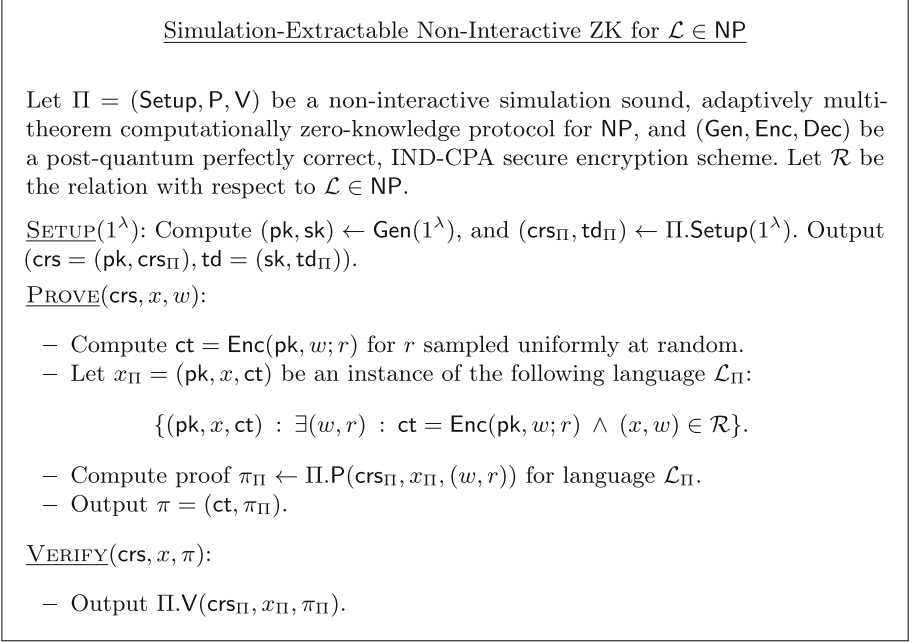


Fig. 1. Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \text{NP}$

Theorem 10 (Post-quantum Simulation-Extractable NIZK for NP in the CRS Model). *Let NP relation \mathcal{R} with corresponding language \mathcal{L} be given.*

Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a non-interactive post-quantum simulation sound, adaptively multi-theorem computationally zero-knowledge protocol for NP. Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a post-quantum perfectly correct, IND-CPA secure encryption scheme.

$(\text{Setup}, \text{P}, \text{V})$ as defined in Fig. 1 will be a non-interactive post-quantum simulation-extractable, adaptively multi-theorem computationally zero-knowledge argument for \mathcal{L} in the common reference string model.

Corollary 2 (Post-quantum Simulation-Extractable NIZK for NP in the CRS Model). *Assuming the polynomial quantum hardness of LWE, there exists a simulation-extractable, adaptively multi-theorem computationally zero-knowledge argument for NP in the common reference string model.*

Proof. This follows from Corollary 1 and Theorem 10.

4.2 Unclonability Definitions

We consider two definitions of unclonability for NIZKs. The first one, motivated by simplicity, informally guarantees that no adversary given honestly proofs for “hard” instances is able to output more than one accepting proof for the same instance.

Definition 3 ((Quantum) Hard Distribution). *Let an NP relation \mathcal{R} be given. $(\mathcal{X}, \mathcal{W})$ is a (quantum) hard distribution over \mathcal{R} if the following properties hold.*

- **Syntax.** $(\mathcal{X}, \mathcal{W})$ is indexable by a security parameter $\lambda \in \mathbb{N}$. For every choice of $\lambda \in \mathbb{N}$, the support of $(\mathcal{X}_\lambda, \mathcal{W}_\lambda)$ is over instance and witness pairs (x, w) such that $x \in \mathcal{L}$, $|x| = \lambda$, and $(x, w) \in \mathcal{R}$.
- **Hardness.** For every polynomial-sized (quantum) circuit family $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_\lambda, \mathcal{W}_\lambda)} [(x, \mathcal{A}_\lambda(x)) \in \mathcal{R}] \leq \text{negl}(\lambda).$$

Definition 4. (Unclonable Security for Hard Instances). *A proof $(\text{Setup}, \text{P}, \text{V})$ satisfies unclonable security for a language \mathcal{L} with corresponding relation $\mathcal{R}_\mathcal{L}$ if for every polynomial-sized quantum circuit family $\{C_\lambda\}_{\lambda \in \mathbb{N}}$, and for every hard distribution $\{\mathcal{X}_\lambda, \mathcal{W}_\lambda\}_{\lambda \in \mathbb{N}}$ over $\mathcal{R}_\mathcal{L}$, there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,*

$$\Pr_{(x,w) \leftarrow (\mathcal{X}_\lambda, \mathcal{W}_\lambda)} \left[\text{V}(\text{crs}, x, \pi_1) = 1 \bigwedge \text{V}(\text{crs}, x, \pi_2) = 1 \left| \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \text{P}(\text{crs}, x, w) \\ \pi_1, \pi_2 \leftarrow C_\lambda(x, \pi) \end{array} \right. \right] \leq \text{negl}(\lambda).$$

We will now strengthen this definition to consider a variant where from any adversary \mathcal{A} that on input a single proof of membership of $x \in \mathcal{L}$ outputs two proofs for x , we can extract a valid witness w for x with high probability. In fact, we can further generalize this definition to a setting where the adversary obtains an even larger number (say $k-1$) input proofs on instances x_1, \dots, x_{k-1} , and outputs k or more proofs. Then we require the extraction of an NP witness corresponding to any proofs that are *duplicated* (i.e. two or more proofs w.r.t. the same instance $x_i \in \{x_1, \dots, x_{k-1}\}$). We write this definition below.

Definition 5 (($k-1$)-to- k -Unclonable Extractable NIZK). *Let security parameter $\lambda \in \mathbb{N}$ and NP relation \mathcal{R} with corresponding language \mathcal{L} be given. Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be given such that Setup, P and V are $\text{poly}(\lambda)$ -size quantum algorithms. We have that for any $(x, w) \in \mathcal{R}$, (crs, td) is the output of Setup on input 1^λ , P receives an instance and witness pair (x, w) along with crs as input and outputs π , and V receives an instance x , crs , and proof π as input and outputs a value in $\{0, 1\}$.*

Π is a non-interactive $(k-1)$ -to- k -unclonable zero-knowledge quantum protocol for language \mathcal{L} if the following holds:

- Π is a quantum non-interactive zero-knowledge protocol for language \mathcal{L} .

- **$(k - 1)$ -to- k -Unclonable with Extraction:** *There exists an oracle-aided polynomial-size quantum circuit \mathcal{E} such that for every polynomial-size quantum circuit \mathcal{A} , for every tuple of $k - 1$ instance-witness pairs $(x_1, \omega_1), \dots, (x_{k-1}, \omega_{k-1}) \in \mathcal{R}$, for every instance x , if there exists a polynomial $p(\cdot)$ such that*

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall i \in [k-1], \pi_i \leftarrow \text{P}(\text{crs}, x_i, w_i) \\ \{\tilde{x}_i, \tilde{\pi}_i\}_{i \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_i, \pi_i\}_{i \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall i \in \mathcal{J}, \mathbf{V}(\text{crs}, x, \tilde{\pi}_i) = 1 \end{array} \right] \geq \frac{1}{p(\lambda)},$$

then there is also a polynomial $q(\cdot)$ such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(x_1, \dots, x_{k-1}, x)} [(x, w) \in \mathcal{R}] \geq \frac{1}{q(\lambda)}.$$

We observe in Definition 5 that we can generically boost the extractor's success probability to $1 - \text{negl}(\lambda)$ with respect to a security parameter λ .

Definition 6 (**$(k - 1)$ -to- k -Unclonable Strong-Extractable NIZK**). *Let security parameter $\lambda \in \mathbb{N}$ and NP relation \mathcal{R} with corresponding language \mathcal{L} be given. Let $\Pi = (\text{Setup}, \text{P}, \mathbf{V})$ be given such that Setup, P and \mathbf{V} are $\text{poly}(\lambda)$ -size quantum algorithms. We have that for any $(x, w) \in \mathcal{R}$, (crs, td) is the output of Setup on input 1^λ , P receives an instance and witness pair (x, w) along with crs as input and outputs π , and \mathbf{V} receives an instance x , crs , and proof π as input and outputs a value in $\{0, 1\}$.*

Π is a non-interactive $(k - 1)$ -to- k -unclonable zero-knowledge quantum protocol for language \mathcal{L} if the following holds:

- Π is a quantum non-interactive zero-knowledge protocol for language \mathcal{L} .
- **$(k - 1)$ -to- k -Unclonable with Strong-Extraction:** *There exists an oracle-aided polynomial-size quantum circuit \mathcal{E} such that for every polynomial-size quantum circuit \mathcal{A} with non-uniform quantum advice aux , for every tuple of $k - 1$ instance-witness pairs $(x_1, \omega_1), \dots, (x_{k-1}, \omega_{k-1}) \in \mathcal{R}$, for every instance x if there is a polynomial $p(\cdot)$ where*

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall i \in [k-1], \pi_i \leftarrow \text{P}(\text{crs}, x_i, w_i) \\ \{\tilde{x}_i, \tilde{\pi}_i\}_{i \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_i, \pi_i\}_{i \in [k-1]}, \text{aux})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall i \in \mathcal{J}, \mathbf{V}(\text{crs}, x, \tilde{\pi}_i) = 1 \end{array} \right] \geq \frac{1}{p(\lambda)},$$

then there is also a polynomial $\text{poly}(\cdot)$ and a negligible function $\text{negl}(\cdot)$ such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(x_1, \dots, x_{k-1}, x, \text{aux}^{\otimes \text{poly}(\lambda)})} [(x, w) \in \mathcal{R}] \geq 1 - \text{negl}(\lambda).$$

We describe two useful lemmas to compare the above definitions.

Lemma 1. *Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a 1-to-2-unclonable with extraction, non-interactive zero-knowledge quantum protocol (Definition 5). Then, Π satisfies Definition 4.*

For a proof of Lemma 1, we refer to the full version [33].

Lemma 2. *Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a $(k - 1)$ -to- k -unclonable with extraction, non-interactive zero-knowledge quantum protocol (Definition 5). Then, Π satisfies Definition 6.*

For a proof of Lemma 2, we refer to the full version [33].

From the above lemmas, we conclude that Definition 5 is the strongest definition. In the following sections, we construct a protocol that satisfies Definition 5.

4.3 Unclonable NIZK Implies Public-Key Quantum Money Mini-scheme

Public-Key Quantum Money Mini-Scheme

Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \text{NP}$. Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be an unclonable non-interactive zero-knowledge protocol for \mathcal{L} .

GEN (1^λ) : Sample a hard instance-witness pair $(x, w) \leftarrow (\mathcal{X}, \mathcal{Y})$, a common reference string $(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, x)$, and a proof $\pi \leftarrow \text{P}(\text{crs}, x, w)$. Output $(\rho_{\mathcal{S}} = \pi, s = (\text{crs}, x))$.

VERIFY $(\rho_{\mathcal{S}}, s)$: Parse $\rho_{\mathcal{S}} = \pi$ and $s = (\text{crs}, x)$. Output $\text{V}(\text{crs}, x, \pi)$.

Fig. 2. Public-Key Quantum Money Mini-Scheme from an Unclonable Non-Interactive Quantum Protocol

Theorem 11. *Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \text{NP}$. Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ satisfy Definition 4. Then $(\text{Setup}, \text{P}, \text{V})$ implies a public-key quantum money mini-scheme as described in Fig. 2.*

We defer the proof to the full version [33].

4.4 Construction and Analysis of Unclonable-Extractable NIZK in CRS Model

Unclonable Non-Interactive ZK for $\mathcal{L} \in \text{NP}$

Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a non-interactive simulation-extractable, adaptively multi-theorem computationally zero-knowledge protocol for NP , Com be a post-quantum perfectly binding, computationally hiding commitment scheme, and $(\text{NoteGen}, \text{Ver})$ be a public-key quantum money scheme. Let \mathcal{R} be the relation with respect to $\mathcal{L} \in \text{NP}$.

SETUP (1^λ) : Sample the common reference string $(\text{crs}_\Pi, \text{td}_\Pi) \leftarrow \Pi.\text{Setup}(1^\lambda)$, and s^*, r^* uniformly at random. Define $c = \text{Com}(s^*; r^*)$ and output $(\text{crs} = (\text{crs}_\Pi, c), \text{td} = \text{td}_\Pi)$.

PROVE (crs, x, w) :

- Compute a quantum note and associated serial number $(\rho_s, s) \leftarrow \text{NoteGen}$.
- Let $x_\Pi = (c, x, s)$ be an instance of the following language \mathcal{L}_Π :

$$\{(c, x, s) : \exists z : (x, z) \in \mathcal{R} \vee c = \text{Com}(s; z)\}.$$

- Compute proof $\pi_\Pi \leftarrow \Pi.\text{P}(\text{crs}_\Pi, x_\Pi, w)$ for language \mathcal{L}_Π .
- Output $\pi = (\rho_s, s, \pi_\Pi)$.

VERIFY (crs, x, π) :

- Check that $\text{Ver}(\rho_s, s)$ outputs 1 and that $\Pi.\text{V}(\text{crs}_\Pi, x_\Pi, \pi_\Pi)$ outputs 1.
- If both checks pass, output 1. Otherwise, output 0.

Fig. 3. Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \text{NP}$

Theorem 12. *Let $k(\cdot)$ be a polynomial. Let NP relation \mathcal{R} with corresponding language \mathcal{L} be given.*

Let $(\text{NoteGen}, \text{Ver})$ be a public-key quantum money mini-scheme and Com be a post-quantum commitment scheme. Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a non-interactive post-quantum simulation-extractable, adaptive multi-theorem computational zero-knowledge protocol for NP .

$(\text{Setup}, \text{P}, \text{V})$ as defined in Fig. 3 will be a non-interactive quantum simulation-extractable, adaptive multi-theorem computationally zero-knowledge, and $(k - 1)$ -to- k -unclonable argument with extraction protocol for \mathcal{L} in the common reference string model (Definition 5).

Proof. Completeness follows from perfect correctness of the public key quantum money scheme, and perfect completeness of Π .

See the full version [33] for proofs of zero-knowledge and simulation extractability.

Let $\Pi.\text{Sim} = (\Pi.\text{Sim}_0, \Pi.\text{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π . We define Sim_0 with oracle access to $\Pi.\text{Sim}_0$ as follows: *Input:* 1^λ .

- (1) Send 1^λ to $\Pi.\text{Sim}_0$. Receive $(\text{crs}_\Pi, \text{td}_\Pi)$ from $\Pi.\text{Sim}_0$.
- (2) Sample s^*, r^* uniformly at random. Define $c = \text{Com}(s^*; r^*)$.
- (3) Output $\text{crs} = (\text{crs}_\Pi, c)$ and $\text{td} = \text{td}_\Pi$.

We define Sim_1 with oracle access to $\Pi.\text{Sim}_1$ as follows:

Input: $\text{crs} = (\text{crs}_\Pi, c)$, $\text{td} = \text{td}_\Pi$, x .

- (1) Sample $(\rho_\$, s) \leftarrow \text{NoteGen}(1^\lambda)$.
- (2) Define $x_\Pi = (c, x, s)$. Send $(\text{crs}_\Pi, \text{td}_\Pi, x_\Pi)$ to $\Pi.\text{Sim}_1$. Receive π_Π from $\Pi.\text{Sim}_1$.
- (3) Output $\pi = (\rho_\$, s, \pi_\Pi)$.

Claim (4.1). Let Ext be as defined earlier, in the current proof of simulation-extractability. There exists a negligible function $\text{negl}(\cdot)$ such that for every polynomial-size quantum circuit \mathcal{B} ,

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{B}^{\text{Sim}_1(\text{crs}, \text{td}, \cdot)}(\text{crs}) \\ w \leftarrow \text{Ext}(\text{crs}, \text{td}, x, \pi)}} [\Pi.\mathcal{V}(\text{crs}_\Pi, x_\Pi, \pi_\Pi) = 1 \wedge x_\Pi \notin Q_\Pi \wedge (x, w) \notin \mathcal{R}] \leq \text{negl}(\lambda)$$

where Q_Π is the list of queries forwarded by Sim_1 to $\Pi.\text{Sim}_1$.

See the full version [33] for proof of Claim 4.1.

Unclonable Extractability. Let $\Pi.\text{Sim} = (\Pi.\text{Sim}_0, \Pi.\text{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π . Let $\Pi.\text{Ext}$ be the simulation-extraction extractor of Π with respect to $\Pi.\text{Sim}$. Let $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ be the simulator, with oracle access to $\Pi.\text{Sim}$, as defined in the proof that Fig. 3 is adaptive multi-theorem computational zero-knowledge. Let Ext be the extractor, based on Sim , as defined in the proof that Fig. 3 is simulation-extractable. We define \mathcal{E} with oracle access to Sim , Ext , and some \mathcal{A} as follows:
Hardwired: x_1, \dots, x_{k-1}, x

- (1) Send 1^λ to Sim_0 . Receive (crs, td) from Sim_0 .
- (2) For $\iota \in [k-1]$: send $(\text{crs}, \text{td}, x_\iota)$ to Sim_1 , and receive π_ι from Sim_1 .
- (3) Send $(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})$ to \mathcal{A} . Receive $\{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]}$ from \mathcal{A} .
- (4) Define j' uniformly at random from $[k]$.
- (5) Output $\text{Ext}(\text{crs}, \text{td}, x, \tilde{\pi}_{j'})$ as w .

Let $\mathcal{A}, (x_1, w_1), \dots, (x_{k-1}, w_{k-1}) \in \mathcal{R}$, x , polynomial $p(\cdot)$, and negligible function $\text{negl}(\cdot)$ be given such that \mathcal{A} outputs more accepting proofs for x than \mathcal{A} received, and yet the extractor \mathcal{E} is unable to extract a valid witness for x from \mathcal{A} . Restated more formally, that is that

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{P}(\text{crs}, x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathcal{V}(\text{crs}, x, \tilde{\pi}_\iota) = 1 \end{array} \right] \geq \frac{1}{p(\lambda)}, \quad (1)$$

and for all polynomials $p'(\cdot)$ (there are infinitely many λ) such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(x_1, \dots, x_{k-1}, x)} [(x, w) \in \mathcal{R}] \leq \frac{1}{p'(\lambda)}. \quad (2)$$

We parse the output of the adversary \mathcal{A} as $\tilde{\pi}_\iota = (\widetilde{\rho_{\mathbb{S}, \iota}}, \widetilde{s}_\iota, \widetilde{\pi_{\Pi, \iota}})$ for all $\iota \in [k]$.

Given Eq. (1), we may be in one of the two following cases: either \mathcal{A} generates two accepting proofs which have the same serial number as an honestly generated proof (for an infinite set of λ), or \mathcal{A} does not (for an infinite set of λ). We consider that either of these two scenarios occur with at least $1/(2p(\lambda))$ probability and show that each reaches a contradiction.

Scenario One

Say that (for an infinite set of λ) \mathcal{A} generates two accepting proofs which have the same serial number as an honestly generated proof with at least $1/(2p(\lambda))$ probability. Symbolically,

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{P}(\text{crs}, x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \text{V}(\text{crs}, x, \tilde{\pi}_\iota) = 1 \\ \text{and } \exists i^* \in [k-1] \exists j^*, \ell^* \in \mathcal{J} \\ \text{s.t. } s_{i^*} = \widetilde{s}_{j^*} = \widetilde{s}_{\ell^*} \end{array} \right] \geq \frac{1}{2p(\lambda)}. \quad (3)$$

Through a hybrid argument, we can get a similar event with fixed indices i^* , j^* , and ℓ^* which belong to their respective sets with an advantage of $1/(2k^3p(\lambda))$. By using the advantage of \mathcal{A} in this game, we can show a reduction that breaks the unforgeability of the quantum money scheme. We will now outline this reduction. Reduction: to unforgeability of quantum money scheme given oracle access to \mathcal{A} .

Hardwired with: $(x_1, w_1), \dots, (x_{k-1}, w_{k-1}), x, i^*, j^*, \ell^*$.

- (1) Compute $(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda)$ where $\text{crs} = (\text{crs}_{\Pi}, c)$ and $\text{td} = \text{td}_{\Pi}$.
- (2) Receive $(\rho_{\mathbb{S}}, s) \leftarrow \text{NoteGen}$ from the challenger.
- (3) Define $\rho_{\mathbb{S}, i^*} = \rho_{\mathbb{S}}$, $s_{i^*} = s$, and $x_{\Pi} = (c, x_{i^*}, s_{i^*})$.
Compute $\pi_{\Pi, \ell} \leftarrow \Pi.\text{P}(\text{crs}_{\Pi}, x_{\Pi}, w_{i^*})$. Define $\pi_{i^*} = (\rho_{\mathbb{S}, i^*}, s_{i^*}, \pi_{\Pi, i^*})$.
- (4) Define $\pi_\iota \leftarrow \text{P}(\text{crs}, x_\iota, w_\iota)$ for $\iota \in [k-1] \setminus \{i^*\}$.
- (5) Send $\{x_\iota, \pi_\iota\}_{\iota \in [k-1]}$ to \mathcal{A} .
- (6) Receive $\{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]}$ from \mathcal{A} .
- (7) Parse $\widetilde{\pi}_{j^*} = (\widetilde{\rho_{\mathbb{S}, j^*}}, \widetilde{s}_{j^*}, \widetilde{\pi_{\Pi, j^*}})$ and $\widetilde{\pi}_{\ell^*} = (\widetilde{\rho_{\mathbb{S}, \ell^*}}, \widetilde{s}_{\ell^*}, \widetilde{\pi_{\Pi, \ell^*}})$.
- (7) Send $(\rho_{\mathbb{S}, j^*}, \rho_{\mathbb{S}, \ell^*})$ to the challenger.

Given the event in Eq. (3) holds (for the afore mentioned fixed indices), then the reduction will return two quantum money states with the same serial number as the challenger sent. With advantage $1/(2k^3p(\lambda))$, the reduction will succeed at breaking unforgeability of the quantum money scheme, thus reaching a contradiction.

Scenario Two.

Alternatively, say that (for an infinite set of λ) \mathcal{A} does not generate two accepting proofs which have the same serial number as an honestly generated

proof with at least $1/(2p(\lambda))$ probability. By the pigeon-hole principle, this means that \mathcal{A} generates an accepting proof with a serial number which is not amongst the ones it received. In summary, we have that

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{P}(\text{crs}, x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathbf{V}(\text{crs}, x, \tilde{\pi}_\iota) = 1 \\ \text{and } \exists j^* \in \mathcal{J} \text{ s.t. } \tilde{s}_{j^*} \notin \{s_\iota\}_{\iota \in [k-1]} \end{array} \right] \geq \frac{1}{2p(\lambda)}. \quad (4)$$

Through an averaging argument, we can get a similar event with a fixed index j^* that belongs to the event's set \mathcal{J} with an advantage of $1/(2kp(\lambda))$. We will now switch to a hybrid where we provide \mathcal{A} with simulated proofs.

Claim (Claim 4.2). There exists a polynomial $q(\cdot)$ such that

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{Sim}_1(\text{crs}, \text{td}, x_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathbf{V}(\text{crs}, x, \tilde{\pi}_\iota) = 1 \\ \text{and } j^* \in \mathcal{J} \\ \text{and } \tilde{s}_{j^*} \notin \{s_\iota\}_{\iota \in [k-1]} \end{array} \right] \geq \frac{1}{q(\lambda)}. \quad (5)$$

We will later see a proof of Sect. 4.4. For now, assuming that this claim holds, by the definition of \mathcal{E} , Eq. (2), and Eq. (5), there exists a polynomial $q'(\cdot)$ such that

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{Sim}_1(\text{crs}, \text{td}, x_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]}) \\ j' \stackrel{\$}{\leftarrow} [k] \\ w \leftarrow \text{Ext}(\text{crs}, \text{td}, x, \tilde{\pi}_{j'})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathbf{V}(\text{crs}, x, \tilde{\pi}_\iota) = 1 \\ \text{and } j^* \in \mathcal{J} \\ \text{and } \tilde{s}_{j^*} \notin \{s_\iota\}_{\iota \in [k-1]} \\ \text{and } (x, w) \notin \mathcal{R} \end{array} \right] \geq \frac{1}{q'(\lambda)}.$$

We will additionally have that $j' = j^*$ with advantage at least $1/(kq'(\lambda))$. Since \mathbf{V} accepts $\tilde{\pi}_{j^*}$ with respect to x , $\Pi.\mathbf{V}$ must accept $\tilde{\pi}_{\Pi, j^*}$ with respect to $\tilde{x}_{\Pi, j^*} = (c, x, \tilde{s}_{j^*})$. Since $\tilde{s}_{j^*} \notin \{s_\iota\}_{\iota \in [k-1]}$, we have that $\Pi.\text{Sim}_1$, through Sim_1 , has not previously received \tilde{x}_{Π, j^*} as a query. As such, we have that

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Sim}_0(1^\lambda) \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \text{Sim}_1(\text{crs}, \text{td}, x_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]}) \\ w \leftarrow \text{Ext}(\text{crs}, \text{td}, \tilde{x}_{j^*}, \tilde{\pi}_{j^*})}} \left[\begin{array}{l} \Pi.\mathbf{V}(\text{crs}_\Pi, (c, x, \tilde{s}_{j^*}), \tilde{\pi}_{\Pi, j^*}) = 1 \\ \text{and } (c, x, \tilde{s}_{j^*}) \notin Q_\Pi \\ \text{and } (x, w) \notin \mathcal{R} \end{array} \right] \geq \frac{1}{kq'(\lambda)} \quad (6)$$

where Q_Π is the set of queries asked through Sim_1 to $\Pi.\text{Sim}_1$. We now define \mathcal{B} with oracle access to \mathcal{A} and Sim_1 ¹:

¹ Here, \mathcal{B} is given oracle access to Sim_1 which has the terms (crs, td) fixed by the output of Sim_0 .

Hardwired: $x_1, \dots, x_{k-1}, x_{j^*}$

Input: $\text{crs} = (\text{crs}_\Pi, c)$

- (1) For $\iota \in [k-1]$: send x_ι to Sim_1 , and receive π_ι from Sim_1 .
- (2) Send $(\text{crs}, \{x_\iota, \pi_\iota\}_{\iota \in [k-1]})$ to \mathcal{A} . Receive $\{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]}$ from \mathcal{A} .
- (3) Output $((c, x, \widetilde{s_{j^*}}, \widetilde{\pi_{j^*}})$.

Given that the event in Eq. (6) holds, then \mathcal{B} contradicts Sect. 4.4. Thus, all that remains to be proven is Sect. 4.4.

See the full version [33] for a proof of Claim 4.2.

By completing the proofs of our claim, we have concluding the proof of our theorem statement.

Corollary 3. *Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, and post-quantum iO exists, there exists a non-interactive adaptive argument of knowledge, adaptive computationally zero-knowledge, and $(k-1)$ -to- k -unclonable argument with extraction protocol for NP in the common reference string model (Definition 5).*

Proof. This follows from Theorem 5, Corollary 2, Theorem 9, and Theorem 12.

We have thus shown that Fig. 3 is an unclonable NIZK AoK in the CRS model as defined according to our proposed unclonability definition, Definition 5.

In the upcoming sections, we will consider unclonable proof systems in the QROM.

5 Unclonable NIZK in the Quantum Random Oracle Model

5.1 A Modified Sigma Protocol

We will begin by introducing a slightly modified sigma protocol. In the coming sections, our construction will involve applying Fiat-Shamir to this modified protocol.

Theorem 13. *Let a post-quantum sigma protocol with unpredictable commitments Π be given. Let \mathcal{R}_Π be an NP relation. Let $\mathcal{R} = \{((x, \mathcal{S}), w) : (x, w) \in \mathcal{R}_\Pi \wedge \mathcal{S} \neq \emptyset\}$. We argue that the following protocol will be a post-quantum sigma protocol with unpredictable commitments:*

- $\text{P.Com}(1^\lambda, (x, \mathcal{S}), w)$: Sends (x, α, s) to \mathcal{V} where $(\alpha, \text{st}) \leftarrow \Pi.\text{P.Com}(1^\lambda, x, w)$ and s is sampled from \mathcal{S} .
- $\text{V.Ch}(1^\lambda, (x, \mathcal{S}), (x, \alpha, s))$: Sends β to \mathcal{P} where $\beta \leftarrow \Pi.\text{V.Ch}(1^\lambda, x, \alpha)$.
- $\text{P.Com}(1^\lambda, (x, \mathcal{S}), w, \text{st}, \beta)$: Sends γ to \mathcal{V} where $\gamma \leftarrow \Pi.\text{P.Prove}(1^\lambda, x, w, \text{st}, \beta)$.
- $\text{V.Ver}(1^\lambda, (x, \mathcal{S}), (x, \alpha, s), \beta, \gamma)$: 1 iff $s \in \text{Support}(\mathcal{S})$ and $\Pi.\text{V.Ver}(1^\lambda, x, \alpha, \beta, \gamma) = 1$.

See the full version [33] for the proof of Theorem 13.

Corollary 4. *The Fiat-Shamir transform applied to the post-quantum sigma protocol defined in Theorem 13 yields a classical post-quantum NIZKAoK Π' in the QROM.*

Proof. This follows by Theorem 13 and Theorem 8.

5.2 Unclonability Definitions

Unclonable NIZKs in the quantum random oracle model are defined analogously to the CRS model – we repeat these definitions in the QRO model for completeness in the full version [33].

5.3 Unclonable NIZK Implies Public-Key Quantum Money Mini-Scheme in QROM

We defer the construction and proof to the full version [33]; below we state our results.

Theorem 14. *Let \mathcal{O} be a quantum random oracle. Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \text{NP}$. Let $\Pi = (\text{P}, \text{V})$ be a 1-to-2 unclonable non-interactive perfectly complete, computationally zero-knowledge protocol for \mathcal{L} in the QRO model.*

Then (P, V) implies a public-key quantum money mini-scheme in the QRO model.

5.4 Construction and Analysis of Unclonable-Extractable NIZK in QROM

We now introduce our construction in Fig. 4 and prove the main theorem of this section.

Theorem 15. *Let $k(\cdot)$ be a polynomial. Let NP relation \mathcal{R} with corresponding language \mathcal{L} be given.*

Let $(\text{NoteGen}, \text{Ver})$ be a public-key quantum money mini-scheme and $\Pi = (\text{P}, \text{V})$ be a post-quantum sigma protocol.

(P, V) as defined in Fig. 4 will be a non-interactive knowledge sound, computationally zero-knowledge, and $(k - 1)$ -to- k -unclonable argument with extraction protocol for \mathcal{L} in the quantum random oracle model.

Proof. Let the parameters and primitives be as given in the theorem statement. We argue that completeness follows from the protocol construction in Fig. 4, and we prove the remaining properties below.

See the full version [33] for complete proofs of argument of knowledge and zero-knowledge properties.

Unclonable NIZK for NP in the QROM

Let \mathcal{O} be a random oracle. Let $\Pi = (\mathsf{P} = (\mathsf{P.Com}, \mathsf{P.Prove}), \mathsf{V} = (\mathsf{V.Ch}, \mathsf{V.Ver}))$ be a post-quantum sigma protocol with unpredictable commitments, and $(\mathsf{NoteGen}, \mathsf{Ver})$ be a public-key quantum money mini-scheme. Let \mathcal{R} be the relation with respect to $\mathcal{L} \in \mathsf{NP}$.

PROVE $^{\mathcal{O}}(x, \omega)$:

- Compute a quantum note and associated serial number $(\rho_{\mathbb{S}}, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.
- Compute $(\alpha, \zeta) \leftarrow \mathsf{P.Com}(x, \omega)$.
- Query \mathcal{O} at (x, α, s) to get β .
- Compute $\gamma \leftarrow \mathsf{P.Prove}(x, \omega, \beta, \zeta)$.
- Output $\pi = (\rho_{\mathbb{S}}, s, \alpha, \beta, \gamma)$.

VERIFY $^{\mathcal{O}}(x, \pi)$:

- Check that $\mathsf{Ver}(\rho_{\mathbb{S}}, s)$ outputs 1.
- Check that \mathcal{O} outputs β when queried at (x, α, s) .
- Output the result of $\mathsf{V.Ver}(x, \alpha, \beta, \gamma)$.

Fig. 4. Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \mathsf{NP}$ in the Quantum Random Oracle Model

Let \mathcal{S} be the distribution of serial numbers as output by $\mathsf{NoteGen}(1^\lambda)$. We define Ext^2 with oracle-access to Ext_{FS} , \mathcal{O} , and some \mathcal{A} as follows:

Hardwired with: \mathcal{S} .

Input: x .

- (1) Given an oracle-query (x, α, s) from \mathcal{A} : send (x, α, s) to \mathcal{O} , receive β from \mathcal{O} , and send β to \mathcal{A} .
- (2) Upon receiving $\pi = (\rho_{\mathbb{S}}, s, \alpha, \beta, \gamma)$ from \mathcal{A} : send $\pi_{FS} = ((x, \alpha, s), \beta, \gamma)$ to Ext_{FS} .
- (3) Output the result of Ext_{FS} as w .

Let Sim_{FS} be the simulator for Π' in Corollary 4 (where Π instantiates Theorem 13). Let \mathcal{R}_{FS} be the relation for Π' with respect to \mathcal{R} . We define Sim with oracle-access to Sim_{FS} and program access to some random oracle \mathcal{O} as follows:
Input: x (ignores any witnesses it may receive).

- (1) Sample $(\rho_{\mathbb{S}}, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.
- (2) Let \mathcal{S} be the distribution where all probability mass is on s .
- (3) Compute $((x, \alpha, s), \beta, \gamma) \leftarrow \Pi.\mathsf{Sim}(x, \mathcal{S})$. Allow $\Pi.\mathsf{Sim}$ to program \mathcal{O} at (x, α, s) to return β .
- (5) Output $\pi = (\rho_{\mathbb{S}}, s, \alpha, \beta, \gamma)$.

² An extractor whose local code is implementable as a simple unitary which allows for straightforward rewinding.

Unclonable Extractability. Let Ext be the quantum circuit of the extractor we defined earlier (in our proof that Fig. 4 is an argument of knowledge). Let Sim be the quantum circuit of the simulator that we defined earlier (in our proof that Fig. 4 is a zero-knowledge protocol). We define a simulator for our extractor, SimExt , which interacts with some \mathcal{A} and has oracle-access to \mathcal{O} as follows:

Hardwired with: x_1, \dots, x_{k-1}, x

- (1) Compute $\pi_\iota \leftarrow \text{Sim}(x_\iota)$ for $\iota \in [k-1]$ where we store all points Sim would program into a list \mathcal{P} .
- (2) Send $\{x_\iota, \pi_\iota\}_{\iota \in [k-1]}$ to \mathcal{A} .
- (3) For every query from \mathcal{A} , if the query is in \mathcal{P} , then reply with the answer from \mathcal{P} . Else, forward the query to \mathcal{O} and send the answer back to \mathcal{A} .

We now define our extractor \mathcal{E} with oracle-access to some \mathcal{A} as follows:

Hardwired with: some choice of x_1, \dots, x_{k-1}, x .

- (1) Instantiates a simulatable and extractable random oracle \mathcal{O} . Runs Ext on \mathcal{O} throughout the interaction with \mathcal{A} (which may involve rewinding, in which case we would rewind \mathcal{A} and repeat the following steps).
- (2) Run $\text{SimExt}^{\mathcal{O}}(x_1, \dots, x_{k-1}, x)$ which interacts with \mathcal{A} .
- (3) Receive $\{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]}$ from \mathcal{A} .
- (4) Samples $\ell \in [k]$ uniformly at random. Send $\tilde{\pi}_\ell$ to Ext .
- (5) Outputs the result of Ext as w .

Let $\mathcal{A}, (x_1, w_1), \dots, (x_{k-1}, w_{k-1}) \in \mathcal{R}$, x , polynomial $p(\cdot)$, and negligible function $\text{negl}(\cdot)$ be given such that \mathcal{A} outputs more accepting proofs for x than \mathcal{A} received, and yet the extractor \mathcal{E} is unable to extract a valid witness for x from \mathcal{A} . Restated more formally, that is that

$$\Pr_{\substack{\mathcal{O} \\ \forall \iota \in [k-1], \pi_\iota \leftarrow \mathcal{P}^{\mathcal{O}}(x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}^{\mathcal{O}}(\{x_\iota, \pi_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathbf{V}^{\mathcal{O}}(x, \tilde{\pi}_\iota) = 1 \end{array} \right] \geq \frac{1}{p(\lambda)}, \quad (7)$$

and for all polynomials $p'(\cdot)$ (there are infinitely many λ) such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(x_1, \dots, x_{k-1}, x)} [(x, w) \in \mathcal{R}] \leq \frac{1}{p'(\lambda)}. \quad (8)$$

We parse the output of the adversary \mathcal{A} as $\tilde{\pi}_\iota = (\widetilde{\rho_{\mathbb{S}, \iota}}, \widetilde{s}_\iota, \widetilde{\alpha}_\iota, \widetilde{\beta}_\iota, \widetilde{\gamma}_\iota)$ for all $\iota \in [k]$.

Given Eq. (7), we may be in one of the two following cases: either \mathcal{A} generates two accepting proofs which have the same serial number as a honestly generated proof (for an infinite set of λ), or \mathcal{A} does not (for an infinite set of λ). We consider that either of these two scenarios occur with at least $1/(2p(\lambda))$ probability and show that each reaches a contradiction.

Scenario One.

Say that (for an infinite set of λ) \mathcal{A} generates two accepting proofs which have the same serial number as an honestly generated proof with at least $1/(2p(\lambda))$ probability. Symbolically,

$$\Pr_{\substack{\forall \iota \in [k-1], \pi_\iota \leftarrow \mathcal{P}^\mathcal{O}(x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}^\mathcal{O}(\{x_\iota, \pi_\iota\}_{\iota \in [k-1]}}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathcal{V}^\mathcal{O}(x, \tilde{\pi}_\iota) = 1 \\ \text{and } \exists i^* \in [k-1] \exists j^*, \ell^* \in \mathcal{J} \\ \text{s.t. } s_{i^*} = \widetilde{s}_{j^*} = \widetilde{s}_{\ell^*} \end{array} \right] \geq \frac{1}{2p(\lambda)}. \quad (9)$$

Through a hybrid argument, we can get a similar event with fixed indices i^* , j^* , and ℓ^* which belong to their respective sets with an advantage of $1/(2k^3p(\lambda))$. By using the advantage of \mathcal{A} in this game, we can show a reduction that breaks the unforgeability of the quantum money scheme. We will now outline this reduction.

Reduction: to unforgeability of quantum money scheme given oracle access to \mathcal{A} and \mathcal{O} .

Hardwired with: $(x_1, w_1), \dots, (x_{k-1}, w_{k-1}), x, i^*, j^*, \ell^*$.

- (1) Receive $(\rho_{\mathbb{S}}, s)$ from the challenger.
- (2) Define $\rho_{\mathbb{S}, i^*} = \rho_{\mathbb{S}}$ and $s_{i^*} = s$. Sample $(\rho_{\mathbb{S}, \iota}, s_\iota) \leftarrow \text{NoteGen}(1^\lambda)$ for $\iota \in [k-1] \setminus \{i^*\}$. Compute $(\alpha_\iota, \zeta_\iota) \leftarrow \text{II.P.Com}(x_\iota, w_\iota)$, query \mathcal{O} at $(x_\iota, \alpha_\iota, s_\iota)$ to get β_ι , compute $\gamma_\iota \leftarrow \text{II.P.Prove}(x_\iota, w_\iota, \beta_\iota, \zeta_\iota)$, and define $\pi_\iota = (\rho_{\mathbb{S}, \iota}, s_\iota, \alpha_\iota, \beta_\iota, \gamma_\iota)$ for $\iota \in [k-1]$.
- (3) Send $\{x_\iota, \pi_\iota\}_{\iota \in [k-1]}$ to \mathcal{A} .
- (4) Receive $\{\tilde{\pi}_\iota\}_{\iota \in [k]}$ from \mathcal{A} .
- (5) Send $(\widetilde{\rho}_{\mathbb{S}, j^*}, \widetilde{\rho}_{\mathbb{S}, \ell^*})$ to the challenger.

Given the event in Eq. (9) holds (for the afore mentioned fixed indices), then the reduction will return two quantum money states with the same serial number as the challenger sent. With advantage $1/(2k^3p(\lambda))$, the reduction will succeed at breaking unforgeability of the quantum money scheme, thus reaching a contradiction.

Scenario Two.

Alternatively, say that (for an infinite set of λ) \mathcal{A} does not generate two accepting proofs which have the same serial number as an honestly generated proof with at least $1/(2p(\lambda))$ probability. By the pigeon-hole principle, this means that \mathcal{A} generates an accepting proof with a serial number which is not amongst the ones it received. In summary, we have that

$$\Pr_{\substack{\forall \iota \in [k-1], \pi_\iota \leftarrow \mathcal{P}^\mathcal{O}(x_\iota, w_\iota) \\ \{\tilde{x}_\iota, \tilde{\pi}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}^\mathcal{O}(\{x_\iota, \pi_\iota\}_{\iota \in [k-1]}}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \iota \in \mathcal{J}, \mathcal{V}^\mathcal{O}(x, \tilde{\pi}_\iota) = 1 \\ \text{and } \exists j^* \in \mathcal{J} \text{ s.t. } \widetilde{s}_{j^*} \notin \{s_\iota\}_{\iota \in [k-1]} \end{array} \right] \geq \frac{1}{2p(\lambda)}. \quad (10)$$

Through an averaging argument, we can get a similar event with a fixed index j^* that belongs to the event's set \mathcal{J} with an advantage of $1/(2kp(\lambda))$. We will now switch to a hybrid where we provide \mathcal{A} with simulated proofs.

Claim (5.1). There exists a polynomial $q(\cdot)$ such that

$$\Pr_{\substack{\{\pi_\ell\}_{\ell \in [k-1]} \leftarrow \text{SimExt}^\mathcal{O}(x_1, \dots, x_{k-1}) \\ \{\tilde{x}_\ell, \tilde{\pi}_\ell\}_{\ell \in [k]} \leftarrow \mathcal{A}^{\text{SimExt}^\mathcal{O}}(\{x_\ell, \pi_\ell\}_{\ell \in [k-1]}}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : \tilde{x}_j = x\} \\ \text{s.t. } |\mathcal{J}| > |\{i : x_i = x\}| \\ \text{and } \forall \ell \in \mathcal{J}, \mathcal{V}^{\text{SimExt}^\mathcal{O}}(x, \tilde{\pi}_\ell) = 1 \\ \text{and } j^* \in \mathcal{J} \\ \text{and } \tilde{s}_{j^*} \notin \{s_\ell\}_{\ell \in [k-1]} \end{array} \right] \geq \frac{1}{q(\lambda)}. \quad (11)$$

We will later see a proof of Sect. 5.4. For now, assuming that this claim holds, we can define an adversary from which Ext can extract a valid witness for x .

Claim (5.2). There exists a polynomial $q'(\cdot)$ such that

$$\Pr_{w \leftarrow \mathcal{E}^\mathcal{A}(x_1, \dots, x_{k-1}, x)} [(x, w) \in \mathcal{R}] \geq \frac{1}{q'(\lambda)}. \quad (12)$$

We will soon see a proof for Sect. 5.4. Meanwhile, if this claim is true, then we will have a direct contradiction with Eq. (8). Thus, all that remains to be proven are the two claims.

See proof of Claim 5.1 and Claim 5.2 in the full version [33].

By completing the proofs of our claims, we have concluding the proof of our theorem statement.

Corollary 5. *Assuming the injective one-way functions exist, and post-quantum $i\mathcal{O}$ exists, there exists a non-interactive knowledge sound, computationally zero-knowledge, and $(k-1)$ -to- k -unclonable with extraction protocol for NP in the quantum random oracle model.*

Proof. This follows from Theorem 9 and Theorem 15.

We have thus shown that Fig. 4 is an unclonable NIZK AoK in the ROM model as defined according to our unclonability definition.

6 Applications

6.1 Unclonable Signatures of Knowledge

Definition 7 (Unclonable Extractable SimExt-secure Signatures of Knowledge). *Let NP relation \mathcal{R} with corresponding language \mathcal{L} be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$. Let a message space \mathcal{M} be given such that it can be indexed by a security parameter $\lambda \in \mathbb{N}$.*

(Setup, Sign, Verify) is an unclonable signature of knowledge of a witness with respect to \mathcal{L} and \mathcal{M} if it has the following properties:

- (Setup, Sign, Verify) is a quantum Sim-Ext signature of knowledge.

- **$(k - 1)$ -to- k -Unclonable with Extraction:** *There exists an oracle-aided polynomial-size quantum circuit \mathcal{E} such that for every polynomial-size quantum circuit \mathcal{A} , for every tuple of $k - 1$ instance-witness pairs $(x_1, \omega_1), \dots, (x_{k-1}, \omega_{k-1}) \in \mathcal{R}$, every $\{m_\iota \in \mathcal{M}_\lambda\}_{\iota \in [k-1]}$, for every (x, m) , if there is a polynomial $p(\cdot)$ where*

$$\Pr_{\substack{(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \forall \iota \in [k-1], \sigma_\iota \leftarrow \text{Sign}(\text{crs}, x_\iota, \omega_\iota, m_\iota) \\ \{\tilde{\sigma}_\iota\}_{\iota \in [k]} \leftarrow \mathcal{A}(\text{crs}, \{x_\iota, m_\iota, \sigma_\iota\}_{\iota \in [k-1]})}} \left[\begin{array}{l} \exists \mathcal{J} \subseteq \{j : (\tilde{x}_j, \tilde{m}_j) = (x, m)\} \\ \text{s.t. } |\mathcal{J}| > |\{i : (x_i, m_i) = (x, m)\}| \\ \text{and } \forall \iota \in \mathcal{J}, \text{Verify}(\text{crs}, x, m, \tilde{\sigma}_\iota) = 1 \end{array} \right] \geq \frac{1}{p(\lambda)},$$

then there is also a polynomial $q(\cdot)$ such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(\{x_\iota, m_\iota\}_{\iota \in [k-1]}, x, m)} [(x, w) \in \mathcal{R}] \geq \frac{1}{q(\lambda)}.$$

Unclonable Signature of Knowledge with CRS

Let $(\text{Setup}, \text{P}, \text{V})$ be non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for NP. Let \mathcal{R} be the relation with respect to $\mathcal{L} \in \text{NP}$.

SETUP (1^λ) : $(\text{crs}, \text{td}) \leftarrow \Pi.\text{Setup}(1^\lambda)$.

SIGN (crs, x, w, m) :

- Let $x_\Pi = (x, m)$ be an instance and $w_\Pi = w$ be its corresponding witness for the following language \mathcal{L}_Π :

$$\{(x, m) : \exists w : (x, w) \in \mathcal{R}\}.$$
- Compute $\pi_\Pi \leftarrow \Pi.\text{P}(\text{crs}, x_\Pi, w_\Pi)$.
- Output $\sigma = \pi_\Pi$.

VERIFY $(\text{crs}, x, m, \sigma)$: Output $\Pi.\text{V}(\text{crs}, (x, m), \pi_\Pi)$.

Fig. 5. Unclonable Signature of Knowledge in CRS model

Theorem 16. *Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for NP (Definition 5).*

(Setup, Sign, Verify) in Fig. 5 is an unclonable-extractable SimExt-secure signature of knowledge (Definition 7).

Corollary 6. *Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, post-quantum iO exists, there exists an unclonable SimExt-secure signature of knowledge (Definition 7).*

Proof. This follows from Corollary 3 and Theorem 16.

6.2 Revocable Anonymous Credentials

Definition 8 (Revocable Anonymous Credentials). $(\text{IssuerKeyGen}, \text{Issue}, \text{VerifyCred}, \text{Revoke}, \text{Prove}, \text{VerRevoke})$ is a revocable anonymous credentials scheme with respect to some set of accesses $\{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$ if it has the following properties:

- **Correctness:** For every sufficiently large $\lambda \in \mathbb{N}$, and every access $\in \mathcal{S}_\lambda$,

$$\Pr_{\substack{(\text{nym}, \text{sk}) \leftarrow \text{IssuerKeyGen}(1^\lambda) \\ \text{cred} \leftarrow \text{Issue}(1^\lambda, \text{nym}, \text{sk}, \text{access})}} [\text{VerifyCred}(1^\lambda, \text{nym}, \text{access}, \text{cred}) = 1] = 1$$

and

$$\Pr_{\substack{(\text{nym}, \text{sk}) \leftarrow \text{IssuerKeyGen}(1^\lambda) \\ \text{cred} \leftarrow \text{Issue}(1^\lambda, \text{nym}, \text{sk}, \text{access}) \\ \text{revnotice} \leftarrow \text{Revoke}(1^\lambda, \text{nym}, \text{sk}, \text{access}) \\ \pi \leftarrow \text{Prove}(1^\lambda, \text{nym}, \text{revnotice}, \text{cred})}} [\text{VerRevoke}(\text{nym}, \text{sk}, \text{access}, \text{revnotice}, \pi) = 1] = 1.$$

- **Revocation:** For every polynomial-size quantum circuit \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for sufficiently large $\lambda \in \mathbb{N}$, and every access $\in \mathcal{M}_\lambda$

$$\Pr_{\substack{(\text{nym}, \text{sk}) \leftarrow \text{IssuerKeyGen}(1^\lambda) \\ \text{cred} \leftarrow \text{Issue}(1^\lambda, \text{nym}, \text{sk}, \text{access}) \\ \text{revnotice} \leftarrow \text{Revoke}(1^\lambda, \text{nym}, \text{sk}, \text{access}) \\ \pi, \text{cred}' \leftarrow \mathcal{A}(1^\lambda, \text{nym}, \text{revnotice}, \text{cred})}} \left[\begin{array}{l} \text{VerRevoke}(1^\lambda, \text{nym}, \text{sk}, \text{access}, \text{revnotice}, \pi) = 1 \\ \wedge \text{VerifyCred}(1^\lambda, \text{nym}, \text{access}, \text{cred}') = 1 \end{array} \right] \leq \text{negl}(\lambda).$$

We now introduce a construction based on unclonable signatures of knowledge.

Theorem 17. Let $(\mathcal{X}, \mathcal{W})$ be a hard-distribution of instance and witness pairs for some NP relation. Let $\{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$ be some set of accesses. Let $(\text{Setup}, \text{Sign}, \text{Verify})$ be an unclonable-extractable *SimExt*-secure signature of knowledge for message space $\{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$ (Definition 7).

$(\text{IssuerKeyGen}, \text{Issue}, \text{VerifyCred}, \text{Revoke}, \text{Prove}, \text{VerRevoke})$ defined in Fig. 6 is a revocable anonymous credentials scheme (Definition 8).

Proof (Proof Sketch of Theorem 17). The correctness of this revocable anonymous credentials scheme follows from the correctness of the unclonable signature of knowledge scheme.

We will now sketch the proof of revocation. Say that there exists an adversary \mathcal{A} , access access , and polynomial $p(\cdot)$ such that, with probability at least $1/p(\lambda)$: (1) π passes the revocation check, and (2) cred' passes the credential check. This means that both π and cred' are valid signatures with respect to the same crs , x , and access that the signature cred was issued under. This satisfies the “if” condition of the unclonability property of the unclonable signature of knowledge.

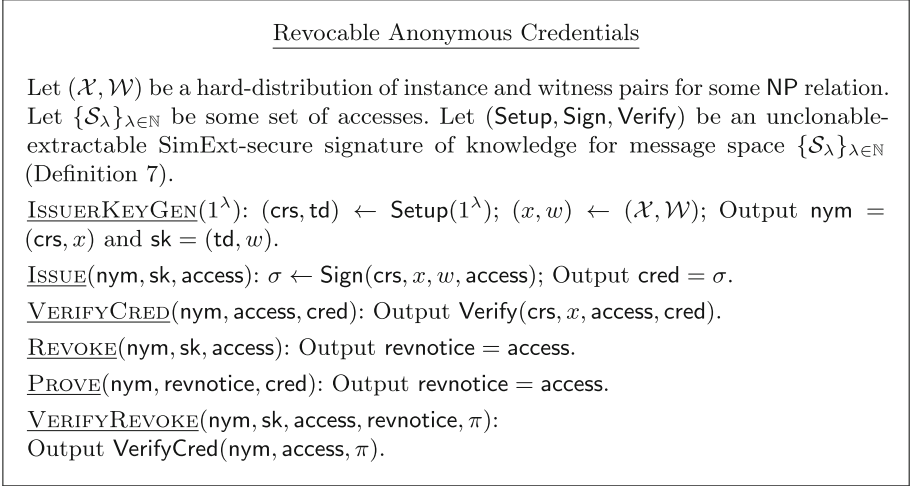


Fig. 6. Revocable Anonymous Credentials

As such, there exists a polynomial $q(\cdot)$ such that the unclonable signature of knowledge's extractor can produce a witness w for x with probability at least $1/q(\lambda)$. However, this contradicts the hardness of the distribution $(\mathcal{X}, \mathcal{W})$. Hence, our protocol must have the revocation property.

Corollary 7. *Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, post-quantum iO exists, and the hardness of NP, there exists a revocable anonymous credentials scheme (Definition 8).*

Proof. This follows from Corollary 6 and Theorem 17.

6.3 Unclonable Anonymous Credentials

We will show that our revocable anonymous credentials construction in Fig. 6 also satisfies a definition of unclonable anonymous credentials. We defer the definitions and proofs to the full version [33].

Theorem 18. *Let $(\mathcal{X}, \mathcal{W})$ be a hard-distribution of instance and witness pairs for some NP relation. Let $\{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$ be some set of accesses. Let $(\text{Setup}, \text{Sign}, \text{Verify})$ be an unclonable-extractable SimExt-secure signature of knowledge for message space $\{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$ (Definition 7).*

(IssuerKeyGen, Issue, VerifyCred) defined in Fig. 6 is an unclonable anonymous credentials scheme.

Corollary 8. *Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, post-quantum iO exists, and the hardness of NP, there exists an unclonable anonymous credentials scheme.*

Acknowledgments. The authors were supported in part by DARPA SIEVE, NSF QIS-2112890, NSF CAREER CNS-2238718, and NSF CNS-2247727. This material is based on work supported by DARPA under Contract No. HR001120C0024. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

References

1. Aaronson, S.: Quantum copy-protection and quantum money. In: Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009. pp. 229–242. IEEE Computer Society (2009). <https://doi.org/10.1109/CCC.2009.42>, <https://doi.org/10.1109/CCC.2009.42>
2. Aaronson, S., Christiano, P.F.: Quantum money from hidden subspaces. *Theory Comput.* **9**, 349–401 (2013). <https://doi.org/10.4086/toc.2013.v009a009>, <https://doi.org/10.4086/toc.2013.v009a009>
3. Aaronson, S., Liu, J., Liu, Q., Zhandry, M., Zhang, R.: New approaches for quantum copy-protection. In: Malkin, T., Peikert, C. (eds.) *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 12825, pp. 526–555. Springer (2021). https://doi.org/10.1007/978-3-030-84242-0_19, https://doi.org/10.1007/978-3-030-84242-0_19
4. Acar, T., Nguyen, L.: Revocation for delegatable anonymous credentials. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings. Lecture Notes in Computer Science*, vol. 6571, pp. 423–440. Springer (2011). https://doi.org/10.1007/978-3-642-19379-8_26, https://doi.org/10.1007/978-3-642-19379-8_26
5. Amos, R., Georgiou, M., Kiayias, A., Zhandry, M.: One-shot signatures and applications to hybrid quantum/classical authentication. In: Makarychev, K., Makarychev, Y., Tulsiani, M., Kamath, G., Chuzhoy, J. (eds.) *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*. pp. 255–268. ACM (2020). <https://doi.org/10.1145/3357713.3384304>, <https://doi.org/10.1145/3357713.3384304>
6. Ananth, P., Kaleoglu, F.: Unclonable encryption, revisited. In: Nissim, K., Waters, B. (eds.) *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 13042, pp. 299–329. Springer (2021). https://doi.org/10.1007/978-3-030-90459-3_11, https://doi.org/10.1007/978-3-030-90459-3_11
7. Ananth, P., Kaleoglu, F., Li, X., Liu, Q., Zhandry, M.: On the feasibility of unclonable encryption, and more. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II. Lecture Notes in Computer Science*, vol. 13508, pp. 212–241. Springer (2022). https://doi.org/10.1007/978-3-031-15979-4_8, https://doi.org/10.1007/978-3-031-15979-4_8
8. Ananth, P., Placa, R.L.L.: Secure software leasing. In: Canteaut, A., Standaert, F. (eds.) *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb,*

- Croatia, October 17-21, 2021, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12697, pp. 501–530. Springer (2021). https://doi.org/10.1007/978-3-030-77886-6_17, https://doi.org/10.1007/978-3-030-77886-6_17
9. Ananth, P., Poremba, A., Vaikuntanathan, V.: Revocable cryptography from learning with errors. In: Rothblum, G.N., Wee, H. (eds.) Theory of Cryptography - 21st International Conference, TCC 2023, Taipei, Taiwan, November 29 - December 2, 2023, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 14372, pp. 93–122. Springer (2023). https://doi.org/10.1007/978-3-031-48624-1_4, https://doi.org/10.1007/978-3-031-48624-1_4
 10. Barhoush, M., Salvail, L.: How to sign quantum messages (2023)
 11. Barhoush, M., Salvail, L.: Powerful primitives in the bounded quantum storage model (2023)
 12. Bartusek, J., Garg, S., Goyal, V., Khurana, D., Malavolta, G., Raizes, J., Roberts, B.: Obfuscation and outsourced computation with certified deletion. Cryptology ePrint Archive, Paper 2023/265 (2023), <https://eprint.iacr.org/2023/265>
 13. Bartusek, J., Khurana, D.: Cryptography with certified deletion. In: Crypto 2023 (to appear) (2023)
 14. Bartusek, J., Khurana, D., Poremba, A.: Publicly-verifiable deletion via target-collapsing functions. In: Crypto 2023 (to appear) (2023)
 15. Belenkiy, M., Camenisch, J., Chase, M., Kohlweiss, M., Lysyanskaya, A., Shacham, H.: Randomizable proofs and delegatable anonymous credentials. In: Halevi, S. (ed.) Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5677, pp. 108–125. Springer (2009). https://doi.org/10.1007/978-3-642-03356-8_7, https://doi.org/10.1007/978-3-642-03356-8_7
 16. Ben-David, S., Sattath, O.: Quantum tokens for digital signatures. CoRR **abs/1609.09047** (2016), <http://arxiv.org/abs/1609.09047>
 17. Ben-David, S., Sattath, O.: Quantum tokens for digital signatures. IACR Cryptol. ePrint Arch. p. 94 (2017), <http://eprint.iacr.org/2017/094>
 18. Broadbent, A., Islam, R.: Quantum encryption with certified deletion. In: Pass, R., Pietrzak, K. (eds.) Theory of Cryptography. pp. 92–122. Springer International Publishing, Cham (2020)
 19. Broadbent, A., Lord, S.: Uncloneable quantum encryption via oracles. In: Flammia, S.T. (ed.) 15th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2020, June 9-12, 2020, Riga, Latvia. LIPIcs, vol. 158, pp. 4:1–4:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPIcs.TQC.2020.4>, <https://doi.org/10.4230/LIPIcs.TQC.2020.4>
 20. Camenisch, J., Kohlweiss, M., Soriente, C.: Solving revocation with efficient update of anonymous credentials. In: Garay, J.A., Prisco, R.D. (eds.) Security and Cryptography for Networks, 7th International Conference, SCN 2010, Amalfi, Italy, September 13-15, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6280, pp. 454–471. Springer (2010). https://doi.org/10.1007/978-3-642-15317-4_28, https://doi.org/10.1007/978-3-642-15317-4_28
 21. Chase, M., Lysyanskaya, A.: On signatures of knowledge. In: Dwork, C. (ed.) Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4117, pp. 78–96. Springer (2006). https://doi.org/10.1007/11818175_5, https://doi.org/10.1007/11818175_5

22. Coiteux-Roy, X., Wolf, S.: Proving erasure. In: IEEE International Symposium on Information Theory, ISIT 2019, Paris, France, July 7-12, 2019. pp. 832–836 (2019). <https://doi.org/10.1109/ISIT.2019.8849661>, <https://doi.org/10.1109/ISIT.2019.8849661>
23. Coladangelo, A., Liu, J., Liu, Q., Zhandry, M.: Hidden cosets and applications to unclonable cryptography. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12825, pp. 556–584. Springer (2021). https://doi.org/10.1007/978-3-030-84242-0_20, https://doi.org/10.1007/978-3-030-84242-0_20
24. Farhi, E., Gosset, D., Hassidim, A., Lutomirski, A., Shor, P.W.: Quantum money from knots. In: Goldwasser, S. (ed.) Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012. pp. 276–289. ACM (2012). <https://doi.org/10.1145/2090236.2090260>, <https://doi.org/10.1145/2090236.2090260>
25. Fu, H., Miller, C.A.: Local randomness: Examples and application. Phys. Rev. A **97**, 032324 (Mar 2018). <https://doi.org/10.1103/PhysRevA.97.032324>, <https://link.aps.org/doi/10.1103/PhysRevA.97.032324>
26. Georgiou, M., Zhandry, M.: Unclonable decryption keys. IACR Cryptol. ePrint Arch. p. 877 (2020), <https://eprint.iacr.org/2020/877>
27. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. **18**(1), 186–208 (1989). <https://doi.org/10.1137/0218012>, <https://doi.org/10.1137/0218012>
28. Gottesman, D.: Uncloneable encryption. Quantum Inf. Comput. **3**(6), 581–602 (2003). <https://doi.org/10.26421/QIC3.6-2>, <https://doi.org/10.26421/QIC3.6-2>
29. Goyal, V., Malavolta, G., Raizes, J.: Unclonable commitments and proofs. IACR Cryptol. ePrint Arch. p. 1538 (2023), <https://eprint.iacr.org/2023/1538>
30. Hiroka, T., Morimae, T., Nishimaki, R., Yamakawa, T.: Quantum encryption with certified deletion, revisited: Public key, attribute-based, and classical communication. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021. pp. 606–636. Springer International Publishing, Cham (2021)
31. Hiroka, T., Morimae, T., Nishimaki, R., Yamakawa, T.: Certified everlasting zero-knowledge proof for QMA. CRYPTO (2022), <https://ia.cr/2021/1315>
32. IBM: Cost of a data breach report 2023. Tech. rep., IBM (2023)
33. Jawale, R., Khurana, D.: Unclonable non-interactive zero-knowledge. IACR Cryptol. ePrint Arch. p. 1532 (2023), <https://eprint.iacr.org/2023/1532>
34. Kane, D.M.: Quantum money from modular forms. CoRR **abs/1809.05925** (2018), <http://arxiv.org/abs/1809.05925>
35. Kitagawa, F., Nishimaki, R.: One-out-of-many unclonable cryptography: Definitions, constructions, and more. IACR Cryptol. ePrint Arch. p. 229 (2023), <https://eprint.iacr.org/2023/229>
36. Kundu, S., Tan, E.Y.Z.: Composably secure device-independent encryption with certified deletion (2020). <https://doi.org/10.48550/ARXIV.2011.12704>, <https://arxiv.org/abs/2011.12704>
37. Liu, Q., Zhandry, M.: Revisiting post-quantum fiat-shamir. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II. Lecture Notes in Computer Science, vol. 11693, pp. 326–355. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_12, https://doi.org/10.1007/978-3-030-26951-7_12

38. Lombardi, A., Schaeffer, L.: A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Paper 2019/279 (2019), <https://eprint.iacr.org/2019/279>, <https://eprint.iacr.org/2019/279>
39. Majenz, C., Schaffner, C., Tahmasbi, M.: Limitations on uncloneable encryption and simultaneous one-way-to-hiding. IACR Cryptol. ePrint Arch. p. 408 (2021), <https://eprint.iacr.org/2021/408>
40. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I. Lecture Notes in Computer Science, vol. 11692, pp. 89–114. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_4, https://doi.org/10.1007/978-3-030-26948-7_4
41. Poremba, A.: Quantum proofs of deletion for learning with errors. Cryptology ePrint Archive, Report 2022/295 (2022), <https://ia.cr/2022/295>
42. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: 40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA. pp. 543–553. IEEE Computer Society (1999). <https://doi.org/10.1109/SFCS.1999.814628>, <https://doi.org/10.1109/SFCS.1999.814628>
43. Santis, A.D., Crescenzo, G.D., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Kilian, J. (ed.) Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2139, pp. 566–598. Springer (2001). https://doi.org/10.1007/3-540-44647-8_33, https://doi.org/10.1007/3-540-44647-8_33
44. Santis, A.D., Crescenzo, G.D., Persiano, G.: Necessary and sufficient assumptions for non-iterative zero-knowledge proofs of knowledge for all NP relations. In: Montanari, U., Rolim, J.D.P., Welzl, E. (eds.) Automata, Languages and Programming, 27th International Colloquium, ICALP 2000, Geneva, Switzerland, July 9-15, 2000, Proceedings. Lecture Notes in Computer Science, vol. 1853, pp. 451–462. Springer (2000). https://doi.org/10.1007/3-540-45022-X_38, https://doi.org/10.1007/3-540-45022-X_38
45. Santis, A.D., Persiano, G.: Zero-knowledge proofs of knowledge without interaction (extended abstract). In: 33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992. pp. 427–436. IEEE Computer Society (1992). <https://doi.org/10.1109/SFCS.1992.267809>, <https://doi.org/10.1109/SFCS.1992.267809>
46. Unruh, D.: Revocable quantum timed-release encryption. In: Nguyen, P.Q., Oswald, E. (eds.) Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8441, pp. 129–146. Springer (2014). https://doi.org/10.1007/978-3-642-55220-5_8, https://doi.org/10.1007/978-3-642-55220-5_8
47. Unruh, D.: Post-quantum security of fiat-shamir. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10624, pp. 65–95. Springer (2017). https://doi.org/10.1007/978-3-319-70694-8_3, https://doi.org/10.1007/978-3-319-70694-8_3

48. Wiesner, S.: Conjugate coding. *SIGACT News* **15**(1), 78–88 (1983). <https://doi.org/10.1145/1008908.1008920>, <https://doi.org/10.1145/1008908.1008920>
49. Zhandry, M.: How to record quantum queries, and applications to quantum indistinguishability. In: Boldyreva, A., Micciancio, D. (eds.) *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II. *Lecture Notes in Computer Science*, vol. 11693, pp. 239–268. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_9, https://doi.org/10.1007/978-3-030-26951-7_9
50. Zhandry, M.: Quantum lightning never strikes the same state twice. In: Ishai, Y., Rijmen, V. (eds.) *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part III. *Lecture Notes in Computer Science*, vol. 11478, pp. 408–438. Springer (2019). https://doi.org/10.1007/978-3-030-17659-4_14, https://doi.org/10.1007/978-3-030-17659-4_14



Unclonable Secret Sharing

Prabhanjan Ananth¹, Vipul Goyal², Jiahui Liu³, and Qipeng Liu⁴

¹ University of California, Santa Barbara, Santa Barbara, USA
prabhanjan@cs.ucsb.edu

² NTT Research, Carnegie Mellon University, Pittsburgh, USA
vipul@cmu.edu

³ Massachusetts Institute of Technology, Cambridge, USA

⁴ University of California, San Diego, San Diego, USA

Abstract. Unclonable cryptography utilizes the principles of quantum mechanics to address cryptographic tasks that are impossible classically. We introduce a novel unclonable primitive in the context of secret sharing, called unclonable secret sharing (USS). In a USS scheme, there are n shareholders, each holding a share of a classical secret represented as a quantum state. They can recover the secret once all parties (or at least t parties) come together with their shares. Importantly, it should be infeasible to copy their own shares and send the copies to two non-communicating parties, enabling both of them to recover the secret.

Our work initiates a formal investigation into the realm of unclonable secret sharing, shedding light on its implications, constructions, and inherent limitations.

- **Connections:** We explore the connections between USS and other quantum cryptographic primitives such as unclonable encryption and position verification, showing the difficulties to achieve USS in different scenarios.
- **Limited Entanglement:** In the case where the adversarial shareholders do not share any entanglement or limited entanglement, we demonstrate information-theoretic constructions for USS.
- **Large Entanglement:** If we allow the adversarial shareholders to have unbounded entanglement resources (and unbounded computation), we prove that unclonable secret sharing is impossible. On the other hand, in the quantum random oracle model where the adversary can only make a bounded polynomial number of queries, we show a construction secure even with unbounded entanglement. Furthermore, even when these adversaries possess only a polynomial amount of entanglement resources, we establish that any unclonable secret sharing scheme with a reconstruction function implementable using Cliffords and logarithmically many T-gates is also unattainable.

1 Introduction

Alice is looking for storage for her sensitive data. She decides to hire multiple independent cloud providers and secret shares her data across them. Later

on, Alice retrieves these shares and reconstructs the data. Everything went as planned. However: what if the cloud providers keep a copy and sell shares of her data to her competitor, Bob? How can Alice make sure that once she retrieves her data, no one else can?

This is clearly impossible in the classical setting. The cloud providers can always keep a copy of the share locally and later, if Bob comes along, sell that copy to Bob. Nonetheless, this problem has been recently studied in the classical setting by a recent work of Goyal, Song, and Srinivasan [GSS21] who introduced the notion of traceable secret sharing (TSS). In TSS, if (a subset of) the cloud providers sell their shares to Bob, they cannot avoid leaving a cryptographic proof of fraud with Bob. Moreover, this cryptographic proof could not have been generated by Alice. Hence, (assuming Bob cooperates with Alice), Alice can sue the cloud providers in court and recover damages. Thus, TSS only acts as a deterrent and indeed, cannot stop the cloud providers from copying the secret.

However, in the quantum setting, the existence of no cloning theorem offers the tantalizing possibility that perhaps one may be able to build an “unclonable secret sharing” (USS) scheme. Very informally, the most basic version of a USS can be described as follows:

- Alice (the dealer) has a classical secret $m \in \{0, 1\}^*$. She hires n cloud providers $\mathcal{P}_1, \dots, \mathcal{P}_n$.
- Alice computes shares (ρ_1, \dots, ρ_n) , which is an n -partite state, from m and sends the share ρ_i to the party \mathcal{P}_i (note that Alice does not need to store any information like a cryptography key on her own).
- Given (ρ_1, \dots, ρ_n) , it is easy to recover m . But given any strict subset of the shares, no information about m can be deduced (i.e., it is an n -out-of- n secret sharing scheme).
- The most important is the unclonability. For every $i \in [n]$, the party \mathcal{P}_i computes a bipartite state $\sigma_{\mathbf{X}_i \mathbf{Y}_i}$. It sends the register \mathbf{X}_i to Bob and \mathbf{Y}_i to Charlie. Assuming that the message m was randomly chosen to be either m_0 or m_1 (where (m_0, m_1) is chosen adversarially), the probability that both Bob and Charlie can guess the correct message must be upper bounded by a quantity negligibly close to $\frac{1}{2}$.

In other words, the parties $\mathcal{P}_1, \dots, \mathcal{P}_n$ must be unable to locally clone their shares such that both sets of shares allow for reconstruction. Indeed, as we mentioned, this is the most basic version of USS. Even this basic setting has a practical significance: the servers which store Alice’s shares may not intentionally communicate her shares with each other, because they belong to companies with conflict of interest; but a malicious Bob may still buy a copy of Alice’s share from each of them.

One can consider more general settings where, e.g., we are interested in threshold (i.e., t -out-of- n) USS or, where a subset of the n parties might collude in attempting to clone their shares. One can also consider the setting where the parties $\mathcal{P}_1, \dots, \mathcal{P}_n$ share some entanglement (allowing them to use quantum teleportation).

Unclonable cryptography leverages the power of quantum information and empowers one to achieve primitives which are clearly impossible in classical cryptography. While a lot of efforts have been made towards various unclonable cryptographic primitives including but not limited to quantum money [BB20, AC12, Zha17, Shm22, LMZ23], copy-protection [Aar09, CLLZ21, AL20], tokenized signatures [BS16, CLLZ21, Shm22] and unclonable encryption (UE) [Got02, BL20, AK21, AKL+22, AKL23], the question of unclonable secret sharing had not been studied prior to our work. Secret sharing is one of the most fundamental primitives in cryptography and as such, we believe that studying unclonable secret sharing is an important step towards laying the foundation of unclonable cryptography. Our contribution lies in initiating a systematic study of USS.

Connection to Unclonable Encryption. The classical counterparts of unclonable encryption and (2-out-of-2) unclonable secret sharing are very similar. For instance, both one-time pad encryption and 2-out-of-2 secret sharing rely on the same ideas in the classical setting. One may wonder if UE and USS share similar a relation. UE resembles standard encryption with one additional property: now ciphertext is unclonable, meaning no one can duplicate a ciphertext into two parts such that both parts can be used separately to recover the original plaintext. At first glance, it might seem like UE directly implies a 2-out-of-2 USS. To secret share m , the dealer (Alice) would generate a secret key sk , and compute ciphertext ρ_{ct} , which encrypts the classical message m . One of the shares will be ρ_{ct} while the other will be sk . Since ρ_{ct} is unclonable, this may prevent two successful reconstructions of the original message.

However, the above intuition does not work if the two parties in (2-out-of-2) USS share entanglement. In UE, the ciphertext ρ_{ct} is a split into two components and sent to Alice and Bob. Later on, the secret key sk is sent (without any modification) to both Alice and Bob. However, in USS, the secret key sk corresponds to the second share and might also be split into two register such that one is sent to Alice and the other to Bob. This split could be done using a quantum register which is entangled with the quantum register used to split the cipher text ρ_{ct} . It is unclear if such an attack can be reduced to the UE setting, where there is no analog of such an entangled register. In fact, we show the opposite. We show that in some settings, USS implies UE, thus showing that USS could be a stronger primitive.

Connection to Instantaneous Non-local Computation. It turns out that the positive results on instantaneous non-local computation imply negative results on USS in specific settings. The problem of instantaneous non-local computation [Vai03, BK11, Spe15, IH08, GC19] is the following: Dave and Eve would like to compute a unitary U on a state $\rho_{\mathbf{X}\mathbf{Y}}$, where Dave has the register \mathbf{X} and Eve has the register \mathbf{Y} . They need to do so by just exchanging one message simultaneously with each other. Non-local computation has connections to the theory of quantum gravity, as demonstrated in some recent works [May19, May22]. Suppose there is a unitary U for which non-local computation is possible then this rules out a certain class of unclonable secret sharing schemes. Specifically, it

disallows certain reconstruction procedures that are functionally equivalent to U . In more detail, consider a USS scheme that is defined as follows: on input a message m , it produces shares on two registers \mathbf{X} and \mathbf{Y} . The reconstruction procedure¹ takes as input the shares and outputs m in both registers \mathbf{X} and \mathbf{Y} . Any non-local computation protocol for such a reconstruction procedure would violate the security of the USS scheme. Investigating both positive and negative results of USS schemes could shed more light on the feasibility of non-local computation. In this work, we adapt and generalize techniques used in the literature on non-local computation to obtain impossibility results for USS.

USS also has connections to position verification, a well-studied notion in quantum cryptography that has connections to problems in fundamental physics. We discuss this in the next section.

1.1 Our Results

In this work, our primary emphasis will be on n -out-of- n unclonable secret sharing schemes as even though they are the simplest, they give rise to numerous intriguing questions. Our results are twofold, as below.




Fig. 1. Relations between USS and UE in the information-theoretic regime.

Results on Information-Theoretic USS. We first examine the connections between USS and UE and constructions of UE in the information-theoretic regime. The first part of our results can be summarized by Fig. 1. In the figure, USS_1 stands for information-theoretic USS, secure against adversarial parties sharing *unbounded* amount of entanglement; we will explain why we call it USS_1 later on. We first show that, even if we restrict adversaries in USS_1 to have a polynomial amount of entanglement, it implies UE.

¹ In general, a reconstruction procedure need not output a copy of the secret twice but using CNOT gates, we can easily transform any reconstruction procedure into one that outputs two copies of the secret.

Theorem 1 (direction (a) in Fig. 1, Sect. 6.3). *Information-theoretic USS that is secure against adversarial parties \mathcal{P} sharing polynomial amount of entanglement implies UE.*

This leads us to ponder whether USS_1 and UE share equivalence, like their classical counterparts do. Perhaps surprisingly, we show that this connection is unlike to hold. We prove that USS_1 does not exist in the information-theoretic setting. Since there is no obvious evidence to refute UE in the IT setting and many candidates were proposed toward information-theoretic UE, our impossibility stands in sharp contrast to UE.

Theorem 2 (direction (b) in Fig. 1, Sect. 6.1). *Information-theoretic USS that is secure against adversarial parties \mathcal{P} sharing unbounded amount of entanglement with each other, does not exist.*

Facing the above impossibility, it seems like USS in the IT regime comes to a dead end. To overcome the infeasibility result, we investigate USS against adversarial parties with specific entanglement configurations. We consider the case where every pair of \mathcal{P}_i and \mathcal{P}_j either shares unbounded entanglement or shares no entanglement. In this case, we can define an entanglement graph, of which an edge (i, j) corresponds to entanglement between \mathcal{P}_i and \mathcal{P}_j . Then, we propose the natural generalization and define USS_d for any $d > 1$:

USS_d : Information-theoretic USS, secure against adversarial parties sharing entanglement whose entanglement graph has at least d connected components.

The above definition captures the case that there are d groups of parties; there is unlimited entanglement between parties in the same group and no entanglement between parties in different groups. This notation is not only for overcoming the barrier, but also has practical interest: parties from different groups are geographically separated or have conflict of interest, maintaining entanglement between them is either too expensive or impossible. Note that the characterization of entanglement is only for adversarial parties, whereas honest execution of the scheme does not need any pre-shared entanglement. We also like to note that aforementioned USS_1 is also captured by the above definition when $d = 1$.

It is easy to see that the existence of USS_d implies USS_{d+1} for any $d \geq 1$, as having less entanglement makes attacking more difficult. However, since USS_1 is impossible, can we construct USS_d for some d ? We complete the picture of USS and UE by presenting the following two theorems.

Theorem 3 (direction (c) in Fig. 1, Sect. 5.2). *UE implies USS_2 in the information-theoretic setting. As a corollary, it implies USS_d for any $d > 1$ in the IT setting.*

Theorem 4 (construction (d) in Fig. 1, Sect. 5.1). *USS_d exists for every $d = \omega(\log \lambda)$ in the information-theoretic setting, where λ is the security parameter.*

Along with Theorem 4, we proved a special XOR lemma of the well-known monogamy-of-entanglement property for BB84 states [BB20, TFKW13], when the splitting adversary is limited to tensor strategies. More precisely, we only consider cloning strategies that apply channels on each individual qubit, but never jointly on two or more qubits. Given a BB84 state, let $p(n)$ be the probability of the optimal tensor cloning strategy, that later two non-communicating parties recover the parity simultaneously. $p(1) = 1/2 + 1/2\sqrt{2}$ was proved in [TFKW13]. In this work, we show that $p(n) = 1/2 + \exp(-\Omega(n))$, which demonstrates a XOR hardness amplification for tensor strategies. We believe the proof of the theorem will be of independent interest, as a more general version of the theorem (that applies to any cloning strategies) will imply UE in the IT setting, resolving an open question on unclonable encryption since [BL20].

These two theorems establish a clear distinction between USS_1 and USS_d for all d greater than 1. Furthermore, the latter theorem illustrates that as the value of d becomes sufficiently large, it becomes feasible to achieve USS_d within the IT setting. Consequently, it implies that, at the very least, certain objectives outlined in Fig. 1 can be constructed.

Lastly, as the final arrow in Fig. 1, does USS_2 or $USS_{\omega(\log \lambda)}$ implies UE?

Remark 1 (direction (e) in Fig. 1). We do not have an answer yet. Nonetheless, we assert that either USS_d does not imply UE, or establishing this implication is as challenging as constructing UE. The latter assertion arises from our existing knowledge of $USS_{\omega(\log \lambda)}$ —demonstrating such an implication should, in turn, furnish us with a means to construct UE within the IT framework.

Results on Computational USS. In this computational regime, adversarial parties are computationally bounded; this in turn implies that the amount of pre-shared entanglement is also computationally bounded. Unlike the comprehensive picture presented in Fig. 1, our understanding here is more intricate. Specifically, as demonstrated in Fig. 2, the feasibility or infeasibility hinges on factors such as the computational complexity of USS schemes and the actual quantity of shared entanglement among malicious parties.

Similar to the IT setting, the implication of USS_1 and UE still works (direction (a) in Fig. 2). What is new here is that we present one impossibility result and one infeasibility result on USS_1 .

Theorem 5 (Informal, impossibility (f) in Fig. 2, Sect. 6.2). *USS whose reconstruction function has only d T gates, can be attacked with adversarial parties sharing $O(2^d)$ qubits of pre-shared entanglement.*

Therefore, when the reconstruction has low T complexity, say $d = \log \lambda$, then such USS does not exist even in the computational regime. Next, we present a construction, in sharp contrast to the impossibility above. Quantum random oracle [BDF+11], models the perfect (and unrealizable) cryptographic hash function. As it should behave as a truly random function, it can not have a small number of T gates.

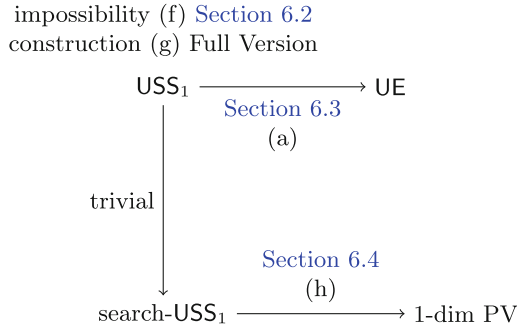


Fig. 2. Relations between USS and UE in the computational regime.

Theorem 6 (construction (g) in Fig. 2, Full Version). *USS that is secure against query-efficient adversarial parties sharing an arbitrary amount of pre-shared entanglement², exists in the quantum random oracle model (QROM).*

As quantum random oracle is not realizable in general, we wonder whether USS_1 can be constructed in the plain model. To the end, we show that USS_1 implies a cryptographic primitive called 1-dimensional position verification that is secure against parties sharing any polynomial amount of entanglement. Position verification represents an actively explored research area. Despite all the ongoing efforts, the development of a construction for position verification within the standard model remains elusive. This underscores the formidable challenge of devising USS_1 , when relying on computational assumptions.

Theorem 7 (direction (h) in Fig. 2, Sect. 6.4). *USS that is secure against adversarial parties having pre-shared entanglement, implies 1-dimensional position verification that is secure against parties sharing the same amount of pre-shared entanglement.*

1.2 Other Related Works

On Secret Sharing of Quantum States. Our work focuses on secret-sharing classical secrets by encoding them into a quantum state to achieve unclonability. One may be curious about the relationship of our new primitive to the existing studies on secret-sharing schemes where the secret messages are *quantum states* to begin with.

In short, all the existing quantum secret sharing schemes fall short of satisfying one crucial property in our model: the requirement of *no or low entanglement* for honest parties. Their unclonability also remains elusive, as they require much more complicated structures on quantum states than ours. We provide a detailed

² The adversary is polynomially bounded in queries but not in the pre-shared entanglement.

discussion below and will carefully incorporate all the discussions into the subsequent version.

In the paper, we consider a model where malicious parties can share some amount of entanglement before attacking the protocol. As illustrated in Fig. 1 and Fig. 2, the amount of entanglement (or more precisely, the entanglement graph) plays an important role in both the construction and barriers of such schemes. Therefore, we do not want the entanglement used in honest shares to scale to the same order or surpass what adversaries can access. Our constructions (Theorem 4 and Theorem 6) are based on unentangled quantum shares of single qubits, thus no entanglement required.

[HBB99] first proposed the idea of using quantum states to secret-share a classical bit. Their idea is to use n -qubit GHZ states for an n -out-of- n secret share scheme. However, an n -qubit GHZ state requires entanglement across n quantum registers, which enforces shareholders to maintain entanglement with each other. A subsequent proposal in [KKI99] followed a similar path but also required a large amount of entanglement. The idea of using quantum state to secret share classical secrets was also discussed by Gottesman [Got00], but they mostly focused on the lower bounds of general schemes (potentially requiring entanglement): for example, how many qubits are required to secret-share one classical bit.

There is another line of works on secret-sharing quantum secrets, including [CGL99, Smi00] and most recently [ÇGLR23] by Çakan et al. Since the goal is to secret-share a quantum state, entanglement is also necessary in these protocols.

2 Technical Overview

In this section, unless otherwise specified, we focus on 2-out-of-2 USS, with Share and Reconstruct. Share takes as input a message m and outputs two shares ρ_0, ρ_1 ; whereas Reconstruct takes two quantum shares and outputs a string. We assume ρ_0, ρ_1 are unentangled. When we consider impossibility results, all arguments mentioned in this overview carry in the same way to the general cases; for constructions, we only require unentangled shares.

2.1 USS₁ Implies UE, UE Implies USS₂

We first examine two directions (directions (a) and (c) in Figs. 1 and 2); that is, how USS₁ implies UE and how UE implies USS₂. We briefly recall the definition of UE: it is a secret key encryption scheme with the additional property: there is no way to split a quantum ciphertext into two parts, both combining with the classical secret key can recover the original plaintext (with probability at least $1/2$ plus negligible).

USS₁ *implies* UE, Sect. 6.3. Given a 2-out-of-2 USS, we now design a UE:

UE.Enc(k, m) takes as input a secret key k and a message,

1. it first produces two shares $(\rho_1, \rho_2) \leftarrow \text{USS.Share}(m)$,

2. it parses $k = (a, b)$, let the unclonable ciphertext be $\text{ct} = (\rho_1, X^a Z^b \rho_2 Z^b X^a)$. In other words, it sends out ρ_1 in clear, while having ρ_2 one-time padded by the key k .

Decryption is straightforward, by unpadding $X^a Z^b \rho_2 Z^b X^a$ and applying **Reconstruct** to (ρ_1, ρ_2) . Correctness and semantic security follows easily. Its unclonability can be based on the unclonability of USS_1 ; indeed, the scheme corresponds to a special strategy of malicious \mathcal{P}_1 and \mathcal{P}_2 . Suppose there exists an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ that violates the above scheme, there exists $(\mathcal{P}_1, \mathcal{P}_2, \mathcal{B}, \mathcal{C})$ that violates the security of USS_1 .

\mathcal{P}_1 and \mathcal{P}_2 share EPR pairs. \mathcal{P}_2 uses the EPR pairs to teleport ρ_2 to \mathcal{P}_1 , with \mathcal{P}_2 having random (a, b) and \mathcal{P}_1 obtaining $(\rho_1, X^a Z^b \rho_2 Z^b X^a)$. As \mathcal{P}_2 only has classical information, it sends (a, b) to both \mathcal{B} and \mathcal{C} , while \mathcal{P}_1 applies \mathcal{A} on $(\rho_1, X^a Z^b \rho_2 Z^b X^a)$ and shares the bipartite state with both \mathcal{B} and \mathcal{C} .

It is not hard to see that the above attacking strategy for USS_1 exactly corresponds to an attack in the UE we proposed above: \mathcal{P}_1 tries to split a ciphertext while \mathcal{P}_2 simply forwards the secret key $k = (a, b)$. Therefore, we can base the unclonability of the UE on that of USS_1 , which completes the first direction.

UE *implies* USS_2 , *Sect. 5.2*. Recall that 2-out-of-2 USS_2 describes adversarial parties who do not share any entanglement. We can simply set up our USS_2 scheme as follows, using UE:

Share(m) takes as input a message m , it samples a key k for UE, and let $|\text{ct}\rangle$ be the unclonable ciphertext of m under k ; the procedure **Share** outputs the first share as $\rho_1 = k$, and the second share as $\rho_2 = |\text{ct}\rangle$.

As there is no entanglement between \mathcal{P}_1 and \mathcal{P}_2 , \mathcal{P}_1 with $\rho_1 = k$ forwards the classical information to both Alice and Bob. In the meantime, \mathcal{P}_2 employs her cloning strategy, which remains entirely independent of the key k . Consequently, the unclonability of our USS_2 aligns with that of UE.

When we generalize the conclusion to n -out-of- n USS_2 , we first secret share the targeted message m into n shares. For any two adjacent parties $\mathcal{P}_i, \mathcal{P}_{i+1}$ and the i -th share, the first part receives the key and the second one gets the unclonable ciphertext. As long as all the malicious parties form at least two connected components (as defined in USS_2), there must be two adjacent parties who do not have entanglement. Thus, we can incur the same logic to prove its unclonability, basing on the unclonability of UE.

2.2 Construction of $\text{USS}_{\omega(\log \lambda)}$

For simplicity, we focus on an n -out-of- n USS, where $n = \omega(\log \lambda)$ and no entanglement is shared between any malicious parties, which is a special case of a general n -out-of- n $\text{USS}_{\omega(\log \lambda)}$, for a larger $n \gg \omega(\log \lambda)$. Our construction is based on the BB84 states. Our scheme first classically secret-shares m into $(n-1)$ shares and encodes each classical share into a single-qubit BB84 state. One party will receive the basis information θ which contains $(n-1)$ basis; every other party will receive a BB84 state for the i -th classical share.

Share(m): it takes as input a secret $m \in \{0, 1\}$,

- it samples m_1, \dots, m_{n-1} conditioned on their parity equals to m ;
- it samples $\theta \in \{0, 1\}^{n-1}$;
- let the first $(n - 1)$ shares be $\rho_i = H^{\theta_i} |m_i\rangle \langle m_i| H^{\theta_i}$ and the last share $\rho_n = |\theta\rangle \langle \theta|$.

Reconstruction of shares is straightforward. After receiving all shares, one uses the basis information θ to recover all the classical shares m_i ; m then is clearly determined by these m_i .

To reason about the unclonability of our protocol, we first recall a theorem on BB84 states, initially proposed by Tomamichel, Fehr, Kaniewski and Wehner [TFKW13] and later adapted in constructing unclonable encryption by Broadbent and Lord [BL20]. We start by considering a cloning game of single-qubit BB84 states.

1. \mathcal{A} receives $H^\theta |x\rangle \langle x| H^\theta$ for uniformly random $x, \theta \in \{0, 1\}$, it applies a channel and produces $\sigma_{\mathbf{BC}}$. Bob and Charlie receive their registers accordingly.
2. Bob \mathcal{B} and Charlie \mathcal{C} apply their POVMs and try to recover x ; they win if and only if both guess x correctly.

Lemma 1 (Corollary 2 when $n = 1$, [BL20]). *No (unbounded) quantum $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins the above game with probability more than 0.855.*

Tomamichel, Fehr, Kaniewski and Wehner [TFKW13] and Broadbent and Lord [BL20] studied parallel repetitions of the above cloning game³. In the parallel repetition, n random and independent BB84 states are generated, which encode an n -bit string x . The goal of cloning algorithms is to guess the n -bit string x simultaneously. They showed that the cloning game follows parallel repetition, meaning that the optimal winning probability in an n -fold parallel repetition game is at most $(0.855)^n$.

Our proposed scheme also prepares these BB84 states in parallel, but hides the secret m as the XOR of the longer secret. Indeed, the XOR repetition of the BB84 cloning game has been a folklore and was considered as a candidate for UE. More specifically, it is conjectured that the following game can not be won by any algorithm with probability more than $1/2 + \exp(-\Omega(n))$:

XOR repetition of BB84 cloning games.

1. \mathcal{A} receives $H^\theta |x\rangle \langle x| H^\theta$ for uniformly random $x, \theta \in \{0, 1\}^n$, it applies a channel and produces $\sigma_{\mathbf{BC}}$. Bob and Charlie receive their register accordingly.
2. Bob \mathcal{B} and Charlie \mathcal{C} apply their POVMs and try to recover $\text{parity}(x)$; they win if and only if both guess correctly.

Although there is no evidence to disprove the bound for the XOR repetition so far, the validity of the bound still remains unknown. In this work, we prove this

³ Indeed, [TFKW13] proved a stronger statement on a different game, which ultimately implied the parallel repetition theorem, shown by [BL20].

bound, when \mathcal{A} is restricted to a collection of strategies. It applies \mathcal{C}_i on the i -th qubit of the BB84 state and get $\sigma_{\mathbf{BC}}^{(i)}$; the final state $\sigma_{\mathbf{BC}} = \bigotimes_i \sigma_{\mathbf{BC}}^{(i)}$. Note that the lemma does not put any constraint on the behaviors of \mathcal{B} or \mathcal{C} .

Lemma 2 (An XOR lemma for BB84 cloning games, Sect. 5.1). *When \mathcal{A} only applies a tensor cloning strategy to prepare $\sigma_{\mathbf{BC}}$, the optimal success probability in the XOR repetition of BB84 games is $1/2 + \exp(-\Omega(n))$.*

Equipped with it, it is straightforward to show the unclonability of our protocol.

A proof for the XOR repetition. Finally, we give a brief recap on the proof for Lemma 2.

For any \mathcal{A} 's tensor strategy with channels \mathcal{C}_i applied on the i -th qubit of a BB84 state, we recall the notation $\sigma_{\mathbf{BC}}^{(i)}$. This is the state produced from the i -th qubit of the BB84 state, when θ_i, x_i was sampled uniformly at random. Let $\sigma_{\mathbf{B}}^{(i,0)}$ be the density matrix, describing the register that will be given to Bob, when $x_i = 0$. We can similarly define $\sigma_{\mathbf{B}}^{(i,1)}$, $\sigma_{\mathbf{C}}^{(i,0)}$ and $\sigma_{\mathbf{C}}^{(i,1)}$. Lemma 1 tells us that, there exists a constant $c > 0$, either

$$\text{TD}(\sigma_{\mathbf{B}}^{(i,0)}, \sigma_{\mathbf{B}}^{(i,1)}) < c \quad \text{or} \quad \text{TD}(\sigma_{\mathbf{C}}^{(i,0)}, \sigma_{\mathbf{C}}^{(i,1)}) < c.$$

This indicates that for every i , either Bob or Charlie can not perfectly tell the value of x_i , regardless of the channel \mathcal{C}_i . Furthermore, as the BB84 state has n qubits, w.l.o.g. we can assume that the above holds for Bob, for at least $n/2$ positions.

In the XOR repetition, Bob eventually will receive $\sigma_{\mathbf{B}}^{(i,m_i)}$. We show that Bob can not tell whether the parity of all m_i is odd or even. More precisely, we will show:

$$\text{TD} \left(\sum_{\substack{m_1, \dots, m_{n-1}: \\ \oplus_i m_i = 0}} \frac{1}{2^{n-2}} \left(\bigotimes_i \sigma_{\mathbf{B}}^{(i,m_i)} \right), \sum_{\substack{m_1, \dots, m_{n-1}: \\ \oplus_i m_i = 1}} \frac{1}{2^{n-2}} \left(\bigotimes_i \sigma_{\mathbf{B}}^{(i,m_i)} \right) \right) \leq c^{n/2}.$$

We connect the trace distance directly to the trace distance of *each pair of states* $\text{TD}(\sigma_{\mathbf{B}}^{(i,0)}, \sigma_{\mathbf{B}}^{(i,1)})$ and demonstrate *an equality* (see Sect. 5.1):

$$\begin{aligned} & \text{TD} \left(\sum_{\substack{m_1, \dots, m_{n-1}: \\ \oplus_i m_i = 0}} \frac{1}{2^{n-2}} \left(\bigotimes_i \sigma_{\mathbf{B}}^{(i,m_i)} \right), \sum_{\substack{m_1, \dots, m_{n-1}: \\ \oplus_i m_i = 1}} \frac{1}{2^{n-2}} \left(\bigotimes_i \sigma_{\mathbf{B}}^{(i,m_i)} \right) \right) \\ &= \prod_i \text{TD} \left(\sigma_{\mathbf{B}}^{(i,0)}, \sigma_{\mathbf{B}}^{(i,1)} \right). \end{aligned}$$

Since every trace distance is bounded by 1 and there are at least $n/2$ terms in the product smaller than c , we conclude the result.

2.3 Impossibility of USS_1

Since USS_1 implies UE, it is natural to consider building UE from USS_1 . Constructing UE in the basic model remained unresolved since [BL20]. Perhaps the connections in the last section provide a new avenue for constructing UE. In this section, we present two impossibility results (referred to as (b) in Fig. 1 and (f) in Fig. 2) that highlight challenges associated with USS_1 .

Information-theoretic USS_1 does not exist, Sect. 6.1. We begin by examining the case of 2-out-of-2 USS_1 with unentangled shares, and our impossibility result extends to the general case. Let us consider two malicious parties, \mathcal{P}_1 and \mathcal{P}_2 , who share an unlimited amount of entanglement. \mathcal{P}_2 receives the initial share, ρ_2 , and teleports it to \mathcal{P}_1 . This action leaves \mathcal{P}_2 with a random one-time pad key, denoted as (a, b) while \mathcal{P}_1 now possesses $(\rho_1, X^a Z^b \rho_2 Z^b X^a)$. Now, \mathcal{P}_1 aims to jointly apply the reconstruction procedure to (ρ_1, ρ_2) , but there's a problem: \mathcal{P}_1 lacks all the necessary information, especially the one-time padded key. To address this challenge, we recall the concept of port-based teleportation [IH08, BK11] to help \mathcal{P}_1 .

Port-based teleportation allows one party to teleport a d -qubit quantum state to another party, while leaving the state in plain. This is certainly impossible without paying any cost, as it contradicts with special relativity. Two parties need to pre-share about $O(d2^d)$ EPR pairs, divided into $O(2^d)$ blocks of d qubits. After the port-based teleportation, the teleported state will be randomly dropped into one of the blocks of \mathcal{P}_2 , while only \mathcal{P}_1 has the classical information about which block consists of the original state.

Equipped with port-based teleportation, \mathcal{P}_1 teleports $(\rho_1, X^a Z^b \rho_2 Z^b X^a)$ to \mathcal{P}_2 ; it has the classical information and specifying the location of the teleported state. \mathcal{P}_2 then runs $\text{Reconstruct} \circ (I \otimes Z^b X^a)$ on every possible block among the pre-shared entanglement, yielding $O(2^d)$ different values; even though most of the execution is useless, the ind-th block will store the correct (classical) answer. Finally, both \mathcal{P}_1 and \mathcal{P}_2 sends all their classical information to Alice and Bob; each of them can independently determine the message. This clearly violates the unclonability of USS_1 . Thus, for any 2-out-of-2 USS_1 whose shares are of length d , there is an attacking strategy that takes time and entanglement of order $\tilde{O}(d2^d)$ and completely breaks its unclonability.

We refer readers to Sect. 6.1 for the proof of a general theorem statement.

Impossibility of computationally secure USS_1 , with low- T Reconstruct , Sect. 6.2. We now focus on the case when the reconstruction circuit can be implemented by Clifford gates and logarithmically many T gates. Denote C to be the reconstruction circuit. That is, on input two shares of the form ρ_1, ρ_2 , the output is the first bit of $C(\rho_1 \otimes \rho_2)C^\dagger = |m\rangle \langle m| \otimes \tau$.

We let \mathcal{P}_2 teleport ρ_2 to \mathcal{P}_1 and they try to compute Reconstruct in a non-local manner. In the previous attack, this is done by leveraging an exponential amount of entanglement. To avoid this and make the attack efficient, we hope that \mathcal{P}_1 can homomorphically compute on the one-time padded data $(\rho_1, X^a Z^b \rho_2 Z^b X^a)$, without decrypting it.

Suppose C is a Clifford circuit. We use the fact that the Clifford group is a normalizer for the Pauli group (specifically, the $X^a Z^b$ operator). Let us assume each ρ_1, ρ_2 is of ℓ qubits. In other words, for any $a, b \in \{0, 1\}^\ell$ and Clifford circuit C , there exists a polynomial-time computable $a', b' \in \{0, 1\}^{2\ell}$ depending only on a, b and C , such that

$$C(\rho_1 \otimes X^a Z^b \rho_2 Z^b X^a)C^\dagger = X^{a'} Z^{b'} C(\rho_1 \otimes \rho_2)C^\dagger Z^{b'} X^{a'}.$$

Here a', b' act as a bigger quantum one-time pad operated on $C(\rho_1 \otimes \rho_2)C^\dagger = |m\rangle \langle m| \otimes \tau$.

Now \mathcal{P}_1 measures the first qubit in the computational basis, yielding $m \oplus a'_1$; whereas \mathcal{P}_2 compute a', b' (and most importantly, a'_1) from its classical information a, b . They send their knowledge to both Alice and Bob, who later simultaneously recover m .

Next, let us consider the more general case where C consists of Clifford gates and t number of T gates. The homomorphic evaluation of Clifford gates are as before. However, the homomorphic evaluation of T gates are handled differently.

Let us consider one single T gate that applies to the first qubit. We consider two identities, for any $x, z \in \{0, 1\}$ and any single-qubit state $|\psi\rangle$

$$\begin{aligned} (i) \quad T(X^x Z^z) |\psi\rangle &= (X^x Z^{x \oplus z} P^x) T |\psi\rangle, \\ (ii) \quad P(X^x Z^z) |\psi\rangle &= (X^x Z^{x \oplus z}) P |\psi\rangle \end{aligned}$$

Suppose, the current state is of the form $X^x Z^z |\psi\rangle$ and we apply $P^x T$ to the state. We would like to show that the resulting state is $X^{a'} Z^{b'} T |\psi\rangle$ for some $a' \in \{0, 1\}, b' \in \{0, 1\}$. We use the above identities:

$$(P^x T)(X^x Z^z) |\psi\rangle \stackrel{\text{From (i)}}{=} P^x (X^x Z^{x \oplus z} P^x) T |\psi\rangle \stackrel{\text{From (ii)}}{=} X^x Z^{x \oplus z} P^{x \oplus x} T |\psi\rangle.$$

Note that $P^2 = P^0 = I$. Thus, if we can learn x ahead, we can successfully homomorphic compute T on the encrypted data. However, in our case, x corresponds to any bit in the one-time pad key a of any stage. \mathcal{P}_1 has no way to learn x . This is where the limitation of low-T gate comes from. Instead of knowing x ahead, each time when a T homomorphic evaluation is needed, one simply guesses x' ; as long as $x = x'$ (which happens with probability $1/2$), we succeed. Thus, \mathcal{P}_1 only guesses all x 's (for each T gate) correctly with probability 2^{-t} . If t is logarithmic, our attack violates the security with inverse polynomial probability; therefore, it rules out computationally secure USS_1 with a low-T Reconstruct procedure.

2.4 Barriers of USS_1 (Implication of PV)

To further demonstrate the challenge of building USS against entangled adversaries, we show that 2-party USS_1 implies a primitive called position verification. Position verification (PV) has remained a vexing problem since its inception [CGMO09].

We briefly introduce the notion of position verification for the 1-dimensional setting: two verifiers on a line will send messages to a prover who claims to be

located at a position between the two verifiers. By computing a function of the verifiers' messages and returning the answers to the verifiers in time, the prover ensures them of its location. However, two malicious provers may collude to impersonate such an honest verifier by standing at the two sides of the claimed position.

We demonstrate that 2-party USS_1 , even with the weaker search-based security, will imply PV: the two verifiers in the position verification protocol will generate secret shares (ρ_0, ρ_1) of a random string s ; then they will each send the messages ρ_0 and ρ_1 respectively to the prover; the prover needs to reconstruct s and send s to both verifiers in time. Any attack against PV can be viewed as a two-stage strategy—one can perfectly turn the first-stage strategy in PV into the shareholders' strategy in USS and the second-stage strategy in PV into the recoverers' strategy in USS.

Despite many efforts, progress on PV in the computational setting against entangled adversaries has unfortunately been slow. We do not even know of any secure computational PV against adversaries with unbounded polynomial amount of entanglement in the plain model, nor any impossibility result. Moreover, some recent advancement in quantum gravity has unveiled some connections between the security of position verification and problems in quantum gravity [May19, May22].

Any progress of USS_1 in the plain model will contribute towards resolving this long-standing open problem and unveil more implications.

3 Preliminaries

3.1 Notations

We assume that the reader is familiar with the basic background from [NC10]. The Hilbert spaces we are interested in are \mathbb{C}^d , for $d \in \mathbb{N}$. We denote the quantum registers with capital bold letters \mathbf{R} , \mathbf{W} , \mathbf{X} , We abuse the notation and use registers in place of the Hilbert spaces they represent. The set of all linear mappings from \mathbf{R} to \mathbf{W} is denoted by $L(\mathbf{R}, \mathbf{W})$, and $L(\mathbf{R})$ denotes $L(\mathbf{R}, \mathbf{R})$. We denote unitaries with capital letters C , E , ... and the set of unitaries on register \mathbf{R} with $U(\mathbf{R})$. We denote the identity operator on \mathbf{R} with $\mathbb{I}_{\mathbf{R}}$; if the register \mathbf{R} is clear from the context, we drop the subscript \mathbf{R} from the notation $\mathbb{I}_{\mathbf{R}}$. We denote the set of all positive semi-definite linear mappings in $L(\mathbf{R}, \mathbf{R})$ with trace 1 (i.e., set of all valid quantum states) by $D(\mathbf{R})$. For a register \mathbf{R} in a multi-qubit system, we denote $\overline{\mathbf{R}}$ to be a register consisting of all the qubits in the system not contained in \mathbf{R} . We denote $\text{Tr}_{\mathbf{R}}(\rho)$ to be the state obtained by tracing out all the registers of ρ except \mathbf{R} . A quantum channel Φ refers to a completely positive and trace-preserving (CPTP) map from a Hilbert space \mathcal{H}_1 to a possibly different Hilbert space \mathcal{H}_2 .

3.2 Unclonable Encryption

Unclonable encryption was originally defined in [BL20] and they considered two security notions, namely search and indistinguishability security, with the latter

being stronger than the former. We consider below a mild strengthening of the indistinguishability security due to [AK21].

Definition 1. *An unclonable encryption scheme UE is a triple of efficient quantum algorithms (UE.KeyGen, UE.Enc, UE.Dec) with the following procedures:*

- $\text{KeyGen}(1^\lambda)$: On input a security parameter 1^λ , returns a classical key sk ⁴.
- $\text{Enc}(\text{sk}, m)$: It takes the key sk and the message m for $m \in \{0, 1\}^{\text{poly}(\lambda)}$ as input and outputs a quantum ciphertext ρ_{ct} .
- $\text{Dec}(\text{sk}, \rho_{ct})$: It takes the key sk and the quantum ciphertext ρ_{ct} , it outputs a quantum state τ .

Correctness. The following must hold for the encryption scheme. For every $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$ and every message m , we must have $\text{Tr}[[|m\rangle\langle m| \text{Dec}(\text{sk}, \text{Enc}(\text{sk}, |m\rangle\langle m|))] \geq 1 - \text{negl}(\lambda)$.

Unclonability. In the rest of the work, we focus on unclonable IND-CPA security. The regular IND-CPA security follows directly from its unclonable IND-CPA security. To define unclonable security, we introduce the following security game.

Definition 2 (Unclonable IND-CPA game). *Let $\lambda \in \mathbb{N}^+$. Consider the following game against the adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$.*

- The adversary \mathcal{A} generates $m_0, m_1 \in \{0, 1\}^{n(\lambda)}$ and sends (m_0, m_1) to the challenger.
- The challenger randomly chooses a bit $b \in \{0, 1\}$ and returns $\text{Enc}(\text{sk}, m_b)$ to \mathcal{A} . \mathcal{A} produces a quantum state $\rho_{\mathcal{BC}}$ on registers \mathcal{B} and \mathcal{C} , and sends the corresponding registers to \mathcal{B} and \mathcal{C} .
- \mathcal{B} and \mathcal{C} receive the key sk , and output bits $b_{\mathcal{B}}$ and $b_{\mathcal{C}}$ respectively.

The adversary wins if $b_{\mathcal{B}} = b_{\mathcal{C}} = b$.

We denote the success probability of the above game by $\text{adv}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}(\lambda)$. We say that the scheme is information-theoretically (resp., computationally) secure if for all (resp., quantum polynomial-time) adversaries $(\mathcal{A}, \mathcal{B}, \mathcal{C})$,

$$\text{adv}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}(\lambda) \leq 1/2 + \text{negl}(\lambda).$$

4 Definitions and Notations

4.1 Unclonable Secret Sharing

An (t, n) -unclonable secret sharing scheme, associated with n parties $\mathcal{P}_1, \dots, \mathcal{P}_n$, consists of the following QPT algorithms:

⁴ In our construction, we require sk being a uniform random string. Such a UE scheme can be constructed in QROM [AKL+22, AKL23].

- $\text{Share}(1^\lambda, 1^n, 1^t, m) \rightarrow \rho_{\mathbf{R}_1 \mathbf{R}_2 \dots \mathbf{R}_n}$: On input security parameter λ , n parties, a secret $m \in \{0, 1\}^*$, output registers $\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_n$.
- $\text{Reconstruct}(\rho_{\mathbf{R}'_{i_1}}, \dots, \rho_{\mathbf{R}'_{i_t}})$: On input shares $\mathbf{R}'_{i_1}, \dots, \mathbf{R}'_{i_t}$, output a secret \hat{m} .

When it is an n -out-of- n USS scheme, we ignore the input 1^t in Share . In the rest of the work, we will focus on constructions with unentangled shares and impossibility results for entangled shares. For sake of clarity, we will use ρ_1, \dots, ρ_n to denote these shares. We require the following properties to hold.

Correctness. We can recover the secret with probability (almost) 1, more formally:

$$\Pr[\text{Reconstruct}(\rho_{i_1}, \dots, \rho_{i_k}) = m \mid (\rho_1, \dots, \rho_n) \leftarrow \text{Share}(1^\lambda, 1^n, m) \cap k \geq t] = 1 - \text{negl}(\lambda).$$

Soundness/Privacy. Given (at most) $(t-1)$ shares, it is information-theoretically impossible/computationally hard to recover the original message. Formally, for any unbounded/QPT \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, for every $m \in \{0, 1\}$, for every $\lambda > 0$, $i_1, \dots, i_{t-1} \in [n]$,

$$\Pr[\mathcal{A}(\rho_{i_1}, \dots, \rho_{i_{t-1}}) = m \mid (\rho_1, \dots, \rho_n) \leftarrow \text{Share}(1^\lambda, 1^n, m)] = \frac{1}{2} + \text{negl}(\lambda).$$

All our schemes satisfy information-theoretic soundness/privacy.

4.2 Indistinguishability-Based Security

In this work, we will mostly focus on the (n, n) -unclonable secret sharing case. For simplicity, we call it n -party USS.

In this section, we define indistinguishability-based security for n -party USS. The security guarantees that for any two messages m_0, m_1 , no two reconstructing parties can simultaneously distinguish between whether the secret is m_0 or m_1 , given their sets of respective cloned shares. Formally, we define the following experiment:

$\text{Expt}_{(\{\mathcal{A}_i\}, \mathcal{B}, \mathcal{C}, \xi)}$:

1. Let ξ be a quantum state on registers $\mathbf{Aux}_1, \dots, \mathbf{Aux}_n$. For every $i \in [n]$, \mathcal{A}_i gets the register \mathbf{Aux}_i .
2. $\text{Adv} = (\{\mathcal{A}_i\}, \mathcal{B}, \mathcal{C}, \xi)$ sends (m_0, m_1) to the challenger such that $|m_0| = |m_1|$.
3. **Share Phase:** The challenger chooses a bit $b \xleftarrow{\$} \{0, 1\}$. It computes $\text{Share}(1^\lambda, 1^n, m_b)$ to obtain (ρ_1, \dots, ρ_n) and sends ρ_i to \mathcal{A}_i .
4. **Challenge Phase:** For every $i \in [n]$, \mathcal{A}_i computes a bipartite state $\sigma_{\mathbf{X}_i \mathbf{Y}_i}$. It sends the register \mathbf{X}_i to \mathcal{B} and \mathbf{Y}_i to \mathcal{C} .
5. \mathcal{B} on input the registers $\mathbf{X}_1, \dots, \mathbf{X}_n$, outputs a bit $b_{\mathcal{B}}$. \mathcal{C} on input the registers $\mathbf{Y}_1, \dots, \mathbf{Y}_n$, outputs a bit $b_{\mathcal{C}}$.
6. Output 1 if $b_{\mathcal{B}} = b$ and $b_{\mathcal{C}} = b$.

Definition 3 (Information-theoretic Unclonable Secret Sharing). An n -party unclonable secret sharing scheme (Share, Reconstruct) satisfies 1-bit unpredictability if for any non-uniform adversary $\text{Adv} = (\{\mathcal{A}_i\}_{i \in [n]}, \mathcal{B}, \mathcal{C}, \xi)$, the following holds:

$$\Pr \left[1 \leftarrow \text{Expt}_{(\{\mathcal{A}_i\}, \mathcal{B}, \mathcal{C}, \xi)} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

Definition 4 (Computational Unclonable Secret Sharing). An n -party unclonable secret sharing scheme (Share, Reconstruct) satisfies 1-bit unpredictability if for any non-uniform quantum polynomial-time adversary $\text{Adv} = (\{\mathcal{A}_i\}_{i \in [n]}, \mathcal{B}, \mathcal{C}, \xi)$, the following holds:

$$\Pr \left[1 \leftarrow \text{Expt}_{(\{\mathcal{A}_i\}, \mathcal{B}, \mathcal{C}, \xi)} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

Claim. Existence of $(n - 1)$ -party USS unconditionally implies n -party USS.

This is straightforward to see, by creating a dummy share.

4.3 Entanglement Graph

We will focus on the setting when there are multiple quantum adversaries with shared entanglement modeled as a graph, that we refer to as an *entanglement graph*. We formally define entanglement graphs below.

Definition 5 (Entanglement Graph). Let ρ be a n -partite quantum state over the registers $\mathbf{X}_1, \dots, \mathbf{X}_n$. Let $\rho[i]$ be the mixed state over register \mathbf{X}_i (i.e., $\rho[i] = \text{Tr}_{\mathbf{X}_i}(\rho)$) and $\rho[i, j]$ be the mixed state over the registers $\mathbf{X}_i, \mathbf{X}_j$ (i.e., $\rho[i, j] = \text{Tr}_{\mathbf{X}_i, \mathbf{X}_j}(\rho)$). An entanglement graph $G = (V, E)$ associated with $(\rho, \mathbf{X}_1, \dots, \mathbf{X}_n)$ is defined as follows:

- G is an undirected graph;
- $V = \{1, 2, \dots, n\}$;
- E contains an edge (u, v) if and only if \mathbf{X}_u and \mathbf{X}_v are entangled; or in other words, there does not exist σ_u, σ_v such that $\rho[u, v] = \sigma_u \otimes \sigma_v$.

Performing non-local operations on a state ρ , over the registers $\mathbf{X}_1, \dots, \mathbf{X}_n$, could change the entanglement graph. For instance, performing arbitrary channels on some \mathbf{X}_i , could remove some edges associated with the node i ; for example, a resetting channel that maps every state to $|0\rangle\langle 0|$. However, on the other hand, performing only unitary operations on each of $\mathbf{X}_1, \dots, \mathbf{X}_n$ is not going to change the entanglement graph.

Unless otherwise specified, we assume that the amount of entanglement shared between the different parties is either unbounded for information-theoretic protocols, or arbitrarily polynomial for computational protocols.

Definition 6. Let $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$ be the set of parties with ρ being the state received by all the parties. That is, ρ is an n -partite quantum state over the registers $\mathbf{X}_1, \dots, \mathbf{X}_n$ such that the party \mathcal{P}_i gets the register \mathbf{X}_i . We say that G is the entanglement graph associated with \mathcal{P} if G is the entanglement graph associated with $(\rho, \mathbf{X}_1, \dots, \mathbf{X}_n)$.

Definition 7 (USS_d). We say an information-theoretic/computational unclonable secret sharing scheme is a secure USS_d scheme, if it has indistinguishability-based security against all unbounded/efficient adversaries with pre-shared entanglement, whose entanglement graph has at least d connect components.

It is not hard to see that, USS₁ is a USS satisfying the regular indistinguishability security.

5 Adversaries with Disconnected Entanglement Graphs

In this section, we give a construction of unclonable secret sharing with security against quantum adversaries with disconnected entanglement graphs.

5.1 USS_{ω(log λ)}: an Information-Theoretic Approach

We present an information-theoretic protocol in the setting when there are $\omega(\log \lambda)$ connected components. For simplicity, we consider the case when there are $(n + 1)$ parties and the entanglement graph does not have any edges. We demonstrate a construction of USS in this setting, where the security scales with n .

1. Share($1^\lambda, 1^{(n+1)}, m \in \{0, 1\}$):
 - (a) Sample uniformly random $r_1, \dots, r_n \leftarrow \{0, 1\}$ conditioned on $\oplus_i r_i = m$.
 - (b) Sample $\theta_1, \dots, \theta_n \leftarrow \{0, 1\}$.
 - (c) For each $i \in [n]$: let the i^{th} share be $\rho_i = H^{\theta_i} |r_i\rangle\langle r_i| H^{\theta_i}$. Let the $(n+1)^{th}$ share be $\rho_{n+1} = (\theta_1, \dots, \theta_n)$.
 - (d) Output $(\rho_1, \dots, \rho_{n+1})$.
2. Reconstruct($\rho_1, \dots, \rho_{n+1}$):
 - (a) Measure ρ_{n+1} in the computational basis to get $(\theta_1, \dots, \theta_n)$.
 - (b) For every $i \in [n]$, apply H^{θ_i} to ρ_i . Measure the resulting state in the computational basis to get r_i .
 - (c) Output $\oplus_i r_i = m$.

Correctness and Soundness. We refer readers to the full version. Note that the soundness only holds for $n = \Omega(\log n)$, i.e., the protocol should have at least $\Omega(\log n)$ shares.

Security. Consider the adversary to be $\text{Adv} = (\{\mathcal{A}_i\}, \mathcal{B}, \mathcal{C}, \xi)$, where ξ is a product state. Henceforth, we omit mentioning $\xi = \xi_1 \otimes \dots \otimes \xi_{n+1}$, where \mathcal{A}_i receives ξ_i , since we can think of ξ_i to be part of the description of \mathcal{A}_i .

For $b \in \{0, 1\}$, let $(\rho_1^{r_1}, \dots, \rho_n^{r_n}, \rho_{n+1}) \leftarrow \text{Share}(1^\lambda, 1^{(n+1)}, b)$, where $\oplus_i r_i = b$ and $\rho_i = H^{\theta_i} |r_i\rangle\langle r_i| H^{\theta_i}$ and $\rho_{n+1} = |\theta_1 \dots \theta_n\rangle\langle \theta_1 \dots \theta_n|$. Suppose upon receiving $\rho_i^{r_i}$, \mathcal{A}_i sends registers $\{\mathbf{X}_i^{r_i}\}$ and $\{\mathbf{Y}_i^{r_i}\}$ respectively to \mathcal{B} and \mathcal{C} . We denote the reduced density matrix on $\mathbf{X}_i^{r_i}$ to be $\sigma_i^{r_i}$ and on $\mathbf{Y}_i^{r_i}$ to be $\zeta_i^{r_i}$. We assume without loss of generality that ρ_{n+1} is given to both \mathcal{B} and \mathcal{C} since it is a computational basis state.

Define \mathcal{S}_B and \mathcal{S}_C as follows:

$$\begin{aligned}\mathcal{S}_B &= \{i \in [n] : \text{TD}(\sigma_i^0, \sigma_i^1) \leq 0.86\} \\ \mathcal{S}_C &= \{i \in [n] : \text{TD}(\zeta_i^0, \zeta_i^1) \leq 0.86\}\end{aligned}$$

We prove the following claims.

Claim. Either $|\mathcal{S}_B| \geq \lceil \frac{n}{2} \rceil$ or $|\mathcal{S}_C| \geq \lceil \frac{n}{2} \rceil$.

Proof. We prove by contradiction; suppose it is not the case. Then there exists an index $i \in [n]$ such that $i \notin \mathcal{S}_B$ and $i \notin \mathcal{S}_C$. That is, $\text{TD}(\sigma_i^0, \sigma_i^1) > 0.86$ and $\text{TD}(\zeta_i^0, \zeta_i^1) > 0.86$, meaning the optimal state distinguishing circuit can distinguish σ_i^0, σ_i^1 with probability at least $0.93 = (1 + 0.86)/2$. Similarly, the optimal distinguishing probability for states ζ_i^0, ζ_i^1 is at least 0.93.

Using this, we design an adversary that violates the unclonable security of single-qubit BB84 states [BL20, Corollary 2]. Let us first recall the security game for the unclonability of single-qubit BB84 states:

1. \mathcal{A} receives $H^\theta|x\rangle\langle x|H^\theta$ for uniformly random $x, \theta \in \{0, 1\}$, it applies a channel and produces $\sigma_{\mathcal{BC}}$. Bob and Charlie receive their register accordingly.
2. Bob \mathcal{B} and Charlie \mathcal{C} apply their POVMs and try to recover x ; they win if and only if both guess x correctly.

Lemma 3 (Corollary 2 when $\lambda = 1$, [BL20]). *No (unbounded) quantum $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins the game with probability more than 0.855.*

We design an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ as follows, with winning probability $0.86 > 0.855$, a contradiction.

- \mathcal{A} receives as input an unknown BB84 state. It runs \mathcal{A}_i on the state to obtain a bipartite state, which it shares with \mathcal{B} and \mathcal{C} .
- \mathcal{B} and \mathcal{C} in the security game of BB84 state will receive θ_i from the challenger.
- \mathcal{B} runs the optimal distinguisher distinguishing σ_i^0 and σ_i^1 . Based on the output of the distinguisher, it outputs its best guess of the challenge bit. Similarly, Charlie runs the optimal distinguisher distinguishing ζ_i^0 and ζ_i^1 . It outputs its best guess of the challenge bit.

By a union bound, the probability that one of \mathcal{B} or \mathcal{C} fails is at most $0.14 = 0.07 \times 2$. Thus, they simultaneously succeed with probability at least 0.86, a contradiction.

Claim. The following holds:

1.

$$\text{TD} \left(\sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 0}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right), \sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 1}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right) \right) \leq 0.86^{|\mathcal{S}_B|}$$

2.

$$\text{TD} \left(\sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 0}} \frac{1}{2^{n-1}} \left(\bigotimes_i \zeta_i^{r_i} \right), \sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 1}} \frac{1}{2^{n-1}} \left(\bigotimes_i \zeta_i^{r_i} \right) \right) \leq 0.86^{|\mathcal{S}_C|}$$

Proof. We prove bullet 1 since bullet 2 follows symmetrically.

$$\begin{aligned} & \text{TD} \left(\sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 0}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right), \sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 1}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right) \right) \\ &= \frac{1}{2} \left\| \sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 0}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right) - \sum_{\substack{r_1, \dots, r_n: \\ \oplus_i r_i = 1}} \frac{1}{2^{n-1}} \left(\bigotimes_i \sigma_i^{r_i} \right) \right\|_1 \\ &= \left\| \bigotimes_i \frac{(\sigma_i^0 - \sigma_i^1)}{2} \right\|_1 \\ &= \prod_i \left\| \frac{(\sigma_i^0 - \sigma_i^1)}{2} \right\|_1 \\ &\leq \prod_{i \in \mathcal{S}_B} \text{TD}(\sigma_i^0, \sigma_i^1) \\ &\leq 0.86^{|\mathcal{S}_B|} \end{aligned}$$

Here $\|\cdot\|_1$ denotes the trace norm. In the above proof, we use the fact that $\|\bigotimes_i \tau_i\|_1 = \prod_i \|\tau_i\|_1$.

Lemma 4. *The above USS scheme satisfies indistinguishability security against any adversaries with no shared entanglement; i.e., it is a secure USS_n scheme (see Definition 7) with $n = \omega(\log \lambda)$.*

Proof. From Sect. 5.1, either $|\mathcal{S}_B| \geq \lceil \frac{n}{2} \rceil$ or $|\mathcal{S}_C| \geq \lceil \frac{n}{2} \rceil$. We will assume without loss of generality that $|\mathcal{S}_B| \geq \lceil \frac{n}{2} \rceil$. From bullet 1 of Sect. 5.1, it holds that \mathcal{B} can successfully distinguish whether it is in the experiment when the challenge bit 0 was used or when the challenge bit 1 was used, with probability at most $\frac{1+\nu(n)}{2}$, for some exponentially small function ν in n . Thus, both \mathcal{B} and \mathcal{C} can only simultaneously distinguish with probability at most $\frac{1+\nu(n)}{2}$. This completes the proof. \square

5.2 USS_d , for $d \geq 2$: from Unclonable Encryption

We present a construction of USS with security against quantum adversaries associated with *any* disconnected entanglement graph. In the construction, we use an information-theoretically secure unclonable encryption scheme, $\text{UE} = (\text{UE.KeyGen}, \text{UE.Enc}, \text{UE.Dec})$. The resulting USS scheme is consequently information-theoretically secure.

1. $\text{Share}(1^\lambda, 1^n, m)$:
 - (a) Sample $r_1, \dots, r_n \leftarrow \{0, 1\}^{|m|}$.
 - (b) For each $i \in [n]$, let $y_i = r_i$; let $y_n = m \oplus \sum_{i=1}^n r_i$.
 - (c) For each $i \in [n]$:
 - (a) Compute $\text{sk}_i \leftarrow \text{UE.KeyGen}(1^\lambda)$. We denote the length of sk_i to be $\ell = \ell(\lambda)$.
 - (b) Compute $|\text{ct}_i\rangle \leftarrow \text{UE.Enc}(\text{sk}_i, y_i)$
 - (d) For each $i \in [n]$: let each share $\rho_i = (\text{sk}_{i-1}, |\text{ct}_i\rangle)$; here we define $\text{sk}_0 = \text{sk}_n$.
 - (e) Output (ρ_1, \dots, ρ_n)
2. $\text{Reconstruct}(\rho_1, \dots, \rho_n)$:
 - (a) For each $i \in [n]$,
 - i. Parse ρ_i as $(\text{sk}_{i-1}, |\text{ct}_i\rangle)$. We define $\text{sk}_n = \text{sk}_0$.
 - ii. Compute $y_i \leftarrow \text{UE.Dec}(\text{sk}_i, |\text{ct}_i\rangle)$
 - (b) Output $m = \sum_{i=1}^n y_i$.

Theorem 8. *The above scheme satisfies indistinguishability-based security against adversaries with any disconnected entanglement graph. More precisely, it is a secure USS_2 scheme (see Definition 7).*

Proof. The correctness of the scheme follows from the correctness of UE decryption.

We now prove the security of the above scheme. Suppose we have an USS adversary $(\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n), \mathcal{B}, \mathcal{C}, \xi)$ who succeeds with probability $\frac{1}{2} + \varepsilon$ in Definition 7, we construct an UE adversary $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ who succeeds with probability $\frac{1}{2} + \varepsilon$ in Definition 2.

Let \mathcal{A} receive as input an n -partite state ξ over the registers $\mathbf{Aux}_1, \dots, \mathbf{Aux}_n$ such that \mathcal{A}_i receives as input the register \mathbf{Aux}_i . Additionally, without loss of generality, we can assume that \mathcal{A} also receives as input the challenge messages (m_0, m_1) , where $|m_0| = |m_1|$. Let $G = (V, E)$ be the entanglement graph associated with $(\xi, \mathbf{Aux}_1, \dots, \mathbf{Aux}_n)$, where, $V = \{1, \dots, n\}$. Since G is disconnected, there exists $i^* \in [n]$ such that $(i^*, i^* + 1) \notin E$. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two subgraphs of G such that $V_1 \cup V_2 = V$, $V_1 \cap V_2 = \emptyset$, $i^* \in V_1$, $i^* + 1 \in V_2$. Moreover, G_1 and G_2 are disconnected with each other. This further means that ξ can be written as $\xi_{G_1} \otimes \xi_{G_2}$, for some states ξ_{G_1}, ξ_{G_2} , such that ξ_{G_1} is over the registers $\{\mathbf{Aux}_i\}_{i \in V_1}$ and ξ_{G_2} is over the registers $\{\mathbf{Aux}_i\}_{i \in V_2}$.

We describe $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ as follows:

Description of \mathcal{A}' . Fix $i^*, (m_0, m_1)$ (as defined above). Upon receiving a quantum state $|\text{ct}^*\rangle$ \mathcal{A}' does the following:

- It prepares quantum states $\xi_{G_1}, (\xi_{G_2})^{\otimes 2^\ell}$.
- It samples $r_i \xleftarrow{\$} \{0, 1\}^{|m_0|}$, where $i \in [n]$, subject to the constraint that $\oplus_i r_i = m_0$.
- It submits $(r_{i^*}, r_{i^*} \oplus m_0 \oplus m_1)$ to the UE challenger and in return, it receives $|\text{ct}^*\rangle$. It sets $|\text{ct}_{i^*+1}\rangle = |\text{ct}^*\rangle$.
- For every $i \in [n]$, generate $\text{sk}_i \leftarrow \text{UE.KeyGen}(1^\lambda)$; let $\text{sk}_{n+1} = \text{sk}_1$.

- For every $i \in [n]$ and $i \neq i^*$, generate $|\text{ct}_i\rangle \leftarrow \text{UE.Enc}(\text{sk}_i, \text{sh}_i)$.
- For every $i \in [n]$ and $i \neq i^* + 1$, define $\rho_i = (\text{sk}_{i-1}, |\text{ct}_i\rangle)$.
- We need to define $\rho_{i^*+1} = (\text{sk}_{i^*}, |\text{ct}_{i^*+1}\rangle)$. However, as sk_{i^*} will only be received by \mathcal{B}' and \mathcal{C}' in the UE security game later, we will enumerate all possible values of sk_{i^*} and the corresponding computation result in the subgraph G_2 .
 - For every $x \in \{0, 1\}^\ell$ (possible value of sk_{i^*}), compute $\{\mathcal{A}_i\}_{i \in V_2}$ on $\{\rho_i\}_{i \in V_2}, \xi_{G_2}$ to obtain two sets of registers $\{\mathbf{X}_i^{(x)}\}_{i \in G_2}$ and $\{\mathbf{Y}_i^{(x)}\}_{i \in G_2}$.
- Compute $\{\mathcal{A}_i\}_{i \in V_1}$ on $\{\rho_i\}_{i \in V_1}$ and ξ_{G_1} to obtain two sets of registers $\{\mathbf{X}_i\}_{i \in G_1}$ and $\{\mathbf{Y}_i\}_{i \in G_1}$.
- Send the registers $\{\mathbf{X}_i\}_{i \in G_1}$ and $\{\mathbf{X}_i^{(x)}\}_{i \in G_2, x \in \{0, 1\}^\lambda}$ to \mathcal{B}' . Send the registers $\{\mathbf{Y}_i\}_{i \in G_1}$ and $\{\mathbf{Y}_i^{(x)}\}_{i \in G_2, x \in \{0, 1\}^\lambda}$ to \mathcal{C}' .

Description of \mathcal{B}' and \mathcal{C}' . \mathcal{B}' upon receiving the secret key k (which is sk_{i^*}), computes \mathcal{B} on $\{\mathbf{X}_i\}_{i \in G_1}$ and $\{\mathbf{X}_i^{(k)}\}_{i \in G_2}$ to obtain a bit $b_{\mathcal{B}}$. \mathcal{C}' is defined similarly. We denote the output of \mathcal{C}' to be $b_{\mathcal{C}}$.

If the challenger of the UE security chooses the bit $b = 0$, then $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ in the above reduction are receiving shares of m_0 ; otherwise, they are receiving shares of m_1 . Thus, the success probability of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ in Definition 7 is precisely the same as the success probability of $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ in Definition 2. \square

6 Impossibilities and Barriers

In this section, we present two impossibility results on USS. Furthermore, we present two implications of USS: namely, unclonable encryption and position verification secure against large amount of entanglement. Since no construction known for the latter two primitives, this further underscores the formidable barriers of building USS.

6.1 Impossibility in the Information-Theoretic Setting

Theorem 9. *Let \mathcal{P} be a set of parties. Information-theoretically secure USS for \mathcal{P} is impossible if the entanglement graph for \mathcal{P} is connected and in particular, there is an edge from P_1 to everyone else.*

Proof. The attack strategy is as follows. The n parties P_1, \dots, P_n pre-share a large amount of entanglement with one another. In the protocol, each P_i receives its share ρ_i .

- *Regular Teleportation Stage:* all parties P_i , where $i \neq 1$ teleport their shares to party P_1 via regular teleportation. Each P_i obtains a measurement outcome (a_i, b_i) .
- Now P_1 holds a state in the following format: $(\mathbb{I} \otimes X^{a_2} Z^{b_2} \otimes \dots \otimes X^{a_n} Z^{b_n}) |\Psi\rangle_{P_1 P_2 \dots P_n}$ which can be represented as mixed states $(\rho_1, X^{a_2} Z^{b_2} \rho_2 X^{a_2} Z^{b_2}, \dots, X^{a_n} Z^{b_n} \rho_n X^{a_n} Z^{b_n})$. That is, quantum one-time padded shares from all other parties and its own share in the clear.

– *Port-Based Teleportation Stage:*

- P_1 performs port-based teleportation for the state $(\mathbb{I} \otimes X^{a_2} Z^{b_2} \otimes \dots \otimes X^{a_n} Z^{b_n}) |\Psi\rangle_{P_1 P_2 \dots P_n}$ to P_2 . P_1 obtains a measurement outcome that stands for some index i_1 . Recall that by the guarantee of port-based teleportation, the index i_1 specifies the register of P_2 that holds the above state in the clear, *without any Pauli errors on top*.
- P_2 will now remove the quantum one time pad information X^{a_2}, Z^{a_2} on its share in the teleported state above. Since P_2 does not have information about i_1 , it simply performs $\mathbb{I} \otimes Z^{a_2} X^{a_2} \otimes \mathbb{I} \dots \otimes \mathbb{I}$ on all exponentially many possible registers that it may receive the teleported state from P_1 .
- Next P_2 performs port-based teleportation with P_3 for *all registers that could possibly hold the state* $(\mathbb{I} \otimes \mathbb{I} \otimes X^{a_3} Z^{b_3} \otimes \dots \otimes X^{a_n} Z^{b_n}) |\Psi\rangle_{P_1 P_2 \dots P_n}$. Thus, P_2 obtains an exponential number of indices about the registers that will receive the teleported states on P_3 's hands.
- P_3 accordingly, applies $\mathbb{I} \otimes \mathbb{I} \otimes Z^{b_3} X^{a_3} \dots \mathbb{I}$ on all the possible registers that can hold the teleported state; performs a port-based teleportation to P_4 with all of these registers and obtains a measurement outcome that has a doubly-exponential number of indices⁵.
- ...
- Finally, P_n receives the teleported states from P_{n-1} and performs $\mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes Z^{b_n} X^{a_n}$ on all of them. One of these registers will hold the state $|\Psi\rangle_{P_1 \dots P_n} = (\rho_1, \dots, \rho_n)$ in the clear. Then P_n performs the reconstruction algorithm on all of these registers to obtain a large number of possible outcomes. One of them will hold the correctly reconstructed secret s .

– *Reconstruction Stage:* now P_n sends all its measurement outcomes to both Bob and Charlie. All other P_i 's send their indices information measured in the port teleportation protocol. Bob and Charlie can therefore find the correct index in P_n 's measurement outcomes that holds s , by following a path of indices.

□

Remark 2. The above strategy can be easily converted into a strategy where the underlying entanglement graph is connected (but may not be a complete graph) and every pair of connected parties share (unbounded) entanglement. The similar idea applies by performing regular teleportation and port-based teleportation via any DFS order of the graph. Thus, we have the following theorem.

Theorem 10. *Let \mathcal{P} be a set of parties. Information-theoretically secure USS for \mathcal{P} is impossible if the entanglement graph for \mathcal{P} is connected.*

6.2 Impossibility with Low T-gates for Efficient Adversaries

Our impossibility result above in the information-theoretic setting requires exponential amount of entanglement between the parties. We also present an attack

⁵ For $P_i, 2 \leq i < n$, the measurement outcome will have its size grow in an exponential tower of height i .

that can be performed by efficient adversaries, albeit on USS schemes with restricted reconstruction algorithms.

We would like to mention that a similar result has already been shown in [Spe15] in the context of instantaneous non-local computation; we rediscovered the following simple attack for unclonable secret sharing. We also extend the attack to an n -party setting whereas [Spe15] considers only 2 parties.

Theorem 11. *Let \mathcal{P} be a set of parties and if the entanglement graph for \mathcal{P} is connected, then there exists an attack using polynomial-time and polynomial amount of entanglement on any USS scheme where the procedure Reconstruct consists of only Clifford gates and $O(\log \lambda)$ number of T gates.*

We refer readers to the full version for the proof.

6.3 USS Implies Unclonable Encryption

Theorem 12. *Unclonable secret sharing with IND-based security against adversaries with (bounded) polynomial amount of shared entanglement and connected pre-shared entanglement graph implies secure unclonable encryption.*

We will first look at the 2-party case, which can be easily extended to the $n(> 2)$ -party case.

Proof. Assume a secure USS = (USS.Share, USS.Reconstruct) with IND-based security, we construct the following UE scheme:

1. $\text{KeyGen}(1^\lambda, 1^{|m|})$: samples a random $\text{sk} \leftarrow \{0, 1\}^{2\ell}$, where $\ell = \ell(\lambda)$ is the number of qubits in each share generated by $\text{USS.Share}(1^\lambda, 1^{|m|}, \cdot)$. Output sk .
2. $\text{Enc}(\text{sk}, m)$:
 - (a) compute $(\rho_1, \rho_2) \leftarrow \text{USS.Share}(1^\lambda, 1^{|m|}, m)$.
 - (b) sample random $(a, b) \leftarrow \{0, 1\}^{2\ell}$. Use them to quantum one-time pad the second share ρ_2 to obtain $X^a Z^b \rho_2 Z^b X^a$.
 - (c) compute $s \leftarrow (a, b) \oplus \text{sk}$
 - (d) Output $\text{ct} = (\rho_1, X^a Z^b \rho_2 Z^b X^a, s)$.
3. $\text{Dec}(\text{ct}, \text{sk})$:
 - (a) parse $\text{ct} = (\rho_1, \rho'_2, s)$;
 - (b) compute $(a, b) \leftarrow s \oplus \text{sk}$;
 - (c) output $m \leftarrow \text{USS.Reconstruct}(\rho_1, X^a Z^b \rho'_2 Z^b X^a)$.

Correctness. The correctness easily follows from the correctness of the underlying USS scheme.

Security. Suppose we have UE adversaries $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ that wins in the IND-based UE security game, we can construct adversary $(\mathcal{A}' = (\mathcal{A}_1, \mathcal{A}_2), \mathcal{B}', \mathcal{C}')$ for the USS IND-based security.

Before receiving the shares from the challenger, \mathcal{A}_1 and \mathcal{A}_2 agrees on a random strong $r \leftarrow \{0, 1\}^{2\ell}$. When receiving the shares, \mathcal{A}_2 teleports its share ρ_2 to \mathcal{A}_1 and obtains Pauli errors (a, b) .

\mathcal{A}_1 gives (ρ_1, r) the UE adversary \mathcal{A} . \mathcal{A}_2 computes $\text{sk}' \leftarrow (a, b) \oplus r$.

In the USS challenge phase, \mathcal{A}_2 sends sk' to both \mathcal{B}' and \mathcal{C}' . The UE adversaries \mathcal{A} has finished giving the bipartite it generated from (ρ_1, r) state $\sigma_{\mathcal{B}, \mathcal{C}}$ to \mathcal{B} and \mathcal{C} .

Then \mathcal{B}' feeds \mathcal{B} with sk' as the secret key in the UE security game (and \mathcal{C}' feeding sk' to \mathcal{C} , respectively), and outputs their output bit $b_{\mathcal{B}}, b_{\mathcal{C}}$ as the answer to USS game. Since the classical part in the unclonable ciphertext is the classical information (a, b) masked by a uniformly random sk , the reduction perfectly simulates the above scheme by first giving the UE adversary \mathcal{A} a uniformly random string r and later feeding \mathcal{B}, \mathcal{C} with $r \oplus (a, b)$.

Extending to n -party case. We can change the scheme to sample a longer $\text{sk} \in \{0, 1\}^{2(n-1)\ell}$ and let the unclonable ciphertext be $(\rho_1, X^{a_2} Z^{b_2} \rho_2 Z^{b_2} X^{a_2}, \dots, X^{a_n} Z^{b_n} \rho_n Z^{b_n} X^{a_n}, s = (a_1, b_1, \dots, a_n, b_n) \oplus \text{sk})$.

In the reduction, when receiving the shares, $\mathcal{A}_i, i \neq 1$ teleports its share ρ_i to \mathcal{A}_1 and obtains Pauli errors (a_i, b_i) . The rest of the reduction follows easily. □

Theorem 13. *Unclonable secret sharing with IND-based security against adversaries with disconnected entanglement graph, where one of the parties receives as a share a quantum state and all other parties receive classical shares (in other words, computational basis states), implies secure unclonable encryption.*

Proof. In the case where only one party has a quantum share, the others classical shares, we can easily modify the above construction to have a UE scheme from USS:

1. $\text{KeyGen}(1^\lambda, 1^{|m|})$: samples a random $\text{sk} \leftarrow \{0, 1\}^{(n-1)\ell}$, where $\ell = \ell(\lambda)$ is the number of qubits/bits in each share generated by $\text{USS.Share}(1^\lambda, 1^{|m|}, \cdot)$. Output sk .
2. $\text{Enc}(\text{sk}, m)$:
 - (a) compute $(\rho_1, y_2, \dots, y_n) \leftarrow \text{USS.Share}(1^\lambda, 1^{|m|}, m)$. y_1, \dots, y_n are binary strings.
 - (b) sample random $\text{sk} \leftarrow \{0, 1\}^{(n-1)\ell}$. Compute $s \leftarrow (y_1, \dots, y_n) \oplus \text{sk}$
 - (c) Output $\text{ct} = (\rho_1, s)$.
3. $\text{Dec}(\text{ct}, \text{sk})$:
 - (a) parse $\text{ct} = (\rho_1, s)$;
 - (b) compute $(y_1, \dots, y_n) \leftarrow s \oplus \text{sk}$;
 - (c) output $m \leftarrow \text{USS.Reconstruct}(\rho_1, y_1, \dots, y_n)$.

Security. Suppose we have an UE adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ that wins in the IND-based UE security game with probability $\frac{1}{2} + \varepsilon$, we construct an adversary $(\mathcal{A}' = (\mathcal{A}_1, \dots, \mathcal{A}_n), \mathcal{B}', \mathcal{C}')$ that wins in the USS IND-based security game with probability $\frac{1}{2} + \varepsilon$. Thus, if the USS scheme is secure then ε has to be negligible. We describe $\mathcal{A}_1, \dots, \mathcal{A}_n$ as follows.

Before receiving the shares from the challenger, $\mathcal{A}_1, \dots, \mathcal{A}_n$ agrees on a random string $r \leftarrow \{0, 1\}^{(n-1)\ell}$.

\mathcal{A}_1 gives (ρ_1, r) to the UE adversary \mathcal{A} . \mathcal{A}_i , for $i \neq 1$, when receiving the classical share y_i from the challenger, computes $\text{sk}'_i \leftarrow y_i \oplus r_i$, where r_i is the $(i-1)$ -th block of length- ℓ string in r .

In the USS challenge phase, each \mathcal{A}_i , for $i \neq 1$, sends sk'_i to both \mathcal{B}' and \mathcal{C}' . \mathcal{A}_1 sends the bipartite state $\sigma_{\mathcal{B}, \mathcal{C}}$ to \mathcal{B}' and \mathcal{C}' , where $\sigma_{\mathcal{B}, \mathcal{C}}$ is the output of \mathcal{A} .

Then \mathcal{B}' feeds \mathcal{B} with $\text{sk}' = (\text{sk}'_2, \dots, \text{sk}'_n)$ as the secret key in the UE security game (and \mathcal{C}' feeding sk' to \mathcal{C} , respectively), and outputs their output bit $b_{\mathcal{B}}, b_{\mathcal{C}}$ as the answer to USS game. Since the classical part in the unclonable ciphertext is the classical information (y_2, \dots, y_n) masked by a uniformly random sk , the reduction perfectly simulates the above scheme by first giving the UE adversary \mathcal{A} a uniformly random string r and later feeding \mathcal{B}, \mathcal{C} with $r \oplus (y_2, \dots, y_n)$. Thus, the advantage of $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ in breaking the USS security game is precisely the same as the advantage of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ breaking the UE security game.

6.4 Search-Based USS Implies Position Verification

The definition of quantum position verification is in the full version.

QPV with Pre-shared Entanglement. In QPV, we also consider different adversarial setup such as: (1) (P_0, P_1) do not have pre-shared entanglement; (2) (P_0, P_1) can share a bounded/unbounded polynomial amount of entanglement; (3) (P_0, P_1) can share unbounded amount of entanglement. We also divide the settings into computational and information-theoretic.

Theorem 14. *2-party USS (computational/IT resp.) with search-based security implies 1-dimensional QPV (computational/IT, resp.), where the two adversarial provers in the QPV protocol pre-share the same amount of entanglement as the two parties in the USS protocol do.*

The following theorem demonstrates from another point of view the barrier of constructing secure protocols against entangled adversaries for USS in the IT setting. Even if we consider computational assumptions, the development in building secure QPV protocols against entangled adversaries has been slow, which indicates further evidence on how challenging USS can be in the entangled setting.

Theorem 15 ([BK11, BCF+14]). *Quantum position verification is impossible in the information theoretic setting if we allow the adversaries to pre-share entanglement.*

We leave the proof to the full version.

References

- [Aar09] Scott Aaronson. “Quantum copy-protection and quantum money”. In: *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE. 2009, pp. 229–242 (cit. on p. 3).
- [AC12] Scott Aaronson and Paul Christiano. *Quantum Money from Hidden Subspaces*. 2012. DOI: <https://doi.org/10.48550/ARXIV.1203.4740>. URL: <https://arxiv.org/abs/1203.4740> (cit. on p. 3).
- [AK21] Prabhanjan Ananth and Fatih Kaleoglu. “Unclonable Encryption, Revisited”. In: *Theory of Cryptography Conference*. Springer. 2021, pp. 299–329 (cit. on pp. 3, 15).
- [AKL+22] Prabhanjan Ananth, Fatih Kaleoglu, Xingjian Li, Qipeng Liu, and Mark Zhandry. “On the feasibility of unclonable encryption, and more”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 212–241 (cit. on pp. 3, 15).
- [AKL23] Prabhanjan Ananth, Fatih Kaleoglu, and Qipeng Liu. “Cloning Games: A General Framework for Unclonable Primitives”. In: *Annual International Cryptology Conference*. Springer. 2023, pp. 66–98 (cit. on pp. 3, 15).
- [AL20] Prabhanjan Ananth and Rolando L. La Placa. *Secure Software Leasing*. 2020. DOI: <https://doi.org/10.48550/ARXIV.2005.05289>. URL: <https://arxiv.org/abs/2005.05289> (cit. on p. 3).
- [BB20] Charles H Bennett and Gilles Brassard. “Quantum cryptography: Public key distribution and coin tossing”. In: *arXiv preprint arXiv:2003.06557* (2020) (cit. on pp. 3, 6).
- [BCF+14] Harry Buhrman, Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, Rafail Ostrovsky, and Christian Schaffner. “Position-based quantum cryptography: Impossibility and constructions”. In: *SIAM Journal on Computing* 43.1 (2014), pp. 150–178 (cit. on p. 27).
- [BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. “Random oracles in a quantum world”. In: *International conference on the theory and application of cryptology and information security*. Springer. 2011, pp. 41–69 (cit. on p. 6).
- [BK11] Salman Beigi and Robert König. “Simplified instantaneous nonlocal quantum computation with applications to position-based cryptography”. In: *New Journal of Physics* 13.9 (2011), p. 093036 (cit. on pp. 3, 12, 27).
- [BL20] Anne Broadbent and Sébastien Lord. “Uncloneable Quantum Encryption via Oracles”. en. In: *Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, 2020. DOI: <https://doi.org/10.4230/LIPICS.TQC.2020.4>. URL: <https://drops.dagstuhl.de/opus/volltexte/2020/12063/> (cit. on pp. 3, 6, 10, 12, 15, 19).
- [BS16] Shalev Ben-David and Or Sattath. *Quantum Tokens for Digital Signatures*. 2016. DOI: <https://doi.org/10.48550/ARXIV.1609.09047>. URL: <https://arxiv.org/abs/1609.09047> (cit. on p. 3).
- [CGL99] Richard Cleve, Daniel Gottesman, and Hoi-Kwong Lo. In: *Phys. Rev. Lett.* 83 (3 July 1999), pp. 648–651. DOI: <https://doi.org/10.1103/PhysRevLett.83.648>. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.648> (cit. on p. 8).
- [ÇGLR23] Alper Çakan, Vipul Goyal, Chen-Da Liu-Zhang, and João Ribeiro. “Computational Quantum Secret Sharing”. In: *arXiv preprint arXiv:2305.00356* (2023) (cit. on p. 8).

- [CGMO09] Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. “Position based cryptography”. In: *Annual International Cryptology Conference*. Springer. 2009, pp. 391–407 (cit. on p. 14).
- [CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. “Hidden Cosets and Applications to Unclonable Cryptography”. In: *Advances in Cryptology – CRYPTO 2021*. Ed. by Tal Malkin and Chris Peikert. Cham: Springer International Publishing, 2021, pp. 556–584. ISBN: 978-3-030-84242-0 (cit. on p. 3).
- [GC19] Alvin Gonzales and Eric Chitambar. “Bounds on instantaneous nonlocal quantum computation”. In: *IEEE Transactions on Information Theory* 66.5 (2019), pp. 2951–2963 (cit. on p. 3).
- [Got00] Daniel Gottesman. “Theory of quantum secret sharing”. In: *Physical Review A* 61.4 (2000), p. 042311 (cit. on p. 8).
- [Got02] Daniel Gottesman. “Uncloneable Encryption”. In: (2002). DOI: <https://doi.org/10.48550/ARXIV.QUANT-PH/0210062>. URL: <https://arxiv.org/abs/quant-ph/0210062> (cit. on p. 3).
- [GSS21] Vipul Goyal, Yifan Song, and Akshayaram Srinivasan. “Traceable secret sharing and applications”. In: *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part III 41*. Springer. 2021, pp. 718–747 (cit. on p. 2).
- [HBB99] Mark Hillery, Vladimír Bužek, and André Berthiaume. In: *Phys. Rev. A* 59 (3 Mar. 1999), pp. 1829–1834. DOI: <https://doi.org/10.1103/PhysRevA.59.1829>. URL: <https://link.aps.org/doi/10.1103/PhysRevA.59.1829> (cit. on p. 8).
- [IH08] Satoshi Ishizaka and Tohya Hiroshima. “Asymptotic teleportation scheme as a universal programmable quantum processor”. In: *Physical review letters* 101.24 (2008), p. 240501 (cit. on pp. 3, 12).
- [KKI99] Anders Karlsson, Masato Koashi, and Nobuyuki Imoto. “Quantum entanglement for secret sharing and secret splitting”. In: *Phys. Rev. A* 59 (1 Jan. 1999), pp. 162–168. DOI: <https://doi.org/10.1103/PhysRevA.59.162>. URL: <https://link.aps.org/doi/10.1103/PhysRevA.59.162> (cit. on p. 8).
- [LMZ23] Jiahui Liu, Hart Montgomery, and Mark Zhandry. “Another Round of Breaking and Making Quantum Money: How to Not Build It from Lattices, and More”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2023, pp. 611–638 (cit. on p. 3).
- [May19] Alex May. “Quantum tasks in holography”. In: *Journal of High Energy Physics* 2019.10 (2019), pp. 1–39 (cit. on pp. 3, 14).
- [May22] Alex May. “Complexity and entanglement in non-local computation and holography”. In: *Quantum* 6 (2022), p. 864 (cit. on pp. 3, 14).
- [NC10] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010 (cit. on p. 14).
- [Shm22] Omri Shmueli. “Public-key Quantum money with a classical bank”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 790–803 (cit. on p. 3).
- [Smi00] Adam D Smith. “Quantum secret sharing for general access structures”. In: *arXiv preprint quant-ph/0001087* (2000) (cit. on p. 8).
- [Spe15] Florian Speelman. “Instantaneous non-local computation of low T-depth quantum circuits”. In: *arXiv preprint arXiv:1511.02839* (2015) (cit. on pp. 3, 24).

- [TFKW13] Marco Tomamichel, Serge Fehr, Jędrzej Kaniewski, and Stephanie Wehner. “A monogamy-of-entanglement game with applications to device-independent quantum cryptography”. In: *New Journal of Physics* 15.10 (Oct. 2013), p. 103002. DOI: <https://doi.org/10.1088/1367-2630/15/10/103002>. URL: <https://doi.org/10.1088%2F1367-2630%2F15%2F10%2F103002> (cit. on pp. 6, 10).
- [Vai03] Lev Vaidman. “Instantaneous measurement of nonlocal variables”. In: *Physical review letters* 90.1 (2003), p. 010402 (cit. on p. 3).
- [Zha17] Mark Zhandry. “Quantum Lightning Never Strikes the Same State Twice”. In: *CoRR* abs/1711.02276 (2017). arXiv: 1711.02276. URL: <http://arxiv.org/abs/1711.02276> (cit. on p. 3).



Improved Quantum Lifting by Coherent Measure-and-Reprogram

Alexandru Cojocaru^{1(✉)}, Juan Garay², Qipeng Liu³, and Fang Song⁴

¹ School of Informatics, University of Edinburgh, Edinburgh, UK
a.cojocaru@ed.ac.uk

² Department of Computer Science and Engineering, Texas A&M University,
College Station, USA
garay@tamu.edu

³ Department of Computer Science and Engineering, UC San Diego, San Diego, USA
qipengliu@ucsd.edu

⁴ Department of Computer Science, Portland State University, Portland, USA
fang.song@pdx.edu

Abstract. We give a tighter lifting theorem for security games in the quantum random oracle model. At the core of our main result lies a novel measure-and-reprogram framework that we call *coherent reprogramming*. This framework gives a tighter lifting theorem for query complexity problems, that only requires purely classical reasoning. As direct applications of our lifting theorem, we first provide a quantum direct product theorem in the average case—i.e., an enabling tool to determine the hardness of solving multi-instance security games. This allows us to derive in a straightforward manner the hardness of various security games, for example (i) the non-uniform hardness of salted games, (ii) the hardness of specific cryptographic tasks such as the multiple instance version of one-wayness and collision-resistance, and (iii) uniform or non-uniform hardness of many other games.

1 Introduction

Hash functions are a fundamental workhorse in modern cryptography. Efficient constructions such as SHA-2 and SHA-3 are widely used in real-world cryptographic applications. To facilitate the analysis of constructions based on hash functions, Bellare and Rogaway [BR93] proposed a framework known as the random oracle model (ROM). Recent development of quantum computing demands re-examining security against potential quantum attackers. The *quantum random oracle model* (QROM) has since been proposed by Boneh *et al.* [BDF+11] as an extension of ROM by taking into account quantum attackers. Various techniques have been developed for analyzing security in the QROM; however, they are often either *ad-hoc* (for specific scenarios) or too involved to apply.

In this paper we revisit a general tool for lifting security from ROM to QROM by Yamakawa and Zhandry [YZ21]. The lifting theorems are applicable to search games in (Q)ROM between a classical challenger interacting with

an adversary (e.g., think of an adversary that aims to find a preimage of 0 in the random oracle, with the challenger querying the random oracle to verify the adversary’s answer). Roughly speaking, the lifting theorems assert that if a search game with a challenger performing a constant number of queries to the random oracle is hard against a classical adversary, then it is also hard against a quantum adversary in the QROM. Specifically, if the challenger performs k queries and if a quantum adversary performs q quantum queries and wins the search game with probability ϵ , then there exists a classical adversary performing only k classical queries winning with probability $\epsilon/(2q+1)^{2k}$.

This tool is particular powerful to establish quantum query lower bounds in the QROM. Let us consider function inversion from above for example; the goal is to find an input x , whose image equals 0. In this case, $k = 1$ and

$$\frac{\epsilon}{(2q+1)^{2k}} = \frac{\epsilon}{(2q+1)^2} \leq \frac{1}{N}.$$

This is because a single query reveals a pre-image of 0 with probability at most $1/N$. Therefore we have $\epsilon \leq (2q+1)^2/N$, which immediately reproduces the tight bound for the famous Grover’s search problem [BBBV97]. However, for a k -search problem whose goal is to find k distinct inputs that all map to 0, the bound derived from [YZ21] is $O((kq)^2/N)^k$ for any quantum algorithm with kq queries, which has a large k^{2k} factor gap from the tight bound $\Theta(q^2/N)^{k1}$. Similar weaknesses appear in a variety of problems involving multiple inputs.

In this work, we derive a new tighter lifting theorem for search games. If the challenger performs k queries and the quantum adversary performs q quantum queries and wins the search game with probability ϵ , then there exists a quantum adversary performing only k quantum queries and winning with probability $\epsilon/2^{2k} \binom{q+k}{k}^2$, improving on the previous lifting theorem by Yamakawa and Zhandry. Let us consider a (kq) -quantum-query algorithm for the previous k -search problem. Our bound (in this case, q in the theorem should be kq) gives the tight bound as below:

$$\frac{\epsilon}{2^{2k} \binom{kq+k}{k}^2} \leq \frac{1}{N^k} \iff \epsilon \leq \frac{2^{2k} \binom{kq+k}{k}^2}{N^k} \leq O\left(\frac{(q+1)^2}{N}\right)^k.$$

To achieve this, we develop a new measure-and-reprogram technique which is a key technical contribution of our work. The technique, which we call *coherent reprogramming*, improves on the recent results on adaptively reprogramming QRO on multiple points by Don *et al.* [DFM20, DFMS22] and Liu and Zhandry [LZ19], yielding tighter reprogramming bounds than the existing measure-and-reprogram proofs. As an immediate consequence, we are able to derive tighter

¹ We believe this is a folklore result that to our knowledge, this bound follows from a result in [CGK+23] (Theorem 3.1). Moreover, we would like to emphasize that since our main result is a strengthening of the lifting Lemma of [YZ21], we can also show that our result concerning the bound of this problem is stronger than the bound derived from [YZ21].

quantum hardness bounds for many applications, such as direct-product theorems, salted security, and non-uniform security.

1.1 Summary of Our Results

Lifting Theorem for Search Games. Our central result is a new lifting theorem for search games that relates (and upper bounds) the success probability of an arbitrary quantum algorithm to the success probability of a quantum algorithm performing a small number of queries to the RO. More formally:

Theorem 1 (Quantum Lifting Theorem (Informal)). *Let \mathcal{G} be a search game with a classical challenger \mathcal{C} that performs at most k queries to the RO, and let \mathcal{A} be a q -quantum query adversary in the game \mathcal{G} (against the k -classical query challenger \mathcal{C}). Then there exists a k -quantum query adversary \mathcal{B} such that:*

$$\Pr[\mathcal{B} \text{ wins } \mathcal{G}] \geq \frac{1}{2^{2k} \binom{q+k}{k}^2} \Pr[\mathcal{A} \text{ wins } \mathcal{G}].$$

Remark 1. Comparing to the lifting theorem in [YZ21], we have a better loss $2^{2k} \binom{q+k}{k}^2$, whereas it is $(2q+1)^k$ in their work. Since the algorithm often makes more queries than the challenger $q \gg k$, it is roughly a $(k!)^2$ save. In [YZ21], they are able to reduce a q -quantum-query algorithm to a k -classical-query algorithm; whereas in this work, we only reduce the number of the queries, with the algorithm \mathcal{B} still making quantum queries. Nonetheless, it does not affect the applications and improvement we have in this work. Our framework handles the case where the challenge is independent of the oracle (similar to the results in [YZ21]). We leave the case of oracle-dependent challenges as an interesting open question.

Coherent Reprogramming. At the core of our main lifting result above lies a new framework for quantum reprogramming which we call *coherent reprogramming*. This new framework has the following advantages:

1. It simplifies the proofs and frameworks of existing quantum reprogramming results;
2. It yields improved tighter reprogramming bounds; and
3. It implies in a straightforward manner several applications in quantum query complexity and cryptography.

In order to present our main coherent reprogramming result, we first need to introduce a few notions. For an oracle H we call $H_{x,y}$ the reprogrammed oracle that behaves almost like the original function H , with the only difference that its value on input x will be y . Similarly, we define the reprogrammed oracle on k inputs $\mathbf{x} = (x_1, \dots, x_k)$ and k corresponding outputs $\mathbf{y} = (y_1, \dots, y_k)$, denoted by $H_{\mathbf{x},\mathbf{y}}$, as the original function H with the only difference that for every input x_i in \mathbf{x} , the corresponding image will be y_i in \mathbf{y} .

Theorem 2 (Coherent Reprogramming (Informal)). *Let H, G be two random oracles. Let \mathcal{A} be any q -quantum query algorithm to the oracle H , and let $\mathbf{x}_o = (x_1, \dots, x_k) \in X^k$ be any k -vector of inputs and $\mathbf{y}_o = (y_1, \dots, y_k) = (G(x_1), \dots, G(x_k))$. Then there exists a simulator algorithm Sim that given oracle access to H, G and \mathcal{A} , simulates the output of \mathcal{A} having oracle access to $H_{\mathbf{x}_o, \mathbf{y}_o}$ (the reprogrammed version of H) with probability:*

$$\Pr_{H, G} [\text{Sim outputs correct } (\mathbf{x}, \mathbf{y})] \geq \frac{1}{2^{2k} \binom{q+k}{k}^2} \cdot \Pr_{H, G} [\mathcal{A}^{H_{\mathbf{x}_o, \mathbf{y}_o}} \text{ outputs correct } (\mathbf{x}, \mathbf{y})].$$

where “correct” is defined with respect to some predicate that can depend on the reprogrammed oracle $H_{\mathbf{x}_o, \mathbf{y}_o}$.

Remark 2. Similar to the comparison between Theorem 1 and [YZ21], the second theorem improves the factor $(2q+1)^k$ in [DFM20] to $2^{2k} \binom{q+k}{k}^2$. Our simulator does not measure and reprogram directly, but does everything coherently (or in superposition).

Next, we show the applications of our lifting theorem in query complexity and cryptography.

Quantum Lifting Theorem with Classical Reasoning. A multi-output k -search game between a challenger and an adversary is defined as follows. The adversary receives k different challenges from the challenger, and at the end of their interaction, the adversary needs to respond with k outputs. If the k outputs (taken together) satisfy some relation R specified by the game, we say the adversary wins the multi-output k -search game. The goal of the *lifting theorem* is to establish the hardness of solving the multi-output k -search game by any general quantum adversary, with only simple classical reasoning. For an arbitrary k -ary relation R , let S_k be the symmetric group on $[k]$ and we define:

$$p(R) := \Pr[\exists \pi \in S_k \mid (y_{\pi(1)}, y_{\pi(2)}, \dots, y_{\pi(k)}) \in R : (y_1, \dots, y_k) \stackrel{\$}{\leftarrow} Y^k].$$

Note that $p(R)$ is a quantity that only depends on the game itself, and can be calculated with only classical reasoning.

Theorem 3 (Quantum Lifting Theorem with Classical Reasoning (Informal)). *For any quantum algorithm \mathcal{A} equipped with q quantum queries to a random oracle $H : X \rightarrow Y$, \mathcal{A} 's success probability to solve the multi-output k -search game as specified by the winning relation R , is bounded by:*

$$\Pr[\mathcal{A} \text{ wins multi-output } k\text{-search game}] \leq 2^{2k} \binom{q+k}{k}^2 \cdot p(R).$$

Our lifting theorem translates into the following quantum hardness results for our applications in query complexity and cryptography.

Direct Product Theorem. We give the first direct product theorem (DPT) in the average case (in the QROM). Previously, only worst-case quantum DPTs were

known [She11, LR13] and were proof-method dependent; until recently, Dong et al. [DLW24] shows the first average-case quantum DPTs for some problems in the QROM. While they are non-tight, our DPTs works for all games in the QROM and proof-method independent.

Our direct product theorem establishes the hardness of solving g independent instances (each instance is associated with an independent oracle) of a game \mathcal{G} given a total of $g \cdot q$ quantum queries:

Theorem 4 (Direct Product Theorem). *For any quantum algorithm \mathcal{A} equipped with $g \cdot q$ quantum queries, \mathcal{A} 's success probability to solve the Direct Product game $\mathcal{G}^{\otimes g}$ with the underlying \mathcal{G} specified by the winning relation R , is bounded by*

$$\Pr[\mathcal{A} \text{ wins } \mathcal{G}^{\otimes g}] \leq \left(2^{2k} \binom{q+k}{k}^2 p(R) \right)^g.$$

Non-uniform Security of Salting. The above theorem directly implies non-uniform security of salting. Non-uniform attacks allow a malicious party to perform heavy computation offline and attack a protocol much more efficiently, using the information in the offline stage. Salting is a generic method that prevent non-uniform attacks against hash functions. Chung et al. [CGLQ20] shows that “salting generically defeats quantum preprocessing attacks”; they show that if a game in the QROM is $\epsilon(q)$ secure, the salted game with salt space $[K]$ is $\epsilon(q) + \frac{Sq}{K}$ secure against a quantum adversary with S -bit of advice. Their bound is non-tight, since when the underlying game is collision-finding, the tight non-uniform security should be $\epsilon(q) + \frac{S}{K}$. Improving the additive factor is an interesting open question and until recently [DLW24] is able to answer this question affirmatively for a limited collection of games.

Using our direct product theorem, we show:

Theorem 5 (Another “Salting Defeats Quantum Preprocessing”). *For any non-uniform quantum algorithm \mathcal{A} equipped with q quantum queries and S -bit of classical advice, \mathcal{A} 's success probability to solve the salted game \mathcal{G}_s with the underlying \mathcal{G} specified by the winning relation R , is bounded by*

$$\Pr[\mathcal{A} \text{ wins } \mathcal{G}_s] \leq 4 \cdot \left(2^{2k} \binom{q+k}{k}^2 p(R) + \frac{S}{K} \right).$$

Our bound is incomparable to that in [CGLQ20]. We are able to improve the additive term from $\frac{Sq}{K}$ to $\frac{S}{K}$, while only able to give an upper bound for $\epsilon(q)$. Even our bound is non-tight in general, it still confirms (on a high level) that the help from classical advice only comes from the following:

- using S -bit advice to store solutions for S random salts;
- if the challenge salt matches with the random salts (with probability $\frac{S}{K}$), the attack succeeds; otherwise, proceed the attack as if there is no advice.

Non-uniform Security. By combining our lifting theorem with the results by Chung *et al.* [CGLQ20], we derive the following results concerning the security (hardness) against non-uniform quantum adversaries with classical advice, for a broader class of games.

Lemma 1 (Security against Quantum Non-Uniform Adversaries (Informal)).

Let \mathcal{G} be any classically verifiable search game specified by the winning relation R . Let $R^{\otimes S}$ be the winning relation of the multi-instance game of \mathcal{G} . Any quantum non-uniform algorithm \mathcal{A} equipped with q quantum queries and S classical bits of advice, can win \mathcal{G} with probability at most:

$$\Pr[\mathcal{A} \text{ wins } \mathcal{G}] \leq 4 \cdot 2^{2k} \binom{S(q+k)}{Sk}^{\frac{2}{S}} \cdot p(R^{\otimes S})^{1/S}.$$

To demonstrate the power of our results, we also apply them to the hardness of three concrete cryptographic tasks: the multiple instance versions of one-wayness, collision resistance and search, as described next. Note that the applications we list below are non-exhaustive, given $p(R)$ is easy to define for almost every game.

Hardness of Multi-image Inversion. Firstly, we can analyze the quantum hardness of inverting k different images of a random oracle $H : [M] \rightarrow [N]$.

Our first result establishes the quantum hardness of multi-image inversion, which is a tight bound as already proven in [CGLQ20], but achieved here in a much simpler way, directly from our quantum lifting theorem.

Lemma 2 (Quantum Hardness of Multi-Image Inversion (Informal)). For any distinct $\mathbf{y} = (y_1, \dots, y_k) \in [N]^k$ and for any q -quantum query algorithm \mathcal{A} whose aim is to invert all the images in \mathbf{y} , the success probability of \mathcal{A} is upper bounded by:

$$\Pr_H[\mathcal{A}(\mathbf{y}) \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = y_i \ \forall i \in [k]] \leq 2^{2k} \binom{q+k}{k}^2 \cdot \frac{k!}{N^k}.$$

Hardness of Multi-Collision Finding. Secondly, we can analyse the quantum hardness of finding k collisions, namely, k inputs that map to the same image of the random oracle $H : [M] \rightarrow [N]$. We can also determine upper bounds for solving the salted version of this task, as well as the hardness of finding a collision for quantum algorithms that are also equipped with advice.

Lemma 3 (Quantum Hardness of Multi-Collision Finding and Salted Multi-Collision Finding (Informal)). For any q -quantum query algorithm \mathcal{A} , the probability of solving the k -multi-collision problem is at most:

$$\Pr_H[\mathcal{A}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_1) = \dots = H(x_k)] \leq \frac{1}{N^{k-1}} \left[\frac{2e(q+k)}{k} \right]^{2k}.$$

Any quantum algorithm \mathcal{A} equipped with q quantum queries and S -bit of classical advice can win the salted multi-collision finding game with salted space $[K]$ with probability at most:

$$\Pr_H[\mathcal{A}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_1) = \dots = H(x_k)] \leq \frac{4}{N^{k-1}} \left[\frac{2e(q+k)}{k} \right]^{2k} + \frac{4S}{K}.$$

The above bounds become $O(q^4/N)$ and $O(q^4/N + S/K)$ respectively, for $k = 2$ (the standard collision-finding). Previous work [YZ21] achieves the same uniform bound, but only achieves $O((Sq)^4/N + S/K)$, due to the extra loss in their lifting theorem.

Hardness of Multi-search. Finally, we also establish a tight bound for finding k distinct inputs that all map to 0 under the random oracle $H : [M] \rightarrow [N]$. This is potentially useful in analyzing proofs-of-work in the blockchain context [GKL15].

Lemma 4 (Quantum Hardness of Multi-Search). *For any q -quantum query algorithm \mathcal{A} whose task is to find k different preimages of 0 of the random oracle, the success probability of \mathcal{A} is upper bounded by:*

$$\Pr_H[\mathcal{A}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = 0 \forall i \in [k]] \leq \left[\frac{4e^2(q+k)^2}{Nk^2} \right]^k.$$

1.2 Related Work

The measure-and-reprogram framework was proposed, and subsequently generalized and improved with tighter bounds in the works of [DFM20, LZ19, DFMS22, GHM21]. A main application of the framework has been in the context of the Fiat-Shamir transformation, with several works establishing its post-quantum security [Cha19, DFMS19, AFK22, AFKR23, GOP+23]. Other cryptographic applications of measure-and-reprogram have been considered in [Kat21, BKS21, ABKK23, JMZ23, DFHS23, KX24]. Finally, applications in query complexity of the framework have been developed in [CGLQ20, YZ21].

2 Preliminaries

Notation. For two vectors $\mathbf{x}, \mathbf{x}' \in X^k$, we say $\mathbf{x} \equiv \mathbf{x}'$ if and only if there exists a permutation σ over the indices $\{1, 2, \dots, k\}$ such that $x'_i = x_{\sigma(i)}$. For a function $H : X \rightarrow Y$ and $\mathbf{x} \in X^k$, $H(\mathbf{x})$ is defined as $(H(x_1), H(x_2), \dots, H(x_k))$. We say $x \in \mathbf{x}$, if $x = x_i$ for some $i \in [k]$.

2.1 Quantum Query Algorithms

We will denote a quantum query algorithm by \mathcal{A} . Let q be the total number of quantum queries of \mathcal{A} . By \mathcal{A}^H we mean that \mathcal{A} has quantum access to the function H .

A quantum oracle query to H will be applied as the unitary $O_H: O_H|x\rangle|y\rangle \rightarrow |x\rangle|y \oplus H(x)\rangle$. Without loss of generality, we assume an algorithm will never perform any measurement (until the very end) and thus the internal state is always pure. We use $|\phi_i^H\rangle$ to denote the algorithm \mathcal{A} 's internal (pure) state right after the i -th query.

$$|\phi_i^H\rangle = O_H U_i \cdots O_H U_1 |0\rangle.$$

Specifically, we have,

- $|\phi_0^H\rangle = |0\rangle$ is the initial state of \mathcal{A} ;
- $|\phi_q^H\rangle$ is the final state of \mathcal{A} .

Without loss of generality, the algorithm will have three registers $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ at the end of the computation, where \mathcal{X} consists of a list of inputs, \mathcal{Y} consists of a list of outputs corresponding to these inputs and some auxiliary information in \mathcal{Z} .

Definition 1 (Reprogrammed Oracle). *Reprogram oracle H to output y on input x , results in the new oracle, defined as:*

$$H_{x,y}(z) = \begin{cases} y, & \text{if } z = x \\ H(z), & \text{otherwise.} \end{cases}$$

We can similarly define a multi-input reprogram oracle $H_{x,\Theta}$ for $x \in X^k$ without duplicate entries and $\Theta \in Y^k$:

$$H_{x,\Theta}(z) = \begin{cases} \Theta_i, & \text{if } z = x_i \\ H(z), & \text{otherwise.} \end{cases}$$

2.2 Quantum Measure-and-Reprogram Experiment

We recall the measure-and-reprogram experiment and the state-of-the-art results here, first proposed by [DFM20] and later adapted by [YZ21].

Definition 2 (Measure-and-Reprogram Experiment). *Let \mathcal{A} be a q -quantum query algorithm that outputs $x \in X^k$ and $z \in \mathcal{Z}$. For a function $H : X \rightarrow Y$ and $\mathbf{y} = (y_1, \dots, y_k) \in Y^k$, define a measure-and-reprogram algorithm $\mathcal{B}[H, \mathbf{y}]$:*

1. For each $j \in [k]$, uniformly pick $(i_j, b_j) \in ([q] \times \{0, 1\}) \cup \{(\perp, \perp)\}$ such that there does not exist $j \neq j'$ such that $i_j = i_{j'} \neq \perp$;
2. Run \mathcal{A}^O where the oracle O is initialized to be a quantumly accessible classical oracle that computes H and when \mathcal{A} makes its i -th query, the oracle is simulated as follows:
 - (a) If $i = i_j$ for some $j \in [k]$, measure \mathcal{A} 's query register to obtain x'_j and do either of the following:
 - i. If $b_j = 0$, reprogram O using (x'_j, y_j) and answer \mathcal{A} 's i_j -th query using the reprogrammed oracle;

- ii. If $b_j = 1$, answer \mathcal{A} 's i_j -th query using oracle before reprogramming and then reprogram O using (x'_j, y_j) ;
- (b) Else, answer \mathcal{A} 's i -th query by just using the oracle O without any measurement or reprogramming;
3. Let $(\mathbf{x} = (x_1, \dots, x_k), z)$ be \mathcal{A} 's output;
4. For all $j \in [k]$ such that $i_j = \perp$, set $x'_j = x_j$
5. Output $\mathbf{x}' := ((x'_1, \dots, x'_k), z)$

We next state the current state-of-the-art quantum measure-and-reprogram result.

Lemma 5 (Quantum Measure-and-Reprogram (adaptation from [DFM20, YZ21])). For any $H : X \rightarrow Y$, for any $\mathbf{x}^* = (x_1^*, \dots, x_k^*) \in X^k$ without duplicated entries, for all $\mathbf{y} = (y_1, \dots, y_k)$ and any relation $R \subseteq X^k \times Y^k \times Z$, we have:

$$\begin{aligned} \Pr[\mathbf{x}' = \mathbf{x}^* \wedge (\mathbf{x}', \mathbf{y}, z) \in R \mid (\mathbf{x}', z) \leftarrow \mathcal{B}[H, \mathbf{y}]] \\ \geq \frac{1}{(2q+1)^{2k}} \Pr[\mathbf{x} = \mathbf{x}^* \wedge (\mathbf{x}, \mathbf{y}, z) \in R \mid (\mathbf{x}, z) \leftarrow \mathcal{A}^{H_{\mathbf{x}^*, \mathbf{y}}}] \end{aligned}$$

where $\mathcal{B}[H, \mathbf{y}]$ is the measure-and-reprogram experiment.

2.3 Predicates and Success Probabilities

Definition 3 (Predicate/Verification Projection/Symmetric Predicate). Let R be a relation on $X^k \times Y^k \times Z$. A predicate $V^H(\mathbf{x}, \mathbf{y}, z)$ parameterized by an oracle H , returns 1 if and only if $(\mathbf{x}, \mathbf{y}, z) \in R$ and $H(x_i) = y_i$ for every $i \in \{1, 2, \dots, k\}$.

Let $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ be the registers that store $\mathbf{x}, \mathbf{y}, z$, respectively. We define Π_V^H as the projection corresponding to V^H :

$$\Pi_V^H |\mathbf{x}, \mathbf{y}, z\rangle = \begin{cases} |\mathbf{x}, \mathbf{y}, z\rangle & \text{if } V^H(\mathbf{x}, \mathbf{y}, z) = 1 \\ 0 & \text{otherwise} \end{cases}.$$

Finally, for any predicate V^H , we are able to establish the success probability using the projection Π_V^H .

Definition 4 (Success Probability). Let \mathcal{A} a quantum query algorithm. Its success probability of outputting $\mathbf{x}, \mathbf{y}, z$ such that $H(\mathbf{x}, \mathbf{y}, z) = 1$ is defined by

$$\Pr[\mathcal{A}^H \rightarrow (\mathbf{x}, \mathbf{y}, z) \text{ and } V^H(\mathbf{x}, \mathbf{y}, z) = 1] = \|\Pi_V^H |\phi_q^H\rangle\|^2.$$

(Recall that $|\phi_q^H\rangle$ is the final state of \mathcal{A} .)

Sometimes, we care about the event that \mathcal{A} outputs a particular \mathbf{x} and still succeeds. For any \mathbf{x}_o , the following probability denotes that \mathcal{A} outputs $\mathbf{x} \equiv \mathbf{x}_o$ and succeeds:

$$\Pr[\mathcal{A}^H \rightarrow (\mathbf{x}, \mathbf{y}, z), \mathbf{x} \equiv \mathbf{x}_o \text{ and } V^H(\mathbf{x}, \mathbf{y}, z) = 1] = \|G_{\mathbf{x}_o} \Pi_V^H |\phi_q^H\rangle\|^2,$$

where $G_{\mathbf{x}_o}$ is defined as the projection that checks whether \mathcal{A} consists of $\mathbf{x} \equiv \mathbf{x}_o$.

3 Coherent Measure-and-Reprogram

In this section, we give our main theorem: the coherent measure-and-reprogram theorem. A main difference between our theorem and the previous measure-and-reprogram theorem [DFM20] is that our simulator needs to make quantum queries, instead of classical queries, which is potentially required by the coherent nature of our simulator and gives tighter reprogramming bounds for many applications. While this makes the simulator slightly more complicated, it yields improved bounds on the various applications that we mention in the next section.

3.1 Main Theorem

We give our main theorem below.

Theorem 6. *Let $H, G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ be two functions $X \rightarrow Y$. Let k be a positive integer (can be a computable function in both n and m). There exists a black-box quantum algorithm $\text{Sim}^{H,G,A}$, satisfying the properties below. Let V^H be any predicate defined over $X^k \times Y^k \times Z$. Let \mathcal{A} be any q -quantum query algorithm to the oracle H . Then for any $\mathbf{x}_o \in X^k$ without duplicate entries and $\mathbf{y}_o = G(\mathbf{x}_o)$, we have,*

$$\Pr_{H,G} \left[\text{Sim}^{H,G,A} \rightarrow (\mathbf{x}, \mathbf{y}, z) \text{ and } \mathbf{x} \equiv \mathbf{x}_o \text{ and } V^{H_{\mathbf{x}_o, \mathbf{y}_o}}(\mathbf{x}, \mathbf{y}, z) = 1 \right] \\ \geq \frac{1}{2^{2k} \binom{q+k}{k}} \cdot \Pr_{H,G} \left[\mathcal{A}^{H_{\mathbf{x}_o, \mathbf{y}_o}} \rightarrow (\mathbf{x}, z) \text{ and } \mathbf{x} \equiv \mathbf{x}_o \text{ and } V^{H_{\mathbf{x}_o, \mathbf{y}_o}}(\mathbf{x}, H_{\mathbf{x}_o, \mathbf{y}_o}(\mathbf{y}), z) = 1 \right].$$

Furthermore, Sim makes exactly k quantum queries to G and has a running time polynomial in n, m, k and the running time of \mathcal{A} .

Before formally defining our simulator, we introduce one more notation: controlled reprogrammed oracle queries. That is, an oracle query will be reprogrammed by a list of input and output pairs in a control register.

Definition 5 (Controlled Reprogrammed Oracle Query). *For every $x \in X, y \in Y$, every $\ell > 0$ and $\mathbf{x} \in X^\ell$ without duplicated entries, $\Theta \in Y^\ell$, controlled reprogrammed oracle O_H^{ctrl} acts as below.*

$$O_H^{\text{ctrl}}|x\rangle|y\rangle|\mathbf{x}, \Theta\rangle = (O_{H_{\mathbf{x}, \Theta}}|x\rangle|y\rangle)|\mathbf{x}, \Theta\rangle$$

Its behavior on $|x\rangle$ with duplicated entries can be arbitrarily defined as long as unitarity is maintained, as this case will never occur in the simulator or our analysis.

We define our simulator used in Theorem 6, as follows:

Definition 6 (Coherent Measure-and-Reprogram Experiment). *Let \mathcal{A} be a $(q+k)$ -quantum query algorithm that outputs $\mathbf{x} = (x_1, \dots, x_k) \in X^k, \mathbf{y} \in Y^k$ and $z \in Z$. We assume \mathbf{y} is always computed by $H(\mathbf{x})$, using the last k queries. For a function $H : X \rightarrow Y$ and $G : X \rightarrow Y$, define a measure-and-reprogram algorithm $B[H, G]$:*

1. Pick a uniformly random subset \mathbf{v} of $[q + k]$, of length k . We have $1 < v_1 < \dots < v_k \leq q + k$. Pick $\mathbf{b} \in \{0, 1\}^k$ uniformly at random.
2. Run \mathcal{A} with an additional control register \mathcal{R} , initialized as empty $|\emptyset\rangle$. Define the following operation U that updates the control register: for x that is not in L ,

$$U|x\rangle|L\rangle_{\mathcal{R}} \leftarrow |x\rangle|L \cup (x, G(x))\rangle_{\mathcal{R}}.$$

Here L is the set of input and output pairs. Since we will only work with basis states $|x\rangle|L\rangle$ whose x is not in L , U clearly can be implemented by a unitary (by assuming that the list is initialized as empty).

3. When \mathcal{A} makes its i -th query,
 - (a) If $i = v_j$ for some $j \in [k]$, do either of the following:
 - i. If $b_j = 0$, update \mathcal{R} using the input register and G , and make the i -th query to H controlled by \mathcal{R} (see O_H^{ctrl} above);
 - ii. If $b_j = 1$, make the i -th query to H controlled by \mathcal{R} and update \mathcal{R} using the input register and G .
 - iii. Before updating the control register, it checks coherently that the input register is not contained in the control register; otherwise, it aborts.
 - (b) Else, answer \mathcal{A} 's i -th query controlled by \mathcal{R} ;
4. Let $(\mathbf{x}, \mathbf{y}, z)$ be \mathcal{A} 's output;
5. Measure \mathcal{R} register to obtain $L = (\mathbf{x}', \Theta')$.
6. Output $(\mathbf{x}, \mathbf{y}, z)$ if $\mathbf{x}' \equiv \mathbf{x}$; otherwise, abort.

At a high level, our simulator resembles that in Definition 2; instead of measuring \mathcal{A} 's queries, we put it into a separate register (a.k.a., measure the queries coherently). With the “controlled reprogrammed oracle query”, we are still able to progressively reprogram the oracle and run the algorithm with (an) updated oracle(s). The ability of coherently measuring and reprogramming, makes all the improvement (mentioned in the later sections) possible.

Proof of Theorem 6. Before we start with the proof, we first recall and introduce some notations. Fix any $\mathbf{x} \in X^k$ without duplicate entries and $\Theta \in Y^k$. Recall that in Sect. 2.1, $|\phi_q^{H_{\mathbf{x}, \Theta}}\rangle$ is the state of the algorithm \mathcal{A} after making all its queries to $H_{\mathbf{x}, \Theta}$. More precisely, it is:

$$|\phi_q^{H_{\mathbf{x}, \Theta}}\rangle = O_{H_{\mathbf{x}, \Theta}} U_q \cdots O_{H_{\mathbf{x}, \Theta}} U_1 |0\rangle.$$

In the next step, we decompose this quantum state, so that each component corresponds to one of the cases in the quantum simulator Definition 6.

The First Query. We start by considering the state up to the first query: $O_{H_{\mathbf{x},\Theta}} U_1|0\rangle$. We insert an additional identity operator and have,

$$\begin{aligned}
 O_{H_{\mathbf{x},\Theta}} U_1|0\rangle &= O_{H_{\mathbf{x},\Theta}} I U_1|0\rangle \\
 &\stackrel{(i)}{=} O_{H_{\mathbf{x},\Theta}} \left(I - \sum_{x_j} |x_j\rangle\langle x_j| + \sum_{x_j} |x_j\rangle\langle x_j| \right) U_1|0\rangle \\
 &= O_{H_{\mathbf{x},\Theta}} \left(I - \sum_{x_j} |x_j\rangle\langle x_j| \right) U_1|0\rangle + O_{H_{\mathbf{x},\Theta}} \left(\sum_{x_j} |x_j\rangle\langle x_j| \right) U_1|0\rangle \\
 &\stackrel{(ii)}{=} O_H \left(I - \sum_{x_j} |x_j\rangle\langle x_j| \right) U_1|0\rangle + \sum_{x_j} O_{H_{x_j,\Theta_j}} |x_j\rangle\langle x_j| U_1|0\rangle \\
 &= \underbrace{O_H U_1|0\rangle}_{(1)} - \sum_{x_j} \underbrace{O_H |x_j\rangle\langle x_j| U_1|0\rangle}_{(2)} + \sum_{x_j} \underbrace{O_{H_{x_j,\Theta_j}} |x_j\rangle\langle x_j| U_1|0\rangle}_{(3)}
 \end{aligned}$$

Above, x_j is enumerated over all entries in \mathbf{x} .

Line (i) follows easily. Line (ii) is due to the fact that, if the query input is not in \mathbf{x} , $H_{\mathbf{x},\Theta}$ is functionally equivalent to H ; similarly, if the query input is x_j , $H_{\mathbf{x},\Theta}$ is functionally equivalent to H_{x_j,Θ_j} .

Next, we look at the three terms (1), (2), (3):

- (1) $O_H U_1|0\rangle$ corresponds to the case that no measurement happens for the first query.
- (2) $O_H |x_j\rangle\langle x_j| U_1|0\rangle$ corresponds to the case that measurement is made at the first query and the query input is x_j ; the oracle is not reprogrammed immediately. In other words, the case $(v_1, b_1) = (1, 1)$ in the simulator.
- (3) $O_{H_{x_j,\Theta_j}} |x_j\rangle\langle x_j| U_1|0\rangle$ corresponds to the case that measurement is made at the first query and the query input is x_j ; the oracle is reprogrammed immediately and used for the first query. In other words, the case $(v_1, b_1) = (1, 0)$ in the simulator.

The Second Query. We do the same: insert an additional identity operator. To make the presentation clearer, we focus only on one term $O_H |x_j\rangle\langle x_j| U_1|0\rangle$; the other cases are simpler.

$$\begin{aligned}
& O_{H_{\mathbf{x},\Theta}} U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& \stackrel{(1)}{=} O_{H_{\mathbf{x},\Theta}} \left(I - \sum_{x_k \neq x_j} |x_k\rangle \langle x_k| + \sum_{x_k \neq x_j} |x_k\rangle \langle x_k| \right) U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& \stackrel{(2)}{=} O_{H_{\mathbf{x},\Theta}} \left(I - \sum_{x_k \neq x_j} |x_k\rangle \langle x_k| \right) U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& \quad + O_{H_{\mathbf{x},\Theta}} \left(\sum_{x_k \neq x_j} |x_k\rangle \langle x_k| \right) U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& \stackrel{(3)}{=} O_{H_{x_j,\Theta_j}} \left(I - \sum_{x_k \neq x_j} |x_k\rangle \langle x_k| \right) U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& \quad + \left(\sum_{x_k \neq x_j} O_{H_{(x_j,x_k),(\Theta_j,\Theta_k)}} |x_k\rangle \langle x_k| \right) U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle \\
& = \underbrace{O_{H_{x_j,\Theta_j}} U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle}_{(1)} - \underbrace{\sum_{x_k \neq x_j} O_{H_{x_j,\Theta_j}} |x_k\rangle \langle x_k| U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle}_{(2)} \\
& \quad + \underbrace{\sum_{x_k \neq x_j} O_{H_{(x_j,x_k),(\Theta_j,\Theta_k)}} |x_k\rangle \langle x_k| U_2 O_H |x_j\rangle \langle x_j| U_1 |0\rangle}_{(3)}
\end{aligned}$$

We explain the equations line by line:

1. This one is straightforward by realizing the summation inside the bracket is an identity operator.
 2. This one follows from the distributive property.
 3. This is the most important one.
 - For the first term, we realize that the oracle will only be applied to inputs that are not in \mathbf{x} , or are equal to x_j . Thus, $H_{\mathbf{x},\Theta}$ is functionally equivalent to H_{x_j,Θ_j} .
 - For the second term, the oracle will only be applied to inputs that are equal to x_k . Thus, $H_{\mathbf{x},\Theta}$ is functionally equivalent to $H_{(x_j,x_k),(\Theta_j,\Theta_k)}$ ².
- (1) corresponds to the case that no measurement happens for the second query, but since the first query is measure-and-reprogrammed, the second query is made with the oracle H_{x_j,Θ_j} . In other words, the case $(v_1, b_1) = (1, 1)$.
- (2) corresponds to the case that measurement is made at the second query and the query input is x_k ; the oracle is not reprogrammed immediately. In other words, the case $(v_1, b_1) = (1, 1)$ and $(v_2, b_2) = (2, 1)$ in the simulator.

² It is also equivalent to H_{x_k,Θ_k} . However, due to our description of the simulator, $H_{(x_j,x_k),(\Theta_j,\Theta_k)}$ is more natural to work with.

- (3) corresponds to the case that measurement is made at the second query and the query input is x_k ; the oracle is reprogrammed immediately. In other words, the case $(v_1, b_1) = (1, 1)$ and $(v_2, b_2) = (2, 0)$ in the simulator.

Generalization to all Queries—State Decomposition. By repeating the same state decomposition up to the first q queries (instead of all $q + k$ queries), we will end up a collection of subnormalized states, who sum up to the original state $|\phi_q^{H_{x,\theta}}\rangle$. These states are parameterized by when the measurement happens (an ordered vector \mathbf{v} such that $1 \leq v_1 \cdots \leq v_t \leq q$), whether these queries are made before or after each reprogramming ($\mathbf{b} \in \{0, 1\}^t$), and $t \in \{0, \dots, q\}$; in the following we will denote these states by $|\phi_{\mathbf{v}, \mathbf{b}}\rangle$. For example, assuming $\mathbf{b} = \mathbf{0}$ (all reprogramming happens immediately), we have,

$$\begin{aligned}
 |\phi_{\mathbf{v}, \mathbf{0}}\rangle &= \sum_{\sigma \in S_t^k} O_{H_{x_{\sigma_1}, \theta_{\sigma_1}}} U_q \cdots O_{H_{x_{\sigma_t}, \theta_{\sigma_t}}} U_{v_t+1} \underbrace{O_{H_{x_{\sigma_t}, \theta_{\sigma_t}}} |x_{\sigma_t}\rangle \langle x_{\sigma_t}| \cdots U_{v_{t-1}+1}}_{\text{stage (t)}} \\
 &\quad \cdots \\
 &\quad \cdot \underbrace{O_{H_{(x_{\sigma_1}, x_{\sigma_2}), (\theta_{\sigma_1}, \theta_{\sigma_2})}} |x_{\sigma_2}\rangle \langle x_{\sigma_2}| U_{v_2} \cdots O_{H_{x_{\sigma_1}, \theta_{\sigma_1}}} U_{v_1+1}}_{\text{stage (2)}} \\
 &\quad \cdot \underbrace{O_{H_{x_{\sigma_1}, \theta_{\sigma_1}}} |x_{\sigma_1}\rangle \langle x_{\sigma_1}| U_{v_1} O_H \cdots O_H U_1 |0\rangle}_{\text{stage (1)}}
 \end{aligned}$$

Here S_t^k denotes all ordered list of length t , with elements in $\{1, \dots, k\}$ without duplication; \mathbf{x}_σ denotes $(x_{\sigma_1}, \dots, x_{\sigma_t})$ and Θ_σ denotes $(\theta_{\sigma_1}, \dots, \theta_{\sigma_t})$. We can similarly define $|\phi_{\mathbf{v}, \mathbf{b}}\rangle$ for all other $\mathbf{b} \in \{0, 1\}^t$, the only difference here is the oracle may not be immediately reprogrammed at the end of each stage. More generally, for each \mathbf{v} of length t and $\mathbf{b} \in \{0, 1\}^t$, we define

$$|\phi_{\mathbf{v}, \mathbf{b}}\rangle = \sum_{\sigma \in S_t^k} |\phi_{\mathbf{v}, \mathbf{b}, \sigma}\rangle,$$

where $|\phi_{\mathbf{v}, \mathbf{b}, \sigma}\rangle$ is the state that is measured-and-reprogrammed according to \mathbf{v}, \mathbf{b} with the order σ , similar to that in the definition of $|\phi_{\mathbf{v}, \mathbf{0}}\rangle$. Thus, we have:

$$|\phi_q^{H_{x,\theta}}\rangle = \sum_{\mathbf{v}, \mathbf{b}} |\phi_{\mathbf{v}, \mathbf{b}}\rangle,$$

Adding the Extra k Queries. We assume the algorithm \mathcal{A} , after the first q queries, already prepares the output \mathbf{x}, z . We will force \mathcal{A} making the last k queries, to generate $\mathbf{y} = H(\mathbf{x})$. Recall the definitions $\Pi_V^{H_{x_0, y_0}}$ and G_{x_0} in Definition 4. By setting $\mathbf{x} = \mathbf{x}_0$ and $\Theta = \mathbf{y}_0$ in the above analysis, the probability on the RHS in the theorem we are proving is equal to:

$$\begin{aligned}
 &\Pr_{H, G} [\mathcal{A}^{H_{x_0, y_0}} \rightarrow (\mathbf{x}, z) \text{ and } \mathbf{x} \equiv \mathbf{x}_0 \text{ and } V^{H_{x_0, y_0}}(\mathbf{x}, H_{x_0, y_0}(\mathbf{x}), z) = 1] \\
 &= \left\| G_{x_0} \Pi_V^{H_{x_0, y_0}} \left| \phi_{q+k}^{H_{x_0, y_0}} \right. \right\|^2.
 \end{aligned}$$

Since G_{x_o} and $\Pi_V^{H_{x_o, y_o}}$ commute (as they are both projections over computational basis), we can assume G_{x_o} is applied to the state first. Even further, as the computation of $\mathbf{y} = H(\mathbf{x})$ and the projection G_{x_o} also commute, we can assume G_{x_o} applies to the state right before the last k queries, which are used to compute \mathbf{y} . Therefore, for every $|\phi_{\mathbf{v}, \mathbf{b}, \sigma}\rangle$, even if $t < k$ (the length of \mathbf{v}), we can measure-and-(immediately)-reprogram exactly $k - t$ locations of the last k queries, and making the random oracle exactly reprogrammed to H_{x_o, y_o} .

Thus, we have:

$$\left| \phi_{q+k}^{H_{x_o, y_o}} \right\rangle = \sum_{\substack{\mathbf{v}, \mathbf{b} \\ |\mathbf{v}|=k}} |\phi_{\mathbf{v}, \mathbf{b}}\rangle, \tag{1}$$

where the RHS has (at most) $2^k \binom{q+k}{k}$ terms.

By Eq. (1), Cauchy-Schwartz and the triangle inequality, we have:

$$\left\| G_{x_o} \Pi_V^{H_{x_o, y_o}} \left| \phi_{q+k}^{H_{x_o, y_o}} \right\rangle \right\|^2 \leq 2^k \binom{q+k}{k} \sum_{\substack{\mathbf{v}, \mathbf{b} \\ |\mathbf{v}|=t}} \left\| G_{x_o} \Pi_V^{H_{x_o, y_o}} |\phi_{\mathbf{v}, \mathbf{b}}\rangle \right\|^2. \tag{2}$$

Finally, to prove the theorem statement, we relate each individual term on the RHS with the behaviors of our simulator B .

Relating Each Term with Our Simulator B. Next, we prove that each term

$\left\| G_{x_o} \Pi_V^{H_{x_o, y_o}} |\phi_{\mathbf{v}, \mathbf{b}}\rangle \right\|^2$ is upper bounded by the probability that when the simulator B picks \mathbf{v}, \mathbf{b} , it succeeds and outputs $\mathbf{x} \equiv \mathbf{x}_o$, which we denote by $p_{x_o, \mathbf{v}, \mathbf{b}}$.

Since the simulator B ensures that (1) no duplicated elements ever in the control register, (2) at the end, the control register only consists of inputs that are outputted by A (which will be \mathbf{x}_o , enforced by G_{x_o}), we have that $p_{x_o, \mathbf{v}, \mathbf{b}}$ is the squared norm of the state $(G_{x_o} \Pi_V^{H_{x_o, y_o}} \otimes I_{\mathcal{R}}) |\psi_{\mathbf{v}, \mathbf{b}}\rangle$, with the state $|\psi_{\mathbf{v}, \mathbf{b}}\rangle$ being:

$$|\psi_{\mathbf{v}, \mathbf{b}}\rangle = \sum_{\sigma \in S_k^k} |\phi_{\mathbf{v}, \mathbf{b}, \sigma}\rangle \otimes |\text{set}\{(x_{o, \sigma_1}, y_{o, \sigma_1}), \dots, (x_{o, \sigma_k}, y_{o, \sigma_k})\}\rangle_{\mathcal{R}}.$$

The only difference between $|\psi_{\mathbf{v}, \mathbf{b}}\rangle$ and $|\phi_{\mathbf{v}, \mathbf{b}}\rangle$ is the extra control register! However, we realize that in this case, when σ is a permutation of $[k]$, the control register is unentangled, making $p_{x_o, \mathbf{v}, \mathbf{b}}$ is equal to $\left\| G_{x_o} \Pi_V^{H_{x_o, y_o}} |\phi_{\mathbf{v}, \mathbf{b}}\rangle \right\|^2$. This is because the set will simply be $\text{set}\{(x_{o, 1}, y_{o, 1}), \dots, (x_{o, k}, y_{o, k})\}$, regardless of what σ is.

Finally, we have the L.H.S. is equal to

$$\begin{aligned} & \Pr_{H, G} \left[\text{Sim}^{H, G, A} \rightarrow (\mathbf{x}, \mathbf{y}, z) \text{ and } \mathbf{x} \equiv \mathbf{x}_o \text{ and } V^{H_{x_o, y_o}}(\mathbf{x}, \mathbf{y}, z) = 1 \right] \\ &= \frac{1}{2^k \binom{q+k}{k}} \sum_{\substack{\mathbf{v}, \mathbf{b} \\ |\mathbf{v}|=k}} p_{x_o, \mathbf{v}, \mathbf{b}} \end{aligned}$$

Thus, combining Eq. (2) and the equation above, we have:

$$\begin{aligned}
 \text{R.H.S.} &= \left\| G_{x_o} \Pi_V^{H_{x_o \cdot y_o}} \left| \phi_q^{H_{x_o \cdot y_o}} \right\rangle \right\|^2 \leq 2^k \binom{q+k}{k} \sum_{\substack{v,b \\ |v|=k}} p_{x_o, v, b} \\
 &\leq \left(2^k \binom{q+k}{k} \right)^2 \cdot \frac{1}{2^{k \binom{q+k}{k}}} \sum_{\substack{v,b \\ |v|=k}} p_{x_o, v, b} \\
 &= \text{L.H.S.}
 \end{aligned}$$

where L.H.S. and R.H.S. denote the left/right-hand side term in the theorem statement. Therefore, we conclude the proof. \square

Lemma 6 (Coherent Measure-and-Reprogram results in Uniform Images). *Consider the Coherent Measure-and-Reprogram Experiment in Definition 6, but where we choose G to be uniformly random. Then, for the measure-and-reprogram algorithm \mathcal{B} , the measurement $L = (x', \Theta')$ of the \mathcal{R} register (in Step 5) will result in uniformly random images Θ' .*

Proof. We will proceed with a proof by induction over the number of quantum queries of \mathcal{A} . In this proof, we will denote by n the total number of queries ($n = q+k$). For $n = 1$, let $v_j = n = 1$. Then, if $b_j = 0$, after updating the register \mathcal{R} using the unitary U (in step 2), the register \mathcal{R} will contain superposition of L sets consisting of a single pair $(x, G(x))$. Then, we perform the query to H controlled by \mathcal{R} using O_H^{ctrl} , which is a query to the reprogrammed H on single points x , modifying accordingly the image register ($y \rightarrow y \oplus H_{x, G(x)}$), but which does not affect the \mathcal{R} register. As G is random oracle, measuring \mathcal{R} will result in a uniform image $\theta' = G(x)$, for some $x \in X$. If $b_j = 1$, we first query using O_H^{ctrl} , which is a query to the original H as L is empty. Then, we update \mathcal{R} using unitary U , resulting in \mathcal{R} containing superposition of L sets consisting of a single pair $(x, G(x))$. As before, measuring \mathcal{R} will result in a uniform image $\theta' = G(x)$, for some $x \in X$. We emphasize that although Definition 6 defines G as an arbitrary function, in the statement of this Lemma, we consider uniformly random G instead of an arbitrary G .

For the inductive step, suppose that up to query $n-1$, the register \mathcal{R} consists of sets L' with uniform images. Let \mathcal{A} make its n -th query. If there does not exist any v_j equal to n then algorithm \mathcal{B} answer \mathcal{A} 's query controlled by \mathcal{R} , reprogramming the oracle with the inputs and outputs pairs in L' , but importantly \mathcal{R} remains unchanged, hence \mathcal{R} contains only uniform images by our inductive hypothesis. Otherwise, suppose there exists j^* such that $v_{j^*} = n$. Then if $b_{j^*} = 0$, we are first going to add in L the pair $(x, G(x))$ if x is not already in L , i.e. $L = L' \cup \{(x, G(x))\}$, otherwise $L = L'$. We are then going to make the controlled query O_H^{ctrl} to the reprogrammed oracle O_{H_L} , which does not affect the register \mathcal{R} . Hence by measuring \mathcal{R} results in either $(x', \theta') \in L'$, which by hypothesis contains uniform image θ' or in $(x, G(x))$, which given that G is a random oracle, also results in a uniform image. Similarly, if $b_{j^*} = 1$ we are first

going to make the controlled query O_H^{ctrl} to the reprogrammed oracle $O_{H'_L}$, then we will update the register \mathcal{R} using the unitary U , which as before will either result in either $L = L' \cup \{(x, G(x))\}$ or $L = L'$. In both cases, by measuring \mathcal{R} we will get a uniform image by using the uniformity of G and the inductive hypothesis. \square

4 Applications

4.1 Query Complexity

We will begin by first introducing the family of (security) games for which we will establish their quantum query complexity, namely the hardness of a quantum adversary to win such games.

Definition 7 (Multi-Output k -Search Game (Single-Instance)). *Let the random oracle $H : [M] \rightarrow [N]$, a distribution over challenges π_H and a winning relation $R_{H, ch}$ defined over Y^k .*

Then we define the multi-output k -search game \mathcal{G} as follows:

1. *Challenger samples randomness ch and sends it to a quantum algorithm \mathcal{A} having (quantum) oracle access to H ;*
2. *Adversary \mathcal{A} gets oracle access to H and outputs $\mathbf{x} := (x_1, \dots, x_k), z$;*
3. *Challenger queries \mathbf{x} to the random oracle, resulting in $\mathbf{y} := (y_1 = H(x_1), \dots, y_k = H(x_k))$ and checks if they satisfy the winning relation:
 $b := (x_1, \dots, x_k, y_1, \dots, y_k, z) \in R_{H, ch}$;*
4. *If $b = 1$, \mathcal{A} wins the \mathcal{G} game.*

We will denote by $\epsilon_{\mathcal{G}}(q)$ the maximum probability over all q -quantum algorithms \mathcal{A} of winning the multi-output k -search game \mathcal{G} .

Our main result is a quantum lifting theorem in the average case, relating the success probability of an arbitrary quantum algorithm to win a multi-output k -search game with the probability of success of a quantum algorithm equipped with exactly k quantum queries.

Theorem 7 (Lifting for Multi-Output k -Search Games). *Let \mathcal{G} be a multi-output k -search game (as defined in Definition 7). Let \mathcal{A} be a q -quantum query adversary in the game \mathcal{G} (against the k -classical query challenger \mathcal{C}). Then there exists a k -quantum query adversary \mathcal{B} against the game such that:*

$$\Pr[\mathcal{B}^{[H]} \text{ wins } \mathcal{G}] \geq \frac{1}{2^{2k} \binom{q+k}{k}^2} \Pr[\mathcal{A}^{[H]} \text{ wins } \mathcal{G}].$$

Proof. We will show that our Coherent Reprogramming result in Theorem 6 implies the lifting theorem. We will now show how to instantiate the coherent reprogramming theorem. Let \mathbf{x}_o be uniformly sampled from X^k . Let $H', G' : X \rightarrow Y$ be two uniform random oracles. Then, it is clear that, as $\mathbf{y}_o = G'(\mathbf{x}_o)$

is also uniform over Y^k , the reprogrammed function $H'_{x_o, y_o} : X \rightarrow Y$ is a uniform random function; this is due to the fact that, as stated in Theorem 6 (invoked here), the tuple x_o has distinct values for its element. We will instantiate the random oracle in the game \mathcal{G} as the function H'_{x_o, y_o} . Assume that in the game \mathcal{G} after receiving the challenge and after performing its q quantum queries to H'_{x_o, y_o} , the adversary \mathcal{A} returns to the Challenger the outcome x . Then, the Challenger queries x to H'_{x_o, y_o} resulting in y and checks if y satisfies the winning relation $R_{H'_{x_o, y_o}, ch}$. Define $V^{H'_{x_o, y_o}}$ as the predicate that outputs 1 if $y \in R_{H'_{x_o, y_o}, ch}$ and 0 else. In this way, we observe that the probability that \mathcal{A} wins the game \mathcal{G} is exactly the RHS of Theorem 6. As a result, by Theorem 6, there must exist an efficient quantum simulator $\text{Sim}^{H', G', \mathcal{A}}$ performing k quantum queries that also wins the game \mathcal{G} . Hence, it suffices to instantiate \mathcal{B} as the simulator Sim . \square

Let L_C represent the set of (classical) queries that a challenger performs during a multi-output k -search game \mathcal{G} (Definition 7). For a quantum query adversary \mathcal{B} against \mathcal{G} , we will denote by L_B the result of measuring its input and output query registers. Now, for the query complexity applications we will need the following stronger lifting theorem, which intuitively additionally guarantees the existence of an algorithm against \mathcal{G} such that at the end of the game, measuring its input and output registers gives us exactly the set of queries of the challenger.

Theorem 8 (Lifting for Search Game with Uniform Images). *Let \mathcal{G} be a multi-output k -search game (as defined in Definition 7). Let \mathcal{A} be a q -quantum query adversary in the game \mathcal{G} (against the k -classical query challenger \mathcal{C}). Then there exists a k -quantum query adversary \mathcal{B} such that L_B is uniform, satisfying:*

$$\Pr[\mathcal{B}^{H'} \text{ wins } \mathcal{G} \text{ and } L_C = L_B] \geq \frac{1}{2^{2k} \binom{q+k}{k}^2} \Pr[\mathcal{A}^{H'} \text{ wins } \mathcal{G}].$$

Proof. The simulator algorithm \mathcal{B} will follow the outline of the algorithm in the proof of Theorem 7, with the only difference that \mathcal{B} will perform an additional step at the end. Namely, after interaction with Challenger \mathcal{C} , compute list of queries of \mathcal{C} as L_C . If any query in L_C has not yet been queried by \mathcal{B} , \mathcal{B} will query them to oracle H . The uniformity of L_B follows directly from Lemma 6. \square

4.2 A New Quantum Lifting Theorem and Direct Product Theorem for Image Relations

Our first quantum lifting result (in Theorem 7) gives a first bound on the quantum hardness of solving any multi-output k -search game \mathcal{G} by relating it to the probability of \mathcal{G} being solved by a quantum algorithm with a small number of quantum queries. In this section we can derive a stronger quantum lifting theorem for the class of relations that only depend on images.

Theorem 9 (Quantum Lifting Theorem for Image Relations). *For any quantum algorithm \mathcal{A} equipped with q quantum queries, \mathcal{A} 's success probability to solve the multi-output k -search game specified by the winning relation R , is bounded by:*

$$\Pr[\mathcal{A}^{|H\rangle} \text{ wins multi-output } k\text{-search game}] \leq 2^{2k} \binom{q+k}{k}^2 \Pr[\exists \text{ perm } \pi \mid (y_{\pi(1)}, y_{\pi(2)}, \dots, y_{\pi(k)}) \in R : (y_1, \dots, y_k) \stackrel{\$}{\leftarrow} Y^k].$$

For simplicity, in the rest of the section, we define $p(R)$ as:

$$p(R) = \Pr[\exists \text{ perm } \pi \mid (y_{\pi(1)}, y_{\pi(2)}, \dots, y_{\pi(k)}) \in R : (y_1, \dots, y_k) \stackrel{\$}{\leftarrow} Y^k].$$

Proof. Let \mathcal{G} be a multi-output k -search game and assume a q -quantum adversary \mathcal{A} sends to the Challenger the answer $\mathbf{x} = (x_1, \dots, x_k)$. Challenger \mathcal{C} will accept if and only if $\mathbf{y} := (H(x_1), \dots, H(x_k)) \in R_{H, ch}$ and if x_i, x_j are pairwise distinct. By Theorem 8 we know there exists a quantum algorithm \mathcal{B} making k quantum queries to H winning the game such that $L_{\mathcal{B}} = L_{\mathcal{C}}$ with success probability at least the success probability of \mathcal{A} multiplied by a factor of $2^{2k} \binom{q+k}{k}^2$. The condition $L_{\mathcal{B}} = L_{\mathcal{C}}$ implies that \mathcal{C} will verify as the images of \mathcal{B} 's answer exactly a permutation of the recorded information in $L_{\mathcal{B}}$. Therefore, due to the property of Theorem 8 that $L_{\mathcal{B}}$ will be uniformly over Y^k , \mathcal{B} 's winning probability will be lower bounded by the probability that there exists a permutation such that for uniformly sampled images from Y^k , the permuted images will belong to our target relation:

$$\Pr[\mathcal{A}^{|H\rangle} \text{ wins multi-output } k\text{-search game}] \leq 2^{2k} \binom{q+k}{k}^2 p(R).$$

□

Next, we show a Direct Product Theorem for Image Relations.

Definition 8 (Direct Product). *Let \mathcal{G} be a multi-output k -search game specified by the winning relation R , with respect to a random oracle $[M] \rightarrow [N]$. Define the following Direct Product $\mathcal{G}^{\otimes g}$:*

- Let H be a random oracle $[g] \times [M] \rightarrow [N]$, and H_i denotes $H(i, \cdot)$;
- Challenger samples ch_i as in \mathcal{G} for $i \in \{1, \dots, g\}$.
- Adversary \mathcal{A} gets oracle access to H and outputs $\mathbf{x}_1, \dots, \mathbf{x}_g, z_1, \dots, z_g$ such that each input in \mathbf{x}_i start with i .
- Challenger computes $b_i := (\mathbf{x}_i, H(\mathbf{x}_i), z_i) \in R_{H_i, ch_i}$;
- If all b_i equal to 1, \mathcal{A} wins the $\mathcal{G}^{\otimes g}$ game.

Theorem 10 (Direct Product Theorem for Image Relations). *For any quantum algorithm \mathcal{A} equipped with gq quantum queries, \mathcal{A} 's success probability to solve the Direct Product $\mathcal{G}^{\otimes g}$ with the underlying \mathcal{G} specified by the winning relation R , is bounded by*

$$\Pr[\mathcal{A}^{|H\rangle} \text{ wins } \mathcal{G}^{\otimes g}] \leq \left(2^{2k} \binom{q+k}{k}^2 p(R) \right)^g.$$

Proof. Let \mathcal{G} be a multi-output k -search game and assume a gq -quantum adversary \mathcal{A} sends to the Challenger the answer $\mathbf{x}_1, \dots, \mathbf{x}_g, z_1, \dots, z_g$. By Theorem 8 we know there exists a quantum algorithm \mathcal{B} making gk quantum queries to H winning the game such that $L_{\mathcal{B}} = L_{\mathcal{C}}$ with success probability at least the success probability of \mathcal{A} multiplied by a factor of $2^{2gk} \binom{gq+gk}{gk}^2$. The condition $L_{\mathcal{B}} = L_{\mathcal{C}}$ implies that \mathcal{C} will verify as the images of \mathcal{B} 's answer exactly a permutation of the recorded information in $L_{\mathcal{B}}$. Moreover, for every image y , its associated input x only belongs to one of the oracles $H(i, \cdot)$; thus, it can only contribute to one of the relation checks R_{H_i, ch_i} . Thus, the permutation of the recorded information can only permute images with respect to the same oracle H_i .

Therefore, due to the property of Theorem 8 that $L_{\mathcal{B}}$ will be uniformly over Y^{gk} , \mathcal{B} 's winning probability will be lower bounded by the probability that there exists a permutation such that for uniformly sampled images from Y^{gk} , the permuted images will belong to our target relation:

$$\begin{aligned} & \Pr[\mathcal{A}^{|H\rangle} \text{ wins } \mathcal{G}^{\otimes g}] \\ & \leq 2^{2gk} \binom{gq+gk}{gk}^2 \Pr[\exists \pi_1, \dots, \pi_g \in \mathcal{S}_k \mid (y_{i, \pi_i(1)}, \dots, y_{i, \pi_i(k)}) \in R : (y_{i,1}, \dots, y_{i,k}) \stackrel{\$}{\leftarrow} Y^k] \\ & \leq \left(2^{2k} \binom{q+k}{k}^2 \Pr[\exists \pi \in \mathcal{S}_k \mid (y_{\pi(1)}, y_{\pi(2)}, \dots, y_{\pi(k)}) \in R : (y_1, \dots, y_k) \stackrel{\$}{\leftarrow} Y^k] \right)^g \\ & \leq \left(2^{2k} \binom{q+k}{k}^2 p(R) \right)^g. \end{aligned}$$

□

In the following section, we will show some of the query complexity and cryptographic applications of our quantum lifting theorems and Direct Product Theorem.

4.2.1 Application 1: Non-uniform Security

Definition 9 (Advice Algorithms). We define an advice (non-uniform) algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ equipped with q queries and advice of length S as follows:

1. $\mathcal{A}_1^H \rightarrow |adv\rangle$: an unbounded algorithm \mathcal{A}_1 outputs the advice $|adv\rangle$ consisting of S qubits;
2. $\mathcal{A}_2^H(|adv\rangle, ch) \rightarrow x$: q -quantum algorithm \mathcal{A}_2 takes as input the quantum advice $|adv\rangle$ and a challenge ch , outputs answer x ;

We define $\epsilon_{\mathcal{G}}^C(q, S)$ as the maximum winning probability over any advice adversary \mathcal{A} equipped with q quantum queries and S classical bits of advice against the classically-verifiable search game \mathcal{G} .

We also consider multi-instance games, similar to Direct Product, except all the instances share the same oracle.

Definition 10 (Multi-Instance Game). Let \mathcal{G} be a multi-output k -search game specified by the winning relation R , with respect to a random oracle $[M] \rightarrow [N]$. Define the following Direct Product $\mathcal{G}_{\text{MIS}}^{\otimes g}$:

- Let H be a random oracle $[M] \rightarrow [N]$;
- Challenger samples ch_i as in \mathcal{G} for $i \in \{1, \dots, g\}$;
- Adversary \mathcal{A} gets oracle access to H and outputs $\mathbf{x}_1, \dots, \mathbf{x}_g, z_1, \dots, z_g$;
- Challenger computes $b_i := (\mathbf{x}_i, H(\mathbf{x}_i), z_i) \in R_{H, ch_i}$;
- If all b_i equal to 1, \mathcal{A} wins the $\mathcal{G}_{\text{MIS}}^{\otimes g}$ game.

From above, we can define $R_{\text{MIS}}^{\otimes g}$ as the winning relation for $\mathcal{G}_{\text{MIS}}^{\otimes g}$.

Lemma 7

(Multi-Output Implies Non-Uniform Classical Advice ([CGLQ20])). Let \mathcal{G} be a search game (as defined in Definition 7). If the maximum winning probability for any quantum algorithm equipped with q quantum queries against $\mathcal{G}_{\text{MIS}}^{\otimes g}$ is $\epsilon_{\mathcal{G}_{\text{MIS}}^{\otimes g}}(q)$, then the maximum probability of any non-uniform adversary equipped with q quantum queries and S -length classical advice against the original game \mathcal{G} is at most:

$$\epsilon_{\mathcal{G}}^C(q, S) \leq 4 \cdot \left[\epsilon_{\mathcal{G}_{\text{MIS}}^{\otimes g}}(Sq) \right]^{\frac{1}{S}}$$

By combining these results with our Quantum Lifting Theorem, we derive the security against advice (non-uniform) quantum algorithms.

Lemma 8 (Security against Advice Quantum Adversaries). Let \mathcal{G} be any multi-output k -search game specified by the winning relation R . Let $\mathcal{G}_{\text{MIS}}^{\otimes g}$ be the multi-instance game and $R_{\text{MIS}}^{\otimes g}$ be the relation. Any non-uniform algorithm \mathcal{A} equipped with q quantum queries and S classical bits of advice can win the game \mathcal{G} with probability at most:

$$\epsilon_{\mathcal{G}}^C(q, S) \leq 4 \cdot 2^{2k} \binom{S(q+k)}{Sk}^{\frac{2}{S}} \cdot p(R_{\text{MIS}}^{\otimes g}).$$

Proof. By combining our (strong) quantum lifting theorem (in Theorem 9) with the two advice results (Lemma 7). \square

4.2.2 Application 2: Salting Against Non-uniform Adversaries

Definition 11 (Salted Game). Let \mathcal{G} be a search game (as defined in Definition 7) specified by a random oracle $H : [M] \rightarrow [N]$, a distribution over challenges π_H and a winning relation $R_{H, ch}$ defined over Y . Then we define the salted version of \mathcal{G} as the game \mathcal{G}_s with salted space $[K]$ defined as follows:

1. The random oracle function is defined as: $G = (H_1, \dots, H_K)$ for K random functions $H_i : [M] \rightarrow [N]$;
2. For any such G , the challenge $ch := (i, ch_i)$ is produced by first sampling uniformly at random $i \in [K]$ and then sampling ch_i according to π_{H_i} ;
3. The winning relation is defined as $R_{G, ch} := R_{H_i, ch_i}$;

We will denote by $\epsilon_{\mathcal{G}_s}(q)$ the maximum probability over all q -quantum algorithms \mathcal{A} of winning the salted game \mathcal{G}_s .

Lemma 9 (Security of Salted Game against Classical Advice). *Let \mathcal{G} be a multi-output k -search game (as defined in Definition 7), specified by a relation R . Let \mathcal{G}_s be the salted game, with salt space $[K]$. Then we have,*

$$\epsilon_{\mathcal{G}_s}^C(q, S) \leq 4 \cdot \frac{S}{K} + 4 \cdot 2^{2k} \binom{q+k}{k}^2 p(R).$$

Proof. By Lemma 7, the non-uniform security is related to the multi-instance game $\mathcal{G}_{s, \text{MIS}}^{\otimes g}$, with salt space $[K]$. The security of the multi-instance game is closely related to the Direct Product, for salted games, as shown in [DLW24] (in the proof of Theorem 4.1). More precisely,

$$\epsilon_{\mathcal{G}_{s, \text{MIS}}^{\otimes g}}(gq)^{1/g} \leq \epsilon_{\mathcal{G}_s^{\otimes g}}(gq)^{1/g} + \frac{g}{K}.$$

Intuitively, the only difference between the multi-instance game and the Direct Product is that, the same salt can be sampled with duplication. The extra factor $\frac{g}{K}$ captures the fact that the salt can be duplicated. Combining with Theorem 10, we have:

$$\begin{aligned} \epsilon_{\mathcal{G}_s}^C(q, S) &\leq 4 \left(\epsilon_{\mathcal{G}_{s, \text{MIS}}^{\otimes S}}(Sq) \right)^{1/S} \\ &\leq 4 \left(\epsilon_{\mathcal{G}_s^{\otimes S}}(Sq)^{1/S} + \frac{S}{K} \right) \\ &\leq 4 \cdot \frac{S}{K} + 4 \cdot 2^{2k} \binom{q+k}{k}^2 p(R). \end{aligned}$$

□

4.2.3 Application 3: Multi-image Inversion

Our first result establishes the quantum hardness of multi-image inversion, which is a tight bound as already proven in [CGLQ20], but achieved here in a much simpler way, directly from our quantum lifting theorem.

Lemma 10 (Quantum Hardness of Multi-Image Inversion).

For any $\mathbf{y} = (y_1, \dots, y_k) \in Y^k = [N]^k$ (without duplicates) and for any q -quantum query algorithm \mathcal{A} whose task is to invert all the images in \mathbf{y} , the success probability of \mathcal{A} is upper bounded by:

$$\begin{aligned} &\Pr_H[\mathcal{A}^{[H]}(\mathbf{y}) \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = y_i \forall i \in [k]] \\ &\leq \left[\frac{4e(q+k)^2}{Nk} \right]^k \end{aligned}$$

Proof. We will show this using our strong quantum lifting theorem for image relations (Theorem 9). Define R as the relation over $[N]^k$, with $H : [M] \rightarrow [N]$ such that: $R = \{y_1, \dots, y_k\}$. Then for each permutation π , we have

$\Pr[(y_{\pi(1)}, \dots, y_{\pi(k)}) \in R \mid (y_1, \dots, y_k) \leftarrow [N]^k] = \frac{1}{N^k}$. Using that the number of permutations π is $k!$ leads to:

$$\Pr_H[\mathcal{A}^{(H)}(\mathbf{y}) \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = y_i \forall i \in [k]] \leq 2^{2k} \binom{q+k}{k}^2 \cdot \frac{k!}{N^k}$$

Using first the inequality $\binom{q+k}{k} \leq \frac{(q+k)^k}{k!}$ and then the Stirling approximation $k! \geq \left(\frac{k}{e}\right)^k$, we get:

$$\begin{aligned} 2^{2k} \binom{q+k}{k}^2 \cdot \frac{k!}{N^k} &\leq \left(\frac{4}{N}\right)^k \cdot k! \cdot \left[\frac{(q+k)^k}{k!}\right]^2 \\ &\leq \left(\frac{4}{N}\right)^k \cdot (q+k)^{2k} \cdot \left(\frac{e}{k}\right)^k = \left[\frac{4e(q+k)^2}{Nk}\right]^k \end{aligned}$$

□

4.2.4 Application 4: Multi-collision Finding and Multi-search

Next, we can determine the quantum hardness of the multi-collision problem, namely finding k different inputs that map to the same output of the random oracle.

Lemma 11 (Quantum Hardness of Multi-Collision Finding and Salted Multi-Collision Finding).

For any q -quantum query algorithm \mathcal{A} , we have the upper bound for solving the k -multi-collision problem:

$$\Pr_H[\mathcal{A}^{(H)}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_1) = \dots = H(x_k)] \leq \frac{1}{N^{k-1}} \left[\frac{2e(q+k)}{k}\right]^{2k}$$

Any quantum algorithm \mathcal{A} equipped with q quantum queries and S -bit of classical advice can win the salted multi-collision finding game with salted space $[K]$ with probability at most:

$$\Pr_H[\mathcal{A}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_1) = \dots = H(x_k)] \leq \frac{4}{N^{k-1}} \left[\frac{2e(q+k)}{k}\right]^{2k} + \frac{4S}{K}.$$

Proof. We will show this using our strong quantum lifting theorem for image relations (Theorem 9). Define $R := \{y, \dots, y\}_y$ the relation over $[N]^k$, where $H : [M] \rightarrow [N]$. Then for each permutation π , we have $\Pr[(y_{\pi(1)}, \dots, y_{\pi(k)}) \in R \mid (y_1, \dots, y_k) \leftarrow [N]^k] = \frac{1}{N^{k-1}}$. As R is permutation invariant, this implies that:

$$\Pr_H[\mathcal{A}^{(H)}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_1) = \dots = H(x_k)] \leq 2^{2k} \binom{q+k}{k}^2 \cdot \frac{1}{N^{k-1}}$$

Using first the inequality $\binom{q+k}{k} \leq \frac{(q+k)^k}{k!}$ and then the Stirling approximation:

$$\begin{aligned} 2^{2k} \binom{q+k}{k}^2 \cdot \frac{1}{N^{k-1}} &\leq N \cdot \left(\frac{4}{N}\right)^k \cdot \left[\frac{(q+k)^k}{k!}\right]^2 \\ &\leq N \left(\frac{4}{N}\right)^k \cdot (q+k)^{2k} \cdot \left(\frac{e}{k}\right)^{2k} = \frac{1}{N^{k-1}} \left[\frac{2e(q+k)}{k}\right]^{2k} \end{aligned}$$

Finally, the security of salted multi-collision against non-uniform quantum adversaries equipped with S bits of advice follows by combining this quantum hardness bound of multi-collision with Lemma 9. \square

Finally, we consider another search application, namely the task of determining k different inputs that all map to 0 under the random oracle. One of the main motivations behind this problem is its relation to the notion of proof-of-work in the blockchain context [GKL15].

Lemma 12 (Quantum Hardness of Multi-Search). *For any q -quantum query algorithm \mathcal{A} whose task is to find different preimages of 0 of a random oracle H , the success probability of \mathcal{A} is upper bounded by:*

$$\Pr_H[\mathcal{A}^{(H)}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = 0 \forall i \in [k]] \leq \left[\frac{4e^2(q+k)^2}{Nk^2}\right]^k$$

Note that this bound is asymptotically tight, as an algorithm with q queries can use q/k queries to find each pre-image (Grover's algorithm), resulting in a probability of $\Theta\left(\left\{\left(\frac{q}{k}\right)^2/N\right\}^k\right)$.

Proof. We will show this using our strong quantum lifting theorem for image relations (Theorem 9). Define $R := \{0, \dots, 0\}$ the relation over $[N]^k$, where $H : [M] \rightarrow [N]$. Then for each permutation π , we have $\Pr[(y_{\pi(1)}, \dots, y_{\pi(k)}) \in R \mid (y_1, \dots, y_k) \leftarrow [N]^k] = \frac{1}{N^k}$. As R is permutation invariant, this implies that:

$$\Pr_H[\mathcal{A}^{(H)}() \rightarrow \mathbf{x} = (x_1, \dots, x_k) : H(x_i) = 0 \forall i \in [k]] \leq 2^{2k} \cdot \binom{q+k}{k}^2 \frac{1}{N^k}$$

Using first the inequality $\binom{q+k}{k} \leq \frac{(q+k)^k}{k!}$ and then the Stirling approximation:

$$\begin{aligned} 2^{2k} \binom{q+k}{k}^2 \cdot \frac{1}{N^k} &\leq \left(\frac{4}{N}\right)^k \cdot \left[\frac{(q+k)^k}{k!}\right]^2 \\ &\leq \left(\frac{4}{N}\right)^k \cdot (q+k)^{2k} \cdot \left(\frac{e}{k}\right)^{2k} = \left[\frac{4e^2(q+k)^2}{Nk^2}\right]^k \end{aligned}$$

\square

Acknowledgements. J.G. was partially supported by NSF SaTC grants no. 2001082 and 2055694. F.S. was partially supported by NSF grant no. 1942706 (CAREER). J.G. and F.S. were also partially supported by Sony by means of the Sony Research Award Program. A.C. acknowledges support from the National Science Foundation grant CCF-1813814, from the AFOSR under Award Number FA9550-20-1-0108 and the support of the Quantum Advantage Pathfinder (QAP) project, with grant reference EP/X026167/1 and the UK Engineering and Physical Sciences Research Council.

References

- [ABKK23] Amit Agarwal, James Bartusek, Dakshita Khurana, and Nishant Kumar. A new framework for quantum oblivious transfer. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 363–394, Cham, 2023. Springer Nature Switzerland.
- [AFK22] Thomas Attema, Serge Fehr, and Michael Klooß. Fiat-shamir transformation of multi-round interactive proofs. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 113–142, Cham, 2022. Springer Nature Switzerland.
- [AFKR23] Thomas Attema, Serge Fehr, Michael Klooß, and Nicolas Resch. The fiat-shamir transformation of $(\gamma_1, \dots, \gamma_\mu)$ -special-sound interactive proofs. *Cryptology ePrint Archive*, Paper 2023/1945, 2023. <https://eprint.iacr.org/2023/1945>.
- [BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997.
- [BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69. Springer, 2011.
- [BKS21] Nir Bitansky, Michael Kellner, and Omri Shmueli. Post-quantum resettably-sound zero knowledge. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography*, pages 62–89, Cham, 2021. Springer International Publishing.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [CGK+23] Alexandru Cojocaru, Juan Garay, Aggelos Kiayias, Fang Song, and Petros Wallden. Quantum Multi-Solution Bernoulli Search with Applications to Bitcoin’s Post-Quantum Security. *Quantum*, 7:944, March 2023.
- [CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 673–684. IEEE, 2020.

- [Cha19] André Chailloux. Tight quantum security of the fiat-shamir transform for commit-and-open identification schemes with applications to post-quantum signature schemes. Cryptology ePrint Archive, Paper 2019/699, 2019. <https://eprint.iacr.org/2019/699>.
- [DFHS23] Jelle Don, Serge Fehr, Yu-Hsuan Huang, and Patrick Struck. On the (in)security of the buff transform. Cryptology ePrint Archive, Paper 2023/1634, 2023. <https://eprint.iacr.org/2023/1634>.
- [DFM20] Jelle Don, Serge Fehr, and Christian Majenz. *The Measure-and-Reprogram Technique 2.0: Multi-round Fiat-Shamir and More*, page 602-631. Springer International Publishing, 2020.
- [DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 356–383, Cham, 2019. Springer International Publishing.
- [DFMS22] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 677–706, Cham, 2022. Springer International Publishing.
- [DLW24] Fangqi Dong, Qipeng Liu, and Kewen Wu. Tight characterizations for pre-processing against cryptographic salting. In *Annual International Cryptology Conference*. Springer, 2024.
- [GHHM21] Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the qrom. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 637–667, Cham, 2021. Springer International Publishing.
- [GKL15] Juan A Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. *Journal of the ACM*, 2015.
- [GOP+23] Chaya Ganesh, Claudio Orlandi, Mahak Panholi, Akira Takahashi, and Daniel Tschudi. Fiat-shamir bulletproofs are non-malleable (in the random oracle model). Cryptology ePrint Archive, Paper 2023/147, 2023. <https://eprint.iacr.org/2023/147>.
- [JMZ23] Haodong Jiang, Zhi Ma, and Zhenfeng Zhang. Post-quantum security of key encapsulation mechanism against cca attacks with a single decapsulation query. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023*, pages 434–468, Singapore, 2023. Springer Nature Singapore.
- [Kat21] Shuichi Katsumata. A new simple technique to bootstrap various lattice zero-knowledge proofs to qrom secure nizks. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021*, pages 580–610, Cham, 2021. Springer International Publishing.
- [KX24] Haruhisa Kosuge and Keita Xagawa. Probabilistic hash-and-sign with retry in the quantum random oracle model. In Qiang Tang and Vanessa Teague, editors, *Public-Key Cryptography – PKC 2024*, pages 259–288, Cham, 2024. Springer Nature Switzerland.
- [LR13] Troy Lee and Jérémie Roland. A strong direct product theorem for quantum query complexity. *computational complexity*, 22:429–462, 2013.
- [LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 326–355, Cham, 2019. Springer International Publishing.

- [She11] Alexander A Sherstov. Strong direct product theorems for quantum communication and query complexity. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 41–50, 2011.
- [YZ21] Takashi Yamakawa and Mark Zhandry. Classical vs quantum random oracles. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, pages 568–597, Cham, 2021. Springer International Publishing.



Adaptive Hardcore Bit and Quantum Key Leasing over Classical Channel from LWE with Polynomial Modulus

Duong Hieu Phan^{1(✉)}, Weiqiang Wen^{1(✉)}, Xingyu Yan^{2(✉)},
and Jinwei Zheng^{1(✉)}

¹ LTCI, Telecom Paris, Institut Polytechnique de Paris, Paris, France
{hieu.phan, weiqiang.wen, jinwei.zheng}@telecom-paris.fr

² State Key Laboratory of Networking and Switching Technology, Beijing University
of Posts and Telecommunications, Beijing 100876, China
yanxy2020@bupt.edu.cn

Abstract. Quantum key leasing, also known as public key encryption with secure key leasing (PKE-SKL), allows a user to lease a (quantum) secret key to a server for decryption purpose, with the capability of revoking the key afterwards. In the pioneering work by Chardouvelis et al. (arXiv:2310.14328), a PKE-SKL scheme utilizing classical channels was successfully built upon the noisy trapdoor claw-free (NTCF) family. This approach, however, relies on the superpolynomial hardness of learning with errors (LWE) problem, which could affect both efficiency and security of the scheme.

In our work, we demonstrate that the reliance on superpolynomial hardness is unnecessary, and that LWE with polynomial-size modulus is sufficient to achieve the same goal. Our approach enhances both efficiency and security, thereby improving the practical feasibility of the scheme on near-term quantum devices. To accomplish this, we first construct a *noticeable* NTCF (NNTCF) family with the adaptive hardcore bit property, based on LWE with polynomial-size modulus. To the best of our knowledge, this is the first demonstration of the adaptive hardcore bit property based on LWE with polynomial-size modulus, which may be of independent interest. Building on this foundation, we address additional challenges in prior work to construct the first PKE-SKL scheme satisfying the following properties: (i) the entire protocol utilizes only classical communication, and can also be lifted to support homomorphism. (ii) the security is solely based on LWE assumption with polynomial-size modulus.

As a demonstration of the versatility of our noticeable NTCF, we show that an efficient proof of quantumness protocol can be built upon it. Specifically, our protocol enables a classical verifier to test the quantumness while relying exclusively on the LWE assumption with polynomial-size modulus.

Keywords: Trapdoor claw-free functions · Adaptive hardcore bit · Secure key leasing · Proofs of quantumness · Learning with errors

1 Introduction

In this article, we mainly focus on a fundamental primitive with a key-revocation capability – public key encryption with secure key leasing (PKE-SKL). Specifically, PKE-SKL refers to the realization of key-revocable PKE functionality, allowing the user/lessor to delegate decryption capability to the server/lessee in the form of a quantum decryption key, whereby once the key is revoked, the lessee loses the ability to decrypt. The PKE-SKL scheme is particularly effective in interactive cryptographic settings involving classical users and quantum servers.

Recently, inspired by secure software leasing in [ALP21], the notion of PKE-SKL was concurrently introduced by Agrawal et al. in [AKN+23] and Ananth et al. in [APV23]. Based on the PKE-SKL scheme, these works subsequently investigated the notion of secure key leasing for several extensions, like identity-based encryption (IBE), attribute-based encryption (ABE), functional encryption (FE), fully homomorphic encryption (FHE), and pseudorandom functions (PF). These key-revocable schemes based on the quantum no-cloning principle enable delegation and revocation of privileges, which is crucial in many cryptographic applications. Unfortunately, both recent works in [AKN+23, APV23] for constructing PKE-SKL have two shortcomings:

- The user and the server must have both quantum capabilities, and the key generation process has to require quantum communication;
- The construction requires *subexponential* hardness of the LWE assumption with superpolynomial modulus.

To address the former issue, Chardouvelis et al. [CGJL23] recently introduced a semi-quantum PKE-SKL scheme, transforming their approach into a scheme with merely classical communication between a classical client and a quantum server. Their work is inspired by the work of classical verification of quantumness from LWE in [BCM+18]. Their construction is mainly based on a powerful cryptographic tool called the LWE-based noisy trapdoor claw-free functions (NTCF) with an adaptive hardcore bit (AHB) property.

However, the PKE-SKL scheme by Chardouvelis et al. [CGJL23] does not address the second issue. Their construction still requires the subexponential hardness of LWE with a superpolynomial modulus. One of the main reasons for this is that their construction relies on NTCF with the AHB property, which in turn depends on the superpolynomial hardness of LWE assumption. This significantly affects the security and the efficiency of PKE-SKL, even making it unfriendly for implementation on near-term quantum devices. Thus, building on these, our main open question is the following:

Can efficient PKE-SKL with completely classical communication be based on the polynomial hardness of standard LWE over polynomially large modulus ?

1.1 Our Results

In this work, we affirmatively solve the above question. Main contributions are summarized in Fig. 1.

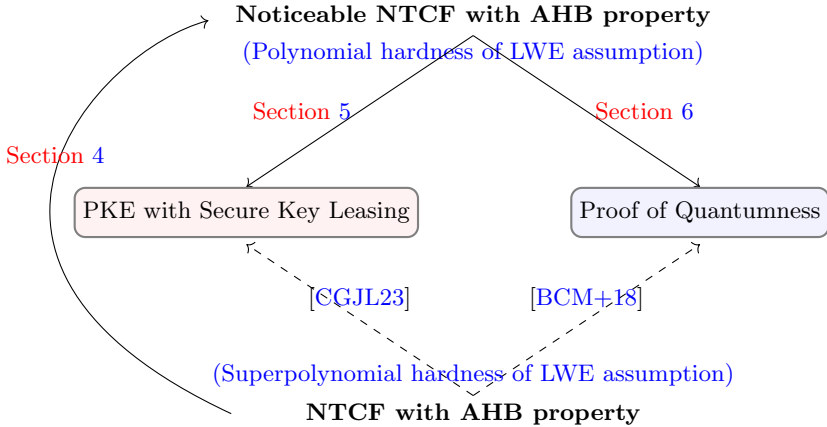


Fig. 1. Outline of main contributions in our work. To achieve a PKE-SKL scheme with a polynomially large modulus, we first improve the NTCF from [BCM+18] and propose a cryptographic primitive called noticeable NTCF (NNTCF). This primitive serves as the core tool for constructing PKE-SKL and can be constructed based on the polynomial hardness of LWE while still retaining the AHB property. We believe that NNTCF may have independent interests. In addition to constructing PKE-SKL schemes using NNTCF with AHB, as an example, we demonstrate a NNTCF-based proof of quantumness protocol to illustrate its versatility.

We show that a modified version of the PKE-SKL scheme [CGJL23] with merely classical communication can be constructed based on the hardness of LWE with polynomial modulus. Informally, we first obtain the following result.

Theorem 1 (Informal). *There exists a secure key leasing scheme for public key encryption with a completely classical lessor, assuming the hardness of LWE with polynomial modulus.*

Specifically, based on LWE with polynomial modulus, we can achieve PKE-SKL introduced in [CGJL23] with the following properties:

1. The protocol only uses polynomial-sized modulus q . This improves both efficiency and security.
2. The protocol executed between a classical lessor and a quantum lessee involves only classical communication, and all deletion certificates are classical.
3. The protocol satisfies a stronger PKE-SKL security described in [CGJL23]. We show that any quantum polynomial-time adversary can only simultaneously provide a valid classical deletion certificate and distinguish ciphertexts with at most negligible probability.

To achieve this target, we realize the *adaptive hardcore bit property from the hardness of LWE with polynomial modulus*, which reduces the modulus from superpolynomial size in [BCM+18] to polynomial size. Besides this, we introduce an important primitive named the noticeable NTCF (NNTCF) family with this property.

Theorem 2 (Informal). *Assuming the hardness of the LWE problem with polynomial modulus, there exists a noticeable NTCF (NNTCF) family with the amplified adaptive hardcore bit property.*

To the best of our knowledge, prior to our work, the NTCF family with adaptive hardcore bit property can only be constructed based on superpolynomial modulus. We believe this noticeable version of NTCF using a smaller modulus may be of independent interest, such as enhancing the security¹ and improving the implementation efficiency of NTCF-based quantum cryptographic protocols: revocable quantum digital signatures [MPY23], proofs of quantumness [BCM+18, BKVV20], quantum delegated computation [Mah18b], certifiable randomness generation [BCM+18] etc. To illustrate this, we present a new proof of quantumness protocol based on NNTCF as an example.

Theorem 3 (Informal). *Assuming the polynomial hardness of the LWE with polynomial modulus, there exists a polynomial-sized proof of quantumness protocol from the NNTCF family.*

Specifically, our NNTCF-based proof of quantumness protocol circumvents the need for a superpolynomial modulus as required in [BCM+18], and fully satisfies both quantum completeness and classical soundness. Namely, the protocol ensures that a quantum polynomial-time prover can succeed with high probability (quantum completeness), while no classical polynomial-time prover can achieve comparable success probability (classical soundness). The soundness relies on the adaptive hardcore bit property of the NNTCF.

1.2 Related Works

Noisy Trapdoor Claw-Free Functions. The concept of noisy trapdoor claw-free functions (NTCF) was first introduced by Brakerski et al. in the proofs of quantumness and certifiable quantum randomness generator [BCM+18], and was further developed by Mahadev within the realms of delegated quantum computing [Mah18b] and quantum homomorphic encryption [Mah18a]. Conceptually, trapdoor claw-free functions (TCFs) consist of a pair of injective functions f_0 and f_1 that share the same image. With access to a secret trapdoor td , it becomes easy to determine the two preimages \mathbf{x}_0 and \mathbf{x}_1 of the same image \mathbf{y} , such that $f_0(\mathbf{x}_0) = f_1(\mathbf{x}_1) = \mathbf{y}$. However, it is computationally difficult to invert f_0, f_1 without the trapdoor td . Such a pair of $(\mathbf{x}_0, \mathbf{x}_1)$ is known as a claw, hence the

¹ Improving the security from the subexponential hardness of LWE assumption to polynomial hardness of LWE assumption.

name is claw-free. This useful cryptographic tool constructed based on the LWE assumption plays a crucial role in quantum-classical interactive proof systems, especially in constraining, describing, and verifying the behavior of untrusted quantum devices. Inspired by these works, LWE-based NTCFs have been applied to many intriguing quantum cryptographic schemes, such as remote state preparation [GV19, GMP23], tests of quantumness [BKVV20, BGKM+23], quantum money [RS19, Shm22], secure quantum extraction [ALP20], public-key deniable encryption [CGV22], quantum copy-protection [CHV23], quantum certified deletion [HMNY21], secure key leasing [AKN+23, APV23, CHV23, MPY23], and secure software leasing [KNY21], etc.

More importantly, the security of LWE-based NTCF requires a very important property – the adaptive hardcore bit (AHB) property, which is widely used in constructing the above cryptographic schemes. The AHB property states that whenever $f_0(\mathbf{x}_0) = f_1(\mathbf{x}_1)$, it is difficult to hold both single preimage (b, \mathbf{x}_b) , as well as a random \mathbf{d} and a bit c such that $c = \mathbf{d}^\top \cdot (\mathbf{x}_0 \oplus \mathbf{x}_1) \bmod 2$. So far, The LWE-based NTCF in [BCM+18, Mah18b] is the only known TCF instance with AHB property, but its security is based on LWE with superpolynomial modulus. In this work, we will consider a noticeable version of NTCF with AHB property that only requires a polynomially large modulus.

Secure Key Leasing/Revocable Cryptography. The notion of secure key leasing (SKL) or key-revocable cryptography is inspired by secure software leasing in [ALP21]. Secure key leasing can be viewed as secure software leasing for decryption algorithms but with stronger security guarantees that the adversary is not restricted from running the software honestly after it is returned. Similar to quantum copy protection schemes, the core idea of SKL is to encode the secret key into a quantum state to prevent it from being copied based on the no-cloning principle. Recently, a couple of works have built PKE-SKL (or called key-revocable PKE) and DSIG-SKL from lattices.

In [AKN+23], Agrawal et al. proposed the notion of public key encryption with secure key leasing. In [APV23], Ananth et al. concurrently introduced the same concept of key-revocable public key encryption. In these two works, key-revocable PKE schemes are constructed based on standard LWE assumption [APV23] or even the mere existence of any PKE scheme [AKN+23]. Independently, for the digital signature primitive, Morimae et al. [MPY23] studied the notions of digital signature with revocable signing keys and digital signature with revocable signatures, assuming the sub-exponential hardness of LWE.

However, the above PKE-SKL works require both the user and the server to possess quantum capabilities and utilize quantum communication. Thus, an interesting question is whether it is possible to transform their schemes into one with classical user and classical communication. To solve this problem and further reduce quantum resources, Chardouvelis et al. [CGJL23] introduced a semi-quantum PKE-SKL scheme in which the user is classical and interacts solely through a classical communication with the quantum server. However, as the construction of this scheme heavily relies on the trapdoor claw-free functions

with AHB property introduced in [BCM+18], the security of the scheme still depends on the sub-exponential time hardness of LWE assumption, necessitating a sub-exponentially large modulus.

To date, all previous works that imply PKE-SKL are designed to achieve quantum/semi-quantum key-revocable cryptography and almost rely on the sub-exponential hardness of LWE. In this paper, inspired by the work of Chardouvelis et al. [CGJL23], our goal is to achieve a PKE-SKL scheme that requires only minimal quantum capabilities (only with classical communication), more desirably from the polynomial hardness of LWE assumption.

NTCF-Based Proofs of Quantumness. A cryptographic proof of quantumness is an interactive protocol that enables classical verifiers to determine whether provers (potentially quantum) is non-classical. To achieve this, [BCM+18] introduced the first groundbreaking proof of quantumness system. This scheme is constructed based on LWE-based NTCF, and its soundness is guaranteed by the adaptive hardcore bit (AHB) property of NTCF. However, a major drawback of this scheme is that the AHB property must rely on the sub-exponential hardness of LWE, requiring the modulus of the scheme to be super-polynomially large. Since then, many methods have been proposed to further simplify NTCF-based proof systems by circumventing the AHB property. For example, [BKVV20] introduced a simple proof of quantumness scheme based on a random oracle model, assuming only the existence of trapdoor claw-free functions. [YZ22] demonstrated a non-interactive quantumness test in the random oracle model. Furthermore, other schemes [KMCVY22, KLVY23, BGKM+23] incorporate NTCF with Bell’s inequality to get rid of dependence on AHB property.

In this work, to avoid relying on the random oracle model or Bell’s inequality, we aim to achieve the AHB property solely based on the polynomial hardness of LWE. This approach will fundamentally and directly enhance the efficiency of the protocol described in [BCM+18]. To the best of our knowledge, no prior research has accomplished this.

1.3 Organization

The remainder of the paper is organized as follows. In Sect. 2 we give the technical overview for our main results. In Sect. 3 we provide cryptographic preliminaries used throughout this work. In Sect. 4 we formalize our definition of noticeable noisy trapdoor claw-free family (NNTCF) and show that its construction can be built from standard LWE assumption with polynomial modulus. Furthermore, we prove our NNTCF still satisfies the adaptive hardcore bit (AHB) property.

In Sect. 5, we describe the construction of the NNTCF-based PKE-SKL scheme, assuming the polynomial hardness of LWE with polynomial modulus.

In Sect. 6, we describe the construction of the NNTCF-based proof of quantumness scheme, assuming the polynomial hardness of LWE.

2 Technical Overview

In this section, we will provide a technical overview of our works described in Fig. 1. We first slightly extend the original NTCF from [BCM+18] to define our noticeable NTCF (NNTCF) in Subsect. 2.1. Notably, our NNTCF family with adaptive hardcore bit property can be built upon LWE with polynomial modulus. In Subsect. 2.2, we will explain how the NNTCF primitive can be used to optimize the PKE-SKL scheme in [CGJL23] such that its ciphertexts' modulus q can be reduced to polynomial size and the security is based on the LWE with polynomial modulus. Finally, in Subsect. 2.3, we will present the main idea of constructing a new proof of quantumness protocol based on our NNTCF. In particular, our scheme is solely based on LWE with polynomial modulus and does not need to rely on random oracle model [BKVV20] or Bell's inequality [KMCVY22, KLVY23, BGKM+23].

Throughout this section, we will try to be consistent with prior works about the notation of parameters for easier comparison and comprehension.

2.1 Noticeable NTCF from Polynomial Hardness of LWE

Before explaining this approach, we need to recall how the LWE can be employed to construct a NTCF in [BCM+18] and why this original LWE-based NTCF requires a superpolynomial-sized modulus.

Recap: LWE-Based NTCF and Its Two Superpolynomial Gaps. Given function $f_{k,b}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{e} + b \cdot \mathbf{A}\mathbf{s}$ defined with standard LWE samples $k = (\mathbf{A}, \mathbf{t} = \mathbf{A}\mathbf{s} + \mathbf{e}_0) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, a NTCF can be informally defined by $f'_{k,b}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{e} + b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e}_0)$ for $b \in \{0, 1\}$. We can see that if \mathbf{e}_0 were 0, $f'_{k,b}(\mathbf{x})$ is the same as $f_{k,b}(\mathbf{x})$, such that $f_{k,1}(\mathbf{x}) = f_{k,0}(\mathbf{x} + \mathbf{s})$. But in fact, \mathbf{e}_0 really won't be 0. In this case, to ensure that $f'_{k,1}(\mathbf{x})$ and $f'_{k,0}(\mathbf{x} + \mathbf{s})$ still appear to be the same, we must strictly constrain the norm of \mathbf{e} . Typically, we can sample \mathbf{e} from a Gaussian distribution with width superpolynomially larger than the Gaussian distributed noise \mathbf{e}_0 , implying that $f'_{k,1}(\mathbf{x})$ is statistically close to $f'_{k,0}(\mathbf{x} + \mathbf{s})$.

Specifically, if $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}_q^m, B_V}$, $\mathbf{e} \leftarrow D_{\mathbb{Z}_q^m, B_P}$ and B_P/B_V is superpolynomial in security parameter λ , the Hellinger statistical distance between $f'_{k,b}(\mathbf{x})$ and $f_{k,b}(\mathbf{x})$, $1 - \exp(-2\pi m B_V/B_P)$, can be bounded by $1 - \text{negl}(\lambda)$. Therefore, \mathbf{e} can be viewed as a flooding noise for \mathbf{e}_0 , which incurs the first superpolynomial gap B_P/B_V .

Next, we explain the second superpolynomial gap B_V/B_L . This gap is incurred by noise flooding used to ensure the adaptive hardcore bit (AHB) property of NTCF, which is briefly introduced below. Given a description of a NTCF described as above, a quantum device can easily set up a claw superposition as $\frac{1}{\sqrt{2}}(|0, \mathbf{x}_0\rangle + |1, \mathbf{x}_1\rangle)$ by creating the state $\sum_{b, \mathbf{x}} |b\rangle |\mathbf{x}\rangle |f'_{k,b}(\mathbf{x})\rangle$ and measuring the last register, where $f'_{k,0}(\mathbf{x}_0) = f'_{k,1}(\mathbf{x}_1)$ and $\mathbf{x}_1 = \mathbf{x}_0 - \mathbf{s} \bmod q$. For the generated state $\frac{1}{\sqrt{2}}(|0, \mathbf{x}_0\rangle + |1, \mathbf{x}_1\rangle)$, performing a computational basis measurement will yield a preimage $(b, \mathbf{x}_b) \in \{0, 1\} \times \mathbb{Z}_q^n$. On the other hand, performing a

Hadamard basis measurement will yield a pair $(c, \mathbf{d}) \in \{0, 1\} \times \{0, 1\}^{n \log q}$ such that \mathbf{d} is uniform random and $c = \mathbf{d}^\top \cdot (\mathbf{x}_0 \oplus \mathbf{x}_1) \bmod 2^2$.

The AHB property asserts that it is not possible to simultaneously obtain both (b, \mathbf{x}_b) and (c, \mathbf{d}) under the LWE assumption. From the Lemma 2, Brakerski et al. have proven that if $(b, \mathbf{x}_b, c, \mathbf{d})$ are given, there exists an efficiently computable function $I_{b, \mathbf{x}_b}(\mathbf{d})$ for random \mathbf{d} can compute a string $\hat{\mathbf{d}}$ such that $\mathbf{d}^\top \cdot (\mathbf{x}_0 \oplus \mathbf{x}_1) = \hat{\mathbf{d}}^\top \cdot \mathbf{s}$, where $\hat{\mathbf{d}} \in \{0, 1\}^n \setminus \{0^n\}$ and $\mathbf{s} \in \{0, 1\}^n$. Thus, the AHB property can be reformulated as stating that it is hard to produce a pair $(c, \hat{\mathbf{d}})$ such that $c = \hat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2$. In other words, the AHB property holds if the distribution $\hat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2$ is statistically close to a uniformly random bit, where $\hat{\mathbf{d}}$ is conditioned on LWE sample.

To prove this, [BCM+18] used the leakage resilience of LWE: Given an LWE instance, any given bit of \mathbf{s} is computationally indistinguishable from a uniformly random bit. This approach replaces the matrix \mathbf{A} with a computationally indistinguishable lossy matrix $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, D_{\mathbb{Z}_q, L})$, where $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ has a large kernel and $\mathbf{F} \leftarrow D_{\mathbb{Z}_q^{m \times n}, B_L}$ is small. Now, the LWE instance $(\mathbf{A}, \mathbf{As} + \mathbf{e}_0)$ is replaced by $(\mathbf{BC} + \mathbf{F}, \mathbf{BCs} + \mathbf{Fs} + \mathbf{e}_0)$. As we know, the choice of $\hat{\mathbf{d}}$ indeed depends upon the LWE sample, which corresponds in the leakage resilience argument to $\hat{\mathbf{d}}$ depending on \mathbf{Cs} . Thus, the core proof of AHB property is to argue that given a sample of the form $(\mathbf{BC} + \mathbf{F}, \mathbf{BCs} + \mathbf{Fs} + \mathbf{e}_0)$, for any fixed $\hat{\mathbf{d}}$, the distribution $\hat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2$ is still statistically close to uniform distribution with overwhelming probability. In other words, we need to show for any fixed $\hat{\mathbf{d}}$ and \mathbf{C} , the joint distribution $(\mathbf{Cs}, \hat{\mathbf{d}} \cdot \mathbf{s} \bmod 2)$ is statistically close to uniform.

To achieve this, their solution relies on \mathbf{s} being a computationally random binary vector, but now the \mathbf{s} is subject to \mathbf{Fs} information leakage. To solve this, they choose \mathbf{e}_0 from a Gaussian distribution with a width sufficiently larger than Gaussian distributed noise \mathbf{F} (i.e., B_V/B_L also be superpolynomial). Since $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, B_V}$ and $\|\mathbf{Fs}\| \leq nB_L\sqrt{m}$, this ensures that \mathbf{e}_0 statistically “floods” the term \mathbf{Fs} . Then, this noise flooding technology could efficiently ensure that $(\mathbf{Cs}, \hat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2)$ is statistically close to uniform.

Noticeable NTCF from Polynomial LWE Assumption. To circumvent the above two superpolynomial flooding noises, we develop a family of noticeable NTCF (NNTCF) endowed with the AHB property from the hardness of standard LWE. The formal definition and construction are described in Sect. 4. Below we elaborate on the high-level idea of reducing B_P/B_V and B_V/B_L to polynomial size, respectively.

- *Circumvent superpolynomial B_P/B_V :* We introduce the concept of a noticeable version of NTCF (NNTCF). Intuitively, “noticeable” here means that we can slightly relax the statistical distance between the two distributions, $f'_{k,1}(\mathbf{x})$ and $f_{k,1}(\mathbf{x})$, in the NTCF. Specifically, we relax the Hellinger

² In fact, the bit c is evaluated by $c = \mathbf{d}^\top \cdot (\mathcal{J}(\mathbf{x}_0) \oplus \mathcal{J}(\mathbf{x}_1))$ in [BCM+18], where $\mathcal{J}(\cdot)$ is the binary representation function. For simplicity, we omit this function in the expression.

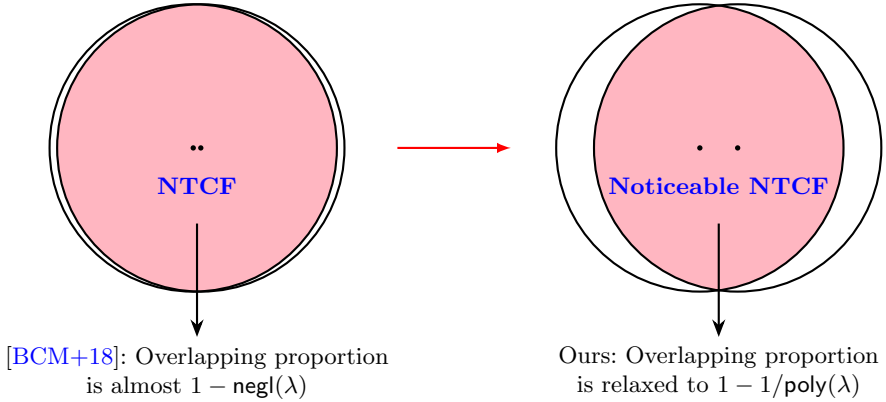


Fig. 2. Schematic representation of noticeable NTCF from original NTCF. The term “noticeable NTCF” emphasizes that the distance we consider is not negligible but noticeable (inverse polynomially small).

distance between $f'_{k,1}(\mathbf{x})$ and $f_{k,1}(\mathbf{x})$ from $\text{negl}(\lambda)$ to $1/\text{poly}(\lambda)$, as shown in Fig. 2. This relaxation allows us to naturally reduce B_P/B_V to a polynomial size. In fact, this relaxation has already been implicitly used in [BKVV20] to simplify the proof of quantumness. Here, we decide to explicitly define this concept to emphasize that the statistical distance between the two aforementioned distributions is not necessarily negligible.

- *Circumvent superpolynomial B_V/B_L :* We illustrate the high-level idea in Fig. 3. To ensure that distribution $(\mathbf{C}\mathbf{s}, \widehat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2)$ is statistically close to the uniform distribution $U(\mathbb{Z}_q^l \times \mathbb{Z}_2)$, [BCM+18] uses the superpolynomial flooding noise \mathbf{e}_0 hides the term $\mathbf{F}\mathbf{s}$. This method is very straightforward, however, we observe that it is not necessary to completely hide the $\mathbf{F}\mathbf{s}$ information, but only to obscure the \mathbf{s} information well. Intuitively, there is no need to hide \mathbf{s} perfectly. We argue that if there is sufficiently high entropy left in \mathbf{s} , then the argument $(\mathbf{C}\mathbf{s}, \widehat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2) \approx_s U(\mathbb{Z}_q^l \times \mathbb{Z}_2)$ still holds. Specifically, we use the refined flooding technique, also known as the gentle flooding approach in [BD20]. The main solution is to apply refined noise flooding to replace the error \mathbf{e}_0 with term $\mathbf{F}\mathbf{e}_0^{(1)} + \mathbf{e}_0^{(2)}$. Refer to [BD20], we set $\mathbf{e}_0^{(1)}$ and $\mathbf{e}_0^{(2)}$ as independent random variables with polynomially large width. Consequently, the term $\mathbf{F}\mathbf{s} + \mathbf{e}_0$ is reformulated as $\mathbf{F}(\mathbf{s} + \mathbf{e}_0^{(1)}) + \mathbf{e}_0^{(2)}$. Building on this, we prove that the AHB property still holds in the NNTCF family. The heart of the proof lies in the fact that we can directly argue that the distribution $(\mathbf{C}\mathbf{s}, \widehat{\mathbf{d}}^\top \cdot \mathbf{s} \bmod 2)$, conditioned on $\mathbf{v} = \mathbf{s} + \mathbf{e}_0^{(1)}$ for any fixed \mathbf{v} , is statistically indistinguishable from the uniform distribution $U(\mathbb{Z}_q^l \times \mathbb{Z}_2)$ (See Lemma 6).

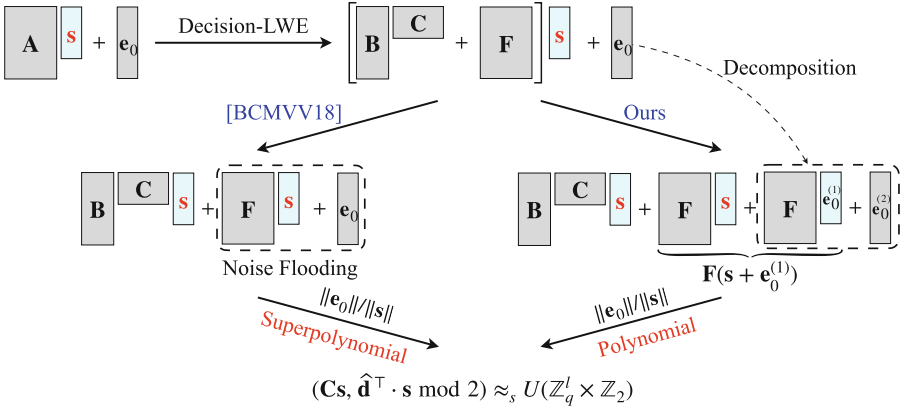


Fig. 3. Summary of circumventing superpolynomial flooding noise in the proof of AHB property of NTCF. $U(\mathbb{Z}_q^l \times \mathbb{Z}_2)$ denotes the density of the uniform distribution over $\mathbb{Z}_q^l \times \mathbb{Z}_2$.

Therefore, both B_P for \mathbf{e} and B_V for \mathbf{e}_0 can be polynomially large, thereby ensuring the LWE-based NNTCF only relies on polynomial hardness of LWE assumption.

2.2 Secret Key Leasing for PKE from LWE-Based NNTCF

Building on the NNTCF family with AHB property, we show how to construct a PKE-SKL scheme with a polynomial modulus. We first review the PKE-SKL scheme described in [CGJL23] from the LWE-based NTCF with AHB property.

Recap: NTCF-based PKE-SKL in [CGJL23]. Their construction is inspired by the “proof of quantumness” construction in [BCM+18]. To obtain a key leasing scheme, the idea is to use the claw superposition in their construction as a quantum decryption key. We describe (a slightly simplified version of) Chardouvelis’s PKE-SKL scheme based on Regev’s two-key PKE and LWE-based NTCF, as illustrated in Fig. 4.

The formal scheme is a parallel repetition of the above scheme. Recall the AHB property of NTCF that no quantum polynomial-time adversary can obtain both (b, \mathbf{x}_b) and (c, \mathbf{d}) . This property will guarantee the security of the PKE-SKL scheme of Fig. 4, i.e., any adversary cannot both provide a valid classical deletion certificate and distinguish ciphertexts (the latter corresponds to the ability to extract \mathbf{x}_b).

Now, we explain why the PKE-SKL construction in Fig. 4 requires a superpolynomial modulus. The primary reason is the presence of four superpolynomial gaps: all ratios B_V/B_L , B_P/B_V , $B_{P'}/B_P$, B_X/B_S need to be superpolynomial in λ . The first two superpolynomial gaps B_V/B_L , B_P/B_V are caused by LWE-based NTCF with AHB property, which has been explained in Sect. 2.1.

- Setup(1^λ) \rightarrow (mpk, sk):
 - Generate a NTCF pair $k = (\mathbf{A}, \mathbf{t} = \mathbf{A}\mathbf{s} + \mathbf{e}_0)$ where $\mathbf{s} \xleftarrow{\$} [B_s]^n$ and $\mathbf{e}_0 \xleftarrow{D_{\mathbb{Z}_q^m, B_V}}$, along with a trapdoor \mathbf{td} , send just the k .
 - Output the master public-key and the trapdoor as (mpk, sk) = (k , \mathbf{td}).
- KeyGen(mpk) \rightarrow (ρ_{sk} , pk):
 - Create a state $|\psi\rangle = \sum_{b, \mathbf{x}} |b\rangle |\mathbf{x}\rangle |f'_{k,b}(\mathbf{x})\rangle$. Measure the last register.
 - Obtain an image \mathbf{y} , where $\mathbf{y} = \mathbf{A}\mathbf{x}_b + \mathbf{e} + b\mathbf{t}$, $\mathbf{e} \xleftarrow{D_{\mathbb{Z}_q^m, B_P}}$, $\mathbf{x}_b \in [B_X]^n$. Generate a claw state: $|\phi\rangle = \frac{1}{\sqrt{2}} \sum_{b \in \{0,1\}} |b, \mathbf{x}_b\rangle = \frac{1}{\sqrt{2}} (|0\rangle |\mathbf{x}_0\rangle + |1\rangle |\mathbf{x}_1\rangle)$.
 - Output public key $\text{pk} = \{k, \mathbf{y}\}$ and quantum decryption key $\rho_{\text{sk}} = |\phi\rangle$.
- Enc(pk, μ) \rightarrow ct: For a message $\mu \in \{0, 1\}$:
 - Sample a binary random vector $\mathbf{r} \in \mathbb{Z}_q^{m \times 1}$ and computes $\text{ct}_1 = \mathbf{r}^\top \mathbf{A}$, $\text{ct}_2 = \mathbf{r}^\top \mathbf{t}$ and $\text{ct}_3 = \mathbf{r}^\top \mathbf{y} + \mathbf{e}' + \mu \cdot \lceil q/2 \rceil$, where $\mathbf{e}' \xleftarrow{D_{\mathbb{Z}_q^m, B_{P'}}}$.
 - Let $\text{ct} := (\text{ct}_1, \text{ct}_2, \text{ct}_3)$ and output ciphertext ct.
- Dec(ρ_{sk} , ct) \rightarrow (μ , ρ_{sk}):
 - Coherently compute $\text{ct}_3 - \text{ct}_1 \cdot \mathbf{x}_b - b \cdot \text{ct}_2$ on the ancilla register, obtain $|\phi\rangle \otimes |\text{ct}_3 - \text{ct}_1 \cdot \mathbf{x}_b - b \cdot \text{ct}_2\rangle \approx |\phi\rangle \otimes |\mu \cdot q/2\rangle$.
 - Measure the last register, return $\mu = 0$ if the outcome is less than $q/4$; Return $\mu = 1$ otherwise. The ρ_{sk} in the first register remains intact.
- Del(ρ_{sk}) \rightarrow cert:
 - Take in the decryption key $|\phi\rangle$, measure it in the Hadamard basis, resulting in cert = (c , \mathbf{d}) as the deletion certificate.
- VerDel(sk, pk, cert) \rightarrow \top/\perp :
 - Use \mathbf{td} to compute claw $(\mathbf{x}_0, \mathbf{x}_1)$. Check if $c = \mathbf{d}^\top \cdot (\mathbf{x}_0 \oplus \mathbf{x}_1) \bmod 2$ holds.
 - If this check passes, output \top ; Otherwise output \perp .

Fig. 4. Core subroutine of PKE-SKL in [CGJL23]

The $B_{P'}/B_P$ must be superpolynomial because the ct_3 in ciphertext needs a flooding noise $\mathbf{e}' \xleftarrow{D_{\mathbb{Z}_q^m, B_{P'}}}$ to flood the term $\mathbf{r}^\top \mathbf{e}$. In the PKE-SKL security game, the $\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{e}$ is given by the adversary who plays the role of the user, hence \mathbf{e} can be related to \mathbf{A} . Therefore we cannot apply a general leftover hash lemma directly to $\mathbf{r}^\top \mathbf{A}$ conditioned on $\mathbf{r}^\top \mathbf{e}$. To make $\mathbf{r}^\top \mathbf{A}$ independently random, they use smudging noise \mathbf{e}' with superpolynomially larger width $B_{P'} \gg B_P$ to flood the term $\mathbf{r}^\top \mathbf{e}$, thereby ensuring $\mathbf{r}^\top \mathbf{A}$ to be independently random under the entropy of \mathbf{r} . This is crucial for the (quantum) extractor to work given a (quantum) distinguisher for distinguishing encryptions of 0 and 1. Due to the AHB property, the successful construction of such an extractor will ensure that any lessee can not decrypt anymore after submitting a deletion certificate, which proves the PKE-SKL security.

Finally, regarding B_X , it is required to be either superpolynomially larger than B_S or equal to modulus q . This condition is to ensure that the two distributions $U([B_X])$ and $U([B_X] + B_S)$ are statistically close, which is crucial for ensuring correct generation of claw superposition as $\frac{1}{\sqrt{2}}(|0, \mathbf{x}_0\rangle + |1, \mathbf{x}_1\rangle)$. Later, when we resolve all three primary gaps B_V/B_L , B_P/B_V and $B_{P'}/B_P$, the mod-




Fig. 5. Summary of circumventing superpolynomial $B_{P'}/B_P$. $U(\mathbb{Z}_q^n)$ denotes the density of the uniform distribution over \mathbb{Z}_q^n .

ulus can then be reduced to be polynomial-size. As a result, we can set $B_X = q$ and the gap B_X/B_S also becomes polynomial.

Therefore, to achieve polynomial-size modulus in the PKE-SKL [CGJL23], it suffices to address the three gaps: B_V/B_L , B_P/B_V and $B_{P'}/B_P$.

Solving Superpolynomial Gaps in NTCF-Based PKE-SKL. Below, we provide high-level ideas for circumventing these gaps.

- *The gaps B_V/B_L and B_P/B_V :* Firstly, we can replace the NTCF with the NNTCF family described in Sect. 2.1, thereby immediately avoiding two superpolynomial gaps B_V/B_L , B_P/B_V . The security of PKE-SKL will then be based on the AHB property of our NNTCF.
- *The gap $B_{P'}/B_P$:* Intuitively, to make $\mathbf{r}^\top \mathbf{A}$ independently random, it is not necessary to completely hide the $\mathbf{r}^\top \mathbf{e}$ information, but only to obscure the \mathbf{r} well. Furthermore, there is no need to statistically hide \mathbf{r} , we only need to properly hide enough information in \mathbf{r} to ensure that $\mathbf{r}^\top \mathbf{A}$ appears independently random, as shown in Fig. 5.

In order to remove the superpolynomially large \mathbf{e}' , our key idea is to perturb \mathbf{e} with another vector $\mathbf{e}_1 \in [-\|\mathbf{e}\|_\infty, \|\mathbf{e}\|_\infty]^m$ before its product with \mathbf{r} . The hope is that many entries of $\mathbf{e} + \mathbf{e}_1$ indexed by some set \mathcal{Z} will become 0's after the random perturbation. As explained before, the \mathbf{e} can be related to \mathbf{A} . Therefore, it is crucial to argue that the set \mathcal{Z} is random and independent from \mathbf{e} , in which case the remaining entropy of \mathbf{r} given $\mathbf{r}^\top (\mathbf{e} + \mathbf{e}_1)$ is sufficient to make $\mathbf{r}^\top \mathbf{A}$ independently random.

Unfortunately, the length of \mathbf{e} is smaller than the infinity norm of \mathbf{e} , so we cannot expect a high probability that $\mathbf{e} + \mathbf{e}_1$ has many 0's. To address this, we need to modify the scheme such that the public key contains more samples. We apply a standard technique that can help increase the number of sam-

ples (i.e., $m' > m$), but will also slightly increase the error size. Concretely, we select $\mathbf{U} \stackrel{\$}{\leftarrow} \{0, 1\}^{m' \times m}$ and derive more samples $(\mathbf{A}', \mathbf{y}')$ with the same secret \mathbf{x}_0 , where $\mathbf{A}' = \mathbf{U}\mathbf{A}$ and $\mathbf{y}' = \mathbf{U}\mathbf{y} = \mathbf{U}(\mathbf{A}\mathbf{x}_0 + \mathbf{e}) = \mathbf{A}'\mathbf{x}_0 + \mathbf{U}\mathbf{e}$. We let $\mathbf{e}'' = \mathbf{U}\mathbf{e}$ denote the new error. Now we can guess \mathbf{e}'' with \mathbf{e}'_1 from $[-\|\mathbf{e}''\|_\infty, \|\mathbf{e}''\|_\infty]^{m'}$ of a larger length. For an appropriate choice of m' , we can ensure that $\mathbf{e}'' + \mathbf{e}'_1$ has sufficiently many 0's. Under this condition, sufficient randomness in \mathbf{r} will be preserved, which allows us to argue that $\mathbf{r}^\top \mathbf{A}'$ is independently random.

Next, we discuss why the decryption functionality and security of the PKE-SKL scheme can still be maintained when switching from NTCF to NNTCF.

Decryption Functionality. For the ciphertext in the PKE-SKL scheme, the correctness of decryption depends on the quantum decryption key ρ_{sk} . As we know, in the NTCF-based PKE-SKL scheme, the decryption key ρ_{sk} is the uniform claw superposition $\frac{1}{\sqrt{2}}(|0, \mathbf{x}_0\rangle + |1, \mathbf{x}_1\rangle)$ quantumly generated by the NTCF.

When we switch from NTCF to NNTCF, the generated claw state may exhibit slight variations. In NNTCF, we relax the Hellinger distance between distributions $f'_{k,b}(\mathbf{x})$ and $f_{k,b}(\mathbf{x})$ from $1 - \text{negl}(\lambda)$ to $1 - 1/\text{poly}(\lambda)$. In this case, once the quantum device measures the last register of state $\sum_{b,\mathbf{x}} |b\rangle |\mathbf{x}\rangle |f'_{k,b}(\mathbf{x})\rangle$, the first two registers will not always produce a uniform claw superposition.

However, we must point out that even if the generated claw state is not a perfect uniform superposition of $(\mathbf{x}_0, \mathbf{x}_1)$, the state can still successfully decrypt with a probability of $1 - \text{negl}(\lambda)$. This is because the decryption operation $\rho_{sk} \otimes |\text{ct}_3 - \text{ct}_1 \cdot \mathbf{x}_b + b \cdot \text{ct}_2\rangle$ is performed coherently. As long as the claw state generated by the NNTCF is still over the two preimages $(0, \mathbf{x}_0)$ and $(1, \mathbf{x}_1)$, regardless of whether their amplitudes differ from $1/\sqrt{2}$, decryption will be successful.

Key Leasing Security. The core of key leasing security is to ensure that the lessee cannot perform decryption after deleting the quantum decryption state ρ_{sk} . The deletion operation requires the lessee to perform a Hadamard measurement on ρ_{sk} to produce a valid deletion certificate (c, \mathbf{d}) such that $c = \mathbf{d}^\top \cdot (\mathbf{x}_0 \oplus \mathbf{x}_1) \bmod 2$. Since the decryption capability corresponds to obtaining the information (b, \mathbf{x}_b) , the security of PKE-SKL will ultimately be guaranteed by the AHB security property of the NTCF. However, as shown in [CGJL23], a single valid deletion certificate is insufficient. For example, the adversary can forge a certificate (c, \mathbf{d}) by randomly picking \mathbf{d} and c . Therefore with $1/2$ probability, the adversary can produce a valid certificate. In this case, the adversary does not need to run a Hadamard measurement on its state and can continue to decrypt ciphertext successfully. Therefore, the security of key leasing needs to be further amplified through a parallel repetition mechanism.

In the parallel repeated NTCF-based PKE-SKL scheme, the lessor is required to prepare many (say, N) independent LWE instances $\{k_i = (\mathbf{A}_i, \mathbf{t}_i = \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})\}_{i \in [N]}$. Correspondingly, the lessee generates its public key $\text{pk} = \{k_i, \mathbf{y}_i\}_{i \in N}$ and secret key $\rho_{sk} = \otimes_{i=1}^N \rho_{sk,i}$, where $\rho_{sk,i} = \frac{1}{\sqrt{2}} \sum_{b_i \in \{0,1\}} |b_i, \mathbf{x}_{i,b_i}\rangle =$

$\frac{1}{\sqrt{2}}(|0, \mathbf{x}_{i,0}\rangle + |1, \mathbf{x}_{i,1}\rangle)$. The deletion certificate now consists of a collection of N responses $\{(c_i, \mathbf{d}_i)\}_{i \in [N]}$ certifying the deletion of ρ_{sk} . The ciphertext is revised as $\text{ct} := (\mathbf{ct}_1, \mathbf{ct}_2, \mathbf{ct}_3)$, where $\mathbf{ct}_1 = [\mathbf{r}^\top \mathbf{A}_1, \dots, \mathbf{r}^\top \mathbf{A}_N]^\top$, $\mathbf{ct}_2 = [\mathbf{r}^\top \mathbf{t}_1, \dots, \mathbf{r}^\top \mathbf{t}_N]^\top$ and $\mathbf{ct}_3 = \langle \mathbf{r}, \sum_{i=1}^N \mathbf{y}_i \rangle + \mathbf{e}' + \mu \cdot \lceil q/2 \rceil$. Then, the decryption is performed in a coherent way as

$$\rho_{\text{sk}} \otimes |\mathbf{ct}_3 - [\mathbf{x}_{1,b_1}, \dots, \mathbf{x}_{N,b_N}] \cdot \mathbf{ct}_1 - [b_1, \dots, b_N] \cdot \mathbf{ct}_2\rangle \approx \rho_{\text{sk}} \otimes |\mu \cdot \lceil q/2 \rceil\rangle.$$

Thus the security of the parallel repeated scheme will naturally depend on an amplified AHB property, i.e., the probability that the adversary can simultaneously obtain $\{(c_i, \mathbf{d}_i)\}_{i \in [N]}$ and $\{(b_i, \mathbf{x}_{i,b_i})\}_{i \in [N]}$ can be approximately bounded by 2^{-N} .

As mentioned previously, our main concern now is whether the security previously based on amplified AHB can be ensured if we replace NTCF with NNTCF. On the positive aspect, our NNTCF still enjoys the AHB property under the polynomial hardness of LWE assumption. On the negative side, the claw state generated with our NNTCF will sometimes lead to failure in the verification. In more detail, by using the NNTCF family, the claw state $\rho_{\text{sk},i}$ corresponding to some \mathbf{y}_i may no longer be a uniform superposition state as $\frac{1}{\sqrt{2}}(|0, \mathbf{x}_{i,0}\rangle + |1, \mathbf{x}_{i,1}\rangle)$. For such a non-uniform superposition claw state, performing a Hadamard measurement will no longer produce a valid certificate (c_i, \mathbf{d}_i) that satisfies $c_i = \mathbf{d}_i \cdot (\mathbf{x}_{i,0} \oplus \mathbf{x}_{i,1}) \bmod 2$, thereby causing the certificate verification algorithm to fail. Here, we need to point out that the generation of non-uniform superposition claw states, as described above, does not affect the overall security of our NNTCF-based PKE-SKL scheme.

Intuitively, the Hellinger distance between the distributions $f'_{k,b}(\mathbf{x})$ and $f_{k,b}(\mathbf{x})$ in NNTCF is $1 - 1/\text{poly}(\lambda)$. Although this is not $1 - \text{negl}(\lambda)$, it is still sufficiently close. Therefore, while we cannot generate a uniform superposition claw state with $1 - \text{negl}(\lambda)$ probability every time, in N independent events, we can use the Chernoff bound to ensure that there are at least $0.78N$ valid deletion certificates. We further show that it suffices to verify a major (e.g., the carefully chosen 78%) proportion of the certificate for the security guarantee. Now suppose it requires passing verification of all certificates over a prefixed size- $0.78N$ set of indices i 's. Under the amplified AHB, one can claim that the advantage of any adversary passing the verification without losing decryption capability is approximately $2^{-0.78N}$. However, in the real protocol, the adversary is available to choose any size- $0.78N$ set of indices, and there are $\binom{N}{0.78N}$ many choices over a set of N indices. Up to a union bound, the advantage of a successful adversary can still be properly bounded.

2.3 Proof of Quantumness from LWE-Based NNTCF

In this subsection, we introduce how to use the NNTCF to construct a proof of quantumness protocol based on the polynomial hardness of LWE problem.

Fix a security parameter λ and a LWE-based NNTCF family. Let \mathcal{P} denote a quantum prover and \mathcal{V} denote a classical verifier. The NNTCF-based proof of

Repeat the following 1-4 steps N times:

1. $\mathcal{V} \rightarrow \mathcal{P}$: Generate a NNTCF pair $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0)$, along with a trapdoor \mathbf{td} , send just the k .
2. $\mathcal{P} \rightarrow \mathcal{V}$: Create the state $|\psi\rangle = \sum_{b,\mathbf{x}} |b\rangle |\mathbf{x}\rangle |f'_{k,b}(\mathbf{x})\rangle$ then measure the last register. Return a string \mathbf{y} .
3. $\mathcal{V} \rightarrow \mathcal{P}$: Send a uniformly random challenge $r \xleftarrow{\$} \{0, 1\}$.
4. $\mathcal{P} \rightarrow \mathcal{V}$: Take in r , perform preimage test ($r = 0$) or equation test ($r = 1$):
 - $r = 0$: Perform standard measurement and return proof $\sigma_0 = (b, \mathbf{x}_b)$.
 - $r = 1$: Perform Hadamard measurement and return proof $\sigma_1 = (c, \mathbf{d})$.
5. \mathcal{V} : Take in $\{(\mathbf{y}_i, r_i, \sigma_{r,i})\}_{i \in [N]}$, use each \mathbf{td}_i to compute $\{(\mathbf{x}_{0,i}, \mathbf{x}_{1,i})\}_{i \in [N]}$. Initialize $\text{count} = 0$. For each proof $\sigma_{r,i}$, perform the following check:
 - If $r = 0$, check the validity of $(b_i, \mathbf{x}_{b_i,i})$; If $r = 1$, check \mathbf{d}_i is non-zero and $c_i = \mathbf{d}_i^\top \cdot (\mathbf{x}_{0,i} \oplus \mathbf{x}_{1,i}) \bmod 2$.
 - If above check passes for $r = 1$, increment the value of count by 1. We let N_1 denote the total number of equation tests. If $N_1 > \frac{1}{4}N$ and the final $\text{count} > 0.75N_1$, output 1, else output \perp .

Fig. 6. Polynomial-sized proof of quantumness from NNTCF (Simple version)

quantumness protocol is described in Fig. 6. Our NNTCF-based proof of quantumness protocol can be viewed as a revised version of the works in [BCM+18] and [BKVV20], while the security is solely based on the AHB property of the NNTCF. Compared to [BCM+18], the protocol construction no longer requires a superpolynomial LWE modulus; compared to [BKVV20], the protocol construction no longer requires the random oracle model (ROM).

Quantum Completeness. Regarding the preimage test, as long as the claw state generated by the NNTCF is still over the two preimages $(0, \mathbf{x}_0)$ and $(1, \mathbf{x}_1)$ with any amplitude, any one of the two measured values will certainly pass the verification. As shown in Fig. 6, there are almost half of the NNTCF instances devoted to the preimage test instead of the equation test. Therefore, we need to correspondingly adapt the number of valid equation tests such that an honest quantum prover can pass the equation test except with negligible probability. Overall, the probability that the quantum prover can successfully pass the protocol described in Fig. 6 is $1 - \text{negl}(\lambda)$.

Classical Soundness. Intuitively, any malicious classical prover will be ruled out as it is required to pass a majority of the equation tests in our protocol. In particular, under the AHB property, conditioned on always passing the preimage test, any classical PPT prover should not be able to subsequently win in the equation test with probability noticeably larger than $\frac{1}{2}$.

As a result, the cheating advantage of any classical PPT adversary to pass a major proportion (say, 75%) of the equation tests should be negligible.

2.4 Open Problems

Our work opens several promising avenues for future research, particularly concerning the NNTCF construction and its potential applications. While we focused primarily on quantum key leasing due to its compatibility with NNTCF, numerous other NTCF-based applications, particularly those involving the adaptive hardcore bit property, could benefit from our findings. We identify significant unexplored directions to extend and generalize our results, which could inspire further advancements in the field. They can be mainly divided into the following three categories.

The first category is about applications based on standard NTCF over classical channels. Within this category, we have successfully improved both the proof of quantumness scheme in [BCM+18] and the key leasing scheme in [CGJL23]. As far as we know, there are more applications within this category such as the certifiable randomness generation protocol [BCM+18] and the semi-quantum money [RS19]. However, these adaptations appear to be more involved and we leave them as future work.

The second category concerns the tasks based on variants of NTCF over classical channels. An example is the quantum delegated computation in [Mah18b], which is based on the extended TCF. It seems more effort would be needed to properly adapt these applications, which can also be interesting for future work.

The last category includes all applications based on (variants of) NTCF that require quantum channels. One example is the revocable quantum digital signature [MPY23]. In this work, to start, we tried to focus on the applications solely over classical channels. However, we believe that the adaptation for this category can be an interesting direction for future research.

3 Preliminaries

3.1 Notions

In this paper, we use λ to denote the security parameter. For positive integer N , let $[N]$ denote the set $\{1, 2, \dots, N\}$. Let \mathbb{Z} be the set of integers and \mathbb{N} be the set of natural numbers. For any $q \geq 2 \in \mathbb{N}$, we let \mathbb{Z}_q denote the ring of integers modulo q . The vectors are denoted by bold lowercase letters (e.g., $\mathbf{x} \in \mathbb{Z}^n$), matrices by bold uppercase letters (e.g., $\mathbf{A} \in \mathbb{Z}^{m \times n}$). We write $\text{negl}(\lambda)$ for any function $f : \mathbb{N} \rightarrow \mathbb{R}_+$ such that for any polynomial p , $\lim_{\lambda \rightarrow \infty} p(\lambda)f(\lambda) = 0$. Let $\text{poly}(n)$ be a polynomial in n . Let $\mathcal{B}(\mathbf{c}, R)$ denote the ball with center \mathbf{c} and radius R . Let the letter D denote a distribution over a finite domain X and f for a density on X , i.e., a function $f : X \rightarrow [0, 1]$ s.t. $\sum_{x \in X} f(x) = 1$. $x \leftarrow D$ indicates that x is sampled from the distribution D , and $x \xrightarrow{\$} X$ indicates that x is sampled uniformly from the set X in random. Let D_X for the set of all densities on X . For any $f \in D_X$, $\text{SUPP}(f)$ is denoted the support of f , $\text{SUPP}(f) = \{x \in X | f(x) > 0\}$.

3.2 Lattices and Lattice Problems

We give some background on lattice in this section. Let $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\} \subset \mathbb{R}^n$, where $m \leq n$ consist of m linearly independent vectors. The m -dimensional lattice generated by the basis \mathbf{B} is

$$A = \mathcal{L}(\mathbf{B}) = \left\{ \mathbf{x} = \sum_{i \in [m]} c_i \mathbf{b}_i, c_i \in \mathbb{Z} \right\}.$$

In the following part, we will introduce discrete Gaussian distribution over a lattice A and some properties of discrete Gaussian distribution. For a full-rank, symmetric, positive definite $n \times n$ matrix Σ , we define the n -dimension Gaussian function of deviation parameter $\sqrt{\Sigma}$ as $\rho_{\sqrt{\Sigma}}(\mathbf{x}) = \exp(-\pi \cdot (\mathbf{x})^T \Sigma^{-1}(\mathbf{x}))$, for any $\mathbf{x} \in \mathbb{R}^n$. Particularly, if Σ is a diagonal matrix and each non-zero term equals r^2 , Gaussian function can be simplified as $\rho_r(\mathbf{x}) = \exp(-\pi \cdot \|\mathbf{x}\|^2/r^2)$.

We recall the discrete Gaussian distribution on the integer lattice \mathbb{Z}^n .

Definition 1 (Discrete Gaussian Distribution). For a full-rank, symmetric, positive definite $n \times n$ matrix Σ , we define the n -dimension discrete Gaussian distribution over the lattice \mathbb{Z}^n , $D_{\mathbb{Z}^n, \sqrt{\Sigma}}$ of standard deviation parameter matrix $\sqrt{\Sigma}$ by

$$\forall \mathbf{x} \in \mathbb{Z}^n : D_{\mathbb{Z}^n, \sqrt{\Sigma}}(\mathbf{x}) = \rho_{\sqrt{\Sigma}}(\mathbf{x}) / \rho_{\sqrt{\Sigma}}(\mathbb{Z}^n),$$

where $\rho_{\sqrt{\Sigma}}(\mathbb{Z}^n) = \sum_{\mathbf{x} \in \mathbb{Z}^n} \rho_{\sqrt{\Sigma}}(\mathbf{x})$.

Now we recall the following lemma about the approximate upper bounds of the vectors selected from discrete Gaussian distribution.

Lemma 1 ([Ban93, Lemma 1.4]). Let $n \in \mathbb{N}, r > 0$, then it holds that

- 1) For any $k > 0$, $Pr[|x| > kr; x \leftarrow \mathcal{D}_{\mathbb{Z}, r}] \leq 2e^{-\frac{k^2}{2}}$;
- 2) for any $k > 1$, $Pr[\|\mathbf{x}\| > kr\sqrt{n}; \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^n, r}] < k^n e^{\frac{n}{2}(1-k^2)}$.

We can now define bounded discrete Gaussian distribution.

Definition 2 (Bounded Gaussian Distribution). For the integer lattice \mathbb{Z}^n , the bound B and the derivation parameter r , the bounded discrete Gaussian distribution is defined by:

$$D_{\mathbb{Z}^n, r, B}(\mathbf{x}) = \begin{cases} \frac{\rho_r(\mathbf{x})}{\sum_{\|\mathbf{x}\| \leq B} \rho_r(\mathbf{x})} & , \text{ if } \|\mathbf{x}\| \leq B, \\ 0 & , \text{ otherwise.} \end{cases}$$

Due to the Lemma 1, when $B > r\sqrt{n}$, the bounded discrete Gaussian distribution $D_{\mathbb{Z}^n, r, B}$ is statistically closed to the discrete Gaussian distribution $D_{\mathbb{Z}^n, r}$.

Definition 3 (LWE Problem). For a security parameter λ , let $n, m, q \in \mathbb{N}$ be integer functions of λ . Let $\chi = \chi(\lambda)$ be a distribution over \mathbb{Z} . The $\text{LWE}_{n, m, q, \chi}$ problem is to distinguish between the distributions $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod q)$ and (\mathbf{A}, \mathbf{u}) , where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. Often we consider the hardness of solving LWE for any function m such that m is at most a polynomial in $n \log q$. This problem is denoted $\text{LWE}_{n, q, \chi}$.

4 Noticeable Noisy Trapdoor Claw-Free Function Family

In this section, we will describe our construction of a noticeable noisy trapdoor claw-free function (NNTCF) family and prove its properties including the adaptive hardcore bit.

4.1 Construction of NNTCF from LWE with Polynomial Modulus

Our construction of NNTCF is similar to the one in [BCM+18]. Let λ be the security parameter. All other parameters are functions of λ as follows: $l = \mathcal{O}(\lambda)$, $n \geq \lambda \cdot l \cdot \lceil \log q \rceil$, $m \geq n \cdot \lceil \log q \rceil$ and $m > 500$, $w = n \lceil \log q \rceil$, $q \geq 8\sigma\sqrt{m}$, and q is a prime, $\sigma_0 \geq n^{\frac{3}{2}}\sqrt{m}$, $150 \cdot m \cdot \sigma_0 \leq \sigma \leq \frac{q}{C_T \sqrt{mn \log q}}$.

Under the above parameters, we describe the noticeable NTCF family \mathcal{F}_{LWE} based on LWE with polynomial modulus. Let $\mathcal{X} = \mathbb{Z}_q^n$ and $\mathcal{Y} = \mathbb{Z}_q^m$. The key space $\mathcal{K}_{\mathcal{F}_{\text{LWE}}}$ is subset of $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$. For $b \in \{0, 1\}$, $x \in \mathcal{X}$ and the key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0)$, the $f_{k,b}(\mathbf{x})$ is given as

$$\forall \mathbf{y} \in \mathcal{Y} : (f_{k,b}(\mathbf{x}))(\mathbf{y}) = D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{y} - \mathbf{A}\mathbf{x} - b \cdot \mathbf{A}\mathbf{s}), \quad (1)$$

Then we show that each of the properties of NNTCF holds. The first two properties are the same as that of LWE-based NTCF in [BCM+18], while the last two properties differ due to the use of polynomial-size modulus.

Efficient Function Generation. On input the security parameter λ , the procedure $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ samples a random $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, together with trapdoor information $\mathbf{T}_{\mathbf{A}}$. This is done using the procedure $\text{GENTRAP}(1^n, 1^m, q)$.

Moreover, the distribution on matrices \mathbf{A} returned by GENTRAP is negligibly close to the uniform distribution on $\mathbb{Z}_q^{m \times n}$.

Next, the sampling procedure selects $\mathbf{s} \in \{0, 1\}^n$ uniformly at random, and a vector $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \sigma_0, \sigma_0\sqrt{m}}$. $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ returns $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and $\text{td}_k = \mathbf{T}_{\mathbf{A}}$.

Trapdoor Injective Pair

(a) *Trapdoor.* For any key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$ and for all $\mathbf{x} \in \mathcal{X}$,

$$\text{SUPP}(f_{k,0}(\mathbf{x})) = \{ \mathbf{A}\mathbf{x} + \mathbf{e} \mid \|\mathbf{e}\| \leq \sigma\sqrt{m} \}, \quad (2)$$

$$\text{SUPP}(f_{k,1}(\mathbf{x})) = \{ \mathbf{A}\mathbf{x} + \mathbf{A}\mathbf{s} + \mathbf{e} \mid \|\mathbf{e}\| \leq \sigma\sqrt{m} \}. \quad (3)$$

The procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ takes as input the trapdoor $t_{\mathbf{A}}$, $b \in \{0, 1\}$, and $\mathbf{y}' \in \mathcal{Y}$, it uses the algorithm INVERT to determine \mathbf{x}' , \mathbf{e}' such that $\mathbf{y}' = \mathbf{A}\mathbf{x}' + \mathbf{e}'$, and returns the element $\mathbf{x}' - b \cdot \mathbf{s} \in \mathcal{X}$. This procedure returns the unique correct outcome provided $\mathbf{y}' = \mathbf{A}\mathbf{x}' + \mathbf{e}'$ for some \mathbf{e}' such that $\|\mathbf{e}'\| \leq \sigma\sqrt{m}$. This condition is satisfied for all $\mathbf{y}' \in \text{SUPP}(f_{k,b}(\mathbf{x}'))$ provided σ is chosen so that $\sigma \leq \frac{q}{C_T \sqrt{mn \log q}}$.

(b) *Injective Pair.* We let \mathcal{R}_k be the set of all pairs $(\mathbf{x}_0, \mathbf{x}_1)$ such that $f_{k,0}(\mathbf{x}_0) = f_{k,1}(\mathbf{x}_1)$. By definition, this occurs if and only if $\mathbf{x}_1 = \mathbf{x}_0 - \mathbf{s} \pmod{q}$, and so \mathcal{R}_k is a perfect matching.

Efficient Range Superposition. For $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$, $b \in \{0, 1\}$ and $\mathbf{x} \in \mathcal{X}$, let

$$(f'_{k,b}(\mathbf{x}))(\mathbf{y}) = D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{y} - \mathbf{A}\mathbf{x} - b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e}_0)). \quad (4)$$

Note that $f'_{k,0}(\mathbf{x}) = f_{k,0}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$. The distributions $f'_{k,1}(\mathbf{x})$ and $f_{k,1}(\mathbf{x})$ are shifted by \mathbf{e}_0 . Given the key k and $\mathbf{x} \in \mathcal{X}$, the densities $f'_{k,0}(\mathbf{x})$ and $f'_{k,1}(\mathbf{x})$ are efficiently computable. For all $\mathbf{x} \in \mathcal{X}$,

$$\text{SUPP}(f'_{k,0}(\mathbf{x})) = \text{SUPP}(f_{k,0}(\mathbf{x})), \quad (5)$$

$$\text{SUPP}(f'_{k,1}(\mathbf{x})) = \{\mathbf{A}\mathbf{x} + \mathbf{e} + \mathbf{A}\mathbf{s} + \mathbf{e}_0 \mid \|\mathbf{e}\| \leq 2\sigma\sqrt{m}\}. \quad (6)$$

- (a) Using that $\sigma \geq \sigma_0 m$, it follows that the norm of the term $\mathbf{e}_0 + \mathbf{e}$ in Eq. (6) is always at most $3\sigma\sqrt{m}$. Therefore, the inversion procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ can be guaranteed to return \mathbf{x} on input $t_{\mathbf{A}}, b \in \{0, 1\}, \mathbf{y} \in \text{SUPP}(f'_{k,b}(\mathbf{x}))$ if we strengthen the requirement on σ to $\sigma \leq \frac{q}{2C_T\sqrt{mn}\log q}$. This strengthened trapdoor requirement also implies that for all $b \in \{0, 1\}, (\mathbf{x}_0, \mathbf{x}_1) \in \mathcal{R}_k$, and $\mathbf{y} \in \text{SUPP}(f'_{k,b}(\mathbf{x}_b)) \cap \text{SUPP}(f'_{k,b\oplus 1}(\mathbf{x}_{b\oplus 1}))$, $\text{INV}_{\mathcal{F}_{\text{LWE}}}(t_{\mathbf{A}}, b\oplus 1, \mathbf{y}) = \mathbf{x}_{b\oplus 1}$.
- (b) The procedure $\text{CHK}_{\mathcal{F}_{\text{LWE}}}$ is identical to the one in [BCM+18]. On input $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0), b \in \{0, 1\}, \mathbf{x} \in \mathcal{X}$, and $\mathbf{y} \in \mathcal{Y}$, if $b = 0$, it computes $\mathbf{e}' = \mathbf{y} - \mathbf{A}\mathbf{x}$. If $\|\mathbf{e}'\| \leq 2\sigma\sqrt{m}$, the procedure returns 1, and 0 otherwise. If $b = 1$, it computes $\mathbf{e}' = \mathbf{y} - \mathbf{A}\mathbf{x} - (\mathbf{A}\mathbf{s} + \mathbf{e}_0)$. If $\|\mathbf{e}'\| \leq 2\sigma\sqrt{m}$, it returns 1, and 0 otherwise.
- (c) The procedure $\text{SAMP}_{\mathcal{F}_{\text{LWE}}}$ is identical to the one in [BCM+18]. We bound the Hellinger distance between the densities $f_{k,b}(\mathbf{x})$ and $f'_{k,b}(\mathbf{x})$. If $b = 0$ they are identical. If $b = 1$, both densities are shifts of $D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}$, where the shifts differ by \mathbf{e}_0 and $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \sigma_0, \sigma_0\sqrt{m}}$. It holds that $H^2(f_{k,1}(\mathbf{x}), f'_{k,1}(\mathbf{x})) \leq 1 - e^{-\frac{9\sqrt{m}\|\mathbf{e}_0\|}{4\sigma}}(1 - 2e^{-\frac{1}{2}m}) \leq 1 - e^{-\frac{9m\sigma_0}{4\sigma}}(1 - 2e^{-\frac{1}{2}m}) \leq 1 - e^{-\frac{3}{200}}(1 - 2e^{-\frac{1}{2}m})$. When $m > 500$, $1 - e^{-\frac{3}{200}}(1 - 2e^{-\frac{1}{2}m}) < \frac{1}{50}$. Therefore, the requirement $E_{x \leftarrow \mathbb{Z}_q^n}[H^2(f_{k,1}(\mathbf{x}), f'_{k,1}(\mathbf{x}))] \leq \frac{1}{50}$ holds.

Finally, it remains to describe the procedure $\text{SAMP}_{\mathcal{F}_{\text{LWE}}}$. At the first step, the procedure creates the following superposition

$$\sum_{\mathbf{e} \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e})} |\mathbf{e}\rangle. \quad (7)$$

At the second step, the procedure creates a uniform superposition over $\mathbf{x} \in \mathcal{X}$, yielding the state

$$(2q)^{-\frac{n}{2}} \sum_{\substack{\mathbf{x} \in \mathcal{X} \\ b \in \{0,1\} \\ \mathbf{e} \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e})} |b, \mathbf{x}\rangle |\mathbf{e}\rangle. \quad (8)$$

At the third step, using the key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, the procedure computes

$$\begin{aligned}
 & (2q)^{-\frac{n}{2}} \sum_{\substack{\mathbf{x} \in \mathcal{X} \\ b \in \{0,1\} \\ \mathbf{e} \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e})|b, \mathbf{x}\rangle|\mathbf{A}\mathbf{x} + \mathbf{e} + b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e}_0)\rangle} \\
 &= (2q)^{-\frac{n}{2}} \sum_{\substack{\mathbf{x} \in \mathcal{X} \\ b \in \{0,1\} \\ \mathbf{y} \in \text{SUPP}(f'_{k,b}(\mathbf{x}))}} \sqrt{f'_{k,b}(\mathbf{x})(\mathbf{y})|b, \mathbf{x}\rangle|\mathbf{y}\rangle}
 \end{aligned} \tag{9}$$

Adaptive Hardcore Bit. Now we show that our NNTCF family also enjoys the adaptive hardcore bit property. We start by providing some useful statements and lemmata. Recall that $\mathcal{X} = \mathbb{Z}_q^n$ and let $w = n \lceil \log q \rceil$. Let $\mathcal{J} : \mathcal{X} \rightarrow \{0, 1\}^w$ be such that $\mathcal{J}(\mathbf{x})$ returns the binary representation of $\mathbf{x} \in \mathcal{X}$. For $b \in \{0, 1\}$, $\mathbf{x} \in \mathcal{X}$, and $\mathbf{d} \in \{0, 1\}^w$, let $I_{b,\mathbf{x}}(\mathbf{d}) \in \{0, 1\}^n$ be the vector whose each coordinate is obtained by taking the inner product mod 2 of the corresponding block of $\lceil \log q \rceil$ coordinates of \mathbf{d} and of $\mathcal{J}(\mathbf{x}) \oplus \mathcal{J}(\mathbf{x} - (-1)^b \mathbf{1})$, where $\mathbf{1} \in \mathbb{Z}_q^n$ is the vector with all its coordinates equal to 1 $\in \mathbb{Z}_q$. There is a useful claim in [BCM+18] that the inner product $\mathbf{d} \cdot \mathcal{J}(\mathbf{x}) \oplus \mathcal{J}(\mathbf{x} - (-1)^b \mathbf{1})$ is exactly equal to $I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s}$, which is recalled as follows.

Lemma 2 (Claim 4.5 in [BCM+18]). *For all $b \in \{0, 1\}$, $\mathbf{x} \in \mathcal{X}$, $\mathbf{d} \in \{0, 1\}^w$ and $\mathbf{s} \in \{0, 1\}^n$ the following equality holds:*

$$\mathbf{d} \cdot (\mathcal{J}(\mathbf{x}) \oplus \mathcal{J}(\mathbf{x} - (-1)^b \mathbf{s})) = I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s}. \tag{10}$$

Note that in [BCM+18], the \mathbf{d} is only required to have one non-zero place in the first and second half as each bit of secret \mathbf{s} is computationally indistinguishable from random. In our case, we consider a relaxed condition on \mathbf{s} , which then requires the string \mathbf{d} to have more non-zero positions. Therefore, for $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0)$, $b \in \{0, 1\}$ and $\mathbf{x} \in \mathcal{X}$, we define the set $G_{k,b,\mathbf{x}}$ as

$$G_{k,b,\mathbf{x}} = \left\{ \mathbf{d} \in \{0, 1\}^w : \text{HW}(I_{b,\mathbf{x}}(\mathbf{d})_{\mathcal{I}_b}) \geq \frac{n}{8} \right\},$$

where $\text{HW}(\cdot)$ represents the Hamming weight and $I_{b,\mathbf{x}}(\mathbf{d})_{\mathcal{I}_b}$ is the concatenation of all the entries indexed by \mathcal{I}_b , which satisfies $\mathcal{I}_b = \{b \frac{n}{2}, \dots, b \frac{n}{2} + \frac{n}{2}\}$. Besides, we also divide $\mathbf{s} = (\mathbf{s}_0, \mathbf{s}_1)$. Here \mathbf{s}_0 is the vector containing the first $\frac{n}{2}$ entries and \mathbf{s}_1 contains the last $\frac{n}{2}$ entries. We define $\hat{G}_{\mathbf{s}_1,0,\mathbf{x}_0} = \hat{G}_{\mathbf{s}_0,1,\mathbf{x}_1} = G_{k,0,\mathbf{x}_0} \cap G_{k,1,\mathbf{x}_1}$. Actually, for all $b \in \{0, 1\}$ and $\mathbf{x} \in \mathcal{X}$, if \mathbf{d} is sampled uniformly at random, $\mathbf{d} \notin \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}_b}$ with probability $e^{-\frac{n}{32} + 1}$.

Theorem 4 (Adaptive hardcore bit). *For m, n, q set the same as Sect. 4.1 and $\sigma_0 \geq n^{\frac{3}{2}} \sqrt{m}$, assume the hardness assumption $\text{LWE}_{q,\sigma_0}^{m,n}$ and $\mathbf{s} \in \{0, 1\}^n$, we define two sets:*

$$\begin{aligned}
 H_{\mathbf{s}} &= \{(b, \mathbf{x}, \mathbf{d}, (\mathbf{d} \cdot (\mathcal{J}(\mathbf{x}) \oplus \mathcal{J}(\mathbf{x} - (-1)^b \mathbf{s})) \bmod 2)) | b \in \{0, 1\}, \mathbf{x} \in \mathbb{Z}_q^n, \\
 &\quad \mathbf{d} \in \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}}\}, \\
 \bar{H}_{\mathbf{s}} &= \{(b, \mathbf{x}, \mathbf{d}, c) | (b, \mathbf{x}, \mathbf{d}, c \oplus 1) \in H_{\mathbf{s}}\}.
 \end{aligned}$$

For any quantum polynomial-time algorithm $\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0, 1\} \times \mathbb{Z}_q^n \times \{0, 1\}^{n \lceil \log(q) \rceil} \times \{0, 1\}$ and the any LWE sample $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0) \leftarrow \mathbf{Gen}(1^\lambda, \mathbf{s}, m, n)$, the negligible difference always exists:

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0) \in H_{\mathbf{s}}] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0) \in \overline{H}_{\mathbf{s}}]| \leq \text{negl}(\lambda)$$

Refer to [BCM+18, Section 4.4.1], it suffices to prove the following lemma, which implies the above theorem.

Lemma 3. *Under the hardness assumption $\text{LWE}_{q, \sigma_0, \mathbf{s}}^{m, n}$, $\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0, 1\} \times \mathbb{Z}_q^n \times \{0, 1\}^{n \lceil \log(q) \rceil} \times \{0, 1\}$ is a quantum polynomial-time algorithm. The following two distributions are computationally indistinguishable:*

$$D_0 = (k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0), (b, x, \mathbf{d}, c) \leftarrow \mathcal{A}(k), I_{b, \mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2)$$

$$D_1 = (k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_0), (b, x, \mathbf{d}, c) \leftarrow \mathcal{A}(k), (\delta_{d \in \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}}} \cdot r) \oplus (I_{b, \mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2))$$

where r is a random bit and $\delta_{d \in \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}}} = 1$ if $d \in \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}}$ and 0 otherwise.

Here, we recall some useful notions such as moderate vector and moderate matrix, together with the lemma of the lower bound of the probability for a uniformly selected matrix to be moderate from [BCM+18].

Definition 4. *For a vector $\mathbf{b} \in \mathbb{Z}_q^n$, we say \mathbf{b} is moderate if there are at least $\frac{n}{4}$ entries of \mathbf{b} has absolute value in the range $(\frac{q}{8}, \frac{3q}{8}]$. A matrix $\mathbf{C} \in \mathbb{Z}_q^{l \times n}$ is moderate if every vector in the spanning space of row vectors of \mathbf{C} , $\text{span}(\mathbf{C})$, is moderate.*

Lemma 4 ([BCM+18, Lemma 4.8]). *Let q be prime and l, n be integers. Then*

$$\Pr_{\mathbf{C} \leftarrow \mathbb{Z}_q^{l \times n}} [\mathbf{C} \text{ is moderate}] \geq 1 - q^l \cdot 2^{-\frac{n}{8}}.$$

Now suppose $\sigma \geq n$, we present our main lemma as follows.

Lemma 5. *Let $\mathbf{C} \in \mathbb{Z}_q^{l \times n}$ be an arbitrary moderate matrix and $\hat{\mathbf{d}} \in \{0, 1\}^n$ be an arbitrary non-zero binary vector satisfying that its hamming weight is at least $\frac{n}{4}$. Let $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow D_{\mathbb{Z}^n, \sigma, \sigma\sqrt{n}}$, where $\sigma = n$. Consider the random variables $\mathbf{v} = \mathbf{C}\mathbf{s} \bmod q$ and $z = \langle \hat{\mathbf{d}}, \mathbf{s} \rangle \bmod 2$ conditioned on $\mathbf{s} + \mathbf{e} = \mathbf{t}$ for any \mathbf{t} fixed. Then statistical distance between the distribution of (\mathbf{v}, z) and the distribution of $U(\mathbb{Z}_q^l \times \mathbb{Z}_2)$ is at most $q^{\frac{l}{2}} \cdot 2^{-\frac{n}{4}}$.*

Proof (Proof of Lemma 5). Let f be the probability density function of (\mathbf{v}, z) and \hat{f} be the Fourier transform over $\mathbb{Z}_q^l \times \mathbb{Z}_2$. It's clear that $\hat{f}(\mathbf{0}, 0) = 1$. Let U denote the density of the uniform distribution over $\mathbb{Z}_q^l \times \mathbb{Z}_2$. It's easy to see that $\hat{U}(\mathbf{0}, 0) = 1$ and $\hat{U}(\hat{\mathbf{v}}, \hat{z}) = 0$ for all $(\hat{\mathbf{v}}, \hat{z}) \neq (\mathbf{0}, 0)$. Then we can compute: $\frac{1}{2} \|f - U\|_1 \leq \sqrt{\frac{q^l}{2}} \|f - U\|_2 = \frac{1}{2} \|\hat{f} - \hat{U}\|_2 =$

$\frac{1}{2} \left(\sum_{(\widehat{\mathbf{v}}, \widehat{\mathbf{z}} \in \mathbb{Z}_q^l \times \mathbb{Z}_2 \setminus (0,0)} \left| \widehat{f}(\widehat{\mathbf{v}}, \widehat{\mathbf{z}}) \right|^2 \right)^{1/2}$, where the first inequality follows from the Cauchy-Schwarz inequality and the second line follows from Parseval's identity. Denote $\omega_{2q} = e^{-\frac{2\pi\sqrt{-1}}{2q}}$, then we can write: $\widehat{f}(\widehat{\mathbf{v}}, \widehat{\mathbf{z}}) = E_{\mathbf{s}} \left[\omega_{2q}^{(2\widehat{\mathbf{v}}^\top \mathbf{C} + q\widehat{\mathbf{z}}\widehat{\mathbf{d}}^\top)\mathbf{s}} \right] = E_{\mathbf{s}} \left[\omega_{2q}^{\mathbf{w}^\top \mathbf{s}} \right] = \prod_{i \in [n]} E_{s_i} \left[\omega_{2q}^{w_i s_i} \right]$, where $\mathbf{w}^\top = 2 \cdot \widehat{\mathbf{v}}^\top \mathbf{C} + q \cdot \widehat{\mathbf{z}} \cdot \widehat{\mathbf{d}}^\top \in \mathbb{Z}_{2q}^n$. To compute $\widehat{f}(\widehat{\mathbf{v}}, \widehat{\mathbf{z}})$, we have:

$$\begin{aligned} & \Pr[s_i | e_i + s_i = t_i, t_i \text{ fixed}, s_i \stackrel{\$}{\leftarrow} \{0, 1\}, e_i \leftarrow D_{\mathbb{Z}, \sigma}] \\ &= \frac{\rho_\sigma(t_i - s_i)}{\rho_\sigma(t_i) + \rho_\sigma(t_i - 1)} = \frac{e^{-\pi(-2s_i t_i + s_i^2)/\sigma^2}}{1 + e^{-\pi(-2t_i + 1)/\sigma^2}} \end{aligned}$$

Therefore,

$$\Pr[s_i | e_i + s_i = t_i, t_i \text{ fixed}, s_i \stackrel{\$}{\leftarrow} \{0, 1\}, e_i \leftarrow D_{\mathbb{Z}, \sigma}] = \begin{cases} \frac{1}{e^{-\pi(-2t_i + 1)/\sigma^2} + 1}, & s_i = 0; \\ \frac{e^{-\pi(-2t_i + 1)/\sigma^2}}{e^{-\pi(-2t_i + 1)/\sigma^2} + 1}, & s_i = 1. \end{cases}$$

$$\text{For } (\widehat{\mathbf{v}}, \widehat{\mathbf{z}}) = (\mathbf{0}, 1), E_{s_i}[\omega_{2q}^{w_i s_i}] = \begin{cases} \frac{1 - e^{-\pi(-2t_i + 1)/\sigma^2}}{e^{-\pi(-2t_i + 1)/\sigma^2} + 1}, & d_i = 1; \\ 1, & d_i = 0. \end{cases}$$

When $d_i = 1$, $E_{s_i}[\omega_{2q}^{w_i s_i}] \leq \frac{1 - e^{(-2\pi\sigma\sqrt{n} - \pi)/\sigma^2}}{1 + e^{(-2\pi\sigma\sqrt{n} - \pi)/\sigma^2}} \leq \frac{1 - e^{-3\pi/\sqrt{n}}}{1 + e^{-3\pi/\sqrt{n}}} \leq \frac{1}{2}$. Since the hamming weight of $\widehat{\mathbf{d}}$ is at least $\frac{n}{4}$, $\widehat{f}(\mathbf{0}, 1) \leq (\frac{1}{2})^{\frac{n}{4}}$.

For $\widehat{\mathbf{v}} \neq \mathbf{0}$, $E_{s_i}[\omega_{2q}^{w_i s_i}] = 1 - \frac{e^{(2\pi t_i - \pi)/\sigma^2}}{1 + e^{(2\pi t_i - \pi)/\sigma^2}} (1 - e^{2\pi\sqrt{-1}w_i/(2q)})$, and for the second term of the formula above,

$$\begin{aligned} & \left| \frac{e^{(2\pi t_i - \pi)/\sigma^2}}{1 + e^{(2\pi t_i - \pi)/\sigma^2}} (1 - e^{2\pi\sqrt{-1}w_i/(2q)}) \right| \geq \frac{2e^{(2\pi t_i - \pi)/\sigma^2}}{1 + e^{(2\pi t_i - \pi)/\sigma^2}} \left| \sin\left(\frac{\pi w_i}{2q}\right) \right| \\ & \geq \frac{2e^{(-2\pi\sigma\sqrt{n} - \pi)/\sigma^2}}{1 + e^{(-2\pi\sigma\sqrt{n} - \pi)/\sigma^2}} \left| \sin\left(\frac{\pi w_i}{2q}\right) \right| \geq 2 \cdot \sin(\pi/8), \end{aligned}$$

the last inequality exists for at least $\frac{n}{4}$, $i \in [n]$ because \mathbf{C} is moderate. In this case $E_{s_i}(\omega_{2q}^{w_i s_i}) \leq 1 - 2 \cdot \sin(\pi/8)$. Hence for $\widehat{\mathbf{v}} \neq \mathbf{0}$, $\widehat{f}(\widehat{\mathbf{v}}, \widehat{\mathbf{z}}) \leq (1 - 2 \cdot \sin(\pi/8))^{\frac{n}{4}} \leq (\frac{1}{4})^{\frac{n}{4}}$. Therefore, $\Delta(f, U) \leq \frac{1}{2} \sqrt{(\frac{1}{2})^{\frac{n}{4}} + 2(q^l - 1) (\frac{1}{4})^{\frac{n}{4}}} \leq \frac{1}{2} \sqrt{2 \cdot q^l \cdot (\frac{1}{2})^{\frac{n}{2}}} \leq q^{\frac{l}{2}} 2^{-\frac{n}{4}}$ \square

Based on the lemma above, the following lemma can be proved by following the same merit of the proof for [BCM+18, Lemma 4.6].

Lemma 6 *Let q be a prime, $l, n \geq 1$ integers, and $\mathbf{C} \in \mathbb{Z}_q^{l \times n}$ a uniformly random matrix. With probability at least $1 - q^l \cdot 2^{-\frac{n}{8}}$ over the choice of \mathbf{C} the following holds. For a fixed \mathbf{C} , all $\mathbf{v} \in \mathbb{Z}_q^l$ and $\widehat{\mathbf{d}} \in \{0, 1\}^n$ with hamming weight larger than $\frac{n}{4}$, the distance of $(\widehat{\mathbf{d}} \cdot \mathbf{s} \bmod 2)$, where \mathbf{s} is uniform in $\{0, 1\}^n$ conditioned on $\mathbf{C}\mathbf{s} = \mathbf{v}$ and $\mathbf{s} + \mathbf{e} = \mathbf{t}$ fixed, where $\mathbf{e} \leftarrow D_{\mathbb{Z}^n, \sigma, \sigma\sqrt{n}}$ is within statistical distance $\mathcal{O}(q^{\frac{3l}{2}} \cdot 2^{-\frac{n}{4}})$ of the uniform distribution $\{0, 1\}$.*

Our idea to circumvent noise flooding in the proof of [BCM+18, Lemma 4.4] is inspired by the Gaussian decomposition technique introduced in [BD20]. In short, to hide \mathbf{s} in $\mathbf{F}\mathbf{s}$, one can decompose $\mathbf{e}_0 = \mathbf{F}\mathbf{e}_0^{(1)} + \mathbf{e}_0^{(2)}$ and use $\mathbf{e}_0^{(1)}$ to hide \mathbf{s} .

Now we can prove the Lemma 3.

Proof (Proof of Lemma 3). We use hybrid arguments to prove the lemma. Here are the six hybrids we introduce.

In the **Hybrid 1**,

$$D^{(1)} = ((\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}_0), (b, x, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}_0), I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2),$$

where $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F}$ and $\mathbf{B} \in \mathbb{Z}_q^{m \times l}$, $\mathbf{C} \in \mathbb{Z}_q^{l \times n}$ and \mathbf{F} is selected from the distribution $D_{\mathbb{Z}_q^{m \times n}, \sigma_{\mathbf{F}}}$, where $\sigma_{\mathbf{F}} = \sqrt{n}$. According to the hardness of $\text{LWE}_{q, \sigma_{\mathbf{F}}}^{m, l}$ assumption, distribution D_0 and $D^{(1)}$ are computationally indistinguishable.

In the **Hybrid 2**,

$$D^{(2)} = ((\tilde{\mathbf{A}}', \tilde{\mathbf{A}}'\mathbf{s} + \mathbf{e}_0), (b, x, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}', \tilde{\mathbf{A}}'\mathbf{s} + \mathbf{e}_0), I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2).$$

The only difference between distribution $D^{(1)}$ and $D^{(2)}$ is that we select $\tilde{\mathbf{A}}' = \mathbf{BC} + \mathbf{F}$ with totally the same parameters with $\tilde{\mathbf{A}}$ in $D^{(1)}$, but abort if $\|\mathbf{F}\| \geq \sigma_{\mathbf{F}}\sqrt{m}$. As the probability of aborting is negligible, the distributions of $D^{(1)}$ and $D^{(2)}$ are statistically indistinguishable.

In the **Hybrid 3**,

$$D^{(3)} = ((\tilde{\mathbf{A}}', \mathbf{BCs} + \mathbf{F}(\mathbf{s} + \mathbf{e}_0^{(1)}) + \mathbf{e}_0^{(2)}), \\ (b, \mathbf{x}, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}', \mathbf{BCs} + \mathbf{F}(\mathbf{s} + \mathbf{e}_0^{(1)}) + \mathbf{e}_0^{(2)}), I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2),$$

where $\mathbf{e}_0^{(1)} \leftarrow D_{\mathbb{Z}^n, \sigma_0^{(1)}}$ with $\sigma_0^{(1)} = n$ and $\mathbf{e}_0^{(2)} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma}}$ with $\Sigma = \sigma_0^2 \mathbf{I}_m - (\sigma_0^{(1)})^2 \mathbf{F}^\top \mathbf{F}$. The distributions of $D^{(3)}$ is identical to that of $D^{(2)}$.

In the **Hybrid 4**,

$$D^{(4)} = ((\tilde{\mathbf{A}}', \mathbf{BCs} + \mathbf{F}(\mathbf{s} + \mathbf{e}_0^{(1)}) + \mathbf{e}_0^{(2)}), \\ (b, \mathbf{x}, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}', \mathbf{BCs} + \mathbf{F}(\mathbf{s} + \mathbf{e}_0^{(1)}) + \mathbf{e}_0^{(2)}), \\ (\delta_{\mathbf{d} \in \hat{C}_{\mathbf{s}b \oplus 1, b, \mathbf{x}}} \cdot r) \oplus (I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2)).$$

According to the Lemma 6, based on the condition of $(\mathbf{s} + \mathbf{e}_0^{(1)})$, \mathbf{Cs} and the hamming weight of $I_{b,\mathbf{x}}(\mathbf{d})$ larger than $\frac{n}{4}$, the distribution of $I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2$ is within statistical distance at most $q^{\frac{3l}{2}} \cdot 2^{-\frac{n}{4}}$ to the uniform distribution over \mathbb{Z}_2 . Since $n = \lambda \cdot l \cdot \log q$, these two distributions are statistically close.

In the **Hybrid 5**,

$$D^{(5)} = ((\tilde{\mathbf{A}}', \tilde{\mathbf{A}}'\mathbf{s} + \mathbf{e}_0), \\ (b, \mathbf{x}, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}', \tilde{\mathbf{A}}'\mathbf{s} + \mathbf{e}_0), \\ (\delta_{\mathbf{d} \in \hat{C}_{\mathbf{s}b \oplus 1, b, \mathbf{x}}} \cdot r) \oplus (I_{b,\mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2)).$$

The distribution of $D^{(5)}$ is identical to that of $D^{(4)}$.

In the **Hybrid 6**,

$$\begin{aligned}
 D^{(6)} &= ((\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}_0), \\
 &\quad (b, \mathbf{x}, \mathbf{d}, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}_0), \\
 &\quad (\delta_{\mathbf{d} \in \hat{G}_{\mathbf{s}_b \oplus 1, b, \mathbf{x}}} \cdot r) \oplus (I_{b, \mathbf{x}}(\mathbf{d}) \cdot \mathbf{s} \bmod 2)).
 \end{aligned}$$

The distribution of $D^{(6)}$ is statistically closed to that of $D^{(5)}$ since the probability of aborting due to the selection of $\tilde{\mathbf{A}}$ is negligible.

Finally, the distribution of $D^{(6)}$ is computationally indistinguishable with the distribution D_1 , according to the hardness $\text{LWE}_{q, \sigma_{\mathbf{F}}}^{m, l}$ assumption. This completes the proof of the lemma. \square

5 Public Key Encryption with Secret Key Leasing from LWE with Polynomial Modulus

5.1 Our PKE-SKL Scheme Description

Here we describe our construction for PKE-SKL over classical channel with single-bit messages from NNTCF family.

Construction 1 (Parallel Repetition Version of our PKE-SKL protocol)

- $\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{sk})$:
 - let $l = \mathcal{O}(\lambda)$; $n = \omega(l \cdot \lceil \log q \rceil)$; q is a $\text{poly}(\lambda)$ -sized prime satisfying $q > 8Bm' \lceil \log q \rceil$ and $m = n \cdot \lceil \log q \rceil$. $\sigma_0 = n^{\frac{3}{2}} \sqrt{m}$, $\sigma = 150 \cdot \sigma_0 \cdot m$, $B = m \cdot (\sigma + \sigma_0) \sqrt{\lambda}$, $m'/B = \omega(n \log q)$, $N = \lambda$.
 - Run $(k_i, \text{td}_i) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$, return $k_i = (\mathbf{A}_i, \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})$ and $\text{td}_i = \mathbf{T}_{\mathbf{A}_i}$.
 - Output $(\text{mpk}, \text{sk}) = (\{k_i\}_{i \in [N]}, \{\text{td}_i\}_{i \in [N]})$.
- $\text{KeyGen}(\text{mpk}) \rightarrow (\rho_{\text{sk}}, \text{pk})$:
 - Take in $\text{mpk} = \{(\mathbf{A}_i, \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})\}_{i=1}^N$. Run $\text{SAMP}_{\mathcal{F}}(k_i, \cdot)$ on a uniform superposition of b_i 's, to obtain the state

$$\bigotimes_{i=1}^N \frac{1}{\sqrt{2q^n}} \sum_{\substack{b_i \in \{0,1\}, \\ \mathbf{x}_i \in \mathbb{Z}_q^n, \\ \mathbf{e}_i \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e}_i) |b_i, \mathbf{x}_i\rangle | \mathbf{e}_i\rangle},$$

then compute the following state:

$$\bigotimes_{i=1}^N \frac{1}{\sqrt{2q^n}} \sum_{\substack{b_i \in \{0,1\}, \\ \mathbf{x}_i \in \mathbb{Z}_q^n, \\ \mathbf{e}_i \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e}_i) |b_i, \mathbf{x}_i\rangle | \mathbf{e}_i + \mathbf{A}_i \mathbf{x}_i + b_i \cdot (\mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})\rangle},$$

- Measure the last register to obtain $\mathbf{y}_i = \mathbf{A}_i \mathbf{x}'_i + \mathbf{e}'_i$, where $\mathbf{x}'_i = \mathbf{x}'_{i,b_i} + b_i \mathbf{s}_i$, $\mathbf{e}'_i = \mathbf{e}_i + b_i \cdot \mathbf{e}_{0,i}$. The resulting post-measurement state constitutes the quantum decryption key:

$$\rho_{sk} = \bigotimes_{i=1}^N \frac{1}{\sqrt{2}} \sum_{b_i \in \{0,1\}} p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i) |b_i, \mathbf{x}'_{i,b_i}\rangle,$$

where $\mathbf{x}'_{i,b_i} = \mathbf{x}'_i - b_i \mathbf{s}_i$ and the value of $p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i)$ satisfying:

- 1) $p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) = \frac{1}{\sqrt{1+e^{(2\pi \langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \pi \|\mathbf{e}_{0,i}\|^2)/\sigma^2}}}$,
 $p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) = \frac{e^{(\pi \langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \frac{\pi}{2} \|\mathbf{e}_{0,i}\|^2)/\sigma^2}}{\sqrt{1+e^{(2\pi \langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \pi \|\mathbf{e}_{0,i}\|^2)/\sigma^2}}}$ if $\mathbf{e}'_i \in \mathcal{S}_0 \cap \mathcal{S}_1 \cap \mathbb{Z}^m$;
 - 2) $p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 0$, $p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 1$ if $\mathbf{e}'_i \in (\mathcal{S}_0 \setminus \mathcal{S}_1) \cap \mathbb{Z}^m$;
 - 3) $p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 1$, $p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 0$ if $\mathbf{e}'_i \in (\mathcal{S}_1 \setminus \mathcal{S}_0) \cap \mathbb{Z}^m$,
- where $\mathcal{S}_0 = \mathcal{B}(0, \sigma\sqrt{m})$ and $\mathcal{S}_1 = \mathcal{B}(\mathbf{e}_0, \sigma\sqrt{m})$.

- Output public key $\text{pk} = \{(\mathbf{A}_i, \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i}, \mathbf{y}_i)\}_{i \in [N]}$ and quantum decryption key ρ_{sk} .
- $\text{Enc}(\text{pk}, \mu) \rightarrow \text{ct}$:
- Take in a message $\mu \in \{0, 1\}$. Select $\mathbf{U}_i \xleftarrow{\$} \{0, 1\}^{m' \times m}$, $\mathbf{R} \leftarrow \{0, 1\}^{m' \times m'}$ and $\hat{\mathbf{e}}_{1,i} \xleftarrow{\$} [-B, B]^{m'}$, where $B = (\sigma + \sigma_0)\sqrt{\lambda} \cdot m$.
 - Let $\hat{\mathbf{e}}_1 = \sum_{i \in [N]} \hat{\mathbf{e}}_{1,i}$. The algorithm computes ciphertexts as follows:

$$\text{ct} = \mathbf{R}\mathbf{A} + \mathbf{R}\mathbf{E}_1 + \mu \cdot \mathbf{G}_{m', N(n+1)+1},$$

where $\mathbf{A} \in \mathbb{Z}_q^{m' \times N(n+1)+1}$ is given as:

$$\mathbf{A} = (\mathbf{U}_1(\mathbf{A}_1 \mathbf{s}_1 + \mathbf{e}_{0,1}) | \mathbf{U}_1 \mathbf{A}_1 \cdots \mathbf{U}_N(\mathbf{A}_N \mathbf{s}_N + \mathbf{e}_{0,N}) | \mathbf{U}_N \mathbf{A}_N | \sum_{i \in [N]} \mathbf{U}_i \mathbf{y}_i),$$

$\mathbf{E}_1 \in \mathbb{Z}_q^{m' \times N(n+1)+1}$ is a matrix with all columns $\mathbf{0}^{m'}$ except the last column which equals $\hat{\mathbf{e}}_1$. $\mathbf{G}_{m', N(n+1)+1}$ is the Gadget matrix.

- Output ciphertext ct .
- $\text{Add}(\text{ct}_1, \text{ct}_2) \rightarrow \text{ct}_{Add}$: On input two ciphertexts ct_1, ct_2 , output $\text{ct}_{Add} = \text{ct}_1 + \text{ct}_2$.
- $\text{Mult}(\text{ct}_1, \text{ct}_2) \rightarrow \text{ct}_{Mult}$: On input two ciphertexts ct_1, ct_2 , output $\text{ct}_{Mult} = \mathbf{G}^{-1}(\text{ct}_1) \cdot \text{ct}_2$.
- $\text{Dec}(\rho_{sk}, \text{ct}) \rightarrow (\mu', \rho'_{sk})$:
- On the input quantum decryption key ρ_{sk} and ciphertext ct , run decryption algorithm in a coherent way as follows:

$$\left(\bigotimes_{i=1}^N \frac{1}{\sqrt{2}} \sum_{b_i \in \{0,1\}} p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i) |b_i, \mathbf{x}'_{b_i,i}\rangle \right) | \underbrace{\mathbf{v}_{inv} \cdot \text{ct} \cdot \mathbf{v}_{sk}}_{y'} | \lceil y' / \lfloor q/2 \rfloor \rceil \rangle, \quad (11)$$

where $\mathbf{v}_{inv} = \mathbf{G}^{-1}(\mathbf{0}^{1 \times (N(n+1))} | \lfloor \frac{q}{2} \rfloor)$ and $\mathbf{v}_{sk} = (-b_1, -(\mathbf{x}'_{b_1,1})^\top, \dots, -b_N, -(\mathbf{x}'_{b_N,N})^\top, 1)^\top$ is a column vector where b_i 's, $\mathbf{x}'_{b_i,i}$'s for $\forall i \in [N]$ are from the corresponding registers in the secret key ρ_{sk} .

- Measure the last register to obtain μ' . Uncompute the register $|\mathbf{y}'\rangle$ with the ciphertexts and b_i, \mathbf{x}'_{b_i} . Then the first N registers consist of ρ'_{sk} . Ideally, $\rho'_{sk} = \rho_{sk}$ holds.
- Del(ρ_{sk}) \rightarrow cert:
- Take in the ρ_{sk} , perform Hadamard operations and obtain

$$|\psi\rangle = \bigotimes_{i=1}^N 2^{-\frac{n \cdot \lceil \log(q) \rceil + 2}{2}} \sum_{\substack{\mathbf{d}_i \in \{0,1\}^{n \cdot \lceil \log(q) \rceil}, \\ b_i \in \{0,1\}, \\ u_i \in \{0,1\}}} (-1)^{\mathbf{d}_i \cdot \mathcal{J}(\mathbf{x}'_{b_i,i}) \oplus u_i b_i} p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i) |u_i\rangle |\mathbf{d}_i\rangle.$$

- Measure this quantum state, thereby resulting in cert = $\{(u_i, \mathbf{d}_i)\}_{i=1}^N \in (\mathbb{Z}_2 \times \mathbb{Z}_2^{n \cdot \lceil \log(q) \rceil})^N$ as the deletion certificate.
- VerDel(sk, pk, cert) \rightarrow \top/\perp :
- Compute $\mathbf{x}'_{b',i} \leftarrow \text{INV}_{\mathcal{F}}(\mathbf{T}_{\mathbf{A}_i}, b_i, \mathbf{y}_i)$ for all $i \in [N]$ and both $b' \in \{0, 1\}$.
 - Check if $\|\mathbf{y}_i - \mathbf{A}_i \mathbf{x}'_{b',i} - b' \cdot \mathbf{A}_i \mathbf{s}_i\|_2 \leq (\sigma + \sigma_0)\sqrt{m}$ for all $i \in [N]$ and $b' \in \{0, 1\}$. If not, output invalid \perp . If yes, continue.
 - Check if $\mathbf{d}_i \in G_{k_i,0,\mathbf{x}'_{0,i}} \cap G_{k_i,1,\mathbf{x}'_{1,i}}$ and $u_i = \mathbf{d}_i^\top \cdot (\mathcal{J}(\mathbf{x}'_{0,i}) \oplus \mathcal{J}(\mathbf{x}'_{1,i})) \bmod 2$. Count the number of the i that passes the checking step and denote this number as N' . If $N' > 0.78N$, output valid \top . Otherwise, output invalid \perp .

The completeness of our protocol is given as follows.

Theorem 5. *The PKE-SKL scheme with classical lessor described in Construction 1 satisfies the correctness property of the PKE-SKL Definition.*

The Theorem 5 follows immediately from the following Lemmas 7 and 8.

Lemma 7 (Correctness of Decryption). *The algorithm Dec in PKE-SKL Construction 1 satisfies decryption correctness, namely,*

$$\Pr \left[\begin{array}{c} (\text{mpk}, \text{td}) \leftarrow \text{Setup}(1^\lambda) \\ \text{Dec}(\rho_{sk}, \text{ct}) = \mu : (\text{pk}, \rho_{sk}) \leftarrow \text{KeyGen}(\text{mpk}) \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, \mu) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Lemma 8 (Correctness of Verifying Deletion). *For $m > 500$, honestly prepared $\{\mathbf{y}_i\}_{i \in [N]}$ and secret key ρ_{sk} , the probability passing the algorithm VerDel is overwhelming.*

To prove the security of our PKE-SKL scheme, we need to show that the lessee should not have any noticeable advantage to distinguish between ciphertexts of messages 0 and 1, after submitting the deletion certificate.

Theorem 6 (Security of our PKE-SKL). *For $\sigma_0 = n^{\frac{3}{2}}\sqrt{m}$, assuming the post-quantum hardness $\text{LWE}_{n,m,q,\sigma_0}$ with polynomial modulus, the Construction 1 satisfies strong γ -SKL security for any noticeable γ .*

Please refer to the full version of this paper for proof.

6 Polynomial-Sized Proof of Quantumness

In this section, we present an NNTCF-based proof of quantumness protocol based on the polynomial hardness of LWE problem without reliance on a random oracle or Bell's inequality. Our proof of quantumness protocol can be viewed as an improved version of the work in [BCM+18], where the soundness is directly guaranteed by the AHB property of the NNTCF.

Let \mathcal{P} denote a quantum prover and \mathcal{V} denote a classical verifier, our proof of quantumness protocol is given in Construction 2.

Construction 2 (Proof of Quantumness based on LWE-based NNTCF)

1. **Setup**(1^λ): Fix a security parameter λ and the NNTCF family \mathcal{F} described by algorithms $(\text{GEN}_{\mathcal{F}}, \text{SAMP}_{\mathcal{F}}, \text{INV}_{\mathcal{F}}, \text{CHK}_{\mathcal{F}})$, assuming the polynomial hardness of LWE. Set $l = \mathcal{O}(\lambda)$, $n = \omega(l \lceil \log q \rceil)$, $m \geq n \cdot \lceil \log q \rceil$ and $m > 500$; $w = n \lceil \log q \rceil$, $q \geq 8\sigma\sqrt{m}$ a prime, $\sigma_0 \geq n^{\frac{3}{2}}\sqrt{m}$, $150 \cdot m \cdot \sigma_0 \leq \sigma \leq \frac{q}{C_T \sqrt{mn \log q}}$. $N = \lambda$. Output $\text{pp} = (n, m, w, q, \sigma_0, \sigma)$.
2. For $i = 1, \dots, N$,
 - i.1. \mathcal{V} : On input the parameters pp , the verifier runs $(k_i = (\mathbf{A}_i, \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i}), \mathbf{T}_{\mathbf{A}_i}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^m, 1^n, \sigma_0, q)$, where $\mathbf{e}_{0,i} \leftarrow D_{\mathbb{Z}^m, \sigma_0, \sigma_0 \sqrt{m}}$, sends k to the prover and keeps the trapdoor $\mathbf{T}_{\mathbf{A}_i}$ private.
 - i.2. \mathcal{P} : On receive the key $k_i = (\mathbf{A}_i, \mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})$, the prover will do the following steps:
 - * Generate the following quantum state:

$$\frac{1}{\sqrt{2q^n}} \sum_{\substack{b_i \in \{0,1\}, \\ \mathbf{x}_i \in \mathbb{Z}_q^n, \\ \mathbf{e}_i \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e}_i)} |b_i, \mathbf{x}_i\rangle |\mathbf{e}_i\rangle$$

and then compute with the key k_i as below:

$$\frac{1}{\sqrt{2q^n}} \sum_{\substack{b_i \in \{0,1\}, \\ \mathbf{x}_i \in \mathbb{Z}_q^n, \\ \mathbf{e}_i \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}^m, \sigma, 2\sigma\sqrt{m}}(\mathbf{e}_i)} |b_i, \mathbf{x}_i\rangle |\mathbf{e}_i + \mathbf{A}_i \mathbf{x}_i + b_i \cdot (\mathbf{A}_i \mathbf{s}_i + \mathbf{e}_{0,i})\rangle, \quad (12)$$

- * Measure the last register to obtain $\mathbf{y}_i = \mathbf{A}_i \mathbf{x}'_i + \mathbf{e}'_i$, where $\mathbf{x}'_{i,b_i} = \mathbf{x}'_i - b_i \mathbf{s}_i$ and $\mathbf{x}'_i = \mathbf{x}'_{i,0} + b_i \mathbf{s}_i$, $\mathbf{e}'_i = \mathbf{e}_i + b_i \cdot \mathbf{e}_{0,i}$ for some fixed \mathbf{e}_i . The resulting post-measurement state is:

$$|\varphi_i\rangle = \frac{1}{\sqrt{2}} \sum_{b_i \in \{0,1\}} p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i) |b_i, \mathbf{x}'_{i,b_i}\rangle,$$

where the value of $p_{b_i}(\mathbf{e}_{0,i}, \mathbf{e}'_i)$ satisfying:

$$\begin{aligned}
1) \quad p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) &= \frac{1}{\sqrt{1+e^{(2\pi\langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \pi\|\mathbf{e}_{0,i}\|^2)/\sigma^2}}}, \\
p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) &= \frac{e^{(\pi\langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \frac{\pi}{2}\|\mathbf{e}_{0,i}\|^2)/\sigma^2}}{\sqrt{1+e^{(2\pi\langle \mathbf{e}'_i, \mathbf{e}_{0,i} \rangle - \pi\|\mathbf{e}_{0,i}\|^2)/\sigma^2}}} \text{ if } \mathbf{e}'_i \in \mathcal{S}_0 \cap \mathcal{S}_1 \cap \mathbb{Z}^m; \\
2) \quad p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) &= 0, \quad p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 1 \text{ if } \mathbf{e}'_i \in (\mathcal{S}_0 \setminus \mathcal{S}_1) \cap \mathbb{Z}^m; \\
3) \quad p_0(\mathbf{e}_{0,i}, \mathbf{e}'_i) &= 1, \quad p_1(\mathbf{e}_{0,i}, \mathbf{e}'_i) = 0 \text{ if } \mathbf{e}'_i \in (\mathcal{S}_1 \setminus \mathcal{S}_0) \cap \mathbb{Z}^m, \\
\text{where } \mathcal{S}_0 &= \mathcal{B}(0, \sigma\sqrt{m}) \text{ and } \mathcal{S}_1 = \mathcal{B}(\mathbf{e}_0, \sigma\sqrt{m}).
\end{aligned}$$

* Output the string \mathbf{y}_i .

i.3. \mathcal{V} : Reply with a uniformly random challenge bit $c_i \xleftarrow{\$} \{0, 1\}$.

i.4. \mathcal{P} : Take in the challenge c_i , do the following tests:

* Preimage test (if $c_i = 0$): Perform a standard basis measurement onto $|\varphi_i\rangle$, return a pair $(b_i, \mathbf{x}'_{i,b_i})$ as the proof σ_{c_i} .

* Equation test (if $c_i = 1$): Perform a Hadamard basis measurement onto $|\varphi_i\rangle$, return a pair (u_i, \mathbf{d}_i) as the proof σ_{c_i} .

3. \mathcal{V} : Take in $\{\mathbf{T}_{\mathbf{A}_i}, \mathbf{y}_i, c_i, \sigma_{c_i}\}_{i \in [N]}$, do the following steps:

- Compute $\mathbf{x}'_{b',i} \leftarrow \text{INV}_{\mathcal{F}}(\mathbf{T}_{\mathbf{A}_i}, b_i, \mathbf{y}_i)$ for all $i \in [N]$ and both $b' \in \{0, 1\}$.
- When $c_i = 0$, check if $\text{CHK}_{\mathcal{F}_{\text{LWE}}}(k_i, b_i, \mathbf{x}'_{i,b_i}, \mathbf{y}_i) = 1$ holds for all $i \in [N]$.
- When $c_i = 1$, check if $\mathbf{d}_i \in G_{k_i,0,\mathbf{x}'_{0,i}} \cap G_{k_i,1,\mathbf{x}'_{1,i}}$ and $u_i = \mathbf{d}_i^T \cdot (\mathcal{J}(\mathbf{x}'_{0,i}) \oplus \mathcal{J}(\mathbf{x}'_{1,i})) \bmod 2$.
- Count the number of i 's that $c_i = j$ for $j \in \{0, 1\}$ and denote this number as N_j . Count the number of the i 's that pass the Equation tests and denote this number as N' . If $N_0 > \frac{1}{4}N, N_1 \geq \frac{1}{4}N$ and $N' > 0.75N_1$, output valid \top . Otherwise, output invalid \perp .

Based on the NNTCF, the correctness and soundness of our protocol are given as follows.

Theorem 7 (Correctness of Our Proof of Quantumness). *Let $\lambda \in \mathbb{N}$ be the security parameter. A QPT prover \mathcal{P} , following the honest strategy in the Construction 2, is accepted with probability $1 - \text{negl}(\lambda)$.*

Theorem 8 (Soundness of Our Proof of Quantumness). *Based on the adaptive hardcore bit property of the NNTCF family \mathcal{F} , the probability for any classical $\tilde{\mathcal{P}}$ to pass the verification in the Construction 2 is negligible.*

Please refer to the full version of this paper for proof.

Acknowledgments. This work was supported in part by the National Key Research and Development Program of China (No. 2022YFB2702700), by the European Union's Horizon Europe research and innovation programme under the project "Quantum Secure Networks Partnership" (QSNP, No. 101114043), and by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute. We thank the anonymous ASIACRYPT referees for corrections and suggestions that improved the presentation of the paper.

References

- AKN+23. Shweta Agrawal, Fuyuki Kitagawa, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Public key encryption with secure key leasing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 581–610. Springer, 2023.

- ALP20. Prabhanjan Ananth and Rolando L La Placa. Secure quantum extraction protocols. In *Theory of Cryptography Conference*, pages 123–152. Springer, 2020.
- ALP21. Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, pages 501–530, Cham, 2021. Springer International Publishing.
- APV23. Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. Revocable cryptography from learning with errors. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 93–122, Cham, 2023. Springer Nature Switzerland.
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.
- BCM+18. Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331. IEEE, 2018.
- BD20. Zvika Brakerski and Nico Döttling. Hardness of lwe on general entropic distributions. In *Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II 30*, pages 551–575. Springer, 2020.
- BGKM+23. Zvika Brakerski, Alexandru Gheorghiu, Gregory D Kahanamoku-Meyer, Eitan Porat, and Thomas Vidick. Simple tests of quantumness also certify qubits. In *Annual International Cryptology Conference*, pages 162–191. Springer, 2023.
- BKV20. Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler proofs of quantumness. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography*, 2020.
- CGJL23. Orestis Chardouvelis, Vipul Goyal, Aayush Jain, and Jiahui Liu. Quantum key leasing for pke and fhe with a classical lessor. *arXiv preprint arXiv:2310.14328*, 2023.
- CGV22. Andrea Coladangelo, Shafi Goldwasser, and Umesh Vazirani. Deniable encryption in a quantum world. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1378–1391, 2022.
- CHV23. Céline Chevalier, Paul Hermouet, and Quoc-Huy Vu. Semi-quantum copy-protection and more. In *Theory of Cryptography Conference*, pages 155–182. Springer, 2023.
- GMP23. Alexandru Gheorghiu, Tony Metger, and Alexander Poremba. Quantum cryptography with classical communication: Parallel remote state preparation for copy-protection, verification, and more. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
- GV19. Alexandru Gheorghiu and Thomas Vidick. Computationally-secure and composable remote state preparation. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1024–1033. IEEE, 2019.
- HMNY21. Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum encryption with certified deletion, revisited: Public key, attribute-based, and classical communication. In *Advances in*

- Cryptography–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part I 27*, pages 606–636. Springer, 2021.
- KLVI23. Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1617–1628, 2023.
- KMCVY22. Gregory D Kahanamoku-Meyer, Soonwon Choi, Umesh V Vazirani, and Norman Y Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18(8):918–924, 2022.
- KNY21. Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Secure software leasing from standard assumptions. In *Theory of Cryptography Conference*, pages 31–61. Springer, 2021.
- Mah18a. Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 332–338. IEEE, 2018.
- Mah18b. Urmila Mahadev. Classical verification of quantum computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267. IEEE, 2018.
- MPY23. Tomoyuki Morimae, Alexander Poremba, and Takashi Yamakawa. Revocable quantum digital signatures. *arXiv preprint arXiv:2312.13561*, 2023.
- RS19. Roy Radian and Or Sattath. Semi-quantum money. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies, AFT '19*, page 132-146. Association for Computing Machinery, 2019.
- Shm22. Omri Shmueli. Public-key quantum money with a classical bank. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 790–803, 2022.
- YZ22. Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 69–74. IEEE, 2022.



Mind the Bad Norms

Revisiting Compressed Oracle-Based Quantum Indistinguishability Proofs

Ritam Bhaumik¹✉, Benoît Cogliati², Jordan Ethan³,
and Ashwin Jha⁴

¹ TII, Abu Dhabi, UAE

bhaumik.ritam@gmail.com

² Thales DIS France SAS, Meudon, France

³ CISA Helmholtz Center for Information Security, Saarbrücken, Germany
jordan.ethan@cispa.de

⁴ Ruhr-Universität Bochum, Bochum, Germany

Abstract. In this work, we revisit the Hosoyamada-Iwata (HI) proof for the quantum CPA security of the 4-round Luby-Rackoff construction and identify a gap that appears to undermine the security proof. We emphasize that this is *not* an attack, and the construction may still achieve the claimed security level. However, this gap raises concerns about the feasibility of establishing a formal security proof for the 4-round Luby-Rackoff construction. In fact, the issue persists even if the number of rounds is increased arbitrarily. On a positive note, we restore the security of the 4-round Luby-Rackoff construction in the *non-adaptive* setting, achieving security up to $2^{n/6}$ superposition queries. Furthermore, we establish the quantum CPA security of the 4-round MistyR and 5-round MistyL constructions, up to $2^{n/5}$ and $2^{n/7}$ superposition queries, respectively, where n denotes the size of the underlying permutation.

Keywords: quantum security · compressed oracle · recording standard oracle with errors · Luby-Rackoff · Misty

1 Introduction

Quantum Security. In symmetric cryptography, it is generally admitted that a doubling of the key length would be sufficient to deter the threat of quantum computers. Indeed, this corresponds to the lowered cost of exhaustive search from Grover’s algorithm. However, in recent years a plethora of results (see for instance [8–12, 20, 26, 28–31]) have shown that this view is too simplistic, and that

The authors would like to thank Akinori Hosoyamada and Tetsu Iwata for their comments on a precursor note that evolved into this paper. Jordan Ethan’s research was conducted within the framework of the French-German Center for Cybersecurity, a collaboration between CISA and LORIA. Ashwin Jha’s work was supported by the German Research Foundation (DFG) within the framework of the Excellence Strategy of the Federal Government and the States – EXC 2092 CASA – 39078197.

© International Association for Cryptologic Research 2025

K.-M. Chung and Y. Sasaki (Eds.): ASIACRYPT 2024, LNCS 15492, pp. 215–247, 2025.

https://doi.org/10.1007/978-981-96-0947-5_8

more efficient distinguishers can be created. This highlights the need to study whether existing security proofs for generic constructions and modes of operation can be extended to the quantum setting, which has received a considerable focus in a series of recent works [4, 5, 7, 16, 23–25, 27, 43, 44].

Pseudorandom Functions and Permutations. Classically, most of the well-known symmetric cryptographic algorithms are constructed as a mode of operation over fixed length primitives that are instantiated with either a pseudorandom¹ permutation (PRP) or function (PRF).

Some well-known examples of generic PRP constructions include the Luby-Rackoff cipher [33], Lai-Massey [32] and the generic Misty ciphers [35, 37]. Of these the former two constructions can be instantiated by any primitive (function or permutation), while the latter solely works with permutations. In general, PRP-based constructions are preferred as they can be directly instantiated with well-analyzed block ciphers. On the other hand, PRF based constructions are usually easier to analyze in security proofs. Indeed, many security proofs involve the boilerplate switching lemma [3, 42]: replace PRP calls with PRF calls with a factor of $O(q^2/2^n)$ per call, where q and n denote the number of queries and output size, respectively. Thus, all of the above mentioned constructions are classically secure birthday-bound PRFs. On the other hand more recent efforts have focused on building beyond-the-birthday-bound secure PRP-to-PRF constructions, starting with the well-known sum of permutations [2, 21] and the truncation of permutation [21] to the more recent encrypted Davis-Meyer [15] and its dual [36]. The analysis of these PRP and PRF constructions lead to a great advancement in the provable security research, mushrooming several new proof techniques such as the H-coefficient technique [22, 40], mirror theory [14, 39, 41] the χ^2 -technique [17], and the recent use of Fourier analysis [18] to prove the exact security of sum of permutations.

The Compressed Oracle. In the quantum setting, however, the research on the security of these well-known constructions is still in the rudimentary stage. While there are some generic attacks on Luby-Rackoff [23, 30] and Misty [19], on the security proofs front the results are still far from tight even in the birthday-bound² regime. Having said that, the situation has changed in recent years, largely due to Zhandry’s compressed oracle technique [44]—an elegant way to lazy sample a random function. Indeed most recent security proofs [5, 23–25] in symmetric cryptography relied on the compressed oracle [44] and its variants respectively introduced by Hosoyamada and Iwata [23] and Chung *et al.* [13].

When proving the indistinguishability of a construction C based on PRFs from a true random function, the proof typically follows these steps:

¹ the fixed-length permutation /function is keyed, efficiently computable, and indistinguishable from a uniform random permutation/function.

² Note that, in the quantum setting birthday-bound refers to the cube-root of the output size.

- Model the random function as a construction with a structure similar to C , but with some of the inputs augmented with adversarial queries to ensure the uniqueness of inputs, thereby guaranteeing the uniformity of outputs.
- Identify “bad events” that occur when the output of intermediate function calls leads to input collisions in subsequent calls.
- Upper-bound the probability of such bad events occurring.
- Establish a one-to-one mapping between intermediate values in both constructions, assuming no bad event has occurred.

It is important to note that ensuring these bad events are described only using inputs and outputs recorded by the compressed oracle is critical to the proof. In particular, certain information may be lost in this process, such as the specific adversarial query or the relationship between input-output pairs belonging to the same query.

1.1 Our Contribution

Our contribution is three-fold. Firstly, we identify some critical issues in some of the previous works in this direction. They relate to the aforementioned one-to-one mapping: most notably, in the 4-round Luby-Rackoff security proof [23], the authors cannot prevent bad collisions without relying on information that is not present in the compressed oracle entries. We also spotted similar flaws in [5, 25, 34].

Secondly, we propose a new security proof for the 4-round Luby-Rackoff construction in the non-adaptive chosen plaintext attack setting: the adversary has to prepare all of its queries in advance, and receive the corresponding outputs at once. By using an artificial dummy database call on all the adversary’s inputs, this allows us to mitigate the issue from [23], since now the database contains all the necessary information to handle the bad events.

Finally, we prove the security of Misty schemes in the quantum setting using the two-domain framework from [5]. In more details, we prove that the 4-round MistyR (resp. 5-round MistyL) construction is secure up to $2^{n/5}$ (resp. $2^{n/7}$) chosen plaintext queries, where n denotes the size of the underlying permutation. We note that, in both cases, this corresponds to the minimum number of rounds to achieve an exponential bound in n , since period-finding attacks based on Simon’s algorithm exist for the 3-round MistyR (resp. 4-round MistyL) constructions [19].

2 Quantum Computing

Throughout, we assume familiarity with the fundamentals of finite dimensional linear algebra and Quantum computing. A comprehensive exposition on these subjects is given in [1, 38]. In this section, we introduce some notation we use later in the paper; an introductory overview of the relevant notions is also available in the full version of this paper [6].

2.1 General Notation

The set of all binary strings, including the empty string ε , is denoted $\{0, 1\}^*$. For some $x, y \in \{0, 1\}^*$, $x\|y$ denotes the concatenation of x and y . For some positive integer m , $[m]$ denotes the set $\{1, \dots, m\}$, and $\{0, 1\}^m$ denotes the set of all m -bit binary strings.

We use the standard Dirac notations. $\langle \cdot | \cdot \rangle$ denotes the inner product over a k -dimensional Hilbert space $\mathcal{H} := \mathbb{C}^k$, and $\|\cdot\|$ denotes the norm. Given an orthonormal basis \mathbf{B} of \mathcal{H} , we sometimes write $\mathcal{H}[\mathbf{B}]$ to emphasize the basis representation of \mathcal{H} . $\mathbf{U}[\mathcal{H}]$ will denote the set of all unitaries on \mathcal{H} . $\text{Tr}(\mathbf{L})$ will denote the trace of a linear operator \mathbf{L} . $\text{Tr}_{\mathcal{H}_1}(\mathbf{L})$ will denote the partial trace on \mathcal{H}_1 of a linear operator \mathbf{L} over the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$. $\mathbf{D}(\mathcal{H})$ will denote the set of all density operators of \mathcal{H} . $\|\mathbf{L}\|_1$ will denote the trace norm of \mathbf{L} .

2.2 Quantum (Non-adaptive) Oracle-Algorithms

In what follows, we define $\mathcal{H}_{in} := \mathbb{C}^{2^m}$, $\mathcal{H}_{out} := \mathbb{C}^{2^n}$. Let \mathcal{H}_{work} and \mathcal{H}_{state} be two finite dimensional complex Hilbert spaces.

Any function $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ can be realized by the unitary mapping $|x, y\rangle$ to $|x, y \oplus f(x)\rangle$ on $\mathcal{H}_{in} \otimes \mathcal{H}_{out}$. Indeed, the oracle access to f , denoted \mathbf{O}_f , is represented by this standard unitary

$$\mathbf{O}_f|x, y\rangle \mapsto |x, y \oplus f(x)\rangle,$$

on the space $\mathcal{H}_{in} \otimes \mathcal{H}_{out}$. To represent a stateful oracle, we simply bestow additional qubits to represent the oracle state. Formally, we define

$$\mathbf{O}_f|x, y, s\rangle \mapsto |x, y + f(x), s'\rangle,$$

on the product space $\mathcal{H}_{\mathbf{O}_f} := \mathcal{H}_{in} \otimes \mathcal{H}_{out} \otimes \mathcal{H}_{state}$, where $\{|x, y, s\rangle\}$ denotes the computational basis of $\mathcal{H}_{\mathbf{O}_f}$. The oracle state space \mathcal{H}_{state} into $\mathcal{H}_{db} \otimes \mathcal{H}_{aux}$, where \mathcal{H}_{db} denotes the internal state which is (possibly transient) and persistent across queries, and \mathcal{H}_{aux} denotes the state space of any ancillary qubits required to compute the function itself. As ancillary qubits are always reset after each query, it is convenient to focus solely on the former (the useful *state*) while disregarding the latter (the ancillary qubits). Indeed, we often drop \mathcal{H}_{aux} from the description and simply consider \mathcal{H}_{db} as the oracle state space.

For any quantum *oracle-algorithm* A that makes q black-box queries to a (possibly stateful) oracle \mathbf{O}_f , we define the interactive game $A^{\mathbf{O}_f}$ to be the sequence of $2q + 1$ unitaries: $\mathbf{U}_q \mathbf{O}_f \dots \mathbf{U}_1 \mathbf{O}_f \mathbf{U}_0$ over the product space $\mathcal{H}_{A^{\mathbf{O}_f}} = \mathcal{H}_{in} \otimes \mathcal{H}_{out} \otimes \mathcal{H}_{work} \otimes \mathcal{H}_{state}$, where it is implicitly understood that \mathbf{U}_i 's operate on $\mathcal{H}_A = \mathcal{H}_{in} \otimes \mathcal{H}_{out} \otimes \mathcal{H}_{work}$ and \mathbf{O}_f operates on $\mathcal{H}_{\mathbf{O}_f}$.

We write $A^{\mathbf{O}_f}[\rho_A \otimes \rho_{\mathbf{O}_f}] = b$ to denote the event that the oracle-aided algorithm A outputs b after making q queries to oracle \mathbf{O}_f , where A and \mathbf{O}_f are initialized in $\rho_A \in \mathbf{D}(\mathcal{H}_A)$ and $\rho_{\mathbf{O}_f} \in \mathbf{D}(\mathcal{H}_{state})$, or jointly as $\rho_{A, \mathbf{O}_f}^0 := \rho_A \otimes \rho_{\mathbf{O}_f}$.

Capturing Non-adaptivity. For any oracle-algorithm A that makes q non-adaptive queries to \mathbf{O}_f , we define the non-adaptive interactive game $A^{\mathbf{O}_f^{\otimes q}}$ to be the unitary $\mathbf{U}_1 \mathbf{O}_f^{\otimes q} \mathbf{U}_0$ on the product space $\mathcal{H}_{in}^{\otimes q} \otimes \mathcal{H}_{out}^{\otimes q} \otimes \mathcal{H}_{work} \otimes \mathcal{H}_{state}$ where it is implicitly understood that $\mathbf{O}_f^{\otimes q}$ operates on $\mathcal{H}_{in}^{\otimes q} \otimes \mathcal{H}_{out}^{\otimes q} \times \mathcal{H}_{state}$, while \mathbf{U}_0 and \mathbf{U}_1 operates on $\mathcal{H}_{in}^{\otimes q} \otimes \mathcal{H}_{out}^{\otimes q} \otimes \mathcal{H}_{work}$.

Indeed the above formalism is analogous to the classical setting, where the non-adaptive algorithm makes all q queries, $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_q) \in (\{0, 1\}^m)^q$, together and receives all q responses, $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_q) \in (\{0, 1\}^n)^q$, together from the oracle. Analogously, in the quantum setting, we have

$$\mathbf{O}_f^{\otimes q} |\mathbf{x}, \mathbf{y}, s\rangle = |\mathbf{x}, \mathbf{y} + f(\mathbf{x}), s'\rangle,$$

where $f(\mathbf{x}) = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_q))$ is simply the pointwise application of f on \mathbf{x} .

2.3 Quantum Distinguishing Games

For any two quantum oracles \mathbf{I} and \mathbf{R} , we define the distinguishing advantage of any quantum distinguisher³ A by

$$\mathbf{Adv}_{\mathbf{I};\mathbf{R}}^{\text{dist}}(A) := |\Pr(A^{\mathbf{I}}[\rho_{A,\mathbf{I}}^0] = 1) - \Pr(A^{\mathbf{R}}[\rho_{A,\mathbf{R}}^0] = 1)|,$$

where $\rho_{A,\mathbf{I}}^0$ and $\rho_{A,\mathbf{R}}^0$ denote the initial state of $A^{\mathbf{I}}$ and $A^{\mathbf{R}}$, respectively.

The Computationally Unbounded Case. For any computationally-unbounded A , it is well known that

$$\mathbf{Adv}_{\mathbf{I};\mathbf{R}}^{\text{dist}}(A) \leq \frac{1}{2} \|\text{Tr}_{\mathcal{H}_{\text{db}}}(\rho_{A,\mathbf{I}}^q) - \text{Tr}_{\mathcal{H}_{\text{db}}}(\rho_{A,\mathbf{R}}^q)\|_1,$$

where $\rho_{A,\mathbf{O}}^q := A^{\mathbf{O}} \rho_{A,\mathbf{O}} A^{\mathbf{O}\dagger}$ is the state after q queries to the oracle at-hand $\mathbf{O} \in \{\mathbf{I}, \mathbf{R}\}$. In addition, without loss of generality, we can assume A to be deterministic, and thus, define the initial state of A , $\rho_A = |\psi_A\rangle\langle\psi_A|$ for some fixed unit vector $|\psi_A\rangle \in \mathcal{H}_A$.

The Quantum IND-CPA Game. Let $F = \{F_K : \{0, 1\}^m \rightarrow \{0, 1\}^n\}_{K \in \mathcal{K}}$ be a family of functions. The IND-qCPA advantage of some distinguisher A against F is defined as

$$\mathbf{Adv}_F^{\text{qcpa}}(A) := \mathbf{Adv}_{\mathbf{O}_{F_K}; \mathbf{O}_f}^{\text{dist}}(A), \tag{1}$$

where K is uniformly distributed over \mathcal{K} , and $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is a uniform random function.

For a non-adaptive distinguisher A , the non-adaptive IND-qCPA advantage is defined analogously as:

$$\mathbf{Adv}_F^{\text{qncpa}}(A) := \mathbf{Adv}_{\mathbf{O}_{F_K}^{\otimes q}; \mathbf{O}_f^{\otimes q}}^{\text{dist}}(A), \tag{2}$$

³ An oracle-algorithm with binary output.

3 Zhandry’s Compressed Oracle

In [44], Zhandry proposed an elegant solution to implement a restricted form of lazy sampling for quantum random oracle, or simply a uniform random function $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$. We will largely follow the Chung-Fehr-Huang-Liao (CFHL) interpretation [13] of the compressed oracle, and its refinement by Bhaumik-Cogliati-Ethan-Jha (BCEJ) [5].

3.1 The Chung-Fehr-Huang-Liao Interpretation

Let \mathcal{Y} denote $\{0, 1\}^n$ and define $\mathcal{C}_{\mathcal{Y}}$ to be the computational basis of the n -qubit space \mathbb{C}^{2^n} . Let $\widehat{\mathcal{Y}}$ denote the dual group of \mathcal{Y} , consisting of all the group homomorphisms $\widehat{y}(z) := (-1)^{y \cdot z}$. It is well-known that $\widehat{\mathcal{Y}}$ is isomorphic to \mathcal{Y} . We assume $\widehat{\mathcal{Y}}$ to be an additive group with the group operation $\widehat{y} + \widehat{z} := \widehat{y \oplus z}$. Naturally, $\widehat{0}$ denotes the identity. For each $\widehat{y} \in \widehat{\mathcal{Y}}$ define

$$|\widehat{y}\rangle := \frac{1}{2^{n/2}} \sum_{z \in \mathcal{Y}} \widehat{y}(z) |z\rangle = \frac{1}{2^{n/2}} \sum_{z \in \mathcal{Y}} (-1)^{y \cdot z} |z\rangle,$$

The set $\mathbf{F}_{\mathcal{Y}} := \{|\widehat{y}\rangle\}$ is referred as the *Fourier* basis of \mathbb{C}^{2^n} , and the mapping $|y\rangle \rightarrow |\widehat{y}\rangle$ is the well-known Hadamard transformation that maps the computational basis $\mathcal{C}_{\mathcal{Y}}$ to Fourier basis $\mathbf{F}_{\mathcal{Y}}$. The reverse basis transformation from $\mathbf{F}_{\mathcal{Y}}$ to $\mathcal{C}_{\mathcal{Y}}$ is given by

$$|y\rangle := \frac{1}{2^{n/2}} \sum_{\widehat{z} \in \widehat{\mathcal{Y}}} \widehat{z}(y) |\widehat{z}\rangle = \frac{1}{2^{n/2}} \sum_{\widehat{z} \in \widehat{\mathcal{Y}}} (-1)^{z \cdot y} |\widehat{z}\rangle.$$

Next, let \mathcal{Z} denote the set $\mathcal{Y} \cup \{\perp\}$ for a special symbol \perp ; similarly $\widehat{\mathcal{Z}}$ will denote $\widehat{\mathcal{Y}} \cup \{\perp\}$. We also choose a corresponding norm-1 vector $|\perp\rangle$ orthogonal to \mathbb{C}^{2^n} , so that the span of both $\mathcal{C}_{\mathcal{Z}} := \{|y\rangle \mid y \in \mathcal{Z}\}$ and $\mathbf{F}_{\mathcal{Z}} := \{|\widehat{y}\rangle \mid \widehat{y} \in \widehat{\mathcal{Z}}\}$ is $\mathbb{C}^{2^{n+1}}$; we’ll call $\mathcal{C}_{\mathcal{Z}}$ and $\mathbf{F}_{\mathcal{Z}}$ the computational basis and Fourier basis respectively of the extended space $\mathbb{C}^{2^{n+1}}$.

Functions and Databases. Let \mathcal{X} denote $\{0, 1\}^m$ for some arbitrary m , and let \mathcal{F} denote the set of m -bit-to- n -bit classical functions $f : \mathcal{X} \rightarrow \mathcal{Y}$. The *quantum truth table* of f is defined as

$$|f\rangle := \bigotimes_{x \in \mathcal{X}} |x\rangle |f(x)\rangle.$$

Let $\widehat{\mathcal{F}}$ denote the set of *Fourier* functions $\widehat{f} : \mathcal{X} \rightarrow \widehat{\mathcal{Y}}$. The quantum truth table of \widehat{f} is defined similarly as

$$|\widehat{f}\rangle := \bigotimes_{x \in \mathcal{X}} |x\rangle |\widehat{f}(x)\rangle.$$

For a subset $\mathcal{S} \subseteq \mathcal{X}$, a function $f : \mathcal{S} \rightarrow \mathcal{Y}$ will be called a *partial function* from \mathcal{X} to \mathcal{Y} . A partial function f can be extended to a function $d_f : \mathcal{X} \rightarrow \mathcal{Z}$ by defining $d_f(y) = \perp$ for all $y \in \mathcal{X} \setminus \mathcal{S}$. We call d_f the *database* representing f , with \perp denoting the cells where f is not defined. (When f is a full function, d_f coincides with f .) The database will also be represented as a quantum truth table

$$|d_f\rangle := \bigotimes_{x \in \mathcal{X}} |x\rangle |d_f(x)\rangle.$$

Similarly we define partial Fourier functions $\widehat{f} : \mathcal{S} \rightarrow \widehat{\mathcal{Y}}$, databases $d_{\widehat{f}} : \mathcal{X} \rightarrow \widehat{\mathcal{Z}}$ representing partial Fourier functions, and their quantum truth tables $|d_{\widehat{f}}\rangle$.

When f and \widehat{f} are clear from context, we'll find it convenient to drop the subscripts and write d_f and $d_{\widehat{f}}$ simply as d and \widehat{d} respectively. We'll write \mathcal{D} (resp. $\widehat{\mathcal{D}}$) to denote the set of all databases $d : \mathcal{X} \rightarrow \mathcal{Z}$ (resp. all Fourier databases $\widehat{d} : \mathcal{X} \rightarrow \widehat{\mathcal{Z}}$). When convenient we will treat a database d as a relation on $\mathcal{X} \times \mathcal{Y}$ and write $(x, y) \in d$ to denote $d(x) = y$; $|d|$ will then denote the size of this relation, i.e., the size of $\{x \in \mathcal{X} \mid d(x) \in \mathcal{Y}\}$.

For any function $f \in \mathcal{F}$, let $\widehat{f} \in \widehat{\mathcal{F}}$ be defined as the map $x \mapsto \widehat{f}(x)$. Then we have

$$|\widehat{f}\rangle = \frac{1}{2^{n2^m/2}} \sum_{g \in \mathcal{F}} (-1)^{f \cdot g} |g\rangle, \tag{3}$$

where $f \cdot g$ is defined as $\sum_{x \in \mathcal{X}} f(x) \cdot g(x)$. Thus, $\{|f\rangle \mid f \in \mathcal{F}\}$ and $\{|\widehat{f}\rangle \mid \widehat{f} \in \widehat{\mathcal{F}}\}$ span the same space (isomorphic to $\mathbb{C}^{2^{2^m}}$). Similarly we can show that $\{|d\rangle \mid d \in \mathcal{D}\}$ and $\{|\widehat{d}\rangle \mid \widehat{d} \in \widehat{\mathcal{D}}\}$ span the same space isomorphic to $\mathbb{C}^{(2^n+1)^{2^m}}$; we call this space the *database space* \mathbb{D} .

Letting $\mathbf{0}$ denote the constant 0^n function and observing that $\mathbf{0} \cdot g = 0$ for any $g \in \mathcal{F}$, we have

$$|\widehat{\mathbf{0}}\rangle = \frac{1}{2^{n2^m/2}} \sum_{g \in \mathcal{F}} |g\rangle,$$

the uniform superposition over all functions in \mathcal{F} .

The Standard Oracle. The *standard oracle* is a stateful oracle with $\mathcal{H}_{db} = \mathcal{H}[\mathcal{F}]$. Given a truth-table representation $|f\rangle$ of a function $f \in \mathcal{F}$, it acts on the adversary registers $|x\rangle|y\rangle$ and the truth-table registers $|f\rangle$ as

$$\mathbf{stO}|x\rangle|y\rangle \otimes |f\rangle = |x\rangle|y \oplus f(x)\rangle \otimes |f\rangle. \tag{4}$$

It is obvious to see that \mathbf{stO} is perfectly indistinguishable with a uniform random function, when the truth table register is initialized in $|\widehat{\mathbf{0}}\rangle$.

If we first put the adversary's response register and the truth-table register in the Fourier basis, we have

$$\mathbf{stO}|x\rangle|\widehat{y}\rangle \otimes |\widehat{f}\rangle = |x\rangle|\widehat{y}\rangle \otimes |\widehat{f} + \widehat{\delta}_{xy}\rangle, \tag{5}$$

where δ_{xy} is the function in \mathcal{F} defined as

$$\delta_{xy}(z) = \begin{cases} y & \text{when } z = x, \\ 0 & \text{otherwise,} \end{cases}$$

and the operations \oplus in \mathcal{F} and $+$ in $\widehat{\mathcal{F}}$ are defined point-wise. We define the operator $\mathbf{O}_{x\widehat{y}}$ on the truth-table register as

$$\mathbf{O}_{x\widehat{y}}|\widehat{f}\rangle := |\widehat{f} + \widehat{\delta}_{xy}\rangle.$$

Then we can write $\mathbf{stO}|x\rangle|\widehat{y}\rangle \otimes |\widehat{f}\rangle = |x\rangle|\widehat{y}\rangle \otimes \mathbf{O}_{x\widehat{y}}|\widehat{f}\rangle$.

The Compressed Oracle. For any $x \in \mathcal{X}$, the *cell compression* unitary \mathbf{comp}_x on \mathbb{C}^{2^n+1} is defined on the basis $\mathbf{F}_{\mathcal{Z}}$ as

$$\mathbf{comp}_x := |\perp\rangle\langle\widehat{0}| + |\widehat{0}\rangle\langle\perp| + \sum_{\widehat{y} \in \widehat{\mathcal{Y}} \setminus \{\widehat{0}\}} |\widehat{y}\rangle\langle\widehat{y}|.$$

The *database compression* unitary \mathbf{comp} on \mathbb{D} is defined as

$$\mathbf{comp} := \bigotimes_{x \in \mathcal{X}} \mathbf{comp}_x.$$

The *compressed oracle* \mathbf{cO} is a stateful oracle with $\mathcal{H}_{db} = \mathbb{D}$. It acts on the adversary’s registers and the oracle’s database registers as

$$\mathbf{cO} := (\mathbf{I}_{\mathcal{H}[\mathcal{X}] \otimes \mathcal{H}[\widehat{\mathcal{Y}}]} \otimes \mathbf{comp}) \circ \mathbf{stO} \circ (\mathbf{I}_{\mathcal{H}[\mathcal{X}] \otimes \mathcal{H}[\widehat{\mathcal{Y}}]} \otimes \mathbf{comp}).$$

For a database \widehat{d} we have

$$\mathbf{cO}|x\rangle|\widehat{y}\rangle \otimes |\widehat{d}\rangle = |x\rangle|\widehat{y}\rangle \otimes \mathbf{cO}_{x\widehat{y}}|\widehat{d}\rangle,$$

where $\mathbf{cO}_{x\widehat{y}} := \mathbf{comp}_x \circ \mathbf{O}_{x\widehat{y}} \circ \mathbf{comp}_x$.

3.2 The Two-Domain Distance Technique

Bhaumik et al. distilled [5] the Chung et al. interpretation [13] for indistinguishability setting and proposed a generic way to represent both the ideal and real world oracles using a single compressed permutation oracle. In addition, they combined it with a result from Hosoyamada and Iwata to get a quantum analog of “identical-up to-bad”, the so-called two-domain distance lemma.

Domain-Restricted Databases. For a subset $\widetilde{\mathcal{X}}$ of \mathcal{X} we will write $\mathcal{D}|_{\widetilde{\mathcal{X}}}$ to denote the set of databases restricted to $\widetilde{\mathcal{X}}$, defined equivalently as $\{d|_{\widetilde{\mathcal{X}}} \mid d \in \mathcal{D}\}$ or the set of databases $d : \widetilde{\mathcal{X}} \rightarrow \mathcal{Z}$. Since \mathcal{D} is a basis of the database space \mathbb{D} , a domain-restricted database space will span a subspace of \mathbb{D} isomorphic to $\mathbb{C}^{(2^n+1)^{|\widetilde{\mathcal{X}}|}}$. We continue to represent elements of $\widetilde{\mathcal{X}}$ as m -bit numbers.

Transition Capacity. For a domain-restricted database-set $\mathcal{D}|_{\tilde{\mathcal{X}}}$, a subset $\mathcal{P} \subseteq \mathcal{D}|_{\tilde{\mathcal{X}}}$ will be called a *database property* on $\mathcal{D}|_{\tilde{\mathcal{X}}}$. We also define the projection

$$\Pi_{\mathcal{P}} := \sum_{d \in \mathcal{P}} |d\rangle\langle d|.$$

For a database $d \in \mathcal{D}|_{\tilde{\mathcal{X}}}$ and an $x \in \tilde{\mathcal{X}}$ define

$$d|^x := \{d' \in \mathcal{D}|_{\tilde{\mathcal{X}}} \mid d'(x') = d(x') \forall x' \in \tilde{\mathcal{X}} \setminus \{x\}\}.$$

In other words, $d|^x$ is the set of databases in $\mathcal{D}|_{\tilde{\mathcal{X}}}$ which are identical to d except (possibly) at x . (Note that since d (resp. x) is also in \mathcal{D} (resp. \mathcal{X}), $d|^x$ is only well-defined when we specify $\mathcal{D}|_{\tilde{\mathcal{X}}}$ as well; however, since $\mathcal{D}|_{\tilde{\mathcal{X}}}$ will usually be clear from the context, for notational convenience we leave the dependence of $d|^x$ on $\mathcal{D}|_{\tilde{\mathcal{X}}}$ implicit.)

For two properties \mathcal{P} and \mathcal{P}' , the *transition capacity* from \mathcal{P} to \mathcal{P}' is defined as

$$\llbracket \mathcal{P} \hookrightarrow \mathcal{P}' \rrbracket := \max_{x \in \tilde{\mathcal{X}}, \hat{y} \in \hat{\mathcal{Y}}, d \in \mathcal{D}|_{\tilde{\mathcal{X}}}} \|\Pi_{\mathcal{P}' \cap d|^x} \circ \mathbf{cO}_{x\hat{y}} \circ \Pi_{\mathcal{P} \cap d|^x}\|.$$

The transition capacity $\llbracket \mathcal{P} \hookrightarrow \mathcal{P}' \rrbracket$ is roughly a measure of an upper bound for how likely it can be that a database in \mathcal{P} will transition into a database in \mathcal{P}' after a single query to \mathbf{cO} .

For a property $\mathcal{P} \subseteq \mathcal{D}|_{\tilde{\mathcal{X}}}$, let \mathcal{P}^c denote its negation, i.e., $\mathcal{D}|_{\tilde{\mathcal{X}}} \setminus \mathcal{P}$. Then we have the following lemma from [5, Transition Capacity Bound].

Lemma 1. *Let $\mathcal{P}, \mathcal{P}'$ be properties on $\mathcal{D}|_{\tilde{\mathcal{X}}}$ such that for every $x \in \tilde{\mathcal{X}}$ and $d \in \mathcal{D}|_{\tilde{\mathcal{X}}}$, we can find a set $\mathcal{S}_{x,d}^{\mathcal{P}^c \hookrightarrow \mathcal{P}'} \subseteq \mathcal{Y}$ satisfying*

$$\mathcal{P}' \cap d|^x \subseteq \{d' \in d|^x \mid d'(x) \in \mathcal{S}_{x,d}^{\mathcal{P}^c \hookrightarrow \mathcal{P}'}\} \subseteq \mathcal{P} \cap d|^x. \quad (6)$$

In other words, for any database $d' \in d|^x$,

$$d' \in \mathcal{P}' \implies d'(x) \in \mathcal{S}_{x,d}^{\mathcal{P}^c \hookrightarrow \mathcal{P}'} \implies d' \in \mathcal{P}.$$

Then we have

$$\llbracket \mathcal{P}^c \hookrightarrow \mathcal{P}' \rrbracket \leq \max_{x \in \tilde{\mathcal{X}}, d \in \mathcal{D}|_{\tilde{\mathcal{X}}}} \sqrt{\frac{10|\mathcal{S}_{x,d}^{\mathcal{P}^c \hookrightarrow \mathcal{P}'}|}{2^n}}.$$

Size-Restricted Properties. For a domain-restricted database-set $\mathcal{D}|_{\tilde{\mathcal{X}}}$, a property $\mathcal{P} \subseteq \mathcal{D}|_{\tilde{\mathcal{X}}}$, and some $i \leq |\tilde{\mathcal{X}}|$, we define

$$\mathcal{P}_{\leq i} := \{d \in \mathcal{P} \mid |d| \leq i\}.$$

Then the transition capacity $\llbracket \mathcal{P}_{\leq i-1}^c \hookrightarrow \mathcal{P}_{\leq i} \rrbracket$ is a measure of the maximum probability of a database outside \mathcal{P} with at most $i-1$ entries changing to a

database in \mathcal{P} after a single application $\mathbf{cO}_{x\hat{y}}$. (Note that $\mathcal{P}_{\leq i-1}^c$ denotes the size-restriction of \mathcal{P}^c , and not the complement of $\mathcal{P}_{\leq i-1}$.)

Let $\perp := \{d_\perp\}$ denote the *empty* property (where d_\perp is the empty database, i.e., the constant- \perp function). Then for \mathcal{P} such that $d_\perp \notin \mathcal{P}$, $\perp = \mathcal{P}_{\leq 0}^c$. We define

$$\left(\perp \overset{q}{\rightsquigarrow} \mathcal{P}\right) := \sum_{i=1}^q \llbracket \mathcal{P}_{\leq i-1}^c \leftrightarrow \mathcal{P}_{\leq i} \rrbracket,$$

the *q-query transition bound* from \perp to \mathcal{P} . In other words, $\left(\perp \overset{q}{\rightsquigarrow} \mathcal{P}\right)$ is a measure of the probability that the empty database changes into a database in \mathcal{P} at any point during q successive queries.

Prefixed Oracle. Fix some $\bar{t} < m$ and write $\mathcal{X} = \mathcal{T} \times \mathcal{I}$, where $\mathcal{T} = \{0, 1\}^{\bar{t}}$ and $\mathcal{I} = \{0, 1\}^{m-\bar{t}}$. For every non-zero $t \leq 2^{\bar{t}}$, any family of functions $\mathbf{p} = (\mathbf{p}_k : \mathcal{I} \rightarrow \mathcal{X})_{k \in [t]}$ is said to be a (\bar{t}, m) -*domain-separator* if for each $k \in [t]$ and for all $x \in \mathcal{I}$, $\mathbf{p}_k(x) \in \{\delta_{\bar{t}}(k) \parallel x : x \in \mathcal{I}\}$, for some fixed injective function $\delta_{\bar{t}} : [t] \rightarrow \mathcal{T}$. Let $\mathbf{p}_k(\mathcal{I}) := \{\mathbf{p}_k(x) : x \in \mathcal{I}\}$ and $\mathbf{p}(\mathcal{I}) := \cup_{k \in [t]} \mathbf{p}_k(\mathcal{I})$.

To any (\bar{t}, m) -domain-separator $\mathbf{p} = (\mathbf{p}_k : \mathcal{I} \rightarrow \mathcal{X})_{k \in [t]}$, we associate the *prefixed-compressed oracle* \mathbf{cOP} which is defined as a family of oracles $\{\mathbf{cOP}^k\}_{k \in [t]}$, where \mathbf{cOP}^k denotes the restriction of \mathbf{cO} to inputs from $\mathbf{p}_k(\mathcal{I}) \subset \mathcal{X}$, i.e., for any $k \in [t]$, $x \in \mathcal{I}$, $\hat{y} \in \hat{\mathcal{Y}}$ and $\hat{d} \in \hat{\mathcal{D}}$, we have

$$\mathbf{cOP}^k |x\rangle|\hat{y}\rangle \otimes |\hat{d}\rangle = |x\rangle|\hat{y}\rangle \otimes \mathbf{cOP}_{x\hat{y}}^k |\hat{d}\rangle,$$

where $\mathbf{cOP}_{x\hat{y}}^k := \mathbf{comp}_{\mathbf{p}_k(x)} \circ \mathbf{O}_{\mathbf{p}_k(x)\hat{y}} \circ \mathbf{comp}_{\mathbf{p}_k(x)}$. Consequently, \mathbf{cOP} can also be viewed as the restriction of \mathbf{cO} to inputs from $\mathbf{p}(\mathcal{I}) \subset \mathcal{X}$.

Two-Domain Systems. Let \mathbf{I} and \mathbf{R} be two stateful oracles with $\mathcal{H}_{in} = \mathcal{H}[\mathcal{I}]$, $\mathcal{H}_{out} = \mathcal{H}[\mathcal{Z}]$, $\mathcal{H}_{ab} = \mathbb{D}$, defined by the sequences of unitaries:

$$\mathbf{I} := \mathbf{F}_t \mathbf{cO}^{\mathbf{I}[t]} \dots \mathbf{cO}^{\mathbf{I}[1]} \mathbf{F}_0, \quad \mathbf{R} := \mathbf{F}_t \mathbf{cO}^{\mathbf{R}[t]} \dots \mathbf{cO}^{\mathbf{R}[1]} \mathbf{F}_0,$$

where with a slight abuse of notations we reuse \mathbf{I} and \mathbf{R} to also denote the corresponding (\bar{t}, m) -domain-separators, and the unitaries $\mathbf{F}_0, \dots, \mathbf{F}_t$ only operate on the input, output and ancillary qubits, if any, needed to compute the function itself. Whenever convenient, we will continue ignoring the ancillary qubits.

Consider a q -query interactive game where a computationally unbounded and deterministic distinguisher A aims to distinguish \mathbf{R} from \mathbf{I} . We emphasize that in such an interactive game with \mathbf{I} or \mathbf{R} , the compressed oracle \mathbf{cO} is invoked a total of $q' := tq$ times. Fix two domains $\tilde{\mathcal{X}}_{\mathbf{I}} = \mathbf{I}(\mathcal{I})$, $\tilde{\mathcal{X}}_{\mathbf{R}} = \mathbf{R}(\mathcal{I})$, and define $\mathcal{D}_{\mathbf{I}} := \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{I}}}$ and $\mathcal{D}_{\mathbf{R}} := \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{R}}}$. Consider properties $\mathcal{B}_{\mathbf{I}} \subseteq \mathcal{D}_{\mathbf{I}} \setminus \perp$ and $\mathcal{B}_{\mathbf{R}} \subseteq \mathcal{D}_{\mathbf{R}} \setminus \perp$, and define $\mathcal{G}_{\mathbf{I}} := \mathcal{D}_{\mathbf{I}} \setminus \mathcal{B}_{\mathbf{I}}$ and $\mathcal{G}_{\mathbf{R}} := \mathcal{D}_{\mathbf{R}} \setminus \mathcal{B}_{\mathbf{R}}$. The central tool of our security proofs will be the following adaptation of [5, Lemma 4]. A proof of this lemma is available in the full version of this paper [6].

Lemma 2 (Two-Domain Distance Lemma). *Suppose we can find a map $h : \mathcal{G}_{\mathbf{I}} \rightarrow \mathcal{G}_{\mathbf{R}}$ such that the following hold:*

- h is a bijection from $\mathcal{G}_{\mathbf{I}}$ to $\mathcal{G}_{\mathbf{R}}$;
- For every $i \in [q'] \cup \{0\}$, $h|_{\mathcal{G}_{\mathbf{I}[\leq i]}}$ is a bijection from $\mathcal{G}_{\mathbf{I}[\leq i]}$ to $\mathcal{G}_{\mathbf{R}[\leq i]}$;
- For every $i \in [q']$, $x \in \mathcal{I}$, $\hat{y} \in \hat{\mathcal{Y}}$, $d \in \mathcal{G}_{\mathbf{I}[\leq i-1]}$, and $d' \in \mathcal{G}_{\mathbf{I}[\leq i]}$,

$$\langle d' | \mathbf{cO}_{x\hat{y}}^{\mathbf{I}[k]} | d \rangle = \langle h(d') | \mathbf{cO}_{x\hat{y}}^{\mathbf{R}[k]} | h(d) \rangle.$$

where $k = t$ if $i = 0 \bmod t$, and $k = i \bmod t$ otherwise.

Then, we have

$$\|\mathrm{Tr}_{\mathbb{D}}(\rho_{A,\mathbf{I}}^q) - \mathrm{Tr}_{\mathbb{D}}(\rho_{A,\mathbf{R}}^q)\|_1 \leq 3 \left(\perp \overset{q'}{\rightsquigarrow} \mathcal{B}_{\mathbf{I}} \right)_{\mathbf{I}} + 3 \left(\perp \overset{q'}{\rightsquigarrow} \mathcal{B}_{\mathbf{R}} \right)_{\mathbf{R}},$$

where $\rho_{A,\mathbf{p}}^q := A^{\mathbf{p}} |\psi_A, d_{\perp}\rangle \langle \psi_A, d_{\perp}| A^{\mathbf{p}\dagger}$ is the state after q queries to the oracle at-hand $\mathbf{p} \in \{\mathbf{I}, \mathbf{R}\}$ for some norm-1 vector $|\psi_A\rangle$ and the empty database $|d_{\perp}\rangle$. The transition bounds $\left(\perp \overset{q'}{\rightsquigarrow} \cdot \right)_{\mathbf{I}}$ and $\left(\perp \overset{q'}{\rightsquigarrow} \cdot \right)_{\mathbf{R}}$ are computed for queries to $\mathbf{cO}^{\mathbf{I}}$ and $\mathbf{cO}^{\mathbf{R}}$, respectively.

When the oracle in use is clear from the context, we will drop the subscripts for the transition bounds and simply write both as $\left(\perp \overset{q'}{\rightsquigarrow} \cdot \right)$. We'll also keep the domain-separator implicit when there's no scope for ambiguity.

3.3 The Hosoyamada-Iwata Interpretation

Hosoyamada and Iwata proposed a slightly different variant of \mathbf{stO} with an aim to characterize and analyze databases in an explicit computational basis with an exact definition of \perp with the help of an ancillary flag bit that signifies if the database entry is defined or not.

Let $\mathcal{S} \subseteq \mathcal{X}$ and $\mathcal{Z} = \{0, 1\} \times \mathcal{Y}$. For any partial function $f : \mathcal{S} \rightarrow \mathcal{Y}$, we associate the database function $d_f : \mathcal{X} \rightarrow \mathcal{Z}$ defined as:

$$d_f(x) := \begin{cases} (1, y) & \text{when } f(x) = y \in \mathcal{Y}, \\ (0, 0^n) & \text{if } f(x) \text{ is undefined,} \end{cases}$$

On comparing this with Zhandry's original interpretation, we see that the \perp in original interpretation corresponds to $(0, 0^n)$ in HI interpretation. As before, we drop the subscripts when f is either clear from the context or inconsequential.

We define the database space as the $2^{(n+1)2^m}$ -dimensional complex Hilbert space $\mathcal{H}_{db} = \mathcal{H}[\mathcal{Z}]$ which is isomorphic to $\mathbb{C}^{2^{(n+1)2^m}}$. Note that not all $d \in \mathcal{Z}$ can be associated with some partial function f . A database $d = ((b_0, \beta_0), \dots, (b_{2^m-1}, \beta_{2^m-1}))$ is said to be *valid* if it satisfies that for each $i \in \{0, 1, \dots, 2^m - 1\}$ such that $b_i = 0$ we have $\beta_i = 0^n$. Indeed, any valid database $((b_0, \beta_0), \dots, (b_{2^m-1}, \beta_{2^m-1}))$ is identified with the set $\{(i, \beta_i) \mid b_i = 1\}$,

which is nothing but the truth table of a partially-defined function from $\{0, 1\}^m$ to $\{0, 1\}^n$. Accordingly, let Π_{valid} be the orthogonal projection onto the vector space spanned by valid databases.

Any database $|d\rangle \in \mathcal{H}[\mathcal{Z}]$ can be equivalently viewed as an array of 2^m cells $|d[0]\rangle \dots |d[2^m - 1]\rangle$. Writing $|d[i]\rangle$ as $|b_i, \beta_i\rangle$ for each $i \in \{0, 1, \dots, 2^m - 1\}$ (where b_i and β_i are respectively the control qubit and the response register of the i -th cell $|d[i]\rangle$ of $|d\rangle$), the standard oracle \mathbf{stO} is now defined as:

$$\mathbf{stO}|i, y\rangle|d\rangle := |i, y + \beta_i\rangle|d\rangle$$

for each $|i, y, d\rangle \in \mathcal{H}_{in} \times \mathcal{H}_{out} \times \mathcal{H}_{db}$. For $|d\rangle$ such that $|d[i]\rangle = |0, 0^n\rangle$, we define $|d \cup (i, \beta)\rangle$ to be the database with $|1, \beta\rangle$ as its i -th cell and identical to $|d\rangle$ in all other cells.

Define the following unitary operators on database cells:

$$\begin{aligned} \mathbf{IH}_0 &:= \mathbf{I}_1 \otimes \mathbf{H}^{\otimes n} & \mathbf{Tg}_0 &:= \mathbf{I}_1 \otimes |0^n\rangle\langle 0^n| + \mathbf{X}(\mathbf{I}_{2^n} - |0^n\rangle\langle 0^n|) \\ \mathbf{cH}_0 &:= |0\rangle\langle 0| \otimes \mathbf{I}_{2^n} + |1\rangle\langle 1| \otimes \mathbf{H}^{\otimes n} \end{aligned}$$

and databases:

$$\mathbf{IH} := \mathbf{IH}_0^{\otimes 2^m} \quad \mathbf{Tg} := \mathbf{Tg}_0^{\otimes 2^m} \quad \mathbf{cH} := \mathbf{cH}_0^{\otimes 2^m}$$

where \mathbf{X} and \mathbf{H} are the well-known flip and Hadamard operators on \mathbb{C} , i.e. in the computational basis:

$$\mathbf{X} := |0\rangle\langle 1| + |1\rangle\langle 0| \quad \mathbf{H} := \frac{1}{2}(|0\rangle\langle 0| + |0\rangle\langle 1| + |1\rangle\langle 0| - |1\rangle\langle 1|)$$

Note that all these operators are Hermitian. Using these, we define the encode and decode operator \mathbf{dec} on databases as follows:

$$\begin{aligned} \mathbf{enc} &:= \mathbf{cH} \circ \mathbf{Tg} \circ \mathbf{IH}; \\ \mathbf{dec} &:= \mathbf{enc}^\dagger = \mathbf{IH} \circ \mathbf{Tg} \circ \mathbf{cH}; \end{aligned}$$

The recording standard oracle \mathbf{RStOE} , due to Hosoyamada and Iwata [23], is defined as:

$$\mathbf{RStOE} := (\mathbf{I}_{2^{m+n}} \otimes \mathbf{enc})\mathbf{stO}(\mathbf{I}_{2^{m+n}} \otimes \mathbf{dec})$$

Thus, \mathbf{RStOE} first decodes the database, then applies \mathbf{stO} on the adversary's registers and the decoded database, and then encodes the database again. Let $|0\rangle$ denote the valid empty database.

Hosoyamada and Iwata proved [23, 25] the following useful propositions.

Proposition 1 (Proposition 1 in [25]). *Suppose that the oracle state is initialized in $|0\rangle$. For any $i \geq 1$, if the oracle state register is measured after i queries, then the resulting database d is valid, and contains at most i entries.*

Proposition 2 (Proposition 2 in [25]). *For any valid database d satisfying $d[i] = |0, 0^n\rangle$, we have*

$$\mathbf{RStOE}|i, y\rangle|d \cup (i, \beta)\rangle = |i, y \oplus \beta\rangle|d \cup (i, \beta)\rangle + |\epsilon_1\rangle; \quad (7)$$

$$\mathbf{RStOE}|i, y\rangle|d\rangle = \sum_{\beta \in \{0,1\}^n} \frac{1}{2^{n/2}} |i, y \oplus \beta\rangle|d \cup (i, \beta)\rangle + |\epsilon_2\rangle; \quad (8)$$

for some $|\epsilon_1\rangle$ and $|\epsilon_2\rangle$ such that $\| |\epsilon_1\rangle \|, \| |\epsilon_2\rangle \| \in O(1/\sqrt{2^n})$.

Although we do not require them in this paper, we remark that [23] gives an exact description of $|\epsilon_1\rangle$ and $|\epsilon_2\rangle$. Intuitively, $|\epsilon_1\rangle$ and $|\epsilon_2\rangle$ can be viewed as the errors introduced in the lazy sampling of a quantum random function due to interference.

Finally, the main technical result used to study the indistinguishability game and bound the advantage is given below.

Proposition 3 (Proposition 3 in [25]). *For each $j \in \{0, 1, \dots, q\}$, let $|\mathbb{R}_j\rangle$ and $|\mathbb{I}_j\rangle$ denote the state vector corresponding to the real and ideal worlds after the j -th query, respectively. Suppose, there exist vectors $|\mathbb{R}_j^{\mathbf{g}}\rangle, |\mathbb{R}_j^{\mathbf{b}}\rangle, |\mathbb{I}_j^{\mathbf{g}}\rangle, |\mathbb{I}_j^{\mathbf{b}}\rangle$ and non-negative reals $\epsilon_{\mathbf{I}}^{(j)}$ and $\epsilon_{\mathbf{R}}^{(j)}$ such that*

1. $|\mathbb{R}_j\rangle = |\mathbb{R}_j^{\mathbf{g}}\rangle + |\mathbb{R}_j^{\mathbf{b}}\rangle, |\mathbb{I}_j\rangle = |\mathbb{I}_j^{\mathbf{g}}\rangle + |\mathbb{I}_j^{\mathbf{b}}\rangle;$
2. $|\mathbb{R}_j^{\mathbf{g}}\rangle\langle\mathbb{R}_j^{\mathbf{g}}| = |\mathbb{I}_j^{\mathbf{g}}\rangle\langle\mathbb{I}_j^{\mathbf{g}}|;$
3. $\| |\mathbb{I}_j^{\mathbf{b}}\rangle \| \leq \| |\mathbb{I}_{j-1}^{\mathbf{b}}\rangle \| + \epsilon_{\mathbf{I}}^{(j)}, \| |\mathbb{R}_j^{\mathbf{b}}\rangle \| \leq \| |\mathbb{R}_{j-1}^{\mathbf{b}}\rangle \| + \epsilon_{\mathbf{R}}^{(j)}.$

Then, for any computationally unbounded and deterministic distinguisher A we have $\| \text{Tr}_{\mathcal{H}_{\text{Idb}}}(\rho_{A,\mathbf{I}}^q) - \text{Tr}_{\mathcal{H}_{\text{Rdb}}}(\rho_{A,\mathbf{R}}^q) \|_1 \leq \sum_{i=1}^q \epsilon_{\mathbf{I}}^{(i)} + \sum_{i=1}^q \epsilon_{\mathbf{R}}^{(i)}$, where $\rho_{A,\mathbf{R}}^q = |\psi_A\rangle\langle\psi_A| \otimes |\mathbf{0}_{\mathbf{R}}\rangle\langle\mathbf{0}_{\mathbf{R}}|$ and $\rho_{A,\mathbf{I}}^q = |\psi_A\rangle\langle\psi_A| \otimes |\mathbf{0}_{\mathbf{I}}\rangle\langle\mathbf{0}_{\mathbf{I}}|$ for some norm-1 vector $\psi_A \in \mathcal{H}_A$ and $|\mathbf{0}_{\mathbf{R}}\rangle$ and $|\mathbf{0}_{\mathbf{I}}\rangle$ denote the all zero database states in the real and ideal worlds respectively.

4 Revisiting IND-qCPA Security of \mathbf{LR}_4

4.1 The Luby-Rackoff Construction

For some $r \geq 1$ and $f_1, \dots, f_r : \{0, 1\}^n \rightarrow \{0, 1\}^n$, we define $g : [r] \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ by the mapping:

$$(i, x_1, x_2) \mapsto (x_2 \oplus f_i(x_1), x_1),$$

and write $g_i(\cdot, \cdot) := g(i, \cdot, \cdot)$. The r -round Luby-Rackoff construction, denoted \mathbf{LR}_r is defined as:

$$(x_1, x_2) \mapsto g_r \circ \dots \circ g_1(x_1, x_2). \quad (9)$$

For all $i \in [r]$, we write (also see Fig. 1):

- $x^{i-1} := (x_1^{i-1}, x_2^{i-1})$ to denote the input to g_i , where $x^0 := x = (x_1, x_2)$, denotes the input to \mathbf{LR}_r .

- (u_i, v_i) to denote the input-output tuple corresponding to f_i .
- $y = (y_1, y_2) := (x_1^r, x_2^r)$ to denote the output of LR_r .

Hosoyamada and Iwata stated [23] the following IND-qCPA security bound for LR_4 .

Theorem 1 (Theorem 3 in [23]). *Suppose $f_1, f_2, f_3, f_4 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are four mutually independent uniform random functions. Then, for any $q \geq 0$, and any quantum adversary A that makes at most q CPA queries, we have*

$$\text{Adv}_{\text{LR}_4}^{\text{qcpa}}(A) = O\left(\sqrt{\frac{q^3}{2^n}}\right).$$

The proof of this theorem uses the HI interpretation of Zhandry’s compressed oracle, the so-called **RStOE**. The high level proof approach is as follows:

1. Simulate the random functions f_1, f_2, f_3, f_4 using independent instances of **RStOE** with the corresponding databases, $d_1, d_2, d_{\mathbf{R}}, d_4$, respectively.
2. The authors then apply a series of hybrids, introducing intermediate constructions between the real construction LR_4 , and the ideal construction, a uniform random function $\Gamma : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. The first of these intermediate constructions is a length-preserving function, that we refer as $\widetilde{\text{LR}}_4$, defined by the mapping (see also Fig. 1):

$$(x_1, x_2) \mapsto g_4 \circ G_3 \circ g_2 \circ g_1(x_1, x_2), \tag{10}$$

where $G_3(x'_1, x'_2) := (F_3(x'_1, x'_2), x'_1)$ for all $(x'_1, x'_2) \in \{0, 1\}^{2n}$. The function $F_3 : \{0, 1\}^{2n} \rightarrow \{0, 1\}^n$ is a uniform random function, to be implemented by an appropriate **RStOE**, say $d_{\mathbf{I}}$.

In this note, we will solely focus on the distance between LR_4 and $\widetilde{\text{LR}}_4$. In fact, showing a negligible distance between the two systems is the technical core of the proof. For the discussion in this paper, it is sufficient to consider the chopped output x_1^3 . So, we drop the application of f_4 . We write $d^{\mathbf{R}} = (d_1, d_2, d_{\mathbf{R}})$ and $d^{\mathbf{I}} = (d_1, d_2, d_{\mathbf{I}})$.

3. In a bid to use Proposition 3 to bound the advantage, the authors iteratively apply Proposition 2 to study the action of each of f_1, f_2, f_3 (only in the real world), and F_3 (only in the ideal world) in that order, followed by the respective uncomputation steps for f_2 and f_1 in that order.
4. The key idea in the proof is the observation that LR_4 and $\widetilde{\text{LR}}_4$ are indistinguishable as long as the inputs to f_3 (res. F_3 in the ideal world) are pairwise distinct across all queries, i.e., the database triple $d^{\mathbf{R}} = (d_1, d_2, d_{\mathbf{R}})$ (res. $d^{\mathbf{I}} = (d_1, d_2, d_{\mathbf{I}})$ in the ideal world) is considered to be *good* if and only if there does not exist distinct database entries $(u_1, v_1), (u'_1, v'_1) \in d_1, (u_2, v_2), (u'_2, v'_2) \in d_2$, and $(u_3, v_3) \in d_{\mathbf{R}}$ (res. $(u_3, x_2^2, v_3) \in d_{\mathbf{I}}$ in the real world) such that $u_1 \oplus v_2 = u'_1 \oplus v'_2 = u_3$. All other database triples are considered *bad*. Let Π_{bad} denote the projection onto the space spanned by bad databases. A key property of good database triples is the fact that they enable a one-to-one correspondence $d^{\mathbf{R}} \mapsto [d^{\mathbf{R}}]_{\mathbf{I}}$ between the real and ideal databases, i.e., the two worlds can be easily shown to behave identically

when the databases remain good throughout the execution. Thus, by setting $|\mathbb{R}_j^b\rangle = \Pi_{bad}|\mathbb{R}_j\rangle$, $|\mathbb{R}_j^g\rangle = |\mathbb{R}_j\rangle - |\mathbb{R}_j^b\rangle$, $|\mathbb{I}_j^b\rangle = \Pi_{bad}|\mathbb{I}_j\rangle$, and $|\mathbb{I}_j^g\rangle = |\mathbb{I}_j\rangle - |\mathbb{I}_j^b\rangle$, we satisfy condition 1 and 2 in Proposition 3.

Now, all that remains is to study the action of each function call, and bound the norm of the bad vectors after each application, assuming that the state is spanned by good databases before the action. In particular, we concentrate on the application of f_1 in the next section, uncovering a flaw in the argumentation that breaks the proof.

4.2 Action of f_1 and the Trivialization of Norm

For any unit vector $|\psi\rangle$ and an arbitrary projection operator Π , we say that $\|\Pi|\psi\rangle\|$ is *trivially bounded* when we simply use the fact that $\|\Pi|\psi\rangle\| \leq 1$.




Fig. 1. 4-round Luby-Rackoff (left) and 4-round Luby-Rackoff with a BIG function (right).

We will study the action in the ideal world, although the same issue lies in the real world application as well. For brevity we assume that the output of f_1 is written on some ancillary register to be used in later actions. By a recursive application of Proposition 2, there exists vectors $|\epsilon_1\rangle$ and $|\epsilon_2\rangle$ such that

$$\begin{aligned} \mathbf{O}_{f_1} |\mathbb{I}_{j-1}^g\rangle &:= \sum_{\substack{x,y,z,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) \neq \perp}} \alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)} |x,y,z\rangle \otimes |d_1(x_1)\rangle \otimes |d^{\mathbf{I}}\rangle \\ &+ \sum_{\substack{x,y,z,\beta,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) = \perp}} \frac{\alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)}}{2^{n/2}} |x,y,z\rangle \otimes |\beta\rangle \otimes |d^{\mathbf{I}} \cup (x_1, \beta)_1\rangle \\ &+ |\epsilon_1\rangle + |\epsilon_2\rangle, \end{aligned}$$

where $|d^{\mathbf{I}} \cup (x_1, \beta)_1\rangle = |d_1 \cup (x_1, \beta)\rangle \otimes |d_2\rangle \otimes |d_{\mathbf{I}}\rangle$ denotes the database that is same as $|d^{\mathbf{I}}\rangle$ except for $d_1(x_1)$ which has been newly defined as β .

In this note we are only concerned with the second summand, denoted $|\mathbb{I}_j^{g,1}\rangle$, which gives the state transition on a fresh input to f_1 starting with a good state. Roughly speaking, a new entry (x_1, β) is recorded in d_1 at the cost of an amplitude factor of $2^{-n/2}$.

Formally, we are interested in the following norm, which is an equivalent representation of [23, (51)]:

$$\begin{aligned} \|I_{bad} |\mathbb{I}_j^{g,1}\rangle\|^2 &= \left\| \sum_{\substack{x,y,z,\beta,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) = \perp \\ d^{\mathbf{I}} \cup (x_1, \beta)_1:\text{bad}}} \frac{\alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)}}{2^{n/2}} |x,y,z\rangle \otimes |\beta\rangle \otimes |d^{\mathbf{I}} \cup (x_1, \beta)_1\rangle \right\|^2 \\ &= \sum_{\substack{x,y,z,\beta,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) = \perp \\ d^{\mathbf{I}} \cup (x_1, \beta)_1:\text{bad}}} \left| \frac{\alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)}}{2^{n/2}} \right|^2 \end{aligned} \quad (11)$$

$$= \sum_{\substack{x,y,z,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) = \perp}} \left| \alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)} \right|^2 \sum_{\substack{\beta \\ d^{\mathbf{I}} \cup (x_1, \beta)_1:\text{bad}}} \frac{1}{2^n} \quad (12)$$

$$\leq O\left(\frac{j}{2^n}\right) \sum_{\substack{x,y,z,d^{\mathbf{I}} \\ d^{\mathbf{I}}:\text{good} \\ d_1(x_1) = \perp}} \left| \alpha_{x,y,z,d^{\mathbf{I}}}^{(j-1)} \right|^2 \quad (13)$$

$$\leq O\left(\frac{j}{2^n}\right), \quad (14)$$

where (13) to (14) follows from the fact that $\|\mathbb{I}_{j-1}\| \leq 1$. However, there is no supporting argument in [23] for (12) to (13). In fact, we claim that

$$\sum_{\beta} \frac{1}{2^n} = O(1). \tag{15}$$

$d^{\mathbf{I}} \cup (x_1, \beta)_1 : \text{bad}$

To bound the summation, we have to estimate the size of the set $\{\beta : d^{\mathbf{I}} \cup (x_2, \beta)_1 \text{ is bad}\}$. Now, $d^{\mathbf{I}} \cup (x_1, \beta)_1$ is bad if and only if there exists distinct database entries $(u'_1, v'_1) \in d_1, (u_2, v_2), (u'_2, v'_2) \in d_2$, and $(u'_1 \oplus v'_2, u'_2, v'_3) \in d_{\mathbf{I}}$ such that

$$x_1 \oplus v_2 = u'_1 \oplus v'_2.$$

Note that, the above predicate is independent of β ! Thus, in the worst case, the predicate is true for all possible values of β which immediately establishes the claim. Once we plug in the bound from (15) in (12), we get

$$\|\Pi_{\text{bad}} \mathbb{I}_j^{g,1}\|^2 = O(1), \tag{16}$$

which clearly trivializes the norm. This completely breaks the security proof, as this revised bound leads to a trivial bound of $O(1)$ on the PRF advantage.

4.3 Do Additional Rounds Help?

One might think that, while this approach does not work for three rounds, maybe it will if we add more rounds, i.e., by considering r -round Luby-Rackoff for $r \geq 4$. Unfortunately, as we show in this section, the “trivialization of norm” seems to be a fundamental issue. We will argue this further for input collision at f_i for any odd $i \in \{1, \dots, r\}$. A similar argument can also be given for any even i .

Consider the database snapshot after $j \geq 2$ queries. Suppose, the adversary makes a query (x_1, x_2) , such that $d_1(x_1) = \perp$, i.e., the database entry corresponding to x_1 is empty, and a new entry (x_1, β) is to be created. Now, if we have distinct $(u'_1, v'_1) \in d_1, (u_2, v_2), (u'_2, v'_2) \in d_2, \dots, (u_{i-1}, v_{i-1}), (u'_{i-1}, v'_{i-1}) \in d_{i-1}, (u'_i, v'_i) \in d_i$, such that

$$u'_i = u'_1 \oplus v'_2 \oplus \dots \oplus v'_{i-1}, \text{ and}$$

$$x_1 \oplus v_2 \oplus \dots \oplus v_{i-1} = u'_1 \oplus v'_2 \oplus \dots \oplus v'_{i-1},$$

then there is a possibility⁴ that this query leads to a collision at the input of f_i . And what’s more, this condition is independent⁵ of β , and thus, a similar trivialization of norms as in (15) would occur in this case as well, rendering this line of argumentation effectively useless.

⁴ We are obviously overcounting by considering all possible combinations of queries. In fact, most of these combinations are never queried by the adversary. However, as of now, there is no effective way to find out the query ordering from database entries.

⁵ This independence only holds corresponding to the badness condition. In a typical execution of LR_r , these variables will obviously depend on β . However, due to the badness condition and the ignorance of query ordering (see the above point), this dependence is lost.

5 Non-adaptive IND-qCPA Security of LR₄

The main reason that the existing Luby-Rackoff proof fails is a lack of global knowledge of adversarial query pattern. At any instant, the compressed oracle only has the information recorded in the database and the current input. Thus, one has to argue as if every possible combination of global inputs are possible which as we showed in Sect. 4 leads to a trivialization of norm in case of LR₄. At the same time, for several other constructions, like TNT and LRWQ, one can still try to reconstruct a moderately global view to achieve some security bound.

THE DUMMY CALL IDEA: In the non-adaptive setting, the adversary makes a single query of the form $x^q = (x_1, \dots, x_q)$. We can employ a single dummy compressed oracle call to record x^q , and then implement the oracle at-hand. Note that the compressed oracle in both the dummy call and actual oracle evaluation can be implemented by a single compressed oracle using the prefixed oracle technique. More formally, suppose \mathbf{O}_f denote the stateful oracle corresponding to the function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^n$, defined as follows:

$$\mathbf{O}_f := \mathbf{F}_{t-1} \mathbf{cO}^{\mathbf{P}^{t-1}} \dots \mathbf{cO}^{\mathbf{P}^1} \mathbf{F}_0,$$

where \mathbf{p} is a (\bar{t}, m) -domain-separator for some $\bar{t} \geq \lceil \log_2 t \rceil$ such that $m \geq \ell q + \bar{t}$. Keep in mind that the unitaries $\mathbf{F}_0, \dots, \mathbf{F}_{t-1}$ only operate on the input, output and ancillary qubits, if any. Then, the q -query variant of \mathbf{O}_f with dummy call is defined to be the sequence

$$(\mathbf{cO}^{\mathbf{P}^t})^\dagger \circ \mathbf{O}_f^{\otimes q} \circ \mathbf{cO}^{\mathbf{P}^t}.$$

In other words, we enclose the original non-adaptive oracle between two compressed oracle calls, which record and erase the global input (x^q, \hat{y}^q) . Note that erasing the dummy call entries is crucial; otherwise, this perturbs the state.

In what follows, we assume the actions of the dummy call are implicit and do not analyze them explicitly. Consequently, we will often focus only on the relevant subspace of the database used in the other actions.

We prove the following IND-qNCPA bound for LR₄.

Theorem 2. *Suppose $f_1, f_2, f_3, f_4 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are three mutually independent uniform random functions. Then, for any $q \geq 0$, and any quantum adversary \mathcal{A} that makes at most q qNCPA queries, we have*

$$\mathbf{Adv}_{\text{LR}_4}^{\text{qncpa}}(\mathcal{A}) = 3\sqrt{\frac{q^6}{2^n}} + 6\sqrt{\frac{q^5}{2^n}}.$$

Proof. Our goal is to bound the distinguishing advantage for any non-adaptive adversary trying to distinguish LR₄ from a uniform random function. First, let $F_3, F_4 : \{0, 1\}^{3n} \rightarrow \{0, 1\}^n$ be two uniform random functions. For $i \in \{3, 4\}$, define

$$G_i(x_1, x_2, x'_1, x'_2) := (x'_2 \oplus F_i(x_1, x_2, x'_1), x'_1),$$

for any $(x_1, x_2, x'_1, x'_2) \in \{0, 1\}^{4n}$. We define the hybrid random function $\widetilde{\text{LR}}_4$ as (see also Fig. 2):

$$\widetilde{\text{LR}}_4(x_1, x_2) := G_4(x_1, x_2, G_3(x_1, x_2, \text{LR}_2(x_1, x_2))).$$

Then, it is easy to see that $\widetilde{\text{LR}}_4$ is indistinguishable to a uniform random function $\Gamma : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. So, it is sufficient to bound the distance between LR_4 and $\widetilde{\text{LR}}_4$. Let $\mathcal{X} = \{0, 1\}^{4+2nq}$, $\mathcal{Y} = \{0, 1\}^n$ and $\Gamma : \mathcal{X} \rightarrow \mathcal{Y}$ be a uniform random




Fig. 2. LR₄ (left) vs the hybrid random function, $\widetilde{\text{LR}}_4$ (right).

function. For each $x_1, x_2, x_3 \in \{0, 1\}^n$, we define

$$\begin{aligned} f_1(x_1) &:= \Gamma(1001\|x_1\|0^{2nq-n}) \\ f_2(x_1) &:= \Gamma(1010\|x_1\|0^{2nq-n}) \\ f_3(x_1) &:= \Gamma(1011\|x_1\|0^{2nq-n}) \\ f_4(x_1) &:= \Gamma(1100\|x_1\|0^{2nq-n}) \\ F_3(x_1, x_2, x_3) &:= \Gamma(1101\|x_1\|x_2\|x_3\|0^{2nq-3n}) \\ F_4(x_1, x_2, x_3) &:= \Gamma(1110\|x_1\|x_2\|x_3\|0^{2nq-3n}) \end{aligned}$$

In addition, we implicitly define the dummy call, denoted *dummy*, to operate over a disjoint⁶ subspace of the database, mapping $2qn$ -bit inputs to n -bit outputs. The exact description of the dummy call is not necessary as the output is never used.

The distinctness of the first four bits ensures that $f_1, f_2, f_3, f_4, F_3, F_4$ are all independent, and they are independent of *dummy* by definition.

The database in the real world is denoted $d_{\mathbf{R}}$ (tracking *dummy*, f_1, f_2, f_3, f_4) and $d_{\mathbf{I}}$ in the ideal world (tracking *dummy*, f_1, f_2, F_3, F_4). Let $\mathcal{D}_{\mathbf{R}}$ (resp. $\mathcal{D}_{\mathbf{I}}$) be the set of all possible choices for $d_{\mathbf{R}}$ (resp. $d_{\mathbf{I}}$).

For some $x = (x_1, x_2, \dots, x_{2q}) \in \mathcal{Y}^{2q}$, let

$$\begin{aligned} [x]_0 &:= 0000\|x & [x_1]_1 &:= 1001\|x_1\|0^{2nq-n} \\ [x_1, x_2, x_3]_5 &:= 1101\|x_1\|x_2\|x_3\|0^{2nq-3n} & [x_1]_2 &:= 1010\|x_1\|0^{2nq-n} \\ [x_1, x_2, x_3]_6 &:= 1110\|x_1\|x_2\|x_3\|0^{2nq-3n} & [x_1]_3 &:= 1011\|x_1\|0^{2nq-n} \\ & & [x_1]_4 &:= 1100\|x_1\|0^{2nq-n} \end{aligned}$$

In addition, for all $k \in [q]$, we write $[x_{2k-1}, x_{2k}]_{0\|k}$ to denote the k -th diblock coordinate (x_{2k-1}, x_{2k}) of x . We will mostly use this view, and thus, view the $2qn$ -bit entry as q separate entries of size $2n$ -bit each, and thus, $d_{\mathbf{R}}([x_{2k-1}, x_{2k}]_{0\|k}) \neq \perp$ (or $d_{\mathbf{I}}([x_{2k-1}, x_{2k}]_{0\|k}) \neq \perp$) is well-defined as long as $d_{\mathbf{R}}([x]_0) \neq \perp$ (res. $d_{\mathbf{I}}([x]_0) \neq \perp$ for some $x = (z, (x_{2k-1}, x_{2k}), z')$ where z and z' are $2(k-1)n$ -bit and $2(q-k)n$ -bit strings.

Define

$$\begin{aligned} \tilde{\mathcal{X}}_{\mathbf{R}} &:= \{[x]_0, [x_1]_1, [x_1]_2, [x_1]_3, [x_1]_4 : x = (x_1, \dots, x_{2q}) \in \mathcal{Y}^{2q}\} \\ \tilde{\mathcal{X}}_{\mathbf{I}} &:= \{[x]_0, [x_1]_1, [x_1]_2, [x_1, x_2, x_3]_5, [x_1, x_2, x_3]_6 : x = (x_1, \dots, x_{2q}) \in \mathcal{Y}^{2q}\} \end{aligned}$$

Then it is easy to see that $\mathcal{D}_{\mathbf{R}} = \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{R}}}$ and $\mathcal{D}_{\mathbf{I}} = \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{I}}}$.

5.1 Bad and Good Databases

Let $\mathcal{B}_{\mathbf{R}}$ be the set of databases $d_{\mathbf{R}}$ satisfying one of the following condition: we can find $(x_1, x_2) \neq (x'_1, x'_2) \in \mathcal{Y}^2$ and $v_1, v_2, v'_1, v'_2 \in \mathcal{Y}$ such that

⁶ Disjoint from the other functions due to the first bit.

- for some $k \notin k' \in [q]$, $d_{\mathbf{R}}([x_1, x_2]_{0\parallel k}) \neq \perp$, $d_{\mathbf{R}}([x'_1, x'_2]_{0\parallel k'}) \neq \perp$;
- $([x_1]_1, v_1), ([x'_1]_1, v_1) \in d_{\mathbf{R}}$;
- $([x_2 \oplus v_1]_2, v_2), ([x'_2 \oplus v'_1]_2, v'_2) \in d_{\mathbf{R}}$;
- $x_1 \oplus v_2 = x'_1 \oplus v'_2$;

or we can find $(x_1, x_2) \neq (x'_1, x'_2) \in \mathcal{Y}^2$ and $v_1, v_2, v_3, v'_1, v'_2, v'_3 \in \mathcal{Y}$ such that

- for some $k \notin k' \in [q]$, $d_{\mathbf{R}}([x_1, x_2]_{0\parallel k}) \neq \perp$, $d_{\mathbf{R}}([x'_1, x'_2]_{0\parallel k'}) \neq \perp$;
- $([x_1]_1, v_1), ([x'_1]_1, v_1) \in d_{\mathbf{R}}$;
- $([x_2 \oplus v_1]_2, v_2), ([x'_2 \oplus v'_1]_2, v'_2) \in d_{\mathbf{R}}$;
- $([x_1 \oplus v_2]_3, v_3), ([x'_1 \oplus v'_2]_3, v'_3) \in d_{\mathbf{R}}$;
- $x_2 \oplus v_1 \oplus v_3 = x'_2 \oplus v'_1 \oplus v'_3$;

Next, let $\mathcal{B}_{\mathbf{I}}$ be the set of databases $d_{\mathbf{I}}$ satisfying one of the the following condition: we can find $(x_1, x_2) \neq (x'_1, x'_2) \in \mathcal{Y}^2$ and $v_1, v_2, v'_1, v'_2 \in \mathcal{Y}$

- for some $k \notin k' \in [q]$, $d_{\mathbf{I}}([x_1, x_2]_{0\parallel k}) \neq \perp$, $d_{\mathbf{I}}([x'_1, x'_2]_{0\parallel k'}) \neq \perp$;
- $([x_1]_1, v_1), ([x'_1]_1, v_1) \in d_{\mathbf{I}}$;
- $([x_2 \oplus v_1]_2, v_2), ([x'_2 \oplus v'_1]_2, v'_2) \in d_{\mathbf{I}}$;
- $x_1 \oplus v_2 = x'_1 \oplus v'_2$;

or we can find $(x_1, x_2) \neq (x'_1, x'_2) \in \mathcal{Y}^2$ and $v_1, v_2, v_3, v'_1, v'_2, v'_3 \in \mathcal{Y}$ such that

- for some $k \notin k' \in [q]$, $d_{\mathbf{I}}([x_1, x_2]_{0\parallel k}) \neq \perp$, $d_{\mathbf{I}}([x'_1, x'_2]_{0\parallel k'}) \neq \perp$;
- $([x_1]_1, v_1), ([x'_1]_1, v_1) \in d_{\mathbf{I}}$;
- $([x_2 \oplus v_1]_2, v_2), ([x'_2 \oplus v'_1]_2, v'_2) \in d_{\mathbf{I}}$;
- $([x_1, x_2, x_1 \oplus v_2]_5, v_5), ([x'_1, x'_2, x'_1 \oplus v'_2]_5, v_5) \in d_{\mathbf{I}}$;
- $x_2 \oplus v_1 \oplus v_3 = x'_2 \oplus v'_1 \oplus v'_3$;

Let $\mathcal{G}_{\mathbf{R}} := \mathcal{D}_{\mathbf{R}} \setminus \mathcal{B}_{\mathbf{R}}$ and $\mathcal{G}_{\mathbf{I}} := \mathcal{D}_{\mathbf{I}} \setminus \mathcal{B}_{\mathbf{I}}$. The above definitions mean that in both $\mathcal{G}_{\mathbf{R}}$ and $\mathcal{G}_{\mathbf{I}}$, each u_3 and u_4 is associated with a unique pair (x_1, x_2) . Then it is easy to see that $\mathcal{G}_{\mathbf{R}}$ and $\mathcal{G}_{\mathbf{I}}$ have an obvious bijection $h : \mathcal{G}_{\mathbf{R}} \rightarrow \mathcal{G}_{\mathbf{I}}$ as follows: for each $d_{\mathbf{R}}$ we define $d_{\mathbf{I}} := h(d_{\mathbf{R}})$ such that

- for each $x \in \mathcal{Y}^{2q}$, $d_{\mathbf{I}}([x]_0) = d_{\mathbf{R}}([x]_0)$. Note that, by definition of the oracle, there will be only one entry of this type in both the worlds;
- for each $u_1 \in \mathcal{Y}$, $d_{\mathbf{I}}([u_1]_1) = d_{\mathbf{R}}([u_1]_1)$;
- for each $u_2 \in \mathcal{Y}$, $d_{\mathbf{I}}([u_2]_2) = d_{\mathbf{R}}([u_2]_2)$;
- for each $u_3, u_4 \in \mathcal{Y}$ such that $d_{\mathbf{R}}([u_3]_3) \neq \perp$ and $d_{\mathbf{R}}([u_4]_4) \neq \perp$, find the unique $(x_1, x_2) \in \mathcal{Y}^2$, and define $d_{\mathbf{I}}([x_1, x_2, u_3]_5) = d_{\mathbf{R}}([u_3]_3)$ and $d_{\mathbf{I}}([x_1, x_2, u_4]_6) = d_{\mathbf{R}}([u_4]_4)$.

Then h satisfies the conditions of Lemma 2. To complete the proof, we show that

$$\left(\perp \overset{4q+2}{\rightsquigarrow} \mathcal{B}_{\mathbf{R}}\right) + \left(\perp \overset{4q+2}{\rightsquigarrow} \mathcal{B}_{\mathbf{I}}\right) \leq 2\sqrt{\frac{q^6}{2^n}} + 4\sqrt{\frac{q^5}{2^n}}.$$

5.2 Sequence of Actions

We ignore the dummy call actions, as the transition from a good to bad database is independent of the output of this operator.

Recall that the q non-adaptive queries can be represented by a single q -fold query to be evaluated sequentially.

ACTION OF f_1 . For $i \in \{4k+2 : 0 \leq k \leq q-1\}$, we bound the transition capacity $\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket$. For any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i-1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([x'_1]_1) \oplus d_{\mathbf{R}}([u'_3]_3) \oplus d_{\mathbf{R}}([u_3]_3) \oplus x_2 \oplus x'_2 \mid \mathbf{E}\},$$

where \mathbf{E} denotes the predicate $d_{\mathbf{R}}([u_3]_3) \neq \perp, d_{\mathbf{R}}([u'_3]_3) \neq \perp, d_{\mathbf{R}}([x, x_2]_{0||*}) \neq \perp, d_{\mathbf{R}}([x'_1, x'_2]_{0||*}) \neq \perp$.

There are at most q choices for (x'_1, x'_2) , $\lceil (i-1)/4 \rceil$ choices for each of u_3 and u'_3 , and at most q choices for x_2 , so $|\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}}| \leq q^2 \lceil (i-1)/3 \rceil^2 \leq q^4$, and from there using Lemma 1 we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^4}{2^n}}, \quad \forall i \in \{4k+2 : 0 \leq k \leq q-1\}. \quad (17)$$

By the same arguments we can also show that

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^4}{2^n}}, \quad \forall i \in \{4k+2 : 0 \leq k \leq q-1\}. \quad (18)$$

ACTION OF f_2 . Next consider the transition capacity $\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket$ for $i \in \{4k+3 : 0 \leq k \leq q-1\}$. For any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i-1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([u'_2]_2) \oplus x_1 \oplus x'_1 \mid \mathbf{E}\},$$

where \mathbf{E} denotes the predicate $d_{\mathbf{R}}([u'_2]_2) \neq \perp, d_{\mathbf{R}}([x_1, x_2]_{0||*}) \neq \perp, d_{\mathbf{R}}([x'_1, x'_2]_{0||*}) \neq \perp$. Again, there are at most $\lceil (i-1)/4 \rceil$ choices for u'_2 and at most q^2 choices for (x_1, x'_1) . Thus, from Lemma 1, we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k+3 : 0 \leq k \leq q-1\}. \quad (19)$$

ACTION OF f_3 (RESP. F_3): For $i \in \{4k+4 : 1 \leq k \leq q-1\}$, for any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i-1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([x_1]_1) \oplus d_{\mathbf{R}}([x'_1]_1) \oplus d_{\mathbf{R}}([u'_3]_3) \oplus x_2 \oplus x'_2 \mid \mathbf{E}\},$$

where \mathbf{E} denotes the predicate $d_{\mathbf{R}}([x_1]_1), d_{\mathbf{R}}([x'_1]_1), d_{\mathbf{R}}([u_3]_3) \neq \perp, d_{\mathbf{R}}([x_1, x_2]_{0||*}) \neq \perp, d_{\mathbf{R}}([x'_1, x'_2]_{0||*}) \neq \perp$. There are at most $\lceil (i-1)/4 \rceil$

choices for u'_3 and at most q^2 choices for $((x_1, x_2), (x'_1, x'_2))$. Since the analysis is identical in both the worlds, by using Lemma 1, we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k+4 : 0 \leq k \leq q-1\} \quad (20)$$

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k+4 : 0 \leq k \leq q-1\} \quad (21)$$

ACTION OF f_4 (RESP. F_4): Since the property $\mathcal{B}_{\mathbf{R}}$ (resp. $\mathcal{B}_{\mathbf{I}}$) is independent of the output of f_4 (resp. F_4) and the database is good right before the action, we have $\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \emptyset$. Thus,

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket = 0, \quad \forall i \in \{4k+5 : 0 \leq k \leq q-1\} \quad (22)$$

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket = 0, \quad \forall i \in \{4k+5 : 0 \leq k \leq q-1\} \quad (23)$$

Summing over the $4q+2$ actions using (17)–(23) gives

$$\left(\perp \overset{4q+2}{\rightsquigarrow} \mathcal{B}_{\mathbf{R}} \right) \leq 2\sqrt{\frac{10q^5}{2^n}} + \sqrt{\frac{10q^6}{2^n}}, \quad \left(\perp \overset{4q+2}{\rightsquigarrow} \mathcal{B}_{\mathbf{I}} \right) \leq 2\sqrt{\frac{10q^5}{2^n}} + \sqrt{\frac{10q^6}{2^n}}. \quad (24)$$

Adding the two inequalities completes the proof of Theorem 2.

5.3 The Problem with the Adaptive Setting

A closer look at the non-adaptive proof serves to show why a similar proof is difficult to achieve in the adaptive setting. The dummy call is used to record all the q non-adaptive queries of the adversary in the database, before LR_4 is applied to each of them sequentially. This enables us to argue that the oracle knows all q queries at the time of each of the subsequent actions (f_1, f_2, f_3 etc.) which in turn helps in upper bounding the bad norm to a non-trivial value.

The proof hinges on the characterisation as bad of any database which has a ‘collision’ on the f input in either of the last two rounds, i.e., collisions on $x_1 \oplus v_2$ or $x_2 \oplus v_1 \oplus v_3$ for different database entries. Specifically, this implies that certain later values of x_1 or x_2 can always make the database go bad irrespective of earlier choices of v_1, v_2 , or v_3 . As a concrete example, recall (from Sect. 5.1) that a database is (also) considered bad if:

- for some $k \neq k' \in [q]$, $d_{\mathbf{R}}([x_1, x_2]_{0\|k}), d_{\mathbf{R}}([x'_1, x'_2]_{0\|k'}) \neq \perp$ (i.e. the adversary has made these two queries).
- $([x_1]_1, v_1), ([x'_1]_1, v'_1) \in d_{\mathbf{R}}$; (f_1 has been evaluated over x_1 and x'_1)
- $([x_2 \oplus v_1]_2, v_2), ([x'_2 \oplus v'_1]_2, v'_2) \in d_{\mathbf{R}}$; (f_2 has been evaluated over $x_2 \oplus v_1$ and $x'_2 \oplus v'_1$)

– $x_1 \oplus v_2 = x'_1 \oplus v'_2$; (there is an input-collision on f_3)

Now, in the context of f_1 's action, comparing the above definition with the previous proofs (specifically see the discussion around (15) and (16)), one can see that conditions 1 and 3 are missing in previous proofs. This is because it is impossible for the oracle to detect the queries made by the adversary, as at any given instant, it can only see the database entries, nothing less and nothing more. As a result, the norm bound becomes trivial. On the other hand, in our case, specifically because condition 1 can be checked at all times (once the dummy call is executed), condition 3 is also well-defined. As a result, as shown in (17) and (18), the norm bound is non-trivial.

At the same time, the dummy call must be erased before the oracle returns an output to the adversary. Otherwise, this perturbs the state, which can be detected by the adversary. So, this approach only works in non-adaptive games which can be modelled as an adversary making a single “big” query (consisting of q usual queries) to the oracle and the oracle returning a single “big” output (consisting of q usual outputs). An adaptive game, on the other hand, does not adhere to such simplifications. More specifically, since future values of x_1 and x_2 are directly under the adversary's control and are not known to the oracle in advance, the amplitude of such events cannot be bounded using known techniques. In the HI framework, this problem appears as the trivialization of the norm (see Sect. 4). In the BCEJ framework, this observation implies that databases can go bad *between* two actions, something that the framework does not account for. In the non-adaptive setting, however, the oracle knows in advance the future values of x_1 and x_2 , and the outputs of f can accordingly be classified as ‘bad’ and bounded at the time of the action of f .

Lastly, we remark that this is not a problem specific to Luby-Rackoff, but is inherent to any proof for which the definition of bad databases is in terms of an input that the adversary can adaptively choose. We have also noticed this error in other proofs. For example, in [5], the security proofs of TNT, LRWQ and LRQ suffer from this problem, and do not hold in the adaptive setting. While for TNT and LRWQ this seems to be more of a definitional problem, since the bad events can be defined directly in terms of the database entries (though possibly leading to a slightly worse bound), for the LRQ proof this looks like a more fundamental issue that does not admit an easy fix. We spotted similar flaws in other works like the proof of LRWQ in [25] and the tight security proof for TNT [34]. While the former seems to be fixable, the latter is again a fundamental issue.

6 IND-qCPA Security of Misty

6.1 The Misty Constructions

For some $r \geq 1$ and $f_1, \dots, f_r : \{0, 1\}^n \rightarrow \{0, 1\}^n$, we define

– $g^L : [r] \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ by the mapping:

$$(i, x_1, x_2) \mapsto (x_2, x_2 \oplus f_i(x_1)),$$

– $g^R : [r] \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ by the mapping:

$$(i, x_1, x_2) \mapsto (x_2 \oplus f_i(x_1), f_i(x_1)),$$

and write $g_i^L(\cdot, \cdot) := g^L(i, \cdot, \cdot)$ and $g_i^R(\cdot, \cdot) := g^R(i, \cdot, \cdot)$.

MistyL Construction: The r -round MistyL, denoted MistyL_r is defined as:

$$(x_1, x_2) \mapsto g_r^L \circ \dots \circ g_1^L(x_1, x_2). \tag{25}$$

MistyR Construction: The r -round MistyR construction, denoted MistyR_r is defined as:

$$(x_1, x_2) \mapsto g_r^R \circ \dots \circ g_1^R(x_1, x_2). \tag{26}$$

For all $i \in [r]$, we write:

- $x^{i-1} := (x_1^{i-1}, x_2^{i-1})$ to denote the input to g_i , where $x^0 := x = (x_1, x_2)$, denotes the input to $\text{Misty}\{\text{L}|\text{R}\}_r$.
- (u_i, v_i) to denote the input-output tuple corresponding to f_i .
- $y = (y_1, y_2) := (x_1^r, x_2^r)$ to denote the output of $\text{Misty}\{\text{L}|\text{R}\}_r$.

6.2 IND-qCPA Security of MistyR

We prove the following IND-qCPA bound for MistyR_4 .

Theorem 3. *Suppose $f_1, f_2, f_3, f_4 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are four mutually independent uniform random functions. Then, for any $q \geq 0$, and any quantum adversary \mathcal{A} that makes at most q queries, we have*

$$\text{Adv}_{\text{MistyR}_4}^{\text{qcpa}}(\mathcal{A}) = O\left(\sqrt{\frac{q^5}{2^n}}\right).$$

Proof. Let $F_3, F_4 : \{0, 1\}^{3n} \rightarrow \{0, 1\}^n$ be two uniform random functions. Define

$$\begin{aligned} G_3^R(x_1, x_2, x'_1, x'_2) &:= (x'_2 \oplus F_3(x_1, x_2, x'_1), F_3(x_1, x_2, x'_1)) \\ G_4^R(x_1, x_2, x'_1, x'_2) &:= (x'_2 \oplus F_4(x_1, x_2, x'_1), F_4(x_1, x_2, x'_1)) \end{aligned}$$

for any $(x_1, x_2, x'_1, x'_2) \in \{0, 1\}^{4n}$. We define the hybrid random function $\widetilde{\text{MistyR}}_4$ as (see also Fig. 3):

$$\widetilde{\text{MistyR}}_4(x_1, x_2) := G_4^L(x_1, x_2, G_3^L(x_1, x_2, \text{MistyR}_2(x_1, x_2))).$$

Then, it is easy to see that $\widetilde{\text{MistyR}}_4$ is indistinguishable to a uniform random function $\Gamma : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. So, it is sufficient to bound the distance between MistyR_4 and $\widetilde{\text{MistyR}}_4$.

Let $\mathcal{X} := \{0, 1\}^{3n+3}$, and let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a $(3n + 3)$ -bit-to- n -bit uniform random function. We implement f through \mathbf{cO} defined over $\mathcal{H}[\mathcal{X}] \otimes \mathcal{H}[\mathcal{Y}] \otimes \mathbb{D}$. For each $x, y, z \in \mathcal{Y}$,

$$\begin{aligned} f_1(x) &= f(000\|x\|0^{2n}) \\ f_2(x) &= f(001\|x\|0^{2n}) \\ f_3(x) &= f(010\|x\|0^{2n}) \end{aligned}$$

$$\begin{aligned} f_4(x) &= f(011\|x\|0^{2n}) \\ F_3(x, y, z) &= f(100\|x\|y\|z) \\ F_4(x, y, z) &= f(101\|x\|y\|z). \end{aligned}$$

The distinctness of the first three bits ensures that $f_1, f_2, f_3, f_4, F_3, F_4$ are all independent, and they can be implemented by the prefix oracle. We do not give the implementation explicitly as it is obvious. This setup allows us to use a single

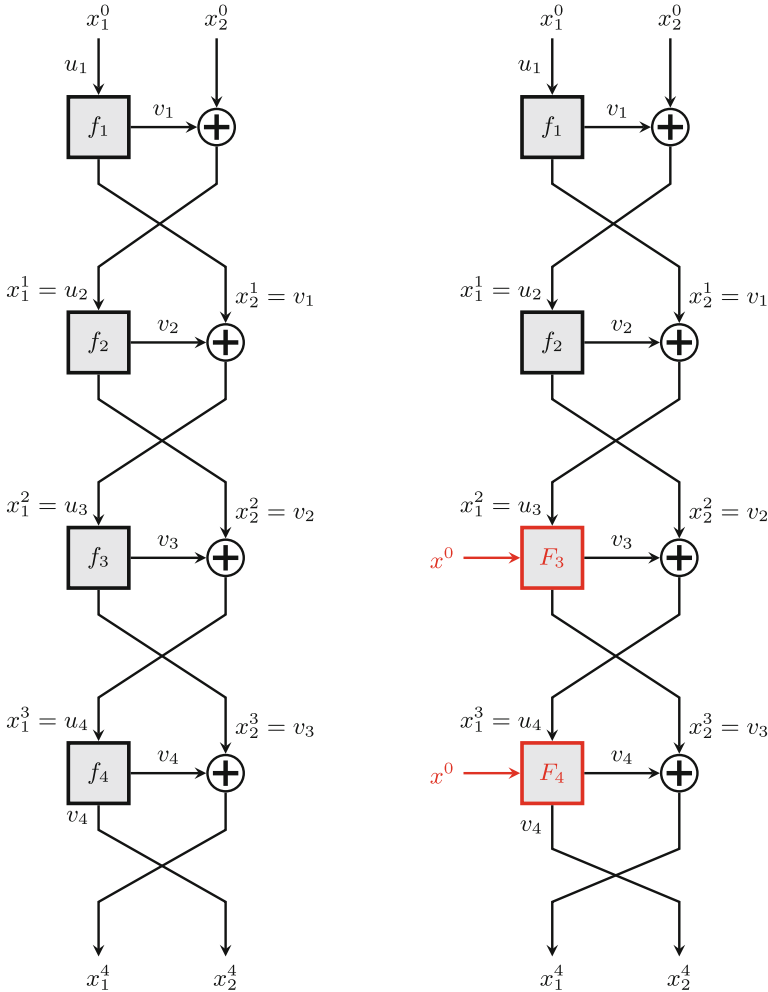


Fig. 3. MistyR_4 (left) vs the hybrid random function, $\widetilde{\text{MistyR}}_4$ (right).

database $d_f : \mathcal{X} \rightarrow \mathcal{Z}$ to keep track of $f_1, f_1, f_2, f_3, f_4, F_3$ and F_4 ; we refer to this database as $d_{\mathbf{R}}$ in the real world (tracking f_1, f_2, f_3 and f_4) and $d_{\mathbf{I}}$ in the ideal world (tracking f_1, f_2, F_3 and F_4). Let $\mathcal{D}_{\mathbf{R}}$ (resp. $\mathcal{D}_{\mathbf{I}}$) be the set of all possible choices for $d_{\mathbf{R}}$ (resp. $d_{\mathbf{I}}$). Let

$$[x]_1 := 000\|x\|0^{2n}, [x]_2 := 001\|x\|0^{2n},$$

$$[x]_3 := 010\|x\|0^{2n}, [x]_4 := 011\|x\|0^{2n}.$$

and define the sets

$$\tilde{\mathcal{X}}_{\mathbf{R}} := \{[x]_1, [x]_2, [x]_3, [x]_4 \mid x \in \mathcal{Y}\},$$

$$\tilde{\mathcal{X}}_{\mathbf{I}} := \{[x]_1, [x]_2, (100\|x\|x'\|y), (101\|x\|x'\|y) \mid x, x', y \in \mathcal{Y}\}.$$

Then it is easy to see that $\mathcal{D}_{\mathbf{R}} = \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{R}}}$ and $\mathcal{D}_{\mathbf{I}} = \mathcal{D}|_{\tilde{\mathcal{X}}_{\mathbf{I}}}$.

Let $\mathcal{B}_{\mathbf{R}}$ be the set of databases $d_{\mathbf{R}}$ satisfying one of the two following conditions: we can find $u_1, u'_1, u_2, u'_2, v_1, v'_1, v_2, v'_2 \in \mathcal{Y}$ such that

1. $([u_1]_1, v_1), ([u'_1]_1, v'_1), ([u_2]_2, v_2), ([u'_2]_2, v'_2) \in d_{\mathbf{R}}$;
2. $v_2 \oplus v_1 = v'_2 \oplus v'_1$;

or we can find $u_1, u'_1, u_2, u'_2, v_1, v'_1, v_2, v'_2, v_3, v'_3 \in \mathcal{Y}$ such that

1. $([u_1]_1, v_1), ([u'_1]_1, v'_1), ([u_2]_2, v_2), ([u'_2]_2, v'_2),$
 $[v_2 \oplus v_1]_3, v_3), ([v'_2 \oplus v'_1]_3, v'_3) \in d_{\mathbf{R}}$;
2. $v_3 \oplus v_2 = v'_3 \oplus v'_2$;

Next, let $\mathcal{B}_{\mathbf{I}}$ be the set of databases $d_{\mathbf{I}}$ satisfying one of the two following conditions: we can find $u_1, u'_1, u_2, u'_2, v_1, v'_1, v_2, v'_2 \in \mathcal{Y}$ such that

1. $([u_1]_1, v_1), ([u'_1]_1, v'_1), ([u_2]_2, v_2), ([u'_2]_2, v'_2) \in d_{\mathbf{I}}$;
2. $v_2 \oplus v_1 = v'_2 \oplus v'_1$;

or we can find $u_1, u'_1, u_2, u'_2, v_1, v'_1, v_2, v'_2, v_3, v'_3 \in \mathcal{Y}$ such that

1. $([u_1]_1, v_1), ([u'_1]_1, v'_1), ([u_2]_2, v_2), ([u'_2]_2, v'_2),$
 $(100\|u_1\|v_1 \oplus u_2\|v_2 \oplus v_1, v_3), (100\|u'_1\|v'_1 \oplus u'_2\|v'_2 \oplus v'_1, v'_3) \in d_{\mathbf{I}}$;
2. $v_3 \oplus v_2 = v'_3 \oplus v'_2$;

Let $\mathcal{G}_{\mathbf{R}} := \mathcal{D}_{\mathbf{R}} \setminus \mathcal{B}_{\mathbf{R}}$ and $\mathcal{G}_{\mathbf{I}} := \mathcal{D}_{\mathbf{I}} \setminus \mathcal{B}_{\mathbf{I}}$. Suppose $d_{\mathbf{R}} \in \mathcal{G}_{\mathbf{R}}$ and $d_{\mathbf{I}} \in \mathcal{G}_{\mathbf{I}}$. Then each u_3 for which there exists v_3 such that $([u_3]_3, v_3) \in d_{\mathbf{R}}$ is associated with a unique pair $([u_1]_1, v_1), ([u_2]_2, v_2) \in d_{\mathbf{R}}$ such that $u_3 = v_1 \oplus v_2$, and each u_4 for which there exists v_4 such that $([u_4]_4, v_4) \in d_{\mathbf{R}}$ is associated with a unique triple $([u_1]_1, v_1), ([u_2]_2, v_2), ([u_3]_3, v_3) \in d_{\mathbf{R}}$ such that $u_3 = v_1 \oplus v_2$ and $u_4 = v_2 \oplus v_3$.

Similarly, each u_3 for which there exist x_1, x_2, v_3 such that $(100\|x_1\|x_2\|u_3, v_3) \in d_{\mathbf{I}}$ is associated with a unique pair $([u_1]_1, v_1), ([u_2]_2, v_2) \in d_{\mathbf{I}}$ such that $u_3 = v_1 \oplus v_2$, and this pair also satisfies $x_1 = u_1, x_2 = v_1 \oplus u_2$; and each u_4 for which there exist x_1, x_2, v_4 such that $(101\|x_1\|x_2\|u_4, v_4) \in d_{\mathbf{I}}$ is associated with a unique triple $([u_1]_1, v_1), ([u_2]_2, v_2), (100\|x_1\|x_2\|u_3, v_3) \in d_{\mathbf{I}}$ such that $u_3 = v_1 \oplus v_2$ and $u_4 = v_2 \oplus v_3$, and this triple also satisfies $x_1 = u_1, x_2 = v_1 \oplus u_2$.

Then we can define the bijection $h : \mathcal{G}_{\mathbf{R}} \rightarrow \mathcal{G}_{\mathbf{I}}$ as follows: for each $d_{\mathbf{R}} \in \mathcal{G}_{\mathbf{R}}$ we define $d_{\mathbf{I}} := h(d_{\mathbf{R}})$ such that

- for each $u_1 \in \mathcal{Y}$, $d_{\mathbf{I}}([u_1]_1) = d_{\mathbf{R}}([u_1]_1)$;
- for each $u_2 \in \mathcal{Y}$, $d_{\mathbf{I}}([u_2]_2) = d_{\mathbf{R}}([u_2]_2)$;
- for each $x_1, x_2 \in \mathcal{Y}$ and the associated (u_3, u_4) , $d_{\mathbf{I}}(100\|x_1\|x_2\|u_3) = d_{\mathbf{R}}([u_3]_3)$ and $d_{\mathbf{I}}(101\|x_1\|x_2\|u_4) = d_{\mathbf{R}}([u_4]_4)$.

Then h satisfies the conditions of Lemma 2. To complete the proof of Theorem 3, we just need to show that $\left(\perp \overset{4q}{\rightsquigarrow} \mathcal{B}_{\mathbf{R}}\right) + \left(\perp \overset{4q}{\rightsquigarrow} \mathcal{B}_{\mathbf{I}}\right) \leq (4 + 2\sqrt{2})\sqrt{10q^5/2^n}$.

Sequence of Actions. Each query by the adversary to its oracle results in a sequence of four queries to f , one each to f_1 , f_2 , and one to f_3 and f_4 in the real world or F_3 and F_4 in the ideal world, in that order. We view the query response phase as a sequence of $4q$ (possibly duplicate) *actions* and analyze the transition capacity at each action.

ACTION OF f_1 : For $i \in \{4k + 1 : 0 \leq k \leq q - 1\}$, we first look at the transition capacity $\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket$. For any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i - 1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([u_1]_1) \oplus d_{\mathbf{R}}([u_2]_2) \oplus d_{\mathbf{R}}([u'_2]_2) \mid d_{\mathbf{R}}([u_1]_1) \neq \perp, \\ d_{\mathbf{R}}([u_2]_2) \neq \perp, d_{\mathbf{R}}([u'_2]_2) \neq \perp\}.$$

There are at most $\lceil (i-1)/4 \rceil^3$ choices for the triple (u_2, u'_1, u'_2) , so $|\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}}| \leq \lceil (i-1)/4 \rceil^3 \leq q^3$, and from there using Lemma 1 we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k + 1 : 0 \leq k \leq q - 1\}. \quad (27)$$

By the same arguments we can also show that

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k + 1 : 0 \leq k \leq q - 1\}. \quad (28)$$

ACTION OF f_2 : Next we look at the transition capacity $\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket$ for $i \in \{4k + 2 : 0 \leq k \leq q - 1\}$. For any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i - 1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([u_1]_1) \oplus d_{\mathbf{R}}([u'_1]_1) \oplus d_{\mathbf{R}}([u'_2]_2) \mid d_{\mathbf{R}}([u_1]_1) \neq \perp, d_{\mathbf{R}}([u'_1]_1) \neq \perp, \\ d_{\mathbf{R}}([u'_2]_2) \neq \perp\} \cup \{d_{\mathbf{R}}([u_3]_3) \oplus d_{\mathbf{R}}([u'_3]_3) \oplus d_{\mathbf{R}}([u'_2]_2) \mid \\ d_{\mathbf{R}}([u_3]_3) \neq \perp, d_{\mathbf{R}}([u'_3]_3) \neq \perp, d_{\mathbf{R}}([u'_2]_2) \neq \perp\}.$$

Again, there are at most $\lceil (i-1)/4 \rceil^3$ choices for each of the triples (u_2, u'_1, u'_2) and (u_3, u'_2, u'_3) , and arguing as before we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \hookrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{20q^3}{2^n}}, \quad \forall i \in \{4k + 2 : 0 \leq k \leq q - 1\}. \quad (29)$$

By the same arguments we can also show that

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket \leq \sqrt{\frac{20q^3}{2^n}}, \quad \forall i \in \{4k+2 : 0 \leq k \leq q-1\}. \quad (30)$$

ACTION OF f_3 (RESP. F_3): Next we look at the transition capacity $\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket$ for $i \in \{4k+3 : 0 \leq k \leq q-1\}$. For any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i-1$ and any $x \in \mathcal{Y}$, we have

$$\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \leftrightarrow \mathcal{B}_{\mathbf{R}}} = \{d_{\mathbf{R}}([u_2]_2) \oplus d_{\mathbf{R}}([u'_2]_2) \oplus d_{\mathbf{R}}([u'_3]_3) \mid d_{\mathbf{R}}([u_2]_2) \neq \perp, \\ d_{\mathbf{R}}([u'_2]_2) \neq \perp, d_{\mathbf{R}}([u'_3]_3) \neq \perp\}.$$

Again, there are at most $\lceil (i-1)/4 \rceil^3$ choices for the pair (u_2, u'_2, u'_3) , and arguing as before we have

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k+3 : 0 \leq k \leq q-1\}. \quad (31)$$

By the same arguments we can also show that

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket \leq \sqrt{\frac{10q^3}{2^n}}, \quad \forall i \in \{4k+3 : 0 \leq k \leq q-1\}. \quad (32)$$

ACTION OF f_4 (RESP. F_4): Finally, for $i \in \{4k : 1 \leq k \leq q\}$, for any $d_{\mathbf{R}}$ with $|d_{\mathbf{R}}| \leq i-1$ (resp. any $d_{\mathbf{I}}$ with $|d_{\mathbf{I}}| \leq i-1$) and any $x \in \mathcal{Y}$, since the property $\mathcal{B}_{\mathbf{R}}$ (resp. $\mathcal{B}_{\mathbf{I}}$) does not depend on $d_{\mathbf{R}}([x]_4)$ (resp. $d_{\mathbf{I}}(101\|x_1\|x_2\|x)$), we have $\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{R}}^c \leftrightarrow \mathcal{B}_{\mathbf{R}}} = \emptyset$ (resp. $\mathcal{S}_{x,d}^{\mathcal{B}_{\mathbf{I}}^c \leftrightarrow \mathcal{B}_{\mathbf{I}}} = \emptyset$). Thus,

$$\llbracket \mathcal{B}_{\mathbf{R}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{R}[\leq i]} \rrbracket = 0, \quad \forall i \in \{4k : 0 \leq k \leq q-1\}, \quad (33)$$

and also,

$$\llbracket \mathcal{B}_{\mathbf{I}[\leq i-1]}^c \leftrightarrow \mathcal{B}_{\mathbf{I}[\leq i]} \rrbracket = 0, \quad \forall i \in \{4k : 0 \leq k \leq q-1\}. \quad (34)$$

Summing over the $4q$ actions using (27)–(34) gives

$$\left(\perp \overset{4q}{\rightsquigarrow} \mathcal{B}_{\mathbf{R}}\right) \leq (2 + \sqrt{2})\sqrt{\frac{10q^5}{2^n}}, \quad \left(\perp \overset{4q}{\rightsquigarrow} \mathcal{B}_{\mathbf{I}}\right) \leq (2 + \sqrt{2})\sqrt{\frac{10q^5}{2^n}}. \quad (35)$$

Adding the two inequalities completes the proof of Theorem 3.

6.3 IND-qCPA Security of MistyL

We prove the following IND-qCPA bound for MistyL₅.

Theorem 4. *Suppose $f_1, f_2, f_3, f_4, f_5 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are five mutually independent uniform random functions. Then, for any $q \geq 0$, and any quantum adversary \mathcal{A} that makes at most q queries, we have*

$$\text{Adv}_{\text{MistyL}_5}^{\text{qcpa}}(\mathcal{A}) = O\left(\sqrt{\frac{q^7}{2^n}}\right).$$

A proof of this theorem is available in the full version of this paper [6].

7 Conclusion

In this work, we uncover a flaw in the proof of quantum security for the Luby-Rackoff, TNT, LRWQ and LRQ constructions. While TNT and LRWQ might still be proven secure (most likely with a degraded bound), the issue in the other cases seems inherent to the proof techniques that were used. In particular, for the technique to work, it is critical that bad databases are only described with information that is actually present in the database. For some constructions, notably the Luby-Rackoff and LRQ constructions, a part of the input to the construction will never appear in the database directly which means that it cannot be used to characterize bad databases. On a positive note, we restore the security of the 4-round Luby-Rackoff construction in the *non-adaptive* setting, and prove the security of the 4-round MistyR and 5-round MistyL constructions.

References

1. A. Ramachandra Rao, P.B.: Linear Algebra. Hindustan Book Agency (2000). <https://doi.org/10.1007/978-93-86279-01-9>
2. Bellare, M., Krovetz, T., Rogaway, P.: Luby-rackoff backwards: Increasing security by making block ciphers non-invertible. In: Nyberg, K. (ed.) Advances in Cryptology - EUROCRYPT '98, Proceeding. Lecture Notes in Computer Science, vol. 1403, pp. 266–280. Springer (1998). <https://doi.org/10.1007/BFB0054132>
3. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT 2006, Proceedings. Lecture Notes in Computer Science, vol. 4004, pp. 409–426. Springer (2006). https://doi.org/10.1007/11761679_25
4. Bhaumik, R., Bonnetain, X., Chailloux, A., Leurent, G., Naya-Plasencia, M., Schrottenloher, A., Seurin, Y.: QCB: Efficient quantum-secure authenticated encryption. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part I. LNCS, vol. 13090, pp. 668–698. Springer, Heidelberg (Dec 2021). https://doi.org/10.1007/978-3-030-92062-3_23
5. Bhaumik, R., Cogliati, B., Ethan, J., Jha, A.: On quantum secure compressing pseudorandom functions. In: Guo, J., Steinfeld, R. (eds.) ASIACRYPT 2023, Part III. LNCS, vol. 14440, pp. 34–66. Springer, Heidelberg (Dec 2023). https://doi.org/10.1007/978-981-99-8727-6_2
6. Bhaumik, R., Cogliati, B., Ethan, J., Jha, A.: Mind the bad norms: Revisiting compressed oracle-based quantum indistinguishability proofs. Cryptology ePrint Archive, Report 2024/1478 (2024), <https://eprint.iacr.org/2024/1478>
7. Boneh, D., Zhandry, M.: Quantum-secure message authentication codes. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 592–608. Springer, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_35
8. Bonnetain, X., Naya-Plasencia, M.: Hidden shift quantum cryptanalysis and implications. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part I. LNCS, vol. 11272, pp. 560–592. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03326-2_19
9. Bonnetain, X., Naya-Plasencia, M., Schrottenloher, A.: On quantum slide attacks. In: Paterson, K.G., Stebila, D. (eds.) SAC 2019. LNCS, vol. 11959, pp. 492–519. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-38471-5_20

10. Bonnetain, X., Naya-Plasencia, M., Schrottenloher, A.: Quantum security analysis of AES. *IACR Trans. Symm. Cryptol.* **2019**(2), 55–93 (2019). <https://doi.org/10.13154/tosc.v2019.i2.55-93>
11. Bonnetain, X., Schrottenloher, A., Sibleyras, F.: Beyond quadratic speedups in quantum attacks on symmetric schemes. In: Dunkelman, O., Dziembowski, S. (eds.) *EUROCRYPT 2022, Part III*. LNCS, vol. 13277, pp. 315–344. Springer, Heidelberg (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_12
12. Chailloux, A., Naya-Plasencia, M., Schrottenloher, A.: An efficient quantum collision search algorithm and implications on symmetric cryptography. In: Takagi, T., Peyrin, T. (eds.) *ASIACRYPT 2017, Part II*. LNCS, vol. 10625, pp. 211–240. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70697-9_8
13. Chung, K.M., Fehr, S., Huang, Y.H., Liao, T.N.: On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In: Canteaut, A., Standaert, F.X. (eds.) *EUROCRYPT 2021, Part II*. LNCS, vol. 12697, pp. 598–629. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77886-6_21
14. Cogliati, B., Dutta, A., Nandi, M., Patarin, J., Saha, A.: Proof of mirror theory for a wide range of ξ_{\max} . In: Hazay, C., Stam, M. (eds.) *Advances in Cryptology - EUROCRYPT 2023, Proceedings, Part IV*. Lecture Notes in Computer Science, vol. 14007, pp. 470–501. Springer (2023). https://doi.org/10.1007/978-3-031-30634-1_16
15. Cogliati, B., Seurin, Y.: EWCDM: An efficient, beyond-birthday secure, nonce-misuse resistant MAC. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part I*. LNCS, vol. 9814, pp. 121–149. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53018-4_5
16. Czajkowski, J., Hülsing, A., Schaffner, C.: Quantum indistinguishability of random sponges. In: Boldyreva, A., Micciancio, D. (eds.) *CRYPTO 2019, Part II*. LNCS, vol. 11693, pp. 296–325. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_11
17. Dai, W., Hoang, V.T., Tessaro, S.: Information-theoretic indistinguishability via the chi-squared method. In: Katz, J., Shacham, H. (eds.) *CRYPTO 2017, Part III*. LNCS, vol. 10403, pp. 497–523. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_17
18. Dinur, I.: Tight indistinguishability bounds for the XOR of independent random permutations by fourier analysis. In: Joye, M., Leander, G. (eds.) *Advances in Cryptology - EUROCRYPT 2024, Proceedings, Part I*. Lecture Notes in Computer Science, vol. 14651, pp. 33–62. Springer (2024). https://doi.org/10.1007/978-3-031-58716-0_2
19. Gouget, A., Patarin, J., Toulemonde, A.: (Quantum) cryptanalysis of misty schemes. In: Hong, D. (ed.) *ICISC 20*. LNCS, vol. 12593, pp. 43–57. Springer, Heidelberg (Dec 2020). https://doi.org/10.1007/978-3-030-68890-5_3
20. Grassi, L., Naya-Plasencia, M., Schrottenloher, A.: Quantum algorithms for the k -xor problem. In: Peyrin, T., Galbraith, S. (eds.) *ASIACRYPT 2018, Part I*. LNCS, vol. 11272, pp. 527–559. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03326-2_18
21. Hall, C., Wagner, D.A., Kelsey, J., Schneier, B.: Building prfs from prps. In: Krawczyk, H. (ed.) *Advances in Cryptology - CRYPTO '98, Proceedings*. Lecture Notes in Computer Science, vol. 1462, pp. 370–389. Springer (1998). <https://doi.org/10.1007/BFB0055742>

22. Hoang, V.T., Tessaro, S.: Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 3–32. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53018-4_1
23. Hosoyamada, A., Iwata, T.: 4-round Luby-Rackoff construction is a qPRP. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part I. LNCS, vol. 11921, pp. 145–174. Springer, Heidelberg (Dec 2019). https://doi.org/10.1007/978-3-030-34578-5_6
24. Hosoyamada, A., Iwata, T.: On tight quantum security of HMAC and NMAC in the quantum random oracle model. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 585–615. Springer, Heidelberg, Virtual Event (Aug 2021). https://doi.org/10.1007/978-3-030-84242-0_21
25. Hosoyamada, A., Iwata, T.: Provably quantum-secure tweakable block ciphers. IACR Trans. Symm. Cryptol. **2021**(1), 337–377 (2021). <https://doi.org/10.46586/tosc.v2021.i1.337-377>
26. Hosoyamada, A., Sasaki, Y., Xagawa, K.: Quantum multicollision-finding algorithm. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part II. LNCS, vol. 10625, pp. 179–210. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70697-9_7
27. Hosoyamada, A., Yasuda, K.: Building quantum-one-way functions from block ciphers: Davies-Meyer and Merkle-Damgård constructions. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part I. LNCS, vol. 11272, pp. 275–304. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03326-2_10
28. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Breaking symmetric cryptosystems using quantum period finding. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 207–237. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53008-5_8
29. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Quantum differential and linear cryptanalysis. IACR Trans. Symm. Cryptol. **2016**(1), 71–94 (2016). <https://doi.org/10.13154/tosc.v2016.i1.71-94>, <https://tosc.iacr.org/index.php/ToSC/article/view/536>
30. Kuwakado, H., Morii, M.: Quantum distinguisher between the 3-round feistel cipher and the random permutation. In: IEEE International Symposium on Information Theory, ISIT 2010, Proceedings. pp. 2682–2685. IEEE (2010). <https://doi.org/10.1109/ISIT.2010.5513654>
31. Kuwakado, H., Morii, M.: Security on the quantum-type even-mansour cipher. In: International Symposium on Information Theory and its Applications, ISITA 2012, Proceedings. pp. 312–316. IEEE (2012), <https://ieeexplore.ieee.org/document/6400943/>
32. Lai, X.: On the Design and Security of Block Ciphers. Ph.D. thesis, ETH Zürich (1992)
33. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM J. Comput. **17**(2), 373–386 (1988). <https://doi.org/10.1137/0217022>
34. Mao, S., Zhang, Z., Hu, L., Li, L., Wang, P.: Quantum security of tnt. Cryptology ePrint Archive, Paper 2023/1280 (2023), <https://eprint.iacr.org/2023/1280>, <https://eprint.iacr.org/2023/1280>
35. Matsui, M.: The first experimental cryptanalysis of the data encryption standard. In: Desmedt, Y. (ed.) CRYPTO'94. LNCS, vol. 839, pp. 1–11. Springer, Heidelberg (Aug 1994). https://doi.org/10.1007/3-540-48658-5_1

36. Mennink, B., Neves, S.: Encrypted Davies-Meyer and its dual: Towards optimal security using mirror theory. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 556–583. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_19
37. Nachev, V., Patarin, J., Treger, J.: Generic attacks on misty schemes. In: Abdalla, M., Barreto, P.S.L.M. (eds.) Progress in Cryptology - LATINCRYPT 2010, Proceedings. Lecture Notes in Computer Science, vol. 6212, pp. 222–240. Springer (2010). https://doi.org/10.1007/978-3-642-14712-8_14
38. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information (10th Anniversary edition). Cambridge University Press (2010)
39. Patarin, J.: Security of random feistel schemes with 5 or more rounds. In: Franklin, M.K. (ed.) Advances in Cryptology - CRYPTO 2004, Proceedings. Lecture Notes in Computer Science, vol. 3152, pp. 106–122. Springer (2004). https://doi.org/10.1007/978-3-540-28628-8_7
40. Patarin, J.: The "coefficients h" technique. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) Selected Areas in Cryptography - SAC 2008, Revised Selected Papers. Lecture Notes in Computer Science, vol. 5381, pp. 328–345. Springer (2008). https://doi.org/10.1007/978-3-642-04159-4_21
41. Patarin, J.: A proof of security in $o(2n)$ for the xor of two random permutations. In: Safavi-Naini, R. (ed.) Information Theoretic Security - ICITS 2008, Proceedings. Lecture Notes in Computer Science, vol. 5155, pp. 232–248. Springer (2008). https://doi.org/10.1007/978-3-540-85093-9_22
42. Shoup, V.: OAEP reconsidered. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 239–259. Springer, Heidelberg (Aug 2001). https://doi.org/10.1007/3-540-44647-8_15
43. Song, F., Yun, A.: Quantum security of NMAC and related constructions - PRF domain extension against quantum attacks. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 283–309. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63715-0_10
44. Zhandry, M.: How to record quantum queries, and applications to quantum indistinguishability. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 239–268. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_9

Symmetric-key Cryptography



Toward Full n -bit Security and Nonce Misuse Resistance of Block Cipher-Based MACs

Wonseok Choi^{1,2(✉)}, Jooyoung Lee^{3(✉)}, and Yeongmin Lee^{4(✉)}

¹ Purdue University, West Lafayette, IN, USA
wonseok@purdue.edu

² Georgia Institute of Technology, Atlanta, GA, USA

³ KAIST, Daejeon, Korea
hicalf@kaist.ac.kr

⁴ DESILO Inc., Seoul, Korea
yeongmin.lee@desilo.ai

Abstract. In this paper, we study the security of MAC constructions among those classified by Chen *et al.* in ASIACRYPT '21. Precisely, $F_{B_2}^{\text{EDM}}$ (or EWCDM as named by Cogliati and Seurin in CRYPTO '16), $F_{B_3}^{\text{EDM}}$, $F_{B_2}^{\text{SoP}}$, $F_{B_3}^{\text{SoP}}$ (all as named by Chen *et al.*) are proved to be fully secure up to 2^n MAC queries in the nonce-respecting setting, improving the previous bound of $\frac{3n}{4}$ -bit security. In particular, $F_{B_2}^{\text{SoP}}$ and $F_{B_3}^{\text{SoP}}$ enjoy graceful degradation as the number of queries with repeated nonces grows (when the underlying universal hash function satisfies a certain property called *multi-xor-collision resistance*). To do this, we develop a new tool, namely, extended Mirror theory for two independent permutations with a wide range of ξ_{\max} including inequalities. We also present matching attacks on $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$ using $O(2^{3n/4})$ MAC queries and $O(1)$ verification query without using repeated nonces.

Keywords: message authentication code · beyond birthday bound security · Mirror theory

1 Introduction

BEYOND BIRTHDAY BOUND MACs. A message authentication code (MAC) is a fundamental symmetric primitive allowing two entities sharing a secret key to verify that a received message originates from one of the two parties and was not modified by an attacker. Most popular MAC constructions are based

Wonseok Choi was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00239620), by AnalytiXIN, and by Sunday Group.

Jooyoung Lee was supported by the MSIT (Ministry of Science and ICT), Korea, under the Convergence security core talent training business support program (IITP- 2024-II221202) supervised by the IITP (Institute of Information & Communications Technology Planning & Evaluation).

This work was done while Yeongmin Lee was a PhD student at KAIST.

© International Association for Cryptologic Research 2025

K.-M. Chung and Y. Sasaki (Eds.): ASIACRYPT 2024, LNCS 15492, pp. 251–279, 2025.

https://doi.org/10.1007/978-981-96-0947-5_9

on block ciphers (e.g., CBC-MAC [2], PMAC [7], and OMAC [20]). At a high level, well-known block cipher-based MAC constructions such as CBC-MAC and PMAC follow the *UHF-then-PRF* design paradigm: a message is first mapped onto a short string through a universal hash function (UHF) and then encrypted through a fixed-input-length PRF to obtain a short tag. This method is simple, deterministic and stateless, yet its security caps at the so-called birthday bound; any collision at the output of the UHF, which translates into a tag collision, is usually enough to break the security of the scheme. The birthday bound security might not be enough, in particular, when the MAC construction is instantiated with a block cipher such as PRESENT [8], LED [17], and GIFT [1] operating on small blocks. A small block length, such as 64 bits, of the underlying primitive can render it a practical attack target when used in modes with birthday-bound security, as was illustrated by the recent attacks on popular communication protocols such as TLS [6].

NONCE-BASED MACs. Authenticated encryption schemes use a nonce (a value that never repeats) to give diversity to encryption of messages. The tag generation can be modeled as a nonce-based MAC in this case. Nonce-based MACs might be designed by a deterministic MAC using the concatenation of a nonce and a message as an input, or the well-known Wegman-Carter (WC) [29,30] construction. Many studies have tried to tweak deterministic MACs to obtain BBB security. They share a similar structural design of doubling the internal state of the hash function [25,31–33]. Better security bounds can be obtained for Wegman-Carter style MACs [4,13,29,30]. The WC construction is based on a universal hash function H and a pseudorandom function (PRF) F , that computes the corresponding tag as

$$T = H_{K_h}(M) \oplus F_K(N)$$

where K is the key for F , K_h is the key for H , and N and M denote a nonce and a message, respectively. It enjoys a powerful security bound when nonces are never repeated. Assuming F_K is a uniformly random function, the adversary can make a forgery with probability at most $v\epsilon$, where v is the number of verification queries and ϵ is the collision probability of H . By assuming ϵ is close to $\frac{1}{2^n}$, WC is secure up to $O(2^n)$ forgery attempts. This paradigm has been widely employed, e.g., in the Poly1305-AES [5] and GMAC [24] standards, and studied in depth [4].

NONCE MISUSE RESISTANCE. Despite the strong security advantages, the WC construction suffers from one major shortcoming: it is vulnerable to *nonce-misuse*. The construction might be seriously attacked if a nonce is repeated even once. For example, in the case of polynomial universal hashing, a repeated nonce can lead to the recovery of the hash key, which allows successful forgeries [18]. It might be challenging to maintain the uniqueness of a nonce in certain environments, for example, when a nonce is chosen from a set of low entropy or when the state of the MAC is reset due to some fault in its implementation. For this reason, there has been a considerable amount of research on constructing nonce-based MACs that provide security under nonce misuse.

1.1 Motivation

EWCDM [13] is based on an n -bit hash function H and an n -bit block cipher E ; it takes as input an n -bit nonce N and a message M , and outputs the corresponding tag as follows.

$$\text{EWCDM}[H, E](N, M) = E_{K_2}(H_{K_h}(M) \oplus E_{K_1}(N) \oplus N)$$

for hash key K_h and block cipher keys K_1 and K_2 . By using two block cipher calls, its security has been proved up to $O(2^{2n/3})$ MAC queries and $O(2^n)$ verification queries. As a variant of EWCDM, Datta *et al.* [15] proposed to replace the second block cipher call of EWCDM by block cipher decryption using the same key; for a nonce $N = N^* \parallel 0^{n/3}$ and a message M ,

$$\text{DWCDM}[H, E](N, M) = E_K^{-1}(H_{K_h}(M) \oplus E_K(N) \oplus N).$$

DWCDM is also secure up to $O(2^{2n/3})$ MAC queries and $O(2^n)$ verification queries.

Notably, Mennink and Neves [22] proved n -bit PRF security of EWCDM, but their proof relied on unverifiable Mirror theory. Recently, Datta *et al.* [14] proved $\frac{3n}{4}$ -bit MAC security of EWCDM and DWCDM using $\frac{3n}{4}$ -bit nonces using verifiable Mirror theory. More precisely, the adversarial advantages against the PRF security of EWCDM and DWCDM are upper bounded by $O(q^{4/3}/2^n)$ and $O(q^{1/3}/2^{n/4})$, respectively, in the nonce-respecting setting, while both constructions are secure up to $O(2^n)$ verification queries.

Dutta *et al.* [16] formalized the faulty nonce model for MAC constructions, where a MAC query is considered *faulty* if it is queried with a repeated nonce. They introduced the nonce-based Enhanced Hash-then-Mask (nEHtM) construction and proved its security up to $O(2^{2n/3})$ MAC queries and $O(2^n)$ verification queries in a nonce-respecting setting. Moreover, nEHtM enjoys graceful security degradation when nonces are misused. For the number of faulty nonces μ , their bound on the forging advantage includes $\mu q/2^n$ and $\mu v/2^n$ terms, where q and v denote the number of MAC queries and the number of verification queries, respectively. Subsequently, Choi *et al.* [10] improved this security bound to $\frac{3n}{4}$ bits when the number of faulty nonces is below $2^{3n/8}$, and also proved graceful security degradation for $\mu \leq 2^{n/2}$. Recently, Chen *et al.* [9] classified nonce-based MAC constructions that use two block cipher calls, one universal hash function call and an arbitrary number of XOR operations, and analyzed their PRF security in the faulty nonce model. Some constructions have been shown to achieve $\frac{3n}{4}$ -bit PRF security. However, the tightness of those constructions still remains open. This line of research raises the following fundamental question:

“Is there a block cipher-based MAC construction using nonces that provides both full n -bit security and nonce misuse resistance?”

1.2 Our Contribution

To affirmatively answer the question, we selected six candidates of nonce-based MAC constructions from [9]; EWCDM (denoted as $F_{B_2}^{\text{EDM}}$ in [9]), while denoted




Fig. 1. MAC constructions $F_{B_2}^{SoP}$ and $F_{B_4}^{EDM}$ based on a universal hash function H and a block cipher E .

EWCDM in this paper), $F_{B_3}^{EDM}$, $F_{B_2}^{SoP}$, $F_{B_3}^{SoP}$, $F_{B_4}^{EDM}$, and $F_{B_5}^{EDM}$. The subscript is defined in [9] based on where the hash value is added. For a nonce N and a message M , $F_{B_2}^{SoP}$ and $F_{B_4}^{EDM}$ compute the corresponding tags as follows:

$$\begin{aligned}
 F_{B_2}^{SoP}[H, E](N, M) &= E_{K_1}(N) \oplus E_{K_2}(N \oplus H_{K_h}(M)), \\
 F_{B_4}^{EDM}[H, E](N, M) &= E_{K_2}(E_{K_1}(N \oplus H_{K_h}(M)) \oplus N)
 \end{aligned}$$

where K_h is a hash key and K_1 and K_2 are block cipher keys (see Fig. 1). We can also prove the security of the following constructions:

$$\begin{aligned}
 F_{B_3}^{EDM}[H, E](N, M) &= \text{EWCDM}[H, E](N, M) \oplus H_{K_h}(M), \\
 F_{B_3}^{SoP}[H, E](N, M) &= F_{B_2}^{SoP}[H, E](N, M) \oplus H_{K_h}(M), \\
 F_{B_5}^{EDM}[H, E](N, M) &= F_{B_4}^{EDM}[H, E](N, M) \oplus H_{K_h}(M),
 \end{aligned}$$

since adding $H_{K_h}(M)$ to the tag does not significantly affect their security proof. Our contribution is summarized as follows:

1. We prove the tightness of the security bounds for 6 MAC schemes using two (independent) block cipher calls except $F_{B_2}^{EDMD}$ from Chen *et al.* [9]. This result will be discussed in more detail in the next part of this section.
2. To prove their security, we generalize state-of-the-art Mirror theory for two independent permutations with equation and inequality systems. To obtain the result, we first prove the Mirror theory when the distinction condition between variables is relaxed. Then, we further formalize the extended Mirror theory by using a new approach: estimates the ratio between the number of solutions to a system of equations and those with the addition of inequalities.

Table 1. Security of MAC constructions where μ is the number of faulty nonces and n is the block size. NR (resp. NM) denotes security in the nonce respecting (resp. misuse) setting. CR and MCR denote xor-collision resistance and multi-xor-collision resistance, respectively.

MAC	NR	NM	Tightness	Hash assumption	References
WC	2^n	0	tight	CR	[30]
EWCDM	$2^{3n/4}$	$2^{n/2}$	-	CR	[13, 14]
$F_{B_3}^{\text{EDM}}$	$2^{3n/4}$	$2^{n/2}$	-	CR	[9]
$F_{B_2}^{\text{SoP}}$	$2^{3n/4}$	$2^{3n/4}$ ($\mu \leq 2^{n/4}$)	-	CR	[9]
$F_{B_3}^{\text{SoP}}$	$2^{3n/4}$	$2^{3n/4}$ ($\mu \leq 2^{n/4}$)	-	CR	[9]
$F_{B_4}^{\text{EDM}}$	$2^{3n/4}$	$2^{3n/4}$ ($\mu < 2^{n/2}$)	tight	CR	[9], Sect. 6
$F_{B_5}^{\text{EDM}}$	$2^{3n/4}$	$2^{3n/4}$ ($\mu < 2^{n/2}$)	tight	CR	[9], Sect. 6
EWCDM	2^n	$2^{n/2}$	tight	CR	Sect. 4
$F_{B_3}^{\text{EDM}}$	2^n	$2^{n/2}$	tight	CR	Sect. 4
$F_{B_2}^{\text{SoP}}$	2^n	$2^n / \mu$ ($\mu \leq 2^{n/2}$) [†]	tight (NR)	MCR	Sect. 5
$F_{B_3}^{\text{SoP}}$	2^n	$2^n / \mu$ ($\mu \leq 2^{n/2}$) [†]	tight (NR)	MCR	Sect. 5

[†] In this paper, we proved the security bound for $\mu \leq 2^{n/4}$, while the same bound is obtained when $2^{n/4} \leq \mu \leq 2^{n/2}$ in a similar way to [16].

3. We also prove multi-xor-collision probability of CBC-MAC is negligible: for any distinct $x_1, \dots, x_k \in \{0, 1\}^*$ and distinct $y_1, \dots, y_k \in \{0, 1\}^n$,

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : H_{K_h}(x_1) \oplus y_1 = \dots = H_{K_h}(x_k) \oplus y_k] \leq \epsilon$$

for a small ϵ . This allow us to prove that $F_{B_2}^{\text{SoP}}$ and $F_{B_5}^{\text{EDM}}$ with the internal hash function instantiated with CBC-MAC achieves n -bit security.

For the tightness of the security bounds, we have the following results (see also Table 1):

1. We prove n -bit MAC security of EWCDM and $F_{B_3}^{\text{EDM}}$ in the nonce respecting setting. More precisely, EWCDM and $F_{B_3}^{\text{EDM}}$ are secure up to $O(2^n)$ MAC queries and $O(2^n)$ verification queries. It is the first concrete proof of n -bit MAC security of EWCDM to the best of our knowledge.
2. We prove that $F_{B_2}^{\text{SoP}}$ and $F_{B_3}^{\text{SoP}}$ are secure up to $O(2^n)$ MAC queries and $O(2^n)$ verification queries in the nonce respecting setting. In addition, we show that $F_{B_2}^{\text{SoP}}$ and $F_{B_3}^{\text{SoP}}$ are secure up to $O(2^n/\mu)$ MAC queries and $O(2^n)$ verification queries when the adversary makes μ faulty queries. Compared to the previous analysis, it enjoys stronger provable security when $\mu \leq O(2^{n/4})$. However, for these constructions, the underlying hash function should have a multi-xor-collision resistance property. As a concrete example, we show that CBC-MAC [19] is multi-xor-collision resistant.
3. We present a matching universal forgery attack on $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$ using $O(2^{3n/4})$ MAC queries and $O(1)$ verification query without using repeated

nonces. Since $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$ are provably secure up to $O(2^{3n/4})$ queries when $\mu < O(2^{n/2})$, they achieve tight $\frac{3n}{4}$ -bit security within the range of μ . The core idea of this attack is to find four query-answer tuples (N_1, M_1, T_1) , (N_2, M_2, T_2) , (N_3, M_3, T_3) , and (N_4, M_4, T_4) satisfying the following conditions:

$$\begin{aligned} N_1 \oplus H_{K_h}(M_1) &= N_2 \oplus H_{K_h}(M_2), \\ T_2 &= T_3, \\ N_3 \oplus H_{K_h}(M_3) &= N_4 \oplus H_{K_h}(M_4), \\ T_1 &= T_4, \\ N_1 \oplus N_2 \oplus N_3 \oplus N_4 &= \mathbf{0}. \end{aligned}$$

By repeating a nonce $O(2^{n/2})$ times, one can find such pairs with high probability. On the other hand, in the nonce-respecting setting, one can choose a well-structured set of nonces. From such pairs, a forgery is made with high probability.

As a proof strategy, we first extend a two-permutation version of Mirror theory to a wider range of ξ_{\max} , and then give a generic extension of Mirror theory for equation systems and Mirror theory for equation and inequality systems.

The main tool of our security proof is Mirror theory, which systematically estimates the number of solutions to a system of equations of the form $X_i \oplus X_j = \lambda_{i,j}$ such that X_1, \dots, X_q are pairwise distinct. Recently, Cogliati *et al.* [12] presented the complete proof of Mirror theory for a wide range of ξ_{\max} , where ξ_{\max} denotes the maximum component size when a system of equations is represented by a graph. However, we cannot directly apply their result to our problem; since our target constructions are based on two independent permutations, all variables are not necessarily pairwise distinct. To address this case, we divide the set of variables \mathcal{V} into \mathcal{V}_1 and \mathcal{V}_2 where $\mathcal{V} = \mathcal{V}_1 \sqcup \mathcal{V}_2$. Then, we estimate the number of solutions to a system of equations such that only the variables in \mathcal{V}_1 (or \mathcal{V}_2) are pairwise distinct. By letting $\mathcal{V}_1 = \mathcal{V}$ and $\mathcal{V}_2 = \emptyset$, one can recover the Mirror theory for a single permutation. Even with n -bit Mirror theory for independent permutations, the security proof is not immediate. It is not trivial to prove MAC security (also called “unforgeability”) from regular Mirror theory. We propose a generic method for deriving extended Mirror theory from a regular Mirror theory. With our modular approach, we can apply regular Mirror theory to the extended Mirror theory, which is much simpler than proving the extended Mirror theory directly.

When it comes to $F_{B_2}^{\text{SoP}}$ and $F_{B_3}^{\text{SoP}}$, the underlying hash function is required to satisfy the multi-xor-collision resistance property. We prove multi-xor-collision resistance of CBC-MAC which is one of ISO standards using the well-known structure graph technique [3, 21, 28]. We believe other MACs of ISO/IEC 9797-1 can be proved similarly since they have the same iteration algorithm.

2 Preliminaries

NOTATION. Throughout this paper, we fix positive integers n to denote the block size. We denote 0^n (i.e., n -bit string of all zeros) by $\mathbf{0}$. The set $\{0, 1\}^n$ is sometimes regarded as a set of integers $\{0, 1, \dots, 2^n - 1\}$ by converting an n -bit string $a_{n-1} \dots a_1 a_0 \in \{0, 1\}^n$ to an integer $a_{n-1}2^{n-1} + \dots + a_1 2 + a_0$. We also identify $\{0, 1\}^n$ with a finite field $\mathbf{GF}(2^n)$ with 2^n elements. For a positive integer q , we write $[q] = \{1, \dots, q\}$.

Given a non-empty finite set \mathcal{X} , $x \leftarrow_{\S} \mathcal{X}$ denotes that x is chosen uniformly at random from \mathcal{X} . $|\mathcal{X}|$ means the number of elements in \mathcal{X} . The set of all permutations of $\{0, 1\}^n$ is simply denoted $\text{Perm}(n)$. For some positive integer m , the set of all functions with domain $\{0, 1\}^n$ and codomain $\{0, 1\}^m$ is simply denoted by $\text{Func}(n, m)$. For a keyed function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ with key space \mathcal{K} and non-empty sets \mathcal{X} and \mathcal{Y} , we will denote $F(K, \cdot)$ by $F_K(\cdot)$ for $K \in \mathcal{K}$. The set of all sequences that consist of b pairwise distinct elements of \mathcal{X} is denoted \mathcal{X}^{*b} . For integers $1 \leq b \leq a$, we will write $(a)_b = a(a-1) \dots (a-b+1)$ and $(a)_0 = 1$ by convention. If $|\mathcal{X}| = a$, then $(a)_b$ becomes the size of \mathcal{X}^{*b} .

When two sets \mathcal{X} and \mathcal{Y} are disjoint, their (disjoint) union is denoted $\mathcal{X} \sqcup \mathcal{Y}$.

HASH FUNCTION. Let \mathcal{K}_h and \mathcal{X} be two non-empty finite sets and $H : \mathcal{K}_h \times \mathcal{X} \rightarrow \{0, 1\}^n$ be the hash function. Then,

1. H is said to be an ϵ -almost xor universal (AXU) hash function, if for any distinct $x, x' \in \mathcal{X}$ and $y \in \{0, 1\}^n$,

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : H_{K_h}(x) \oplus H_{K_h}(x') = y] \leq \epsilon.$$

2. H is said to be an (k, ϵ) -almost xor universal (AXU) hash function, if for any distinct $x_1, \dots, x_k \in \mathcal{X}$ and distinct $y_1, \dots, y_k \in \{0, 1\}^n$,

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : H_{K_h}(x_1) \oplus y_1 = \dots = H_{K_h}(x_k) \oplus y_k] \leq \epsilon.$$

BLOCK CIPHER. Let $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an n -bit block cipher with key space \mathcal{K} . We will consider an information-theoretic distinguisher \mathcal{A} that makes oracle queries to E , and returns a single bit. The advantage of \mathcal{A} in breaking the prp security of E is defined as

$$\mathbf{Adv}_E^{\text{prp}}(\mathcal{A}) = |\Pr [K \leftarrow_{\S} \mathcal{K} : \mathcal{A}^{E_K} = 1] - \Pr [P \leftarrow_{\S} \text{Perm}(n) : \mathcal{A}^P = 1]|.$$

We define $\mathbf{Adv}_E^{\text{prp}}(q, t)$ as the maximum of $\mathbf{Adv}_E^{\text{prp}}(\mathcal{A})$ over all the distinguishers against E making at most q queries and running in time at most t . When considering information-theoretic security, we will drop the parameter t .

NONCE-BASED PSEUDORANDOM FUNCTION. Let $F : \mathcal{K} \times \mathcal{N} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a nonce-based keyed function with key space \mathcal{K} and nonce space \mathcal{N} . We will consider an information-theoretic distinguisher \mathcal{A} that makes oracle queries to F , and returns a single bit. The advantage of \mathcal{A} in breaking the prf security of

F , i.e., in distinguishing F_K where $K \leftarrow_{\S} \mathcal{K}$ from the random oracle Rand , is defined as

$$\mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) = |\Pr [K \leftarrow_{\S} \mathcal{K} : \mathcal{A}^{F_K} = 1] - \Pr [\mathcal{A}^{\text{Rand}} = 1]|.$$

We define $\mathbf{Adv}_F^{\text{prf}}(\mu, q, t)$ as the maximum of $\mathbf{Adv}_F^{\text{prf}}(\mathcal{A})$ over all the distinguishers against F making at most q queries, at most μ faulty queries and running in time at most t . We also denote $\mathbf{Adv}_F^{\text{prf}}(q, t)$ for $\mathbf{Adv}_F^{\text{prf}}(0, q, t)$. When we consider information theoretic security, we will drop the parameter t .

NONCE-BASED MACs. Given four non-empty sets \mathcal{K} , \mathcal{N} , \mathcal{M} , and \mathcal{T} , a nonce-based keyed function with key space \mathcal{K} , nonce space \mathcal{N} , message space \mathcal{M} and tag space \mathcal{T} is simply a function $F : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \rightarrow \mathcal{T}$. Stated otherwise, it is a keyed function whose domain is a cartesian product $\mathcal{N} \times \mathcal{M}$. We denote $F_K(N, M)$ for $F(K, N, M)$.

For $K \in \mathcal{K}$, let Auth_K be the MAC oracle which takes as input a pair $(N, M) \in \mathcal{N} \times \mathcal{M}$ and returns $F_K(N, M)$, and let Ver_K be the verification oracle which takes as input a triple $(N, M, T) \in \mathcal{N} \times \mathcal{M} \times \mathcal{T}$ and returns 1 (“accept”) if $F_K(N, M) = T$, and 0 (“reject”) otherwise. We assume an adversary queries the two oracles Auth_K and Ver_K for a secret key $K \in \mathcal{K}$.

A (μ, q, v, t) -adversary against the nonce-based MAC-security of F is an adversary \mathcal{A} with oracle access to oracles Auth_K and Ver_K , making at most q MAC queries to Auth oracle, at most μ faulty queries, at most v verification queries to Ver oracle, and running in time at most t . We say that \mathcal{A} forges if any of its queries to Ver_K returns 1. The advantage of \mathcal{A} against the nonce-based MAC security of F is defined as

$$\mathbf{Adv}_F^{\text{mac}}(\mathcal{A}) = \Pr [K \leftarrow_{\S} \mathcal{K} : \mathcal{A}^{\text{Auth}_K, \text{Ver}_K} \text{ forges}].$$

where the probability is also taken over the random coins of \mathcal{A} , if any. \mathcal{A} is not allowed to ask a verification query (N, M, T) to Ver_K if a previous query (N, M) to Auth_K returned T . When $\mu = 0$, we say that \mathcal{A} is nonce-respecting, otherwise, \mathcal{A} is said nonce-misusing. However, the adversary is allowed to repeat nonces in its verification queries.

We define $\mathbf{Adv}_F^{\text{mac}}(\mu, q, v, t)$ as the maximum of $\mathbf{Adv}_F^{\text{mac}}(\mathcal{A})$ over all (μ, q, v, t) -adversaries. We also define $\mathbf{Adv}_F^{\text{mac}}(q, v, t)$ as the maximum of $\mathbf{Adv}_F^{\text{mac}}(\mathcal{A})$ over all $(0, q, v, t)$ -adversaries. When we consider information-theoretic security, we will drop the parameter t .

We obtain an upper bound for the forging advantage of F in terms of distinguishing advantage, where the ideal world is comprised of a random oracle Rand and the reject oracle Rej that always returns 0 for any verification query. For any (μ, q, v, t) -adversary \mathcal{A} , $\mathbf{Adv}_F^{\text{mac}}(\mathcal{A})$ is upper bounded by

$$\max_{\mathcal{A}} |\Pr [K \leftarrow_{\S} \mathcal{K} : \mathcal{A}^{\text{Auth}_K, \text{Ver}_K} = 1] - \Pr [\mathcal{A}^{\text{Rand, Rej}} = 1]|.$$

2.1 Coefficient-H Technique

We will use Patarin’s coefficient-H technique. The goal of this technique is to upper bound the adversarial distinguishing advantage between a real construc-

tion and its ideal counterpart. In the ideal and the real worlds, an information-theoretic adversary \mathcal{A} is allowed to make q queries to certain oracles (with the same oracle interfaces), denoted \mathcal{S}_0 and \mathcal{S}_1 , respectively. The interaction between the adversary \mathcal{A} and the oracle determines a “transcript” $\tau \in \Omega^q$; it contains all the information obtained by \mathcal{A} during the interaction. We call a transcript τ *attainable* if the probability of obtaining τ in the ideal world is non-zero.

We partition the set of attainable transcripts Θ into a set of “good” transcripts Θ_{good} such that the probabilities of obtaining some transcript $\tau \in \Theta_{\text{good}}$ are close in the real world and the ideal world, and a set Θ_{bad} of “bad” transcripts such that the probability of obtaining any $\tau \in \Theta_{\text{bad}}$ is small in the ideal world. The coefficient-H technique is summarized in the following lemma.

Lemma 1. *Let $\Theta = \Theta_{\text{good}} \sqcup \Theta_{\text{bad}}$ be a partition of the set of attainable transcripts, where there exists a non-negative ϵ_1 such that for any $\tau \in \Theta_{\text{good}}$,*

$$\frac{p_{\mathcal{S}_1}^q(\tau)}{p_{\mathcal{S}_0}^q(\tau)} \geq 1 - \epsilon_1,$$

and there exists ϵ_2 such that $\sum_{\tau \in \Theta_{\text{bad}}} p_{\mathcal{S}_0}^q(\tau) \leq \epsilon_2$. Then,

$$\sum_{\tau \in \Theta} \max \{0, p_{\mathcal{S}_0}^q(\tau) - p_{\mathcal{S}_1}^q(\tau)\} \leq \epsilon_1 + \epsilon_2.$$

Proof. We have

$$\begin{aligned} \sum_{\tau \in \Theta} \max \{0, p_{\mathcal{S}_0}^q(\tau) - p_{\mathcal{S}_1}^q(\tau)\} &= \sum_{\substack{\tau \in \Theta \\ p_{\mathcal{S}_0}^q(\tau) > p_{\mathcal{S}_1}^q(\tau)}} (p_{\mathcal{S}_0}^q(\tau) - p_{\mathcal{S}_1}^q(\tau)) \\ &= \sum_{\substack{\tau \in \Theta \\ p_{\mathcal{S}_0}^q(\tau) > p_{\mathcal{S}_1}^q(\tau)}} p_{\mathcal{S}_0}^q(\tau) \left(1 - \frac{p_{\mathcal{S}_1}^q(\tau)}{p_{\mathcal{S}_0}^q(\tau)}\right) \\ &\leq \sum_{\tau \in \Theta_{\text{good}}} p_{\mathcal{S}_0}^q(\tau) \epsilon_1 + \sum_{\tau \in \Theta_{\text{bad}}} p_{\mathcal{S}_0}^q(\tau) \\ &\leq \epsilon_1 + \epsilon_2. \end{aligned}$$

□

3 Mirror Theory

Patarin’s Mirror theory [26, 27] has been a valuable tool for proving PRF security and MAC security. However, the original proof provided by Patarin is complex and hard to verify, containing several gaps. Recently, Cogliati *et al.* [12] presented the complete proof of Mirror theory for a wide range of ξ_{max} . Nevertheless, there are limitations when it comes to proving the security of our target MACs using the Mirror theory in [12]. This is because the Mirror theory focuses on a single permutation. To address this limitation, we refine the Mirror theory to cover constructions based on two independent permutations, allowing us to analyze the security of two permutation-based constructions. Additionally, we need to extend the Mirror theory to include inequalities for MAC security. This extended version is known as “Extended Mirror theory”.

3.1 Extended Mirror Theory for Two Independent Permutations

The goal of this section is to compute a lower bound of the number of solutions to a certain type of system of equations and inequalities.

We consider a system of equations and inequalities $\gamma = (\gamma^=, \gamma^\neq)$, which is divided into a system of equations $\gamma^=$ and a system of inequalities γ^\neq . A set of variables \mathcal{V} is partitioned into $\mathcal{V}_1 \sqcup \mathcal{V}_2$. Intuitively, variables in \mathcal{V}_1 come from one permutation and ones in \mathcal{V}_2 are results of the other permutation. In this section, we assume that they are arbitrarily partitioned. So, the variables in \mathcal{V}_1 (or \mathcal{V}_2) should be distinct. We use the notion $X \sim Y$ to indicate that X and Y belong to the same subset meaning that X and Y are distinct elements within that subset. Additionally, we impose the following constraint on γ : If $X \sim Y$, then $X \neq Y$.

Fix a positive integer c . For $1 \leq i \leq c$ and a positive integer $\xi_i > 1$, the system of equations as $\gamma^=$ is represented as:

$$\gamma^= : \begin{cases} X_{1,0} \oplus X_{1,1} = \lambda_{1,1}, \dots, X_{1,0} \oplus X_{1,\xi_1-1} = \lambda_{1,\xi_1-1}, \\ \vdots \\ X_{c,0} \oplus X_{c,1} = \lambda_{c,1}, \dots, X_{c,0} \oplus X_{c,\xi_c-1} = \lambda_{c,\xi_c-1} \end{cases}$$

where $\lambda_{\alpha,i} \in \{0, 1\}^n$ for $1 \leq \alpha \leq c$ and $0 \leq i \leq \xi_\alpha - 1$. The set of variables on $\gamma^=$ is denoted as $\mathcal{V}^=$ and we define $\mathcal{V}_1^= \stackrel{\text{def}}{=} \mathcal{V}^= \cap \mathcal{V}_1$ and $\mathcal{V}_2^= \stackrel{\text{def}}{=} \mathcal{V}^= \cap \mathcal{V}_2$. We also define $\mathcal{V}^\neq \stackrel{\text{def}}{=} \mathcal{V} \setminus \mathcal{V}^=$, $\mathcal{V}_1^\neq \stackrel{\text{def}}{=} \mathcal{V}^\neq \cap \mathcal{V}_1$ and $\mathcal{V}_2^\neq \stackrel{\text{def}}{=} \mathcal{V}^\neq \cap \mathcal{V}_2$. The set of variables $\mathcal{V}^=$ consists of c components, and for $i \in [c]$, the i -th component takes form of $\{X_{i,0}, \dots, X_{i,\xi_i-1}\}$. The largest number of components is denoted as ξ_{\max} , where $\xi_{\max} = \max_{i \in [c]} \{\xi_i\}$.

We separately establish a system of inequalities with γ^\neq . For a non-negative integer v , we denote

$$\gamma^\neq : \begin{cases} X'_1 \oplus X'_2 \neq \lambda'_1, \\ X'_3 \oplus X'_4 \neq \lambda'_2, \\ \vdots \\ X'_{2v-1} \oplus X'_{2v} \neq \lambda'_v \end{cases}$$

where $\lambda'_i \in \{0, 1\}^n$ for $1 \leq i \leq v$. Note that a variable that appears in both systems of equations and inequalities can be represented by a single symbol. However, to clearly distinguish between the system of equations and the system of inequalities, we use separate symbols for those in $\gamma^=$ and those in γ^\neq . The equivalence between variables is indicated using the relation \sim_{eq} . Specifically, for some i , X'_i can be identified as an element of $\mathcal{V}^=$ or another element of \mathcal{V}^\neq . This identification is publicly known and can be denoted as a relation \sim_{eq} , i.e., $X'_i \sim_{\text{eq}} X_{j,k} \Leftrightarrow X'_i = X_{j,k}$ and $X'_i \sim_{\text{eq}} X'_j \Leftrightarrow X'_i = X'_j$.

In this section, we express the system of equations and inequalities with relation \sim and \sim_{eq} ; denoted as $\Gamma \stackrel{\text{def}}{=} (\gamma^=, \gamma^\neq, \sim, \sim_{\text{eq}})$. $h(\Gamma)$ denotes the number of solutions to γ subject to the above constraints.

In this work, we focus on a system Γ for which at least one solution exists. To ensure a solution, the system must satisfy the non-degeneracy properties outlined below:

1. $\lambda_{\alpha,i} \neq 0$ for all $\alpha \in [c]$ and $i \in [\xi_\alpha - 1]$ such that $X_{\alpha,0} \sim X_{\alpha,i}$.
2. $\lambda_{\alpha,i} \neq \lambda_{\alpha,j}$ for all $\alpha \in [c]$ and distinct $i, j \in [\xi_\alpha - 1]$ such that $X_{\alpha,i} \sim X_{\alpha,j}$.
3. There is no (α, β, i, j) such that γ^- contains $X_{\alpha,i} \oplus X_{\alpha,j} = \lambda'_\beta$ and γ^{\neq} contains $X_{\alpha,i} \oplus X_{\alpha,j} \neq \lambda'_\beta$.

We refer to any system Γ satisfying the above properties as a *nice* system. The following theorem provides a lower bound of $h(\Gamma)$ for a nice system Γ .

Theorem 1. *Let Γ be a nice system over $\{0, 1\}^n$ such that the number of equations is q and the number of inequalities is v . Suppose the number of variables in the largest component of γ^- is ξ_{\max} . If $\xi_{\max}^2 n + \xi_{\max} \leq 2^{n/2}$, $q \xi_{\max}^2 \leq \frac{2^n}{12}$ and $q + v \leq 2^{n-1}$, one has*

$$h(\Gamma) \geq \frac{(2^n - 2)^{|\mathcal{V}_1|} (2^n - 2)^{|\mathcal{V}_2|}}{2^{nq}} \left(1 - \frac{2v}{2^n} \right).$$

The proof of Theorem 1 is involved, so to enhance clarity, we begin by focusing on a subset of Γ that contains only the equations, as outlined in Sect. 3.2. Subsequently, we extend the theorem to include inequalities, which will be discussed in Sect. 3.3.

3.2 Mirror Theory with Equations

Set Representation. We start by establishing a fixed system Γ . Additionally, we define a new partition of \mathcal{V} as $\mathcal{V} = C_1 \sqcup \dots \sqcup C_c$, with each C_i being a set of variables defined as $C_i = \{X_{i,0}, \dots, X_{i,\xi_i-1}\}$. Let $\mathcal{F} = \{C_1, \dots, C_c\}$ represent a family of the sets C_i . In this context, we introduce a label function denoted as $\Lambda : \mathcal{V} \times \mathcal{V} \rightarrow \{0, 1\}^n \cup \{\perp\}$ defined as follows: when both $X_{i,j}$ and $X_{i,k}$ are within the same set C_i , $\Lambda(X_{i,j}, X_{i,k})$ returns $\lambda_{i,j} \oplus \lambda_{i,k}$ by letting $\lambda_{i,0} = 0$. Otherwise, it returns \perp . In this section, we fix \mathcal{F} and Λ .

Now, for any family of sets $\mathcal{G} = \{A_1, \dots, A_a\}$ where A_i is a subset of \mathcal{V} , and a given label function \mathcal{L} , we define the terminologies used in our proof:

- $N(\mathcal{G})$ represents the total number of variables in \mathcal{G} , i.e., $N(\mathcal{G}) = \sum_{1 \leq i \leq |\mathcal{G}|} |A_i|$.
- $N_1(\mathcal{G})$ and $N_2(\mathcal{G})$ denote the number of variables of \mathcal{G} contained in \mathcal{V}_1 and \mathcal{V}_2 , respectively. In other words,

$$N_1(\mathcal{G}) = \sum_{1 \leq i \leq |\mathcal{G}|} |A_i \cap \mathcal{V}_1|,$$

$$N_2(\mathcal{G}) = \sum_{1 \leq i \leq |\mathcal{G}|} |A_i \cap \mathcal{V}_2|.$$

- For a variable $v \in \mathcal{V}$,

$$N_v(\mathcal{G}) = \begin{cases} N_1(\mathcal{G}) & \text{if } v \in \mathcal{V}_1, \\ N_2(\mathcal{G}) & \text{if } v \in \mathcal{V}_2. \end{cases}$$

- We denote $h(\mathcal{G}, \mathcal{L})$ as the number of assignments to \mathcal{G} according to the label function \mathcal{L} while all the variables in \mathcal{V}_1 (resp. \mathcal{V}_2) should take on different values. Specifically, $h(\mathcal{F}, \Lambda)$ is equivalent to $h(\Gamma)$.
- $M(\mathcal{G})$ is the maximum number of components within the family, i.e., $M(\mathcal{G}) = \max_{1 \leq i \leq |\mathcal{G}|} \{|A_i|\}$.

We first estimate the number of solutions for a system of equations. Let $\Gamma^=$ be a system of equations $\gamma^=$ with relation \sim and $h(\Gamma^=)$ be the number of solutions to $\Gamma^=$. We can prove the following theorem.

Theorem 2. *Let $\Gamma^=$ be a nice system over $\{0, 1\}^n$ such that the number of equations is q . Suppose the number of variables in the largest component of $\Gamma^=$ is ξ_{\max} . If $\xi_{\max}^2 n + \xi_{\max} \leq 2^{n/2}$ and $q \xi_{\max}^2 \leq \frac{2^n}{12}$, one has*

$$h(\Gamma^=) \geq \frac{(2^n - 2)^{|\mathcal{V}_1^=|} (2^n - 2)^{|\mathcal{V}_2^=|}}{2^{nq}}.$$

Proof. Let \mathcal{G} be a sub-family of \mathcal{F} . For any set $S \in \mathcal{G}$, we claim that

$$h(\mathcal{G}, \Lambda) \geq 2^n h(\mathcal{G} \setminus \{S\}, \Lambda) \prod_{i=N_1(\mathcal{G} \setminus \{S\})+1}^{N_1(\mathcal{G})} \left(1 - \frac{i+1}{2^n}\right) \prod_{i=N_2(\mathcal{G} \setminus \{S\})+1}^{N_2(\mathcal{G})} \left(1 - \frac{i+1}{2^n}\right).$$

If $|S| = 1$, it means S contains only one element, say v , i.e., $S = \{v\}$. The claim is obvious since

$$h(\mathcal{G}, \Lambda) = h(\mathcal{G} \setminus \{S\}, \Lambda) \times (2^n - N_v(\mathcal{G}) + 1). \tag{1}$$

Next, suppose $|S| \geq 2$. We first consider the case that $N(\mathcal{G}) \leq 2^{\frac{n}{2}}$. We have

$$h(\mathcal{G}, \Lambda) \geq h(\mathcal{G} \setminus \{S\}, \Lambda) \times (2^n - N_1(S) \times N_1(\mathcal{G} \setminus \{S\}) - N_2(S) \times N_2(\mathcal{G} \setminus \{S\})).$$

In order to prove the claim, it is enough to show that

$$1 - \frac{N_1(S) \times N_1(\mathcal{G} \setminus \{S\})}{2^n} - \frac{N_2(S) \times N_2(\mathcal{G} \setminus \{S\})}{2^n} \geq \prod_{i=N_1(\mathcal{G} \setminus \{S\})+1}^{N_1(\mathcal{G})} \left(1 - \frac{i+1}{2^n}\right) \prod_{i=N_2(\mathcal{G} \setminus \{S\})+1}^{N_2(\mathcal{G})} \left(1 - \frac{i+1}{2^n}\right). \tag{2}$$

The above inequality is represented by

$$1 - \frac{ar + bs}{2^n} \geq \left(1 - \frac{a+2}{2^n}\right) \dots \left(1 - \frac{a+r+1}{2^n}\right) \left(1 - \frac{b+2}{2^n}\right) \left(1 - \frac{b+s+1}{2^n}\right).$$

This can be shown by induction on r and s . For $r = 1$ and $s = 1$, the inequality holds since

$$\begin{aligned} \left(1 - \frac{a+2}{2^n}\right) \left(1 - \frac{b+2}{2^n}\right) &\leq 1 - \frac{a+b}{2^n} - \frac{4}{2^n} \left(1 - \frac{(a+2)(b+2)}{2^{n+2}}\right) \\ &\leq 1 - \frac{a+b}{2^n}. \end{aligned}$$

The last inequality holds since $a, b \leq 2^{\frac{n}{2}}$ and $n \geq 2$. If $r \geq s$ we obtain

$$\begin{aligned} \left(1 - \frac{a(r-1) + bs}{2^n}\right) \left(1 - \frac{a+r+1}{2^n}\right) &\leq 1 - \frac{ar+bs}{2^n} - \frac{r+1}{2^n} + \frac{(a(r-1) + bs)(a+r+1)}{2^{2n}} \\ &\leq 1 - \frac{ar+bs}{2^n} - \frac{r+1}{2^n} \left(1 - \frac{(a+b)(a+r+1)}{2^n}\right) \\ &\leq 1 - \frac{ar+bs}{2^n} \end{aligned}$$

since $a+b \leq 2^{n/2}$ and $a+r+1 \leq 2^{n/2}$. If $r < s$, similarly, we have

$$\left(1 - \frac{ar+b(s-1)}{2^n}\right) \left(1 - \frac{b+s+1}{2^n}\right) \leq 1 - \frac{ar+bs}{2^n}.$$

By applying induction hypothesis for r and s , the Eq. (2) holds.

For an element $v \in S \in \mathcal{G}$, we denote \mathcal{G}_{-v} as a family of partitions deleting v , i.e., $\mathcal{G}_{-v} = (\mathcal{G} \setminus \{S\}) \cup \{S \setminus \{v\}\}$. We state the following lemma.

Given a set $S \in \mathcal{G}$, $v, w \in S$ and a label function \mathcal{L} , we define $\delta_{S,\mathcal{L}}(v, w)$ as the number of 2-subsets $\{a, b\}$ of S such that $a \sim v$ and $b \sim w$ with $\mathcal{L}(a, b) = \mathcal{L}(v, w)$. We define

$$\delta_{\mathcal{G},\mathcal{L}}(v, w) \stackrel{\text{def}}{=} \sum_{S \in \mathcal{G}} \delta_{S,\mathcal{L}}(v, w), \quad \Delta_{\mathcal{G},\mathcal{L}} \stackrel{\text{def}}{=} \max_{S \in \mathcal{G}} \max_{(v,w) \in S^{*2}} \delta_{S,\mathcal{L}}(v, w).$$

Then, we estimate the lower bound of $h(\mathcal{G}, \Lambda)$.

Lemma 2. *Suppose the maximum $\Delta_{\mathcal{G},\Lambda}$ is attained for $v, v' \in S \in \mathcal{G}$. If $2^{\frac{n}{2}} \leq N(\mathcal{G}) \leq \frac{2^n}{12\xi_{\max}^2}$, we have*

$$h(\mathcal{G}, \Lambda) \geq h(\mathcal{G}_{-v}, \Lambda) \left(1 - \frac{N_v(\mathcal{G}) + 1}{2^n}\right)$$

The proof of Lemma 2 is deferred to the full version of this paper [11].

When $2^{\frac{n}{2}} \leq N(\mathcal{G}) \leq \frac{2^n}{12\xi_{\max}^2}$, the claim holds by Lemma 2. By iterating the inequality, we conclude that

$$h(\mathcal{F}, \Lambda) \geq (2^n)^c \prod_{i=1}^{|\mathcal{V}_1|} \left(1 - \frac{i+1}{2^n}\right) \prod_{i=1}^{|\mathcal{V}_2|} \left(1 - \frac{i+1}{2^n}\right) \geq \frac{(2^n - 2)^{|\mathcal{V}_1|} (2^n - 2)^{|\mathcal{V}_2|}}{2^{nq}}.$$

□

3.3 Generalization of Extended Mirror Theory

Mirror theory is later generalized to *extended* Mirror theory [15, 16], by including inequalities in the system. The extended Mirror theory systematically estimates the number of solutions to a system of equations and inequalities. On the other hand, the goal of this section is slightly different: we will estimate *the ratio* between two quantities:

1. The number of solutions to a system of equations.
2. The number of solutions to a system of equations and inequalities.

This approach separates the counting of inequalities from (equations-only) Mirror theory, eliminating the need for developing the Extended Mirror theory each time whenever there is an improvement of Mirror theory.

When a given system Γ is nice, we can compute a lower bound on the ratio

$$\frac{h(\Gamma)}{h(\Gamma^=)}$$

as follows.

Lemma 3. *Let Γ be a nice system over $\{0, 1\}^n$ such that the number of equations is q and the number of inequalities is v . If $q + v \leq 2^{n-1}$, one has*

$$\frac{h(\Gamma)}{h(\Gamma^=)(2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq|}(2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq|}} \geq 1 - \frac{2v}{2^n}.$$

Proof (of Lemma 3). Let $\Gamma_0 = \Gamma^=$ and $\Gamma_i = \Gamma_{i-1} \sqcup \{X'_{2i-1} \oplus X'_{2i} \neq \lambda'_i\}$ for $i \in [v]$. We additionally define $\Gamma'_i = \Gamma_{i-1} \sqcup \{X'_{2i-1} \oplus X'_{2i} = \lambda'_i\}$ for $i \in [v]$. Then, we have

$$h(\Gamma_{i+1}) = h(\Gamma_i) - h(\Gamma'_{i+1}). \tag{3}$$

If both X'_{2i-1} and X'_{2i} are in $\mathcal{V}^=$, then Γ'_{i+1} contradicts, i.e., $h(\Gamma'_{i+1}) = 0$ since Γ_{i+1} is nice. Thus, $h(\Gamma_{i+1}) = h(\Gamma_i)$.

Now, we suppose that X'_{2i-1} or X'_{2i} is not in $\mathcal{V}^=$. The number of possible assignments of distinct values outside $\mathcal{V}^=$ to the variables in \mathcal{V}^\neq is $(2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq|}(2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq|}$. Among these assignments, it violates the inequality conditions when $X'_{2i-1} \oplus X'_{2i} = \lambda'_i$ for each $i \in [v]$. These assignments are at most

$$A \stackrel{\text{def}}{=} \max \left\{ (2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq| - 1} (2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq|}, (2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq|} (2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq| - 1} \right\}.$$

Therefore, we have

$$\frac{h(\Gamma)}{h(\Gamma^=)} \geq (2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq|} (2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq|} - vA$$

which means

$$\frac{h(\Gamma)}{h(\Gamma^=)(2^n - |\mathcal{V}_1^=|)_{|\mathcal{V}_1^\neq|}(2^n - |\mathcal{V}_2^=|)_{|\mathcal{V}_2^\neq|}} \geq 1 - \frac{2v}{2^n}.$$

It concludes the proof. □

By combining Theorem 2 and Lemma 3, Theorem 1 can be proved as

$$\begin{aligned} h(\Gamma) &\geq h(\Gamma^-)(2^n - |\mathcal{V}_1^-|)_{|\mathcal{V}_1^-|} (2^n - |\mathcal{V}_2^-|)_{|\mathcal{V}_2^-|} \left(1 - \frac{2v}{2^n}\right) \\ &\geq \frac{(2^n - 2)_{|\mathcal{V}_1^-|} (2^n - 2)_{|\mathcal{V}_2^-|}}{2^{nq}} \cdot (2^n - |\mathcal{V}_1^-|)_{|\mathcal{V}_1^-|} (2^n - |\mathcal{V}_2^-|)_{|\mathcal{V}_2^-|} \left(1 - \frac{2v}{2^n}\right) \\ &\geq \frac{(2^n - 2)_{|\mathcal{V}_1|} (2^n - 2)_{|\mathcal{V}_2|}}{2^{nq}} \left(1 - \frac{2v}{2^n}\right). \end{aligned}$$

4 Security of EWCDM and $F_{B_3}^{\text{EDM}}$

In this section, we consider EWCDM[H, E] and $F_{B_3}^{\text{EDM}}$ [H, E] based on an n -bit ϵ -AXU hash function H and an n -bit block cipher E . For given n -bit nonce N and a message M , the user receives a tag as

$$\text{EWCDM}[H, E](N, M) = E_{K_2}(H_{K_h}(M) \oplus E_{K_1}(N) \oplus N)$$

and

$$F_{B_3}^{\text{EDM}}[H, E](N, M) = E_{K_2}(H_{K_h}(M) \oplus E_{K_1}(N) \oplus N) \oplus H_{K_h}(M)$$

by a hash key K_h and block cipher keys K_1 and K_2 . The goal of this section is to prove the security of EWCDM[H, E] and $F_{B_3}^{\text{EDM}}$ [H, E]. As a result, we have the following theorem.

Theorem 3. *Let $n \geq 30$, $\epsilon > 0$, $H : \mathcal{K}_h \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be an ϵ -AXU hash function, and $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. Let q, v, t be nonnegative integers such that $q + v \leq 2^{n-1}$. Then, one has*

$$\text{Adv}_{\text{EWCDM}[H, E]}^{\text{mac}}(q, v, t) \leq \frac{6q}{2^n} + \frac{q^2\epsilon}{2^n} + \frac{6v}{2^n} + v\epsilon + 2\text{Adv}_E^{\text{prp}}(q + v, t + t').$$

where t' is the time complexity necessary to compute E for $q + v$ times.

Since adding $H_{K_h}(M)$ to the tag does not make any significant difference, the MAC security of $F_{B_3}^{\text{EDM}}$ follows immediately.

Corollary 1. *Let $n \geq 30$, $\epsilon > 0$, $H : \mathcal{K}_h \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be an ϵ -AXU hash function, and $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. Let q, v, t be nonnegative integers such that $q + v \leq 2^{n-1}$. Then, one has*

$$\text{Adv}_{F_{B_3}^{\text{EDM}}[H, E]}^{\text{mac}}(q, v, t) \leq \frac{6q}{2^n} + \frac{q^2\epsilon}{2^n} + \frac{6v}{2^n} + v\epsilon + 2\text{Adv}_E^{\text{prp}}(q + v, t + t').$$

where t' is the time complexity necessary to compute E for $q + v$ times.

4.1 Proof of Theorem 3

We assume that the adversary is deterministic and never repeats a prior query. Assume further that the adversary never makes a redundant query. Up to the prp-security of E , keyed block ciphers E_{K_1} and E_{K_2} can be replaced by truly random permutations P_1 and P_2^{-1} , respectively. The cost of this replacement is upper bounded by

$$2\text{Adv}_E^{\text{prp}}(q + v, t + t').$$

The resulting construction denotes $\text{EWCDM}^*[H]$.

At the end of the interaction between an adversary and the oracle, additional information is freely given to an adversary, and a transcript is defined as a pair of query-answer pairs and additional information K_h . In the real world, K_h is the hash key used in EWCDM. In the ideal world, K_h is uniformly randomly chosen after the end of the interaction between an adversary and the oracle. Without loss of generality, we rearrange query indices so that verification queries come after MAC queries. Let Θ be the set of all attainable transcripts in the ideal world and $\tau = (\tau_m, \tau_v, K_h) \in \Theta$ be a transcript where τ_m and τ_v denote the list of MAC queries and the list of verification queries, i.e.,

$$\begin{aligned} \tau_m &= \{(N_1, M_1, T_1), \dots, (N_q, M_q, T_q)\}, \\ \tau_v &= \{(N_{q+1}, M_{q+1}, T_{q+1}, b_{q+1}), \dots, (N_{q+v}, M_{q+v}, T_{q+v}, b_{q+v})\}. \end{aligned}$$

From a transcript τ , \mathcal{A} can compute $X_i = H_{K_h}(M_i) \oplus N_i$ for $i \in [q + v]$ before outputting its decision bit.

This proof utilizes the extended Mirror theory stated in Theorem 1 and the coefficient-H technique stated in Lemma 1. The core of the security proof is to estimate the number of possible ways of fixing evaluations P_1 and P_2 in a way that

$$X_i = P_1(N_i) \oplus P_2(T_i)$$

for $i = 1, \dots, q$ and

$$X_i \neq P_1(N_i) \oplus P_2(T_i)$$

for $i = q + 1, \dots, q + v$. We will identify $\mathcal{V}_1 = \{P_1(N_i)\}$ and $\mathcal{V}_2 = \{P_2(T_i)\}$ with as sets of variables. We also define $\mathcal{V} = \mathcal{V}_1 \sqcup \mathcal{V}_2$. Then we can construct the system of equations Γ_τ as defined in Sect. 3. To satisfy the conditions in Theorem 1, we must first define bad events on a transcript τ , and then we can apply the extended Mirror theory to each transcript that the bad event does not happen.

Defining and Bounding Bad Events. A transcript $\tau = (\tau_m, \tau_v, K_h)$ is defined as *bad* if one of the following condition holds.

- $\text{bad}_1 \Leftrightarrow$ there exists $(i_1, \dots, i_n) \in [q]^{*n}$ such that $T_{i_1} = \dots = T_{i_n}$.
- $\text{bad}_2 \Leftrightarrow$ there exists $(i, j) \in [q]^{*2}$ such that $T_i = T_j$ and $X_i = X_j$.

- $\text{bad}_3 \Leftrightarrow$ there exists $(i, j) \in [q] \times [q + 1, q + v]$ such that $N_i = N_j, T_i = T_j$, and $X_i = X_j$.

If a transcript τ is not bad, then it will be called a *good* transcript. The probability that the bad event occurs is obtained as follows:

- Since the tag is random in the ideal world, we have

$$\Pr[\text{bad}_1] = \frac{\binom{q}{n}}{(2^n)^{n-1}} \leq \left(\frac{2q}{2^n}\right)^n \leq \frac{2q}{2^n}$$

since $q \leq 2^{n-1}$ and $\Pr[\text{bad}_2] \leq \frac{q^2 \epsilon}{2^n}$.

- For each $j \in [q + 1, q + v]$, there is at most one $i \in [q]$ such that $N_i = N_j$. For such pair (i, j) , one has $\Pr[X_i = X_j] \leq \epsilon$. Therefore, we have

$$\Pr[\text{bad}_3] \leq v\epsilon.$$

Therefore, we have

$$\Pr[\text{bad}] \leq \Pr[\text{bad}_1] + \Pr[\text{bad}_2] + \Pr[\text{bad}_3] \leq \frac{2q}{2^n} + \frac{q^2 \epsilon}{2^n} + v\epsilon. \tag{4}$$

Good Transcript Analysis. For a good transcript τ and its system Γ_τ , by assuming nonces are not repeated, we observe that

- Γ_τ is nice by $\neg(\text{bad}_2 \vee \text{bad}_3)$;
- $\xi_{\max} \leq n + 1$ and $\xi_{\max}^2 n + \xi_{\max} \leq n(n + 1)^2 + n + 1 \leq 2^{n/2}$ since $n \geq 30$ by $\neg\text{bad}_1$.

Henceforth, we can apply Theorem 1 and then we have

$$h(\Gamma_\tau) \geq \frac{(2^n - 2)^{|\mathcal{V}_1|} (2^n - 2)^{|\mathcal{V}_2|}}{2^{nq}} \left(1 - \frac{2v}{2^n}\right).$$

Furthermore, we see that

$$\mathbf{p}_{\mathcal{S}_0}^{q+v}(\tau) = \frac{1}{|\mathcal{K}_h|} \cdot \frac{1}{2^{nq}}, \quad \mathbf{p}_{\mathcal{S}_1}^{q+v}(\tau) = \frac{1}{|\mathcal{K}_h|} \cdot \frac{h(\Gamma_\tau)}{(2^n)^{|\mathcal{V}_1|} (2^n)^{|\mathcal{V}_2|}}.$$

From the above, one has

$$\begin{aligned} \frac{\mathbf{p}_{\mathcal{S}_1}^{q+v}(\tau)}{\mathbf{p}_{\mathcal{S}_0}^{q+v}(\tau)} &= \frac{h(\Gamma_\tau) 2^{nq}}{(2^n)^{|\mathcal{V}_1|} (2^n)^{|\mathcal{V}_2|}} \\ &\geq \frac{(2^n - 2)^{|\mathcal{V}_1|} (2^n - 2)^{|\mathcal{V}_2|}}{(2^n)^{|\mathcal{V}_1|} (2^n)^{|\mathcal{V}_2|}} \left(1 - \frac{2v}{2^n}\right) \\ &= \frac{(2^n - |\mathcal{V}_1|)_2 (2^n - |\mathcal{V}_2|)_2}{(2^n)_2 (2^n)_2} \left(1 - \frac{2v}{2^n}\right) \\ &\geq \left(1 - \frac{q + v}{2^n}\right)^4 \left(1 - \frac{2v}{2^n}\right) \\ &\geq 1 - \frac{4q}{2^n} - \frac{6v}{2^n} \end{aligned} \tag{5}$$

since $|\mathcal{V}_1|, |\mathcal{V}_2| \leq q + v$.

Plugging (4) and (5) to Lemma 1, we conclude that

$$\|p_{S_0}^{q+v}(\cdot) - p_{S_1}^{q+v}(\cdot)\| \leq \frac{6q}{2^n} + \frac{q^2\epsilon}{2^n} + \frac{6v}{2^n} + v\epsilon.$$

5 Security of $F_{B_2}^{\text{SoP}}$ and $F_{B_3}^{\text{SoP}}$

In this section, we consider $F_{B_2}^{\text{SoP}}[H, E]$ and $F_{B_3}^{\text{SoP}}[H, E]$ based on an n -bit (n, ϵ_n) -AXU hash function H and an n -bit block cipher E . For given n -bit nonce N and a message M , the user receives a tag as

$$E_{K_1}(N) \oplus E_{K_2}(H_{K_h}(M) \oplus N)$$

for $F_{B_2}^{\text{SoP}}[H, E]$, and

$$E_{K_1}(N) \oplus E_{K_2}(H_{K_h}(M) \oplus N) \oplus H_{K_h}(M)$$

for $F_{B_3}^{\text{SoP}}[H, E]$ by a hash key K_h and block cipher keys K_1 and K_2 . This section aims to prove the security of $F_{B_2}^{\text{SoP}}[H, E]$ and $F_{B_3}^{\text{SoP}}$. As a result, we have the following theorem and corollary.

Theorem 4. *Let $\epsilon > 0$ and $n \geq 32$. Let $H : \mathcal{K}_h \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be an ϵ -AXU hash function and (n, ϵ_n) -AXU hash function, and $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. Let μ, q, v, t be nonnegative integers such that $n\mu \leq 2^{n/4}$, $12(\mu + n + 2)^2q \leq 2^n$ and $q + v \leq 2^{n-1}$. Then, one has*

$$\begin{aligned} \text{Adv}_{F_{B_2}^{\text{SoP}}[H, E]}^{\text{mac}}(\mu, q, v, t) &\leq \binom{q}{n} \epsilon_n + 2\mu q \epsilon + \mu^2 \epsilon + \frac{\mu^2}{2^n} + \frac{q^2 \epsilon}{2^n} + \frac{4q}{2^n} + v\epsilon + \frac{6v}{2^n} \\ &\quad + 2\text{Adv}_E^{\text{prp}}(q + v, t + t') \end{aligned}$$

where t' is the time complexity necessary to compute E for $q + v$ times.

Since adding $H_{K_h}(M)$ to the tag does not make any significant difference, the MAC security of $F_{B_3}^{\text{SoP}}$ follows immediately.

Corollary 2. *Let $\epsilon > 0$ and $n \geq 32$. Let $H : \mathcal{K}_h \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be an ϵ -AXU hash function and (n, ϵ_n) -AXU hash function, and $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. Let μ, q, v, t be nonnegative integers such that $n\mu \leq 2^{n/4}$, $12(\mu + n + 2)^2q \leq 2^n$ and $q + v \leq 2^{n-1}$. Then, one has*

$$\begin{aligned} \text{Adv}_{F_{B_3}^{\text{SoP}}[H, E]}^{\text{mac}}(\mu, q, v, t) &\leq \binom{q}{n} \epsilon_n + 2\mu q \epsilon + \mu^2 \epsilon + \frac{\mu^2}{2^n} + \frac{q^2 \epsilon}{2^n} + \frac{4q}{2^n} + v\epsilon + \frac{6v}{2^n} \\ &\quad + 2\text{Adv}_E^{\text{prp}}(q + v, t + t') \end{aligned}$$

where t' is the time complexity necessary to compute E for $q + v$ times.

We claim that CBC-MAC is a multi-xor-collision resistant hash function. More specifically, for any distinct $M_1, \dots, M_n \in (\{0, 1\}^n)^*$ and distinct $Y_1, \dots, Y_n \in \{0, 1\}^n$, the following holds:

$$\begin{aligned} \Pr[\pi \leftarrow_{\$} \text{Perm}(n) : \text{CBC-MAC}[\pi](X_1) \oplus Y_1 = \dots = \text{CBC-MAC}[\pi](X_1) \oplus Y_1] \\ \leq \frac{1}{(2^n - n(\ell + 1))^{n-1}} + \left(\frac{(n\ell + 1)^2}{2^n} \right)^n \end{aligned}$$

when $n(\ell + 1) \leq 2^{n-1}$. At the last of this section, we prove the claim in Lemma 4. When the underlying hash function is instantiated with CBC-MAC, we have the following corollary.

Corollary 3. *Let $\epsilon > 0$ and $n \geq 32$. Let $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. Let μ, q, v, t be nonnegative integers such that $n\mu \leq 2^{n/4}$, $12(\mu + n + 2)^2 q \leq 2^n$ and $q + v \leq 2^{n-1}$. Let $n(\ell + 1) \leq 2^{n-1}$ where ℓ be the maximum block length of MAC queries. Then, one has*

$$\begin{aligned} \text{Adv}_{F_{B_2}^{\text{SoP}}}^{\text{mac}}[\text{CBC-MAC}, E](\mu, q, v, t) \leq \frac{q(n\ell + 1)^2}{2^n} + 2\mu q\epsilon + \mu^2\epsilon + \frac{\mu^2}{2^n} + \frac{q^2\epsilon}{2^n} + \frac{4q}{2^n} + v\epsilon \\ + \frac{6v}{2^n} + 2\text{Adv}_E^{\text{prp}}(q + v, t + t') \end{aligned}$$

and

$$\begin{aligned} \text{Adv}_{F_{B_3}^{\text{SoP}}}^{\text{mac}}[\text{CBC-MAC}, E](\mu, q, v, t) \leq \frac{q(n\ell + 1)^2}{2^n} + 2\mu q\epsilon + \mu^2\epsilon + \frac{\mu^2}{2^n} + \frac{q^2\epsilon}{2^n} + \frac{4q}{2^n} + v\epsilon \\ + \frac{6v}{2^n} + 2\text{Adv}_E^{\text{prp}}(q + v, t + t') \end{aligned}$$

where t' is the time complexity necessary to compute E for $q + v$ times.

5.1 Proof of Theorem 4

Similarly to the proof of Theorem 3, we assume that the adversary is deterministic and never makes a redundant query. Up to the prp-security of E , keyed block ciphers E_{K_1} and E_{K_2} can be replaced by truly random permutations P_1 and P_2 , respectively. The cost of this replacement is upper bounded by

$$2\text{Adv}_E^{\text{prp}}(q + v, t + t').$$

The resulting construction denotes $F_{B_2}^{\text{SoP}*}[H]$. At the end of the interaction, additional information K_h is freely given to an adversary. Without loss of generality, we rearrange query indices so that verification queries come after MAC queries.

Let Θ be the set of all attainable transcripts in the ideal world and $\tau = (\tau_m, \tau_v, K_h) \in \Theta$ be a transcript where τ_m and τ_v denote the list of MAC queries and the list of verification queries, i.e.,

$$\begin{aligned} \tau_m &= \{(N_1, M_1, T_1), \dots, (N_q, M_q, T_q)\}, \\ \tau_v &= \{(N_{q+1}, M_{q+1}, T_{q+1}, b_{q+1}), \dots, (N_{q+v}, M_{q+v}, T_{q+v}, b_{q+v})\}. \end{aligned}$$

From a transcript τ , \mathcal{A} can compute $X_i = H_{K_h}(M_i) \oplus N_i$ for $i \in [q + v]$ before outputting its decision bit.

The core of the security proof is to estimate the number of possible ways of fixing evaluations P_1 and P_2 in a way that

$$T_i = P_1(N_i) \oplus P_2(X_i)$$

for $i = 1, \dots, q$ and

$$T_i \neq P_1(N_i) \oplus P_2(X_i)$$

for $i = q + 1, \dots, q + v$. We will identify $\mathcal{V}_1 = \{P_1(N_i)\}$ and $\mathcal{V}_2 = \{P_2(X_i)\}$ with as sets of variables. We also define $\mathcal{V} = \mathcal{V}_1 \sqcup \mathcal{V}_2$. Then we can construct the system of equations Γ_τ as defined in Sect. 3. To satisfy the conditions in Theorem 1, we must first define bad events on a transcript τ , and then we can apply the extended Mirror theory to each transcript that the bad event does not happen.

Defining and Bounding Bad Events. A transcript $\tau = (\tau_m, \tau_v, K_h)$ is defined as *bad* if one of the following condition holds.

- $\text{bad}_1 \Leftrightarrow$ there exists $(i_1, \dots, i_n) \in [q]^{*n}$ where N_{i_1}, \dots, N_{i_n} are all distinct such that $X_{i_1} = \dots = X_{i_n}$.
- $\text{bad}_2 \Leftrightarrow$ there exists $(i, j) \in [q]^{*2}$ such that $N_i = N_j$ and $X_i = X_j$.
- $\text{bad}_3 \Leftrightarrow$ there exists $(i, j) \in [q]^{*2}$ such that $N_i = N_j$ and $T_i = T_j$.
- $\text{bad}_4 \Leftrightarrow$ there exists $(i, j) \in [q]^{*2}$ such that $X_i = X_j$ and $T_i = T_j$.
- $\text{bad}_5 \Leftrightarrow$ there exists $(i, j, k) \in [q]^{*3}$ such that $N_i = N_j$ and $X_j = X_k$.
- $\text{bad}_6 \Leftrightarrow$ there exists $(i, j) \in [q] \times [q + 1, q + v]$ such that $N_i = N_j$, $X_i = X_j$ and $T_i = T_j$.

If a transcript τ is not bad, then it will be called a *good* transcript. Now, we upper bound the probability happens *bad* in the ideal world by the following:

1. Since H is (n, ϵ_n) -AXU hash function, we have

$$\Pr[\text{bad}_1] \leq \binom{q}{n} \epsilon_n.$$

2. By symmetry, we can assume that $i < j$, which means that N_j is a faulty nonce. For each MAC query using a faulty nonce, there are at most μ other queries using the same nonce. So, the number of pairs (i, j) such that $i < j$ and $N_i = N_j$ is at most μ^2 . For each of such pair (i, j) , the probability that $X_i = X_j$ is ϵ . Therefore, we have

$$\Pr[\text{bad}_2] \leq \mu^2 \epsilon.$$

Similarly, we can show that

$$\Pr[\text{bad}_3] \leq \frac{\mu^2}{2^n}$$

and

$$\Pr [\text{bad}_4] \leq \frac{q^2\epsilon}{2^n}$$

3. The number of indices j such that $N_i = N_j$ is at most 2μ . So, the number of choices of (j, k) is at most $2\mu q$. For each of such pairs, the probability that $X_j = X_k$ is at most ϵ . Therefore, we have

$$\Pr [\text{bad}_5] \leq 2\mu q\epsilon.$$

4. Suppose bad_3 does not occur. When an adversary makes a verification query (N_j, M_j, T_j) , there is one MAC query (N_i, M_i, T_i) such that $N_i = N_j$ and $T_i = T_j$. For each of such pairs, the probability that $X_i = X_j$ is at most ϵ . Therefore, we have

$$\Pr [\text{bad}_6 \mid \neg\text{bad}_3] \leq v\epsilon.$$

To sum up, we have

$$\begin{aligned} \Pr [\text{bad}] &= \Pr [\text{bad}_1] + \Pr [\text{bad}_2] + \Pr [\text{bad}_3] + \Pr [\text{bad}_4] + \Pr [\text{bad}_5] + \Pr [\text{bad}_6 \mid \neg\text{bad}_3] \\ &\leq \binom{q}{n} \epsilon_n + 2\mu q\epsilon + \mu^2\epsilon + \frac{\mu^2}{2^n} + \frac{q^2\epsilon}{2^n} + v\epsilon. \end{aligned} \tag{6}$$

Good Transcript Analysis. For a good transcript τ and its system of equations Γ_τ , we observe that

- Γ_τ is nice by $\neg(\text{bad}_2 \vee \text{bad}_3 \vee \text{bad}_4 \vee \text{bad}_5 \vee \text{bad}_6)$. Since $\neg(\text{bad}_2 \vee \text{bad}_5)$, for any component $\{X_{i,0}, \dots, X_{i,\xi_i-1}\}$, $X_{i,0} \not\sim X_{i,j}$ for $1 \leq j \leq \xi_i - 1$, which means the first condition holds. The second and the third conditions are satisfied by $\neg(\text{bad}_3 \vee \text{bad}_4)$ and $\neg\text{bad}_6$.
- By $\neg(\text{bad}_1 \vee \text{bad}_5)$, $\xi_{\max} \leq \max\{\mu + 1, n + 1\}$. Therefore, we have

$$\xi_{\max}^2 n + \xi_{\max} \leq n(\mu + n + 2)^2 + \mu + n + 2 \leq 2^{n/2}$$

since $n \geq 32$ and $n\mu \leq 2^{n/4}$. We also have $q\xi_{\max}^2 \leq (\mu + n + 2)^2 q \leq \frac{2^n}{12}$.

Henceforth, we can apply Theorem 1 and then we have

$$h(\Gamma_\tau) \geq \frac{(2^n - 2)^{|\mathcal{V}_1|} (2^n - 2)^{|\mathcal{V}_2|}}{2^{nq}} \left(1 - \frac{2v}{2^n}\right).$$

Furthermore, we see that

$$\mathbf{p}_{S_0}^{q+v}(\tau) = \frac{1}{|\mathcal{K}_h|} \cdot \frac{1}{2^{nq}}, \quad \mathbf{p}_{S_1}^{q+v}(\tau) = \frac{1}{|\mathcal{K}_h|} \cdot \frac{h(\Gamma_\tau)}{(2^n)^{|\mathcal{V}_1|} (2^n)^{|\mathcal{V}_2|}}.$$

From the above, since $|\mathcal{V}_1|, |\mathcal{V}_2| \leq q + v$, one has

$$\begin{aligned}
 \frac{\mathbf{p}_{S_1}^{q+v}(\tau)}{\mathbf{p}_{S_0}^{q+v}(\tau)} &= \frac{h(\Gamma_\tau)2^{nq}}{(2^n)^{|\mathcal{V}_1|}(2^n)^{|\mathcal{V}_2|}} \\
 &\geq \frac{(2^n - 2)^{|\mathcal{V}_1|}(2^n - 2)^{|\mathcal{V}_2|}}{(2^n)^{|\mathcal{V}_1|}(2^n)^{|\mathcal{V}_2|}} \left(1 - \frac{2v}{2^n}\right) \\
 &= \frac{(2^n - |\mathcal{V}_1|)_2(2^n - |\mathcal{V}_2|)_2}{(2^n)_2(2^n)_2} \left(1 - \frac{2v}{2^n}\right) \\
 &\geq \left(1 - \frac{q+v}{2^n}\right)^4 \left(1 - \frac{2v}{2^n}\right) \\
 &\geq 1 - \frac{4q}{2^n} - \frac{6v}{2^n}.
 \end{aligned} \tag{7}$$

Plugging (6) and (7) to Lemma 1, we conclude that

$$\|\mathbf{p}_{S_0}^{q+v}(\tau) - \mathbf{p}_{S_1}^{q+v}(\tau)\| \leq \binom{q}{n} \epsilon_n + 2\mu q \epsilon + \mu^2 \epsilon + \frac{\mu^2}{2^n} + \frac{q^2 \epsilon}{2^n} + \frac{4q}{2^n} + v \epsilon + \frac{6v}{2^n}.$$

5.2 Multi-xor-Collision Probability of CBC-MAC

We state an example of a multi-xor-collision resistant hash function. We consider CBC-MAC $[\pi]$ based on pseudorandom permutation π . For a permutation π and a message $M = (M[1], \dots, M[m]) \in (\{0, 1\}^n)^m$ with m blocks, the tag is given by

$$\text{CBC-MAC}[\pi](M) = X[m]$$

where $X[i] = \pi(X[i - 1] \oplus M[i])$ for $i \in [m]$ and $X[0] = 0$. We will show that CBC-MAC is a $(n, \frac{2}{2^{(n-1)^2}})$ -AXU hash function.

We fix n distinct messages $M_1, \dots, M_n \in (\{0, 1\}^n)^*$ and n distinct strings $Y_1, \dots, Y_n \in \{0, 1\}^n$ throughout this section. We use m_i to denote the block length of M_i and let $\ell = \max_{i \in [n]} \{m_i\}$. For simplicity, we assume that the length of each message is a multiple of n .

We define an n -multi-collision event

$$\text{Coll}_\pi \Leftrightarrow \text{CBC-MAC}[\pi](M_1) \oplus Y_1 = \dots = \text{CBC-MAC}[\pi](M_n) \oplus Y_n.$$

Equivalently, the collision event is regarded as

$$\text{Coll}_\pi \Leftrightarrow \text{CBC-MAC}[\pi](M_1 \parallel Y_1) = \dots = \text{CBC-MAC}[\pi](M_n \parallel Y_n).$$

We bound the probability of Coll_π by the following lemma:

Lemma 4. *With the above notations, suppose that $n(\ell + 1) \leq 2^{n-1}$. Then, we have*

$$\Pr[\pi \leftarrow_{\S} \text{Perm}(n) : \text{Coll}_\pi] \leq \frac{1}{(2^n - n(\ell + 1))^{n-1}} + \left(\frac{(n\ell + 1)^2}{2^n}\right)^n.$$

Proof. Let $\mathcal{M} = (M_1 \parallel Y_1, \dots, M_n \parallel Y_n)$ and $m = \sum_{i=1}^n m_i + n$. We first represent a relation of internal outputs through the computation of CBC-MAC via the structure graph. The intermediate values will be defined as sequences over a two-dimensional index set. Each index is a pair where the first element of the pair corresponds to the message number and the second element is the block number of that message. We define the index set

$$\mathcal{I} = \{(r, i) \mid r \in [k], i \in [m_r]\}$$

and the dictionary order \prec on it as follows: $(r, i) \prec (s, j)$ if $r < s$ or $r = s$ and $i < j$. We also consider the index set $\mathcal{I}_0 = \mathcal{I} \cup \{(r, 0) \mid r \in [q]\}$ and the natural extension of the order \prec on \mathcal{I}_0 .

For any $\pi \in \text{Perm}(n)$, we build the structure graph G_π , which is a directed graph (V, E) as follows:

- For any $\pi \in \text{Perm}(n)$, we denote the intermediate values for each message as

$$X_\pi[r, i] = \pi(X_\pi[r, i - 1] \oplus M_r[i])$$

for $(r, i) \in \mathcal{I}$ and $X_\pi[r, 0] = \mathbf{0}$ for $r \in [q]$.

- From this $X[r, i]$'s, we define the mapping $[\cdot]_\pi : \mathcal{I}_0 \rightarrow \mathcal{I}_0$ as $[(r, i)]_\pi = \min \{(s, j) \mid X_\pi[s, j] = X_\pi[r, i]\}$ where the minimum is determined through the dictionary order. Now the structure graph $G_\pi = (V, E)$ is given by

$$\begin{aligned} V &= \{[(r, i)]_\pi \mid (r, i) \in \mathcal{I}_0\}, \\ E &= \{([(r, i - 1)]_\pi, [(r, i)]_\pi; M_r[i]) \mid r \in [q], i \in [m_r]\}. \end{aligned}$$

Note that $[(r, 0)]_\pi = (1, 0)$ for $r \in [n]$.

We define a binary function Iszero such that for a structure graph G_π , $\text{Iszero}(G_\pi) = 1$ if the vertex $(1, 0)$ has positive in-degree, otherwise it maps to 0. We say that G_π has a collision in a vertex z if there exist u and v such that $e_1 \stackrel{\text{def}}{=} (u, z; L_u), e_2 \stackrel{\text{def}}{=} (v, z; L_v) \in E$. Then, we must have $X[u] \oplus X[v] = L_u \oplus L_v$. For all collisions, the collection of those linear equations is denoted \mathcal{L} . Let $\text{rank}(G_\pi)$ denote the rank of \mathcal{L} . We define the accident of a structure graph G_π as $\text{Acc}(G_\pi) \stackrel{\text{def}}{=} \text{rank}(G_\pi) + \text{Iszero}(G_\pi)$.

Coll_π occurs if and only if $\text{Acc}(G_\pi) \geq n - 1$ since the last blocks of all messages are pairwise distinct. Moreover, at least $n - 1$ accidents occur at a vertex $(1, m_1)$. Similarly to Proposition 2 in [21], we have

$$\Pr[\pi \leftarrow_{\S} \text{Perm}(n) : \text{Coll}_\pi] \leq \frac{A}{(2^n - m)^{n-1}} + \left(\frac{m^2}{2^n}\right)^n$$

where A is the number of all structure graphs with $n - 1$ accidents and satisfying Coll_π . It is easy to see that $A \leq 1$ since no collision can occur except the vertex $(1, m_1)$. Therefore, we have

$$\begin{aligned} \Pr[\pi \leftarrow_{\S} \text{Perm}(n) : \text{Coll}_\pi] &\leq \frac{1}{(2^n - m)^{n-1}} + \left(\frac{m^2}{2^n}\right)^n \\ &\leq \frac{1}{(2^n - n(\ell + 1))^{n-1}} + \left(\frac{(n\ell + 1)^2}{2^n}\right)^n \end{aligned}$$

since $m \leq n(\ell + 1)$. □

6 Matching Attack on $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$

In this section, we present a universal forging attack on $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$ with probability $\frac{1}{2}$ using $O(2^{3n/4})$ queries in the nonce-respecting setting. For given n -bit nonce N and a message M , a tag is computed as

$$\begin{aligned} F_{B_4}^{\text{EDM}}[H, E](N, M) &= E_{K_2}(E_{K_1}(N \oplus H_{K_h}(M)) \oplus N), \\ F_{B_5}^{\text{EDM}}[H, E](N, M) &= F_{B_4}^{\text{EDM}}[H, E](N, M) \oplus H_{K_h}(M) \end{aligned}$$

where a hash key K_h and block cipher keys K_1 and K_2 (see Fig. 1). To ease the notation, we show an attack on $F_{B_4}^{\text{EDM}}$ below, but the same idea is easily mounted to $F_{B_5}^{\text{EDM}}$. The attack is described in Algorithm 1 that outputs a valid forgery for a target message $M \in \{0, 1\}^n$.

To compute the success probability of the attack, we analyze the probability of obtaining a specific value for the leftmost bit of the hash difference between two messages. Let $M, M' \in \{0, 1\}^n$ be distinct two messages. For a randomly selected hash key K_h , we assume that at least one bit of $H_{K_h}(M) \oplus H_{K_h}(M')$ is 1 with a high probability. Without loss of generality, we say

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : H_{K_h}(M) \oplus H_{K_h}(M') = 1 \parallel *] \approx \frac{1}{2}. \tag{8}$$

In the following, we state that PolyHash [23] and CBC-MAC satisfy the above property. For input $M \in \{0, 1\}^n$, $\text{Poly}_{K_h} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined as $\text{Poly}_{K_h}(M) = M \cdot K_h$ and $\text{CBC-MAC}[E_{K_h}](M) : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined as

$$\text{CBC-MAC}[E_{K_h}](M) = E_{K_h}(M).$$

Then, we have

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : (M \oplus M') \cdot K_h = 1 \parallel *] \approx \frac{1}{2}$$

and

$$\Pr [K_h \leftarrow_{\S} \mathcal{K}_h : E_{K_h}(M) \oplus E_{K_h}(M') = 1 \parallel *] \approx \frac{1}{2}.$$

Based on the above analysis, we then evaluate the success probability of the attack described in Algorithm 1.

Algorithm 1: A universal forgery attack on $F_{B_4}^{\text{EDM}}$ and $F_{B_5}^{\text{EDM}}$

Input: A target message $M \in \{0, 1\}^n$ **Output:** A set of forgeries \mathcal{F}

```

1  $\mathcal{F} \leftarrow \emptyset$ 
   // First Phase
2  $M' \leftarrow_{\S} \{0, 1\}^n \setminus \{M\}$ 
3  $\text{Used} \leftarrow \emptyset$ 
4 for  $i \leftarrow 0$  to  $2^{3n/4} - 1$  do
5    $N_i \leftarrow 0^{n/4} \parallel \langle i \rangle_{3n/4}$ 
6    $N'_i \leftarrow 1 \parallel \langle i \rangle_{3n/4} \parallel 0^{n/4-1}$ 
7    $T_i \leftarrow \mathcal{O}(N_i, M)$ 
8    $T'_i \leftarrow \mathcal{O}(N'_i, M')$ 
9    $\text{Used} \leftarrow \text{Used} \cup \{N_i, N'_i\}$ 
10  $\mathcal{Y} \leftarrow \emptyset$ 
11 if  $\exists(i, j, k, l)$  such that  $(N_i \oplus N_j \oplus N'_k \oplus N'_l = \mathbf{0}) \wedge (T_i = T_j) \wedge (T'_k = T'_l)$ 
   then
12    $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{N_i \oplus N'_k\}$ 
13 if  $\mathcal{Y} = \emptyset$  then
14   return  $\perp$ 
   // Second Phase
15 for  $i \leftarrow 0$  to  $2^{n/2} - 1$  do
16    $\bar{N}_i \leftarrow \{0, 1\}^n \setminus \text{Used}$ 
17    $\bar{T}_i \leftarrow \mathcal{O}(\bar{N}_i, M')$ 
18   if  $\exists(i, j)$  such that  $\bar{T}_i = \bar{T}_j$  then
19     for  $Y \in \mathcal{Y}$  do
20        $T \leftarrow \mathcal{O}(\bar{N}_i \oplus Y, M)$ 
21        $\mathcal{F} \leftarrow \mathcal{F} \cup \{\bar{N}_j \oplus Y, M, T\}$ 
22     return  $\mathcal{F}$ 
23 return  $\perp$ 

```

Theorem 5. Let \mathcal{A}^* be an adversary running Algorithm 1. Then,

$$\text{Adv}_{F_{B_4}^{\text{EDM}}}^{\text{mac}}(\mathcal{A}^*) \approx \frac{1}{4}$$

where the error is exponentially small.

Proof. We argue that \mathcal{A}^* can find at least one pair (i, j, k, l) with a high probability. Suppose that $H_{K_h}(M) \oplus H_{K_h}(M') = N_i \oplus N'_k = N_j \oplus N'_l$. Then, it holds

$$\begin{aligned} T_i = T_j &\Leftrightarrow E_{K_1}(N_i \oplus H_{K_h}(M)) \oplus N_i = E_{K_1}(N_j \oplus H_{K_h}(M)) \oplus N_j \\ &\Leftrightarrow E_{K_1}(N'_k \oplus H_{K_h}(M')) \oplus N'_k = E_{K_1}(N'_l \oplus H_{K_h}(M')) \oplus N'_l \\ &\Leftrightarrow T'_k = T'_l \end{aligned}$$

For each quadruple (i, j, k, l) , the probability that $T_i = T_j$ and $T'_k = T'_l$ is $\frac{1}{2^n}$ if $H_{K_h}(M) \oplus H_{K_h}(M') = N_i \oplus N'_k = N_j \oplus N'_l$. Otherwise, the probability is $\frac{1}{2^{2n}}$.

For $2^{n-1} \leq y \leq 2^n - 1$, there are $q' = 2^{n/2+1}$ tuples of indices (i, j) such that $N_i \oplus N'_j = \langle y \rangle_n$. If $H_{K_h}(M) \oplus H_{K_h}(M') = \langle y \rangle_n$ for some y , \mathcal{A}^* can find a quadruple (i, j, k, l) such that $N_i \oplus N'_k = N_j \oplus N'_l = \langle y \rangle_n$ with overwhelming probability since the expected number is $\frac{\binom{q'}{2}}{2^n} \geq 1$. So, the probability \mathcal{Y} contains the real hash difference is, by (8),

$$\Pr [K_h \leftarrow_{\$} \mathcal{K}_h : H_{K_h}(M) \oplus H_{K_h}(M') = 1 \parallel *] \approx \frac{1}{2} \tag{9}$$

and the expected size of \mathcal{Y} is

$$1 + (2^n - 1) \frac{\binom{q'}{2}}{2^{2n}} \leq 5.$$

Once the hash difference is found, one can compute a forgery by finding a tag collision in the second phase. Roughly, the attack samples $O(2^{n/2})$ fresh nonces and computes tags with the chosen message M' , where the collision probability of tags is given as

$$\frac{\binom{2^{\frac{n}{2}}}{2}}{2^n} \approx \frac{1}{2}.$$

Combined with (9), the probability of successful forgery is approximately $\frac{1}{4}$. The above attack works when $q \leq 2^{\frac{3n}{4}+2}$ with a constant number of verification queries. □

References

1. Banik, S., Pandey, S.K., Peyrin, T., Sasaki, Y., Sim, S.M., Todo, Y.: GIFT: A Small Present. In: Fischer, W., Homma, N. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2017. LNCS, vol. 10529, pp. 321–345. Springer (2017). https://doi.org/10.1007/978-3-319-66787-4_16
2. Bellare, M., Kilian, J., Rogaway, P.: The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences* **61**(3), 362–399 (2000)
3. Bellare, M., Pietrzak, K., Rogaway, P.: Improved security analyses for cbc macs. In: *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 14-18, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3621, pp. 527–545. Springer (2005). https://doi.org/10.1007/11535218_32
4. Bernstein, D.J.: Stronger Security Bounds for Wegman-Carter-Shoup Authenticators. In: Cramer, R. (ed.) *Advances in Cryptology - EUROCRYPT 2005*. LNCS, vol. 3494, pp. 164–180. Springer (2005). https://doi.org/10.1007/11426639_10





5. Bernstein, D.J.: The Poly1305-AES message-authentication code. In: Gilbert, H., Handschuh, H. (eds.) Fast Software Encryption. LNCS, vol. 3557, pp. 32–49. Springer (2005). https://doi.org/10.1007/11502760_3
6. Bhargavan, K., Leurent, G.: On the Practical (In-)Security of 64-bit Block Ciphers: Collision Attacks on HTTP over TLS and OpenVPN. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS. pp. 456–467. ACM (2016). <https://doi.org/10.1145/2976749.2978423>
7. Black, J., Rogaway, P.: A Block-Cipher Mode of Operation for Parallelizable Message Authentication. In: Knudsen, L.R. (ed.) Advances in Cryptology - EUROCRYPT 2002. LNCS, vol. 2332, pp. 384–397. Springer (2002). https://doi.org/10.1007/3-540-46035-7_25, <https://iacr.org/archive/eurocrypt2002/23320380/pmac.pdf>
8. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., Seurin, Y., Vikkelsoe, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbauwhede, I. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer (2007). https://doi.org/10.1007/978-3-540-74735-2_31
9. Chen, Y.L., Mennink, B., Preneel, B.: Categorization of Faulty Nonce Misuse Resistant Message Authentication. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology - ASIACRYPT 2021. LNCS, vol. 13092, pp. 520–550. Springer (2021). https://doi.org/10.1007/978-3-030-92078-4_18
10. Choi, W., Lee, B., Lee, Y., Lee, J.: Improved security analysis for nonce-based enhanced hash-then-mask MACs. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology - ASIACRYPT 2020. LNCS, vol. 12491, pp. 697–723. Springer (2020). https://doi.org/10.1007/978-3-030-64837-4_23
11. Choi, W., Lee, J., Lee, Y.: Toward full n -bit security and nonce misuse resistance of block cipher-based MACs. Cryptology ePrint Archive, Paper 2024/731 (2024), <https://eprint.iacr.org/2024/731>
12. Cogliati, B., Dutta, A., Nandi, M., Patarin, J., Saha, A.: Proof of Mirror Theory for a Wide Range of ξ_{\max} . In: Hazay, C., Stam, M. (eds.) Advances in Cryptology - EUROCRYPT 2023. LNCS, vol. 14007, pp. 470–501. Springer (2023). https://doi.org/10.1007/978-3-031-30634-1_16
13. Cogliati, B., Seurin, Y.: EWCDM: An Efficient, Beyond-Birthday Secure, Nonce-Misuse Resistant MAC. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology - CRYPTO 2016. LNCS, vol. 9814, pp. 121–149. Springer (2016). https://doi.org/10.1007/978-3-662-53018-4_5
14. Datta, N., Dutta, A., Dutta, K.: Improved Security Bound of (E/D)WCDM. IACR Transactions on Symmetric Cryptology **Issue 4**, 138–176 (2021). <https://doi.org/10.46586/tosc.v2021.i4.138-176>
15. Datta, N., Dutta, A., Nandi, M., Yasuda, K.: Encrypt or Decrypt? To Make a Single-Key Beyond Birthday Secure Nonce-Based MAC. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology - CRYPTO 2018. LNCS, vol. 10991, pp. 631–661. Springer (2018). https://doi.org/10.1007/978-3-319-96884-1_21
16. Dutta, A., Nandi, M., Talnikar, S.: Beyond Birthday Bound Secure MAC in Faulty Nonce Model. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology - EUROCRYPT 2019. LNCS, vol. 11476, pp. 437–466. Springer (2019). https://doi.org/10.1007/978-3-030-17653-2_15

17. Guo, J., Peyrin, T., Poschmann, A., Robshaw, M.J.B.: The LED Block Cipher. In: Preneel, B., Takagi, T. (eds.) *Cryptographic Hardware and Embedded Systems - CHES 2011*. LNCS, vol. 6917, pp. 326–341. Springer (2011). https://doi.org/10.1007/978-3-642-23951-9_22
18. Handschuh, H., Preneel, B.: Key-recovery attacks on universal hash function based mac algorithms. In: *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008*. Lecture Notes in Computer Science, vol. 5157, pp. 144–161. Springer (2008). https://doi.org/10.1007/978-3-540-85174-5_9, <https://iacr.org/archive/crypto2008/51570145/51570145.pdf>
19. Information technology – Security techniques – Message Authentication Codes (MACs) – Part 1: Mechanisms using a block cipher. Standard, International Organization for Standardization (Mar 2011)
20. Iwata, T., Kurosawa, K.: OMAC: One-Key CBC MAC. In: Johansson, T. (ed.) *Fast Software Encryption*. LNCS, vol. 2887, pp. 129–153. Springer (2003). https://doi.org/10.1007/978-3-540-39887-5_11, <https://iacr.org/archive/fse2003/28870137/28870137.pdf>
21. Jha, A., Nandi, M.: Revisiting structure graphs: Applications to cbc-mac and emac. *Journal of Mathematical Cryptology* **10**(3-4), 157–180 (2016)
22. Mennink, B., Neves, S.: Encrypted Davies-Meyer and Its Dual: Towards Optimal Security Using Mirror Theory. In: Katz, J., Shacham, H. (eds.) *Advances in Cryptology - CRYPTO 2017*. LNCS, vol. 10403, pp. 556–583. Springer (2017). https://doi.org/10.1007/978-3-319-63697-9_19
23. Minematsu, K., Iwata, T.: Building blockcipher from tweakable blockcipher: Extending fse 2009 proposal. In: *Cryptography and Coding: 13th IMA International Conference, IMACC 2011, Oxford, UK, December 12-15, 2011*. Proceedings 13. pp. 391–412. Springer (2011)
24. Morris J. Dworkin: Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication **800-38D** (Nov 28 2007)
25. Naito, Y.: Blockcipher-Based MACs: Beyond the Birthday Bound Without Message Length. In: *ASIACRYPT (3)*. pp. 446–470. Springer (2017). https://doi.org/10.1007/978-3-319-70700-6_16
26. Patarin, J.: Introduction to Mirror Theory: Analysis of Systems of Linear Equalities and Linear Non Equalities for Cryptography. IACR Cryptology ePrint Archive, Report 2010/287 (2010), available at <https://eprint.iacr.org/2010/287>
27. Patarin, J.: Mirror Theory and Cryptography. IACR Cryptology ePrint Archive, Report 2016/702 (2016), available at <https://eprint.iacr.org/2016/702>
28. Pietrzak, K.: A tight bound for emac. In: *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II* 33. pp. 168–179. Springer (2006)
29. Shoup, V.: On Fast and Provably Secure Message Authentication Based on Universal Hashing. In: Kobitz, N. (ed.) *Advances in Cryptology - CRYPTO '96*. LNCS, vol. 1109, pp. 313–328. Springer (1996). https://doi.org/10.1007/3-540-68697-5_24
30. Wegman, M.N., Carter, J.L.: New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences* **22**, Issue 3, 265–279 (1981)
31. Yasuda, K.: The sum of CBC MACs is a secure PRF. In: *Cryptographers' Track at the RSA Conference*. pp. 366–381. Springer (2010)

32. Yasuda, K.: A New Variant of PMAC: Beyond the Birthday Bound. In: Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference. Lecture Notes in Computer Science, vol. 6841, p. 593. Springer (2011). https://doi.org/10.1007/978-3-642-22792-9_34, <https://www.iacr.org/archive/crypto2011/68410593/68410593.pdf>
33. Zhang, L., Wu, W., Sui, H., Wang, P.: 3kf9: Enhancing 3GPP-MAC beyond the Birthday Bound. In: ASIACRYPT. vol. 7658, pp. 296–312. Springer (2012). https://doi.org/10.1007/978-3-642-34961-4_19, <https://www.iacr.org/archive/asiacrypt2012/76580291/76580291.pdf>



General Practical Cryptanalysis of the Sum of Round-Reduced Block Ciphers and ZIP-AES

Antonio Flórez-Gutiérrez¹ , Lorenzo Grassi² , Gregor Leander² ,
Ferdinand Sibleyras¹, and Yosuke Todo¹ 

¹ NTT Social Informatics Laboratories, Tokyo, Japan

{antonio.florez,yosuke.todo}@ntt.com

² Ruhr University Bochum, Bochum, Germany

{lorenzo.grassi,gregor.leander}@rub.de

Abstract. We introduce a new approach between classical security proofs of modes of operation and dedicated security analysis for known cryptanalysis families: General Practical Cryptanalysis. This allows us to analyze generically the security of the sum of two keyed permutations against known attacks. In many cases (of course, not all), we show that the security of the sum is strongly linked to that of the composition of the two permutations. This enables the construction of beyond-birthday bound secure low-latency PRFs by cutting a known-to-be-secure block cipher into two equal parts. As a side result, our general analysis shows an inevitable difficulty for the key recovery based on differential-type attacks against the sum, which leads to a correction of previously published attacks on the dedicated design Orthros.


1 Introduction

Symmetric primitives are used to encrypt most of our sensitive data in virtually all applications. Block ciphers are arguably the most studied primitives.


Overhead of Modes. In order to encrypt actual data, primitives have to be used in a mode-of-operation. As a consequence of block ciphers being the most studied primitives, the majority of symmetric-key cryptographic schemes are built as block cipher modes. The advantage of using primitives in a mode-of-operation instead of directly designing an (authenticated) encryption is obvious: a well-designed mode comes with a proof that reduces its security to the security of the primitive. Using such a mode with a well-understood (block) cipher results in a secure scheme. One example is the counter-mode, where a pseudo-random function (PRF) is constructed by encrypting a counter. Indeed, AES-CRT is a frequently used scheme for encryption. In this paper, we instead focus on the sum of two block ciphers. Given two pseudo-random permutations (PRPs) (or independent block ciphers) E_k and E'_k , the sum $E_k(x) \oplus E'_k(x)$ is a secure PRF.

However, modes have a significant overhead. For example, AES-CRT is only secure only up to the birthday bound. For better security, modes with two (or more) calls to the block cipher are required. Turning our focus to the sum-of-PRP construction, we wonder whether it is necessary that both parts are secure PRPs. This question was already posed by the dedicated PRF Orthros [4], which consists of the sum of two specific keyed permutations that would not be secure block ciphers individually. A similar approach was taken in [48], where AES-PRF is proposed as a round-reduced instance of the EDMD construction presented in [47]. The security of AES-PRF required dedicated cryptanalysis to explain why known attacks do not apply. Interestingly, the authors of [48] state that the sum construction seems more risky than the EDMD construction, an opinion we clearly object to as explained below. The main difference with AES-PRF and Orthros is that we are interested in a more general approach.

Link to Composition. As an example, consider a differential attack on the sum construction. One would typically consider an input difference α that would be input to both parts and try to find the most probable output differences β and γ for the individual parts, leading finally to an output difference of $\beta \oplus \gamma$.



The starting point for our work is the observation that the probability for this event, assuming the independence of the parts, is the same as the probability of the following differential trail on the composition of E_1^{-1} and E_2 .



That is, at least intuitively, *the sum construction is as secure as the composition with respect to differential distinguishers*. Ideally, we might hope for a result stating that if $E_2 \circ E_1^{-1}$ is a secure (strong) PRP, then $E_1 \oplus E_2$ is a secure PRF. Before discussing why this is not actually true, let us elaborate on how useful such a statement would be. Such a statement would allow us to *take any secure block cipher, split it into two parts, and obtain a secure PRF*. This would (i) remove the overhead of having two calls to a secure cipher (ii) remove the need for dedicated cryptanalysis as done in Orthros and (iii) result in a PRF with roughly half the latency of the corresponding block cipher.

The problem is, as mentioned, the result is wrong. The easiest example is to take E_1^{-1} to be identity. Then, the resulting scheme is the classical feed-forward construction for which distinguishing attacks exist with square-root complexity. So the main question was if and how this statement might be corrected without losing the great advantages it would provide.

Latency. Latency is an especially important fundamental criterion for the design of symmetric primitives. Indeed, compared to other performance criteria, low latency is much harder to achieve. In a nutshell, asking for a minimal latency cipher is asking about the minimal amount of computation necessary to obtain a secure cipher - a question as fundamental as it is open. Besides being a fundamental property, low latency ciphers have important applications, with memory encryption being one of the most prominent. There are a few dedicated low-latency designs, e.g. PRINCE [16], PRINCEv2 [17], MANTIS [7], QARMA [2], QARMAv2 [3], and SPEEDY [43]. While all these designs use different ideas, their latency seems to converge. Differences in latency are mainly due to different security margins. Substantially improving latency with another block cipher design seems hard if not impossible, which means the possibility of essentially halving the latency with the sum of permutations construction is very enticing.

Our Contribution. It turns out it is possible to show that a practically identical statement holds to an extent. For this, we introduce a new approach which lies between general security reduction on modes of operation and dedicated security analysis of a specific primitive. Specifically, we compare, without analyzing the inside of each component, the security of the sum of two components with their composition. We name this approach General Practical Cryptanalysis.

We show that for many attack families, distinguishers on the sum construction are related to distinguishers on the composition. In the case of the two main attack families, differential and linear distinguishers, as well as their variants, their behaviors are very similar. In particular, (i) differential and linear trails have the same probability/correlation in $E_1 \oplus E_2$ as in $E_2 \circ E_1^{-1}$ or $E_2^{-1} \circ E_1$ and (ii) differential-linear and boomerang distinguishers on $E_2 \circ E_1^{-1}$ are equivalent to differential-and-linear and second order differential distinguishers on $E_1 \oplus E_2$. Of course, there are exceptions; for example, the sum construction is only as strong as the strongest part against the integral attack.

An attack on a symmetric primitive is, in most cases, built from a distinguisher and a key-recovery part. Equally interesting as the results on distinguishers is, therefore, to understand how one can add key-recovery rounds to the different distinguishers on the sum construction. Returning to the example of differential cryptanalysis, it is intuitively clear that adding key recovery at the end is unpromising. Adding key recovery at the top is also more difficult than for the composition, as one has to control both branches simultaneously. We argue that this is only possible under strict conditions. As an interesting side result, our general findings imply that the previous differential attack on Orthros published in [44] must be reviewed.

This novel practical general approach leads to our main result: with respect to the most important attack vectors (with the exceptions mentioned above), the sum $E_1 \oplus E_2$ is as secure as the composition $E_2 \circ E_1^{-1}$. Taking a secure block cipher and splitting it into equal parts, with some additional analysis to cover the exceptions, leads to a PRF that is secure against all known attacks. Of course, this does not rule out the existence of new attacks, but this is the case for all new symmetric primitives.

Instances. To showcase the power and flexibility of our approach, we give a concrete instance in Sect. 4: ZIP-AES, a variant built as the sum of two 5-round AES. This results in a secure PRF with half the latency of AES-CTR and twice the security in terms of data complexity. When implemented with AES-NI, as inverse rounds are more costly, it does not achieve half the latency, but still provides slightly better running times, as detailed in Sect. 4.3.

We finally mention that a ZIP cipher based on a 64-bit lightweight block cipher is promising, e.g., ZIP-GIFT in Sect. 5. The resulting PRF is secure up to the entire 2^{64} blocks, which is enough for all practical cases, while the counter mode of such a 64-bit block cipher can be broken with only 2^{32} blocks of data complexity. Again, not only would security double, but the latency would also be halved, and therefore, it would be very competitive with the dedicated low-latency designs mentioned above.

2 Preliminaries

2.1 Known Attacks on Symmetric Primitives

We work a lot with linear and differential attacks and their variants. We expect the reader to be familiar with them and use this section to fix our notation.

Differential Cryptanalysis [13]. Differential attacks use pairs of plaintexts with a well-chosen difference. For a function $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$, a given input difference $\alpha \in \mathbb{F}_2^n$, and an output difference $\beta \in \mathbb{F}_2^m$, we denote by

$$\text{Prob}(\alpha \xrightarrow{F} \beta) = \frac{|\{x \in \mathbb{F}_2^n \mid F(x) \oplus F(x \oplus \alpha) = \beta\}|}{2^n}$$

the probability that the difference α results in the difference β . Given two (or more) functions $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ and $G : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^\ell$, a differential trail or characteristic for $G \circ F$ also includes an intermediate difference γ . Its probability is usually estimated by multiplying the probabilities

$$\text{Prob}(\alpha \xrightarrow{F} \gamma \xrightarrow{G} \beta) \simeq \text{Prob}(\alpha \xrightarrow{F} \gamma) \cdot \text{Prob}(\gamma \xrightarrow{G} \beta),$$

which can be justified if F and G are key-alternating ciphers with independent round keys and considering the average probability over all keys. From now on, we adopt this independence assumption. Without assumptions, it holds that

$$\text{Prob}(\alpha \xrightarrow{G \circ F} \beta) = \sum_{\gamma} \text{Prob}(\alpha \xrightarrow{F} \gamma \xrightarrow{G} \beta),$$

which is referred to as a differential in contrast to a differential trail.

Linear Cryptanalysis [46]. A linear approximation is a linear combination of input and output bits of the cipher. The main measure of its quality is its correlation. Given a function F , an input mask α , and output mask β , it's

$$\text{cor}_F(\alpha, \beta) = \text{Prob}_x(\langle \beta, F(x) \rangle = \langle \alpha, x \rangle) - \text{Prob}_x(\langle \beta, F(x) \rangle \neq \langle \alpha, x \rangle).$$

Again, given two functions, a linear trail for the composition is specified by an input mask α , an intermediate mask γ , and an output mask β , and its correlation contribution is formally defined as $\text{cor}_F(\alpha, \gamma)\text{cor}_G(\gamma, \beta)$. The set of all linear trails sharing the same input and output masks is often called linear hull. This definition is motivated by the fact that

$$\text{cor}_{G \circ F}(\alpha, \beta) = \sum_{\gamma} \text{cor}_F(\alpha, \gamma)\text{cor}_G(\gamma, \beta).$$

Similarly, given a Boolean function $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, its correlation is

$$\text{cor}(f) = \text{Prob}_x(f(x) = 0) - \text{Prob}_x(f(x) = 1).$$

Differential-Linear Cryptanalysis. The data complexity is given by the autocorrelation, which for an input difference δ and output mask α is defined as

$$\text{Aut}_F(\delta, \alpha) = \text{Prob}_x(\langle \alpha, F(x) \oplus F(x \oplus \delta) \rangle = 0) - \text{Prob}_x(\langle \alpha, F(x) \oplus F(x \oplus \delta) \rangle = 1).$$

In most cases, it is infeasible to obtain all trails in a linear hull or a differential. Hence, security arguments are often based on bounding the probability or correlation of trails. We mainly stick to this approach in this work.

2.2 The Sum-of-PRPs

Constructing PRFs from PRPs is a well-studied topic from a provable security perspective. The sum-of-PRPs construction is a well-known research topic. It was initially introduced by Bellare et al. at EUROCRYPT 1998 [9]. The first proof of its security was given by Lucks at EUROCRYPT 2000 [45], where he proved a suboptimal security bound up to $2^{2n/3}$ queries. This was improved by Bellare and Impagliazzo [8] to $2^n/n$. Finally, with the introduction of the H-coefficient technique, Patarin [49] proved the optimal full n -bit security, and Dutta et al. in [26] filled some gaps in Patarin's proof. Very recently, Dinur [24], using Fourier-analysis, proved optimal bounds for the general case of the sum of permutations and the multi-user setting. A good survey of the state of the art of this and other constructions is given in the later paper as well as in [39].

Complementing this line of work, some recent work has focused on the question of constructing a public function from public (i.e., non-keyed) permutations. This setting requires the notion of indifferenciability and is technically more

involved. After several attempts that turned out to be flawed or non-optimal, the work of Gunsing et al. finally settled the result at CRYPTO 2023 [34].

Despite the general usefulness of constructing a pseudo-random function, there was for a long time no practical cryptanalysis discussion against this construction, mainly because there were no practical instances that have been used or even proposed. The first concrete design was, to the best of our knowledge, Orthros [4]. Motivated by the fact that the output of each pseudo-random permutation is not visible to the attacker, the authors used the so-called proof-then-prune approach [38] to realize an efficient pseudo-random function by reducing the rounds of the two parts. This significantly improved the latency of the resulting scheme but required dedicated cryptanalysis. As discussed below, getting this analysis right is more difficult than usual, in particular when considering differential-type attacks with key recovery.

To capture all designs derived by summing two not necessarily pseudo-random permutations, we give the following general definition.

Definition 1 ($P \oplus Q$). *Let P, Q be two families of permutations, indexed by the keys k_p, k_q in the sets \mathcal{P} and \mathcal{Q} , respectively:*

$$(x, k_p) \in \mathbb{F}_2^n \times \mathcal{P} \mapsto P_{k_p}(x) \in \mathbb{F}_2^n, \quad (x, k_q) \in \mathbb{F}_2^n \times \mathcal{Q} \mapsto Q_{k_q}(x) \in \mathbb{F}_2^n.$$

We define the $P \oplus Q$ construction as the following family of functions:

$$\begin{aligned} P \oplus Q : \mathbb{F}_2^n \times \mathcal{P} \times \mathcal{Q} &\rightarrow \mathbb{F}_2^n \\ (x, (k_p, k_q)) &\mapsto P_{k_p}(x) \oplus Q_{k_q}(x). \end{aligned}$$

Unlike in provable security analysis, it is not assumed that P and Q are pseudo-random permutations. In other words, P and Q are not necessarily secure block ciphers with sound security claims on their own. Our objective is to reveal whether $P \oplus Q$ enhances the practical security in the context of cryptanalysis.

3 General Practical Cryptanalysis of $P \oplus Q$

This section discusses the resistance of the $P \oplus Q$ construction against well-known attack families, and compares it to compositions of P , Q and their inverses. As stated above, for our arguments, we make the usual assumption on the independence of rounds and therefore multiply probabilities over multiple rounds. While for attacks, this tends to lead to flawed complexity estimations, for security arguments there is currently no alternative technique avoiding this.

3.1 Differential Cryptanalysis

Differential Characteristic Equivalence. The differential trails of the parallel construction $P \oplus Q$ are tightly linked to those of the sequential construction $Q \circ P^{-1}$, as shown by the following result:




Fig. 1. Differential and linear trail equivalence

Proposition 1. *Let P, Q be two keyed permutations over \mathbb{F}_2^n , and let $F := P \oplus Q$ and $S := Q \circ P^{-1}$. For each differential trail with probability p traversing F , there is a trail traversing S with the same probability p .*

Proof. Given $\delta_I, \delta_O \in \mathbb{F}_2^n$, we consider the differential $\delta_I \xrightarrow{F} \delta_O$. All its trails take the same form given by the choice of $\gamma \in \mathbb{F}_2^n$ and have probability

$$p = \text{Prob}(\delta_I \xrightarrow{P} \gamma) \cdot \text{Prob}(\delta_I \xrightarrow{Q} \gamma \oplus \delta_O).$$

Since $\text{Prob}(\delta_I \xrightarrow{P} \gamma) = \text{Prob}(\gamma \xrightarrow{P^{-1}} \delta_I)$, p is also the probability of the differential trail $\gamma \xrightarrow{P^{-1}} \delta_I \xrightarrow{Q} \gamma \oplus \delta_O$ traversing S . \square

The left diagram in Fig. 1 shows the trail equivalence between $P \oplus Q$ and $Q \circ P^{-1}$.

Aggregating the Trails. While individual trails of $P \oplus Q$ and $Q \circ P^{-1}$ are equivalent (and thus both have the same maximum differential trail probability), it is hard to compare the resulting differential probabilities when adding up all the trail probabilities in a differential. We can try to compare the expected differential probability (EDP) of both constructions:

$$\begin{aligned} \text{Prob}(\delta_I \xrightarrow{P \oplus Q} \delta_O) &= \sum_{\gamma} \text{Prob}(\gamma \xrightarrow{P^{-1}} \delta_I) \cdot \text{Prob}(\delta_I \xrightarrow{Q} \gamma \oplus \delta_O), \\ \text{Prob}(\delta_I \xrightarrow{Q \circ P^{-1}} \delta_O) &= \sum_{\gamma} \text{Prob}(\delta_I \xrightarrow{P^{-1}} \gamma) \cdot \text{Prob}(\gamma \xrightarrow{Q} \delta_O). \end{aligned}$$

However, we quickly realize that both sums cover sets of differential trails which are non-equivalent, which makes further analysis difficult. Indeed, in the case of $P \oplus Q$, the sum covers all trails $\gamma \xrightarrow{P^{-1}} \delta_I \xrightarrow{Q} \gamma \oplus \delta_O$ for all γ , and $\delta_I \xrightarrow{P^{-1}} \gamma \xrightarrow{Q} \delta_O$ in the case of $Q \circ P^{-1}$. Therefore, the maximum expected differential probability (MEDP) is not necessarily identical.

Taka et al. studied this effect on multiple-branch-based designs and investigated the differential clustering effect on Orthros [51]. They focused on several γ , evaluated the clustering effect on each branch for each γ , and combined them. On the other hand, in general, we do not expect either $P \oplus Q$ or $Q \circ P^{-1}$ to have a stronger clustering effect because the number of terms in both sums is the same. More importantly, the clustering inside P and Q is exactly the same in both cases. We also note that if P and Q are almost the same structure, $\text{Prob}(\delta_I \xrightarrow{P \oplus Q} 0)$ is expected to be high, but so will be $\text{Prob}(\delta_I \xrightarrow{Q \circ P^{-1}} \delta_I)$.

On Key Recovery in Differential Cryptanalysis. Regarding the key recovery based on the differential attack, $P \oplus Q$ appears to be more resilient than $Q \circ P^{-1}$. More precisely, we find an inevitable difficulty in mounting an effective key-recovery attack on $P \oplus Q$.

The most common strategy for the key-recovery attack is to append key-recovery rounds to the differential distinguisher. We construct a differential distinguisher and append key-recovery rounds for attacking more rounds. The data complexity depends on the probability of the differential distinguisher, since the key-recovery rounds are deterministic under each key guess. We now consider two possible key-recovery strategies: it is added to the output or input.

Key Recovery on the Output Side. The output is $P(x) \oplus Q(x)$, where $P(x)$ and $Q(x)$ are unknown to the attacker. It is unlikely to add key recovery unless the attacker can compute at least part of (differences in) $P(x)$ or $Q(x)$. We suppose P and Q contain almost the same rounds. This implies that the key-recovery part can cover half of the total round when we attack the composition. As long as this is not the case, adding key-recovery at the output is not possible.

Key Recovery on the Input Side. Key recovery on the input side seems more natural because the attacker knows or even chooses the inputs to P and Q . We consider a differential key-recovery attack on $F := (P_2 \circ P_1) \oplus (Q_2 \circ Q_1)$, where the input differences to P_2 and Q_2 are fixed to δ_P and δ_Q , respectively. Therefore, we exploit a high differential probability $p = \text{Prob}((\delta_P, \delta_Q) \xrightarrow{P_2 \oplus Q_2} \delta_O)$ with key recovery on P_1 and Q_1 . Conventionally, the data complexity can be p^{-1} in the optimal case, but we show such a strategy does not work.

Proposition 2. *Let $F = (P_2 \circ P_1) \oplus (Q_2 \circ Q_1)$. We consider a differential key-recovery attack where the input differences of P_2 and Q_2 are fixed to δ_P and δ_Q , respectively, and the output difference is δ_O . The necessary key material from P_1 and Q_1 is guessed. Such an attack works only when*

$$\text{Prob}((\delta_P, \delta_Q) \xrightarrow{P_2 \oplus Q_2} \delta_O) \cdot \text{Prob}(\delta_P \xrightarrow{Q_1 \circ P_1^{-1}} \delta_Q) > 2^{-n}.$$

Proof. Let us count the number of input pairs X, X' to $P \oplus Q$ that produce a difference of δ_P after P_1 and δ_Q after Q_1 simultaneously.

$$\begin{aligned} T &= |\{(X, X') \mid P_1(X) \oplus P_1(X') = \delta_P \text{ and } Q_1(X) \oplus Q_1(X') = \delta_Q\}| \\ &= |\{(x, x \oplus \delta_P) \mid Q_1 \circ P_1^{-1}(x) \oplus Q_1 \circ P_1^{-1}(x \oplus \delta_P) = \delta_Q\}| \\ &= 2^n \cdot \text{Prob}(\delta_P \xrightarrow{Q_1 \circ P_1^{-1}} \delta_Q) \end{aligned}$$

Observing that the expected data complexity for the distinguisher is at least the inverse of the probability of the differential and at most T , i.e.

$$\text{Prob}((\delta_P, \delta_Q) \xrightarrow{P_2 \oplus Q_2} \delta_O)^{-1} < T$$

leads to the claimed result. □

In practice, the attacker would choose a differential trail given by $\delta_P \xrightarrow{P_2} \gamma$ and $\delta_Q \xrightarrow{Q_2} \gamma \oplus \delta_O$ and estimate the probability of the resulting distinguisher as

$$\text{Prob}((\delta_P, \delta_Q) \xrightarrow{P_2 \oplus Q_2} \delta_O) \approx \text{Prob}(\delta_P \xrightarrow{P_2} \gamma) \cdot \text{Prob}(\delta_Q \xrightarrow{Q_2} \gamma \oplus \delta_O).$$

The usual condition $\text{Prob}(\delta_P \xrightarrow{P_2} \gamma) \cdot \text{Prob}(\delta_Q \xrightarrow{Q_2} \gamma \oplus \delta_O) > 2^{-n}$ is not sufficient for an attack to be possible. If

$$\text{Prob}(\gamma \xrightarrow{P_2^{-1}} \delta_P) \cdot \text{Prob}(\delta_P \xrightarrow{Q_1 \circ P_1^{-1}} \delta_Q) \cdot \text{Prob}(\delta_Q \xrightarrow{Q_2} \gamma \oplus \delta_O) < 2^{-n},$$

there may be no pairs satisfying the differential characteristic.

Review of the Differential Key-Recovery Attack against Orthros in [44]. Proposition 2 implies that the data complexity of a differential key-recovery attack must be estimated carefully. In a nice paper at Africacrypt 2022, Li, Sun, and Wang proposed differential cryptanalysis against round-reduced Orthros. Their attacks add a 1-round key recovery to the input side of both branches. Specifically, they prepared pairs of chosen plaintexts whose differences take the form

$$(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \delta_1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \delta_2, 0, 0, 0, \delta_3, 0).$$

Branch 1 requires three nibble difference transitions in the Sbox layer: $\delta_1 \xrightarrow{S} 0x2$, $\delta_2 \xrightarrow{S} 0x2$, and $\delta_3 \xrightarrow{S} 0x8$. Similarly, branch 2 requires $\delta_1 \xrightarrow{S} 0x8$, $\delta_2 \xrightarrow{S} 0x1$, and $\delta_3 \xrightarrow{S} 0x2$. Excluding these first S-box layers, the differential probability on each branch is estimated as 2^{-64} and 2^{-48} , so the total probability is $p = 2^{-112}$. They finally estimated the data complexity as 2^{115} based on their attack framework.

Proposition 2 implies that a key-recovery attack is possible only when

$$p \cdot \text{Prob}(0x2 \xrightarrow{S \circ S^{-1}} 0x8) \cdot \text{Prob}(0x2 \xrightarrow{S \circ S^{-1}} 0x1) \cdot \text{Prob}(0x8 \xrightarrow{S \circ S^{-1}} 0x2) > 2^{-128}.$$

This probability highly depends on the key (difference) involved in $S \circ S^{-1}$. The detailed review is shown in the full version [28]. We notice that the probability

is zero for more than half of the keys in each Sbox. Therefore, it is a weak-key attack whose fraction of weak keys is $5/16 \times 7/16 \times 5/16 \approx 2^{-4.55}$.

We assume that one of the weak keys is used. Since the attacker does not know which (weak) key is used, the attacker must fully activate corresponding 12-bit inputs. Among 12-bit active inputs, we can construct about 2^{24} pairs. However, given a fixed key, the number of pairs satisfying input differences of both branches is limited. In some (weak) keys, the number is only 8 (see the full version [28] for details). Therefore, to observe differential characteristics with $p = 2^{-112}$, we need at least 2^{109} texts in addition to the 12-bit active. As a result, the attacker must use at least $2^{109+12} = 2^{121}$ chosen plaintexts to lead a valid key-recovery attack for all keys belonging to the weak keys, which is more than 2^{115} by the analysis of [44].

Remark 1. Assuming that the keys in three active S-boxes are identical in P and Q , the input differences of the two branches must be the same because $\text{Prob}(\delta_P \xrightarrow{S \circ S^{-1}} \delta_Q) = 0$ for $\delta_P \neq \delta_Q$. In other words, to lead the key-recovery attack that is valid for all keys, it is necessary to construct differential characteristics whose input differences are equal in both branches.

3.2 Linear Cryptanalysis

Linear Characteristic Equivalence. Similarly to the differential cryptanalysis, the linear trails of $P \oplus Q$ are equivalent to those of the sequential construction $Q^{-1} \circ P$, as shown in the right diagram of Fig. 1, and by the following result:

Proposition 3. *Let P, Q be two keyed permutations over \mathbb{F}_2^n , and let $F := P \oplus Q$ and $S^* := Q^{-1} \circ P$. For each linear trail with correlation c traversing F , there is a linear trail with the same correlation c traversing S^* .*

Proof. Consider any masks $\alpha, \gamma, \beta \in \mathbb{F}_2^n$, let $c = \text{cor}_P(\gamma, \beta) \text{cor}_Q(\gamma \oplus \alpha, \beta)$ be the correlation of a linear trail through F . Again, notice that

$$\text{cor}_P(\gamma, \beta) \text{cor}_Q(\gamma \oplus \alpha, \beta) = \text{cor}_P(\gamma, \beta) \text{cor}_{Q^{-1}}(\beta, \gamma \oplus \alpha).$$

Thus, c is the correlation of the linear trail $\gamma \xrightarrow{P} \beta \xrightarrow{Q^{-1}} \gamma \oplus \alpha$ traversing S^* . \square

Similar to differential cryptanalysis, while individual trails or characteristics are equivalent, it is hard to compare the resulting linear approximation correlation when adding up the trail correlation contributions:

$$\begin{aligned} \text{cor}_F(\alpha, \beta) &= \sum_{\gamma} \text{cor}_P(\gamma, \beta) \text{cor}_Q(\gamma \oplus \alpha, \beta), \\ \text{cor}_{S^*}(\alpha, \beta) &= \sum_{\gamma} \text{cor}_P(\alpha, \gamma) \text{cor}_Q(\beta, \gamma). \end{aligned}$$

It is possible that the largest correlations of the linear approximations F and S^* are not the same, due to differences in the clustering effect for both constructions.

About Sequential Applications. One peculiar aspect of Propositions 1 and 3 is that the sequential function with equivalent trails or characteristics differs between the differential ($S := Q \circ P^{-1}$) and linear ($S^* := Q^{-1} \circ P$) cases. This occurs because differential trails traversing F must coincide in the input of the two branches (the output differentials can be added) while linear trails must coincide in the output of the two branches (the input masks can be added).

However, in the case in which $\mathcal{Q} = \{P^{-1} | P \in \mathcal{P}\}$, then the compositions $S = S^*$ conform the same set of permutations, and F has the same differential and linear characteristics as $P_1 \circ P_2$, where $P_1, P_2 \in \mathcal{P}$. This is the ZIP-design strategy we employ in Sects. 4 and 5.

We note that the behavior of both constructions is not necessarily the same when it comes to trail clustering (so that the maximum differential probability or correlation may still differ). Again, all the clustering that happens within P and/or Q is equivalent in S, S^* , and $P \oplus Q$. Thus, even so our argument does not cover all possible clustering, it covers more than done in both attacks in the vast majority of cases.

On Key Recovery in Linear Cryptanalysis. Unlike with differential cryptanalysis, it is possible to mount linear key-recovery attacks on $P \oplus Q$. While it is not possible on the output side due to the irreversibility of the XOR operation, it is possible on the input side. Indeed, assume that the branches can be written as $P = P_2 \circ P_1$ and $Q = Q_2 \circ Q_1$. We are given a linear approximation of $P_2 \oplus Q_2$, and we want to perform key recovery over P_1 and Q_1 . As long as the combined size of the necessary key guesses to determine the parity of the input masks to P_2 and Q_2 is small enough, it is possible to perform key recovery on both P_1 and Q_1 simultaneously without increasing the data complexity. Linear cryptanalysis is a known plaintext attack, so the cryptanalyst does not need to control internal values in either branch and, most notably, does not need to control both branches at the same time, which is the impediment to differential key-recovery attacks). In summary, linear key-recovery attacks over the first few rounds of P and Q can be carried out in the same manner as on an iterative block cipher. Thus, assuming that differential and linear distinguishers cover the same number of rounds, linear cryptanalysis may lead to stronger attacks.

3.3 Differential-Linear Cryptanalysis

We next look at differential-linear cryptanalysis. First, we investigate how the autocorrelation of $P \oplus Q$ is related to the properties of P and Q , and we find the following straightforward result:

Proposition 4. *Let P, Q be keyed permutations over \mathbb{F}_2^n and let $F = P \oplus Q$. Let $\delta \in \mathbb{F}_2^n$ be an input difference, and let $\alpha \in \mathbb{F}_2^n$ be an output linear mask. Then*

$$\text{Aut}_F(\delta, \alpha) = \text{Aut}_P(\delta, \alpha) \cdot \text{Aut}_Q(\delta, \alpha).$$

Proof. From the definition of the autocorrelation:

$$\begin{aligned} \text{Aut}_F(\delta, \alpha) &= \text{cor}(\langle \alpha, F(x) \rangle \oplus \langle \alpha, F(x \oplus \delta) \rangle) \\ &= \text{cor}(\langle \alpha, P(x) \rangle \oplus \langle \alpha, Q(x) \rangle \oplus \langle \alpha, P(x \oplus \delta) \rangle \oplus \langle \alpha, Q(x \oplus \delta) \rangle) \\ &= \text{cor}(\langle \alpha, P(x) \rangle \oplus \langle \alpha, P(x \oplus \delta) \rangle \oplus \langle \alpha, Q(x) \rangle \oplus \langle \alpha, Q(x \oplus \delta) \rangle). \end{aligned}$$

Assuming the independence of both halves of the expression (or, alternatively, that $\text{cor}(\langle \alpha, P(x) \rangle \oplus \langle \alpha, Q(x) \rangle)$ is negligible), we deduce:

$$\begin{aligned} \text{Aut}_F(\delta, \alpha) &= \text{cor}(\langle \alpha, P(x) \rangle \oplus \langle \alpha, P(x \oplus \delta) \rangle) \cdot \text{cor}(\langle \alpha, Q(x) \rangle \oplus \langle \alpha, Q(x \oplus \delta) \rangle) \\ &= \text{Aut}_P(\delta, \alpha) \cdot \text{Aut}_Q(\delta, \alpha) \end{aligned}$$

from the piling-up-lemma [46]. □

We note two important differences between this result and the ones for differential and linear distinguishers. It describes the behavior of a whole differential-linear distinguisher without singling out an individual trail. However, the autocorrelation cannot generally be related to that on the composition of P, Q or their inverses, and relies just on the product of the autocorrelations for P and Q . This does not make a large difference for constructions in which the logarithm of the maximum autocorrelation decreases linearly with the number of rounds, but it may create a gap when this exponent decreases very quickly.

Practical Strategies for Finding DL Distinguishers. The autocorrelation of F is computed as the multiplication of the autocorrelations of P and Q having the same input difference and output mask. On the other hand, in practice a DL distinguisher is found by studying a trail perspective.

Traditionally, a cipher is separated into two parts, so that a differential trail is considered over the first part and a linear trail over the second. Let $P = P_l \circ P_d$ and $Q = Q_l \circ Q_d$, where differentials $\delta \xrightarrow{P_d} \delta_P$ and $\delta \xrightarrow{Q_d} \delta_Q$ and linear approximations on $\alpha_P \xrightarrow{P_l} \beta$ and $\alpha_Q \xrightarrow{Q_l} \beta$ are known. We consider the composition $S := P_l^{-1} \circ Q_l \circ Q_d \circ P_d^{-1}$. Then, the differential-linear distinguishers $\delta \xrightarrow{F} \beta$ and $\delta_P \xrightarrow{S} \alpha_P$ are expected to have the same autocorrelation, assuming that these trails are dominant and independent. When P_d and P_l are iterations of the round function and Q_d and Q_l are iterations of the inverse round function, $P \oplus Q$ is equivalent to the composition.

On the other hand, we can consider truncated differentials, $(\delta_P, \delta_Q) \in U_P \times U_Q$, instead of a single differential trail. As mentioned later, the behavior of the truncated differential is different in $P \oplus Q$ and the composition. Moreover, the differential-linear hull aggregates multiple intermediate masks instead of a single intermediate mask. When we switch differential trails into linear trails, we also have the so-called independence assumption issue. In particular, the strategy above has two different switches for each side of P and Q . Considering such a complicated situation, it is preferable to analyze the autocorrelation of each branch rather than optimistically trusting the relationship to the composition.

On Key Recovery in Differential-Linear Cryptanalysis. Considering the differential-linear key recovery, a similar problem arises to the one shown in the differential key recovery: it is necessary to control input differences in both branches simultaneously, which puts a limitation on the usable distinguishers.

Proposition 5. *Let $F = (P_2 \circ P_1) \oplus (Q_2 \circ Q_1)$. We consider a differential-linear key-recovery attack, where the input differences of P_2 and Q_2 are δ_P and δ_Q , respectively. The output linear mask is α . The necessary key material from P_1 and Q_1 is guessed. Such an attack works only when*

$$(\text{Aut}_{P_2}(\delta_P, \alpha) \cdot \text{Aut}_{Q_2}(\delta_Q, \alpha))^{-2} < 2^n \cdot \text{Prob}(\delta_P \xrightarrow{Q_1 \circ P_1^{-1}} \delta_Q).$$

Proof. Let δ_P and δ_Q be fixed input differences of P_2 and Q_2 , respectively. Let α be the output linear mask. Therefore, assuming the input pairs to P_2 and Q_2 already satisfy δ_P and δ_Q , the necessary number of pairs is estimated as $(\text{Aut}_{P_2}(\delta_P, \alpha) \cdot \text{Aut}_{Q_2}(\delta_Q, \alpha))^{-2}$. The number of available pairs satisfying δ_P and δ_Q at the same time is expected as $2^n \cdot \text{Prob}(\delta_P \xrightarrow{Q_1 \circ P_1^{-1}} \delta_Q)$. Therefore, when this number is less than $(\text{Aut}_{P_2}(\delta_P, \alpha) \cdot \text{Aut}_{Q_2}(\delta_Q, \alpha))^{-2}$, the attacker cannot collect enough pairs to complete the attack. \square

Review of the DL Key-Recovery Attack against Orthros in [44]. We again review the existing attack against Orthros proposed at [44]. It also presents differential-linear cryptanalysis. It uses a differential-linear distinguisher whose autocorrelation is 2^{-46} . They also estimated the data complexity to be 2^{95} chosen plaintexts.

This has the same problem as the key recovery in differential attacks, i.e., the attack is a weak-key attack and requires a higher data complexity than their estimation. The key-recovery structure is the same as the differential case. Therefore, the fraction of weak keys is $2^{-4.55}$. From 12-bit active inputs, there are weak keys, where the number of pairs satisfying input differences of both branches is only 8. Therefore, to recover any weak key, we need at least $2^{46 \times 2} / 8 \times 2^{12} = 2^{101}$ chosen plaintexts, which is more than 2^{95} by the previous estimation.

3.4 Differential-and-Linear Key-Recovery Attack

In previous sections, we have noted that attacks which require the adversary to control an input difference in both branches are difficult to turn into key-recovery attacks. On the other hand, linear attacks lend themselves well to key recovery because of the known-plaintext nature. We next introduce a hybrid key-recovery attack which uses a differential-linear distinguisher on one of the branches and a linear distinguisher on the other. On the differential-linear branch, the key recovery can be performed because the attacker can control the input difference by choosing plaintexts as in a standard differential or differential-linear attack. On the linear branch, the attacker only needs to establish the parity of the input linear mask, so it does not interfere with the key recovery on the other branch.

Let us describe this situation in more detail (see Fig. 2). P is divided into $P = P_2 \circ P_1$. Key recovery will be carried out on P_1 while a differential-linear

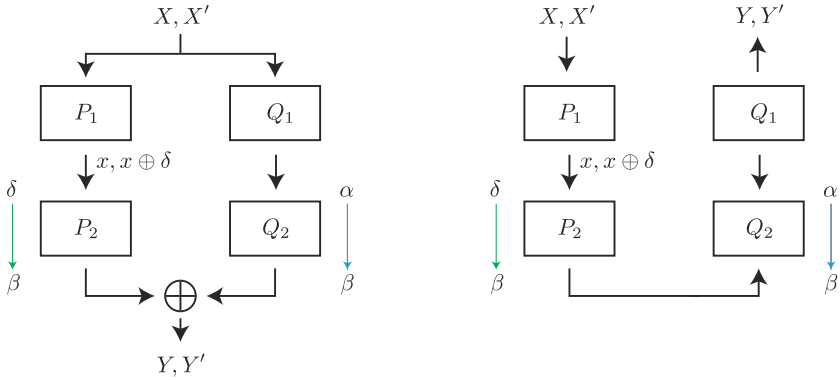


Fig. 2. The differential-and-linear key-recovery attack on $P \oplus Q$ (left) and differential-linear key-recovery attack on $Q^{-1} \circ P$.

distinguisher is considered on P_2 with input difference δ , output mask β , and autocorrelation c_1 . Q is also divided into $Q = Q_2 \circ Q_1$ where Q_1 is reserved for key recovery, and a linear approximation with masks α and β and correlation c_2 is considered for Q_2 . We note that the roles of P and Q can be exchanged.

By guessing parts of the key in P_1 and Q_1 , the attacker can compute the following parity from arbitrary X .

$$\langle \beta, F(X) \rangle \oplus \langle \beta, F(P_1^{-1}(P_1(X) \oplus \delta)) \rangle \oplus \langle \alpha, Q_1(X) \rangle \oplus \langle \alpha, Q_1(P_1^{-1}(P_1(X) \oplus \delta)) \rangle.$$

Thus, by querying enough plaintexts, the attacker can obtain the experimental correlation.

We will first determine the correlation of this function, and then we will briefly describe the key-recovery attack algorithm. For the former, we note that we can, by expanding F , rearrange the formula as follows:

$$\begin{aligned} & \langle \beta, P_2(P_1(X)) \rangle \oplus \langle \beta, P_2(P_1(X) \oplus \delta) \rangle \oplus \\ & \langle \alpha, Q_1(X) \rangle \oplus \langle \beta, Q_2(Q_1(X)) \rangle \oplus \\ & \langle \alpha, Q_1(P_1^{-1}(P_1(X) \oplus \delta)) \rangle \oplus \langle \beta, Q_2(Q_1(P_1^{-1}(P_1(X) \oplus \delta))) \rangle \end{aligned}$$

From the assumptions on the distinguishers for P_2 and Q_2 , the correlation of the first line is c_1 , and the correlations of the second and third lines are c_2 . As a result, and from the piling-up lemma, we deduce that the correlation for the whole expression is $c_1 \cdot c_2^2$, which means an attack can be mounted with data complexity $c_1^{-2}c_2^{-4}$.

We next sketch the key recovery algorithm for this attack. Using a key guess in P_1 , the attacker can use structures to construct pairs (X, X') so that $P_1(X) \oplus P_1(X') = \delta$ in the same way they would for a differential or a differential-linear attack, and at the same cost. Once these pairs (X, X') are constructed, a guess of part of the key in Q_1 enables the attacker to determine the values of $\langle \alpha, Q_1(X) \rangle$

and $\langle \alpha, Q_1(X') \rangle$. With these, and for each key guess, the attacker can compute the experimental correlations of

$$\langle \beta, F(X) \rangle \oplus \langle \beta, F(X') \rangle \oplus \langle \alpha, Q_1(X) \rangle \oplus \langle \alpha, Q_1(X') \rangle,$$

where X and X' are constructed so that $P_1(X) \oplus P_1(X') = \delta$. We verified our assumption and validity of our key-recovery attack by using ZIP-AES introduced in the next section. In detail, we discuss it in the full version [28].

Interestingly, again this kind of attack is related to a cryptanalysis on the composition of P and Q (see the right diagram of Fig. 2). Indeed, we notice that the differential-linear distinguisher on P_2 and the linear approximation of Q_2 can be combined into a differential-linear distinguisher on $Q_2^{-1} \circ P_2$. Furthermore, the whole key-recovery attack corresponds to a differential-linear key-recovery attack on $Q^{-1} \circ P$ guessing the same key material. However, we note that the autocorrelation of the differential-linear distinguisher on the composition may be larger, because the intermediate mask β is not fixed, while in the case of the attack on F the mask β has to be fixed by the attacker.

3.5 Truncated Differential Cryptanalysis

A variant of classical differential cryptanalysis is truncated differential cryptanalysis [41], in which the attacker can predict only part of the difference between pairs of texts. When considering truncated differentials cryptanalysis, the parallel construction $F := P \oplus Q$ seems to offer a security that is *hardly* comparable with any sequential construction and thus may require a dedicated analysis, which is also to be expected when compared to differential-linear attacks.

Firstly, the parallel and sequential constructions involving inverse permutations become hardly comparable as truncated differentials do not propagate backwards so that truncated differential characteristics in P generally differ from characteristics in P^{-1} .

Secondly, if we consider the sequential construction $S := Q \circ P^{-1}$ then a truncated differential attack works as such for any linear subspaces $\mathcal{U}, \mathcal{V}, \mathcal{W}$:

$$\begin{aligned} &\text{Prob}(P^{-1}(x) \oplus P^{-1}(x \oplus \alpha) \in \mathcal{V} \mid x \in \mathbb{F}_2^n, \alpha \in \mathcal{U}) = p \\ &\quad \text{Prob}(Q(x) \oplus Q(x \oplus \beta) \in \mathcal{W} \mid x \in \mathbb{F}_2^n, \beta \in \mathcal{V}) = q \\ \implies &\text{Prob}(S(x) \oplus S(x \oplus \alpha) \in \mathcal{W} \mid x \in \mathbb{F}_2^n, \alpha \in \mathcal{U}) \geq p \cdot q. \end{aligned}$$

On the other hand, Proposition 6 shows how to mount a truncated differential attack on $P \oplus Q$:

Proposition 6. *Let P, Q be two keyed permutations over \mathbb{F}_2^n , and let $F := P \oplus Q$. Let $\mathcal{U}_P, \mathcal{U}_Q, \mathcal{V}_P, \mathcal{V}_Q \subseteq \mathbb{F}_2^n$ be four non-trivial linear subspaces such that $\mathcal{U}_P \cap \mathcal{U}_Q$ is non-empty. Assume that the following truncated differentials hold with probabilities $p, q \in (0, 1]$ respectively:*

$$\begin{aligned} &\text{Prob}(P(x) \oplus P(x \oplus \alpha) \in \mathcal{V}_P \mid x \in \mathbb{F}_2^n, \alpha \in \mathcal{U}_P) = p, \\ &\text{Prob}(Q(x) \oplus Q(x \oplus \beta) \in \mathcal{V}_Q \mid x \in \mathbb{F}_2^n, \beta \in \mathcal{U}_Q) = q. \end{aligned}$$

Then:

$$\text{Prob}(\mathbf{F}(x) \oplus \mathbf{F}(x \oplus \gamma) \in \mathcal{V}_P \oplus \mathcal{V}_Q \mid x \in \mathbb{F}_2^n, \gamma \in \mathcal{U}_P \cap \mathcal{U}_Q) \geq p \cdot q.$$

We note that \oplus denotes the sum of binary vector subspaces, which may not necessarily be a direct sum. Obviously, if $\mathcal{V}_P \oplus \mathcal{V}_Q = \mathbb{F}_2^n$ is the full space, the last probability is equal to 1, making the truncated differential to be meaningless. This is not the case for \mathbf{S} .

Proof. Let $x \in \mathbb{F}_2^n$. We know that $P(x) \oplus P(x \oplus \gamma) \in \mathcal{V}_P$ with probability p over $\gamma \in \mathcal{U}_P$, and that $Q(x) \oplus Q(x \oplus \gamma) \in \mathcal{V}_Q$ with probability q over $\gamma \in \mathcal{U}_Q$. Assuming that both events are statistically independent of each other, over $\gamma \in \mathcal{U}_P \cap \mathcal{U}_Q$, the probability that they both occur at the same time is $p \cdot q$. Since \mathcal{V}_P and \mathcal{V}_Q are vector subspaces, we have

$$\mathbf{F}(x) \oplus \mathbf{F}(x \oplus \gamma) = P(x) \oplus Q(x) \oplus P(x \oplus \gamma) \oplus Q(x \oplus \gamma) \in \mathcal{V}_P \oplus \mathcal{V}_Q,$$

which concludes the proof. □

As shown in Proposition 6, an interesting constraint to find a truncated differential attack on $P \oplus Q$ is to find two linear subspaces \mathcal{U}_P and \mathcal{U}_Q such that both $\mathcal{U}_P \cap \mathcal{U}_Q$ is not empty and $\mathcal{V}_P \oplus \mathcal{V}_Q$ is not the full space \mathbb{F}_2^n . As a result, even if we find two truncated differentials, where p and q are high enough, it does not always guarantee a non-trivial truncated differential on \mathbf{F} .

Based on this, we encourage to pay particular attention when arguing the security against truncated differentials.

On Key Recovery in Truncated Differential Attacks. Extending a truncated differential distinguisher into a key recovery presents the same problems discussed in Sect. 3.1 for the analogous case of differential cryptanalysis.

Proposition 7. *Let $\mathbf{F} = (P_2 \circ P_1) \oplus (Q_2 \circ Q_1)$. We consider a key-recovery attack, where the truncated input differences of P_2 and Q_2 are in the affine subspace \mathcal{U}_P and \mathcal{U}_Q respectively, and the key involved in P_1 and Q_1 is guessed. When N pairs are needed for the distinguishing attacks based on the truncated differential to succeed, $(\mathcal{U}_P, \mathcal{U}_Q) \xrightarrow{P_2 \oplus Q_2} \mathcal{V}$, such an attack works only when*

$$2^n \cdot |\mathcal{U}_P| \cdot \text{Prob}(\mathcal{U}_P \xrightarrow{Q_1 \circ P_1^{-1}} \mathcal{U}_Q) > N.$$

As the input of P_2 , the number of pairs satisfying the truncated differential is $2^n \cdot |\mathcal{U}_P|$. To mount the key recovery, the attacker needs to find pairs that satisfy the truncated differential in the input of Q_2 simultaneously. Therefore, the number of pairs we can collect is $2^n \cdot |\mathcal{U}_P| \cdot \text{Prob}(\mathcal{U}_P \xrightarrow{Q_1 \circ P_1^{-1}} \mathcal{U}_Q)$. If this value is less than N , it is insufficient to execute the key-recovery attack.

Impossible (Truncated) Differentials. An impossible (truncated) differential [11] is a (truncated) differential that holds with probability 0. In general, the existence of impossible differentials for the composition does not imply the existence of non-trivial¹ impossible differentials for $F := P \oplus Q$.

Assuming $\text{Prob}(\delta_I \xrightarrow{Q^{-1} \circ P} \delta_O) = 0$, let \mathcal{V}_P and \mathcal{V}_Q denote a subset satisfying $\text{Prob}(\delta_I \xrightarrow{P} \mathcal{V}_P) = \text{Prob}(\delta_O \xrightarrow{Q} \mathcal{V}_Q) = 1$, and $\mathcal{V}_P \cap \mathcal{V}_Q = \emptyset$. In contrast, assuming $\text{Prob}(\delta_I \xrightarrow{F} \delta_O) = 0$, it implies $\text{Prob}(\delta_I \xrightarrow{P} \mathcal{V}_P) = \text{Prob}(\delta_I \xrightarrow{Q} \mathcal{V}_Q) = 1$, and $\mathcal{V}_P \cap (\mathcal{V}_Q \oplus \delta_O) = \emptyset$. The former can choose both input differences for P and Q arbitrarily. The latter restricts them to be the same, but we can add arbitrary δ_O to \mathcal{V}_Q . While it finally depends on case by case, probably, the former is easier to find impossible differentials than the latter.

3.6 Algebraic and Integral Attacks

The security of $P \oplus Q$ against algebraic attacks does not seem much better than the most secure between P and Q against this family of cryptanalysis. In this section, we formulate the cipher as a polynomial on the key and input bits. More precisely, we interpret the cipher as a multivariate polynomial of the n input bits of x with coefficients that are functions of the key k ,

$$F(k, x) := \bigoplus_{u \in \mathbb{F}_2^n} f_u(k) x^u .$$

The degree of F is defined as the highest degree monomial with a non-zero coefficient, that is, $\text{deg}(F) := \max_u \{ \text{wt}(u) \mid f_u(k) \neq 0 \}$, where $\text{wt}(u)$ denotes the Hamming weight of u . Since the attacker usually exploits the weakest bit, or more generally component function, the *minimum degree* is more important than the degree: $\text{minDeg}(F) := \min_{\beta} \text{deg}(\langle \beta, F(k, x) \rangle)$. However, in terms of security, we rather look at non-constant coefficients only, as any monomial that is key-independent distinguishes the function from random. Therefore, we define $\widetilde{\text{deg}}$ and $\widetilde{\text{minDeg}}$ as follows:

$$\begin{aligned} \widetilde{\text{deg}}(F) &:= \max_u \{ \text{wt}(u) \mid f_u(k) \text{ is not constant} \}, \\ \widetilde{\text{minDeg}}(F) &:= \min_{\beta} \widetilde{\text{deg}}(\langle \beta, F(k, x) \rangle). \end{aligned}$$

Proposition 8. *Let P, Q be keyed permutations over \mathbb{F}_2^n and $F := P \oplus Q$, then:*

$$\widetilde{\text{minDeg}}(F) = \min_{\beta} \max \{ \widetilde{\text{deg}}(\langle \beta, P \rangle), \widetilde{\text{deg}}(\langle \beta, Q \rangle) \} .$$

Proof. Let k_P and k_Q in \mathcal{K}_P and \mathcal{K}_Q , respectively, and let:

$$\langle \beta, P(k_P, x) \rangle := \bigoplus_{u \in \mathbb{F}_2^n} p_{\beta, u}(k_P) x^u , \quad \langle \beta, Q(k_Q, x) \rangle := \bigoplus_{u \in \mathbb{F}_2^n} q_{\beta, u}(k_Q) x^u .$$

¹ If $F(x)$ belongs to \mathcal{U} with probability 1 for each $x \in \mathcal{V}$, then $F(x) \in \mathcal{U}^c$ with probability 0, where \cdot^c is the complimentary subspace.

Given $k := k_P || k_Q \in \mathcal{K}_P \times \mathcal{K}_Q$, summing the polynomials for P and Q :

$$\langle \beta, F(k, x) \rangle = \bigoplus_{u \in \mathbb{F}_2^n} f_{\beta,u}(k) x^u = \bigoplus_{u \in \mathbb{F}_2^n} (p_{\beta,u}(k_P) + q_{\beta,u}(k_Q)) x^u.$$

So we have $f_{\beta,u} = p_{\beta,u} + q_{\beta,u}$ defined on inputs $k \in K_P \times K_Q$. Note that $f_{\beta,u}$ is constant if and only if $p_{\beta,u}$ **and** $q_{\beta,u}$ are constant. Therefore, we conclude by:

$$\begin{aligned} \widetilde{\text{deg}}(\langle \beta, F \rangle) &= \max_u \{ \text{wt}(u) : p_{\beta,u} \text{ is not constant } \text{or} \ q_{\beta,u} \text{ is not constant} \} \\ &= \max_u \{ \max \{ \text{wt}(u) : p_{\beta,u} \text{ is not constant} \}, \max \{ \text{wt}(u) : q_{\beta,u} \text{ is not constant} \} \} \\ &= \max \{ \widetilde{\text{deg}}(\langle \beta, P \rangle), \widetilde{\text{deg}}(\langle \beta, Q \rangle) \}. \end{aligned}$$

□

To show that a cipher is secure against algebraic attacks often involves arguing that the cipher reaches a high degree. Proposition 8 shows that $P \oplus Q$ can only reach a high degree if either P or Q reaches it. Thus, integral attacks could be one of the most powerful attacks on $P \oplus Q$. Indeed, if a cipher has a degree d then the cipher is vulnerable to an integral attack for any linear subspace with dimension $d + 1$. In particular, if P has degree d greater than Q , then any dimension $d + 1$ linear subspace will allow an integral attack on both P and Q simultaneously, so on $P \oplus Q$ as well.

A similar statement holds for the stronger arguments against integral attacks as given in [37]. Again, to argue for full resistance against integral cryptanalysis either P or Q already has to be fully resistant.

On Key Recovery in Integral Attacks. On the other hand, we cannot expect a strong integral key-recovery attack. Usually, the integral key-recovery attack focuses on the ciphertext side, but it is impossible in $P \oplus Q$. In [27], Ferguson et al. added one-round key recovery to the plaintext side, but it requires almost the full code book even for one-branch analysis. Besides, we must control the input of both branches in $P \oplus Q$. As discussed above, such a key recovery is difficult because the inputs of both branches are unlikely to take sets satisfying higher-order differences simultaneously after applying each key-recovery round from the common plaintext set.

The cube attack [25] is another possible key-recovery strategy. It is possible only when $f_{\beta,u}(k)$ is a very sparse polynomial. A common block cipher, where subkey is XORed every round, tends to have complicated polynomials, and the feature is used to guarantee the lower bound of the degree or the integral resistance property in [36,37]. Therefore, the cube attack is unlikely in such ciphers unless $\widetilde{\text{minDeg}}(F)$ is insufficient.

Zero-Correlation Linear. Instead of considering the zero-correlation linear [15] explicitly, we first consider the link between the zero-correlation and integral

[14, 50]. When we have zero-correlation linear on F , we also have an integral distinguisher on F . Therefore, if F is secure enough against the integral, it should also be secure against the zero-correlation linear.

It is also possible to find the zero-correlation linear directly. However, because of the analogous argument of the impossible differential, we do not suppose that the sum is weaker than the composition against the zero-correlation linear.

3.7 Second-Order Differential Cryptanalysis

We look at attacks exploiting independent differential properties of P and Q . Interestingly, this distinguisher on $P \oplus Q$ is linked to the Boomerang distinguisher [52] on $Q^{-1} \circ P$, as depicted in Fig. 3.

Assuming we have two independent differential transitions that are $\text{Prob}(\delta_P \xrightarrow{P} \delta'_P) = p$ and $\text{Prob}(\delta_Q \xrightarrow{Q} \delta'_Q) = q$, then for some x :

$$\begin{cases} P(x) \oplus P(x \oplus \delta_P) = \delta'_P, & P(x \oplus \delta_Q) \oplus P(x \oplus \delta_Q \oplus \delta_P) = \delta'_P, \\ Q(x) \oplus Q(x \oplus \delta_Q) = \delta'_Q, & Q(x \oplus \delta_P) \oplus Q(x \oplus \delta_P \oplus \delta_Q) = \delta'_Q \end{cases} \\ \implies F(x) \oplus F(x \oplus \delta_P) \oplus F(x \oplus \delta_Q) \oplus F(x \oplus \delta_Q \oplus \delta_P) = 0.$$

With the usual independent assumptions, this happens with probability $(p \cdot q)^2$ for a random x when $F = P \oplus Q$. Therefore, such a second-order differential requires about $4(p \cdot q)^{-2}$ queries to F .

We review the same differential transitions on $S = Q^{-1} \circ P$ and perform the following boomerang attack. For some x ,

$$\begin{cases} P(x) \oplus P(x \oplus \delta_P) = \delta'_P, \\ P(S^{-1}(S(x) \oplus \delta_Q)) \oplus P(S^{-1}(S(x) \oplus \delta_Q) \oplus \delta_P) = \delta'_P, \\ Q(S(x)) \oplus Q(S(x) \oplus \delta_Q) = \delta'_Q, \\ Q(S(x \oplus \delta_P)) \oplus Q(S(x \oplus \delta_P) \oplus \delta_Q) = \delta'_Q, \end{cases} \\ \implies S^{-1}(S(x) \oplus \delta_Q) \oplus S^{-1}(S(x \oplus \delta_P) \oplus \delta_Q) = \delta_P.$$

This well-known Boomerang holds with a probability of $(p \cdot q)^2$ with some independent assumptions. It requires about $4(p \cdot p^*)^{-2}$ queries to S and S^{-1} .

Note that the relationship above ignores some independent issues when switching differential trails. For example, although $\delta_P = \delta_Q$ is a meaningful parameter for the Boomerang distinguisher on $Q^{-1} \circ P$, it is meaningless on $P \oplus Q$. Due to different independent issues, the resulting Boomerang probability on S and the 2nd order differential probability on $P \oplus Q$ differ. On the other hand, when p and q are reasonably high, that is a natural setting in real cryptanalysis, we would observe a similar feature in both cases.

On Key Recovery in 2nd-Order Differential Attacks. When considering key recovery, we observe a similar difficulty to that of differential key recovery. Let $P = P_2 \circ P_1$ and $Q = Q_2 \circ Q_1$. Assuming that there is a non-negligible

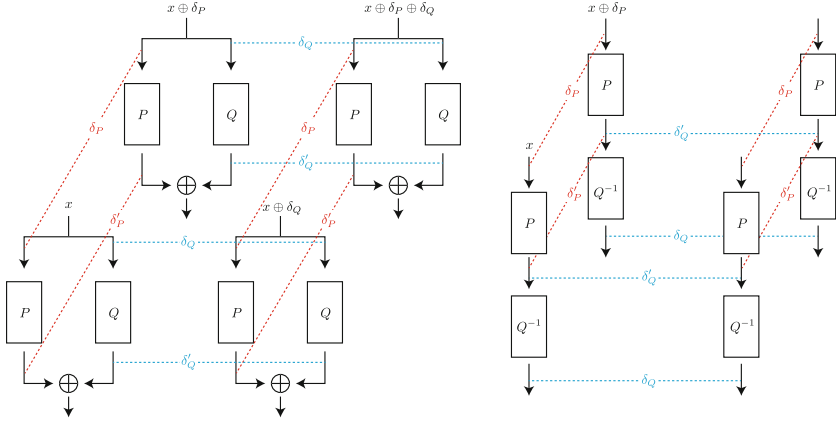


Fig. 3. 2nd-order differential on $P \oplus Q$ (left) and Boomerang on $P^{-1} \circ Q$ (right).

2nd-order differential distinguisher on $P_2 \oplus Q_2$. We apply the key recovery to P_1 and Q_1 . Let (x_1, y_1) , (x_2, y_2) , (x_3, y_3) , and (x_4, y_4) be the input of (P_2, Q_2) . Then, a quartet satisfying $x_1 \oplus x_2 = x_3 \oplus x_4 = \delta_P$ and $y_1 \oplus y_3 = y_2 \oplus y_4 = \delta_Q$ is constructed by $y_1 = Q_1 \circ P_1^{-1}(x_1)$, $x_2 = x_1 \oplus \delta_P$, $y_2 = Q_1 \circ P_1^{-1}(x_2)$, $y_3 = y_1 \oplus \delta_Q$, $x_3 = P_1 \circ Q_1^{-1}(y_3)$, $x_4 = x_3 \oplus \delta_P$, and

$$y_4 = Q_1 \circ P_1^{-1}(x_4) = y_2 \oplus \delta_Q .$$

In general, $Q_1 \circ P_1^{-1}(x_4) = y_2 \oplus \delta_Q$ does not hold with a probability of 1.

3.8 Meet-in-the-Middle Attacks

The meet-in-the-middle (MitM) attack [23] is another of the typical cryptanalysis of keyed symmetric primitives. In a traditional meet-in-the-middle attack, the adversary obtains a plaintext-ciphertext pair, and aims to extract the key faster than through an exhaustive search. The attacker guesses part of the key on the plaintext side and part of the key on the ciphertext side, and constructs two tables: one consists of all possible partial encryptions of the plaintext and the other of all possible partial decryptions of the ciphertext. When a collision between both tables is found, a candidate for both key guesses is obtained.

When applying this approach to the $P \oplus Q$ construction, we note that no information about the outputs of both branches can be obtained directly from the ciphertext. Thus, any MitM attack would require guessing part of one of the branches. However, by xoring the known ciphertext, this is equivalent to guessing part of an internal state of $Q^{-1} \circ P$, which is an ineffective guessing strategy in a MitM attack.

The DS-MitM attack [22] is an extension of the Meet-in-the-Middle attack and consists of the distinguisher and key recovery. When the distinguisher covers the initial few rounds in both branches, the key recovery requires the inverse

query but there is no such query in the PRF. When the distinguisher covers the last few rounds in both branches, it involves the output of the PRF. Therefore, the parameter size of the distinguisher significantly increases. Consequently, using the distinguisher in either the inside of P or that of Q is promising, but then, such an attack is very similar to the attack against the composition, $Q^{-1} \circ P$ too.

3.9 Summary and Other Attacks

In this section, we analyzed differential, linear, differential-linear, differential-and-linear key recovery, (impossible) truncated differential, algebraic and integral, zero-correlation linear, the 2nd-order differential, and the MitM attacks. Some of them are strongly linked to the cryptanalysis against the composition.

When we mount a key recovery, where we need to control differences in two branches simultaneously, it is more difficult than the corresponding analysis against the composition. Notably, linear key recovery and differential-and-linear key recovery are promising attack strategies against the sum structure because they are friendly to key recovery, but they are strongly linked to linear key recovery and differential-linear key recovery against the composition.

Other well-known attacks exist. For example, Boomerang [52] or Yo-Yo [10] attacks require adaptive chosen-plaintext-ciphertext attacks. However, the sum structure does not provide the decryption query, so applying these attacks is non-trivial. Note that an amplified Boomerang [40] and Rectangle [12] attacks are a chosen-plaintext variant of the Boomerang attack. However, it contains a probability of 2^{-2n} because the intermediate state size is $2n$ bits. Thus, it is unlikely that those attacks are applicable.

4 The ZIP Structure: Designing PRF in Light Work

Respecting the discussions in Sect. 3, we introduce the *ZIP structure*, which is defined as follows:

Definition 2 (ZIP structure). *Let $E = E_1 \circ E_0$ be a secure iterative block cipher. We define the ZIP construction of E as the following family of functions $E_0 \oplus E_1^{-1} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$. We suppose E_0 and E_1 contain almost the same rounds.*

The ZIP structure has three advantages:

- We can inherit many cryptanalysis results against E .
- Since the resulting primitive is a pseudo-random function, it derives beyond-birthday security in some modes of operation.
- On performance, the latency is about half of the original block cipher.

Of course, the discussion in Sect. 3 never shows that the ZIP structure has the same security as the original block cipher against all attack strategies. In particular, algebraic (integral), differential-linear, and truncated differential have to be carefully analyzed, but it is not as hard work as designing it from scratch.

In a practical application, the ZIP structure can achieve beyond-birthday security in some modes of operation while keeping the throughput in the case we use the original block cipher. It is useful in a wide situation. Moreover, its half latency is promising in several practical applications such as memory encryption or communication over the 5G and the beyond 5G as discussed in [1].

In this section, we focus on the ZIP-AES as an example.

4.1 ZIP-AES: A Concrete Instantiation via AES-128

AES-128. The Advanced Encryption Standard [21] is a SPN scheme designed by Daemen and Rijmen, and based on the Wide-Trail design strategy [19, 20]. Focusing on AES-128, the key size is of 128 bits, and the number of rounds is 10. Each AES-round $R_{\text{AES}} : \mathbb{F}_{2^8}^{4 \times 4} \rightarrow \mathbb{F}_{2^8}^{4 \times 4}$ applies three operations besides the key-addition to the state x , that is, $x \mapsto R_{\text{AES}}(x) := MC \circ SR \circ SB(x)$. An additional AddRoundKey operation is applied at the input of the first round, and the last MixColumns operation is omitted (we denote a round without MC as \hat{R}_{AES}). We refer to [21] for the details of the key-schedule.

The ZIP-AES PRF. We define the ZIP-AES as

$$\forall x \in \mathbb{F}_{2^8}^{4 \times 4} : \quad \text{ZIP-AES}_5(x) := \text{AES}_5(x) \oplus \text{AES}_5^{-1}(x),$$

where AES_5 denotes 5 encryption rounds of AES-128

$$\text{AES}_5(\cdot) = AK \circ \underbrace{MC \circ SR \circ SB}_{R_{\text{AES}}} \circ \dots \circ AK \circ \underbrace{MC \circ SR \circ SB}_{R_{\text{AES}}} \circ AK(\cdot)$$

including the final MC in the last round as well, and where AES_5^{-1} denotes 5 decryption rounds of AES-128

$$\text{AES}_5^{-1}(\cdot) = AK^{-1} \circ \underbrace{(MC \circ SR \circ SB)^{-1}}_{R_{\text{AES}}^{-1}} \circ AK^{-1} \circ \dots \circ \underbrace{(MC \circ SR \circ SB)^{-1}}_{R_{\text{AES}}^{-1}} \circ AK^{-1}(\cdot)$$

where $(MC \circ SR \circ SB)^{-1}(\cdot) := SB^{-1} \circ SR^{-1} \circ MC^{-1}(\cdot)$, and including the initial MC^{-1} in the first round as well.

Regarding the sub-keys, let $k_0 = \mathbf{k}, k_1, k_2, \dots, k_{10} \in \mathbb{F}_{2^8}^{4 \times 4}$ be the sub-keys generated by the key-schedule of AES-128, where $\mathbf{k} \in \mathbb{F}_{2^8}^{4 \times 4}$ is the whitening key.

- AES_5 is instantiated with $k_0, k_1, k_2, k_3, k_4, k_5$;
- AES_5^{-1} is instantiated with $k_6, k_7, k_8, k_9, k_{10}, 0^{128}$.

We claim that ZIP-AES is a 128-bit secure pseudo-random function.

Design Rationale and Modified Versions of ZIP-AES. Before going on, we briefly discuss some technical choices regarding ZIP-AES, with particular attention both at the MixColumns operation at the end of AES_5 , and at the inverse MixColumns operation at the beginning of AES_5^{-1} . As we pointed out, the final *MC* operation is omitted in AES. However, we decided to keep it for ZIP-AES.

This choice is necessary considering our motivation: ZIP-AES shares many cryptanalysis results to the original AES. As mentioned in Sect. 3, $P \oplus Q$ and $Q \circ P^{-1}$ shares the same differential characteristic, and $P \oplus Q$ and $Q^{-1} \circ P$ shares the same linear trail. If there is no inverse MixColumns in the beginning of AES_5^{-1} , the inverse MixColumns is missing between Q and P^{-1} in $Q \circ P^{-1}$. Similarly, if there is no MixColumns in the last of AES_5 , the MixColumns is missing between Q^{-1} and P in $Q^{-1} \circ P$. In other words, such a construction corresponds to the variant of AES, where the MixColumns is omitted in the 5th round, which is clearly more insecure than the AES.

In practice, in order to prove this fact, we consider these variants of ZIP-AES, in which the final *MC* operation for AES and/or the initial MC^{-1} operation for AES^{-1} are omitted in the full version [28]. In there, we show that these modified versions are (much) weaker against attacks such as truncated differentials and mixture differentials with respect to the ZIP-AES defined here.

4.2 Security Analysis of ZIP-AES

In this section, we present our security analysis of ZIP-AES. Our results show that the strongest attack against it is the integral attack, which can distinguish up to 4+4 rounds (namely, ZIP-AES_{4,4}) from a PRF. All other attacks (including classical linear and differential attacks, truncated differentials, mixture differentials, and so on) can only cover a smaller number of rounds. Moreover, in the full version [28], we also show that the attacks against AES-PRF_{1,r} and AES-PRF_{2,r} for any $r \geq 1$ proposed in [48] work against ZIP-AES_{1,r} and ZIP-AES_{2,r} as well.

Unbalanced Variants. For the follow-up, we introduce “reduced-round variants” of ZIP-AES defined as $\text{ZIP-AES}_{r_0, r_1}(x) := \text{AES}_{r_0}(x) \oplus \text{AES}_{r_1}^{-1}(x)$. We encourage to analyze its security with particular attention to the case $r_0 = r_1 \geq 2$, in order to better evaluate ZIP-AES’s resistance against attacks.

Differential and Linear Attacks. In the case of differential cryptanalysis, we have seen in Proposition 1 that, given two independent keyed permutations P, Q , then for each differential characteristic (trail) with probability p traversing $P \oplus Q$, there is a differential characteristic with the same probability p traversing $Q \circ P^{-1}$. Due to the wide-trail design strategy, it is well known that any differential characteristic over 4-round AES has a probability of at least 2^{-150} . This means that ZIP-AES_{2,2} does *not* admit any differential characteristic with probability lower than 2^{-150} . Based on this, we claim that ZIP-AES_{5,5} is secure against differential distinguishers and key-recovery attacks.

We have an analogous argument for linear cryptanalysis, differential-and-linear key recovery, and the 2nd order differential attacks.

Differential-Linear Attacks. The differential-linear distinguisher (autocorrelation) is estimated as the product of each branch’s autocorrelation. In [35], the authors evaluated the autocorrelation of the AES. They are 1 , $2^{-7.66}$, $2^{-31.66}$, and $2^{-55.66}$, for 2, 3, 4, and 5 rounds, respectively. Although there are no references in the AES inverse, we expect the autocorrelations to be similar, considering the well-aligned structure of the AES. Then, the autocorrelation of ZIP-AES_{5,5} is expected as $2^{-55.66 \times 2}$, which is unlikely to be observed with 2^{128} , full code-book, queries. In practice, the input difference and output mask must be the same in both branches. Such a restriction does not allow us to use the optimal autocorrelation for both branches simultaneously. We verified this observation by using ZIP-AES_{3,3}. When we used the 3-round differential-linear distinguisher shown in [35] in the left branch, we could not observe a significant autocorrelation in the right branch. Therefore, we expect that the autocorrelation is worse than the squared value of the best autocorrelation of each branch. In detail, see the full version [28].

Integral Attacks. Following [32], we introduce the following subspaces of $\mathbb{F}_{2^8}^{4 \times 4}$: the diagonal subspace \mathcal{D}_i , in which the i -th diagonal for $i \in \{0, 1, 2, 3\}$ is active and all the others are constant; the column subspace $\mathcal{C}_i := SR(\mathcal{D}_i)$, in which the i -th column for $i \in \{0, 1, 2, 3\}$ is active and all the others are constant; the anti-diagonal subspace $\mathcal{ID}_i := SR(\mathcal{C}_i)$, in which the i -th anti/inverse diagonal for $i \in \{0, 1, 2, 3\}$ is active and all the others are constant; the mixed subspace $\mathcal{M}_i := MC(\mathcal{ID}_i)$.

As it is well known [27, 29, 42], the following integral attacks hold

$$\bigoplus_{x \in \mathcal{D}_i \oplus \alpha} \text{AES}_4(x) = \bigoplus_{x \in \mathcal{M}_i \oplus \beta} \text{AES}_4^{-1}(x) = 0$$

for each $i \in \{0, 1, 2, 3\}$ and for any $\alpha, \beta \in \mathbb{F}_{2^8}^{4 \times 4}$. It follows that for each $i, j \in \{0, 1, 2, 3\}$:

$$\bigoplus_{x \in (\mathcal{D}_i \oplus \mathcal{M}_j) \oplus \alpha} \text{ZIP-AES}_{4,4}(x) = 0$$

for each $\alpha \in \mathbb{F}_{2^8}^{4 \times 4}$, where $\dim(\mathcal{D}_i \oplus \mathcal{M}_j) = 8$ – the dimension is considered at byte level. Therefore, we have the integral distinguisher by using 2^{64} chosen plaintexts.

Since no other integral distinguisher is known for 5 or more rounds of AES, and since appending a key recovery to the plaintext side is not easy (see Sect. 3 for more details), we claim that ZIP-AES_{5,5} is secure against integral attacks.

Truncated Differential and Subspace Trail Attacks. With respect to the previous attacks and distinguishers, truncated differential requires a more dedicated analysis, since it is not possible to reduce the security of $F := P \oplus Q$ to the one of any sequential construction (see Sect. 3.5 for more details).

We first re-call some results regarding the subspace trails presented in [32]. Given $\mathcal{D}_I := \bigoplus_{i \in I} \mathcal{D}_i$, $\mathcal{C}_I := \bigoplus_{i \in I} \mathcal{C}_i$, $\mathcal{ID}_I := \bigoplus_{i \in I} \mathcal{ID}_i$, $\mathcal{M}_I := \bigoplus_{i \in I} \mathcal{M}_i$ for each $I \subseteq \{0, 1, 2, 3\}$, we have that

Table 1. Practical tests on ZIP-AES over $\mathbb{F}_{2^8}^{4 \times 4}$. In the table, we assume $|I| = |I'| = 3$ and $|J| = 2$ (P \equiv Practical – Prob. \equiv Probability).

# Rounds	Input Subspace	Output Subspace	ZIP-AES P-Prob.	PRF Prob.
1 + 1	\mathcal{C}_i	$\mathcal{D}_i \cap \mathcal{M}_i$	1	2^{-64}
2 + 2	\mathcal{C}_i	\mathcal{C}_I	$2^{-32} + 2^{-52.8}$	2^{-32}
2 + 2	$\mathcal{M}_J \cap \mathcal{D}_I$	$\mathcal{C}_{I'}$	$2^{-32} + 2^{-53.7}$	2^{-32}

- $\mathcal{D}_{i,i+2} = \mathcal{I}\mathcal{D}_{i,i+2}$ for each $i \in \{0, 1, 2, 3\}$,
- for each $I, J \subseteq \{0, 1, 2, 3\}$: $\dim(\mathcal{C}_I \cap \mathcal{M}_J) = \dim(\mathcal{C}_I \cap \mathcal{D}_J) = |I| \cdot |J|$,
- for each $I, J \subseteq \{0, 1, 2, 3\}$ with $|I| + |J| \leq 4$: $\mathcal{D}_I \cap \mathcal{M}_J = \mathcal{I}\mathcal{D}_I \cap \mathcal{M}_J = \emptyset$,

where $|I|$ and $|J|$ represent the cardinality of I and J respectively.

Let $\text{AES}_r(\cdot)$ be r rounds of AES. For each $x \in \mathbb{F}_{2^8}^{4 \times 4}$, and for each $I, J \subseteq \{0, 1, 2, 3\}$, the following truncated differentials hold:

$$\begin{aligned} \text{Prob}(\text{AES}_1(x) \oplus \text{AES}_1(x \oplus \delta) \in \mathcal{C}_I \mid \delta \in \mathcal{D}_I) &= 1, \\ \text{Prob}(\text{AES}_1(x) \oplus \text{AES}_1(x \oplus \delta) \in \mathcal{M}_I \mid \delta \in \mathcal{C}_I) &= 1, \\ \text{Prob}(\text{AES}_2(x) \oplus \text{AES}_2(x \oplus \delta) \in \mathcal{M}_I \mid \delta \in \mathcal{D}_I) &= 1, \\ \text{Prob}(\text{AES}_3(x) \oplus \text{AES}_3(x \oplus \delta) \in \mathcal{M}_J \mid \delta \in \mathcal{D}_I) &= 2^{8 \cdot |I| \cdot (|J| - 4)}. \end{aligned}$$

We refer to [6, 31] for truncated differentials up to 6-round AES.

Truncated Differentials for ZIP-AES_{1,1}. Since $\text{Prob}(\text{AES}_1(x) \oplus \text{AES}_1(x \oplus \delta) \in \mathcal{M}_i \mid \delta \in \mathcal{C}_i) = \text{Prob}(\text{AES}_1^{-1}(x) \oplus \text{AES}_1^{-1}(x \oplus \delta) \in \mathcal{D}_i \mid \delta \in \mathcal{C}_i) = 1$, the following truncated differentials on ZIP-AES_{1,1} holds:

$$\text{Prob}(\text{ZIP-AES}_{1,1}(x) \oplus \text{ZIP-AES}_{1,1}(x \oplus \delta) \in \mathcal{D}_i \oplus \mathcal{M}_i \mid \delta \in \mathcal{C}_i) = 1.$$

For comparison, note that $\text{Prob}(\Pi(x) \oplus \Pi(x \oplus \delta) \in \mathcal{D}_i \oplus \mathcal{M}_i \mid \delta \in \mathcal{C}_i) = 2^{-64}$ for a PRF Π over $\mathbb{F}_{2^8}^{4 \times 4}$.

Truncated Differentials for ZIP-AES_{2,2}: a Negative Result. Due to the existence of probability-1 truncated differentials for both 2-round AES and 2-round AES^{-1} , corresponding to $R^2(\mathcal{D}_I \oplus \alpha) = \mathcal{M}_I \oplus \beta$ and $R^{-2}(\mathcal{M}_J \oplus \alpha') = \mathcal{D}_J \oplus \beta'$, it could seem natural to combine them in order to set up a truncated differential for ZIP-AES_{2,2}, defined via an initial subspace $\mathcal{D}_I \cap \mathcal{M}_J$ and a final subspace $\mathcal{M}_I \oplus \mathcal{D}_J$. However, a problem arises, since

- $\mathcal{D}_I \cap \mathcal{M}_J$ contains only the zero-element for each I, J with $|I| + |J| \leq 4$, and
- $\mathcal{D}_J \oplus \mathcal{M}_I$ is the full space $\mathbb{F}_{2^8}^{4 \times 4}$ for each I, J with $|I| + |J| \geq 4$,

due to the results listed before. For this reason, it seems impossible to set up a probability-1 truncated differential for ZIP-AES_{2,2} via this strategy.

Truncated Differentials for ZIP-AES_{r,r} with $r \geq 2$: Practical Results. At the same time, probabilistic truncated differential distinguishers for ZIP-AES_{r,r} with

$r \geq 2$ exist. Our practical results for ZIP-AES is summarized in Table 1.² As it is possible to observe, for all the considered cases, the probability that a truncated differential distinguisher holds for ZIP-AES $_{r,r}$ with $r \in \{2, 3\}$ is only slightly higher than the corresponding probability for a generic PRF.

Conclusion. Based on our practical tests, we conjecture that if a bias between the probability for ZIP-AES $_{r,r}$ for $r \geq 4$ and a generic PRF exists, it would be too small for being useful in practice. Together with the fact that extending a distinguisher that ends with \mathcal{C}_I with $|I| \geq 2$ by 1 round is **not** possible, we claim that ZIP-AES $_{5,5}$ is secure against truncated differential distinguishers.

Mixture Differential Attacks (and More). A powerful attack on round-reduced AES is the mixture differential cryptanalysis [30]. Given two plaintexts p_0, p_1 in the same column space $\mathcal{C}_I \oplus \gamma \subseteq \mathbb{F}_{2^8}^{4 \times 4}$, let $p'_0, p'_1 \in \mathcal{C}_I \oplus \gamma$ be two new texts obtained by carefully swapping the *generating variables* of p_0, p_1 . Independently of the values of the round-keys, the difference between p_0 and p_1 after 2-round AES is equal to the corresponding difference of p'_0, p'_1 , that is,

$$\text{AES}_2(p_0) \oplus \text{AES}_2(p_1) = \text{AES}_2(p'_0) \oplus \text{AES}_2(p'_1). \tag{1}$$

This is also known as the *integral mixture distinguisher* [33]. Moreover, p_0 and p_1 belong to the same coset of a mixed space \mathcal{M}_J after 4-round AES if and only if p_0 and p_1 satisfy the same property, that is, $\forall J \subseteq \{0, 1, 2, 3\}$:

$$\text{AES}_4(p_0) \oplus \text{AES}_4(p_1) \in \mathcal{M}_J \iff \text{AES}_4(p'_0) \oplus \text{AES}_4(p'_1) \in \mathcal{M}_J.$$

Similar distinguishers hold in the backward direction. (A variant of such distinguisher – the exchange attack [6] – is discussed in the full version [28]).

(Deterministic) Mixture Integral Distinguishers for ZIP-AES $_{2,2}$: a Negative Result. At the current state, it does **not** seem possible to set up an integral mixture distinguisher for ZIP-AES $_{2,2}$, that is,

$$\text{ZIP-AES}_{2,2}(p_0) \oplus \text{ZIP-AES}_{2,2}(p_1) \neq \text{ZIP-AES}_{2,2}(p'_0) \oplus \text{ZIP-AES}_{2,2}(p'_1)$$

in general, where $p_0, p_1, p'_0, p'_1 \in \mathcal{C}_I \oplus \gamma$ for $I \subseteq \{0, 1, 2, 3\}$, and where p'_0 and p'_1 are constructed by carefully swapping the generating variables of p_0, p_1 in the same way described in [30]. The problem arises from the fact that generating variables of p_0, p_1 and the ones of $MC^{-1}(p_0), MC^{-1}(p_1)$ are different.

(Probabilistic) Mixture Differential Distinguishers for ZIP-AES $_{2,2}$. Having said that, it is possible to set up a *probabilistic* mixture differential distinguisher for ZIP-AES $_{2,2}$ by exploiting the following result.

² The truncated differentials are not affected by the details (as the degree) of the S-Box. Hence, we also provide practical results for small-scale ZIP-AES (that is, AES over $\mathbb{F}_{2^4}^{4 \times 4}$ as presented in [18]) in the full version [28].

Table 2. Performance comparison on the counter mode.

	cycle-per-byte						counter
	16B	32B	256B	2KB	16KB	128KB	
AES	3.56	1.84	0.51	0.36	0.34	0.34	integer
AES-PRF	3.63	1.94	0.55	0.39	0.37	0.37	integer
ZIP-AES	2.96	1.58	0.53	0.41	0.39	0.39	integer
AES	3.53	1.81	0.47	0.35	0.34	0.33	gray code
AES-PRF	3.57	1.88	0.51	0.36	0.34	0.34	gray code
ZIP-AES	2.90	1.61	0.47	0.34	0.33	0.33	gray code

Lemma 1. *Let $p_0, p_1 \in \mathcal{C}_i \oplus \alpha$. Let $p'_0, p'_1 \in \mathcal{C}_i \oplus \alpha$ be defined as the mixture couples generated by p_0 and p_1 such that Eq. (1) holds. For any $I \subseteq \{0, 1, 2, 3\}$ with $|I| = 3$:*

$$\text{Prob}\left(\text{ZIP-AES}_{2,2}(p_0) \oplus \text{ZIP-AES}_{2,2}(p_1) \oplus \text{ZIP-AES}_{2,2}(p'_0) \oplus \text{ZIP-AES}_{2,2}(p'_1) \in \mathcal{D}_I\right) \geq 2^{-16}.$$

For comparison, $\text{Prob}(\Pi(p_0) \oplus \Pi(p_1) \oplus \Pi(p'_0) \oplus \Pi(p'_1) \in \mathcal{D}_I) = 2^{-32}$ for a PRF Π over $\mathbb{F}_{2^8}^{4 \times 4}$.

See the full version [28] for the proof of Lemma 1.

At the current state, it does not seem possible to extend the previous distinguisher for more rounds of ZIP-AES. For this reason, we conjecture that ZIP-AES_{5,5} is secure against such an attack.

4.3 Performance Evaluation

We implemented the counter mode of ZIP-AES to measure the performance. For the comparison, we also implemented the counter modes of AES-128 and AES-PRF-128 [48]. All measurements were taken on a single core of Intel Core i7-1185G7 (Tiger Lake) with Turbo Boost and Hyperthreading disabled, and averaged over $100000 \times \frac{4096}{\text{byte}}$ repetitions, where *byte* denotes the processing data size in bytes. All subkeys are pre-computed, and the process is measured when the IV and plaintext are given in a byte array. The counter mode uses the 64-bit IV and 64-bit counter for the top and bottom halves of the input, respectively.

Table 2 (top 3 rows) summarizes the cycle-per-byte of each cipher for each size of processing message. As expected, ZIP-AES performs better than AES and AES-PRF for small data because the latency for one block processing is lower. On the other hand, when we encrypt more than 2KB, ZIP-AES performs worse than AES and AES-PRF. The reason is that AESENC performs $AK^{-1} \circ MC^{-1} \circ SR^{-1} \circ SB^{-1}$ and is not the straightforward AES inverse round function. AES-NI consists of six instructions:

- AESENC performs $AK \circ MC \circ SR \circ SB$.

- AESENCLAST performs $AK \circ SR \circ SB$.
- AESDEC performs $AK^{-1} \circ MC^{-1} \circ SR^{-1} \circ SB^{-1}$.
- AESDECLAST performs $AK^{-1} \circ SR^{-1} \circ SB^{-1}$.
- AESIMC performs MC^{-1} . It is prepared to prepare subkeys for decryption.
- AESKEYGENASSIST assists to create round keys.

To perform AES_5^{-1} , we first use AESIMC and then use AESDEC. Unfortunately, AESIMC of the AES-NI is worse than the other main instructions. For example, on Tiger Lake CPU, the latency and throughput of the main four instructions are 3 and 0.5, respectively, but the latency and throughput of AESIMC are 6 and 1, respectively. The overhead by AESIMC is not negligible for long data.

To solve the overhead issue, we replace an integer counter with a gray code counter. In the gray code, the counting up is implemented by one XOR with a counter-dependent value. Notably, the counting up and MC^{-1} (and the whitening key XORing) is commutative. Given the IV, we first prepare the counter for AES_5 and prepare the counter for AES_5^{-1} by applying MC^{-1} . Then, we perform each counting up independently by one XOR. Then, we can avoid AESIMC for every block. Modern CPUs can perform XOR instructions in 3 ports, and the XOR instruction is executed with the AES instruction in parallel. Therefore, the overhead can be negligible. Table 2 (bottom 3 rows) summarizes each cycle-per-byte, where the counter is implemented by the gray code. We notice that the overhead of ZIP-AES for the long data can be resolved, and the performance is competitive with the case of AES and AES-PRF.

5 Future Work: Other ZIP Ciphers and Modes

In addition to ZIP-AES, one can consider several ZIP ciphers. Although we did not discuss it in this paper, we are interested in ZIP-AES-256; does it successfully derive the 256-bit secure PRF? Another interesting instance is the ZIP cipher using the 64-bit block cipher, e.g., ZIP-GIFT, instantiated by GIFT-64 [5].

GIFT-64 consists of 28 rounds. So, ZIP-GIFT consists of 14-round GIFT-64 and 14-round inverse GIFT-64. Unlike ZIP-AES, we do not provide a detailed analysis, and it is left as an open problem. As a reference, the following is a related analysis for GIFT-64. For the integral attack, in [37], the integral resistance property is guaranteed in 12-round GIFT-64, and the best integral distinguisher is up to 10 rounds. Therefore, ZIP-GIFT also guarantees integral resistance property. In [53], the autocorrelation is evaluated in GIFT-64, where the squared autocorrelation is $2^{-57.22}$ in 12 rounds. Therefore, the autocorrelation of ZIP-GIFT would be low enough.

Besides looking into more ZIP ciphers, it is promising to apply the general practical cryptanalysis to other structures. In particular, the feed-forward EDMD structure used in [47, 48] to construct AES-PRF is a natural candidate to check which attack vectors link to AES and which do not. Another example is the generalization of the sum of two permutations, i.e., a sum of several permutations. There is already a concrete instance that has been designed, i.e., Gleeok [1] named after the multiple head dragon.

Finally, it is worth investigating if the new differential-and-linear attack that we introduced and linked to a differential-linear attack on the composition, is applicable to Orthros.

Acknowledgements. The authors would like to thank anonymous reviewers for their useful feedback in preparing the final version of this paper and Jan Vorloeper for performing the practical tests regarding the truncated differential on ZIP-AES.

This work has been funded in parts by the German Research foundation (DFG) within the framework of the Excellence Strategy of the Federal Government and the States - EXC 2092 CaSa - 39078197 and by the ERC project 101097056 (SYMTRUST).

References

1. Anand, R., Banik, S., Caforio, A., Ishikawa, T., Isobe, T., Liu, F., Minematsu, K., Rahman, M., Sakamoto, K.: Gleeok: A family of low-latency prfs and its applications to authenticated encryption. *IACR TCHES* **2024**(2), 545–587 (2024)
2. Avanzi, R.: The QARMA block cipher family. almost MDS matrices over rings with zero divisors, nearly symmetric even-mansour constructions with non-involutory central rounds, and search heuristics for low-latency s-boxes. *IACR ToSC* **2017**(1), 4–44 (2017)
3. Avanzi, R., Banik, S., Dunkelman, O., Eichlseder, M., Ghosh, S., Nageler, M., Regazzoni, F.: The qarmav2 family of tweakable block ciphers. *IACR ToSC* **2023**(3), 25–73 (2023)
4. Banik, S., Isobe, T., Liu, F., Minematsu, K., Sakamoto, K.: Orthros: A low-latency PRF. *IACR ToSC* **2021**(1), 37–77 (2021)
5. Banik, S., Pandey, S.K., Peyrin, T., Sasaki, Y., Sim, S.M., Todo, Y.: GIFT: A Small Present - Towards Reaching the Limit of Lightweight Encryption. In: Fischer, W., Homma, N. (eds.) *CHES 2017*. LNCS, vol. 10529, pp. 321–345. Springer (2017)
6. Bardeh, N.G., Rønjom, S.: The Exchange Attack: How to Distinguish Six Rounds of AES with $2^{88.2}$ Chosen Plaintexts. In: Galbraith, S.D., Moriai, S. (eds.) *ASIACRYPT 2019, Part III*. LNCS, vol. 11923, pp. 347–370. Springer (2019)
7. Beierle, C., Jean, J., Kölbl, S., Leander, G., Moradi, A., Peyrin, T., Sasaki, Y., Sasdrich, P., Sim, S.M.: The SKINNY family of block ciphers and its low-latency variant MANTIS. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part II*. LNCS, vol. 9815, pp. 123–153. Springer (2016)
8. Bellare, M., Impagliazzo, R.: A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. *IACR Cryptol. ePrint Arch.* p. 24 (1999)
9. Bellare, M., Krovetz, T., Rogaway, P.: Luby-rackoff backwards: Increasing security by making block ciphers non-invertible. In: Nyberg, K. (ed.) *EUROCRYPT '98*. LNCS, vol. 1403, pp. 266–280. Springer (1998)
10. Biham, E., Biryukov, A., Dunkelman, O., Richardson, E., Shamir, A.: Initial observations on skipjack: Cryptanalysis of skipjack-3xor. In: Tavares, S.E., Meijer, H. (eds.) *SAC'98*. LNCS, vol. 1556, pp. 362–376. Springer (1998)
11. Biham, E., Biryukov, A., Shamir, A.: Cryptanalysis of skipjack reduced to 31 rounds using impossible differentials. In: Stern, J. (ed.) *EUROCRYPT '99*. LNCS, vol. 1592, pp. 12–23. Springer (1999)
12. Biham, E., Dunkelman, O., Keller, N.: The rectangle attack - rectangling the serpent. In: Pfitzmann, B. (ed.) *EUROCRYPT 2001*. LNCS, vol. 2045, pp. 340–357. Springer (2001)

13. Biham, E., Shamir, A.: Differential cryptanalysis of des-like cryptosystems. In: Menezes, A., Vanstone, S.A. (eds.) CRYPTO '90. LNCS, vol. 537, pp. 2–21. Springer (1990)
14. Bogdanov, A., Leander, G., Nyberg, K., Wang, M.: Integral and multidimensional linear distinguishers with correlation zero. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 244–261. Springer (2012)
15. Bogdanov, A., Rijmen, V.: Linear hulls with correlation zero and linear cryptanalysis of block ciphers. *Des. Codes Cryptogr.* **70**(3), 369–383 (2014)
16. Borghoff, J., Canteaut, A., Güneysu, T., Kavun, E.B., Knezevic, M., Knudsen, L.R., Leander, G., Nikov, V., Paar, C., Rechberger, C., Rombouts, P., Thomsen, S.S., Yalçın, T.: PRINCE - A low-latency block cipher for pervasive computing applications (full version). *IACR Cryptol. ePrint Arch.* p. 529 (2012)
17. Bozilov, D., Eichlseder, M., Knezevic, M., Lambin, B., Leander, G., Moos, T., Nikov, V., Rasoolzadeh, S., Todo, Y., Wiemer, F.: Princev2 - more security for (almost) no overhead. In: Dunkelman, O., Jr., M.J.J., O'Flynn, C. (eds.) SAC 2020. LNCS, vol. 12804, pp. 483–511. Springer (2020)
18. Cid, C., Murphy, S., Robshaw, M.J.B.: Small Scale Variants of the AES. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 145–162. Springer (2005)
19. Daemen, J., Rijmen, V.: The Wide Trail Design Strategy. In: Honary, B. (ed.) 8th IMA. LNCS, vol. 2260, pp. 222–238. Springer (2001)
20. Daemen, J., Rijmen, V.: Security of a Wide Trail Design. In: Menezes, A., Sarkar, P. (eds.) INDOCRYPT 2002. LNCS, vol. 2551, pp. 1–11. Springer (2002)
21. Daemen, J., Rijmen, V.: The Design of Rijndael - The Advanced Encryption Standard (AES), Second Edition. *Information Security and Cryptography*, Springer (2020)
22. Demirci, H., Selçuk, A.A.: A meet-in-the-middle attack on 8-round AES. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 116–126. Springer (2008)
23. Diffie, W., Hellman, M.E.: Exhaustive cryptanalysis of the NBS data encryption standard. *Computer* **10**(6), 74–84 (1977)
24. Dinur, I.: Tight indistinguishability bounds for the XOR of independent random permutations by fourier analysis. In: Joye, M., Leander, G. (eds.) EUROCRYPT 2024, Part I. LNCS, vol. 14651, pp. 33–62. Springer (2024)
25. Dinur, I., Shamir, A.: Cube attacks on tweakable black box polynomials. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 278–299. Springer (2009)
26. Dutta, A., Nandi, M., Saha, A.: Proof of mirror theory for $\xi_{\max} = 2$. *IEEE Trans. Inf. Theory* **68**(9), 6218–6232 (2022)
27. Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D.A., Whiting, D.: Improved Cryptanalysis of Rijndael. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 213–230. Springer (2000)
28. Flórez-Gutiérrez, A., Grassi, L., Leander, G., Sibleyras, F., Todo, Y.: General practical cryptanalysis of the sum of round-reduced block ciphers and zip-aes, full version of this paper
29. Gilbert, H.: A Simplified Representation of AES. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part I. LNCS, vol. 8873, pp. 200–222. Springer (2014)
30. Grassi, L.: Mixture Differential Cryptanalysis: a New Approach to Distinguishers and Attacks on round-reduced AES. *IACR ToSC* **2018**(2), 133–160 (2018)
31. Grassi, L., Rechberger, C.: Truncated Differential Properties of the Diagonal Set of Inputs for 5-Round AES. In: Nguyen, K., Yang, G., Guo, F., Susilo, W. (eds.) ACISP 2022. LNCS, vol. 13494, pp. 24–45. Springer (2022)
32. Grassi, L., Rechberger, C., Rønjom, S.: Subspace Trail Cryptanalysis and its Applications to AES. *IACR ToSC* **2016**(2), 192–225 (2016)

33. Grassi, L., Schafneger, M.: Mixture Integral Attacks on Reduced-Round AES with a Known/Secret S-Box. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) INDOCRYPT 2020. LNCS, vol. 12578, pp. 312–331. Springer (2020)
34. Guning, A., Bhaumik, R., Jha, A., Mennink, B., Shen, Y.: Revisiting the indistinguishability of the sum of permutations. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part III. LNCS, vol. 14083, pp. 628–660. Springer (2023)
35. Hadipour, H., Derbez, P., Eichlseder, M.: Revisiting differential-linear attacks via a boomerang perspective with application to aes, ascon, clefia, skinny, present, knot, twine, warp, lblock, simeck, and SERPENT. In: Reyzin, L., Stebila, D. (eds.) CRYPTO 2024, Part IV. LNCS, vol. 14923, pp. 38–72. Springer (2024)
36. Hebborn, P., Lambin, B., Leander, G., Todo, Y.: Lower bounds on the degree of block ciphers. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part I. LNCS, vol. 12491, pp. 537–566. Springer (2020)
37. Hebborn, P., Lambin, B., Leander, G., Todo, Y.: Strong and tight security guarantees against integral distinguishers. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part I. LNCS, vol. 13090, pp. 362–391. Springer (2021)
38. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 15–44. Springer (2015)
39. Jha, A., Nandi, M.: A survey on applications of h-technique: Revisiting security analysis of PRP and PRF. *Entropy* **24**(4), 462 (2022)
40. Kelsey, J., Kohno, T., Schneier, B.: Amplified boomerang attacks against reduced-round MARS and serpent. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 75–93. Springer (2000)
41. Knudsen, L.R.: Truncated and higher order differentials. In: Preneel, B. (ed.) FSE 2nd. LNCS, vol. 1008, pp. 196–211. Springer (1994)
42. Knudsen, L.R., Rijmen, V.: Known-Key Distinguishers for Some Block Ciphers. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 315–324. Springer (2007)
43. Leander, G., Moos, T., Moradi, A., Rasoolzadeh, S.: The SPEEDY family of block ciphers engineering an ultra low-latency cipher from gate level for secure processor architectures. *IACR TCHES* **2021**(4), 510–545 (2021)
44. Li, M., Sun, L., Wang, M.: Automated key recovery attacks on round-reduced orthros. In: Batina, L., Daemen, J. (eds.) AFRICACRYPT 2022. LNCS, vol. 13503, pp. 189–213. Springer Nature Switzerland (2022)
45. Lucks, S.: The sum of prps is a secure PRF. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 470–484. Springer (2000)
46. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseth, T. (ed.) EUROCRYPT '93. LNCS, vol. 765, pp. 386–397. Springer (1993)
47. Mennink, B., Neves, S.: Encrypted davies-meyer and its dual: Towards optimal security using mirror theory. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 556–583. Springer (2017)
48. Mennink, B., Neves, S.: Optimal PRFs from Blockcipher Designs. *IACR ToSC* **2017**(3), 228–252 (2017)
49. Patarin, J.: Introduction to mirror theory: Analysis of systems of linear equalities and linear non equalities for cryptography. *IACR Cryptol. ePrint Arch.* p. 287 (2010)
50. Sun, B., Liu, Z., Rijmen, V., Li, R., Cheng, L., Wang, Q., AlKhazaimi, H., Li, C.: Links among impossible differential, integral and zero correlation linear cryptanalysis. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 95–115. Springer (2015)

51. Taka, K., Ishikawa, T., Sakamoto, K., Isobe, T.: An efficient strategy to construct a better differential on multiple-branch-based designs: Application to orthros. In: Rosulek, M. (ed.) *CT-RSA 2023*. LNCS, vol. 13871, pp. 277–304. Springer (2023)
52. Wagner, D.A.: The boomerang attack. In: Knudsen, L.R. (ed.) *FSE '99*. LNCS, vol. 1636, pp. 156–170. Springer (1999)
53. Wang, S., Liu, M., Hou, S., Lin, D.: Differential-linear cryptanalysis of GIFT family and GIFT-based ciphers. *IACR Communications in Cryptology* **1**(1) (2024)



Robust AE With Committing Security

Viet Tung Hoang¹(✉) and Sanketh Menda²

¹ Department of Computer Science, Florida State University, Tallahassee, USA

tvhoang@cs.fsu.edu

² Cornell Tech, New York, USA

<https://cs.fsu.edu/~tvhoang/>, <https://snkth.com>

Abstract. There has been a recent interest to develop and standardize Robust Authenticated Encryption (Robust AE) schemes. NIST, for example, is considering an Accordion mode (a wideblock tweakable blockcipher), with Robust AE as a primary application. On the other hand, recent attacks and applications suggest that encryption needs to be committing. Indeed, committing security is also a design consideration in the Accordion mode. Yet it is unclear how to build a Robust AE with committing security.

In this work, we give a modular solution for this problem. We first show how to transform any wideblock tweakable blockcipher TE to a Robust AE scheme SE that commits just the key. The overhead is cheap, just a few finite-field multiplications and blockcipher calls. If one wants to commit the entire encryption context, one can simply hash the context to derive a 256-bit subkey, and uses SE on that subkey. The use of 256-bit key on SE only means that it has to rely on AES-256 but doesn't require TE to have 256-bit key.

Our approach frees the Accordion designs from consideration of committing security. Moreover, it gives a big saving for several key-committing applications that don't want to pay the inherent hashing cost of full committing.

Keywords: Robust authenticated encryption · committing security

1 Introduction

Authenticated Encryption (AE) is widely used in practice to provide data privacy and authenticity. Yet standard AE schemes such as GCM are both fragile and inflexible. On the one hand, if some misuse happens, say nonce repetition, then security is completely broken. On the other hand, for standard AE schemes, the ciphertext C must be sufficiently longer than the message M , so that forgeries will never happen in practice. In particular the *ciphertext expansion* $\tau = |C| - |M|$ is typically 128 bits. But several applications, such as Voice-over-IP or IoT, demand shorter expansion (say 64 bits, or even 32 bits) to minimize latency or energy consumption. One cannot simply truncate the tag of a standard AE

scheme because forgeries will happen sooner or later, and once a forgery happens, all security guarantees are voided.¹

ROBUST AE. There has been a long line of work to deal with the situation above [2, 3, 18, 26], culminating in the Robust AE notion of Hoang, Krovetz, and Rogaway [18]. Instead of using a fixed expansion τ , Robust AE explicitly accepts a user-defined choice of τ for each encryption. If τ is small then forgeries of course will happen, but just occasionally, roughly once per 2^τ attempts. Robust AE also provides misuse resistance for nonce reuse [26] and protects against releases of unverified decrypted messages [2]. Because of such strong guarantees, standard agencies are actively seeking Robust AE schemes for standardization. The UK National Cyber Security Centre, for example, recently release their own Robust AE schemes [9]. NIST is also considering an Accordion mode for a wideblock tweakable blockcipher (TBC), with Robust AE as a primary application.

THE NEED FOR COMMITTING SECURITY. Robust AE aims to provide the best security possible, but it only covers privacy and authenticity. Recent attacks, such as the Partitioning-Oracle attack on password-based encryption [20], highlight the need for encryption to be committing. This guarantee is also needed by several recent applications, such as Facebook’s Message Franking [16], Amazon Cloud encryption, Subscribe with Google [1], or TLS Oracle [22]. Most applications require committing just the key K , but some need to commit all the four inputs (K, N, A, M) of encryption. Due to such demand, it is desirable to build a scheme that provides both Robust AE and committing security. Indeed, committing security is also a property that NIST are considering in the call for the Accordion mode.

OBSTACLES. Unfortunately, existing Robust AE schemes such as AEZ [18] or HCTR2 [14] do not offer (full) committing security. The obvious reason is their hashing of the associated data via a universal hash instead of a collision-resistant one. However, there is a subtle, quantitative reason. In particular, suppose that we only need s bits of expansion and target s bits of committing security.² In prior constructions of committing AE schemes [1, 4, 5, 11], the common approach is to make the tag a commitment of (K, N, A) , and the message will be committed due to decryption correctness. However, this approach doesn’t work for Robust AE schemes. First, Robust AE can only be realized via the Encode-then-Encipher (EtE) paradigm [7]: encode the message with s -bit redundancy (say padding it with 0^s) and then encipher it with a wideblock TBC. The EtE method has no tag, defeating the prior approach of building committing AE.

¹ If the expansion is short and an adversary can obtain decrypted messages, standard AE schemes are *inherently* insecure because the encryption algorithm makes just one pass over data. Specifically, two ciphertexts of the same prefix would decrypt to two messages of the same prefix.

² Generic attacks [5] show that we can at best hope for s -bit committing security given s -bit expansion. Moreover, given that committing attacks are offline, we want to achieve s bits of committing security instead of a “birthday-bound” $s/2$ bits of security.

Next, birthday attacks suggest that the commitment must be at least $2s$ bits, but we only have s -bit space.

Given the birthday attacks, at the first glance, it seems that to provide s -bit of committing security, we are doomed to use at least $2s$ bits of expansion. However, a closer inspection reveals that the attacks only require that the ciphertext must be at least $2s$ bits. Thus as long as messages are at least s bits long³ then the attacks do not apply. This gives a way out of the impossibility, but it is unclear how to exploit this opening.

RELATED WORK. Chen et al. [13] show that AEZ [18] offers 64-bit key-committing security for 128-bit expansion, assuming that the underlying tweakable blockcipher is modeled as ideal. (The result is tight, with a matching attack.) However, this security guarantee is weak, because given 128-bit expansion, we want 128-bit security, not merely 64-bit. Moreover, committing security has never been a design goal of AEZ, and thus the security by accident here gives no insight on how one should build a committing Robust AE scheme.

A very recent work by Bellare and Hoang [5] considers adding s bits of committing security to a base AE scheme using just s bits of expansion (assuming that messages are at least s bits). Their work only deals with tag-based AE schemes and thus doesn't apply to Robust AE. However, implicitly their work contains a technical tool that is central to our construction. We will elaborate later how to use their ideas for our setting.

CONTRIBUTIONS. We initiate the study of committing Robust AE. By extending the definitions of Bellare and Hoang [4], we formalize two notions of committing security: (1) the CMT notion that commit all inputs (K, N, A, τ, M) of the encryption algorithm, and (2) the CMT-1 notion that commits just the key K .

Achieving CMT security demands that one hashes the associated data A with a cryptographic hash function, such as SHA-2 or SHA-3. While this cost is $O(1)$ in theory, the actual relative overhead is huge for small data. We therefore consider a modular route. First, we focus on building a CMT-1 Robust AE scheme SE that is enough for most applications and doesn't have to pay the hashing penalty. Next, for applications that demand the full CMT security, one can *non-intrusively* add CMT security to SE via the Hash-then-Encrypt (HtE) transform of Hoang and Bellare [4]: first hash (K, N, A, τ) to derive a subkey L , and then encrypt with key L , the empty nonce, the empty AD, and expansion τ . Using HtE means that SE needs to use 256-bit key, but this aligns well with (i) NIST's requirement that an Accordion mode must support 256-bit key, and (ii) the fact that our CMT-1 construction has to use a blockcipher of 256-bit key anyway.

For CMT-1 security, we build a transform EwC that turns any wideblock TBC TE to another $\overline{\text{TE}}$ such that using the latter in the EtE method provides CMT-1 security. While EwC uses a 256-bit key, it doesn't require the base TE

³ This assumption is reasonable. Indeed, existing Robust AE schemes can't encrypt tiny messages, as AES-based wideblock TBC can only efficiently encipher messages of at least 128 bits.

to have a 256-bit key. This allows us to leverage a wealth of existing wideblock TBC constructions and future Accordion submissions. The overhead of EwC is cheap, just a few finite-field multiplications and blockcipher calls. The underlying blockcipher E uses a 256-bit key, and thus can be instantiated via AES-256 or Rijndael-256. Moreover, EwC only uses E in the forward direction, which saves code size in hardware implementation.

If the underlying blockcipher E has 256-bit block length (such as Rijndael-256), then EwC would encipher messages of at least 512 bits, and provides s bits of CMT-1 security when used in EtE with s bits expansion. If E only has 128-bit block length (such as AES-256) then the CMT-1 security of EwC is more nuanced, because implicitly it takes a parameter $\ell < 128$, and enciphers messages at least $256 + \ell$ bits. Using EwC $[\ell]$ in EtE still allows one to use any expansion s but only guarantees to deliver CMT-1 security when $s > \ell$, and in that case it provides just $\ell - 8$ bits of CMT-1 security.

At the bird’s-eye view, EwC is a four-round (unbalanced) Feistel-like structure. (See Fig. 15 for an illustration.) The first and last rounds, following Naor and Reingold [24], are based on an AXU hash function. Since the input length of the AXU is short, it amounts to just one or two finite-field multiplications for each hashing. The second round uses TE. The third round uses a collision-resistant PRF H and a *committing concealer*, a new primitive that we will discuss later. The function H only needs to deal with short inputs and has to produce a 256-bit output. Thus if the blockcipher E has 256-bit block length, we can directly instantiate H via the Davies-Meyer construction, meaning $H(K, M) = E_K(M) \oplus M$. If E has just 128-bit blockcipher, we show how to build H via “doubling” Davies-Meyer, meaning $H(K, M) = (E_K(U) \oplus U) \parallel (E_K(V) \oplus V)$, where $U = M \parallel 0$ and $V = M \parallel 1$.

COST. The overhead of EwC is listed in Table 1. For each instantiation of the blockcipher (AES-256 or Rijndael-256), we list two values for the number of blockcipher calls because (i) if one only uses the standalone EwC for key-committing security, we can ignore the cost of subkey generation since the subkeys can be cached, but (ii) if one uses EwC with the HtE transform for full committing security, then the cost of subkey generation must be included. While the subkey generation seems expensive (say six AES calls), these blockcipher calls are fully parallelizable. On platforms with vector AES instructions, these AES calls would take almost the same amount of time as couple AES calls.

Table 1. The overhead of the EwC transform. The third column shows the number of blockcipher calls; two numbers are given, one includes the cost of subkey generation and another doesn’t. The last column shows the number of finite-field multiplications.

Block length n	Blockcipher	Blockcipher calls	Mults in $\text{GF}(2^n)$
128	AES-256	10 (or 4 if subkeys are cached)	4
256	Rijndael-256	6 (or 2 if subkeys are cached)	2

Implementing a performant version of EwC is tricky. The problem comes from having three AES keys (one for the wideblock TBC, another for the Davies-Meyer, and yet another for the committing concealer). Overall, that generates a lot of AES subkeys, and it is tricky to ensure that we have no register spill. See Fig. 16 for experimental data for the overhead of EwC on HCTR2. The cost is significant for small data, but becomes negligible for messages bigger than 1 KB.

COMMITTING CONCEALER. Central to our EwC transform is a new primitive that we call *committing concealer*. Committing concealer can be viewed as a blockcipher but the security requirement is different. Traditionally, we want a blockcipher to be a strong PRP; if we build it from a Feistel network, we need at least four rounds. For committing concealer, we only need it to be a *one-time* strong PRP (meaning that the adversary can make just a single query per user), and thus we can build it from just two-round Feistel. Still, we want committing concealer to commit the key if used only on the set of messages that are encoded with s -bit redundancy.

If we have a blockcipher E of 256-bit block length, one can directly use it as a committing concealer; the key-committing property would be justified by modeling E as an ideal cipher. However, it is tricky if E only has 128-bit block length. The core idea is implicitly in the recent work of Bellare and Hoang [5]. Technically, they need to commit messages up to $m < n$ bits to produce a commitment of $(m+n)$ -bit length and nearly n bits of binding security. They give a construction via the SIV paradigm [26], but alternatively, their construction can be viewed as padding the message with zeros, and then enciphering it with a committing concealer. Their (implicit) committing concealer is based on a two-round unbalanced Feistel, where the round functions are implemented via the Davies-Meyer construction on a blockcipher of block length n . See Fig. 13 for an illustration.

THE USE OF IDEAL-CIPHER MODEL. If one uses our EwC transform in the EtE construction, one would have Robust AE security in the standard model. But the committing security has to be justified in the ideal-cipher model. The use of idealized models in cryptographic constructions has always been a controversial issue, and sometimes it can stand in the way towards adoption. The UK National Cyber Security Centre, for example, do not target committing security in their Robust AE schemes due to concerns about the use of the ideal-cipher model [9]. We argue that as long as the Robust AE is still justified in the standard model, we can only gain by additionally providing committing security (even in the ideal-cipher model).

2 Preliminaries

2.1 Notation and Terminology

Let ε denote the empty string. For a string x we write $|x|$ to refer to its bit length, and $x[i : j]$ is the bits i through j (inclusive) of x , for $1 \leq i \leq j \leq |x|$.

<p><u>Game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$</u></p> <p>$v \leftarrow 0$; $b \leftarrow_{\\$} \{0, 1\}$; $b' \leftarrow_{\\$} \mathcal{A}^{\text{NEW, EVAL}}$</p> <p>Return $(b' = b)$</p> <p><u>NEW()</u></p> <p>$v \leftarrow v + 1$; $K_v \leftarrow_{\\$} \mathcal{K}$</p> <p>$f_v \leftarrow_{\\$} \text{Func}(\text{Dom}, \text{Rng})$</p>	<p><u>EVAL(i, M)</u></p> <p>If $i \notin \{1, \dots, v\}$ return \perp</p> <p>$C_1 \leftarrow F(K_i, M)$; $C_0 \leftarrow f_i(M)$</p> <p>Return C_b</p>
---	---

Fig. 1. Game defining (multi-user) PRF security of F .

By $\text{Func}(\text{Dom}, \text{Rng})$ we denote the set of all functions $f : \text{Dom} \rightarrow \text{Rng}$. We use \perp as a special symbol to denote rejection, and it is assumed to be outside $\{0, 1\}^*$. If X is a finite set, we let $x \leftarrow_{\$} X$ denote picking an element of X uniformly at random and assigning it to x . For an integer $n \geq 1$, let $\{0, 1\}^{\leq n}$ denote the set of all bit strings whose length is at most n , and let $\{0, 1\}^{\geq n}$ denote the set of all bit strings whose length is at least n .

2.2 Some Standard Primitives

COLLISION RESISTANCE. Let $H : \text{Dom} \rightarrow \text{Rng}$ be a function. A collision for H is a pair (X_1, X_2) of distinct points in Dom such that $H(X_1) = H(X_2)$. For an adversary \mathcal{A} , define its advantage in breaking the collision resistance of H as

$$\text{Adv}_H^{\text{coll}}(\mathcal{A}) = \Pr[(X_1, X_2) \text{ is a collision for } H]$$

where the probability is over $(X_1, X_2) \leftarrow_{\$} \mathcal{A}$.

AXU HASH. Let $G : \mathcal{K} \times \mathcal{M} \rightarrow \{0, 1\}^n$ be a keyed hashed function. We say that G is *c-almost XOR-universal* (*c-AXU*) if for any distinct $X, Y \in \mathcal{M}$ and any $\Delta \in \{0, 1\}^n$,

$$\Pr_{K \leftarrow_{\$} \mathcal{K}}[G_K(X) \oplus G_K(Y) = \Delta] \leq \frac{c}{2^n} .$$

PRF. For a function $F : \mathcal{K} \times \text{Dom} \rightarrow \text{Rng}$ and an adversary \mathcal{A} , we define the advantage of \mathcal{A} in breaking the (multi-user) PRF security of F as

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_F^{\text{prf}}(\mathcal{A})] - 1 ,$$

where game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$ is shown in Fig. 1.

PRP. For a blockcipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and an adversary \mathcal{A} , we define the advantage of \mathcal{A} in breaking the (multi-user) strong-PRP security of E as

$$\text{Adv}_E^{\pm \text{prp}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_E^{\pm \text{prp}}(\mathcal{A})] - 1 ,$$

where game $\mathbf{G}_E^{\pm \text{prp}}(\mathcal{A})$ is shown in Fig. 2.

<p><u>Game $\mathbf{G}_E^{\pm\text{prp}}(\mathcal{A})$</u></p> <p>$v \leftarrow 0$; $b \leftarrow_{\\$} \{0, 1\}$ $b' \leftarrow_{\\$} \mathcal{A}^{\text{NEW, ENC, DEC}}$ Return ($b' = b$)</p> <p><u>NEW()</u></p> <p>$v \leftarrow v + 1$; $K_v \leftarrow_{\\$} \{0, 1\}^k$ $\Pi_v \leftarrow_{\\$} \text{Perm}(n)$</p>	<p><u>ENC(i, M)</u></p> <p>If $i \notin \{1, \dots, v\}$ return \perp $C_1 \leftarrow E(K_i, M)$; $C_0 \leftarrow \Pi_i(M)$ Return C_b</p> <p><u>DEC(i, C)</u></p> <p>If $i \notin \{1, \dots, v\}$ return \perp $M_1 \leftarrow E^{-1}(K_i, C)$; $M_0 \leftarrow \Pi_i^{-1}(C)$ Return M_b</p>
--	---

Fig. 2. Game defining (multi-user) strong PRP security of E . Here $\text{Perm}(n)$ denotes the set of all permutations in $\{0, 1\}^n$

<p><u>Game $\mathbf{G}_{\text{TE}}^{\pm\widetilde{\text{prp}}}(\mathcal{A})$</u></p> <p>$v \leftarrow 0$; $b \leftarrow_{\\$} \{0, 1\}$ $b' \leftarrow_{\\$} \mathcal{A}^{\text{NEW, ENC, DEC}}$ Return ($b' = b$)</p> <p><u>NEW()</u></p> <p>$v \leftarrow v + 1$; $K_v \leftarrow_{\\$} \{0, 1\}^k$ For $T \in \mathcal{T}$ do $\Pi_{v,T} \leftarrow_{\\$} \text{LP}(\mathcal{M})$</p>	<p><u>ENC(i, T, M)</u></p> <p>If $i \notin \{1, \dots, v\}$ return \perp $C_1 \leftarrow \text{TE.Enc}(K_i, T, M)$; $C_0 \leftarrow \Pi_{i,T}(M)$ Return C_b</p> <p><u>DEC(i, T, C)</u></p> <p>If $i \notin \{1, \dots, v\}$ return \perp $M_1 \leftarrow \text{TE.Dec}(K_i, T, C)$; $M_0 \leftarrow \Pi_{i,T}^{-1}(C)$ Return M_b</p>
--	--

Fig. 3. Game defining (multi-user) strong tweakable-PRP security of TE . Here $\text{LP}(\mathcal{M})$ denote the set of permutations π on \mathcal{M} that are length-preserving, meaning $|\pi(M)| = |M|$ for every $M \in \mathcal{M}$.

(WIDEBLOCK) TWEAKABLE BLOCKCIPHER. A *tweakable blockcipher* (TBC) TE consists of two deterministic algorithms TE.Enc and TE.Dec , and is associated with a key space \mathcal{K} , a message space \mathcal{M} , and a tweak space \mathcal{T} . The enciphering algorithm TE.Enc takes as input a key $K \in \mathcal{K}$, a message $M \in \mathcal{M}$, a tweak $T \in \mathcal{T}$, and outputs a ciphertext $C \leftarrow \text{TE.Enc}(K, T, M)$. The deciphering algorithm TE.Dec takes as input (K, T, C) and produces $M \leftarrow \text{TE.Dec}(K, T, C)$. For correctness, we require that deciphering reverses enciphering, meaning that if $C \leftarrow \text{TE.Enc}(K, T, M)$ then $\text{TE.Dec}(K, T, C) = M$.

In this paper, we consider *wideblock* TBC, meaning that the message space consists of messages of different length, say $\mathcal{M} = \{0, 1\}^{\geq m}$. We require that enciphering preserves the message length, meaning that $|C| = |M|$.

Define the advantage of an adversary \mathcal{A} breaking the strong tweakable-PRP security of TE as

$$\text{Adv}_{\text{TE}}^{\pm\widetilde{\text{prp}}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{\text{TE}}^{\pm\widetilde{\text{prp}}}(\mathcal{A})] - 1 ,$$

where game $\mathbf{G}_{\text{TE}}^{\pm\widetilde{\text{prp}}}(\mathcal{A})$ is shown in Fig. 3.

2.3 Robust Authenticated Encryption

SYNTAX. An *robust authenticated encryption* (RAE) scheme SE consists of two deterministic algorithms SE.Enc and SE.Dec; it is associated with a nonce space \mathcal{N} , a message space \mathcal{M} , key space \mathcal{K} , and an expansion space \mathcal{X} . The encryption algorithm takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \{0, 1\}^*$, a message $M \in \mathcal{M}$, and an expansion $\tau \in \mathcal{X}$, and returns a ciphertext $C \leftarrow \text{SE.Enc}(K, N, A, M, \tau)$ such that $|C| = |M| + \tau$. The decryption algorithm takes as input (K, N, A, C, τ) and returns either a message $M \in \mathcal{M}$ or a leakage $L \notin \mathcal{M}$. The correctness requirement says that decryption reverses encryption, namely if $C \leftarrow \text{SE.Enc}(K, N, A, M, \tau)$ then $\text{SE.Dec}(K, N, A, C, \tau)$ returns M .

We say that SE is *tidy* [23] if $M \leftarrow \text{SE.Dec}(K, N, A, C, \tau)$ and $M \in \mathcal{M}$ imply that $\text{SE.Enc}(K, N, A, M, \tau)$ returns C . Combining correctness and tidiness means that functions $\text{SE.Enc}(K, N, A, \cdot, \tau)$ and $\text{SE.Dec}(K, N, A, \cdot, \tau)$ are the inverse of each other. The schemes we consider will be tidy.

Standard AE schemes are a special case of RAE schemes where the set \mathcal{X} is a singleton, meaning that the expansion is a constant, say 128. In general, RAE schemes support a large range of expansion values, typically $\{0, 1, \dots, 128\}$, and the expansion value can dynamically change within the same session.

RAE SECURITY. Let SE be an RAE scheme of message space \mathcal{M} . Define the advantage of an adversary \mathcal{A} breaking the RAE security of SE with respect to a (stateful) simulator Sim as

$$\text{Adv}_{\text{SE, Sim}}^{\text{rae}}(\mathcal{A}) = 2 \cdot \Pr[\mathbf{G}_{\text{SE, Sim}}^{\text{rae}}(\mathcal{A})] - 1 ,$$

where game $\mathbf{G}_{\text{SE, Sim}}^{\text{rae}}(\mathcal{A})$ is defined in Fig. 4. To prevent trivial attacks, we forbid the adversary from first querying $C \leftarrow \text{ENC}(i, N, A, M, \tau)$ and then querying $\text{DEC}(i, N, A, C, \tau)$.

In the game above, the simulator Sim is only called on invalid ciphertexts to simulate the decryption leakage, but is *not* given any information on messages of encryption queries. The simulator is stateful and explicitly maintains its state st .

ENCODE-THEN-ENCIPHER (ETE). Hoang, Krovetz, and Rogaway [18] show that to achieve RAE security, one has to use the Encode-then-Enciphering (EtE) paradigm of Bellare and Rogaway [7]. We now recall the details of the EtE construction. Let $\text{pad} : \{0, 1\}^* \times \mathcal{X} \rightarrow \{0, 1\}^*$ be a padding scheme such that (1) for any $\tau \in \mathcal{X}$, the function $\text{pad}(\cdot, \tau)$ is injective, and let $\text{unpad}(\cdot, \tau) : \{0, 1\}^* \rightarrow \{0, 1\}^* \cup \{\perp\}$ be its inverse, and (2) If $Y \leftarrow \text{pad}(X, \tau)$ then $|Y| = |X| + \tau$. For example, we can let $\text{pad}(M, \tau) = M \parallel 0^\tau$. Let TE be a wideblock TBC with tweak space $\mathcal{N} \times \{0, 1\}^* \times \mathcal{X}$. The scheme EtE[TE] is specified in Fig. 5; it has nonce space \mathcal{N} and expansion space \mathcal{X} . Informally, to encrypt M under (K, N, A, τ) , we pad M with $\text{pad}(\cdot, \tau)$ and then encipher it with tweak (N, A, τ) . On decryption, we first recover $V \leftarrow \text{TE.Dec}(K, T, C)$ and check the unpadding. If the unpadding fails then we model the decryption leakage as $L \leftarrow (\perp, V)$. (Note that since we require $L \notin \mathcal{M}$, the invalidity symbol has to be included.)

<p>Game $\mathbf{G}_{\text{SE}, \text{Sim}}^{\text{rae}}(\mathcal{A})$</p> <p>$v \leftarrow 0$; $b \leftarrow_{\\$} \{0, 1\}$; $st \leftarrow \varepsilon$ $b' \leftarrow_{\\$} \mathcal{A}^{\text{NEW, ENC, DEC}}$</p> <p>Return $(b' = b)$</p> <hr/> <p>$\text{NEW}()$</p> <p>$v \leftarrow v + 1$; $K_v \leftarrow_{\\$} \{0, 1\}^k$ For $(N, A, \tau) \in \mathcal{N} \times \{0, 1\}^* \times \mathcal{X}$ do $\Pi_{v, N, A, \tau} \leftarrow_{\\$} \text{Inj}(\tau)$</p>	<p>$\text{ENC}(i, N, A, M, \tau)$</p> <p>If $i \notin \{1, \dots, v\}$ return \perp $C_1 \leftarrow \text{SE.Enc}(K_i, N, A, M, \tau)$ $C_0 \leftarrow \Pi_{i, N, A, \tau}(M)$ Return C_b</p> <hr/> <p>$\text{DEC}(i, N, A, C, \tau)$</p> <p>If $i \notin \{1, \dots, v\}$ return \perp $M_1 \leftarrow \text{SE.Dec}(K_i, N, A, C, \tau)$ $M_0 \leftarrow \Pi_{i, N, A, \tau}^{-1}(C)$ If $M_0 = \perp$ then $(M_0, st) \leftarrow_{\\$} \text{Sim}(i, N, A, C, \tau, st)$ Return M_b</p>
--	--

Fig. 4. Game defining (multi-user) RAE security of SE with respect to a simulator Sim. Here $\text{Inj}(\tau)$ denote the set of injective functions $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ such that $|f(M)| = |M| + \tau$ for all messages M .

<p>$\text{EtE}[\text{TE}].\text{Enc}(K, N, A, M, \tau)$</p> <p>$V \leftarrow \text{pad}(M, \tau)$; $T \leftarrow (N, A, \tau)$ $C \leftarrow \text{TE.Enc}(K, T, V)$ Return C</p>	<p>$\text{EtE}[\text{TE}].\text{Dec}(K, N, A, C, \tau)$</p> <p>$T \leftarrow (N, A, \tau)$ $V \leftarrow \text{TE.Dec}(K, T, C)$; $M \leftarrow \text{unpad}(V, \tau)$ If $M \neq \perp$ then return M Else return (\perp, V) // Decryption leakage</p>
--	---

Fig. 5. The EtE method, with decryption leakage on invalid ciphertexts.

3 Committing Security for RAE Schemes

In this section, we give a definitional treatment of committing security for RAE schemes. The definitions are a straightforward extension of the work of Bellare and Hoang [4] for standard AE schemes. The main issue is whether to restrict adversaries from generating collisions on the same, or possibly different, expansions τ_1 and τ_2 . It is easy to see that allowing different ones is a strictly stronger security goal, and so we opt for it.

3.1 Definitions

COMMITTING SECURITY FOR RAE SCHEMES. For an adversary \mathcal{A} , we define its advantage in breaking the committing security of an RAE scheme SE as

$$\text{Adv}_{\text{SE}}^{\text{cmt}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{SE}}^{\text{cmt}}(\mathcal{A})],$$

where game $\mathbf{G}_{\text{SE}}^{\text{cmt}}(\mathcal{A})$ is defined in Fig. 6. Informally, committing security means that an adversary cannot produce a ciphertext collision. This generalizes the notion CMT-4 security of Bellare and Hoang [4].

Game $\mathbf{G}_{SE}^{\text{cmt}}(\mathcal{A})$
 $((K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2)) \leftarrow^* \mathcal{A}$
 $C_1 \leftarrow \text{SE.Enc}(K_1, N_1, A_1, M_1, \tau_1)$; $C_2 \leftarrow \text{SE.Enc}(K_2, N_2, A_2, M_2, \tau_2)$
Return $((C_1 = C_2) \wedge (K_1, N_1, A_1, M_1, \tau_1) \neq (K_2, N_2, A_2, M_2, \tau_2))$

Fig. 6. Game defining committing security, encryption-based style.

Game $\mathbf{G}_{SE}^{\text{cmtd}}(\mathcal{A})$
 $(C, (K_1, N_1, A_1, \tau_1), (K_2, N_2, A_2, \tau_2)) \leftarrow^* \mathcal{A}$
 $M_1 \leftarrow \text{SE.Dec}(K_1, N_1, A_1, C_1, \tau_1)$; $M_2 \leftarrow \text{SE.Dec}(K_2, N_2, A_2, C, \tau_2)$
Return $((M_1 \in \mathcal{M}) \wedge (M_2 \in \mathcal{M}) \wedge (K_1, N_1, A_1, \tau_1) \neq (K_2, N_2, A_2, \tau_2))$

Fig. 7. Game defining committing security, decryption-based style. Here \mathcal{M} is the message space.

The definition above uses an encryption-based style where the adversary specifies the messages and the game encrypts them to compare the ciphertexts. Alternatively, we can define a decryption-based one as follows. Define

$$\mathbf{Adv}_{SE}^{\text{cmtd}}(\mathcal{A}) = \Pr[\mathbf{G}_{SE}^{\text{cmtd}}(\mathcal{A})] ,$$

where game $\mathbf{G}_{SE}^{\text{cmtd}}(\mathcal{A})$ is defined in Fig. 7. Informally, this means that a ciphertext cannot be properly decrypted under two different contexts (K_1, N_1, A_1, τ_1) and (K_2, N_2, A_2, τ_2) .

RELATIONS. Following Bellare and Hoang [4], we show that CMT and CMT-D security are equivalent.

▷ CMTD \rightarrow CMT: First we show that CMTD implies CMT. Let SE be an RAE scheme with message space \mathcal{M} . Consider an adversary \mathcal{A}_e that attacks the CMT security of SE. We now construct an adversary \mathcal{A}_d attacking the CMTD security of SE. It runs \mathcal{A}_e to get $(K_1, N_1, A_1, M_1, \tau_1)$ and $(K_2, N_2, A_2, M_2, \tau_2)$. It then computes $C \leftarrow \text{SE.Enc}(K_1, N_1, A_1, M_1, \tau_1)$, and outputs $(C, (K_1, N_1, A_1, \tau_1), (K_2, N_2, A_2, \tau_2))$.

For analysis, without loss of generality, assume that \mathcal{A}_e outputs distinct tuples $(K_1, N_1, A_1, M_1, \tau_1)$ and $(K_2, N_2, A_2, M_2, \tau_2)$. Suppose that \mathcal{A}_e wins its game, meaning that $\text{SE.Enc}(K_2, N_2, A_2, M_2, \tau_2)$ is also C . From the correctness of SE, we have $\text{SE.Dec}(K_i, N_i, A_i, C, \tau_i) = M_i \in \mathcal{M}$ for each $i \in \{1, 2\}$. If $(K_1, N_1, A_1, \tau_1) = (K_2, N_2, A_2, \tau_2)$ then $M_1 = M_2$, which is a contradiction. Hence $(K_1, N_1, A_1, \tau_1) \neq (K_2, N_2, A_2, \tau_2)$, thus \mathcal{A}_d also wins its game. Therefore,

$$\mathbf{Adv}_{SE}^{\text{cmtd}}(\mathcal{A}_d) \geq \mathbf{Adv}_{SE}^{\text{cmt}}(\mathcal{A}_e) .$$

▷ CMT \dashrightarrow CMTD: Conversely, we show that for tidy schemes, CMT implies CMTD. Let SE be a tidy RAE scheme with message space \mathcal{M} . Consider an adversary \mathcal{A}_d that attacks the CMTD security of SE. We now construct

Game $\mathbf{G}_{\text{SE}}^{\text{cmt-1}}(\mathcal{A})$
 $((K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2)) \leftarrow_{\$} \mathcal{A}$
 $C_1 \leftarrow \text{SE.Enc}(K_1, N_1, A_1, M_1, \tau_1)$; $C_2 \leftarrow \text{SE.Enc}(K_2, N_2, A_2, M_2, \tau_2)$
 Return $((C_1 = C_2) \wedge (K_1 \neq K_2))$

Fig. 8. Game defining CMT-1 security.

an adversary \mathcal{A}_e that attacks the CMT security of SE. It runs \mathcal{A}_d to get $(C, (K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2))$, and gets $M_i \leftarrow \text{SE.Dec}(K_i, N_i, A_i, C, \tau_i)$ for each $i \in \{1, 2\}$. It then outputs $((K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2))$.

For analysis, without loss of generality, assume that \mathcal{A}_d outputs distinct tuples $(K_1, N_1, A_1, \tau_1), (K_2, N_2, A_2, \tau_2)$. Suppose that \mathcal{A}_d wins its game, meaning that $M_1 \in \mathcal{M}$ and $M_2 \in \mathcal{M}$. Since SE is tidy, we have $\text{SE.Enc}(K_i, N_i, A_i, M_i, \tau_i) = C$ for each $i \in \{1, 2\}$, and thus \mathcal{A}_e also wins its game. Hence

$$\text{Adv}_{\text{SE}}^{\text{cmt}}(\mathcal{A}_e) \geq \text{Adv}_{\text{SE}}^{\text{cmt-d}}(\mathcal{A}_d) .$$

CMT-1 SECURITY. The notions CMT and CMTD above commit the entire context (K, N, A, M, τ) . Many applications however only need to commit just the key. Following Bellare and Hoang [4], define the advantage of an adversary breaking the CMT-1 of an RAE scheme SE as

$$\text{Adv}_{\text{SE}}^{\text{cmt-1}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{SE}}^{\text{cmt-1}}(\mathcal{A})] ,$$

where game $\mathbf{G}_{\text{SE}}^{\text{cmt-1}}(\mathcal{A})$ is defined in Fig. 8.

DISCUSSION. Note that for CMT security, in the special case that the message is empty, the ciphertext is a τ -bit commitment of the AD. This means that CMT security requires hashing AD by a collision-resistant hash function such as SHA-512 or SHA-3. While this overhead is constant, it's expensive for small messages. In contrast, CMT-1 security only needs to commit a short input (namely the key), and we can use, for example, the Davies-Meyer construction with very low overhead.

LOWER BOUNDS. RAE schemes cannot achieve commitment security for small expansion. For example, if an adversary is allowed to choose $\tau_1 = \tau_2 = 0$, then any ciphertext will decrypt to some message under any context. More formally, let SE be a tidy RAE scheme of expansion space \mathcal{X} . Let λ be the minimum value in \mathcal{X} , and let ℓ be the smallest message length that SE supports. Prior generic committing attacks on standard AE schemes do apply to RAE schemes if we restrict to λ -bit expansion, and treat decryption leakage as a symbol $\perp \notin \mathcal{M}$. In particular, from the attacks of Bellare and Hoang [5], one can at best hope for $\min\{\lambda, (\lambda + \ell)/2\}$ bits of CMT/CMT-1 security for SE.

The term $\lambda + \ell$ is also the smallest message length of the underlying wideblock TBC TE of SE. Practical constructions of TE typically require that $\lambda + \ell$ to

$\text{HtE}[H, \text{SE}].\text{Enc}(K, N, A, M, \tau)$ $L \leftarrow H(K, N, A, \tau)$ $C \leftarrow \text{SE}.\text{Enc}(L, \varepsilon, \varepsilon, M, \tau)$ Return C	$\text{HtE}[H, \text{SE}].\text{Dec}(K, N, A, C, \tau)$ $L \leftarrow H(K, N, A, \tau)$ $M \leftarrow \text{SE}.\text{Dec}(K, \varepsilon, \varepsilon, C, \tau)$ Return M
---	---

Fig. 9. The HtE transform.

be reasonably large, say 128, because otherwise TE has to include a Format-Preserving Encryption scheme [6], which is expensive. For our schemes in Sect. 6, $\lambda + \ell$ is even larger, say 344 for the AES-based instantiation (if we want 80-bit committing security), or 512 for the Rijndael-256-based one.

3.2 From CMT-1 to CMT Security

We extend the Hash-then-Encrypt (HtE) transform of Bellare and Hoang [4] for RAE. This transform turns a CMT-1 secure scheme SE to a CMT-secure one $\text{HtE}[H, \text{SE}]$,

THE HtE TRANSFORM. Let SE be an RAE scheme of key space $\{0, 1\}^k$, nonce space $\{\varepsilon\}$, and expansion space \mathcal{X} . Let $H : \mathcal{K} \times (\mathcal{N} \times \{0, 1\}^* \times \mathcal{X}) \rightarrow \{0, 1\}^k$ be a (keyed) hash function. The code of $\text{HtE}[H, \text{SE}]$ is specified in Fig. 9. The idea is simple. We first hash $L \leftarrow H(K, N, A, \tau)$ to derive a subkey L , and then run SE to encrypt M with the subkey key L and tweak $(\varepsilon, \varepsilon, \tau)$. Note that the AD A is only processed once, because we use SE with the empty AD. The overhead of HtE is essentially optimal, since the hashing of AD is required for achieving CMT security. Therefore, in this paper, we only focus on building CMT-1 secure RAE scheme.

CMT SECURITY OF HtE. The following result shows that if H is collision-resistant then HtE promotes CMT-1 security to CMT one.

Proposition 1. *Let SE be an RAE scheme of key space $\{0, 1\}^k$, nonce space \mathcal{N} , and expansion space \mathcal{X} . Let $H : \{0, 1\}^k \times (\mathcal{N} \times \{0, 1\}^* \times \mathcal{X}) \rightarrow \{0, 1\}^k$ be a (keyed) hash function. For any adversary \mathcal{A} , we can build adversaries \mathcal{B}_1 and \mathcal{B}_2 such that*

$$\text{Adv}_{\text{HtE}[H, \text{SE}]}^{\text{cmt}}(\mathcal{A}) \leq \text{Adv}_H^{\text{coll}}(\mathcal{B}_1) + \text{Adv}_{\text{SE}}^{\text{cmt-1}}(\mathcal{B}_2) .$$

Adversary \mathcal{B}_1 has the same running time and uses the same amount of resource as \mathcal{A} . Adversary \mathcal{B}_2 runs \mathcal{A} and then runs H on \mathcal{A} 's inputs.

Proof. We first describe the adversaries \mathcal{B}_0 and \mathcal{B}_1 . Adversary \mathcal{B}_1 runs

$$((K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2)) \leftarrow^s \mathcal{A} .$$

It then outputs $((K_1, (N_1, A_1, \tau_1)), (K_2, (N_2, A_2, \tau_2)))$.

Adversary \mathcal{B}_2 also runs

$$((K_1, N_1, A_1, M_1, \tau_1), (K_2, N_2, A_2, M_2, \tau_2)) \leftarrow^s \mathcal{A} .$$

Let $L_i \leftarrow H(K_i, (N_i, A_i, \tau_i))$ for every $i \in \{1, 2\}$. Adversary \mathcal{B}_2 then outputs $((L_1, \varepsilon, \varepsilon, M_1, \tau_1), (L_2, \varepsilon, \varepsilon, M_2, \tau_2))$.

For analysis, let $C_i \leftarrow \text{HtE}[H, \text{SE}].\text{Enc}(K_i, N_i, A_i, M_i, \tau_i)$ for each $i \in \{1, 2\}$. Note that $C_i = \text{SE}.\text{Enc}(L_i, \varepsilon, \varepsilon, M_i, \tau_i)$. Assume that \mathcal{A} succeeds in creating a collision, meaning that $C_1 = C_2$. If $L_1 = L_2$ then \mathcal{B}_1 also creates a collision. Suppose that $L_1 \neq L_2$. Then \mathcal{B}_2 creates a collision. Hence

$$\text{Adv}_{\text{HtE}[H, \text{SE}]}^{\text{cmt}}(\mathcal{A}) \leq \text{Adv}_H^{\text{coll}}(\mathcal{B}_1) + \text{Adv}_{\text{SE}}^{\text{cmt}^{-1}}(\mathcal{B}_2)$$

as claimed. □

HtE PRESERVES RAE SECURITY. The following result shows that if H is a PRF then HtE preserves the RAE security.

Proposition 2. *Let SE be an RAE scheme of key space $\{0, 1\}^k$, nonce space $\{\varepsilon\}$, and expansion space \mathcal{X} . Let $H : \{0, 1\}^k \times (\mathcal{N} \times \{0, 1\}^* \times \mathcal{X}) \rightarrow \{0, 1\}^k$ be a (keyed) hash function. Consider an adversary \mathcal{A} of q queries, with at most B queries per (user, nonce, AD, expansion). For any simulator Sim_b , we can build adversaries \mathcal{B}_1 and \mathcal{B}_2 and a simulator Sim_a such that*

$$\text{Adv}_{\text{HtE}[H, \text{SE}], \text{Sim}_a}^{\text{rae}}(\mathcal{A}) \leq \text{Adv}_H^{\text{prf}}(\mathcal{B}_1) + \text{Adv}_{\text{SE}, \text{Sim}_b}^{\text{rae}}(\mathcal{B}_2) .$$

The running time of Sim_a is about the same as that of Sim_b . Adversary \mathcal{B}_1 makes q queries. Its running time is about that of \mathcal{A} plus the cost of using SE to encrypt/decrypt \mathcal{A} 's queries. Adversary \mathcal{B}_2 has about the same running time as \mathcal{A} and also makes q queries of the total length as \mathcal{A} , but it makes only B queries per user.

Proof. We first construct the simulator Sim_a . It maintains a counter v for the current number of users and a state st^* for Sim_b , and lazily maintains a map Tbl to translate a tuple (i, N, A, τ) to a user u . On (i, N, A, C, τ, st) , it parses st into (v, st^*, Tbl) . It then checks if $\text{Tbl}[i, N, A, \tau]$ is defined. If not then it increments v and sets $\text{Tbl}[i, N, A, \tau] \leftarrow v$. In any case, it gets $u \leftarrow \text{Tbl}[i, N, A, \tau]$ and runs $(M, st^*) \leftarrow \text{Sim}_b(u, \varepsilon, \varepsilon, C, \tau, st^*)$. It then update its own state $st \leftarrow (v, \text{Tbl}, st^*)$, and returns (M, st) .

Consider the following sequence of games. Game G_0 corresponds to game $\mathbf{G}_{\text{HtE}[H, \text{SE}], \text{Sim}_a}^{\text{rae}}(\mathcal{A})$ with challenge bit 1. Game G_1 is identical to game G_0 , except that instead of using $H(K_i, \cdot, \cdot, \cdot)$ to derive the subkeys for user i , we use a truly random function $f_i : \mathcal{N} \times \{0, 1\}^* \times \mathcal{X} \rightarrow \{0, 1\}^k$. To bound the gap between G_0 and G_1 , we construct an adversary \mathcal{B}_1 attacking the (multi-user) PRF security of H . It runs \mathcal{A} and simulates game G_0 , but each call to $H(K_i, \cdot)$ is replaced by a corresponding call to $\text{EVAL}(i, \cdot)$. Then

$$\text{Adv}_H^{\text{prf}}(\mathcal{B}_1) = \Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})] .$$

Next, game G_2 corresponds to game $\mathbf{G}_{\text{HtE}[H, \text{SE}], \text{Sim}_a}^{\text{rae}}(\mathcal{A})$ with challenge bit 0. To bound the gap between G_1 and G_2 , we construct an adversary \mathcal{B}_2 attacking the (multi-user) RAE security of SE. It runs \mathcal{A} . For each encryption

query (i, N, A, M, τ) of \mathcal{A} , it calls $C \leftarrow \text{ENC}(u, \varepsilon, \varepsilon, M, \tau)$ for the effective user $u = (i, N, A, \tau)$, and returns C to \mathcal{A} . (This means \mathcal{B}_2 must lazily maintain a map from $\mathbb{N} \times \mathcal{N} \times \{0, 1\}^* \times \mathcal{X} \rightarrow \mathbb{N}$ to translate (i, N, A, τ) to an integer u .) Likewise, for each decryption query (i, N, A, C, τ) of \mathcal{A} , it returns $\text{DEC}(u, \varepsilon, \varepsilon, C, \tau)$ for the effective user $u = (i, N, A, \tau)$. Hence

$$\text{Adv}_{\text{SE}, \text{Sim}_b}^{\text{rae}}(\mathcal{B}_2) = \Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})] .$$

Summing up,

$$\begin{aligned} \text{Adv}_{\text{HttE}[H, \text{SE}], \text{Sim}_a}^{\text{rae}}(\mathcal{A}) &= \Pr[G_0(\mathcal{A})] - \Pr[G_2(\mathcal{A})] \\ &= (\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]) + (\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]) \\ &= \text{Adv}_H^{\text{prf}}(\mathcal{B}_1) + \text{Adv}_{\text{SE}, \text{Sim}_b}^{\text{rae}}(\mathcal{B}_2) . \end{aligned}$$

This concludes the proof. \square

INSTANTIATION. To have strong committing security, the key length k needs to be 256-bit. If the nonce length is fixed then one can instantiate $H(K, N, A, \tau)$ via $\text{SHA-512}(K \| N \| A \| [\tau]_{16})[1 : 256]$ or $\text{SHA-3}(K \| N \| A \| [\tau]_{16})[1 : 256]$, where $[\tau]_{16}$ is a 16-bit representation of the number of τ . We stress that one should avoid using SHA-256, because of the extension attack.

4 Fast Collision-Resistant PRF From Blockcipher

Recall that in CMT-1 security, we want to commit a short string (namely the key). This doesn't require a fully-fledged collision-resistant hash like SHA-512 or SHA-3. Instead, one can use cheaper constructions like Davies-Meyer applied to AES, as first suggested in the context of commitment security in [4]. Running Davies-Meyer on a blockcipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ however can only provide an n -bit commitment. In this section, we investigate how to produce a λ -bit commitment, where $n \leq \lambda \leq 2n$. Specifically, we want to build a collision-resistant PRF $H : \{0, 1\}^k \times \{0, 1\}^{n-1} \rightarrow \{0, 1\}^\lambda$ on top of E with $\lambda/2$ bits of security. Below, we show how to do that by "doubling" Davies-Meyer.

THE DOUBLE DAVIES-MEYER HASH. The code of the hash function $\text{DM2}[E, \lambda] : \{0, 1\}^k \times \{0, 1\}^{n-1} \rightarrow \{0, 1\}^\lambda$ is given in Fig. 10. Informally, to hash M with key K , we run two Davies-Meyer, one with $(K, M \| 0)$, and another with $(K, M \| 1)$, and then truncate the concatenated output. For $\lambda = n$, we recover the conventional Davies-Meyer construction.

COLLISION RESISTANCE OF DM2. The following result confirms that $\text{DM2}[E, \lambda]$ has $\lambda/2$ -bit collision resistance if E is modeled as an ideal cipher.

Proposition 3. *Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher that we will model as an ideal cipher. Let $n \leq \lambda \leq 2n$. Then for any adversary \mathcal{A} that makes at most q ideal-cipher queries,*

$$\text{Adv}_{\text{DM2}[E, \lambda]}^{\text{coll}}(\mathcal{A}) \leq \frac{8q^2}{2^\lambda} + \frac{2}{2^n} .$$

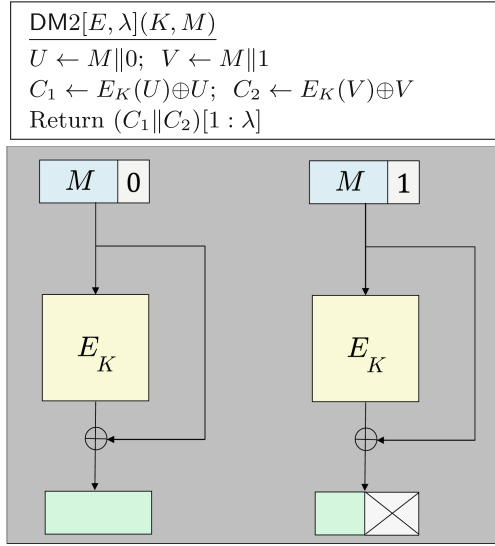


Fig. 10. The DM2 hashing.

Proof. Without loss of generality, assume that the adversary does not make redundant queries. That is, it does not repeat prior queries, and if it queries $C \leftarrow E_K(P)$ then later it will not query $E_K^{-1}(C)$, and if it queries $P \leftarrow E_K^{-1}(C)$ then it will not later query $E_K(P)$. For each query $C \leftarrow E(K, P)$ we store a log entry $(K, P, C \oplus P)$. Likewise, for each query $P \leftarrow E^{-1}(K, C)$, we store a log entry $(K, P, C \oplus P)$. For an n -bit string P , let \bar{P} denote the string obtained by flipping the last bit of P .

If a query results in a log entry (K, P, C) and there is no prior entry (K, \bar{P}, C^*) then we immediately grant the adversary a free query $E_K(\bar{P})$ and store the corresponding log entry. These free queries can only help the adversary. As a result, if we sort the log entries according to their querying order, then the i th query is the granted ones, for every even $i \leq 2q$.

Without loss of generality, assume that the right-hand side of the claimed bound is smaller than 1; otherwise the bound is moot. That is, $q \leq 2^{n-1}$. Thus for each entry (K, P, C) , conditioning on prior entries, the value X is uniformly chosen from a set of at least $2^n - q \geq 2^{n-1}$ members. Let $r = \lambda - n$.

Suppose that \mathcal{A} outputs (K_1, M_1, K_2, M_2) . Let Bad be the event that there are entries $(K_1, M_1 \parallel 0, X_1), (K_1, M_1 \parallel 1, X_1^*), (K_2, M_2 \parallel 0, X_2), (K_2, M_2 \parallel 1, X_2^*)$ in the logs. We now bound the advantage of the adversary depending on whether Bad happens.

IF Bad DOES NOT HAPPEN. If Bad does not happen, because of the symmetry and the way we grant free queries, without loss of generality, suppose that there is no entry $(K_1, M_1 \parallel 0, X_1)$. In that case, the chance that $E(K_1, M_1 \parallel 0) \oplus (M_1 \parallel 0) = E(K_2, M_2 \parallel 0) \oplus (M_2 \parallel 0)$ is at most 2^{1-n} .

IF Bad HAPPENS. By symmetry, without loss of generality, assume that among the four corresponding entries, $(K_1, M_1 \| 0, X_1)$ happens first (meaning that it is a non-granted query). For every (i, j) such that i is odd and $1 \leq i < j \leq 2q$, let $\text{Bad}_{i,j}$ be the event that the query of $(K_1, M_1 \| 0, X_1)$ is the i th query, and the first query of $(K_1, M_2 \| 0, X_2)$ and $(K_2, M_2 \| 1, X_2^*)$ is the j th query. Then

$$\text{Bad} = \bigcup \text{Bad}_{i,j} .$$

Note that there are at most

$$(2q - 1) + (2q - 3) + \cdots + 1 = q^2$$

pairs (i, j) . Fix one such pair. We now bound the adversary's advantage assuming that $\text{Bad}_{i,j}$ happens. We consider the following cases.

Case 1: The j th query creates the entry $(K_2, M_2 \| 0, X_2)$. Then $X_2 = X_1$ with probability at most 2^{1-n} . Moreover, conditioning on $X_2 = X_1$, because the entries $(K_1, M_1 \| 1, X_1^*)$ and $(K_2, M_2 \| 1, X_2^*)$ corresponds to the granted queries, the conditional probability that $X_2^*[1 : r] = X_1^*[1 : r]$ is at most 2^{1-r} . Summing up over at most q^2 pairs (i, j) , the chance that this case happens is at most $4q^2/2^{n+r} = 4q^2/2^\lambda$.

Case 2: The j th query creates the entry $(K_2, M_2 \| 1, X_2)$. Then the entries $(K_1, M_1 \| 1, X_1^*)$ and $(K_2, M_2 \| 0, X_2)$ corresponds to the granted queries. Thus the chance that $(X_1^*[1 : r], X_2) = (X_2^*[1 : r], X_1)$ is at most $4/2^\lambda$. Summing up over at most q^2 pairs (i, j) , the chance that this case happens is at most $4q^2/2^\lambda$.

WRAPPING UP. Summing up all cases,

$$\text{Adv}_{\text{DM2}[E,\lambda]}^{\text{coll}}(\mathcal{A}) \leq \frac{8q^2}{2^\lambda} + \frac{2}{2^n}$$

as claimed. □

PRF SECURITY OF DM2. The following result shows that if we model E as a good PRF then DM2 is also a good PRF.

Proposition 4. *Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher. Let $n \leq \lambda \leq 2n$. Then for any adversary \mathcal{A} that makes at most q queries, we can construct an adversary \mathcal{B} of about the same time and $2q$ queries such that*

$$\text{Adv}_{\text{DM2}[E,\lambda]}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_E^{\text{prf}}(\mathcal{B}) .$$

Proof. Without loss of generality, assume that \mathcal{A} does not repeat a prior query. Consider the following sequence of games. Game G_0 is game $\mathbf{G}_{\text{DM2}[E,\lambda]}^{\text{prf}}(\mathcal{A})$ with challenge bit 1. Game G_1 is identical to game G_0 , except that each call to $E(K_i, \cdot)$ is replaced with a corresponding call to a truly random function $f_i : \{0, 1\}^n \rightarrow \{0, 1\}^n$. To bound the gap between the two games, we construct an adversary \mathcal{B} attacking the (multi-user) PRF security of E as follows. It runs \mathcal{A} and simulates game G_0 , but each call to $E(K_i, \cdot)$ is replaced by a corresponding call to $\text{EVAL}(i, \cdot)$. Then

$$\text{Adv}_E^{\text{prf}}(\mathcal{B}) = \Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})] .$$

Let G_2 be game $\mathbf{G}_{\text{DM2}[E,\lambda]}^{\text{prf}}(\mathcal{A})$ with challenge bit 0. We now bound the gap between G_1 and G_2 for a computationally unbounded adversary \mathcal{A} . Without loss of generality, assume that \mathcal{A} is deterministic and never repeats a prior query. Note that in game G_1 , thanks to the domain separation in DM2, each f_i is never called on the same input twice, and thus effectively, in game G_1 , for each EVAL call, adversary \mathcal{A} receives a truly random answer. Likewise, in game G_2 , for each EVAL call, adversary \mathcal{A} receives a truly random answer. Hence

$$\Pr[G_1(\mathcal{A})] = \Pr[G_2(\mathcal{A})] .$$

Summing up,

$$\begin{aligned} \mathbf{Adv}_{\text{DM2}[E,\lambda]}^{\text{prf}}(\mathcal{A}) &= \Pr[G_0(\mathcal{A})] - \Pr[G_2(\mathcal{A})] \\ &= (\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]) + (\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]) \\ &= \mathbf{Adv}_E^{\text{prf}}(\mathcal{B}) . \end{aligned}$$

This concludes the proof. \square

5 Committing Concealer

In this section, we formalize a new primitive that we call a *committing concealer*. Functionality wise, a committing concealer $\mathbf{C} : \mathcal{K} \times \{0, 1\}^m \rightarrow \{0, 1\}^m$ is simply a blockcipher on $\{0, 1\}^m$, with \mathbf{C}^{-1} denoting its inverse. But traditionally a blockcipher is only secure if it is a strong PRP, but we will weaken this security goal in order to allow more efficient constructions. Looking ahead to our Feistel-based approach to committing concealers, we'll show a weaker security goal that allows us to get by with a two-round Feistel network, rather than the four rounds that would be required to achieve security as a strong PRP [21].

HIDING SECURITY. Our weaker security goal is what we call hiding security. It requires a committing concealer be a *one-time* strong PRP, meaning that an adversary is allowed only a single query per user. In particular, define the advantage of an adversary \mathcal{A} breaking the hiding security of \mathbf{C} as

$$\mathbf{Adv}_{\mathbf{C}}^{\text{hide}}(\mathcal{A}) = \Pr[\mathbf{G}_{\mathbf{C}}^{\text{hide}}(\mathcal{A})] ,$$

where game $\mathbf{G}_{\mathbf{C}}^{\text{hide}}(\mathcal{A})$ is defined in Fig. 11.

BINDING SECURITY. Let $s \leq m$ be an integer, and let $\{0, 1\}^{\leq m-s}$ denote the set of bit strings whose length is at most $m - s$. Let $\text{encode} : \{0, 1\}^{\leq m-s} \rightarrow \{0, 1\}^m$ be a function. Define the binding advantage of an adversary \mathcal{A} against \mathbf{C} with respect to encode as

$$\mathbf{Adv}_{\mathbf{C}, \text{encode}}^{\text{bind}}(\mathcal{A}) = \Pr[\mathbf{G}_{\mathbf{C}, \text{encode}}^{\text{bind}}(\mathcal{A})] ,$$

where game $\mathbf{G}_{\mathbf{C}, \text{encode}}^{\text{bind}}(\mathcal{A})$ is defined in Fig. 12. Informally, we want the ciphertext of \mathbf{C} to be a commitment of the key, if we restrict to messages in $\{\text{encode}(X) \mid X \in \{0, 1\}^{\leq m-s}\}$.

<p><u>Game $\mathbf{G}_C^{\text{hide}}(\mathcal{A})$</u> $v \leftarrow 0$; $b \leftarrow_s \{0, 1\}$; $Used \leftarrow \emptyset$ $b' \leftarrow_s \mathcal{A}^{\text{NEW, ENC, DEC}}$ Return $(b' = b)$</p> <p><u>NEW()</u> $v \leftarrow v + 1$; $K_v \leftarrow_s \{0, 1\}^k$</p>	<p><u>ENC(i, M)</u> If $i \notin \{1, \dots, v\} \setminus Used$ return \perp $C_1 \leftarrow C(K_i, M)$; $C_0 \leftarrow_s \{0, 1\}^{ M }$ $Used \leftarrow Used \cup \{i\}$; Return C_b</p> <p><u>DEC(i, C)</u> If $i \notin \{1, \dots, v\} \setminus Used$ return \perp $M_1 \leftarrow C^{-1}(K_i, C)$; $M_0 \leftarrow_s \{0, 1\}^{ M }$ $Used \leftarrow Used \cup \{i\}$; Return M_b</p>
--	--

Fig. 11. Game defining hiding security of C . The game maintains a set $Used$ of users that \mathcal{A} has queried.

<p><u>Game $\mathbf{G}_{C, \text{encode}}^{\text{bind}}(\mathcal{A})$</u> $(K_1, M_1, K_2, M_2) \leftarrow_s \mathcal{A}$ Return $(K_1 \neq K_2) \wedge (C(K_1, \text{encode}(M_1)) = C(K_2, \text{encode}(M_2)))$</p>

Fig. 12. Game defining binding security of C .

RELATION TO PRIOR WORK. Bellare and Hoang [5] recently consider how to add committing security to a standard AE scheme without expanding the ciphertext length. Their method requires that the scheme is tag-based, meaning that the ciphertext consists of a ciphertext core C^* and a tag T , and one can recover the message from just C^* (but of course without authenticity guarantees). As such, their method doesn't work for SIV constructions [26] (because the tag is needed for decryption), or EtE constructions (because there is no tag).

The work of Bellare and Hoang relies on an invertible PRF (IPF) that is also collision-resistant. Their construction of collision-resistant IPF encodes the message and then enciphers it with what is implicitly a committing concealer. In Sect. 5.2 we will study this committing concealer construction.

5.1 Committing Concealer from Ideal Cipher

As a warmup, we show that a blockcipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ can be used directly as a committing concealer if we model E as an ideal cipher. Still, if we want E to have strong binding security, the block length n needs to be at least 256 bits, meaning that we can't instantiate E from AES.

HIDING SECURITY. If E is modeled as a strong PRP then it obviously has good hiding security. We state the formal result for completeness.

Proposition 5. *Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher. Then for any adversary \mathcal{A} ,*

$$\mathbf{Adv}_E^{\text{hide}}(\mathcal{A}) \leq \mathbf{Adv}_E^{\pm\text{PRP}}(\mathcal{A}) .$$

BINDING SECURITY. The following result, from Bellare and Hoang [5], shows that if we model E as an ideal cipher then it also has good binding security, for any encoding mechanism. Due to the term $q^2/2^n$, if we aim for strong binding security, we need $n \geq 256$, meaning we need to instantiate E from, say Rijndael-256.

Proposition 6 ([5]). *Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher that we model as an ideal cipher. Let $\text{encode} : \{0, 1\}^{\leq n-s} \rightarrow \{0, 1\}^n$ be a function. Then for any adversary \mathcal{A} that makes at most q ideal-cipher queries,*

$$\text{Adv}_{E, \text{encode}}^{\text{bind}}(\mathcal{A}) \leq \frac{4q}{2^s} + \frac{2q^2}{2^n} .$$

5.2 Committing Concealer from Two-Round Feistel

In this section, we show how to build a committing concealer FF from a two-round Feistel network. The round functions are built on top of a blockcipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ that we can instantiate via AES.

THE FF CONSTRUCTION. Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher. Let $\ell < n$ be an integer. Define $\text{pad}(\cdot, 0) : \{0, 1\}^n \rightarrow \{0, 1\}^n$ via $\text{pad}(X, 0) = X[1 : n - 1]||0$, and define $\text{pad}(\cdot, 1) : \{0, 1\}^\ell \rightarrow \{0, 1\}^n$ via $\text{pad}(Y, 1) = Y||1^{n-\ell}$. Note that pad is a domain separation in the sense that $\text{pad}(X, 0) \neq \text{pad}(Y, 1)$ for any X, Y .

The committing concealer $\text{FF}[E, \ell]$ has message space $\{0, 1\}^{n+\ell}$. It is a two-round unbalanced Feistel network, where the left-hand side is n bits, and the right-hand side is ℓ bits. The round functions are based on the Davies-Meyer construction of E . See Fig. 13 for the code and also an illustration. This committing concealer is implicit in the recent work of Bellare and Hoang [5].

HIDING SECURITY OF FF. The following result shows that FF has good hiding security, assuming that E is a good PRF.

Proposition 7. *Let n, ℓ be integers such that $\ell < n$. Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher. Then for an adversary \mathcal{A} that makes at most q queries, we can construct an adversary \mathcal{B} of about the same time and $2q$ queries, with two queries per user, such that*

$$\text{Adv}_{\text{FF}[E, \ell]}^{\text{hide}}(\mathcal{A}) \leq \text{Adv}_E^{\text{prf}}(\mathcal{B}) .$$

Proof. Consider the following sequence of games. Game G_0 coincides with game $\mathbf{G}_C^{\text{hide}}(\mathcal{A})$ with challenge bit 1. Game G_1 is identical to game G_0 , except that each $E(K_i, \cdot)$ is replaced by a truly random function $f_i : \{0, 1\}^n \rightarrow \{0, 1\}^n$. To bound the gap between the two games, we construct an adversary \mathcal{B} attacking the (multi-user) PRF security of E as follows. It runs \mathcal{A} and simulates game G_0 , but each call to $E(K_i, \cdot)$ is replaced by a corresponding call to $\text{EVAL}(i, \cdot)$. Then

$$\text{Adv}_E^{\text{prf}}(\mathcal{B}) = \Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})] .$$

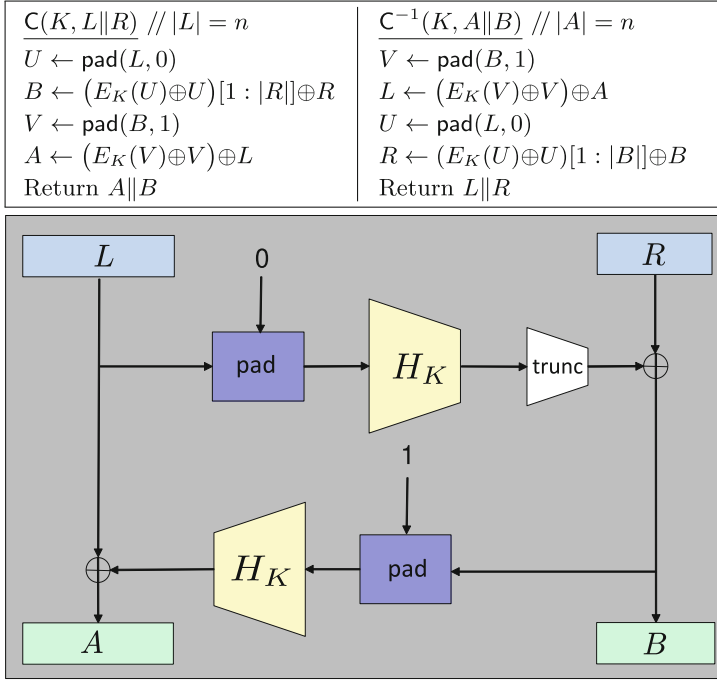


Fig. 13. The committing concealer $C = \text{FF}[E, \ell]$. In the illustration, the function H denotes the Davies-Meyer construction on E , meaning $H(K, M) = E_K(M) \oplus M$.

$\underline{\text{Game } G_2(\mathcal{A})}$ $v \leftarrow 0; \text{Used} \leftarrow \emptyset; b' \leftarrow \$_s \mathcal{A}^{\text{New, ENC, DEC}}$ Return $(b' = 1)$ $\underline{\text{ENC}(i, L\ R)}$ If $i \notin \{1, \dots, v\} \setminus \text{Used}$ return \perp $U \leftarrow \text{pad}(L, 0); X \leftarrow \$_s \{0, 1\}^n$ $B \leftarrow (X \oplus U)[1 : R] \oplus R$ $V \leftarrow \text{pad}(B, 1); Y \leftarrow \$_s \{0, 1\}^n$ $A \leftarrow (Y \oplus V) \oplus L$ $\text{Used} \leftarrow \text{Used} \cup \{i\};$ Return $A\ B$	$\underline{\text{NEW}()}$ $v \leftarrow v + 1$ $\underline{\text{DEC}(i, A\ B)}$ If $i \notin \{1, \dots, v\} \setminus \text{Used}$ return \perp $V \leftarrow \text{pad}(B, 1); Y \leftarrow \$_s \{0, 1\}^n$ $L \leftarrow (Y \oplus V) \oplus A$ $U \leftarrow \text{pad}(L, 0); X \leftarrow \$_s \{0, 1\}^n$ $R \leftarrow (X \oplus U)[1 : B] \oplus B$ $\text{Used} \leftarrow \text{Used} \cup \{i\};$ Return $L\ R$
---	--

Fig. 14. Game G_2 in the proof of Proposition 7.

Note that in game G_1 , each f_i is never run on the same input twice, thanks to the domain separation in FF and the requirement that the adversary can make a single query per user. Then we can rewrite game G_1 as game G_2 in Fig. 14, and the two games are equivalent. Note that effectively, in game G_2 , each ENC/DEC query returns a truly random answer. Thus game G_2 is equivalent to game

$\mathbf{G}_C^{\text{hide}}(\mathcal{A})$ with challenge bit 0. Summing up,

$$\begin{aligned} \mathbf{Adv}_{\text{FF}[E,\ell]}^{\text{hide}}(\mathcal{A}) &= \Pr[G_0(\mathcal{A})] - \Pr[G_2(\mathcal{A})] \\ &= (\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})]) + (\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})]) \\ &= \mathbf{Adv}_E^{\text{prf}}(\mathcal{B}) . \end{aligned}$$

This concludes the proof. □

BINDING SECURITY OF FF. Let $1 \leq t \leq n$. Let $\text{encode} : \{0, 1\}^{\leq n-t} \rightarrow \{0, 1\}^{n+\ell}$ be a function such that $\text{encode}(X)$ must end with $0^{\ell+1}$. The following result of Bellare and Hoang [5] shows that $\text{FF}[E, \ell]$ has about $(\ell - \log_2(n))$ bits of binding security in the ideal-cipher model.

Proposition 8 ([5]). *Let n, ℓ be integers such that $n \geq 32$ and $\ell < n$. Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher that we will model as an ideal cipher. Let encode be as above. Then for an adversary \mathcal{A} that makes at most q ideal-cipher queries,*

$$\mathbf{Adv}_{\text{FF}[E,\ell],\text{encode}}^{\text{bind}}(\mathcal{A}) \leq \frac{4(n + \ell)q + 5}{2^\ell} .$$

6 A Committing Transform for Wideblock TBC

Let TE be a wideblock TBC with message space $\{0, 1\}^{\geq k}$, key space $\{0, 1\}^k$, and tweak space \mathcal{T} . Our goal is to turn it into a wideblock TBC $\overline{\text{TE}}$ of the same tweak space and key space such that using $\overline{\text{TE}}$ in the EtE transform gives a scheme of both CMT-1 and RAE security. We achieve this via the Encipher-with-Concealer (EwC) transform below.

THE EwC TRANSFORM. Let C be a committing concealer of message space $\{0, 1\}^m$ and key space $\{0, 1\}^k$. Let $G : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ be a c -AXU hash function, with $n \leq k$. Let $H : \{0, 1\}^k \times \{0, 1\}^r \rightarrow \{0, 1\}^k$ be a collision-resistant PRF, with $r \leq n$. For an integer $i \in \{0, \dots, 2^r - 1\}$, let $[i]_r$ denote the r -bit representation of i . For two strings X and Y with $|X| \leq |Y|$, we write $X \oplus Y$ to denote $(X \parallel 0^{|Y|-|X|}) \oplus Y$. The code of the transform $\text{EwC}[\text{TE}, \text{C}, H, G]$ is given in Fig. 15. The scheme has message space $\{0, 1\}^{\geq m+k}$.

Informally, we use a four-round Feistel-like structure, where the right-hand side is m bits. Following Naor and Reingold [24], the first and last rounds are implemented via the AXU hash G (whose output is padded with 0's). In the second round, we encipher the intermediate left-hand side U via $V \leftarrow \text{TE.Enc}(L, T, U)$, where the subkey L is derived via $L \leftarrow H_K([2]_r)$. In the third round, we derive a one-time key $R \leftarrow H_I(U[1 : r] \oplus V[1 : r])$ for C , where $I \leftarrow H_K([3]_r)$, and use C to encipher the intermediate right-hand side.

Since we only need to use the r -bit prefix of the output of TE for C , on long messages, the evaluation of TE and C can be parallelized, for off-the-shelf

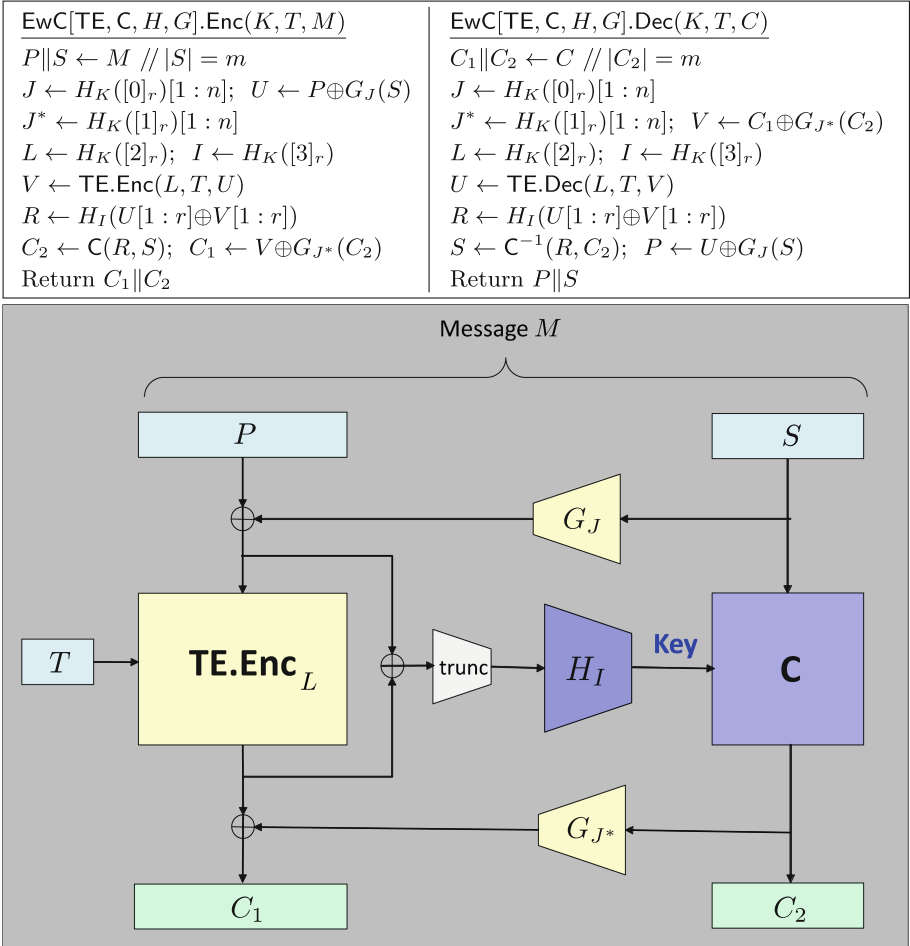


Fig. 15. The EwC transform. We omit the derivation of subkeys J, J^*, L, I in the illustration.

constructions of TE such as AEZ [18] or HCTR2 [14]. Conversely, on decryption, this allows fast rejection of invalid ciphertexts, which is important to resist denial-of-service attacks. The subkeys J, J^*, L, I can be cached and the cost of their derivation can be ignored if EwC is used in a stand-alone way. Still, if we compose it with the HtE transform in Sect. 3.2 then we have to account for this key derivation. However, in that case, the overhead of EwC is negligible compared to the hashing cost in HtE.

DISCUSSION. Structurally, EwC resembles the Hash-CTR-Hash (HCH) method [10] but there are nuances in the design. Here the message is split into two *uneven* halves, one of $|M| - O(1)$ bits, and the other just $O(1)$ bits. HCH runs the universal hash on the big half, meaning the hashing cost is $\Theta(|M|)$. In

contrast, EwC runs the universal hash on the small half, and thus the hashing cost is merely $O(1)$. On the other hand, while both have to encrypt $\Theta(|M|)$ bits with their base encryption schemes, EwC has to use the expensive TE but HCH only needs to run the cheap CTR.

INSTANTIATIONS. If we want to use AES, we can instantiate H from the DM2 construction in Sect. 4, with AES-256 as the underlying blockcipher, and instantiate G from GHASH or POLYVAL [17]. This means $k = 256$ and $r = 127$ and $n = 128$ and $c = 2$. The committing concealer C can be built from the FF construction in Sect. 5.2, again on AES-256. If we want to achieve around $\ell - 8$ bits of CMT-1 security, with $\ell < 128$, the input length of C should be $m = \ell + 128$. However, in EtE, if we want CMT-1 security, the minimum expansion must be $\ell + 1$ bits. Moreover, EwC only uses AES-256 in the forward direction. The overhead of EwC (assuming that the subkeys are cached) is four multiplications in $\text{GF}(2^{128})$ and four AES-256 calls plus an AES key setup. If AES-NI is available, the AES cost is approximately three sequential AES-256, because a good implementation can hide the key setup cost, and running two parallel AES calls in DM2 has the same cost as one.⁴

If we instead have a blockcipher of 256-bit block length, say Rijndael-256, the instantiation is much simpler. In particular, we can instantiate H from the Davies-Meyer construction of Rijndael-256, and C directly from Rijndael-256, meaning $k = m = r = 256$. Moreover, we can pick $n = 256$ and instantiate $G_J(X)$ as $X \times J$, where \times denotes the finite-field multiplication in $\text{GF}(2^{256})$, meaning $c = 1$. Thus the overhead of EwC in this case is two multiplications in $\text{GF}(2^{256})$ and two sequential Rijndael-256 calls plus the Rijndael-256 key setup cost (that can be hidden with a good implementation if AES-NI is available).⁵

For both instantiations, in EtE, we can pad with either 10^* or 0^* .

CMT-1 SECURITY OF EtE[EwC]. Suppose that C has good binding security with respect to an encoding $\text{encode} : \{0, 1\}^{\leq m-s} \rightarrow \{0, 1\}^m$. Define the following padding mechanism in EtE. If we have a message M and want to pad $\tau \geq s$ bits to it, we parse M to $M_1 \| M_2$, with $|M_2| = m - \tau$, and then output $M_1 \| \text{encode}(M_2)$. (We assume that one can efficiently recover the message X from $\text{encode}(X)$ and $|X|$, so that EtE is decryptable under the padding above.) For example, if $\text{encode}(X) = X \| 0^{m-|X|}$, the padding mechanism simply adds 0^τ to the message. The following result shows that EtE[EwC] has good CMT-1 security. Intuitively, we have a commitment chain $K \xrightarrow{H} I \xrightarrow{H} R \xrightarrow{C} C_2$, and thus C_2 is a commitment of K .

⁴ If we count the cost of subkey generation, we need six extra AES calls (instead of eight). In particular, since J and J^* are 128-bit long, each only needs one AES call (instead of two). These six AES are fully parallel, so running them costs as much as one AES call in AES-NI platforms.

⁵ If we count the cost of subkey generation, we need four extra (parallel) Rijndael-256 calls.

Theorem 9. Let $\overline{\text{TE}} = \text{EwC}[\text{TE}, \text{C}, H, G]$ be as above. For any adversary \mathcal{A} , we can construct adversaries \mathcal{B}_1 and \mathcal{B}_2 such that

$$\mathbf{Adv}_{\text{EtE}[\overline{\text{TE}}]}^{\text{cmt-1}}(\mathcal{A}) \leq \mathbf{Adv}_H^{\text{coll}}(\mathcal{B}_1) + \mathbf{Adv}_{\text{C,encode}}^{\text{bind}}(\mathcal{B}_2) .$$

Proof. We first construct adversary \mathcal{B}_1 . It runs

$$((K, N, A, P \| S, \tau), (K^*, N^*, A^*, P^* \| S^*, \tau^*)) \leftarrow_s \mathcal{A} .$$

It then runs $\text{EtE}[\overline{\text{TE}}]$ on $(K, N, A, P \| S, \tau)$ to obtain the subkeys I and $R \leftarrow H(I, X)$. It also runs $\text{EtE}[\overline{\text{TE}}]$ on $(K^*, N^*, A^*, P^* \| S^*, \tau^*)$ to obtain (I^*, R^*, X^*) . If $I = I^*$ then it outputs $((K, [3]_r), (K^*, [3]_r))$. Otherwise, it outputs $((I, X), (I^*, X^*))$.

Next, we construct \mathcal{B}_2 . It runs

$$((K, N, A, P \| S, \tau), (K^*, N^*, A^*, P^* \| S^*, \tau^*)) \leftarrow_s \mathcal{A} .$$

It then obtains subkeys R and R^* , and outputs $((R, S), (R^*, S^*))$.

For analysis, let $C_1 \| C_2$ and $C_1^* \| C_2^*$ be the corresponding ciphertexts. Suppose that \mathcal{A} can create a ciphertext collision, meaning $C_1 = C_1^*$ and $C_2 = C_2^*$ and $K \neq K^*$. If $R = R^*$ then \mathcal{B}_1 also creates a collision. If $R \neq R^*$ then \mathcal{B}_2 creates a collision. Hence

$$\mathbf{Adv}_{\text{EtE}[\overline{\text{TE}}]}^{\text{cmt-1}}(\mathcal{A}) \leq \mathbf{Adv}_H^{\text{coll}}(\mathcal{B}_1) + \mathbf{Adv}_{\text{C,encode}}^{\text{bind}}(\mathcal{B}_2) .$$

This concludes the proof. \square

STRONG TWEAKABLE-PRP SECURITY OF EwC. The following result shows that EwC is a strong tweakable-PRP. The proof is deferred to Sect. 8.

Theorem 10. Let $\overline{\text{TE}} = \text{EwC}[\text{TE}, \text{C}, H, G]$ be as above. For any adversary \mathcal{A} making q queries, with at most B queries per user, we can construct adversaries $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{B}_3 such that

$$\mathbf{Adv}_{\overline{\text{TE}}}^{\pm\text{PRP}}(\mathcal{A}) \leq 2\mathbf{Adv}_H^{\text{PRF}}(\mathcal{B}_1) + \mathbf{Adv}_{\overline{\text{TE}}}^{\pm\text{PRP}}(\mathcal{B}_2) + \mathbf{Adv}_{\text{C}}^{\text{hide}}(\mathcal{B}_3) + \frac{6cqB}{2^r} + \frac{6qB}{2^m} .$$

7 Performance

SCHEMES. We start from HCTR2 [14] as our baseline wideblock tweakable block cipher. We use HCTR2 on AES-256 since the Accordion call targets 256-bit (key-recovery) security, but note that one can use EwC on HCTR2-AES-128 as well. We implement EwC with the committing concealer C being the FF construction from Sect. 5.2, the hash H being the DM2 construction from Sect. 4, the AXU hash G being POLYVAL [17], and set $\ell = 120$.

EXPERIMENTAL SETUP. We evaluated the schemes on a range of message lengths from 64 bytes to 16384 bytes. For each message length, after warming up the




Fig. 16. Performance of EwC transform. The graph shows CPU cycles-per-byte (y-axis, lower is better) for encrypting messages of varying sizes (x-axis, in bytes) on a x86_64 processor.

operation 2048 invocations, we measured the elapsed time in cycles for 2048 invocations. We divided this elapsed time by the number of encrypted bytes to compute the number of *cycles per byte*. To minimize variance, we repeated this eight times, checked that the standard deviation across the repetitions was less than 0.05 cycles per byte, and took the mean.

Our benchmarking program used the implementation accompanying the paper [14] and was executed on an Intel i7-1360P, on a specified core running at 2.4 GHz with frequency scaling disabled.

RESULTS. The statistics is given in Fig. 16. Overall, the overhead is significant for small data, but becomes negligible for messages bigger than 1 KB.

8 Proof of Theorem 10

8.1 A Technique to Simplify Game-Based Proofs

Our proof relies on a novel use of the H-coefficient technique [12, 25] to simplify game-based proofs [8]. This technique is generic, and in this section, we will elaborate on its details.

THE SETTING. Suppose we have a construction based on an indistinguishability-based primitive Π , such as a PRF, an encryption scheme, or in our case, a committing concealer. The security notion of this construction involves bounding the gap between the real game G_0 (where Π is used) and an ideal game G_2 . The standard approach is to (i) define an intermediate game G_1 where Π is replaced by its ideal reference, (ii) give a reduction to bound the gap between G_0 and G_1 , and (iii) bound the (information-theoretic) gap between G_1 and G_2 . This approach doesn't work if in step (ii), the reduction only works as long as G_0 doesn't set a flag bad. For example, in our case, we have to derive the key for the committing concealer, and the key is uniformly random only when G_0 doesn't set bad.

To deal with the situation above, define G_1 as the analogue of G_0 where Π is replaced by its ideal reference; this means that G_1 also includes a flag bad.

The reduction still works as usual, but keeps track of the flag **bad**. If **bad** is set, then the reduction returns 1, suggesting that it's interacting with the real world. Otherwise, it follows the guess of the adversary.

In this case, the reduction advantage Δ doesn't bound $\Pr[G_0] - \Pr[G_1]$. Instead, $\Pr[G_0] - \Pr[G_1] \leq \Delta + \Pr[G_1 \text{ sets bad}]$. This means that we need to bound (1) the chance that G_1 sets **bad**, and also (2) the gap $\Pr[G_1] - \Pr[G_2]$. Here we only consider information-theoretic bounds, meaning the adversary is computationally unbounded, and thus can be assumed to be deterministic. There are many techniques [12, 15, 19, 25] for simplifying the analysis in (2), but none considers (1). In this section, we show how to extend the H-coefficient technique [12, 25] to *simultaneously* bound both (1) and (2). That is, not only can we simplify the analysis of (1), but we can kill two birds with one stone.

THE H-COEFFICIENT TECHNIQUE. Following [19], it is convenient to consider interactions of a distinguisher \mathcal{A} with an abstract system \mathbf{S} which answers \mathcal{A} 's queries. This system takes inputs and produces outputs, and is randomized and possibly stateful. The interaction between an adversary \mathcal{A} and a system \mathbf{S} defines a *transcript* $\theta = ((u_1, v_1), \dots, (u_q, v_q))$ containing the ordered sequence of query-answer pairs. Let $\text{ps}(\theta)$ be the probability that if the adversary queries u_1, \dots, u_q in that order, the answers will be v_1, \dots, v_q respectively.

In the H-coefficient technique, one wants to bound the distinguishing advantage of a real system \mathbf{S}_{real} and an ideal one $\mathbf{S}_{\text{ideal}}$. The adversary's interactions with those systems define transcripts $\mathcal{T}_{\text{real}}$ and $\mathcal{T}_{\text{ideal}}$, respectively. The following result bounds the distinguishing advantage of \mathcal{A} .

Lemma 11. [12, 25] *Suppose we can partition the set of valid transcripts for the ideal system into good and bad ones. Further, suppose that there exists a constant $\epsilon \geq 0$ such that $1 - \frac{\text{ps}_{\text{real}}(\theta)}{\text{ps}_{\text{ideal}}(\theta)} \leq \epsilon$ for every good transcript θ . Then, the advantage of \mathcal{A} in distinguishing \mathbf{S}_{real} and $\mathbf{S}_{\text{ideal}}$ is at most $\epsilon + \Pr[\mathcal{T}_{\text{ideal}} \text{ is bad}]$.*

APPLICATION TO OUR SETTING. Recall that we have two games G_1 and G_2 , and we need to bound (1) $\Pr[G_1 \text{ sets bad}]$ and (2) $\Pr[G_1] - \Pr[G_2]$. To use the H-coefficient technique, we view G_1 as the real system \mathbf{S}_{real} , and G_2 as the ideal system $\mathbf{S}_{\text{ideal}}$, and define what's meant for transcripts to be bad. Then one can use Lemma 11 to bound (2). Our key idea here is that it's often possible to extend the definition of bad transcripts (say adding some certain conditions, or revealing some extra information when the adversary finishes querying) so that if G_1 sets **bad** then the transcript must be bad. In that case, the following result shows that the *same* bound of Lemma 11 can be used to bound (1) as well. In fact, the actual bound is even slightly better.

Lemma 12. *Suppose we can partition the set of valid transcripts for the ideal system into good and bad ones such that if G_1 sets **bad** then the transcript must be bad. Further, suppose that there exists a constant $\epsilon \geq 0$ such that $1 - \frac{\text{ps}_{\text{real}}(\theta)}{\text{ps}_{\text{ideal}}(\theta)} \leq \epsilon$ for every good transcript θ . Then*

$$\Pr[G_1 \text{ sets bad}] \leq \epsilon + (1 - \epsilon) \Pr[\mathcal{T}_{\text{ideal}} \text{ is bad}] .$$

Proof. Since G_1 's setting `bad` will lead to a bad transcript,

$$\Pr[G_1 \text{ sets bad}] \leq \Pr[\mathcal{T}_{\text{real}} \text{ is bad}] \leq 1 - \Pr[\mathcal{T}_{\text{real}} \text{ is good}] . \tag{1}$$

Recall that for any good transcript θ ,

$$\text{ps}_{\text{real}}(\theta) \geq (1 - \epsilon)\text{ps}_{\text{ideal}}(\theta) .$$

Summing this over all good transcripts,

$$\Pr[\mathcal{T}_{\text{real}} \text{ is good}] \geq (1 - \epsilon) \Pr[\mathcal{T}_{\text{ideal}} \text{ is good}] . \tag{2}$$

By combining Eq. (1) and Eq. (2), we obtain

$$\begin{aligned} \Pr[G_1 \text{ sets bad}] &\leq 1 - (1 - \epsilon) \Pr[\mathcal{T}_{\text{ideal}} \text{ is good}] \\ &= 1 - (1 - \epsilon)(1 - \Pr[\mathcal{T}_{\text{ideal}} \text{ is bad}]) \\ &= \epsilon + (1 - \epsilon) \Pr[\mathcal{T}_{\text{ideal}} \text{ is bad}] . \end{aligned}$$

This concludes the proof. □

DISCUSSION. Let Bad_{real} be the event that G_1 sets `bad`. Our approach requires extending the definition of bad transcripts so that if Bad_{real} happens, the resulting transcript will be bad. Effectively, this requires bounding the probability of an extra event $\text{Bad}_{\text{ideal}}$ in the *ideal* system when we deal with $\Pr[\mathcal{T}_{\text{ideal}} \text{ is bad}]$. This doesn't mean we gain nothing, because events in the ideal system are often much easier to analyze.

8.2 Proof of Theorem 10

Without loss of generality, assume that the adversary doesn't repeat prior queries. Moreover, if it queries $C \leftarrow \text{ENC}(i, T, M)$ then it won't later query $\text{DEC}(i, T, C)$, and vice versa.

Consider the following sequence of games. Game G_0 coincides to game $\mathbf{G}_{\text{TE}}^{\pm\text{prp}}(\mathcal{A})$ with challenge bit 1. Game G_1 is identical to game G_0 , except that each call to $H(K_i, \cdot)$ is replaced by a corresponding call to a truly random function $f_i : \{0, 1\}^r \rightarrow \{0, 1\}^k$. To bound the gap between the two games, we construct an adversary \mathcal{D}_1 attacking the (multi-user) PRF security of H as follows. It runs \mathcal{A} and simulates game G_0 , but each call to $H(K_i, \cdot)$ is replaced by a corresponding call to $\text{EVAL}(i, \cdot)$. Then

$$\text{Adv}_H^{\text{prf}}(\mathcal{D}_1) = \Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})] .$$

Then in game G_1 , the subkeys J_i, J_i^*, L_i, I_i will be uniformly random. In game G_2 , instead of using $H(I_i, \cdot)$, we use truly random functions $g_i : \{0, 1\}^r \rightarrow \{0, 1\}^k$. To bound the gap between G_2 and G_1 , we construct an adversary \mathcal{D}_2 attacking the (multi-user) security of H as follows. It runs \mathcal{A} and simulates game

<p>Game $G_4(\mathcal{A}), G_5(\mathcal{A})$</p> <p>$v \leftarrow 0; b' \leftarrow \\$_\mathcal{A}^{\text{NEW, ENC, DEC}}$ Return ($b' = 1$)</p> <hr/> <p>$\text{ENC}(i, T, P \ S)$</p> <p>$U \leftarrow P \oplus G(J_i, S); V \leftarrow \Pi_{i, T}(U)$ $X \leftarrow U[1 : r] \oplus V[1 : r]; R \leftarrow H(I_i, X)$ If $X \in \text{Dom}_i$ then $\text{bad} \leftarrow \text{true}$ $\text{Dom}_i \leftarrow \text{Dom}_i \cup \{X\}$ $C_2 \leftarrow \mathcal{C}(R, S); C_2 \leftarrow \\$_\mathcal{A}^{\{0, 1\}^m}$ $C_1 \leftarrow V \oplus G(J_i^*, C_2)$ Return $C_1 \ C_2$</p>	<p>$\text{NEW}()$</p> <p>$v \leftarrow v + 1; \text{Dom}_v \leftarrow \emptyset$ $J_v, J_v^* \leftarrow \\$_\mathcal{A}^{\{0, 1\}^n}; I_v \leftarrow \\$_\mathcal{A}^{\{0, 1\}^k}$ For $T \in \mathcal{T}$ do $\Pi_{v, T} \leftarrow \\$_\mathcal{A}^{\text{LP}(\mathcal{M})}$</p> <hr/> <p>$\text{DEC}(i, T, C_1 \ C_2)$</p> <p>$V \leftarrow C_1 \oplus G(J_i^*, C_2); U \leftarrow \Pi_{i, T}^{-1}(V)$ $X \leftarrow U[1 : r] \oplus V[1 : r]; R \leftarrow H(I_i, X)$ If $X \in \text{Dom}_i$ then $\text{bad} \leftarrow \text{true}$ $\text{Dom}_i \leftarrow \text{Dom}_i \cup \{X\}$ $S \leftarrow \mathcal{C}^{-1}(R, C_2); S \leftarrow \\$_\mathcal{A}^{\{0, 1\}^m}$ $P \leftarrow U \oplus G(J_i, S)$ Return $P \ S$</p>
--	--

Fig. 17. Games G_4 and G_5 in the proof of Theorem 10. Game G_5 contains the highlighted code but game G_4 does not.

G_1 , but each call to $H(I_i, \cdot)$ is replaced by a corresponding call to $\text{EVAL}(i, \cdot)$. Then

$$\mathbf{Adv}_H^{\text{prf}}(\mathcal{D}_2) = \Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})] .$$

So far we have two adversaries attacking the PRF security of H . We can unify them to an adversary \mathcal{B}_1 as follow: adversary \mathcal{B}_1 picks a coin $b \leftarrow \$_\mathcal{A}^{\{0, 1\}}$, and runs \mathcal{D}_b . Then

$$\mathbf{Adv}_H^{\text{prf}}(\mathcal{B}_1) = \frac{1}{2} \mathbf{Adv}_H^{\text{prf}}(\mathcal{D}_1) + \frac{1}{2} \mathbf{Adv}_H^{\text{prf}}(\mathcal{D}_2) ,$$

and thus

$$\Pr[G_0(\mathcal{A})] - \Pr[G_2(\mathcal{A})] = 2 \cdot \mathbf{Adv}_H^{\text{prf}}(\mathcal{B}_1) .$$

Let \mathcal{M} be the domain of TE , and let $\text{LP}(\mathcal{M})$ denote the set of permutations on \mathcal{M} that are length-preserving, meaning $|\pi(M)| = |M|$ for every $M \in \mathcal{M}$. Game G_3 is identical to game G_2 , but calls to $\text{TE.Enc}(L_i, T, \cdot)$ and $\text{TE.Dec}(L_i, T, \cdot)$ are replaced by corresponding calls to $\Pi_{i, T} \leftarrow \$_\mathcal{A}^{\text{LP}(\mathcal{M})}$ and its inverse. To bound the gap between the two games, we construct an adversary \mathcal{B}_2 attacking the (multi-user) strong tweakable-PRP security of TE as follows. It runs \mathcal{A} and simulates game G_2 , but calls to calls to $\text{TE.Enc}(L_i, T, \cdot)$ and $\text{TE.Dec}(L_i, T, \cdot)$ are replaced by corresponding calls to $\text{ENC}(i, T, \cdot)$ and $\text{DEC}(i, T, \cdot)$. Then

$$\mathbf{Adv}_{\text{TE}}^{\pm \widetilde{\text{PRP}}}(\mathcal{B}_2) = \Pr[G_2(\mathcal{A})] - \Pr[G_3(\mathcal{A})] .$$

Game G_4 is specified in Fig. 17. It is the same as G_3 with some bookkeeping, and thus

$$\Pr[G_4(\mathcal{A})] = \Pr[G_3(\mathcal{A})] .$$

We are now in the setting of Sect. 8.1, as the keys for \mathcal{C} are uniformly only when G_4 doesn't set bad . Let G_5 be identical to game G_4 , except that the answers

of C and C^{-1} are replaced by uniformly random strings. The code of game G_5 is specified in Fig. 17. To bound the gap between the two games, we construct an adversary \mathcal{B}_3 attacking the hiding security of C as follows. It runs \mathcal{A} and simulates game G_4 . For each call to $C(R, \cdot)$, it creates a new user u and makes a corresponding call to $\text{ENC}(u, \cdot)$. Likewise, for each call to $C^{-1}(R, \cdot)$, it creates a new user u^* and makes a corresponding call to $\text{DEC}(u^*, \cdot)$. If the simulated game sets **bad** then \mathcal{B}_3 returns 1, indicating that it's in the real world. Otherwise, it returns the same guess as \mathcal{A} .

Let d be the challenge bit of game $\mathbf{G}_C^{\text{hide}}(\mathcal{B}_3)$. Then on the one hand,

$$\Pr[\mathbf{G}_C^{\text{hide}}(\mathcal{B}_3) \Rightarrow \text{true} \mid d = 1] \geq \Pr[G_4(\mathcal{A})] .$$

On the other hand,

$$\Pr[\mathbf{G}_C^{\text{hide}}(\mathcal{B}_3) \Rightarrow \text{false} \mid d = 0] \leq \Pr[G_5(\mathcal{A})] + \Pr[G_5(\mathcal{A}) \text{ sets bad}] .$$

Subtracting, we obtain

$$\Pr[G_4(\mathcal{A})] - \Pr[G_5(\mathcal{A})] \leq \mathbf{Adv}_C^{\text{hide}}(\mathcal{B}_3) + \Pr[G_5(\mathcal{A}) \text{ sets bad}] .$$

Let G_6 be game $\mathbf{G}_{\text{TE}}^{\pm\text{prp}}(\mathcal{A})$ with challenge bit 0. Using the technique in Sect. 8.1, we obtain the following result; the proof is given in the full version of the paper.

Lemma 13. *Let G_5 and G_6 be as above. Then*

$$\Pr[G_5(\mathcal{A}) \text{ sets bad}] + (\Pr[G_5(\mathcal{A})] - \Pr[G_6(\mathcal{A})]) \leq \frac{2qB}{2^m} + \frac{6cqB}{2^r} .$$

Summing up,

$$\begin{aligned} \mathbf{Adv}_{\text{TE}}^{\pm\text{prp}}(\mathcal{A}) &= \Pr[G_0(\mathcal{A})] - \Pr[G_6(\mathcal{A})] \\ &= \sum_{i=0}^5 \Pr[G_i(\mathcal{A})] - \Pr[G_{i+1}(\mathcal{A})] \\ &\leq 2\mathbf{Adv}_H^{\text{prf}}(\mathcal{B}_1) + \mathbf{Adv}_{\text{TE}}^{\pm\text{prp}}(\mathcal{B}_2) + \mathbf{Adv}_C^{\text{hide}}(\mathcal{B}_3) + \frac{6cqB}{2^r} + \frac{2qB}{2^m} . \end{aligned}$$

Acknowledgement. We thank the Asiacrypt 2024 reviewers for their valuable comments. Viet Tung Hoang was supported in part by NSF grant CNS-2046540 (CAREER).

References

1. A. Albertini, T. Duong, S. Gueron, S. Kölbl, A. Luykx, and S. Schmiege. How to abuse and fix authenticated encryption without key commitment. In K. R. B. Butler and K. Thomas, editors, *USENIX Security 2022*, pages 3291–3308. USENIX Association, Aug. 2022.

2. E. Andreeva, A. Bogdanov, A. Luykx, B. Mennink, N. Mouha, and K. Yasuda. How to securely release unverified plaintext in authenticated encryption. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 105–125. Springer, Heidelberg, Dec. 2014.
3. E. Andreeva, A. Bogdanov, A. Luykx, B. Mennink, E. Tischhauser, and K. Yasuda. Parallelizable and authenticated online ciphers. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 424–443. Springer, Heidelberg, Dec. 2013.
4. M. Bellare and V. T. Hoang. Efficient schemes for committing authenticated encryption. In O. Dunkelman and S. Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 845–875. Springer, Heidelberg, May/June 2022.
5. M. Bellare and V. T. Hoang. Succinctly-committing authenticated encryption. In *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, 2024*.
6. M. Bellare, T. Ristenpart, P. Rogaway, and T. Stegers. Format-preserving encryption. In M. J. Jacobson Jr., V. Rijmen, and R. Safavi-Naini, editors, *SAC 2009*, volume 5867 of *LNCS*, pages 295–312. Springer, Heidelberg, Aug. 2009.
7. M. Bellare and P. Rogaway. Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography. In T. Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 317–330. Springer, Heidelberg, Dec. 2000.
8. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
9. P. Campbell. GLEVIAN and VIGORNIAN: Robust beyond-birthday aead modes. Cryptology ePrint Archive, Paper 2023/1379, 2023. <https://eprint.iacr.org/2023/1379>.
10. D. Chakraborty and P. Sarkar. HCH: A new tweakable enciphering scheme using the hash-encrypt-hash approach. In R. Barua and T. Lange, editors, *INDOCRYPT 2006*, volume 4329 of *LNCS*, pages 287–302. Springer, Heidelberg, Dec. 2006.
11. J. Chan and P. Rogaway. On committing authenticated-encryption. In V. Atluri, R. Di Pietro, C. D. Jensen, and W. Meng, editors, *ESORICS 2022, Part II*, volume 13555 of *LNCS*, pages 275–294. Springer, Heidelberg, Sept. 2022.
12. S. Chen and J. P. Steinberger. Tight security bounds for key-alternating ciphers. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, Heidelberg, May 2014.
13. Y. L. Chen, A. Flórez-Gutiérrez, A. Inoue, R. Ito, T. Iwata, K. Minematsu, N. Mouha, Y. Naito, F. Sibleyras, and Y. Todo. Key committing security of AEZ. The Third NIST Workshop on Block Cipher Modes of Operation, 2023.
14. P. Crowley, N. Huckleberry, and E. Biggers. Length-preserving encryption with HCTR2. Cryptology ePrint Archive, Paper 2021/1441, 2021. <https://eprint.iacr.org/2021/1441>.
15. W. Dai, V. T. Hoang, and S. Tessaro. Information-theoretic indistinguishability via the chi-squared method. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 497–523. Springer, Heidelberg, Aug. 2017.
16. P. Grubbs, J. Lu, and T. Ristenpart. Message franking via committing authenticated encryption. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 66–97. Springer, Heidelberg, Aug. 2017.

17. S. Gueron, A. Langley, and Y. Lindell. AES-GCM-SIV: specification and analysis. *IACR Cryptol. ePrint Arch.*, 2017.
18. V. T. Hoang, T. Krovetz, and P. Rogaway. Robust authenticated-encryption AEZ and the problem that it solves. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 15–44. Springer, Heidelberg, Apr. 2015.
19. V. T. Hoang and S. Tessaro. Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 3–32. Springer, Heidelberg, Aug. 2016.
20. J. Len, P. Grubbs, and T. Ristenpart. Partitioning oracle attacks. In M. Bailey and R. Greenstadt, editors, *USENIX Security 2021*, pages 195–212. USENIX Association, Aug. 2021.
21. M. Luby and C. Rackoff. How to construct pseudo-random permutations from pseudo-random functions (abstract). In H. C. Williams, editor, *CRYPTO’85*, volume 218 of *LNCS*, page 447. Springer, Heidelberg, Aug. 1986.
22. Z. Luo, Y. Jia, Y. Shen, and A. Kate. Proxying is enough: Security of proxying in TLS oracles and AEAD context unforgeability. *Cryptology ePrint Archive*, Paper 2024/733, 2024. <https://eprint.iacr.org/2024/733>.
23. C. Namprempre, P. Rogaway, and T. Shrimpton. Reconsidering generic composition. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, Heidelberg, May 2014.
24. M. Naor and O. Reingold. On the construction of pseudorandom permutations: Luby-Rackoff revisited. *Journal of Cryptology*, 12(1):29–66, Jan. 1999.
25. J. Patarin. The “coefficients H” technique (invited talk). In R. M. Avanzi, L. Keliher, and F. Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, Heidelberg, Aug. 2009.
26. P. Rogaway and T. Shrimpton. A provable-security treatment of the key-wrap problem. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 373–390. Springer, Heidelberg, May/June 2006.

Author Index

A

Ananth, Prabhajan 129

B

Bhaumik, Ritam 215

C

Choi, Wonseok 251

Cogliati, Benoît 215

Cojocaru, Alexandru 65, 158

E

Ethan, Jordan 215

F

Flórez-Gutiérrez, Antonio 280

G

Garay, Juan 65, 158

Goyal, Vipul 129

Grassi, Lorenzo 280

H

Hoang, Viet Tung 312

J

Jawale, Ruta 94

Jha, Ashwin 215

K

Khurana, Dakshita 94

L

Leander, Gregor 280

Lee, Jooyoung 251

Lee, Yeongmin 251

Liu, Jiahui 129

Liu, Qipeng 129, 158

M

Menda, Sanketh 312

Montgomery, Hart 33

Morimae, Tomoyuki 3

P

Phan, Duong Hieu 185

S

Sharif, Shahed 33

Sibleyras, Ferdinand 280

Song, Fang 65, 158

T

Todo, Yosuke 280

W

Wen, Weiqiang 185

Y

Yamada, Shogo 3

Yamakawa, Takashi 3

Yan, Xingyu 185

Z

Zheng, Jinwei 185