20th Conference on the Theory of Quantum Computation, Communication and Cryptography

TQC 2025, September 15–19, 2025, Indian Institute of Science, Bengaluru, India

Edited by

Bill Fefferman

LIPICS

Editors

Bill Fefferman
University of Chicago, IL, USA
wjf@uchicago.edu

ACM Classification 2012
Theory of computation → Quantum computation theory; Theory of computation → Quantum complexity theory; Theory of computation → Quantum information theory; Theory of computation → Cryptographic protocols; Theory of computation → Cryptographic primitives; Hardware → Quantum computation; Computer systems organization → Quantum computing; Hardware → Quantum error correction and fault tolerance

ISBN 978-3-95977-392-8

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at https://www.dagstuhl.de/dagpub/978-3-95977-392-8.

Publication date
September, 2025

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists all publications of this volume in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at https://portal.dnb.de.

Digital Object Identifier: 10.4230/LIPIcs.TQC.2025.0

ISBN 978-3-95977-392-8 **ISSN 1868-8969** **https://www.dagstuhl.de/lipics**

LIPIcs – Leibniz International Proceedings in Informatics

ISSN 1868-8969

https://www.dagstuhl.de/lipics

Contents

Regular Papers

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman
LIPICS Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Preface

The 20th Conference on The Theory of Quantum Computation, Communication and Cryptography (TQC) was hosted by the Indian Institute of Science, Bengaluru in India, and held from September 15 to September 19, 2025. The TQC conference is a leading annual international conference for students and researchers working in the theoretical aspects of quantum information science. The scientific objective of TQC is to bring together the theoretical quantum information science community to present and discuss the latest advances in the field.

Areas of interest for TQC include, but are not restricted to: quantum algorithms, models of quantum computation, quantum complexity theory, the simulation of quantum systems, quantum cryptography, quantum communication, quantum information theory, quantum estimation and measurement, quantum error correction and fault-tolerant quantum computing, the intersection of quantum information and condensed-matter theory, and the intersection of quantum information and machine learning.

A list of the previous editions of TQC follows:

- TQC 2024, Okinawa Institute for Science and Technology, Japan
- TQC 2023, University of Aveiro, Portugal
- TQC 2022, University of Illinois at Urbana-Champaign, USA
- TQC 2021, University of Latvia, Latvia (virtual conference)
- TQC 2020, University of Latvia, Latvia (virtual conference)
- TQC 2019, University of Maryland, USA
- TQC 2018, University of Technology Sydney, Australia
- TQC 2017, Université Pierre et Marie Curie, France
- TQC 2016, Freie Universität Berlin, Germany
- TQC 2015, Université libre de Bruxelles, Brussels, Belgium
- TQC 2014, National University of Singapore, Singapore
- TQC 2013, University of Guelph, Canada
- TQC 2012, University of Tokyo, Japan
- TQC 2011, Universidad Complutense de Madrid, Spain
- TQC 2010, University of Leeds, UK
- TQC 2009, Institute for Quantum Computing, University of Waterloo, Canada
- TQC 2008, University of Tokyo, Japan
- TQC 2007, Nara Institute of Science and Technology, Nara, Japan
- TQC 2006, NTT R&D Center, Atsugi, Kanagawa, Japan

The conference consisted of invited talks, contributed talks and a poster session. The invited talks were given by André Chailloux (French Institute for Research in Computer Science and Automation), Stacey Jeffery (CWI), Rajendra Kumar (IIT Delhi), and Hayata Yamasaki (University of Tokyo).

Submissions were solicited for two tracks: With Proceedings (talk and proceedings) and Without Proceedings (talk only). There were 375 submissions. The program committee selected 77 submissions for talks, including 12 to be published in the With Proceedings track. I wish to thank the members of the Program Committee and all subreviewers for their incredible work which helped to form an excellent program. Also I wish to thank the Local Organizing Committee for all their efforts in organizing the conference and the Steering Committee for their guidance, as well as for maintaining the conference's high standards. Last but not least, I thank the authors of all the TQC 2025 submissions.

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman
LIPICS Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Conference Organization

Organizing Committee

Local organizers in Bengaluru

- C. M. Chandrashekar, IISc, Bengaluru
- Baladitya Suri, IISc, Bengaluru
- Navin Kashyap, IISc, Bengaluru
- Henry Sukumar, C-DAC, Bengaluru
- Krishnakumar Sabapathy, Fujitsu Research, Bengaluru

National organizers

- Anirban Pathak, JIIT, Noida
- Prabha Mandayam, IIT Madras
- Sudhir Kamath, DRDO-IARCOE
- Manik Banik, SN Bose National Centre for Basic Sciences, Kolkata
- Saibal K. Pal , SAG-DRDO Delhi

International organizers

- Lídia del Rio, Squids and University of Zurich
- Nuriya Nurgalieva, Squids and University of Zurich

Program Committee

- Bill Fefferman (University of Chicago, Program Chair)
- Simon Apers (Université Paris Cité)
- Aleksandrs Belovs (University of Latvia)
- Adam Bene Watts (University of Calgary)
- Michael Beverland (IBM Quantum)
- Nikolas Breuckmann (University of Bristol)
- Anne Broadbent (University of Ottawa)
- Anthony Chen (Simons Institute, Berkeley)
- Matthias Christandl (University of Copenhagen)
- Andrea Coladangelo (University of Washington)
- Alexander Dalzell (Amazon)
- Abhinav Deshpande (IBM Quantum)
- Di Fang (Duke University)
- Omar Fawzi (Inria and École Normale Supérieure de Lyon)
- Sevag Gharibian (Paderborn University)
- Daniel Grier (University of California San Diego)
- Vojtech Havlicek (IBM Quantum)
- Jiaqing Jiang (Caltech)
- Tamara Kohler (Stanford)
- Aleksander Kubica (Yale)
- Richard Kueng (Johannes Kepler University Linz)
- Srijita Kundu (University of Waterloo)

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman
LIPICS Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- Ludovico Lami (Scuola Normale Superiore, Pisa)
- Cécilia Lancien (Institut Fourier Grenoble and CNRS)
- Nicholas Laracuente (Indiana University Bloomington)
- Felix Leditzky (University of Illinois Urbana-Champaign)
- Troy Lee (University of Technology Sydney)
- Jiaqi Leng (Simons Institute, Berkeley)
- Jiahui Liu (Fujitsu Research)
- Saeed Mehraban (Tufts University)
- Tomoyuki Morimae (Yukawa Institute for Theoretical Physics, Kyoto University)
- Hui Khoon Ng (National University of Singapore)
- Nelly Ng (Nanyang Technological University)
- Harumichi Nishimura (Nagoya University)
- Changhun Oh (Korea Advanced Institute of Science and Technology)
- Māris Ozols (University of Amsterdam, QuSoft)
- David Pérez-García (Universidad Complutense de Madrid)
- Alexander Poremba (MIT)
- Luke Schaeffer (University of Waterloo)
- Norbert Schuch (University of Vienna)
- Thomas Schuster (Caltech)
- Christian Schaffner (University of Amsterdam, QuSoft)
- Kunal Sharma (IBM Quantum)
- Makrand Sinha (University of Illinois Urbana-Champaign)
- Graeme Smith (University of Waterloo)
- Mehdi Soleimanifar (Caltech)
- Fang Song (Portland State University)
- Daniel Stilck França (University of Copenhagen)
- Sergii Strelchuk (Oxford University)
- Marco Tomamichel (Centre for Quantum Technologies, National University of Singapore)
- Dave Touchette (Université de Sherbrooke)
- Benjamin Villalonga (Google)
- Dominic Williamson (University of Sydney)
- Freek Witteveen (University of Copenhagen and QuSoft)
- Takashi Yamakawa (NTT Social Informatics Laboratories)
- Jiayu Zhang (Zhongguancun Laboratory, Beijing)
- Ruizhe Zhang (Simons Institute, Berkeley)
- Sisi Zhou (Perimeter Institute)

Steering Committee

- Kai-Min Chung, Academia Sinica
- Steve Flammia, AWS Center for Quantum Computing
- Min-Hsiu Hsieh, Hon Hai (Foxconn)
- Shelby Kimmel, Middlebury College
- François Le Gall, Nagoya University (chair)
- Frederic Magniez, CNRS (co-chair)
- Kae Nemoto, OIST
- Lídia del Rio, Squids and University of Zurich

Subreviewers

- Tejas Acharya
- Avantika Agarwal
- Francesco Albarelli
- Richard Allen
- Prabhanjan Ananth
- Jonas Anderson
- Eric Anschuetz
- Harriet Apel
- Roy Araiza
- Mateus Araujo
- Srinivasan Arunachalam
- Vahid Asadi
- Nikita Astrakhantsev
- Brandon Augustino
- Kaniuar Bacho
- Andrew Baczewski
- Joonwoo Bae
- Mohsen Bagherimehrab
- Akshay Bansal
- Nouédyn Baspin
- Joao Basso
- Jessica Bavaresco
- Emily Beatty
- Niel de Beaudrap
- Jacob Beckey
- Daniel Belkin
- Maxim van den Berg
- Thiago Bergamaschi
- Pablo Bermejo
- Kishor Bharti
- Archishna Bhattacharyya
- Andreas Bluhm
- Anselm Blumer
- Pablo Bonilla
- John Bostanci
- Sami Boulebnane
- Arthur Braida
- Lukas Brenner
- Adam Burchardt
- Maddie Cain
- Alper Cakan
- John Calsamiglia
- Charles Cao
- Ningping Cao
- Matthias C. Caro
- Joseph Carolan
- Enrique Cervero
- Ulysse Chabaud
- Shouvanik Chakrabarti
- Shantanav Chakraborty
- Rohit Chatterjee
- Boyang Chen
- Jielun Chen
- Senrui Chen
- Xinan Chen
- Zherui Chen
- Bin Cheng
- Jinglei Cheng
- Nai-Hui Chia
- Chi-Ning Chou
- Cristina Cirstoiu
- Baptiste Claudon
- Richard Cleve
- Alexandru Cojocaru
- Arjan Cornelissen
- Eleanor Crane
- Laura Cui
- Eric Culf
- Jakub Czartowski
- Shaun Datta
- Idris Delsol
- Zhiyan Ding
- Jordan Docter
- Yangjing Dong
- João Doriguello
- Arpit Dua
- Benoît Dubus
- Arkopal Dutt
- David Elkouss
- Alex Essery
- Kun Fang
- Cameron Foreman
- Honghao Fu
- Francois Le Gall
- Thomas Galley
- Ray Ganardi
- Tuvia Gefen
- Ian George
- Alexandru Gheorghiu
- Soumik Ghosh
- Amin Shiraz Gilani
- Andras Gilyen

- Filippo Girardi
- Matthew L. Goh
- Eli Goldin
- Louis Golowich
- Weiyuan Gong
- Guillermo González-garcía
- David Gosset
- Ashutosh Goswami
- Gilad Gour
- Matthew Gray
- Sabee Grewal
- Sander Gribling
- Dmitry Grinko
- Andi Gu
- Shouzhen Gu
- Aditya Gulati
- Jakob Günther
- Andrew Guo
- Francisco Escudero Gutiérrez
- Casper Gyurik
- Jonas Haferkamp
- Oliver Hahn
- Thomas Hahn
- Yassine Hamoudi
- Erkka Happasalo
- Dylan Harley
- Robin Harper
- Aram Harrow
- Atsuya Hasegawa
- Jing Yan Haw
- Ryu Hayakawa
- Zhiyang He
- Markus Heinrich
- Paul Hermouet
- Bence Hetenyi
- Minki Hhan
- Timo Hillman
- Christoph Hirche
- Taiga Hiroka
- Zahra Honjani
- Peter Hoyer
- Chung-Yun Hsieh
- Hong-Ye Hu
- Yanglin Hu
- Austin Hulse
- Nick Hunter-Jones
- Mark Myers II
- Luca Innocenti
- Joseph T. Iosue
- Sandy Irani
- Vishnu Iyer
- Dale Jacobs
- Rimika Jaiswal
- Samuel Jaques
- Stacey Jeffery
- Tomas Jochym-o'connor
- Stephen Jordan
- Sanad Kadu
- Gregory Kahanamoku-meyer
- Michael Kastoryano
- Marie Kempkes
- Sumeet Khatri
- Tanuj Khattar
- Chloe Kim
- Shelby Kimmel
- Robbie King
- William Kretschmer
- Hari Krovi
- Alexander Kulpe
- Niraj Kumar
- Rajendra Kumar
- Philippe Lamontagne
- Martin Larocca
- Seok-Hyung Lee
- Su-un Lee
- Itai Leigh
- Brian Lester
- Laura Lewis
- Andrew Li
- Bowen Li
- Jianqiang Li
- Jiawei Li
- Joseph Li
- Longcheng Li
- Lvzhou Li
- Xingjian Li
- Zhi Li
- Daniel Liang
- Xiao Liang
- Romi Lifshitz
- Timothy Lim
- Huiping Lin
- Yao-Ting Lin
- Diyi Liu
- Junyu Liu
- Li Liu

- Qipeng Liu
- Yinchen Liu
- Yupan Liu
- Zhengwei Liu
- Zhenhuan Liu
- Robin Lorenz
- Benjamin Lovitz
- Angus Lowe
- Chuhan Lu
- Xi Lu
- Angelo Lucia
- Jingquan Luo
- Cosmo Lupo
- Henry Ma
- Mark Ma
- Muzhou Ma
- Julio Magdalena
- Swarnadeep Majumder
- Daniel Malz
- Salvatore Mandra
- Zachary Mann
- John Martin
- Victor Martinez
- Kunal Marwaha
- Dmitri Maslov
- Kasra Masoudi
- Kieran Mastel
- Kaname Matsue
- Alex May
- Campbell Mclauchlan
- Saeed Mehraban
- Qiang Miao
- David Miloschewsky
- Shintaro Minagawa
- Michele Minervini
- Arjun Mirani
- Akimasa Miyake
- Masayuki Miyamoto
- Yin Mo
- Milad Moazami
- Ankith Mohan
- Wai-Keong Mok
- Léo Monbroussou
- Arsalan Motamedi
- Hamoon Mousavi
- Masih Mozakka
- Alexander Mueller-hermes
- Garazi Muguraza
- Anthony Munson
- Saachi Mutreja
- Long My
- Shivam Nadimpalli
- Shlok Nahar
- Preksha Naik
- Yoshifumi Nakata
- Giacomo Nannicini
- Ashwin Nayak
- Ion Nechita
- Barak Nehoran
- Jon Nelson
- Joshua Nevin
- Iu-Iong Ng
- Hongkang Ni
- Stuart Nicholls
- Junhong Nie
- Harold Nieuwboer
- Ryo Nishimaki
- Bryan O'Gorman
- Leo Orshansky
- Aadil Oufkir
- Connor Paddock
- Carlos Palazuelos
- Nikhil Pappu
- Natalie Parham
- Hakop Pashayan
- Dhrumil Patel
- Yash Patel
- Christopher Pattison
- Yuxiang Peng
- Tristan Philippe
- Stephen Piddock
- Yoann Pietri
- Pierre Pocreau
- Supartha Podder
- Abhinav Prem
- Timothy Proctor
- James Purcell
- Luowen Qian
- Minglong Qin
- Susan Qin
- Yihui Quek
- Marco Tulio Quintino
- Rebecca Radebold
- Seyoon Ragavan
- Michael Ragone
- Mizanur Rahaman

- Justin Raizes
- Joel Rajakumar
- Ronak Ramachandran
- Navneeth Ramakrishnan
- Sujit Rao
- Marco-Olivier Renou
- Denis Rochette
- Jérémie Roland
- Gregory Rosenthal
- Ingo Roth
- Cambyse Rouzé
- Baptiste Royer
- Roberto Rubboli
- Dorian Rudolph
- Manuel S. Rudolph
- Adrián Pérez Salinas
- Robert Salzmann
- Shengqi Sang
- Samuel Scalet
- Louis Schatzki
- Alexander Schmidhuber
- Sayantan Sen
- Zhong-Xia Shang
- Changpeng Shao
- Omar Shehab
- Yuki Shirakawa
- Omri Shmueli
- Oles Shtanko
- Nadish de Silva
- Thais Lima Silva
- Sophia Simon
- Sam Slezak
- William Slofsta
- Joseph Slote
- Kevin Smith
- Thomas Smith
- Rolando Somma
- Jeongrak Son
- Thomas Steckmann
- Anna Steffinlongo
- David Stephen
- Arthur Strauss
- Georgios Styliaris
- Yuan Su
- Ivan Šupić
- Jacopo Surace
- Daiki Suruga
- Yudai Suzuki
- Ryan Sweke
- Dániel Szabó
- Amnon Ta-shma
- Mostafa Taheri
- Yasuhiro Takahashi
- Yuki Takeuchi
- Suguru Tamaki
- Ernest Tan
- Xinyu Tan
- Er-Cheng Tang
- Eugene Tang
- Seiichiro Tani
- Barbara Terhal
- Supanut Thanasilp
- Thomas Theurer
- Quan Le Thien
- Ryan Tiew
- Sydney Timmerman
- Erickson Tjoa
- Kabir Tomer
- Yu Tong
- Allan Tosta
- Chung-En Tsai
- Kento Tsubouchi
- Takahiro Tsunoda
- Devashish Tupkary
- Jordi Tura
- Varun Upreti
- Michael Vasmer
- Almudena Carrera Vazquez
- Jevgenijs Vihrovs
- Tatiana Vovk
- Quoc-Huy Vu
- Rafael Wagner
- Mattia Walschaers
- Michael Walter
- Jiasu Wang
- Qisheng Wang
- Samson Wang
- Wenyuan Wang
- Xinzhao Wang
- Yunkai Wang
- James Watson
- Zack Weinstein
- Albert Werner
- Adam Wills
- Henrik Wilming
- Marek Winczewski

- Ronald de Wolf
- Ramona Wolf
- Lewis Wooltorton
- Jiawei Wu
- Peixue Wu
- Ya-Dong Wu
- Yue Wu
- Ziyi Xie
- Qian Xu
- Yijia Xu
- Shogo Yamada
- Hayata Yamasaki
- Gengzhi Yang
- Hongshun Yao
- Penghui Yao
- Jinmin Yi
- Chao Yin
- Hualei Yin
- Nelly NG Huei Ying
- Theodore J. Yoder
- Duyal Yolcu
- Satoshi Yoshida
- Nobuyuki Yoshioka
- Peter Yuen
- Allen Zang
- Pei Zeng
- Wei Zhan
- Chenyi Zhang
- Jiaqi Zhang
- Qing Zhang
- Xingjian Zhang
- Andrew Zhao
- Haimeng Zhao
- Mingnan Zhao
- Yuming Zhao
- Jerry Zheng
- Yufan Zheng
- Lai Zhou
- Leo Zhou
- Shuo Zhou
- You Zhou
- Shuchen Zhu
- Wei Zi
- Sebastian Zur
- Michael Zurel

Best Paper Prizes

The Program Committee selected as the TQC 2025 Best Paper Prize:

- *Polylog-time- and constant-space-overhead fault-tolerant quantum computation with quantum low-density parity-check codes*, by Shiro Tamiya, Masato Koashi, Hayata Yamasaki

The Program Committee selected as the TQC 2025 Best Student Paper Prize:

- *Quantum Purity Amplification: Optimality and Efficient Algorithm*, by Zhaoyi Li, Honghao Fu, Takuya Isogawa, Caio Silva, Isaac Chuang

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman
LIPICS Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Quantum Search with In-Place Queries

Blake Holman
Sandia National Laboratories, Albuquerque, NM, USA
Purdue University, West Lafeyette, IN, USA

Ronak Ramachandran
The University of Texas at Austin, TX, USA

Justin Yirka
Sandia National Laboratories, Albuquerque, NM, USA
The University of Texas at Austin, TX, USA

Abstract

Quantum query complexity is typically characterized in terms of XOR queries $|x, y\rangle \mapsto |x, y \oplus f(x)\rangle$ or phase queries, which ensure that even queries to non-invertible functions are unitary. When querying a permutation, another natural model is unitary: in-place queries $|x\rangle \mapsto |f(x)\rangle$.

Some problems are known to require exponentially fewer in-place queries than XOR queries, but no separation has been shown in the opposite direction. A candidate for such a separation was the problem of inverting a permutation over N elements. This task, equivalent to unstructured search in the context of permutations, is solvable with $O(\sqrt{N})$ XOR queries but was conjectured to require $\Omega(N)$ in-place queries.

We refute this conjecture by designing a quantum algorithm for Permutation Inversion using $O(\sqrt{N})$ in-place queries. Our algorithm achieves the same speedup as Grover's algorithm despite the inability to efficiently uncompute queries or perform straightforward oracle-controlled reflections.

Nonetheless, we show that there are indeed problems which require fewer XOR queries than in-place queries. We introduce a subspace-conversion problem called Function Erasure that requires 1 XOR query and $\Theta(\sqrt{N})$ in-place queries. Then, we build on a recent extension of the quantum adversary method to characterize exact conditions for a decision problem to exhibit such a separation, and we propose a candidate problem.

2012 ACM Subject Classification Theory of computation → Quantum query complexity

Keywords and phrases Quantum algorithms, query complexity, quantum complexity theory, quantum search, Grover's algorithm, permutation inversion

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.1

Related Version *Full Version*: `https://arxiv.org/abs/2504.03620v1`

Funding Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This work is supported by a collaboration between the US DOE and other Agencies. BH and JY acknowledge this work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research in Quantum Computing, Fundamental Algorithmic Research for Quantum Utility, with support also acknowledged from Fundamental Algorithm Research for Quantum Computing.
Ronak Ramachandran: Supported by the NSF AI Institute for Foundations of Machine Learning (IFML).
Justin Yirka: This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator.

Acknowledgements We especially thank John Kallaugher for helpful conversations and mentorship. RR and JY thank Scott Aaronson for stimulating discussions. JY thanks Bill Fefferman for suggesting his conjecture on PERMUTATION INVERSION.

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 1; pp. 1:1–1:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Quantum algorithms are typically developed and characterized in terms of query complexity. The strongest promises of quantum advantage over classical computation come from unconditional separations proved in terms of black-box queries, including Shor's period-finding algorithm and Grover's search algorithm. Understanding the nuances of the query model is therefore essential for advancing quantum algorithm design and sculpting quantum advantages.

Given an arbitrary Boolean function f, the standard query model in quantum computation is defined by XOR oracles S_f, also known as "standard oracles", which map basis states $|x\rangle|y\rangle$ to $|x\rangle|y \oplus f(x)\rangle$. Other common models, such as phase oracles, are known to be equivalent. The use of XOR oracles goes back to the early days of quantum computation [23, 20, 21, 18, 17] and even reversible computation [14, 15, 36, 16]. XOR oracles embed potentially irreversible functions in a reversible way, ensuring that all queries are unitary. This enables quantum query complexity to encompass arbitrary Boolean functions and offers a standard input-output format for using one algorithm as a sub-routine in another.

Other oracle models for quantum computation have been studied, but most abandon unitarity [37, 26, 41, 30, 27, 39, 32, 33] or provide query access to quantum functions with no analog in classical query complexity, e.g. general unitaries [3, 5].

When querying a permutation, there is another natural oracle model: an *in-place* oracle P_f which maps $|x\rangle$ to $|f(x)\rangle$. These oracles have been called in-place [22, 9], erasing [1, 2], and minimal [28, 8].[1] Just like XOR oracles, in-place oracles can be directly studied and compared in both quantum and classical computation.

In-place oracles were first studied in the quantum setting by Kashefi, Kent, Vedral, and Banaszek [28]. They showed several results comparing XOR oracles and in-place oracles, including a proof that $\Theta(\sqrt{N})$ queries to an XOR oracle are required to simulate an in-place query to the same permutation. Around the same time, Aaronson [1] proved that SET COMPARISON, an approximate version of the COLLISION problem, requires an exponential number of XOR queries but only a constant number of in-place queries.

These oracles relate to multiple topics in quantum algorithms and complexity theory. Aaronson's lower bound for the collision problem [1] was partially inspired by the desire to separate the in-place and XOR query models. [28] observed that a constant number of in-place queries is sufficient to solve RIGID GRAPH ISOMORPHISM, a necessary subcase for solving general GRAPH ISOMORPHISM. An identical protocol was later generalized to define the concept of QSAMPLING, which is sufficient to solve SZK, by Aharonov and Ta-Shma [4]. These ideas inspired pursuing lower bounds on the INDEX ERASURE problem [40, 7, 31], ruling out potential algorithms for GRAPH ISOMORPHISM using XOR oracles. Fefferman and Kimmel [22] showed an oracle separation of QMA and QCMA relative to randomized in-place oracles. Also, the expressive power of in-place oracles relates to the conjectured existence of one-way permutations [15, p. 926]. Additionally, because in-place oracles are not self-inverse, they offer a setting in which to study computation with inverse-free gate sets [19].

In-place oracles outperform XOR oracles in every established separation between the two query models, but it is conjectured that the oracles are incomparable, each better-suited for certain tasks. Aaronson [2] raised proving such a separation as an open problem. Fefferman

[1] Unfortunately, "permutation oracle" has been used to refer to any oracle which embeds a permutation. Following a suggestion by John Kallaugher, we have found it convenient in conversation to refer to "xoracles" and "smoracles" (for "small oracles").

and Kimmel [22] conjectured that inverting a permutation over N elements, a task which requires only $O(\sqrt{N})$ queries to an XOR oracle, requires $\Omega(N)$ queries to an in-place oracle. PERMUTATION INVERSION is formally as hard as unstructured search [34], so this conjecture effectively predicts that the speedup of Grover's algorithm [25] is impossible with an in-place oracle.

Results

We refute the conjecture of [22] by designing a new quantum algorithm that solves PERMUTATION INVERSION with $O(\sqrt{N})$ queries to an in-place oracle, recovering the same speedup as Grover's search algorithm.

We additionally apply this algorithm to tightly characterize the ability of XOR and in-place oracles to simulate each other. Then, we change focus and make progress towards showing the desired separation. We introduced a subspace-conversion problem that requires 1 XOR query and exponentially-many in-place queries.

Finally, we propose a candidate decision problem that can be solved with $O(\sqrt{N})$ queries to an XOR oracle and that we conjecture requires $\Omega(N)$ queries to an in-place oracle. We then apply recent advances in the quantum adversary bound to define a new class of adversary matrices which must be used if such a decision-problem separation exists.

1.1 Quantum Search

Unstructured search, famously solved by Grover's algorithm with $O(\sqrt{N})$ queries to an XOR oracle, is one of the most well-studied problems in quantum query complexity. The first non-trivial quantum lower bound was for unstructured search [17]. Later work modifying the query model, for instance by introducing noise or faults into queries, focused on unstructured search [37, 41, 30, 27, 39, 5].

In-place oracles are only defined for bijections (see Section 2). Restricted to permutations, the unstructured search problem is equivalent to PERMUTATION INVERSION [34].[2]

▶ **Definition 1.** *Given query access to a permutation π on $[N] = \{0, \ldots, N-1\}$, the* PERMUTATION INVERSION *problem is to output $\pi^{-1}(0)$.*

The choice to invert 0 can of course be replaced with any element. It is also straightforward to define a related decision problem, for example, deciding if $\pi^{-1}(0)$ is odd or even.

Like general unstructured search, PERMUTATION INVERSION has been a frequent target for new lower bound techniques. It can be solved with $O(\sqrt{N})$ queries to an XOR oracle using Grover's algorithm. Ambainis [6] applied his new quantum adversary method to show that $\Omega(\sqrt{N})$ queries to an XOR oracle are in fact required to solve the problem. Nayak [34] gave an alternative proof by showing the problem is as hard as general unstructured search. Rosmanis [38] also reproduced this tight lower bound using the compressed oracle technique on random permutations. As for in-place oracles, [22] proved that $\Omega(\sqrt{N})$ in-place queries are needed to solve PERMUTATION INVERSION. Belovs and Yolcu [13] later applied their advancements on the quantum adversary method to reprove the same lower bound. We add to this sequence of work, studying PERMUTATION INVERSION in Section 3 to give the following result.

[2] The reductions between PERMUTATION INVERSION and unstructured search are entirely classical. So the reductions hold using either XOR oracles or in-place oracles, although some quantum garbage registers may differ.

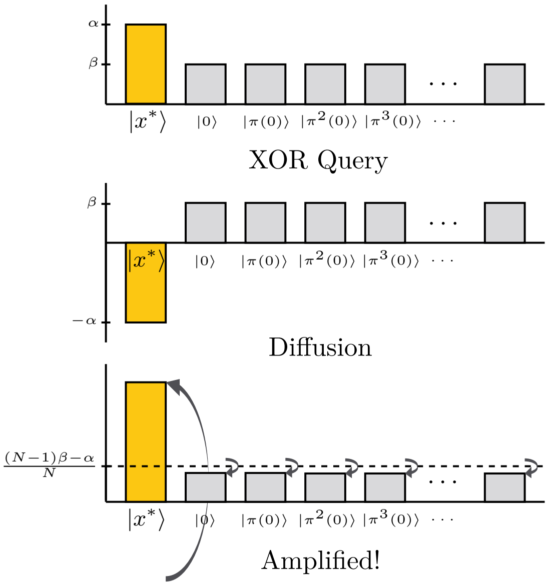

Figure 1 (Color) Illustration of how one iteration of Grover's search algorithm amplifies $|x^*\rangle := |\pi^{-1}(0)\rangle$. Amplitudes are ordered according to the permutation in order to match Figure 3 later.

▶ **Theorem 2.** *For a permutation π on $[N]$,* Permutation Inversion *can be solved with* $\mathrm{O}(\sqrt{N})$ *in-place queries to π.*

Thus, we refute the conjecture that $\Omega(N)$ in-place queries are required, and we show the $\Omega(\sqrt{N})$ lower bound [22, 12] is tight.

Grover's Algorithm

Before we sketch our algorithm, we first recall Grover's algorithm for unstructured search [25] in the context of Permutation Inversion. Grover's algorithm repeatedly alternates between using XOR queries to negate the amplitude of $|\pi^{-1}(0)\rangle$ and using the "Grover Diffusion operator" to reflect all amplitudes about the average, steadily amplifying $|\pi^{-1}(0)\rangle$ on every iteration. In other words, the algorithm alternates between the oracle-dependent reflection $I - 2\,|\pi^{-1}(0)\rangle\langle\pi^{-1}(0)|$ and the diffusion reflection

$$\mathrm{D} = I - 2\,|s\rangle\langle s|\,, \tag{1}$$

where $|s\rangle$ is the uniform superposition $\frac{1}{\sqrt{N}}\sum|i\rangle$. This is illustrated in Figure 1.

In-place oracles seem at odds with oracle-dependent reflections, since reflections – like XOR queries – are self-inverse, but inverting an in-place query is equivalent to inverting the underlying permutation, which would solve Permutation Inversion. With this in mind, it would be natural to conjecture, as [22] did, that no Grover-style speedup is possible using in-place oracles.

A New Algorithm

Let $x^* := \pi^{-1}(0)$ be the "marked item" to be found. Our algorithm starts with an equal superposition over $[N]$ along with an ancilla register and a "flag" qubit: $\frac{1}{\sqrt{N}}\sum|i\rangle\,|0^n\rangle\,|0\rangle$. The algorithm repeatedly iterates over steps *Mark*, *Shift*, and *Diffuse the Difference*. The intuition behind these steps is as follows.

- *Mark:* For every basis state $i \in [N]$, make a copy of i and query the oracle. Then, conditioned on the output of $\pi(i)$ being 0, flip the flag qubit from $|0\rangle$ to $|1\rangle$.
  (The *Mark* step cannot be used to implement Grover's algorithm as usual because the query answer remains in the ancilla register, as garbage, until the next step.)
- *Shift:* In the $|1\rangle$-flagged branch, all amplitude is concentrated on $|x^*\rangle$, while in the $|0\rangle$-flagged branch, the amplitude is spread evenly over all basis states except $|x^*\rangle$.
  In only the $|0\rangle$-flagged branch of the superposition, query the oracle to shift the amplitude of each basis state forward according to π (perform a controlled in-place query to π). This shifts amplitude from $|i\rangle$ onto $|\pi(i)\rangle$, and in particular, from $|\pi^{-1}(x^*)\rangle$ onto $|x^*\rangle$.
- *Diffuse the Difference:* The two branches are now such that if they are interfered to produce two branches, one branch which adds amplitudes and another branch which subtracts amplitudes, then the amplitude on $|x^*\rangle$ would be above average in the former branch and below average in the latter branch.
  Perform the standard Grover diffusion operator (Equation (1)) controlled on the flag qubit being the $|-\rangle$ state, which reflects the "difference branch" about its average amplitude. This results in the amplitude on $|x^*\rangle$ being similarly amplified in both branches. In fact, we find the branches are inverse-exponentially close to each other, and that after the t-th iteration, the overall state is effectively

  $$|\psi_t\rangle = \left(\alpha_t |x^*\rangle + \sum_{i \in [N] \setminus \{x^*\}} \beta_t |i\rangle \right) |0^n\rangle |0\rangle ,$$

  where α_t increases by approximately $1/\sqrt{N}$ each iteration.

These steps are repeated $\mathrm{O}(\sqrt{N})$ times to increase the amplitude on $|x^*\rangle$ until there is a constant probability of measuring it. Each iteration uses a constant number of in-place queries, so the overall query complexity is $\mathrm{O}(\sqrt{N})$. For more intuition, see a circuit diagram in Figure 2 and an illustration in Figure 3 similar to Figure 1 above.

In Section 2.1, we give a construction for the controlled in-place query necessary for the *Shift* step of the algorithm. This construction differs significantly from the analogous construction for XOR oracles.

▶ **Lemma 3.** *There exists a unitary circuit making 1 in-place query to π which for all $x \in [N]$ maps*

$$|a\rangle |x\rangle |y\rangle \mapsto \begin{cases} |a\rangle |x\rangle |y\rangle & \text{when } a = 0 \\ |a\rangle |\pi(x)\rangle |y\rangle & \text{when } a = 1 \end{cases},$$

where y is the image under π of some fixed point, such as $y = \pi(0)$.

Note that although y depends on the oracle π, it is independent of the query x. So while y is garbage, it is effectively negligible. Because it is never entangled with the input register, the garbage can be safely measured and erased. See Section 2.1 for more details.

1.2 Simulating Other Oracles

In Section 4, we tightly characterize the ability of XOR and in-place oracles to simulate each other. We do so by applying our new algorithm to give new upper bounds and by developing a novel lower bound. The contents of Section 4 are deferred to the Full Version. For a permutation π on $[N]$, Grover's algorithm can be used to simulate an XOR query to π^{-1}, an in-place query to π^{-1}, or an in-place query to π using $\mathrm{O}(\sqrt{N})$ XOR queries to π, and this complexity is known to be tight [28]. We show how to use our new algorithm to perform the analogous simulations using $\mathrm{O}(\sqrt{N})$ queries to an in-place oracle. The constructions are non-trivial due to the inability of in-place oracles to uncompute garbage. The simulations are approximate with inverse-exponential error due to the error in our algorithm for PERMUTATION INVERSION.

Next, we prove that our simulations are tight by giving matching lower bounds. Inspired by [28], we prove this by arguing that if few in-place queries could simulate an XOR query, then we could violate the lower bound of [22] for performing unstructured search.

▶ **Theorem 4.** *For a permutation π on $[N]$, $\Omega(\sqrt{N})$ in-place queries to π are necessary to approximately simulate an* XOR *query to π.*

Given that an XOR query to π can be implemented using 1 XOR query to π, Theorem 4 makes this the first task known to require more in-place queries than XOR queries. We improve on this in the next section.

We can summarize all upper and lower bounds above as follows.

▶ **Corollary 5** (Summary of relationships)**.** *For a permutation π on $[N]$, $\Theta(\sqrt{N})$ queries to any one of an in-place oracle for π, an in-place oracle for π^{-1}, an* XOR *oracle for π, or an* XOR *oracle for π^{-1} are necessary and sufficient to approximately simulate any one of the others.*

1.3 A Subspace-Conversion Separation

In Section 5 we improve the unitary-implementation separation given in the previous section to a subspace-conversion separation. The contents of Section 5 are deferred to the Full Version.

INDEX ERASURE is the task of generating the state $\frac{1}{\sqrt{N}}\sum_{x\in[N]}|f(x)\rangle$ given queries to f. It was introduced by Shi [40] and formalized as a state-generation task by Ambainis, Magnin, Roetteler, and Roland [7]. As noted by [40], solving INDEX ERASURE would imply solutions to SET EQUALITY and GRAPH ISOMORPHISM. Similar work on QSampling [4] suggests many more applications. INDEX ERASURE requires $\Omega(\sqrt{N})$ XOR queries [7, 31] but just 1 in-place query, so the problem seems to capture key differences between the models.

We define the converse problem, FUNCTION ERASURE.

▶ **Definition 6.** *Given query access to a function f,* FUNCTION ERASURE *is the subspace-conversion problem of transforming any superposition of the form $\sum \alpha_x |x\rangle |f(x)\rangle$ to $\sum \alpha_x |x\rangle$.*

A state-conversion problem requires implementing an algorithm which, given an oracle to function f, maps an input $|\psi_f\rangle$ to output $|\phi_f\rangle$. A subspace-conversion problem simply generalizes this to multiple input-output pairs for each oracle function f. We discuss the details of unitary-implementation, subspace-conversion, and other types of problems in Section 5.

FUNCTION ERASURE can trivially be solved with 1 XOR query to f. Then by Corollary 5, $\mathrm{O}(\sqrt{N})$ in-place queries are sufficient. Finally, we show how FUNCTION ERASURE and one additional in-place query are sufficient to simulate an XOR query. To avoid violating Theorem 4, this implies $\Omega(\sqrt{N})$ queries are necessary.

▶ **Theorem 7.** *For a permutation π on $[N]$, $\Theta(\sqrt{N})$ in-place queries to π are necessary and sufficient for* FUNCTION ERASURE*.*

Theorem 7 makes FUNCTION ERASURE the first coherent subspace-conversion problem known to require fewer XOR queries than in-place queries. This improves on the new unitary-implementation separation from the previous section.

1.4 Lower Bounds

The first works to study in-place oracles proved that there are problems which can be solved with asymptotically fewer queries to in-place oracles than to the corresponding XOR oracles [28, 1]. They left open the question of whether a separation could be shown in the opposite direction, making the two oracles formally incomparable, or whether in-place oracles are generically superior to XOR oracles. Our main result (Theorem 2) refutes one conjectured path towards constructing a problem for which XOR oracles are better than in-place oracles. Our study of Function Erasure demonstrates the first problem which provably requires fewer queries to an XOR oracle than an in-place oracle, although it is a subspace-conversion problem instead of a decision problem. In Section 6, we consider the possibility of improving this to a decision-problem separation.

Conventional Lower Bound Techniques

Section 6.1 is deferred to the Full Version. There, we discuss how common quantum lower bound techniques, the polynomial method [10] and the unweighted adversary method [6], fail to prove the desired separation. We show that under these techniques, any lower bound on the number of in-place queries implies the same lower bound on the number of XOR queries, making these techniques unable to prove a separation where XOR oracles outperform in-place oracles.

A Candidate Decision Problem

In Section 6.2, we introduce a new problem, Embedded PermInv, which can be solved with $\Theta(\sqrt{N})$ queries to an XOR oracle and which we conjecture requires $\Omega(N)$ queries to an in-place oracle. As we discuss, the problem is designed to embed an injection from $[N^2]$ to $[N]$ into a bijection on $[N^2]$, which we believe circumvents algorithms using in-place oracles. The idea behind this problem builds on the "Simon's problem with garbage" proposed by Aaronson [2].

Techniques for a Decision-Problem Separation

Finally, in Section 6.3, we briefly discuss the potential for more sophisticated lower bound methods to prove a decision-problem separation, including for our candidate Embedded PermInv. A full exposition is given in the Appendix of the full version of this article.

The recent extension of the quantum adversary method by Belovs and Yolcu [13] applies to arbitrary linear transformations, including in-place oracles. The adversary bound is an optimization problem over *adversary matrices* such that the optimal value equals the quantum query complexity for a given problem. Of course, the difficulty with the adversary method is to design a "good" adversary matrix exhibiting a tight bound.

We introduce a special class of feasible solutions which we call *extended adversary matrices*. We show, with some technical caveats, that there exists an XOR query advantage over in-place oracles for a decision problem if and only if it is witnessed by extended adversary matrices. Then, for our candidate problem Embedded PermInv, we are able to remove these caveats and state that if our conjectured separation is true, then it must be witnessed by extended adversary matrices.

1.5 Open Problems

A list of open problems is given in the full version of this article.

2 Quantum Oracles

As stated previously, the standard query model in quantum computation and classical reversible computation is the XOR oracle. Other common models, such as the phase oracle, are equivalent. For a function f, an XOR oracle S_f maps $|x\rangle |y\rangle \mapsto |x\rangle |y \oplus f(x)\rangle$, where $\oplus$ denotes bitwise XOR with queries encoded in binary.

When querying an invertible function, there is another natural unitary query model.[3] An in-place oracle P_π maps $|x\rangle \mapsto |\pi(x)\rangle$.

Here we list several basic identities given by [28].

1. Given query access to both π and π^{-1}, standard and in-place oracles are equivalent. More precisely, P_π can be simulated using 1 query to S_π and 1 query to either of $\mathrm{S}_{\pi^{-1}}, \mathrm{P}_{\pi^{-1}}$. Similarly, S_π can be simulated using 1 query to P_π and 1 query to either of $\mathrm{S}_{\pi^{-1}}, \mathrm{P}_{\pi^{-1}}$. So, the interesting case is when we can query π but cannot query its inverse.
2. XOR oracles are self-inverse, $\mathrm{S}_\pi = (\mathrm{S}_\pi)^\dagger$, but generally $(\mathrm{S}_\pi)^\dagger \neq \mathrm{S}_{\pi^{-1}}$. In contrast, generally $\mathrm{P}_\pi \neq (\mathrm{P}_\pi)^\dagger$ but it does hold that $(\mathrm{P}_\pi)^\dagger = \mathrm{P}_{\pi^{-1}}$.
3. $\Theta(\sqrt{N})$ queries to an XOR oracle S_π can be used to simulate a query to P_π. The upper bound is due to Grover's search algorithm. The lower bound follows by observing that a circuit for P_π querying S_π can be inverted to give a circuit for $\mathrm{P}_{\pi^{-1}}$ querying $(\mathrm{S}_\pi)^\dagger = \mathrm{S}_\pi$, which would solve PERMUTATION INVERSION, which requires $\Omega(\sqrt{N})$ queries to S_π.

The XOR query model was motivated by two needs. First is the need to embed non-invertible functions in a reversible query. Second is that because XOR oracles are self-inverse, they enable uncomputing. An early criticism of reversible computation by Landauer [29] was that in order to maintain reversibility, a computation would need to retain intermediate work until the end, only deferring the cost of information erasure instead of avoiding it. To the contrary, Bennett showed that any circuit can efficiently be made into a reversible one that uncomputes any intermediate work and gives its original output in the form of an XOR query [14, 15]. Given a garbage-producing reversible circuit, first apply the circuit, then copy the desired output into a new register using XOR, and then apply the circuit in reverse, gate-by-gate, to uncompute all intermediate steps, leaving only the input and the copied output. Moreover, such a gate-by-gate reversal works when one algorithm is used as a black-box subroutine for another, since given a black-box following this XOR-model, it is self-inverse. So full algorithms, including subroutines, can indeed be reversed gate-by-gate. Besides these two reasons, XOR oracles simply appeared natural at the time quantum computing was formalized. As far as we are aware, in-place oracles have not been studied in the classical reversible computing literature. There have been just a few references to alternative classical reversible implementations of 1-to-1 functions [36, 16]. So quantum computation, which is based on reversible operations, later inherited the XOR model. At the same time, the ability to uncompute enabled quantum interference [23, 18]. Many early results also only involved binary functions, and other results were motivated more by ensuring quantum computers could implement tasks such as error-reduction and subroutines ($\mathsf{BQP}^{\mathsf{BQP}} = \mathsf{BQP}$ [17]) rather than questioning the query model.

[3] We restrict our study to bijections, and without loss of generality to permutations on $[N]$. A similar oracle which queries an injection would still be reversible, but it would be an isometry rather than a unitary. Our algorithm seems to require a bijection since is uses the oracle's previous outputs as its next inputs.

One more important feature of XOR oracles is that for a function f, the complexity of implementing S_f using reversible operations is at most a constant multiplicative factor more than for the general, irreversible circuit implementing f [14]. For in-place oracles, no construction is known for efficiently transforming an irreversible circuit for permutation π into a reversible circuit for P_π. In fact, the widely believed existence of one-way permutations implies that there exist permutations for which this is impossible. This is because given a reversible implementation of P_π, inverting the circuit gate-by-gate gives $(\mathrm{P}_\pi)^\dagger = \mathrm{P}_{\pi^{-1}}$ with exactly the same circuit size, whereas one-way permutations should have different complexities than their inverses. This may limit the practical instantiation of in-place oracles, although they may lead to useful insights in other ways.

2.1 Controlled In-Place Oracle

The proof of Lemma 3 is deferred to the Full Version.

3 Permutation Inversion

In this section, we prove our main result, that Permutation Inversion (Definition 1) can be solved with $\Theta(\sqrt{N})$ queries to an in-place oracle.

Proof of Theorem 2. The lower bound was proved by Fefferman and Kimmel [22] and later reproved by [13]. To prove the upper bound, we give an algorithm.

Algorithm. For convenience, we assume $N = 2^n$ and identify the integers $[N]$ by their binary representations in $\{0,1\}^n$. We denote the target element $\pi^{-1}(0)$ by x^*.

First, query P_π once to check whether $\pi(0)$ is 0, and terminate early with answer 0 if it is. Otherwise, initialize three registers to the state $|\psi_0\rangle := \frac{1}{\sqrt{N}} \sum_{i=1}^{N} |i\rangle_{\mathcal{A}} |0^n\rangle_{\mathcal{B}} |0\rangle_{\mathcal{C}}$, where $\mathcal{A}$ and $\mathcal{B}$ are each $n = \log N$ qubits and $\mathcal{C}$ is one qubit. Then, repeat the following steps $T = \mathrm{O}(\sqrt{N})$ times:

(1) *Mark*
XOR register $\mathcal{A}$ into $\mathcal{B}$, and apply P_π to $\mathcal{B}$.
Controlled on $\mathcal{B}$ being $|0^n\rangle$, apply NOT to $\mathcal{C}$, flagging the branch where $\mathcal{A}$ contains x^*.

(2) *Shift (and Clean Up)*
Controlled on $\mathcal{C}$ being $|0\rangle$, apply P_π to $\mathcal{A}$.
Controlled on $\mathcal{C}$ being $|0\rangle$, XOR $\mathcal{A}$ into $\mathcal{B}$, resetting $\mathcal{B}$ to $|0^n\rangle$.

(3) *Diffuse the Difference*
Controlled on $\mathcal{C}$ being $|-\rangle$, apply the diffusion operator to $\mathcal{A}$.
The diffusion operator $\mathrm{D} := 2H^{\otimes n} |0^n\rangle\langle 0^n| H^{\otimes n} - I$ is the same used in Grover's algorithm [25], equivalent to a reflection about the uniform superposition.

(4) Optional: *Measure*
Measure $\mathcal{C}$. If $|1\rangle$ is observed then abort and report failure.

Finally, measure register $\mathcal{A}$ and output the result. See Figure 2 for a circuit diagram of one iteration of the algorithm and Figure 3 for an illustration of the effect.

Below, we will find that each *Measure* step aborts with probability $1/N$. So, these intermediate measurements could be omitted and the qubit reused as it is, and the quantum union bound [24, 35] implies the overall success probability would decrease by at most $\sqrt{T/N} = \mathrm{O}\big(N^{-1/4}\big)$. For now, we include the optional *Measure* step to simplify the analysis.

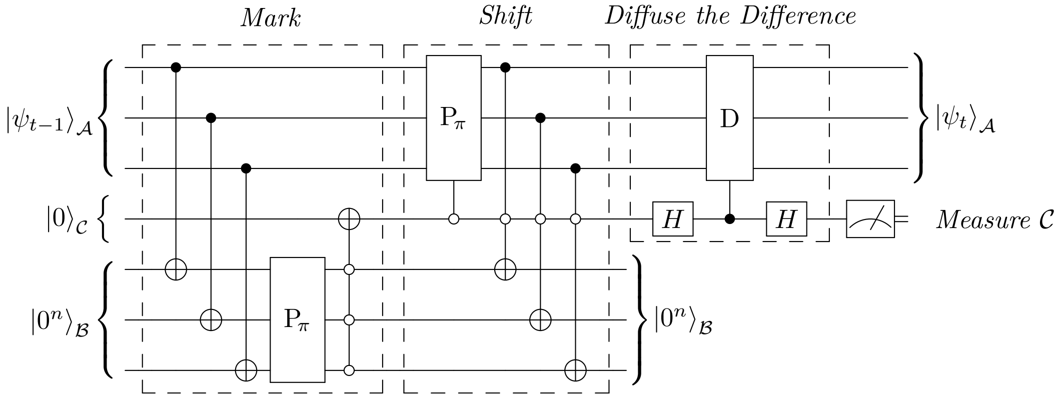
Mark
Shift
Diffuse the Difference
$|\psi_{t-1}\rangle_{\mathcal{A}}$
$|0\rangle_{\mathcal{C}}$
$|0^n\rangle_{\mathcal{B}}$
P_π
P_π
D
H
H
$|\psi_t\rangle_{\mathcal{A}}$
Measure $\mathcal{C}$
$|0^n\rangle_{\mathcal{B}}$

Figure 2 One iteration of our Permutation Inversion algorithm. D is the standard diffusion operator. • denotes an operation controlled on $|1\rangle$ and ◦ denotes an operation controlled on $|0\rangle$.

Analysis. Now we prove that our algorithm succeeds with high probability.

We use $|\psi_t\rangle$ to denote the state after t iterations. We will show by induction that after each iteration, if the algorithm did not terminate early, then the state is of the form

$$|\psi_t\rangle = \left(\alpha_t\,|x^*\rangle + \sum_{i\in[N]\setminus\{x^*\}} \beta_t\,|i\rangle\right)_{\mathcal{A}} \otimes |0^n\rangle_{\mathcal{B}}\,|0\rangle_{\mathcal{C}} \tag{2}$$

for some real values α_t, β_t. In particular, all $|i\rangle$ for $i \neq x^*$ share the same amplitude. The transformation from $|\psi_{t-1}\rangle$ to $|\psi_t\rangle$ is illustrated in Figure 3.

By construction, the initial state $|\psi_0\rangle$ is the uniform superposition, with $\alpha_0 = \beta_0 = \frac{1}{\sqrt{N}}$. Next, the t-th iteration begins with the state

$$|\psi_{t-1}\rangle = \left(\alpha_{t-1}\,|x^*\rangle + \sum_{i\in[N]\setminus\{x^*\}} \beta_{t-1}\,|i\rangle\right) |0^n\rangle\,|0\rangle\,.$$

For ease of notation, we will drop the subscripts so that α, β implicitly refer to $\alpha_{t-1}, \beta_{t-1}$. After the *Mark* step, the state will be

$$|\psi'_{t-1}\rangle = \alpha\,|x^*\rangle\,|0^n\rangle\,|1\rangle + \sum_{i\in[N]\setminus\{x^*\}} \beta\,|i\rangle\,|\pi(i)\rangle\,|0\rangle\,.$$

After the *Shift (and Clean Up)* step, the state will be

$$\begin{aligned}|\psi''_{t-1}\rangle &= \alpha\,|x^*\rangle\,|0^n\rangle\,|1\rangle + \sum_{i\in[N]\setminus\{x^*\}} \beta\,|\pi(i)\rangle\,|0^n\rangle\,|0\rangle \\ &= \alpha\,|x^*\rangle\,|0^n\rangle\,|1\rangle + \sum_{i\in[N]\setminus\{0\}} \beta\,|i\rangle\,|0^n\rangle\,|0\rangle\,.\end{aligned}$$

As the name suggests, this step shifts amplitudes within the summation off of $|0\rangle$ and onto $|x^*\rangle$.

Next, to prepare for the *Diffuse the Difference* step, we rewrite register $\mathcal{C}$ in the Hadamard basis. The state is equivalent to

$$\begin{aligned}|\psi''_{t-1}\rangle = \quad & \frac{1}{\sqrt{2}}\left[(\beta+\alpha)\,|x^*\rangle + \sum_{i\in[N]\setminus\{0,x^*\}} \beta\,|i\rangle\right]|0^n\rangle\,|+\rangle \\ & + \frac{1}{\sqrt{2}}\left[(\beta-\alpha)\,|x^*\rangle + \sum_{i\in[N]\setminus\{0,x^*\}} \beta\,|i\rangle\right]|0^n\rangle\,|-\rangle\,.\end{aligned}$$

Next, the *Diffuse the Difference* step applies the diffusion operator D controlled on $\mathcal{C}$ being $|-\rangle$. The diffusion operator can be viewed as reflecting every amplitude about the average amplitude. This results in

$$\begin{aligned}|\psi'''_{t-1}\rangle &= \frac{1}{\sqrt{2}}\left[(\beta+\alpha)\,|x^*\rangle + \sum_{i\in[N]\setminus\{0,x^*\}} \beta\,|i\rangle\right]|0^n\rangle\,|+\rangle \\ &+ \frac{1}{\sqrt{2}}\Bigg[\left(\beta+\alpha-\frac{2(\beta+\alpha)}{N}\right)|x^*\rangle + \left(2\beta - \frac{2(\beta+\alpha)}{N}\right)|0\rangle \\ &\qquad\qquad + \sum_{i\in[N]\setminus\{0,x^*\}}\left(\beta - \frac{2(\beta+\alpha)}{N}\right)|i\rangle\Bigg]|0^n\rangle|-\rangle\,.\end{aligned}$$

Returning register $\mathcal{C}$ to the standard basis, we see

$$\begin{aligned}|\psi'''_{t-1}\rangle = &\left[\left(\beta+\alpha-\frac{\beta+\alpha}{N}\right)|x^*\rangle + \sum_{i\in[N]\setminus\{x^*\}}\left(\beta-\frac{\beta+\alpha}{N}\right)|i\rangle\right]|0^n\rangle\,|0\rangle \\ &+ \left[\frac{\beta+\alpha}{N}\,|x^*\rangle - \left(\beta - \tfrac{\beta+\alpha}{N}\right)|0\rangle + \sum_{i\in[N]\setminus\{0,x^*\}} \frac{\beta+\alpha}{N}\,|i\rangle\right]|0^n\rangle\,|1\rangle\,.\end{aligned}$$

The amplitude on $|x^*\rangle$ is now larger than the original amplitude α.

Finally, for the sake of analysis, we choose to measure $\mathcal{C}$ and abort if $|1\rangle$ is observed. We will handle the failure case later. For now, we postselect on having observed $|0\rangle$. This results in the final (normalized) state

$$\begin{aligned}|\psi_t\rangle &= \sqrt{\frac{N}{N-1}}\left[\left(\beta+\alpha-\frac{\beta+\alpha}{N}\right)|x^*\rangle + \sum_{i\in[N]\setminus\{x^*\}}\left(\beta-\frac{\beta+\alpha}{N}\right)|i\rangle\right]|0^n\rangle\,|0\rangle \\ &= \left[\sqrt{\frac{N-1}{N}}\,(\beta+\alpha)\,|x^*\rangle + \sum_{i\in[N]\setminus\{x^*\}}\left(\sqrt{\frac{N-1}{N}}\beta - \frac{1}{\sqrt{N}\sqrt{N-1}}\alpha\right)|i\rangle\right]|0^n\rangle\,|0\rangle\,.\end{aligned}$$

at the end of the t-th iteration. This state has the form we claimed, with

$$\alpha_t = \sqrt{\frac{N-1}{N}}\,(\beta_{t-1}+\alpha_{t-1}) \quad\text{and}\quad \beta_t = \sqrt{\frac{N-1}{N}}\beta_{t-1} - \frac{1}{\sqrt{N}\sqrt{N-1}}\alpha_{t-1},$$

concluding our induction.

The above recurrence lets us write a closed form for α_t and β_t:

$$\begin{bmatrix}\alpha_t\\ \beta_t\end{bmatrix} = \begin{bmatrix}\sqrt{\frac{N-1}{N}} & \sqrt{\frac{N-1}{N}} \\ \frac{-1}{\sqrt{N}\sqrt{N-1}} & \sqrt{\frac{N-1}{N}}\end{bmatrix}\begin{bmatrix}\alpha_{t-1}\\ \beta_{t-1}\end{bmatrix} = \begin{bmatrix}\sqrt{\frac{N-1}{N}} & \sqrt{\frac{N-1}{N}} \\ \frac{-1}{\sqrt{N}\sqrt{N-1}} & \sqrt{\frac{N-1}{N}}\end{bmatrix}^t\begin{bmatrix}\frac{1}{\sqrt{N}}\\ \frac{1}{\sqrt{N}}\end{bmatrix}.$$

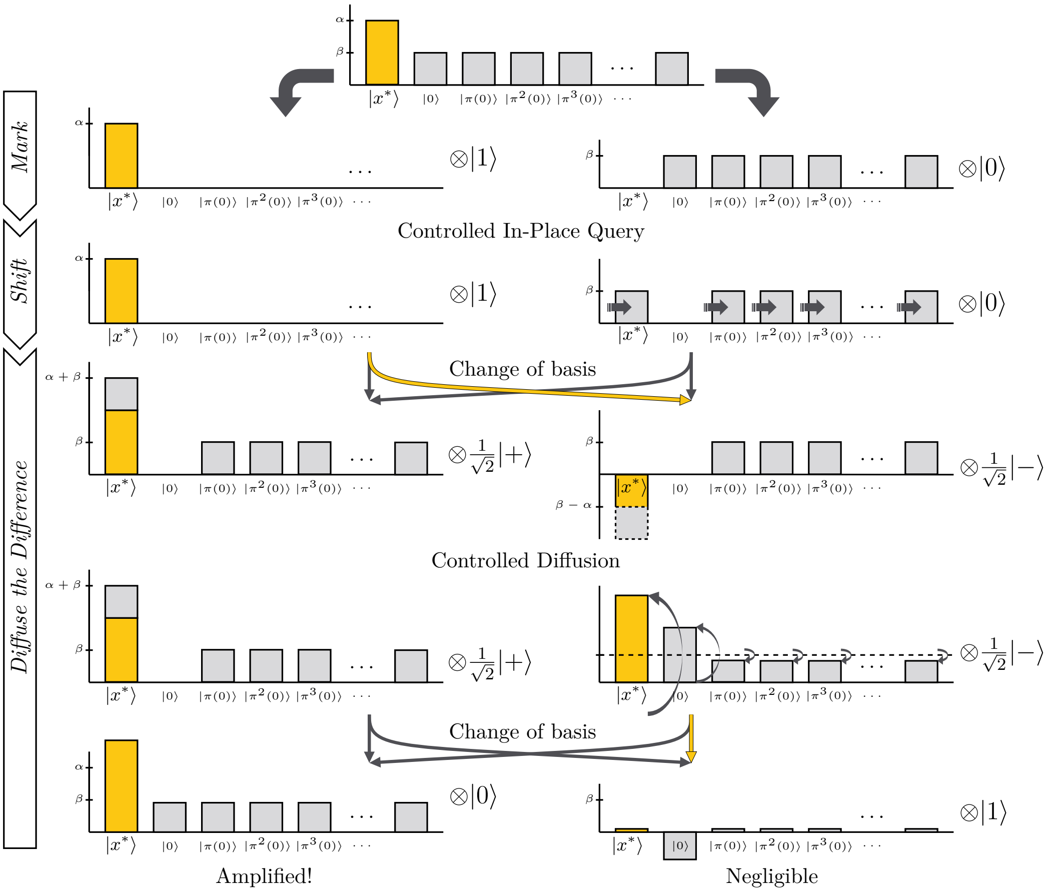
Mark
Shift
Diffuse the Difference
α
β
$|x^*\rangle$ $|0\rangle$ $|\pi(0)\rangle$ $|\pi^2(0)\rangle$ $|\pi^3(0)\rangle$ $\cdots$
$\otimes|1\rangle$
$\otimes|0\rangle$
Controlled In-Place Query
Change of basis
α + β
β − α
$\otimes\frac{1}{\sqrt{2}}|+\rangle$
$\otimes\frac{1}{\sqrt{2}}|-\rangle$
Controlled Diffusion
Change of basis
Amplified!
Negligible

Figure 3 (Color) Illustration of how amplitudes change in each iteration of the algorithm. Register $\mathcal{B}$ is left implicit (note it is unentangled with $\mathcal{A}$ and $\mathcal{C}$ by the end of the *Shift* step). Each iteration begins with the nearly uniform superposition from Equation (2). The *Mark* step queries π and creates a marked branch and an unmarked branch, illustrated in two columns. The *Shift* step makes a query in only the unmarked branch, shifting amplitude onto $|x^*\rangle$. The *Diffuse the Difference* step is controlled on $|-\rangle$, so we first rewrite the basis of $\mathcal{C}$, rearranging amplitudes accordingly. Black and yellow arrows indicate positive and negative contributions. The diffusion operator reflects all amplitudes about their mean. A final change of basis leaves a state almost entirely entangled with $|0\rangle$ and with increased amplitude on x^*.

For a diagonalizable matrix $M = ADA^{-1}$, we know $M^t = AD^tA^{-1}$, so we can diagonalize the above matrix to find

$$\alpha_t = \frac{1}{\sqrt{N^{t+1}}}\frac{1}{2i}\left[\left(\sqrt{N-1}+i\right)^{t+1} - \left(\sqrt{N-1}-i\right)^{t+1}\right].$$

Rewriting the expression in polar form, this is equivalent to

$$\alpha_t = \frac{1}{\sqrt{N^{t+1}}}\frac{1}{2i}\left[\sqrt{N^{t+1}}e^{i(t+1)\theta} - \sqrt{N^{t+1}}e^{-i(t+1)\theta}\right] \quad \text{for} \quad \theta = \arctan\left(\frac{1}{\sqrt{N-1}}\right).$$

Finally, the identity $\frac{z-\bar{z}}{2i} = \operatorname{Im}(z) = \sin(\phi)$ for $z = e^{i\phi}$ yields

$$\alpha_t = \sin\left[(t+1)\arctan\left(\frac{1}{\sqrt{N-1}}\right)\right].$$

We want to find the value of t that maximizes α_t. Setting

$$t^* = \frac{\frac{\pi}{2}}{\arctan\left(\frac{1}{\sqrt{N-1}}\right)} - 1$$

achieves $\alpha_{t^*} = 1$. The series expansion of this formula shows t^* is asymptotically $\frac{\pi}{2}\sqrt{N}+\mathrm{O}(1)$, as desired. However, t must be an integer, so we set the number of iterations to $T = \lfloor t^* \rfloor$. Observe that $\sin(x)$ increases as x approaches $\frac{\pi}{2}$, so it is sufficient to lower bound $\alpha_{t^*-1} \leq \alpha_T$. Substituting and then simplifying, we find

$$\alpha_{t^*-1} = \sin\left[\frac{\pi}{2} - \arctan\left(\frac{1}{\sqrt{N-1}}\right)\right] = \cos\left[\arctan\left(\frac{1}{\sqrt{N-1}}\right)\right] = \sqrt{1-\frac{1}{N}}.$$

So, given the algorithm never terminates early, it outputs $|x^*\rangle$ with probability at least $|\alpha_T|^2 \geq 1 - 1/N$.

Finally, we handle the possibility of the algorithm terminating early. In each iteration, given $|\psi'''_{t-1}\rangle$, the probability of measuring $|1\rangle$ is $\left(\alpha^2 + (N-1)\beta^2\right)/N = 1/N$. Therefore, in $T = \mathrm{O}(\sqrt{N})$ iterations, the probability of aborting is at most a negligible $T/N = \mathrm{O}(1/\sqrt{N})$.

Overall, we have that our algorithm aborts with probability at most $\mathrm{O}(1/\sqrt{N})$, while if it does not abort, then it fails to find $|x^*\rangle$ with probability at most $\mathrm{O}(1/N)$. We conclude that with $T = \frac{\pi}{2}\sqrt{N} + \mathrm{O}(1)$ queries to P_π, we can solve Permutation Inversion with probability $1 - \mathrm{O}(1/\sqrt{N})$. ◀

4 Simulating Other Oracles

This section is omitted due to space constraints and appears in the Full Version.

5 A Subspace-Conversion Separation

This section is omitted due to space constraints and appears in the Full Version.

6 Lower Bounds

In this section, we consider avenues for improving our separations with XOR oracles outperforming in-place oracles to demonstrate a decision-problem separation.

In Section 6.1, we explain the limitations of conventional lower bound techniques for showing that fewer XOR queries are required for a task than in-place queries. In Section 6.2, we introduce a candidate decision problem which we conjecture exhibits such a separation. Then in Section 6.3, we explore recently developed tools for proving lower bounds for arbitrary oracles, including in-place oracles. We develop exact conditions for a decision problem to exhibit a separation. Further details are given in the Appendix of the Full Version.

6.1 Conventional Lower Bound Techniques

In pursuit of proving a decision-problem separation with XOR oracles outperforming in-place oracles, we begin by considering standard tools. Two techniques have dominated quantum query complexity: the polynomial method and the (basic) adversary method. See the thesis of Belovs [11, Chap. 3] for an excellent survey of these tools. Unfortunately, we find that these two methods are insufficient for proving the desired separation.

The remainder of this section is omitted and appears in the full version of this article.

6.2 A Candidate Decision Problem

In this section, we introduce a decision problem called EMBEDDED PERMINV which can be solved with $\Theta(\sqrt{N})$ XOR queries and which we conjecture requires $\Omega(N)$ in-place queries.

Earlier, we showed that in-place query algorithms can achieve the same query complexity as XOR oracles for PERMUTATION INVERSION. As noted in Footnote 3, our algorithm appears to crucially rely on the fact that it is inverting a permutation rather than an injection. The algorithm uses the image of the permutation from one iteration as the input in the next. Now that our goal is to find a problem for which in-place queries are less useful than XOR queries, we leverage this limitation. (Below, S_i is the symmetric group of degree i.)

▶ **Definition 8** (Promised Permutation Inversion). *Given query access to a permutation f on $[N^2] = \{0, \dots, N^2 - 1\}$, the decision problem* EMBEDDED PERMINV : $S_{N^2} \to \{0,1\}$ *is defined by*

$$\text{EMBEDDED PERMINV}(f) = \begin{cases} 1 & \textit{if } f^{-1}(0) \leq N, \textit{ and} \\ 0 & \textit{otherwise.} \end{cases}$$

In effect, this problem embeds an injection from $[N] \to [N^2]$ into a bijection on $[N^2]$, with the promise that an algorithm only needs to search over $[N]$. This problem is inspired by a candidate proposed by Aaronson [2] which was a version of Simon's problem with garbage appended to each query. When querying an XOR oracle, it is easy to copy the desired part of any answer and then uncompute with an additional query, allowing the garbage to be ignored. In contrast, it is unclear how to uncompute or erase the garbage with an in-place oracle, which would prevent interference. Here, instead of Simon's problem we focus on PERMUTATION INVERSION, and we formalize the idea of appending garbage as embedding an injection into a bijection.

▶ **Lemma 9.** EMBEDDED PERMINV *can be decided with at most $\Theta(\sqrt{N})$ XOR queries.*

The proof of Lemma 9 is deferred to the Full Version.

It is unclear how to solve EMBEDDED PERMINV as efficiently as the above algorithm when using in-place queries. Simply querying $\sum |x\rangle \mapsto \sum |f(x)\rangle$ would be useless. One can instead consider algorithms that involve mapping $|x\rangle |0\rangle \mapsto |x\rangle |f(x)\rangle$. Any such query $x \in [N]$ will lead to an unknown element $f(x) \in [N^2]$. Since $f(x)$ may not be in $[N]$, this (a priori) unknown element seems useless for finding the pre-image $f^{-1}(0) \in [N]$. Moreover, in-place oracles are not self-inverse and do not readily allow uncomputing queries. So the image is both useless to keep around and not readily uncomputable using in-place queries. We conjecture this task as a candidate for which XOR oracles outperform in-place oracles.

▶ **Conjecture 10.** EMBEDDED PERMINV *requires at least* $\Omega(N)$ *queries to an in-place oracle.*

Note that even a classical algorithm can solve the problem with N queries by simply querying every element of $[N]$. Also note that while an exponential query separation is possible for Simon's with garbage, the largest separation possible with EMBEDDED PERMINV is polynomial. We hope that the structure of the problem makes a separation more tractable.

6.3 Sketch of Techniques for a Decision Problem Separation

Here, we briefly we explore applying a recent version of the quantum adversary bound to prove the desired decision-problem separation. A full exposition is given in the Appendix of the Full Version.

Quantum query complexity can be characterized by the adversary method. This method has been used to develop several different adversary bounds or adversary theorems in different contexts. For example, prior work derived adversary bounds in the XOR oracle model. In general, an adversary bound for a decision problem $\phi : D \to \{0,1\}$ is an optimization problem such that the optimum is a lower bound on the query complexity. Belovs and Yolcu [13] recently developed a new version of the adversary bound that applies to arbitrary linear transformations. In fact, [13, Section 10] specifically observed this includes in-place oracles in addition to XOR oracles. Moreover, the bound of [13] is tight, meaning the optimum value of the optimization problem corresponds to the optimum query complexity and vice versa.

One caveat is that the lower bound of [13] is for *Las Vegas* query complexity, a quantum analog of the expected number of queries needed for a zero-error algorithm, in contrast to the usual notion of bounded-error complexity. So, our results in this section are primarily focused on Las Vegas complexity. But, for the special case of EMBEDDED PERMINV, we are able to extend the analysis to bounded-error complexity.

The optimization problem in the adversary bound developed by [13] is specifically an optimization over *adversary matrices* Γ. The optimal choice of adversary matrix then corresponds to the optimal query algorithm. In other versions of the adversary method, adversary matrices have been restricted to nonnegative values (the positive weight method) or to general real numbers (the negative weights method). For a decision problem $\phi : D \to \{0,1\}$, previous methods have nearly always restricted Γ such that an entry $\Gamma[f,g]$ indexed by problem instances f and g satisfies that if $\phi(f) = \phi(g)$, then $\Gamma[f,g] = 0$. But, one feature of this new version of the adversary method is that it removes that restriction: we are free to assign nonzero values to *all* entries of Γ.

We call these matrices, with nonzero entries corresponding to problem instances with the same answer, *extended* adversary matrices. We show that, just as negative-weight adversary matrices are necessary to prove tight lower bounds for certain problems, these "extended" adversary matrices are necessary to prove the desired decision-problem separation with XOR oracles outperforming in-place oracles. In other words, if we use only tools from the negative-weight adversary bound to construct adversary matrices Γ, then we cannot prove our desired query separation.

▶ **Theorem** (Informal statement). *For a decision problem* $\phi : D \to \{0,1\}$*, the Las Vegas query complexity using* XOR *oracles is asymptotically less than the Las Vegas query complexity using in-place oracles if and only if optimizing over extended adversary matrices witnesses it.*

Again, the above statement is in terms of Las Vegas complexity instead of the more typical bounded-error complexity. But, for our candidate problem Embedded PermInv introduced in the previous section, we are able to extend the statement to bounded-error complexity

▶ **Theorem** (Informal statement). *For the decision problem* Embedded PermInv*, the bounded-error query complexity using* xor *oracles is asymptotically less than the bounded-error query complexity using in-place oracles if and only if optimizing over extended adversary matrices witnesses it.*

See the Appendix of the Full Version for details. In sum, we considerably narrow down what techniques could possibly prove an $\Omega(N)$ lower bound on Embedded PermInv. Although we rule out the polynomial and unweighted adversary methods, the new adversary method of Belovs and Yolcu [13] is tight, so that if such a lower bound is possible, then it is witnessed by adversary matrices. By the above theorem, we see that any lower bound stronger than $\Omega(\sqrt{N})$ must use this new class of *extended* adversary matrices.

References

1 Scott Aaronson. Quantum lower bound for the collision problem. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 635–642. ACM, 2002. doi:10.1145/509907.509999.

2 Scott Aaronson. Open problems related to quantum query complexity. *ACM Transactions on Quantum Computing*, 2(4), 2021. doi:10.1145/3488559.

3 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. doi:10.4086/toc.2007.v003a007.

4 Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 20–29. ACM, 2003. doi:10.1145/780542.780546.

5 Gorjan Alagic, Chen Bai, Alexander Poremba, and Kaiyan Shi. On the two-sided permutation inversion problem. *IACR Communications in Cryptology*, 1(1), 2024. doi:10.62056/a0qj89n4e.

6 Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 4(64):750–767, 2002. doi:10.1006/jcss.2002.1826.

7 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, 2011. doi:10.1109/CCC.2011.24.

8 Alp Atici. Comparative computational strength of quantum oracles, 2004. arXiv:quant-ph/0312107v3.

9 Roozbeh Bassirian, Bill Fefferman, and Kunal Marwaha. On the power of nonstandard quantum oracles. In Omar Fawzi and Michael Walter, editors, *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023)*, volume 266 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:25. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.TQC.2023.11.

10 Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. doi:10.1145/502090.502097.

11 Aleksandrs Belovs. *Applications of Adversary Method in Quantum Query Algorithms*. PhD thesis, University of Latvia, 2014. URL: https://dspace.lu.lv/dspace/handle/7/4854.

12 Aleksandrs Belovs. Variations on quantum adversary, 2015. arXiv:1504.06943v1.

13 Aleksandrs Belovs and Duyal Yolcu. One-way ticket to Las Vegas and the quantum adversary, 2023. arXiv:2301.02003v1.

14 Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, 1973. doi:10.1147/rd.176.0525.

15 Charles H. Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21:905–940, 1982. doi:10.1007/bf02084158.

16 Charles H. Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989. doi:10.1137/0218053.

17 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. doi:10.1137/S0097539796300933.

18 Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. doi:10.1137/S0097539796300921.

19 Adam Bouland and Tudor Giurgica-Tiron. Efficient universal quantum compilation: An inverse-free Solovay-Kitaev algorithm, 2021. arXiv:2112.02040v1.

20 David Deutsch. Quantum theory, the Church–Turing principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400(1818):97–117, 1985. doi:10.1098/rspa.1985.0070.

21 David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proc. R. Soc. Lond. A*, 439(1907):553–558, 1992. doi:10.1098/rspa.1992.0167.

22 Bill Fefferman and Shelby Kimmel. Quantum vs. classical proofs and subset verification. In Igor Potapov, Paul Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018)*, volume 117 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.MFCS.2018.22.

23 Richard P. Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, 1986. doi:10.1007/BF01886518.

24 Jingliang Gao. Quantum union bounds for sequential projective measurements. *Phys. Rev. A*, 92(5):052331, 2015. doi:10.1103/PhysRevA.92.052331.

25 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. doi:10.1145/237814.237866.

26 Aram Harrow and David Rosenbaum. Uslessness for an oracle model with internal randomness. *Quantum Info. Comput.*, 14(7&8):608–624, 2013. doi:10.26421/QIC14.7-8-5.

27 Atsuya Hasegawa and François Le Gall. An optimal oracle separation of classical and quantum hybrid schemes. In Sang Won Bae and Heejin Park, editors, *33rd International Symposium on Algorithms and Computation (ISAAC 2022)*, volume 248 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ISAAC.2022.6.

28 Elham Kashefi, Adrian Kent, Vlatko Vedral, and Konrad Banaszek. Comparison of quantum oracles. *Phys. Rev. A*, 65:050304, 2002. doi:10.1103/PhysRevA.65.050304.

29 R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961. doi:10.1147/rd.53.0183.

30 Cedric Yen-Yu Lin and Han-Hsuan Lin. Upper bounds on quantum query complexity inspired by the Elitzur–Vaidman bomb tester. *Theory of Computing*, 12(18):1–35, 2016. doi:10.4086/toc.2016.v012a018.

31 Nathan Lindzey and Ansis Rosmanis. A tight lower bound for non-coherent index erasure. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 59:1–59:37. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ITCS.2020.59.

32 David Rasmussen Lolck, Maňinska, and Manaswi Paraashar. Quantum advantage with a faulty oracle, 2024. arXiv:2411.04931v1.

33 Anand Natarajan and Chinmay Nirkhe. A distribution testing oracle separation between QMA and QCMA. *Quantum*, 8:1377, 2024. doi:10.22331/q-2024-06-17-1377.

34 Ashwin Nayak. Inverting a permutation is as hard as unordered search. *Theory of Computing*, 7(2):19–25, 2011. doi:10.4086/toc.2011.v007a002.

35 Ryan O'Donnell and Ramgopal Venkateswaran. The quantum union bound made easy. In *2022 Symposium on Simplicity in Algorithms (SOSA)*, pages 314–320. SIAM, 2022. doi:10.1137/1.9781611977066.25.

36 Asher Peres. Reversible logic and quantum computers. *Phys. Rev. A*, 32:3266–3276, 1985. doi:10.1103/PhysRevA.32.3266.

37 Oded Regev. Impossibility of a quantum speed-up with a faulty oracle. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming (ICALP)*, pages 773–781. Springer Berlin Heidelberg, 2008. doi:10.1007/978-3-540-70575-8_63.

38 Ansis Rosmanis. Tight bounds for inverting permutations via compressed oracle arguments, 2022. arXiv:2103.08975v2.

39 Ansis Rosmanis. Quantum search with noisy oracle, 2023. arXiv:2309.14944v1.

40 Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems, 2001. arXiv:quant-ph/0112086v1.

41 Kristan Temme. Runtime of unstructured search with a faulty Hamiltonian oracle. *Phys. Rev. A*, 90:022310, 2014. doi:10.1103/PhysRevA.90.022310.

A Quantum Cloning Game with Applications to Quantum Position Verification

Léo Colisson Palais
Laboratoire Jean Kuntzmann, Université Grenoble Alpes, France

Llorenç Escolà-Farràs
TU Eindhoven, The Netherlands

Florian Speelman
QuSoft & Informatics Institute, University of Amsterdam, The Netherlands

Abstract

We introduce a quantum cloning game in which k separate collaborative parties receive a classical input, determining which of them has to share a maximally entangled state with an additional party (referee). We provide the optimal winning probability of such a game for every number of parties k, and show that it decays exponentially when the game is played n times in parallel. These results have applications to quantum cryptography, in particular in the topic of quantum position verification, where we show security of the routing protocol (played in parallel), and a variant of it, in the random oracle model.

2012 ACM Subject Classification Security and privacy → Cryptography

Keywords and phrases Quantum position verification, cloning game, random oracle, parallel repetition

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.2

Related Version *Previous Version*: https://arxiv.org/abs/2410.22157

Funding *Léo Colisson Palais*: European Union (ERC, ASC-Q, 101040624) and the Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL programme.
Llorenç Escolà-Farràs: Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL programme.
Florian Speelman: Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL programme.

1 Introduction

Non-local correlations have extensively been studied in the field of quantum information theory, see e.g. [12]. Bell [9] originally showed that distant parties sharing quantum resources can reproduce correlations that could never be attained by any classical theory. Often, non-local correlations are studied as *non-local games*, which provide an operational framework for understanding them. These games are interesting per se from a fundamental point of view, since they give rise to understanding the underlying essence of nature, but they additionally lead to applications such as secure key distribution [1], certified randomness [33], reduced communication complexity [14], self-testing [32, 37], and computation [5].

A vast literature in non-local games covers the scenario where a classical referee sends questions to non-communicating collaborative parties, and their task is to produce answers according to a certain publicly-known predicate, where the questions and answers are all *classical*. The best-known non-local game is the CHSH game [18]. Non-locality has also been investigated in terms of supersets of non-local games, called *monogamy-of-entanglement (MoE) games* [38], where a quantum referee sends the same classical question to the players and the parties have to guess the (classical) outcome of a referee's quantum measurement

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 2; pp. 2:1–2:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

(depending on the question). MoE games have been used to provide security proofs for the quantum cryptographic primitives device-independent quantum key distribution [10] and quantum position verification [28]. Such games were later generalized under the name of *extended non-local games* [27].

Here, we introduce the concept of the *quantum cloning game*, played by k distant parties and a quantum referee. The referee publically announces a party, i.e., sends the same classical question to all the players, and the chosen party has to end up with the maximally entangled (EPR) state with the referee. At the beginning of the game, the players are allowed to share any quantum state with the referee. In this work, we show that the optimal winning probability for players using any quantum resources is given by $\frac{1}{2}+\frac{1}{2k}$, converging to $\frac{1}{2}$ for a large number of players. We analyze the game when it is played n times in parallel, showing an exponential decay in n of the optimal winning probability. Additionally, the quantum cloning game can be generalized to any arbitrary quantum state instead of an EPR state, and we provide its optimal winning probability.

We show that these results have applications in quantum position verification (QPV), which is a cryptographic primitive consisting of verifying the location of an untrusted party. Securely implementing this primitive is unachievable using only classical information, because a general attack exists even when using computational assumptions [17]. Due to the no-cloning theorem [41] the general classical attack does not apply if quantum information is used instead [28, 31], however, a general quantum attack exists which requires exponential entanglement [13, 8]. This means that hope for protocols secure against reasonable amounts of entanglement is alive, and indeed there has been much analysis on attacks on specific protocols [2, 28, 29, 34, 16, 36, 21, 22, 25, 20], and security analysis under extra assumptions [30, 24], such as the random oracle model [39]. A generic 1-dimensional (the main ideas generalize to higher dimensions) QPV protocol is described in the following way: two verifiers V_0 and V_1, placed on the left and right of an untrusted prover P, supposedly at the position *pos*, send quantum and classical messages to P at the speed of light, and he has to pass a challenge and reply correctly to them at the speed of light as well, if so, the verifiers *accept*, and if any of them receives a wrong answer or the timing does not correspond with the time it would have taken for light to travel back from the honest prover, the verifiers *reject*. The time consumed by the prover to perform the challenge is assumed to be negligible, and the verifiers are assumed to have perfectly-synchronized clocks.

In this work, we consider the *routing* QPV protocol [28], which has an appealing simple form: the prover has to return a received qubit to one of the verifiers, where the choice of verifier is a function of the classical information sent by the verifiers [28]. Besides the theoretical interest of this protocol, it is also an appealing candidate for free-space quantum position verification, when the quantum messages can travel with the vacuum speed of light, since the hardware of the prover could hypothetically be as simple as a mirror or an optical switch. Despite theoretical work on this protocol [15, 20, 11, 3, 6], there were gaps left in our understanding relative to measurement-based QPV protocol variants: namely the security of parallel repetition of this protocol against unentangled attackers and attackers who pre-share a linear (in the security parameter) amount of entangled qubits, and its security in the random-oracle model against arbitrary adversaries. As an application of the quantum cloning game, we show the security of the routing protocol in these scenarios.

2 Preliminaries

For $k \in \mathbb{N}$, we will denote $[k] := \{0, \dots, k-1\}$. Let $\mathcal{H}$, $\mathcal{H}'$ be finite-dimensional Hilbert spaces, we denote by $\mathcal{B}(\mathcal{H}, \mathcal{H}')$ the set of bounded operators from $\mathcal{H}$ to $\mathcal{H}'$ and $\mathcal{B}(\mathcal{H}) = \mathcal{B}(\mathcal{H}, \mathcal{H})$. Denote by $\mathcal{S}(\mathcal{H})$ the set of quantum states on $\mathcal{H}$, i.e. $\mathcal{S}(\mathcal{H}) = \{\rho \in \mathcal{B}(\mathcal{H}) \mid \rho \geq 0, \mathrm{Tr}[\rho] = 1)\}$. A pure state will be denoted by a ket $|\psi\rangle \in \mathcal{H}$. The maximally entangled state or EPR pair is $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. We denote the identity matrix by $\mathbb{I}$. For $M \in \mathcal{B}(\mathcal{H})$, we denote its Schatten ∞-norm by $\|M\|$. We will use the notation $1, \dots, \cancel{i}, \dots, k-1$ to denote $1, \dots, i-1, i+1, \dots, k-1$.

▶ **Definition 1.** *Let $N \in \mathbb{N}$. Two permutations $\pi, \pi' : [N] \to [N]$ are said to be* orthogonal *if $\pi(i) \neq \pi'(i)$ for all $i \in [N]$.*

▶ **Lemma 2** (Lemma 2 in [38])**.** *Let $\Pi^1, \dots, \Pi^N$ be projectors acting on a Hilbert space $\mathcal{H}$. Let $\{\pi_k\}_{k \in [n]}$ be a set of mutually orthogonal permutations. Then,*

$$\left\| \sum_{i \in [N]} \Pi^i \right\| \leq \sum_{k \in [N]} \max_{i \in [N]} \left\| \Pi^i \Pi^{\pi_k(i)} \right\|. \tag{1}$$

▶ Remark 3. There always exist a set of N permutations of $[N]$ that are mutually orthogonal, an example is the N cyclic shifts.

▶ **Lemma 4** (Lemma 1 in [38])**.** *Let $A, B, L \in \mathcal{B}(\mathcal{H})$ such that $AA^\dagger \succeq B^\dagger B$. Then it holds that $\|AL\| \geq \|BL\|$.*

3 k-party quantum cloning game

In the following definition, we introduce the quantum cloning game.

▶ **Definition 5.** *The k-party* quantum cloning game, *shortly denoted by QCG_k, consists of a referee R with associated Hilbert space $\mathcal{H}_R = \mathbb{C}^2$ and k collaborative distant parties (players) $P_0, \dots, P_{k-1}$. Before the game starts, the parties prepare a joint quantum state of arbitrary dimension between themselves and the referee. During the game, the referee sends $x \in [k]$, drawn uniformly at random, to all the collaborative parties. The players win the game if and only if the party P_x (holding a qubit register P_x) ends up sharing the maximally entangled state with the referee, i.e. if a projection onto $|\Phi^+\rangle_{RP_x}$ yields the correct outcome.*

See Figure 1 for a schematic representation of the QCG_k. Intuitively, in such a game, the referee publically announces which party has to create an entangled state with herself.

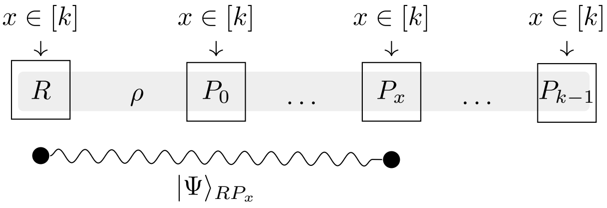

Figure 1 Schematic representation of the k-party quantum cloning game, where $|\Psi\rangle_{RP_x} = |\Phi^+\rangle_{RP_x}$, where the gray-shaded region represents the shared state ρ. If $|\Psi\rangle_{RP_x}$ is arbitrary, this represents a Ψ-QCG_k.

A strategy S for the QCG_k is described by a quantum state $\rho \in \mathcal{S}(\mathcal{H}_R \otimes \mathcal{H}_{P_0E_0} \otimes \cdots \otimes \mathcal{H}_{P_{k-1}E_{k-1}})$, where, for $i \in [k]$, registers P_i are of the same dimension as $\mathcal{H}_R$ and E_i are auxiliary systems of arbitrary dimension that each party possess, and completely positive trace-preserving (CPTP) maps $\{\mathcal{E}^x_{P_iE_i \to P_i}\}_x$, where the subscript $P_iE_i \to P_i$ indicates that the map has input and output registers P_iE_i and P_i, respectively, i.e. $\mathcal{E}^x_{P_iE_i \to P_i} : \mathcal{B}(\mathcal{H}_{P_iE_i}) \to \mathcal{B}(\mathcal{H}_{P_i})$. The winning probability of such a game, given the strategy S, is provided by

$$\omega(\mathrm{QCG}_k, S) = \frac{1}{k} \sum_{x \in [k]} \mathrm{Tr}\left[|\Phi^+\rangle\langle\Phi^+|_{RP_x} \mathrm{Tr}_{P_0 \ldots \not{P}_x \ldots P_{k-1}} \left[\mathbb{I}_R \bigotimes_{i \in [k]} \mathcal{E}^x_{P_iE_i \to P_i}(\rho) \right] \right]. \tag{2}$$

The optimal winning probability of such games is given by

$$\omega^*(\mathrm{QCG}_k) = \sup_S \omega(\mathrm{QCG}_k, S), \tag{3}$$

where the supremum is taken over all the possible strategies over all possible Hilbert spaces. The following theorem gives the optimal winning probability of this game for every number of parties k.

▶ **Theorem 6.** *For every $k \in \mathbb{N}$, the optimal winning probability of the* QCG_k *is given by*

$$\omega^*(\mathrm{QCG}_k) = \frac{1}{2} + \frac{1}{2k}. \tag{4}$$

Intuitively, this game cannot be perfectly won since, otherwise, it would be possible to have maximal entanglement between the referee and each of the parties, and this is not possible since entanglement is *monogamous* [19]. In the proof, see below, the key part is to show that the optimal winning probability is attainable by the actions of the players being *independent* of x, intuitively, each party acts as if they were chosen to reproduce the maximally entangled state with the referee. In addition, in the proof, we show that the optimal value can be attained by preparing an initial state ρ where, locally, each of the parties holds a qubit and no further actions taken by the players, i.e. their local actions are described by the identity channel ($\mathbb{I}_{P_i}$). We then specify a strategy by providing a quantum state, since any local actions are independent of x, they can be absorbed in the quantum state. More precisely, the optimal winning probability for the QCG_k can be attained by the strategy given by the (pure) quantum state

$$|\varphi\rangle = \sqrt{\frac{2}{k(k+1)}} \sum_{x \in [k]} |\Phi^+\rangle_{RP_x} |0\rangle_{P_0 \ldots \not{P}_x \ldots P_{k-1}}. \tag{5}$$

Note that other natural multi-party entangled states that have been widely studied in the literature, such as the GHZ state ([26]) $|GHZ\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$ and the W state ([23]) $|W\rangle = \frac{1}{\sqrt{3}}(|001\rangle + |010\rangle + |100\rangle)$, and their respective generalizations to arbitrary dimensions, as well as the strategy of "guessing" which party has to reproduce the quantum state, e.g. guessing $x = 0$, given by preparing the state $|\Phi^+\rangle_{VP_0}|0\rangle_{P_1} \ldots |0\rangle_{k-1}$, provide significantly suboptimal winning probabilities. For 2-players, $\omega^*(\mathrm{QCG}_2) = \frac{3}{4}$, and

$$\omega^*(\mathrm{QCG}_k) \xrightarrow{k \to \infty} \frac{1}{2}, \tag{6}$$

which converges to the value attained by the strategy given by preparing the state $|0\rangle_R|0\rangle_{P_0} \ldots |0\rangle_{P_{k-1}}$, showing that when k increases even unentangled states allow for a near-optimal winning probability.

Proof. A strategy S for the QCG$_k$ is described by a quantum state $\rho \in \mathcal{S}(\mathcal{H}_R \otimes \mathcal{H}_{P_0E_0} \otimes \ldots \otimes \mathcal{H}_{P_{k-1}E_{k-1}})$, where, for $i \in [k]$, registers P_i are of the same dimension as $\mathcal{H}_R$ and E_i are auxiliary systems of arbitrary dimension that each party possesses, and unitary transformations $U = \{U^x_{P_iE_i}\}_x$, acting on the registers in the subscripts (due to the Stinespring dilation of the quantum channels, we restrict our attention to unitary transformations $U^x_{P_iE_i}$ instead of quantum channels $\mathcal{E}^x_{P_iE_i \to P_i}$). Let d be the dimension of the above (total) Hilbert space, which we denote by $\mathcal{H}_d$. Then, the winning probability of the QCG$_k$, given the strategy S on a d-dimensional Hilbert space, is provided by

$$\begin{aligned}
&\omega(\mathrm{QCG}_k, S, d)\\
&=\frac{1}{k}\sum_{x\in[k]} \mathrm{Tr}\left[\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)\left(\left(\mathbb{I}_R\otimes U^x_{P_iE_i}\otimes\ldots\otimes U^x_{P_iE_i}\right)\rho(\mathbb{I}_R\otimes U^x_{P_iE_i}\otimes\ldots\otimes U^x_{P_iE_i})^\dagger\right)\right]\\
&=\frac{1}{k}\sum_{x\in[k]} \mathrm{Tr}\left[\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_iE_i}\otimes\ldots\otimes U^{x\dagger}_{P_iE_i}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)(\mathbb{I}_R\otimes U^x_{P_iE_i}\otimes\ldots\otimes U^x_{P_iE_i})\rho\right],
\end{aligned}$$

where in the last equation we used cyclicity of the trace. For a specific choice of unitary transformations $U = \{U^x_{P_iE_i}\}_x$, the optimal winning probability is given by

$$\begin{aligned}
&\omega^*(\mathrm{QCG}_k, U, d)\\
&= \sup_{\rho\in\mathcal{S}(\mathcal{H}_d)}\frac{1}{k}\sum_{x\in[k]} \mathrm{Tr}\left[\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_iE_i}\otimes\ldots\otimes U^{x\dagger}_{P_iE_i}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)(\mathbb{I}_R\otimes U^x_{P_iE_i}\otimes\ldots\otimes U^x_{P_iE_i})\rho\right]\\
&= \frac{1}{k}\|\sum_{x\in[k]}\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_iE_i}\otimes\ldots\otimes U^{x\dagger}_{P_iE_i}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)(\mathbb{I}_R\otimes U^x_{P_iE_i}\otimes\ldots\otimes U^x_{P_iE_i})\|\\
&= \frac{1}{k}\|\sum_{x\in[k]}\left(\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_xE_x}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\right)(\mathbb{I}_R\otimes U^x_{P_xE_x})\right)\bigotimes_{i\neq x\in[k]}U^{x\dagger}_{P_iE_i}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\bigotimes_{i\neq x\in[k]}U^x_{P_iE_i}\|\\
&= \frac{1}{k}\|\sum_{x\in[k]}\left(\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_xE_x}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\right)(\mathbb{I}_R\otimes U^x_{P_xE_x})\right)\bigotimes_{i\neq x\in[k]}U^{x\dagger}_{P_iE_i}U^x_{P_iE_i}\|,
\end{aligned}$$

Notice that, since $\{U^x_{P_iE_i}\}_x$ are unitary matrices, $U^{x\dagger}_{P_iE_i}U^x_{P_iE_i} = \mathbb{I}_{P_iE_i}$, moreover, $\mathbb{I}_{P_iE_i} = U^{i\dagger}_{P_iE_i}U^i_{P_iE_i}$, then we can use $U^{x\dagger}_{P_iE_i}U^x_{P_iE_i} = U^{i\dagger}_{P_iE_i}U^i_{P_iE_i}$, and therefore

$$\begin{aligned}
&\omega^*(\mathrm{QCG}_k, U, d)\\
&\quad= \frac{1}{k}\|\sum_{x\in[k]}\left(\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_xE_x}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\right)\left(\mathbb{I}_R\otimes U^{x\dagger}_{P_xE_x}\right)\right)\bigotimes_{i\neq x\in[k]}U^{i\dagger}_{P_iE_i}U^i_{P_iE_i}\|\\
&\quad= \frac{1}{k}\|\sum_{x\in[k]}\left(\mathbb{I}_R\bigotimes_{i\neq\in[k]}U^{i\dagger}_{P_iE_i}\right)\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)\left(\mathbb{I}_R\bigotimes_{i\in[k]}U^i_{P_iE_i}\right)\|\\
&\quad= \frac{1}{k}\|\left(\mathbb{I}_R\bigotimes_{i\neq\in[k]}U^{i\dagger}_{P_iE_i}\right)\left(\sum_{x\in[k]}|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)\left(\mathbb{I}_R\bigotimes_{i\in[k]}U^i_{P_iE_i}\right)\|\\
&\quad= \frac{1}{k}\|\sum_{x\in[k]}|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\|\\
&\quad= \sup_{\rho\in\mathcal{S}(\mathcal{H}_d)}\frac{1}{k}\sum_{x\in[k]}\mathrm{Tr}\left[\left(|\Phi^+\rangle\langle\Phi^+|_{RP_x}\otimes\mathbb{I}_{E_x}\bigotimes_{i\neq x\in[k]}\mathbb{I}_{P_iE_i}\right)\rho\right] = \omega^*(\mathrm{QCG}_k, d) \qquad (7)
\end{aligned}$$

where in the fourth equality we used that the Schatten ∞-norm is unitarily invariant, i.e. $\|V * W\| = \| * \|$ for unitary matrices V and W, and $\omega^*(\mathrm{QCG}_k, d)$ denotes the optimal winning probability if the dimension of the total initial Hilbert space is d. Equation (7) shows that, given a Hilbert space $\mathcal{H}_R \otimes \mathcal{H}_{P_0E_0} \otimes \ldots \otimes \mathcal{H}_{P_{k-1}E_{k-1}}$, the optimal winning probability can be attained by an optimal quantum state independently of the actions of the players after knowing x, i.e. the optimal winning probability is independent of $\{U^x_{P_iE_i}\}_x$ and they can apply $\{\mathbb{I}^x_{P_iE_i}\}_x$. We are going to see that, actually, the optimal winning probability can be attained by each of the parties possessing a qubit (2-dimensional Hilbert space), i.e. by the total Hilbert space being $\mathcal{H}_{2^k} = \bigotimes_{i\in[k]} \mathbb{C}^2$. From (7),

$$\begin{aligned}
\omega^*(\mathrm{QCG}_k) &= \sup_{d\in\mathbb{N}} \omega^*(\mathrm{QCG}_k, d) = \sup_{d\in\mathbb{N}} \frac{1}{k} \| \left(\sum_{x\in[k]} |\Phi^+\rangle\langle\Phi^+|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \right) \bigotimes_{i\in[k]} \mathbb{I}_{E_i} \| \\
&= \sup_{d\in\mathbb{N}} \frac{1}{k} \| \sum_{x\in[k]} |\Phi^+\rangle\langle\Phi^+|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \| \| \bigotimes_{i\in[k]} \mathbb{I}_{E_i} \| \\
&= \sup_{d\in\mathbb{N}} \frac{1}{k} \| \sum_{x\in[k]} |\Phi^+\rangle\langle\Phi^+|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \| \\
&= \sup_{\rho\in\mathcal{S}(\mathcal{H}_{2^k})} \frac{1}{k} \sum_{x\in[k]} \mathrm{Tr}\Big[\Big(|\Phi^+\rangle\langle\Phi^+|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \Big) \rho \Big],
\end{aligned} \tag{8}$$

where, in the arguments of the supremums, the dependence on d is implicit in the auxiliary spaces, which, together with the registers P_i and V, fully describe the total Hilbert space, and thus its dimension.

In order to provide the explicit value for the optimal winning probability, we have that, from (8),

$$\omega^*(\mathrm{QCG}_k) = \frac{1}{k} \| \sum_{x\in[k]} |\Phi^+\rangle\langle\Phi^+|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \| = \frac{1}{2} + \frac{1}{2k}, \tag{9}$$

where the last equation is obtained by direct computation. ◀

3.1 Quantum cloning game with any target state

The concept of QCG_k can be generalized to the case where, instead of the parties having to reproduce EPR pairs with the referee, the state that has to be reproduced is an arbitrary-fixed state, i.e. the referee's Hilbert space $\mathcal{H}_R$ is now of arbitrary dimension, and on input x the party P_x has to generate a given state $|\Psi\rangle_{RP_x}$. Here, the dimension of the registers P_i is the same for all $i \in [k]$. We will refer to such a game as a k-party *quantum cloning game with target state* $|\Psi\rangle$, in short denoted by Ψ-QCG_k, see Figure 1. Notice that this game becomes trivial if the target state $|\Psi\rangle_{RP}$ is a tensor product state. In the following theorem, we provide the optimal winning probability for any Ψ-QCG_k for every number of parties k and for any target state $|\Psi\rangle$.

▶ **Theorem 7.** *The optimal winning probability for every* Ψ-QCG_k is given by

$$\omega^*(\Psi\text{-}\mathrm{QCG}_k) = \frac{1}{k} \| \sum_{x\in[k]} |\Psi\rangle\langle\Psi|_{RP_x} \otimes \mathbb{I}_{P_0\ldots\cancel{P}_x\ldots P_{k-1}} \|. \tag{10}$$

Along the lines of the proof of Theorem 6, the key idea relies on showing that the optimal winning probability can be attained by the actions of the players being independent on x.

Proof. The result follows from the proof of Theorem 6 by repeating the same steps, replacing $|\Phi^+\rangle_{VP_x}$ by $|\Psi\rangle_{VP_x}$, and from (8), we obtain (10). ◀

4 Parallel repetition of QCG_k

A case of particular interest is given when QCG_k is played n times in parallel, denoted by $\mathrm{QCG}_k^{\times n}$. Specifically, we will analyze QCG_2 where now the two collaborative parties, who we rename as Alice and Bob, will receive $x = x_0 \dots x_{n-1} \in \{0,1\}^n$. We denote by $R_0 \dots R_{n-1}$, $A_0 \dots A_{n-1}$ and $B_0 \dots B_{n-1}$ the final (qubit) registers of the referee, Alice and Bob, respectively. The players win if at the end of the game Alice is able to create the maximally entangled state with the referee in all her registers such that $x_i = 0$, and analogously for Bob in all his registers such that $x_i = 1$. See Figure 2 for a schematic representation.

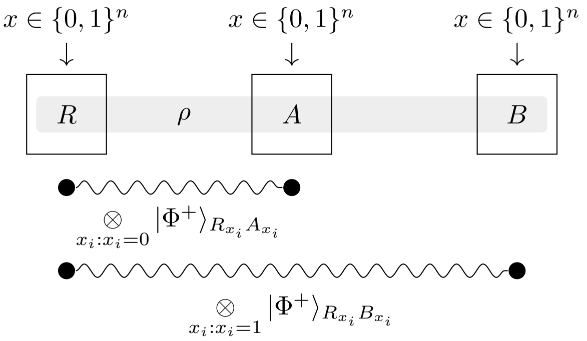
$x \in \{0,1\}^n$
$x \in \{0,1\}^n$
$x \in \{0,1\}^n$
R
ρ
A
B
$\bigotimes_{x_i : x_i = 0} |\Phi^+\rangle_{R_{x_i} A_{x_i}}$
$\bigotimes_{x_i : x_i = 1} |\Phi^+\rangle_{R_{x_i} B_{x_i}}$

Figure 2 Schematic representation of the n-fold parallel repetition of the 2-party quantum cloning game. The gray-shaded region represents the tripartie state ρ that Alice and Bob prepare.

Similarly as before, at the beginning of the game the three parties are allowed to share any arbitrary quantum state and, upon receiving the classical information, Alice and Bob can apply CPTP maps $\{\mathcal{E}^x_{A_0 \dots A_{n-1} E_A \to A_0 \dots A_{n-1}}\}_x$ and $\{\mathcal{E}^x_{B_0 \dots B_{n-1} E_B \to B_0 \dots B_{n-1}}\}_x$, where E_A and E_B are arbitrary auxiliary systems that Alice and Bob possess, respectively.

In the following theorem, we state that the optimal winning probability decays exponentially with the number of parallel repetitions n.

▶ **Theorem 8.** *The optimal winning probability for n parallel repetitions of the QCG_2 is such that*

$$\left(\frac{3}{4}\right)^n \leq \omega^*(\mathrm{QCG}_2^{\times n}) \leq \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n. \tag{11}$$

The key idea of the proof relies on combining ideas used in the proof of Theorem 7 together with Proposition 4.3 in [35], which was also used in [38] to prove parallel repetition for monogamy-of-entanglement games.

Proof. A strategy S_n for the n-parallel repetition of QCG_2 is described by a quantum state $\rho \in \mathcal{S}(\mathcal{H}_R \otimes \mathcal{H}_{A_0 \dots A_{n-1} E_A} \otimes \mathcal{H}_{B_0 \dots B_{n-1} E_B})$, where, for $i \in [n]$, registers A_i and B_i are of the same dimension as $\mathcal{H}_R$ and E_A and E_B are auxiliary systems of arbitrary dimension that each party possess, and unitary transformations $\{U^x_{A_0 \dots A_{n-1} E_A}\}_x$ and $\{V^x_{B_0 \dots B_{n-1} E_B}\}_x$, acting on the registers in the subscripts (due to the Stinespring dilation of the quantum channels, we restrict our attention to unitary transformations). For $x = x_0 \dots x_{n-1} \in \{0,1\}^n$, let $Q_{x_i} = A_i$ if $x_i = 0$ and $Q_{x_i} = B_i$ if $x_i = 1$, and we use the shorthand notation $R = R_0 \dots R_{n-1}$, $A = A_0 \dots A_{n-1}$ and $B = B_0 \dots B_{n-1}$. Then, the winning probability of this game, given the strategy S_n, is provided by

$$\begin{aligned}
&\omega(\mathrm{QCG}_2^{\times n}, S_n)\\
&=\frac{1}{2^n}\sum_{x\in\{0,1\}^n}\mathrm{Tr}\left[\left(\left(\bigotimes_{i\in[n]}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x_i}}\otimes\mathbb{I}_{Q_{1-x_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})\rho(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})^\dagger\right]\\
&=\frac{1}{2^n}\sum_{x\in\{0,1\}^n}\mathrm{Tr}\left[(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x\dagger}_{BE_B})\left(\left(\bigotimes_{i\in[n]}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x_i}}\otimes\mathbb{I}_{Q_{1-x_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})\rho\right]\\
&\leq\frac{1}{2^n}\|\sum_{x\in\{0,1\}^n}(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x\dagger}_{BE_B})\left(\left(\bigotimes_{i\in[n]}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x_i}}\otimes\mathbb{I}_{Q_{1-x_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})\|.
\end{aligned}$$

Denote

$$M^x := (\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x\dagger}_{BE_B})\left(\left(\bigotimes_{i\in[n]}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x_i}}\otimes\mathbb{I}_{Q_{1-x_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B}), \tag{12}$$

then,

$$\omega(\mathrm{QCG}_2^{\times n}, S_n)\leq\frac{1}{2^n}\|\sum_{x\in\{0,1\}^n}M^x\|\leq\frac{1}{2^n}\sum_{k\in[2^n]}\max_{x,x'}\|M^xM^{x'}\|, \tag{13}$$

where we used Lemma 2, and $x'=\pi_k(x)$, for $\{\pi_k\}_k$ being a set of mutually orthogonal permutations. Fix x and x', and let $\mathcal{T}$ be the set of indices where x and x' differ, i.e. $\mathcal{T}=\{i\mid x_i\neq x'_i\}$, and let $t=|\mathcal{T}|$. Let $\mathcal{T}_A=\{i\in\mathcal{T}\mid x_i=0\}$, and $t_A:=|\mathcal{T}_A|$, then we have that

$$\begin{aligned}
&M^x\preceq M^x_A\\
&:=(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x\dagger}_{BE_B})\\
&\qquad\left(\left(\bigotimes_{i\in\mathcal{T}_A}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x_i}}\otimes\mathbb{I}_{Q_{1-x_i}}\right)\otimes\left(\bigotimes_{i\in[n]\setminus\mathcal{T}_A}\mathbb{I}_{R_iQ_{x_i}Q_{1-x_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})\\
&=(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x\dagger}_{BE_B})\left(\left(\bigotimes_{i\in\mathcal{T}_A}|\Phi^+\rangle\langle\Phi^+|_{R_iA_i}\bigotimes_{i\in[n]\setminus\mathcal{T}_A}\mathbb{I}_{R_iA_iE_A}\right)\otimes\mathbb{I}_{BE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^x_{BE_B})\\
&=(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x'\dagger}_{BE_B})\left(\left(\bigotimes_{i\in\mathcal{T}_A}|\Phi^+\rangle\langle\Phi^+|_{R_iA_i}\bigotimes_{i\in[n]\setminus\mathcal{T}_A}\mathbb{I}_{R_iA_iE_A}\right)\otimes\mathbb{I}_{BE_B}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^{x'}_{BE_B}),
\end{aligned}$$

where in the last equality we used that $V^{x\dagger}_{BE_B}V^{x\dagger}_{BE_B}=\mathbb{I}_{BE_B}=V^{x'\dagger}_{BE_B}V^{x'}_{BE_B}$. Similarly,

$$\begin{aligned}
&M^{x'}\preceq M^{x'}_B\\
&:=(\mathbb{I}_R\otimes U^{x'\dagger}_{AE_A}\otimes V^{x'\dagger}_{BE_B})\\
&\qquad\left(\left(\bigotimes_{i\in\mathcal{T}_A}|\Phi^+\rangle\langle\Phi^+|_{R_iQ_{x'_i}}\otimes\mathbb{I}_{Q_{1-x'_i}}\right)\otimes\left(\bigotimes_{i\in[n]\setminus\mathcal{T}_A}\mathbb{I}_{R_iQ_{x'_i}Q_{1-x'_i}}\right)\otimes\mathbb{I}_{E_AE_B}\right)(\mathbb{I}_R\otimes U^{x'}_{AE_A}\otimes V^{x'}_{BE_B})\\
&=(\mathbb{I}_R\otimes U^{x\dagger}_{AE_A}\otimes V^{x'\dagger}_{BE_B})\left(\left(\bigotimes_{i\in\mathcal{T}_A}|\Phi^+\rangle\langle\Phi^+|_{R_iB_i}\bigotimes_{i\in[n]\setminus\mathcal{T}_A}\mathbb{I}_{R_iB_iE_B}\right)\otimes\mathbb{I}_{AE_A}\right)(\mathbb{I}_R\otimes U^x_{AE_A}\otimes V^{x'}_{BE_B}),
\end{aligned}\tag{14}$$

By Lemma 4,

$$\|M^xM^{x'}\|\leq\|M^x_AM^{x'}_B\|, \tag{15}$$

then

$$\begin{aligned}
&M_A^x M_B^{x'}\\
&= (\mathbb{I}_R \otimes U^{x\dagger}_{AE_A} \otimes V^{x'\dagger}_{BE_B})\left(\left(\bigotimes_{i\in\mathcal{T}_A} |\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iA_iE_A}\right) \otimes \mathbb{I}_{BE_B}\right)(\mathbb{I}_R \otimes U^{x}_{AE_A} \otimes V^{x'}_{BE_B})\\
&\cdot (\mathbb{I}_R \otimes U^{x\dagger}_{AE_A} \otimes V^{x'\dagger}_{BE_B})\left(\left(\bigotimes_{i\in\mathcal{T}_A} |\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iB_iE_B}\right) \otimes \mathbb{I}_{AE_A}\right)(\mathbb{I}_R \otimes U^{x}_{AE_A} \otimes V^{x'}_{BE_B}).
\end{aligned}$$

We have that $(\mathbb{I}_R \otimes U^{x}_{AE_A} \otimes V^{x'}_{BE_B})(\mathbb{I}_R \otimes U^{x\dagger}_{AE_A} \otimes V^{x'\dagger}_{BE_B}) = \mathbb{I}_{RAEABE_B}$, and, since the Schatten ∞-norm is unitarily invariant,

$$\begin{aligned}
&\|M_A^x M_B^{x'}\|\\
&= \left\|\left(\bigotimes_{i\in\mathcal{T}_A} |\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iA_iE_A} \otimes \mathbb{I}_{BE_B}\right)\left(\bigotimes_{i\in\mathcal{T}_A} |\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iB_iE_B} \otimes \mathbb{I}_{AE_A}\right)\right\|\\
&= \left\|\left(\bigotimes_{i\in\mathcal{T}_A} \big(|\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \otimes \mathbb{I}_{B_i}\big)\big(|\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \otimes \mathbb{I}_{A_i}\big)\right) \bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iA_iB_i} \otimes \mathbb{I}_{E_AE_B}\right\|\\
&= \left\|\bigotimes_{i\in\mathcal{T}_A} \big(|\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \otimes \mathbb{I}_{B_i}\big)\big(|\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \otimes \mathbb{I}_{A_i}\big)\right\| \left\|\bigotimes_{i\in[n]\setminus\mathcal{T}_A} \mathbb{I}_{R_iA_iB_i} \otimes \mathbb{I}_{E_AE_B}\right\|\\
&= \prod_{i\in\mathcal{T}_A} \left\|\big(|\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \otimes \mathbb{I}_{B_i}\big)\big(|\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \otimes \mathbb{I}_{A_i}\big)\right\|\\
&= 2^{-t_A},
\end{aligned} \tag{16}$$

where we used that, for every i,

$$\left\|\big(|\Phi^+\rangle\langle\Phi^+|_{R_iA_i} \otimes \mathbb{I}_{B_i}\big)\big(|\Phi^+\rangle\langle\Phi^+|_{R_iB_i} \otimes \mathbb{I}_{A_i}\big)\right\| = 2^{-1}. \tag{17}$$

Without loss of generality, assume $t_A \geq t/2$, then, combining (15) and (16), we have that

$$\|M^x M^{x'}\| \leq \|M_A^x M_B^{x'}\| \leq 2^{-\frac{t}{2}}. \tag{18}$$

In order to apply the bound in Lemma 4, consider the set of permutations given by $\pi_k(x) = x \oplus k$, where $x, k \in \{0,1\}^n$ (they are such that they have the same Hamming distance). There are $\binom{n}{i}$ permutations with Hamming distance i. Then, we have

$$\omega(\mathrm{QCG}_2^{\times n}, S_n) \leq \frac{1}{2^n}\sum_{k\in[2^n]} \max_{x,x'} \|M^x M^{x'}\| \leq \frac{1}{2^n}\sum_{t=0}^{n}\binom{n}{t}2^{-\frac{t}{2}} = \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n. \tag{19}$$

◀

5 Application to QPV in the No-PE and BE(m) models

In this section, we analyze the security of the *routing* QPV protocol, originally introduced in [28]. A round of this protocol, see Figure 3 for a schematic representation, is described as follows:

1. The verifier V_0 selects a qubit $|\phi\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$, and the verifier V_1 selects $x \in \{0,1\}$, both picked uniformly at random. They send $|\phi\rangle$ and x (at time $t = 0$) so that they arrive at the same time ($t = 1$) at *pos*.
2. Upon receiving the information sent by V_0 and V_1, the prover sends the qubit $|\phi\rangle$ to the verifier V_x.

3. If $|\phi\rangle$ arrives at the time consistent with *pos* ($t = 2$), and a projective measurement performed by V_x on the state sent by V_0 leads to the correct outcome, the verifiers accept. Otherwise, they reject.

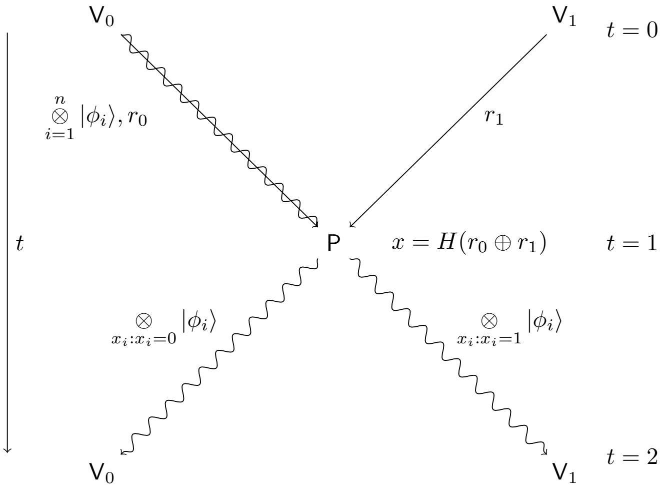
V_0
V_1
$t = 0$
$\bigotimes_{i=1}^{n} |\phi_i\rangle, r_0$
r_1
t
P
$x = H(r_0 \oplus r_1)$
$t = 1$
$\bigotimes_{x_i : x_i = 0} |\phi_i\rangle$
$\bigotimes_{x_i : x_i = 1} |\phi_i\rangle$
V_0
V_1
$t = 2$

Figure 3 Schematic representation of the (H, n)-routing QPV protocol. If r_0 is an empty bit string, and $x = r_1$, this figure represents the n-parallel repetition of the routing QPV protocol. The time arrow is represented by t.

The most general attack to the routing protocol, pictured in Figure 4, consists of having two attackers Alice (A) and Bob (B), located between V_0 and P, and between P and V_0, respectively. Before $t = 0$, the attackers agree on a strategy and might prepare an entangled state. After $t = 0$, Alice (A_0) and Bob (B_0) intercept the information sent from their closest verifier, respectively. Due to timing constraints, they are allowed to perform one round of simultaneous communication. After communicating (after $t = 1$), Alice (A_1) and Bob (B_1) answer to their closest verifier, respectively.

Here, we analyze security within three attack models: (i) the *No Pre-shared Entanglement* (No-PE) model [13], where adversaries do not (pre-)share any entanglement before the protocol's execution; (ii) the *Bounded-Entanglement* BE(m) model, where adversaries pre-share at most m entangled qubits; and (iii) the *Random Oracle Model* (ROM), where attackers (pre-)share any amount of entanglement before the protocol's execution. We formalize the concept of security, given an attack model $\mathcal{M}$, as follows:

▶ **Definition 9.** *The routing protocol is said to be* α-sound *in the* $\mathcal{M}$ *model if, for any attackers acting according to such an attack model, the verifiers* accept *with probability at most* α.

The security of a variation of this protocol, the f-routing QPV protocol, where the classical information x is split into two bit strings, each sent from each verifier, and the qubit has to be routed according to the outcome of a boolean function f on those bit strings has been studied in the BE(m) model [11, 6, 7]. The authors of these works showed that the f-routing QPV protocol remains secure as long as m is at most linear in the size of the bit strings. However, unlike other protocols [13, 38, 4] the security of the routing QPV protocol in the No-PE model was never analyzed. The No-PE assumption is necessary to obtain non-trivial bounds, since there is a perfect attack if the attackers pre-share entanglement [28].

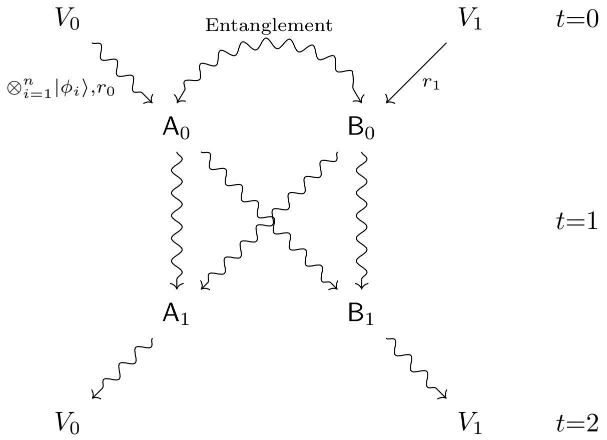

Figure 4 Schematic representation of a generic attack to the (H, n)-routing protocol (and in particular, to the routing protocol).

We show security in the No-PE model, providing the tight result, summarized in the following proposition:

▶ **Proposition 10.** *In the No-PE model, the routing QPV protocol is $\frac{3}{4}$-sound. Moreover, this is optimal.*

The intuition behind Proposition 10 relies on the fact that the most general attack can be reduced to a QCG$_2$. Consider the purified version of the routing protocol, which is equivalent to the original version, and where the only difference relies on V_0, instead of sending the qubit $|\phi\rangle$, prepares the state $|\Phi^+\rangle$ and keeps a register for herself and sends the other register to the prover. Then, as seen in Figure 4, the most general attack to the routing QPV protocol is to place an adversary between V_0 and the prover, Alice, and another adversary between the prover and V_1, Bob. In the No-PE model, we can simplify it further, as Alice intercepts the qubit sent by V_0, applies an arbitrary quantum operation to it, and possibly some ancillary systems she possesses. She keeps a part of it and sends the other to Bob. On the other side, Bob intercepts x, copies it and sends the copy to Alice. Since they share no entanglement, any quantum operation that Bob could perform as a function of x can be included in Alice's operation. After one-round of simultaneous communication, Alice and Bob share a tripartite state with V_0, and their task is that the party designated by x has to end up with a maximally entangled state with the V_0. By Theorem 7, even if Alice and Bob can share any state with the referee (in this case V_0), they can succeed with at most probability $\frac{3}{4}$.

On the other hand, to show optimality, consider the attack where (i) at the beginning of the protocol Alice prepares the 3-qubit state $\frac{1}{\sqrt{3}}(|\Phi^+\rangle_{A_0A}|0\rangle_B + |\Phi^+\rangle_{A_0B}|0\rangle_A)$, (ii) intercepts $|\phi\rangle$ and performs a Bell measurement on the intercepted state and her register A_0, immediately she applies the teleportation corrections to both of her registers A and B, (iii) she keeps register A and sends register B to Bob, (iv) in the meantime, Bob intercepts and broadcasts x, after receiving the information from their fellow attacker, if $x = 0$, Alice sends her register (A) to V_0, whereas if $x = 1$, Bob sends his register (B) to V_1. This attack has a winning probability of $\frac{3}{4}$.

An analogous reduction applies when the routing QPV protocol is executed n times in parallel, and therefore, its security can be reduced to the n-parallel repetition of QCG$_2$:

▶ **Proposition 11.** *In the No-PE model, the routing QPV protocol executed n times in parallel is* $(\frac{1}{2}+\frac{1}{2\sqrt{2}})^n$*-sound.*

A direct consequence of Lemma 5.3 in [8] implies, similarly as in [38], security for the routing protocol executed in parallel for attackers who pre-share a linear amount of qubits:

▶ **Corollary 12.** *In the* $BE(m)$ *model, the routing QPV protocol executed n times in parallel is* $\left(2^m(\frac{1}{2}+\frac{1}{2\sqrt{2}})^n\right)$*-sound.*

In particular, the above soundness is exponentially small in n if $m < n\log\left((\frac{1}{2}+\frac{1}{2\sqrt{2}})^{-1}\right) \simeq 0.228n$.

6 Application to QPV in the random oracle model

Consider the n-parallel repetition of the routing QPV protocol but instead of V_1 sending $x \in \{0,1\}^n$, V_0 and V_1 send $r_0, r_1 \in \{0,1\}^\ell$, for $\ell \in \mathbb{N}$, to the prover, respectively. Then, the x used in the rest of the protocol is computed via $x = H(r_0 \oplus r_1)$, for a given hash function $H : \{0,1\}^\ell \to \{0,1\}^n$. We will denote this variation as (H,n)-routing QPV protocol, see Figure 3 for a schematic representation.

To provide security in the quantum random oracle model against adversaries sharing an arbitrary amount of entanglement, we use some techniques introduced in [40]. A quantum random oracle is defined as a fixed function $H\colon \{0,1\}^\ell \to \{0,1\}^n$ that is sampled uniformly at random from the set of functions from ℓ bits to n bits[1]. The parties are not given the full description of H directly, but they are given oracle access to H, in the sense that they have access to a special gate implementing the unitary $U_H\colon |r\rangle|b\rangle \to |r\rangle|b \oplus H(r)\rangle$. We denote the number of queries made by the adversary by q. As a proof technique, an oracle can also be reprogrammed, where security of the protocol is shown by first studying a variant where the gate applied by the oracle may change over time. A typical setting is where we change a single entry of the oracle: we denote by $H[r \mapsto x]$ the new oracle that behaves like H except that $H(r) = x$. The chances of distinguishing whether H has been reprogrammed or not can be bounded using [40], which informally states that if we can distinguish whether the oracle has been reprogrammed or no, then we have queried it on r before it has been reprogrammed (for completeness, see Lemma 14). In the following theorem, we show security of the (H,n)-routing protocol in the ROM.

▶ **Theorem 13.** *If the (possibly entangled) attackers Alice and Bob perform at most q queries to the (quantum) random oracle H, the* (H,n)*-routing QPV protocol is* ϵ*-sound, with*

$$\epsilon = 2q2^{-\frac{\ell}{2}} + \left(\frac{1}{2}+\frac{1}{2\sqrt{2}}\right)^n. \tag{20}$$

In particular, ϵ *is negligible if q and* ℓ *scale polynomially with n.*

The starting idea of the proof follows [40], where we send Bell pairs instead of single qubits in order to make the input state independent of x, the output of the oracle. Then we reprogram this oracle *after* adversaries share their state to ensure x is truly random and independent of

[1] This can be done by simply sampling a large table T of size 2^ℓ, and outputting $T[r]$ when queried on input r. Note that this sampling procedure is not efficient: while having an efficient oracle [42] is sometimes required, for instance when working with composable security and computationally bounded distinguishers, or when reducing to problems that are hard only for bounded adversaries, in our case we do not need this additional property since we do a reduction to a problem that is hard even for unbounded adversaries.

the state shared by malicious parties at time $t = 1$. Finally, we realize that we can rewrite this into an instance of Theorem 8. In the proof of Theorem 13, we will rely on this lemma by Unruh:

▶ **Lemma 14** ([40, Lemma 3]). *Let $(\mathcal{A}_1, \mathcal{A}_2)$ be oracle algorithms sharing state between invocations that perform at most q queries to H. Let C be an oracle algorithm that on input (j, r) does the following: Run $\mathcal{A}_1^H(r)$ until the j-th query to H, then measure the argument of that query in the computational basis, and output the measurement outcome (or $\perp$ if no j-th query occurs). Let:*

$$P_{\mathcal{A}}^1 := \Pr_{\substack{H \overset{\$}{\leftarrow} (\{0,1\}^\ell \to \{0,1\}^n) \\ r \overset{\$}{\leftarrow} \{0,1\}^l, \mathcal{A}_1^H(r) \\ b' \to \mathcal{A}_2^H(r, H(r))}} [\, b' = \mathsf{Accept} \,], \tag{21}$$

$$P_{\mathcal{A}}^2 := \Pr_{\substack{H \overset{\$}{\leftarrow} (\{0,1\}^\ell \to \{0,1\}^n) \\ r \overset{\$}{\leftarrow} \{0,1\}^l, x \overset{\$}{\leftarrow} \{0,1\}^n \\ H' := H[r \mapsto x], \mathcal{A}_1^H(r) \\ b' \to \mathcal{A}_2^{H'}(r, x)}} [\, b' = \mathsf{Accept} \,], \tag{22}$$

$$P_C := \Pr_{\substack{H \overset{\$}{\leftarrow} (\{0,1\}^\ell \to \{0,1\}^n) \\ r \overset{\$}{\leftarrow} \{0,1\}^l, j \overset{\$}{\leftarrow} \{1, \ldots, q\} \\ x' \to C^H(j, x)}} [\, x = x' \,]. \tag{23}$$

Then, $|P_{\mathcal{A}}^1 - P_{\mathcal{A}}^2| \le 2q\sqrt{P_C}$.

Proof (of Theorem 13). To prove this theorem, we must show that the probability that the verifiers accept in a malicious run of the protocol is lower bounded by ϵ, i.e., if we denote by $\mathsf{V}_0 \leftrightsquigarrow \mathsf{A} \leftrightsquigarrow \mathsf{B} \leftrightsquigarrow \mathsf{V}_1$ the output of the verifiers (Accept or Reject) at the end of a protocol involving a malicious Alice A and a malicious Bob B, we want to show that

$$\Pr[\, \mathsf{V}_0 \leftrightsquigarrow \mathsf{A} \leftrightsquigarrow \mathsf{B} \leftrightsquigarrow \mathsf{V}_1 = \mathsf{Accept} \,] \le \epsilon. \tag{24}$$

We prove this by defining a series of games, where the probability of accepting each game is close to the probability of accepting the next game. By ensuring that the first game corresponds to the real protocol, and that the probability of the last game can easily be computed, we can bound ϵ by transitivity.

Game_1. This game is defined as the real protocol, i.e. $\mathsf{Game}_1 := \mathsf{V}_0 \leftrightsquigarrow \mathsf{A} \leftrightsquigarrow \mathsf{B} \leftrightsquigarrow \mathsf{V}_1$. Therefore, we trivially have:

$$\Pr[\, \mathsf{V}_0 \leftrightsquigarrow \mathsf{A} \leftrightsquigarrow \mathsf{B} \leftrightsquigarrow \mathsf{V}_1 = \mathsf{Accept} \,] = \Pr[\, \mathsf{Game}_1 = \mathsf{Accept} \,]. \tag{25}$$

Game_2. Is like Game_1, except that each $|\phi_i\rangle$ is replaced with one half of a Bell pair. Similarly, instead of projecting on $|\phi_i\rangle$, the verifier will do a Bell measurement between the state sent by the prover and its corresponding half of Bell pair, accepting only if the outcome is $(0, 0)$. This trick is often used in literature, hence we skip the computations. Hence

$$\Pr[\, \mathsf{Game}_1 = \mathsf{Accept} \,] = \Pr[\, \mathsf{Game}_2 = \mathsf{Accept} \,]. \tag{26}$$

Game_3. Is like Game_2, except that at time $t = 1$, one samples the random bit string $x \overset{\$}{\leftarrow} \{0,1\}^n$, and reprogram the oracle to implement $H' := H[r_0 \oplus r_1 \mapsto x]$ (i.e., A_1 and B_1 will have oracle access to H' instead of H). Note that the simulators will use this value of

$x = H'(r_0 \oplus r_1)$ instead of $H(r_0 \oplus r_1)$ to perform the verification at the end. Intuitively, the only way to distinguish this game from the previous game is if the adversary managed to query $H(r_0 \oplus r_1)$ before $t = 0$ and after, but this is highly unlikely since neither A_0 nor B_0 know both r_0 and r_1 (and remember that they cannot query the oracle more than q times, so they cannot just evaluate the oracle on all inputs). This intuition is formalized thanks to Lemma 14: if we define $\mathcal{A}_1(r)$ as the execution of the protocol in Game_2 until $t = 1$ (which is the same as in Game_3), except that r_2 is chosen as $r_2 \coloneqq r_1 \oplus r$, and $\mathcal{A}_2(r, x)$ as the execution of the protocol after time $t = 1$, we can remark that (using notations from Lemma 14):

$$P^1_{\mathcal{A}} = \Pr[\mathsf{Game}_2 = \mathsf{Accept}]. \tag{27}$$

since sampling (r_1, r_2) uniformly at random is strictly equivalent to sampling (r_1, r) randomly and then defining $r_2 \coloneqq r \oplus r_1$. Similarly, we also have:

$$P^2_{\mathcal{A}} = \Pr[\mathsf{Game}_3 = \mathsf{Accept}]. \tag{28}$$

The remaining part is to bound P_C. To compute P_C, we need to bound the probability of querying $H(r)$ during the first part of the protocol on j-th query. But when A_0 does this query, we know that it must be independent of r since all inputs of A_0 are independent of r (if not, we could break non-signaling). Similarly, queries made by B_0 are independent of r: the exact same argument does not hold since $r_2 = r_1 \oplus r$ does depend on r… but this is only a very superficial dependency, since we could have exactly the same probability distributions of r_1 and r_2 by sampling instead r_2 randomly and $r_1 = r_2 \oplus r$, making r_2 independent of r now. Hence, the j-th query is independent of r, so the best probability of it being equal to r is lower bounded by $P_C \leq \frac{1}{2^\ell}$. Hence, using Lemma 14, we have:

$$|\Pr[\,\mathsf{Game}_2 = \mathsf{Accept}\,] - \Pr[\,\mathsf{Game}_3 = \mathsf{Accept}\,]| \tag{29}$$

$$\overset{(28)}{=} |P^1_{\mathcal{A}} - P^2_{\mathcal{A}}| \overset{(14)}{\leq} 2q\sqrt{P_C} \leq 2q2^{-\ell/2}. \tag{30}$$

Game_4. Now, we can realize that all operations in Game_3 until $t = 1$ is independent of x. So let us call $|\psi\rangle_{RP_AP_B}$ the (purification) of the state owned by the verifier (consisting in a list of qubits part of a shared Bell pair), Alice and Bob (we also include in P_B the message sent by A_0 to B_1, and similarly in P_A the message sent by B_0 to A_1). Additionally, we also include in $|\psi\rangle$ the (exponentially large) definition of H, r_0 and r_1 in both registers P_A and P_B, one copy for each party. Then, we define the referee operation R as the identity, P_A as the map that runs A_1, simulating the query to H' using the table H, r_0 and r_1 that are part of $|\psi\rangle$, and x that is given as an input to P_A, and we define similarly P_B simulating B_1. We define then Game_4 as the game $\mathrm{QCG}_2^{\times n}$, i.e. the parallel repetition of the 2-party quantum cloning game (Definition 5), involving the shared state $|\psi\rangle$, the referee R, and the two parties P_A and P_B. This game is exactly like Game_3 as we simulate exactly the same process, just grouping differently the various circuits involved. Hence, $\Pr[\,\mathsf{Game}_4 = \mathsf{Accept}\,] = \Pr[\,\mathsf{Game}_3 = \mathsf{Accept}\,]$. But using Theorem 8, we have:

$$\Pr[\,\mathsf{Game}_4\,] = \omega^*(\mathrm{QCG}_2^{\times n}) \overset{(8)}{\leq} \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n. \tag{31}$$

Hence, we can combine all the above equations to obtain:

$$\Pr[\,\mathsf{V}_0 \leftrightsquigarrow \mathsf{A} \leftrightsquigarrow \mathsf{B} \leftrightsquigarrow \mathsf{V}_1 = \mathsf{Accept}\,] = \Pr[\,\mathsf{Game}_1 = \mathsf{Accept}\,] \leq \left(\frac{1}{2} + \frac{1}{2\sqrt{2}}\right)^n + 2q2^{-\ell/2}, \tag{32}$$

concluding the proof. ◀

7 Discussion

We have introduced the concept of the k-party quantum cloning game and provided the optimal winning probability for any number of parties. The parallel repetition for the two-party version was studied, showing an exponential decay of the optimal winning probability. We applied the above results to show security of the routing QPV protocol in the *No Pre-shared Entanglement* and *Bounded-Entanglement* models, as well as in the *Random Oracle Model*. The tightness of Theorem 8 remains an open question, either by showing a strategy attaining the value (11), or if strong parallel repetition holds and actually the optimal value is $\left(\frac{3}{4}\right)^n$ (or neither of them). Closing this gap would imply knowing what is the optimal security for the routing protocol in the No-PE model, and would further tighten its security in the BE(m) and random-oracle models.

References

1 Antonio Acín, Nicolas Brunner, Nicolas Gisin, Serge Massar, Stefano Pironio, and Valerio Scarani. Device-Independent Security of Quantum Cryptography against Collective Attacks. *Physical Review Letters*, 98(23):230501, June 2007. `doi:10.1103/PhysRevLett.98.230501`.

2 Tomothy Spiller Adrian Kent, William Munro and Raymond Beausoleil. Tagging systems. us patent nr 2006/0022832, 2006.

3 Rene Allerstorfer, Harry Buhrman, Alex May, Florian Speelman, and Philip Verduyn Lunel. Relating non-local quantum computation to information theoretic cryptography. *Quantum*, 8:1387, June 2024. `doi:10.22331/q-2024-06-27-1387`.

4 Rene Allerstorfer, Harry Buhrman, Florian Speelman, and Philip Verduyn Lunel. On the role of quantum communication and loss in attacks on quantum position verification, 2022. `arXiv:2208.04341`.

5 Janet Anders and Dan E. Browne. Computational Power of Correlations. *Physical Review Letters*, 102(5):050502, February 2009. `doi:10.1103/PhysRevLett.102.050502`.

6 Vahid Asadi, Richard Cleve, Eric Culf, and Alex May. Linear gate bounds against natural functions for position-verification, 2024. `doi:10.22331/q-2025-01-21-1604`.

7 Vahid Asadi, Eric Culf, and Alex May. Rank lower bounds on non-local quantum computation. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 11:1–11:18, 2025. `doi:10.4230/LIPIcs.ITCS.2025.11`.

8 Salman Beigi and Robert König. Simplified instantaneous non-local quantum computation with applications to position-based cryptography. *New Journal of Physics*, 13(9):093036, September 2011. `doi:10.1088/1367-2630/13/9/093036`.

9 J. S. Bell. On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1:195–200, November 1964. `doi:10.1103/PhysicsPhysiqueFizika.1.195`.

10 Charles H. Bennett and Gilles Brassard. Quantum cryptography: public key distribution and coin tossing. *Proc. IEEE Int. Conf. on Computers, Systems, and Signal Process (Bangalore) (Piscataway, NJ: IEEE)*, 1984.

11 Andreas Bluhm, Matthias Christandl, and Florian Speelman. A single-qubit position verification protocol that is secure against multi-qubit attacks. *Nature Physics*, pages 1–4, 2022.

12 Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality. *Rev. Mod. Phys.*, 86:419–478, April 2014. `doi:10.1103/RevModPhys.86.419`.

13 Harry Buhrman, Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, Rafail Ostrovsky, and Christian Schaffner. Position-based quantum cryptography: Impossibility and constructions. *SIAM Journal on Computing*, 43(1):150–178, January 2014. `doi:10.1137/130913687`.

14 Harry Buhrman, Richard Cleve, Serge Massar, and Ronald de Wolf. Nonlocality and communication complexity. *Reviews of Modern Physics*, 82(1):665–698, March 2010. `doi:10.1103/RevModPhys.82.665`.

15 Harry Buhrman, Serge Fehr, Christian Schaffner, and Florian Speelman. The garden-hose model. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science - ITCS '13*. ACM Press, 2013. doi:10.1145/2422436.2422455.

16 Kaushik Chakraborty and Anthony Leverrier. Practical position-based quantum cryptography. *Physical Review A*, 92(5), November 2015. doi:10.1103/physreva.92.052304.

17 Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. Position based cryptography. In *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference*, volume 5677 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2009. doi:10.1007/978-3-642-03356-8_23.

18 John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett. 23*, 1969.

19 Valerie Coffman, Joydip Kundu, and William K. Wootters. Distributed entanglement. *Physical Review A*, 61(5), April 2000. doi:10.1103/physreva.61.052306.

20 Sam Cree and Alex May. Code-routing: a new attack on position-verification. *Quantum*, 7, 1079, 2022. doi:10.22331/q-2023-08-09-1079.

21 Kfir Dolev. Constraining the doability of relativistic quantum tasks. *arXiv preprint*, 2019. arXiv:1909.05403.

22 Kfir Dolev and Sam Cree. Non-local computation of quantum circuits with small light cones. *arXiv preprint*, 2022. arXiv:2203.10106.

23 W. Dür, G. Vidal, and J. I. Cirac. Three qubits can be entangled in two inequivalent ways. *Phys. Rev. A*, 62:062314, November 2000. doi:10.1103/PhysRevA.62.062314.

24 Fei Gao, Bin Liu, and QiaoYan Wen. Quantum position verification in bounded-attack-frequency model. *SCIENCE CHINA Physics, Mechanics & Astronomy*, 59(11):1–11, 2016.

25 Alvin Gonzales and Eric Chitambar. Bounds on instantaneous nonlocal quantum computation. *IEEE Transactions on Information Theory*, 66(5):2951–2963, 2019.

26 Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. Going beyond bell's theorem. In Menas Kafatos, editor, *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, pages 69–72. Springer Netherlands, Dordrecht, 1989. doi:10.1007/978-94-017-0849-4_10.

27 Nathaniel Johnston, Rajat Mittal, Vincent Russo, and John Watrous. Extended nonlocal games and monogamy-of-entanglement games. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2189):20160003, May 2016. doi:10.1098/rspa.2016.0003.

28 Adrian Kent, William J. Munro, and Timothy P. Spiller. Quantum tagging: Authenticating location via quantum information and relativistic signaling constraints. *Physical Review A*, 84(1), July 2011. doi:10.1103/physreva.84.012326.

29 Hoi-Kwan Lau and Hoi-Kwong Lo. Insecurity of position-based quantum-cryptography protocols against entanglement attacks. *Physical Review A*, 83(1), January 2011. doi:10.1103/physreva.83.012322.

30 Jiahui Liu, Qipeng Liu, and Luowen Qian. Beating Classical Impossibility of Position Verification. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 100:1–100:11, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.100.

31 Robert A. Malaney. Location-dependent communications using quantum entanglement. *Phys. Rev. A*, 81:042319, April 2010. doi:10.1103/PhysRevA.81.042319.

32 Dominic Mayers and Andrew Yao. Self Testing Quantum Apparatus. *Quantum Info. Comput.*, 4(4):273–286, July 2004. URL: http://dl.acm.org/citation.cfm?id=2011827.2011830.

33 S. Pironio, A. Acín, S. Massar, A. Boyer de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by Bell's theorem. *Nature*, 464(7291):1021–1024, April 2010. doi:10.1038/nature09008.

34 Jérémy Ribeiro and Frédéric Grosshans. A tight lower bound for the bb84-states quantum-position-verification protocol, 2015. doi:10.48550/arXiv.1504.07171.

35 Christian Schaffner. Cryptography in the bounded-quantum-storage model, 2007. arXiv:0709.0289.

36 Florian Speelman. Instantaneous non-local computation of low T-depth quantum circuits. In *11th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2016)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.

37 Ivan Šupić and Joseph Bowles. Self-testing of quantum systems: a review. *arXiv:1904.10042 [quant-ph]*, April 2019. arXiv: 1904.10042. arXiv:1904.10042.

38 Marco Tomamichel, Serge Fehr, Jedrzej Kaniewski, and Stephanie Wehner. A monogamy-of-entanglement game with applications to device-independent quantum cryptography. *New Journal of Physics*, 15(10):103002, October 2013. doi:10.1088/1367-2630/15/10/103002.

39 Dominique Unruh. Quantum position verification in the random oracle model. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, pages 1–18, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

40 Dominique Unruh. Quantum Position Verification in the Random Oracle Model. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, Lecture Notes in Computer Science, pages 1–18, Berlin, Heidelberg, 2014. Springer. doi:10.1007/978-3-662-44381-1_1.

41 William K. Wootters and Wojciech Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.

42 Mark Zhandry. How to Record Quantum Queries, and Applications to Quantum Indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, Lecture Notes in Computer Science, pages 239–268. Springer International Publishing, 2019. doi:10.1007/978-3-030-26951-7_9.

Mixing Time of Quantum Gibbs Sampling for Random Sparse Hamiltonians

Akshar Ramkumar
Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA, USA

Mehdi Soleimanifar
Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA, USA

Abstract

Providing evidence that quantum computers can efficiently prepare low-energy or thermal states of physically relevant interacting quantum systems is a major challenge in quantum information science. A newly developed quantum Gibbs sampling algorithm [11] provides an efficient simulation of the detailed-balanced dissipative dynamics of non-commutative quantum systems. The running time of this algorithm depends on the mixing time of the corresponding quantum Markov chain, which has not been rigorously bounded except in the high-temperature regime. In this work, we establish a polylog(n) upper bound on its mixing time for various families of random $n \times n$ sparse Hamiltonians at any constant temperature. We further analyze how the choice of the jump operators for the algorithm and the spectral properties of these sparse Hamiltonians influence the mixing time. Our result places this method for Gibbs sampling on par with other efficient algorithms for preparing low-energy states of quantumly easy Hamiltonians.

2012 ACM Subject Classification Theory of computation → Quantum information theory; Theory of computation → Design and analysis of algorithms

Keywords and phrases Quantum algorithms, quantum Gibbs sampling, mixing time analysis

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.3

Related Version *Extended Version*: `arXiv preprint` [32]

Supplementary Material *Software (Source Code)*: `https://github.com/countableclouds/lindbladian_gap`, archived at `swh:1:dir:9db70c0155cd182d780b7562eca062cc64a051ad`

Funding *Akshar Ramkumar*: AR acknowledges the Caltech SURF program and Zhuang Tang and Gebing Yi for their generous support.
Mehdi Soleimanifar: MS is supported by AWS Quantum Postdoctoral Scholarship and the National Science Foundation NSF CAREER award CCF-2048204.

Acknowledgements We thank Jiaoyang Huang, Luca Nashabeh, John Preskill, and Yongtao Zhan for valuable discussions. We are grateful to Chi-Fang Chen for insightful discussions and proposing this project in its early stages. We also thank Hongrui Chen and the authors of [12] for informing us about the overlap between our results and those in Section IV of their work. Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

1 Introduction

One of the main anticipated applications of quantum computers is the simulation and characterization of quantum systems in condensed matter physics [40], quantum chemistry [29], and high-energy physics [30, 4]. The problem of simulating the dynamics (time evolution) of an interacting quantum system under a local or sparse Hamiltonian $\boldsymbol{H}$ has largely been addressed, with efficient algorithms [22, 27, 5, 28, 20] that scale well with the number of

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 3; pp. 3:1–3:23
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

particles, simulation time, and required precision. However, the ability of quantum computers to evaluate the static features of quantum systems, such as their ground state or thermal properties, is less understood.

In this work, we focus on preparing the Gibbs (thermal) state $\boldsymbol{\rho}_\beta = \frac{e^{-\beta \boldsymbol{H}}}{\mathrm{Tr}(e^{-\beta \boldsymbol{H}})}$ of a quantum system, which represents the equilibrium state when the system is in contact with a thermal bath at a fixed temperature β^{-1}. This computational problem, known as Gibbs sampling or "cooling," is valuable not only for simulating thermodynamic properties but also as a subroutine in quantum algorithms for optimization and learning [7, 2, 6]. However, to prepare the Gibbs state, quantum computers face challenges. In general, it is not believed that estimating the low-temperature properties of quantum systems can be solved efficiently by a quantum computer in the worst-case [24]. Fortunately, it has been hypothesized that this worst-case hardness of finding low-temperature states implied by arguments from complexity theory is due to pathological Hamiltonians, which are not apparent in many physical systems that normally occur in nature. This hypothesis is substantiated by the empirical success of natural cooling, such as using refrigerators, in reaching thermal equilibrium.

Quantum Gibbs sampling. Aiming to mimic nature's cooling processes, a series of recent works have introduced quantum Markov Chain Monte Carlo (MCMC) algorithms, or quantum Gibbs samplers [11, 10, 36, 42, 31, 23, 43, 16, 19], as promising alternatives for tackling a range of classically intractable low-temperature simulation tasks on quantum computers. These algorithms are designed to replicate the success of classical Markov chains in preparing Gibbs states for classical Hamiltonians. The analysis of classical MCMC algorithms relies on the principle of detailed balance; however, achieving this in the quantum setting has been challenging and was only recently addressed by an algorithm in [11]. Part of the difficulty arises from a conflict between the finite energy resolution σ_E achievable by efficient quantum algorithms and the seemingly strict requirement to precisely distinguish energy levels to satisfy detailed balance. In this work, we focus primarily on this algorithm, referring to it as the CKG algorithm or the quantum Gibbs sampler when the context is clear. We give a detailed review of this algorithm in Section 4.1.3 and Appendix 4.2.1.

The Gibbs sampling algorithm provides a fully general method for preparing Gibbs states by evolving an initial state $\boldsymbol{\rho}_0$ under a Lindbladian $\mathcal{L}_\beta$, which is efficiently implementable on a quantum computer and produces the state $\boldsymbol{\rho}_t = e^{\mathcal{L}_\beta t}[\boldsymbol{\rho}_0]$ after time t. The runtime of the quantum Gibbs sampler is governed by the *mixing time* of the corresponding quantum Markov chain, which is roughly the time required for $\boldsymbol{\rho}_t$ to approach the Gibbs state $\boldsymbol{\rho}_\beta$. This in turn is bounded by the spectral gap $\lambda_{\mathrm{gap}}(\mathcal{L}_\beta)$ of the Lindbladian by

$$t_{\mathrm{mix}}(\mathcal{L}_\beta) \leq \frac{\mathcal{O}(\beta\|\boldsymbol{H}\| + \log(n))}{\lambda_{\mathrm{gap}}(\mathcal{L}_\beta)}.$$

The spectral gap is defined here to be $\lambda_{\min}$, the smallest eigenvalue of $-\mathcal{L}_\beta$ for any eigenvector other than the fixed point ρ_β. Bounding the spectral gap, therefore, proves not only that $\mathcal{L}_\beta$ has a unique fixed point, but also quantifies the rate of convergence. The mixing time varies based on the quantum system in question. Bounding this mixing time is challenging without access to fault-tolerant quantum computers, as we cannot run and benchmark the algorithm directly, making theoretical analysis essential. However, such analysis is hindered by a lack of technical tools for two key reasons. Firstly, the theory of convergence of quantum Markov chains is new, unlike the very mature twin field for classical Markov chains. Secondly, the Markov chain described by the algorithm is considerably complex, and depends on several parameters that we will discuss in more detail shortly: an energy resolution σ_E, a series of jump operators $\boldsymbol{A}^a$ for $a \in [M]$, and the inverse temperature β. The space of possibilities makes the algorithm's performance more difficult to characterize.

This motivates the identification of quantum systems whose mixing times are tractable for analysis yet exhibit rich features that provide insights into the performance of the quantum Gibbs sampler for more general non-commuting Hamiltonians. In line with this, the mixing time of the CKG algorithm has recently been bounded for local Hamiltonians, showing a polynomial scaling with system size at high enough temperatures [33].

Mixing time of sparse Hamiltonians. In this work, we consider an alternative approach by characterizing the mixing time of a family of *sparse* Hamiltonians of the form

$$\boldsymbol{H} = \sum_{i,j\in[n]} H_{ij} \left|e_i\right\rangle \left\langle e_j\right| . \tag{1}$$

Such an operator can be understood as the Hamiltonian on a graph $G = (V, E)$ with $n = |V|$ vertices indexed by basis states $|e_i\rangle$, $i \in [n]$ and a set of edges E connecting vertices with $H_{ij} \neq 0$. When non-zero entries H_{ij} are all equal to 1, the Hamiltonian $\boldsymbol{H}$ corresponds to the $n \times n$ adjacency matrix of the n-vertex graph. We define the degree d of the graph G as the sparsity of the underlying Hamiltonian and refer to Hamiltonians with constant or slowly increasing degrees $d = \mathrm{polylog}(n)$ as *sparse*. Note that any $\log(n)$-qubit Hamiltonian that consists of $m = \mathrm{polylog}(n)$ terms each acting locally on $\kappa = O(1)$ qubits is a sparse Hamiltonian with degree $d \leq m2^{\kappa} \leq \mathrm{polylog}(n)$. However, not all sparse Hamiltonians admit local qubit encodings.

Having defined sparse Hamiltonians, we now consider the dissipative dynamics of the system induced by a set of M jump operators $\boldsymbol{A}^a = \sum_{i,j\in[n]} A^a_{ij} \left|e_i\right\rangle \left\langle e_j\right| , \quad a \in [M]$. We will soon explain how the jump operators $\boldsymbol{A}^a$ relate to the Lindbladian $\mathcal{L}_\beta$. Briefly, the resulting dynamics can be understood as a combination of two processes: a continuous-time quantum walk of a single particle on the graph of states due to the coherent evolution of the Hamiltonian $\boldsymbol{H}$, which is combined with stochastic jumps on the graph determined by the jump operators $\boldsymbol{A}^a$.

Our interest in bounding the mixing time of the sparse Hamiltonians is multifaceted:

(1) **Single-particle dynamics.** As stated earlier, bounding the mixing time of general interacting multipartite Hamiltonians is a challenging task. However, for simple choices of graphs G, the mixing time of the quantum Gibbs sampler may be easier to analyze, potentially leading to relevant techniques for tackling the case of interacting particles. In fact, we can think of the dynamics induced by the Hamiltonian $\boldsymbol{H}$ (1) as the dynamics of a single-particle hopping on the graph G. This single-particle evolution on path graphs or grids is commonly analyzed in the tight-binding model in condensed matter physics. That being said, even in the simplified case of a single particle, the Hamiltonian $\boldsymbol{H}$ is non-commuting, characterizing a continuous-time quantum walk that can yield exponential quantum advantage for certain oracular problems on graphs such as the glued trees [13].

(2) **Chaotic Hamiltonians.** Our additional motivation for studying random sparse Hamiltonians stems from the fact that their spectra exhibit many of the same characteristics as *chaotic* Hamiltonians, such as the SYK model [34, 25, 26] and random p-spin models [37, 41]. Understanding whether chaotic Hamiltonians have a fast mixing time as they approach their thermal and low-energy states is a fundamental question in the study of quantum chaos [8, 1]. As a concrete step toward addressing this problem, we identify key spectral properties of random sparse Hamiltonians that can ensure a fast mixing time.

(3) **Algorithmic applications.** Preparing quantum Gibbs states, and more broadly computing the matrix exponential of sparse matrices such as the adjacency or Laplacian of a graph, is a fundamental subroutine in solving various graph and optimization problems. For instance, the Estrada index – defined as the trace of the matrix exponential of a graph's adjacency matrix – measures subgraph centrality and provides structural insights [18]. Computing the matrix exponential is also related to matrix inversion and linear system solvers [35]. Moreover, quantum Gibbs sampling has been applied to solving semidefinite programs (SDPs) in optimization problems [21, 7, 6, 3], offering quantum speedups for these problems.

2 Our main results

Motivated by these considerations, we investigate the mixing time of quantum Gibbs samplers for sparse Hamiltonians and different choices of jump operators. Our study addresses two key questions regarding the performance of the quantum Gibbs sampler for sparse Hamiltonians. First, we ask

What choices of jump operators lead to a fast mixing time?

After exploring the effects of different jump operators $\boldsymbol{A}^a$, we then focus on the spectral properties of sparse Hamiltonians to understand:

What spectral property of the Hamiltonian determines its mixing time?

Answering these questions allows us to provide broad and intuitive insights on how the quantum Gibbs sampler operates for general families of sparse Hamiltonians.

2.1 Choice of jumps: graph-local vs unitary design

A natural set of jump operators for a given $n \times n$ Hamiltonian on a graph G are $\boldsymbol{A}^a = \frac{1}{\sqrt{n}}|e_a\rangle\langle e_a|$ or similar operators supported on a few neighboring vertices of G. Importantly, these are not "local" in the sense of multi-particle Hamiltonians, which refers to being composed of terms that act on a small number of qubits – often also geometrically close to one another. Utilizing graph-local jump operators also significantly simplifies the structure of the Lindbladian and the analysis of mixing times for certain graph families.

Moving beyond graph-local jumps, the Lindbladian $\mathcal{L}_\beta$ of the quantum Gibbs sampler can still be efficiently implemented on a quantum computer with a much broader class of jumps. This is possible as long as each jump $\boldsymbol{A}^a$ is efficiently implementable, the set of jumps M includes both $\boldsymbol{A}^a$ and its adjoint $\boldsymbol{A}^{a\dagger}$, and $\sum_{a\in[M]} \|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\|_\infty = 1$ (due to this normalization condition, we will sometimes speak of the jumping distribution $\mathcal{A}$, from which the jump operators $\boldsymbol{A}^a$ are sampled with probability $\|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\|$). This raises the question of whether there is an advantage in using *non-local* jumps that have a bounded spectral norm, or if more structured local jumps are sufficient to achieve a fast-mixing quantum MCMC. After all, *classical* continuous-time random walks are typically considered with local jumps on the graph vertices. However, in the context of graphs, we will see that the structured nature of graph-local jumps offers no advantage, but rather seems to cause a slowdown of the resulting algorithm.

Graph-local jumps. To this end, in the next theorem, we establish tight bounds on the spectral gap of the Lindbladian for cyclic graphs for graph-local jumps $\boldsymbol{A}^a$, with an approach similar to the one used in [38] to bound the spectral gap of a Davies generator.

▶ **Theorem 2.1** (Spectral gap of cyclic graphs with local jumps). *Fix temperature β^{-1}. There exists some constant energy resolution σ_E for which the spectral gap of the CKG Lindbladian $\mathcal{L}_\beta$ for a cyclic graph with n vertices with jump operators $\boldsymbol{A}^a = \frac{1}{\sqrt{n}}|e_a\rangle\langle e_a|$ is asymptotically $\Theta(n^{-3})$.*

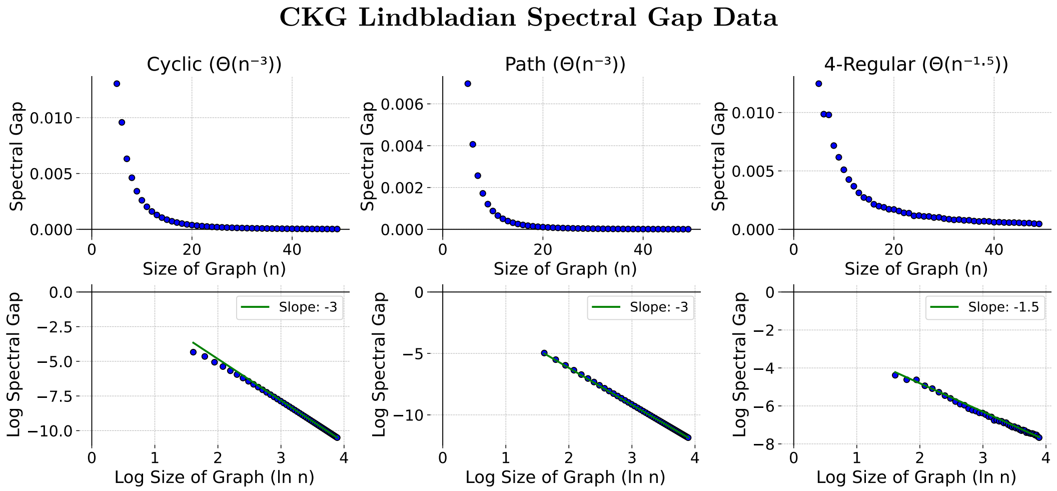

Figure 1 Linear (above) and log-log (below) graphs of spectral gap with respect to system size. Gaps of ten random 4-regular graphs were averaged for each data point. For the cyclic graphs (one-dimensional lattices), the proven asymptotic decay aligns closely with the data.

In addition to theoretical analysis, we also generated data for cyclic graphs, path graphs, and random d-regular graphs with n vertices, as shown in Figure 1. These numerical results suggest spectral gaps of $o(n^{-1})$ for generic sparse graphs with graph-local jumps. We observed that increasing the constant d does improve the spectral gap decay, though it never improved past the asymptotic decay $O(n^{-1})$.

These results are all suboptimal, since for an $n \times n$ Hamiltonian $\boldsymbol{H}$, we expect an efficient result would be polynomial in the number of qubits, i.e. polylog(n) rather than poly(n). The poor performance can be attributed to two factors. (1) The operators $\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a$ have L^1 norm $\frac{1}{n}$. (2) In the energy basis, many entries of $\boldsymbol{A}^a$ are highly correlated.

The first drawback effectively scales the Lindbladian down by $\frac{1}{n}$, since the L^1 norm of $\boldsymbol{A}^\dagger\boldsymbol{A}$ can be as high as 1 when the operator norm $\|\boldsymbol{A}^\dagger\boldsymbol{A}\| = \frac{1}{n}$. However, the chosen jump operators are projectors, so their L^1 and operator norms are equal. In both Theorem 2.1 and in the data, a spectral gap even worse than $\frac{1}{n}$ is observed. This is due to the second drawback. The aforementioned correlations lead to off-diagonal terms in the Lindbladian, which in general have the potential to dampen the spectral gap, and in the case of the cyclic graph provably do so. It appears that more generally, the biases of an ensemble of local jumps can introduce off-diagonal terms to the Lindbladian that decrease the spectral gap. The same harmful correlations appear to exist in higher degree graphs in addition to cyclic ones, though to a lesser extent as evidenced by the improved spectral gap.

Unitary design jumps. To address some of these shortcomings, we next consider non-local jump operators, each independently drawn (along with its adjoint pair) according to a unitary 1-design $\mathcal{D}(U(n))$ on n vertices. More precisely, we define

▶ **Definition 2.2.** *A set of jump operators $\{\boldsymbol{A}^a : a \in [M]\}$ is* ***drawn from a 1-design jumping distribution*** *if it is obtained by sampling $M/2$ jump operators i.i.d from a unitary 1-design $\mathcal{D}(U(n))$, normalizing each by $\frac{1}{\sqrt{M}}$, and including these operators along with their adjoint.*

We include the adjoint of each randomly chosen jump since the CKG Lindbladian requires the set of jump operators to be closed under adjoint, $\{\boldsymbol{A}^a : a \in [M]\} = \{\boldsymbol{A}^{a\dagger} : a \in [M]\}$. When n is a power of 2, the unitary 1-design can be constructed as a tensor product of random Pauli operators on $\log_2(n)$ qubits, in which case the jumps are self-adjoint and can be sampled and implemented efficiently. The efficiency of our results on a general system relies on the ability to efficiently implement some unitary 1-design.

As we will see, in our application, this 1-design sampling is effectively equivalent to sampling from a Haar-random distribution. This approach improves on the results given for graph-local jumps, and is able to achieve an efficient algorithm in the number of qubits for a graph (running time $\mathrm{polylog}(n)$) for Gibbs sampling. This improved performance is in part because all the eigenvalues of a Haar random unitary have magnitude 1. Hence, it avoids the problem of $\boldsymbol{A}^\dagger \boldsymbol{A}$ having a relatively small L^1 norm given the constraint on its operator norm $\|\boldsymbol{A}^\dagger \boldsymbol{A}\|$. These jumps also avoid the second problem encountered for the graph-local jumps: Since the number of degrees of freedom of randomness is very large over the ensemble, any form of bias is mitigated. Indeed, the resulting Lindbladian over the full ensemble has no off-diagonal terms resulting from correlated elements of the jump operator.

Our results extend beyond cyclic graphs to any graph of bounded degree $d = O(1)$ where $\|\boldsymbol{H}\| \leq d$ at constant temperature, or more generally when $\beta\|\boldsymbol{H}\| = O(1)$. We refer to these sparse Hamiltonians as bounded degree and formally define them as:

▶ **Definition 2.3.** *A* ***bounded degree*** *system is a sequence of temperatures $\beta(n)^{-1}$ and Hamiltonians $\boldsymbol{H}(n)$ for which $\beta(n)\|\boldsymbol{H}(n)\|$ is bounded from above by a constant independent of system size.*

▶ **Theorem 2.4** (Constant spectral gap of Lindbladian in bounded degree systems)**.** *Let $\beta(n)^{-1}$ be a sequence of temperatures and $\boldsymbol{H}(n)$ a sequence of $n \times n$ Hamiltonians such that $\beta(n)\|\boldsymbol{H}(n)\| = O(1)$.*

With any constant probability $1 - \xi$, the spectral gap of a Lindbladian $\mathcal{L}_\beta$ with $\sigma_E = \beta^{-1}$ and M jump operators sampled from a 1-design jumping distribution for some $M = \Theta(\log(n))$, is bounded below by a constant, i.e. $\lambda_{gap} = \Omega(1)$.

Assume access to an efficient block-encoding of $\boldsymbol{H}(n)$. Then as a consequence and in the same setup, the Gibbs state of $\boldsymbol{H}$ can be prepared with error ϵ in trace distance, in time $\mathrm{poly}(\log(n), \log(\epsilon^{-1}))$.

While the examples of bounded degree Hamiltonians we consider are mostly graphs, the above theorem applies to preparing the Gibbs state for any Hamiltonian at a temperature β^{-1} such that $\beta^{-1} = O(\|\boldsymbol{H}\|)$.

2.2 Mixing time from spectral profile

Theorem 2.4 demonstrates that bounded-degree Hamiltonians with non-local jumps exhibit fast mixing times. However, it leaves open the case of Hamiltonians with unbounded degrees, such as those with $d \leq \mathrm{polylog}(n)$. More formally, we define

▶ **Definition 2.5.** *An* ***unbounded degree*** *system is a sequence of temperatures $\beta(n)^{-1}$ and Hamiltonians $\boldsymbol{H}(n)$ for which $\lim_{n\to\infty} \beta(n)\|\boldsymbol{H}(n)\| = \infty$. However, we still assume that $\beta\|\boldsymbol{H}\| = \mathrm{polylog}(n)$, polynomial in the number of qubits.*

In our next result, we show that again selecting the jumping distribution (including adjoints) to be M samples from a unitary 1-design for sufficiently large M and choosing the energy resolution $\sigma_E = \beta^{-1}$ yields an algorithm whose efficiency depends on its low-energy spectrum. In particular, the runtime scales inverse polynomially with the fraction of eigenvalues $\delta(n)$ of $\boldsymbol{H}$ that are within $O(\beta^{-1})$ of the minimum eigenvalue.

▶ **Theorem 2.6** (Spectral gap of Lindbladian in unbounded degree systems). *Let $\beta(n)^{-1}$ be a sequence of temperatures and $\boldsymbol{H}(n)$ a sequence of $n \times n$ Hamiltonians, and let $\delta(n)$ be the fraction of eigenvalues of $\boldsymbol{H}(n)$ within $O(\beta^{-1})$ of λ_{min}.*

With any constant probability $1-\xi$, the spectral gap of a Lindbladian $\mathcal{L}_\beta$ with $\sigma_E = \beta^{-1}$ and M jump operators sampled from a 1-design jumping distribution for some $M = \Theta(\delta(n)^{-2}\log(n)\log(\beta\|\boldsymbol{H}\|))$, is lower bounded by $\Omega(\delta(n))$.

Assume access to an efficient block-encoding of $\boldsymbol{H}(n)$. As a consequence, and in the same setup, the Gibbs state of $\boldsymbol{H}$ can be prepared with error ϵ in trace distance, in time $\mathrm{poly}(\delta(n)^{-1}, \log(n), \log(\epsilon^{-1}))$.

Note that if $\beta\|\boldsymbol{H}\|$ is bounded, then $\delta(n) = 1$. This result therefore generalizes Theorem 2.4.

2.3 Explicit examples of random sparse Hamiltonians

Having established a sufficient spectral condition for the fast mixing of random ensembles of sparse Hamiltonians, we now give explicit examples that satisfy this criterion. We also give one example, the hypercube, which does not, and for which local jumps in place of unitary design jumps achieve an exponential speedup. This example elucidates the potential of structured local jumps for speedups, in contrast to the case of the cyclic graph in which structured graph-local jumps yielded a slowdown.

Random regular graphs. The first example is when $\boldsymbol{H}$ is the adjacency matrix of a randomly selected $\log(n)$-regular graph, with $\mathrm{polylog}(n)$ random 1-design jumps. In Section 4.4.1, we prove using Theorem 2.6 that this ensemble has a Lindbladian spectral gap of $\Omega(\log(n)^{-3/4})$ at constant temperature. This yields a polynomial algorithm to prepare the Gibbs state, given access to an efficient block-encoding of $\boldsymbol{H}$.

Random signed Pauli ensemble. The second example is the family of sparse Hamiltonians considered in [9], composed of random Pauli strings with random sign coefficients given by $\boldsymbol{H}_{PS} = \sum_{j=1}^{m} \frac{r_j}{\sqrt{m}}\boldsymbol{\sigma}_j$, where $\boldsymbol{\sigma}$ is a random Pauli string on n_0 qubits (such that the size of Hamiltonian is $n \times n$ for $n = 2^{n_0}$), each r_j is sampled randomly from $\{-1, 1\}$, and $m = O(\frac{n_0^5}{\epsilon^4})$ for a parameter $1 \geq \epsilon \geq 2^{-o(n_0)}$. We show in Section 4.4.2 that the CKG Lindbladian has a spectral gap of $\Omega(\epsilon^{-3/2})$ when we choose $M = \widetilde{O}\left(n_0^2\epsilon^{-3}\right)$ unitary 1-design jumps, inverse temperature $\beta = O(\epsilon^{-1})$, and $\sigma_E = \beta^{-1}$.

Hypercubes. The final example is the family of hypercubes. A hypercube with 2^d vertices and degree d can be interpreted as a Hamiltonian on d qubits $\sum_i \boldsymbol{X}_i$. At constant temperature, only an exponentially small fraction of eigenvalues lie near the minimum eigenvalue. As a result, the spectral profile implies a poor mixing time with unitary design jumps.

However, we show in Theorem 4.5 that by choosing local jumps $\frac{1}{\sqrt{d}}\boldsymbol{Z}_i$, the spectral gap is $\frac{1}{d}$, yielding an efficient algorithm for Gibbs sampling. Hypercubes therefore provide an example where not graph-local jumps, but rather local jumps on the qubits, ensure fast mixing. The crucial feature of local jumps that improves mixing time is that a local jump $\boldsymbol{A}^a$ on a local Hamiltonian $\boldsymbol{H}$ satisfies $\|[\boldsymbol{A}^a, \boldsymbol{H}]\| = O(1)$ – i.e., the jump operator only jumps between nearby eigenstates. This property is not held by graph-local jumps in general, so they displayed no improvement in the studied cases. In general, local jumps are the strongest candidates for fast-mixing on local Hamiltonians, though for which classes of local Hamiltonians fast-mixing can be achieved is still largely open. For most interesting classes of local Hamiltonians, the condition of Theorem 2.6 is unlikely to hold.

3 Proof sketch

3.1 Graph-local jumps

The proof of the mixing time for graph-local jumps in the cyclic graph involves two steps: First, in Appendix A, we derive a general expression for the CKG Lindbladian in the energy basis given by equation (5).

We then utilize this expression along with the fully known spectrum of the cyclic graph to show that the Lindbladian is block-diagonal in the energy basis of this graph, as demonstrated in Appendix B. One of these blocks corresponds to a classical Markov chain on the diagonal entries of the state in the energy basis, for which we establish a spectral gap lower bound using the canonical path method. For the remaining $n-1$ blocks, we apply the Gershgorin circle theorem to bound their eigenvalues.

3.2 Unitary design jumps

Bounded-degree systems. To establish a lower bound on the spectral gap of the Lindbladian $\mathcal{L}_\beta$ with unitary 1-design jumps, we consider a decomposition of the form $\mathcal{L}_\beta = \mathcal{L}_\mu + \delta\mathcal{L}$, where $\mathcal{L}_\mu = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}(U(n))}[\mathcal{L}_\beta]$ is the expected Lindbladian with the expectation taken over a single jump operator sampled from a unitary 1-design distribution $\mathcal{D}(U(n))$, and $\delta\mathcal{L}$ represents the remainder term. Due the quadratic form of the Lindbladian given in expression (3) we see that $\mathbb{E}_{\boldsymbol{A}\sim\mathcal{H}(U(n))}[\mathcal{L}_\beta] = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}(U(n))}[\mathcal{L}_\beta]$. Here, $\mathcal{H}(U(n))$ is a Haar random distribution over jump operators.

Note that a CKG Lindbladian must have jump operators in adjoint pairs, so a Lindbladian with a single jump operator will not satisfy detailed balance. However, the expected Lindbladian over one Haar random jump operator is equal to the expected Lindbladian over the adjoint pair of a Haar random jump operator, by linearity of expectation.

The proof of Theorem 2.4 proceeds by first showing that this expected Lindbladian $\mathbb{E}_{\boldsymbol{A}\sim\mathcal{H}(U(n))}[\mathcal{L}_\beta]$ has a constant spectral gap, as long as $\beta(n)\|\boldsymbol{H}(n)\|$ is bounded by a constant as a function of system size. As before, we call such systems *bounded degree*, since for constant β bounded degree graphs satisfy the required property. Indeed, if such a system has degree bounded by d, it must have spectrum in $[-d,d]$ by Gershgorin's circle theorem, since every row consists of zeros and at most d ones. Adding phases to the edges of these bounded degree graphs remains feasible by a similar argument, so there is no constraint of stoquasticity.

The result is stated formally as follows:

▶ **Lemma 3.1** (Constant spectral gap of average Lindbladian for bounded degree systems). *Let $\beta(n)^{-1}$ be a sequence of temperatures and $\boldsymbol{H}(n)$ be a sequence of $n\times n$ Hamiltonians such that $\beta\|\boldsymbol{H}\| = O(1)$. The spectral gap of $\mathcal{L}_\mu = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}(U(n))}\,[\mathcal{L}_\beta]$, the expected CKG Lindbladian with energy resolution $\sigma_E = \beta^{-1}$ over the ensemble of one jump operator sampled from a unitary 1-design, is asymptotically $\Omega(1)$.*

To establish Lemma 3.1, we show that for any system, the average Lindbladian over a Haar random ensemble of jump operators decomposes as $\mathcal{L}_\mu = \mathcal{L}_{\text{classical}} + \mathcal{L}_{\text{dephasing}}$. For a density matrix in the energy basis, the evolution of $\mathcal{L}_{\text{classical}}$ is a classical continuous Markov chain of the diagonal. The evolution of this classical Lindbladian maps the diagonal, in the limit, to the Gibbs distribution. The spectral gap of $\mathcal{L}_{\text{classical}}$ can be analyzed with the large suite of techniques for classical Markov chains.

Meanwhile, $\mathcal{L}_{\text{dephasing}}$ damps the off-diagonal terms of the density matrix. In the limit as $t\to\infty$, the state therefore converges to a classical distribution on the diagonal in the energy basis, with no off-diagonal terms, as desired. The operator $\mathcal{L}_{\text{dephasing}}$ diagonalizes in the energy basis of density matrices, with each off-diagonal element decaying at an independent rate. It is therefore simple to analyze as well. In summary, Lemma 3.1 establishes that the $\lambda_{\text{gap}}(\mathcal{L}_\mu) = \min(\lambda_{\text{gap}}(\mathcal{L}_{\text{classical}}), \lambda_{\text{min}}(\mathcal{L}_{\text{dephasing}})) = \Omega(1)$.

However, this result does not imply that any given jump sampled from the unitary 1-design (with its adjoint pair) would yield a gapped Lindbladian. As a result, it does not yet yield an efficient Gibbs sampling algorithm. To obtain such a result in Theorem 2.4, we demonstrate that the remainder term $\delta\mathcal{L} = \mathcal{L}_\beta - \mathcal{L}_\mu$ has a small spectral norm when a Lindbladian is constructed from a sufficiently large number of jumps M, rather than just one. In particular, a Lindbladian sampled with $\Theta(\log(n))$ normalized jumps from any 1-design concentrates closely to its expectation, thereby establishing a spectral gap lower bound. Since this lower bound applies to any graph at constant temperature β^{-1} with bounded degree, it applies to the periodic lattices, path graphs, and k-regular graphs discussed in the previous section.

Unbounded degree systems. In the context of unbounded degree systems, 1-design unitaries can no longer, in general, achieve an algorithm that is efficient in $\log(n)$ at constant temperature. Indeed, $\mathcal{L}_\mu = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}(U(n))}[\mathcal{L}_\beta] = \mathcal{L}_{\text{classical}} + \mathcal{L}_{\text{dephasing}}$ does not necessarily have a constant spectral gap in general, as it did in the case of bounded degree systems. However, we may establish a condition on the spectrum of $\boldsymbol{H}$, with which we can recover a lower bound for the spectral gap:

▶ **Lemma 3.2** (Spectral gap of average Lindbladian for unbounded systems)**.** *Let $\boldsymbol{H}(n)$ be a sequence of $n \times n$ Hamiltonians. For some C, let $\delta(n)$ be the proportion of eigenvalues λ_j of $\boldsymbol{H}$ such that $\beta^{-1}(\lambda_j - \lambda_{min}) \leq C$. The spectral gap of $\mathcal{L}_\mu = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}(U(n))}[\mathcal{L}_\beta]$, the expected CKG Lindbladian over the Haar random unitary ensemble of its jump operator at temperature β^{-1} with $\sigma_E = \Theta(\beta^{-1})$, is asymptotically $\Omega(\delta(n))$.*

The lemma expresses that if λ_{min} is within $O(\beta^{-1})$ of $\delta(n)$ of the eigenvalues, the spectral gap is at least $\delta(n)$. Similarly to Theorem 2.4, using this result to obtain an efficient Gibbs sampling algorithm amounts to showing that a Lindbladian with enough independently sampled jump operators shares a similar asymptotic spectral gap to the average Lindbladian, using a concentration bound. When the average spectral gap from the above lemma is $\delta(n)$, the number of jump operators to concentrate around the expectation increases to $\delta(n)^{-2}$, along with an overhead of $\log(n)\log(\beta\|\boldsymbol{H}\|)^2$. This result is captured in Theorem 2.6, and results in an algorithm with runtime $\text{poly}(\delta(n)^{-1}, \log(n), \log(\epsilon^{-1}))$ for Gibbs sampling, where ϵ is the error in trace distance. This runtime bound relies on the standing assumption in this paper that $\log(\beta\|\boldsymbol{H}\|) = \text{poly}(\log(n))$.

4 Technical details

4.1 The quantum Gibbs sampler

4.1.1 Lindbladian evolution

The recently proposed CKG quantum MCMC algorithm addresses the problem of finding thermal states by imitating thermodynamic processes [11, 10]. In this process, a system of particles evolves in contact with a thermal bath at some fixed temperature β^{-1}. Due to interactions with the bath, the system is described by a probabilistic mixture of quantum states $\boldsymbol{\rho}$. This state evolves in time, by approximation, with Markovian *dissipative* dynamics, $\frac{\mathrm{d}\boldsymbol{\rho}}{\mathrm{d}t} = \mathcal{L}_\beta[\boldsymbol{\rho}]$, given in terms of an operator $\mathcal{L}_\beta$ known as the Lindbladian. This operator involves a coherent term $\boldsymbol{B}$ that describes the interaction among the particles in the system. There is also a term in $\mathcal{L}_\beta$ that is specified by a series of *Lindblad operators* $\boldsymbol{L}^j$ that drive the dissipative transitions. The dynamics of the coherent term $\boldsymbol{B}$ are reversible, while the dissipative transitions drive all states toward some "stationary state". These transitions can be understood as perturbations from the bath, and as all states converge to the Gibbs state, the information of the system is leaking via these perturbations to the bath. The expression, in terms of $\boldsymbol{B}$ and $\boldsymbol{L}^j$, is:

$$\mathcal{L}_\beta[\cdot] = -i[\boldsymbol{B}, \cdot] + \sum_j \left(\boldsymbol{L}^j (\cdot) \boldsymbol{L}^{j\dagger} - \frac{1}{2}\{\boldsymbol{L}^{j\dagger}\boldsymbol{L}^j, \cdot\} \right).$$

The summands $\boldsymbol{L}^j(\cdot)\boldsymbol{L}^{j\dagger}$ are termed the transition part of the Lindbladian, and $-\frac{1}{2}\{\boldsymbol{L}^{j\dagger}\boldsymbol{L}^j, \cdot\}$ are the decay part of the Lindbladian. The choice of the Lindbladian operator $\mathcal{L}_\beta$ can vary depending on the precise nature of interactions between the system and the bath. However, to prepare the Gibbs (thermal) state at temperature β^{-1}, the Lindbladian should satisfy

$$\frac{\mathrm{d}\boldsymbol{\rho}_\beta}{\mathrm{d}t} = \mathcal{L}_\beta[\boldsymbol{\rho}_\beta] = 0 \quad \text{where} \quad \boldsymbol{\rho}_\beta := \mathrm{e}^{-\beta \boldsymbol{H}}/\mathrm{Tr}(\mathrm{e}^{-\beta \boldsymbol{H}}), \tag{2}$$

and moreover $\boldsymbol{\rho}_\beta$ should be the unique stationary state of the Lindbladian. The long-term evolution of the system under this Lindbladian, as a result, would converge to the Gibbs state of the Hamiltonian $\boldsymbol{H}$ at temperature $1/\beta$.

4.1.2 Detailed balance

To ensure that the Lindbladian $\mathcal{L}_\beta$ converges to a state $\boldsymbol{\rho}_\beta$, [11] designs a Lindbladian that satisfies Kubo-Martin-Schwinger (KMS) detailed balance with respect to $\boldsymbol{\rho}_\beta$. KMS detailed balance is one of several ways of quantizing the notion of classical detailed balance for Markov chains. KMS detailed balance of $\mathcal{L}_\beta$ is self-adjointness with respect to the inner product

$$\langle \sigma_1, \sigma_2 \rangle_{\rho_\beta^{-1}} = \mathrm{Tr}(\sigma_1^\dagger \boldsymbol{\rho}_\beta^{-1/2} \sigma_2 \boldsymbol{\rho}_\beta^{-1/2}). \tag{KMS Inner Product}$$

In particular, it is equivalent to the relation that

$$\mathcal{L}_\beta[\cdot] = \boldsymbol{\rho}_\beta^{1/2} \mathcal{L}_\beta^\dagger \left[\boldsymbol{\rho}_\beta^{-1/2} (\cdot) \boldsymbol{\rho}_\beta^{-1/2} \right] \boldsymbol{\rho}_\beta^{1/2} \tag{Detailed Balance}$$

where $\mathcal{L}_\beta^\dagger$ is the adjoint Lindbladian with respect to the Hilbert-Schmidt inner product $\langle \sigma_1, \sigma_2 \rangle = \mathrm{Tr}(\sigma_1^\dagger \sigma_2)$. The adjoint operator $\mathcal{L}_\beta^\dagger$, in the Heisenberg picture, describes the dynamics of observables under evolution by $\mathcal{L}_\beta$. The Lindbladian evolution is described by some quantum channel and therefore the observable I must always be fixed by $\exp(\mathcal{L}_\beta^\dagger)$. This implies that $\mathcal{L}_\beta^\dagger[I] = 0$. The detailed balance formula thereby implies that $\mathcal{L}_\beta[\boldsymbol{\rho}_\beta] = 0$, as desired. Note that KMS detailed balance can be dually described as the self-adjointness of $\mathcal{L}_\beta^\dagger$ with respect to the inner product $\langle \sigma_1, \sigma_2 \rangle_{\rho_\beta} = \mathrm{Tr}(\sigma_1^\dagger \boldsymbol{\rho}_\beta^{1/2} \sigma_2 \boldsymbol{\rho}_\beta^{1/2})$.

4.1.3 Construction and parameters

The quantum Gibbs sampler in [11] constructs a Lindbladian that satisfies two properties:

1. $\mathcal{L}_\beta$ satisfies detailed balance with respect to $\boldsymbol{\rho}_\beta$, and therefore $\mathcal{L}_\beta[\boldsymbol{\rho}_\beta] = 0$.
2. The dynamics of $\mathcal{L}_\beta$ can be efficiently implemented.

Their Lindbladian, which we term the CKG Lindbladian, can be simulated on a quantum computer with a cost per unit time $t = 1$ roughly equal to that of simulating the Hamiltonian dynamics of $\boldsymbol{H}$. The CKG Lindbladian is closely related to the Davies generator, which is a physically motivated Lindbladian that satisfies detailed balance, but that is not efficiently implementable in general. A full description of both Lindbladians are given in the Appendix.

The Gibbs sampling algorithm evolves an initial state $\boldsymbol{\rho}_0$ according to the efficiently implemented Lindbladian $\mathcal{L}_\beta$, and produces the state $\boldsymbol{\rho}_t = e^{\mathcal{L}_\beta t}[\boldsymbol{\rho}_0]$ after time period t. The *mixing time* is roughly the time that it takes for the state $\boldsymbol{\rho}_t$ to approach the Gibbs state $\boldsymbol{\rho}_\beta$. That is, $\mathrm{e}^{\mathcal{L}_\beta t_{\mathrm{mix}}}[\boldsymbol{\rho}_0] \approx \boldsymbol{\rho}_\beta$. The efficiency of the algorithm therefore scales linearly with the unit time simulation cost and the mixing time. The algorithm has several parameters in the Lindbladian's construction. In addition to the *inverse temperature* β, the algorithm specifies an *energy resolution* σ_E. A salient feature of [11]'s construction is that it can achieve detailed balance even though the algorithm only probes the energies of the Hamiltonian $\boldsymbol{H}$ with approximate precision. σ_E quantifies this level of precision. The cost of the Lindbladian simulation depends linearly on σ_E^{-1}, but increasing the precision may also improve the mixing time. Taking $\sigma_E \to 0$ for absolute precision recovers the Davies generator – when distinguishing the energies of the system exactly is infeasible, this Lindbladian cannot be simulated efficiently.

A set of *jump operators* $\boldsymbol{A}^a$ must also be specified for the Lindbladian. These operators are decomposed by frequency and reassembled in a particular way to construct the Lindblad operators that help $\mathcal{L}_\beta$ satisfy detailed balance. They must appear in adjoint pairs: i.e., if $\boldsymbol{A} \in \{\boldsymbol{A}^a\}$, then $\boldsymbol{A}^\dagger \in \{\boldsymbol{A}^a\}$. The cost of simulation scales with the cost of implementing the oracle $|a\rangle \to |a\rangle \otimes \boldsymbol{A}^a$. In particular, the jump operators must be normalized when implemented for the algorithm, satisfying $\sum_a \|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\| \leq 1$. CKG Lindbladians are linear in their jump operators – if $\mathcal{L}_1$ has one jump operator $\boldsymbol{A}^1$ and $\mathcal{L}_2$ has one jump operator $\boldsymbol{A}^2$, then a Lindbladian $\mathcal{L}$ with jump operators $\boldsymbol{A}^1$ and $\boldsymbol{A}^2$ satisfies $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2$. If $\mathcal{L}_\beta$ was constructed from jumps $\boldsymbol{A}^a$, then jump operators $\sqrt{s}\boldsymbol{A}^a$ produce the Lindbladian $s\mathcal{L}_\beta$, scaling the mixing time by s. So we may therefore assume that $\sum_a \|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\| = 1$ exactly, since renormalizing can only improve the spectral gap. In its normalized form, the set of jump operators can be understood as a jumping distribution over $\boldsymbol{A}^a$ which we will notate $a \sim \mathcal{A}$, where each is sampled with probability $\|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\|$.

4.2 Mathematical description

We begin with a description of the Davies generator, which is the limit of the CKG Lindbladian as $\sigma_E \to 0$. This generator was developed from a physical approximation of an open thermalizing quantum system, but at low temperatures it is unphysical and can be hard to implement. We then generalize the notions to the implementable CKG Lindbladian.

4.2.1 Davies generator

In the description of the Davies generator for a given system $\boldsymbol{H}$, there is a coherent term and there are jump operators $\boldsymbol{A}^a$. The $\boldsymbol{A}^a$ terms must appear in adjoint pairs in the Davies generator. The dissipative part of the Lindbladian is expressed as follows:

$$\mathcal{L}_\beta[\cdot] = \sum_{a \in [M]} \int_{-\infty}^{\infty} \gamma(\omega) \left(\boldsymbol{A}^a_\omega (\cdot) \boldsymbol{A}^\dagger_\omega - \frac{1}{2}\{\boldsymbol{A}^\dagger_\omega \boldsymbol{A}_\omega, \cdot\} \right) d\omega, \tag{3}$$

where $\boldsymbol{A}^a_\omega$ is the Operator Fourier Transform (OFT) of jump operator $\boldsymbol{A}^a$:

$$\boldsymbol{A}^a_\omega = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\boldsymbol{H}t} \boldsymbol{A}^a e^{-i\boldsymbol{H}t} e^{-i\omega t} dt.$$

The Davies' generator chooses Lindblad operators $\sqrt{\gamma(\omega)}\boldsymbol{A}^a_\omega$, each of which selects the energy transitions, or Bohr frequencies, in $\boldsymbol{A}^a$ that are precisely ω. Because it requires certainty in energy, by Heisenberg's uncertainty principle of energy and time, in the general case simulating the evolution of the Davies generator efficiently is infeasible. In the above, γ is some function satisfying $\gamma(\omega) = \gamma(\omega) = e^{-\beta\omega}\gamma(-\omega)$. The Lindblad operators are scaled by $\gamma(\omega)$ precisely to satisfy KMS detailed balance. Since A^a_ω represents jumps with Bohr frequency ω, the functional equation of γ ensures a desired ratio of jumps with Bohr frequency ω and $-\omega$. We choose the Metropolis filter, $\gamma(\omega) = \min(1, e^{-\beta\omega})$, though another common filter $\gamma(\omega) = \frac{1}{1+e^{-\beta\omega}}$ for "Glauber dynamics" could also be used for the same results.

The Davies generator satisfies detailed balance with respect to $\boldsymbol{\rho}_\beta$, the thermal state. In some presentations of the Davies generator, it contains a coherent term $-i[\boldsymbol{H}, \cdot]$. If this term is included, the generator does not satisfy detailed balance, so we do not follow this convention. However, the term does not affect the fixed point of the generator, since $\boldsymbol{\rho}_\beta$ commutes with $\boldsymbol{H}$ and therefore $-i[\boldsymbol{H}, \boldsymbol{\rho}_\beta] = 0$.

4.2.2 CKG Lindbladian

The CKG Lindbladian is defined almost identically to the Davies generator, but is altered slightly so that it still obeys detailed balance, but is efficiently implementable.

$$\mathcal{L}_\beta[\cdot] = \underbrace{-i[\boldsymbol{B}, \cdot]}_{\text{coherent term}} + \sum_{a\in[M]} \int_{-\infty}^{\infty} \gamma(\omega) \left(\underbrace{\hat{\boldsymbol{A}}^a(\omega)(\cdot)\hat{\boldsymbol{A}}^a(\omega)^\dagger}_{\text{transition term}} - \underbrace{\frac{1}{2}\{\hat{\boldsymbol{A}}^a(\omega)^\dagger\hat{\boldsymbol{A}}^a(\omega), \cdot\}}_{\text{decay term}} \right) d\omega,$$

where $\hat{\boldsymbol{A}}^a(\omega)$ is now the Gaussian-supported OFT of jump operator $\boldsymbol{A}^a$:

$$\hat{\boldsymbol{A}}^a(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\boldsymbol{H}t}\boldsymbol{A}^a e^{-i\boldsymbol{H}t} e^{-i\omega t} f(t) dt.$$

To ensure that the jump operators do not have infinite precision in energy, a Gaussian supported OFT is performed instead to obtain $\hat{\boldsymbol{A}}^a(\omega)$, which selects a Gaussian band energies of around ω.

Here, $f(t) = e^{-\sigma_E^2 t^2}\sqrt{\sigma_E\sqrt{2/\pi}}$, with Fourier transform $\hat{f}(\omega) = \frac{1}{\sqrt{\sigma_E\sqrt{2\pi}}}\exp(-\frac{\omega^2}{4\sigma_E^2})$. As a result, the operator $\hat{\boldsymbol{A}}^a(\omega)$ can be shown to be equal to $\sum_\nu \hat{f}(\omega - \nu)\boldsymbol{A}^a_\nu$. The function $f(t)$ was chosen so that its squared Fourier transform $\hat{f}^2(\omega)$ is a Gaussian with standard deviation σ_E, which features prominently in the Lindbladian (since it consists of quadratic terms in $\hat{\boldsymbol{A}}^a(\omega)$). Taking $\sigma_E = \Theta(\beta^{-1})$ yields an efficient simulation algorithm with the assumption of a block-encoding of $\boldsymbol{H}$ and a block-encoding for the jump operators $\sum_{a\in[M]} |a\rangle \otimes \boldsymbol{A}^a$, so σ_E is taken to be on the order of β^{-1} in this paper.

Since $\boldsymbol{A}^a(\omega)$ is a noisy decomposition of $\boldsymbol{A}^a$ into frequencies, it is not immediately clear whether there is a choice of function $\gamma(\omega)$ for which they can be recombined to achieve detailed balance. Indeed, as shown in [11], there is! The choice of $\gamma(\omega)$ is such that the transition part of $\mathcal{L}_\beta$, the summand $\sum_{a\in[M]}\int_{-\infty}^{\infty}\gamma(\omega)\hat{\boldsymbol{A}}^a(\omega)(\cdot)\hat{\boldsymbol{A}}^a(\omega)^\dagger d\omega$, still satisfies KMS detailed balance. [11] proved that there is a unique choice of $\boldsymbol{B}$, up to translation by a scalar, such that $-i[\boldsymbol{B}, \cdot] - \frac{1}{2}\sum_{a\in[M]}\int_{-\infty}^{\infty}\gamma(\omega)\{\hat{\boldsymbol{A}}^a(\omega)^\dagger\hat{\boldsymbol{A}}^a(\omega), \cdot\}d\omega$ also satisfies detailed balance. For the Davies generator, this coherent term $\boldsymbol{B}$ is simply 0 (or corresponds to a Lamb shift that commutes with the Hamiltonian), and the decay term by itself already satisfies detailed balance. $\boldsymbol{B}$ can be expressed in general as:

$$\boldsymbol{B} = \sum_{a\in[M]}\sum_{\nu_1,\nu_2} \frac{\tanh(-\beta(\nu_1 - \nu_2)/4)}{2i}(\boldsymbol{A}^a_{\nu_2})^\dagger \boldsymbol{A}^a_{\nu_1}.$$

The choice of γ for this algorithm, for which the filter is efficiently implementable, is $\gamma(\omega) = \exp\left(-\beta \max\left(\omega + \frac{\beta\sigma_E^2}{2}, 0\right)\right)$. As $\sigma_E \to 0$ it converges to the Metropolis filter of the Davies generator. In particular, this γ is precisely Metropolis filter for the Davies generator shifted by $\beta\sigma_E^2$. The CKG Lindbladian requires a choice of γ for which $\alpha = \gamma * g$ satisfies the functional equation that was originally satisfied by γ in the Davies generator.

We also note that $\mathcal{L}_\beta$ is bounded in operator norm.

▶ **Lemma 4.1.** *Consider the CKG Lindbladian $\mathcal{L}_\beta$ with temperature β^{-1}, using the Metropolis filter, and with jump operators $\boldsymbol{A}^a$ for which $\sum_a \|\boldsymbol{A}^{a\dagger}\boldsymbol{A}^a\| \leq 1$. This Lindbladian satisfies $\|\mathcal{L}_\beta\|_{\infty\to\infty} = O(\log(\beta\|\boldsymbol{H}\|))$, where $\|\cdot\|_{\infty\to\infty}$ is the operator norm of $\mathcal{L}_\beta$, with respect to the operator norm on the input and output vector spaces.*

Proof. The result follows from citing Proposition B.2 in [11] to bound the operator norm of the coherent term, and bounding the transition and decay terms manually. This proof is described in the arXiv version [32]. ◀

4.3 Spectral gap

Since the quantum MCMC algorithm was proposed recently, numerical and analytic characterizations of algorithm are limited. As for classical Markov chains, it has been shown that the mixing time of the algorithm can be characterized by the spectral gap $\lambda_{\text{gap}}(\mathcal{L}_\beta)$ of the Lindbladian. If the first eigenvalue $\lambda_1 = 0$ corresponds to eigenvector ρ_β, then the spectral gap is $\lambda_{\text{gap}}(\mathcal{L}_\beta) = \min_{j>1} |\lambda_j|$ [11]. Lindbladians are in general negative semidefinite like classical Markov chain generators, so $\lambda_{\text{gap}}(\mathcal{L}_\beta) = \min_j(-\lambda_{\text{j}})$. More precisely, it holds that

$$\frac{\Omega(1)}{\lambda_{\text{gap}}(\mathcal{L}_\beta)} \leq t_{\text{mix}}(\mathcal{L}) \leq \frac{\log\left(\left\|\rho_\beta^{-1/2}\right\|\right)}{\lambda_{\text{gap}}(\mathcal{L}_\beta)} \leq \frac{\mathcal{O}(\beta\|\boldsymbol{H}\| + \log(\dim(\boldsymbol{H})))}{\lambda_{\text{gap}}(\mathcal{L}_\beta)}. \tag{4}$$

In particular, analytically bounding this spectral gap from below is sufficient to prove that ρ_β is the unique fixed point, and for obtaining an upper bound on the mixing time. For so-called rapid mixing, in which the mixing time is logarithmic in the number of qubits, the spectral gap bound often does not suffice. For our purposes of proving efficiency in the number of qubits, however, this issue is moot.

4.4 Unbounded degree systems

We now prove efficient Gibbs sampling results for certain unbounded sparse Hamiltonians.

4.4.1 Random $\log(n)$-regular graphs

With high probability at constant temperature, a randomly selected $d = \log(n)$-regular graph, with poly(d) random 1-design jumps, has a Lindbladian spectral gap of $\Omega(d^{-3/4})$. This gives a polynomial algorithm to prepare the Gibbs state for most such graphs at constant temperature.

The gap of $\Omega(d^{-3/4})$ arises because a random d-regular graph, for $d \to \infty$, has one eigenvalue at d and the rest distributed from $-2\sqrt{d-1}$ to $2\sqrt{d-1}$ in a distribution that converges to a (normalized) semicircle. This semicircular distribution frequently appears in random matrix theory, for instance in the Gaussian unitary ensemble (GUE), which models the spectrum of many chaotic quantum systems. When the spectrum of a quantum system indeed follows this distribution, it implies that $\delta(n) = \Omega(d^{-3/4})$ of the eigenvalues lie within a constant of the minimum eigenvalue.

▶ **Theorem 4.2.** *With any constant probability $1-\xi$, for a randomly selected $d=\log(n)$-regular graph, there are $\delta=\Omega(d^{-3/4})$ eigenvalues within $O(1)$ of the minimum eigenvalue.*

As an immediate result of Theorem 4.2 and Theorem 2.6, we obtain an algorithm polynomial in d to prepare the Gibbs state of a d-regular graph. To prove the corollary, we use Theorem 2 in [17]. For this context, the following statement suffices.

▶ **Theorem 4.3.** *Let degree $d=\log(n)$. For sufficiently large n, there exists some $D>0$ such that for any interval $I\subset\mathbb{R}$, $0<\alpha<1$, and $0<\epsilon<\alpha$ such that $|I|>Dd^{-\alpha+\epsilon}$, with probability $1-o(n^{-1})$ over all random d-regular graphs, $|\delta(n)-\mu|<d^{-\epsilon}|I|$, where $\mu=\int_I\rho_{sc}(x)dx$ and $\delta(n)$ is the fraction of eigenvalues of the d-regular graph in $I\sqrt{d-1}$.*

In the above, ρ_{sc} is the asymptotic distribution as $d\to\infty$ of a random d-regular graph is the semicircular distribution with radius $2\sqrt{d-1}$, $\rho_{sc}(x)=\frac{1}{2\pi(d-1)}\sqrt{4(d-1)-x^2}$. From this result, Theorem 4.2 follows. As described more explicitly in [32], the estimated density ρ_{sc} near the bottom of the semicircle can be estimated. Then, using the theorem, this can be related to the fraction of eigenvalues near the minimum eigenvalue as well, proving the result.

4.4.2 Pauli String Ensemble

We now mention another ensemble of Hamiltonians studied by [9] in the context of low-energy state preparation. In [9], efficient low energy state preparation with phase estimation is demonstrated under the same conditions as our efficient Gibbs sampling in Theorem 2.6. Indeed, if many eigenvectors are close to the ground-state energy, as we require, then performing phase estimation on the maximally mixed state has a high probability of measuring a low-energy state, so low-energy state preparation is possible as well. They study the following ensemble of Hamiltonians on n_0 qubits, $\boldsymbol{H}_{PS}=\sum_{j=1}^{m}\frac{r_j}{\sqrt{m}}\boldsymbol{\sigma}_j$, where $\boldsymbol{\sigma}$ is a random Pauli string on n_0 qubits, each r_j is sampled randomly from $\{-1,1\}$, and $m=\left\lfloor c_2\frac{n_0^5}{\epsilon^4}\right\rfloor$. The parameter ϵ satisfies $\epsilon\geq 2^{-n_0/c_1}$, and c_1,c_2 are absolute constants. The resulting spectrum is again close enough to a semicircular distribution to obtain an efficient Gibbs sampler for certain temperatures that depend on ϵ. As ϵ decreases, Gibbs sampling becomes efficient for even larger values of β (lower temperatures), since the ensemble's spectrum converges closer to a perfect semicircular distribution at the edge of the spectrum.

Using the results in their paper, we establish that Gibbs sampling is efficient in n_0 for certain values of ϵ and corresponding temperatures β^{-1}.

▶ **Theorem 4.4.** *Say that $\boldsymbol{H}(n_0)$ is sampled from the ensemble $\boldsymbol{H}_{PS}$ on n_0 qubits, with $\epsilon=2^{-o(n_0)}$ and $\epsilon\leq 1$. With any constant probability $1-\xi$, for sufficiently large n_0, $\delta=\Omega(\epsilon^{3/2})$ fraction of the eigenvalues lie within $O(\epsilon)$ of the minimum eigenvalue.*

Proof. We utilize two results from [9]. Firstly, they argue that $\Pr[\|\boldsymbol{H}(n_0)\|\geq 2(1+\epsilon)]\leq\exp(-c_2n_0)$ when $m\geq\frac{n_0^3}{\epsilon^4}$, which is satisfied in this case. With an arbitrary constant probability for sufficiently large n_0, therefore, $\|\boldsymbol{H}(n_0)\|\leq 4$, since $\epsilon\leq 1$. The second result is that with probability $1-\exp(-c_3n_0^{1/3})$, at least $\Omega(\epsilon^{3/2})$ of the eigenvalues satisfy $\lambda_i\leq(1-\epsilon)\lambda_{\min}$ where c_3 is an absolute constant. With any large constant probability, we therefore have that $|\lambda_i-\lambda_{\min}|\leq\epsilon\lambda_{\min}\leq 4\epsilon=O(\epsilon)$ for $\Omega(\epsilon^{-3/2})$ of the eigenvalues. ◀

By Theorem 2.6, we obtain a Gibbs sampling algorithm that is $\text{poly}(\epsilon^{-1},n_0)$ to prepare the Gibbs state at inverse temperature ϵ^{-1}. We may rephrase this result in terms of β. For any polynomially large β, it provides a Pauli string ensemble of Hamiltonians, H_{PS} with $\epsilon=\beta^{-1}$, for which with high likelihood preparing the Gibbs state is efficient in n_0, assuming access to a block-encoding of the Hamiltonian of interest.

4.4.3 Hypercube graphs

For hypercube with varying dimension at a constant temperature, using unitary 1-design jumps would yield an exponentially large runtime. The spectrum of a hypercube with dimension d and 2^d vertices consists of the integers $-d, -d+2\ldots, d-2, d$. The eigenvalue j has multiplicity $\binom{d}{\frac{d+j}{2}}$. In particular, for any constant C, only an exponentially small fraction of the eigenvalues $\delta(d)$ lie below $-d+C$. This leads to a naive algorithm with at worst exponential complexity in d.

However, a better result can be obtained considering the hypercube as a system of d qubits. The graph with dimension d has 2^d vertices, which can be considered length d bitstrings. Then, the adjacency matrix is the sum of Pauli X operators on each qubit, $\sum_{i=1}^{d} X_i$, since the hypercube has an edge between any two bitstrings of Hamming distance 1. Choosing d jump operators as $\frac{1}{\sqrt{d}}Z_a$, the mixing time can be improved to $\mathrm{poly}(d, \log(\epsilon^{-1}))$:

▶ **Theorem 4.5** (Spectral Gap for Hypercube with Local Jumps). *For fixed β^{-1}, there exists some energy resolution σ_E such that the spectral gap of the CKG Lindbladian $\mathcal{L}_\beta$ for a d-dimensional hypercube with jump operators $\boldsymbol{A}^a = \frac{1}{\sqrt{d}}Z_a$, is asymptotically $\Omega(d^{-1})$.*

Proof. The proof of this statement is given in the arXiv version [32] by showing that the Lindbladian, with this choice of jump operators, is the product of independent Lindbladians on each qubit. Each of their spectral gaps can then be calculated explicitly. ◀

In the case of the hypercube, the local jump operators $\frac{1}{\sqrt{d}}Z_i$ only jump between eigenstates whose eigenvalues differ by 1. This vastly improves the performance of the classical Markov chain and dephasing Markov chain within the Lindbladian. However, the Lindbladian does not consist only of these two terms, as it did in the limit of independently sampled 1-design jumps. Off-diagonal terms do exist, and the presence of Z_a for *every* index is necessary to ensure that these off-diagonal terms do not completely eliminate the spectral gap. In some way, there must be "enough uncorrelated" local energy jumps to dampen these off-diagonal terms. For more complicated local Hamiltonians, it is not clear how correlations may be suppressed while still maintaining locality.

5 Connection to previous work

Our results show that for a Hamiltonian $\boldsymbol{H}$ with temperature β^{-1} such that some $\delta(n)$ fraction of the eigenstates are within $O(\beta^{-1})$ of the ground-state energy, the CKG quantum Gibbs sampler with 1-design jumps efficiently prepares the Gibbs state with trace distance at most ϵ. The running time scales polynomially with $\delta(n)^{-1}$, $\beta\|\boldsymbol{H}\|$, $\log(\epsilon^{-1})$, and the complexity of the block encoding of $\boldsymbol{H}$. This result is a baseline test that shows the CKG algorithm performs as well as other methods for preparing low-energy states of Hamiltonians. Indeed, our spectral condition is precisely the same as a condition that ensures easy quantum phase estimation of a near-ground state. Namely, performing quantum phase estimation on the maximally mixed state can prepare a random eigenstate, and with probability $\delta(n)$ it is within $O(\beta^{-1})$ of the minimum eigenvalue. Obtaining $O(\delta(n)^{-1})$ samples and taking the minimum energy can therefore prepare a near ground-state eigenvector. This approach is the basis of the previous analysis of random sparse Hamiltonians in [9].

Moreover, in [14], a quantum algorithm is presented that prepares the Gibbs state with a complexity that scales as $\mathrm{poly}(\frac{n}{\mathcal{Z}(\beta)}, \log(\epsilon^{-1}))$. If $\delta(n)$ of the eigenstates are within $O(\beta^{-1})$ of the ground-state energy, then $\frac{n}{\mathcal{Z}(\beta)} = \Omega(\delta(n)^{-1})$, and therefore under such conditions, this algorithm efficiently prepares the Gibbs states as well. Effectively, the CKG Gibbs sampler with "generic" 1-design jumps performs the same as previously developed algorithms – the algorithm is at least as powerful, but a potential advantage in cooling must arise from a smart (i.e., local and unbiased) choice of jump operators.

After completing our work, we also became aware of [12], where, among other contributions, the authors derive a new, efficient quantum Gibbs sampler algorithm that utilizes jump operators sampled from a Clifford-random circuit. This Gibbs sampler is shown to exhibit a spectral gap bound under the same condition on the spectral density considered in Theorem 2.6. In comparison, we show that under the conditions of Theorem 2.6, the spectral gap of the CKG Lindbladian with an ensemble of 1-design jumps is bounded with high probability.

Finally, our conditions on the spectrum and the structure of random unitary design jumps resemble previous works on chaotic Hamiltonians that apply the Eigenstate Thermalization Hypothesis (ETH) to prove the fast mixing of dissipative dynamics [8, 15]. In particular, in [8], the proposed algorithm implements a "rounded" Davies generator, yielding a physical Lindbladian that block-diagonalizes into components consisting of small-energy transitions. They propose their own version of ETH that relies on jump operators, for small Bohr frequencies ω, having independent Gaussian-distributed entries. The assumption that these entries are independent for the result is very strong, allowing them to conclude that their jump operators are both local *and* that distinct energy transitions are completely uncorrelated.

Our work shows fast mixing unconditionally for quantumly easy Hamiltonians, replacing the local jumps and ETH assumption for the rounded Davies generator with 1-design jump operators for the CKG Lindbladian. A similar ETH assumption to [8] would also yield fast-mixing for the CKG Lindbladian with local jumps, but more generally some approach must be taken to provably mitigate the correlations induced by implementing local jumps, in contrast with 1-design jump operators.

References

1 Eric R Anshuetz, Chi-Fang Chen, Bobak T Kiani, and Robbie King. Strongly interacting fermions are non-trivial yet non-glassy. *arXiv preprint arXiv:2408.15699*, 2024. doi:10.48550/arXiv.2408.15699.

2 Joran van Apeldoorn and András Gilyén. Improvements in quantum SDP-solving with applications. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 99:1–99:15, 2019. arXiv: 1804.05058 doi:10.4230/LIPIcs.ICALP.2019.99.

3 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. Earlier version in FOCS'17. arXiv: 1705.01843 doi:10.22331/q-2020-02-14-230.

4 Christian W. Bauer, Zohreh Davoudi, A. Baha Balantekin, Tanmoy Bhattacharya, Marcela Carena, Wibe A. de Jong, Patrick Draper, Aida El-Khadra, Nate Gemelke, Masanori Hanada, Dmitri Kharzeev, Henry Lamm, Ying-Ying Li, Junyu Liu, Mikhail Lukin, Yannick Meurice, Christopher Monroe, Benjamin Nachman, Guido Pagano, John Preskill, Enrico Rinaldi, Alessandro Roggero, David I. Santiago, Martin J. Savage, Irfan Siddiqi, George Siopsis, David Van Zanten, Nathan Wiebe, Yukari Yamauchi, Kübra Yeter-Aydeniz, and Silvia Zorzetti. Quantum simulation for high-energy physics. *PRX Quantum*, 4:027001, May 2023. doi:10.1103/PRXQuantum.4.027001.

5 Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114(9):090502, 2015. arXiv: 1412.4687 doi:10.1103/PhysRevLett.114.090502.

6 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 27:1–27:14, 2019. arXiv: 1710.02581 doi:10.4230/LIPIcs.ICALP.2019.27.

7 Fernando G. S. L. Brandão and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 415–426, 2017. arXiv: 1609.05537 doi:10.1109/FOCS.2017.45.

8 Chi-Fang Chen and Fernando G. S. L. Brandão. Fast thermalization from the eigenstate thermalization hypothesis, 2021. doi:10.48550/arXiv.2112.07646.

9 Chi-Fang Chen, Alexander M. Dalzell, Mario Berta, Fernando G. S. L. Brandão, and Joel A. Tropp. Sparse random Hamiltonians are quantumly easy. *Phys. Rev. X*, 14:011014, February 2024. doi:10.1103/PhysRevX.14.011014.

10 Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation. arXiv: 2303.18224, 2023.

11 Chi-Fang Chen, Michael J Kastoryano, and András Gilyén. An efficient and exact noncommutative quantum Gibbs sampler. *arXiv preprint arXiv:2311.09207*, 2023. doi:10.48550/arXiv.2311.09207.

12 Hongrui Chen, Bowen Li, Jianfeng Lu, and Lexing Ying. A randomized method for simulating Lindblad equations and thermal state preparation. *arXiv preprint arXiv:2407.06594*, 2024. doi:10.48550/arXiv.2407.06594.

13 Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the 35th ACM Symposium on the Theory of Computing (STOC)*, pages 59–68, 2003. arXiv: quant-ph/0209131 doi:10.1145/780542.780552.

14 Anirban Narayan Chowdhury and Rolando D. Somma. Quantum algorithms for Gibbs sampling and hitting-time estimation. *Quantum Information and Computation*, 17(1&2):41–64, 2017. arXiv: 1603.02940 doi:10.26421/QIC17.1-2.

15 Zhiyan Ding, Chi-Fang Chen, and Lin Lin. Single-ancilla ground state preparation via Lindbladians. *Phys. Rev. Res.*, 6:033147, August 2024. doi:10.1103/PhysRevResearch.6.033147.

16 Zhiyan Ding, Bowen Li, and Lin Lin. Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition. *arXiv preprint arXiv:2404.05998*, 2024. doi:10.48550/arXiv.2404.05998.

17 Ioana Dumitriu and Soumik Pal. Sparse regular random graphs: Spectral density and eigenvectors. *The Annals of Probability*, 40(5), September 2012. doi:10.1214/11-aop673.

18 Ernesto Estrada and Juan A. Rodríguez-Velázquez. Subgraph centrality in complex networks. *Phys. Rev. E*, 71:056103, May 2005. doi:10.1103/PhysRevE.71.056103.

19 András Gilyén, Chi-Fang Chen, Joao F Doriguello, and Michael J Kastoryano. Quantum generalizations of Glauber and Metropolis dynamics. *arXiv preprint arXiv:2405.20322*, 2024. doi:10.48550/arXiv.2405.20322.

20 András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics [full version], 2018. arXiv: 1806.01838

21 Fernando G.S L. Brandão, Richard Kueng, and Daniel Stilck França. Faster quantum and classical SDP approximations for quadratic binary optimization. *Quantum*, 6:625, January 2022. doi:10.22331/q-2022-01-20-625.

22 Jeongwan Haah, Matthew B. Hastings, Robin Kothari, and Guang Hao Low. Quantum algorithm for simulating real time evolution of lattice Hamiltonians. In *Proceedings of the 59th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 350–360, 2018. arXiv: 1801.03922 doi:10.1109/FOCS.2018.00041.

23 Jiaqing Jiang and Sandy Irani. Quantum Metropolis sampling via weak measurement. *arXiv preprint arXiv:2406.16023*, 2024. doi:10.48550/arXiv.2406.16023.

24 Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local Hamiltonian problem. In Kamal Lodaya and Meena Mahajan, editors, *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science*, pages 372–383, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. doi:10.48550/arXiv.quant-ph/0406180.

25 A Kitaev. Hidden correlations in the Hawking radiation and thermal noise, talk given at KITP. *Santa Barbara, California*, 12, 2015.

26 Alexei Kitaev. A simple model of quantum Holography (part 2). *Entanglement in strongly-correlated quantum matter*, page 38, 2015.

27 Guang Hao Low and Isaac L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118(1):010501, 2017. arXiv: 1606.02685 doi:10.1103/PhysRevLett.118.010501.

28 Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. arXiv: 1610.06546 doi:10.22331/q-2019-07-12-163.

29 Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, 2020. arXiv: 1808.10402 doi:10.1103/RevModPhys.92.015003.

30 John Preskill. Simulating quantum field theory with a quantum computer. arXiv: 1811.10085, 2018.

31 Patrick Rall, Chunhao Wang, and Pawel Wocjan. Thermal state preparation via rounding promises. *Quantum*, 7:1132, October 2023. doi:10.22331/q-2023-10-10-1132.

32 Akshar Ramkumar and Mehdi Soleimanifar. Mixing time of quantum Gibbs sampling for random sparse Hamiltonians. arXiv: 2411.04454, 2024.

33 Cambyse Rouzé, Daniel Stilck França, and Álvaro M Alhambra. Efficient thermalization and universal quantum computing with quantum Gibbs samplers. *arXiv preprint arXiv:2403.12691*, 2024. doi:10.48550/arXiv.2403.12691.

34 Subir Sachdev and Jinwu Ye. Gapless spin-fluid ground state in a random quantum Heisenberg magnet. *Physical Review Letters*, 70(21):3339–3342, May 1993. doi:10.1103/physrevlett.70.3339.

35 Sushant Sachdeva and Nisheeth K Vishnoi. Matrix inversion is as easy as exponentiation. *arXiv preprint arXiv:1305.0526*, 2013. doi:10.48550/arXiv.1305.0526.

36 Oles Shtanko and Ramis Movassagh. Preparing thermal states on noiseless and noisy programmable quantum processors. *arXiv preprint arXiv:2112.14688*, 2021. doi:10.48550/arXiv.2112.14688.

37 Brian Swingle and Mike Winer. Bosonic model of quantum Holography. *Physical Review B*, 109(9):094206, 2024. doi:10.48550/arXiv.2311.01516.

38 Kristan Temme. Lower bounds to the spectral gap of Davies generators. *Journal of Mathematical Physics*, 54(12), December 2013. doi:10.1063/1.4850896.

39 Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015. arXiv: 1501.01571 doi:10.1561/2200000048.

40 Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak, and Matthias Troyer. Solving strongly correlated electron models on a quantum computer. *Phys. Rev. A*, 92:062318, December 2015. doi:10.1103/PhysRevA.92.062318.

41 Michael Winer, Richard Barney, Christopher L Baldwin, Victor Galitski, and Brian Swingle. Spectral form factor of a quantum spin glass. *Journal of High Energy Physics*, 2022(9):1–47, 2022. doi:10.48550/arXiv.2203.12753.

42 Pawel Wocjan and Kristan Temme. Szegedy walk unitaries for quantum maps. *Communications in Mathematical Physics*, 402(3):3201–3231, 2023. doi:10.48550/arXiv.2107.07365.

43 Daniel Zhang, Jan Lukas Bosse, and Toby Cubitt. Dissipative quantum Gibbs sampling. *arXiv preprint arXiv:2304.04526*, 2023. doi:10.48550/arXiv.2304.04526.

A CKG Lindbladian in the energy basis

We consider a quantum system consisting of basis states $|e_i\rangle$ and a Hamiltonian $\boldsymbol{H}$. We choose some jump operators $\boldsymbol{A}^a$ and denote $A^a_{lm} = \langle l| \boldsymbol{A}^a |m\rangle$. Notate the energy eigenstates as $|j\rangle$ with energy E_j. We independently calculate the three parts of the Lindbladian: the transition term $\mathcal{L}_t$, the decay term $\mathcal{L}_d$, and coherent term $\mathcal{L}_c$ so that $\mathcal{L}_\beta = \mathcal{L}_c + \mathcal{L}_t + \mathcal{L}_d$. First, as mentioned above, the OFT of the jump operator $\boldsymbol{A}^a$ is $\hat{\boldsymbol{A}}^a(\omega) = \sum_{lm} A^a_{lm} \hat{f}(\omega - \nu_{lm})|l\rangle\langle m|$, where $\nu_{lm} = E_l - E_m$. To represent superoperators as linear maps, we vectorize operators with respect to the basis of operators $|m_1\rangle\langle m_2|$. In particular, $|\mathbf{m}\rangle$ with $\mathbf{m} = (m_1, m_2)$ will notate the basis operator $|m_1\rangle\langle m_2|$. Now, we may expand $\mathcal{L}_t, \mathcal{L}_d, \mathcal{L}_c$. The full calculations are performed in the arXiv version of this paper [32], yielding the following expressions:

$$
\begin{aligned}
\langle \mathbf{l}| \mathcal{L}_t |\mathbf{m}\rangle &= \sum_a A^a_{l_1 m_1} \overline{A^a_{l_2 m_2}} \theta(\nu_{l_1 m_1}, \nu_{l_2 m_2}), \\
\langle \mathbf{l}| \mathcal{L}_d |\mathbf{m}\rangle &= -\frac{1}{2}\left(\delta_{l_1 m_1} \sum_{a,j} \overline{A^a_{j m_2}} A^a_{j l_2} \theta(\nu_{j m_2}, \nu_{j l_2}) + \delta_{l_2 m_2} \sum_{a,j} \overline{A^a_{j l_1}} A^a_{j m_1} \theta(\nu_{j l_1}, \nu_{j m_1}) \right), \\
\langle \mathbf{l}| \mathcal{L}_c |\mathbf{m}\rangle &= \frac{1}{2}\Bigg(\delta_{l_1 m_1} \tanh(\beta \nu_{m_2 l_2}/4) \sum_{a,j} \overline{A^a_{j m_2}} A^a_{j l_2} \theta(\nu_{j m_2}, \nu_{j l_2}) - \\
&\qquad \delta_{l_2 m_2} \tanh(\beta \nu_{l_1 m_1}/4) \sum_{a,j} \overline{A^a_{j l_1}} A^a_{j m_1} \theta(\nu_{j l_1}, \nu_{j m_1}) \Bigg).
\end{aligned} \tag{5}
$$

Note that they can all be taken in terms of $\alpha(\nu) = \theta(\nu, \nu)$ using the identity $\theta(\nu_1, \nu_2) = \alpha(\frac{\nu_1+\nu_2}{2}) \exp\left(-\frac{(\nu_1-\nu_2)^2}{8\sigma_E^2}\right)$.

B Graph-Local Jumps for Cyclic Graphs

In this section we prove Theorem 2.1. Consider a cyclic graph with n vertices with adjacency matrix $\boldsymbol{H}$, and eigenvectors $|j\rangle$.

The eigenbasis of a cyclic graph consists of vectors $|j\rangle = n^{-1/2} \sum_a \zeta_n^{-aj} |e_a\rangle$ with eigenvalues $2\cos\left(\frac{2\pi j}{n}\right)$. The jump operators on the graph are chosen to be graph-local $\boldsymbol{A}^a = n^{-1/2}|e_a\rangle\langle e_a|$, and therefore have coefficients $A^a_{lm} = n^{-3/2}\zeta_n^{a(l-m)}$. Now we observe that $\sum_a \zeta_n^{a(i-j)} = n\delta_{ij}$. We therefore have the relation that

$$
\sum_a A^a_{l_1 m_1} \overline{A^a_{l_2 m_2}} = n^{-3} \sum_a \zeta_n^{a((l_1-m_1)-(l_2-m_2))} = n^{-2}\delta_{(l_1-m_1)(l_2-m_2)}.
$$

As computed more explicitly in the arXiv version [32], we compute the components of the Lindbladian with these jump operators:

$$
\begin{aligned}
\langle \mathbf{l}| \mathcal{L}_t |\mathbf{m}\rangle &= n^{-2}\theta(\nu_{l_1 m_1}, \nu_{l_2 m_2})\delta_{(l_1-m_1)(l_2-m_2)}. \\
\langle \mathbf{l}| \mathcal{L}_d |\mathbf{m}\rangle &= -\frac{n^{-2}}{2}\delta_{l_1 m_1}\delta_{l_2 m_2}(\alpha(\nu_{j m_2}) + \alpha(\nu_{j m_1})) \\
\langle \mathbf{l}| \mathcal{L}_c |\mathbf{m}\rangle &= 0.
\end{aligned}
$$

The above formulae imply that the Lindbladian is block diagonal in the eigenbasis. The coherent term vanishes and the decay term is fully diagonal. Setting $k = m_1 - m_2$ and $k' = l_1 - l_2$, the transition term $\langle \mathbf{l}| \mathcal{L}_t |\mathbf{m}\rangle$ is nonzero only if $k = k'$, due to the factor

$\delta_{(l_1-m_1)(l_2-m_2)} = \delta_{(l_1-l_2)(m_1-m_2)}$. There is therefore one block corresponding to each k, which we will denote $\mathcal{L}^k$. The block for $k = 0$ is the classical block of the Markov chain on the diagonal entries of the state. Finding its spectral gap, and then lower bounding the eigenvalues of the remaining $n - 1$ blocks, yields a bound for the spectral gap of the Lindbladian.

B.1 Spectral Gap Lower Bound

We will first show that the spectral gap of the classical block is asymptotically $\Omega(n^{-1})$. We have by explicit calculation that

$$\begin{aligned}\langle \mathbf{l}| \mathcal{L}_t^0 |\mathbf{m}\rangle &= \frac{1}{n^2}\alpha(\nu_{lm})\\ \langle \mathbf{l}| \mathcal{L}_d^0 |\mathbf{m}\rangle &= -\frac{1}{n^2}\delta_{lm}\sum_j \alpha(\nu_{jm}).\end{aligned}$$

We will use the canonical path bound, a standard technique in the theory of classical Markov chains, to establish lower bounds on their spectral gaps. The canonical path lemma fixes a "canonical" path between each pair of vertices on a graph and obtains a corresponding spectral graph bound. For our purposes we let the canonical path between any two vertices to be the edge joining them, which obtains the following statement:

▶ **Lemma B.1.** *Say L^0 is a Markov chain generator with stationary state σ. Then, the spectral gap satisfies the following bound:*

$$\lambda \geq \min_{(l,m)} \frac{L^0_{lm}}{\sigma_l}.$$

Applying this bound in this case, and noting that the stationary state of this Markov chain is ρ_{ll}, we obtain the lower bound

$$\lambda \geq \min_{l\neq m} \frac{\alpha(\nu_{lm})n^{-2}}{\rho_{ll}}. \tag{6}$$

The first equality holds because every canonical path is length 1, so the only path containing the edge (l, m) is γ_{lm}. We may upper bound ρ_{ll} with $\rho_{ll} \leq \frac{e^{2\beta}}{\sum_i E_i} \leq n^{-1}\frac{e^{2\beta}}{e^{-2\beta}} = n^{-1}e^{4\beta}$. Moreover, $|\nu_{lm}| \leq 4$ since all energies lie in $[-2, 2]$, so α is bounded below by a positive constant C that is independent of n. We conclude that $\lambda \geq \min_{l\neq m} \frac{\alpha(\nu_{lm})n^{-2}}{\rho_{ll}} \geq \frac{Cn^{-2}}{e^{4\beta}n^{-1}} = \Omega(n^{-1})$, as desired.

Now, for the blocks with $k \neq 0$, we utilize the Gershgorin bound on the columns of the kth block of $-\mathcal{L}_\beta$, $-\mathcal{L}^k$, which states that the eigenvalues of $\mathcal{L}^k$ must be larger than $\min_{\mathbf{m}} \left[\langle \mathbf{m}|\mathcal{L}^k|\mathbf{m}\rangle - \sum_{\mathbf{l}\neq\mathbf{m}} |\langle \mathbf{l}|\mathcal{L}^k|\mathbf{m}\rangle|\right]$. This approach is very similar to the one outlined in [38] for the Davies generator. By explicitly calculating and bounding this term in the arXiv version of this paper-[32], we obtain that the eigenvalues of $\mathcal{L}^k$ are $\Omega(n^{-3})$. This completes the lower bound, showing that all the spectral gap in full must be $\Omega(n^{-3})$.

B.2 Spectral Gap Upper Bound

To prove the upper bound on the spectral gap, we consider the row vector v of length n^2, that is 1 on the indices that correspond to the block $k = 1$, and 0 elsewhere. As an operator, it takes the value 1 on one offdiagonal with a fixed $l_1 - l_2 = k$. When calculating

$(v\mathcal{L}_\beta)_{\mathbf{m}} = (\mathcal{L}_\beta^\dagger v)_{\mathbf{m}}$, we obtain the same formula as found in the Gershgorin bound calculation above – indeed, the values on the diagonal are all positive, while the off-diagonal values are negative:

$$\begin{aligned}(\mathcal{L}_\beta^\dagger v)_{\mathbf{m}} =& \frac{1}{n^2}\sum_{\mathbf{l}}\left(\frac{\alpha(\nu_{l_1m_1})+\alpha(\nu_{l_2m_2})}{2} - \alpha\left(\frac{\nu_{l_1m_1}+\nu_{l_2m_2}}{2}\right)\right)\\ &+\frac{1}{n^2}\sum_{\mathbf{l}}\alpha\left(\frac{\nu_{l_1m_1}+\nu_{l_2m_2}}{2}\right)\left(1-\exp\left(\frac{-(\nu_{l_1m_1}-\nu_{l_2m_2})^2}{8\sigma_E^2}\right)\right). \end{aligned} \tag{7}$$

The previous lower bound shows that each of these values is nonnegative. Since $l_1 - l_2 = m_1 - m_2 = 1$, $\nu_{l_1m_1} - \nu_{l_2m_2} = \nu_{l_1m_1} - \nu_{l_2m_2}$ is $O(n^{-1})$. The first term in the expression, since it is composed of summands that are second differences in α, is n terms that are $O(n^{-2})$ scaled by n^{-2} – it is therefore $O(n^{-3})$. The summands of the second term can be similarly estimated to be $O(n^{-2})$, so it is also $O(n^{-3})$. The terms of $(\mathcal{L}_\beta^\dagger v)_{\mathbf{m}}$ are therefore nonnegative and are at most Cn^{-3} for some constant C.

To prove the upper bound, we make use of the inner product $\langle\ ,\ \rangle_{\rho_\beta^{-1}}$ with respect to which $\mathcal{L}_\beta^\dagger$ is self-adjoint. Note that $\langle |i_1\rangle\langle i_2|, |j_1\rangle\langle j_2|\rangle_{\rho_\beta^{-1}} = \delta_{i_2j_1}\delta_{i_1j_2}(\rho_\beta)_{i_1i_1}^{1/2}(\rho_\beta)_{i_2i_2}^{1/2} \geq 0$. Hence, when $\langle\ ,\ \rangle_{\rho_\beta^{-1}}$ is expressed in the energy basis as vMw for a matrix M, M has nonnegative elements. We therefore may upper bound $\langle(\mathcal{L}_\beta^\dagger v), v\rangle_{\rho_\beta^{-1}}$ by $Cn^{-3}\langle v, v\rangle$, since the coefficients of $(\mathcal{L}_\beta^\dagger v)$ and v are nonnegative and $(\mathcal{L}_\beta^\dagger v)$ is dominated by $Cn^{-3}v$ for some $C > 0$. We conclude that $\frac{\langle(\mathcal{L}_\beta^\dagger v),v\rangle_{\rho_\beta^{-1}}}{\langle v,v\rangle_{\rho_\beta^{-1}}}$ is $O(n^{-3})$. Since $\mathcal{L}_\beta^\dagger$ is self-adjoint with respect to this inner product, we obtain that v has Rayleigh quotient $O(n^{-3})$. v is also orthogonal to I, since $\langle v, \rho_\beta\rangle_{\rho_\beta^{-1}} = \mathrm{Tr}(v\rho_\beta^{1/2}I\rho_\beta^{1/2}) = \mathrm{Tr}(v\rho_\beta) = 0$, where the last equality holds since as an operator v is zero along the diagonal. v has no overlap with I, the fixed point of $\mathcal{L}_\beta^\dagger$, and therefore its Rayleigh quotient is an upper bound on the spectral gap. The spectral gap must therefore also be $O(n^{-3})$. This completes the proof of Theorem 2.1.

C Bounded Degree Systems with 1-Design Jumps

In this section we prove Lemma 3.1 and Theorem 2.4, demonstrating an improvement over the result in Theorem 2.1 for local jumps in cyclic graphs.

Proof of Lemma 3.1. To prove Lemma 3.1, we make use of the expressions (5) for the transition, decay, and coherent parts of a general Lindbladian, but with simply one Haar random jump (or equivalently any 1-design since the second moments of the operators are equal). The transition term is

$$\langle\mathbf{l}|\,\mathcal{L}_t\,|\mathbf{m}\rangle = A_{l_1m_1}\overline{A_{l_2m_2}}\theta(\nu_{l_1m_1},\nu_{l_2m_2}).$$

The expectation of the product $A_{l_1m_1}\overline{A_{l_2m_2}}$ is zero if $l_1 \neq l_2$ or $m_1 \neq m_2$. The expectation of the norm squared of an element, on the other hand, is n^{-1}. By explicit calculation, we therefore obtain

$$\begin{aligned}\mathbb{E}\left[\langle\mathbf{l}|\,\mathcal{L}_t\,|\mathbf{m}\rangle\right] &= \frac{\delta_{l_1l_2}\delta_{m_1m_2}}{n}\alpha(\nu_{l_1m_1})\\ \mathbb{E}\left[\langle\mathbf{l}|\,\mathcal{L}_d\,|\mathbf{m}\rangle\right] &= -\frac{1}{2}\left(\frac{\delta_{l_1m_1}\delta_{l_2m_2}}{n}\sum_j(\alpha(\nu_{jm_2})+\alpha(\nu_{jm_1}))\right),\\ \mathbb{E}\left[\langle\mathbf{l}|\,\mathcal{L}_c\,|\mathbf{m}\rangle\right] &= 0.\end{aligned}$$

The final Lindbladian $\mathcal{L}_\mu$ is therefore completely diagonal except for a "classical block" $\mathcal{L}^0$ of indices $|(m,m)\rangle$, whose off-diagonal terms are populated by the elements of $\mathcal{L}_t$. The spectral gap of this Lindbladian is therefore the minimum of the values along the diagonal, which are all positive, and the spectral gap of the classical block.

As in Section B, we use the Lemma B.1 to bound the spectral gap of the classical block. Using this lemma, with the canonical path being the edge between a pair of vertices, we obtain the bound

$$\lambda \geq \min_{l\neq m} \frac{\alpha(\nu_{lm})n^{-1}}{\rho_{ll}}. \tag{8}$$

We may upper bound ρ_{ll} with $\rho_{ll} \leq \frac{e^{-E_{\min}\beta}}{\sum_i E_i} \leq n^{-1}\frac{e^{-E_{\min}\beta}}{e^{-E_{\max}\beta}} = n^{-1}e^{-O(1)} = O(n^{-1})$ due to the fact that $\beta\|\boldsymbol{H}\| = O(1)$. Similarly, since α is the convolution of a Gaussian of radius $\sigma_E = \beta^{-1}$ with $\gamma(\omega) = \exp\left(-\beta \max\left(\omega + \frac{\beta\sigma_E^2}{2}, 0\right)\right)$, the assumption that $\beta\|\boldsymbol{H}\| = O(1)$ again yields that α evaluated at ν_{lm} is $\Omega(1)$. Indeed, within $O(\beta^{-1})$ of any value of ν_{lm}, $\gamma(\omega)$ is $\Omega(1)$, and as a consequence $\alpha(\nu_{lm}) = \Omega(1)$. This yields a lower bound on λ of $\min_{l\neq m}\frac{\Omega(n^{-1})}{O(n^1)} = \Omega(1)$.

Now, we lower bound the diagonal elements outside of the classical block. Since each such element is of the form

$$\mathbb{E}\left[\langle \mathbf{m}|\,\mathcal{L}_d\,|\mathbf{m}\rangle\right] = -\frac{1}{2}\left(\frac{1}{n}\sum_j \alpha(\nu_{jm_2}) + \frac{1}{n}\sum_j \alpha(\nu_{jm_1})\right)$$

and we have already established that each α term is $\Omega(1)$, so the resulting diagonal values are all $\Omega(1)$. We conclude that the spectral gap of the Lindbladian is $\Omega(1)$. ◀

Proof of Theorem 2.4. We construct our Lindbladian by sampling $M = \Theta(\log(n))$ unnormalized jumps $\boldsymbol{A}^a$ from the 1-design $\mathcal{D}(U(n))$ as in Definition 2.2, each with a corresponding Lindbladian $\mathcal{L}_a$ (which has one jump $\boldsymbol{A}^a$ along with its adjoint, normalized by 2). Then, we want to prove that with high probability, $\mathcal{L}_\beta = \frac{2}{M}\sum_{a=1}^{M/2}\mathcal{L}_a$, the Lindbladian with all M of these jumps now normalized by the number of jumps, has spectral gap bounded by a constant.

To prove the result, we shall make use of the matrix Bernstein's inequality for our concentration bound:

▶ **Lemma C.1** (cf. [39]). *Say $X_1, \ldots, X_N$ are independent random $d \times d$ Hermitian matrices, such that $\mathbb{E}[X_i] = 0$ and $\|X_i\| \leq R$. Define $Y = \frac{1}{N}\sum_{i=1}^N X_i$, and say that $N\mathbb{E}[Y^2] \leq T$. Then* $\Pr(\|Y\| \geq t) \leq 2d\exp(-\frac{3}{2}\frac{Nt^2}{3T+Rt})$.

Call $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$, where $\mathcal{L}_\mu = \mathbb{E}_{\boldsymbol{A}\sim\mathcal{D}}[\mathcal{L}_\beta]$. Each of these operators has zero expectation. We have that $\delta\mathcal{L} = \frac{2}{M}\sum_{a=1}^{M/2}\delta\mathcal{L}_a$ is precisely the discrepancy between our Lindbladian $\mathcal{L}_\beta = \frac{2}{M}\sum_{a=1}^{M/2}\mathcal{L}_a$ and the expected Lindbladian $\mathcal{L}_\mu$. We will apply Bernstein's inequality for $X_a = \delta\mathcal{L}_a$, $Y = \delta\mathcal{L}$, and $N = \frac{M}{2}$. By Lemma 4.1, the CKG Lindbladian has operator norm $O(\log(\beta\|\boldsymbol{H}\|))$, which is $O(1)$ in this regime. We denote this upper bound $\frac{R}{2}$. We can now verify the condition $N\mathbb{E}[Y^2] \leq T$ in the statement of Lemma C.1, since we have that $\|\delta\mathcal{L}_a\| \leq \|\mathcal{L}_a\| + \|\mathcal{L}_\mu\| \leq R = \Theta(1)$, and

$$\frac{M}{2}\left\|\mathbb{E}\left[\left(\frac{\sum_{a=1}^{M/2}\delta\mathcal{L}_a}{M/2}\right)^2\right]\right\| \leq \frac{2}{M}\sum_{a=1}^{M/2}\|\delta\mathcal{L}_a\|^2 \leq R^2.$$

Every operator that satisfies detailed balance is Hermitian in some fixed basis, and therefore each $\mathcal{L}_a$, as well as $\mathcal{L}_\mu$, can be considered Hermitian. By linearity, the same holds true for $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$. The operators $\delta\mathcal{L}_a$ therefore satisfy the conditions of the matrix Bernstein inequality, and so their average $\delta\mathcal{L}$ satisfies

$$\Pr\left(\|\delta\mathcal{L}\| \geq t\right) \leq 2n^2 \exp\left(-\frac{3}{4}\frac{Mt^2}{3R^2 + Rt}\right),$$

where the n^2 is due to an overhead of the dimension of the Lindbladian.

For any constant t, there exists an $M = \Theta(\log(n))$ such that the term inside the exponential is at least $3\log(n)$, since R is a constant. The probability that $\|\delta\mathcal{L}\| \leq t$ is then arbitrarily close to 1 for sufficiently large n and choice of $M = \Theta(\log(n))$. By Lemma 3.1, $\mathcal{L}_\mu$ has a constant spectral gap bounded below by some C. By Weyl's theorem, the eigenvalues of $\mathcal{L}_\mu + \delta\mathcal{L}$ may differ by at most t from those of $\mathcal{L}_\mu$. Choosing $t \leq \frac{C}{2}$, it follows that $\mathcal{L}_a$, with any constant probability, has constant spectral gap. ◀

D Unbounded Degree Systems with 1-Design Jumps

Proof of Lemma 3.2. We follow the proof of Lemma 3.1. As in Lemma 3.1, we may obtain the following bound on the classical block:

$$\lambda \geq \min_{l \neq m} \frac{\alpha(\nu_{lm})n^{-1}}{\rho_{ll}} = \frac{\alpha(\nu_{lm})\left(\frac{1}{n}\sum_j \exp(-\beta E_j)\right)}{\exp(-\beta E_l)}. \tag{9}$$

An explicit calculation, as in [32], therefore yields that

$$\lambda = \Omega(\delta(n))$$

by assumption that $\delta(n)$ of the eigenvalues are within $O(\beta^{-1})$ of $\lambda_{\min}$. Now bounding the diagonal elements outside of the classical block, we see that

$$\mathbb{E}\left[\langle\mathbf{m}|\,\mathcal{L}_d\,|\mathbf{m}\rangle\right] = -\frac{1}{2}\left(\frac{1}{n}\sum_j \alpha(\nu_{jm_2}) + \frac{1}{n}\sum_j \alpha(\nu_{jm_1})\right).$$

Again, since $\delta(n)$ of eigenvalues are within $O(\beta^{-1})$ of $\lambda_{\min}$, the above sum is $\Omega(\delta(n))$, as desired. ◀

Proof of Theorem 2.6. We follow the proof of Theorem 2.4. Defining once again $\mathcal{L}_a$ to be the Lindbladian with one jump operator $\boldsymbol{A}^a$ and its adjoint (normalized by 2), and defining $\delta\mathcal{L}_a = \mathcal{L}_a - \mathcal{L}_\mu$, we can apply the matrix Bernstein's inequality to obtain

$$\Pr\left(\|\delta\mathcal{L}\| \geq t\right) \leq n^2 \exp\left(-\frac{3}{4}\frac{Mt^2}{3R^2 + Rt}\right),$$

since all the conditions of the inequality are again satisfied.

By Lemma 3.1, $\mathcal{L}_\mu$ has a constant spectral gap bounded below by some $C\delta(n)$. Selecting t to be $\frac{C}{2}\delta(n)$, there exists an $M = \Theta(\delta(n)^{-2}\log(\beta\|\boldsymbol{H}\|)^2\log(n))$ for which the value inside the exponential is at least $3\log(n)$. The probability that $\|\delta\mathcal{L}\| \leq t$ is therefore above any constant probability for sufficiently large n. By Weyl's theorem, the eigenvalues of $\mathcal{L}_\mu + \delta\mathcal{L}$ may differ by at most $t \leq \frac{C}{2}\delta(n)$. It follows that $\mathcal{L}_a$, with any constant probability, has spectral gap $\Omega(\delta(n))$. ◀

Optimal Locality and Parameter Tradeoffs for Subsystem Codes

Samuel Dai
Department of Physics, Northeastern University, Boston, MA, USA

Ray Li
Math & CS Department, Santa Clara University, CA, USA

Eugene Tang
Departments of Math and Physics, Northeastern University, Boston, MA, USA

Abstract

We study the tradeoffs between the locality and parameters of subsystem codes. We prove lower bounds on both the number and lengths of interactions in any D-dimensional embedding of a subsystem code. Specifically, we show that any embedding of a subsystem code with parameters $[[n, k, d]]$ into $\mathbb{R}^D$ must have at least M^* interactions of length at least ℓ^*, where

$$M^* = \Omega(\max(k, d)), \quad \text{and} \quad \ell^* = \Omega\left(\max\left(\frac{d}{n^{\frac{D-1}{D}}}, \left(\frac{kd^{\frac{1}{D-1}}}{n} \right)^{\frac{D-1}{D}} \right) \right).$$

We also give tradeoffs between the locality and parameters of commuting projector codes in D-dimensions, generalizing a result of Dai and Li [8]. We provide explicit constructions of embedded codes that show our bounds are optimal in both the interaction count and interaction length.

2012 ACM Subject Classification Hardware → Quantum error correction and fault tolerance; Theory of computation → Error-correcting codes

Keywords and phrases Quantum Error Correcting Code, Locality, Subsystem Code, Commuting Projector Code

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.4

Related Version *Full Version*: `arXiv:2503.22651` [9]

Funding *Ray Li*: Supported by NSF grant CCF-2347371.

Acknowledgements We thank Andy Liu, Shouzhen Gu, Zhiyang He for helpful feedback on our manuscript.

1 Introduction

Quantum computing necessitates the manipulation of fragile states of information. The most promising way towards large-scale fault-tolerant quantum computing involve the extensive use of quantum error-correcting codes (QECCs). Physical implementations of quantum computing hardware naturally favor architectures which are local in 2 or 3 spatial dimensions – architectures where the qubits are embedded in 2 or 3 dimensions, and interactions occur only between qubits that are spatially nearby. On the other hand, it has long been known that the constraint of spatial locality places severe limitations on the parameters of QECCs. For example, the Bravyi-Terhal [7] and Bravyi-Poulin-Terhal (BPT) [6] bounds state that a commuting projector code whose constraints are local in D-dimensions necessarily have code parameters satisfying, respectively,

$$d = O(n^{\frac{D-1}{D}}), \qquad \text{and} \qquad kd^{\frac{2}{D-1}} = O(n). \tag{1}$$

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 4; pp. 4:1–4:22
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

These bounds suggest that there are tradeoffs between better code performance and the cost of non-local implementation. Consequently, the *locality* of a QECC becomes another key factor to consider when choosing a code for applications.

What is the quantitative tradeoff between locality and code quality? This problem was initially investigated by Baspin and Krishna [3], who asked, for a quantum low-density parity-check (qLDPC) code in D-dimensions, how many "long-range" interactions must there be, and how long must those interactions be? Baspin and Krishna gave bounds for D-dimensional codes which are nearly optimal in certain parameter settings. For 2-dimensional codes, Dai and Li [8] improved the bounds to be tight across all parameter regimes and also gave matching constructions that saturate the upper bounds (see also Hong, et al. [11], who considered the special case $k = 1, d = \sqrt{n}$ for 2-dimensional codes). Dai and Li showed that an $[[n, k, d]]$ quantum code embedded in 2 spatial dimensions must have $\Omega(M^*)$ interactions of length $\Omega(\ell^*)$, where

$$M^* = \max(k, d), \qquad \text{and} \qquad \ell^* = \max(\frac{d}{\sqrt{n}}, \sqrt[4]{\frac{kd^2}{n}}). \tag{2}$$

Both the interaction count M^* and interaction length ℓ^* are tight in strong ways.

In this paper, we study the locality versus parameter tradeoffs for quantum *subsystem* codes. Bravyi [5] showed that the BPT bound could be violated by the use of local subsystem codes, providing 2D-local subsystem codes with parameters $k, d = \Theta(\sqrt{n})$. Subsystem codes are nevertheless constrained by locality. Bravyi [5] showed that a $[[n, k, d]]$ subsystem code whose gauge generators are local in a D-dimensional lattice embedding satisfies

$$d = O(n^{\frac{D-1}{D}}), \qquad \text{and} \qquad kd^{\frac{1}{D-1}} = O(n). \tag{3}$$

While previous work has made it clear that outperforming local quantum codes requires copious amounts of long-ranged interactions, it is not a priori clear whether the same requirements hold for subsystem codes. Is it possible that small violations of locality in the gauge generators suffice to define subsystem codes parametrically better than those allowed by Bravyi's bound? More concretely:

► **Question 1.** *How much non-locality is required for a subsystem code to exceed Bravyi's bound?*

We address Question 1 by demonstrating that subsystem codes, like their commuting projector counterparts, require an extensive number of long-ranged interactions to surpass Bravyi's bound. We also provide constructions of subsystem codes that show our lower bounds are tight in strong ways. Additionally, we also generalize the results of Dai and Li [8] from 2-dimensions to D-dimensions. Our work establishes optimal bounds on interaction lengths and counts for embeddings of both commuting projector codes and subsystem codes in any number of dimensions.

1.1 Main Result

We study subsystem codes whose gauge generators are not necessarily local. Our main result is a lower bound on the number and length of interactions in any D-dimensional embedding of a $[[n, k, d]]$ subsystem code. Formally, a D-dimensional embedding is a mapping of the code's n physical qubits into $\mathbb{R}^D$, such that any two qubits are at distance at least 1.

▶ **Theorem 2** (Main Result for Subsystem Codes). *For any $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a nontrivial*[1] *$[[n, k, d]]$ subsystem code with $kd^{\frac{1}{D-1}} \geq c_1 n$ or $d \geq c_1 n^{\frac{D-1}{D}}$ must have at least M^* interactions of length ℓ^*, where*

$$M^* = c_0 \cdot \max(k, d), \qquad \text{and} \qquad \ell^* = c_0 \cdot \max\left(\frac{d}{n^{\frac{D-1}{D}}}, \left(\frac{kd^{\frac{1}{D-1}}}{n}\right)^{\frac{D-1}{D}}\right). \tag{4}$$

Prior to this work, no such bounds of this form were known for subsystem codes aside from Bravyi's original bound. While such bounds were known for commuting projector codes, our bound shows that a locality versus parameter trade-off also holds for subsystem codes. Our result also generalizes Bravyi's bound [5], not only in that we (optimally) address the number and length of long-range interactions, but also in that we handle more general embeddings. Bravyi's bound [5] considers only embeddings onto a $n^{1/D} \times \cdots \times n^{1/D}$ lattice, but our bound applies to arbitrary embeddings, even those not constrained to a $O(n^{1/D}) \times \cdots \times O(n^{1/D})$ box (see Section 3 for further discussion).

Like for stabilizer codes, subsystem codes beyond the "local regime" – above the BPT bound for stabilizer codes, or above the Bravyi bound for subsystem codes – need copious amounts of non-locality. In particular, the number of required long-range interactions $\Omega(\max(k, d))$ is the same for both subsystem and stabilizer codes. Additionally, for codes with a large number k of logical qubits, the required length of the long range interactions is similar. For example, a 2-dimensionally embedded asymptotically good subsystem code (with $k, d = \Omega(n)$) needs $M^* = \Omega(n)$ interactions of length $\ell^* = \Omega(\sqrt{n})$ – the worst possible case – just as for stabilizer codes. Our results show that, compared to stabilizer codes, subsystem codes do not offer substantial improvements in locality outside of the "local regime," though they can offer some quantitative improvements in the interaction length.

We also provide matching constructions that show M^* and ℓ^* are optimal in strong ways (see Figure 2). An asymptotically good qLDPC code [14, 12] has $O(M^*) = O(\max(k, d))$ interactions of any length (see Theorem 1.3 of [8]). Since a stabilizer code can also be trivially regarded as a subsystem code, this shows that our bounds are tight in terms of interaction count. For optimality in the interaction length, we exhibit subsystem codes embedded in D-dimensions where all interactions are of length at most $O(\ell^*)$ (Theorem 21).

1.2 Generalizing [8] to D-dimensions

We also show that the bounds in [8] can be generalized to D-dimensional embeddings and to commuting projector codes.

▶ **Theorem 3** (Generalization of [8] to D-dimensions). *For any $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a nontrivial $[[n, k, d]]$ commuting projector code with $kd^{\frac{2}{D-1}} \geq c_1 n$ or $d \geq c_1 n^{\frac{D-1}{D}}$ must have at least M^* interactions of length ℓ^*, where*

$$M^* = c_0 \cdot \max(k, d), \qquad \text{and} \qquad \ell^* = c_0 \cdot \max\left(\frac{d}{n^{\frac{D-1}{D}}}, \left(\frac{kd^{\frac{2}{D-1}}}{n}\right)^{\frac{D-1}{2D}}\right). \tag{5}$$

[1] Nontrivial here simply means that $k > 0$.

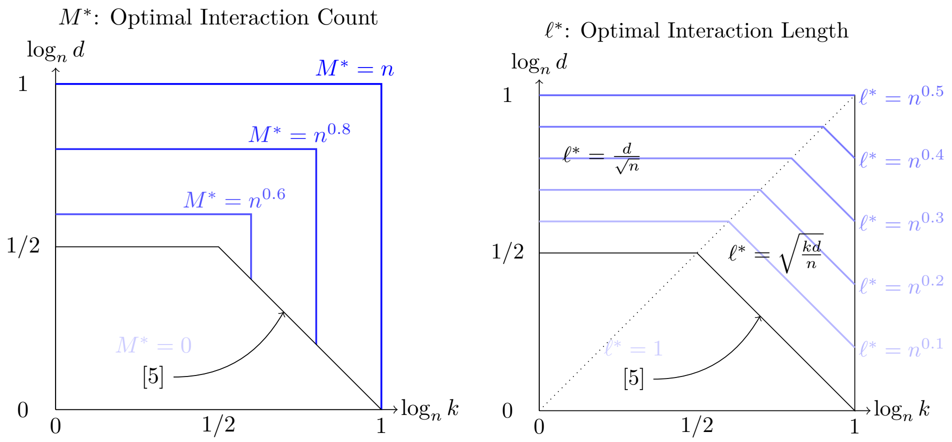

Figure 1 The (asymptotically) optimal interaction count and length for subsystem codes in 2D: A $[[n, k, d]]$ subsystem code need at least $\Omega(M^*)$ interactions of length $\Omega(\ell^*)$, where M^* is plotted on the left and ℓ^* is plotted on the right. Above, we plot the contours of k vs. d tradeoffs for various values of the Interaction Count or Interaction Length. Everywhere, big-O is suppressed for clarity.

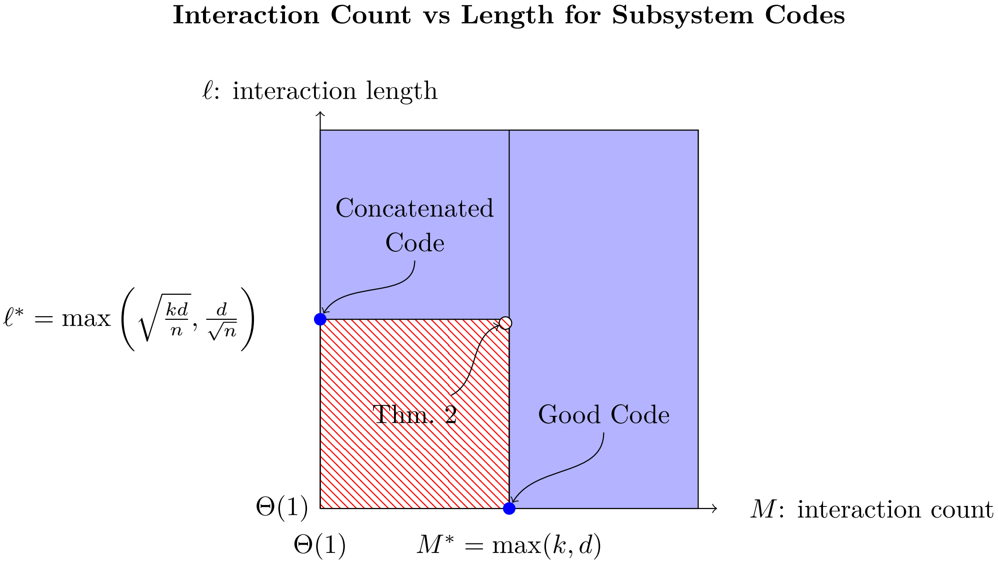
Interaction Count vs Length for Subsystem Codes
ℓ: interaction length
Concatenated Code
$\ell^* = \max\left(\sqrt{\frac{kd}{n}}, \frac{d}{\sqrt{n}}\right)$
Thm. 2
Good Code
$\Theta(1)$
$\Theta(1)$
$M^* = \max(k, d)$
M: interaction count

Figure 2 Schematic diagram illustrating the optimality of our lower bounds for all n, k, d: A point (M, L) represents that there is a code with $O(M)$ interactions of length $\omega(L)$. Blue shaded region is achievable, red lined region is unachievable. Our lower bound shows that (M, ℓ) with $M \le o(M^*)$ and $\ell \le o(\ell^*)$ is impossible, where M^* and ℓ^* are the optimal interaction count and length, respectively, given by Theorem 2. There is a construction (good qLDPC code) with $O(M^*)$ interactions of any length, and another construction (concatenated local code, Theorem 21) with zero interactions of length $\omega(\ell^*)$.

We now compare our works to prior works. First, we note that, when setting $D = 2$, our bound matches the bounds in [8] up to the implied constant. For $D > 2$ dimensions, the only prior bounds we are aware of are due to Baspin and Krishna [3]. Theirs match our bounds up to polylog factors when $d \ge \sqrt{kn}$ and when $k = \Theta(n)$, and we improve their bounds in the remaining parameter regimes. We also generalize their results; our results hold for all commuting projector codes, whereas theirs hold only for qLDPC codes.

Similar to Theorem 2 and [8], our M^* and L^* in Theorem 3 are optimal up to constant factors. Any asymptotically good qLDPC code [14, 12] is a commuting projector code with at most $O(M^*) = O(\max(k,d))$ interactions of any length. Further, we exhibit stabilizer codes embedded in D dimensions, all of whose interactions are of length at most $O(\ell^*)$; see Theorem 23.

1.3 Organization of the Paper

We divide the proof of Theorem 2 into two parts:

▶ **Theorem 4** (Main Result – Part 1). *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a nontrivial $[[n,k,d]]$ subsystem or commuting projector code with $d \geq c_1 n^{\frac{D-1}{D}}$ must have at least $c_0 d$ interactions of length at least $c_0 \frac{d}{n^{\frac{D-1}{D}}}$.*

▶ **Theorem 5** (Main Result – Part 2). *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a $[[n,k,d]]$ subsystem code with $kd^{\frac{1}{D-1}} \geq c_1 n$ must have at least $c_0 k$ interactions of length at least $c_0 (\frac{kd^{\frac{1}{D-1}}}{n})^{\frac{D-1}{D}}$.*

Theorem 4 implies Theorem 2 in the regime where $d \geq k$, and Theorem 5 implies Theorem 2 in the regime when $d \leq k$. Note that Theorem 4 also implies Theorem 3 when $d \geq \sqrt{kn}$. The remaining case of Theorem 3 is when $d \leq \sqrt{kn}$, which we prove in Theorem 6.

▶ **Theorem 6** (Generalization of [8] when $d \leq \sqrt{kn}$). *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a $[[n,k,d]]$ commuting projector code with $kd^{\frac{2}{D-1}} \geq c_1 n$ must have at least $c_0 k$ interactions of length $c_0 (\frac{kd^{\frac{2}{D-1}}}{n})^{\frac{D-1}{2D}}$.*

2 Preliminaries

Notation and Definitions

We use standard Landau notation $O(\cdot), \Omega(\cdot), \Theta(\cdot), o(\cdot), \omega(\cdot)$. We also use the notations $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$, which are variants of $O(\cdot)$ and $\Omega(\cdot)$, respectively, that ignore logarithmic factors. For example, $f(n) = \tilde{O}(h(n))$ means that there exists an integer k such that $f(n) = O(h(n)\log^k n)$. For a set S, we write $S^{\leq D} \overset{\text{def}}{=} S \cup S^2 \cup \cdots \cup S^D$.

In $\mathbb{R}^D$, *distance* refers to Euclidean (ℓ_2) distance unless otherwise specified. We sometimes also use the ℓ_∞-*distance* of two points $(x,y),(x',y') \in \mathbb{R}^D$, which is $\max(|x-x'|,|y-y'|)$. A *grid tiling* is a division of $\mathbb{R}^d$ given by axis aligned hyperplanes equally spaced at a fixed distance w. Throughout, a *box* is always a set of the form $[a_1,b_1] \times \cdots \times [a_D,b_D]$. In particular, boxes contain their boundary and are axis-parallel. A *cube* is a box all of whose side lengths are equal: $b_1 - a_1 = b_2 - a_2 = \cdots = b_D - a_D$.

An *embedded set* in $\mathbb{R}^D$ is a finite set $Q \subset \mathbb{R}^D$ with pairwise (ℓ_2) distance at least 1. A function $f : \mathbb{R}^D \to \mathbb{N}$ is *finitely supported* if $f(x) \neq 0$ for finitely many $x \in \mathbb{R}^D$. For a finitely supported function $f : \mathbb{R}^D \to \mathbb{N}$ and a region $R \subset \mathbb{R}^D$, define, by abuse of notation, $f(R) = \sum_{i \in R; f(i) \neq 0} f(i)$. We will be primarily concerned with the finitely supported function given by Definition 8.

2.1 Quantum codes

We associate the pure states of a qubit with unit vectors in $\mathbb{C}^2$ and pure n-qubit states with unit vectors in $(\mathbb{C}^2)^{\otimes n}$. Let $\mathcal{P}$ denote the (single-qubit) Pauli group, which consists of the Pauli matrices $\mathsf{I}, \mathsf{X}, \mathsf{Y}, \mathsf{Z}$, and their scalar multiples by $\{\pm 1, \pm i\}$. The n-qubit Pauli group is $\mathcal{P}_n = \mathcal{P}^{\otimes n}$. Given $P \in \mathcal{P}_n$, its *weight* $|P|$ is the number of tensor components not equal to I.

A *quantum error-correcting code* $\mathcal{C}$ is a subspace of $(\mathbb{C}^2)^{\otimes n}$. The parameter n is called the *(block) length* of the code. We define the *dimension* k of the code to be $k = \log_2(\dim \mathcal{C})$.

Stabilizer Codes

A *stabilizer group* $\mathcal{S}$ is an abelian subgroup of the n-qubit Pauli group $\mathcal{P}_n$ that does not contain $-\mathsf{I}$. A *stabilizer code* $\mathcal{Q} = \mathcal{Q}(\mathcal{S}) \subseteq (\mathbb{C}^2)^{\otimes n}$ is defined to be the subspace of states left invariant under the action of the stabilizer group $\mathcal{S}$, i.e. $\mathcal{Q} = \{|\psi\rangle : \ \mathsf{S}|\psi\rangle = |\psi\rangle, \ \forall \mathsf{S} \in \mathcal{S}\}$. Being an abelian group, we can describe $\mathcal{S}$ by $n-k$ independent generators $\{\mathsf{S}_1, ..., \mathsf{S}_{n-k}\}$, where k is the *dimension* of the code. The *distance* d is the minimum weight of an error $E \in \mathcal{P}_n$ that maps a codeword in $\mathcal{Q}$ to another codeword. A quantum code $\mathcal{Q}$ with distance d can correct up to $d-1$ qubit erasures.

Subsystem Codes

A *subsystem code* is a choice of decomposition of a stabilizer code $\mathcal{C}$ into a tensor product $\mathcal{C} = \mathcal{A} \otimes \mathcal{B}$, where $\mathcal{A} \cong (\mathbb{C}^2)^{\otimes k}$ and $\mathcal{B} \cong (\mathbb{C}^2)^{\otimes g}$ are the *logical* and *gauge* parts of $\mathcal{C}$, respectively. The dimension k of a subsystem code is defined as the number of qubits encoded in its logical subsystem $\mathcal{A}$. One can view a subsystem code as a stabilizer code that can encode $k+g$ logical qubits, but only k of the logical qubits are actually used to protect information.

We can define a subsystem code by starting with a stabilizer code $\mathcal{S}$ given by $n-k-g$ independent stabilizer generators, with $k+g$ logical qubits associated with $k+g$ pairs of logical operators $\bar{X}_1, \bar{Z}_1, \ldots, \bar{X}_{k+g}, \bar{Z}_{k+g}$. The first k logical qubits are used to encode information, and the last g logical qubits are called *gauge qubits*. The *gauge group* of the subsystem is the group $\mathcal{G} = \langle \mathcal{S}, \bar{X}_{k+1}, \bar{Z}_{k+1}, \ldots, \bar{X}_{k+g}, \bar{Z}_{k+g} \rangle$. Given the gauge group $\mathcal{G}$, the code's stabilizer group $\mathcal{S}$ can be recovered as the center of $\mathcal{G}$, so a subsystem code is uniquely defined by its gauge group. Any stabilizer code can be equivalently regarded as a subsystem code whose gauge group is abelian, so stabilizer codes form a subset of subsystem codes.

For subsystem codes, we make a distinction between bare logical operations, which act trivially on the gauge qubits, and dressed logical operators, which may not. Formally, (non-trivial) bare logical operators are elements of $\mathsf{C}(\mathcal{G}) \setminus \mathcal{G}$, where $\mathsf{C}(\mathcal{G})$ denotes the centralizer of $\mathcal{G}$, and (non-trivial) dressed logical operators are elements of $\mathsf{C}(\mathcal{S}) \setminus \mathcal{G}$. Note that for stabilizer codes there is no distinction. The distance d of a subsystem code is defined as the minimum weight of a non-trivial dressed logical operator, i.e., $d = \min_{P \in \mathsf{C}(\mathcal{G}) \setminus \mathcal{G}} |P|$. We will sometime denote a subsystem code $\mathcal{C}$ with n physical qubits, k logical qubits, distance d, and g gauge qubits by $\mathcal{C} = [[n, k, d, g]]$.

Commuting Projector Codes

A commuting projector code $\mathcal{C} \subseteq (\mathbb{C})^{\otimes n}$ is a subspace defined by a set of pairwise commuting projections $\{\Pi_1, \ldots, \Pi_m\}$. The code $\mathcal{C}$ is the subspace of states left invariant by all projections Π_i. Every stabilizer code is also a commuting projector code where the defining projections are of Pauli type, i.e., $\Pi_i = (\mathsf{I} + P_i)/2$, for some Pauli operator $P_i \in \mathcal{P}_n$. For the purposes of establishing our locality bounds, the only properties we need of commuting projector codes is the fact that all the properties of correctable sets for stabilizer codes, i.e., those listed in Lemma 11, continue to hold without modification for commuting projector codes [10]. Finally, we note that while stabilizer codes can be considered a subset of both subsystem and commuting projector codes, there is no direct relation between subsystem and commuting projector codes themselves.

Quantum codes in D dimensions

Given a finite set S, an *embedding* of S into $\mathbb{R}^D$ is a map $\iota : S \to \mathbb{R}^D$ such that $\|\iota(s_i)-\iota(s_j)\| \geq 1$ for all distinct $s_i, s_j \in S$. The image $\iota(S)$ is then said to be an *embedded set*. Throughout, the embedding map will usually be implicit. We identify the qubits Q of a quantum code with a D-dimensional embedded set, which we continue to call $Q \subset \mathbb{R}^D$ by abuse of notation. By further abuse of notation, we refer to $Q \subset \mathbb{R}^D$ as the embedding of the qubits Q. When the set of qubits Q is understood, given a subset $V \subset Q$, we write $\overline{V} \overset{\text{def}}{=} Q \setminus V$ to denote the *complement* of V in Q.

▶ **Definition 7** (Interactions). *Given an embedding $Q \subset \mathbb{R}^D$ of a quantum code C, either a commuting projector code or a subsystem code,* interactions *of the code are defined with respect to a specific set of generators for that code. In the case that C is a commuting projector code with defining projections $\{\Pi_1, \cdots, \Pi_m\}$, we say that a pair of qubits $q, p \in Q$ define an* interaction *if p and q are both in the support of some projection Π_i. Similarly, if C is a subsystem code with a set of generators $\{G_1, \cdots, G_m\}$ for its gauge group, we say that $p, q \in Q$ define an* interaction *if p and q are both in the support of some gauge generator G_i. In both cases, the* length *of an interaction (p, q) is defined to be the ℓ_2 distance between p and q.*

Interactions are always defined with respect to a particular set of generators (either projector or gauge), but throughout we assume that the generator set is fixed (but otherwise arbitrary), and thus the set of interactions is fixed as well. Note that for subsystem codes, interactions are always defined with respect to gauge generators *and not* stabilizer generators*. In particular, since each stabilizer generator is generally a product of multiple gauge generators, it is possible for a subsystem code to have local gauge generators, but non-local stabilizer generators. Indeed, it is only in such cases that a separation in the locality bounds for stabilizer and subsystem codes is possible.*

In the proofs of our results, we typically consider a fixed interaction length ℓ. We then refer to an interaction as *bad* if the length is at least ℓ and as *good* if the length is less than ℓ. Intuitively, good interactions are easier to deal with for our proof; they are effectively local, and can be treated in a similar to how local interactions are treated in the original proofs of the BPT and Bravyi bounds. To control the number of bad interactions, we introduce the following function that counts the number of bad interactions that a particular qubit participates in:

▶ **Definition 8** (Interaction counter). *Given a quantum code with a D-dimensional embedding $Q \subset \mathbb{R}^D$, let $f_{\geq \ell} : \mathbb{R}^D \to \mathbb{N}$ denote the* interaction counting function*, where $f_{\geq \ell}(q)$ for $q \in Q$ equals the number of interactions of length at least ℓ that qubit q participates in, and $f_{\geq \ell}(\cdot) = 0$ outside of Q.*

Correctable sets

Like in previous works [2, 3, 1, 8], we analyze the limitations of quantum codes using correctable sets. Intuitively, a subset $U \subset Q$ of qubits is correctable if the code can correct the erasure of the qubits in U. We state the definition of a correctable set for completeness, though we only interface with the definition indirectly using Lemmas 11, 12, and 13.

▶ **Definition 9** (Correctable set). *Let $U \subset Q$ be a subset of qubits in a quantum code, and let $\overline{U} = Q \backslash U$. Let $\mathcal{D}[\overline{U}]$ and $\mathcal{D}[Q]$ denote the space of density operators associated with the sets of qubits $\overline{U}$ and Q respectively. The set U is* correctable *if there exists a recovery channel $\mathcal{R} : \mathcal{D}[\overline{U}] \to \mathcal{D}[Q]$ such that for any code state ρ, we have $\mathcal{R}(\mathrm{Tr}_U(\rho)) = \rho$.*

For an embedded set of qubits $Q \subset \mathbb{R}^D$, we will say that a region $R \subset \mathbb{R}^D$ is correctable if the subset of all qubits contained in R is a correctable subset. As an abuse of terminology, we will often refer to subsets of qubits as regions. For stabilizer and commuting projector codes, a region being correctable is equivalent to having no non-trivial logical operators supported on that region. For subsystem codes, a region $U \subset Q$ is correctable if and only if no non-trivial *dressed* logical operators are supported on U. Note that U being correctable also implies that no non-trivial *bare* logical operators are supported on U, but the converse does not necessarily hold. This motivates the following definition.

▶ **Definition 10** (Dressed-Cleanable). *If there are no non-trivial* bare *logical operators supported on a region $U \subset Q$, then we say that U is* dressed-cleanable *[15].*

We use the following notions in Lemma 11 to reason about correctable sets in subsystem codes and commuting projector codes. In a quantum code with qubits Q, say sets $U_1, \ldots, U_\ell \subset Q$ are *decoupled* if there are no interactions between two distinct U_i's. For a set $U \subset Q$, let $\partial U = \partial_+ U \cup \partial_- U$ be the *boundary* of U, where $\partial_+ U$ denotes the *outer boundary* of U, the set of qubits outside U that have an interaction with U, and $\partial_- U = \partial_+ \overline{U}$ is the *inner boundary* of U.

▶ **Lemma 11** ([7, 10, 15]). *Let Q be the qubits of a $[[n, k, d]]$ commuting projector or subsystem code $\mathcal{C}$.*

1. **Subset Closure:** *Let $U \subset Q$ be a correctable set. Then any subset $W \subset U$ is correctable.*
2. **Distance Property:** *Let $U \subset Q$ with $|U| < d$. Then U is correctable.*
3. **Union Lemma:** *Let $U_1, \ldots, U_\ell$ be decoupled, and let each U_i be correctable. If $\mathcal{C}$ is a subsystem code, then $\bigcup_{i=1}^{\ell} U_i$ is dressed-cleanable. If $\mathcal{C}$ is a commuting projector code, then $\bigcup_{i=1}^{\ell} U_i$ is correctable.*
4. **Expansion Lemma:** *Let $U, T \subset Q$ be correctable sets such that $T \supset \partial U$. Then $T \cup U$ is correctable.*

A key point in Lemma 11 is that the union lemma differs for commuting projector and subsystem codes. For subsystem codes, the union of decoupled and correctable sets is *not necessarily correctable* – only dressed-cleanable [15]. In general, being dressed-cleanable is weaker than being correctable. One of the major problems with generalizing Theorem 3 from commuting projector codes to subsystem codes is that the union lemma for subsystem codes only allows the conclusion that the union of correctable sets is dressed-cleanable. This version of the union lemma is too weak to adapt the original proof of Theorem 3 to subsystem codes. Instead, we take an alternative approach in proving Theorem 4 which is based solely on the expansion lemma.

The usefulness of reasoning about correctable sets is that the sizes of the correctable sets in a quantum code directly give bounds on the parameters:

▶ **Lemma 12** (AB Lemma – Implicit in [5], Section VIII)**.** *Suppose that the qubits Q of a $[[n,k,d]]$ subsystem code can be partitioned as $Q = A \sqcup B$. If A is dressed-cleanable, then $k \leq |B|$.*

▶ **Lemma 13** (ABC Lemma [6])**.** *Suppose that the qubits Q of a $[[n,k,d]]$ commuting projector code can be partitioned as $Q = A \sqcup B \sqcup C$. If A and B are correctable, then $k \leq |C|$.*

2.2 Geometric Lemmas

In this section, we give two lemmas about D-dimensional embeddings of sets. The first, Lemma 14, allows us to generalize our results from lattice embeddings to arbitrary embeddings.

▶ **Lemma 14** (Point Density)**.** *Let $R \subseteq \mathbb{R}^D$ be a box with side lengths $L_1 \geq \cdots \geq L_D$. Suppose $Q \subseteq \mathbb{R}^D$ is an embedded set. Then*

$$|R \cap Q| \leq \frac{2^D}{\mathrm{vol}(B_D)} \prod_{i=1}^{D} (1 + L_i) \qquad (6)$$

where B_D is the unit ball in $\mathbb{R}^D$.

We refer the reader to [9] for the proof.

The second lemma, Lemma 15, utilizes the probabilistic method to generate a grid tiling that allows us to maintain a convenient distribution of the qubits and bad interactions in our embeddings (see Figure 3).

▶ **Lemma 15** (Tiling Lemma)**.** *Let $X, Y \subseteq \mathbb{R}^D$ be two multi-sets. Let w and ℓ be positive integers with $w \geq 4\ell$. There exists a tiling of $\mathbb{R}^D$ using hypercubes of side length w such that:*

1. *at most a $(4\ell D/w)^2$ fraction of points in X are within ℓ^∞-distance 2ℓ of a codimension-2 face of some hypercube,*
2. *at most a $8\ell D/w$ fraction of points in Y are within ℓ^∞-distance 2ℓ of a codimension-1 face of any hypercube.*

We refer the reader to [9] for the proof.

3 Proof of Theorem 4

We now prove Theorem 4, which covers the $d \geq k$ case of Theorem 2, our lower bound for subsystem codes. This also covers the $d \geq \sqrt{kn}$ case of Theorem 3, our generalization of [8] to D-dimensions.

▶ **Theorem** (Theorem 4, restated)**.** *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a nontrivial $[[n,k,d]]$ subsystem or commuting projector code with $d \geq c_1 n^{\frac{D-1}{D}}$ must have at least $c_0 d$ interactions of length at least $c_0 \frac{d}{n^{\frac{D-1}{D}}}$.*

As mentioned after the statement of Lemma 11, the Union Lemma for subsystem codes is substantially weaker than the corresponding result for commuting projector codes. Without the ability to conclude that the union of correctable sets remains correctable, we cannot directly generalize the techniques previously employed in the proofs of the generalized BPT bound, which required alternating applications of the expansion and union lemmas [2, 8].

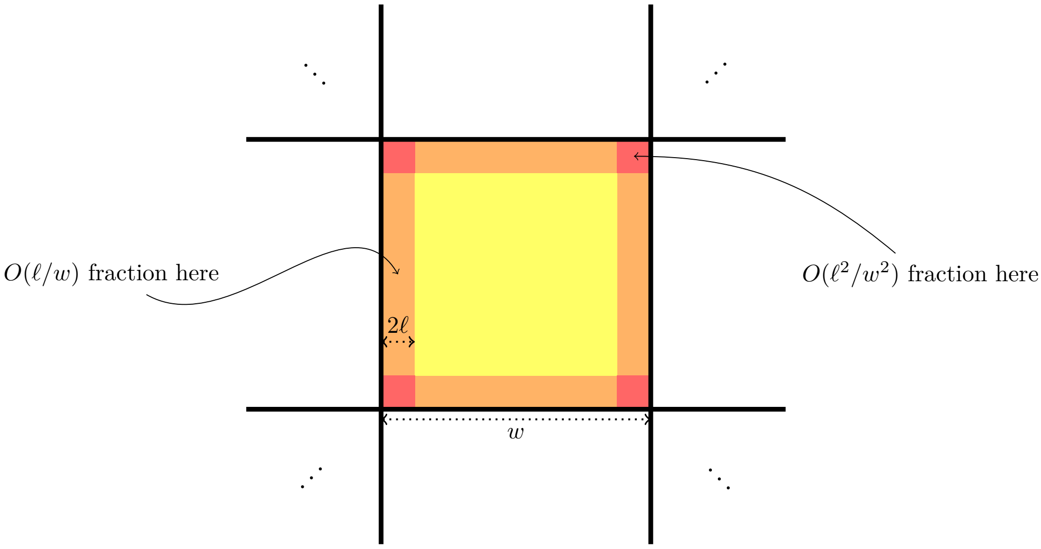

Figure 3 Tiling Lemma: for fixed sets of points X and Y and a random width-w grid tiling, we expect a $O(\ell^2/w^2)$ fraction of X to be within a $O(\ell)$ of a grid codimension-2 face, and a $O(\ell/w)$ fraction of Y to be within $O(\ell)$ of a codimension-1 face.

This poses a challenge for subsystem codes since we only obtain a dressed-cleanable set after the union lemma, and there is no straightforward way to continue with the expansion lemma, which requires correctable sets.

In view of these challenges, we take an alternative approach to the proof of Theorem 4 by repeatedly – and exclusively – applying the expansion lemma to a carefully crafted subset of qubits in order to grow our correctable region. To do this, we grow our correctable region by sweeping across our set of qubits $Q \subseteq \mathbb{R}^D$ one dimension at a time, changing directions and moving into a new dimension whenever the expansion process in unable to continue in the previous dimensions. The brunt of the proof is showing that this process never gets stuck, and that we are eventually able to grow our correctable set without obstruction to encompass the entire set of qubits Q. In this way, we are able to bypass the usage of the union lemma altogether.

We point out that the usual way of doing this expansion [7, 6, 8] – starting with a $d^{1/D} \times \cdots \times d^{1/D}$ box and repeatedly adding the boundary – does not work for general D-dimensional embeddings. For general embeddings, the qubits may not be constrained to an $O(n^{1/D}) \times \cdots \times O(n^{1/D})$ box, so, at some step, the uncontrolled expansion boundaries could contain more than d qubits, in which case we cannot apply the expansion lemma. For example, consider qubits embedded in 2-dimensions in a $n^{2/3} \times n^{2/3}$ square, with $\Omega(n)$ qubits distributed within $n^{1/3}$ of the box's boundary, and the remainder randomly distributed in the square. If we expand a rectangle from anywhere inside the square, applications of the expansion lemma fail when we approach the boundary. The easiest D-dimensional embeddings to realize are ones on a lattice structure contained in a $O(n^{1/D}) \times \cdots \times O(n^{1/D})$ box; our general statement shows that more creative embeddings like the one above cannot save on locality.

The general expansion process is quite involved in D-dimensions. We refer the reader to Appendix A for an extended exposition of the 2-dimensional case as an illustration of the basic idea of the proof, as well as the full proof in $D-$dimensions.

4 Proof of Theorem 5

We now prove Theorem 5, which covers the $k \geq d$ case of Theorem 2, our lower bound for subsystem codes.

▶ **Theorem** (Theorem 5, restated)**.** *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a $[[n,k,d]]$ subsystem code with $kd^{\frac{1}{D-1}} \geq c_1 n$ must have at least $c_0 k$ interactions of length at least $c_0(\frac{kd^{\frac{1}{D-1}}}{n})^{\frac{D-1}{D}}$.*

First, we state two lemmas that help us find large correctable sets in a D-dimensional embedding of a quantum code. The first is a generalization of the "holographic principle" for error correction in [5], which shows that the area, rather than the volume, governs the size of correctable sets. Recall that $f_{\geq \ell}(V)$ counts the number of interactions involving qubits in V with length at least ℓ (see Definition 8).

▶ **Lemma 16** (Holographic Principle)**.** *Suppose we have a $[[n,k,d]]$ quantum code (either commuting projector or subsystem) with an embedding $Q \subset \mathbb{R}^D$, and suppose $\ell \leq \frac{1}{8\sqrt{D}} d^{1/D}$. Let $V \subset Q$ be the subset of qubits contained in a box with sides of length at most*

$$w_0 \overset{\text{def}}{=} \left(\frac{\text{vol}(B_D)}{2 \cdot 4^{D+1} D} \cdot \frac{d}{\ell} \right)^{\frac{1}{D-1}}. \tag{7}$$

If $f_{\geq \ell}(V) \leq d/10$, then V is correctable.

We refer the reader to [9] for the formal proof.

If we divide $\mathbb{R}^D$ into cubes using Lemma 15, most cubes will be "good" in that they contain $\ll d$ long-ranged interactions (see footnote 5 of [8]). Lemma 16 then says that all the good cubes are correctable. How do we handle the cubes with large numbers of bad qubits? In [8], the solution was to further subdivide the bad cubes into sufficiently small rectangles, most of which then contains a sufficiently small number of bad qubits. The same strategy works in D-dimensions:

▶ **Lemma 17** (Subdivision)**.** *Let w, ℓ and d_1 be positive real numbers. Let $f : \mathbb{R}^D \to \mathbb{N}$ be a finitely supported function. Let R be a $h \times w^{D-1}$ box, with $h \geq 5\ell$ and $f(R) \geq d_1$. Then there exists a division of R by hyperplanes orthogonal to x_1 into boxes $R_1, \ldots, R_m$ such that:*
1. *Each box has dimensions $h_i \times w^{D-1}$, with $h_i \geq 5\ell$.*
2. *Each R_i satisfies either (i) $f(R_i) \leq d_1$ or (ii) has $h_i \leq 10\ell$.*
3. *The number of boxes m is at most $\frac{2f(R)}{d_1}$.*

We refer the reader to [9] for the formal proof.

Finally, we collect a few inequalities that we will frequently use in the proof of Theorems 5 and 6.

▶ **Lemma 18.** *Let*

$$w_0 = \left(\frac{\text{vol}(B_D)}{2 \cdot 4^{D+1} D} \cdot \frac{d}{\ell} \right)^{\frac{1}{D-1}}, \qquad \textit{and} \qquad c = \frac{\text{vol}(B_D)^{\frac{1}{D}}}{400\alpha D} \tag{8}$$

for some $\alpha \geq 1$. Suppose that ℓ satisfies $\ell \leq c d^{\frac{1}{D}}$. Then we have
1. $\dfrac{2^D}{\text{vol}(B_D)} (2w_0)^{D-1} \ell = \dfrac{d}{16D}$,

2. $w_0 \geq 100\alpha D\ell$,
3. $w_0 \geq \dfrac{1}{90\sqrt{D}}\left(\dfrac{d}{\ell}\right)^{\frac{1}{D-1}}$.

We are now ready to prove Theorem 5.

Proof of Theorem 5. We give a proof by contradiction, where we first assume that we have an embedding of a subsystem code with k logical qubits that has few long interactions, and then show that the code's dimension must actually be less than k. With hindsight, choose $c_0 = \frac{\mathrm{vol}(B_D)^{\frac{1}{D}}}{400D}$. Note that we have $c_0 \leq 1/(200D) \leq 1/400$ for $D \geq 2$. Choose $c_1 = (1/c_0)^{\frac{D}{D-1}}$. Suppose we have a $[[n,k,d]]$ subsystem code with a D-dimensional embedding $Q \subset \mathbb{R}^D$ satisfying $kd^{\frac{1}{D-1}} \geq c_1 n$. Let

$$\ell = c_0\left(\frac{kd^{\frac{1}{D-1}}}{n}\right)^{\frac{D-1}{D}}, \tag{9}$$

and note that we have $1 \leq \ell \leq c_0 d^{\frac{1}{D}} < \frac{1}{8\sqrt{D}}d^{\frac{1}{D}}$, where the first upper bound follows from $k \leq n$, and the lower bound from $kd^{\frac{1}{D-1}} \geq c_1 n$.

Now assume for the sake of contradiction that the embedding $Q \subset \mathbb{R}^D$ has at most $c_0 k$ interactions of length $\geq \ell$. Call an interaction *long* if its length is at least ℓ and *short* otherwise. Call a qubit $v \in Q$ *bad* if it participates in a long interaction and *good* otherwise. Then the function $f_{\geq\ell}(v)$ counts the number of long interactions that the qubit v participates in. By assumption, there are at most $c_0 k$ long interactions, so the total number of bad qubits is at most $\sum_{v\in Q} f_{\geq\ell}(v) \leq 2c_0 k$. Now we construct a division of $\mathbb{R}^D$ into $\mathcal{A} \sqcup \mathcal{B}$ that outlines the partition of the qubits $Q = A \sqcup B$. Let

$$w_0 = \left(\frac{\mathrm{vol}(B_D)}{2\cdot 4^{D+1}D}\cdot\frac{d}{\ell}\right)^{\frac{1}{D-1}} \tag{10}$$

as in Lemma 16. It follows from Lemma 18(2) that $w_0 \geq 100D\ell$. Apply Lemma 15 with $Y = Q$ (and with X arbitrary). This produces a tiling of $\mathbb{R}^D$ into cubes $\{S_m\}_{m\in\mathbb{Z}^D}$ of side length w_0, where at most $\frac{8D\ell}{w_0}n$ qubits of Q are within ℓ_∞ distance 2ℓ of some codimension-1 face of some cube. We call a cube S_m *good* if $f_{\geq\ell}(S_m) < d/10$ and *bad* otherwise. Now apply Lemma 17, with $d_1 = d/10$, to decompose each bad cube S_m into boxes $R_{m,1},\cdots,R_{m,n_m}$. All boxes obtained in this way will also be called *bad*. This process results in a division of $\mathbb{R}^D$ into good cubes and bad boxes. It follows from Lemma 17 (item 3) that the total number of bad boxes is no more than

$$\sum_{m:S_m\,\mathrm{bad}} \frac{2f_{\geq\ell}(S_m)}{d/10} \leq \sum_m \frac{2f_{\geq\ell}(S_m)}{d/10} \leq \frac{20}{d}\sum_m f_{\geq\ell}(S_m) \leq \frac{40}{d}c_0 k < \frac{k}{10d}. \tag{11}$$

Now we define the division $\mathcal{A} \sqcup \mathcal{B}$ as follows:

- $\mathcal{B}$ is the set of all points within ℓ_∞ distance 2ℓ of some codimension-1 face of either a good cube S_m or a bad box $R_{m,i}$.
- $\mathcal{A}$ is the set of points not in $\mathcal{B}$.

Note that we can perturb the tiling slightly in order to ensure that no qubits lie on the boundary of any subregion of $\mathcal{A}$ or $\mathcal{B}$.

Having constructed the division, we will now construct a corresponding partition of qubits $Q = A \sqcup B$ such that A is dressed-cleanable and $|B| < k$. This will give us our desired contradiction from Lemma 12. We define the partition $Q = A \sqcup B$ as follows:

- A is the set of all good qubits in region $\mathcal{A}$.
- B is the set of all remaining qubits. These are either good qubits in region $\mathcal{B}$ or bad qubits.

It remains to check that A is dressed-cleanable and that $|B| < k$. We refer the reader to [9] for the formal proof. ◀

5 Proof of Theorem 6

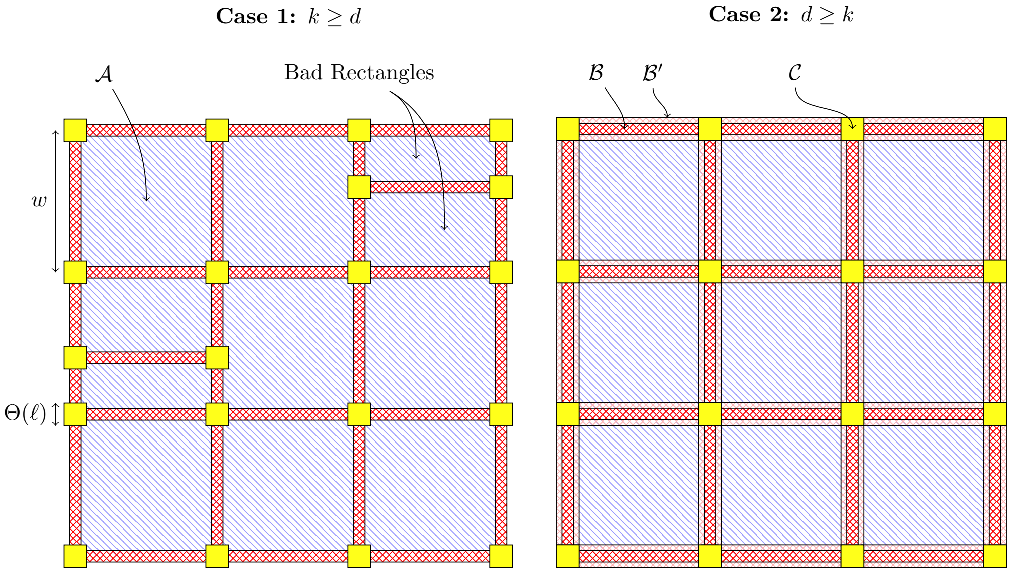

Figure 4 The division of the plane into regions $\mathcal{A}$ (lined blue), $\mathcal{B}$ (red and pink crosshatch), and $\mathcal{C}$ (solid yellow) for the proof of Theorem 6. The region $\mathcal{B}'$ (pink crosshatch) is also indicated in the figure on the right. These regions inform our qubit division $Q = A \sqcup B \sqcup C$. We use this division in different ways for the cases $k \geq d$ and $d \geq k$. When $k \geq d$ (left), we ignore $\mathcal{B}'$, and also subdivide any bad squares into bad rectangles. When $d \geq k$ (right), there are no bad squares, but we need to explicitly consider the region $\mathcal{B}'$.

We now prove Theorem 6, which, together with Theorem 4, yields Theorem 3, our generalization of the main result of [8] to case of D-dimensional embeddings.

▶ **Theorem** (Theorem 6, restated). *For all $D \geq 2$, there exist constants $c_0 = c_0(D) > 0$ and $c_1 = c_1(D) > 0$ such that the following is true: Any D-dimensional embedding of a $[[n,k,d]]$ commuting projector code with $kd^{\frac{2}{D-1}} \geq c_1 n$ must have at least $c_0 k$ interactions of length $c_0(\frac{kd^{\frac{2}{D-1}}}{n})^{\frac{D-1}{2D}}$.*

Proof. For $d \geq \sqrt{kn}$, the result follows from Theorem 4, so it suffices to consider the case where $d \leq \sqrt{kn}$. With hindsight, choose $c_0 = \frac{\mathrm{vol}(B_D)^{\frac{1}{D}}}{800D^2}$, and let $c_1 = (1/c_0)^{\frac{2D}{D-1}}$. Note that $c_0 \leq 1/(400D^2) \leq 1/400$. Suppose we have a $[[n,k,d]]$ commuting projector code with a D-dimensional embedding $Q \subset \mathbb{R}^D$ satisfying $kd^{\frac{2}{D-1}} \geq c_1 n$. Let

$$\ell = c_0 \left(\frac{kd^{\frac{2}{D-1}}}{n} \right)^{\frac{D-1}{2D}}. \tag{12}$$

Note that we have $1 \le \ell \le c_0 d^{\frac{1}{D}} \le \frac{1}{8\sqrt{D}} d^{\frac{1}{D}}$, where the lower bound follows from $kd^{\frac{2}{D-1}} \ge c_1 n$ and the first upper bound from the $k \le n$.

Now assume for the sake of contradiction that the embedding $Q \subset \mathbb{R}^D$ has at most $c_0 \max(k,d)$ interactions of length $\ge \ell$. Call an interaction *long* if its length is at least ℓ and *short* otherwise. Call a qubit $v \in Q$ *bad* if it participates in a long interaction and *good* otherwise. The function $f_{\ge \ell}(v)$ counts the number of long interactions that the qubit v participates in. By assumption, the total number of long interactions is at most $c_0 \max(k,d)$, so the total number of bad qubits is at most $\sum_{v \in Q} f_{\ge \ell}(v) \le 2c_0 \max(k,d)$. We will construct a division of $\mathbb{R}^D$ into subsets $\mathcal{A} \sqcup \mathcal{B} \sqcup \mathcal{C}$ that will inform the partition of the qubits $Q = A \sqcup B \sqcup C$. Let

$$w_0 = \left(\frac{\mathrm{vol}(B_D)}{2 \cdot 4^{D+1} D} \frac{d}{\ell} \right)^{\frac{1}{D-1}}, \tag{13}$$

as in Lemma 16. Note that it follows from Lemma 18 that $w_0 \ge 200 D^2 \ell$ and $w_0 \ge \frac{1}{90\sqrt{D}} (d/\ell)^{\frac{1}{D-1}}$.

Apply Lemma 15 with $X = Q$ and with Y as the multiset where each qubit v appears with multiplicity $f_{\ge \ell}(v)$. This gives a partition of $\mathbb{R}^D$ into a set of cubes $\{S_m\}_{m \in \mathbb{Z}^D}$ of side length w_0. By construction, at most $\frac{16 D^2 \ell^2}{w_0^2} n$ qubits of Q are within ℓ_∞ distance 2ℓ of a codimension-2 face of some cube, and at most $\frac{8D\ell}{w_0} \cdot 2c_0 \max(k,d)$ bad interactions involve a qubit within ℓ_∞ distance 2ℓ of a codimension-1 face of some cube. We call a cube S_m *good* if $f_{\ge \ell}(S_m) < d/10$ and *bad* otherwise. Now apply Lemma 17 to decompose each bad cube into boxes $R_{m,1}, \cdots, R_{m,n_m}$. All boxes obtained by subdividing a bad cube will also be called bad. This process results in a division of $\mathbb{R}^D$ into good cubes and bad boxes. By Lemma 17 (item 3), the total number of bad boxes is no more than

$$\sum_{m : S_m \,\mathrm{bad}} \frac{2 f_{\ge \ell}(S_m)}{d/10} \le \sum_m \frac{2 f_{\ge \ell}(S_m)}{d/10} \le \frac{20}{d} \sum_m f_{\ge \ell}(S_m) \le \frac{40}{d} c_0 \max(k,d) < \frac{\max(k,d)}{10d}. \tag{14}$$

Now we define the division $\mathcal{A} \sqcup \mathcal{B} \sqcup \mathcal{C}$ as follows:

- $\mathcal{C}$ is the set of all points within ℓ_∞ distance 2ℓ of some codimension-2 face of a good cube S_m or a bad box $R_{m,i}$.
- $\mathcal{B}$ is the set of all points *not* already in $\mathcal{C}$ and within ℓ_∞ distance ℓ of some codimension-1 face of a good cube S_m or bad box $R_{m,i}$.
- $\mathcal{B}' \subset \mathcal{B}$ is the set of all points *not* already in $\mathcal{C}$ and within ℓ_∞ distance 2ℓ of some codimension-1 face of a good cube S_m.
- $\mathcal{A}$ is the set of points not in $\mathcal{B}$ or $\mathcal{C}$.

Note that we can perturb the tiling slightly in order to ensure that no qubits lie on the boundary of any subregion of $\mathcal{A}$, $\mathcal{B}$, $\mathcal{B}'$, or $\mathcal{C}$.

Having defined the division of $\mathbb{R}^D$, we will now construct our partition of qubits $Q = A \sqcup B \sqcup C$. A sketch of the high-level ideas are as follows: We aim to define our qubit partition with the goal of having A, B be correctable, and $|C| < k$. This will lead to the desired contradiction using Lemma 13. There are two cases to consider, depending on whether $k \ge d$ or $k \le d$. When $k \ge d$, we have $2c_0 k \ll k$ bad qubits, which can then be directly placed in C without affecting the requirement that $|C| < k$. When $k \le d$, we have $2c_0 d \ll d$ bad qubits, and the set of all bad qubits is itself correctable. Our chosen division of $\mathbb{R}^D$ implies that very few bad qubits can interact with the qubits in $\mathcal{B}$, so the union lemma suggests that we can add almost all of the bad qubits to $\mathcal{B}$ while preserving its correctability. We now continue with our proof, divided into the two cases $k \ge d$ and $k \le d$.

Case 1: $k \geq d$. We define the partition of qubits $Q = A \sqcup B \sqcup C$ as follows:

- C is the set of qubits in region $\mathcal{C}$, along with all bad qubits.
- B is the set of all good qubits in region $\mathcal{B}$
- A is the set of all good qubits in region $\mathcal{A}$.

It's clear that this is indeed a partition of Q. It remains to show that A, B are correctable, and that $|C| < k$. This gives us a contradiction with the fact that A and B are correctable through Lemma 13. We refer the reader to [9] for the formal proof.

Case 2: $d \geq k$. From equation (13), the total number of bad boxes is at most $\max(k, d)/(10d) = 1/10$, which is less than 1. It follows that there are no bad boxes in this case, only good cubes. We define the partition of qubits $Q = A \sqcup B \sqcup C$ in this case as follows:

- C consists of the set of qubits in region $\mathcal{C}$, together with all qubits participating in a long interaction with a qubit in $\mathcal{B}'$ (including the bad qubits in $\mathcal{B}'$).
- B is the set of good qubits in region $\mathcal{B}$, together with the bad qubits not in C.
- A is the set of good qubits in region $\mathcal{A}$.

It's clear that this is a partition of the qubits in Q. It remains to check that A, B are correctable and $|C| < k$. This gives us our desired contradiction using Lemma 13. We refer the reader to [9] for the formal proof.

We have obtained a contradiction in both the $d \leq k$ and $d \geq k$ cases, and this completes the proof of the theorem. ◀

6 Construction for Upper Bounds

The bounds derived in Theorems 2 and 3 are tight in both the interaction count M^* and the interaction length ℓ^*. Tightness is shown by constructing explicit examples of embedded codes which saturate the interaction count or length. For the interaction count, it suffices to consider an asymptotically good quantum low-density parity-check (qLDPC) code [14, 12], which has $O(M^*) = O(\max(k, d))$ interactions of any length. Since a stabilizer code can also be regarded as a subsystem code with zero gauge qubits, this shows that both Theorem 2 and 3 are tight in terms of interaction count. This is covered by Theorem 1.3 of [8].

We now describe constructions that show the interaction length is optimal in Theorem 2 and Theorem 3. In both cases, we construct a code that saturates the bound for interaction length by concatenating an asymptotically good qLDPC code with a geometrically local code which saturates the Bravyi and BPT bounds, respectively.

6.1 Subsystem codes

We start by showing the interaction length for subsystem codes (Theorem 2) is optimal. We will define a concatenated subsystem code composed of an asymptotically good qLDPC code, together with a subsystem code which is geometrically local in D-dimensions and saturates the Bravyi bound. For the local subsystem code, we employ the "wire code" construction of Baspin and Williamson [4].

▶ **Theorem 19** (Wire code [4]). *For all $D \geq 2$, there exists an $\varepsilon > 0$ such that, for all positive integers n there exists a subsystsem code with parameters $[[n, \geq \varepsilon n^{\frac{D-1}{D}}, \geq \varepsilon n^{\frac{D-1}{D}}]]$ that has a set of gauge generators that are $O(1)$-local in a D-dimensional embedding.*

The concatenation procedure for subsystem codes is formally identical to the process for stabilizer codes. Namely, if $\mathcal{C}_1 = [[n_1, k_1]]$ and $\mathcal{C}_2 = [[n_2, k_2]]$ are subsystem codes, then their concatenation $\mathcal{C}_2 \circ \mathcal{C}_1$ is defined using n_2 blocks of the inner code $\mathcal{S}_1$ and k_1 copies of the outer code $\mathcal{C}_2$. Let q_{ij} be the ith logical qubit of the jth $\mathcal{S}_1$ block. Then the concatenated code is defined by replacing the jth physical qubit of the ith $\mathcal{C}_2$ block with q_{ij}.

▶ **Lemma 20** (Concatenated Subsystem Codes). *Let $\mathcal{C}_1 = [[n_1, k_1, d_1, g_1]]$ and $\mathcal{C}_2 = [[n_2, k_2, d_2, g_2]]$ be two subsystem codes. Then there exists a subsystem code $\mathcal{C} = \mathcal{C}_2 \circ \mathcal{C}_1 = [[n_1 n_2, k_1 k_2, d \geq d_1 d_2, k_1 g_2 + n_2 g_1]]$, called the concatenation of $\mathcal{C}_2$ and $\mathcal{C}_1$*

Theorems 19 and 20 together give the desired construction. We state the result below and refer the reader to [9] for the formal proof.

▶ **Theorem 21** (Optimality of Interaction Length for Subsystem Codes). *For all $D \geq 2$, there exists a constant $c_1 = c_1(D) > 0$ such that the following holds: for all $n, k, d > 0$ with $k, d \leq n$ satisfying $kd^{\frac{1}{D-1}} \geq c_1 n$ or $d \geq c_1 n^{\frac{D-1}{D}}$, there exists an $[[n, \geq k, \geq d]]$ subsystem code with a D-dimensional embedding containing no interactions of length at least*

$$\ell = \max\left(\frac{d}{n^{\frac{D-1}{D}}}, \left(\frac{kd^{\frac{1}{D-1}}}{n}\right)^{\frac{D-1}{D}}\right). \tag{15}$$

6.2 Commuting Projector Codes

We now show the interaction length in Theorem 3 is optimal. The construction is very similar to the one used in Theorem 1.3 of [8], except generalized to D-dimensions. In 2D, the surface code offers a simple and natural candidate for a geometrically local code that saturates the BPT bound. In higher dimensions, we instead use the family of "subdivided codes" constructed by Lin, Wills and Hsieh [13].

▶ **Theorem 22** (Subdivided code [13]). *For all $D \geq 2$, there exists an $\varepsilon > 0$ such that, for all positive integers n there exists a stabilizer code with parameters $[[n, \geq \varepsilon n^{\frac{D-2}{D}}, \geq \varepsilon n^{\frac{D-1}{D}}]]$ that has a set of stabilizer generators that are $O(1)$-local in a D-dimensional embedding.*

The optimality of the interaction length follows by concatenating a good qLDPC code with the subdivided code. We refer the reader to [9] for the formal proof.

▶ **Theorem 23** (Optimality of Interaction Length for Stabilizer Codes). *For all $D \geq 2$, there exists a constant $c_1 = c_1(D) > 0$ such that the following holds: for all $n, k, d > 0$ with $k, d \leq n$ satisfying either $kd^{\frac{2}{D-1}} \geq c_1 \cdot n$ or $d \geq c_1 \cdot n^{\frac{D-1}{D}}$, there exists a $[[n, \geq k, \geq d]]$ quantum stabilizer code with a D-dimensional embedding containing no interactions of length at least*

$$\ell = \max\left(\frac{d}{n^{\frac{D-1}{D}}}, \left(\frac{kd^{\frac{2}{D-1}}}{n}\right)^{\frac{D-1}{2D}}\right). \tag{16}$$

References

1 Nouédyn Baspin, Venkatesan Guruswami, Anirudh Krishna, and Ray Li. Improved rate-distance trade-offs for quantum codes with restricted connectivity. *Quantum Science and Technology*, 10(1):015021, 2024.

2 Nouédyn Baspin and Anirudh Krishna. Connectivity constrains quantum codes. *Quantum*, 6:711, 2022.

3 Nouédyn Baspin and Anirudh Krishna. Quantifying nonlocality: How outperforming local quantum codes is expensive. *Physical Review Letters*, 129(5):050505, 2022.

4 Nouédyn Baspin and Dominic Williamson. Wire codes. *arXiv preprint arXiv:2410.10194*, 2024.

5 Sergey Bravyi. Subsystem codes with spatially local generators. *Physical Review A*, 83(1), January 2011. doi:10.1103/physreva.83.012320.

6 Sergey Bravyi, David Poulin, and Barbara Terhal. Tradeoffs for reliable quantum information storage in 2D systems. *Physical Review Letters*, 104(5):050503, 2010. doi:10.1103/PhysRevLett.104.050503.

7 Sergey Bravyi and Barbara Terhal. A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes. *New Journal of Physics*, 11(4):043029, 2009. doi:10.1088/1367-2630/11/4/043029.

8 Samuel Dai and Ray Li. Locality vs quantum codes. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 677–688, 2025.

9 Samuel Dai, Ray Li, and Eugene Tang. Optimal locality and parameter tradeoffs for subsystem codes. *arXiv preprint arXiv:2503.22651*, 2025.

10 Jeongwan Haah and John Preskill. Logical-operator tradeoff for local quantum codes. *Physical Review A*, 86(3), September 2012. doi:10.1103/physreva.86.032308.

11 Yifan Hong, Matteo Marinelli, Adam M Kaufman, and Andrew Lucas. Long-range-enhanced surface codes. *Physical Review A*, 110(2):022607, 2024.

12 Anthony Leverrier and Gilles Zémor. Quantum tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE, 2022.

13 Xingjian Li, Ting-Chun Lin, and Min-Hsiu Hsieh. Transform arbitrary good quantum ldpc codes into good geometrically local codes in any dimension. *arXiv preprint arXiv:2408.01769*, 2024.

14 Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 375–388, 2022.

15 Fernando Pastawski and Beni Yoshida. Fault-tolerant logical gates in quantum error-correcting codes. *Physical Review A*, 91(1), January 2015. doi:10.1103/physreva.91.012305.

A Proof of Theorem 4

In this appendix, we first give a detailed sketch of the proof of Theorem 4 for the case of 2-dimensional embeddings. This is meant as a simplified illustration of the proof idea in the general D-dimensional case, which may be difficult to visualize. We then give the full proof in D-dimensions.

A.1 Detailed Sketch of Theorem 4 in 2-dimensions

Let Q be a 2-dimensional embedding of a subsystem code $\mathcal{C} = [[n,k,d]]$ satisfying $d \geq 32\sqrt{n}$. Without loss of generality, we may assume that all of our qubits are contained in the interior of a large square $[\ell, A-\ell] \times [\ell, A-\ell]$, for some integer $A \in \mathbb{N}$. We remove a border of width ℓ so that we do not have to worry about edge cases later in our proof. Suppose for the sake of contradiction that there exists at most $d/4$ interactions of length at least $\ell = \frac{d}{32\sqrt{n}}$. Note that our choice of code parameters ensures that we have $1 \leq \ell \leq \sqrt{n}/32$. We will call any qubit participating in a length $\geq \ell$ interaction a bad qubit. Let B be the set of all bad qubits. Then by assumption, $|B| \leq d/2$. Given any region $R \subset \mathbb{R}^2$, we will say that R is correctable if and only if $Q \cap R$ is correctable.

The basic idea of the proof is as follows. We wish to show that, given our assumption on the number of long interactions, a small correctable region can be iteratively grown without bound using the expansion lemma. Eventually the correctable region will encompass the entire set of qubits Q, which is a contradiction with our initial assumption that the code is non-trivial. We do this by first expanding along the x_1-direction, and then switching to the

x_2-direction whenever we are unable to continue in the x_1-direction. The details for the two cases are given below.

A.1.1 Expansion in the x_1-direction

Let us imagine starting with a region of the form of a vertical strip $V[a_1] = [0, a_1] \times \mathbb{R}$, for some $a_1 > 0$. If $a_1 < \ell$, then $V[a_1]$ contains no qubits by assumption, so it is vacuously correctable. Now, given an existing correctable region of the form $V[a_1]$, we wish to apply the expansion lemma to obtain a larger correctable set of the form $V[a_1 + \ell]$. This will be possible provided that the number of qubits in the boundary of $V[a_1]$ is sufficiently small so that the boundary set itself is correctable. We will formalize this requirement by saying that a number $a_1 \in \mathbb{R}$ is "good" if the density of qubits around the line $x = a_1$ is low. Formally, $a_1 \in \mathbb{R}$ is a *good* x_1-coordinate if the set $[a_i - \ell, a_1 + \ell] \times \mathbb{R}$ contains at most $\ell\sqrt{n}$ qubits, and *bad* otherwise.

The boundary of $V[a_1]$ consists of all qubits within a distance ℓ of the line $x = a_1$, together with a subset of the bad qubits. Therefore the boundary is a subset of $B \cup [a_1 - \ell, a_1 + \ell] \times \mathbb{R}$. If a_1 is good, then by definition this subset contains at most

$$d/2 + \ell\sqrt{n} = d/2 + d/32 < d \tag{17}$$

qubits, and is hence correctable. It follows by expansion and subset closure that if $V[a_1]$ is correctable and a_1 is good, then $V[a_1 + \ell] \subset V[a_1] \cup \partial V[a_1]$ is also correctable. This gives us an easy way to grow the sets $V[a_1]$. However, this process cannot continue indefinitely, since there is no guarantee at each step that the new coordinate $a_1 + \ell$ is good. When $a_1 + \ell$ is a bad coordinate, our expansion process gets stuck. To get unstuck, we will fill in the stretch around the bad coordinate $a_1 + \ell$ by expanding in the x_2-direction. After this gap containing bad x_1-coordinates has been filled, we can continue expanding in the x_1-direction until we reach our next obstacle.

A.1.2 Stuck in the x_1-direction, expand along the x_2-direction

Now we formalize the process of expanding in the x_2 direction. Given a bad coordinate a_1, we will define $\mathsf{Next}(a_1)$ to be the "next available" good coordinate. More precisely, we define

$$\mathsf{Next}(a_1) = \inf G_{>a_1} + \gamma, \tag{18}$$

where $G_{>a_1}$ is the set of all good coordinates larger than a_1, and $\gamma > 0$ is a sufficiently small value so that we actually have $\mathsf{Next}(a_1) \in G_{>a_1}$.[2] Given a set $V[a_1]$ where a_1 is good but $a_1 + \ell$ is bad, let us define $V[a_1, a_2]$ to be

$$V[a_1, a_2] = V[a_1] \cup ([a_1, \mathsf{Next}(a_1 + \ell)] \times [0, a_2]) = ([0, a_1] \times \mathbb{R}) \cup ([a_1, \mathsf{Next}(a_1 + \ell)] \times [0, a_2]). \tag{19}$$

We will call sets $V[a_1, a_2]$ defined this way *legal* sets. Our goal is to show that given a correctable legal set $V[a_1, a_2]$, we can always apply expansion in the x_2-direction to obtain a new correctable legal set $V[a_1, a_2 + \ell]$.

[2] The small constant γ is necessary since the set of good coordinates is an open set, and as such does not contain its own infimum. A bit of thought reveals that it is sufficient to take γ smaller than the minimum element of $\bigcup_{q \neq q'} \{|q_1 - q'_1|, ||q_1 - q'_1| - 2\ell|\}$, where q_1, q'_1 denote the x_1-coordinates of qubits $q, q' \in Q$.

The key observation now is that the additional block $[a_1, \mathsf{Next}(a_1+\ell)] \times [0, a_2]$ in $V[a_1, a_2]$ must be thin in the x_1-direction. Indeed, the entire interval $[a_1+\ell, \mathsf{Next}(a_1+\ell)-\gamma]$ consists of bad coordinates. Any strip of width 2ℓ centered around a bad coordinate contains at least $\ell\sqrt{n}$ qubits, and since we only have a total of n qubits, the set $[a_1+\ell, \mathsf{Next}(a_1+\ell)-\gamma] \times \mathbb{R}$ can fit at most $\lfloor \sqrt{n}/\ell \rfloor$ such strips inside. It follows that we have

$$\mathsf{Next}(a_1+\ell) - \gamma - a_1 - \ell < 2\ell \left\lfloor \frac{\sqrt{n}}{\ell} \right\rfloor + 2\ell, \tag{20}$$

which implies

$$\mathsf{Next}(a_1+\ell) - a_1 < 2\sqrt{n} + 4\ell. \tag{21}$$

Now, consider the boundary of a correctable legal set $V[a_1, a_2]$. This will be a subset of

$$\begin{aligned} B \cup & \underbrace{[a_1-\ell, a_1+\ell] \times \mathbb{R}}_{S_1} \\ & \cup \underbrace{[\mathsf{Next}(a_1+\ell)-\ell, \mathsf{Next}(a_1+\ell)+\ell] \times \mathbb{R}}_{S_2} \\ & \cup \underbrace{[a_1, \mathsf{Next}(a_1+\ell)] \times [a_2-\ell, a_2+\ell]}_{S_3}. \end{aligned} \tag{22}$$

Since a_1 and $\mathsf{Next}(a_1+\ell)$ are good coordinates, we have

$$|Q \cap S_1| + |Q \cap S_2| \le 2\ell\sqrt{n} \tag{23}$$

by assumption. By Lemma 14, the number of qubits in the subset S_3 is bounded above by

$$|Q \cap S_3| \le \frac{4}{\pi}(2\ell+1)(2\sqrt{n}+4\ell+1) \le 9\ell\sqrt{n}. \tag{24}$$

Therefore the boundary of $V[a_1, a_2]$ contains at most

$$|B| + |Q \cap S_1| + |Q \cap S_2| + |Q \cap S_3| \le \frac{d}{2} + 11\ell\sqrt{n} = \frac{27}{32}d < d \tag{25}$$

qubits, and is hence correctable. It follows by expansion and subset closure that if $V[a_1, a_2]$ is a correctable legal set, then $V[a_1, a_2+\ell]$ is again a correctable legal set. We can therefore continue expanding in the x_2-direction until we reach $V[a_1, A] = V[\mathsf{Next}(a_1+\ell)]$. This is a vertical strip with a good boundary coordinate $\mathsf{Next}(a_1+\ell)$. We can therefore return to the previous case and proceed to expand along the x_1-direction again. This process can now continue indefinitely, alternating between the x_1 and x_2-directions whenever we get stuck again. In this way, starting from an initial correctable set, say the vacuously correctable region $V[\ell/2]$, we can iteratively grow our correctable region without bound to encompass the entire set of qubits Q. Thus we arrive at our desired contradiction.

A.2 Proof of Theorem 4 in D-dimensions

For notational convenience, we grow the correctable region from low coordinates to high coordinates for this proof. For a set S, we write $S^{\le D} = S \cup S^2 \cup \cdots \cup S^D$.

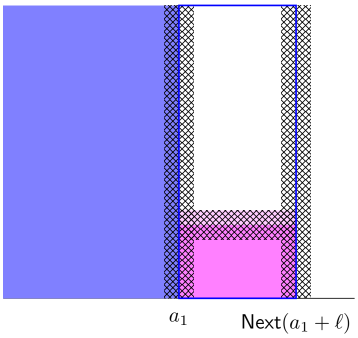

Figure 5 Sketch of the expansion process in 2 dimensions. The blue region is $V[a_1]$ and the pink region is $[a_1, \mathsf{Next}(a_1+\ell)] \times [0, a_2]$. The crosshatched region (labeled F in the proof of Theorem 4) contains the boundary of $V[a_1, a_2]$.

Proof. With hindsight, set $c_1 = 2(6^D D/\operatorname{vol}(B_D))^{D/(D-1)}$. Suppose n, k, d are code parameters with $d \geq c_1 n^{\frac{D-1}{D}}$. With hindsight, choose

$$\ell = \frac{\operatorname{vol}(B_D)}{6^D D n^{(D-1)/D}} d, \tag{26}$$

so that $1 < \ell < \frac{n^{1/D}}{4^D D}$. Let Q be a D-dimensional embedding of a $[[n, k, d]]$ subsystem code $\mathcal{C}$ with $d \geq k$, and suppose there are at most $d/4$ interactions of length at least ℓ.

Let B denote the set of qubits participating in long range interactions, so that $|B| \leq d/2$. Without loss of generality, we may assume all the qubits are in the box $[\ell, A-\ell]^D$ for some large integer A. Choose a sufficiently small constant $\gamma > 0$. With hindsight, γ less than the minimum nonzero element of $\bigcup_{i\in[D], q\neq q'} \{|q_i - q'_i|, ||q_i - q'_i| - 2\ell|\}$ suffices.

For $i = 1, \ldots, D-1$, call real number $x_i \in \mathbb{R}$ *i-good* if $\mathbb{R}^{i-1} \times [x_i - \ell, x_i + \ell] \times \mathbb{R}^{D-i}$ has at most $\ell n^{(D-1)/D}$ qubits and *i-bad* otherwise. By convention, we will consider every real number to be i-good when $i = D$. An i-bad number represents an i-th-coordinate-value where we would "get stuck" and need to start expanding in a different direction.

For $i = 1, \ldots, D-1$ and an i-bad real number x, let $\mathsf{Next}_i(x)$ be the maximum value such that all values in the interval $[x, \mathsf{Next}_i(x) - \gamma]$ are i-bad ($\mathsf{Next}_i(x)$ is undefined if x is i-good). Note, as γ is sufficiently small, that $\mathsf{Next}_i(x)$ is always i-good when it is defined.

For $a_1, \ldots, a_i \in \mathbb{R}$, with $i \leq D$, let

$$\begin{aligned} V[a_1, \ldots, a_i] = {} & [0, a_1] \times \mathbb{R}^{D-1} \\ & \cup [a_1, \mathsf{Next}_1(a_1+\ell)] \times [0, a_2] \times \mathbb{R}^{D-2} \\ & \cup [a_1, \mathsf{Next}_1(a_1+\ell)] \times [a_2, \mathsf{Next}_2(a_2+\ell)] \times [0, a_3] \times \mathbb{R}^{D-3} \\ & \quad \vdots \qquad\qquad \vdots \qquad\qquad \vdots \\ & \cup [a_1, \mathsf{Next}_1(a_1+\ell)] \times \cdots \times [a_{i-1}, \mathsf{Next}_{i-1}(a_{i-1}+\ell)] \times [0, a_i] \times \mathbb{R}^{D-i}. \end{aligned} \tag{27}$$

Call such a set $V[a_1, a_2, \ldots, a_i]$ *legal* if (i) a_j is j-good for all $j \in \{1, \ldots, i\}$ and (ii) $a_j + \ell$ is j-bad for $j \in \{1, \ldots, i-1\}$. We grow a large correctable set of the above form using the following four properties:

1. (If possible, expand in ith dimension): *If $i \leq D$, the set $V[a_1, \ldots, a_i]$ is correctable and legal, and $V[a_1, \ldots, a_{i-1}, a_i + \ell]$ is legal, then $V[a_1, \ldots, a_{i-1}, a_i + \ell]$ is correctable.*

We apply the expansion lemma. Note that the boundary is a subset of

$$\begin{aligned} F \overset{\text{def}}{=} B \\ &\cup \left([a_1-\ell, a_1+\ell] \times \mathbb{R}^{D-1}\right) \\ &\cup \left([\mathsf{Next}_1(a_1+\ell)-\ell, \mathsf{Next}_1(a_1+\ell)+\ell] \times \mathbb{R}^{D-1}\right) \\ &\qquad\vdots \qquad\qquad \vdots \qquad\qquad \vdots \\ &\cup \left(\mathbb{R}^{i-2} \times [a_{i-1}-\ell, a_{i-1}+\ell] \times \mathbb{R}^{D-i+1}\right) \\ &\cup \left(\mathbb{R}^{i-2} \times [\mathsf{Next}_{i-1}(a_{i-1}+\ell)-\ell, \mathsf{Next}_{i-1}(a_{i-1}+\ell)+\ell] \times \mathbb{R}^{D-i+1}\right) \\ &\cup \left(\mathbb{R}^{i-1} \times [a_i-\ell, a_i+\ell] \times \mathbb{R}^{D-i}\right). \end{aligned} \tag{28}$$

Since $V[a_1,\dots,a_i]$ is legal, the values a_j and $\mathsf{Next}(a_j+\ell)$ are j-good for $j \in \{1,\dots,i\}$ and $j \in \{1,\dots,i-1\}$, respectively. By definition of i-good, each set in the above union, other than B, has at most $\ell n^{(D-1)/D}$ qubits. Thus, F has at most $d/2 + (2D-1)\ell n^{1/D} < d$ qubits, so F is correctable. It follows that $V[a_1,\dots,a_i] \cup F$ is correctable, and by Subset Closure, $V[a_1,\dots,a_{i-1},a_i+\ell]$ is correctable.

2. (Stuck in i-th dimension, start in $(i+1)$-th-dimension): *If $V[a_1,a_2,\dots,a_i]$ is correctable and legal, and $V[a_1,\dots,a_{i-1},a_i+\ell]$ is not legal, then $V[a_1,\dots,a_i,0]$ is correctable and legal.*
The set $V[a_1,\dots,a_i,0]$ is correctable by definition as $V[a_1,\dots,a_i,0] = V[a_1,\dots,a_i]$. It is legal also by definition, as we assume a_i is i-good but $a_i+\ell$ is i-bad, and 0 is trivially $(i+1)$-good.

3. (Expand in D-th dimension): *If $V[a_1,a_2,\dots,a_D]$ is correctable and legal, then $V[a_1,\dots,a_{D-1},a_D+\ell]$ is correctable and legal.*
It is legal as every real number is D-good by definition. For correctable, we again use the expansion lemma. Following part 1, the boundary is a subset of

$$\begin{aligned} F \overset{\text{def}}{=} B \\ &\cup \left([a_1-\ell, a_1+\ell] \times \mathbb{R}^{D-1}\right) \\ &\cup \left([\mathsf{Next}_1(a_1+\ell)-\ell, \mathsf{Next}_1(a_1+\ell)+\ell] \times \mathbb{R}^{D-1}\right) \\ &\qquad\qquad\qquad \vdots \qquad\qquad \vdots \qquad\qquad \vdots \\ &\cup \left(\mathbb{R}^{D-2} \times [a_{D-1}-\ell, a_{D-1}+\ell] \times \mathbb{R}\right) \\ &\cup \left(\mathbb{R}^{D-2} \times [\mathsf{Next}_{D-1}(a_{D-1}+\ell)-\ell, \mathsf{Next}_{D-1}(a_{D-1}+\ell)+\ell] \times \mathbb{R}\right) \\ &\cup [a_1, \mathsf{Next}_1(a_1+\ell)] \times \cdots \times [a_{D-1}, \mathsf{Next}_{D-1}(a_{D-1}+\ell)] \times [a_D-\ell, a_D+\ell]. \end{aligned} \tag{29}$$

As in part 1, we have $|B| \le d/2$, and all but the last set in the union above has size at most $\ell n^{(D-1)/D}$. We now bound the size of the last set in F. Since $[a_j+\ell, \mathsf{Next}_j(a_j+\ell)-\gamma]$ is j-bad for all j, a counting argument yields

$$\mathsf{Next}_j(a_j+\ell) - (a_j+\ell) \le 2n^{1/D} + 2\ell. \tag{30}$$

To see this, pack strips of width 2ℓ in the jth dimension into $\mathbb{R}^{j-1} \times [a_j+\ell, \mathsf{Next}_j(a_j+\ell)] \times \mathbb{R}^{D-j}$. Each strip has at least $\ell n^{1/D}$ qubits by definition of being j-bad. There are at most n qubits, so there are at most $n/(\ell n^{(D-1)/D})$ packed strips, so the total width satisfies

$$\mathsf{Next}_j(a_j+\ell) - (a_j+\ell) \le \frac{2\ell n}{\ell n^{(D-1)/D}} + 2\ell = 2n^{1/D} + 2\ell. \tag{31}$$

We conclude $\mathsf{Next}_j(a_j+\ell)-a_j \le 2n^{1/D}+3\ell$. Hence, the last box has at most

$$(2n^{1/D}+3\ell+1)^{D-1}(2\ell+1) < \frac{6^D}{\mathrm{vol}(B_D)}\ell n^{(D-1)/D} \tag{32}$$

qubits inside by Lemma 14. The total boundary thus has at most

$$\frac{d}{2}+(2D-2)\ell n^{(D-1)/D}+\frac{6^D}{\mathrm{vol}(B_D)}\ell n^{(D-1)/D} < d \tag{33}$$

qubits. It follows that F is correctable, so $V[a_1,\dots,a_D]\cup F$ is correctable, and by Subset Closure, $V[a_1,\dots,a_{D-1},a_D+\ell]$ is correctable.

4. (Finish $(i+1)$-th-dimension, get unstuck in i-th dimension): *If $V[a_1,\dots,a_{i+1}]$ is correctable and legal, and $A-\ell \le a_{i+1} < A$, then $V[a_1,\dots,a_{i-1},\mathsf{Next}_i(a_i+\ell)]$ is correctable.* The set is correctable because $V[a_1,\dots,a_{i-1},\mathsf{Next}_i(a_i+\ell)]$ equals $V[a_1,\dots,a_i,\infty]$, which equals $V[a_1,\dots,a_{i+1}]$. It is legal because $\mathsf{Next}_i(a_i)$ is not j-bad by definition of Next_i.

We can repeatedly apply these properties to get that the set of all qubits is correctable by induction. Here are the details. Let $\vec{t_1},\vec{t_2},\dots,$ be the enumeration of the $A_{\text{tot}}=(A+1)+(A+1)^2+\cdots+(A+1)^D$ tuples in $\{0,1,2,\dots,A\}^{\le D}$, in lexicographical order. Define the *lexicographical index* of a region $V[a_1,\dots,a_i]$ as the largest α such that t_α is lexicographically less than or equal to $(a_1,\dots,a_i)\in\mathbb{R}^{\le D}$. We prove by induction that, for all $r\le A_{\text{tot}}$, there exists a correctable and legal set with lexicographical index at least r. For the base case, $V[0]=\emptyset$ is clearly correctable and legal. For the induction step, suppose we have a correctable and legal set $V[a_1,\dots,a_i]$ with lexicographical index r. The above items shows that we can find a region with strictly larger lexicographical index: If $a_i\ge A$, either $i=1$, in which case we are done, or $i\ge 2$ and we apply item 4 – the lexicographical index increases because $\mathsf{Next}_i(a_i+\ell)-a_i\ge\ell\ge 1$. Otherwise, if $i=D$, we apply item 3. Otherwise, if $V[a_1,\dots,a_i+\ell]$ is legal, we apply item 1, and if not, we apply item 2. This completes the induction.

Since the entire set of qubits Q is correctable, an application of the AB Lemma (12) with $A=Q$ and $B=\emptyset$ implies that $k=0$. This contradicts our assumption that the code is nontrivial. ◀

Towards a Complexity-Theoretic Dichotomy for TQFT Invariants

Nicolas Bridges
Department of Mathematics, Purdue University, West Lafayette, IN, USA

Eric Samperton
Departments of Mathematics and Computer Science, Purdue University, West Lafayette, IN, USA

Abstract

We show that for any fixed $(2+1)$-dimensional TQFT over $\mathbb{C}$ of either Turaev-Viro-Barrett-Westbury or Reshetikhin-Turaev type, the problem of (exactly) computing its invariants on closed 3-manifolds is either solvable in polynomial time, or else it is #P-hard to (exactly) contract certain tensors that are built from the TQFT's fusion category. Our proof is an application of a dichotomy result of Cai and Chen [J. ACM, 2017] concerning weighted constraint satisfaction problems over $\mathbb{C}$. We leave for future work the issue of reinterpreting the conditions of Cai and Chen that distinguish between the two cases (*i.e.* #P-hard tensor contractions vs. polynomial time invariants) in terms of fusion categories. We expect that with more effort, our reduction can be improved so that one gets a dichotomy directly for TQFTs' invariants of 3-manifolds rather than more general tensors built from TQFTs' fusion categories.

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Theory of computation → Problems, reductions and completeness; Theory of computation → Quantum complexity theory

Keywords and phrases Complexity, topological quantum field theory, dichotomy theorems, constraint satisfaction problems, tensor categories

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.5

Related Version *Full Version*: https://arxiv.org/abs/2503.02945

Funding Both authors supported in part by NSF CCF grant 2330130.

1 Introduction

1.1 Main results

Quantum computation – especially *topological* quantum computation – motivates a number of complexity-theoretic questions concerning TQFT invariants of manifolds, particularly in dimensions 2 and 3. One of the most central is to classify "anyonic systems" according to whether or not they are powerful enough to (approximately) encode arbitrary quantum circuits over qubits. Anyons that are powerful in this way are important because (in theory) it should be possible to build fault tolerant quantum computers using them [12, 10]. We refer the reader to [23, 19] for a broad review of the mathematical side of these matters and Subsection 1.3 for more discussion. For now, we simply note that the Property F conjecture of Naidu and Rowell is currently the only concrete, published formulation of a proposed (partial) answer to this classification question that we know. The conjecture is surprisingly easy to formulate: the possible braidings of n copies of a simple anyon X in a unitary modular tensor category $\mathcal{B}$ generate only finitely many unitaries for each n (and, hence, are not "braiding universal" for quantum computation) if and only if the square of the quantum dimension of X is an integer $d_X^2 \in \mathbb{Z}$ [17].

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 5; pp. 5:1–5:21
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this work, we will not attack the Property F conjecture directly. However, our main result has a similar spirit and shares the same motivations. See Subsection 1.3 for some discussion of these points.

▶ **Theorem 1.**

(a) *Fix a spherical fusion category $\mathcal{C}$ over $\mathbb{C}$, presented skeletally with all data given as algebraic numbers over $\mathbb{Q}$. Then $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$–the problem of contracting tensor networks defined from $\mathcal{C}$ – is either solvable in polynomial time or $\#\mathsf{P}$-hard. Moreover, if M is a closed, oriented, triangulated 3-manifold (treated as computational input), then either the problem of computing the Turaev-Viro-Barrett-Westbury invariant $|M|_{\mathcal{C}} \in \mathbb{C}$ is solvable in (classical) polynomial time or $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$ is $\#\mathsf{P}$-hard.*

(b) *Fix a modular fusion category $\mathcal{B}$ over $\mathbb{C}$, presented skeletally with all data given as algebraic numbers over $\mathbb{Q}$. Then $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{B}})$–the problem of contracting tensor networks built from $\mathcal{B}$ – is either solvable in polynomial time or $\#\mathsf{P}$-hard. Moreover, if M is a closed, oriented 3-manifold encoded via a surgery diagram (treated as computational input), then either the problem of computing the Reshetikhin-Turaev invariant $\tau_{\mathcal{B}}(M) \in \mathbb{C}$ is solvable in (classical) polynomial time or $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{B}})$ is $\#\mathsf{P}$-hard.*

Two routine points of clarification are due.

First, we note that all fusion and modular categories over $\mathbb{C}$ admit finite skeletal presentations using algebraic numbers over $\mathbb{Q}$. This is because the defining equations for the skeletal data are all algebraic over $\mathbb{Q}$. In particular, since we are interested in how the complexity of $|M|_{\mathcal{C}}$ or $\tau_{\mathcal{B}}(M)$ depends on variable M for *fixed* $\mathcal{C}$ or $\mathcal{B}$, there is no harm in assuming that $\mathcal{C}$ and $\mathcal{B}$ are encoded in this way.

Second, as usual in computational 3-manifold topology, to say that a problem whose input is a triangulated 3-manifold is solvable in polynomial time means that there exists an algorithm to solve the problem that runs in time polynomial in the size of the triangulation. For a 3-manifold presented via integral surgery on a link diagram in S^3, the algorithm must run in time jointly polynomial in the crossing number of the link diagram, the number of components of the link, and the absolute values of the surgery coefficients.[1]

We refer the reader to [7] for further elaboration of both of these matters.

We now explain briefly the meaning and importance of dichotomy theorems within complexity theory. Of course, it is an infamous open problem to show that $\mathsf{P} \neq \mathsf{NP}$ (the two complexity classes might be equal, but most experts do not expect this to be the case). To establish this inequality it is necessary and sufficient to show that there exists an NP-complete problem with no polynomial-time algorithm. Intriguingly, Ladner showed that if $\mathsf{P} \neq \mathsf{NP}$, then there exist problems in NP that are neither in P nor NP-complete [16]. These are usally referred to as NP-*intermediate*. In other words, an NP-intermediate problem is a problem in NP that is neither in P nor NP-hard. Intuitively, Ladner's theorem shows that if one considers a family of decision problems, then it need not be the case that every problem in the family is either "easy" (that is, in P) or "hard" (that is, NP-hard)–there could be problems that have intermediate complexity. When a given family of problems has the property that none of the problems has intermediate complexity, then one says that the family satisfies a *dichotomy theorem*. The archetypical dichotomy theorem was established by Schaefer, who showed that Boolean satisfiability problems in generalized conjunctive form (where the clauses are taken from a finite set of constraints) satisfy a dichotomy theorem (with respect to the set of constraints) [21].

[1] In particular, we might understand the surgery coefficients as being expressed in unary, not binary. (If we used the latter, then we would not be able to build a triangulation from a surgery diagram in polynomial time.)

In the case of our Theorem 1, we interpret (2+1)-d TQFT invariants as generalized (*i.e.* $\mathbb{C}$-valued instead of $\mathbb{N}$-valued) counting problems parametrized by spherical fusion categories $\mathcal{C}$ and modular fusion categories $\mathcal{B}$ (the categories are analogs of the allowed local constraints in Schaefer's dichotomy). Our results establish the dichotomy that either a function of the type $M \mapsto |M|_{\mathcal{C}} \in \mathbb{C}$ or $M \mapsto \tau_{\mathcal{B}}(M) \in \mathbb{C}$ is "easy" to compute (polynomial time) or else it is "hard" to contract certain tensors built from the category $\mathcal{C}$ or $\mathcal{B}$ (#P-hard). In this way, there are no (2+1)-d TQFTs whose tensors are of "intermediate" complexity. In fact, we conjecture that the same can be said directly of the TQFT's *invariants of 3-manifolds per se*. Let us expound on these points now.

1.2 Mapping the dichotomy

Whether or not $M \mapsto |M|_{\mathcal{C}}$ or $M \mapsto \tau_{\mathcal{B}}(M)$ is computable in polynomial time depends on $\mathcal{C}$ and $\mathcal{B}$. Having established Theorem 1 – which only asserts the *existence* of a dichotomy – it is natural to wonder where one should draw the line between easy and hard. Better yet, ideally, one would like to be able to prove that the dichotomy of Theorem 1 is *effective*, meaning, given $\mathcal{C}$ or $\mathcal{B}$, there exists a polynomial-time algorithm to decide precisely when the category falls into the easy case (here "polynomial-time" means in the size of the skeletalization of $\mathcal{C}$ or $\mathcal{B}$). Our proof of Theorem 1 relies on the main result of Cai and Chen's work [2], which establishes a dichotomy theorem for a generalized type of "solution counting" to constraint satisfaction problems $\#\mathsf{CSP}(\mathcal{F})$ with a fixed "$\mathbb{C}$-weighted constraint family" $\mathcal{F}$. We carefully define $\#\mathsf{CSP}(\mathcal{F})$ in Subsection 2.1, but here we note that not only do Cai and Chen prove that for every choice of constraint family $\mathcal{F}$, $\#\mathsf{CSP}(\mathcal{F})$ is either #P-hard or computable in polynomial time – they also provide three necessary and sufficient conditions that *characterize precisely* which $\mathcal{F}$ allow for polynomial time solutions to $\#\mathsf{CSP}(\mathcal{F})$. These conditions are called "block orthogonality," "Mal'tsev" and "Type Partition". Our proof of Theorem 1 consists in converting a spherical fusion category $\mathcal{C}$ or modular fusion category $\mathcal{B}$ into an appropriate constraint family $\mathcal{F}_{\mathcal{C}}$ or $\mathcal{F}_{\mathcal{B}}$ such that computing $|M|_{\mathcal{C}}$ or $\tau_{\mathcal{B}}(M)$ is equivalent to computing an instance (depending on M) of a problem in $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$ or $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{B}})$, respectively. In particular, for the constraint families $\mathcal{F}_{\mathcal{C}}$ and $\mathcal{F}_{\mathcal{B}}$ we shall build, it should be possible to interpret the three conditions of Cai and Chen directly in terms of the categories $\mathcal{C}$ and $\mathcal{B}$. It is beyond the scope of the present work to attempt to accomplish this. However, we believe this is an important problem, since it should shed light on variations of the Property F conjecture related to anyon classification, as we explain in the next subsection.

Let us now address the more important deficiency of Theorem 1, alluded to at the end of the previous subsection: it would be better to get an outright dichotomy for 3-manifold invariants, and not just general tensors derived from a fusion category. To this end, Theorem 1 can be understood as a first step towards proving the following more desirable result.

▶ **Conjecture 2.**

(a) *Fix a spherical fusion category $\mathcal{C}$ over $\mathbb{C}$, presented skeletally with all data given as algebraic numbers over $\mathbb{Q}$. If M is a closed, oriented, triangulated 3-manifold (treated as computational input), then computing the Turaev-Viro-Barrett-Westbury invariant $|M|_{\mathcal{C}} \in \mathbb{C}$ is either solvable in (classical) polynomial time or is #P-hard.*

(b) *Fix a modular fusion category $\mathcal{B}$ over $\mathbb{C}$, presented skeletally with all data given as algebraic numbers over $\mathbb{Q}$. If M is a closed, oriented 3-manifold encoded via a surgery diagram (treated as computational input), then the problem of computing the Reshetikhin-Turaev invariant $\tau_{\mathcal{B}}(M) \in \mathbb{C}$ is either solvable in (classical) polynomial time or is #P-hard.*

See Subsection 3.4 for some discussion of how we expect one might get started on proving this conjecture – in particular, the relevance of holant problems [3].

1.3 Implications for "anyon classification"

Theorem 1 and Conjecture 1 assert that for certain precise formulations of the problem of "anyon classification," whatever the "type" is for a given modular fusion category $\mathcal{B}$, it can only be one of two things, with no "intermediate" cases. In order to explain this more carefully, we pause to note the many ways one can make the problem of anyon classification precise, and situate our result exactly in this milieu. Moving from the more "purely mathematical" to the more "applied" end of the spectrum, "anyon classification" could mean any of the following precise problems:

1. Algebraic classification of unitary modular fusion categories (MFCs) up to ribbon tensor equivalence.
   - Much of the literature on fusion categories can be considered as contributing to this problem.
   - Presumably one would be satisfied with a solution to this problem "modulo finite group theory."
2. Algebraic classification of simple objects X in unitary MFCs $\mathcal{B}$ according to whether or not the braid group representations $B_n \to U(\mathrm{End}_{\mathcal{B}}(X^{\otimes n}))$ have finite image, dense image, or something else.
   - The Property F conjecture is of course directly related to this matter.
   - One can generalize this question to consider mapping class group representations of higher genus surfaces with different types of anyons on them.
3. Complexity-theoretic classification of MFCs according to how easy or hard it is to *exactly compute* their Reshetikhin-Turaev 3-manifold invariants (as algebraic numbers over $\mathbb{Q}$).
   - **Our Theorem 1 is situated here – almost!** We have established a dichotomy of the form either "3-manifold invariants easy" or "tensors in the category are hard to contract". Our results represent a non-trivial step towards the desired dichotomy of Conjecture 2: "invariants easy" or "invariants hard".
   - One can ask this question for *restricted classes of 3-manifolds* (such as "knots in S^3" or "links in S^3" or "integer homology spheres"), and the classification may change [7].
4. Complexity-theoretic classification of MFCs according to how easy or hard it is to "*approximate*" their Reshetikhin-Turaev 3-manifold invariants.
   - There are different types of approximations one might ask for. *A priori*, each type should be understood as giving a different version of this question.
   - "Exactly compute" is one way to "approximate."
   - Pioneering works of Freedman, Kitaev, Larsen and Wang show that for certain approximation schemes, there exists unitary MFCs $\mathcal{B}$ for which the ability to approximate their 3-manifold invariants is equivalent in power to BQP (bounded-error quantum polynomial time) [9, 10]. In particular, their work established the original paradigm for topological quantum computation via anyon braiding.
   - Kuperberg showed that results for one type of approximation can have important implications for other types of approximations [13]. In particular, the kinds of approximations that a quantum computer can efficiently make for Reshetikhin-Turaev invariants are (in general/worst case) not precise enough to do anything useful for distinguishing 3-manifolds even if their invariants are promised to be unequal by a large amount.

5. Complexity-theoretic classification of MFCs according to whether or not they support universal quantum computation with braiding *and adaptive anyonic charge measurements*.
   - Quantum computation using braidings and charge measurements of anyons in a fixed unitary MFC is often called “topological quantum computing with adaptive charge measurement.”
   - Our entirely subjective opinion is that this is the most important flavor of anyon classification, at least when considered from the perspective of the goal of actually building a universal, fault-tolerant quantum computer.
   - Even unitary MFCs whose anyons all have Property F (and, hence, are not universal via braiding alone) can be universal when braiding is supplemented with charge measurements, see e.g. [6].
   - While adaptive charge measurement is generally considered fault-tolerant for topological reasons, unlike the case of braiding-only topological quantum computing, the amplitudes with which one performs a quantum computation in this paradigm are not (normalizations of) Reshetikhin-Turaev invariants of 3-manifolds.

There are known relations between these different classification problems. For example, on one hand, if X is an anyon such that $B_n \to U(\mathrm{End}_{\mathcal{B}}(X^{\otimes n}))$ is dense, then the Solovay-Kitaev theorem implies that $\mathcal{B}$ supports universal topological quantum computation via braiding (without needing adaptive charge measurement). On the other hand, if X has Property F, then it is known that braiding with X is never powerful enough to encode all of BQP in its braidings. This latter point was the main motivation for the Property F conjecture in the first place, since one would like to rule out the “obviously” un-useful anyons easily.

To understand the potential usefulness of Theorem 1 or Conjecture 2, it is perhaps helpful to pull on the thread of these motivations for the Property F conjecture a bit more so that we can compare and contrast.

On one hand, there is no “unconditional” implication known between classification problems (2) and (4) above in either direction, except if we condition on properties in a way we have already mentioned, namely: if an anyon has Property F, then it is definitely not braiding universal, while if an anyon has dense braidings, then it is universal. This is not “unconditional” in the sense that as far as problem (2) is concerned, there is a third case that remains to be addressed: anyons with braidings that are neither dense nor have Property F. Do they even exist? If so, what are we to make of them? Are they universal or not? Maybe sometimes they are and sometimes they are not? Conversely, if an anyon is braiding universal, does it necessarily have dense braid group representations? These are interesting questions worth pursuing, but it could require quite a bit of effort to resolve each of them.

On the other hand, there is an “unconditional” connection between (3) and (4) in at least one direction: (3) is simply the special case of (4) where the type of “approximation” is chosen to be “exact computation.” So classification problem (3) might be understood as a warm-up to the version of problem (4) where the type of approximation is not “exact”, but is instead the kind of approximation relevant to topological quantum computing. (For the sake of space, we refrain from precisely defining this type of approximation here; see the intro discussions of [13] or [20].) The key technical issue that this perspective highlights is the following: even categories whose anyons all have property F (and thus are not braiding universal) can have $\#\mathsf{P}$-hard invariants [11, 14, 15]. Hence, more work needs to be done to properly understand the relationship between the BQP-universality of anyon braidings in a given modular fusion category $\mathcal{B}$ and $\#\mathsf{P}$-hardness of (exactly) computing $\tau_{\mathcal{B}}(M)$ on 3-manifolds M. At the end of the day this is not so different from the situation between (2) and (4).

However, we conclude this discussion by noting that it is conceivable there exists a very tight connection between anyon classification problems (3) and (5) (while there is essentially no way to relate (2) and (5)). Indeed, one might reasonably guess that #P-hardness for the exact calculation of invariants implies that topological quantum computing with adaptive charge measurements is always sufficient to generate BQP-universal topological gates. This guess would be consistent with all known examples. We plan to explore these matters in future work.

1.4 Outline

Subsection 2.1 briefly reviews the definiton of #CSP($\mathcal{F}$), as well as Cai and Chen's dichotomy theorem for these problems. Subsection 2.2 contains the proof of part (a) of Theorem 1. Subsection 2.3 contains some preliminary results about the graphical calculus in fusion categories needed for the proof of part (b). The proof of part (b) itself is relegated to Appendix A for the sake of space. Section 3 contains some further discussion.

2 Proof of Theorem 1

2.1 Cai and Chen's dichotomy for weighted CSPs

Before proving either part of our main theorem, we review the definition of #CSP($\mathcal{F}$), following [2]:

- We fix a finite set $D = \{1, \dots, d\}$ called the *domain* (which, by an abuse of notation, we will suppress from the notation #CSP($\mathcal{F}$)).
- We fix a *($\mathbb{C}$-valued) weighted constraint family* $\mathcal{F} = \{f_1, \dots, f_h\}$, where each f_i is a $\mathbb{C}$-valued function $f_i : D^{r_i} \to \mathbb{C}$ for some $r_i \geq 1$ called the *arity* of f_i. We assume all the values that the f_i assume are encoded as algebraic numbers over $\mathbb{Q}$.
- An *instance* I of #CSP($\mathcal{F}$) consists of a tuple $\mathbf{x} = (x_1, \dots, x_n)$ of variables over D (which will be suppressed in our notation) and a set I of tuples $(f, i_1, \dots, i_r)$ in which f is an r-ary function from $\mathcal{F}$ and $i_1, \dots, i_r \in \{1, \dots, n\}$ are indices of the variables in $\mathbf{x}$.
- The *output* of #CSP($\mathcal{F}$) on instance I is the algebraic number $Z(I) \in \mathbb{C}$ given by

$$Z(I) \overset{\text{def}}{=} \sum_{\mathbf{x} \in D^n} F_I(\mathbf{x}),$$

where

$$F_I(\mathbf{x}) \overset{\text{def}}{=} \prod_{(f, i_1, \dots, i_r) \in R} f(x_{i_1}, \dots, x_{i_r}).$$

The main result of [2] is

▶ **Theorem 3** (Thm. 1, [2]). *Given any constraint set $\mathcal{F}$ as above, #CSP($\mathcal{F}$) is either computable in polynomial time or #P-hard.*

2.2 Proof of Theorem 1(a)

Proof. Let $\mathcal{C}$ be a spherical fusion category over $\mathbb{C}$. To prove part (a) of Theorem 1, it suffices – thanks to Theorem 3 – to build a domain $D_{\mathcal{C}}$ and weighted constraint set $\mathcal{F}_{\mathcal{C}}$ with the following property: there exists a polynomial time algorithm that converts a triangulated 3-manifold M into an instance I_M of #CSP($\mathcal{F}_{\mathcal{C}}$) such that

$$Z(I_M) = |M|_{\mathcal{C}}.$$

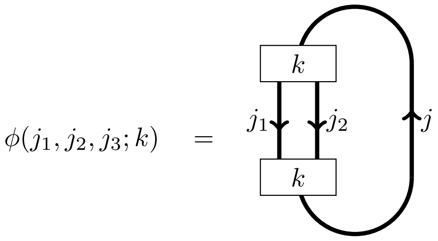

Figure 1 A 3j+1k-symbol.

Readers already familiar with the state-sum formula for TVBW invariants will notice that the definition of $Z(I)$ is quite similar in spirit. The goal of the present proof is simply to make this similarity precise.

To this end, let us recall the state-sum formula for $|M|_{\mathcal{C}}$:

$$|M|_{\mathcal{C}} = \mathcal{D}^{-2|V_M|} \sum_{\substack{L:E_M \to \mathrm{Irr}(\mathcal{C}) \\ F_L:F_M \to \mathcal{N} \text{ consistent w/ } L}} \frac{\prod_{e\in E_M} \dim(L(e))^2 \prod_{t\in T_M} |t^L|}{\prod_{f\in F_M} |f^L|}$$

where our notation is as follows:

- V_M is the ordered list of vertices in the triangulation M and $\mathcal{D}$ is the total quantum dimension of $\mathcal{C}$.
- E_M is the set of edges in the triangulation M and $\mathrm{Irr}(\mathcal{C})$ is set of simple objects in the given skeletalization of $\mathcal{C}$.
- F_M is the set of faces in the triangulation M and $\mathcal{N}$ is the set of labels of the trivalent Hom spaces

  $\mathrm{Hom}(k, i \otimes j) = \mathrm{span}\Big\{$ 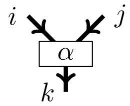 $\Big\}_{\alpha=1,\ldots,N^k_{ij}}$

  and $|f^L|$ is the 3j+1k-symbol obtained by evaluating the face f with a given labeling L of the edges and faces of M (See Figure 1).
- T_M is the set of tetrahedra in the triangulation M, and $|t^L|$ is the 6j+4k-symbol obtained by evaluating the tetrahedron t with a given labeling L of the edges and faces of M, where we take into account whether the orientation of t given by the orientation of M matches the induced orientation given by the ordering of the vertices. This will be made more precise below.

Since we assume $\mathcal{C}$ is not necessarily multiplicity-free, then instead of 6j-symbols, we will be using so-called 6j+4k-symbols $\begin{bmatrix} j_1 & j_2 & j_3 & k_{1,2} & k_{2,3} \\ j & j_{12} & j_{23} & k_{12,3} & k_{1,23} \end{bmatrix}^{\pm}$, which are defined by the contraction of a specific colored graph.

$$\begin{bmatrix} j_1 & j_2 & j_{12} & k_{1,2} & k_{2,3} \\ j_3 & j & j_{23} & k_{12,3} & k_{1,23} \end{bmatrix}^+ \overset{\text{def}}{=}$$

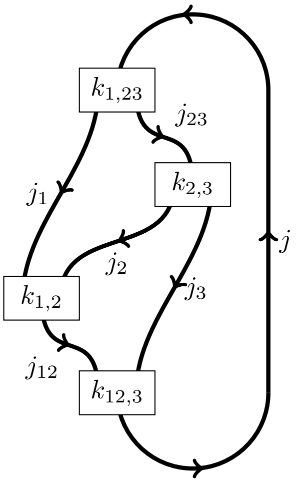

This defines the "positive" 6j+4k-symbols. We also define a "negative" version of the 6j+4k-symbols. We will call them negative 6j+4k-symbols since they correspond to negatively-oriented tetrahedra with respect to the standard orientation on $\mathbb{R}^3$:

$$\begin{bmatrix} j_1 & j_2 & j_{12} & k_{1,2} & k_{2,3} \\ j_3 & j & j_{23} & k_{12,3} & k_{1,23} \end{bmatrix}^- \overset{\text{def}}{=}$$

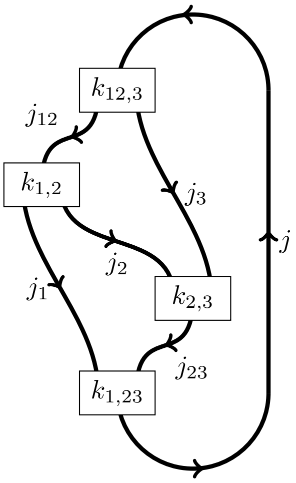

Here, $j_1, j_2, j_3, j, j_{12}, j_{23}$ are simple objects, $k_{1,2} \in \{0, \ldots, N_{j_1 j_2}^{j_{12}}\}$, $k_{2,3} \in \{0, \ldots, N_{j_2 j_3}^{j_{23}}\}$, $k_{12,3} \in \{0, \ldots, N_{j_{12} j_3}^{j}\}$, and $k_{1,23} \in \{0, \ldots, N_{j_1, j_{23}}^{j}\}$.

We now have enough to identify our domain and weighted constraint set. Define

$$D_{\mathcal{C}} \overset{\text{def}}{=} \operatorname{Irr}(\mathcal{C}) \sqcup \mathcal{N} \sqcup \{*\}.$$

Now extend the 6j+4k symbols to be 10-ary functions on our domain $D_{\mathcal{C}}$ in the "trivial" way:

$$\Delta^+(x_1, \ldots, x_{10}) \overset{\text{def}}{=} \begin{cases} \begin{bmatrix} x_1 & x_2 & x_5 & x_7 & x_8 \\ x_3 & x_4 & x_6 & x_9 & x_{10} \end{bmatrix}^+ & \text{if } x_1, \ldots, x_6 \in \operatorname{Irr}(\mathcal{C}),\ x_7, \ldots, x_{10} \in \mathcal{N}, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\Delta^-(x_1, \ldots, x_{10}) \overset{\text{def}}{=} \begin{cases} \begin{bmatrix} x_1 & x_2 & x_5 & x_7 & x_8 \\ x_3 & x_4 & x_6 & x_9 & x_{10} \end{bmatrix}^- & \text{if } x_1, \ldots, x_6 \in \operatorname{Irr}(\mathcal{C}),\ x_7, \ldots, x_{10} \in \mathcal{N}, \\ 0 & \text{otherwise.} \end{cases}$$

We similarly define 4-ary functions on our domain using the 3j+1k-symbols ϕ by taking

$$\Phi^{-1}(x_1,x_2,x_3,x_4) \overset{\text{def}}{=} \begin{cases} \phi(x_1,x_2,x_3;x_4)^{-1} & \text{if } x_1,x_2,x_3 \in \operatorname{Irr}(\mathcal{C}),\ x_4 \in \mathcal{N}, \\ 0 & \text{otherwise.} \end{cases}$$

And we define 1-ary functions using the quantum dimensions of simple objects:

$$\mathrm{d}^2(x) \overset{\text{def}}{=} \begin{cases} \dim(x)^2 & \text{if } x \in \operatorname{Irr}(\mathcal{C}), \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we define a 1-ary function to encode the total quantum dimension of $\mathcal{C}$:

$$\mathcal{D}^{-2} \overset{\text{def}}{=} \mathcal{D}^{-2}(x) \overset{\text{def}}{=} \begin{cases} \left(\sum_{j\in\operatorname{Irr}(\mathcal{C})} \dim(j)^2\right)^{-1} & \text{if } x = *, \\ 0 & \text{otherwise.} \end{cases}$$

Using these functions, we define our weighted constraint family

$$\mathcal{F}_{\mathcal{C}} \overset{\text{def}}{=} \{\Delta^{\pm}, \Phi^{-1}, \mathrm{d}^2, \mathcal{D}^{-2}\}.$$

Our next goal is to describe how to convert a triangulation M of an oriented manifold into an instance I_M of $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$.

The data of M is comprised of:

- An ordered list of vertices $\{v_1,\dots,v_a\}$.
- A list of oriented edges $\{e_1(v_{1_1},v_{2_1}),\dots,e_b(v_{b_1},v_{b_2})\}$ where $e_i(v_{i_1},v_{i_2})$ means that e_i is an edge connecting v_{i_1} to v_{i_2}. (Note that these orientations are chosen arbitrarily.)
- A list of oriented faces $\{f_1(e_{1_1},e_{2_1},e_{3_1}),\dots f_c(e_{c_1},e_{c_2},e_{c_3})\}$ where $f_i(e_{i_1},e_{i_2},e_{i_3})$ means that f_i is a face whose boundary consists of the edges e_{i_1}, e_{i_2}, and e_{i_3}. (Note that the orientations of the faces need to be consistent with the edge orientations in any way.)
- A list of tetrahedra $\{t_1,\dots,t_d\}$ where $t_i = t_i(f_{i_1},\dots,f_{i_4})$ means that t_i is a tetrahedron with faces given by f_{i_1}, …, f_{i_4}.
- To encode the orientation of M, each tetrahedron t_i is endowed with a sign $+$ or $-$ to indicate the local orientation inside that tetrahedron.[2]

For a given triangulation M as described, define a tuple

$$\mathbf{x}_M \overset{\text{def}}{=} (x_1,\dots,x_a,y_1,\dots,y_b,z_1,\dots,z_c)\}$$

that has a variable for each vertex, edge and face in M. We now describe how to build the desired instance I_M of $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$. It will be clear from the construction that the mapping $M \mapsto I_M$ can be done in polynomial time in the size of M.

First we put the functions $\mathcal{D}^{-2}(x_1),\dots,\mathcal{D}^{-2}(x_a)$ and $\mathrm{d}^2(y_1),\dots,\mathrm{d}^2(y_b)$ in I_M for every vertex and edge of M. For each face $f_j(e_{j_1},e_{j_2},e_{j_3})$, we include $\Phi^{-1}(y_{j_1},y_{j_2},y_{j_3},z_j)$. Finally, for each tetrahedron t_i, we include either

$$\Delta^{+}(y_{j_1},\dots,y_{j_6},z_{i_1},\dots,z_{i_4})),$$

[2] These signs must assemble to give a $\{\pm\}$-valued 0-cocycle on the dual cellulation. This condition could be easily checked, but for our purposes it is simply part of the data structure of M, and so this condition can be assumed to be met as a promise. This condition is not necessary for our proof (although it is necessary for the proof that $|M|_{\mathcal{C}}$ is an invariant of M).

or

$$\Delta^{-}(y_{j_1},\ldots,y_{j_6},z_{i_1},\ldots,z_{i_4})),$$

where t_i has faces $f_{i_1}(e_{j_1},e_{j_2},e_{j_5})$, $f_{i_2}(e_{j_5},e_{j_3},e_{j_4})$, $f_{i_3}(e_{j_3},e_{j_2},e_{j_6})$, and $f_{i_4}(e_{j_6},e_{j_1},e_{j_4})$. To determine whether we should include Δ^+ or Δ^- for t_i, we check if the orientation of t_i given by the orientation of M matches the induced orientation by the ordering of the vertices; if they match, then we use Δ^+, and otherwise we use Δ^-.

This I_M defines an instance of $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$ that computes the Turaev-Viro invariant for M. Indeed, plugging in the definitions of our constraint functions, we get

$$Z(I_M) = \sum_{\mathbf{x}\in D^{a+b+c}} \prod_{v=1}^{a} \mathcal{D}^{-2}(x_v) \prod_{e=1}^{b} \mathrm{d}^2(y_e) \prod_{F_i=(E_{i_1},E_{i_2},E_{i_3})} \Phi^{-1}(y_{i_1},y_{i_2},y_{i_3},z_i) \prod_{T_i=(F_{i_1},\ldots,F_{i_4})} \Delta^{\eta_i}(y_{j_1},\ldots,y_{j_6},z_{i_1},\ldots,z_{i_4}))$$

where $\eta_i = +$ if the orientation of T_i given by the orientation of M matches the induced orientation by the ordering of the vertices, and $\eta_i = -$ otherwise. *A priori*, $Z(I_M)$ includes a sum over more types of labelings than the state-sum formula for $|M|_{\mathcal{C}}$. However, because of how we have chosen to define the functions in the constraint family $\mathcal{F}_{\mathcal{C}}$, all of these additional terms in the sum vanish. To see this, first note that when $x_1,\ldots,x_v \neq *$, the entire term is 0. In particular, this means all non-zero terms have a common factor of the global quantum dimension to the $-a$ power, and hence we can pull it out as the normalizing factor. Similarly, when the edges are not labeled by elements of the domain $D_{\mathcal{C}}$ that are not simple objects of $\mathcal{C}$, or the faces are not labeled with multiplicities, the terms are zero. We have furthermore arranged so that when the labeling of the edges and faces is not admissible, the 6j+4k-symbol for that term vanishes. Therefore, the only surviving terms in the sum $Z(I_M)$ are the $\mathbf{x}$ which define admissible labelings of the edges and faces of M. It is then straightforward to see $Z(I_M) = |M|_{\mathcal{C}}$ recovers the Turaev-Viro invariant. ◀

2.3 Graphical calculus preliminaries

Before proving part (b) of Theorem 1, we establish a technical result about the graphical calculus in a spherical fusion category. The result is likely well-known to experts, but does not appear in the literature anywhere that we are aware. We begin by reviewing what we need of closed trivalent graphs in S^2.

For our purposes, a *(closed) trivalent graph* Γ in $\mathbb{R}^2$ (or $S^2 = \mathbb{R}^2 \cup \{\infty\}$) has:

- A finite collection V of vertices $v_1, v_2, \ldots, v_n$, where $v_i \in \mathbb{R} \times \{i\}$
- A collection E of directed edges $e_1, e_2, \ldots, e_k$ in S^2

subject to the conditions that

- Each edge e_i is either a loop disjoint from V or an arc connecting two (not necessarily distinct) vertices, with an interior that is disjoint from all vertices in V.
- If v is a vertex, then there is an open disk neighborhood $D(v)$ so that $D(v) \cap \Gamma$ has three arcs (coming from intersections of $D(v)$ with E, not necessarily distinct) emanating from v with one arc parallel to the vector $\langle 0,1\rangle$, one arc parallel to the vector $\langle 1,1\rangle$, and one arc parallel to the vector $\langle -1,1\rangle$.

Such a graph Γ is closed in the sense that there are no vertices that are involved in precisely one half-edge.

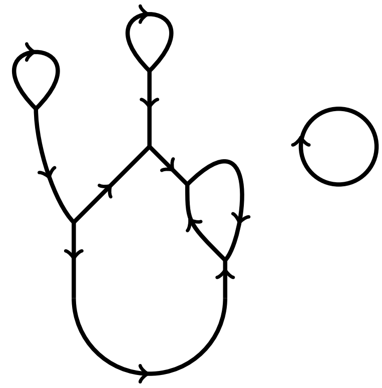

Figure 2 A trivalent graph in $\mathbb{R}^2$ which exhibits all three types of edges.

We will also need to consider closed trivalent graphs which have crossings. We will call these *crossed (closed) trivalent graphs.* Crossed trivalent graphs are trivalent graphs where they are allowed to have finitely many double points in the interior of its edges. These double points must indicate which segment of the edge crosses "over" or "under" the other.

Recall that computing the Reshetikhin-Turaev invariant $\tau_{\mathcal{B}}(M)$ of a 3-manifold M presented by a surgery diagram involves a process where we interpret a coloring of the components of the diagram by simple objects of $\mathcal{B}$ as an endomorphism of the tensor unit $\mathbf{1}$ of $\mathcal{B}$. Since $\mathbf{1}$ is itself simple, such a coloring gives rise to an endomorphism $\mathbf{1} \to \mathbf{1}$, which in turn can be identified with a complex number because $\mathrm{End}(\mathbf{1}) = \mathbb{C}$. The invariant $\tau_{\mathcal{B}}(M)$ is then (roughly) the sum of all of these numbers over all choices of colorings of the surgery diagram of M. For any single coloring, the complex number associated to it can generally be understood as the result of a sequence of tensor contractions on a tensor induced by that coloring. Moreover, this sequence of tensor contractions can be represented diagramatically, using a small number of standard diagrammatic operations that are determined from the data of the modular fusion category $\mathcal{B}$. We call these operations Circle Removal, Tadpole Trim, Bubble Pop, F-Move, and Vertex Spiral (aka "bending" moves); see Figures 3-5. We also allow ourselves to reverse edge orientations. To compute the complex number associated to a colored surgery link, one must identify a sequence of these operations that simplifies the diagram to the empty diagram. The desired number will then be a product of numbers determined from the operations in the sequence and the given coloring.

Our proof of Theorem 1(b) essentially revolves around two key observations.

First, a kind of uniformity: given a surgery description of M, there exists a *single sequence* of diagrammatic operations that can be used to evaluate *all* colorings of the surgery link to complex numbers. This uniformity is, more-or-less, what will make it possible to encode $M \mapsto \tau_{\mathcal{B}}(M)$ as an instance I_M of #$\mathsf{CSP}(\mathcal{F})$ for an appropriately chosen $\mathcal{F}$.

Second: such a uniform sequence of operations can be identified in polynomial time from the surgery description of M. This will imply that the reduction $M \mapsto I_M$ can be performed in polynomial time. We make this point more precise now.

▶ **Lemma 4.** *If $\Gamma \subset S^2$ is a closed trivalent graph embedded in S^2, then there is a polynomial time algorithm (in the size of the encoding of Γ) to construct a sequence of embedded graphs $\Gamma_0, \Gamma_1, \ldots, \Gamma_l$ where $\Gamma_0 = \Gamma$, $\Gamma_l = \emptyset$ such that each Γ_{i+1} is related to Γ_i by one of the diagrammatic operations in Figures 3-5 or edge orientation reversals.*

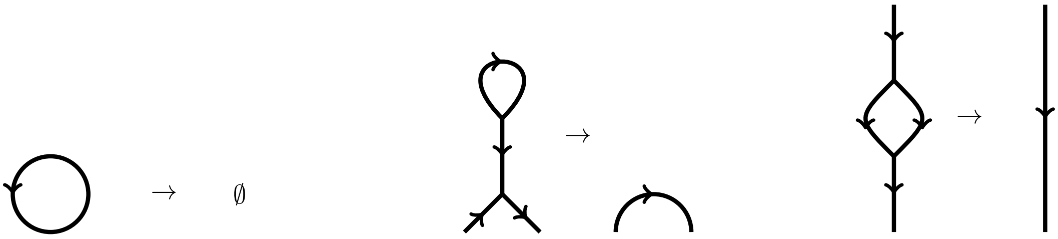

Figure 3 A circle removal, tadpole trim, and bubble pop, respectively.

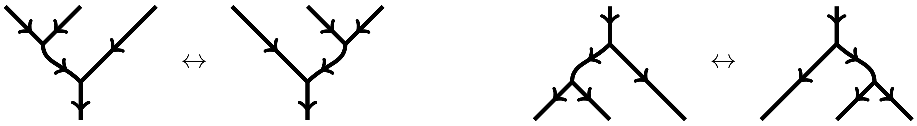

Figure 4 The F-moves.

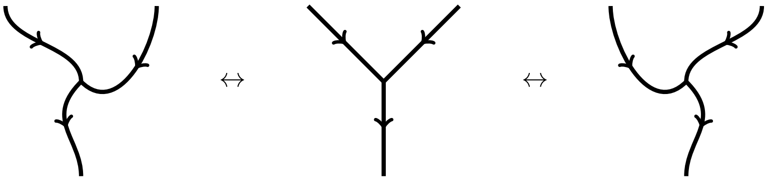

Figure 5 Vertex spiral (aka "bending").

Proof. A simple greedy algorithm suffices. We sketch the idea.

Begin by greedily choosing a complementary region R of Γ, *i.e.* a connected component R of $S^2 \setminus \Gamma$. Note that we may identify the boundary edges and vertices of R in polynomial time. Suppose R has k unique edges and l unique vertices on its boundary. Now simplify and update Γ according to the following cases.

1. If $(k, l) = (0, 0)$, then $\Gamma = \emptyset$, and so we terminate.
2. If $(k, l) = (1, 0)$, then the boundary of R is a circle in Γ, which we remove as in Figure 3.
3. If $(k, l) = (1, 1)$, then the boundary of R is part of a tadpole, which we trim as in Figure 3, but possibly only after first applying an appropriate set of vertex spirals as in Figure 5 and edge orientation reversals.
4. If $(k, l) = (2, 2)$, then the boundary of R is part of a bubble, which we pop as in Figure 3, but possibly only after first applying an appropriate set of vertex spirals and edge orientation reversals.
5. Otherwise, $k = l > 2$. Greedily pick an edge e on the boundary of R. After perhaps first applying up to two vertex spirals and 5 edge orientation reversals, we can arrange so that around e, Γ looks like one of the four diagrams in Figure 4, with e the edge in the middle. Apply the available F-move around e. The complementary regions of the resulting graph are naturally in bijection with the regions of the previous graph (see Figure 6 for an example). Let R' be the region of the new graph associated with R. If R' has $k = l > 2$ edges on its boundary, then repeat what we just did, but with R' and the new graph, instead of R; otherwise, R' has $k = l = 2$ edges, and we pop the bubble as in case (3).

Repeat this process of greedily picking a region R and proceeding as in the above cases. Each step of identifying an R and carrying through the appropriate case takes polynomial time, and, moreover, reduces the number of complementary regions of Γ by 1. Since there are at most a polynomial number of complementary regions to begin with, the entire procedure takes place in polynomial time. ◀

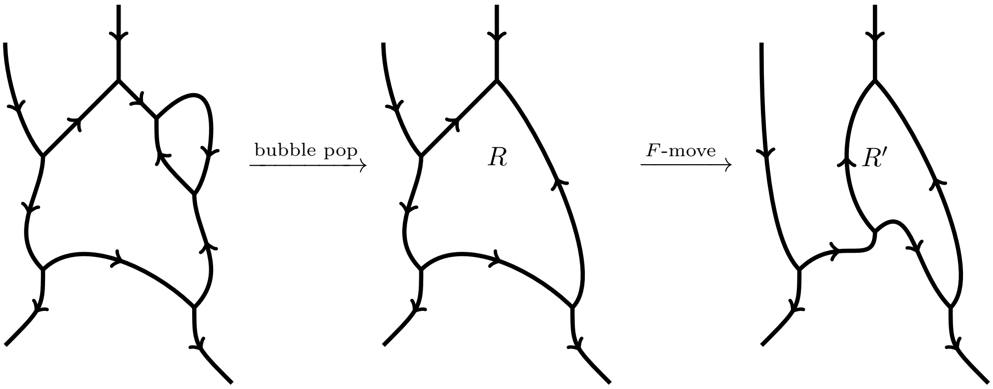

Figure 6 An example of a portion of the algorithm.

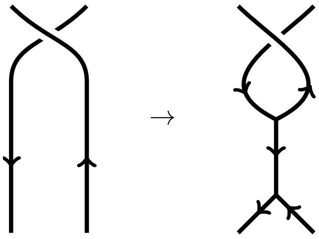

Figure 7 Inserting trivalent vertices to resolve the identity. Following up with R-moves reduces us to planar diagrams.

Lemma 4 only involves *planar* trivalent graphs, while the proof of part (b) of Theorem 1 needs crossings. Indeed, when we compute Reshetikhin-Turaev invariants, a 3-manifold is encoded by a framed link diagram L (which can be assumed to be in plat position, as in Figure 8). Fortunately, for modular fusion categories we have "R-moves" that – together with a "resolution of the identity" trick shown in Figure 7 – allow us to reduce to the planar case, and thereby use Lemma 4. The basic strategy is then not so different from the proof of part (a), but it requires that we introduce new variables for every crossing and then every step of the algorithm of the lemma. The details are found in Appendix A.

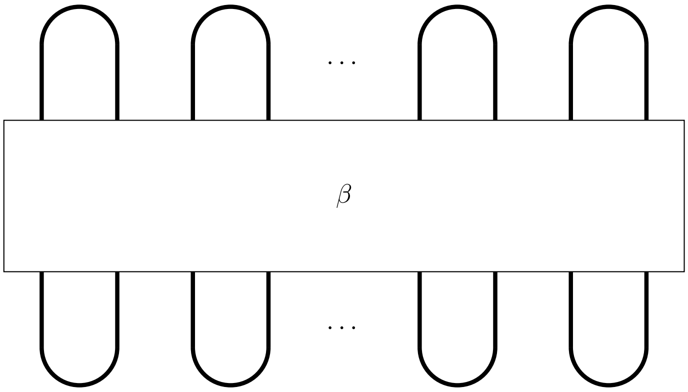

Figure 8 A blackboard-framed link L in standard plat position is determined by a braid word $\beta \in B_{2k}$.

3 Discussion

3.1 Unitarity

We note that none of our results depend on the unitarity of the spherical fusion category $\mathcal{C}$ or the modular fusion category $\mathcal{B}$. This is not surprising, since a choice of unitary structure is not necessary to define the TQFT invariants of closed 3-manifolds from $\mathcal{C}$ or $\mathcal{B}$, and so, as far as exact calculation of invariants is concerned, such a choice will not affect any dichotomies. Nevertheless, a unitary structure is certainly needed in order to do topological quantum computation with the TQFT determined by $\mathcal{C}$ or $\mathcal{B}$, since, for such applications, one needs the TQFT to be unitary. *A priori*, a specific choice of unitary structure might affect the BQP-universality of braiding (with or without adaptive measurement); however, *a posteriori*, this is not the case because Reutter showed that unitarizable fusion categories admit *unique* unitary structures [18] (we thank Milo Moses for reminding us of this reference).

3.2 TQFTs in other dimensions

We furthermore note that the same strategy we used for the proof of Theorem 1(a) should work more generally to prove that any fully-extended $(d+1)$-dimensional TQFT in any dimension will satisfy a similar dichotomy involving its invariants of closed $(d+1)$-dimensional manifolds, so long as the TQFT is defined using a state-sum formula based on finite combinatorial-algebraic data. In particular, similar dichotomies should be possible for $(3+1)$-dimensional TQFTs based on spherical fusion 2-categories [8] or lattice gauge theories based on finite groups (sometimes called Dijkgraaf-Witten theories) in arbitrary dimension.

3.3 Alternative proof strategies

Building on the previous point, one might try to give an alternative proof of Theorem 1(b) by using the $(3+1)$-dimensional Crane-Yetter TQFT based on the modular fusion category $\mathcal{B}$ [5]. To put it more carefully, it is known that the Reshetikhin-Turaev invariant $\tau_{\mathcal{B}}(M)$ can be computed by choosing a triangulated 4-manifold Y whose boundary is $\partial Y = M$, and computing an appropriate state-sum invariant of Y (similar to the Turaev-Viro invariant of a triangulated 3-manifold) [4]. So one could try to prove a dichotomy for the $(2+1)$-dimensional surgery-invariant case of Reshetikhin-Turaev in a simpler way by instead proving a dichotomy for the $(3+1)$-dimensional triangulation-invariant case of Crane-Yetter. Accomplishing this requires using the fact that given a surgery diagram for a 3-manifold M, one can build a triangulated 4-manifold Y with $\partial Y = M$ in polynomial time (for example, cf. [1]).

In the opposite direction, it is known that for a spherical fusion category $\mathcal{C}$, $|M|_{\mathcal{C}} = \tau_{\mathcal{Z}(\mathcal{C})}(M)$, where $\mathcal{Z}(\mathcal{C})$ is the Drinfeld center of $\mathcal{C}$. If one were able to efficiently convert a triangulation of a 3-manifold M into a surgery presentation of the same manifold, then part (b) of Theorem 1 (or Conjecture 2) would immediately imply part (a) of the same.

3.4 Towards Conjecture 2

The current gap between Theorem 1 and Conjecture 2 is explained by a rather simple and undesirable property of our proof of the former: our reductions $M \mapsto I_M$ are not "surjective" from 3-manifold encodings to instances of $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$ or $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{B}})$. For example, it would be consistent with our results for there to exist a spherical fusion category $\mathcal{C}$ such that $|M|_{\mathcal{C}}$ is computable in polynomial-time, and yet, the problem $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{C}})$ is still #P-hard (although we consider this unlikely).

To establish an outright dichotomy theorem for TQFT invariants via Cai and Chen's Theorem 3, we would need to arrange our choices of $\mathcal{F}_\mathcal{C}$ and $\mathcal{F}_\mathcal{B}$ with more care so that every instance I of $\#\mathsf{CSP}(\mathcal{F}_\mathcal{C})$ or $\#\mathsf{CSP}(\mathcal{F}_\mathcal{B})$ that is not of the form I_M satisfies two properties: first, it can be identified in polynomial time as an instance that is not of the form I_M, and, second, $Z(I)$ can be computed in polynomial time. This seems difficult to arrange, as it is not clear how to choose the constraint families $\mathcal{F}_\mathcal{C}$ and $\mathcal{F}_\mathcal{B}$ so that instances can "self-report" as not being of the form I_M.

It is instructive to compare TQFT invariants with "holant problems" as defined in [3] and inspired by the "holographic reductions" of [22]. Holant problems are a kind of generalization of counting CSPs that impose more structure on the way in which the individual functions comprising an instance are "wired together." Intuitively, an instance of $\#\mathsf{CSP}(\mathcal{F})$ has no locality constraints on its variables, other than that the constraint functions $f \in \mathcal{F}$ have bounded arity (assuming $\mathcal{F}$ is finite). An instance of a holant problem, on the other hand, has a set of variables that are determined by the *edges of a graph* with constraint functions assigned to the *vertices*. The Turaev-Viro-Barrett-Westbury invariant of closed 3-manifolds determined by a spherical fusion category $\mathcal{C}$ can be seen as generalization of this idea, with variables assigned to both the edges *and* faces of a 3-dimensional triangulation, and constraints assigned to the tetrahedra. We expect it should be possible to formulate TVBW invariants of triangulated 3-manifolds directly as instances of holant problems using a similar construction as in the proof of Theorem 1(a). Of course, even if one could achieve this, such a reduction from TVBW invariants to holant problems would – *a priori* – suffer in the same way as our current reduction to $\#\mathsf{CSP}(\mathcal{F}_\mathcal{C})$. Thus, it seems likely that proving Conjecture 2 will require substantially new ideas. Nevertheless, it appears that there could be much to gain by attempting to import what has been learned about holant problem dichotomies to TQFT invariants of 3-manifolds.

References

1 Rhuaidi Antonio Burke. Practical Software for Triangulating and Simplifying 4-Manifolds. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:23, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2024.29.

2 Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. *Journal of the ACM (JACM)*, 64(3):1–39, 2017.

3 Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Computational complexity of holant problems. *SIAM Journal on Computing*, 40(4):1101–1132, 2011.

4 Louis Crane and Louis H Kauffman. Evaluating the crane-yetter invariant. *Quantum topology*, 3:131–8, 1993.

5 Louis Crane and David N Yetter. A categorical construction of 4d tqfts. *arXiv preprint hep-th/9301062*, 1993.

6 Shawn X Cui and Zhenghan Wang. Universal quantum computation with metaplectic anyons. *Journal of Mathematical Physics*, 56(3), 2015.

7 Colleen Delaney, Clément Maria, and Eric Samperton. An Algorithm for Tambara-Yamagami Quantum Invariants of 3-Manifolds, Parameterized by the First Betti Number. In Oswin Aichholzer and Haitao Wang, editors, *41st International Symposium on Computational Geometry (SoCG 2025)*, volume 332 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:15, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2025.38.

8 Christopher L Douglas and David J Reutter. Fusion 2-categories and a state-sum invariant for 4-manifolds. *arXiv preprint arXiv:1812.11933*, 2018.

9 Michael H Freedman, Alexei Kitaev, and Zhenghan Wang. Simulation of topological field theories by quantum computers. *Communications in Mathematical Physics*, 227:587–603, 2002.

10 Michael H Freedman, Michael Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, 227:605–622, 2002.

11 Robion Kirby and Paul Melvin. Local surgery formulas for quantum invariants and the Arf invariant. *Geometry & Topology Monographs*, 7, 2004.

12 Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006.

13 Greg Kuperberg. How hard is it to approximate the Jones polynomial? *Theory of Computing*, 11(6):183–219, 2015.

14 Greg Kuperberg and Eric Samperton. Computational complexity and 3–manifolds and zombies. *Geometry & Topology*, 22(6):3623–3670, 2018.

15 Greg Kuperberg and Eric Samperton. Coloring invariants of knots and links are often intractable. *Algebraic & Geometric Topology*, 21(3):1479–1510, 2021.

16 Richard E Ladner. On the structure of polynomial time reducibility. *Journal of the ACM (JACM)*, 22(1):155–171, 1975.

17 Deepak Naidu and Eric C Rowell. A finiteness property for braided fusion categories. *Algebras and representation theory*, 14:837–855, 2011.

18 David Reutter. Uniqueness of unitary structure for unitarizable fusion categories. *Comm. Math. Phys.*, 397(1):37–52, 2023. doi:10.1007/s00220-022-04425-7.

19 Eric Rowell and Zhenghan Wang. Mathematics of topological quantum computing. *Bulletin of the American Mathematical Society*, 55(2):183–238, 2018.

20 Eric Samperton. Topological quantum computation is hyperbolic. *Communications in Mathematical Physics*, 402(1):79–96, 2023.

21 Thomas J Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226, 1978.

22 Leslie G Valiant. Holographic algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008.

23 Zhenghan Wang. *Topological quantum computation*, volume 112 of *CBMS Regional Conference Series in Mathematics*. Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2010. doi:10.1090/cbms/112.

A Proof of Theorem 1(b)

Proof. Let $\mathcal{B}$ be a modular fusion category over $\mathbb{C}$. As in our proof of part(a) of Theorem 1, it suffices to build a domain $D_{\mathcal{B}}$ and weighted constraint set $\mathcal{F}_{\mathcal{B}}$ for which there is a polynomial time algorithm to encode a surgery diagram for a 3-manifold M into an instance I_M of $\#\mathsf{CSP}(\mathcal{F}_{\mathcal{B}})$ such that

$$Z(I_M) = \tau_{\mathcal{B}}(M)$$

Let us recall the formula for $\tau_{\mathcal{B}}(M)$:

$$\tau_{\mathcal{B}}(M) = p_-^{\frac{\sigma(L)-m-1}{2}} p_+^{\frac{-\sigma(L)-m-1}{2}} \sum_{\mathrm{col}:\{1,\ldots,m\}\to\mathrm{Irr}(\mathcal{B})} \left(\prod_{j=1}^{m} \dim(\mathrm{col}(j))\right) |L^{\mathrm{col}}|$$

where our notation is as follows:

- L is the surgery link diagram that defines the 3-manifold M. L consists of m components (labeled $1, 2, \ldots, m$) and, for convenience is endowed with the blackboard framing.[3] We will also assume, for convenience, that L is in "standard plat position." This means that all of the local minima of the diagram (which, recall, is a picture in the xy-plane) occur below all crossings, all local maxima of the diagram occur above all crossings, and the sets of cups and caps have no "nesting".[4] This means L is entirely determined by a braid word. See Figure 8.
- $\sigma(L)$ is the signature of L: the number of positive eigenvalues of the linking matrix minus the number of negative eigenvalues. (This can be computed in polynomial time from L.)
- $p_\pm = \sum_{i \in \mathrm{Irr}(\mathcal{B})} \theta_i^\pm (\dim(i))^2$ are the Gauss sums of $\mathcal{B}$.
- $|L^{\mathrm{col}}|$ is the evaluation of the colored ribbon graph defined by coloring the components of L by a function $\mathrm{col} : \{1, \ldots, m\} \to \mathrm{Irr}(\mathcal{B})$.

We now define

$$D_{\mathcal{B}} \overset{\mathrm{def}}{=} \mathrm{Irr}(\mathcal{B}) \sqcup \mathcal{N} \sqcup \{*\}$$

where $\mathcal{N}$ is the set of labels of the trivalent Hom space as in the proof of Theorem 1(a). As hinted at in Lemma 4, we need $\mathbb{C}$-valued constraint functions on the domain $D_{\mathcal{B}}$ that implement bubble pops, tadpole removals, F-moves, *etc.* We define them as follows.

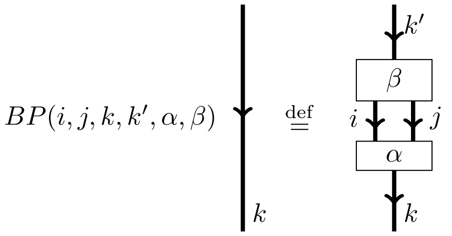

where $i, j, k, k' \in \mathrm{Irr}(\mathcal{B})$, $\alpha \in \{1, \ldots, N_{ij}^k\}$, and $\beta \in \{1, \ldots, N_{ij}^{k'}\}$. These implement the bubble pop.

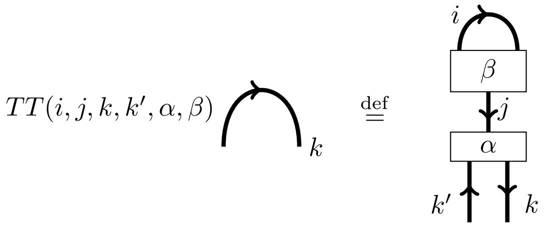

where $i, j, k, k' \in \mathrm{Irr}(\mathcal{B})$, $\alpha \in \{1, \ldots, N_{k'^*k}^j\}$, $\beta \in \{1, \ldots, N_{i^*i}^j\}$. These implement tadpole trims. There is an upside-down version of the tadpole trim, which we denote $\widetilde{TT}$.

[3] To justify this convenience, simply apply a Reidemeister move of type 1 to each the components of L so that the blackboard framing agrees with the desired integral surgery coefficients. This can be done in polynomial time because we encode the surgery coefficents in unary.

[4] This convenience can be justified by the fact that any link diagram can be put in standard plat position in polynomial time by simply applying a sequence of Reidemeister 2 moves.

We define functions VS^{left} by requiring

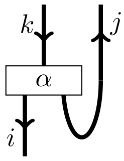

$$= \sum_{\beta\in\{1,\dots,N^i_{kj^*}\}} VS^{\text{left}}(i,j,k,\alpha,\beta)$$

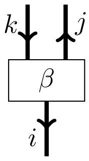

where $i,j,k \in \operatorname{Irr}(\mathcal{B})$ and $\alpha \in \{1,\dots,N^k_{ij}\}$. These implement the "left" vertex spin. We similarly define VS^{right} coefficients to implement the right vertex spin. There are upside-down versions of these vertex spirals. We denote them by $\widetilde{VS}^{\text{left}}$ and $\widetilde{VS}^{\text{right}}$, respectively.

We need the F-matrices, which have coefficients F^+ that satisfy the equation

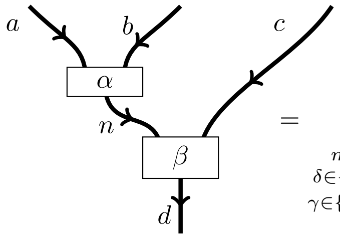

$$= \sum_{\substack{m\in\operatorname{Irr}(\mathcal{B})\\ \delta\in\{1,\dots,N^m_{bc}\}\\ \gamma\in\{1,\dots,N^d_{am}\}}} F^+(a,b,c,d,m,n,\alpha,\beta,\delta,\gamma)$$

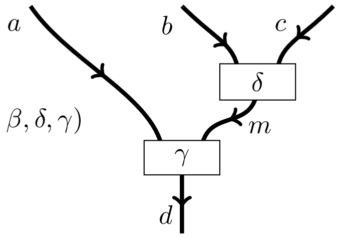

where $a,b,c,d,m \in \operatorname{Irr}(\mathcal{B})$, $\alpha \in \{1,\dots,N^n_{ab}\}$, and $\beta \in \{1,\dots,N^d_{mc}\}$. The inverse F-matrix has coefficients that we denote by F^-. We also need to include the matrix coefficients implementing the upside-down version of this picture. We call these $G^\pm$, respectively.

We then extend the above to functions on our domain:

$$F^\pm(x_1,x_2,\dots,x_{10}) \overset{\text{def}}{=} \begin{cases} F^\pm(x_1,x_2,\dots,x_{10}) & \text{if } x_1,\dots,x_6 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_7,\dots,x_{10} \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$G^\pm(x_1,x_2,\dots,x_{10}) \overset{\text{def}}{=} \begin{cases} G^\pm(x_1,x_2,\dots,x_{10}) & \text{if } x_1,\dots,x_6 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_7,\dots,x_{10} \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$BP(x_1,x_2,\dots,x_6) \overset{\text{def}}{=} \begin{cases} BP(x_1,x_2,\dots,x_6) & \text{if } x_1,x_2,x_3,x_4 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_5,x_6 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$TT(x_1,x_2,\dots,x_6) \overset{\text{def}}{=} \begin{cases} TT(x_1,x_2,\dots,x_6) & \text{if } x_1,x_2,x_3,x_4 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_5,x_6 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$\widetilde{TT}(x_1,x_2,\dots,x_6) \overset{\text{def}}{=} \begin{cases} \widetilde{TT}(x_1,x_2,\dots,x_6) & \text{if } x_1,x_2,x_3,x_4 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_5,x_6 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$VS^*(x_1,x_2,x_3,x_4,x_5) \overset{\text{def}}{=} \begin{cases} VS^*(x_1,x_2,x_3,x_4,x_5) & \text{if } x_1,x_2,x_3 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_4,x_5 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

$$\widetilde{VS}^*(x_1,x_2,x_3,x_4,x_5) \overset{\text{def}}{=} \begin{cases} \widetilde{VS}^*(x_1,x_2,x_3,x_4,x_5) & \text{if } x_1,x_2,x_3 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_4,x_5 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

for $* \in \{\text{left}, \text{right}\}$.

We also need to implement braidings, which are described diagrammatically via R-moves. Recall the definition of the R-symbols: for $i,j,k \in \operatorname{Irr}(\mathcal{B})$ and $\alpha \in \{1,\dots,N^k_{ji}\}$, they satisfy

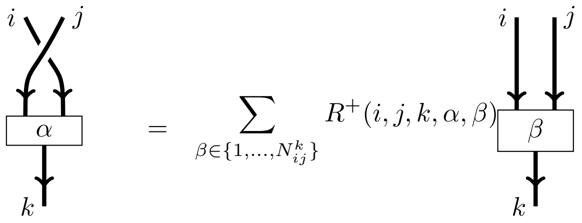

i
j
α
k
=
$\sum_{\beta\in\{1,\dots,N_{ij}^k\}} R^+(i,j,k,\alpha,\beta)$
i
j
β
k

The inverse R-symbol $R^-(i,j,k,\alpha,\beta)$ is similarly defined to describe the inverse braiding. We turn these R-symbols into 5-ary functions on the entire domain $D_{\mathcal{B}}$ in the same trivial way as before, namely

$$R^{\pm}(x_1,x_2,x_3,x_4,x_5) \overset{\text{def}}{=} \begin{cases} R^{\pm}(x_1,x_2,x_3,x_4,x_5) & \text{if } x_1,x_2,x_3 \in \operatorname{Irr}(\mathcal{B}) \text{ and } x_4,x_5 \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$$

Finally, we define 1-ary dimension functions, 1-ary Gauss sum functions (and their inverses), 1-ary dual functions, and 2-ary Kronecker delta functions as follows (respectively):

$$d(x) \overset{\text{def}}{=} \begin{cases} \dim(x) & \text{if } x \in \operatorname{Irr}(\mathcal{B}) \\ 0 & \text{otherwise} \end{cases}$$

$$p_{\pm}^{1/2}(x) \overset{\text{def}}{=} \begin{cases} \left(\sum_{j\in\operatorname{Irr}(\mathcal{B})} \theta_j^{\pm 1} \dim(j)^2\right)^{1/2} & \text{if } x = * \\ 0 & \text{otherwise} \end{cases}$$

$$p_{\pm}^{-1/2}(x) \overset{\text{def}}{=} \begin{cases} (p_{\pm}(x))^{-1} & \text{if } x = * \\ 0 & \text{otherwise} \end{cases}$$

$$\delta(x_1,x_2) \overset{\text{def}}{=} \begin{cases} \delta_{x_1,x_2} & \text{if } x_1,x_2 \in \operatorname{Irr}(\mathcal{B}) \\ 0 & \text{otherwise} \end{cases}$$

With all of this, we define our weighted constraint family $\mathcal{F}_{\mathcal{B}}$ to be

$$\mathcal{F}_{\mathcal{B}} \overset{\text{def}}{=} \{F^{\pm}, G^{\pm}, BP, TT, \widetilde{TT}, VS^{\text{left}}, VS^{\text{right}}, \widetilde{VS}^{\text{left}}, \widetilde{VS}^{\text{right}}, R^{\pm}, d, p_{\pm}^{1/2}, p_{\pm}^{-1/2}, \delta\}.$$

We reiterate that all of this is computed independently of M, and can simply be considered as part of what it means to "have the data" of $\mathcal{B}$.

We conclude our proof by describing how to encode a surgery presentation of M into an instance I_M of #CSP$(\mathcal{F}_{\mathcal{B}})$. In addition to the conveniences described above, we assume that the surgery link diagram L is oriented, embedded in $\mathbb{R}\times[-1,2]$, and is given by plat closure of a braid word $b_1b_2\cdots b_n$ so that each crossing corresponding to b_i lies in $\mathbb{R}\times(\frac{i-1}{n},\frac{i}{n})$ and the only maxima or minima lie in $\mathbb{R}\times([-1,0)\cup(1,2])$.

In order to describe the variables that will be involved in the instance I_M we want, we first describe a polynomial time algorithm to replace L with a planar trivalent graph $\Gamma_L \subset S^2$:

1. At each crossing b_i, insert trivalent vertices to resolve the identity (see Figure 7) so that there is a vertex directly adjacent to the crossing, resulting in a crossed trivalent graph.
2. Perform an R-move for each crossing, resulting in a trivalent graph in $\mathbb{R}^2 \subset S^2$ (after potentially scaling so the vertices lie in $\mathbb{R}\times\mathbb{Z}$).

We can now identify the tuple of variables (valued in the domain $D_{\mathcal{B}}$) that will be involved in our desired instance I_M. Recall that m is the number of components of L, $\sigma(L)$ is the signature, and n is the number of crossings. Let $\Gamma_0 = \Gamma_L, \Gamma_1, \dots, \Gamma_l = \emptyset$ be the sequence of graphs provided by Lemma 4, and define N_v to be the sum of the number of vertices in all of these graphs. Similarly, define N_e to be the sum of the number of edges in the graphs. Then define a tuple of variables associated to the instance I_M by

$$\mathbf{x}_M \overset{\text{def}}{=} (x_0, x_1, \ldots, x_m, y_1, \ldots, y_{|\sigma(L)|}, z_1, \ldots, z_n, u_1, \ldots, u_{N_v}, w_1, \ldots, w_{N_e}).$$

To determine the functions involved in I_M, we simply need to keep appropriate account of the sequence of diagrammatic operations involved in taking L to Γ_L and then taking Γ_L to $\emptyset$. For each diagrammatic operation in the algorithms above, we will define a list of functions in $\mathcal{F}_{\mathcal{B}}$ which account for the contribution of those operations to $\tau_{\mathcal{B}}(M)$.

The first operations are those involved in step (1) of the process of taking L to Γ_L, and involve the insertion of "resolutions of the identity" at every crossing of L. To account for these operations, we define a set Init that is a list of Kronecker delta functions, each of which pairs up edges which used to belong to the same link component before the operation. For example, in Figure 7, the top-right edge of the upper vertex and the bottom-right edge of the lower vertex were a part of the same link component before the operation, so we introduce a Kronecker delta function between the associated variables. The middle edge is free to vary, as in the original algorithm there is a sum over this edge.

We then account for the operations in step (2) of the process of taking L to Γ_L, all of which are R-moves. Each operation happens locally on the diagram, so we define lists R_i for $1 \leq i \leq n$ that are given by $R^{\pm}(w_{i_1}, w_{i_2}, w_{i_3}, u_{i_4}, u_{i_5})$, where we use + or − depending on the strand that crosses over. The variables $w_{i_1}, w_{i_2}, w_{i_3}$ are the edges in the trivalent vertex in the relevant order, u_{i_4} is the labeling of the vertex before the R-move, and u_{i_5} is the labeling of the vertex after the operation.

The next operations are given using the algorithm of Lemma 4. The algorithm provides an ordered list of operations $O_1, \ldots, O_l$, where upon completion of the final operation O_l, the graph Γ_l is empty. Each operation here is local, so we need only consider the local changes when defining our list. Consider operation O_i in this sequence:

- If O_i is a circle removal, define a list C_i which contains the relevant d function.
- If O_i is a tadpole trim, define a list T_i which contains the relevant TT or $\widetilde{TT}$ function.
- If O_i is a bubble pop, we define a list B_i which contains the relevant BP function.
- If O_i is a vertex spiral, we define a list V_i which contains the relevant VS^{left} or VS^{right} (or their upside-down versions).
- If O_i is an F-move, we define a list F_i which contains the relevant $F^{\pm}$ or $G^{\pm}$ functions.
- if O_i is an orientation reversal where w_{i_1} is the variable associated to the edge before the operation, and w_{i_2} is the variable associated to the edge after the operation, then define a list O_i which contains the Kronecker delta function $\delta(w_{i_1}^*, w_{i_2})$.

For each list, we append Kronecker delta functions δ on all edges or vertices which are held constant before and after performing the given operation. E.g. if the edge associated to w_{22} will be held constant after an F-move, and will then be associated with w_{100} in the trivalent graph associated to after the operation, then we introduce $\delta(w_{22}, w_{100})$ to the list F associated to the operation. If $\sigma(L) \geq 0$, we then define the set I_M by:

$$\begin{aligned} I_M = \{&p_+^{-1/2}(x_0), p_-^{-1/2}(x_0), p_+^{-1/2}(x_1), \ldots, p_+^{-1/2}(x_m), p_-^{-1/2}(x_1), \ldots, p_-^{-1/2}(x_m), \\ &p_-^{1/2}(y_1), \ldots, p_-^{1/2}(y_{|\sigma(L)|}), p_+^{-1/2}(y_1), \ldots, p_+^{-1/2}(y_{|\sigma(L)|}), d(x_1), \ldots d(x_m)\} \\ &\cup \text{Init} \cup R_1 \cup \cdots \cup R_n \cup \bigcup_{p_1, \ldots, p_6} C_{p_1} \cup T_{p_2} \cup B_{p_3} \cup V_{p_4} \cup F_{p_5} \cup O_{p_6} \end{aligned}$$

where $p_1, p_2, \ldots, p_6$ each run over all $1, \ldots, l$ which were defined by the algorithm above. If $\sigma(L) < 0$, then replace the elements $p_-^{1/2}(y_1), \ldots, p_-^{1/2}(y_{|\sigma(L)|})$ with the inverses $p_-^{-1/2}(y_1)$, ..., $p_-^{-1/2}(y_{|\sigma(L)|})$. By virtue of Lemma 4, the construction of I_M can be carried out in polynomial time in the size of M.

We now explain why this choice of instance I_M defines an output which computes the Reshetikhin-Turaev invariant $\tau_{\mathcal{B}}(M)$. The $p_+^{\pm 1/2}$ and $p_-^{\pm 1/2}$ functions at the beginning implement the normalizing factor since these functions are constant. This guarantees they may be factored out of the sum as global factors. The d functions implement the product of dimensions we see in the sum.

Notice that if at any point a coloring is not admissible, the term in I_M is 0. It is thus clear that all circle removals, edge orientation reversals, bubble pops, and tadpole trims are correctly implemented, so we just need to check that the F-moves, R-moves, and vertex spirals are correct.

In the standard algorithm, when an F-move occurs, there are three additional variables introduced in the summation: one for the interior edge and two for each interior vertex. These additional variables are introduced here as well, since we are summing over all edges that occur throughout the algorithm, the new edge which is created in the F-move will contribute to the sum, while the rest are held constant due to the inclusion of the Kronecker delta functions.

Similarly, we see that R-moves and vertex spirals are correctly implemented. Note that also there are no extraneous variables in the sum since the algorithm guarantees that the graph will become empty, and our instance is completely determined by how the algorithm behaves. ◀

Quantum SAT Problems with Finite Sets of Projectors Are Complete for a Plethora of Classes

Ricardo Rivera Cardoso
RCQI, Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia

Alex Meiburg
Perimeter Institute for Theoretical Physics, Waterloo, Canada
Institute for Quantum Computing, University of Waterloo, Canada

Daniel Nagaj
RCQI, Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia

Abstract

Previously, all known variants of the Quantum Satisfiability (QSAT) problem – consisting of determining whether a k-local (k-body) Hamiltonian is frustration-free – could be classified as being either in P; or complete for NP, MA, or QMA$_1$. Here, we present new qubit variants of this problem that are complete for BQP$_1$, coRP, QCMA, PI(coRP, NP), PI(BQP$_1$, NP), PI(BQP$_1$, MA), SoPU(coRP, NP), SoPU(BQP$_1$, NP), and SoPU(BQP$_1$, MA). Our result implies that a complete classification of quantum constraint satisfaction problems (QCSPs), analogous to Schaefer's dichotomy theorem for classical CSPs, must either include these 13 classes, or otherwise show that some are equal. Additionally, our result showcases two new types of QSAT problems that can be decided efficiently, as well as the first nontrivial BQP$_1$-complete problem.

We first construct QSAT problems on qudits that are complete for BQP$_1$, coRP, and QCMA. These are made by restricting the finite set of Hamiltonians to consist of elements similar to H_{init}, H_{prop}, and H_{out}, seen in the circuit-to-Hamiltonian transformation. Usually, these are used to demonstrate hardness of QSAT and *Local Hamiltonian* problems, and so our proofs of hardness are simple. The difficulty lies in ensuring that all Hamiltonians generated with these three elements can be decided in their respective classes. For this, we build our Hamiltonian terms with high-dimensional data and clock qudits, ternary logic, and either monogamy of entanglement or specific clock encodings. We then show how to express these problems in terms of qubits, by proving that any QCSP can be reduced to a qubit problem while maintaining the same complexity – something not believed possible classically. The remaining six problems are obtained by considering "sums" and "products" of some of the QSAT problems mentioned here. Before this work, the QSAT problems generated in this way resulted in complete problems for PI and SoPU classes that were trivially equal to NP, MA, or QMA$_1$. We thus commence the study of these new and seemingly nontrivial classes.

While [Meiburg, 2021] first sought to prove completeness for coRP, BQP$_1$, and QCMA, we note that those constructions are flawed. Here, we rework them, provide correct proofs, and obtain improvements on the required qudit dimensionality.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Quantum complexity theory, quantum satisfiability, circuit-to-Hamiltonian, pairwise union of classes, pairwise intersection of classes

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.6

Related Version *Full Version*: `https://arxiv.org/abs/2506.07244` [14]

Funding RRC and DN were supported by projects APVV-22-0570 (DeQHOST) and VEGA No. 2/0164/25 (QUAS). AM was supported by the Natural Sciences and Engineering Research Council of Canada through grant number RGPIN-2019-04198. IQC and the Perimeter Institute are supported in part by the Government of Canada through ISED and the Province of Ontario.

Acknowledgements We thank Alex Grilo and Dorian Rudolph for helpful discussions, and the authors of Ref. [12] for the inspiration of Figure 8 in the full version of this paper.

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 6; pp. 6:1–6:24
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Many of the interesting and puzzling phenomena in many-body physics occurs at the ground state of materials. One way to study quantum systems in this state is through their ground state energy, as this quantity can be used to provide information about physical and chemical properties of the system. It is thus of great interest to calculate or even estimate this quantity. This task is embodied by the k-LOCAL HAMILTONIAN (k-LH) problem. Specifically, given a *k-local* (k-body) Hamiltonian – an operator of the form $H = \sum_i h_i$ where each h_i acts on at most k qubits – and two numbers $a, b \in \mathbb{R}$ with $b - a \geq 1/poly(n)$, this problem consists of distinguishing between the cases where H has an eigenvalue less than a or greater than b. Kitaev [30] showed that k-LH with $k \geq 5$ (and later improved to $k \geq 2$ [28]) is unlikely to be decided efficiently with a classical or quantum computer. In complexity theory terms, k-LH with $k \geq 2$ is QMA-complete.[1]

The LH problem is considered a "weak" quantum constraint satisfaction problem (QCSP) as states with energy less than a do not necessarily minimize the energy of each h_i. For this reason, LH is often compared to MAX-k-SAT instead of the "strong" CSP k-SAT. Due to the immense importance of SAT in classical complexity and other hard sciences, Bravyi [6] defined the QUANTUM k-SAT (k-QSAT) problem. Given a set of k-local projectors (also referred as *clauses* or *constraints*) and a number $b \in \mathbb{R}$, this problem consists of distinguishing between the cases where there exists a state that simultaneously lies in the null space of all projectors, or for all states, the penalty incurred by violations of the constraints is greater than b.[2] Bravyi showed that 2-QSAT on qubits is in P while k-QSAT with $k \geq 4$ (and later improved to $k \geq 3$ [22]) is QMA_1-complete when using the Clifford+T gate set $\mathcal{G}_8 = \{H, \text{CNOT}, T\}$.[3]

Interestingly, these two problems have in common that they are in P for a certain k but appear to become much harder for $k+1$: LH is in P for $k \leq 1$ and becomes QMA-complete for $k > 1$, while QSAT is in P for $k \leq 2$ and $\mathsf{QMA}_1^{\mathcal{G}_8}$-complete for $k > 2$. This is not entirely surprising since the Hamiltonians considered in the problems have no restriction other than their locality, and perhaps the difficulty lies in deciding "unphysical" Hamiltonians. Following this line of thought, others have considered variations of these problems where the h_i are drawn from more realistic and relevant sets that satisfy some property or correspond to a physical model. To name a few, these may be stoquastic [7], commuting [9], fermionic [31], bosonic [41], or from models like the Heisenberg [39] and Bose-Hubbard [15]. In addition, one might also consider placing restrictions on the geometry of the problem [33, 23, 2, 26, 36].

In a landmark result, Cubitt and Montanaro [18] showed that any LH problem where the h_i are drawn from a finite set of at most 2-local qubit Hermitian matrices can be classified as being either in P, NP-complete, StoqMA-complete, or QMA-complete.[4] As decision problems in the latter three classes are not known to be efficiently solvable in either classical or quantum computers, they showed that the only Hamiltonians of this type for which the LH problem can be solved efficiently are those with only 1-local terms. This is significant,

[1] The class QMA can be thought of as the quantum analog of NP, or more accurately MA since the class has probabilistic acceptance and rejection.

[2] Alternatively, this problem can be defined with local Hamiltonians instead of projectors, in which case, the problem is equivalent to determining whether the Hamiltonian is frustration-free.

[3] QMA_1 is the one-sided error variation of QMA with perfect completeness, i.e. instances for which the answer is "yes" (in this case frustration-free Hamiltonians) are accepted with certainty. The notation $\mathcal{G}_8$ stems from Ref. [4] and denotes the Clifford-cyclotomic gate set of degree of 8. The reason why it is necessary to specify the gate set for classes with perfect completeness is discussed in Section A.2.

[4] StoqMA is the class of problems equivalent to estimating the ground state energy of the transverse-field Ising model [8].

as many relevant Hamiltonians in nature can be approximated by 2-local Hamiltonians of this type (e.g. all those supported on Pauli operators like Heisenberg and Ising spin glass models), and it is then likely that estimating their ground state energy efficiently lies outside of reach. Moreover, their result has led to a much larger repertoire of problems from which to construct reductions and potentially show the complexity of other computational problems.

Prior to our work, all known QSAT problems with finite or infinite sets of local interactions could be classified as being either in P, NP-complete, MA-complete, or QMA_1-complete, but this list is not known to be exhaustive in either case. The fact that QSAT has resisted classification can be attributed to two factors. First, is that since most relevant instances of QSAT can be decided classically (2-QSAT is in P), there is a lack of interest to search for a classification of QSAT problems with $k > 2$. This is unlike in the LH problem where most relevant instances were hard (2-LH is QMA-complete), motivating the study of Cubitt and Montanaro. Second, is the fact that QSAT problems are usually complete for classes that are harder to work with as they seem to depend on gate sets. In this work, it is our goal to concretize the implications that such a theorem may have, and hence motivate its study.

1.1 Summary of results

Our main result establishes that the QSAT problem SLCT-QSAT is $\mathsf{BQP}_1^{\mathcal{G}_8}$-complete. However, as the construction and analysis of this problem is contrived, we first show that the simpler and less optimized version of this problem, LCT-QSAT, is also complete for this class.

▶ **Theorem 1.** *The problem* LINEAR-CLOCK-TERNARY-QSAT (LCT-QSAT) *with* 4*-local clauses acting on* 17*-dimensional qudits is* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete.*

An interesting feature of this problem, and one that may be of independent interest, is that this problem makes clever use of the principle of monogamy of entanglement to strongly constrain the structure of input instances, facilitating the task of deciding whether they are frustration-free.[5] Unfortunately, this trick comes at a price of high qudit dimensionality. Our main result shows that by relaxing the constraint on the instance's structure and instead study the instances more closely, we can obtain a similar problem with the same complexity but with reduced qudit dimensionality.

▶ **Theorem 2.** *The problem* SEMILINEAR-CLOCK-TERNARY-QSAT (SLCT-QSAT) *with* 4*-local clauses acting on* 6*-dimensional qudits is* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete.*

Recently, among many other interesting results, Rudolph [37] demonstrated that $\mathsf{BQP}_1^{\mathcal{G}_{2^i}} = \mathsf{BQP}_1^{\mathcal{G}_{2^j}}$ for any $i, j \in \mathbb{N}$. In other words, any problem in BQP_1 using a Clifford-cyclotomic gate set of degree 2^i can be perfectly simulated with one of degree 2^j for all $i, j \in \mathbb{N}$. For us, this then implies that:

▶ **Corollary 3.** *The problems* LCT-QSAT *and* SLCT-QSAT *are* BQP_1*-complete with any gate set* $\mathcal{G}_{2^l}$ *with* $l \in \mathbb{N}$.

Subsequently, by performing slight modifications to the clauses of SLCT-QSAT, we also obtain QCMA-complete and coRP-complete problems:

▶ **Theorem 4.** *The problem* WITNESSED SLCT-QSAT *with* 4*-local clauses acting on* 8*-dimensional qudits is* QCMA*-complete.*

[5] This construction is the most faithful to those considered by Meiburg in Ref. [32].

▶ **Theorem 5.** *The problem* Classical SLCT-QSAT *with* 5*-local clauses acting on* 8*-dimensional qudits is* coRP*-complete.*

Then, using a similar application of monogamy of entanglement as in LCT-QSAT, we demonstrate that we can reduce any QCSP on qudits to another one on qubits.

▶ **Theorem 6** (informal)**.** *Every QCSP $\mathcal{C}$ on qudits is equivalent in difficulty to some other QCSP $\mathcal{C}'$ on qubits.*

▶ **Corollary 7.** *Together, Theorems 2 and 4–6 imply:*

1. SLCT-QSAT$_2$ *is a* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete problem on qubits with* 48*-local clauses.*
2. Witnessed SLCT-QSAT$_2$ *is a* QCMA*-complete problem on qubits with* 48*-local clauses.*
3. Classical SLCT-QSAT$_2$ *is a* coRP*-complete problem on qubits with* 60*-local clauses.*

We refer to these problems by the same name as before, except that we now add a subindex to represent that the problem refers to the qubit version, e.g. SLCT-QSAT$_2$ is the QSAT problem that results from the reduction of SLCT-QSAT.

Finally, there is a notion of *direct product* "$\otimes$" and *direct sum* "$\oplus$" (Definitions 17 and 18) for both CSPs and QCSPs, which we use to show that there are six new QSAT problems that are complete for classes $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$, where A and B are themselves complexity classes. $\mathsf{PI}(A, B)$ stands for the *pairwise intersection of classes* (Definition 11), and $\mathsf{SoPU}(A, B)$ for the *star of pairwise union of classes* (Definition 12). Roughly, these two classes correspond to the sets of problems that can be expressed as the intersection and union (respectively) of a problem in A and a problem in B.[6] We show:

▶ **Theorem 8.** *Let "$\otimes$" and "$\oplus$" denote the* direct product *and* direct sum *for quantum constraint satisfaction problems. Pairwise combinations of the four* QSAT *problems –* 3-SAT*,* Classical SLCT-QSAT$_2$*,* SLCT-QSAT$_2$*, and* Stoquastic 6-SAT *– yield the following complete problems:*

1. Classical SLCT-QSAT$_2$ $\otimes$ 3-SAT *is* $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$*-complete.*
2. Classical SLCT-QSAT$_2$ $\oplus$ 3-SAT *is* $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$*-complete.*
3. SLCT-QSAT$_2$ $\otimes$ 3-SAT *is* $\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$*-complete.*
4. SLCT-QSAT$_2$ $\oplus$ 3-SAT *is* $\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$*-complete.*
5. SLCT-QSAT$_2$ $\otimes$ Stoquastic 6-SAT *is* $\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$*-complete.*
6. SLCT-QSAT$_2$ $\oplus$ Stoquastic 6-SAT *is* $\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$*-complete.*

Finally, given that the QSAT problems in Corollary 7 and Theorem 8 consist of finite sets of projects with $\mathcal{O}(1)$-local qubit clauses, and similarly 2-SAT, 3-SAT, Stoquastic 6-SAT, and 3-QSAT (which are respectively in P, NP-complete, MA-complete and $\mathsf{QMA}_1^{\mathcal{G}_8}$-complete), our results imply that:

▶ **Corollary 9.** *A complete classification theorem for strong QCSPs with $\mathcal{O}(1)$-local clauses acting on qubits must either include at least* 13 *classes, or otherwise indicate that some of these are equal.*

The relationship between the 13 classes mentioned here is shown in Figure 1.

[6] These classes are not to be confused with $A \cap B$ and $A \cup B$. $A \cap B$ corresponds to the set of problems that are in both A and B, while $A \cup B$ corresponds to those that are in either A or B.

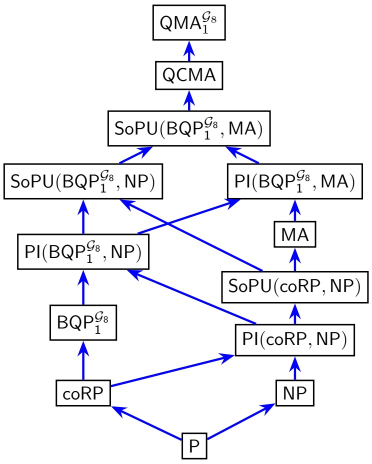
$\mathsf{QMA}_1^{\mathcal{G}_8}$
QCMA
$\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$
$\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$
$\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$
MA
$\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$
$\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$
$\mathsf{BQP}_1^{\mathcal{G}_8}$
$\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$
coRP
NP
P

Figure 1 The classes for which we now have a complete strong QCSP, and their corresponding inclusions. In this work, we show completeness for quantum complexity classes with perfect completeness using the Clifford+T gate set $\mathcal{G}_8 = \{H, T, \text{CNOT}\}$. Rudolph's result [37] further strengthens ours by showing that $\mathsf{BQP}_1^{\mathcal{G}_8} = \mathsf{BQP}_1^{\mathcal{G}_{2^l}}$ for all $l \geq 1$. We discuss some of the inclusions in this figure in Section 2.6 and Section A.2.

2 Contributions

In this section, we summarize the main ideas and proof techniques related to the results presented in Section 1.1. In particular, we detail the main roadblocks in the construction of each QSAT problem, and how we overcome them. The full proofs of the statements here can be found in the full version of the text [14].

Section A covers the notation and background information used here. For the rest of this section, we fix the gate set $\mathcal{G}_8$ and omit the superscript when referring to BQP_1 and QMA_1, except when needed for emphasis.

2.1 BQP_1-complete problem

The goal of the construction is to design a QSAT problem that can encode the computation of any quantum circuit in BQP_1, while keeping all its instances solvable in quantum polynomial-time with perfect completeness and bounded soundness. We define the problem using projectors Π_{init}, $\Pi_{prop,U}$, and Π_{out} similar to P_{init}, $P_{prop,U}$, and P_{out} defined in Equation (6).[7] To see why our projectors must differ from the original ones, consider the QSAT problem built with $\{P_{init}, P_{prop,U}, P_{out}, P_{start}, P_{clock}, P_{end}\}$. Showing that the problem is BQP_1-hard is straightforward, as we can encode the circuit that computes the answer to a BQP_1 problem in a similar way as that shown in Section A.3. This time however, all data particles in the instance should be initialized, instead of having free particles whose role is to accommodate a witness state. The difficulty lies in demonstrating that every instance generated with a polynomial number of these projectors can also be decided in BQP_1. There is a fundamental and a practical limitation for this:

[7] The projectors P_{start}, P_{clock}, and P_{stop} associated with the clock encoding remain unchanged and are integrated into the definitions of Π_{init}, Π_{prop}, and Π_{out}.

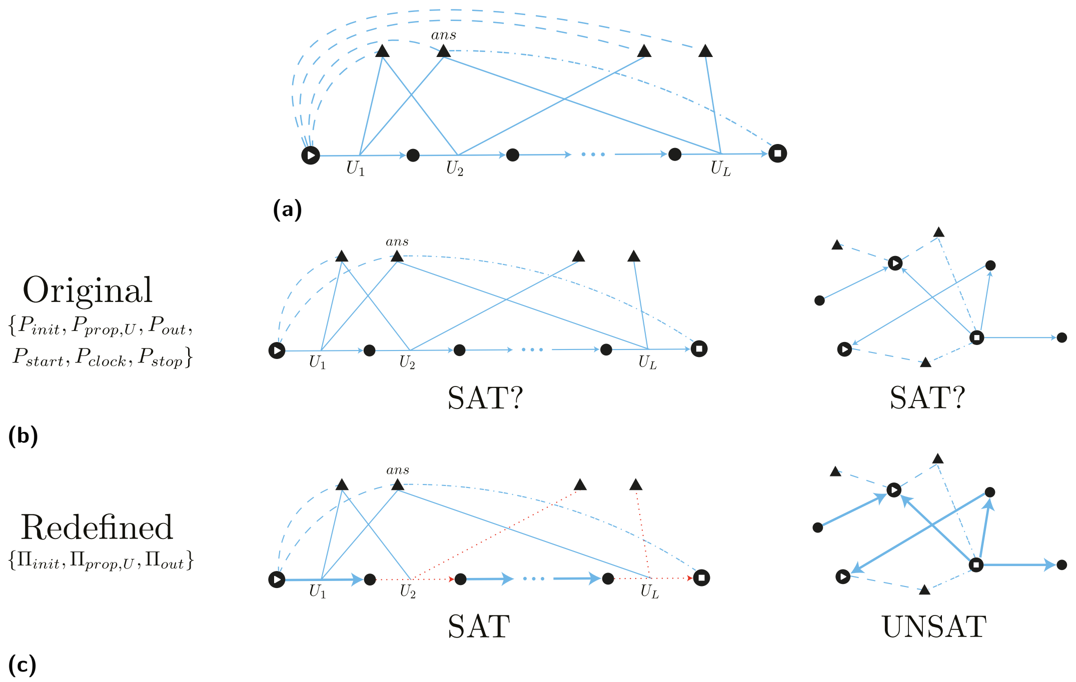
ans
U_1
U_2
U_L
(a)
Original
$\{P_{init}, P_{prop,U}, P_{out}, P_{start}, P_{clock}, P_{stop}\}$
ans
U_1
U_2
U_L
SAT?
SAT?
(b)
Redefined
$\{\Pi_{init}, \Pi_{prop,U}, \Pi_{out}\}$
ans
U_1
U_2
U_L
SAT
UNSAT
(c)

Figure 2 (a) A typical instance that encodes the computation of a BQP_1 circuit $U_L \dots U_1$. The satisfiability of the instance can also be decided in BQP_1. (b) Examples of troublesome instances whose satisfiability is not known to be decidable with a BQP_1 algorithm. (c) The above instances recast with the new set of projectors $\{\Pi_{init}, \Pi_{prop,U}, \Pi_{out}\}$. The bold blue arrows represent the $\Pi_{prop,U}$ clauses which now also indicate the particles should be maximally entangled, and the dotted red arrows those that are connected to undefined logical qudits. With these projectors, their satisfiability can be more easily decided. The left instance is satisfiable due to the undefined clauses, while the one on the right is unsatisfiable, as any potential satisfying state violates monogamy of entanglement. The instance in (a) has the same meaning/satisfiability with either set of projectors.

- Instances which encode the computation of a QMA_1 problem, e.g. the instance in Figure 4 and the left instance in Figure 2b, are valid inputs. This is problematic since it is unknown how to decide these instances in BQP_1 (and doing so would show that $\mathsf{BQP}_1 = \mathsf{QMA}_1$).
- Input instances may form intricate structures complicating the task of deciding if a satisfying state exists, e.g. the right instance in Figure 2b.

We define the projectors Π_{init}, Π_{prop}, and Π_{out} to address these two difficulties (see Figure 2c). Importantly, these projectors do not significantly alter the proof that the problem is BQP_1-hard and can proceed as mentioned. Now, let us briefly discuss how we overcome both difficulties.

Instances like those in Figure 4, which have a proper structure and uninitialized data particles, are prototypical examples of QMA instances. These "free" particles give one the freedom to guess if there exists a state they can be in such that the instance can be satisfied (or equivalently be provided with such a state which we verify). To address this issue, we remove the need to guess a satisfying state by introducing a new *undefined* basis state $|?\rangle$ (making the data particles 3-dimensional), such that setting the free data particles to this state always results in a satisfiable instance. More specifically, we achieve this by defining $\Pi_{prop,U}$ so that if any data particle in the clause is in state $|?\rangle$, the clause is satisfied without needing to apply the associated unitary.[8] Then, for these instances, the satisfying state is given by a truncated version of the history state (without a witness) since the computation is no longer required to elapse past the first $\Pi_{prop,U}$ clause acting on an undefined state. We say the instance is now "trivially satisfiable" as its structure alone suffices to determine its satisfiability.

To determine the satisfiability of intricate instances, the projectors are now also defined to leverage the principle of monogamy of entanglement. Each clock particle is equipped with two 2-dimensional auxiliary subspaces *CA* and *CB* (making them 12-dimensional) and the $\Pi_{prop,U}$ clauses are then defined to require that the *CB* subspace of the predecessor clock particle forms a $|\Phi^+\rangle$ Bell pair with the *CA* subspace of its successor. Then, if a *CA* or *CB* subspace is required to form more than one Bell pair, the principle of monogamy of entanglement states that only one of these clauses can be satisfied, and so the instance is unsatisfiable. Therefore, instances that are not deemed unsatisfiable because of this reason must form one-dimensional chains with a unique "time" direction. Finally, to guarantee that Π_{init} and Π_{out} only act on the ends of the chain, these make use of a new *endpoint* particle consisting of a single two-dimensional space *EC* and require that it also forms a Bell pair with either the *CA* (for Π_{init}) or *CB* (for Π_{out}) subspace of a clock particle. These modifications thus yield a 17-dimensional local Hilbert space: a 3-dimensional data subspace, plus a 2-dimensional endpoint subspace, plus a 12-dimensional clock subspace.

Although these modifications do not get rid off all difficulties, a comprehensive analysis of the resulting instances can be used to demonstrate that a hybrid algorithm can determine the satisfiability status of all input instances. Briefly, the classical part of the algorithm evaluates the structure of the clauses in the instance and concludes whether it is trivially unsatisfiable, trivially satisfiable, or is one requiring the assistance of a quantum subroutine. Trivially unsatisfiable instances are those whose clause arrangement imply one or several clauses cannot be simultaneously satisfied, like those that violate monogamy of entanglement. On the other hand, trivially satisfiable instances are those whose clauses do not create any conflicts but whose structure is simple enough that the satisfying state can be inferred, like those with a proper structure and uninitialized data particles. We show that the only type of instances that are not in either one of these cases, are those like Figure 2a which express the computation of a quantum circuit on initialized ancilla qubits. For these instances, the classical algorithm makes use of a quantum subroutine that executes the quantum circuit expressed by the instance, while simultaneously measuring the eigenvalues of relevant projectors. The measurement outcomes indicate whether the instance should be accepted or rejected.

2.2 Reducing the qudit dimensionality

This section argues that even by removing the projectors that demand successive clock (or endpoint) particles must be entangled with each other, the satisfiability of instances remains the same. Specifically, we argue that the propagation rules, the choice of clock encoding, and the requirement to maintain a consistent clock register state at all times suffice to show that any instance in which the clock particles are not arranged linearly and do not point in the same direction is unsatisfiable. Consequently, there is no longer a need for auxiliary

[8] Although the data particles are 3-dimensional and the unitaries are gates from a set designed to act on qubits, these cause no conflicts as the gates will never act on undefined data particles.

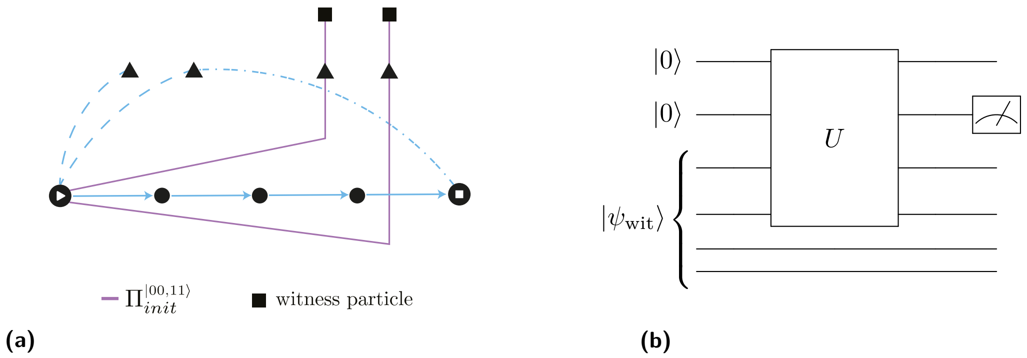

Figure 3 (a) Toy example of an input "quantum" instance with a TACC of length $L = 4$, acting on four logical qudits and two witness qudits. Although not illustrated, the Π_{prop} clauses are assumed to have unitaries $U_1, \ldots, U_4$ which act only on the logical qudits of the instance. These unitaries define a circuit $U = U_4U_3U_2U_1$. (b) Quantum circuit representing the instance on the left.

subspaces or endpoint particles. Together, these results show that while the use of monogamy of entanglement in the construction does facilitate some proofs, it is not crucial for the construction. Removing these elements reduce the local dimension from 17 down to 6.

The main challenge in this construction stems from the weaker constraints that the Π_{init} and Π_{out} clauses set instead of the endpoint particles. In summary, instances with more than a single Π_{init}/Π_{out} pair may now be satisfiable. Part of the proof of this section requires showing that if such sub-instances are potentially satisfiable, they can be further separated into smaller linear instances, each with a single Π_{init}/Π_{out} pair. Each of these smaller pieces is then satisfied by a history state, while the clauses connecting them together (arranged in any shape) can be satisfied trivially. For this reason, we have used the term *semilinear* in the name of the resulting problem.

2.3 QCMA-complete problem

The construction from Section 2.2 can be modified to generate a $\mathsf{QCMA}_1^{\mathcal{G}_8}$-complete problem. Moreover, since $\mathsf{QCMA}_1^{\mathcal{G}_8} = \mathsf{QCMA}$ [27], this results in a QCMA-complete problem. Although there are already many problems known to be complete for this class [19, 42, 21, 24, 40], none of them are strong QCSPs.[9]

In Section 2.1, we argued that the unconstrained or "free" logical qudits of an instance allowed one to guess what state of these qudits (the witness state) might satisfy the instance. This freedom made the problem more difficult and thus not likely contained in BQP_1. For this reason, we introduced the undefined state $|?\rangle$, which simplified these instances and made them decidable in BQP_1. In this construction, we seek to construct a problem that sits in between these two classes so it is QCMA-complete. To accomplish this, we desire to have "free" logical qudits to accommodate a witness state that helps verify whether the instance is satisfiable, but have some sort of constraint to demand that the state is classical.[10]

In practice, creating these constraints is challenging since any superposition of two satisfying states will also satisfy the clause. Instead, we set the constraints such that if there exists a quantum witness state that is part of a satisfying state, there is also a classical

[9] While Ref. [40] also defines a QCMA-complete QSAT problem, it requires additional promise conditions.

[10] We continue using the undefined state for logical qudits whose initial state is not constrained so the difficulty of the problem does not become QMA.

witness state. Loosely, we accomplish this by defining new *witness qudits* and create a new constraint $\Pi^{|00,11\rangle}_{init}$ that connects a witness qudit with a logical one, and require that they are both either $|00\rangle$, $|11\rangle$, or in a superposition of the two.[11] In this way, the two qudits are partially "free" as there is some freedom to their state, yet posses some desired structure. Importantly, we ensure that the witness qudits do not form part of the computation after this initial point.

To see why this leads to the desired effect, consider the toy instance of Figure 3 and suppose there exists a state $|\psi_{\text{wit}}\rangle$ of the four "free" qudits that leads to a satisfying state. Observe that to satisfy the $\Pi^{|00,11\rangle}_{init}$ clauses, this state must be of the form $|\psi_{\text{wit}}\rangle = (\alpha_{00}\,|0000\rangle + \alpha_{01}\,|0011\rangle + \alpha_{10}\,|1100\rangle + \alpha_{11}\,|1111\rangle)_{L_1,W_1,L_2,W_2}$ with $\sum_{b\in\{0,1\}^2}|\alpha_b|^2 = 1$, which we can rewrite in a more convenient form as $|\psi_{\text{wit}}\rangle = \sum_{b\in\{0,1\}^2} \alpha_b\,|b\rangle_L \otimes |b\rangle_W$. Then, a state that satisfies all clauses of the instance is the history state

$$\begin{aligned} |\psi_{hist}\rangle &= \frac{1}{\sqrt{5}}\sum_{t=0}^{4}\left[U_t \dots U_0\,|00\rangle \otimes |\psi_{\text{wit}}\rangle\right] \otimes |\underbrace{d\dots d}_{t}\,a_t\,\underbrace{r\dots r}_{4-t}\rangle \\ &= \frac{1}{\sqrt{5}}\sum_{t=0}^{4}\sum_{b\in\{0,1\}^2} \alpha_b\,|\xi_b^t\rangle \otimes |b\rangle_W \otimes |\underbrace{d\dots d}_{t}\,a_t\,\underbrace{r\dots r}_{4-t}\rangle\,, \end{aligned} \tag{1}$$

where $|\xi_b^t\rangle := U_t \dots U_0\,|00\rangle \otimes |b\rangle_L$. Now, let us argue that there is also a classical witness that leads to a satisfying history state. First, observe that any basis state $|b\rangle_L \otimes |b\rangle_W$ with $|\alpha_b| > 0$ from the decomposition of the witness satisfies the $\Pi^{|00,11\rangle}_{init}$ clauses. Consequently, the history state above but with initial state $|00\rangle \otimes |b\rangle_L \otimes |b\rangle_W$ satisfies the $\Pi^{|0\rangle}_{init}$, $\Pi^{|00,11\rangle}_{init}$, and Π_{prop} clauses of the instance. Finally, to show that this state also satisfies the Π_{out} clause, recall that this clause is satisfied if at time $t = 4$, the probability that the second qubit yields outcome "1" when measured is 1. As shown in Equation (3), this probability can be written as

$$\Pr(\text{outcome } 1) = \sum_{b,b'\in\{0,1\}^2} \alpha_b\alpha_{b'}^*\,\langle\xi_{b'}^4| \otimes \langle b'|\,\Pi^{(1)}\,|\xi_b^4\rangle \otimes |b\rangle = \sum_{b\in\{0,1\}^2} |\alpha_b|^2\,\langle\xi_b^4|\,\Pi^{(1)}\,|\xi_b^4\rangle$$

where $\Pi^{(1)} := |1\rangle\langle 1|_2 \otimes I_{rest}$, and in the last equality we observed that $\langle b'|b\rangle = \delta_{b,b'}$. Then, by the assumption that the instance is satisfiable, it must be that $\langle\xi_b^4|\,\Pi^{(1)}\,|\xi_b^4\rangle = 1$ for all basis states $|b\rangle$ with $|\alpha_b| > 0$. This can also be understood as the probability that at the end of the circuit, the second qubit yields outcome "1" when the witness is the basis state $|b\rangle \otimes |b\rangle$. Therefore, the history state of Equation (1) with classical witness $|\psi_{\text{wit}}\rangle = |b\rangle_L \otimes |b\rangle_W$ also satisfies all clauses of the instance.

The remaining parts of the proof require showing that all new possible qudit connections with new $\Pi^{|00,11\rangle}_{init}$ clause can still be handled, as well as demonstrate perfect completeness and bounded soundness of the hybrid algorithm. For the latter, the majority of the arguments from the BQP_1 construction also directly apply here.

2.4 coRP-complete problem

In Section 2.1, we mentioned that the satisfiability of some instances is decided through a quantum circuit. In particular, this circuit was used to verify the satisfiability of the simultaneous Π_{prop} clauses and final Π_{out} clauses. For the latter, the circuit executed

[11] The local Hilbert space is then 8-dimensional, as it composed of a 3-dimensional data subspace, a 3-dimensional clock subspace, and a 2-dimensional witness subspace.

the quantum circuit $U_T \dots U_1$ on input $|0\rangle^{\otimes q}$, measured some of the qubits, and accepted or rejected depending on the measurement outcomes. Intuitively, to generate the coRP-complete problem, we would like to replace the universal quantum circuit by a universal classical *reversible* circuit $R = R_T \dots R_1$ (reversibility is needed since the best potentially satisfying state is still a *quantum* history state) and introduce randomness into the instance by initializing p ancilla qubits to $|+\rangle$. Then, for these new sub-instances, we could analogously verify the Π_{out} clauses by sampling a bitstring $b \in \{0,1\}^p$, evaluating the circuit R on input $(0^q, b)$ and deciding on its satisfiability based on the final state of the bits. While this idea is close to the actual construction, for reasons mentioned in the full version of the text, it is not sufficient to decide all instances.

For the construction to work, we also incorporate elements of the QCMA problem from the previous section. Namely, we modify the $\Pi_{init}^{|00,11\rangle}$ clause so it initializes both a witness (now referred to as *auxiliary* qudit as we remove the freedom) and a logical qudit to the maximally entangled state $|\Phi^+\rangle$. This new clause is denoted $\Pi_{init}^{|\Phi^+\rangle}$.

Again using the toy example of Figure 3 (replacing the $\Pi_{init}^{|00,11\rangle}$ clauses by $\Pi_{init}^{|\Phi^+\rangle}$ clauses and the unitaries U_i by reversible classical gates R_i), let us illustrate how this construction allows us to verify the satisfiability of Π_{out} clauses. If the instance is satisfiable, the satisfying state must be the history state

$$\begin{aligned}|\psi_{hist}\rangle &= \frac{1}{\sqrt{5}} \sum_{t=0}^{4} \left[R_t \dots R_0 |00\rangle \otimes |\Phi^+\rangle^{\otimes 2}\right] \otimes |\underbrace{d \dots d}_{t} a_t \underbrace{r \dots r}_{4-t}\rangle \\ &= \frac{1}{\sqrt{5}} \sum_{t=0}^{4} \sum_{b \in \{0,1\}^2} \frac{1}{2} |\xi_b^t\rangle \otimes |b\rangle_{Aux} \otimes |\underbrace{d \dots d}_{t} a_t \underbrace{r \dots r}_{4-t}\rangle ,\end{aligned}$$

where in the second line we observed that $|\Phi^+\rangle^{\otimes p} = 2^{-\frac{p}{2}} \sum_{b \in \{0,1\}^p} |b\rangle_L \otimes |b\rangle_{Aux}$ for any $p \in \mathbb{N}$, and defined $|\xi_b^t\rangle := R_t \dots R_1 |00\rangle \otimes |b\rangle_L$. The Π_{out} clause is satisfied if at time $t = 4$, the probability that the second qubit yields outcome "1" when measured is 1. This probability is given by

$$\Pr(\text{outcome } 1) = \frac{1}{4} \sum_{b,b' \in \{0,1\}^2} \langle \xi_{b'}^4 | \otimes \langle b' | \Pi^{(1)} |\xi_b^4\rangle \otimes |b\rangle = \frac{1}{4} \sum_{b \in \{0,1\}^2} \langle \xi_b^4 | \Pi^{(1)} |\xi_b^4\rangle ,$$

from where it is evident that if the instance is satisfiable, $\langle \xi_b^4 | \Pi^{(1)} |\xi_b^4\rangle = 1$ for all $b \in \{0,1\}^2$. Hence, it is possible to verify the Π_{out} clause by sampling one of the strings b, running circuit R on input $(0^2, b)$, and measuring the state of the second qubit.

Another important consideration is that the classical reversible gate set must be chosen with care. Although not covered in this version of the paper, we usually desire that $\mathcal{G}$ is a gate set such that all gates in the set change the basis states upon application, and so $V(\Pi_i)$ of Equation (2) can be implemented perfectly with gates from this set. Here, only the first property is relevant. We choose $\mathcal{G} = \{X, (X \otimes X \otimes X)\text{Toffoli}\}$, which clearly satisfies this property and is also a universal gate set for reversible classical computation. As a consequence, the QSAT problem of this section has 5-local clauses since the Π_{prop} clauses may use a Toffoli. This is the best locality we can achieve as it is also well known that any universal gate set for reversible quantum computation must include a 3-bit gate.

2.5 Universality of qubits for QCSPs

In previous sections, we showed that there are QSAT problems acting on qudits that are complete $\mathsf{BQP}_1^{\mathcal{G}_8}$, QCMA, and coRP. Here, we refine these statements and show that there are QSAT problems on qubits (albeit with higher locality) that are also complete for these classes. To achieve this, we show that any QCSP on qudits can be reduced to another QSAT problem on qubits using little computational power. We note that this section, apart from some changes in the exposition, stems directly from Ref. [32].

At first glance, this statement may seem trivial as operations on qubits are universal for quantum computation, i.e. we can emulate a d-qudit with a $\lceil \log_2(d) \rceil$ qubits and carry out unitaries on those qubits. For our QSAT problems, it is true that any instance retains its satisfiability status when expressed in terms of qubits. However, it is not clear if all input instances generated with these new qubit clauses are contained within this class. For a successful reduction, we must have both.

For an even more explicit example, let us first represent the basis clock states using qubits as: $|r\rangle := |00\rangle$, $|a\rangle := |01\rangle$, and $|d\rangle := |10\rangle$. The $\Pi_{start} = |r\rangle\langle r|$ clause (defined exactly as P_{init} in Equation (7)) can now be written as $\Pi_{start} = |00\rangle\langle 00| + |11\rangle\langle 11|$, where the last term is to prevent the fourth basis state $|11\rangle$, which did not exist before. The clause $(|00\rangle\langle 00| + |11\rangle\langle 11|)_{1,2} + (|00\rangle\langle 00| + |11\rangle\langle 11|)_{2,3}$ acting on three qubits is now valid and is satisfied by the state $|0_1 0_2 1_3\rangle$. This state, however, presents some ambiguity: either we have $|r\rangle$ on qubits 1 and 2, or $|a\rangle$ on qubits 2 and 3. In general, decomposing all clauses into qubits and considering all input instances that may occur adds a significant level of complexity to the problem, making it difficult to determine if it remains in the same class. Moreover, we remark that this is not only particular to our QSAT problems, but in fact applies to all CSPs and QCSPs defined on qudits or non-Boolean variables! In general, the issue is that we cannot ensure that the new qubit clauses are applied to qubits in a consistent fashion based on its parent qudit problem. For example, a qubit clause might treat a particular qubit as "qubit 1" of a previous d-qudit, while another clause might refer to the same qubit as "qubit 2". Moreover, the qubit clauses could also "mix and match", combining "qubit 1" from one previous d-qudit with "qubit 2" from another d-qudit (as in the example with Π_{start}). Overall, these lead to constraints that were unrealizable in the parent qudit problem.

Our main result of this section shows that with a more clever reduction than directly decomposing a d-qudit into $\lceil \log_2(d) \rceil$ qubits, we can guarantee that a satisfiable/unsatisfiable instance on qubits maps to one on qudits with the same satisfiability status. This is something that is not known to be possible classically! More formally, we show that

▶ **Theorem 10** (Theorem 6; formal). *For any QCSP $\mathcal{C}$ on d-qudits, there is another QCSP $\mathcal{C}'$ on qubits, and AC^0 circuits f and g, such that f reduces $\mathcal{C}$ to $\mathcal{C}'$, and g reduces $\mathcal{C}'$ to $\mathcal{C}$. If $\mathcal{C}$ is k-local, then $\mathcal{C}'$ can be chosen to be $4 \cdot 2^{\lceil \log_2(\lceil \log_2(d) \rceil) \rceil} k$ local (that is, $O(\log(d))$ times larger.)*

The main idea behind the proof is that in the quantum world, we can fix the issues mentioned above by again using monogamy of entanglement to bind together our constituent qubits into ordered, entangled larger systems. Ultimately, each clause in the resulting qubit problem incorporates new projectors that force a particular ordering of qubits, and any two clauses that try to "mix and match", or use the same set of qubits but with different ordering, are necessarily frustrated.

If in Theorem 10 we do not require the reductions to be in AC^0, and instead allow P-reductions, a locality of $4\lceil \log_2(d) \rceil k$ suffices. This is used in Corollary 7.

2.6 Direct sum and direct product problems

There is a notion of *direct sum* (denoted by "$\oplus$") and *direct product* (denoted by "$\otimes$") on CSPs and QCSPs that allow us to define the remaining six complete problems. To be clear, these are operations on languages themselves, not on instances. For example, we can talk about the languages 3-Colorable $\oplus$ 4-SAT and 3-Colorable $\otimes$ 4-SAT. Although this notion appears to be quite natural, we were unable to find many sources discussing such ideas – possibly because the classical theory is not as exciting, for reasons we also discuss. The relevant task here is to demonstrate that sum and product QCSPs inherit completeness properties from their constituents. In this way, we are able to construct QCSPs that are complete for PI and SoPU classes, defined as follows.

▶ **Definition 11** (Pairwise intersection of classes). *If $\mathcal{C}_1$ and $\mathcal{C}_2$ are two complexity classes (any sets of languages), then* $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$ *is the class that denotes the* pairwise intersection *of $\mathcal{C}_1$ and $\mathcal{C}_2$. In other words, it is the class of languages that can be written as the intersection, i.e. the logical AND, of a language in $\mathcal{C}_1$ and a language in $\mathcal{C}_2$.*

▶ **Definition 12** (Star of pairwise unions of classes). *If $\mathcal{C}_1$ and $\mathcal{C}_2$ are two complexity classes, then their* star of pairwise unions*, denoted* $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$*, is a complexity class defined as follows: for each language $L_1 \in \mathcal{C}_1$ and $L_2 \in \mathcal{C}_2$, let d be a fresh symbol that is not in the alphabet of L_1 or L_2. Then, the language $(dL_1|dL_2)^*$ is in* $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$. $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$ *is the closure of all such languages under* L *(logspace reductions).*

Definition 12 merits a brief explanation. For a pair of languages L_1 and L_2, what do the strings in the language $L := (dL_1|dL_2)^*$ look like? Given an input string like $d010011d101101d101001$, it will belong to L if and only if each of the three bitstrings $\{010011, 101101, 101001\}$ belongs to either L_1 or L_2. If C is a complexity class powerful enough to break apart the individual bitstrings from the d-delimited string, as well decide both L_1 and L_2, then $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2) \in C$.

We begin discussing that there are CSPs that are complete for these classes, and then extend this to the quantum setting since the latter follows almost identically.

2.6.1 Direct product of constraint satisfaction problems

To begin, it is useful to recall the precise definition of a CSP. A *constraint satisfaction problem* is a triple (V, D, C), where $V = \{v_1, \ldots, v_n\}$ is a finite set of variables, each taking a value from the domain D. If the domain is $D = \{0, 1\}$, then we have a Boolean CSP, and can be generalized to dits if D is instead $D = \{0, \ldots, d\}$. C is a set of constraints, where each constraint $c \in C$ restricts the values that a subset of the variables may take.

Now, let L_1 and L_2 be two CSPs with domains D_1 and D_2, and allowed constraints C_1 and C_2, respectively.

▶ **Definition 13** (Direct product of CSPs). *Given the CSPs L_1 and L_2, their* direct product *$L_1 \otimes L_2$ is a CSP whose domain is the Cartesian product $D_1 \times D_2$. Each constraint $c_i \in C_1$ (resp. C_2) of locality k leads to a constraint c'_i in $L_1 \otimes L_2$, also of locality k, as follows. A tuple $(v_1, v_2, \ldots, v_k) \in (D_1 \times D_2)^k$, where each entry $v_i = (v_{i,1}, v_{i,2})$, belongs to c'_i if the tuple $(v_{i,1}, \ldots, v_{k,1})$ belongs to c_i. Each constraint in $L_1 \otimes L_2$ arises this way from a constraint in L_1 or L_2.*

The goal of this subsection is to show that when one CSP is complete for a complexity class A, and another is complete for a class B, the product problem is complete for the complexity class $\mathsf{PI}(A, B)$. Formally,

▶ **Theorem 14** (Completeness of direct products for PI). *Let M be a set of functions closed under composition with local functions, and closed under concatenations (i.e. if some f, g : $\Sigma_1^* \to \Sigma_2^*$ are each in M, then $h : x \to f(x)g(x)$ is as well). Let L_1 be a CSP complete under*

M-reductions for a class $\mathcal{C}_1$, and likewise L_2 be complete for $\mathcal{C}_2$. Assume that each of $\mathcal{C}_1$ and $\mathcal{C}_2$ are closed under reductions by local functions, closed under intersections, and contain the language All *of all strings, Σ^*. Then, the direct product $L_1 \otimes L_2$ is complete under M-reductions for* $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$.

The assumptions in this theorem are mild and satisfied even for AC^0-reductions and most complexity classes. In other words, this theorem essentially states that if we have CSPs complete for "reasonable" classes $\mathcal{C}_1$ and $\mathcal{C}_2$, the product CSP is complete for $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$.

2.6.2 Direct sum of constraint satisfaction problems

If direct products let us express (informally) a "two-input logical AND" of two CSPs, then direct sums let us express "unbounded-fanin AND of fanin-2 ORs".

▶ **Definition 15** (Direct sum of CSPs). *Given the CSPs L_1 and L_2, their* direct sum *$L_1 \oplus L_2$ is a CSP whose domain is the disjoint union $D_1 \uplus D_2$. Each constraint in $L_1 \oplus L_2$ is either of the form $c_i \cup \left(D_2^k\right)$, where $c_i \in C_1$ is a constraint of locality k; or it is $c_i \cup \left(D_1^k\right)$ for some $c_i \in C_2$.*

To better understand this definition, consider an instance of $L_1 \oplus L_2$ with a single connected component, and assume that it is satisfiable.[12] The definition of the problem and this assumption imply that any satisfying state must either have all variables set to values from D_1, or all of them must be from D_2. Then, for a general instance of $L_1 \oplus L_2$, solving the problem amounts to identifying all of the connected components, and for each one determine whether it can be satisfied entirely from values of D_1 or D_2. The instance is satisfiable iff all components are as well.

One might expect that, by analogy with the direct product, the sum of CSPs should then be complete for the pairwise union of two classes, $\mathsf{PU}(A, B)$. This would be true if we only had to worry about problems that formed a single connected component, which is not the case. This is why we must define the "star of pairwise unions" as in Definition 12. The goal of this section is to demonstrate this fact formally.

▶ **Theorem 16** (Completeness of direct sums for SoPU). *Let be M a set of functions closed under composition with logspace-computable functions (such as the set of logspace functions themselves,* FL*). Let L_1 be a CSP complete under M-reductions for a class $\mathcal{C}_1$, and likewise L_2 be M-complete for $\mathcal{C}_2$. Assume that each of $\mathcal{C}_1$ and $\mathcal{C}_2$ are closed under M-reductions, and contain the language* None *of no strings, $\emptyset$. Then, the direct sum $L_1 \oplus L_2$ is complete under M-reductions for* $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$.

2.6.3 Quantum sums and products

The constructions above transfer in a very natural way to the quantum setting. Now, instead of domains that are a Cartesian product or disjoint union, the Hilbert spaces are a tensor product or direct sum. The clauses are accordingly built as tensor products and direct sums.

▶ **Definition 17** (Direct product of QCSPs). *Given the QCSPs L_1 and L_2, their* direct product *$L_1 \otimes L_2$ is a QCSP whose domain is the tensor product Hilbert space $D_1 \otimes D_2$. Each operator H_i in L_1 leads to an operator $H_i \otimes I$, a tensor product with the identity, and likewise for L_2.*

[12] A *connected component* of a CSP is a connected component in the graph for that CSP, where the vertices are variables, and there is an edge between variables if they share a constraint.

▶ **Definition 18** (Direct sum of QCSPs). *Given the QCSPs L_1 and L_2, their* direct sum *$L_1 \uplus L_2$ is a QCSP whose domain is the direct sum Hilbert space $D_1 \oplus D_2$. Each operator H_i in L_1 leads to an operator $H_i \oplus 0$, a direct sum with the 0 operator, requiring that a frustration-free state lies in the null space of H_i* or *the right half of the direct sum (or a linear combination). Likewise for operators in L_2.*

These have the same essential properties as the direct product and sum for classical CSPs, where we can produce product and sum instances that are satisfiable iff both (resp. either) of the original instances are satisfiable.

Given the discussion above on languages and strings of symbols, one might think that we must talk about quantum states and concatenations of strings of qubits. This is not the case. The strings of symbols are just the encoding of the constraints, which are classical data even for a QCSP. The only quantum-specific requirements involve checking that tensor products or embeddings of satisfying states yield another satisfying state; and the appropriate converse properties. These follow directly from the definition of tensor products and direct sums.

2.6.4 Basic class properties

Here, we state some basic properties of general $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$ classes.

▶ **Lemma 19.** *If the class B includes the language* All *of all strings, then $A \subseteq \mathsf{PI}(A, B)$. Similarly, if B includes the language* None *of no strings, then $A \subseteq \mathsf{SoPU}(A, B)$.*

▶ **Lemma 20.** PI *and* SoPU *respect the inclusion order of complexity classes. That is, $A \subseteq C$ and $B \subseteq D$ implies $\mathsf{PI}(A, B) \subseteq \mathsf{PI}(C, D)$ and $\mathsf{SoPU}(A, B) \subseteq \mathsf{SoPU}(C, D)$.*

SoPU generally leads to a more powerful class than PI, that is:

▶ **Lemma 21.** *If classes A and B are closed under reductions by local functions, then $\mathsf{PI}(A, B) \subseteq \mathsf{SoPU}(A, B)$.*

This is apparent from the definition of these classes since SoPU is also required to compute the AND of multiple inputs. It is also true that PI and SoPU do not increase the power of classes by combining a class A with something weaker. Formally:

▶ **Lemma 22.** *If A is closed under intersection, and $B \subseteq A$, then $\mathsf{PI}(A, B) \subseteq A$. Moreover, if A is closed under logspace reductions, unions, and delimited concatenation, then $\mathsf{SoPU}(A, B) \subseteq A$.*

2.7 New complete problems

As mentioned previously, while the notion of product and sum of constraint problems seems natural, classical constraint problems do not seem to offer such a rich theory. This is due to the fact that most classes with complete CSPs are contained within each other. Indeed, Allender *et al.*'s refinement of Schaefer's dichotomy theorem states that all Boolean CSPs are either in co-NLOGTIME; or are complete for L, NL, ⊕L, P or NP under AC^0 reductions [3]. With the exception of NL and ⊕L, all possible pairs from this list have an obvious containment relation, so the only nontrivial consequence would be that there exists a CSP, on a domain of size four, that is $\mathsf{PI}(\oplus\mathsf{L}, \mathsf{NL})$-complete under AC^0 reductions. However, under the more common P-reductions, the complexity of these problems becomes either in P or NP-complete.[13] Then,

[13] The same is true for CSPs defined on qudits [43].

since $\mathsf{P} \subseteq \mathsf{NP}$ and these classes meet all properties discussed in Lemma 22, it follows that products or sums of these problems result in complexity classes that are trivially equal to NP. For example, 2-SAT $\oplus$ 3-SAT and 2-SAT $\otimes$ 3-SAT are complete problems for $\mathsf{PI}(\mathsf{P}, \mathsf{NP})$ and $\mathsf{SoPU}(\mathsf{P}, \mathsf{NP})$, but these classes are trivially equal to NP.

This is no longer the case in this work. The seven classes we have discussed so far which have complete CSPs are P, coRP, BQP_1, NP, QCMA, and QMA_1. Importantly, these all have the closure properties discussed in this section so far: union, intersection, logspace reductions, and delimited concatenation; and they all include the trivial problems ALL and NONE. Among these classes, most pairs $\{A, B\}$ have $A \subseteq B$, in which case $\mathsf{PI}(A, B) = \mathsf{SoPU}(A, B) = B$. However, there are three pairs that are not known to contain each other, these are: $\mathsf{coRP} \overset{?}{\subseteq} \mathsf{NP}$, $\mathsf{BQP}_1 \overset{?}{\subseteq} \mathsf{NP}$, and $\mathsf{BQP}_1 \overset{?}{\subseteq} \mathsf{MA}$. Each of these pairs leads to two new classes $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$, that are not obviously equal to some other known class. Together, we obtain six more complexity classes with complete QCSPs.

2.7.1 Relations to other classes

Notably, the pair coRP and NP involves only classical classes, and accordingly there is more theory already developed around them.

From the lemmas stated earlier in this section, one can show that $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{MA}$. In addition to this, we can relate $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$ to the class $\mathsf{DP} := \mathsf{PI}(\mathsf{NP}, \mathsf{coNP})$ studied in Ref. [34]. This class forms the second layer of the *boolean hierarchy* BH, i.e. $\mathsf{DP} = \mathsf{BH}_2$ [13]. Since $\mathsf{coRP} \subseteq \mathsf{coNP}$, Lemma 20 tells us that $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{DP}$. On the other hand, $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ does not obviously lie in the Boolean hierarchy. If the class was a simple pairwise union (instead of the "*star* of pairwise unions"), it would lie in BH_3 – the pairwise union of DP and NP. However, it seems unlikely that $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ falls within this class, as doing so would require showing that one could condense the long list of checks required to decide a SoPU problem down to only two queries. In this line of thought, we know that queries to an NP oracle do not need to depend on each other adaptively, so $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ is contained in $\mathsf{P}^{||\mathsf{NP}} = \mathsf{P}^{\mathsf{NP}[\log]}$, studied in Refs. [11, 25].[14]

These classes are also related to two interesting collapse statements. First, observe that if $\mathsf{P} = \mathsf{RP}$ (derandomization), then $\mathsf{coRP} = \mathsf{P} \subseteq \mathsf{NP}$ and so $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) = \mathsf{SoPU}(\mathsf{NP}, \mathsf{NP}) = \mathsf{NP}$. Moreover, an even weaker version of derandomization where $\mathsf{NP} = \mathsf{MA}$ would also lead to a collapse. Here, since $\mathsf{coRP} \subseteq \mathsf{MA}$, we have $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{MA}, \mathsf{NP}) = \mathsf{SoPU}(\mathsf{NP}, \mathsf{NP}) = \mathsf{NP}$. Second, we see that if $\mathsf{NP} = \mathsf{coNP}$ (concise refutations), then $\mathsf{coRP} \subseteq \mathsf{coNP} = \mathsf{NP}$.

For the PI and SoPU classes that involve BQP_1, NP, or MA, it seems difficult to state other inclusions, besides the fact that they lie above BQP_1 and below QCMA.

3 Discussion and open questions

Perhaps the most interesting points of discussion are the implications of Corollary 9. In the latter case, if a complete classification theorem for QSAT problems shows that there are fewer than 13 classes, this would present exciting implications as equalities between some of these classes tackle many interesting and open questions (see Figure 1). This is true even for adjacent classes. For instance, $\mathsf{P} = \mathsf{coRP}$ would imply that probabilistic algorithms

[14] $\mathsf{P}^{||\mathsf{NP}}$ is the class of problems that can be solved by a P machine with polynomially many nonadaptive NP queries, or alternatively, logarithmically many adaptive queries.

with perfect completeness can be derandomized, and $\mathsf{QCMA} = \mathsf{QMA}_1$ would imply that any quantum-verifiable problem (with perfect completeness) could be verified using a classical witness state. Even for the PI and SoPU classes defined here, we have that if $\mathsf{PI}(A, B) = A$, then $B \subseteq A$. As mentioned in Section 2.7, it is expected through derandomization conjectures that some of these classes are in fact equal to each other. Even if this classification theorem proves any of these conjectures, it would be a great result since such proofs have eluded us for many decades. In the former case of Corollary 9, a classification showing that there are more than 13 classes would be a stark contrast with classical CSPs, which can be completely classified as being either in P or NP-complete [38, 43]. This would highlight the more rich and complex panorama of strong QCSPs, and establish a larger repertoire of problems from which to construct reductions and potentially describe the complexity of other problems.

This last point also raises the question whether there could be other classes with complete QSAT problems. Considering those corresponding to polynomial-time computation and verification, we think that this is unlikely. For example, we have not mentioned complete QCSPs for BPP, BQP, or QMA. Since $\mathsf{coRP} \subseteq \mathsf{BPP}$, $\mathsf{BQP}_1 \subseteq \mathsf{BQP}$, and $\mathsf{QMA}_1 \subseteq \mathsf{QMA}$, there are clearly strong QCSPs in these classes. However, the challenge lies in proving their hardness: as shown in Section A.3, these proofs usually require encoding a probabilistic circuit into an instance of this problem. As is also shown there, perfect completeness is critical for the construction, and thus does not work for a circuit with two-sided error. Adressing this would require a different technique.[15] A positive resolution could arise if these classes admit a scheme that boosts their acceptance probabilities to 1. Jordan *et al.* [27] showed that this was possible for QCMA (demonstrating that $\mathsf{QCMA} = \mathsf{QCMA}_1$), but whether this is possible for BPP, BQP, or QMA remains an open question. Another set of classes we have not considered, are those with no error. Little is known about these classes as they appear to be extremely difficult to work with since the perfect soundness requirement implies that no incorrect instance is ever accepted. Besides classes related to polynomial-time computation and verification, there could be other classes with complete QCSPs. After all, the complexity class landscape is vast.

In Theorems 2 and 5, we describe two new types of QSAT problems that can be solved efficiently with a quantum or probabilistic classical computer. Unfortunately, the projectors used in these problems are artifacts built to achieve these results and do not immediately correspond to QSAT problems of interest, even in the qubit case. Recent developments in the fields of quantum chemistry [5], high-energy physics [35] and nuclear physics [10, 16, 17] have shown that 3- or 4-local Hamiltonians are sometimes necessary to explain emergent physics. The QSAT problems for these Hamiltonians are not immediately tractable as $k \geq 3$, so it would be exciting to determine if these problems, or others, fall within these complexity classes. We hope that having demonstrated that such problems exist, our results inspire others to search for more relevant cases.

Finally, Theorem 8 adds an additional six classes to the set of classes with strong QCSP complete problems. Beyond the inclusions shown in Figure 1, little is known about them. It would thus be interesting to investigate how these classes relate to others, and which other problems fall within them.

[15] Another technique is to reduce an already known hard problem into an instance of the target problem. For the LH problem, this is done via perturbation theory gadgets [28, 39, 18]. However, these gadgets rely on approximations and hence do not preserve perfect completeness.

References

1 Scott Aaronson. On perfect completeness for QMA. *Quantum Inf. Comput.*, 9(1&2):81–89, 2009. doi:10.26421/QIC9.1-2-5.

2 Dorit Aharonov, Oded Kenneth, and Itamar Vigdorovich. On the Complexity of Two Dimensional Commuting Local Hamiltonians. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*, volume 111 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:21, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2018.2.

3 Eric Allender, Michael Bauland, Neil Immerman, Henning Schnoor, and Heribert Vollmer. The complexity of satisfiability problems: Refining schaefer's theorem. *Journal of Computer and System Sciences*, 75(4):245–254, June 2009. doi:10.1016/j.jcss.2008.11.001.

4 Matthew Amy, Andrew N. Glaudell, Shaun Kelso, William Maxwell, Samuel S. Mendelson, and Neil J. Ross. Exact synthesis of multiqubit clifford-cyclotomic circuits. In *Reversible Computation - 16th International Conference, RC 2024, Toruń, Poland, July 4-5, 2024, Proceedings*, volume 14680 of *Lecture Notes in Computer Science*, pages 238–245. Springer, 2024. doi:10.1007/978-3-031-62076-8_15.

5 Alberto Baiardi, Michał Lesiuk, and Markus Reiher. Explicitly correlated electronic structure calculations with transcorrelated matrix product operators. *Journal of Chemical Theory and Computation*, 18(7):4203–4217, 2022. doi:10.1021/acs.jctc.2c00167.

6 Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-sat, 2006. arXiv:quant-ph/0602108.

7 Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. Merlin-arthur games and stoquastic complexity, 2006. arXiv:quant-ph/0611021.

8 Sergey Bravyi and Matthew Hastings. On complexity of the quantum ising model. *Communications in Mathematical Physics*, 349(1):1–45, 2017. doi:10.1007/s00220-016-2787-4.

9 Sergey Bravyi and Mikhail Vyalyi. Commutative version of the local hamiltonian problem and common eigenspace problem. *Quantum Info. Comp.*, 5:187–215, May 2005. doi:10.26421/QIC5.3-2.

10 J. H. Busnaina, Z. Shi, Jesús M. Alcaine-Cuervo, Cindy X. Yang, I. Nsanzineza, E. Rico, and C. M. Wilson. Native three-body interactions in a superconducting lattice gauge quantum simulator, 2025. arXiv:2501.13383.

11 Samuel R. Buss and Louise Hay. On truth-table reducibility to sat. *Information and Computation*, 91(1):86–102, March 1991. doi:10.1016/0890-5401(91)90075-d.

12 Chris Cade, Marten Folkertsma, and Jordi Weggemans. Complexity of the guided local hamiltonian problem: Improved parameters and extension to excited states, 2024. arXiv:2207.10097.

13 Jin-Yi Cai, Thomas Gundermann, Juris Hartmanis, Lane A. Hemachandra, Vivian Sewelson, Klaus Wagner, and Gerd Wechsung. The boolean hierarchy i: Structural properties. *SIAM Journal on Computing*, 17(6):1232–1252, December 1988. doi:10.1137/0217078.

14 Ricardo Rivera Cardoso, Alex Meiburg, and Daniel Nagaj. Quantum sat problems with finite sets of projectors are complete for a plethora of classes, 2025. arXiv:2506.07244.

15 Andrew M. Childs, David Gosset, and Zak Webb. The bose-hubbard model is qma-complete. *Theory Comput.*, 11:491–603, 2015. doi:10.4086/TOC.2015.V011A020.

16 Alexander Y. Chuang, Huan Q. Bui, Arthur Christianen, Yiming Zhang, Yiqi Ni, Denise Ahmed-Braun, Carsten Robens, and Martin W. Zwierlein. Observation of a halo trimer in an ultracold bose-fermi mixture, 2024. arXiv:2411.04820.

17 R. Cruz-Torres, D. Nguyen, F. Hauenstein, A. Schmidt, S. Li, D. Abrams, H. Albataineh, S. Alsalmi, et al. Probing few-body nuclear dynamics via ^{3}H and ^{3}He $(e, e^{'}p)$pn cross-section measurements. *Phys. Rev. Lett.*, 124:212501, 2020. doi:10.1103/PhysRevLett.124.212501.

18 Toby Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016.

19 Sevag Gharibian and Jamie Sikora. Ground state connectivity of local hamiltonians. *ACM Trans. Comput. Theory*, 10(2):8:1–8:28, 2018. doi:10.1145/3186587.

20 Brett Giles and Peter Selinger. Exact synthesis of multiqubit clifford+t circuits. *Phys. Rev. A*, 87:032332, March 2013. doi:10.1103/PhysRevA.87.032332.

21 David Gosset, Jenish C. Mehta, and Thomas Vidick. QCMA hardness of ground space connectivity for commuting Hamiltonians. *Quantum*, 1:16, July 2017. doi:10.22331/q-2017-07-14-16.

22 David Gosset and Daniel Nagaj. Quantum 3-sat is QMA_1-complete. *SIAM Journal on Computing*, 45(3):1080–1128, 2016. doi:10.1137/140957056.

23 Sean Hallgren, Daniel Nagaj, and Sandeep Narayanaswami. The local hamiltonian problem on a line with eight states is qma-complete. *Quantum Inf. Comput.*, 13(9-10):721–750, 2013. doi:10.26421/QIC13.9-10-1.

24 Ying hao Chen. 2-local hamiltonian with low complexity is qcma, 2019. arXiv:1909.03787.

25 Lane A. Hemachandra. The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3):299–322, December 1989. doi:10.1016/0022-0000(89)90025-1.

26 Sandy Irani and Jiaqing Jiang. Commuting local hamiltonian problem on 2d beyond qubits, 2023. arXiv:2309.04910.

27 Stephen P. Jordan, Hirotada Kobayashi, Daniel Nagaj, and Harumichi Nishimura. Achieving perfect completeness in classical-witness quantum merlin-arthur proof systems. *Quantum Info. Comput.*, 12(5-6):461–471, 2012. doi:10.26421/QIC12.5-6-7.

28 Julia Kempe, Alexei Kitaev, and Oded Regev. *The Complexity of the Local Hamiltonian Problem*, pages 372–383. Springer, 2004. doi:10.1007/978-3-540-30538-5_31.

29 A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191, 1997.

30 Alexei Y. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate studies in mathematics*. American Mathematical Society, 2002. URL: https://bookstore.ams.org/gsm-47/.

31 Yi-Kai Liu, Matthias Christandl, and F. Verstraete. Quantum computational complexity of the n-representability problem: Qma complete. *Phys. Rev. Lett.*, 98:110503, 2007. doi:10.1103/PhysRevLett.98.110503.

32 Alex Meiburg. Quantum constraint problems can be complete for BQP, QCMA, and more, 2021. arXiv:2101.08381.

33 Roberto Oliveira and Barbara M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Info. Comput.*, 8(10):900–924, 2008. doi:10.26421/QIC8.10-2.

34 C.H. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, April 1984. doi:10.1016/0022-0000(84)90068-0.

35 Francesco Petiziol, Mahdi Sameti, Stefano Carretta, Sandro Wimberger, and Florian Mintert. Quantum simulation of three-body interactions in weakly driven quantum systems. *Phys. Rev. Lett.*, 126:250504, June 2021. doi:10.1103/PhysRevLett.126.250504.

36 Asad Raza, Jens Eisert, and Alex B. Grilo. Complexity of geometrically local stoquastic hamiltonians, 2024. arXiv:2407.15499.

37 Dorian Rudolph. Towards a universal gateset for QMA_1, 2024. arXiv:2411.02681.

38 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC '78, pages 216–226, New York, NY, USA, 1978. Association for Computing Machinery. doi:10.1145/800133.804350.

39 Norbert Schuch and Frank Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature physics*, 5(10):732–735, 2009. doi:10.1038/nphys1370.

40 Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable Local Hamiltonian Problems with Implications to Heuristic Ansatz State Preparation and the Quantum PCP Conjecture. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, volume 310 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:24, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2024.10.

41 Tzu-Chieh Wei, Michele Mosca, and Ashwin Nayak. Interacting boson problems can be qma hard. *Phys. Rev. Lett.*, 104:040501, 2010. doi:10.1103/PhysRevLett.104.040501.

42 Pawel Wocjan, Dominik Janzing, and Thomas Beth. Two qcma-complete problems. *Quantum Inf. Comput.*, 3(6):635–643, 2003. doi:10.26421/QIC3.6-7.

43 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.

A Notation and background

A.1 Notation

For a bitstring x, let $|x|$ denote the number of bits in x.

A *promise problem* $A = (A_{yes}, A_{no})$ is a computational problem consisting of two non-intersecting sets $A_{yes}, A_{no} \subseteq \{0,1\}^*$ where given an instance $x \in \{0,1\}^*$ (promised to be in one of the two sets), one is tasked to determine if $x \in A_{yes}$ (x is a yes-instance) or $x \in A_{no}$ (x is a no-instance).[16] If $A_{yes} \cup A_{no} = \{0,1\}^*$, then A is called a *language*. For an instance x, we let $n = |x|$ denote the size of x.

For some complexity classes, we specify the gate set used. Here, we use the Clifford-cyclotomic gate sets $\mathcal{G}_m$ defined in Ref. [4]. Specifically, we only consider those that are a power of two. These are: $\mathcal{G}_2 := \{X, \text{CNOT}, \text{Toffoli}, H \otimes H\}$, $\mathcal{G}_4 := \{X, \text{CNOT}, \text{Toffoli}, \zeta_8 H\}$, and for $l \geq 3$, $\mathcal{G}_{2^l} := \{H, \text{CNOT}, T_{2^l}\}$. Here, $T_{2^l} = \text{diag}(1, \zeta_{2^l})$ where $\zeta_{2^l} = e^{2\pi i/2^l}$ is a primitive 2^l-th root of unity.

In all quantum circuits considered here, we let $U_0 = I$ be a dummy unitary used for convenience. The same is true for classical circuits Q and classical reversible circuits R. For circuits that decide computational problems, we let *ans* denote the qubit that when measured provides this decision. We accept the instance if the qubit is measured and yields outcome "1", and reject otherwise. Usually, *ans* is the first ancilla qubit of the circuit.

For a circuit U_n that decides an instance x with $|x| = n$, we denote U_x as the circuit where the instance x is encoded into it and the inputs are only ancilla qubits in the $|0\rangle$ state.

A.2 Classes with perfect completeness

This version of the paper assumes familiarity with basic complexity classes; for a detailed introduction, we refer the reader to the full version of the paper [14]. Here, we only discuss a variation of probabilistic complexity classes with *perfect completeness*.

These are the classes where the acceptance probability of yes-instances is equal to one, and they are one of the two types of classes with *one-sided error*. Although these classes appear to be similar to their two-sided error variation, quantum complexity classes with one-sided error require a more precise treatment as they are not known to be independent

[16] The asterisk over the set is known as the *Kleene star* and is used to represent strings of any finite size.

of the gate set used. Indeed, the Solovay-Kitaev theorem [29] used to resolve this issue for classes with two-sided error only works for approximate equivalence of universal gate sets and not perfect equivalence. Thus, for classes with one-sided error (with some exceptions), one must specify the gate set used by the quantum circuits. This is not the case for classical complexity classes as it is known that every classical circuit using gate set $\mathcal{G}$ can be perfectly simulated by another circuit using a universal gate set $\mathcal{G}'$.

Given this discussion, we can then define one-sided error classes as follows:

▶ **Definition 23** (Classes with perfect completeness). *Let $\mathcal{C}$ be a complexity class with two-sided error. The variant of this class with perfect completeness is defined in a similar way to $\mathcal{C}$ except for the following differences:*

1. *For a promise problem A, the acceptance probability must be exactly 1 when $x \in A_{yes}$.*
2. *If $\mathcal{C}$ is a quantum complexity class, the gate set $\mathcal{G}$ used by the quantum circuits $\{U_n\}$ must be specified.*

The class is generally denoted as $\mathcal{C}_1$, or $\mathcal{C}_1^{\mathcal{G}}$ if it is a quantum complexity class.

This sensibility to the gate set in quantum complexity classes is the reason why, in Theorems 1 and 2, we explicitly state that LCT-QSAT and SLCT-QSAT are complete for BQP_1 with the particular choice of gate set $\mathcal{G}_8$. It also presents other complications. To see this, consider BQP. It is evident that $\mathsf{BQP}_1^{\mathcal{G}} \subseteq \mathsf{BQP}$ for any arbitrary gate set $\mathcal{G}$, and also that $\mathsf{P} \subseteq \mathsf{BQP}$. However, is it true that $\mathsf{P} \subseteq \mathsf{BQP}_1^{\mathcal{G}}$? Fortunately, one can show that for the Clifford+T gate set (i.e. $\mathcal{G}_8$) used in this paper, the class $\mathsf{BQP}_1^{\mathcal{G}_8}$ follows the intuitive containment of classes.

Interestingly, Jordan *et al.* [27] showed that if the circuits that decide a QCMA problem consist of gates with a succinct representation (e.g. $\mathcal{G}_8$), the acceptance probability of yes-instances can be amplified additively to be exactly 1. In other words, they showed that $\mathsf{QCMA} \subseteq \mathsf{QCMA}_1^{\mathcal{G}_8}$, concluding that $\mathsf{QCMA}_1^{\mathcal{G}_8} = \mathsf{QCMA}$. This explains why in Theorem 4 we state that the problem WITNESSED SLCT-QSAT is QCMA-complete. To this day, it remains an open question whether a similar scheme can also work for BQP and QMA. In the case of QMA, it seems this is not the case as one can show that there exists an oracle for which $\mathsf{QMA} \neq \mathsf{QCMA}_1$ [1]. However, a similar claim was made about QCMA and QCMA_1.

A.3 k-QSAT & the Circuit-to-Hamiltonian transformation

Here, we introduce QUANTUM k-SAT (denoted here as k-QSAT) as presented by Gosset and Nagaj in Ref. [22]. We present relevant parts of the proofs showing that k-QSAT is contained in QMA_1 for any constant k, and QMA_1-hard for $k \geq 6$. While Bravyi's [6] original work demonstrates hardness for $k \geq 4$, we choose to present this slightly weaker result for brevity, but also to introduce our clock encoding and notation useful for the rest of this paper.

As we are working to prove the inclusion and hardness of this problem for a class requiring perfect completeness, it is necessary to specify the gate set used by the quantum circuits. For reasons discussed below, we choose $\mathcal{G}_8$. In addition, we also have to be wary that all operations can be performed with perfect accuracy using gates from this set and all measurements are in the computational basis. For this purpose, Gosset and Nagaj introduce the following set of projectors.

▶ **Definition 24** (Perfectly measurable projectors). *Let $\mathcal{P}$ be the set of projectors such that every matrix element in the computational basis is of the form $\frac{1}{4}(a + ib + \sqrt{2}c + i\sqrt{2}d)$ for all $a, b, c, d \in \mathbb{Z}$.*

The (promise) problem k-QSAT can be defined as follows.

▶ **Definition 25** (k-QSAT). *Given an integer n and an instance x consisting of a collection of projectors $\{\Pi_i\} \subset \mathcal{P}$ where each Π_i acts nontrivially on at most k qubits, the problem consists on deciding whether (1) there exists an n-qubit state $|\psi_{sat}\rangle$ such that $\Pi_i|\psi_{sat}\rangle = 0$ for all i, or (2) for every n-qubit state $|\psi\rangle$, $\Sigma_i \langle\psi|\Pi_i|\psi\rangle \geq 1/poly(n)$. We are promised that these are the only two cases. We output "YES" if (1) is true, or "NO" otherwise.*

One can think of this problem as being presented with a list of constraints or *clauses* (the projectors Π_i) and tasked with distinguishing between the following cases: (1) there exists a state that satisfies all constraints (a *satisfying state*), or (2) any possible state induces a violation of the constraints greater than $1/poly(n)$. The promise sets the conditions for classifying instances as either $x \in A_{yes}$ or $x \in A_{no}$.[17]

A.3.1 In QMA_1

Suppose we are presented with a witness state $|\psi_{\text{wit}}\rangle$ and a k-QSAT instance composed of projectors $\{\Pi_i\}$. The quantum algorithm that decides whether this state satisfies all projectors Π_i consists of simply measuring the eigenvalues of all projectors on this state. Then, if all measured eigenvalues are 0, we conclude that all projectors are satisfied by the state and output "YES". Otherwise, we reject.

Specifically, we measure the eigenvalue of a projector Π_i by applying the unitary

$$V(\Pi_i) = \Pi_i \otimes X + (I - \Pi_i) \otimes I, \tag{2}$$

to the witness and an additional ancilla qubit in the state $|0\rangle$, followed by a measurement of the ancilla in the computational basis. Here, X denotes the Pauli-X gate. The probability that $|\psi_{\text{wit}}\rangle$ does not satisfy projector Π_i (obtain outcome "1") is given by

$$p_i = \langle\psi_{\text{wit}}|\, \Pi_i \,|\psi_{\text{wit}}\rangle\,. \tag{3}$$

Defining the acceptance probability as the probability that all measurements produce outcome "0", and assuming $V(\Pi_i)$ can be implemented perfectly with gate set $\mathcal{G}$, one can show that this algorithm meets the completeness and soundness conditions of QMA_1, concluding that k-QSAT is contained in this class.

As mentioned, to support this claim, it is necessary to demonstrate that $V(\Pi_i)$ can be implemented perfectly using gate set $\mathcal{G}_8$. This follows from the fact that the projectors Π_i are from the set $\mathcal{P}$ together with a theorem by Giles and Selinger [20].

A.3.2 QMA_1-hard

Now, we discuss elements of the proof demonstrating that k-QSAT is QMA_1-hard when $k \geq 6$ and for any gate set $\mathcal{G}$ that is universal for quantum computation.

The idea is to demonstrate that any instance x of an arbitrary promise problem in QMA_1 can be transformed or *reduced* in polynomial time into an instance x' of k-QSAT, where the answer to both problems is the same for all instances. Furthermore, we also need to show that all projectors of the resulting k-QSAT instance act on at most 6 qubits.

[17] Without the promise, the problem seems to become harder, as it requires distinguishing between the case where the projectors are satisfiable, and the case where they are not but the violation induced by some states could be exponentially close to zero. Without the promise, the problem is most likely not contained in QMA_1.

Let $U_x = U_L \dots U_1$ with $U_i \in \mathcal{G}$ and $L = poly(n)$ be the QMA_1 verification circuit where given an instance x of a problem $A = (A_{yes}, A_{no})$, U_x decides whether $x \in A_{yes}$ or $x \in A_{no}$. The input to the circuit consists of the p-qubit witness state $|\psi_{\text{wit}}\rangle$, and a q-qubit ancilla register D (referred to as the *data* register) initialized to the state $|0\rangle^{\otimes q}$, where p and q are two polynomials in $n = |x|$. Additionally, let the answer be obtained by measuring the ancilla qubit *ans* in the computational basis. The goal of the reduction is to engineer a set of 6-local projectors such that they are uniquely satisfied by the state encoding the evaluation of the circuit U on $|\phi_0\rangle := |0\rangle^{\otimes q} \otimes |\psi_{\text{wit}}\rangle$ at all steps of the computation. This state is appropriately known as the (computational) *history state* and is given by

$$|\psi_{hist}\rangle := \frac{1}{\sqrt{L+1}} \sum_{t=0}^{L} U_t \dots U_0 \, |\phi_0\rangle_D \otimes |C_t\rangle_C \, . \tag{4}$$

Here, we have introduced a *clock* register C acting on a new (not yet specified) Hilbert space used to keep track of the current step in the computation. This history state can be defined in many ways depending on the implementation of the states $|C_t\rangle$. In this paper, we choose a clock encoding acting on $\mathcal{H}_{clock} = (\mathbb{C}^3)^{\otimes L+1}$, consisting of the *ready* state $|r\rangle$, the *active* state $|a\rangle$, and the *dead* state $|d\rangle$. The clock progresses as

$$\begin{aligned} |C_0\rangle &= |a_0 r_1 r_2 \dots r_L\rangle \, , \\ |C_1\rangle &= |d_0 a_1 r_2 \dots r_L\rangle \, , \\ &\vdots \\ |C_L\rangle &= |d_0 d_1 d_2 \dots a_L\rangle \, . \end{aligned} \tag{5}$$

The projectors that allow us to build the required 6-QSAT instance act on both of these Hilbert spaces and are given by

$$\begin{aligned} P^{(i)}_{init} &:= |1\rangle\langle 1|_i \otimes |a\rangle\langle a|_0 \, , \\ P^{(i)}_{out} &:= |0\rangle\langle 0|_i \otimes |a\rangle\langle a|_L \, , \\ P^{(i)}_{prop,U} &:= \frac{1}{2}\big[I^{\otimes 2} \otimes |ar\rangle\langle ar| + I^{\otimes 2} \otimes |da\rangle\langle da| - U \otimes |da\rangle\langle ar| - U^\dagger \otimes |ar\rangle\langle da| \big], \end{aligned} \tag{6}$$

which receive an index to specify its action on a given particle. Moreover, $P_{prop,U}$ acts on clock qudits $i-1$ and i. Observe that P_{init} and P_{out} act on a single data and clock particle, while $P_{prop,U}$ acts on two data qubits and two clock particles. As each clock particle can be represented by two qubits, albeit a bit wastefully, it is evident that these projectors are at most 6-local (on qubits). Other clock encodings may lead to different locality.[18]

Each projector in Equation (6) penalizes states that do not meet certain requirements. (Initialization) P_{init} requires that when clock particle 0 is in the state $|a\rangle$, data qubit i is initialized to $|0\rangle$. (Computational propagation) $P_{prop,U}$ requires that as clock particles i and $i+1$ transition from $|ar\rangle$ to $|da\rangle$, U is applied to two qubits of the data register. (Readout)

[18] In Ref. [6], Bravyi employs a four-state clock encoding, $2L+1$ clock basis states, and an additional propagation projector. This allows interactions between either two clock particles at a time or one clock particle and two data qubits, resulting in 4-local projectors. However, this comes at a cost of increased clock particle dimensionality.

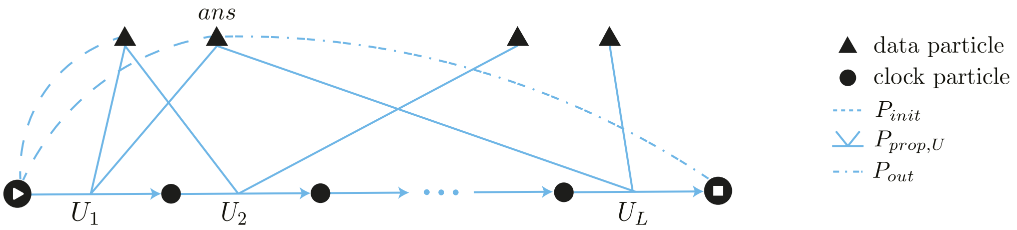

Figure 4 Representation of a 6-QSAT instance which encodes a QMA_1 verification circuit $U = U_L \dots U_1$. For simplicity, we let U act on four data qubits: two ancilla qubits (those present in P_{init} clauses), and two for the witness state (uninitialized ones). The ancilla measured at the end of the computation is labeled *ans*. The leftmost and rightmost clock particles are marked with "start" and "stop" icons, indicating the action of P_{start} and P_{stop} clauses, respectively. The P_{clock} clauses are shown as arrows on top of $P_{prop,U}$ lines, representing the clock progression.

Finally, P_{out} requires that when clock qudit L is in the state $|a\rangle$, data qubit i is in the state $|1\rangle$.[19] Aside from these projectors, one also has to define

$$\begin{aligned} P_{start} &:= |r\rangle\langle r|_0 \,, \\ P_{stop} &:= |d\rangle\langle d|_L \,, \\ P^{(i)}_{clock} &:= |r\rangle\langle r|_i \otimes (I - |r\rangle\langle r|)_{i+1} + |a\rangle\langle a|_i \otimes (I - |r\rangle\langle r|)_{i+1} + |d\rangle\langle d|_i \otimes |r\rangle\langle r|_{i+1} \,, \end{aligned} \tag{7}$$

which are at most 4-local projectors requiring that the clock states have the form described in Equation (5). Furthermore, the six types of projectors of Equations (6) and (7) are of the form given in Definition 24 and are hence projectors from $\mathcal{P}$, as required. Finally, using these projectors, the instance that encodes the verifier circuit $U = U_L \dots U_1$ is given by

$$\begin{aligned} H_{init} &:= \sum_{b \in ancilla} P^{(b)}_{init}, \\ H_{prop} &:= \sum_{t=1}^{L} P^{(t)}_{prop,U_t} \\ H_{out} &:= P^{(ans)}_{out}, \\ H_{clock} &:= P_{start} + P_{stop} + \sum_{c \in C} P^{(c)}_{clock}. \end{aligned}$$

We illustrate this instance in Figure 4. The set of projectors that define this k-QSAT instance are the individual terms of the sum. They are often grouped into positive semi-definite terms as above for historical reasons. Briefly, the H_{init} term requires that all ancilla qubits from register D are initialized to $|0\rangle$, leaving the data qubits for the witness state "free" or uninitialized. H_{prop} defines a clock register of $L+1$ particles and requires that as time progresses from $t-1$ to t, U_t is applied to the data qubits. H_{out} requires that at the end of the computation *ans* is measured to be "1". Finally, H_{clock} requires that we obtain a running clock register progressing as shown in Equation (5). Together, H_{init}, H_{prop}, and H_{clock} require that if there exists a state satisfying all of their projectors, the state must mimic the evaluation of the quantum circuit $U = U_L \dots U_1$ on the state $|\phi_0\rangle$. This is the history state

[19] Unlike Bravyi [6] and Meiburg [32], we define P_{out} so it is satisfied when the logical qubit is in the state $|1\rangle$, and not $|0\rangle$.

of Equation (4) with the clock encoding of Equation (5). Moreover, if the verification circuit U accepts yes-instances with certainty, the history state also satisfies H_{out} and is thus the unique ground state of the 6-local Hamiltonian $H = H_{init} + H_{prop} + H_{out} + H_{clock}$.

This concludes the transformation of the circuit into local Hamiltonians. Completing the proof that 6-QSAT is QMA_1-hard requires showing that, if $x \in A_{yes}$, then x' has a frustration-free ground state, and if $x \in A_{no}$, then the ground state energy of H is not too low. Proving these is beyond the scope of this section.

Uniformity Testing When You Have the Source Code

Clément L. Canonne
The University of Sydney, Australia

Robin Kothari
Google Quantum AI, Santa Barbara, CA, USA

Ryan O'Donnell
Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

We study quantum algorithms for verifying properties of the output probability distribution of a classical or quantum circuit, given access to the source code that generates the distribution. We consider the basic task of uniformity testing, which is to decide if the output distribution is uniform on $[d]$ or ε-far from uniform in total variation distance. More generally, we consider identity testing, which is the task of deciding if the output distribution equals a known hypothesis distribution, or is ε-far from it. For both problems, the previous best known upper bound was $O(\min\{d^{1/3}/\varepsilon^2, d^{1/2}/\varepsilon\})$. Here we improve the upper bound to $O(\min\{d^{1/3}/\varepsilon^{4/3}, d^{1/2}/\varepsilon\})$, which we conjecture is optimal.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Quantum computation theory

Keywords and phrases distribution testing, uniformity testing, quantum algorithms

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.7

Related Version *Full Version*: `https://arxiv.org/abs/2411.04972`

Funding *Clément L. Canonne*: Supported by an ARC DECRA (DE230101329).
Ryan O'Donnell: Supported by ARO grant W911NF2110001 and by a gift from Google Quantum AI.

1 Introduction

For 30 years we have known that quantum computers can solve certain problems significantly faster than any known classical algorithm. Traditionally, most of the research in this area has focused on decision problems (like SAT) or function problems (like Factoring), where for each possible input there is a unique "correct" output. However, we have also found that quantum computers can yield speedups for the task of *sampling* from certain probability distributions. Prominent examples include boson sampling [1] and random circuit sampling [8]. Sampling tasks have seemed more natural for NISQ-era quantum computation, and indeed many of the first candidate experimental demonstrations of quantum advantage have been for sampling problems [6].

One of the downsides of sampling problems is the challenge of *verifying* the output of an algorithm, whether classical or quantum, that claims to sample from a certain distribution. As a simple example, consider a classical or quantum algorithm that implements a supposed hash function with output alphabet $[d] := \{1, \dots, d\}$. The algorithm designer claims that the output distribution of this hash function is uniform on $[d]$. If $\mathbf{p}$ denotes the actual output distribution of the algorithm, and $\mathbf{u}_d$ denotes the uniform distribution on $[d]$, then we would like to test whether $\mathbf{p} = \mathbf{u}_d$, and reject the claim if $\mathbf{p}$ is in fact ε-far from $\mathbf{u}_d$ in total variation distance, meaning $\frac{1}{2}\|\mathbf{p} - \mathbf{u}_d\|_1 > \varepsilon$. (We will also consider other distance measures in this work, since the complexity of the testing task is sensitive to this choice.)

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 7; pp. 7:1–7:20
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This verification task is called “uniformity testing” (in total variation distance) and its complexity is well studied in the classical literature. If we only have access to samples from $\mathbf{p}$, but are not allowed to inspect the algorithm that produces these samples, it is known that $\Theta(d^{1/2}/\varepsilon^2)$ samples are necessary and sufficient to solve this problem; there are various classical algorithms that achieve this bound (starting with that of [28]; see, e.g., [11] for a detailed survey and discussion), and it is also not possible to do better with a quantum algorithm. But what if – as in the examples above – we *do* have access to the algorithm that produces $\mathbf{p}$? Can we improve on this complexity if we have access to the “source code” of the algorithm?

Having the source code

To clarify, the “source code” for a classical randomized sampling algorithm means a randomized circuit (with no input) whose output is one draw from $\mathbf{p}$. More generally, the “source code” for a quantum sampling algorithm means a unitary quantum circuit (with all input qubits fixed to $|0\rangle$) which gives one draw from $\mathbf{p}$ when some of its output bits are measured in the standard basis and the rest are discarded.[1] The simplest way to use the code C for $\mathbf{p}$ is to run it, obtaining one sample. If C has size S, then getting one sample this way has cost S. Another way to use the code C is to deterministically compute all its output probabilities; this gives one perfect information about $\mathbf{p}$, but has cost bound 2^S. But quantum computing has suggested a third way to use the code: “running it in reverse”. For example, Grover’s original algorithm [18] can be seen as distinguishing two possibilities for $\mathbf{p}$ on $[2]$, namely $\mathbf{p}_1 = 0$ or $\mathbf{p}_1 = 1/N$, while using only $O(N^{1/2})$ forwards/backwards executions of C. The total cost here is $O(N^{1/2}) \cdot S$, the same as the cost for $O(N^{1/2})$ samples.

We suggest that the utility of “having the source code” for distribution testing problems remains notably underexplored. Indeed, there is significant room for improvment in the bounds for even the most canonical of all such problems: uniformity testing. Our main theorem is the following:

▶ **Theorem 1.** *There is a computationally efficient quantum algorithm for uniformity testing with the following guarantees: given $\varepsilon \geq 1/\sqrt{d}$, the algorithm makes $O(d^{1/3}/\varepsilon^{4/3})$ uses of “the code” for an unknown distribution $\mathbf{p}$ over $[d]$, and distinguishes with probability at least $.99$ between*

$$(1)\ \mathbf{p} = \mathbf{u}_d, \qquad \textit{and} \qquad (2)\ \mathrm{d_{TV}}(\mathbf{p}, \mathbf{u}_d) > \varepsilon. \tag{1}$$

The main idea behind this theorem is to combine very careful classical probabilistic analysis with a black-box use of Quantum Mean Estimation (QME) [19, 9, 21, 25, 20, 22]; see Section 2 for further discussion. Table 1 below compares our result to prior work on the problem. Table 1 has two columns because it seems that different algorithms are necessary depending on how d and ε relate. (Interestingly, this is not the case in the classical no-source-code model.) Thus combining our new result with that of [24], the best known upper bound becomes $O(\min\{d^{1/3}/\varepsilon^{4/3}, d^{1/2}/\varepsilon\})$. We remark that although [24]’s algorithm/analysis is already simple, we give an alternative simple algorithm and analysis achieving $O(d^{1/2}/\varepsilon)$ in Section A, employing the classical analysis + QME approach used in the proof of our main theorem.

[1] This is sometimes termed the “purified quantum query access model”, and is the most natural and general model. The “quantum string oracle”, referenced later in Table 1, refers to a situation in which one assumes a very specific type of source code for $\mathbf{p}$ (thus making algorithmic tasks easier). See Section 3 for details and [7] for a thorough discussion.

Lower bounds?

As for lower bounds (holding even in the quantum string oracle model): complexity $\Omega(1/\varepsilon)$ is necessary even in the case of constant $d = 2$, following from work of [26]; and, [12] showed a lower bound of $\Omega(d^{1/3})$ even in the case of constant ε, by reduction from the collision problem [2]. For reasons discussed in Section 2, we make the (somewhat bold) conjecture that our new upper bound is in fact tight for all d and ε:

▶ **Conjecture 2.** *Any algorithm that distinguishes $\mathbf{p} = \mathbf{u}_d$ from $\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{u}_d) > \varepsilon$ with success probability at least .99 requires $\Omega(\min\{d^{1/3}/\varepsilon^{4/3}, d^{1/2}/\varepsilon\})$ uses of the code for $\mathbf{p}$. (Moreover, we conjecture this lower bound in the stronger quantum string oracle model.)*

Identity testing

Several prior works in this area have also studied the following natural generalization of uniformity testing: testing identity of the unknown distribution $\mathbf{p}$ to a known hypothesis distribution $\mathbf{q}$. An example application of this might be when $\mathbf{q}$ is a Porter–Thomas-type distribution arising as the ideal output of a random quantum circuit. Luckily, fairly recent work has given a completely generic reduction from *any* fixed identity testing problem to the uniformity testing problem; see [16], or [11, Section 2.2.3]. We can therefore immediately extend our new theorem to the general identity-testing setting:

▶ **Corollary 3.** *There is a computationally efficient quantum algorithm for identity testing to a reference distribution $\mathbf{q}$ over $[d]$ with the following guarantees: The algorithm makes $O(\min(d^{1/3}/\varepsilon^{4/3}, d^{1/2}/\varepsilon))$ uses of "the code" for an unknown distribution $\mathbf{p}$ over $[d]$, and distinguishes with probability at least .99 between*

$$(1)\ \mathbf{p} = \mathbf{q}, \qquad \textit{and} \qquad (2)\ \mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{q}) > \varepsilon. \tag{2}$$

(For completeness, we verify in Section C that the blackbox reduction does indeed carry through in our setting, preserving access to "the code".)

More fine-grained results

In proving our main theorem, we will in fact prove a strictly stronger version, one which is more fine-grained in two ways:

(1) *Tolerance:* Not only does our test accept with high probability when $\mathbf{p} = \mathbf{u}_d$, it also accepts with high probability when $\mathbf{p}$ is sufficiently close to $\mathbf{u}_d$.

Table 1 "Sample" complexity for uniformity testing with respect to total variation distance.

Reference	Large ε regime	Small ε regime	Access model
[28, 4]	$\Theta(d^{1/2}/\varepsilon^2)$		Classical, no source code
[10]	$O(d^{1/3})$ for $\varepsilon = \Theta(1)^*$		Quantum string oracle
[12]	$O(d^{1/3}/\varepsilon^2)$		Quantum string oracle
[15]		$O(d^{1/2}/\varepsilon) \cdot \log(d/\varepsilon)^3 \log\log(d/\varepsilon)$	Source code
[24]		$O(d^{1/2}/\varepsilon)$	Source code
This work	$O(d^{1/3}/\varepsilon^{4/3})$ for $\varepsilon \geq \frac{1}{\sqrt{d}}$		Source code

*The work states a bound of $O(d^{1/3}/\varepsilon^{4/3})$, but adds that ε must be constant.

(2) *Stricter distance measure.* Not only does our test reject with high probability when $\mathrm{d_{TV}}(\mathbf{p}, \mathbf{u}_d) > \varepsilon$, it also rejects with high probability when $\mathrm{d_H}(\mathbf{p}, \mathbf{u}_d) > \varepsilon$. (This is stronger, since $\mathrm{d_{TV}}(\mathbf{p}, \mathbf{q}) \le \mathrm{d_H}(\mathbf{p}, \mathbf{q})$ always.)

To elaborate, recall the below chain of inequalities, which also includes KL- and χ^2-divergence. (We review probability distance measures in Section 3.)

$$\mathrm{d_{TV}}(\mathbf{p}, \mathbf{q})^2 \le \mathrm{d}^2_{\mathrm{H}}(\mathbf{p}, \mathbf{q}) \le \mathrm{KL}(\mathbf{p} \,||\, \mathbf{q}) \le \chi^2(\mathbf{p} \,||\, \mathbf{q}). \tag{3}$$

The strictly stronger version of Theorem 1 that we prove is:

▶ **Theorem 4.** *There is a computationally efficient quantum algorithm for uniformity testing with the following guarantees: For $1/d \le \theta \le 1$, the algorithm makes $O(d^{1/3}/\theta^{2/3})$ uses of "the code" for an unknown distribution $\mathbf{p}$ over $[d]$, and distinguishes with probability at least $.99$ between*

$$(1)\ \chi^2(\mathbf{p} \,||\, \mathbf{u}_d) \le .99\theta \ \ \textit{and} \ \ \|\mathbf{p}\|_\infty \le 100/d, \qquad \textit{and} \qquad (2)\ \mathrm{d}^2_{\mathrm{H}}(\mathbf{p}, \mathbf{u}_d) > \theta. \tag{4}$$

We remark that most prior works on uniformity testing [10, 12, 15, 24] also had some additional such fine-grained aspects, beyond what is stated in Table 1.

Additional results

Speaking of χ^2-divergence, we mention two additional results we prove at the end of our work. These results additionally inform our Conjecture 2.

First, as mentioned earlier, in Section A we give an alternative proof of the $O(d^{1/2}/\varepsilon)$ upper bound of [24], and – like in that work – our result is tolerant with respect to χ^2-divergence. That is, we prove the strictly stronger result that for $\theta \le 1/d$, one can use the code $O(d^{1/2}/\theta^{1/2})$ times to distinguish $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) \le c\theta$ from $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) > \theta$ (for some constant $c > 0$).

Second, recall that $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$ can be as large as d. For example, $\chi^2(\mathbf{u}_S \,||\, \mathbf{u}_d) = \frac{d}{r} - 1$ for any set $S \subseteq [d]$ of size r. Thus it makes sense to consider the uniformity testing problem even with respect to a χ^2-divergence threshold θ that exceeds 1. In Section B we show (albeit only in the quantum string oracle model) that for $\theta \ge 1$, one can use the code $O(d^{1/3}/\theta^{1/3})$ times to distinguish $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) \le c\theta$ from $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) > \theta$, and this is optimal.

2 Technical overview of our proof

Our main algorithm is concerned with achieving the best possible ε-dependence for uniformity testing while maintaining a d-dependence of $d^{1/3}$; in this way, it is best compared with the older works of [10, 12], the latter of which achieves complexity $O(d^{1/3}/\varepsilon^2)$, as well as the classical (no-source-code) algorithm achieving complexity $O(d^{1/2}/\varepsilon^2)$. In fact, all four algorithms here are almost the same (except in terms of the number of samples they use). Let us describe our viewpoint on this common methodology.

We consider the algorithm as being divided into two Phases, and we may as well assume each Phase uses n samples. Phase 1 will have two properties:

- It will make n *black-box* draws from $\mathbf{p}$ (i.e., the source code is not used in Phase 1).
- Using these draws, Phase 1 will end by constructing a certain "random variable" – in the technical sense of a function $Y : [d] \to \mathbb{R}$.
- The mean of this random variable Y, vis-a-vis the unknown distribution $\mathbf{p}$, will ideally be close to $\chi^2(\mathbf{p} \,||\, \mathbf{u}) = d \cdot \|\mathbf{p} - \mathbf{u}_d\|_2^2$. That is, ideally $\mu := \mathbb{E}_{\mathbf{p}}[Y] = \sum_{j=1}^{d} \mathbf{p}_j Y(j) \approx \chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$.

Phase 2 then performs a *mean estimation* algorithm on Y (vis-a-vis $\mathbf{p}$) to get an estimate of μ and therefore of $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$. Ideally, the resulting overall algorithm is not just a uniformity tester, but a χ^2-divergence-from-uniformity *estimator*. This could then be weakened to a TV-distance uniformity tester using the inequality $\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{u}_d)^2 \leq \chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$.

The mean estimation algorithm used in Phase 2 differs depending on whether one has the source code or not. In the classical (no source code) model, one simply uses the naive mean estimation algorithm based on n more black-box samples; by Chebyshev's inequality, this will (with high probability) give an estimate of μ to within $\pm O(\sigma/n^{1/2})$, where $\sigma \coloneqq \mathrm{stddev}_{\mathbf{p}}[Y] = \sqrt{\sum_{j=1}^{d}(Y(j) - \mu)^2}$. In the case of a quantum tester with the source code access, we can use a *Quantum Mean Estimation* (QME) routine; in particular, the one from [22] will (with high probability) yield an estimate of μ to within $\pm O(\sigma/n)$.[2]

A subtle aspect of this overall plan is that the mean μ and standard deviation σ of Y *are themselves random variables* (in the usual sense), where the randomness comes from Phase 1. Thus it is natural to analyze $\mathbb{E}_{\text{Phase 1}}[\mu]$ and $\mathbb{E}_{\text{Phase 1}}[\sigma]$. Of course, these depend on the definition of Y, which we now reveal: $Y(j) = \frac{d}{n}X_j - 1$, where X_j denotes the number of times $j \in [d]$ was drawn in Phase 1. The point of this definition of Y is that a short calculation implies

$$\mathbb{E}_{\text{Phase 1}}[\mu] = \chi^2(\mathbf{p} \,||\, \mathbf{u}_d); \tag{5}$$

that is, the random variable μ is an *unbiased estimator* for our quantity of interest, the χ^2-divergence of $\mathbf{p}$ from $\mathbf{u}_d$. This is excellent, because although the algorithm does not see μ at the end of Phase 1, it will likely get a good estimate of it at the end of Phase 2... so long as (the random variable) σ is small.

We therefore finally have two sources of uncertainty about our final error (in estimating $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$):

1. Although $\mathbb{E}_{\text{Phase 1}}[\mu] = \chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$, the random variable μ may have fluctuated around its expectation at the end of Phase 1. One way to control this would be to bound $\mathrm{Var}_{\text{Phase 1}}[\mu]$ (and then use Chebyshev).
2. The Phase 2 mean estimation incurs an error proportional to σ. One way to control this would be to bound $\mathbb{E}_{\text{Phase 1}}[\sigma^2]$ (and then use Markov to get a high-probability bound on σ^2, and hence σ).

The quantities controlling the error here, $\mathrm{Var}_{\text{Phase 1}}[\mu]$ and $\mathbb{E}_{\text{Phase 1}}[\sigma^2]$, are explicitly calculable symmetric polynomials in $\mathbf{p}_1, \dots, \mathbf{p}_d$ of degree at most 4, depending on n. In principle, then, one can relate these quantities to $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) = d \cdot \|\mathbf{p} - \mathbf{u}_d\|_2^2$ itself, and derive a bound on how big n must be to (with high probability) get a good estimate of $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$.

In the classical (no source code) case, this methodology is a way to obtain the $O(d^{1/2}/\varepsilon^2)$ sample complexity, adding to the number of existing classical sample-optimal algorithms for the task. (This method in particular has some potential useful applications; e.g., one could consider decoupling the number of samples used in Phases 1 and 2 to, e.g., obtain tradeoffs for memory-limited settings). On one hand, with this method one can give a very compressed proof of the $O(d^{1/2}/\varepsilon^2)$ that, factoring out routine calculations, fits in half a

[2] This QME routine was not available at the time of [10, 12] which had to make do with Quantum Approximate Counting [9] – essentially, QME for Bernoulli random variables. But this is not the source of our improvement; one can obtain our main theorem with only a $(\mathrm{polylog}\, d)$-factor loss using just Quantum Approximate Counting.

page (see, e.g., [27, Sec. 10]). On the other hand, one has to execute the calculations and estimations with great care, lest one would obtain a suboptimal result (there is a reason it took 8 years[3] to get the optimal quadratic dependence on ε [17, 28]).

In the case when source code is available, so that one can use the QME algorithm, how well does this methodology fare? On one hand, QME gives a quadratic improvement over naive classical mean estimation, meaning one can try to use signficantly fewer samples in Phase 2. But when one balances out the sample complexity between the two Phases, it implies one is using fewer samples in Phase 1, and hence one gets worse concentration of μ around its mean in Phase 1. So the calcuations become more delicate.

2.1 Heuristic calculations

Instead of diving into complex calculations, let's look at some heuristics. First, let's consider how the algorithm proceeds in the case when $\mathbf{p}$ really is the uniform distribution $\mathbf{u}_d$. In this case, as long as we're in a scenario where $n \ll d^{1/2}$, we will likely get all distinct elements in Phase 1, meaning that X_j will be 1 for exactly n values of j and X_j will be 0 otherwise. Then $Y(j)$ will be $\frac{d}{n} - 1$ for n values of j and will be -1 otherwise. This indeed means $\mu = \mathbb{E}_{\mathbf{p}}[Y] = \frac{1}{d}\sum_{j=1}^{d} Y(j) = 0 = \|\mathbf{p} - \mathbf{u}_d\|_2$ with *certainty* in Phase 1. This is very good; we get no error out of Phase 1. However QME in Phase 2 will not perfectly return the value $\mu = 0$; rather, it will return something in the range $\pm O(\sigma/n)$, where $\sigma = \sqrt{\frac{1}{d}\sum_{j=1}^{d}(Y(j) - 0)^2} = \sqrt{\frac{d}{n} - 1} \sim d^{1/2}/n^{1/2}$. Thus the value returned by QME may well be around $d^{1/2}/n^{3/2}$, which from the algorithm's point of view is consistent with $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \approx d^{1/2}/n^{3/2}$. Thus the algorithm will only become confident that $\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{u}_d)^2 \lessapprox d^{1/2}/n^{3/2}$, and hence it can only confidently accept in the case $\mathbf{p} = \mathbf{u}_d$ provided $d^{1/2}/n^{3/2} \lessapprox \varepsilon^2$; i.e., $n \gtrapprox d^{1/3}/\varepsilon^{4/3}$. We thereby see that with this algorithm, a uniformity testing upper bound of $O(d^{1/3}/\varepsilon^{4/3})$ is the *best* we can hope for. If one also believes that this algorithm might be optimal (and it *has* been the method of choice for essentially all previously known results), then this could possibly be taken as evidence for our Conjecture 2.

At this point, one might try to prove that complexity $O(d^{1/3}/\varepsilon^{4/3})$ *is* achievable; so far we have only argued that with this many samples, the algorithm will correctly accept when $\mathbf{p} = \mathbf{u}_d$ (with high probability). Again, before jumping into calculations, one might try to guess the "hardest" kind of ε-far distributions one might face, and try to work out the calculations for these cases. The hardest distributions in the classical case (i.e., the ones that lead to the matching $\Omega(d^{1/2}/\varepsilon^2)$ lower bound) are very natural: they are the $\mathbf{p}$'s in which half of the elements $j \in [d]$ have $\mathbf{p}_j = \frac{1+2\varepsilon}{d}$ and half have $\mathbf{p}_j = \frac{1-2\varepsilon}{d}$. Assuming this is the "worst case", one can calculate what $\mathrm{Var}_{\text{Phase 1}}[\mu]$ and $\mathbb{E}_{\text{Phase 1}}[\sigma^2]$ will be, and the calculations turn out just as desired. That is, with $n = O(d^{1/3}/\varepsilon^{4/3})$, these two error quantities can be shown to be suffciently small so that the overall algorthm will correctly become confident that $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) = d \cdot \|\mathbf{p} - \mathbf{u}_d\|_2^2 \leq 4d \cdot \mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{u}_d)^2$ significantly exceeds ε^2, and hence the algorithm can correctly reject.

Everything therefore looks good, but there is a fly in the ointment. Even though this particular $\mathbf{p}$ with its values of $\frac{1\pm 2\varepsilon}{d}$ seems like the "hardest" distribution to face, one still has to reason about all possible $\mathbf{p}$'s with $\mathrm{d}_{\mathrm{TV}}(\mathbf{p}, \mathbf{u}_d)$. And when one does the calculations of $\mathrm{Var}[\mu]$ and $\mathbb{E}[\sigma^2]$ as prescribed by the standard methodology, the plan ends up *failing.*

[3] Technically, it took more than 8 years, as the proof of [28] was later shown to have a flaw: so the tight dependence had to wait until [4]. See [11, Section 2.3] for a discussion.

Specifically one gets too much error for $\mathbf{p}$'s that have somewhat "heavy" elements, meaning $\mathbf{p}_j$'s with $\mathbf{p}_j \gg 1/d$. The prior works [10, 12] cope with this failure by taking more samples; i.e., setting $n = O(d^{1/3}/\varepsilon^c)$ for $c > 4/3$ (specifically, [12] achieves $c = 2$). But our goal is to show that this is unnecessary – that the algorithm itself works, even though the standard and natural way of analyzing it fails.

In short, the reason the standard analysis of the algorithm fails is due to "rare events" that are caused by heavy elements in $\mathbf{p}$. These j's with $\mathbf{p}_j \gg 1/d$ may well still have $\mathbf{p}_j \ll 1/n$ (for our desired $n = O(d^{1/3}/\varepsilon^{4/3})$), and thus be drawn only rarely in Phase 1. The major difficulty is that when they *are* drawn, they generate a *very* large contribution to σ^2, causing $\mathbb{E}_{\text{Phase 1}}[\sigma^2]$ to be "misleadingly large". That is, when there are heavy elements, σ^2 may have the property of typically being much smaller than its expectation. Thus controlling the QME error using the *expected* value of σ^2 is a bad strategy.

Perhaps the key insight in our analysis is to show: *In those rare Phase 1 outcomes when σ^2 is unusually large, μ is* also *unusually large compared to its expectation.* The latter event is helpful, because if μ ends up much bigger than its expectation, we can tolerate a correspondingly worse error-bar from QME. In short, we show that the rare *bad* outcomes for σ^2 coincide with the rare *good* outcomes for μ.

In order to make this idea work out quantitatively, we (seem to) need to weaken our ambitions and get something a bit worse than a χ^2-divergence-from-uniform estimation algorithm, in two ways. (This is fine, as our main goal is just a non-tolerant uniformity tester with respect to TV.) First, rather than insisting that we accept with high probability when $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) \leq .99\theta$ and reject with high probability when $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) > \theta$, we need to only require rejection when $\mathrm{d}_{\mathrm{H}}^2(\mathbf{p}, \mathbf{u}_d) > \theta$. The reason is that the rare large values of σ^2 that we face are only comparable with the larger value $\mathrm{d}_{\mathrm{H}}^2(\mathbf{p}, \mathbf{u}_d)$, and not with $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d)$.[4]

As for the second weakening we need to make: We explicitly add to our tester a check that the value of $\max_j\{X_j\}$ arising after Phase 1 is not too large. Roughly speaking, this extra test ensures that there are no very heavy elements. (Of course, this is satisfied when $\mathbf{p} = \mathbf{u}_d$, so we don't mind adding this test.) The reason we need to add this check is so that we can bound the quadratic expression $\sum_{j=1}^d X_j^2$ (which enters into the value of σ^2) by $\max_j\{X_j\} \cdot \sum_{j=1}^d X_j$; in turn, once $\max_j\{X_j\}$ is checked to be small, this expression can be bounded by the linear quantity $\sum_{j=1}^d X_j$, which can be related to μ. It is by relating σ^2 to μ in this way that we are able to show the correlation between rare events – that when σ^2 is big, μ is also big.

To conclude, we apologize to the reader for writing a "technical overview" whose length is nearly comparable to that of the actual proof itself. While we tried to make our argument as streamlined and concise as possible, we felt that it was worth conveying the ideas and detours which led us there, and which, while now hidden, motivated the final proof.

3 Preliminaries

3.1 Probability distances

Throughout, log and ln the binary and natural logarithms, respectively. We identify a probability distribution $\mathbf{p}$ over $[d] = \{1, 2, \ldots, d\}$ with its probability mass function (pmf), or,

[4] We remark that this χ^2-versus-Hellinger-squared dichotomy is quite reminiscent of the one that occurs in classical works on identity testing, such as [4].

equivalently, a vector $\mathbf{p} \in \mathbb{R}^d$ such that $\mathbf{p}_i \geq 0$ for all i and $\sum_{i=1}^{d} \mathbf{p}_i = 1$. For a subset $S \subseteq [d]$, we accordingly let $\mathbf{p}(S) = \sum_{i \in S} \mathbf{p}_i$. The *total variation distance* between two distributions $\mathbf{p}, \mathbf{q}$ over $[d]$ is defined as

$$\mathrm{d_{TV}}(\mathbf{p}, \mathbf{q}) = \sup_{S \subseteq [d]} \{\mathbf{p}(S) - \mathbf{q}(S)\} = \frac{1}{2}\|\mathbf{p} - \mathbf{q}\|_1 \in [0, 1], \tag{6}$$

where the last equality is from Scheffé's lemma. By Cauchy–Schwarz, this gives us the relation

$$\frac{1}{2}\|\mathbf{p} - \mathbf{q}\|_2 \leq \mathrm{d_{TV}}(\mathbf{p}, \mathbf{q}) \leq \frac{\sqrt{d}}{2}\|\mathbf{p} - \mathbf{q}\|_2. \tag{7}$$

We will in this paper also consider other notions of distance between probability distributions: the *squared Hellinger distance*, defined as

$$\mathrm{d_H^2}(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{d} (\sqrt{\mathbf{p}_i} - \sqrt{\mathbf{q}_i})^2 = \|\sqrt{\mathbf{p}} - \sqrt{\mathbf{q}}\|_2^2 \in [0, 2]. \tag{8}$$

(Some texts normalize this by a factor of $\frac{1}{2}$; we do not do so, as it makes our statements cleaner.) The *chi-squared divergence* is then defined as

$$\chi^2(\mathbf{p} \mid\mid \mathbf{q}) = \sum_{i=1}^{d} \frac{(\mathbf{p}_i - \mathbf{q}_i)^2}{\mathbf{q}_i} = \left(\sum_{i=1}^{d} \frac{\mathbf{p}_i^2}{\mathbf{q}_i}\right) - 1\,, \tag{9}$$

while the *Kullback–Leibler divergence* (least relevant to us, but quite common in the literature), in nats, is defined as

$$\mathrm{KL}(\mathbf{p} \mid\mid \mathbf{q}) = \sum_{i=1}^{d} \mathbf{q}_i \ln \frac{\mathbf{q}_i}{\mathbf{p}_i}\,. \tag{10}$$

As mentioned in Equation (3), we have the following well known [14] chain of inequalities:

$$\mathrm{d_{TV}}(\mathbf{p}, \mathbf{q})^2 \leq \mathrm{d_H^2}(\mathbf{p}, \mathbf{q}) \leq \mathrm{KL}(\mathbf{p} \mid\mid \mathbf{q}) \leq \chi^2(\mathbf{p} \mid\mid \mathbf{q}). \tag{11}$$

Moreover, for the special case of the uniform distribution $\mathbf{u}_d$ over $[d]$, we have

$$\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) = d \cdot \|\mathbf{p} - \mathbf{u}_d\|_2^2\,. \tag{12}$$

3.2 Distribution access models

For a probability distribution $\mathbf{p}$ on $[d]$, we say a unitary $U_\mathbf{p}$ is a *synthesizer* for $\mathbf{p}$ if for some k

$$U_\mathbf{p} \left|0^k\right\rangle = \sum_{i \in [d]} \sqrt{p_i} \left|i\right\rangle \left|\psi_i\right\rangle, \tag{13}$$

where the $|\psi_i\rangle$'s are normalized states often called "garbage states". Note that any classical randomized circuit using S gates that samples from $\mathbf{p}$ can be converted to a synthesizer $U_\mathbf{p}$ in a purely black-box way with gate complexity $O(S)$. (See [22] for details and a more thorough discussion of synthesizers.)

In this paper, we say an algorithm makes t uses of "the code for $\mathbf{p}$" to mean that we use (as a black box) the unitaries $U_\mathbf{p}$, $U_\mathbf{p}^\dagger$, and controlled-$U_\mathbf{p}$ a total of t times in the algorithm.

Each of these unitaries is easy to construct given an explicit gate decomposition of $U_{\mathbf{p}}$ with the same gate complexity up to constant factors.

The *quantum string oracle*, which is used in many prior works, is a specific type of source code for $\mathbf{p}$. Here we have standard quantum oracle access to an m-bit string $x \in [d]^m$ for some m. For any symbol $i \in [d]$, the probability $\mathbf{p}_i$ is defined as the frequency with which that symbol appears in x, i.e., $\mathbf{p}_i = \frac{1}{m}|\{j : x_j = i\}|$. Note that calling this oracle on the uniform superposition over m gives us a synthesizer for $\mathbf{p}$. When a randomized sampler for $\mathbf{p}$ is converted to a synthesizer, we get a quantum string oracle, but quantum string oracles are not as general as arbitrary synthesizers. For example, all probabilities described by a quantum string oracle will be integer multiples of $\frac{1}{m}$, whereas an arbitrary synthesizer has no such constraint.

3.3 Quantum Mean Estimation

When we use QME, we will have the source code for some distribution $\mathbf{p}$ on $[d]$, and we will also have explicitly constructed some (rational-valued) random variable $Y : [d] \to \mathbb{Q}$ (say, simply as a table). From this, one can easily generate code that outputs a sample from Y (i.e., outputs $Y(\boldsymbol{j})$ for $\boldsymbol{j} \sim [d]$), using the code for $\mathbf{p}$ just one time. We will then use the following QME result from [22]:

▶ **Theorem 5.** *There is a computationally efficient quantum algorithm with the following guarantee: Given the source code for a random variable Y, as well as parameters n and δ, the algorithm uses the code $O(n \log(1/\delta))$ times and outputs an estimate $\widehat{\boldsymbol{\mu}}$ such that $\Pr[|\widehat{\mu} - \mu| > \sigma/n] \leq \delta$, where $\mu = \mathbb{E}[Y]$ and $\sigma = \mathrm{stddev}[Y]$.*

4 Algorithm in the Large Distance Regime

In this section, we establish Theorem 1, our main technical contribution. We do this by proving the strictly stronger Theorem 4, which we restate more formally:

▶ **Theorem 6.** *For any constant $B > 0$, there exists a computationally efficient quantum algorithm (Algorithm 1) with the following guarantees: on input $\frac{1}{d} \leq \theta \leq 1$, it makes $O(d^{1/3}/\theta^{2/3})$ uses (where the hidden constant depends on B) of "the code" for an unknown probability distribution $\mathbf{p}$ over $[d]$, and satisfies*

1. *If $\chi^2(\mathbf{p} \,||\, \mathbf{u}_d) \leq .99\theta$ and $\|\mathbf{p}\|_\infty \leq B/d$, then the algorithm will* accept *with probability at least .99.*
2. *If $\mathrm{d}_{\mathrm{H}}^2(\mathbf{p}, \mathbf{u}_d) \geq \theta$, then the algorithm will* reject *with probability at least .99.*

Proof. Let us start by recording the following inequalities that we will frequently use:

$$n = \lceil cd^{1/3}/\theta^{2/3} \rceil, \ \theta \geq 1/d \quad \implies \quad c/\theta \leq n \leq cd. \tag{14}$$

We begin with a simple lemma regarding the check on Section 4:

▶ **Lemma 7.** *If $\|\mathbf{p}\|_\infty \leq B/d$, then Section 4 will* reject *with probability at most .001. Conversely, if $\|\mathbf{p}\|_\infty > 2L/n$, then Section 4 will* reject *with probability at least .999.*

Proof. Let $\boldsymbol{X}_j \sim \mathrm{Bin}(n, p_j)$ denote the number of times j is drawn. The second ("conversely") part of of the proposition follows from a standard Chernoff bound. As for the first part, suppose $\|\mathbf{p}\|_\infty \leq B/d$. Now on one hand, if $n \leq d^{.99}/B$, so that $L = 100$, we have

$$\Pr[\mathrm{Bin}(n, p_j) \geq 100] \leq \binom{n}{100} p_j^{100} \leq ((en/100)p_j)^{100} \leq (e/(100d^{.01}))^{100} \leq .001/d, \tag{15}$$

Algorithm 1 for the large distance regime.

```
Require: Parameter $\frac{1}{d} \leq \theta \leq 1$, constant $B \geq 1$.
1: Let $c = c(B)$ and let $C = C(c)$ be sufficiently large, and let $L$ be defined as
       $L := \begin{cases} 100 & \text{if } n \leq d^{.99}/B, \\ Bc \ln d & \text{if } n > d^{.99}/B. \end{cases}$
2: Set $n := \lceil c d^{1/3}/\theta^{2/3} \rceil$.
3: Make $n$ draws $\boldsymbol{J}_1, \dots, \boldsymbol{J}_n$, and let $\boldsymbol{X}_j = \sum_{t=1}^n \mathbb{1}_{\{J_t = j\}}$ be the number of times $j \in [d]$ is seen.
4: if $\boldsymbol{X}_j \geq L$ for any $j$ then reject
5: Do QME with $Cn$ "samples" on the random variable $\boldsymbol{Y}$ defined by $\boldsymbol{Y}_j = \frac{d}{n} X_j - 1$, obtaining $\widehat{\boldsymbol{\mu}}$.
6: if $\widehat{\boldsymbol{\mu}} \leq .995\theta$ then accept
7: else reject
```

and thus $\boldsymbol{X}_j < 100$ for all j except with probability at most .001, as desired. Otherwise, $L = Bc \ln d$, and since $\mathbb{E}[\boldsymbol{X}_j] \leq Bn/d \leq Bc$, the desired result follows from a standard Chernoff and union bound (provided c is large enough). ◀

From this, we conclude:

- In Case (1), Line 4 rejects with probability at most .001.
- In Case (2), we may assume $\|\mathbf{p}\|_\infty \leq 2L/n$ and $\|\boldsymbol{X}\|_\infty \leq L$, else Line 4 rejects with probability $\geq .999$. Call this observation (♢).

Now to begin the QME analysis, write $p_j = \frac{1+\varepsilon_j}{d}$, where $\varepsilon_j \in [-1, d-1]$, and let $\boldsymbol{\mu} = \sum_{j=1}^d p_j \boldsymbol{Y}_j$, the mean of $\boldsymbol{Y}$ (from QME's point of view). Writing $\eta := \mathrm{d}_{\mathrm{H}}^2(p, \mathbf{u}_d)$, our first goal will be to show:

$$\text{In Case (1),} \qquad \boldsymbol{\mu} \leq .991\theta \qquad \text{except with probability at most .001;} \tag{16}$$

$$\text{In Case (2),} \qquad \boldsymbol{\mu} \geq .997\eta \qquad \text{except with probability at most .002.} \tag{17}$$

Starting with Equation (16), a short calculation (using $\sum_{j=1}^d \varepsilon_j = 0$) shows

$$\boldsymbol{\mu} = \operatorname*{avg}_{t=1}^{n}\{\varepsilon_{\boldsymbol{J}_t}\} \quad \implies \quad \mathbb{E}[\boldsymbol{\mu}] = \frac{1}{d}\sum_{j=1}^d \varepsilon_j^2 = \chi^2(p \,||\, \mathbf{u}_d) \quad \implies \quad \mathbb{E}[\boldsymbol{\mu}] \leq .99\theta \text{ in Case (1).} \tag{18}$$

Also in Case (1) we get from Equation (18) that

$$\mathrm{Var}[\boldsymbol{\mu}] = \frac{1}{n}\mathrm{Var}_{\boldsymbol{j}\sim p}[\varepsilon_{\boldsymbol{j}}] \leq \frac{1}{n}\mathbb{E}_{\boldsymbol{j}\sim p}[\varepsilon_{\boldsymbol{j}}^2] \leq \frac{B}{nd}\sum_{j=1}^n \varepsilon_j^2 = \frac{B}{n}\chi^2(p \,||\, \mathbf{u}_d) \leq \frac{.99B\theta}{n} \leq \frac{B\theta^2}{c}, \tag{19}$$

the last inequality using Equation (14). Combining the preceding two inequalities and using Chebyshev, we indeed conclude Equation (16) (provided $c = c(B)$ is sufficiently large).

Towards Equation (17), let $b \geq 2$ be a certain universal constant to be chosen later, and say that $j \in [d]$ is *light* if $p_j \leq b/d$ (i.e., $\varepsilon_j \leq b - 1$), *heavy* otherwise. We will write

$$\boldsymbol{\mu}_1 = \operatorname*{avg}_{t=1}^{n}\{\varepsilon_{\boldsymbol{J}_t} : \boldsymbol{J}_t \text{ heavy}\} \geq 0, \quad \boldsymbol{\mu}_2 = \operatorname*{avg}_{t=1}^{n}\{\varepsilon_{\boldsymbol{J}_t} : \boldsymbol{J}_t \text{ light}\} \qquad (\text{so } \boldsymbol{\mu} = \boldsymbol{\mu}_1 + \boldsymbol{\mu}_2), \tag{20}$$

and also observe

$$\eta = \mathrm{d}_{\mathrm{H}}^2(p, \mathbf{u}_d) = \frac{1}{d}\sum_{j=1}^{d}(\sqrt{1+\varepsilon_j}-1)^2 \le \frac{1}{d}\sum_{j=1}^{d}\min\{|\varepsilon_j|, \varepsilon_j^2\} \le \frac{1}{d}\sum_{\text{heavy } j}\varepsilon_j + \frac{1}{d}\sum_{\text{light } j}\varepsilon_j^2 =: \eta_1+\eta_2. \quad (21)$$

Let us now make some estimates. First:

$$p_{\text{heavy}} \coloneqq \sum_{j \text{ heavy}} p_j = \frac{1}{d}\sum_{j \text{ heavy}}(1+\varepsilon_j) \ge \eta_1. \quad (22)$$

Also, similar to our Case (1) estimates we have

$$\mathbb{E}[\boldsymbol{\mu}_2] = \frac{1}{d}\sum_{\text{light } j}(\varepsilon_j^2+\varepsilon_j) = \eta_2 - \eta_1 \quad \text{(where we used } \sum_{j=1}^{d}\varepsilon_j = 0\text{)}, \quad (23)$$

and

$$\begin{aligned}
&\mathrm{Var}[\boldsymbol{\mu}_2] \\
&= \frac{1}{n}\mathrm{Var}_{\boldsymbol{j}\sim p}[\mathbb{1}_{\boldsymbol{j}\text{ light}}\cdot\varepsilon_{\boldsymbol{j}}] \le \frac{1}{n}\mathbb{E}_{\boldsymbol{j}\sim p}[\mathbb{1}_{\boldsymbol{j}\text{ light}}\cdot\varepsilon_{\boldsymbol{j}}^2] \le \frac{b}{nd}\sum_{j \text{ light}}\varepsilon_j^2 \\
&= \frac{b}{n}\eta_2 \le \frac{b}{c}\theta\eta_2 \le \frac{b}{c}\eta_2\eta \text{ (in Case (2))}. \quad (24)
\end{aligned}$$

We will now establish Equation (17); in fact, we we even will show the following very slightly stronger fact:

$$\text{In Case (2),} \qquad \boldsymbol{\mu} \ge .997(\eta_1+\eta_2) \ge .997\eta \quad \text{except with probability at most .002.} \quad (25)$$

We divide into two subcases:

Case (2a): $\eta_1 \le .001\eta_2$. In this case we have $\eta_2 \ge \frac{1}{1.001}(\eta_1+\eta_2)$, and $\mathbb{E}[\boldsymbol{\mu}_2] \ge .999\eta_2$ from Equation (23). Since Equation (24) implies $\mathrm{Var}[\boldsymbol{\mu}_2] \le 1.001\frac{b}{c}\eta_2^2$, Chebyshev's inequality tells us that $\boldsymbol{\mu}_2 \ge .998\eta_2$ except with probability at most .001 (provided c is large enough). But then $\boldsymbol{\mu} \ge \boldsymbol{\mu}_2 \ge \frac{.998}{1.001}(\eta_1+\eta_2)$, confirming Equation (25).

Case (2b): $\eta_1 > .001\eta_2$. In this case we have $\eta_1 \ge \frac{.001}{1.001}(\eta_1+\eta_2) \ge .0009(\eta_1+\eta_2)$. We now use that heavy j have $\varepsilon_j \ge b-1$ to observe that

$$\boldsymbol{\mu}_1 = \operatorname*{avg}_{t=1}^{n}\{\varepsilon_{\boldsymbol{J}_t} : \boldsymbol{J}_t \text{ heavy}\} \ge (b-1)\cdot(\text{fraction of } \boldsymbol{J}_t\text{'s that are heavy}) = (b-1)\cdot\frac{\mathrm{Bin}(n, p_{\text{heavy}})}{n} \quad (26)$$

(in distribution). We see that $\mathbb{E}[\boldsymbol{\mu}_1] \ge (b-1)p_{\text{heavy}}$, and moreover concentration of Binomials and Equation (22) imply that

$$\boldsymbol{\mu}_1 \ge \frac{1}{2}(b-1)p_{\text{heavy}} \ge \frac{1}{2}(b-1)\eta_1 \text{ except with probability at most .001}, \quad (27)$$

provided that $p_{\text{heavy}}n$ is a sufficiently large constant. But we can indeed ensure this by taking c sufficient large: by Equation (22), being in Case (2b), and Equation (14), it holds that

$$p_{\text{heavy}}n \ge \eta_1 n \ge .0009(\eta_1+\eta_2)n \ge .0009\eta n \ge .0009\theta n \ge .0009c. \quad (28)$$

At the same time, Equation (23) certainly implies $\mathbb{E}[\boldsymbol{\mu}_2] \geq -\eta_1$, and Equation (24) implies $\mathrm{Var}[\boldsymbol{\mu}_2] \leq \frac{b}{c}\eta_2(\eta_1+\eta_2) \leq \frac{1000\cdot 1001 b}{c}\eta_1^2$ (using Case (2b)). Thus Chebyshev implies

$$\boldsymbol{\mu}_2 \geq -1.1\eta_1 \text{ except with probability at most .001,} \tag{29}$$

provided c is large enough. Combining Equations (27) and (29) yields

$$\boldsymbol{\mu} = \boldsymbol{\mu}_1+\boldsymbol{\mu}_2 \geq (\tfrac{b-1}{2}-1.1)\eta_1 \geq .0009(\tfrac{b-1}{2}-1.1)(\eta_1+\eta_2) \text{ except with probability at most .002,} \tag{30}$$

which verifies Equation (25) provided b is a large enough constant.

We have now verified the properties of $\boldsymbol{\mu}$ claimed in Equations (16) and (25). Next we analyze the random variable $\boldsymbol{\sigma}^2$ that represents the variance of $\boldsymbol{Y}$ (from QME's point of view). Our goal will be to show:

$$\text{In Case (1),} \qquad \boldsymbol{\sigma}^2/(Cn)^2 \leq 10^{-6}\cdot\theta^2 \qquad \text{except with probability at most .001,} \tag{31}$$

$$\text{In Case (2),} \qquad \boldsymbol{\sigma}^2/(Cn)^2 \leq 10^{-6}\cdot\boldsymbol{\mu}^2 \qquad \text{except with probability at most .001.} \tag{32}$$

Together with Equations (16) and (25), these facts are sufficient to complete the proof of the theorem, by the QME guarantee of Theorem 5.

We have:

$$\boldsymbol{\sigma}^2 \coloneqq \sum_{j=1}^d p_j\boldsymbol{Y}_j^2 - \boldsymbol{\mu}^2 = (d/n)^2\sum_{j=1}^d p_j\boldsymbol{X}_j^2 - (\boldsymbol{\mu}+1)^2 \leq (d/n)^2\sum_{j=1}^d p_j\boldsymbol{X}_j^2 = \boldsymbol{\sigma}_S^2 + \boldsymbol{\sigma}_{S^c}^2, \tag{33}$$

where we've defined $\boldsymbol{\sigma}_S^2 \coloneqq (d/n)^2\sum_{j\in S}p_j\boldsymbol{X}_j^2$ and $S^c = [d]\setminus S$. We will be making two different choices for S later, but we will always assume

$$S \supseteq \{j : j \text{ light}\}, \quad \text{which implies } \sum_{j\in S}\varepsilon_j \leq 0 \tag{34}$$

(the implication because $\sum_{j=1}^d \varepsilon_j = 0$ and S^c contains only j's with $\varepsilon_j \geq b-1 \geq 0$). Now since $\mathbb{E}[\boldsymbol{X}_j^2] = np_j(1-p_j)+(np_j)^2 \leq np_j+(np_j)^2$, we have

$$\mathbb{E}[\boldsymbol{\sigma}_S^2] \leq (d^2/n)\sum_{j\in S}p_j^2 + d^2\sum_{j\in S}p_j^3 \tag{35}$$

$$\leq d/n + (2/n)\sum_{j\in S}\varepsilon_j + (1/n)\sum_{j\in S}\varepsilon_j^2 + 1/d + (3/d)\sum_{j\in S}\varepsilon_j + (3/d)\sum_{j\in S}\varepsilon_j^2 + (1/d)\sum_{j\in S}\varepsilon_j^3 \tag{36}$$

$$\leq (5cd/n)\left(1+\frac{1}{d}\sum_{j\in S}\varepsilon_j + \frac{1}{d}\sum_{j\in S}\varepsilon_j^2\right) + \frac{1}{d}\sum_{j\in S}\varepsilon_j^3 \tag{37}$$

(where the last inequality used $1/d \leq c/n \leq (c-1)d/n$ from Equation (14)). Using Equation (34) to drop the term of Equation (37) that's linear in the ε_j's, we thereby conclude

$$\mathbb{E}[\boldsymbol{\sigma}_S^2/(Cn)^2] \leq \mathbb{E}[\boldsymbol{\sigma}_S^2/n^2] \leq (5cd/n^3)\left(1+\frac{1}{d}\sum_{j\in S}\varepsilon_j^2\right) + (d^{1/2}/n^2)\left(\frac{1}{d}\sum_{j\in S}\varepsilon_j^2\right)^{3/2} \tag{38}$$

$$\leq (5\theta^2/c^2)(1+\eta_S) + \frac{\theta^{4/3}}{c^2d^{1/6}}\eta_S^{3/2}, \tag{39}$$

where $\eta_S := \frac{1}{d}\sum_{j\in S}\varepsilon_j^2$. In Case (1) we select $S = [d]$, so $\eta_S = \chi^2(p \,||\, \mathbf{u}_d) \leq .99\theta \leq \theta \leq 1$, and the above bound gives

$$\text{Case (1)} \implies \mathbb{E}[\boldsymbol{\sigma}^2/(Cn)^2] \leq 10\theta^2/c^2 + \frac{\theta^{17/6}}{c^2 d^{1/6}} \leq \cdot 10^{-9} \cdot \theta^2 \tag{40}$$

(provided c is large enough). Now Equation (31) follows by Markov's inequality.

In Case (2) we select $S = \{j : j \text{ light}\}$, so $\eta_S = \eta_2$ and we conclude (using obvious notation)

$$\text{Case (2)} \implies \mathbb{E}[\boldsymbol{\sigma}^2_{\text{light}}/(Cn)^2] \leq (5\theta^2/c^2)(1+\eta_2) + \frac{\theta^{4/3}}{c^2 d^{1/6}}\eta_2^{3/2} \leq .4 \cdot 10^{-9} \cdot (\eta_1 + \eta_2)^2, \tag{41}$$

(provided c large enough), where we used $\theta \leq \eta \leq \eta_1 + \eta_2$ and also $\theta \leq 1$. We now complete the bounding of $\boldsymbol{\sigma}^2$ in Case (2) by two different strategies:

Case (2.i): $n > d^{.99}/B$. In this case, $L = Bc \ln d$, and ($\diamondsuit$) tells us $\|\mathbf{p}\|_\infty \leq 2L/n$, so we have

$$\|\mathbf{p}\|_\infty \leq \frac{2Bc\ln d}{n} \leq \frac{2B^2 c \ln d}{d^{.99}}. \tag{42}$$

Now returning to Equation (37), we get

$$\mathbb{E}[\boldsymbol{\sigma}^2_{\text{heavy}}/(Cn)^2] \leq \frac{5cd}{C^2n^3} + \frac{5cd}{C^2n^3}\left(1 + \varepsilon_{\max} + \frac{n}{5cd}\varepsilon_{\max}^2\right) \cdot \frac{1}{d}\sum_{j \text{ heavy}} \varepsilon_j \tag{43}$$

$$\leq \frac{5\theta^2}{(Cc)^2} + \frac{5cd^2}{C^2n^3}\left(\|\mathbf{p}\|_\infty + \frac{n}{5c}\cdot\|\mathbf{p}\|_\infty^2\right)\eta_1 \leq \frac{5\theta^2}{(Cc)^2} + \frac{14B^6c^2\ln^2 d}{C^2 d^{1.96}}\eta_1, \tag{44}$$

where we used Equation (42) and $n > d^{.99}/B$. We can again bound the first expression in Equation (44) as $\frac{5\theta^2}{(Cc)^2} \leq 10^{-6}\cdot(\eta_1+\eta_2)^2$. As for the second expression, either $\eta_1 = 0$ (there are no heavy j's) or else $\eta_1 \geq \frac{b-1}{d}$ (there is at least one heavy j). In either case, we have $\eta_1 \leq \frac{d}{b-1}\eta_1^2 \leq d\eta_1^2$, so we can bound this second expression by

$$\frac{14B^6c^2\ln^2 d}{C^2 d^{.96}}\eta_1^2 \leq .4 \cdot 10^{-9} \cdot (\eta_1+\eta_2)^2 \tag{45}$$

where we used $C = C(c)$ sufficiently large (and we could have taken $C = 1$ were willing to assume d sufficiently large). Putting this bound together with Equation (41) we obtain:

$$\text{Case (2.i)} \implies \mathbb{E}[\boldsymbol{\sigma}/(Cn)^2] \leq .8 \cdot 10^{-9} \cdot (\eta_1+\eta_2)^2 \leq \tfrac{.8}{.997}\cdot 10^{-9}\cdot\boldsymbol{\mu}^2 \leq \cdot 10^{-9}\cdot\boldsymbol{\mu}^2, \tag{46}$$

using Equation (25). Equation (32) now follows (in this Case (2.i)) by Markov's inequality.

Case (2.ii): $n \leq d^{.99}/B$. In this case we use a different strategy. Recall from Equation (33) that

$$\boldsymbol{\sigma}^2 \leq (d/n)^2\sum_{j=1}^d p_j\boldsymbol{X}_j^2 \leq (d/n)^2\|\boldsymbol{X}\|_\infty\sum_{j=1}^d p_j\boldsymbol{X}_j = (d/n)\|\boldsymbol{X}\|_\infty(1+\boldsymbol{\mu}). \tag{47}$$

By ($\diamondsuit$) we may assume $\|\boldsymbol{X}\|_\infty \leq L = 100$, the equality because we are in Case (2.ii). Thus

$$\boldsymbol{\sigma}^2/(Cn)^2 \leq \boldsymbol{\sigma}^2/n^2 \leq 100(d/n^3)(1+\boldsymbol{\mu}) \leq \frac{100\theta^2}{c^3} + \frac{100\theta^2}{c^3}\boldsymbol{\mu} \leq 10^{-6}\cdot\boldsymbol{\mu}^2. \tag{48}$$

(provided c large enough), where we used $\theta \leq \eta \leq \frac{1}{.997}\boldsymbol{\mu}$ (from Equation (25)) and also $\theta \leq 1$. This verifies Equation (32) in Case (2.ii), completing the proof. ◀

References

1 Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013. doi:10.4086/toc.2013.v009a004.

2 Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, July 2004. doi:10.1145/1008731.1008735.

3 Jayadev Acharya, Clément L. Canonne, Yanjun Han, Ziteng Sun, and Himanshu Tyagi. Domain compression and its application to randomness-optimal distributed goodness-of-fit. In *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 3–40. PMLR, July 2020. URL: http://proceedings.mlr.press/v125/acharya20a.html.

4 Jayadev Acharya, Constantinos Daskalakis, and Gautam C. Kamath. Optimal Testing for Properties of Distributions. In *Advances in Neural Information Processing Systems 28*, pages 3577–3598. Curran Associates, Inc., 2015.

5 Andris Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007. doi:10.1137/S0097539705447311.

6 Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, October 2019. doi:10.1038/s41586-019-1666-5.

7 Aleksandrs Belovs. Quantum algorithms for classical probability distributions. In *Proceedings of the 27th Annual European Symposium on Algorithms (ESA)*, pages 50–59. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.1007/978-3-030-19955-5_5.

8 Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, April 2018. doi:10.1038/s41567-018-0124-x.

9 Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum counting. In *Proceedings of the 25th Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 820–831. Springer–Verlag, 1998. doi:10.1007/bfb0055105.

10 Sergey Bravyi, Aram Harrow, and Avinatan Hassidim. Quantum algorithms for testing properties of distributions. *Transactions on Information Theory*, 57(6):3971–3981, 2011. doi:10.1109/TIT.2011.2134250.

11 Clément L. Canonne. Topics and techniques in distribution testing: A biased but representative sample. *Foundations and Trends® in Communications and Information Theory*, 19(6):1032–1198, 2022. doi:10.1561/0100000114.

12 Sourav Chakraborty, Eldar Fischer, Arie Matsliah, and Ronald de Wolf. New Results on Quantum Property Testing. In *30th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, volume 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 145–156, Dagstuhl, Germany, 2010. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.FSTTCS.2010.145.

13 Ilias Diakonikolas and Daniel M. Kane. A new approach for testing properties of discrete distributions. In *57th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2016*. IEEE Computer Society, 2016.

14 Alison Gibbs and Francis Su. On choosing and bounding probability metrics. *International Statistical Rreview*, 70(3):419–435, 2002.

15 András Gilyén and Tongyang Li. Distributional Property Testing in a Quantum World. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:19, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. `doi:10.4230/LIPIcs.ITCS.2020.25`.

16 Oded Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:15, 2016. URL: `http://eccc.hpi-web.de/report/2016/015`.

17 Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. Technical Report TR00-020, Electronic Colloquium on Computational Complexity (ECCC), 2000.

18 Lov Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the* 28*th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 212–219. ACM, New York, 1996. `doi:10.1145/237814.237866`.

19 Lov Grover. A framework for fast quantum mechanical algorithms. In *Proceedings of the* 30*th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 53–62. ACM, New York, 1998. `doi:10.1145/276698.276712`.

20 Yassine Hamoudi. *Quantum Algorithms for the Monte Carlo Method.* PhD thesis, Université de Paris, 2021.

21 Stefan Heinrich. Quantum summation with an application to integration. *Journal of Complexity*, 18(1):1–50, 2002. `doi:10.1006/jcom.2001.0629`.

22 Robin Kothari and Ryan O'Donnell. *Mean estimation when you have the source code; or, quantum Monte Carlo methods*, pages 1186–1215. Society for Industrial and Applied Mathematics, January 2023. `doi:10.1137/1.9781611977554.ch44`.

23 Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(2):29–36, 2005. `doi:10.4086/toc.2005.v001a002`.

24 Jingquan Luo, Qisheng Wang, and Lvzhou Li. Succinct quantum testers for closeness and k-wise uniformity of probability distributions. *IEEE Trans. Inf. Theory*, 70(7):5092–5103, 2024.

25 Ashley Montanaro. Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A*, 471(2181):20150301, 20, 2015. `doi:10.1098/rspa.2015.0301`.

26 Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, STOC '99, pages 384–393, New York, NY, USA, 1999. Association for Computing Machinery. `doi:10.1145/301250.301349`.

27 Ryan O'Donnell and John Wright. Learning and testing quantum states via probabilistic combinatorics and representation theory. In *Current developments in mathematics 2021*, pages 43–94. International Press, Somerville, MA, 2023.

28 Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008. `doi:10.1109/TIT.2008.928987`.

A Algorithm in the Small Distance Regime

In this appendix, we provide an alternative (and arguably simpler) proof of the main result of [24]:

▶ **Theorem 8.** *There is a computationally efficient quantum algorithm (Algorithm 2) for uniformity testing with the following guarantees: it takes* $O(d^{1/2}/\varepsilon)$ *"samples" from an* *unknown probability distribution* $\mathbf{p}$ *over* $[d]$*, and distinguishes with probability at least* $2/3$ *between (1)* $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq \frac{\varepsilon^2}{144}$*, and (2)* $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) > \varepsilon^2$*.*

This in turn will follow from the more general result on tolerant ℓ_2 closeness testing, where one is given access to the source code for *two* unknown probability distributions $\mathbf{p}, \mathbf{q}$ over $[d]$, and one seeks to distinguish $\|\mathbf{p} - \mathbf{q}\|_2 \leq c \cdot \tau$ from $\|\mathbf{p} - \mathbf{q}\|_2 \geq \tau$.

▶ **Theorem 9.** *There is a computationally efficient quantum algorithm (Algorithm 2) for closeness testing with the following guarantees: it takes* $O(1/\tau)$ *"samples" from two unknown probability distributions* $\mathbf{p}, \mathbf{q}$ *over* $[d]$*, and distinguishes with probability at least* $2/3$ *between (1)* $\|\mathbf{p} - \mathbf{q}\|_2 \leq \frac{\tau}{12}$*, and (2)* $\|\mathbf{p} - \mathbf{q}\|_2 > \tau$*.*

Theorem 8 can then be obtained as a direct corollary by setting $\tau = \varepsilon/\sqrt{d}$, recalling that when $\mathbf{q}$ is the uniform distribution $\mathbf{u}_d$, ℓ_2 distance and χ^2 divergence are equivalent:

$$\|\mathbf{p} - \mathbf{u}_d\|_2^2 = \sum_{i=1}^{d} (\mathbf{p}_i - 1/d)^2 = \frac{1}{d} \sum_{i=1}^{d} \frac{(\mathbf{p}_i - 1/d)^2}{1/d} = \frac{1}{d} \chi^2(\mathbf{p} \mid\mid \mathbf{u}_d)$$

We emphasize that the result of Theorem 9 itself is not new, as a quantum algorithm achieving the same sample complexity (in the same access model) was recently obtained by [24].[5] However, our algorithm differs significantly from the one in [24], and we believe it to be of independent interest for several reasons:

- it is *conceptually very simple*: (classically) hash the domain down to *two* elements, and use QME to estimate the bias of the resulting Bernoulli;
- it neatly *separates the quantum and classical aspects* of the task, only using QME (as a blackbox) in a single step of the algorithm;
- in contrast to the algorithm of [24], it *decouples the use of the source code from* $\mathbf{p}$ *and* $\mathbf{q}$, allowing one to run our algorithm when the accesses to the two distributions are on different machines, locations, or even will be granted at different points in time (i.e., one can run part of the algorithm using the source code for $\mathbf{p}$, and, one continent and a year apart, run the remaining part on the now-available source code for $\mathbf{q}$ without needing $\mathbf{p}$ anymore).

The idea behind Theorem 9 is relatively simple: previous work (in the classical setting) showed that hashing the domain from d to a much smaller $d' \ll d$ could yield sample-optimal testing algorithms in some settings, e.g., when testing under privacy bandwidth, or memory constraints. Indeed, while this "domain compression" reduces the total variation distance by a factor $\Theta(\sqrt{d'/d})$, this shrinkage is, in these settings, balanced by the reduction in domain size. The key insight in our algorithm is then to (1) use this hashing with respect to ℓ_2 distance, not total variation distance, and show that one can in this case get a two-sided guarantee in the distance (low-distortion embedding) instead of a one-sided one; and (2) compress the domain all the way to $d' = 2$, so that one can then invoke the QME algorithm to simply estimate the bias of a coin to an additive $\pm\tau$, a task for which a quantum quadratic speedup is well known.

Proof of Theorem 9. As mentioned above, a key building block of our algorithm is the following "binary hashing lemma," a simple case of the domain compression primitive of [3]:

[5] Technically, [24]'s result can be seen as slightly stronger, in that it allows to test $\|\mathbf{p} - \mathbf{q}\|_2 \leq (1-\gamma)\tau$ vs. $\|\mathbf{p} - \mathbf{q}\|_2 \mathbf{u}_d > \tau$, for arbitrarily small constant $\gamma > 0$.

▶ **Lemma 10** (Random Binary Hashing (Lemma 2.9 and Remark 2.4 of [11]). *Let $\mathbf{p}, \mathbf{q} \in \Delta(d)$. Then, for every $\alpha \in [0, 1/2]$,*

$$\Pr_S[\,|\mathbf{p}(S) - \mathbf{q}(S)| \geq \alpha \|\mathbf{p} - \mathbf{q}\|_2\,] \geq \frac{1}{12}(1 - 4\alpha^2)^2\,,$$

where $S \subseteq [d]$ is a uniformly random subset of $[d]$.

Given our goal of tolerant testing, we also require a converse to Lemma 10, stated and proven below:

▶ **Lemma 11.** *Let $\mathbf{p}, \mathbf{q} \in \Delta(d)$. Then, for every $\beta \in [1/2, \infty)$,*

$$\Pr_S[\,|\mathbf{p}(S) - \mathbf{q}(S)| \geq \beta \|\mathbf{p} - \mathbf{q}\|_2\,] \leq \frac{1}{4\beta^2}\,,$$

where $S \subseteq [d]$ is a uniformly random subset of $[d]$.

Proof. As in the proof of Lemma 10, we write $\delta := \mathbf{p} - \mathbf{q} \in \mathbb{R}^d$ and $\mathbf{p}(S) - \mathbf{q}(S) = \frac{1}{2}Z$, where $Z := \sum_{i=1}^d \delta_i \xi_i$ for $\xi_1, \dots, \xi_d$ i.i.d. Rademacher. We will use the following fact established in the proof of this lemma, which we reproduce for completeness:

$$\mathbb{E}\big[Z^2\big] = \sum_{1 \leq i,j \leq d} \delta_i \delta_j \mathbb{E}[\xi_i \xi_j] = \sum_{i=1}^d \delta_i^2 = \|\delta\|_2^2\,. \tag{49}$$

By Markov's inequality, we then have

$$\Pr_S[\,|\mathbf{p}(S) - \mathbf{q}(S)| > \beta \|\mathbf{p} - \mathbf{q}\|_2\,] = \Pr_S\big[\,Z^2 > 4\beta^2 \mathbb{E}\big[Z^2\big]\,\big] \leq \frac{1}{4\beta^2}$$

concluding the proof. ◀

While the above two lemmas allow us to obtain a slightly more general result than in the theorem statement by keeping α, β as free parameters, for concreteness, set $\alpha := 1/(2\sqrt{2})$ and $\beta = 4$. This implies the following:

- If $\|\mathbf{p} - \mathbf{q}\|_2 \geq \tau$, then
$$\Pr_S\left[\left|\mathbf{p}(S) - \frac{|S|}{d}\right| \geq \frac{\tau}{\sqrt{8}}\right] \geq \frac{1}{48}$$
- If $\|\mathbf{p} - \mathbf{q}\|_2 \leq \frac{\tau}{12}$, then
$$\Pr_S\left[\left|\mathbf{p}(S) - \frac{|S|}{d}\right| \geq \frac{\tau}{\sqrt{9}}\right] \leq \frac{1}{64}\,.$$

where $S \subseteq [d]$ is a uniformly random subset of $[d]$. This allows us to distinguish between the two cases with only $O(1)$ repetitions:

Algorithm 2 QME+Binary Hashing Tester.

```
Set $T = O(1)$, $\delta := \frac{1}{600}$, $\tau := \frac{1/48+1/64}{2}$.          ▷ $\delta \leq \frac{1}{3}(\frac{1}{48} - \frac{1}{64})$.
for $t = 1$ to $T$ do
    Pick a u.a.r. subset $S_t \subseteq [d]$ (independently of previous iterations)
    Estimate $\mathbf{p}(S_t), \mathbf{q}(S_t)$ by $\hat{p}_t, \hat{q}_t$ to within $\pm\frac{\tau}{100}$ with error probability $\delta$.   ▷ QME
    if $|\hat{p}_t - \hat{q}_t| \leq \frac{\varepsilon}{\sqrt{8d}}$ then $b_t \leftarrow 0$
    else $b_t \leftarrow 1$
return accept if $\frac{1}{T}\sum_{t=1}^T b_t \leq \tau$          ▷ Estimate of the probability accept
```

A standard analysis shows that, for T a sufficiently large constant, with probability at least $2/3$ the estimate $\frac{1}{T}\sum_{t=1}^{T} b_t$ will be within an additive $\delta + \frac{1}{1000}$ of the corresponding value (either $1/48$ or $1/64$), in which case the output is correct. The total number of samples required is T times the sample of the Quantum Mean Estimation call on Line 4, which is $O(1/\tau)$: the complexity of getting a $O(\tau)$-additive estimate of the mean of a Bernoulli random variable with high (constant) probability. This concludes the proof. ◀

B Algorithm in the Giant Distance Regime

In this appendix, we show that, in the (stronger) quantum string oracle model, one can perform tolerant uniformity testing with respect to χ^2 divergence in the "very large parameter regime," that is, to distinguish $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq c\theta$ from $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) > \theta$ for $\theta \geq 1$:

▶ **Theorem 12.** *There is a computationally efficient quantum algorithm for uniformity testing with the following guarantees: For $\theta \geq 1$, the algorithm makes $O(d^{1/3}/\theta^{1/3})$ calls to the quantum string oracle for an unknown distribution $\mathbf{p}$ over $[d]$, and distinguishes with probability at least .99 between*

$$(1)\ \chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq c \cdot \theta, \qquad \text{and} \qquad (2)\ \chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) > \theta\,, \tag{50}$$

where $c > 0$ is an absolute constant. Moreover, this query complexity is optimal.

Note that, as discussed in the introduction, this result does not imply anything in terms of total variation distance, as the latter is always at most 1; however, we believe this result to be of interest for at least two reasons: (1) it is in itself a reasonable (and often useful) testing question, when total variation distance is not the most relevant distance measure, and implies, for instance, testing $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq c \cdot \theta$ from $\mathrm{KL}(\mathbf{p} \mid\mid \mathbf{u}_d) > \theta$; and (2) one can show that this complexity is tight, by a reduction to the θ-to-1 collision problem, which provides additional evidence for Conjecture 2.

Proof. The main ingredient of the proof is the following lemma, which guarantees that taking $N = \Theta(d/\theta)$ from the unknown distribution $\mathbf{p}$ is enough to obtain (with high constant probability) a multiset of elements with, in one case, no collisions, and in the other at least one collision:

▶ **Lemma 13.** *For $\theta \geq 1$, there exists a constant $c \in (0,1)$ such that taking N i.i.d. samples from an unknown $\mathbf{p}$ over $[d]$ results in a multiset S satisfying the following with probability at least .99:*

- *If $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq c \cdot \theta$, then all elements in S are distinct;*
- *If $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \geq \theta$, then at least two elements in S are identical;*

as long as $1601 \cdot \frac{d}{\theta} \leq N \leq \frac{1}{10c} \cdot \frac{d}{\theta}$. (In particular, taking $c := \frac{1}{16010}$ suffices to ensure such a choice of N is possible.)

Before proving this lemma, we describe how it implies our stated complexity upper bound. Lemma 13 guarantees that we can reduce our testing problem to that of deciding if, given oracle access to a string of size $N = \Theta(\sqrt{d/\theta})$, whether all the elements in it are distinct. This problem is solved by Ambainis' element distinctness quantum-walk algorithm [5] using $O(N^{2/3}) = O(d^{1/3}/\theta^{1/3})$ quantum queries.

Proof of Lemma 13. Suppose we take N i.i.d. samples $X_1, \dots, X_N$ from $\mathbf{p}$, and count the number Z of collisions among them:

$$Z := \sum_{1 \leq i < j \leq N} \mathbb{1}_{\{X_i = X_j\}}$$

Letting $\delta \coloneqq \mathbf{p} - \mathbf{u}_d$ and $\operatorname{pow}_t(x) \coloneqq \sum_{i=1}^d x_i^t$ for all integer $t \geq 0$ and vector $x \in \mathbb{R}^d$ (so that $\delta_i = \mathbf{p}_i - 1/d$ for all i), we have, $\operatorname{pow}_1(\delta) = 0$, and

$$\operatorname{pow}_2(\delta) = \|\mathbf{p} - \mathbf{u}_d\|_2^2 = \frac{1}{d}\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d)$$

Now, it is not hard to verify that $\mathbb{E}[Z] = \binom{N}{2}\|\mathbf{p}\|_2^2 = \binom{N}{2}(\operatorname{pow}_2(\delta) + 1/d)$, and

$$\begin{aligned} \operatorname{Var}[Z] &= \binom{N}{2}\|\mathbf{p}\|_2^2\Big(1 - \|\mathbf{p}\|_2^2\Big) + 6\binom{N}{3}\Big(\|\mathbf{p}\|_3^3 - \|\mathbf{p}\|_2^4\Big) \\ &\leq \mathbb{E}[Z] + 6\binom{N}{3}\left(\operatorname{pow}_3(\delta) + \frac{3}{d}\operatorname{pow}_2(\delta)\right) \end{aligned} \tag{51}$$

From this, we get, setting $\tau \coloneqq \sqrt{\theta/d} \geq 1/\sqrt{d}$:

- If $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \leq c \cdot \theta$, then $\operatorname{pow}_2(\delta) \leq c^2 \cdot \tau^2$, and as long as $N \leq \frac{1}{10c\tau}$ we have $\binom{N}{2}(c^2 \cdot \tau^2 + 1/d) \leq 1/100$, so that by Markov's inequality

$$\Pr[\, Z \geq 1\,] \leq \Pr[\, Z \geq 100\mathbb{E}[Z]\,] \leq \frac{1}{100}$$

- If $\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) \geq \theta$, then $\operatorname{pow}_2(\delta) \geq \tau^2$, and by Chebyshev's inequality and Equation (51)

$$\begin{aligned} \Pr[\, Z = 0\,] \leq \Pr[\, |Z - \mathbb{E}[Z]| \geq \mathbb{E}[Z]\,] &\leq \frac{1}{\mathbb{E}[Z]} + \frac{4}{N} \cdot \frac{\operatorname{pow}_3(\delta) + \frac{3}{d}\operatorname{pow}_2(\delta)}{(\operatorname{pow}_2(\delta) + 1/d)^2} \\ &\leq \frac{2}{N(N-1)\tau^2} + \frac{4}{N} \cdot \frac{\operatorname{pow}_2(\delta)^{3/2} + \frac{3}{d}\operatorname{pow}_2(\delta)}{\operatorname{pow}_2(\delta)^2} \\ &\leq \frac{3}{N^2\tau^2} + \frac{4}{N\tau} + \frac{12}{Nd\tau^2} \\ &\leq \frac{3}{N^2\tau^2} + \frac{4}{N\tau} + \frac{12}{N} \qquad (\tau \geq 1/\sqrt{d}) \\ &\leq \frac{3}{N^2\tau^2} + \frac{16}{N\tau} \end{aligned}$$

  which is at most $\frac{1}{100}$ for $N \geq \frac{1601}{\tau}$.

This proves the lemma. ◀

This concludes the proof of the upper bound part of Theorem 12. To conclude, it only remains to show that this is, indeed, optimal. For this, we need a lower bound of [23], which generalized a lower bound of Aaronson and Shi [2]:

▶ **Theorem 14** ([23]). *Let $d > 0$ and $r \geq 2$ be integers such that $r|d$, and let $f : [d] \to [d]$ be a function to which we have quantum oracle access. Then deciding if f is 1-to-1 or r-to-1, promised that one of these holds, requires $\Omega((d/r)^{1/3})$ quantum queries.*

When we view this function as a quantum string oracle for a probability distribution, the function being 1-to-1 corresponds to the uniform distribution on $[d]$. In the other case, the distribution is uniform on a subset of size $[d/r]$, for any $r \geq \theta + 1$ dividing d. An easy calculation shows that the second distribution is at χ^2 divergence

$$\chi^2(\mathbf{p} \mid\mid \mathbf{u}_d) = \sum_{i \in [d]} \left(\frac{\mathbf{p}_i^2}{1/d}\right) - 1 = d \cdot \frac{r^2}{d^2} \cdot \frac{d}{r} - 1 = r - 1 \geq \theta, \tag{52}$$

from uniform, which completes the proof. ◀

C Reduction from Identity to Uniformity Testing

As mentioned in the introduction, there is a known reduction from identity to uniformity testing, due to Goldreich [16] and inspired by [13]: which, in a blackbox way, converts an instance of uniformity testing (in total variation distance) with reference distribution $\mathbf{q}$ over $[d]$ and distance parameter ε to an instance of uniformity testing over $[4d]$ and distance parameter $\varepsilon/4$. (Here, we follow the exposition and parameter setting of [11, Section 2.2.3].)

To be able to use it in our setting, all we need to check is that this blackbox reduction $\Phi_{\mathbf{q}}$ preserves access to "the code": that is, given the code $C_{\mathbf{p}}$ for a probability distribution $\mathbf{p}$ over $[d]$, that we can efficiently have access to the code $C_{\mathbf{p}'}$ for the resulting distribution $\mathbf{p}' = \Phi_{\mathbf{q}}(\mathbf{p})$ over $[4d]$. To do so, note that $\Phi_{\mathbf{q}}$ is the composition of 3 successive mappings,

$$\Phi_{\mathbf{q}} = \Phi_{\mathbf{q}}^{(1)} \circ \Phi_{\mathbf{q}}^{(2)} \circ \Phi_{\mathbf{q}}^{(3)}$$

where $\Phi_{\mathbf{q}}^{(3)} \colon [d] \to [d]$, $\Phi_{\mathbf{q}}^{(2)} \colon [d] \to [d+1]$, and , $\Phi_{\mathbf{q}}^{(2)} \colon [d+1] \to [4d]$. So it suffices to show that each of these 3 mappings does preserve access to the code generating a sample from the resulting distribution.

- The first, $\Phi_{\mathbf{q}}^{(3)}$, is the easier, as it consists only in mixing its input with the uniform distribution:

  $$\Phi_{\mathbf{q}}^{(3)}(\mathbf{p}) = \frac{1}{2}\mathbf{p} + \frac{1}{2}\mathbf{u}_d$$

  for which a circuit can be easily obtained, given a circuit for $\mathbf{p}$.
- The second, $\Phi_{\mathbf{q}}^{(2)}$, "rounds down" the probability of each of the d elements of the domain, and sends the remaining probability mass to a $(d+1)$-th new element:

  $$\Phi_{\mathbf{q}}^{(2)}(\mathbf{p})_i = \begin{cases} \frac{\lfloor 4d\mathbf{q}_i \rfloor}{4d\mathbf{q}_i} \cdot \mathbf{p}_i, & i \in [d] \\ 1 - \sum_{i=1}^{d} \frac{\lfloor 4d\mathbf{q}_i \rfloor}{4d\mathbf{q}_i} \cdot \mathbf{p}_i, & i = d+1 \end{cases}$$

  This corresponds to adding to the circuit $C_{\mathbf{p}}$ for $\mathbf{p}$ a "postprocessing circuit" which, if the output of $C_{\mathbf{p}}$ is i, outputs i with probability $\frac{\lfloor 4d\mathbf{q}_i \rfloor}{4d\mathbf{q}_i}$ (and $d+1$ otherwise).
- The third, $\Phi_{\mathbf{q}}^{(1)}$, assumes that the reference distribution $\mathbf{q}$ is "grained" (namely, all its probabilities are positive multiples of $1/(4d)$), which will be the case after the first two mappings[6] fully known). Having partitioned $[4d]$ in sets $S_1, \dots, S_d$ where

  $$|S_i| = 4d \cdot \mathbf{q}_i \geq 1$$

  and $\Phi_{\mathbf{q}}^{(1)}$ is given by

  $$\Phi_{\mathbf{q}}^{(3)}(\mathbf{p})_i = \sum_{j=1}^{d} \frac{\mathbf{p}_i}{|S_i|} \mathbb{1}_{\{j \in S_i\}}, \qquad i \in [4d]\,.$$

  This corresponds to adding to the circuit $C_{\mathbf{p}}$ for $\mathbf{p}$ a "postprocessing circuit" which, if the output of $C_{\mathbf{p}}$ is i, outputs an element of S_i uniformly at random. (Importantly, $S_1, \dots, S_d$ are uniquely determined by $\mathbf{q}$, and do not depend on $\mathbf{p}$ or $C_{\mathbf{p}}$ at all.)

To summarize, each of these three mappings can be implemented to provide, given a circuit $C_{\mathbf{p}}$ for $\mathbf{p}$, a circuit $C_{\mathbf{p}'}$ for the output $\mathbf{p}'$, so that altogether the reduction can be implemented in a way which preserves access to "the code."

[6] Specifically, when chaining the three mappings, the reference distribution called $\mathbf{q}$ here is actually $\Phi_{\mathbf{q}}^{(2)} \circ \Phi_{\mathbf{q}}^{(3)}(\mathbf{q})$.

Self-Testing in the Compiled Setting via Tilted-CHSH Inequalities

Arthur Mehta

Department of Mathematics and Statistics, University of Ottawa, Canada

Connor Paddock

Department of Mathematics and Statistics, University of Ottawa, Canada

Lewis Wooltorton

Department of Mathematics, University of York, UK
Quantum Engineering Centre for Doctoral Training, H. H. Wills Physics Laboratory and
Department of Electrical & Electronic Engineering, University of Bristol, UK
Inria, ENS de Lyon, LIP, France

Abstract

This work investigates the family of extended tilted-CHSH inequalities in the single-prover cryptographic compiled setting. In particular, we show that a quantum polynomial-time prover can violate these Bell inequalities by at most negligibly more than the violation achieved by two non-communicating quantum provers. To obtain this result, we extend a sum-of-squares technique to monomials with arbitrarily high degree in the Bob operators and degree at most one in the Alice operators. We also introduce a notion of partial self-testing for the compiled setting, which resembles a weaker form of self-testing in the bipartite setting. As opposed to certifying the full model, partial self-testing attempts to certify the reduced states and measurements on separate subsystems. In the compiled setting, this is akin to the states after the first round of interaction and measurements made on that state. Lastly, we show that the extended tilted-CHSH inequalities satisfy this notion of a compiled self-test.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography

Keywords and phrases Compiled Bell scenarios, self-testing

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.8

Related Version *Full Version*: `https://arxiv.org/abs/2406.04986` [24]

Funding *Arthur Mehta*: NSERC Alliance Consortia Quantum grants, reference number: ALLRP 578455 – 22 and the NSERC Discovery Grants Program.
Connor Paddock: Digital Horizon Europe project FoQaCiA, Foundations of quantum computational advantage, grant no. 101070558, and the Natural Sciences and Engineering Research Council of Canada (NSERC).
Lewis Wooltorton: Engineering and Physical Sciences Research Council (EPSRC Grant No. EP/SO23607/1) and the European Union's Horizon Europe research and innovation programme under the project "Quantum Security Networks Partnership" (QSNP, grant agreement No. 101114043).

Acknowledgements The authors thank Simon Schmidt, Ivan Šupić, Anand Natarajan, and Tina Zhang for helpful discussions. We also thank the anonymous TQC reviewers for valuable feedback.

1 Introduction

In a bipartite Bell scenario, two non-communicating provers receive inputs x and y and reply with outputs a and b to a verifier. The collection of probabilities of observing outcomes (a, b) given (x, y) determines a correlation $\mathsf{p} = \{p(a, b|x, y)\}$. Bell's celebrated theorem implies that if the provers are permitted to share an entangled quantum state and make

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 8; pp. 8:1–8:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

local quantum measurements, called a bipartite (quantum) model, then certain correlations have no realization by a classical (or local hidden-variable) model [7]. The distinction between quantum and classical correlations is often explored through Bell inequalities. A Bell inequality is a linear inequality on the set of correlations which is satisfied by all classical correlations. Hence, these inequalities can be violated by certain models using quantum entanglement, realizing correlations that are not classical. The quantum value of a Bell inequality refers to the largest violation achievable by a bipartite (quantum) model. A prominent example is the Clauser-Horne-Shimony-Holt (CHSH) inequality, where the classical bound is 2, but the quantum value is $2\sqrt{2}$ [11].

Due to their ability to witness these non-classical effects, Bell inequality violations play a major role in areas like device-independent cryptography [1, 15, 30, 31, 33], protocols for verifiable delegated quantum computation [32, 17, 14], and in the study of multiprover interactive proofs (MIPs) and the variant MIP* with entangled provers [12], also called nonlocal games. Many of the key applications of Bell inequalities rely on a remarkable property known as *self-testing* [22, 23, 35, 34]. Informally, a Bell inequality is a self-test for an ideal bipartite (quantum) model Q if there exist local isometries which transform any employed bipartite model Q′ achieving maximum Bell violation into the ideal model Q. It is well-known that the CHSH inequality is a self-test for the bipartite model employing a maximally entangled state on two qubits, along with the Pauli σ_x and σ_z measurements, among others [22]. Another prominent example is the family of tilted-CHSH inequalities [2, 35, 4], which self-test partially entangled two-qubit states, and were integral in the work of Coladangelo, Goh, and Scarani who employed them as part of a protocol to self-test any pure bipartite entangled state [13].

Despite the enormous success of self-testing, a practical drawback is the requirement of multiple non-communicating quantum provers. Recently, a number of cryptographic approaches have been proposed that replace the non-communication assumption with computational assumptions [19, 26, 18]. This makes the setting more practical by having a single quantum prover, rather than multiple. One new and prominent approach is the Kalai-Lombardi-Vaikuntanathan-Yang (KLVY) compilation procedure introduced in [19], which transforms a 2-prover 1-round Bell scenario into a 1-prover 2-round scenario with a single computationally bounded prover. The core ingredient in the KLVY compilation procedure is quantum homomorphic encryption (QHE), which emulates, to a certain extent, the non-communication between the rounds of interaction. In the compiled game, the inputs to the prover happens sequentially. In the first round, the prover obtains an encryption χ of the input x from the verifier. Without breaking the security, the prover cannot distinguish between encryptions of different inputs. The prover performs a polynomial time quantum circuit on χ, and then returns an output α to the verifier. In the second round, the information about x has already been "hidden" from the prover, so the verifier can send input y in the plain (i.e. unencrypted) to the prover, upon which the prover can perform a measurement and return outcome b to the verifier. The verifier checks for a Bell inequality violation (across many such interactions) using the values of x, the decryption of α, along with (y, b). QHE has two key features that makes this resemble the bipartite setting. Firstly, it allows the first round quantum prover to perform measurements as they would have in the bipartite setting, without knowing the input. Secondly, the encryption ensures that no classical polynomial-time prover can violate a Bell inequality by more than an negligible amount (see Section 3 details). Both of these are non-trivial and were the subject of [19].

In a follow-up work, Natarajan and Zhang showed that the maximal quantum violation of the CHSH inequality in the compiled setting is bounded by the maximal violation in the bipartite setting, up to negligible factors in the security parameter [27]. Subsequent works have analyzed the quantum soundness of the KLVY compilation procedure for other multiprover scenarious, including all 2-player XOR nonlocal games [16], Bell inequalities tailored to maximally entangled bipartite states [6], delegated quantum computation with a single-device [27, 25], and even in the study of contextuality [3]. Despite these advancements, many results have yet to be reproduced in the compiled setting. Our work takes another step in growing the list of protocols that will function as desired in the compiled setting.

Upper bounding compiled Bell violations

As mentioned, the compiled value of a Bell inequality is always at least the quantum value. This is because any bipartite (quantum) model can be implemented with homomorphic encryption via a *correctness* property of the QHE scheme used in the procedure. On the other hand, establishing upper bounds on the largest violation possible in the compiled setting is challenging, as general techniques for bounding these violations depend on the spatial separation between the two provers. Nonetheless, upper bounds on the violations of a certain Bell inequalities in the compiled setting can be verified using the sum-of-squares (SOS) technique [27, 16, 6]. The SOS approach is a powerful method and has been used extensively to upper bound Bell inequality violations and the values of nonlocal games in the bipartite setting. Informally, this technique relates the maximum compiled value η, of a Bell functional I, to a decomposition of the Bell operator or Bell polynomial S as a sum of Hermitian squares, $\eta\mathbb{I} - S = \sum_i P_i^\dagger P_i$. Before our work, progress was made on realizing this approach in the compiled setting, however, there were some limitations. In particular, it was required that the polynomials P_i involved in the decomposition were at most degree two in both Alice's and Bob's observables, restricting the technique to Bell inequalities with an SOS decomposition of this form; this excludes, for example, the family of tilted-CHSH inequalities.

Our first result extends the SOS technique to a larger family of Bell polynomials. More specifically, we extend the pseudo-expectation techniques in [27, 16] to allow for evaluations on polynomial terms P_i that consist of arbitrary monomials in the algebra generated by Bob's observables. In Theorem 3 we prove that an extended pseudo-expectation will be positive on the corresponding Hermitian square $P_i^\dagger P_i$ for any such term P_i. Consequently, we show that for any Bell inequality with an SOS decomposition in which P_i are of the form $P_i = \sum_j \gamma_j (A_x)^{k_j} w_j(B)$ for some $\gamma_j \in \mathbb{C}$, $k_j \in \{0, 1\}$ and $w_j(B)$ being arbitrary monomials in Bob's observables, η is an upper-bound on the maximum compiled quantum value. Our extension captures a wide class of Bell inequalities including tilted-CHSH, enabling us to bound the compiled value of the tilted-CHSH inequalities, by the quantum value and a negligible function of security parameter, see Theorem 5 for details.

A compiled self-testing result

Our second contribution is a concept of self-testing in the compiled setting. One of the main obstacles to deriving self-testing results in the compiled setting is the lack of techniques for extracting any algebraic relations on the measurement operators acting under the encryption. Nevertheless, it remains possible to derive relations on the observables in the second round. With this in mind, we consider a partial notion of self-testing that applies to the measurements made by the prover in the second round. In particular, our definition only requires the existence of an isometry robustly certifying the ideal post-measurement state after the first round, and the action of the measurements made in the second.

As our final result, we provide an example by showing violations of the compiled tilted-CHSH inequalities satisfy this notion of partial self-testing. This family of inequalities was introduced by Acín, Massar, and Pironio [2], and the Bell functionals take the form $\alpha_\theta\langle A_0\rangle + \langle A_0B_0\rangle + \langle A_0B_1\rangle + \langle A_1B_0\rangle - \langle A_1B_1\rangle$, where $\langle A_xB_y\rangle$, $\langle A_x\rangle$ denote the expectation of measurements corresponding to settings $X = x$, $Y = y$, and $\alpha_\theta \in \mathbb{R}$. Notably, they are tailored to robustly self-test the two qubit states $\cos(\theta)|00\rangle + \sin(\theta)|11\rangle$ [35, 4], and were used as part of a more complex protocol to obtain self-testing for all pure bipartite entangled states [13]. The work of Barizien, Sekatski, and Bancal [5] extended this family to include extra degrees of freedom in Bob's measurements, which we will refer to as "extended" tilted-CHSH inequalities.

We apply Theorem 3 to the SOS decomposition for the extended tilted-CHSH inequalities presented in [5]. Specifically, in Theorem 5 we prove that the maximum quantum value achieved for any of the extended tilted-CHSH functionals is preserved by the KLVY compilation procedure. Then in Theorem 12 we use this same decomposition to prove that this family of games is a compiled self-test according to Definition 11.

Related work

A recent work of [20] implies that the compiled value of any 2-prover Bell scenario is bounded by the largest violation possible among so-called *commuting operator* models. However, unlike some previous results, such as [6, 16], the upper bound in [20] lacks a dependence on the security parameter λ, making it unclear how the compiled value is related to the quantum value at fixed security parameters. Hence, results such as ours, which obtain a bound on the compiled value that depends negligibly on the security parameter, remain of great importance. Furthermore, [20] also considers a notion of self-testing in the compiled setting, however, due to their methods the results are in terms of commuting operator self-tests (as defined in [29, Proposition 7.8]) and only hold in the limit of the security parameter $\lambda \to \infty$.

Another related work is [26], which presents a protocol for certifying that an unknown computationally bounded device has prepared a maximally entangled pair of qubits, and whether a measurement was performed on each qubit in either the computational or Hadamard basis. The techniques used to prove our compiled self-test have similarities to those of [26], particularly in the choice of isometry (see Definition 11) and proof structure, which in turn resembles self-testing techniques in the bipartite setting [4]. There are however some key differences. Firstly, [26] certifies the preparation of a maximally entangled state by the device before any measurements are made. While our results are tailored to the more general class of partially entangled states, we only make statements about the post-measurement states after each round. It is an interesting open question if our results can be extended in this way (see Section 4.1 for more details), and statements weaker than certifying the prepared state could also be possible. For example, can a compiled self-test be used to show the prepared state must have been entangled? Another significant difference to [26] is that the self-testing protocol in this work strongly resembles the bipartite case, owing to the compilation procedure mapping bipartite nonlocal scenarios to single prover scenarios. Our main result can therefore be interpreted as translating a self-testing statement in the Bell scenario to one in the compiled Bell scenario. On the other hand, the authors of [26] describe their approach as more "custom", guided by the available cryptographic primitives, and pose the open question of finding a general procedure for translating self-testing results from the nonlocal setting. We showed this is possible for the special case of titled-CHSH inequalities.

Future outlook

Moving forward, we consider several natural directions for following up on this work:

1. Tilted-CHSH inequalities were an integral component of the self-testing for all pure bipartite entangled states [13]. Building off of our work on compiled tilted-CHSH inequalities, a natural question is whether similar results can be obtained in the compiled setting.
2. It would be desirable to understand the fundamental limitations of our notion of self-testing and other similar notions such as the computational self-testing given in [26]. Furthermore, is a finer notion of self-testing in the compiled setting that characterizes both Alice's and Bob's operators and the initial state possible without specifying the underlying QHE scheme? Moreover, is every self-test in the standard Bell scenario also a compiled self-test, and vice-versa?
3. Many current techniques for bounding the value of compiled nonlocal games/Bell inequalities can be obtained using some variant of the sum-of-squares decomposition approach. Given our improvements to this approach outlined in Theorem 3, it is possible to search for valid decompositions which include arbitrary words in Bob's operators. Is it possible to use this approach to give a limited variant of the NPA hierarchy [28] in the compiled setting?

2 Background

2.1 Mathematical notation

Throughout the article, Hilbert spaces are denoted by $\mathcal{H}$, and are assumed to be finite-dimensional unless explicitly stated otherwise. Elements of $\mathcal{H}$ are denoted by $|v\rangle \in \mathcal{H}$, where the inner product $\langle u|v\rangle$ for $|v\rangle, |u\rangle \in \mathcal{H}$ is linear in the second argument and defines the vector norm $\||v\rangle\| = \sqrt{\langle v|v\rangle}$. Quantum pure states are the norm 1 elements of $\mathcal{H}$. In this work, $\mathbb{B}(\mathcal{H})$ denotes the unital †-algebra of bounded linear operators on $\mathcal{H}$ with norm $\|M\|_{\text{op}}^2 = \sup_{|v\rangle\in\mathcal{H},|v\rangle\neq 0}\langle v|M^\dagger M|v\rangle/\langle v|v\rangle$. We also write $\|A\|_2 = \sqrt{\text{tr}(A^\dagger A)}$ to denote the Schatten 2-norm for $A \in \mathbb{B}(\mathbb{C}^d) \cong M_d(\mathbb{C})$. The unit in $\mathbb{B}(\mathcal{H})$ is denoted by $\mathbb{I}$, and we write $|M| = \sqrt{M^\dagger M}$ for the positive part of $M \in \mathbb{B}(\mathcal{H})$. Given a finite set $\mathcal{A}$, a collection of positive operators $\{M_a \geq 0 : a \in \mathcal{A}\}$ with the property that $\sum_{a\in\mathcal{A}} M_a = \mathbb{I}$, is called a POVM over $\mathcal{A}$. When the operators in a POVM are orthogonal projections, we call it a PVM. Given a random variable X, which takes values $X = x \in \mathcal{X}$ according to a distribution $\mu : \mathcal{X} \to \mathbb{R}_{\geq 0}$ such that $\sum_{x\in X}\mu(x) = 1$, we denote the expectation of X by $\mathbb{E}[X] = \sum_{x\in\mathcal{X}}\mu(x)\cdot x$. For $a, b \in \mathbb{R}$ and $\delta > 0$, $a \approx_\delta b$ is short for $|a - b| \leq \delta$. A function negl : $\mathbb{N} \to \mathbb{R}$ is called negligible if for all $k \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for every $n \geq N$ it holds that $\text{negl}(n) \leq \frac{1}{n^k}$.

2.2 Bell scenarios, inequalities, and violations

Before we discuss compiled Bell inequalities, let us recall the bipartite case. Here we let $\mathcal{A}, \mathcal{B}, \mathcal{X}$, and $\mathcal{Y}$ be finite sets, with $|\mathcal{A}| = m_A$, $|\mathcal{B}| = m_B$, $|\mathcal{X}| = n_A$, and $|\mathcal{Y}| = n_B$. A bipartite Bell scenario is described by the tuple $\mathcal{S} = (\mathcal{A}, \mathcal{B}, \mathcal{X}, \mathcal{Y}, \pi)$, where $\pi : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}_{\geq 0}$ is a distribution over the measurement settings. In a scenario, each party receives an input $x \in \mathcal{X}$ (resp. $y \in \mathcal{Y}$) sampled according to π, and returns outputs $a \in \mathcal{A}$ (resp. $b \in \mathcal{B}$). The parties are non-communicating, and therefore cannot coordinate their outputs. The behaviour of the provers is characterized by a correlation, a set of conditional probabilities

$\mathsf{p} = \{p(a,b|x,y) : a \in \mathcal{A}, b \in \mathcal{B}, x \in \mathcal{X}, y \in \mathcal{Y}\}$, which is realized by an underlying physical theory or model. In the quantum setting, we allow the provers to share a bipartite quantum state, and say the correlation p is realized by a **bipartite (quantum) model**

$$\mathrm{Q} = \big(\mathcal{H}_A, \mathcal{H}_B, \{\{M_{a|x}\}_{a\in\mathcal{A}}\}_{x\in\mathcal{X}}, \{\{N_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}}, |\Psi\rangle_{AB}\big), \tag{1}$$

where $\mathcal{H}_A$ and $\mathcal{H}_B$ are Hilbert spaces, $\{M_{a|x}\}_{a\in\mathcal{A}}$ and $\{N_{b|y}\}_{b\in\mathcal{B}}$ are POVMs on $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively, and $|\Psi\rangle_{AB}$ is a vector state in $\mathcal{H}_A \otimes \mathcal{H}_B$. More generally, a correlation p is quantum (or an element of $C_{\mathrm{q}}(n_A, n_B, m_A, m_B)$) if there exists a bipartite model Q for which p can be realized via the Born rule as $p(a,b|x,y) = \langle\Psi|M_{a|x} \otimes N_{b|y}|\Psi\rangle$. We denote the class of bipartite (quantum) models by $\mathcal{Q}(n_A, n_B, m_A, m_B)$. From now on we will refer to such models simply as **bipartite models**.

In contrast to the set of quantum correlations, we have the collection of local correlations $C_{\mathrm{loc}}(n_A, n_B, m_A, m_B)$. These are the correlations $\{p(a,b|x,y)\}$ for which there exists a **classical model**, that is a probability distribution μ_k and a local distributions $p_k^A(a|x)$ and $p_k^B(b|y)$ such that $p(a,b|x,y) = \sum_k \mu_k\, p_k^A(a|x)\, p_k^B(b|y)$. We let $\mathrm{C} = (\mu_k, \{p_k^A\}, \{p_k^B\})$ denote a **classical model** and let $\mathcal{L}(n_A, n_B, m_A, m_B)$ denote the class of all classical models. In what follows we consider Bell scenarios where $n_A = n_B = n$, and $m_A = m_B = m$. With this notation Bell's theorem [7] states that $C_{\mathrm{loc}}(2,2)$ is a strict subset of $C_{\mathrm{q}}(2,2)$.

Given a Bell scenario $\mathcal{S}$, one can consider a linear (or Bell) functional on the set of correlations

$$I = \sum_{a\in\mathcal{A},b\in\mathcal{B},x\in\mathcal{X},y\in\mathcal{Y}} w_{abxy}\, p(a,b|x,y), \tag{2}$$

for coefficients $w_{abxy} \in \mathbb{R}$. A **Bell inequality** is a functional I and a bound $\eta > 0$ such that $I \leq \eta$ for all $\mathsf{p} \in C_{\mathrm{loc}}(n,m)$. Given a functional I, the classical value is the maximal value achieved by the classical correlations $\mathsf{p} \in C_{\mathrm{loc}}(n,m)$. We denote this value by $\eta^{\mathrm{L}} := \sup_{\mathsf{p}\in C_{\mathrm{loc}}(m,n)} I$. The quantum value for I is the maximal value achieved by the set of quantum correlations $\mathsf{p} \in C_{\mathrm{q}}(m,n)$, and we denote the quantum value on I by $\eta^{\mathrm{Q}} := \sup_{\mathsf{p}\in C_{\mathrm{q}}(m,n)} I$. Hence, a Bell violation occurs whenever there is a $\mathsf{p} \in C_{\mathrm{q}}(m,n)$ for which $I > \eta^{\mathrm{L}}$. A violation of a Bell inequality by non-communicating provers employing a quantum model is an indication of entanglement between provers.

Typically when η^{L} is known for a given I, the main challenge is finding an upper bound on η^{Q}. In this case, one often considers the **Bell operator**[1] $S = \sum_{abxy} w_{abxy}\, M_{a|x} \otimes N_{b|y}$, and $\langle S\rangle = \langle\Psi|S|\Psi\rangle$ its quantum expectation with respect to $|\Psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$. Since bipartite models with separable quantum states generate the classical correlations $C_{\mathrm{q}}(m,n)$, $\langle\Psi|S|\Psi\rangle \leq \eta^{\mathrm{L}}$ whenever $|\Psi\rangle$ is separable (unentangled). However, it's possible that there could be entangled states for which $\langle\Psi'|S|\Psi'\rangle > \eta^{\mathrm{L}}$. Hence, given a Bell operator S, we can recover the maximum classical and quantum values $\eta^{\mathrm{L}} = \sup_{\mathrm{C}\in\mathcal{L}(n,m)}\langle S\rangle$ and $\eta^{\mathrm{Q}} = \sup_{\mathrm{Q}\in\mathcal{Q}(n,m)}\langle S\rangle$ respectively. Technically, we have not fixed the dimensions of the Bell operator as we want to consider any finite-dimensional model. Hence, the supremum is implicitly over all finite-dimensional Hilbert spaces $\mathcal{H}_A \otimes \mathcal{H}_B$.

An approach to establishing upper bounds on $\langle S\rangle$ is using sum-of-squares techniques. Let S be a Bell operator and $\eta' > 0$. The shifted Bell operator $\eta'\mathbb{I} - S$ admits a **sum-of-squares** (SOS) decomposition if there exists a set of polynomials $\{P_i\}_{i\in\mathcal{I}}$ in the elements $\{M_{a|x}, N_{b|y} : a \in \mathcal{A}, b \in \mathcal{B}, x \in \mathcal{X}, y \in \mathcal{Y}\}$ satisfying $\eta'\mathbb{I} - S = \sum_{i\in\mathcal{I}} P_i^\dagger P_i$. The existence of

[1] For a more mathematically rigorous treatment of Bell operators and the SOS approach consult [16].

an SOS decomposition for the operator $\eta'\mathbb{I} - S$ implies that $\eta'\mathbb{I} - S$ is positive, and therefore η' is an upper bound on the maximum quantum value of $\langle S \rangle$. Additionally, if η' is achievable by a bipartite model, then we write $\eta' = \eta^{\mathrm{Q}}$. In this case, the *shifted* Bell operator is $\bar{S} = \eta^{\mathrm{Q}}\mathbb{I} - S$, and observing $\langle \Psi|\bar{S}|\Psi\rangle = 0$ implies the constraints $P_i|\Psi\rangle = 0$ for all $i \in \mathcal{I}$; these constraints can often be used to infer the algebraic structure (rigidity) of the measurements $\{M_{a|x}\}_{a\in\mathcal{A},x\in\mathcal{X}}, \{N_{b|y}\}_{b\in\mathcal{B},y\in\mathcal{Y}}$ which achieve $\langle S\rangle = \eta^{\mathrm{Q}}$.

3 Compiled Bell scenarios

The compilation procedure of a Bell scenario is essentially the same as the procedure for compiling nonlocal games outlined in [19]. Let $\mathcal{S} = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, \pi)$ be a 2-prover Bell scenario and fix a quantum homomorphic encryption scheme with *security against quantum distinguishers* and *correctness with respect to auxiliary input*. Readers unfamiliar with QHE schemes and these properties can refer to Definition 14 found in the appendix.

A **compiled Bell scenario** is the following 2-round single-prover scenario. To setup, the verifier samples a secret key $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$. Then, the verifier samples a pair of inputs $(x, y) \in \mathcal{X} \times \mathcal{Y}$ according to the distribution $\pi : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}_{\geq 0}$, and encrypts the first input as the ciphertext $\chi \leftarrow \mathsf{Enc}(\mathsf{sk}, x)$.

1. The verifier sends the ciphertext χ to the prover. The prover replies with a ciphertext α encoding their output. The verifier decrypts obtaining outcome $a \leftarrow \mathsf{Dec}(\mathsf{sk}, \alpha)$ from $\mathcal{A}$.
2. The verifier sends the sampled (plaintext) input $y \in \mathcal{Y}$ to the prover, who replies with another outcome $b \in \mathcal{B}$.

In the compiled scenario, for a chosen security parameter λ, the prover prepares an initial quantum polynomial time (QPT) preparable state $|\Psi^{(\lambda)}\rangle \in \widetilde{\mathcal{H}}^{(\lambda)}$ where $\widetilde{\mathcal{H}}^{(\lambda)}$ is a single Hilbert space (see Definition 13 for details on efficient quantum procedures). Then, the first round of the protocol is characterized by a family of POVMs $\{\{\widetilde{M}^{(\lambda)}_{\alpha|\chi}\}_{\alpha\in\bar{\mathcal{A}}}\}_{\chi\in\bar{\mathcal{X}}}$ and unitaries $\{U^{(\lambda)}_{\alpha,\chi}\}_{\alpha\in\bar{\mathcal{A}},\chi\in\bar{\mathcal{X}}}$, where $\bar{\mathcal{X}}$ and $\bar{\mathcal{A}}$ are the set of all valid ciphertexts of the first round input and output, respectively. Unlike in the bipartite setting, we must account for unitary operations applied to the post-measurement state in the first round. With this in mind, we denote the sub-normalized post-measurement state given the measurement over ciphertext χ and encrypted outcome α by

$$U^{(\lambda)}_{\alpha,\chi}\widetilde{M}^{(\lambda)}_{\alpha|\chi}|\Psi^{(\lambda)}\rangle =: |\Psi^{(\lambda)}_{\alpha|\chi}\rangle. \tag{3}$$

Note that these vectors are sub-normalized. In particular, the probability of obtaining $\alpha \in \bar{\mathcal{A}}$ given $\chi \in \bar{\mathcal{X}}$ is given by $\langle \Psi^{(\lambda)}_{\alpha|\chi}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle$. In the second round, the device makes a POVM measurement $\{\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}}$, where the resulting conditional probability is given by

$$\langle \Psi^{(\lambda)}|\widetilde{M}^{(\lambda)\dagger}_{\alpha|\chi}U^{(\lambda)\dagger}_{\alpha,\chi}N^{(\lambda)}_{b|y}U^{(\lambda)}_{\alpha,\chi}\widetilde{M}^{(\lambda)}_{\alpha|\chi}|\Psi^{(\lambda)}\rangle = \langle \Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle, \tag{4}$$

for a fixed, $\lambda \in \mathbb{N}$, $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$, ciphertexts $\chi \in \bar{\mathcal{X}}$, $\alpha \in \bar{\mathcal{A}}$, and plaintexts $y \in \mathcal{Y}$, $b \in \mathcal{B}$.

To summarize, for a fixed QHE scheme, $\lambda \in \mathbb{N}$, a **compiled (quantum) model** is given by a tuple

$$\widetilde{\mathrm{Q}}^{(\lambda)} = (\widetilde{\mathcal{H}}^{(\lambda)}, \{|\Psi^{(\lambda)}_{\alpha|\chi}\rangle\}_{\alpha\in\bar{\mathcal{A}},\chi\in\bar{\mathcal{X}}}, \{\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}}), \tag{5}$$

where all the relevant measurements and states are obtained by some QPT procedure. We remark that one can consider a description of the model which includes the initial state $|\Psi^{(\lambda)}\rangle$ and the operators $\{U^{(\lambda)}_{\alpha,\chi}\widetilde{M}^{(\lambda)}_{\alpha|\chi}\}_{\alpha\in\bar{\mathcal{A}},\chi\in\bar{\mathcal{X}}}$, rather than the post-measurement states $|\Psi^{(\lambda)}_{\alpha|\chi}\rangle$.

Hence, $\widetilde{\mathrm{Q}}^{(\lambda)}$ is really a coarse description of a quantum model in the compiled setting. The joint distribution of the outcomes after both rounds is given by

$$p^{(\lambda)}(a,b|x,y) = \mathop{\mathbb{E}}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha:\mathsf{Dec}(\alpha)=a} \langle\Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle. \tag{6}$$

Note that the marginal distribution $p^{(\lambda)}(a|x)$ obtained from Equation (6) will be independent of the second input y due to the sequential nature of the protocol. However, the marginal $p^{(\lambda)}(b|y,x)$ currently depends on x. The aim of what follows is to establish a computational independence between this distribution and the inputs x. To do so we will need to consider the distributions of the decrypted outputs and appeal to the security promise of the QHE scheme. Specifically, we require a key lemma which has appeared in several works [27, 16, 6]. We borrow a version from [20] and we refer the reader to the reference for the proof.

▶ **Lemma 1** ([20], Proposition 4.6)**.** *Let $\widetilde{\mathrm{Q}}^{(\lambda)}$ be a compiled quantum model, and $\mathcal{N}^{(\lambda)} = w(\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B},y\in\mathcal{Y}})$ be a monomial in the measurement operators $\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B},y\in\mathcal{Y}}$, where $\lambda\in\mathbb{N}$ is the security parameter for a fixed QHE scheme. Then, for any two QPT sampleable distributions $\mathcal{D}_1,\mathcal{D}_2$ over plaintext inputs $x\in\mathcal{X}$ there exists a negligible function* $\mathrm{negl}(\lambda)$ *of the security parameter λ such that the following holds*

$$\left| \mathop{\mathbb{E}}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)} \mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_1} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha\in\bar{\mathcal{A}}} \langle\Psi^{(\lambda)}_{\alpha|\chi}|\mathcal{N}^{(\lambda)}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle - \mathop{\mathbb{E}}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)} \mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_2} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha\in\bar{\mathcal{A}}} \langle\Psi^{(\lambda)}_{\alpha|\chi}|\mathcal{N}^{(\lambda)}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle \right|$$
$$\leq \mathrm{negl}(\lambda).$$

The approximate no-signalling conditions from Alice to Bob can then be seen by applying Lemma 1 to the monomials of degree 1 in the QPT measurement operators $\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B},y\in\mathcal{Y}}$, since

$$\left| \mathop{\mathbb{E}}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha\in\bar{\mathcal{A}}} \langle\Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle - \mathop{\mathbb{E}}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x')=\chi} \sum_{\alpha\in\bar{\mathcal{A}}} \langle\Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle \right| \leq \mathrm{negl}(\lambda) \tag{7}$$

holds for all $b\in\mathcal{B}, y\in\mathcal{Y}$ and $x,x'\in\mathcal{X}$ with $x\neq x'$.

In the above statements, the measurements are completely general, and the states are sub-normalized vectors. The following lemma shows that when considering the compiled value, we can assume that the states and measurement operators in the compiled strategy are pure and projective.

▶ **Lemma 2.** *Let $\mathcal{H}^{'(\lambda)}$ be the Hilbert space of the device, and $\{\{\rho^{(\lambda)}_{\alpha|\chi}\}_{\alpha\in\bar{\mathcal{A}}}\}_{\chi\in\bar{\mathcal{X}}}$ be a family of QPT-preparable sub-normalized states on $\mathcal{H}^{'(\lambda)}$ after the first round. Let $\{\{N^{'(\lambda)}_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}}$ be a family of QPT-implementable POVMs on $\mathcal{H}^{'(\lambda)}$, which induce the behaviour $p^{(\lambda)}(\alpha,b|\chi,y) = tr[N^{'(\lambda)}_{b|y}\rho^{(\lambda)}_{\alpha|\chi}]$. Then there exists a Hilbert space $\mathcal{H}^{(\lambda)}$, a family of QPT-preparable sub-normalized states $\{\{|\Psi^{(\lambda)}_{\alpha|\chi}\rangle\}_{\alpha\in\bar{\mathcal{A}}}\}_{\chi\in\bar{\mathcal{X}}}$ in $\mathcal{H}^{(\lambda)}$, and a family of QPT-implementable PVMs $\{\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}}$ on $\mathcal{H}^{(\lambda)}$ which satisfy*

$$\langle\Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle = p^{(\lambda)}(\alpha,b|\chi,y), \quad \forall\alpha\in\bar{\mathcal{A}},\ \chi\in\bar{\mathcal{X}},\ b\in\mathcal{B},\ y\in\mathcal{Y}. \tag{8}$$

See Section A.2 for the proof of Lemma 2.

We say a compiled model $\widetilde{\mathrm{Q}}^{(\lambda)} = (\widetilde{\mathcal{H}}^{(\lambda)}, \{|\Psi^{(\lambda)}_{\alpha|\chi}\rangle\}_{\alpha\in\bar{\mathcal{A}},\chi\in\bar{\mathcal{X}}}, \{\{N^{(\lambda)}_{b|y}\}_{b\in\mathcal{B}}\}_{y\in\mathcal{Y}})$ is pure and projective whenever the states $|\Psi^{(\lambda)}_{\alpha|\chi}\rangle$ are all pure and the measurements $N^{(\lambda)}_{b|y}$ are all projective (i.e. PVMS).

3.1 Quantum bounds for compiled inequalities

A compiled (quantum) model $\widetilde{\mathrm{Q}}^{(\lambda)}$ describes the correlations $\mathsf{p}^{(\lambda)} = \{p^{(\lambda)}(a,b|x,y)\}_{a\in\mathcal{A},b\in\mathcal{B},x\in\mathcal{X},y\in\mathcal{Y}}$ observed in a compiled Bell scenario. A **compiled Bell functional** is a linear functional $I^{(\lambda)}$ evaluated on correlations realized by compiled models. That is

$$I^{(\lambda)} = \sum_{abxy} w_{abxy} \mathop{\mathbb{E}}_{\substack{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)\\ \chi:\mathsf{Enc}(x)=\chi}} \sum_{\alpha:\mathsf{Dec}(\alpha)=a} \langle\Psi^{(\lambda)}_{\alpha|\chi}|N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle. \tag{9}$$

By the properties of the compilation procedure [19, Theorem 3.2], Bell inequalities are preserved under compilation (up to negligible error). In particular, for large security parameter, efficient classical provers cannot violate a Bell inequality by much more than they could in the (bipartite) scenario. From now on, we will suppress the security parameter $\lambda \in \mathbb{N}$ along with the expectation over secret keys $\mathbb{E}_{\mathsf{sk}\leftarrow\mathsf{Gen}(1^\lambda)}$ and simply write the expectation for a fixed key. In particular, we express the compiled model as $\widetilde{\mathrm{Q}}$ and Equation (9) as

$$I = \sum_{abxy} w_{abxy} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha:\mathsf{Dec}(\alpha)=a} \langle\Psi_{\alpha|\chi}|N_{b|y}|\Psi_{\alpha|\chi}\rangle.$$

We now turn our attention to the maximum value I can take in the compiled setting with an efficient quantum prover. The results of [19] imply that an efficient quantum prover can achieve the same violation in the bipartite setting. However, the existence of a quantum compiled behavior which exceeds the maximal quantum Bell violation in the bipartite case (by more than negligible factors) has not been ruled out. Nonetheless, in several cases (like the CHSH inequality and more generally all XOR games [16]) we know that the quantum compiled behavior cannot exceed the value η^{Q} by more than negligible amounts. One technique for establishing such bounds was introduced in [27] and uses SOS techniques to bound the quantum violation of the compiled Bell functional.

3.2 Extending the pseudo-expectations

Our approach builds off the methods used in [27] and [16]. To explain this approach we recall that a pseudo-expectation is a unital, linear map from a subspace $\mathcal{T}$ of the algebra generated by $\{M_{a|x}, N_{b|y}\}_{a\in\mathcal{A},x\in\mathcal{X},b\in\mathcal{B},y\in\mathcal{Y}}$ to the complex numbers, $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}} : \mathcal{T} \to \mathbb{C}$, which is determined by a compiled quantum model $\widetilde{\mathrm{Q}}$. In the case $n = m = 2$, it suffices to define the pseudo-expectation $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}$ on the observables $A_x = \sum_{a\in\{0,1\}}(-1)^a M_{a|x}$, $B_y = \sum_{b\in\{0,1\}}(-1)^b N_{b|y}$ and require that they are mapped to their expectations in the compiled scenario[2]. We further assume that all measurements are projective (cf. Lemma 2). In previous works, the definition of the pseudo-expectation had been restricted to monomials consisting of at most one Alice and one Bob observable as outlined below:

[2] Though in the following we define $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}$ for $n = m = 2$, this can be directly extended to arbitrary Bell scenarios by defining $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}$ on the POVM elements $M_{a|x}, N_{b|y}$ in an analogous way.

$$\begin{aligned}
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[A_x B_y] &:= \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha (-1)^{\mathsf{Dec}(\alpha)} \langle \Psi_{\alpha|\chi} | B_y | \Psi_{\alpha|\chi} \rangle, \\
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[A_x A_{x'}] &:= \delta_{x,x'}, \\
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[B_y B_{y'}] &:= \mathop{\mathbb{E}}_{x\in\mathcal{X}} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle \Psi_{\alpha|\chi} | B_y B_{y'} | \Psi_{\alpha|\chi} \rangle, \\
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[A_x] &:= \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha (-1)^{\mathsf{Dec}(\alpha)} \langle \Psi_{\alpha|\chi} | \Psi_{\alpha|\chi} \rangle, \\
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[B_y] &:= \mathop{\mathbb{E}}_{x\in\mathcal{X}} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle \Psi_{\alpha|\chi} | B_y | \Psi_{\alpha|\chi} \rangle, \\
\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\mathbb{I}] &:= 1,
\end{aligned} \tag{10}$$

where $\mathbb{E}_{x\in\mathcal{X}}$ denotes the expectation according to an arbitrary fixed distribution over $\mathcal{X}$. This is already sufficient to handle known SOS decompositions for a variety of well-studied Bell inequalities whenever the polynomials are expressed in the basis $\{\mathbb{I}, A_x, B_y\}_{x\in\mathcal{X}, y\in\mathcal{Y}}$. However, there are Bell inequalities, such as the tilted-CHSH inequality [4, 5], for which no known SOS decomposition exists in the basis $\{\mathbb{I}, A_x, B_y\}_{x\in\mathcal{X}, y\in\mathcal{Y}}$.

The contribution of this section is to expand the definition of the pseudo-expectation to the basis encompassing all monomials in A_x, B_0, B_1, for a fixed $x \in \mathcal{X}$, in a way that is approximately non-negative on Hermitian squares. This allows us to handle more general SOS decompositions, and in particular, the tilted-CHSH inequalities. Let $w(A_x, B_0, B_1)$ be a monomial in the elements $\{A_x, B_0, B_1\}$. Importantly, x is fixed, and we do not consider monomials of the form $A_0 A_1 B_y$ for example. Let $\bar{w}$ be the canonical form of w under the relations $[A_x, B_y] = 0$, $(B_y)^2 = (A_x)^2 = \mathbb{I}$, where all A_x terms are commuted to the left. Since we only consider one value of x, these will all be of the form $(A_x)^i \bar{w}(B_0, B_1)$ for some $i \in \{0, 1\}$, where the monomial $\bar{w}(B_0, B_1)$ cannot be reduced further. We then define the pseudo-expectation

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}\big[w(A_x, B_0, B_1)\big] := \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}\big[(A_x)^i \bar{w}(B_0, B_1)\big]. \tag{11}$$

For the case $i = 0$, we define

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}\big[\bar{w}(B_0, B_1)\big] := \mathop{\mathbb{E}}_{x\in\mathcal{X}} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle \Psi_{\alpha|\chi} | \bar{w}(B_0, B_1) | \Psi_{\alpha|\chi} \rangle, \tag{12}$$

and for the case $i = 1$,

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}\big[A_x \bar{w}(B_0, B_1)\big] := \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha (-1)^{\mathsf{Dec}(\alpha)} \langle \Psi_{\alpha|\chi} | \bar{w}(B_0, B_1) | \Psi_{\alpha|\chi} \rangle. \tag{13}$$

From the above definitions, we next state the main result of this section, which can be applied generally to any polynomial expressible in the basis $\{A_x, B_0, B_1\}$.

► **Theorem 3.** *Let $\{A_x\}_{x\in\mathcal{X}}$ and $\{B_y\}_{y\in\mathcal{Y}}$ be binary observables, and let*

$$P = \sum_i \gamma_i (A_x)^{k_i} w_i(B_0, B_1), \tag{14}$$

where $\gamma_i \in \mathbb{C}$, $k_i \in \{0, 1\}$ and each $w_i(B_0, B_1)$ is any monomial in the algebra of $\{B_0, B_1\}$. Then there exists a negligible function $\mathrm{negl}(\lambda)$ *of the security parameter $\lambda \in \mathbb{N}$ such that*

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[P^\dagger P] \geq -\mathrm{negl}(\lambda). \tag{15}$$

Furthermore, for a given Bell functional I, and a compiled model $\widetilde{\mathrm{Q}}$, $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}(I)$ is the expected value of the compiled model $\widetilde{\mathrm{Q}}$ on I.

The proof can be found in Section A.2.

3.3 Quantum bounds for compiled tilted-CHSH expressions

We now present the family of extended tilted-CHSH type expressions and their SOS decompositions discovered in [5]. Let $\theta \in (0, \pi/4]$, $\phi \in \big(\max\{-2\theta, -\pi + 2\theta\}, \min\{2\theta, \pi - 2\theta\}\big) \setminus \{0\}$, and $t_{\theta,\phi} \in \mathbb{R}$ such that

$$\frac{1}{t_{\theta,\phi}^2} = \frac{\sin^2(2\theta)}{\tan^2(\phi)} - \cos^2(2\theta). \tag{16}$$

From here, we define the following expressions:

$$\begin{aligned} S_{\theta,\phi} &:= A_0 \otimes \frac{B_0 + B_1}{\cos(\phi)} + t_{\theta,\phi}^2 \left[\sin(2\theta)\, A_1 \otimes \frac{B_0 - B_1}{\sin(\phi)} + \cos(2\theta)\, \mathbb{I} \otimes \frac{B_0 + B_1}{\cos(\phi)}\right], \\ \eta_{\theta,\phi}^{\mathrm{Q}} &:= 2(1 + t_{\theta,\phi}^2). \end{aligned} \tag{17}$$

We also let $I_{\theta,\phi}$ denote the corresponding Bell functional, and recall the following result.

▶ **Lemma 4** ([5], Section 3.2.1). *Let $\theta \in (0, \pi/4]$, $\phi \in \big(\max\{-2\theta, -\pi + 2\theta\}, \min\{2\theta, \pi - 2\theta\}\big) \setminus \{0\}$, $t_{\theta,\phi}$ be given by Equation (16) and $S_{\theta,\phi}, \eta_{\theta,\phi}^{\mathrm{Q}}$ be defined in Equation (17). Define the following polynomials:*

$$\begin{aligned} N_0 &:= A_0 \otimes \mathbb{I} - \mathbb{I} \otimes \frac{B_0 + B_1}{2\cos(\phi)}, \\ N_1 &:= A_1 \otimes \mathbb{I} - \sin(2\theta)\, \mathbb{I} \otimes \frac{B_0 - B_1}{2\sin(\phi)} - \cos(2\theta)\, A_1 \otimes \frac{B_0 + B_1}{2\cos(\phi)}. \end{aligned} \tag{18}$$

Then the shifted Bell operator $\bar{S}_{\theta,\phi} = \eta_{\theta,\phi}^{\mathrm{Q}} \mathbb{I} - S_{\theta,\phi}$ admits the SOS decomposition

$$\bar{S}_{\theta,\phi} = N_0^\dagger N_0 + t_{\theta,\phi}^2 N_1^\dagger N_1. \tag{19}$$

Using the decomposition in Lemma 4, it was shown in [5] that the inequality $\langle S_{\theta,\phi} \rangle \leq \eta_{\theta,\phi}^{\mathrm{Q}}$ self-tests the partially entangled state $|\psi_\theta\rangle = \cos(\theta)|00\rangle + \sin(\theta)|11\rangle$ and the measurements

$$\begin{aligned} &A_0 = \sigma_Z, \;\; A_1 = \sigma_X, \\ &B_y = \cos(\phi)\, \sigma_Z + (-1)^y \sin(\phi)\, \sigma_X, \;\; y \in \{0, 1\}, \end{aligned} \tag{20}$$

where σ_Z, σ_X are the Pauli operators. Notably, by setting $\phi = \mu_\theta$ where $\tan(\mu_\theta) = \sin(2\theta)$, this family encompasses what are most commonly referred to as "tilted-CHSH inequalities" given by the Bell operator

$$T_\theta = \alpha_\theta A_0 \otimes \mathbb{I} + A_0 \otimes (B_0 + B_1) + A_1 \otimes (B_0 - B_1), \tag{21}$$

where $\alpha_\theta = 2/\sqrt{1 + 2\tan^2(2\theta)}$ [2, 35, 4]. Compared to the SOS decompositions for T_θ from [4], the decomposition of [5] is expressed in the basis for which our extended pseudo-expectation is well defined (cf. Theorem 3), allowing us to provide bounds on the compiled value of T_θ, and more generally the family $S_{\theta,\phi}$.

▶ **Theorem 5.** *Let $\theta \in (0, \pi/4]$, $\phi \in \big(\max\{-2\theta, -\pi + 2\theta\}, \min\{2\theta, \pi - 2\theta\}\big) \setminus \{0\}$, and let $S_{\theta,\phi}$ be the extended tilted-CHSH expression with quantum bound $\eta_{\theta,\phi}^{\mathrm{Q}}$, given by Equation (17). Then the maximum quantum value of the corresponding compiled Bell inequality is given by $\eta_{\theta,\phi}^{\mathrm{Q}} + \mathrm{negl}(\lambda)'$, where $\mathrm{negl}(\lambda)'$ is a negligible function of the security parameter.*

Proof. We evaluate the pseudo-expectation on the shifted Bell expression $\bar{S}_{\theta,\phi}$:

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\bar{S}_{\theta,\phi}] = \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[N_0^\dagger N_0] + \lambda_{\theta,\phi}^2 \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[N_1^\dagger N_1], \tag{22}$$

where we used the decomposition in Lemma 4. The polynomial N_0 is expressed in the basis $\{A_0, B_0, B_1\}$, and we find by Theorem 3 that

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[N_0^\dagger N_0] \geq -\mathrm{negl}(\lambda). \tag{23}$$

Similarly, N_1 is expressed in the basis $\{A_1, B_0, B_1\}$, and we see by Theorem 3 that $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[N_1^\dagger N_1] \geq -\mathrm{negl}(\lambda)$. Putting these together, we obtain

$$\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\bar{S}_{\theta,\phi}] \geq -\mathrm{negl}(\lambda)(1 + \lambda_{\theta,\phi}^2) =: -\mathrm{negl}(\lambda)', \tag{24}$$

which implies $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[S_{\theta,\phi}] \leq \eta_{\theta,\phi}^{\mathrm{Q}} + \mathrm{negl}(\lambda)'$ as desired, where $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[S_{\theta,\phi}]$ is the expected value of the compiled Bell inequality. ◀

▶ Remark 6. The extension of the $S_{\theta,\phi}$ family presented in [5, Section 3.2.3] self-tests the state $|\psi_\theta\rangle$ along with the more general measurements

$$\begin{aligned} A_0 &= \sigma_Z, \; A_1 = \sigma_X, \\ B_0 &= \cos(\phi)\,\sigma_Z + \sin(\phi)\,\sigma_X, \\ B_1 &= \cos(\omega)\,\sigma_Z + \sin(\omega)\,\sigma_X, \end{aligned} \tag{25}$$

for $\phi \in (-2\theta, 0)$ and $\omega \in (0, 2\theta)$. This family of Bell inequalities can also be compiled under our definition of the pseudo-expectation. This is because each SOS polynomial is given in the basis $\{A_x, B_0, B_1\}$ for a fixed x, and we can apply Theorem 3 directly as was done in Theorem 5. We omit the explicit proof of this for brevity.

4 Self-testing in the compiled setting

Recall that a bipartite (quantum) model Q, consists of a shared state $|\Psi\rangle$, along with local POVM measurements $\{M_{a|x}\}$ and $\{N_{b|y}\}$ for Alice and Bob, respectively. Given a Bell expression I, the inequality $I \leq \eta^{\mathrm{Q}}$ self-tests an ideal bipartite model Q* if any optimal bipartite model is essentially the same as Q*, modulo some physically irrelevant degrees of freedom. This is more formally stated in terms or the existence of local isometries which maps the employed model to the ideal one. When small errors are permitted, one considers the following definition of robust self-testing.

▶ **Definition 7** (Bipartite self-test). *The inequality $I \leq \eta^{\mathrm{Q}}$ is a self-test for a bipartite model $\mathrm{Q}^* = \big(\{P_{a|x}\}, \{Q_{b|y}\}, |\phi\rangle\big)$ if there exist a non-negative function $f(\epsilon)$ such that $f(\epsilon) \to 0$ as $\epsilon \to 0$, such that for any bipartite model $\mathrm{Q} = \big(\{M_{a|x}\}, \{N_{b|y}\}, |\Psi\rangle\big)$ achieving $I \geq \eta^{\mathrm{Q}} - \epsilon$ for $\epsilon \geq 0$, there exists a Hilbert space $\mathcal{H}_{\mathrm{aux}}$, an auxiliary state $|\zeta\rangle \in \mathcal{H}_{\mathrm{aux}}$ and local isometries V_A and V_B, such that defining $V : \mathcal{H}_A \otimes \mathcal{H}_B \to \mathbb{C}^d \otimes \mathbb{C}^d \otimes \mathcal{H}_{\mathrm{aux}}$, $V = V_A \otimes V_B$, the following is satisfied for all x, y, a, b:*

$$\big\| V_A \otimes V_B (M_{a|x} \otimes N_{b|y}) |\Psi\rangle - (P_{a|x} \otimes Q_{b|y}) |\phi\rangle \otimes |\zeta\rangle \big\| \leq f(\epsilon).$$

In the bipartite setting, one could consider the situation where Alice measures first using a POVM $\{P_{a|x}\}$, collapsing the state to a post-measurement state $\rho_{a|x}$ on Bob's subsystem $\mathcal{H}_B$, upon which Bob performs his measurement, resulting in the application of the POVM element $Q_{b|y}$. With this in mind, we consider the setting where the only relevant features of the model are those from Bob's (resp. Alice's) perspective. In particular, subsystem A is traced out following the recorded measurement of outcome of a given x.

▶ **Definition 8** (Partial model). *Given a bipartite model* $\mathrm{Q} = \big(\{M_{a|x}\}, \{N_{b|y}\}, |\Psi\rangle\big)$*, we define the partial model of* Q *by* $\mathrm{Q}' = (\{N_{b|y}\}, \{\rho_{a|x}\})$ *where*

$$\rho_{a|x} = tr_A[(M_{a|x} \otimes \mathbb{I}_B)|\Psi\rangle\langle\Psi|]. \tag{26}$$

We note that $\rho_{a|x}$ *will generally be mixed. When each* $\rho_{a|x}$ *is pure, we say that* Q *has a pure partial model, denoted by* $\mathrm{Q}' = (\{N_{b|y}\}, \{|\phi_{a|x}\rangle\})$*.*

Symmetrically, given a bipartite model one can consider a (pure) partial model on $\mathcal{H}_A$ by tracing out subsystem B. However, because our motivation is the compiled setting, we will focus on the partial models on $\mathcal{H}_B$. Furthermore, we remark that the notion of pure partial models is not vacuous. In particular, the optimal bipartite model for the CHSH inequality has a pure partial model on $\mathcal{H}_B$ [10]. With the notion of a partial quantum model, we define the notion of a partial (or one-sided) self-test for a bipartite model.

▶ **Definition 9** (Partial self-test). *The inequality* $I \leq \eta^{\mathrm{Q}}$ *is a partial self-test for a bipartite model* $\mathrm{Q}^* = \big(\{P_{a|x}\}, \{Q_{b|y}\}, |\phi\rangle\big)$ *with a pure partial model* $\big(\{Q_{b|y}\}, \{|\phi_{a|x}\rangle\}\big)$ *if there exists a non-negative function* $f(\epsilon)$ *such that* $f(\epsilon) \to 0$ *as* $\epsilon \to 0$*, such that for any partial quantum model* $\mathrm{Q} = \big(\{N_{b|y}\}, \{\rho_{a|x}\}\big)$ *achieving* $I \geq \eta^{\mathrm{Q}} - \epsilon$ *for* $\epsilon \geq 0$*, there exist a Hilbert space* $\mathcal{H}_{\mathrm{aux}}$*, a collection of auxiliary states* $\{\sigma_{a|x}\}$ *and an isometry* $V : \mathcal{H}_B \to \mathbb{C}^d \otimes \mathcal{H}_{\mathrm{aux}}$ *such that the following is satisfied for all* x, y, a, b*:*

$$\begin{aligned}&\big\|VN_{b|y}\rho_{a|x}N_{b|y}V^\dagger - Q_{b|y}|\phi_{a|x}\rangle\langle\phi_{a|x}|Q_{b|y} \otimes \sigma_{a|x}\big\|_2 \leq f(\epsilon)\\ &\qquad\qquad \textit{and} \quad \big\|V\rho_{a|x}V^\dagger - |\phi_{a|x}\rangle\langle\phi_{a|x}| \otimes \sigma_{a|x}\big\|_2 \leq f(\epsilon),\end{aligned}$$

Give the symmetry of $\mathcal{H}_A$ and $\mathcal{H}_B$ in the bipartite case, one can define a notion of partial self-test for either subsystem. Given a bipartite self-test, one can check that tracing out either subsystem results in a partial self-test. We leave it as an open question as to whether a partial self-test (say over $\mathcal{H}_A$ and over $\mathcal{H}_B$) implies that the correlation is a bipartite self-test.

4.1 Compiled self-tests from partial models

There are two main difficulties with self-testing in the compiled setting. Firstly, the *correctness with respect to auxiliary systems* property of the compiler (see *Property (1)* in Definition 14) only guarantees that a QPT prover can prepare states (possibly mixed) $\rho_{a|x}$ over $\mathcal{H}_B$ that are negligible in trace distance from the post measurement states $P_{a|x}|\Psi\rangle\langle\Psi|P_{a|x}/p(a|x)$ of the ideal bipartite model Q. This puts a fundamental constraint on our ability to exactly describe the set of ideal models in the compiled setting. Secondly, unlike in the nonlocal setting, it is not clear how to extract information about the measurements and states in the first round due to the homomorphic evaluation of the measurements and preparation of the states. To address these challenges we introduce the compiled counter-part of a partial quantum model.

Recall that a compiled (quantum) model $\widetilde{\mathrm{Q}}$ consists of a family of post-measurement states for "Alice" $|\widetilde{\phi}_{\alpha|\chi}\rangle$, which correspond to the state of the device following the encrypted question χ, and encrypted answer α, and a POVM $\{N_{b|y}\}$ employed by "Bob". One could also consider a more general compiled quantum model, which includes a description of the initial state and Alice's operators. The point of taking the coarser model is that it allows us to introduce the notion of the *compiled-counterpart* of a bipartite model Q, which relates the post-measurement information in the bipartite setting with another bipartite model that resembles a compiled model.

▶ **Definition 10** (Compiled-counterpart model). *Given a pure partial model* Q'*, the compiled-counterpart model of* Q' *is the pure partial model* $\widetilde{\mathrm{Q}}^{(\lambda)} = (\{|\widetilde{\phi}^{(\lambda)}_{\alpha|\chi}\rangle\}, \{Q^{(\lambda)}_{b|y}\})$ *satisfying the following conditions for all* $\lambda \in \mathbb{N}$*:*

$$|\widetilde{\phi}^{(\lambda)}_{\alpha|\chi}\rangle = |\phi_{a|x}\rangle, \quad \textit{for all } \ \mathsf{sk} : \mathsf{Gen}(1^{\lambda}) = \mathsf{sk},\ \chi : \mathsf{Enc}(x, \mathsf{sk}) = \chi,\ \alpha : \mathsf{Dec}(\alpha, \mathsf{sk}) = a.$$

$$N^{(\lambda)}_{b|y} = Q_{b|y}, \quad \textit{for all } b, y.$$

We remark that the compiled counterpart need not be an actual compiled model. For example, it is not required to satisfy the QPT conditions needed of a compiled model. Instead it is a model that resembles an idealized version of an honest implementation of a partial model under homomorphic encryption. We proceed with a definition of self-testing in the compiled setting that resembles partial self-testing in the bipartite setting in the context of these compiled-counterparts.

▶ **Definition 11** (Compiled self-test). *Let* I *denote a Bell expression with an optimal pure partial model* Q^**. The inequality* $I \leq \eta^{\mathrm{Q}}$ *is a compiled self-test for the corresponding compiled-counterpart* $\widetilde{\mathrm{Q}}^* = (\{|\widetilde{\phi}_{\alpha|\chi}\rangle\}, \{Q_{b|y}\})$*, if there exists a non-negative function* $f(\epsilon)$ *such that* $f(\epsilon) \to 0$ *as* $\epsilon \to 0$*, such that for every pure and projective compiled model* $\widetilde{\mathrm{Q}} = \big(\{|\Psi_{\alpha|\chi}\rangle\}, \{N_{b|y}\}\big)$ *that achieves* $I \geq \eta^{\mathrm{Q}} - \epsilon$ *for some* $\epsilon \geq 0$*, there exists a negligible function* $\mathrm{negl}(\lambda)$*, an isometry* $V : \tilde{\mathcal{H}} \to \mathbb{C}^d \otimes \mathcal{H}_{\mathrm{aux}}$*, and auxiliary states* $|\mathsf{aux}_{\alpha|\chi}\rangle \in \mathcal{H}_{\mathrm{aux}}$*, which satisfy the following for all* x, b, y*:*

$$\mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha} \big\| V|\Psi_{\alpha|\chi}\rangle - |\widetilde{\phi}_{\alpha|\chi}\rangle \otimes |\mathsf{aux}_{\alpha|\chi}\rangle \big\|^2 \leq \mathrm{negl}(\lambda) + f(\epsilon), \quad \textit{and} \tag{27a}$$

$$\mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_{\alpha} \big\| V N_{b|y}|\Psi_{\alpha|\chi}\rangle - Q_{b|y}|\widetilde{\phi}_{\alpha|\chi}\rangle \otimes |\mathsf{aux}_{\alpha|\chi}\rangle \big\|^2 \leq \mathrm{negl}(\lambda) + f(\epsilon). \tag{27b}$$

Equation (27a) is a statement about the provers state after the first round. It asserts that, given a question x and answer a, the post-measurement state is negligibly close to that of an ideal prover implementing the honest bipartite model. To see this concretely, suppose the right hand side was exactly equal to zero. Then we have the equality $V|\Psi_{\alpha|\chi}\rangle = |\widetilde{\phi}_{\alpha|\chi}\rangle \otimes |\mathsf{aux}_{\alpha|\chi}\rangle$ for all χ such that $\mathsf{Enc}(x) = \chi$ and all α. Substituting $|\widetilde{\phi}_{\alpha|\chi}\rangle$ for the states $|\phi_{a|x}\rangle$ from Definition 10, we obtain

$$V|\Psi_{\alpha|\chi}\rangle = |\phi_{a|x}\rangle \otimes |\mathsf{aux}_{\alpha|\chi}\rangle \tag{28}$$

whenever $\mathsf{Enc}(x) = \chi$ and $\mathsf{Dec}(\alpha) = a$. That is, the post-measurement states are equal to the target states up an isometry. Therefore, we interpret (27a) as an approximate version of Equation (28), which accounts for a finite size security parameter λ and small errors in the Bell violation ϵ. Equation (27b) is the analogous statement including the measurements in the second round. We remark that if V could depend on the question x and answer a, (27a) would trivially hold regardless of the compiled Bell violation, since the states $|\phi_{a|x}\rangle$ could be prepared directly. It is therefore essential to enforce the same isometry is applied for all a and x. Furthermore, (27b) captures several existing self-testing results in the compiled setting. For example those presented in [27, Lemma 34], [16, Theorem 3.6] and [16, Eqs. 98 and 103]. Our proposed definition then goes further by also certifying the states after the first round but before Bob's measurements, as captured by (27a).

It is natural to ask if Definition 11 is the strongest form of self-testing possible in this scenario, or if one can also certify the initial state $|\Psi\rangle$ before Alice's measurements. An initial guess would be to show there exists an isometry V satisfying

$$V|\Psi\rangle \approx_{\mathrm{negl}(\lambda)} |\phi\rangle \otimes |\mathsf{aux}\rangle, \tag{29}$$

where $|\phi\rangle$ is the ideal bipartite entangled state. However, on its own this statement is not very useful: such an isometry always exists, namely, one which ignores $|\Psi\rangle$ and prepares $|\phi\rangle$ directly. A possible way around is to demand the same V also satisfies (27). At a glance, this suggests certifying the initial state alone is not meaningful in the single prover setting; one always needs to also consider the measurements. This contrasts the two prover setting, where self-testing statements made only about the state are known [34] and non-trivial due to the space-like separation of the provers. Another question worth asking is if the assumption of having a pure projective models $\widetilde{\mathrm{Q}}$ can be relaxed in the definition Definition 11.

4.2 Compiled self-test for tilted-CHSH inequalities

Our final result is that the extended tilted-CHSH Bell inequalities are compiled self-tests according to Definition 11. In particular, we have the following result.

▶ **Theorem 12.** *Let $\theta \in (0, \pi/4]$, $\phi \in \big(\max\{-2\theta, -\pi + 2\theta\}, \min\{2\theta, \pi - 2\theta\}\big) \setminus \{0\}$, and let $I_{\theta,\phi}$ be the generalized tilted-CHSH functional with quantum bound $\eta^{\mathrm{Q}}_{\theta,\phi}$ according to Equation* (17)*. Then the inequality $I_{\theta,\phi} \leq \eta^{\mathrm{Q}}_{\theta,\phi}$ is a compiled self-test for the compiled-counter part of* (20) *according to Definition 11.*

The proof is reminiscent of the approach in [4], and includes similar calculations to those used in [27, 6] which establish rigidity statements in the compiled setting. Given the length of the proof, we refer the reader to the longer version of this work [24] for all the details.

References

1 Antonio Acin, Nicolas Brunner, Nicolas Gisin, Serge Massar, Stefano Pironio, and Valerio Scarani. Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters*, 98:230501, 2007. `doi:10.1103/PhysRevLett.98.230501`.

2 Antonio Acín, Serge Massar, and Stefano Pironio. Randomness versus nonlocality and entanglement. *Physical Review Letters*, 108:100402, March 2012. `doi:10.1103/PhysRevLett.108.100402`.

3 Atul Singh Arora, Kishor Bharti, Alexandru Cojocaru, and Andrea Coladangelo. A computational test of contextuality and, even simpler proofs of quantumness. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1106–1125. IEEE, October 2024. `doi:10.1109/focs61266.2024.00073`.

4 Cédric Bamps and Stefano Pironio. Sum-of-squares decompositions for a family of Clauser-Horne-Shimony-Holt-like inequalities and their application to self-testing. *Physical Review A*, 91:052111, May 2015. `doi:10.1103/PhysRevA.91.052111`.

5 Victor Barizien, Pavel Sekatski, and Jean-Daniel Bancal. Custom Bell inequalities from formal sums of squares. *Quantum*, 8:1333, May 2024. `doi:10.22331/q-2024-05-02-1333`.

6 Matilde Baroni, Quoc-Huy Vu, Boris Bourdoncle, Eleni Diamanti, Damian Markham, and Ivan Šupić. Quantum bounds for compiled XOR games and d-outcome CHSH games. *arXiv:2403.05502*, 2024.

7 John S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.

8 Zvika Brakerski. Quantum FHE (almost) as secure as classical. In *Annual International Cryptology Conference*, pages 67–95. Springer, 2018.

9 Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In *Annual Cryptology Conference*, pages 609–629. Springer, 2015.

10 J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23:880–884, 1969. `doi:10.1103/PhysRevLett.23.880`.

11 John Clauser, Michael Horne, Abner Shimony, and Richard Holt. Proposed experiment to test local hidden-variable theories. *Physical review letters*, 23(15):880, 1969.

12 Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 236–249. IEEE, 2004.

13 Andrea Coladangelo, Koon Tong Goh, and Valerio Scarani. All pure bipartite entangled states can be self-tested. *Nature Communications*, 8(1), May 2017. doi:10.1038/ncomms15485.

14 Andrea Coladangelo, Alex B Grilo, Stacey Jeffery, and Thomas Vidick. Verifier-on-a-leash: New schemes for verifiable delegated quantum computation, with quasilinear resources. *Theory of Computing*, 20(1):1–87, 2024.

15 Roger Colbeck. *Quantum and Relativistic Protocols For Secure Multi-Party Computation.* PhD thesis, University of Cambridge, 2007. Also available as *arXiv:0911.3814*.

16 David Cui, Giulio Malavolta, Arthur Mehta, Anand Natarajan, Connor Paddock, Simon Schmidt, Michael Walter, and Tina Zhang. A computational Tsirelson's theorem for the value of compiled XOR games. *arXiv preprint arXiv:2402.17301*, 2024.

17 Alex Grilo. A simple protocol for verifiable delegation of quantum computation in one round. In *icalp2019*, pages 28:1–28:13, 2019. doi:10.4230/LIPIcs.ICALP.2019.28.

18 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, August 2022. doi:10.1038/s41567-022-01643-7.

19 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.

20 Alexander Kulpe, Giulio Malavolta, Connor Paddock, Simon Schmidt, and Michael Walter. A bound on the quantum value of all compiled nonlocal games. *arXiv:2408.06711*, 2024.

21 Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *SIAM Journal on Computing*, 52(6):FOCS18–189, 2020.

22 Dominic Mayers and Andrew Yao. Self testing quantum apparatus. *Quantum Information and Computation*, 4:273–286, 2004. doi:10.26421/QIC4.4-3.

23 Matthew McKague, Tzyh Haur Yang, and Valerio Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(45):455304, 2012.

24 Arthur Mehta, Connor Paddock, and Lewis Wooltorton. Self-testing in the compiled setting via tilted-CHSH inequalities. *arXiv preprint arXiv:2406.04986*, 2024.

25 Tony Metger, Anand Natarajan, and Tina Zhang. Succinct arguments for QMA from standard assumptions via compiled nonlocal games. *arXiv:2404.19754*, 2024.

26 Tony Metger and Thomas Vidick. Self-testing of a single quantum device under computational assumptions. *Quantum*, 5:544, 2021.

27 Anand Natarajan and Tina Zhang. Bounding the quantum value of compiled nonlocal games: from CHSH to BQP verification. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1342–1348. IEEE, 2023.

28 Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10(7):073013, 2008. doi:10.1088/1367-2630/10/7/073013.

29 Connor Paddock, William Slofstra, Yuming Zhao, and Yangchen Zhou. An operator-algebraic formulation of self-testing. *Annales Henri Poincaré*, 25(10):4283–4319, October 2023. doi:10.1007/s00023-023-01378-y.

30 S. Pironio, A. Acin, S. Massar, A. Boyer de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by Bell's theorem. *Nature*, 464:1021–1024, 2010. doi:10.1038/nature09008.

31 Stefano Pironio, Antonio Acin, Nicolas Brunner, Nicolas Gisin, Serge Massar, and Valerio Scarani. Device-independent quantum key distribution secure against collective attacks. *New Journal of Physics*, 11(4):045021, 2009. doi:10.1088/1367-2630/11/4/045021.

32 Ben Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496(7446):456–460, 2013.

33 Valerio Scarani. The device-independent outlook on quantum physics. *Acta Physica Slovaca*, 62(4):347–409, 2012.

34 Ivan Šupić and Joseph Bowles. Self-testing of quantum systems: a review. *Quantum*, 4:337, September 2020. doi:10.22331/q-2020-09-30-337.

35 Tzyh Haur Yang and Miguel Navascués. Robust self-testing of unknown quantum systems into any entangled two-qubit states. *Physical Review A*, 87:050102, May 2013. doi:10.1103/PhysRevA.87.050102.

A Appendix

A.1 Efficient quantum circuits and homomorphic encryption

To define a quantum homomorphic encryption scheme we require the following concepts from quantum cryptography.

▶ **Definition 13.** *A procedure $\mathcal{P}$ is quantum polynomial time (QPT) if:*

1. *there exists a uniform logspace family of quantum circuits that implement $\mathcal{P}$, and*
2. *the runtime of the circuit is polynomial in the number of qubits and the security parameter $\lambda \in \mathbb{N}$.*

A family of quantum states $\mathcal{F}$ is QPT (preparable) if there is a QPT $\mathcal{P}$ for preparing $\mathcal{F}$.

We now define a quantum homomorphic encryption (QHE) scheme. A formal definition of QHE first appeared in [9]. We follow the description of QHE outlined in [19, 8]:

▶ **Definition 14.** *A quantum homomorphic encryption scheme* Q *for a family of circuits $\mathcal{C}$ consists of a security parameter $\lambda \in \mathbb{N}$ and the following algorithms:*

(i) *A PPT algorithm* Gen *which takes as input a unary encoding 1^λ of the security parameter $\lambda \in \mathbb{N}$ and outputs a secret key* sk.

(ii) *A PPT algorithm* Enc *which takes as input the secret key* sk *and a plaintext $x \in \{0,1\}^n$ and produces a ciphertext $\chi \in \{0,1\}^k$.*

(iii) *A QPT algorithm* Eval *which takes as input a classical description of a quantum circuit* $\mathsf{C} : \mathcal{H} \otimes (\mathbb{C}^2)^{\otimes n} \to (\mathbb{C}^2)^{\otimes m}$ *from $\mathcal{C}$, a quantum plaintext $|\Psi\rangle \in \mathcal{H}$ on a Hilbert space, a ciphertext χ, and evaluates a quantum circuit* $\mathsf{Eval}_{\mathsf{C}}(|\Psi\rangle \otimes |0\rangle^{\mathrm{poly}(\lambda,n)}, \chi)$ *producing a ciphertext $\alpha \in \{0,1\}^\ell$.*

(iv) *A QPT algorithm* Dec *which takes as input ciphertext α, and secret key* sk, *and produces a quantum state $|\Psi'\rangle$.*

Although the existence of algorithms (i)-(iv) defines a QHE scheme, we consider several additional important properties a scheme may or may not possess:

1. (Correctness with auxiliary input). For every security parameter $\lambda \in \mathbb{N}$, secret key $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$, classical circuit $\mathsf{C} : \mathcal{H}_A \otimes (\mathbb{C}^2)^{\otimes n} \to \{0,1\}^m$, quantum state $|\Psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$, plaintext $x \in \{0,1\}^n$ ciphertext $\chi \leftarrow \mathsf{Enc}(x, \mathsf{sk})$, the following procedures produce states with negligible trace distance with respect to λ:
   a. Starting from the pair $(x, |\Psi\rangle_{AB})$, run the quantum circuit C on register A, outputting the classical string $a \in \{0,1\}^m$ along with the contents of register B.
   b. Starting from $(\chi, |\Psi\rangle_{AB})$, run the circuit $\mathsf{Eval}_C(\cdot)$ on register A, obtaining ciphertext $\alpha \in \{0,1\}^\ell$, output $a' = \mathsf{Dec}(\alpha, \mathsf{sk})$ along with the contents of register B.

2. (Security against efficient quantum distinguishers). Fix a secret key $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$. Any quantum polynomial time adversary $\mathfrak{A}$ with access to $\mathsf{Enc}(\cdot, \mathsf{sk})$ (but does not know sk) cannot distinguish between ciphertexts $\chi \leftarrow \mathsf{Enc}(x_0, \mathsf{sk})$ and $\chi' \leftarrow \mathsf{Enc}(x_1, \mathsf{sk})$ with non-negligible probability in λ, where x_0 and x_1 are any plaintexts chosen by the adversary. That is

$$|\Pr[\mathfrak{A}^{\mathsf{Enc}(x_0,\mathsf{sk})}(x_0) = 1] - \Pr[\mathfrak{A}^{\mathsf{Enc}(x_0,\mathsf{sk})}(x_1) = 1]| \leq \mathrm{negl}(\lambda),$$

for all pairs (x_0, x_1).

The KLVY compilation procedure requires schemes that satisfy (1) and (2). QHE schemes satisfying (1) and (2) have been described in [21, 8].

A.2 Proofs

Proof of Lemma 2. Let $V_y^{(\lambda)} : \mathcal{H}'^{(\lambda)} \to \mathbb{C}^{|\mathcal{B}|} \otimes \mathcal{H}'^{(\lambda)}$ be the isometry defined by

$$V_y^{(\lambda)}|\phi\rangle = \sum_{b \in \mathcal{B}} |b\rangle \otimes \sqrt{N'^{(\lambda)}_{b|y}}|\phi\rangle, \ \forall |\phi\rangle \in \mathcal{H}'^{(\lambda)}. \tag{30}$$

Furthermore, let $U_y^{(\lambda)}$ be the unitary which satisfies $U_y^{(\lambda)}(|0\rangle \otimes |\phi\rangle) = V_y^{(\lambda)}|\phi\rangle$ for all $|\phi\rangle \in \mathcal{H}'^{(\lambda)}$. Define the projectors,

$$\tilde{N}^{(\lambda)}_{b|y} := U_y^{(\lambda)\dagger}(|b\rangle\langle b| \otimes \mathbb{I})U_y^{(\lambda)}. \tag{31}$$

Since $|\mathcal{B}|$ is constant with respect to λ and each $N'^{(\lambda)}_{b|y}$ is QPT, the resulting PVMs $\{\tilde{N}^{(\lambda)}_{b|y}\}_{b \in \mathcal{B}}$ are QPT for every $y \in \mathcal{Y}$. For the sub-normalized states, let $|\tilde{\Psi}_{\alpha|\chi}\rangle \in \mathcal{H}'^{(\lambda)} \otimes \tilde{\mathcal{H}}^{(\lambda)}$ be any purification[3] of $\rho^{(\lambda)}_{\alpha|\chi}$ with $\mathcal{H}'^{(\lambda)} \cong \tilde{\mathcal{H}}^{(\lambda)}$, and define

$$|\Psi^{(\lambda)}_{\alpha|\chi}\rangle := |0\rangle \otimes |\tilde{\Psi}^{(\lambda)}_{\alpha|\chi}\rangle \in \mathbb{C}^{|\mathcal{X}|} \otimes \mathcal{H}'^{(\lambda)} \otimes \tilde{\mathcal{H}}^{(\lambda)} =: \mathcal{H}^{(\lambda)}. \tag{32}$$

Again, since $|\mathcal{X}|$ is constant with respect to λ the (sub-normalized) states $|\Psi^{(\lambda)}_{\alpha|\chi}\rangle$ are QPT-preparable. Now, extend each $\tilde{N}^{(\lambda)}_{b|y}$ to act trivially on the purifying system $\tilde{\mathcal{H}}^{(\lambda)}$ by defining $N^{(\lambda)}_{b|y} := \tilde{N}^{(\lambda)}_{b|y} \otimes \mathbb{I}$. We observe

$$\begin{aligned} tr\left[N^{(\lambda)}_{b|y}|\Psi^{(\lambda)}_{\alpha|\chi}\rangle\langle\Psi^{(\lambda)}_{\alpha|\chi}|\right] &= tr\left[\tilde{N}^{(\lambda)}_{b|y} tr_{\tilde{Q}}[|\Psi^{(\lambda)}_{\alpha|\chi}\rangle\langle\psi^{(\lambda)}_{\alpha|\chi}|]\right] \\ &= tr[\tilde{N}^{(\lambda)}_{b|y}(|0\rangle\langle 0| \otimes \rho^{(\lambda)}_{\alpha|\chi})] \\ &= tr\left[\left(|b\rangle\langle b| \otimes \mathbb{I}\right)U_y^{(\lambda)}\left(|0\rangle\langle 0| \otimes \rho_{\alpha|\chi}\right)U_y^{(\lambda)\dagger}\right] \\ &= tr\left[\sqrt{N'^{(\lambda)}_{b|y}}\rho^{(\lambda)}_{\alpha|\chi}\sqrt{N'^{(\lambda)}_{b|y}}\right] = p^{(\lambda)}(\alpha, b|\chi, x), \end{aligned} \tag{33}$$

where $\tilde{Q}$ denotes the purifying system $\tilde{\mathcal{H}}^{(\lambda)}$. Since $\tilde{\mathcal{H}}^{(\lambda)}$ has the same dimensions as $\mathcal{H}^{(\lambda)}$, the PVMs $N^{(\lambda)}_{b|y}$ are indeed QPT. ◀

[3] Strictly speaking, since $\rho^{(\lambda)}_{\alpha|\chi}$ is sub-normalized, $|\tilde{\Psi}^{(\lambda)}_{\alpha|\chi}\rangle$ is equal to the purification of $\rho^{(\lambda)}_{\alpha|\chi}/tr[\rho^{(\lambda)}_{\alpha|\chi}]$ weighted by $tr[\rho^{(\lambda)}_{\alpha|\chi}]$, whenever $tr[\rho^{(\lambda)}_{\alpha|\chi}] > 0$.

Proof of Theorem 3. To begin, we write

$$\begin{aligned}\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[P^\dagger P] &= \sum_{ij} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[(A_x)^{k_i} w_i(B_0,B_1)(A_x)^{k_j} w_j(B_0,B_1)] \\ &= \sum_{ij} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[(A_x)^{k_i+k_j} \bar{w}_{ij}(B_0,B_1)],\end{aligned} \tag{34}$$

where we used the linearity of $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\cdot]$ in the first line, and in the second line we used the fact that $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[w] = \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\bar{w}]$ (where $\bar{w}$ is the canonical form of the monomial w), and defined $\bar{w}_{ij}$ to be the canonical form of $w_i w_j$. We now need to consider two types of terms. First, when $k_i \oplus k_j = 0$, we apply the definition in Equation (12) in conjunction with Lemma 1 to write

$$\begin{aligned}&\sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\bar{w}_{ij}(B_0,B_1)] \\ &= \mathop{\mathbb{E}}_{x'\in\mathcal{X}} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x')=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \left(\sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \bar{w}_{ij}(B_0,B_1) \right) |\Psi_{\alpha|\chi}\rangle \\ &\approx_{\mathrm{negl}(\lambda)} \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \left(\sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \bar{w}_{ij}(B_0,B_1) \right) |\Psi_{\alpha|\chi}\rangle \\ &= \sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \bar{w}_{ij}(B_0,B_1) |\Psi_{\alpha|\chi}\rangle.\end{aligned} \tag{35}$$

When $k_i \oplus k_j = 1$, we can apply Equation (13) directly. Putting these two together, we observe

$$\begin{aligned}&\sum_{ij} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[(A_x)^{k_i+k_j} \bar{w}_{ij}(B_0,B_1)] \\ &= \sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[\bar{w}_{ij}] + \sum_{ij:k_i\oplus k_j=1} \gamma_i^* \gamma_j \tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[A_x \bar{w}_{ij}] \\ &\approx_{\mathrm{negl}(\lambda)} \sum_{ij:k_i\oplus k_j=0} \gamma_i^* \gamma_j \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \bar{w}_{ij}(B_0,B_1) |\Psi_{\alpha|\chi}\rangle \\ &+ \sum_{ij:k_i\oplus k_j=1} \gamma_i^* \gamma_j \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha (-1)^{\mathsf{Dec}(\alpha)} \langle\Psi_{\alpha|\chi}| \bar{w}_{ij}(B_0,B_1) |\Psi_{\alpha|\chi}\rangle \\ &= \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \sum_{ij} (-1)^{\mathsf{Dec}(\alpha)\cdot(k_i+k_j)} \gamma_i^* \gamma_j w_i(B_0,B_1) w_j(B_0,B_1) |\Psi_{\alpha|\chi}\rangle \\ &= \mathop{\mathbb{E}}_{\chi:\mathsf{Enc}(x)=\chi} \sum_\alpha \langle\Psi_{\alpha|\chi}| \left| \sum_i (-1)^{\mathsf{Dec}(\alpha)\cdot k_i} \gamma_i w_i(B_0,B_1) \right|^2 |\Psi_{\alpha|\chi}\rangle \geq 0.\end{aligned} \tag{36}$$

Where in the fifth line, we used the fact that Bob's observables satisfy the canonical relations, so we can always replace the canonical monomial $\bar{w}_{ij}$ with $w_i w_j$. The final line is obtained by noting the square inside the expectation. We therefore conclude $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[P^\dagger P] \approx_{\mathrm{negl}(\lambda)} h$ for some $h \geq 0$, which implies $|\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[P^\dagger P] - h| \leq \mathrm{negl}(\lambda)$, and $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}[P^\dagger P] \geq h - \mathrm{negl}(\lambda) \geq -\mathrm{negl}(\lambda)$ as required. Lastly, it is straightforward to verify from the definition that for any Bell functional I we have $\tilde{\mathbb{E}}_{\widetilde{\mathrm{Q}}}(I)$ recovers the expected value under $\widetilde{\mathrm{Q}}$. ◀

Efficient Quantum Pseudorandomness from Hamiltonian Phase States

John Bostanci
Columbia University, New York, NY, USA

Jonas Haferkamp
Harvard University, Cambridge, MA, USA

Dominik Hangleiter
QuICS, University of Maryland & NIST, College Park, MD, USA
Simons Institute for the Theory of Computing, University of California at Berkeley, CA, USA

Alexander Poremba
Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Quantum pseudorandomness has found applications in many areas of quantum information, ranging from entanglement theory, to models of scrambling phenomena in chaotic quantum systems, and, more recently, in the foundations of quantum cryptography. Kretschmer (TQC '21) showed that both pseudorandom states and pseudorandom unitaries exist even in a world without classical one-way functions. To this day, however, all known constructions require classical cryptographic building blocks which are themselves synonymous with the existence of one-way functions, and which are also challenging to implement on realistic quantum hardware.

In this work, we seek to make progress on both of these fronts simultaneously – by decoupling quantum pseudorandomness from classical cryptography altogether. We introduce a quantum hardness assumption called the *Hamiltonian Phase State* (HPS) problem, which is the task of decoding output states of a random instantaneous quantum polynomial-time (IQP) circuit. Hamiltonian phase states can be generated very efficiently using only Hadamard gates, single-qubit Z rotations and CNOT circuits. We show that the hardness of our problem reduces to a worst-case version of the problem, and we provide evidence that our assumption is plausibly *fully quantum*; meaning, it cannot be used to construct one-way functions. We also show information-theoretic hardness when only few copies of HPS are available by proving an approximate t-design property of our ensemble. Finally, we show that our HPS assumption and its variants allow us to *efficiently* construct many pseudorandom quantum primitives, ranging from pseudorandom states, to quantum pseudoentanglement, to pseudorandom unitaries, and even primitives such as public-key encryption with quantum keys.

2012 ACM Subject Classification Theory of computation → Cryptographic primitives

Keywords and phrases Quantum pseudorandomness, quantum phase states, quantum cryptography

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.9

Related Version *Full Version*: https://arxiv.org/abs/2410.08073

Funding *John Bostanci*: Supported under AFORS (award FA9550-21-1-036) and NSF CAREER (award CCF2144219).
Jonas Haferkamp: Supported by the Harvard Quantum Initiative postdoctoral fellowship.
Dominik Hangleiter: Supported by the QuICS Hartree fellowship.
Alexander Poremba: Supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under contract number DE-SC0012704.

Acknowledgements This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by NSF QLCI Grant No. 2016245.

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 9; pp. 9:1–9:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Pseudorandomness [34, 68] is ubiquitous in theoretical computer science and has found applications in many areas, ranging from cryptography, to computational complexity, to the study of randomized algorithms, and even to combinatorics. The celebrated result of Håstad, Impagliazzo, Levin, and Luby [44] shows that one can construct a *pseudorandom generator* from any one-way function – a function that is easy to evaluate but computationally hard to invert. Pseudorandom generators can then in turn be used to construct more advanced cryptographic primitives, such as *pseudorandom functions* [35], i.e., keyed families of functions that appear random to any computationally bounded observer. This fact has elevated the notion of a one-way function as the minimal assumption in all of theoretical cryptography. One-way functions are typically built from well-studied mathematical conjectures, such as the hardness of factoring [61] and discrete logarithms [54], decoding error correcting codes [13, 4], or finding short vectors in high-dimensional lattices [59]. More advanced cryptographic primitives (which are believed to lie beyond what is generically possible to construct from any one-way function), such as public-key encryption, tend to require highly structured assumptions which are more susceptible to algorithmic attacks – particularly by quantum computers [65], which has led to the design of *post-quantum assumptions* [3].

In quantum cryptography, there has recently been a significant interest in so-called "fully quantum" cryptographic primitives (occasionally referred to as *MicroCrypt* primitives) which are potentially *weaker* than the conventional minimal assumptions used in classical cryptography. Here, the notion of *quantum pseudorandomness* has emerged as the natural quantum analogue of pseudorandomness in the classical world [46, 50, 2]. In particular, Ji, Liu and Song [46] proposed the notion of pseudorandom states [46] and pseudorandom unitaries as the natural quantum analogues of pseudorandom generators [44] and pseudorandom functions [35], respectively. The work of Kretschmer [50, 51] has shown that such fully quantum cryptographic primitives can exist in a world in which no classical cryptography exists – including one-way functions. At the same time, quantum pseudorandomness has applications in many areas of quantum information, ranging from entanglement theory [2, 16, 32], quantum learning theory [70], to models of scrambling phenomena in chaotic quantum systems [49, 31], and, more generally, even in the foundations of quantum cryptography [46, 50, 51, 57, 5, 17, 15, 48, 10].

Limitations of existing constructions. Despite strong evidence that MicroCrypt primitives such as pseudorandom quantum states and pseudorandom unitaries lie "below" one-way functions [50, 51], all known constructions implicitly make use of one-way functions (or other assumptions which are themselves synonymous with the existence of one-way functions) [46, 19, 55]. This naturally begs the question:

Is it possible to construct fully quantum primitives, including quantum pseudorandomness, from quantum rather than classical hardness assumptions?

Instantiating fully quantum primitives from a concrete and well-founded quantum hardness assumption (rather than from the existence of one-way functions) has remained a long standing open problem [6, 57].

Moreover, the fact that quantum pseudorandom states and unitaries are built from classical one-way functions makes them nearly impossible to realize on realistic quantum hardware. In some sense, this is inherent because cryptographic pseudorandom functions are highly complex by design [11], and therefore require a massive computational overhead to implement coherently. As a result, this severely limits the potential of using quantum pseudorandomness in practical applications; for example in the context of entanglement theory [2, 16], or when studying the emergence of thermal equilibria in isolated many-body systems [32], or when modeling scrambling phenomena in chaotic quantum systems [49]. A second limitation of existing pseudorandom constructions is therefore also the notion of quantum efficiency, which begs the question:

Are there more efficient constructions of quantum pseudorandomness which can be implemented on realistic quantum hardware?

Making progress on both of these questions would not only lead to new insights in the foundations of quantum cryptography and the study of quantum hardness assumptions more generally, but also make quantum pseudorandomness more useful in practice. To this day, however, no concrete fully quantum hardness assumption has been explored in an attempt to answer this question.

Towards a fully quantum assumption. In order to plausibly claim that quantum pseudorandomness and other fully quantum cryptographic primitives exist in a world in which classical cryptography does not, we must construct these primitives from new assumptions that do not themselves imply classical cryptography.

The history of cryptography has taught us that finding good and well-founded cryptographic assumptions is not at all an easy task – even entirely plausible assumptions have often found surprising attacks [66, 26, 12]. What makes a new cryptographic assumption reasonable? While no widely agreed upon standards exist [36], the conventional belief is to use assumptions

- which are rooted in a well-studied problem (ideally, a problem that has already been analyzed for many years) and which seems intractable in the worst case;
- for which there is a natural notion of what constitutes a "random instance" of the problem; moreover, such an instance can always be efficiently generated;
- for which there is evidence of average-case hardness, ideally in the form of a worst-case to average-case reduction;
- which can be connected to other assumptions or computational tasks that have been studied over the years, and
- which have enough structure to enable interesting cryptographic primitives.

A natural candidate for constructing quantum pseudorandomness (and other fully quantum cryptographic primitives) is via *random quantum circuits.* In fact, the computational pseudorandomness of random quantum circuits appears to be a folklore conjecture and is widely believed among many quantum computer scientists. As we are unaware of a concrete technical conjecture, we provide such a formulation here.

▶ **Conjecture 1** (Random quantum circuits give rise to pseudorandom unitaries). *Consider n-qubit random quantum circuits with m gates defined by repeating the following process m times independently at random: Draw a random pair (i, j) of qubits and apply a gate from a universal gate set $\mathsf{G} \subset SU(4)$ to the qubits i and j. Then, there exist univeral constants $c > 0$ and $C_{\mathsf{G}} > 0$ (depending on the gate set G) such that random quantum circuits with $m \geq C_{\mathsf{G}} n^c$ gates form ensembles of pseudorandom unitaries.*

We note that many other possible formulations (e.g. with specific geometric architectures or only regarding pseudorandom states) are also possible. Indeed, if Conjecture 1 holds even with exponential security, then Ref. [63] implies that a simple ensemble of random quantum circuits in a 1D architecture of depth polylog(n) is also pseudorandom.

Conjecture 1 can be seen as a direct quantum analogue of a claim that was first proposed by Gowers [37] who conjectured that random reversible quantum circuits form pseudorandom permutations on bitstrings. This conjecture has inspired multiple recent works in classical cryptography. For instance, it was recently proven by He and O'Donnell [42] that the Luby-Rackoff [52] construction of pseudorandom permutations from pseudorandom functions can be implemented with reversible permutations. Random reversible circuits have recently also inspired entirely new approaches for constructing program obfuscation schemes [25].

Gowers originally conjectured the emergence of pseudorandomness when attempting to prove that random quantum circuits converge quickly to ensembles of t-wise independent permutations [37] (this bound was further improved later on towards an optimal scaling [43, 24, 28]. In fact, this property can itself be viewed as evidence for pseudorandomness. It turns out that random quantum circuits satisfy an analogous property by converging nearly optimally towards approximate t-designs [28, 20, 40, 63].

However, we currently do not have rigorous evidence for Conjecture 1; for example, in terms of a worst-to-average reduction for a corresponding learning problem. Moreover, and maybe more importantly, it is unclear how one would use unstructured random quantum circuits to construct more advanced quantum cryptographic primitives. A similar situation arises for general one-way functions, which require additional structure to build more advanced cryptographic applications, such as public-key encryption. It could very well be the case that random quantum circuits are simply *too mixing* to be a useful in the context of quantum cryptography. A natural way forward is to search for a sweet spot – an ensemble of random quantum circuits that is sufficiently structured to permit the construction of interesting cryptographic primitives but which, at the same time, is sufficiently mixing to guarantee security.

2 Our contributions

In this work, we simultaneously address the two major open problems in the field of quantum pseudorandomness and propose the first well-founded and fully quantum hardness assumption. To this end, we follow the strategy sketched above, and propose a family of quantum states which we call *Hamiltonian Phase States*. These states are a family of quantum states which are "maximally quantum" in the sense that the state has support on all bitstrings with amplitudes equal in magnitude, but varying phases. Hamiltonian Phase States are generated by a family of commuting *instantaneous quantum polynomial-time* (IQP) circuits which generalize the X programs proposed by Shepherd and Bremner [64]. The corresponding circuits are highly structured in that they are generated by a Hamiltonian with only Z-type terms applied to the all-$|+\rangle$ state. This structure makes them amenable to rigorous analysis [47, 22, 38]. At the same time, these circuits are also believed to be sufficiently mixing and hard to simulate classically [64, 21, 23, 41]. Moreover, since Hamiltonian phase states can be generated by a commuting Hamiltonian, they admit an efficient implementation in practice.

Phase states are a natural direction to look at in the search for a fully-quantum cryptographic assumption with sufficient amounts of structure. On the one hand, this is because of their quantum advantage properties. On the other hand, the (quantum) learnability of different ensembles of phase states has been studied extensively in recent work [7].

There, the authors give optimal bounds for the sample complexity of learning many families of phase states from quantum samples, as well as upper bounds on the time complexity. Importantly, there are families of phase states generated using a small number of (long-range)

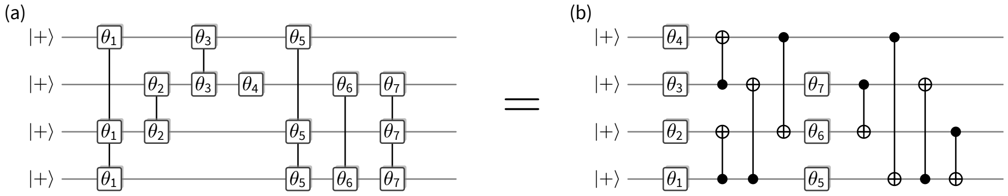

Figure 1 Hamiltonian Phase States (HPS) are generated by sequentially applying Ising-type rotations around angles θ_i to the state $|+^n\rangle = H^{\otimes n}|0^n\rangle$. (a) Example of a HPS on 4 qubits. Connected boxes at sites i, j, k with angle θ represent the unitary $\exp(i\theta Z_i Z_j Z_k)$. (b) HPS can be implemented using only single-qubit Z rotations interlaced with CNOT circuits.

gates, which cannot be learned from polynomially many samples. Following this, our proposed cryptographic assumption is that Hamiltonian Phase States are hard to learn, given quantum samples and classical side information.

Moreover, the known constructions for pseudorandom states with useful cryptographic applications are based on phase states [46]. These are generated using a single-bit output quantum-secure pseudorandom function family $\{f_k\}_k$ with

$$|\phi_k\rangle \propto \sum_x \omega_q^{f_k(x)} |x\rangle \,, \tag{1}$$

where ω_q is a q-th root of unity, for example $q = 2$ [19]. Because these states are based on a classical assumption, they require the reversible implementation of a classical PRF which requires a large number of Toffoli gates. These are extremely expensive in standard fault-tolerant constructions. However, the results of Refs. [46, 7] suggest that a more natural family of phase states which is generated by a quantum circuit with a small number of expensive gates can also yield quantum pseudorandomness. This would require gates affecting a large number of qubits, since low-degree phase states can be learned efficiently. As we show below, in spite of having terms with high support, the Hamiltonian Phase States can be generated highly efficiently using only local Z-rotations and CNOT gates.

2.1 Hamiltonian Phase States

Let $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$ be a binary matrix and let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$ be a set of uniformly random angles in the interval $[0, 2\pi)$ according to some discretization into $q = \mathsf{poly}(n)$ parts. We consider phase states of the form

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}}\right) H^{\otimes n}|0^n\rangle \,. \tag{2}$$

where, for $i \in [m]$, we denote the i-th row of $\mathbf{A}$ by $(\mathbf{A}_{i1}, \dots, \mathbf{A}_{in})$ and let

$$\bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}} = \mathsf{Z}^{\mathbf{A}_{i1}} \otimes \dots \otimes \mathsf{Z}^{\mathbf{A}_{in}} \qquad \text{for} \qquad \mathsf{Z}^0 = \mathbb{I},\;\; \mathsf{Z}^1 = \mathsf{Z}.$$

We call these states *Hamiltonian Phase States* since they can naturally be prepared as the result of a time evolution under a sparse Ising Hamiltonian. We also call the matrix $\mathbf{A}$ the *architecture* of the states, as it specifies the overall structure/location of the Ising terms. Hamiltonian Phase States with a single fixed angle $\theta_i \equiv \theta$ have been studied as a means to demonstrate verified quantum advantage, when measured in the X basis, under the name X-programs [64]. A Hamiltonian Phase State is therefore a generalized version of an X program parameterized by the pair $(\mathbf{A}, \boldsymbol{\theta})$. X programs with $\theta = \pi/8$ have the interesting property that its Fourier coefficients can be computed efficiently classically, but at the same time the simulation of such X-programs is believed to be classically intractable [64, 21, 23, 41].

Our cryptographic assumption rests on the apparent hardness of *learning* Hamiltonian Phase States (or generalized X programs), which was highlighted in recent work [7]. Concretely, our quantum computational assumption amounts to the conjecture that our ensemble of Hamiltonian phase states satisfies the following two properties:

- Random Hamiltonian Phase States are *hard to invert* in the following sense: given $|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle^{\otimes t}$, for any $t = \mathsf{poly}(n)$, it is computationally difficult to reverse-engineer the angles $\boldsymbol{\theta}$ and architecture $\mathbf{A}$. This means that the ensemble $\{|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle\}_{\boldsymbol{\theta},\mathbf{A}}$ gives rise to a so-called *one-way state generator* (OWSG).
- Random Hamiltonian Phase States are *hard to distinguish from Haar random states* in the following sense: given $|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle^{\otimes t}$, for any $t = \mathsf{poly}(n)$, it is computationally difficult to distinguish $|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle^{\otimes t}$ from $|\Psi\rangle^{\otimes t}$, where $|\Psi\rangle$ is a Haar random state. This means that the ensemble $\{|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle\}_{\boldsymbol{\theta},\mathbf{A}}$ gives rise to a so-called *pseudorandom state generator* (PRSG).

We can call the two assumptions above the search (respectively, decision) variant of *Hamiltonian Phase State* assumption ($\mathsf{HPS}_{n,m,q,\chi}$). Here, $n, m \in \mathbb{N}$ are circuit parameters, q is a discretization parameter for the interval $[0, 2\pi)$, and χ is a distribution over the choice of matrix $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$; typically, χ is chosen to be the uniform distribution.

There is some evidence that $\mathsf{HPS}_{n,m,q,\chi}$ is a reasonable assumption for constructing pseudorandom states. Brakerski and Shmueli [19] show that the states $DH^{\otimes n}|0^n\rangle$ form a state t-design when the diagonal operator D consists of a $2t$-wise independent binary phase operator. Previously, Nakata, Koashi and Murao [58] also showed that the states $DH^{\otimes n}|0^n\rangle$, where D is a diagonal operator composed of appropriate diagonal gates with random phases, form a t-design. Starting from this intuition, we now provide rigorous evidence for the hardness of the $\mathsf{HPS}_{n,m,q,\chi}$ assumption.

2.2 Overview of our Results

In this work, we establish $\mathsf{HPS}_{n,m,q,\chi}$ as a well-founded quantum computational assumption. Specifically, we address each of the meta-criteria we mentioned before:

- (Evidence of worst-case hardness) The learnability of ensembles of phase states has been studied extensively in recent work [7], and has been found to have exponential time complexity in the worst case (despite only having polynomial sample complexity).
- (Notion of a random instance) A random Hamiltonian phase state $|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle$, e.g., as in Equation (2), is naturally defined in terms of a random binary matrix $\mathbf{A} \leftarrow \mathbb{Z}_q^{m\times n}$ and a random set of angles $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$. Hence, it can be efficiently generated by a simple quantum circuit comprising $O(m/n \cdot n^2)$ CNOT gates, n Hadamard gates, and m single-qubit Z rotations.
- (Evidence of average-case hardness) Our learning task admits a worst-case to average-case reduction. We separately show how to re-randomize the architecture and the set of angles. Thus, the hardness of our problem reduces to a worst-case version.
- (Relation to other problems) We draw a connection between the task of learning Hamiltonian Phase states and the security of classical *Goppa codes* and the well-known McEliece cryptosystem.

- (Cryptographic applications) Hamiltonian Phase states have a sufficient amount of structure which suffices to construct a number of interesting cryptographic primitives which we sketch in detail in Section 2.3.

Finally, we also provide evidence that $\mathsf{HPS}_{n,m,q,\chi}$ is plausibly fully-quantum and does not allow one to construct one-way functions. In particular, we note that the result of [50] indicates that the idealized versions of any assumption that yields *only* pseudorandom states can not be used to build one-way functions in a black-box way. We further discuss the implications of the fact that HPS states are state t-designs on this reduction, noting that the resulting concentration properties by themselves rule out one-way function constructions that do not simulateneously measure many copies of the HPS state. For all of the primitives we construct, in addition to just constructing these primitives from our hardness assumption, we argue that constructing them from our hardness assumption yields more efficient and practical implementations of these primitives (if and when fault-tolerant quantum computers become widely available).

2.3 Applications

In this section, we give an overview of all the applications which are enabled by the HPS assumption. Besides the natural application of constructing efficient one-way state generators and pseudorandom state generators, which essentially follow by definition of our assumption, we also construct a number of other interesting applications that are relevant in quantum information science more broadly.

Quantum Trapdoor Functions and Public-Key Encryption with Quantum Public Keys

Recent work of Coladangelo [29] introduced the notion of a quantum trapdoor function (QTF). This primitive is essentially a variant of a one-way state generator that also features a secret trapdoor which makes inversion possible. QTFs are interesting in the sense that they *almost* enable public-key encryption: two parties can communicate classical messages over a quantum channel without ever exchanging a shared key in advance – the only caveat being that this requires the public keys to be quantum states [29]. Using a construction based on binary-phase states, Coladangelo [29] showed that quantum trapdoor functions exist, if post-quantum one-way functions exist. However, to this day, it remains unclear how to construct QTFs from assumptions which are potentially weaker than one-way functions, such as the existence of pseudorandom states.

In the full version, we show how to construct QTFs from our (decisional) HPS assumption, which yields the first construction of QTFs from an assumption which is plausibly weaker than that of one-way functions. We believe that this application strongly highlights the versatility of Hamiltonian Phase states in the context of quantum cryptography; for example, it is far less clear how to construct QTFs from other, less structured, assumptions such as genuinely random quantum circuits via Conjecture 1.

Quantum Pseudoentanglement

The notion of pseudoentanglement [2, 16] has found many applications in quantum physics, for example to study the emergence of thermal equilibria in isolated many-body systems [32]. Pseudoentangled states have also been viewed as a potential tool for probing computational aspects of the AdS/CFT correspondence, which physicists believe may shed insight onto the behavior of black holes in certain simplified models of the universe. We note that it is currently not known how to construct these from *any* assumption other than one-way functions. In the full version, we give a construction of pseudoentangled states from our HPS assumption, which yields the first construction of pseudoentanglement from an assumption which is plausibly weaker than that of one-way functions. Our proof sheds new light on the entanglement properties of random IQP circuits more generally.[1] Therefore, we believe that this contribution is of independent interest. Moreover, as we point out in the next section, our construction is also highly efficient and could enable implementations of quantum pseudoentanglement in practical scenarios.

Pseudorandom Unitaries

Pseudorandom unitaries are families of unitaries that are indistinguishable from Haar random unitaries in the presence of computationally bounded adversaries. They are widely considered the most powerful fully-quantum primitive, and there has been a long line of work towards constructing them from the existence of one-way functions [5, 18, 55, 27], eventually resulting in the most recent breakthrough result by Ma and Huang [53].

The result of [53] show that the ensemble of unitaries, colloquially known as the PFC-ensemble [55], form an approximation to a Haar random unitary. However, this construction is not well suited for the HPS assumption, which, in some sense, provides a pseudorandom *diagonal* unitary. In the full version, we provide a plausible construction of efficient pseudorandom unitaries from an natural assumption which is directly related to our HPS assumption: alternating applications of HPS unitaries and Hadamards.

2.4 Physical Implementations

Hamiltonian Phase States with m terms on n qubits can be generated very efficiently compared to phase states constructed from pseudorandom functions: to prepare a HPS, we require only a layer of Hadamard gates, followed by $\lceil m/n \rceil$ alternating layers of single-qubit Z rotations and CNOT circuits. To see this, we observe two facts. First,

$$\mathsf{CNOT}_{k,l} e^{i\theta Z_l} = e^{i\theta Z_k Z_l} \mathsf{CNOT}_{k,l}, \tag{3}$$

where $\mathsf{CNOT}_{k,l}$ is controlled on qubit k and targeted on qubit l. Second,

$$\mathsf{CNOT}_{k,l} \left|+^n\right\rangle = \left|+^n\right\rangle. \tag{4}$$

More generally, if C is a circuit comprised of CNOT gates, then $C\left|x\right\rangle = \left|\mathbf{C}\cdot x\right\rangle$, where $\mathbf{C} \in \mathrm{GL}(n, \mathbb{Z}_2)$ is an invertible binary matrix. Given an HPS, decompose its architecture matrix $\mathbf{A} \in \mathbb{Z}_2^{n\times m}$ into $\lceil m/n \rceil$ submatrices of n rows (except for the last one), and suppose each of those submatrices has full rank. The HPS $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ can then be prepared as

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \left(C_{\lceil m/n \rceil} \prod_{i=(\lceil m/n \rceil -1)n}^{m} e^{i\theta_i Z_i} \right) \cdots \left(C_1 \prod_{i=1}^{n} e^{i\theta_i Z_i} \right) \left|+^n\right\rangle \tag{5}$$

where the CNOT circuits C_k are chosen such that the first n rows of $\mathbf{A}$ are given by $\mathbf{C}_{\lceil m/n \rceil} \cdots \mathbf{C}_1$, the second n rows by $\mathbf{C}_{\lceil m/n \rceil} \cdots \mathbf{C}_2$, and so on, see Figure 1 for an example. If the rank condition above is not satisfied, decompose $\mathbf{A}$ into the minimal number ℓ of submatrices with full rank, and proceed as above. The smallest meaningful example of such

[1] To the best of our knowledge, such bounds for random IQP circuits were previously not known.

states – with n random, linearly independent terms – can thus be prepared using a single layer of rotations and a CNOT circuit. By the fact that $\mathrm{GL}(n, \mathbb{Z}_2)$ is a group, uniformly random architecture matrices $\mathbf{C}_i$ and phases θ_i generate a uniformly random HPS.

This protocol for the implementation of HPS is also interesting from an early fault-tolerance perspective, since there are quantum codes in which all required operations are transversal, yielding a highly efficient fault-tolerant implementation. To see this, consider a $q = 2^d$-fold discretization of the unit circle. Now, we observe that there are d-dimensional CSS codes with a transversal $Z^{1/2^{d-1}}$ gate such as the $[[2^d - 1, 1, 3]]$ simplex code [69]. By the fact that they are CSS codes, they also admit a transversal CNOT gate between code blocks. This means that HPS can be prepared using transversal in-block $Z^{1/2^{d-1}}$ gates as well as inter-block CNOT gates, making them amenable to implementations in early fault-tolerant architectures such as reconfigurable atom arrays [14], or trapped ion processors [62, 60] in which arbitrary inter-block connectivities can be achieved.

3 Hamiltonian Phase State Assumption

In this section, we give a formal definition of our hardness assumption. Recall that an n-qubit Hamiltonian Phase State is of the form

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}}\right) H^{\otimes n}|0^n\rangle$$

where $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$ is a binary matrix and $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$ is a set of angles in the interval $[0, 2\pi)$. To avoid matters of precision, we introduce a discretization parameter $q \in \mathbb{N}$ with $q = \mathsf{poly}(n)$ and partition the interval $[0, 2\pi)$ into q parts via the set

$$\Theta_q := \left\{\frac{2\pi k}{q} \;:\; k \in \{0, 1, \dots, q-1\}\right\}.$$

We now introduce two variants of our hardness assumption.

3.1 Search Variant

Our first variant considers a search problem. Roughly speaking, it says that given many copies of a random Hamiltonian phase state, it is computationally difficult to reverse-engineer its architecture and its angles. Therefore, our assumption says that an ensemble of Hamiltonian Phase states forms a one-way state generator [57].

We now give a formal definition.

▶ **Definition 2** (Search HPS). *Let $n \in \mathbb{N}$ denote the security parameter, and let m and q be integers (possibly depending on n). Let χ be a distribution with support over matrices in $\mathbb{Z}_2^{m\times n}$. Then, the (search) Hamiltonian Phase State assumption ($\mathsf{HPS}_{n,m,q,\chi}$) states that, for any number of copies $t = \mathsf{poly}(n)$ and for any efficient quantum algorithm $\mathcal{A}$,*

$$\Pr\left[1 \leftarrow \mathsf{Ver}(\mathbf{A}', \boldsymbol{\theta}', |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle) \;:\; \substack{\mathbf{A}\sim\chi,\, \boldsymbol{\theta}\sim\Theta_q^m \\ (\mathbf{A}', \boldsymbol{\theta}')\leftarrow\mathcal{A}(|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t})}\right] \leq \mathsf{negl}(n)\,,$$

where $\mathsf{Ver}(\mathbf{A}', \boldsymbol{\theta}', |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle)$ denotes the algorithm which applies the projective measurement

$$\{|\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle\langle\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}|, I - |\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle\langle\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}|\}$$

onto $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ and outputs 1, if the measurement succeeds, and outputs 0 otherwise. We say that a quantum algorithm solves the (search) $\mathsf{HPS}_{n,m,q,\chi}$ problem if it runs in time $\mathsf{poly}(n,m,\log q)$ and succeeds with probability at least $1/\mathsf{poly}(n,m,\log q)$.

An alternative but equivalent formulation of the security property is to say that it is computationally difficult to find a state $|\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle$ which has non-vanishing fidelity with the input state, on average over the choice of architecture and set of angles.

3.2 Decision Variant

Our second variant considers a decision problem. Roughly speaking, it says that given many copies of a random Hamiltonian phase state, it is computationally difficult to distinguish it from many copies of a Haar state. Therefore, our (decision) assumption says that an ensemble of Hamiltonian Phase states forms a pseudorandom state generator [46, 57].

▶ **Definition 3** (Decision HPS). *Let $n \in \mathbb{N}$ denote the security parameter, and let m and q be integers (possibly depending on n). Let χ be a distribution with support over matrices in $\mathbb{Z}_2^{m\times n}$. Then, the (decision) Hamiltonian Phase State assumption ($\mathsf{HPS}_{n,m,q,\chi}$) states that, for any number of copies $t = \mathsf{poly}(n)$ and for any efficient quantum distinguisher $\mathcal{D}$,*

$$\left|\Pr\left[1 \leftarrow \mathcal{D}(|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}) \;:\; {\scriptstyle \mathbf{A}\sim\chi \atop \boldsymbol{\theta}\sim\Theta_q^m}\right] - \Pr\left[1 \leftarrow \mathcal{D}(|\Psi\rangle^{\otimes t}) \;:\; {\scriptstyle |\Psi\rangle\sim\mathrm{Haar}(2^n)}\right]\right| \leq \mathsf{negl}(n)\,,$$

We say that a quantum algorithm solves the (decision) $\mathsf{HPS}_{n,m,q,\chi}$ problem if it runs in time $\mathsf{poly}(n,m,\log q)$ and succeeds with probability at least $1/\mathsf{poly}(n,m,\log q)$.

4 Evidence for Average-Case Hardness and Full Quantumness

In this section, we give several pieces of evidence for the security of the $\mathsf{HPS}_{n,m,q,\chi}$ assumption, as well as evidence that it is a fully quantum assumption.

First, in Section 4.1, we show two worst-to-average-case reductions for the $\mathsf{HPS}_{n,m,q,\chi}$ problem, and also discuss the limitations of those reductions. To this end, we first show that if the Hamiltonian architecture matrix $\mathbf{A}$ is publicly known, then there is a worst-to-average-case reduction for the angles $\theta \in \Theta_q$. Second, we show that for $m = n$, and any fixed set of angles, there is a worst-to-average-case reduction over the architecture matrices $\mathbf{A}$.

Then, we show that if χ is the uniform distribution of $m \times n$ binary matrices and $q > 2t$, Hamiltonian Phase States with $m \gtrsim nt^2$ random terms form approximate state t-designs in Section 4.2. This shows that given less than $\Omega(\sqrt{m/n})$ many copies, HPS are information-theoretically indistinguishable from Haar-random states. It also implies that the Hamiltonian Phase States contain an exponentially large set of almost orthogonal states. This implies that Hamiltonian phase states are fast mixing, giving additional evidence that the learning problem is computationally hard.

In the full version, we discuss algorithms for learning phase states with public and secret architecture matrices. In particular, we give a sample-optimal (but exponential-time) algorithm for solving the HPS problem using pretty good measurements [9, 56] and a simple algorithm that uses classical shadows [1, 45]. Moreover, we also discuss why the HPS assumption is fully quantum. To this end, we give evidence against the possibility of building one-way functions from HPS.

4.1 Worst-Case to Average-Case Reduction

We begin providing evidence for the security of our assumption by showing how in different regimes learning the parameters of the HPS problem of a fixed (worst-case) instance can be reduced to learning a random instance. Our evidence will treat the angles θ and the Hamiltonian architecture matrix $\mathbf{A}$ separately. Specifically, we will show two types of worst-to-average-case reductions. First, we will show that in a certain regime of m, n, given a copy of a HPS instance, a quantum algorithm can efficiently generate a random HPS with the same angles and architecture dimensions. Second, we will fix the Hamiltonian architecture $\mathbf{A}$ and show that given a copy of a HPS instance and its architecture $\mathbf{A}$, a quantum algorithm can generate a random HPS with the same architecture but uniformly random angles. Our worst-to-average-case reductions are therefore similar to those for, say, the Learning with Errors (LWE) problem [59], with different levels of public knowledge.

Reduction for the architecture for $m \leq n$

First, we observe that for any fixed choice of angles $\boldsymbol{\theta}$, the Hamiltonian architecture can be re-randomized if $m \leq n$ and χ is the uniform distribution over full-rank matrices $\mathcal{R}(m,n) := \{\mathbf{A} \in \mathbb{Z}_2^{m\times n} \,|\, \mathrm{rank}(\mathbf{A}) = \min(m,n)\}$. Notice that the restriction to Hamiltonian architectures with full rank is not too significant, since the probability that a uniformly random $\mathbb{Z}_2^{m\times n}$ matrix has full rank with probability[2] $\prod_{k=1}^{\min(m,n)}(1-2^{-k}) \geq 0.288$ [67]. The basic idea of the reduction is to apply a circuit composed of uniformly random CNOT gates to the given HPS instance. In the parameter regime we consider, this will have the effect of completely scrambling the Hamiltonian architecture to a uniformly random one with the same choices of m, n and subject to the full-rank constraint.

▶ **Lemma 4** (Worst-to-average-case reduction for the architecture). *Suppose there exists an algorithm $\mathcal{A}$ that runs in time T and solves the (search) $\mathsf{HPS}_{n,m,q,\chi}$ problem with probability ϵ in the average case, where χ is the uniform distribution over $\mathcal{R}(m,n)$ and $m \leq n$. Then, there exists an algorithm which runs in time $T \cdot \mathsf{poly}(n)$ and inverts Hamiltonian phase states $|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle^{\otimes t}$ with probability ϵ for a worst-case choice of architecture $\mathbf{C} \in \mathcal{R}(m,n)$, uniformly random angles $\boldsymbol{\theta}$, and for any number of copies $t = \mathsf{poly}(n)$. Here, $\mathcal{R}(m,n) = \{\mathbf{A} \in \mathbb{Z}_2^{m\times n} | \mathrm{rank}(\mathbf{A}) = \min(m,n)\}$ is the set of full-rank binary $m \times n$ matrices.*

Proof. Consider the reduction $\mathcal{B}$ which, on input $|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle^{\otimes t}$, does the following:

1. $\mathcal{B}$ samples a uniformly random invertible matrix $\mathbf{R} \sim \mathrm{GL}(n, \mathbb{Z}_2)$.
2. $\mathcal{B}$ runs the average-case solver $\mathcal{A}$ on the input

$$(U_{\mathbf{R}}\,|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle)^{\otimes t}.$$

where $U_{\mathbf{R}}$ is the n-qubit unitary transformation given by $U_{\mathbf{R}} : |x\rangle \mapsto |\mathbf{R}^{-1} \cdot x\rangle$, for $x \in \{0,1\}^n$. Finally, $\mathcal{B}$ outputs whatever $\mathcal{A}$ outputs.

Note that $U_{\mathbf{R}}$ is a quantum circuit composed just of CNOT gates and therefore efficiently implementable. Because the average-case solver $\mathcal{A}$ runs in time T, it follows that the reduction $\mathcal{B}$ runs in time $T \cdot \mathsf{poly}(n)$.

[2] See `https://math.mit.edu/~dav/genlin.pdf`, and this Stackexchange post for a proof.

Next, we show that $\mathcal{B}$ also succeeds with probability ϵ. By assumption, the worst-case instance $|\Phi^{\mathbf{C}}_{\boldsymbol{\theta}}\rangle^{\otimes t}$ consists of structured phase states states

$$|\Phi^{\mathbf{C}}_{\boldsymbol{\theta}}\rangle = \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)H^{\otimes n}|0^n\rangle\,,$$

where $\mathbf{C}\in\mathcal{R}(m,n)$ and $\boldsymbol{\theta}$ is a tuple of random angles $\boldsymbol{\theta}=(\theta_1,\dots,\theta_m)\in\Theta_q^m$. To complete the proof, it suffices to show that $U_{\mathbf{R}}\,|\psi^{\mathbf{C}}_{\boldsymbol{\theta}}\rangle)^{\otimes t}$ is distributed exactly as in the $\mathsf{HPS}_{n,n,q,\chi}$ problem, where χ is the uniform distribution over $\mathcal{R}(m,n)$. First, we make the following key observation: it follows from unitarity of $U_{\mathbf{R}}$ that

$$U_{\mathbf{R}}\,|\Phi^{\mathbf{C}}_{\boldsymbol{\theta}}\rangle = \left(U_{\mathbf{R}}\exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger}\right)U_{\mathbf{R}}H^{\otimes n}|0^n\rangle\,.$$

Because $U_{\mathbf{R}}$ is an invertible matrix, it leaves the state $H^{\otimes n}|0^n\rangle$ invariant, and thus we have $U_{\mathbf{R}}H^{\otimes n}|0^n\rangle = H^{\otimes n}|0^n\rangle$. Next, we study the action of $U_{\mathbf{R}}$ onto tensor products of Pauli operators. We find that for any index $i\in[n]$:

$$\begin{aligned}
U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger} &= \sum_{x\in\{0,1\}^n}\langle x|\,U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger}\,|x\rangle\cdot|x\rangle\langle x|\\
&= \sum_{x\in\{0,1\}^n}\langle\mathbf{R}x|\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)|\mathbf{R}x\rangle\cdot|x\rangle\langle x|\\
&= \sum_{x\in\{0,1\}^n}(-1)^{\sum_{j=1}^{n}\mathbf{C}_{ij}(\mathbf{R}x)_j}|x\rangle\langle x|\\
&= \sum_{x\in\{0,1\}^n}(-1)^{\sum_{j=1}^{n}(\mathbf{C}\cdot\mathbf{R})_{ij}x_j}|x\rangle\langle x|\\
&= \sum_{x\in\{0,1\}^n}\langle x|\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}\right)|x\rangle\cdot|x\rangle\langle x|\\
&= \bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}.
\end{aligned}$$

Because $U_{\mathbf{R}}$ is acting on a matrix exponential of a diagonal matrix, it follows that

$$\begin{aligned}
U_{\mathbf{R}}\exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger} &= \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\,U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger}\right)\\
&= \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}\right).
\end{aligned}$$

Finally, we observe that for $m=n$, $\mathcal{R}(m,n)=\mathrm{GL}(n,\mathbb{Z}_2)$, which is a group. Because $\mathbf{C}\in\mathrm{GL}(n,\mathbb{Z}_2)$ it follows that $\mathbf{C}\cdot\mathbf{R}$ is uniformly distributed whenever $\mathbf{R}\sim\mathrm{GL}(n,\mathbb{Z}_2)$. Putting everything together, it follows that $U_{\mathbf{R}}\,|\psi^{\mathbf{C}}_{\boldsymbol{\theta}}\rangle)^{\otimes t}$ is distributed precisely as in the $\mathsf{HPS}_{n,n,q,\chi}$ problem, and thus $\mathcal{B}$ succeeds with probability ϵ. The claim for $m\le n$ follows from the fact that in that case $\mathbf{C}$ is a submatrix of a $\mathrm{GL}(n,\mathbb{Z}_2)$ matrix. ◀

4.2 Hamiltonian Phase States Form Approximate State Designs

In this subsection we show that the states in the HPS ensemble form approximate state designs if $m \geq Cn$ for a constant $C > 0$. It will be convenient to view HPS as a random walk of depth m on the diagonal group. We will therefore slightly adjust the notation. Consider the following probability distribution ν on the diagonal subgroup of $SU(2^n)$: Draw a uniformly random bitstring $\boldsymbol{A}_1 \in \{0,1\}^n$ and a uniformly random angle $\theta \in [0, 2\pi)$ and apply $e^{\mathrm{i}\theta \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{1j}}}$. We can draw m such diagonal unitaries independently and multiply them. The resulting probability measure is denoted by ν^{*m}.

We will first show that $e^{\mathrm{i}\sum_{i=1}^m \theta_i \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{ij}}}$ is an approximate t-design on the diagonal group. More precisely, we prove the following theorem:

▶ **Theorem 5.** *For $m \geq 2t(2nt+\log(1/\varepsilon))$ the random unitary $e^{\mathrm{i}\theta_i \sum_{i=1}^m \bigotimes_j Z^{\boldsymbol{A}_{ij}}}$ with random A_{ij} and θ_i is a ε-approximate diagonal t-design. Moreover, the same bound holds if θ_i is drawn uniformly from $\{2\pi k/q\}_{k=1}^q$, where q is an integer satisfying $q > 2t$.*

We provide a proof of Theorem 5 in the full version. The proof of Theorem 5 is remarkably simple in comparison to the derivations of similar results for random quantum circuits [28, 20, 40]. Additionally, the constants in Theorem 5 are unusually small: In stark contrast the constants in these results are north of 10^{13}. A similar result was obtained in Ref. [39] for the related random Pauli rotations $e^{\mathrm{i}\theta P}$ for a random $\theta \in (0, 2\pi]$ and a random Pauli string P.

Theorem 5 almost directly implies the following corollary:

▶ **Corollary 6.** *For $m \geq 2t(2nt + \log(1/\varepsilon))$ the state ensemble defined by $|\Phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle = U|+^n\rangle$ for U drawn from ν^{*m} (or ν_q^{*m} for $q > 2t$) is a $\varepsilon + O(t^2/2^n)$-approximate state t-design.*

As a consequence no algorithm with access to t copies can distinguish the states $|\phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle$ from Haar random. In particular, this rules out a large class of natural attacks which make use of a small number of samples. Prominent examples in classical cryptanalysis are linear attacks (2-wise independence rules this out), and differential attacks (t-wise independence rules out $\log_2(t)$ differential attacks). Moreover, the fact that HPS with sufficiently many terms can generate arbitrary state t designs makes it seem unlikely even that there is a distinguishing algorithm using just a few more than t samples. This would mean that there is a sharp transition in the complexity of distinguishing HPS states from uniform. Thus, the t-design property gives evidence for the security of the HPS assumption.

As a consequence we can also show that HPS contains many almost orthogonal states, yielding additional evidence for the HPS assumption:

▶ **Corollary 7.** *Let $m = 100nt$, $\delta = 1 - 2^{-n/8}$ and $t \leq 2^{n/2}$. For any fixed state $|\psi\rangle$, we have with probability $1 - 2^{-\Omega(nt)}$ over the matrix $\boldsymbol{A}$ that*

$$\Pr_{\boldsymbol{\theta}}\left[|\langle\psi| \exp\left(\sum_{i=1}^m \mathrm{i}\theta_i \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{ij}} \right) |+^n\rangle|^2 \geq 1 - \delta \right] \leq 2^{-\Omega(nt)}. \tag{6}$$

We defer the proofs of Theorem 5 and Corollary 7 to the full version of the paper.

4.3 Algorithms for Learning Hamiltonian Phase States

Recall that our (search) HPS assumption can be thought of as a state discrimination task. The goal is to recover the architecture $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$ and the set of angles $\boldsymbol{\theta} \in \Theta_q^m$ given many copies of a random Hamiltonian phase state from the ensemble

$$\mathcal{E} = \left\{ |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i} \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}} \right) |+^n\rangle \right\}_{\mathbf{A}\in\mathbb{Z}_2^{m\times n},\, \boldsymbol{\theta}=(\theta_1,\dots,\theta_m)\in\Theta_q^m}.$$

In this section, we consider various learning algorithms for the (search) HPS probem. We observe that the HPS problem does in fact have polynomial quantum sample complexity, and can thus be solved information-theoretically. However, as we also observe, all known learning algorithms have exponential time complexity, which suggests that the HSP problem cannot be solved efficiently. We distinguish between the *private-key* and *public-key* setting: the former is essentially the learning task from Definition 2, whereas in the latter we further assume that the learner also has access to the architecture matrix $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$. We provide evidence that the learning tasks remains hard even if we reveal additional information about $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$ and the goal is simply to guess the angles $\boldsymbol{\theta}$.

Sample complexity of HPS and hypothesis selection. While we believe that HPS is a computationally hard problem, it can be solved information-theoretically with only polynomially many samples. In full generality, the problem of finding a fixed state ρ_j among many hypothesis states $\rho_1, \dots, \rho_M$ is called quantum hypothesis testing. Currently, the best known general algorithm is threshold search as described in [8, Theorem 1.5] requires $n\log^2(M)$ copies improving over the bound from Ref. [1]. For the HPS problem this implies an upper bound on the sample complexity of $O(n\log^2(q^m 2^{nm})) = O(n^3 m^2 \log(q))$. As the fidelities for pure states are PSD observables of rank 1, we can also use the shadow tomography protocol of Ref. [45]. Given a secret state $|\Phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle$ allows us to estimate the fidelities of all the $M = q^m 2^{nm}$ phase states up to an error of ε from $O(\log(M)/\varepsilon^2) = O(mn\log(q)/\varepsilon^2)$ samples. Then, a solver can simply list all estimated fidelities and pick the state with the largest overlap up to an error of ε.

We expect these bounds to be tight in the regime where $m \leq O(n^{\log(q)})$. For $m \to \infty$ better bounds are available at least for $q = 2^d$. In this case, the HPS instance generated b unitaries in the dth level of the Clifford hierarchy and it was proven in Ref. [7, Theorem 15] that for any state of the form

$$\exp\left(\mathrm{i} \sum_{y\in\{0,1\}^n} a_y \bigotimes_{j=1}^{n} Z^{y_j} \right) |+^n\rangle \tag{7}$$

with $a_y \in \mathbb{Z}$ a circuit description can be learned with $O(n^d)$ copies using only measurements in the standard basis.

Learning algorithms for HPS with a public architecture. In the special case when the architecture is public, our HPS assumption does in fact admit an optimal[3] but nevertheless exponential-time learning algorithm.

[3] Here, we mean an algorithm that achieves the optimal success probability for a given number of copies.

We consider the following state discrimination task, where the goal is to recover the set of angles $\boldsymbol{\theta}$ given many copies from the ensemble

$$\mathcal{E}_{\mathbf{A}} = \left\{ |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i} \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}} \right) |+^n\rangle \right\}_{\boldsymbol{\theta}=(\theta_1,\ldots,\theta_m)\in\Theta_q^m}$$

where the matrix $\mathbf{A} \in \mathbb{Z}_2^{m\times n}$ is a random but fixed *architecture* which is known to the learner. This fits exactly into the framework of the pretty good measurement (PGM) [9, 56]. The ensemble $\mathcal{E}$ now turns out to be *geometrically uniform* because it can be written as $\mathcal{E}_{\mathbf{A}} = \{U_{\boldsymbol{\theta}}^{\mathbf{A}} |+^n\rangle\}_{\boldsymbol{\theta}=(\theta_1,\ldots,\theta_m)}$ where $\{U_{\boldsymbol{\theta}}^{\mathbf{A}}\}_{\boldsymbol{\theta}}$ is an Abelian group of matrices. Eldar and Forney [30] showed that the PGM is optimal for all geometrically uniform ensembles, which implies that it is also optimal for our variant of the HPS problem. Nevertheless, despite the optimality, the best known algorithm for implementing pretty good measurements has exponential-time complexity in the size of the ensemble [33]. Consequently, we believe that the HPS problem remains computationally intractable, even if the architecture is public.

References

1 Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pages 325–338, 2018.

2 Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum Pseudoentanglement. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:21, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2024.2.

3 Gorjan Alagic, David Cooper, Quynh Dang, Thinh Dang, John M. Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl A. Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, and Daniel Apon. Status report on the third round of the nist post-quantum cryptography standardization process, 2022-07-05 04:07:00 2022. doi:10.6028/NIST.IR.8413.

4 Michael Alekhnovich. More on average case vs approximation complexity. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, page 298, USA, 2003. IEEE Computer Society.

5 Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In *Theory of Cryptography Conference*, pages 237–265. Springer, 2022.

6 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 208–236, Cham, 2022. Springer Nature Switzerland.

7 Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. Optimal algorithms for learning quantum phase states, 2023. arXiv:2208.07851.

8 Costin Bădescu and Ryan O'Donnell. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1398–1411, 2021.

9 H. Barnum and E. Knill. Reversing quantum dynamics with near-optimal quantum and classical fidelity, 2000. arXiv:quant-ph/0004088.

10 Rishabh Batra and Rahul Jain. Commitments are equivalent to one-way state generators. *arXiv preprint arXiv:2404.03220*, 2024.

11 Amos Beimel, Tal Malkin, and Noam Mazor. Structural lower bounds on black-box constructions of pseudorandom functions. Cryptology ePrint Archive, Paper 2024/1104, 2024. URL: https://eprint.iacr.org/2024/1104.

12 Ward Beullens. Breaking rainbow takes a weekend on a laptop. Cryptology ePrint Archive, Paper 2022/214, 2022. URL: https://eprint.iacr.org/2022/214.

13 Avrim Blum, Merrick Furst, Michael Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology — CRYPTO' 93*, pages 278–291, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.

14 Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626(7997):58–65, February 2024. doi:10.1038/s41586-023-06927-3.

15 John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the uhlmann transformation problem, 2023. arXiv:2306.13073.

16 Adam Bouland, Bill Fefferman, Soumik Ghosh, Tony Metger, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Public-Key Pseudoentanglement and the Hardness of Learning Ground State Entanglement Structure. In Rahul Santhanam, editor, *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:23, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2024.21.

17 Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. *arXiv preprint arXiv:2209.04101*, 2022.

18 Zvika Brakerski and Nir Magrafta. Real-valued somewhat-pseudorandom unitaries. *arXiv preprint arXiv:2403.16704*, 2024.

19 Zvika Brakerski and Omri Shmueli. (pseudo) random quantum states with binary phase, 2019. arXiv:1906.10611.

20 Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.

21 M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2010. doi:10.1098/rspa.2010.0301.

22 Michael J. Bremner, Bin Cheng, and Zhengfeng Ji. IQP Sampling and Verifiable Quantum Advantage: Stabilizer Scheme and Classical Security, August 2023. arXiv:2308.07152.

23 Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical Review Letters*, 117(8), August 2016. doi:10.1103/physrevlett.117.080501.

24 Alex Brodsky and Shlomo Hoory. Simple permutations mix even better. *Random Structures & Algorithms*, 32(3):274–289, 2008.

25 Ran Canetti, Claudio Chamon, Eduardo Mucciolo, and Andrei Ruckenstein. Towards general-purpose program obfuscation via local mixing. *Cryptology ePrint Archive*, 2024.

26 Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. Cryptology ePrint Archive, Paper 2022/975, 2022. URL: https://eprint.iacr.org/2022/975.

27 Chi-Fang Chen, Adam Bouland, Fernando GSL Brandão, Jordan Docter, Patrick Hayden, and Michelle Xu. Efficient unitary designs and pseudorandom unitaries from permutations. *arXiv preprint arXiv:2404.16751*, 2024.

28 Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan. Incompressibility and spectral gaps of random circuits. *arXiv preprint arXiv:2406.07478*, 2024.

29 Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions, 2023. arXiv:2302.12821.

30 Yonina C. Eldar and G. David Forney Jr. On quantum detection and the square-root measurement, 2000. arXiv:quant-ph/0005132.

31 Netta Engelhardt, Åsmund Folkestad, Adam Levine, Evita Verheijden, and Lisa Yang. Cryptographic censorship, 2024. arXiv:2402.03425.

32 Xiaozhou Feng and Matteo Ippoliti. Dynamics of pseudoentanglement, 2024. arXiv:2403.09619.

33 András Gilyén, Seth Lloyd, Iman Marvian, Yihui Quek, and Mark M. Wilde. Quantum algorithm for petz recovery channels and pretty good measurements. *Physical Review Letters*, 128(22), June 2022. doi:10.1103/physrevlett.128.220502.

34 Oded Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, USA, 2006.

35 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, August 1986. doi:10.1145/6490.6503.

36 Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. Cryptology ePrint Archive, Paper 2015/907, 2015. URL: https://eprint.iacr.org/2015/907.

37 W Timothy Gowers. An almost m-wise independent random permutation of the cube. *Combinatorics, Probability and Computing*, 5(2):119–130, 1996.

38 David Gross and Dominik Hangleiter. Secret extraction attacks against obfuscated IQP circuits, December 2023. arXiv:2312.10156.

39 Jeongwan Haah, Yunchao Liu, and Xinyu Tan. Efficient approximate unitary designs from random pauli rotations. *arXiv preprint arXiv:2402.05239*, 2024.

40 Jonas Haferkamp. Random quantum circuits are approximate unitary t-designs in depth $o(nt^{5+o(1)})$. *Quantum*, 6:795, 2022.

41 Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95(3):035001, July 2023. doi:10.1103/RevModPhys.95.035001.

42 William He and Ryan O'Donnell. Pseudorandom permutations from random reversible circuits. *arXiv preprint arXiv:2404.14648*, 2024.

43 Shlomo Hoory, Avner Magen, Steven Myers, and Charles Rackoff. Simple permutations mix well. *Theoretical computer science*, 348(2-3):251–261, 2005.

44 Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.

45 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

46 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018.

47 Gregory D. Kahanamoku-Meyer. Forging quantum data: Classically defeating an IQP-based quantum test. *Quantum*, 7:1107, September 2023. doi:10.22331/q-2023-09-11-1107.

48 Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 968–978, 2024.

49 Isaac H. Kim and John Preskill. Complementarity and the unitarity of the black hole s-matrix. *Journal of High Energy Physics*, 2023(2), February 2023. doi:10.1007/jhep02(2023)233.

50 William Kretschmer. Quantum pseudorandomness and classical complexity. *arXiv preprint arXiv:2103.09320*, 2021.

51 William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602, 2023.

52 Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

53 Fermi Ma and Robert Huang. How to construct random unitaries, 2024. In preparation. Preliminary version available at https://fermima.com/pru.pdf. URL: https://fermima.com/pru.pdf.

54 Ralph C. Merkle. Secure communications over insecure channels. *Commun. ACM*, 21(4):294–299, April 1978. doi:10.1145/359460.359473.

55 Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries, 2024. arXiv:2404.12647.

56 Ashley Montanaro. Pretty simple bounds on quantum state discrimination, 2019. arXiv:1908.08312.

57 Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint arXiv:2210.03394*, 2022.

58 Yoshifumi Nakata, Masato Koashi, and Mio Murao. Generating a statet-design by diagonal quantum circuits. *New Journal of Physics*, 16(5):053043, May 2014. doi:10.1088/1367-2630/16/5/053043.

59 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.

60 Ben W. Reichardt, David Aasen, Rui Chao, Alex Chernoguzov, Wim van Dam, John P. Gaebler, Dan Gresh, Dominic Lucchetti, Michael Mills, Steven A. Moses, Brian Neyenhuis, Adam Paetznick, Andres Paz, Peter E. Siegfried, Marcus P. da Silva, Krysta M. Svore, Zhenghan Wang, and Matt Zanner. Demonstration of quantum computation and error correction with a tesseract code, September 2024. arXiv:2409.04628.

61 R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, February 1978. doi:10.1145/359340.359342.

62 C. Ryan-Anderson, N. C. Brown, C. H. Baldwin, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, N. Hewitt, C. Holliman, C. V. Horst, J. Johansen, D. Lucchetti, T. Mengle, M. Matheny, Y. Matsuoka, K. Mayer, M. Mills, S. A. Moses, B. Neyenhuis, J. Pino, P. Siegfried, R. P. Stutz, J. Walker, and D. Hayes. High-fidelity and Fault-tolerant Teleportation of a Logical Qubit using Transversal Gates and Lattice Surgery on a Trapped-ion Quantum Computer, April 2024. arXiv:2404.16728.

63 Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth, July 2024. arXiv:2407.07754.

64 Dan Shepherd and Michael J. Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 465(2105):1413–1439, May 2009. doi:10.1098/rspa.2008.0443.

65 Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. doi:10.1137/s0097539795293172.

66 P.W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, November 1994. doi:10.1109/SFCS.1994.365700.

67 N. J. A. Sloane and Hieronymus Fischer. Decimal expansion of $\prod_{k>=1}(1-1/2^k)$. OEIS Entry A048651. URL: https://oeis.org/A048651.

68 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7, 2012. doi:10.1561/0400000010.

69 Bei Zeng, Hyeyoun Chung, Andrew W. Cross, and Isaac L. Chuang. Local unitary versus local Clifford equivalence of stabilizer and graph states. *Phys. Rev. A*, 75(3):032325, March 2007. doi:10.1103/PhysRevA.75.032325.

70 Haimeng Zhao, Laura Lewis, Ishaan Kannan, Yihui Quek, Hsin-Yuan Huang, and Matthias C. Caro. Learning quantum states and unitaries of bounded gate complexity, 2023. arXiv:2310.19882.

Hamiltonian Locality Testing via Trotterized Postselection

John Kallaugher
Sandia National Laboratories, Albuquerque, NM, USA

Daniel Liang
Portland State University, OR, USA

Abstract

The (tolerant) Hamiltonian locality testing problem, introduced in [Bluhm, Caro, Oufkir '24], is to determine whether a Hamiltonian H is ε_1-close to being k-local (i.e. can be written as the sum of weight-k Pauli operators) or ε_2-far from any k-local Hamiltonian, given access to its time evolution operator and using as little total evolution time as possible, with distance typically defined by the normalized Frobenius norm. We give the tightest known bounds for this problem, proving an $\mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2-\varepsilon_1)^5}}\right)$ evolution time upper bound and an $\Omega(1/(\varepsilon_2-\varepsilon_1))$ lower bound. Our algorithm does not require reverse time evolution or controlled application of the time evolution operator, although our lower bound applies to algorithms using either tool.

Furthermore, we show that if we *are* allowed reverse time evolution, this lower bound is tight, giving a matching $\mathrm{O}(1/(\varepsilon_2-\varepsilon_1))$ evolution time algorithm.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases quantum algorithms, property testing, hamiltonians

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.10

Related Version *Full Version*: https://arxiv.org/abs/2505.06478

Funding Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research in Quantum Computing, Fundamental Algorithmic Research for Quantum Utility, with support also acknowledged from Fundamental Algorithm Research for Quantum Computing. This work was performed, in part, at the Center for Integrated Nanotechnologies, an Office of Science User Facility operated for the U.S. Department of Energy (DOE) Office of Science. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. DOE's National Nuclear Security Administration under contract DE-NA-0003525. The views expressed in the article do not necessarily represent the views of the U.S. DOE or the United States Government.
Daniel Liang: Supported by US NSF Award CCF-222413.

Acknowledgements We would like to thank Vishnu Iyer and Justin Yirka for getting us started on this problem. We would also like to thank Francisco Escudero Gutiérrez, Srinivasan Arunachalam, and Fang Song for insightful feedback and comments.

1 Introduction

When dealing with large or expensive-to-measure objects, learning the entire object may be too costly. Property testing algorithms instead attempt to distinguish between the object having a given property, or being far from any object with the property. More generally, one can consider *tolerant testing*, where one attempts to distinguish between the object being

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 10; pp. 10:1–10:20
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

within ε_1-close to having a property, or being at least ε_2-far from any object with the property. Such algorithms have been extensively studied in quantum and classical settings (see [18] for an overview of the quantum case), but [6] was the first to consider it for Hamiltonians accessed via their time evolution operator e^{-iHt}. In this setting the natural measure of cost is *total evolution time*, $\sum_j t_j$ where the j^{th} application of the time evolution operator is e^{-iHt_j}.[1]

The property they considered was k-locality, a problem initially raised (but not studied) in [18, Section 7] as well [19]. A Hamiltonian H is k-local if and only if it can be written as $\sum_j H_j$, where each H_j operates on only k qubits. Such locality constraints (perhaps even geometrically locality constraints) are considered to be physically relevant. Local Hamiltonians also appear to be theoretically relevant, as nearly all general learning algorithms for Hamiltonians assume that the Hamiltonian is local, whether they use the time evolution operator [15, 14, 5], or copies of the Gibbs state [2, 4]. Local Hamiltonians are also conducive to efficient simulation on quantum computers, using the technique of Trotterization to break up the Hamiltonian into local quantum gate operations [16]. Finally, local Hamiltonians play an important role in quantum complexity theory, such as QMA-completeness and the Quantum PCP conjecture [1].

The initial version of [6] gave an $\mathrm{O}\big(n^{k+1}/(\varepsilon_2)^3\big)$ evolution time algorithm when distance is measured by the *normalized* (divided by $2^{n/2}$ for a Hamiltonian acting on n qubits) Frobenius norm, improved in [12] to $\mathrm{O}\big((\varepsilon_2-\varepsilon_1)^{-7}\big)$ and then in a later version of [6] to $\mathrm{O}\big((\varepsilon_2-\varepsilon_1)^{-2.5}\varepsilon_2^{-0.5}\big)$.[2][3] This left open the question: how hard is locality testing? Is it possible to achieve linear (a.k.a. Heisenberg) scaling in $1/\varepsilon$ for evolution time, and is such a scaling optimal in all error regimes? In this work we make progress towards resolving the complexity of this problem, improving the best known upper and lower bounds. Our algorithm is based on a technique we refer to as *Trotterized post-selection*, in which we suppress the effect of local terms in the Hamiltonian evolution by repeatedly evolving for a short time period and post-selecting on the non-local part of the time evolution operator.

1.1 Our Results

Our main result is a improved upper bound for the Hamiltonian locality testing problem. As with past works, our algorithm is also time-efficient and non-adaptive, though it does requires n qubits of quantum memory, like [12, 3].

▶ **Theorem 1.** *Let $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$, $\delta \in (0,1)$, and $k \in \mathbb{N}$. There is an algorithm that distinguishes whether an n-qubit Hamiltonian H is (1) within ε_1 of some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, with probability $1-\delta$. The algorithm uses $\mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2-\varepsilon_1)^7}}\log(1/\delta)\right)$ non-adaptive queries to the time evolution operator with $\mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2-\varepsilon_1)^5}}\log(1/\delta)\right)$ total evolution time.*

We pair it with the first lower bound in the tolerant testing setting. While our upper bound uses only forward time evolution and does not require controlled application of e^{-itH}, our lower bound also applies to algorithms using either of these tools.

[1] Another cost measure that can be considered is total query count, the number of individual applications of the time evolution operator. Our algorithm also uses the fewest number of queries of any known algorithm.

[2] The original [6] algorithm only worked in the intolerant setting of $\varepsilon_1 = 0$.

[3] [12] was later subsumed by [3], which gives an $\mathrm{O}\big((\varepsilon_2-\varepsilon_1)^{-3}\big)$ analysis.

▶ **Theorem 2.** *Let $0 \le \varepsilon_1 < \varepsilon_2 \le 1$ and $k \in \mathbb{N}$. Then any algorithm that can distinguish whether an n-qubit Hamiltonian H is (1) within ε_1 of some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, must use $\Omega\left(\frac{1}{\varepsilon_2-\varepsilon_1}\right)$ evolution time in expectation to achieve constant success probability.*

▶ Remark 3. [6, Theorem 3.6] gives a hardness result for the *unnormalized* Frobenius norm (as well as other Schatten norms) in the *non-tolerant* setting that scales as $\Omega\left(\frac{2^{n/2}}{\varepsilon}\right)$. Once normalized, this also gives a $\Omega\left(\frac{1}{\varepsilon}\right)$ lower bound. However, this hardness result only holds for exponentially small ε, due to the fact that the "hard" Hamiltonian in [6, Lemma 3.2] no longer has $\|H\|_\infty \le 1$ when the *unnormalized* Frobenius distance to k-local is super-constant. Therefore Theorem 2 is, to the authors' knowledge, the first lower bound that works for arbitrary values of ε, in addition to being the first for the tolerant setting. Our proof is also considerably simpler, and still extends to all of the distance measures considered in [6] and more.

Finally, we show that, when reverse time evolution and controlled operations are allowed, it is possible to saturate this lower bound even in the tolerant case (proof in the appendix).

▶ **Theorem 4.** *Let $0 \le \varepsilon_1 < \varepsilon_2 \le 1$, $\delta \in (0,1)$, and $k \in \mathbb{N}$. There is an algorithm that tests whether an n-qubit Hamiltonian H is (1) ε_1-close to some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, with probability $1-\delta$. The algorithm uses $\mathrm{O}\left(\frac{\log(1/\delta)}{(\varepsilon_2-\varepsilon_1)^2}\right)$ non-adaptive queries to the time evolution operator and its inverse, with $\mathrm{O}\left(\frac{\log(1/\delta)}{\varepsilon_2-\varepsilon_1}\right)$ total evolution time.*

2 Proof Overview

2.1 Upper Bound

For simplicity, we will consider the intolerant case ($\varepsilon_1 = 0$, $\varepsilon_2 = \varepsilon$) for this proof overview; the same techniques apply in the tolerant case but require somewhat more care. First we start with the intuition behind the algorithm of [12, 3].

We will need the fact that the space of $2n$ qubit states $\mathbb{C}^{2^{2n}}$ has the Bell basis $(|\sigma_P\rangle)_P$, where P spans the n-fold Paulis, $|\sigma_{I^{\otimes n}}\rangle$ is the maximally entangled state $\frac{1}{\sqrt{2^n}}\sum_{x\in\{0,1\}^n}|x\rangle \otimes |x\rangle$, and $|\sigma_P\rangle = (I^{\otimes n} \otimes P)|\sigma_{I^{\otimes n}}\rangle$. Therefore, for any unitary U, if we apply $I^{\otimes n} \otimes U$ to $|\sigma_{I^{\otimes n}}\rangle$ and then measure in the Bell basis, we are able to sample from the (squared) Pauli spectrum[4] of U (the squares of the Pauli decomposition coefficients always sum to 1 for a unitary [17]).

For any Hamiltonian H, the closest k-local Hamiltonian is given by dropping all of the non-local Paulis from its Pauli decomposition. Therefore, as by the first-order Taylor series expansion,

$$e^{-iHt} \approx I^{\otimes n} - iHt$$

for small enough t, we can set $U = e^{-iH\cdot t}$ in the aforementioned procedure, and if H is ε-far from local we will sample a non-local Pauli term with $\approx (t\cdot\varepsilon)^2$ probability. Conversely, if H is local we should sample no non-local terms, giving us a distinguishing algorithm if the process is repeated $\mathrm{O}\big((t\cdot\varepsilon)^{-2}\big)$ times, for a total time evolution of $\mathrm{O}\big(t^{-1}\cdot\varepsilon^{-2}\big)$.

[4] That is, α_P^2 when U is written as $\sum_P \alpha_P P$.

So ideally we would like t to be $\Theta(1/\varepsilon)$ and only repeat a constant number of times, leading to a total time evolution of $\mathrm{O}\big(\varepsilon^{-1}\big)$, which would be optimal by Theorem 2.

Unfortunately, these higher-order terms in the Taylor series cannot be ignored at larger values of t. As we have $\|H\|_\infty \leq 1$, we can bound the k^{th} order term of the Taylor series expansion of H by $\mathrm{O}\big(t^k\big)$, and so we must set t to be at most $\Theta(\varepsilon)$, resulting in the total time evolution of $\mathrm{O}\big(\varepsilon^3\big)$ obtained in previous work [12, 3].

To evade this barrier, we will instead show that it is possible to (approximately) simulate evolving by $H_{>k}$, which is composed of only the *non-local* terms of the Pauli decomposition of H. Note that if H is k-local, this is 0, while if it is not, $H_>k$ is the difference between H and the closest k-local Hamiltonian. Suppose we could evolve by the time evolution operator of this Hamiltonian. Then performing the Bell sampling procedure from before would return $|\sigma_{I^{\otimes n}}\rangle$ with probability

$$\begin{aligned}
&\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes e^{-iH_{>k}t}\right)|\sigma_{I^{\otimes n}}\rangle\right|^2\\
&=\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes\left(\sum_{\ell=0}^{\infty}(H_{>k})^{\ell}\frac{(it)^{\ell}}{\ell!}\right)\right)|\sigma_{I^{\otimes n}}\rangle\right|^2\\
&=\left|1+\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes\left(\sum_{\ell=2}^{\infty}(H_{>k})^{\ell}\frac{(it)^{\ell}}{\ell!}\right)\right)|\sigma_{I^{\otimes n}}\rangle\right|^2\\
&=1-\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes(H_{>k})^2\right)|\sigma_{I^{\otimes n}}\rangle+\sum_{\ell=3}^{\infty}\mathrm{O}\left(t^{\ell}\cdot\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes(H_{>k})^{\ell}\right)|\sigma_{I^{\otimes n}}\rangle\right|\right)
\end{aligned}$$

as H contains no identity term.

To tame this infinite series, imagine that $\|H_{>k}\|_\infty \leq 1$ (we will eventually evolve by a related operator A that *does* satisfy $\|A\|_\infty \leq 1$). Then we have

$$\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes(H_{>k})^{\ell}\right)|\sigma_{I^{\otimes n}}\rangle\right|\leq\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes(H_{>k})^2\right)|\sigma_{I^{\otimes n}}\rangle$$

for all integers $\ell \geq 2$, so as long as t is a sufficiently small constant, we have that $\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes e^{-iH_{>k}t}\right)|\sigma_{I^{\otimes n}}\rangle\right|^2$ is at least

$$1-0.99\cdot\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n}\otimes(H_{>k})^2\right)|\sigma_{I^{\otimes n}}\rangle=1-0.99\cdot\mathrm{Tr}\left((H_{>k})^2\right)/2^n,$$

where $\mathrm{Tr}\left((H_{>k})^2\right)/2^n=\varepsilon^2$ is exactly the squared normalized Frobenius distance of H from being k-local. So if we apply $e^{-iH_{>k}t}$ with $t=\Theta(1)$, we are left with a $\approx\varepsilon^2$ probability of sampling a non-local Pauli term if H is non-local, and are guaranteed to measure identity if H is local (as then $e^{-iH_{>k}\cdot t}$ is the identity). This means we can distinguish locality from non-locality with $\mathrm{O}\big(\varepsilon^{-2}\big)$ repetitions, requiring $\mathrm{O}\big(\varepsilon^{-2}\big)$ total evolution time.[5]

Now, we cannot actually apply $e^{-iH_{>k}t}$. However, when starting at $|\sigma_{I^{\otimes n}}\rangle$, we can approximate it up to $t=\Theta(1)$ by the use of a process reminiscent of the Elitzur-Vaidman bomb-tester [9] and Quantum Zeno effect [10], which we refer to as *Trotterized post-selection.*

Let D be the subspace of Bell states corresponding to non-local Paulis *or* identity and let Π_D be the projector onto that subspace. Starting with $|\sigma_{I^{\otimes n}}\rangle$ once again, we apply $I^{\otimes n}\otimes e^{-iHt'}$ for $t'=\mathrm{O}(\varepsilon)$, measure with $\{\Pi_D, I^{\otimes 2n}-\Pi_D\}$, and then post-select on the measurement result Π_D. We then repeat our application of $I^{\otimes n}\otimes e^{-iHt'}$ and post-selection, for $\mathrm{O}(1/t')$ iterations, provided our post-selection succeeds each time.

[5] Unfortunately, even with access to the time evolution operator of $H_{>k}$ we cannot set t to the optimal $\Theta(1/\varepsilon)$, as we lose control of the higher-order terms of the Taylor expansion.

As we start with $|\sigma_{I^{\otimes n}}\rangle$, then make small adjustments (i.e., $e^{-iHt} \approx I^{\otimes 2n}$ for small t), the chance of failing the post-selection is small: only $\mathrm{O}(\varepsilon^2)$ at each iteration, and so as long as we only use $\mathrm{O}(1/\varepsilon)$ iterations, we will succeed with probability $1 - \mathrm{O}(\varepsilon)$. Now, as we are taking small steps, we can approximate each iteration of $\Pi_D\left(I^{\otimes n} \otimes e^{-iH\cdot\mathrm{O}(\varepsilon)}\right)\Pi_D$ as

$$\Pi_D\left(I^{\otimes n} \otimes e^{-iH\cdot\mathrm{O}(\varepsilon)}\right)\Pi_D = \Pi_D\left(I^{\otimes n} \otimes \sum_{\ell=0}^{\infty} H^\ell \frac{(-i)^\ell\,\mathrm{O}(\varepsilon^\ell)}{\ell!}\right)\Pi_D = e^{-iA\cdot\mathrm{O}(\varepsilon)} + R$$

where we define $A \coloneqq \Pi_D(I^{\otimes n} \otimes H)\Pi_D$ and choose some $\|R\|_\infty \le \mathrm{O}(\varepsilon^2)$.[6]

Now, in general, $A \ne I^{\otimes n} \otimes H_{>k}$, but as long as H has no identity term in its Pauli decomposition[7], by construction $A|\sigma_{I^{\otimes n}}\rangle = (I^{\otimes n} \otimes H_{>k})\,|\sigma_{I^{\otimes n}}\rangle$, and so $\langle\sigma_{I^{\otimes n}}|A^2|\sigma_{I^{\otimes n}}\rangle = \langle\sigma_{I^{\otimes n}}|I \otimes (H_{>k})^2\,|\sigma_{I^{\otimes n}}\rangle$. Combined with the fact that $\|A\|_\infty = \|\Pi_D\,(I^{\otimes n} \otimes H)\,\Pi_D\|_\infty \le \|H\|_\infty \le 1$, we can argue that, if we iterate t/t' times

$$\begin{aligned}
\langle\sigma_{I^{\otimes n}}|\prod_{i=1}^{t/t'} e^{-iA\cdot t'}|\sigma_{I^{\otimes n}}\rangle &= \langle\sigma_{I^{\otimes n}}|e^{-iA\cdot t}|\sigma_{I^{\otimes n}}\rangle \\
&= \langle\sigma_{I^{\otimes n}}|\left(\sum_{\ell=0}^{\infty} A^\ell \frac{(-it)^\ell}{\ell!}\right)|\sigma_{I^{\otimes n}}\rangle \\
&= 1 - t^2\langle\sigma_{I^{\otimes n}}|H_{>k}^2|\sigma_{I^{\otimes n}}\rangle + \mathrm{O}(t^3\cdot\varepsilon^2)
\end{aligned}$$

where the final inequality follows from the fact that for all $k > 2$,

$$\left|\langle\sigma_{I^{\otimes n}}|A^k|\sigma_{I^{\otimes n}}\rangle\right| \le \|A\|_\infty^{k-2}\langle\sigma_{I^{\otimes n}}|A^2|\sigma_{I^{\otimes n}}\rangle \le \langle\sigma_{I^{\otimes n}}|\left(I^{\otimes n} \otimes (H_{>k})^2\right)|\sigma_{I^{\otimes n}}\rangle = \varepsilon^2.$$

So as our method based on access to the time evolution operator of $H_{>k}$ only required distinguishing between $\langle\sigma_{I^{\otimes n}}|H_{>k}|\sigma_{I^{\otimes n}}\rangle$ being $\Theta(\varepsilon^2)$ and 0 we can emulate it with access to e^{-iAt} without losing too much accuracy, as long as we take t to be a small enough constant. We can therefore test locality with a total time evolution of $\mathrm{O}(\varepsilon^{-2})$.

2.2 Lower Bound

To prove the lower bound, it suffices to show that for any k there exists Hamiltonians H_1 and H_2 such that a query to the time t evolution of H_1 and H_2 differ in diamond distance by at most $\mathrm{O}((\varepsilon_2 - \varepsilon_1)t)$, with H_1 ε_1-close to being k-local and H_2 ε_2-far from being k-local.

We achieve this by considering the weight-k Pauli $Z_{1:k}$ that is Z on the first k qubits, and identity on the last $n-k$ qubits. We then set $H_1 \coloneqq \varepsilon_1 Z_{1:k}$ and $H_2 \coloneqq \varepsilon_2 Z_{1:k}$. Because $Z_{1:k}$ is diagonal, so is $e^{-i\varepsilon Z_{1:k}\cdot t}$, making it straightforward to bound the diamond distance of the two time evolution operators by $\mathrm{O}(t(\varepsilon_2 - \varepsilon_1))$. By the sub-additivity of diamond distance, the total time evolution required to distinguish the two Hamiltonians with constant probability is therefore at least $\Omega\left((\varepsilon_2 - \varepsilon_1)^{-1}\right)$.

[6] Note that the Π_D on the right does nothing besides make A obviously Hermitian, assuming our invariant of our post-selection succeeding.

[7] We can assume this without loss of generality, as our algorithm never uses controlled application of $e^{-iH\cdot t}$, and so any identity term would manifest as an undetectable global phase.

3 Preliminaries

3.1 Quantum Information

A Hamiltonian on n-qubits is a $2^n \times 2^n$ Hermitian matrix. The time evolution operator of a Hamiltonian H for time $t \geq 0$ is the unitary matrix

$$e^{-iHt} := \sum_{k=0}^{\infty} H^k (-i)^k \frac{t^k}{k!}.$$

We define the n-qubit Pauli matrices to be $\mathcal{P}^{\otimes n} := \{I, X, Y, Z\}^{\otimes n}$, where $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. For any Pauli P, we denote the locality $|P|$ to be the number of non-identity terms in the tensor product. Let the Frobenius inner product between matrices A and B be $\langle A, B \rangle := \operatorname{Tr}(A^\dagger B)$. The orthogonality of Pauli matrices under the Frobenius inner product is implied by the fact that any product of Paulis is another Pauli (up to sign) and the fact that among them only the identity has non-zero trace. Given a matrix $A = \sum_{P \in \mathcal{P}^{\otimes n}} \alpha_P P$, the locality of A is the largest $|P|$ such that $\alpha_P \neq 0$. If A is a Hamiltonian (i.e., Hermitian) then all α_P are real-valued. The *normalized* Frobenius norm is given by

$$\|A\|_2 = \sqrt{\frac{\langle A, A \rangle}{2^n}} = \sqrt{\frac{\operatorname{Tr}(A^\dagger A)}{2^n}} = \sqrt{\sum_{P \in \mathcal{P}^{\otimes n}} |\alpha_P|^2},$$

and will be used as our distance to k-locality, in keeping with the previous literature [6, 12, 3]. The other important norm will be the (unnormalized) spectral norm $\|A\|_\infty$, which is the largest singular value of A. For any matrix A, $\|A\|_2 \leq \|A\|_\infty$, recalling that $\|\cdot\|_2$ is the *normalized* Frobenius norm. As a form of normalization and to be consistent with the literature, we will assume that $\|H\|_\infty \leq 1$ for any Hamiltonian referenced. We will also WLOG assume that $\operatorname{Tr}(H) = 0$ for any Hamiltonian, since it does not affect the time evolution unitary beyond a global phase, and so as our algorithms do not use controlled application of the unitary, they cannot be affected by it.

We define $A_{>k} := \sum_{|P|>k} \alpha_P P$ and subsequently $A_{\leq k} := \sum_{|P|\leq k} \alpha_P P$. By the orthogonality of the Pauli matrices under the Frobenius inner product, $A_{\leq k}$ is the k-local Hamiltonian that is closest to A with distance $\|A - A_{\leq k}\|_2 = \|A_{>k}\|_2$.

Let $B = \{|\Phi^+\rangle, |\Phi^-\rangle, |\Psi^+\rangle, |\Psi^-\rangle\}$ denote the set containing the four Bell states. We will view $B^{\otimes n}$ as a basis of $\mathbb{C}^{2^n} \otimes \mathbb{C}^{2^n}$, in which for each copy of B, one qubit is assigned to the left register and one to the right. Note that, up to phase, every state in $B^{\otimes n}$ is equal to $(I^{\otimes n} \otimes P)|\Phi^+\rangle^{\otimes n}$ for a unique $P \in \mathcal{P}^{\otimes n}$. We will write $|\sigma_P\rangle$ for this basis element. As an example,

$$|\Phi^+\rangle^{\otimes n} = |\sigma_{I^{\otimes n}}\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle \otimes |x\rangle.$$

If $U = \sum_{P \in \mathcal{P}^{\otimes n}} \alpha_P P$ is a unitary matrix, then by Parseval's identity, $\sum_{P \in \mathcal{P}^{\otimes n}} |\alpha_P|^2 = 1$, i.e. $|\alpha_P|^2$ gives a probability distribution over the Paulis. Applying $I^{\otimes n} \otimes U$ to the state $|\sigma_{I^{\otimes n}}\rangle = |\Phi^+\rangle^{\otimes n}$ and measuring in the Bell basis $B^{\otimes n}$ allows one to sample from this distribution [17].

For a quantum channel that takes as input an n-qubit state, we will let the diamond norm refer to $\|\Lambda\|_\diamond := \max_\rho \|(I^{\otimes n} \otimes \Lambda)(\rho)\|_1$ where the maximization is over all $2n$-qubit states ρ. The diamond distance famously characterizes the maximum statistical distinguishability (i.e., induced trace distance) between quantum channels [21, Section 9.1.6], even with ancillas.

3.2 Probability

▶ **Fact 5** (Multiplicative Chernoff Bound). *Suppose $X_1, \dots, X_m$ are independent Bernoulli random variables. Let X denote their sum and let $\mu \coloneqq \mathbb{E}[X]$. Then for any $t > 0$*

$$\Pr[X \leq (1-t)\mu] \leq e^{-t^2\mu/2}.$$

We will not need a particularly tight form of this bound, so for ease of analysis we state the following (loose) corollary.

▶ **Corollary 6.** *Suppose $X_1, \dots, X_m$ are i.i.d. Bernoulli random variables with probability p, and*

$$m = \frac{2}{p}\left(d + \log(1/\delta)\right).$$

Then

$$\Pr\left[\sum_{i=1}^{m} X_i < d\right] \leq \delta.$$

Proof. Let $\mu \coloneqq \mathbb{E}[\sum_{i=1}^m X_i] = mp$ and let $\gamma \coloneqq 1 - \frac{d}{\mu}$. By the Multiplicative Chernoff Bound,

$$\begin{aligned}\Pr\left[\sum_{i=1}^{m} X_i < d\right] &= \Pr\left[\sum_{i=1}^{m} X_i < (1-\gamma)\mu\right] \\ &\leq \exp\left(-\frac{\mu}{2}\gamma^2\right) = \exp\left(-\frac{\mu}{2} - \frac{d^2}{2\mu} + d\right) \leq \exp\left(-\frac{mp}{2} + d\right).\end{aligned}$$

Hence, as long as

$$m \geq \frac{2\log(1/\delta) + 2d}{p},$$

then $\sum_{i=1}^m X_i \leq d$ with probability at most δ. ◀

▶ **Fact 7** (Bernstein's inequality). *Suppose $X_1, \dots, X_n$ are independent Bernoulli random variables. Let X denote their sum and let μ and σ^2 be the expectation and variance of X respectively. Then for $t \in (0, n)$*

$$\Pr[X - \mu \geq t] \leq e^{-\frac{\frac{t^2}{2}}{\sigma^2 + \frac{t}{3}}} \text{ and } \Pr[X - \mu \leq -t] \leq e^{-\frac{\frac{t^2}{2}}{\sigma^2 + \frac{t}{3}}}.$$

4 Upper Bound

We will frequently use the truncation of the Taylor series of the matrix exponential to analyze our algorithm. The following will allow us to then bound the error of the truncation.

▶ **Fact 8** ([8, Lemma F.2]). *If $\lambda \in \mathbb{C}$ then $\left|\sum_{k=\ell}^{\infty} \frac{\lambda^k}{k!}\right| \leq \frac{|\lambda|^\ell}{\ell!} e^{|\lambda|}$.*

▶ **Corollary 9.** *For n-qubit Hamiltonian H with $\|H\|_\infty \leq 1$, the first order Taylor series expansion of the matrix exponential gives*

$$e^{-iHt} = I^{\otimes n} - iHt + \frac{e^t \cdot t^2}{2} R$$

for $\|R\|_\infty \leq 1$.

Proof. By the triangle inequality and the fact that $\|H^k\|_\infty \le \|H\|_\infty \le 1$ for $k \ge 1$:

$$\|e^{-iHt} - (I^{\otimes n} - iHt)\|_\infty = \left\| \sum_{k=2}^{\infty} (-i)^k \frac{H^k t^k}{k!} \right\|_\infty \le \sum_{k=2}^{\infty} \frac{\|H^k\|_\infty t^k}{k!} \le \sum_{k=2}^{\infty} \frac{t^k}{k!} \le \frac{e^t \cdot t^2}{2},$$

using Fact 8 at the end. Setting $R := \frac{2}{e^t \cdot t^2}\left(e^{-iHt} - (I^{\otimes n} - iHt)\right)$ completes the proof. ◀

We also prove the related fact to bound the real and imaginary terms.

▶ **Fact 10.** *If $\lambda \in \mathbb{C}$ then*

$$\left| \sum_{k=\ell}^{\infty} \frac{\lambda^{2k}}{(2k)!} \right| \le \frac{|\lambda|^{2\ell}}{(2\ell)!} \cosh(|\lambda|)$$

and

$$\left| \sum_{k=\ell}^{\infty} \frac{\lambda^{2k+1}}{(2k+1)!} \right| \le \frac{|\lambda|^{2\ell+1}}{(2\ell+1)!} \cosh(|\lambda|).$$

Proof.

$$\left| \sum_{k=\ell}^{\infty} \frac{\lambda^{2k}}{(2k)!} \right| \le \sum_{k=\ell}^{\infty} \frac{|\lambda^{2k}|}{(2k)!} = |\lambda|^{2\ell} \sum_{k=0}^{\infty} \frac{|\lambda|^{2k}}{(2k+2\ell)!} \le \frac{|\lambda|^{2\ell}}{(2\ell)!} \sum_{k=0}^{\infty} \frac{|\lambda|^{2k}}{(2k)!} = \frac{|\lambda|^{2\ell}}{(2\ell)!} \cosh(|\lambda|)$$

and

$$\begin{aligned} \left| \sum_{k=\ell}^{\infty} \frac{\lambda^{2k+1}}{(2k+1)!} \right| &\le \sum_{k=\ell}^{\infty} \frac{|\lambda^{2k+1}|}{(2k+1)!} = |\lambda|^{2\ell+1} \sum_{k=0}^{\infty} \frac{|\lambda|^{2k}}{(2k+2\ell+1)!} \\ &\le \frac{|\lambda|^{2\ell+1}}{(2\ell+1)!} \sum_{k=0}^{\infty} \frac{|\lambda|^{2k}}{(2k)!} = \frac{|\lambda|^{2\ell+1}}{(2\ell+1)!} \cosh(|\lambda|). \end{aligned}$$

◀

4.1 Algorithm

▶ **Definition 11.** *We will use D to denote the subspace of $\mathbb{C}^{2^n} \otimes \mathbb{C}^{2^n}$ spanned by $|\sigma_P\rangle$ for Pauli strings P that are either the identity or are* not *k-local, and Π_D to denote the projector onto D. We define $A := \Pi_D\,(I^{\otimes n} \otimes H)\,\Pi_D$.*

We start by giving an algorithm that returns a Bernoulli random variable $X \in \{0, 1\}$, where $\mathbb{E}[X]$ approximates the distance of H from being k-local. It does so by iteratively applying $e^{-i\alpha H}$ sandwiched by $\{\Pi_D, I^{\otimes 2n} - \Pi_D\}$ measurements.

Algorithm 1 Hamiltonian Locality Estimator via Trotterized Postselection.

```
Start with $|\phi\rangle = |\sigma_{I^{\otimes n}}\rangle$.
for $\frac{50}{\sqrt{\varepsilon_2^2 - \varepsilon_1^2}}$ iterations do
    Apply $(I^{\otimes n} \otimes e^{-i\alpha H}$ to $|\phi\rangle$ for $\alpha = \frac{\varepsilon_2^2 - \varepsilon_1^2}{100\varepsilon_2}$.
    Measure $|\phi\rangle$ with the projectors $\Pi_D, I^{\otimes 2n} - \Pi_D$, terminating and returning $\perp$ if the
    result is $I^{\otimes 2n} - \Pi_D$.
end for
Measure $|\phi\rangle$ in the Bell basis, returning 0 if the result is $|\sigma_{I^{\otimes n}}\rangle$ and 1 otherwise.
```

Let $\alpha := \frac{\varepsilon_2^2-\varepsilon_1^2}{100\varepsilon_2}$ be the step-size used in Algorithm 1, $t := \frac{\sqrt{\varepsilon_2^2-\varepsilon_1^2}}{2\varepsilon_2}$ be the total time evolution used in Algorithm 1, and let $m := t/\alpha = \frac{50}{\sqrt{\varepsilon_2^2-\varepsilon_1^2}}$ be the number of iterations used in Algorithm 1. In our analysis will frequently use the fact that $\alpha \leq \frac{\varepsilon_2}{100} \leq \frac{1}{100}$ and $t \leq 0.5$ to simplify higher-order terms.

▶ Remark 12. While we attempted to keep the constants in the algorithm reasonable, no attempt was made to optimize them. We observe that t should remain $\Theta\left(\frac{\sqrt{\varepsilon_2^2-\varepsilon_1^2}}{\varepsilon_2}\right)$ for optimal scaling, but α can be made arbitrarily small to (marginally) improve the constants in the total time evolution used. This has a cost in the total number of queries used, scaling roughly proportional to α^{-1}.

First we show that the final state of the Trotterized postselection algorithm corresponds to evolving $|\sigma_{I^{\otimes n}}\rangle$ by e^{-iAt}, with a bounded error term. There are two main sources of error: (1) the error from higher-order terms in the respective Taylor series of $e^{-iA\alpha}$ and $\Pi_D\left(I^{\otimes n}\otimes e^{-iH\alpha}\right)\Pi_D$ not matching and (2) the error from post-selection causing normalization issues. The following technical lemma allows us to tackle the error from (1). This is done by showing that $e^{-itA} = \Pi_D\left(I^{\otimes n}\otimes e^{-itH}\right)\Pi_D \pm \mathrm{O}\left(\alpha^2\right)$ for sufficiently small α. By chaining these together, the triangle inequality will eventually show in Lemma 14 that the accumulated error is then at most $\mathrm{O}\left(\alpha^2 m\right) = \mathrm{O}(\alpha t)$.

▶ **Lemma 13.** *Let $H = \sum_{P\in\mathcal{P}^{\otimes n}} \alpha_P P$ be any Hamiltonian with $\|H\|_\infty \leq 1$. Then,*

$$\Pi_D(I^{\otimes n}\otimes e^{-i\alpha H})\Pi_D = e^{-i\alpha A} + \eta$$

where $\|\eta\|_\infty \leq e^\alpha \cdot \alpha^2$.

Proof. By Taylor expanding the complex exponential of $e^{-i\alpha H}$ and applying Corollary 9, we get

$$\begin{aligned}\Pi_D(I^{\otimes n}\otimes e^{-i\alpha H})\Pi_D &= \Pi_D\left(I^{\otimes n}\otimes\left(I^{\otimes n} - i\alpha H + \frac{e^\alpha\cdot\alpha^2}{2}R\right)\right)\Pi_D \\ &= I^{\otimes 2n} - i\alpha A + \frac{e^t\cdot\alpha^2}{2}R'\end{aligned}$$

where $\|R'\|_\infty \leq \|I^{\otimes n}\otimes R\|_\infty = \|R\|_\infty \leq 1$.

Next, we observe that $\|A\|_\infty \leq \|I^{\otimes n}\otimes H\|_\infty = \|H\|_\infty \leq 1$ and that A is Hermitian by symmetry. We can then Taylor expand $e^{-i\alpha A}$ to get

$$e^{-i\alpha A} = I^{\otimes 2n} - i\alpha A + \frac{e^\alpha\cdot\alpha^2}{2}Q$$

where $\|Q\|_\infty \leq 1$. By the triangle inequality, the difference

$$\eta := \Pi_D(I^{\otimes n}\otimes e^{-i\alpha H})\Pi_D - e^{-i\alpha A}$$

between these two linear transformations satisfies

$$\|\eta\|_\infty \leq \|R'\|_\infty\cdot\frac{e^\alpha\cdot\alpha^2}{2} + \|Q\|_\infty\cdot\frac{e^\alpha\cdot\alpha^2}{2} \leq e^\alpha\cdot\alpha^2. \qquad ◀$$

Luckily, the error from (2) is mostly a non-issue, using a process similar to the Elitzur-Vaidman bomb [9]: by taking small steps between applications of Π_D, we ensure that we are barely changing our system, and so the post-selection nearly always succeeds. This also means that the normalization error can be suppressed to be arbitrarily small, at the cost of linearly increasing the number of times we have to query the time evolution operator. Using these facts together, we show that Algorithm 1 approximately applies the time evolution operator of A.

▶ **Lemma 14.** *Algorithm 1 terminates before the final measurement with probability at most $\frac{99}{98}\alpha t$. If it does not, $|\phi\rangle = e^{-iAt}|\sigma_{I^{\otimes n}}\rangle + |\Delta\rangle$ just before the final measurement, with $\||\Delta\rangle\|_2 \leq \frac{7}{4}\alpha t$.*

Proof. Note that the algorithm can only be terminated early if, in one of the loop iterations, the measurement in Algorithm 1 returns $I^{\otimes 2n} - \Pi_D$. At the start of the iteration $|\phi\rangle = |\sigma_{I^{\otimes n}}\rangle \in D$. Since $|\phi\rangle$ remains within D after each successful iteration, by Taylor expanding the exponential, and applying Corollary 9 to obtain a suitable R with $\|R\|_\infty \leq 1$, the probability of failure at each iteration is at most

$$\begin{aligned}
&\left\|(I^{\otimes 2n} - \Pi_D)\left(I^{\otimes n} \otimes e^{-iH\alpha}\right)\Pi_D|\phi\rangle\right\|_2^2 \\
&= \left\|(I^{\otimes 2n} - \Pi_D)\left(I^{\otimes n} \otimes \left(I^{\otimes n} - i\alpha H + \frac{\alpha^2}{2}e^{\alpha}R\right)\right)|\phi\rangle\right\|_2^2 \\
&= \left\|(I^{\otimes 2n} - \Pi_D)\left(-i\alpha(I^{\otimes n} \otimes H) + \frac{\alpha^2}{2}e^{\alpha}(I^{\otimes n} \otimes R)\right)|\phi\rangle\right\|_2^2 \\
&\leq \left(\alpha\|H\|_\infty + \frac{\alpha^2 e^{\alpha}}{2}\|R\|_\infty\right)^2 \\
&\leq \left(1 + \alpha e^{\alpha} + \frac{\alpha^2}{4}e^{2\alpha}\right)\alpha^2 \\
&< \frac{99}{98}\alpha^2
\end{aligned}$$

where the third line follows from $|\phi\rangle \in D$, the fourth from the triangle inequality combined with the definition of the spectral norm, and the final line from $\alpha \leq 0.01$. By a union bound over the m iterations, the first part of the lemma follows, noting that $t \coloneqq \alpha \cdot m$.

For the second part pertaining to accuracy, first we note that in each iteration, if the measurement in Algorithm 1 does *not* make the algorithm terminate, the iteration had the effect of taking $|\phi\rangle \in D$ to

$$\Pi_D\left(I^{\otimes n} \otimes e^{-i\alpha H}\right)|\phi\rangle = \Pi_D\left(I^{\otimes n} \otimes e^{-i\alpha H}\right)\Pi_D|\phi\rangle,$$

normalized to length 1. After the m iterations of the loop of Algorithm 1, $|\phi\rangle$ is then

$$\prod_{i=1}^{m}\Pi_D\left(I^{\otimes n} \otimes e^{-i\alpha H}\right)\Pi_D|\sigma_{I^{\otimes n}}\rangle$$

normalized to length 1. By Lemma 13, before normalization this is equivalent to

$$\prod_{i=1}^{m}\left(e^{-i\alpha A} + \eta\right)|\sigma_{I^{\otimes n}}\rangle = \left(\sum_{k=0}^{m}\binom{m}{k}e^{-i\alpha A(m-k)} \cdot \eta^k\right)|\sigma_{I^{\otimes n}}\rangle$$

for $\|\eta\|_\infty \leq \alpha^2 e^{\alpha}$. The distance of the un-normalized vector from $e^{-iAt}|\sigma_{I^{\otimes n}}\rangle$ is then

$$\begin{aligned}
\left\|e^{-iAt}|\sigma_{I^{\otimes n}}\rangle - \prod_{i=1}^{m}\left(e^{-iAt} + \eta\right)|\sigma_{I^{\otimes n}}\rangle\right\|_2 &= \left\|\left(\sum_{k=1}^{m}\binom{m}{k}e^{-i\alpha A(m-k)} \cdot \eta^k\right)|\sigma_{I^{\otimes n}}\rangle\right\|_2 \\
\leq \sum_{k=1}^{m} m^k\|\eta\|_\infty^k \leq \sum_{k=1}^{m}\left(m\alpha^2 e^{\alpha}\right)^k \leq \sum_{k=1}^{\infty}\left(m\alpha^2 e^{\alpha}\right)^k &= m\alpha^2 e^{\alpha}\frac{1}{1 - m\alpha^2 e^{\alpha}} = \alpha t e^{\alpha}\frac{1}{1 - \alpha t e^{\alpha}}.
\end{aligned}$$

Finally, to bound the error introduced by normalization, for each $r \in [m]$, write $|\phi_r\rangle := \prod_{i=1}^{r} \Pi_D(I^{\otimes n} \otimes e^{-i\alpha H})\Pi_D|\sigma_{I^{\otimes n}}\rangle$ for the projected state at iteration r. We note that, by the same argument proving that the probability of the measurement at any given step returning the $I^{\otimes 2n} - \Pi_D$ result is at most $\frac{99}{98}\alpha^2$, $|\phi_r\rangle$ is separated from $e^{-iAt}|\phi_{r-1}\rangle$ by an *orthogonal* vector of length at most $\sqrt{\frac{99}{98}}\alpha\|e^{-iAt}|\phi_{r-1}\rangle\|_2 = \sqrt{\frac{99}{98}}\alpha\||\phi_{r-1}\rangle\|_2$. Therefore,

$$\||\phi_r\rangle\|_2 \geq \||\phi_{r-1}\rangle\|_2\sqrt{1-\frac{99}{98}\alpha^2} \geq \||\phi_{r-1}\rangle\|_2 - 0.6\frac{99}{98}\alpha^2$$

where the last inequality follows from the fact that $1-\sqrt{1-x} \leq 0.6x$ for $x \in [0, \frac{5}{9}]$ and $\frac{99}{98}\alpha^2 < \frac{5}{9}$. The total additional error from the normalization is then at most $\frac{297}{490}\alpha^2 m = \frac{297}{490}\alpha t$. By the triangle inequality, the total distance from $e^{-iAt}|\sigma_{I^{\otimes n}}\rangle$ is at most

$$\frac{297}{490}\alpha t + t\alpha e^{\alpha}\frac{1}{1-\alpha t e^{\alpha}} \leq \frac{7}{4}\alpha t. \qquad ◀$$

We now show that (approximately) applying e^{-iAt} instead of $I^{\otimes n} \otimes e^{-iHt}$ allows us to suppress the higher-order terms that were preventing us from increasing the evolution time t when testing for locality. We will need the following results that let us characterize the individual terms of the Taylor expansion.

▶ **Fact 15.** *For any matrix M, $\langle\sigma_P|(I \otimes M)|\sigma_Q\rangle = \frac{\mathrm{Tr}(PMQ)}{2^n}$.*

Proof.

$$\langle\sigma_P|(I\otimes M)|\sigma_Q\rangle = \frac{1}{2^n}\sum_{x,y\in\{0,1\}^n}\left(\left(\langle x|\otimes\langle x|P\right)\left(|y\rangle\otimes MQ|y\rangle\right)\right)$$

$$= \frac{1}{2^n}\sum_{x,y\in\{0,1\}^n}\langle x|y\rangle\cdot\langle x|PMQ|y\rangle = \frac{1}{2^n}\sum_{x\in\{0,1\}^n}\langle x|PMQ|x\rangle = \frac{\mathrm{Tr}(PMQ)}{2^n} \qquad ◀$$

▶ **Lemma 16.** $\langle\sigma_{I^{\otimes n}}|A|\sigma_{I^{\otimes n}}\rangle = 0$.

Proof.

$$\langle\sigma_{I^{\otimes n}}|A|\sigma_{I^{\otimes n}}\rangle = \langle\sigma_{I^{\otimes n}}|\Pi_D\left(I^{\otimes n}\otimes H\right)\Pi_D|\sigma_{I^{\otimes n}}\rangle = \langle\sigma_{I^{\otimes n}}|I^{\otimes n}\otimes H|\sigma_{I^{\otimes n}}\rangle$$

$$= \frac{1}{2^n}\mathrm{Tr}\left(H\right) = 0 \qquad \text{(Fact 15)}$$

recalling that we have assumed that $\mathrm{Tr}(H) = 0$. ◀

▶ **Lemma 17.** *For $k \geq 2$, $|\langle\sigma_{I^{\otimes n}}|A^k|\sigma_{I^{\otimes n}}\rangle| \leq \langle\sigma_{I^{\otimes n}}|A^2|\sigma_{I^{\otimes n}}\rangle = \|H_{>k}\|_2^2$.*

Proof. The first inequality follows because $\|A\|_\infty \leq \|H\|_\infty \leq 1$, and the fact that H is Hermitian and so A is too, meaning that every eigenvalue of A^k is non-increasing in magnitude as a function of k, and non-negative when k is even.

For the second equality, we observe that

$$A|\sigma_{I^{\otimes n}}\rangle = \Pi_D\left(I^{\otimes n}\otimes H\right)\Pi_D|\sigma_{I^{\otimes n}}\rangle = \Pi_D\left(I^{\otimes n}\otimes H\right)|\sigma_{I^{\otimes n}}\rangle = \left(I^{\otimes n}\otimes H_{>k}\right)|\sigma_{I^{\otimes n}}\rangle,$$

as H has no identity component. By Fact 15,

$$\langle\sigma_{I^{\otimes n}}|A^2|\sigma_{I^{\otimes n}}\rangle = \langle\sigma_{I^{\otimes n}}|I^{\otimes n}\otimes(H_{>k})^2|\sigma_{I^{\otimes n}}\rangle = \frac{1}{2^n}\mathrm{Tr}\left((H_{>k})^2\right) = \|H_{>k}\|_2^2. \qquad ◀$$

Combining Lemmas 14, 16, and 17, we are able to give bounds on the acceptance probability of Algorithm 1 (assuming it does not terminate early) based on how close or far H is from being k-local. This gives us an algorithm for testing locality, through repetition of Algorithm 1 and concentration of measure.

▶ **Lemma 18.** *Let $\varepsilon := \|H_{>k}\|_2$. The probability that Algorithm 1 outputs 1, conditioned on not terminating early, is at least $\varepsilon^2t^2\left(1-\frac{3}{10}\varepsilon^2t^2\right)-\frac{7}{2}\varepsilon\alpha t^2$ and no more than $\varepsilon^2t^2\left(1+\frac{1}{10}t^2\right)+\frac{287}{80}\varepsilon\alpha t^2+\frac{49}{1600}\varepsilon_2\alpha t^2$.*[8]

Proof. At the end of Algorithm 1 (assuming it did not terminate early), the final state lies in D. By Lemma 14 and the definition of the final measurement, the probability that the algorithm outputs 1 is the squared length of the component of $|\psi\rangle := e^{-iAt}|\sigma_{I^{\otimes n}}\rangle + |\Delta\rangle$ along the complement of $|\sigma_{I^{\otimes n}}\rangle$, for some Δ such that $\||\Delta\rangle\|_2 \le 2\alpha t$. So by the triangle inequality, $\Pr[X=1]$ is in the range[9]

$$\left(\left(\sqrt{1-|\langle\sigma_{I^{\otimes n}}|e^{-iAt}|\sigma_{I^{\otimes n}}\rangle|^2}-\||\Delta\rangle\|_2\right)^2,\left(\sqrt{1-|\langle\sigma_{I^{\otimes n}}|e^{-iAt}|\sigma_{I^{\otimes n}}\rangle|^2}+\||\Delta\rangle\|_2\right)^2\right).$$

To analyze $\left|\langle\sigma_{I^{\otimes n}}|e^{-iAt}|\sigma_{I^{\otimes n}}\rangle\right|$, we note that because A is Hermitian, $\langle\sigma_{I^{\otimes n}}|A^k|\sigma_{I^{\otimes n}}\rangle$ is real-valued for all $k\ge 0$. By splitting up the Taylor expansion of the matrix exponential into real and imaginary terms, we see that

$$\begin{aligned}&\left|\langle\sigma_{I^{\otimes n}}|e^{-iAt}|\sigma_{I^{\otimes n}}\rangle\right|^2=\left|\langle\sigma_{I^{\otimes n}}|\left(\sum_{m=0}^{\infty}(-i)^m\frac{A^mt^m}{m!}\right)|\sigma_{I^{\otimes n}}\rangle\right|^2\\&=\left|\langle\sigma_{I^{\otimes n}}|\left(\sum_{m=0}^{\infty}(-1)^m\frac{A^{2m}t^{2m}}{(2m)!}\right)|\sigma_{I^{\otimes n}}\rangle\right|^2+\left|\langle\sigma_{I^{\otimes n}}|\left(\sum_{m=0}^{\infty}(-1)^{m+1}\frac{A^{2m+1}t^{2m+1}}{(2m+1)!}\right)|\sigma_{I^{\otimes n}}\rangle\right|^2.\end{aligned}$$

Analyzing the first term, we see that

$$\begin{aligned}&\left|\langle\sigma_{I^{\otimes n}}|\left(\sum_{m=0}^{\infty}(-1)^m\frac{A^{2m}t^{2m}}{(2m)!}\right)|\sigma_{I^{\otimes n}}\rangle\right|\\&=\left|\langle\sigma_{I^{\otimes n}}|\left(I^{\otimes 2n}-\frac{t^2}{2}A^2+\sum_{m=2}^{\infty}(-1)^m\frac{A^{2m}t^{2m}}{(2m)!}\right)|\sigma_{I^{\otimes n}}\rangle\right|\\&=\left|\frac{\mathrm{Tr}(I^{\otimes n})}{2^n}-\frac{t^2}{2}\langle\sigma_{I^{\otimes n}}|A^2|\sigma_{I^{\otimes n}}\rangle+\langle\sigma_{I^{\otimes n}}|\left(\sum_{m=2}^{\infty}(-1)^m\frac{A^{2m}t^{2m}}{(2m)!}\right)|\sigma_{I^{\otimes n}}\rangle\right| &&\text{(Fact 15)}\\&=\left|1-\frac{\varepsilon^2t^2}{2}+\sum_{m=2}^{\infty}(-1)^m\langle\sigma_{I^{\otimes n}}|\frac{A^{2m}t^{2m}}{(2m)!}|\sigma_{I^{\otimes n}}\rangle\right| &&\text{(Lemma 17)}\\&=1-\frac{\varepsilon^2t^2}{2}+\eta_{\mathrm{real}}\end{aligned}$$

where $|\eta_{\mathrm{real}}|\le\frac{\varepsilon^2t^4}{24}\cosh(t)\le\frac{\varepsilon^2t^4}{20}$ by Fact 10, Lemma 17, the triangle inequality, and the fact that $t\le\frac{1}{2}$.

[8] The ε_2 in the $\frac{49}{1600}\varepsilon_2\alpha t^2$ term of the upper bound is intended and *not* a typo.

[9] One might think to use $1-\left|\langle\sigma_{I^{\otimes n}}|\left(e^{-iAt}|\sigma_{I^{\otimes n}}\rangle+|\Delta\rangle\right)\right|^2$ followed by the triangle inequality, but this actually leads to a lossy analysis of the number of queries used.

Then, for the second term, we have

$$\begin{aligned}
\eta_{\text{imaginary}} &:= \left| \langle \sigma_{I^{\otimes n}} | \left(\sum_{m=0}^{\infty} (-1)^m \frac{A^{2m+1} t^{2m+1}}{(2m+1)!} \right) | \sigma_{I^{\otimes n}} \rangle \right| \\
&= \left| \langle \sigma_{I^{\otimes n}} | \left(A + \sum_{m=1}^{\infty} (-1)^{m+1} \frac{A^{2m+1} t^{2m+1}}{(2m+1)!} \right) | \sigma_{I^{\otimes n}} \rangle \right| \\
&= \left| \langle \sigma_{I^{\otimes n}} | \left(\sum_{m=1}^{\infty} (-1)^m \frac{A^{2m+1} t^{2m+1}}{(2m+1)!} \right) | \sigma_{I^{\otimes n}} \rangle \right| && \text{(Lemma 16)} \\
&\leq \varepsilon^2 \sum_{m=1}^{\infty} \frac{t^{2m+1}}{(2m+1)!} && \text{(Lemma 17)} \\
&\leq \varepsilon^2 \frac{t^3}{6} \cosh(t) \leq \frac{1}{10} \varepsilon^2 t^2. && \text{(Fact 10)}
\end{aligned}$$

Since

$$\left| \langle \sigma_{I^{\otimes n}} | e^{-iAt} | \sigma_{I^{\otimes n}} \rangle \right|^2 = \left(1 - \frac{\varepsilon^2 t^2}{2} + \eta_{\text{real}} \right)^2 + \eta_{\text{imaginary}}^2,$$

we can upper bound it by $\left(1 - \frac{\varepsilon^2 t^2}{2} + |\eta_{\text{real}}|\right)^2 + \eta_{\text{imaginary}}^2$ and, as $\eta_{\text{imaginary}} \geq 0$, lower bound it by $\left(1 - \frac{\varepsilon^2 t^2}{2} - |\eta_{\text{real}}|\right)^2$.

We can therefore upper bound the probability of Algorithm 1 accepting by

$$\begin{aligned}
&\left(\sqrt{1 - |\langle \sigma_{I^{\otimes n}} | e^{-iAt} | \sigma_{I^{\otimes n}} \rangle|^2} + \| |\Delta\rangle \|_2 \right)^2 \\
&\leq \left(\sqrt{1 - \left(1 - \frac{\varepsilon^2 t^2}{2} - |\eta_{\text{real}}|\right)^2} + \frac{7}{4} \alpha t \right)^2 && \text{(Lemma 14)} \\
&\leq \left(\sqrt{\varepsilon^2 t^2 + 2|\eta_{\text{real}}|} + \frac{7}{4} \alpha t \right)^2 \\
&\leq \varepsilon^2 t^2 + 2|\eta_{\text{real}}| + \frac{7}{2} \alpha t \sqrt{\varepsilon^2 t^2 + \frac{1}{10} \varepsilon^2 t^4} + \frac{49}{16} \alpha^2 t^2 \\
&\leq \varepsilon^2 t^2 \left(1 + \frac{1}{10} t^2\right) + \frac{287}{80} \varepsilon \alpha t^2 + \frac{49}{1600} \varepsilon_2 \alpha t^2 && \left(t \leq 0.5,\ \alpha \leq \frac{\varepsilon_2}{100}\right)
\end{aligned}$$

and lower bound it by

$$\begin{aligned}
&\left(\sqrt{1 - |\langle \sigma_{I^{\otimes n}} | e^{-iAt} | \sigma_{I^{\otimes n}} \rangle|^2} - \| |\Delta\rangle \|_2 \right)^2 \\
&\geq \left(\sqrt{1 - \left(1 - \frac{\varepsilon^2 t^2}{2} + |\eta_{\text{real}}|\right)^2 - \eta_{\text{imaginary}}^2} - \| |\Delta\rangle \|_2 \right)^2 \\
&\geq \varepsilon^2 t^2 - \left(\frac{\varepsilon^2 t^2}{2} + |\eta_{\text{real}}| \right)^2 - |\eta_{\text{imaginary}}|^2 - \frac{7}{2} \varepsilon \alpha t^2 \\
&\geq \varepsilon^2 t^2 \left(1 - \frac{3}{10} \varepsilon^2 t^2\right) - \frac{7}{2} \varepsilon \alpha t^2.
\end{aligned}$$

◀

▶ **Theorem 1.** *Let $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$, $\delta \in (0,1)$, and $k \in \mathbb{N}$. There is an algorithm that distinguishes whether an n-qubit Hamiltonian H is (1) within ε_1 of some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, with probability $1-\delta$. The algorithm uses $\mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2-\varepsilon_1)^7}}\log(1/\delta)\right)$ non-adaptive queries to the time evolution operator with $\mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2-\varepsilon_1)^5}}\log(1/\delta)\right)$ total evolution time.*

Proof. By Lemma 18 the output of Algorithm 1, conditioned on succeeding, is a Bernoulli random variable X_i with bounded expectation. That is, when $\varepsilon \geq \varepsilon_2$ then

$$\mathbb{E}[X_i] \geq \varepsilon_2^2 t^2 \left(1 - \frac{3}{10}\varepsilon_2^2 t^2\right) - \frac{7}{2}\varepsilon_2 \alpha t^2$$

and when $\varepsilon \leq \varepsilon_1$ then

$$\mathbb{E}[X_i] \leq \varepsilon_1^2 t^2 \left(1 + \frac{1}{10} t^2\right) + \frac{287}{80}\varepsilon_1 \alpha t^2 + \frac{49}{1600}\varepsilon_2 \alpha t^2.$$

Let

$$\tau := \frac{1}{2}\left[\varepsilon_2^2 t^2 \left(1 - \frac{3}{10}\varepsilon_2^2 t^2\right) - \frac{7}{2}\varepsilon_2 \alpha t^2 + \varepsilon_1^2 t^2 \left(1 + \frac{1}{10} t^2\right) + \frac{287}{80}\varepsilon_1 \alpha t^2 + \frac{49}{1600}\varepsilon_2 \alpha t^2\right]$$

then be the halfway point these two values, and our decision threshold. And for convenience let

$$\xi := \frac{1}{2}\left[\varepsilon_2^2 t^2 \left(1 - \frac{3}{10}\varepsilon_2^2 t^2\right) - \frac{7}{2}\varepsilon_2 \alpha t^2 - \varepsilon_1^2 t^2 \left(1 + \frac{1}{10} t^2\right) - \frac{287}{80}\varepsilon_1 \alpha t^2 - \frac{49}{1600}\varepsilon_2 \alpha t^2\right]$$

be a lower bound on the distance from τ to our bounds on $\mathbb{E}[X_i]$. Observe that $\varepsilon_1 < \varepsilon_2 \leq 1$, $\varepsilon_2 \alpha = \frac{\varepsilon_2^2 - \varepsilon_1^2}{100}$ and $t = \frac{\sqrt{\varepsilon_2^2-\varepsilon_1^2}}{2\varepsilon_2}$ so:

$$\frac{9}{80}\frac{(\varepsilon_2^2 - \varepsilon_1)^2}{\varepsilon_2^2} \leq \frac{1}{2}(\varepsilon_2^2 - \varepsilon_1^2)t^2 - \frac{1}{5}\varepsilon_2^2 t^4 \leq \xi \leq \frac{\varepsilon_2^2 - \varepsilon_1^2}{2}t^2 \leq \frac{\varepsilon_2^2 t^2}{2}.$$

Now say that we have i.i.d samples $\{X_1, \ldots, X_s\}$ from *successful* runs of Algorithm 1 for s to be determined and let $X := \sum_{i=1}^s X_i$. If $\varepsilon \geq \varepsilon_2$, then by Bernstein's inequality the probability that $X \leq s\tau$ is at most:

$$\begin{aligned}
\Pr\left[\sum_{i=1}^s X_i \leq s\tau\right] &= \Pr\Big[X - \mathbb{E}[X] \leq s\tau - \mathbb{E}[X]\Big] \\
&\leq \exp\left[-\frac{\frac{(s\tau - \mathbb{E}[X])^2}{2}}{s\,\mathbb{E}[X]\,(1-\mathbb{E}[X]) + \frac{\mathbb{E}[X]-s\tau}{3}}\right] \\
&\leq \exp\left[-\frac{(s\tau - \mathbb{E}[X])^2}{2\left(s\,\mathbb{E}[X] + \frac{\mathbb{E}[X]-s\tau}{3}\right)}\right] \\
&\leq \exp\left[-\frac{s\xi^2}{2\left(\varepsilon_2^2 t^2 + \frac{\xi}{3}\right)}\right] \\
&\leq \exp\left[-\frac{3s\xi^2}{7\varepsilon_2^2 t^2}\right] \\
&\leq \exp\left[-\frac{s}{46.5}\frac{(\varepsilon_2^2 - \varepsilon_1^2)^3}{\varepsilon_2^4}\right]
\end{aligned}$$

where the fourth line follows due to the expression in the exponential being monotonically increasing with respect to $\mathbb{E}[X] \in (\tau, 1]$. Likewise, if $\varepsilon \leq \varepsilon_1$ then the probability that $X \geq s\tau$ is at most:

$$\begin{aligned}
\Pr\left[\sum_{i=1}^{s} X_i \geq s\tau\right] &= \Pr\left[X - \mathbb{E}[X] \geq s\tau - \mathbb{E}[X]\right] \\
&\leq \exp\left[-\frac{\frac{(s\tau - \mathbb{E}[X])^2}{2}}{s\,\mathbb{E}[X]\left(1 - \mathbb{E}[X]\right) + \frac{s\tau - \mathbb{E}[X]}{3}}\right] \\
&\leq \exp\left[-\frac{\left(s\tau - \mathbb{E}[X]\right)^2}{2\left(s\,\mathbb{E}[X] + \frac{s\tau - \mathbb{E}[X]}{3}\right)}\right] \\
&\leq \exp\left[-\frac{s\xi^2}{2\left(\varepsilon_1^2 t^2\left(1 + \frac{1}{10}t^2\right) + \frac{287}{80}\varepsilon_1 \alpha t^2 + \frac{49}{1600}\varepsilon_2 \alpha t^2 + \frac{\xi}{3}\right)}\right] \\
&\leq \exp\left[-\frac{s\xi^2}{2\left(\varepsilon_2^2 t^2\left(1 + \frac{1}{40} + \frac{287}{800} + \frac{49}{16000} + \frac{1}{6}\right)\right)}\right] \\
&\qquad\qquad\qquad\qquad (\varepsilon_1 < \varepsilon_2,\ t \leq 0.5,\ \alpha \leq \frac{\varepsilon_2}{100}) \\
&\leq \exp\left[-\frac{s}{55.9}\frac{(\varepsilon_2^2 - \varepsilon_1^2)^3}{\varepsilon_2^4}\right]
\end{aligned}$$

where the fourth line also follows due to the expression in the exponential being monotonically decreasing with respect to $\mathbb{E}[X] \in [0, \tau)$. Therefore, setting

$$s = 55.9\frac{\varepsilon_2^4}{(\varepsilon_2^2 - \varepsilon_1^2)^3}\ln(2/\delta)$$

suffices for us to succeed at distinguishing the two cases with probability at most $1 - \delta/2$.

Algorithm 1 has an $\frac{99}{98}\alpha t < \frac{99}{19600}\frac{(\varepsilon_2^2 - \varepsilon_1^2)^{3/2}}{\varepsilon_2^2} \leq \frac{99}{19600}$ chance of failure. By applying Corollary 6,

$$s' = \frac{2}{1 - \frac{99}{19600}}\left(s + \ln(2/\delta)\right) \leq 115\frac{\varepsilon_2^4}{(\varepsilon_2^2 - \varepsilon_1^2)^3}\ln(2/\delta)$$

suffices to achieve s successful runs with probability $1 - \delta/2$. By a union bound, we will correctly differentiate the two cases with probability at least $1 - \delta$.

The total time complexity used is then

$$\begin{aligned}
s't &\leq 115\frac{\varepsilon_2^4}{(\varepsilon_2^2 - \varepsilon_1^2)^3}\ln(2/\delta) \cdot \frac{\sqrt{\varepsilon_2^2 - \varepsilon_1^2}}{2\varepsilon_2} \leq 58\frac{\varepsilon_2^3}{\left((\varepsilon_2 - \varepsilon_1)(\varepsilon_2 + \varepsilon_1)\right)^{5/2}}\log(2/\delta) \\
&\leq 58\sqrt{\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^5}}\log(2/\delta) = \mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^5}}\log(1/\delta)\right),
\end{aligned}$$

with a total number of queries of

$$\begin{aligned}
s'm = \frac{s't}{\alpha} &\leq 58\frac{\varepsilon_2^3}{(\varepsilon_2^2 - \varepsilon_1^2)^{5/2}}\log(2/\delta) \cdot \frac{100\varepsilon_2}{\varepsilon_2^2 - \varepsilon_1^2} \leq 5800\frac{\varepsilon_2^4}{(\varepsilon_2^2 - \varepsilon_1^2)^{7/2}}\log(2/\delta) \\
&\leq 5800\sqrt{\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^7}} = \mathrm{O}\left(\sqrt{\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^7}}\right).
\end{aligned}$$

◀

References

1 Dorit Aharonov, Itai Arad, and Thomas Vidick. The quantum pcp conjecture, 2013. arXiv:1309.7495.

2 Anurag Anshu, Srinivasan Arunachalam, Tomotaka Kuwahar, and Mehdi Soleimanifar. Sample-efficient learning of interacting quantum systems. *Nature Physics*, 17:931–935, August 2021. doi:10.1038/s41567-021-01232-0.

3 Srinivasan Arunachalam, Arkopal Dutt, and Francisco Escudero Gutiérrez. Testing and learning structured quantum hamiltonians, 2024. arXiv:2411.00082.

4 Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. Learning quantum hamiltonians at any temperature in polynomial time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 1470–1477, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649619.

5 Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. Structure learning of hamiltonians from real-time evolution. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1037–1050, 2024. doi:10.1109/FOCS61266.2024.00069.

6 Andreas Bluhm, Matthias C. Caro, and Aadil Oufkir. Hamiltonian Property Testing, 2024. arXiv:2403.02968.

7 Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum Amplitude Amplification and Estimation, 2002. doi:10.1090/conm/305/05215.

8 Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018. doi:10.1073/pnas.1801723115.

9 Avshalom C. Elitzur and Lev Vaidman. Quantum mechanical interaction-free measurements. *Foundations of Physics*, 23:987–997, 1993. doi:10.1007/BF00736012.

10 P Facchi and S Pascazio. Quantum zeno dynamics: mathematical and physical aspects. *Journal of Physics A: Mathematical and Theoretical*, 41(49):493001, October 2008. doi:10.1088/1751-8113/41/49/493001.

11 Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. Low-Stabilizer-Complexity Quantum States Are Not Pseudorandom. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 64:1–64:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2023.64.

12 Francisco Escudero Gutiérrez. Simple algorithms to test and learn local Hamiltonians, 2024. arXiv:2404.06282.

13 J. Haah, R. Kothari, R. O'Donnell, and E. Tang. Query-optimal estimation of unitary channels in diamond distance. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 363–390, Los Alamitos, CA, USA, November 2023. IEEE Computer Society. doi:10.1109/FOCS57990.2023.00028.

14 Jeongwan Haah, Robin Kothari, and Ewin Tang. Learning quantum hamiltonians from high-temperature gibbs states and real-time evolutions. *Nature Physics*, 20:1027–1031, June 2024. doi:10.1038/s41567-023-02376-x.

15 Hsin-Yuan Huang, Yu Tong, Di Fang, and Yuan Su. Learning many-body hamiltonians with heisenberg-limited scaling. *Phys. Rev. Lett.*, 130:200403, May 2023. doi:10.1103/PhysRevLett.130.200403.

16 Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996. doi:10.1126/science.273.5278.1073.

17 Ashley Montanaro and Tobias J. Osborne. Quantum boolean functions, 2010. arXiv:0810.2435.

18 Ashley Montanaro and Ronald de Wolf. *A Survey of Quantum Property Testing.* Number 7 in Graduate Surveys. Theory of Computing Library, 2016. doi:10.4086/toc.gs.2016.007.

19 Adrian She and Henry Yuen. Unitary Property Testing Lower Bounds by Polynomials. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 96:1–96:17, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2023.96.

20 Ramgopal Venkateswaran and Ryan O'Donnell. Quantum Approximate Counting with Nonadaptive Grover Iterations. In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, volume 187 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 59:1–59:12, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.STACS.2021.59.

21 Mark M Wilde. *Quantum Information Theory.* Cambridge university press, 2 edition, 2017.

A Lower Bound

We will utilize the following fact about diamond distance of unitaries that will make calculations easier, at a loss of some constant factors.

▶ **Fact 19** ([13, Proposition 1.6])**.** *For all unitaries U and V of equal dimension,*

$$\frac{1}{2}\|U-V\|_\diamond \le \min_{\theta\in[0,2\pi)} \|e^{i\theta}U - V\|_\infty \le \|U-V\|_\diamond.$$

We now show our lower bound for k-locality testing, simply by showing that the statistical distance of the resulting unitaries (i.e., diamond distance) only grows linearly with time.

▶ **Definition 20.** *For $0 \le k \le n$, we define*

$$Z_{1:k} := \bigotimes_{i=1}^{k} Z \otimes \bigotimes_{j=k+1}^{n} I$$

to be the tensor product of Z on the first k qubits and identity on the last $n-k$ qubits.

▶ **Lemma 21.** *For $0 \le \varepsilon_1 \le \varepsilon_2$*

$$\|e^{-iZ_{1:k}\varepsilon_1 t} - e^{-iZ_{1:k}\varepsilon_2 t}\|_\diamond \le 2(\varepsilon_1 - \varepsilon_2)t.$$

Proof. Since $Z_{1:k}$ is diagonal with ±1 entries, $e^{-iZ_{1:k}\varepsilon t}$ is diagonal with entries $e^{\mp i\varepsilon t}$. Therefore, the eigenvalues of $e^{i\theta}\cdot e^{-iZ_{1:k}\varepsilon_1 t} - e^{-iZ_{1:k}\varepsilon_2 t}$ can be directly calculated, giving us

$$\begin{aligned}
&\min_{\theta\in[0,2\pi)} \|e^{i\theta}\cdot e^{-iZ_{1:k}\varepsilon_1 t} - e^{-iH\varepsilon_2 t}\|_\infty \\
&= \min_{\theta\in[0,2\pi)} \max\left(|e^{i(\theta-\varepsilon_1 t)} - e^{-i\varepsilon_2 t}|, |e^{i(\theta+\varepsilon_1 t)} - e^{i\varepsilon_2 t}|\right) \\
&= \min\left(|e^{-i\varepsilon_1 t} - e^{-i\varepsilon_2 t}|, |e^{-i\varepsilon_1 t} + e^{-i\varepsilon_2 t}|\right) \\
&= 2\min\left(\left|\sin\left(\frac{(\varepsilon_2-\varepsilon_1)t}{2}\right)\right|, \left|\cos\left(\frac{(\varepsilon_2-\varepsilon_1)t}{2}\right)\right|\right) \\
&\le (\varepsilon_2-\varepsilon_1)t,
\end{aligned}$$

where one of $\theta \in \{0,\pi\}$ minimizes the value via symmetry. By Fact 19, $\|e^{-iZ_{1:k}\varepsilon_1 t} - e^{-iZ_{1:k}\varepsilon_2 t}\|_\diamond \le 2(\varepsilon_1-\varepsilon_2)t$.[10] ◀

[10] A direct calculation of the diamond distance will give an upper bound of $(\varepsilon_2-\varepsilon_1)t$, without the factor of 2 from Fact 19. See [14, Proof of Proposition 1.6].

▶ Remark 22. Lemma 21 easily extends to the scenario where one is allowed to make calls to the inverse oracle, controlled versions of the oracle, the complex conjugate of the oracle, and any combination of these augmentations, as the diamond distance between the corresponding unitaries can be bounded as a function of time evolution.

We are now ready to prove our tolerant locality testing lower bound by reducing to Lemma 21.

▶ **Theorem 2.** *Let $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$ and $k \in \mathbb{N}$. Then any algorithm that can distinguish whether an n-qubit Hamiltonian H is (1) within ε_1 of some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, must use $\Omega\left(\frac{1}{\varepsilon_2 - \varepsilon_1}\right)$ evolution time in expectation to achieve constant success probability.*

Proof. Observe that for any $k' > k$, $H_1 := \varepsilon_1 Z_{1:k'}$ is within ε_1 of being k-local and $H_2 := \varepsilon_2 Z_{1:k'}$ is likewise ε_2-far from being k-local. $\|H_1\|_\infty \leq \|H_2\|_\infty \leq 1$ is also satisfied. Let t_i be the time evolution for each query in our algorithm. By Lemma 21, the diamond distance between the time evolution of these two cases is at most $2(\varepsilon_2 - \varepsilon_1)t_i$ for each query. By the sub-additivity of diamond distance, a total time evolution of $\sum_i t_i = \Omega\big((\varepsilon_2 - \varepsilon_1)^{-1}\big)$ is required to distinguish H_1 and H_2 with constant probability. ◀

▶ Remark 23. Theorem 2 also holds when the distance to k-locality is determined by operator norm $\|\cdot\|_\infty$, any *normalized* schatten p-norm $\|X\|_p := \frac{1}{2^{n/p}} \mathrm{Tr}\left(|X|^p\right)^{\frac{1}{p}}$, or any Pauli decomposition p-norm $\|X\|_{\mathrm{Pauli},p} := \left(\sum_{P \in \mathcal{P}^{\otimes n}} |\alpha_P|^p\right)^{\frac{1}{p}}$ for $X = \sum_{P \in \mathcal{P}^{\otimes n}} \alpha_P P$, improving upon that of [6, Theorem 3.6]. This is simply because the distance of $\varepsilon Z_{1:k'}$ (for $k' > k$) from being k-local is exactly ε for all of these distance measures.

B Optimal Tolerant Testing with Inverse Queries

In this section we augment the tolerant testing algorithm in [12, 3], with amplitude estimation to get an optimal tolerant tester when given access to controlled versions of the forward and reverse time evolution.[11]

We begin with the following crucial result of Gutiérrez.

▶ **Lemma 24** ([3, Lemma 3.1]). *Let $0 \leq \varepsilon_1 \leq \varepsilon_2 \leq 1$. Let $\alpha := \frac{\varepsilon_2 - \varepsilon_1}{3c}$ and H be an n-qubit Hamiltonian with $\|H\|_\infty = 1$. Define $U := e^{-iH\alpha}$, and let $U_{>k}$ be $U|\sigma_{I^{\otimes n}}\rangle$ projected onto onto the space spanned by $\{(I \otimes P)|\sigma_{I^{\otimes n}}\rangle : P \in \{I, X, Y, Z\}^{\otimes n}, |P| > k\}$. We have that if H is ε_1-close to being k-local, then*

$$\|U_{>k}\|_2^2 \leq \left((\varepsilon_2 - \varepsilon_1)\frac{2\varepsilon_1 + \varepsilon_2}{9c}\right)^2,$$

and if H is ε_2-far from being k-local, then

$$\|U_{>k}\|_2^2 \geq \left((\varepsilon_2 - \varepsilon_1)\frac{\varepsilon_1 + 2\varepsilon_2}{9c}\right)^2.$$

We also cite the following result of [11], which itself follows as a corollary of the celebrated Quantum Amplitude Estimation [7, Theorem 12] result.

[11] Using the multiplicative error form from [20] should allow for one to remove the need for controlled access while remaining non-adaptive, though it causes the constants to blow-up.

▶ **Lemma 25** (Quantum Amplitude Estimation [11, Corollary 29]). *Let Π be a projector and $|\psi\rangle$ be an n-qubit pure state such that $\langle\psi|\Pi|\psi\rangle = \eta$. Given access to the unitary transformations $R_\Pi = 2\Pi - I$ and $R_\psi = 2|\psi\rangle\langle\psi| - I$, there exists a quantum algorithm that outputs $\widehat{\eta}$ such that*

$$|\widehat{\eta} - \eta| \leq \xi$$

with probability at least $\frac{8}{\pi^2}$. The algorithm makes no more than $\pi\frac{\sqrt{\eta(1-\eta)+\xi}}{\xi}$ calls to the controlled versions of R_Π and R_ψ.

In particular, this implies that if we have (controlled) query access to U, U^* for some unitary U, and a known state $|\phi\rangle$, we can estimate $\eta = \|\Pi U|\phi\rangle\|_2^2$ to ζ accuracy by defining $|\psi\rangle := U|\phi\rangle$ and implementing R_ψ with controlled applications of U.

We are now ready to state the algorithm, which can be seen as the algorithm of [12, 3] augmented with Lemma 25.

▶ **Theorem 4.** *Let $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$, $\delta \in (0,1)$, and $k \in \mathbb{N}$. There is an algorithm that tests whether an n-qubit Hamiltonian H is (1) ε_1-close to some k-local Hamiltonian or (2) ε_2-far from all k-local Hamiltonians, with probability $1-\delta$. The algorithm uses $\mathrm{O}\left(\frac{\log(1/\delta)}{(\varepsilon_2-\varepsilon_1)^2}\right)$ non-adaptive queries to the time evolution operator and its inverse, with $\mathrm{O}\left(\frac{\log(1/\delta)}{\varepsilon_2-\varepsilon_1}\right)$ total evolution time.*

Proof. Let $U := e^{-iH\alpha}$ as in Lemma 24. We apply Lemma 24 with Π the projector onto the space spanned by $\{(I \otimes P)|\sigma_{I^{\otimes n}}\rangle : P \in \{I, X, Y, Z\}^{\otimes n}, |P| > k\}$ to estimate $\|U_{>k}\|_2^2$. Observe that the absolute difference between the two terms in Lemma 24 is

$$\left((\varepsilon_2 - \varepsilon_1)\frac{\varepsilon_1 + 2\varepsilon_2}{9c}\right)^2 - \left((\varepsilon_2 - \varepsilon_1)\frac{2\varepsilon_1 + \varepsilon_2}{9c}\right)^2 = \frac{(\varepsilon_2 - \varepsilon_1)^3(\varepsilon_2 + \varepsilon_1)}{27c^2}.$$

Therefore, we can distinguish the two cases to constant success probability by estimating $\eta = \|U_{>k}\|_2^2$ to error $\zeta = \frac{(\varepsilon_2-\varepsilon_1)^3(\varepsilon_2+\varepsilon_1)}{54c^2}$. By Lemma 25, the number of queries is then no more than

$$\begin{aligned}
&\pi\frac{\sqrt{(\varepsilon_2 - \varepsilon_1)^2(\varepsilon_1 + 2\varepsilon_2)^2/(81c^2) + (\varepsilon_2 - \varepsilon_1)^3(\varepsilon_1 + \varepsilon_2)/(54c^2)}}{(\varepsilon_2 - \varepsilon_1)^3(\varepsilon_1 + \varepsilon_2)/(54c^2)}\\
=&\frac{54\pi c}{(\varepsilon_2 - \varepsilon_1)^2}\frac{\sqrt{(\varepsilon_1 + 2\varepsilon_2)^2/81 + (2\varepsilon_2 - 2\varepsilon_1)(2\varepsilon_1 + 2\varepsilon_2)/216}}{\varepsilon_1 + \varepsilon_2}\\
\leq&\frac{54\pi c}{(\varepsilon_2 - \varepsilon_1)^2}\frac{\sqrt{(2\varepsilon_1 + 2\varepsilon_2)^2/81 + (2\varepsilon_1 + 2\varepsilon_2)^2/216}}{\varepsilon_1 + \varepsilon_2}\\
\leq&\frac{54\pi c}{(\varepsilon_2 - \varepsilon_1)^2}\frac{\sqrt{11(2\varepsilon_1 + 2\varepsilon_2)^2/648}}{\varepsilon_1 + \varepsilon_2}\\
\leq&\frac{3\sqrt{22}\pi c}{(\varepsilon_2 - \varepsilon_1)^2}.
\end{aligned}$$

Since the Hamiltonian is applied for $\alpha := \frac{\varepsilon_2-\varepsilon_1}{3c}$ for each query, the total evolution of the Hamiltonian is at most

$$\frac{3\sqrt{22}\pi c}{(\varepsilon_2 - \varepsilon_1)^2}\frac{\varepsilon_2 - \varepsilon_1}{3c} = \frac{\sqrt{22}\pi}{\varepsilon_2 - \varepsilon_1}.$$

By standard error reduction, we can reduce the constant failure probability to at most δ using $\log(1/\delta)$ repetitions.

Finally, observe that constructing R_Π (and its controlled version), as in Lemma 25 is free, as Π is a known projector onto the low locality Paulis. On the other hand, R_ψ requires us to take (a version of) the Grover Diffusion operator $D := 2|0\rangle\langle 0| - I$ and conjugate it by U. This is the step that requires access to $U^\dagger := e^{iH\alpha}$. ◀

Since this matches the lower bound of Theorem 2, Theorem 4 is optimal.

Quantum Catalytic Space

Harry Buhrman
Quantinuum London, UK
QuSoft, Amsterdam, The Netherlands

Marten Folkertsma
CWI, Amsterdam, The Netherlands
QuSoft, Amsterdam, The Netherlands

Ian Mertz
Charles University, Prague, Czech Republic

Florian Speelman
University of Amsterdam, The Netherlands
QuSoft, Amsterdam, The Netherlands

Sergii Strelchuk
University of Oxford, UK

Sathyawageeswar Subramanian
University of Cambridge, UK

Quinten Tupker
CWI, Amsterdam, The Netherlands
QuSoft, Amsterdam, The Netherlands

Abstract

Space complexity is a key field of study in theoretical computer science. In the quantum setting there are clear motivations to understand the power of space-restricted computation, as qubits are an especially precious and limited resource.

Recently, a new branch of space-bounded complexity called catalytic computing has shown that reusing space is a very powerful computational resource, especially for subroutines that incur little to no space overhead. While quantum catalysis in an information theoretic context, and the power of "dirty" qubits for quantum computation, has been studied over the years, these models are generally not suitable for use in quantum space-bounded algorithms, as they either rely on specific catalytic states or destroy the memory being borrowed.

We define the notion of catalytic computing in the quantum setting and show a number of initial results about the model. First, we show that quantum catalytic logspace can always be computed quantumly in polynomial time; the classical analogue of this is the largest open question in catalytic computing. This also allows quantum catalytic space to be defined in an equivalent way with respect to circuits instead of Turing machines. We also prove that quantum catalytic logspace can simulate log-depth threshold circuits, a class which is known to contain (and believed to strictly contain) quantum logspace, thus showcasing the power of quantum catalytic space. Finally we show that both unitary quantum catalytic logspace and classical catalytic logspace can be simulated in the one-clean qubit model.

2012 ACM Subject Classification Theory of computation → Circuit complexity; Theory of computation → Complexity classes

Keywords and phrases quantum computing, quantum complexity, space-bounded algorithms, catalytic computation, one clean qubit

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.11

Related Version *Full Version*: `https://arxiv.org/abs/2506.16324`

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 11; pp. 11:1–11:24
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Funding *Marten Folkertsma*: Supported by the Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL programme.
Ian Mertz: Supported by the Grant Agency of the Czech Republic under the grant agreement no. 24-10306S and by the Center for Foundations of Contemporary Computer Science (Charles Univ. project UNCE 24/SCI/008).
Florian Speelman: Supported by the Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL program, and the project Divide and Quantum "D&Q" NWA.1389.20.241 of the program "NWA-ORC", which is partly funded by the Dutch Research Council (NWO).
Sergii Strelchuk: Supported by a Royal Society University Research Fellowship and the EPSRC (RoaRQ), Investigation 005 [grantreference EP/W032635/1].
Sathyawageeswar Subramanian: Supported by a Royal Commission for the Exhibition of 1851 Research Fellowship.
Quinten Tupker: Supported by the Dutch National Growth Fund (NGF), as part of the Quantum Delta NL program.

Acknowledgements We thank [24] for pointing us to [3]. Part of this work was done while HB was at CWI, IM was at University of Warwick, SS (1) was at Cambridge University and University of Warwick, and SS (2) was at University of Warwick. SS (2) acknowledges support from the Royal Society University Research Fellowship.

1 Introduction

Space is one of the cornerstones of theoretical computer science, and the study of space-bounded computations has been crucial in the development of complexity theory. Investigating logspace computations revealed the limits of efficient computation under memory constraints and has led to striking results such as Savitch's theorem [38] and $\mathsf{NL} = \mathsf{coNL}$ [25, 41]. Logspace reductions are essential in classifying problems as NL-complete or P-complete, and leading to techniques for efficient parallelization and algorithm design.

Many graph and database problems rely on logspace techniques, making them relevant for query optimization, data retrieval, and formal verification. Furthermore, logspace computations have practical applications in streaming algorithms, embedded systems, cryptography, and model checking, where minimizing memory usage is critical.

The emergence of quantum computing has led to remarkable theoretical speedups over the best known classical algorithms. The promise of exponential computational advantage in using principles of quantum mechanics to process information comes with formidable experimental challenges of building and maintaining quantum computers that can implement long sequences of coherent operations. This led to a renewed interest in the structure of quantum space.

1.1 Space in quantum computation

Understanding the true extent of the power of quantum computing in a variety of space-constrained settings is a major challenge. In contrast to the classical setting where adding a reasonable amount of extra memory to support computations is routinely achievable, producing and maintaining multiple qubits is exceptionally difficult due to several fundamental physical, engineering, and scalability issues. Qubits are fragile and susceptible to decoherence, and maintaining long coherence times becomes significantly harder as the number of qubits increases. Furthermore, quantum error rates scale with the number of qubits, making fault-tolerant quantum computing a major challenge. In the quantum computational setting, space thus comes at a premium, and increasing the amount of space available for computation requires overcoming fundamental challenges to reduce error rates, increase control precision, and maintain entanglement across multiple systems, to name but a few.

The characterization of quantum logspace (QL) and the study of the computational power of bounded-error quantum logarithmic space (BQL) and its relationship to classical complexity classes was first done by Watrous [44], where it was established that $\mathsf{BQL} \subseteq \mathsf{P}$. This showed that any problem solvable in quantum logspace with bounded error is also solvable in polynomial time by a classical deterministic machine. In later work, Watrous [43] showed that $\mathsf{QSPACE}(s) \subseteq \mathsf{SPACE}[O(s^2)]$ for all $s \geq \log n$, even when the quantum machine is allowed to err with probability arbitrarily close to $1/2$; this confirms that quantum logspace computations remain simulable within polynomial space, and is consistent with classical space complexity results such as Savitch's theorem. His work also established that quantum logspace can efficiently solve certain algebraic problems, including the *group word problem for solvable groups*, which lacks efficient classical logspace algorithms [43].

These above obstacles prompted the search for extra ingredients which could lift restricted models of quantum computation (for example – realized by quantum circuits which are classically efficiently simulatable) to regain the power of universal quantum computation. These extra ingredients (e.g. magic state injection) are usually studied in the context of unrestricted space and there has as of yet been no attempt to investigate them under space restrictions.

On the other hand, there have been several notable results that illuminate various properties of quantum logspace. One of the earliest findings shows that any quantum computation that can be performed with logarithmic space can also be efficiently simulated using matchgate circuits of polynomial width, and vice versa [26]. Following this characterisation, there have been a series of further results indicating that quantum logspace describes a non-trivial class of computations. Ta-Shma [42] showed that given a matrix with a bounded condition number, a quantum logspace algorithm can efficiently approximate its inverse or solve linear systems. Girish, Raz, and Zhan [21] described a quantum logspace algorithm to compute powers of an matrix with bounded norm and prove that deterministic logspace is equal to reversible logspace. Recently, it was shown by the same authors that the class of decision problems solvable by a quantum computer in logspace admits an efficient verification procedure [22]; moreover, they also show that every language in BQL has an (information-theoretically secure) streaming proof with a quantum logspace prover and a classical logspace verifier. This hints at a curious interplay between the powers of classical and quantum logspace.

1.2 Catalysis and space

Catalysis is a concept well-studied in the context of quantum information and is widely recognized for its counterintuitive abilities to enable (state) transformations that are otherwise infeasible (see survey by Lipka et al. [30]). A related concept, known as catalytic embedding, was recently introduced in the context of circuit synthesis as an alternative to traditional gate approximation methods in quantum circuit design [4]. Here the goal is to implement a desired unitary operation *more efficiently* (e.g., with fewer gates, lower depth, or using a restricted gate set) than would be possible without assistance. It involves a specific, known, and often small catalyst state that is chosen to aid a particular unitary implementation.

These foregoing lines of work focus on the idea that a specific unitary may be implemented more efficiently if a special state (i.e. catalyst) is available, often discussing resource theories, and do not dwell on complexity theoretic implications.

In this work, we initiate the complexity-theoretic study of the effect of catalytic space in quantum computations. Much like magic state injection is able to promote and increase quantum computational power in the space-unrestricted setting, the presence of a catalyst in the form of an extra register of quantum memory – albeit memory that already contains some stored quantum information – holds a similar promise for space-bounded quantum computations. The notion of catalytic space can be regarded as a theoretical model of qubit reusal.

The first step towards a rigorous study of catalytic logspace quantum computations is to formalize the model and means of interaction with the catalytic space. Identifying new computational capabilities endowed by the presence of a catalyst in the form of additional quantum memory, which however contains an arbitrary unknown quantum state, appears to be a significantly more challenging task due to the nature of quantum information and the inherent limitations of quantum resources. For example, any framework for quantum catalytic space must incorporate the possibility of entanglement and its inherent limitations (e.g. monogamy) between the catalytic memory and the rest of the work space. It has to further account for the irreversibile nature of quantum measurement.

Remarkably, it was recently shown that the addition of a similar notion of catalytic space has major implications even in the classical logspace setting. Buhrman et al. [10] introduced a model of space, called *catalytic computing*, which studies the power of "imperfect" memory. In addition to the usual Turing machine work tape, a catalytic machine is equipped with a much larger *catalytic* work tape, which is filled with an arbitrary initial string τ and which must be reset to the configuration τ at the end of its computation.

The setting of most interest to us is *catalytic logspace* (CL), wherein a logspace machine is given access to a polynomial size catalytic tape. On the positive side, [10] showed that such machines have significantly greater power than traditional logspace, capturing the additional power of both non-determinism (NL) and randomness (BPL); in fact, they showed that CL can simulate the much larger class of *logarithmic-depth threshold circuits* (TC^1). On the negative side, they also showed that CL can be simulated by *(zero-error) randomized polynomial-time machines* (ZPP), which are strongly believed to be much weaker than e.g. polynomial space.

Since then, many works have studied classical catalytic space from a variety of angles, including further results on the power of CL [12, 1, 2] augmenting catalytic machines with other resources such as randomness or non-determinism [11, 15, 12, 29], considering non-uniform models such as catalytic branching programs or catalytic communication complexity [36, 13, 37], analyzing the robustness of classical catalytic machines to alternate conditions [9, 8, 23], and so on. Many properties of catalytic computation have emerged that appear ripe for use in the quantum setting, such as *reversibility* [18, 12], *robustness* [23, 20], and *average-case runtime bounds* [10].

Perhaps most important to motivate our current study, the utility of classical catalytic computation has been strikingly demonstrated in its use as a subroutine in an ordinary space-bounded computation: avoiding linear blowups in space when solving many instances of a problem. The most impactful result is the Tree Evaluation algorithm of Cook and Mertz [14], which was the key piece in Williams' recent breakthrough on time and space [45]. Catalytic subroutines of this kind are even more relevant in the quantum setting, as they may lead to a persistent reduction of the qubit count when executing a quantum algorithm.

1.3 Summary of results

In this paper we initiate the systematic study of catalytic techniques in the quantum setting. To this end we codify a concrete definition of quantum catalytic space (QCSPACE), explore the degrees to which the definition is robust, and establish the relationship of quantum catalytic logspace (QCL) to various classical and quantum complexity classes.

Our main technical contribution is to show that, somewhat surprisingly, quantum Turing machines and quantum circuits are equivalent even in the catalytic space setting:

▶ **Theorem 1.** *Let L be a language, and let $s := s(n)$ and $c := c(n)$. Then L is computable by a quantum catalytic Turing machine with work space $O(s)$ and catalytic space $O(c)$ iff L is computable by a family of quantum catalytic circuits with work space $O(s)$ and catalytic space $O(c)$.*

While this translation is straightforward in other settings, QCL has no *a priori* polynomial time bound, and so there is no obvious way to define the length of a catalyic circuit without running into trouble. However, we prove that the result of Buhrman et al. [10] which shows that CL takes polynomial time on average can be strengthened in the quantum case, to show that QCL *always* takes polynomial time without any error:

▶ **Theorem 2.** $\mathsf{QCL} \subseteq \mathsf{EQP}$

We find Theorem 2 intriguing for many reasons. Naturally it is exciting to be able to solve the "holy grail" of catalytic computing in the quantum setting. The story of classical catalytic computing has been the ability of clever algorithms to circumvent the resetting condition of the catalytic tape and use it for powerful purposes, but Theorem 2 shows that conversely, the additional power of quantum techniques in such algorithms does not offset the additional restrictiveness of resetting a quantum state. Quantum computation is a model fundamentally built on reversible instructions, with the one exception being the final measurement with which we obtain our answer; Theorem 2 shows that this measurement is a massive obstruction to reversibility, as having access to such a huge resource with only the reversible restriction – something which is taken care of in the intermediate computation already – gives less power than we initially assumed.

In terms of class containments, we focus on two questions: the relationship of quantum and classical catalytic space, and the relationship of catalytic space to the one-clean qubit model (DQC_1), a pre-existing object of study in quantum complexity which bears a strong resemblance to catalysis. We show that, while $\mathsf{CL} \subseteq \mathsf{QCL}$ is surprisingly out of reach at the moment, this can be shown for an important subclass of CL, one which captures the strongest known classical containment:

▶ **Theorem 3.** $\mathsf{TC}^1 \subseteq \mathsf{QCL}$

As a consequence, we show that TC^1 constitutes a natural class of functions for which catalysis gives additional power to quantum computation.

We also show that unitary QCL ($\mathsf{Q_UCL}$) and classical CL are both contained in DQC_1:

▶ **Theorem 4.** $\mathsf{BQ_UCL} \subseteq \mathsf{DQC}_1$

▶ **Theorem 5.** $\mathsf{CL} \subseteq \mathsf{DQC}_1$

Note that we use a version of DQC_1 defined using a logspace controller instead of a polynomial time controller as may also be done. These results show how much of the power of DQC_1 comes from avoiding the limitation of the resetting condition on the "dirty" work space.

1.4 Open problems

We identify a number of interesting avenues to further explore the power of quantum catalytic space, and understand its relation to various (quantum) complexity classes.

QCL subroutines

Remarkably, classical catalytic subroutines can already be used to achieve analogous space savings in QCL. Is it possible to identify genuinely quantum subroutines to achieve savings beyond those attained by classical generalizations? This is not so straightforward because the subset of qubits being reused in a catalytic subroutine could become entangled with qubits that cannot be accessed by the subroutine. Therefore, there might be a non-trivial and inaccessible reference system with respect to which the catalytic property must hold. While we show the presence of such an inaccessible reference system does not change the model we define, designing quantum catalytic subroutines (cf. classical results in [14, 45]) stands out as a fertile direction for future work.

QNC^1 vs QCL

Starting with Barrington's Theorem [5], a landmark result in space complexity, a classical line of work [6, 10] has shown that polynomial-size formulas over many different gatesets can be computed using only logarithmic space, using a reversible, algebraic characterization of computation. Such a result in the quantum case, i.e. $\mathsf{QNC}^1 \subseteq \mathsf{QL}$, appears far out of reach, as this would imply e.g. novel derandomizations in polynomial time. However, such techniques are also key to the study of catalytic computation, and so perhaps we can show QNC^1 or a similar quantum circuit class is contained in QCL. This would give a clear indication of the power of quantumness in catalytic computation.

QCL vs DQC_1

While we seem to find that $\mathsf{Q_UCL}$ or QCL without intermediate measurements is contained in DQC_1, it is unclear if this still holds when we allow intermediate measurements.

QCL with errors

One aspect of our results which is discordant with the usual mode of quantum computation is that we require the catalytic tape be *exactly* reset by the computation. On the other hand, many basic primitives in quantum computing, such as converting between gatesets, can introduce errors into the computation, and in practice even the ambient environment can be assumed to cause such issues. Thus it seems natural to study the power of QCL when we allow a small, potentially exponentially small, trace distance between the initial and final catalytic states. This model is well-understood in the classical world [23, 20], but it would be interesting to see whether our techniques can be made robust to this small error or, to the contrary, whether this slight relaxation is enough to overcome the barriers in our work, chiefly the inability to show $\mathsf{CL} \subseteq \mathsf{QCL}$.

2 Preliminaries

2.1 Quantum computation

For this work we will consider complex *Hilbert* spaces $\mathcal{H} \cong \mathbb{C}^d$ of dimension d, that will form the state space for a quantum system. Multiple quantum systems are combined by taking the tensor product of their Hilbert spaces, such as $\mathcal{H}_1 \otimes \mathcal{H}_2$. We will often write $\mathcal{H}_s$ to denote the Hilbert space $\left(\mathbb{C}^2\right)^{\otimes s}$ of s qubits, where the dimension is given by function $d(\mathcal{H}_s) = 2^s$. We will also often use the abbreviation $[n] = \{1, \ldots, n\}$. Below, we recall some of the important background required for this article, referring the reader to [32] for more details.

▶ **Definition 6** (Quantum states). *A* pure quantum states *is a unit vector of the Hilbert space $|\psi\rangle \in \mathcal{H}$, with the normalization condition $\langle\psi|\psi\rangle = 1$. We also make use of more general states represented by* density matrices *ρ which are positive semi definite operators on a Hilbert space with unit trace, $Tr[\rho] = 1$. Density matrices describe* mixed states *which, beyond pure quantum states, can also capture classical uncertainty. In other words, they correspond to classical mixtures of pure quantum states. The density matrix of a pure state is $\rho = |\psi\rangle\langle\psi|$. Given an ensemble of states $\{|\psi_i\rangle\}$ and corresponding probabilities $\{p_i\}$, with $p_i \geq 0$ and $\sum_i p_i = 1$, it can be represented by a mixed state of the form $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$. We will denote the set of mixed states a Hilbert space $\mathcal{H}$ by $D(\mathcal{H})$.*

▶ **Definition 7** (Quantum channels). *A* quantum channel *is a linear operator that maps density matrices to density matrices, $\Phi : D(\mathcal{H}_1) \to D(\mathcal{H}_2)$ (also known as superoperators or CPTP maps). It is also required to have two additional properties: 1) it must be completely positive; and 2) it must be trace preserving. We denote the set of channels from $D(\mathcal{H})$ to itself by $\mathcal{C}(D(\mathcal{H}))$.*

We denote the identity channel on d qubits by $\mathcal{I}_d$, or just $\mathcal{I}$ when d is clear from context. The *Choi matrix* of a channel Φ that acts on an input space $\mathcal{H}$ of dimension d is defined by the action of Φ on the first register of a maximally entangled state in $\mathcal{H} \otimes \mathcal{H}$

$$J(\Phi) := (\Phi \otimes \mathcal{I}_d)\left(\frac{1}{d}\sum_{i,j=1}^{d} |i\rangle\langle j| \otimes |i\rangle\langle j|\right) = \frac{1}{d}\sum_{i,j=1}^{d} \Phi\Big(|i\rangle\langle j|\Big) \otimes |i\rangle\langle j|\,.$$

▶ **Definition 8.** *The trace distance between two density matrices $\rho, \sigma \in D(\mathcal{H})$ is defined by:*

$$||\rho - \sigma||_1 = Tr[\sqrt{(\rho-\sigma)^\dagger(\rho-\sigma)}],$$

where $A^\dagger$ denotes the conjugate transpose of the matrix $A^\dagger = \bar{A}^T$.

It is well known that no physical process can increase the trace distance between two states:

▶ **Lemma 9** (Contractivity under CPTP maps [32, Theorem 9.2]). *Let $\Phi \in \mathcal{C}(D(\mathcal{H}))$ and $\rho, \sigma \in D(\mathcal{H})$ then the trace distance between ρ and σ can not increase under application of Φ:*

$$||\Phi(\rho) - \Phi(\sigma)||_1 \leq ||\rho - \sigma||_1$$

2.1.1 Quantum Turing machines

Our fundamental computation model in quantum computing will be the quantum analogue of Turing machines [17, 7], which we define informally below.

▶ **Definition 10** (Quantum Turing machine). *A* quantum Turing machine *is a classical Turing machine with an additional quantum tape and quantum register. The quantum register does not affect the classical part of the machine in any way, except in that the qubits in the quantum register can be measured in the computational basis. On doing so, the values read from the measurement are copied into the classical registry, from where they can be used to affect the operation of the machine. The quantum Turing machine can perform any gate from its quantum gate set on its quantum registry. We assume this gate set is fixed and universal. Finally, the tape head on the quantum tape can swap qubits between the quantum registry and the position that the quantum tape head is located at. This applies a two-qubit* SWAP *gate.*

We define a number of complexity classes with respect to efficient computation by quantum Turing machines [7, 35][1].

▶ **Definition 11** (BQP). BQP *is the set of all languages $L = (L_{yes}, L_{no}) \subset \{0,1\}^* \times \{0,1\}^*$ for which there exists a quantum Turing machine M using $t = \mathsf{poly}(n)$ time such that for every input $x \in L$ of length $n = |x|$,*

- *if $x \in L_{yes}$ then the probability that M accepts input x is $\geq c$,*
- *if $x \in L_{no}$ then the probability that M accepts input x is $\leq s$.*

▶ **Definition 12** (BQL). BQL *is the set of all languages $L = (L_{yes}, L_{no}) \subset \{0,1\}^* \times \{0,1\}^*$ for which there exists a quantum Turing machine M using $r = O(\log(n))$ quantum and classical space such that for every input $x \in L$ of length $n = |x|$,*

- *if $x \in L_{yes}$ then the probability that M accepts input x is $\geq c$,*
- *if $x \in L_{no}$ then the probability that M accepts input x is $\leq s$.*

The completeness and soundness parameters in both the above definitions can be chosen to be $c = 2/3$ and $s = 1/3$ without affecting the set of languages.

▶ **Definition 13** (EQP). EQP *is the set of all languages $L = (L_{yes}, L_{no}) \subset \{0,1\}^* \times \{0,1\}^*$ for which there exists a quantum Turing machine M using $t = \mathsf{poly}(n)$ time such that for every input $x \in L$ of length $n = |x|$,*

- *if $x \in L_{yes}$ then M outputs one with certainty on measurement,*
- *if $x \in L_{no}$ then M output zero with certainty on measurement.*

▶ Remark 14. Note that the definition of EQP is gateset dependent; this is due to the fact that quantum gatesets only allow universality up to approximation, which means that if a quantum complexity class requires perfect soundness and completeness, as does EQP, it also has to be gateset dependent.

2.1.2 Quantum circuits

We may also define quantum complexity classes using uniform quantum circuits. For this we use similar definitions to those provided by [19], which readers may refer to for more details.

▶ **Definition 15.** *Let $s := s(n), t := t(n), k := k(n)$, let $\mathcal{K}$ be a family of machines, and let $\mathcal{G}$ be a set of k-local operators. A $\mathcal{K}$-uniform space-s time-t family of quantum circuits over $\mathcal{G}$ is a set $\{Q_x\}_{x \in \{0,1\}^n}$, where each Q_x is a sequence of tuples $\langle i, g, j_1 \dots j_k \rangle \in [t] \times \mathcal{G} \times [s]^k$ such that there is a deterministic TM $M \in \mathcal{K}$ which, on input $x \in \mathcal{X}$, outputs a description of Q_x.*

[1] We do not attempt to provide an exhaustive list of references to the vast literature on this topic, and refer the interested reader to the Complexity Zoo for such a list.

The execution of Q_x consists of initializing a vector $|\psi\rangle$ to $|0^s\rangle$ within $\mathcal{H}_s$ and applying, for each step $i \in [t]$ in order, each gate g to qubits $j_1 \dots j_k$ such that $\langle i, g, j_1 \dots j_k\rangle \in Q_x$. The output of Q_x is the value obtained by measuring the first qubit at the end of the computation.

If $\mathcal{G}$ consists of unitary operators, we call these unitary circuits and call each g a gate. If $\mathcal{G}$ additionally consists of measurements together with postprocessing and feed forward by (classical) $\mathcal{K}$-machines, we call these general circuits and call each g a channel.

It is known that polynomial-time uniform general quantum circuits over n qubits with $\mathsf{poly}(n)$ gates can be used to provide an alternative definition of BQP [46]. Similarly, logspace uniform general quantum circuits of logarithmic width can be used as an alternative to define classes such as BQL [19].

2.2 Catalytic computation

We finally recall the known classical definitions of catalytic classical computation.

▶ **Definition 16** ([10]). *A catalytic Turing Machine with space $s := s(n)$ and catalytic space $c := c(n)$ is a Turing Machine M with a work tape of length s and a catalytic tape of length c. We require that for any $\tau \in \{0,1\}^c$, if we initialize the catalytic tape to τ, then on any given input x, the execution of M on x halts with τ on the catalytic tape.*

This definition gives rise to a natural complexity class $\mathsf{CSPACE}[s, c]$, which is a variant of the ordinary class $\mathsf{SPACE}[s]$. The most well-studied variant is *catalytic logspace*, where s is logarithmic and c is polynomial.

▶ **Definition 17.** *We define $\mathsf{CSPACE}[s, c]$ to be the class of all functions f for which there exists a catalytic Turing Machine M with space s and catalytic space c such that on input x, $M(x) = f(x)$. We further define catalytic logspace as*

$$\mathsf{CL} := \bigcup_{k\in\mathbb{N}} \mathsf{CSPACE}(k \log n, n^k)$$

3 Quantum catalytic space

The first goal of this paper is to find a proper definition of quantum catalytic space. There are many choices that have to be made in the model, but we begin with our general definition up front, leaving questions of machine model, uniformity, gateset, and initial catalytic tapes. These will be discussed and clarified in the rest of this section.

▶ **Definition 18** (Quantum catalytic machine). *A quantum catalytic machine with work space $s := s(n)$, catalytic space $c := c(n)$, uniformity $\mathcal{K}$, gateset $\mathcal{G}$, and catalytic set $\mathcal{A}$ is a $\mathcal{K}$-uniform quantum machine M with operations from $\mathcal{G}$ acting on two Hilbert spaces, $\mathcal{H}_s$ and $\mathcal{H}_c$, of dimensions 2^s and 2^c respectively. The latter space, called the catalytic tape, will be initialized to some $\rho \in \mathcal{A} \subseteq D(\mathcal{H}_c)$. We require that for any $\rho \in \mathcal{A}$, if we initialize the catalytic tape to state ρ, then on any given input $x \in \{0,1\}^n$, the execution of $M(x)$ halts with ρ on the catalytic tape. Furthermore, we require that the output state on the worktape is independent of the catalytic state ρ.*[2] *The final action of the machine can be represented by a quantum channel $\Phi_x : |0\rangle\langle 0| \otimes \rho \mapsto \eta_x \otimes \rho$, for any catalytic state ρ and input $x \in \{0,1\}^n$, and some output state η.*

[2] We justify this final requirement in Lemma 42.

This gives rise to the following complexity classes:

▶ **Definition 19** (Quantum catalytic complexity). $\mathsf{QCSPACE}[s, c]$ *is the class of Boolean functions which can be decided with probability 1 by a quantum catalytic machine with work memory s and catalytic memory c.*

$\mathsf{BQCSPACE}[s, c]$ *is the class of Boolean functions which can be decided with probability 2/3 by a quantum catalytic machine with work memory s and catalytic memory c.*

We further specify to the case of quantum catalytic logspace:

▶ **Definition 20** (Quantum catalytic logspace).

$$\mathsf{QCL} = \bigcup_{k \in \mathbb{N}} \mathsf{QCSPACE}[k \log n, n^k]$$

$$\mathsf{BQCL} = \bigcup_{k \in \mathbb{N}} \mathsf{BQCSPACE}[k \log n, n^k]$$

3.1 Machine model

We begin by defining the two natural choices of base model for quantum catalytic machines, namely *Turing machines* and *circuits*.

▶ **Definition 21** (Quantum catalytic Turing machine). *A quantum catalytic Turing machine is defined as in Definition 18 with quantum Turing machines as our machine model. We write* $\mathsf{QCSPACEM}$ *(respectively* $\mathsf{BQCSPACEM}$*,* QCLM*, and* BQCLM*) to refer to* $\mathsf{QCSPACE}$ *with quantum Turing machines.*

▶ **Definition 22** (Quantum catalytic circuits). *A quantum catalytic circuit is defined as in Definition 18 with time-$2^{O(s)}$ quantum circuits as our machine model. We write* $\mathsf{QCSPACEC}$ *(respectively* $\mathsf{BQCSPACEC}$*,* QCLC*, and* BQCLC*) to refer to* $\mathsf{QCSPACE}$ *with quantum catalytic circuits.*

Given that CL and related classes are defined in terms of (classical) Turing machines, the option of circuits seems surprising and perhaps unnatural. For example, Definition 22 imposes a time bound as part of its definition, while for CL there is no known containment in polynomial time. For quantum circuits and Turing machines without access to the catalytic tape, a simple equivalence has been known for a long time [46]; however, Definition 22 only allows for circuits of length $2^{O(s)}$, while a generic transformation on $s + c$ qubit registers would give a circuit of length $2^{O(s+c)}$, i.e. requiring an exponential overhead.

The main result of this paper is to show that these models are in fact equivalent:

▶ **Theorem 23.** *For $s = \Omega(\log n), c = 2^{O(s)}$*

$$\mathsf{QCSPACEM}[O(s), O(c)] = \mathsf{QCSPACEC}[O(s), O(c)]$$

$$\mathsf{BQCSPACEM}[O(s), O(c)] = \mathsf{BQCSPACEC}[O(s), O(c)]$$

For the rest of this section we will deal with all auxiliary issues, namely the choice of catalytic tapes and gateset, for quantum circuits alone; while all proofs can be made to hold for quantum Turing machines without much issue, this is also obviated by Theorem 23, which we will prove in Section 4.

3.2 Catalytic tapes

We now move to discussing the choice of initial catalytic tapes $\mathcal{A}$. Perhaps the most immediate choice would be to put no restrictions on $\mathcal{A}$ and allow our catalytic tapes to come from the set of all density matrices in $D(\mathcal{H}_c)$; this will ultimately be our definition.

▶ **Definition 24.** *We fix the catalytic set in Definition 18 to be* $\mathcal{A} = D(\mathcal{H}_c)$.

While this is a natural option, encompassing every possible state on c qubits, there are other choices one can make. We propose four natural options – density matrices and three others – and show that all four are equivalent, thus justifying our choice.

▶ **Definition 25.** *We define the following catalytic sets:*

- Density *is the set of all density matrices* $\rho \in D(\mathcal{H}_c)$.
- Pure *is the set of all pure states* $|\psi\rangle \in \mathcal{H}_c$.
- PauliProd $= \{|\mathrm{PP}\rangle : |\mathrm{PP}\rangle = \bigotimes_{i=1}^{c} |\phi\rangle_i\}$ *is the set of tensor products of eigenstates of the single-qubit Pauli operators, where* $|\phi\rangle_i \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |\circlearrowright\rangle, |\circlearrowleft\rangle\} \subset \mathcal{H}_2$.
- EPR $= \{\frac{1}{\sqrt{2^c}} \sum_{i=0}^{2^c-1} |i\rangle |i\rangle\} \subset \mathcal{H}_c \otimes \mathcal{H}_c$ *is the unique state of c EPR pairs, where the catalytic tape will be formed of one half of each EPR pair; the other halves are retained as a reference system which cannot be operated on by the quantum circuit. the quantum circuit is of the form* $Q_x = \tilde{Q}_x \otimes \mathcal{I}_c$*, acting as the Identity on the second set of halves of the EPR pairs that is inaccessible to the circuit.*

▶ Remark 26. We briefly comment on the fourth set, i.e. EPR. Using classical catalytic techniques as a subroutine has proven to be very useful, for instance in giving an algorithm for tree evaluation in $\mathcal{O}(\log n \log(\log n))$ space [14]. One can also consider using analogous quantum catalytic techniques as subroutines for quantum computations, albeit this does not appear straightforward due to inherent quantum limitations. We will see that this complication can be effectively modeled by considering the initial state of the catalytic tape to be the halves of c EPR pairs.

We will now prove that the four classes of quantum catalytic circuits with initial catalytic states restricted to one of the four sets $D(\mathcal{H}_c)$, $\mathcal{H}_c$, PauliProd, and EPR respectively, are all equivalent. For this we first require the following lemma.

▶ **Lemma 27.** *Any* $2^d \times 2^d$ *complex matrix can be written as a linear combination of rank-1 outer products of states from* PauliProd *over d qubits.*

In other words, the complex span of the set of d-qubit tensor products of Pauli eigenstates equals the set of $2^d \times 2^d$ complex matrices.

Proof. Note that all four Pauli matrices can be written as a linear combination of two of the Pauli eigenstates:

$$I = |0\rangle\langle 0| + |1\rangle\langle 1|, \quad X = |+\rangle\langle +| - |-\rangle\langle -|,$$
$$Z = |0\rangle\langle 0| - |1\rangle\langle 1|, \quad Y = |\circlearrowright\rangle\langle\circlearrowright| - |\circlearrowleft\rangle\langle\circlearrowleft|.$$

The four Pauli matrices form a basis for 2×2 complex matrices. Consequently, Pauli strings of length d – i.e., tensor products of d Pauli matrices – form a basis for $2^d \times 2^d$ matrices. ◀

Now we can state the theorem:

▶ **Theorem 28.** *Let* QCC_A *denote quantum catalytic circuits with initial catalytic tapes coming from A. Then The following four classes of quantum catalytic circuits are equivalent:*

$$\mathsf{QCC}_{\mathsf{Density}} = \mathsf{QCC}_{\mathsf{Pure}} = \mathsf{QCC}_{\mathsf{PauliProd}} = \mathsf{QCC}_{\mathsf{EPR}}$$

Proof. First note the obvious implications: for any quantum catalytic circuit Φ,

$$\begin{aligned} \Phi \in \mathsf{QCC}_{\mathsf{Density}} &\implies \Phi \in \mathsf{QCC}_{\mathsf{Pure}} \\ \Phi \in \mathsf{QCC}_{\mathsf{Pure}} &\implies \Phi \in \mathsf{QCC}_{\mathsf{PauliProd}} \end{aligned}$$

these follow due to the fact that $\mathsf{PauliProd} \subset \mathsf{Pure} \subset \mathsf{Density}$. To finish the proof, we will further show the following two implications.

$$\begin{aligned} &(1) \quad \Phi \in \mathsf{QCC}_{\mathsf{PauliProd}} \implies \Phi \otimes \mathcal{I}_c \in \mathsf{QCC}_{\mathsf{EPR}} \\ &(2) \quad \Phi \otimes \mathcal{I}_c \in \mathsf{QCC}_{\mathsf{EPR}} \implies \Phi \in \mathsf{QCC}_{\mathsf{Density}} \end{aligned}$$

We first prove implication (1). Let Φ be a circuit from $\mathsf{QCC}_{\mathsf{PauliProd}}$ and consider the action of $\Phi \otimes \mathcal{I}_c$ (where the Identity operator acts on the inaccessible halves of the EPR pairs) on the state $\frac{1}{2^c} |0\rangle \langle 0| \sum_{i,j} |i\rangle \langle j| \otimes |i\rangle \langle j|$:

$$\Phi \otimes \mathcal{I}_c \left(\frac{1}{2^c} |0\rangle \langle 0| \sum_{i,j} |i\rangle \langle j| \otimes |i\rangle \langle j| \right) = \frac{1}{2^c} \sum_{i,j} \Phi\Big(|0\rangle \langle 0| \otimes |i\rangle \langle j| \Big) \otimes |i\rangle \langle j| ,$$

because Φ being a channel is a linear operator. By Lemma 27, $|i\rangle \langle j|$ can be written as a linear combination of rank-1 projectors onto $\mathsf{PauliProd}$ states. Since Φ is catalytic with respect to $\mathsf{PauliProd}$, it follows that

$$\frac{1}{2^c} \sum_{i,j} \Phi\Big(|0\rangle \langle 0| \otimes |i\rangle \langle j| \Big) \otimes |i\rangle \langle j| = \eta \otimes \frac{1}{2^c} \sum_{i,j} |i\rangle \langle j| \otimes |i\rangle \langle j| ,$$

for some state in $\eta \in D(\mathcal{H}_s)$. This shows that $\Phi \in \mathsf{QCC}_{\mathsf{EPR}}$.

Implication (2) requires a similar approach. Let $\tilde{\Phi} \in \mathsf{QCC}_{\mathsf{EPR}}$, then we can write $\tilde{\Phi} = \Phi \otimes \mathcal{I}_c$. For a given input state $|0\rangle \langle 0| \in \mathcal{H}_s$ the action of $\Phi \otimes \mathcal{I}_c$ must satisfy

$$\Phi \left(\frac{1}{2^c} \sum_{i,j} |0\rangle \langle 0| \otimes |i\rangle \langle j| \right) \otimes |i\rangle \langle j| = \eta \otimes \frac{1}{2^c} \sum_{i,j} |i\rangle \langle j| \otimes |i\rangle \langle j| ,$$

for some state in $\eta \in D(\mathcal{H}_s)$. Since the catalytic state of c EPR pairs is returned perfectly unaffected for every choice of input state, the effective channel of Φ can also be written as a tensor product channel: $\Phi = \Gamma_s \otimes \Xi_c{}^3$, with the action of Ξ_c being

$$\frac{1}{2^c} \sum_{i,j} \Xi_c \Big(|i\rangle \langle j| \Big) \otimes |i\rangle \langle j| = \frac{1}{2^c} \sum_{i,j} |i\rangle \langle j| \otimes |i\rangle \langle j| \ .$$

[3] It seems that the catalyst does not offer any improvement, because we can write Φ as a tensor product of the action on the logspace clean qubits and the action of the catalyst, however this does not need to hold. Only the action as a whole is writable as a tensor product, it might actually consist of intermediate steps that are not of tensor product form, therefor Γ_s might only have an efficient circuit description in the presence of a catalyst.

Note that although the effective channel factorises into a tensor product across the work and catalytic registers, without the catalytic tape much larger circuits may be required to implement Γ_c. Moving forward, this implies that the Choi matrix of Ξ_c is

$$J(\Xi_c) = \sum_{i,j} \Xi_c\Big(|i\rangle\langle j|\Big) \otimes |i\rangle\langle j| = \sum_{i,j} |i\rangle\langle j| \otimes |i\rangle\langle j| = J(\mathcal{I}),$$

and therefore the effective channel Ξ_c is the identity channel. This gives that for any state $\rho \in \mathcal{H}_c$ it must hold that on input $|0\rangle\langle 0|$, the channel Φ must act as follows:

$$\Phi(|0\rangle\langle 0| \otimes \rho) = \eta \otimes \rho \qquad ◀$$

▶ Remark 29. In the proof that these channel definitions are equivalent we actually showed that any channel under one definition also furnishes an instance of the other definitions. This means that they are also operationally equivalent. These equivalence proofs therefore have to hold for any type of machine model that has to adhere to the same restrictions of resetting the input state in the catalytic space. In particular it also holds for quantum Turing machines.

3.3 Gateset

When discussing quantum circuits, a fundamental issue is the underlying gate set. Unlike the classical case, unitary operations form a continuous space, and finite-sized circuits over finite gate sets cannot implement arbitrary unitaries. However, there do exist finite gate sets of constant locality (that is, fan-in) which are quantum universal, in the sense that any n-qubit unitary may be approximated to any desired precision ϵ in ℓ_2-distance by a product of $l = O(\mathsf{poly}\log\frac{1}{\epsilon})$ gates from the universal gate set; this is the celebrated Solovay-Kitaev theorem [27, 16, 32]. From the standpoint of complexity classes, Nishimura and Ozawa [34] also showed that polynomial-time quantum Turing machines are exactly equivalent to finitely generated uniform quantum circuits.

We note that Definitions 15 and 22 do not make reference any fixed universal gate set. A potential issue that arises in this regard is that the complexity class being defined may depend in an intricate way on the chosen universal gate set, since it may not be possible to perfectly reset every initial catalytic state under our uniformity and resource constraints. If we relax the notion of catalyticity to mean that the initial catalytic state only has to be reset to within ϵ trace distance at the end of the computation, one can use the Solovay-Kitaev theorem to see that every choice of gate set leads to the same complexity class in Definition 22. This interesting model resembles classical catalytic space classes with small errors in resetting, and we leave it as an open question to determine how it relates to the exact resetting model.

Returning to our setting that requires the quantum catalytic machine to perfectly reset the catalytic space to its initial state at the end of the computation, we will restrict out attention to the case of universal quantum gate sets that are infinite (for the complexity theoretic properties of circuit families over such gate sets, see e.g. [33]). In this case, our definition is robust to the choice of gate set since any unitary may be implemented exactly by finite-sized circuits over such gate sets. Consequently changing the gate set does not change the set of catalytic states that can be reset exactly by the machine. This results in well-defined catalytic complexity classes independent of the specific choice of gate set.

3.4 Uniformity

Similar to gatesets, the question of uniformity is quite subjective, as different uniformity conditions will lead to different levels of expressiveness for our machines.

▶ **Definition 30.** *We fix the uniformity in Definition 18 to be $\mathcal{K} = \mathsf{SPACE}[O(s)]$.*

We choose $\mathsf{SPACE}[O(s)]$ as it is the largest class of classical machines a $\mathsf{QCSPACE}[s, c]$ machine should seemingly contain by default. Thus we believe the choice of $\mathsf{SPACE}[O(s)]$-uniformity is best suited to removing classical uniformity considerations from taking the forefront of the discussion regarding quantum catalytic space.

The question of how uniformity affects the power of QCSPACE is left to future work; we only comment briefly here on natural alternative choices. Perhaps the most immediate would be to consider $\mathsf{CSPACE}[s, c]$ uniformity, as it mirrors our quantum machine. As we will see later, it is not clear how to prove $\mathsf{QCSPACE}[s, c]$ contains $\mathsf{CSPACE}[s, c]$ directly, an interesting technical challenge that would be rendered moot by building it into the uniformity. Similarly we avoid P-uniformity because it is not known, and even strongly disbelieved, that CL contains P.

4 QCL upper bounds

In this section we will finally return to the question of our quantum machine model, showing that Turing machines and circuits are equivalent. One major stepping stone is to show that quantum catalytic Turing machines adhere to a polynomial runtime bound for *all* possible initializations of the catalytic tape.

Before all else, a remark is in order as to why such a restriction should hold for a seemingly stronger model, i.e. QCLM, when it is not in fact known for CL. While quantum catalytic space has access to more powerful computations, i.e. quantum operations, it also has the much stronger restriction of resetting arbitrary density matrices rather than arbitrary bit strings. This restriction gives rise to a much stronger upper bound argument, and in fact rules out one of the main techniques available to classical Turing machines, namely compression arguments (see c.f. [18, 12]).

4.1 Polynomial average runtime bound

We begin by showing an analogue of the classical result of [10], i.e. the average runtime of a quantum catalytic machine for a random initial catalytic state ρ is polynomial in the number of work qubits. We note that the runtime of a quantum Turing machine need not be a deterministic function of the input; M has access to quantum states and intermediate measurements, from which it is possible to generate randomness which might influence the time that machine takes to halt.

▶ **Definition 31.** *Given a quantum catalytic Turing machine M, a fixed input $x \in \{0, 1\}^n$, and an initial catalytic tape ρ, we denote by $T(M, x, \rho)$ the distribution of runtimes of M on input x and initial catalytic tape ρ.*

For an averaging argument to hold, we need to have a quantum notion of non-overlapping configuration graphs.

▶ **Lemma 32.** *Let M be a quantum catalytic Turing machine, and let $\{\tau_i\}_i$ form an orthonormal basis for $D(\mathcal{H}_c)$. For all i and t, let $\rho_{i,t}$ be the density matrix describing the state of the classical tape, quantum tape, and internal state of M at time step t on initial catalytic tape τ_i. Then if M is absolutely halting, all elements of the set $\{\rho_{i,t}\}_{i,t}$ are orthogonal.*

Proof. We first consider the states $\rho_{i,t}$ for a fixed i. Assume instead that there exists some times t and t' where the states are not orthogonal. This means that the state at time step t can be written as a superposition between the state in time step t' and the state

$\rho_{i,t} = p\rho_{i,t'} + (1-p)\eta$ for some $p > 0$. This forms a loop in the configuration graph where part of the state is back at time step t'. The amplitude of the part of the state in this loop will shrink over time, but never go to zero. The part of the state that is stuck in the loop will never reach the halting state, therefore this is in contradiction with the assumption that the quantum Turing machine is absolutely halting.

Next we consider the states $\rho_{i,t}$ for different i. By definition of a quantum Turing machine, the transformations M can apply to the entire state of the machine is given by some quantum channel. By Lemma 9 we know that the trace distance between the entire state of the machine for separate instances of the catalytic tape can only decrease by this quantum channel. Therefore we know that if two instances start out to be orthogonal and end to be orthogonal, they have to remain orthogonal through the entire calculation. ◀

▶ **Lemma 33.** *Let M be a quantum catalytic Turing machine with work space s and catalytic space c, let $\{\rho_i\}_i$ form an orthonormal basis for $D(\mathcal{H}_c)$, and define $T_{max}(M, x, \rho)$ to be the maximum runtime of machine M on input x on starting catalytic tape ρ. Then*

$$\mathbb{E}_i[T_{max}(M, x, \rho_i)] \leq 2^{O(s)}$$

Proof. Our catalytic machine is defined by a $\mathsf{SPACE}[O(s)]$ machine, defined by a tape of length $O(s)$ and an internal machine of size $O(1)$, which acts on $\mathcal{H}_s$ and $\mathcal{H}_c$, which can be addressed into using $\log s$ and $\log c$ bits respectively. Since these quantities plus the Hilbert spaces $\mathcal{H}_s$ and $\mathcal{H}_c$ define the dimensionality of our machine, by Lemma 32 we have that

$$\sum_{\rho \in \{\rho_i\}} T_{max}(M, x, \rho) \leq \mathcal{O}(2^{2(s+c+O(s)+O(1)+\log s+\log c)})$$

and therefore the lemma follows because $|\{\rho_i\}| \leq 2^{2c}$ and $2(s + O(s) + O(1) + \log s + \log c) = O(s)$. ◀

This already gives us a nice containment for our $\mathsf{QCSPACE}[s, c]$ classes.

▶ **Corollary 34.** $\mathsf{QCLM} \subseteq \mathsf{ZQP}$

▶ **Corollary 35.** $\mathsf{BQCLM} \subseteq \mathsf{BQP}$

4.2 Equal running times

We now take a further leap, showing that the initial catalytic tape does not affect the (distribution of the) runtime of our machine M for a fixed input x.

We can first show that given M and only one single copy of a state $\eta \in \mathcal{H}_c$, this probability distribution can be approximated up to arbitrary precision for any x.

▶ **Lemma 36.** *Given catalytic Turing machine M and a single copy of a quantum state $\eta \in \mathcal{H}_c$, $T(M, x, \eta)$ can be approximated up to arbitrary precision for any x.*

Proof. Because M is a quantum catalytic Turing machine it has to reset the quantum state initialized in its catalytic tape perfectly. Therefore we can use the following approach: first fix some input x, then run the catalytic machine given x as input and η on its catalytic tape and record the running time. When the machine halts, η should be returned in the catalytic tape. This means the test can be performed again given the same inputs. This test can be run arbitrarily often giving an arbitrary approximation to $T(M, x, \eta)$. ◀

This gives us the following observation about states with different halting times:

▶ **Lemma 37.** *Let M be a quantum catalytic Turing machine, and let $\rho_1, \rho_2 \in D(\mathcal{H}_c)$. Assume there exists $x \in \{0,1\}^n$ such that $T(M, x, \rho_1) \neq T(M, x, \rho_2)$. Then $||\rho_1 - \rho_2||_1 = 1$, where $||\cdot||_1$ is the trace distance.*

Proof. The Helstrom bound states that the optimal success probability of any state discrimination protocol given one copy of an unnown state is:

$$P_{success} = \frac{1}{2} + \frac{1}{2} \cdot ||\rho_1 - \rho_2||_1$$

By Lemma 36, we know that $T(M, x, \rho)$ can be approximated to any precision with only one copy of ρ. Given a copy of either ρ_1 or ρ_2 at random, one can estimate $T(M, x, \rho)$ and perfectly discriminate between the cases $\rho = \rho_1$ and $\rho = \rho_2$ giving a protocol with $P_{success} = 1$. Therefore it follows that

$$\frac{1}{2} + \frac{1}{2}||\rho_1 - \rho_2||_1 = 1$$

and hence $||\rho_1 - \rho_2||_1 = 1$. ◀

Lemma 37 is sufficient to show that the halting time of a quantum catalytic Turing machine is independent of the initial state in the catalytic tape:

▶ **Theorem 38.** *Let M be a quantum catalytic Turing machine with s-qubit work space and c-qubit catalytic space, and let $x \in \{0,1\}^n$. Then there exists some value $t := t(n)$ such that $T(M, x, \rho) = t$ for all $\rho \in D(\mathcal{H}_c)$.*

Proof. Assume for contradiction that there exist ρ_1, ρ_2 such that $T(M, x, \rho_1) \neq T(M, x, \rho_2)$. By Lemma 37 it holds that $||\rho_1 - \rho_2||_1 = 1$. Consider the state $\rho' = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$, and note that only one of $T(M, x, \rho') = T(M, x, \rho_1)$ or $T(M, x, \rho') = T(M, x, \rho_2)$ can hold, by transitivity. Without loss of generality, let us assume $T(M, x, \rho') = T(M, x, \rho_2)$, thereby $T(M, x, \rho') \neq T(M, x, \rho_1)$ and so $||\rho' - \rho_1||_1 = 1$ by Lemma 37. However, by definition we have that

$$||\rho' - \rho_1||_1 = ||(\frac{1}{2}\rho_1 + \frac{1}{2}\rho_2) - \rho_1||_1 = \frac{1}{2}$$

which is a contradiction. ◀

Putting Lemma 33 and Theorem 38 together immediately shows that the runtime of M is bounded by a polynomial in n for every input x and initial catalytic state ρ:

▶ **Theorem 39.** *Let M be a quantum catalytic Turing machine with work space s and catalytic space c. Then the maximum halting time is bounded by $2^{\mathcal{O}(s)}$.*

This strengthens Corollary 34 to remove the randomness in the output probability; this is the quantum equivalent of showing CL ∈ P, considered the holy grail of open problems in classical catalytic computing:

▶ **Corollary 40.** QCLM ⊆ EQP

4.3 Turing machines and circuits

We finally prove Theorem 23 and show the equivalence of our two definitions of quantum catalytic machines. To do this, we observe, without proof, that Theorem 38 extends to any *classical observable feature* of the initial catalytic state by the same proof. We will apply this to one other aspect, namely the transition applied at a given timestep t:

▶ **Lemma 41.** *Let M be a quantum catalytic Turing machine, and let $x \in \{0,1\}^n$. Then for every time t, there exists a fixed operation g applied by M at time t for every $\rho \in \mathcal{H}_c$.*

This is sufficient to prove Theorem 23:

Proof of Theorem 23. We only prove the equivalence between QCSPACEC and QCSPACEM; the same proof applies to BQCSPACEC and BQCSPACEM. Certainly QCSPACEC$[s,c]$ is contained in QCSPACEM$[O(s),O(c)]$, since QCSPACEC circuits are SPACE$[O(s)]$ uniform and can be directly simulated by a QCSPACEM machine.

Conversely, given a QCSPACEM$[s,c]$ machine M, we wish to find an equivalent quantum catalytic circuit in QCSPACEC$[O(s),O(c)]$. For this, we transform the transition function of the quantum Turing machine into a quantum channel; since the transition only takes a finite number of (qu)bits as input, this can be always be done, and we have our transitions act on the same space $\mathcal{H}_s \otimes \mathcal{H}_c$ as M. Then, by using a method similar to that from the proof of Lemma 56, to make the machine oblivious, the tape head movement of the quantum Turing machine will be fixed. If our circuit is the transition function channel copied to all locations where the tape heads end up, we completely simulate the quantum Turing machine. We know that $T_{max}(M,x,\rho)$ is always at most $2^{O(s)}$ for a machine M by Theorem 39, and so the number of such transition function channels is also at most $2^{O(s)}$. Therefore, we can simulate M using a quantum circuit of length $2^{O(s)}$ as claimed. ◀

As an afterword, we also resolve one other aspect of our initial definition of quantum catalytic space, namely the requirement that the output state be the same for every initial catalytic state. As mentioned above, Lemma 36 extends to all classically observable characteristics, but a similar argument clearly holds for approximating the output state as well:

▶ **Lemma 42.** *Given catalytic Turing machine M and a single copy of a quantum state $\eta \in \mathcal{H}_c$, the output qubit $|\phi_{out}\rangle$ can be approximated up to arbitrary precision for any x.*

Thus again we can appeal to the inststinguishability of nearby catalytic states to claim that $|\phi_{out}\rangle$ must be equal for all inital $|\tau\rangle$.

References

1 Aryan Agarwala and Ian Mertz. Bipartite matching is in catalytic logspace. *Electron. Colloquium Comput. Complex.*, TR25-048, 2025. URL: https://eccc.weizmann.ac.il/report/2025/048/.

2 Yaroslav Alekseev, Yuval Filmus, Ian Mertz, Alexander Smal, and Antoine Vinciguerra. Catalytic computing and register programs beyond log-depth. *Electron. Colloquium Comput. Complex.*, TR25-055, 2025. URL: https://eccc.weizmann.ac.il/report/2025/055/.

3 Noga Alon. Problems and results in extremal combinatorics—i. *Discrete Mathematics*, 273(1):31–53, 2003. EuroComb'01. doi:10.1016/S0012-365X(03)00227-9.

4 Matthew Amy, Matthew Crawford, Andrew N Glaudell, Melissa L Macasieb, Samuel S Mendelson, and Neil J Ross. Catalytic embeddings of quantum circuits. *arXiv preprint arXiv:2305.07720*, 2023.

5 David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1. *Journal of Computer and System Sciences (J.CSS)*, 38(1):150–164, 1989. doi:10.1016/0022-0000(89)90037-8.

6 Michael Ben-Or and Richard Cleve. Computing algebraic formulas using a constant number of registers. *SIAM Journal on Computing (SICOMP)*, 21(1):54–58, 1992. doi:10.1137/0221006.

7 Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. doi:10.1137/s0097539796300921.

8 Sagar Bisoyi, Krishnamoorthy Dinesh, Bhabya Rai, and Jayalal Sarma. Almost-catalytic computation. *CoRR*, abs/2409.07208, 2024. doi:10.48550/arXiv.2409.07208.

9 Sagar Bisoyi, Krishnamoorthy Dinesh, and Jayalal Sarma. On pure space vs catalytic space. *Theor. Comput. Sci.*, 921:112–126, 2022. doi:10.1016/J.TCS.2022.04.005.

10 Harry Buhrman, Richard Cleve, Michal Koucký, Bruno Loff, and Florian Speelman. Computing with a full memory: Catalytic space. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pages 857–866, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2591796.2591874.

11 Harry Buhrman, Michal Koucký, Bruno Loff, and Florian Speelman. Catalytic space: Nondeterminism and hierarchy. *Theory Comput. Syst.*, 62(1):116–135, 2018. doi:10.1007/S00224-017-9784-7.

12 James Cook, Jiatu Li, Ian Mertz, and Edward Pyne. The structure of catalytic space: Capturing randomness and time via compression. In *ACM Symposium on Theory of Computing (STOC)*, 2025.

13 James Cook and Ian Mertz. Trading time and space in catalytic branching programs. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPIcs*, pages 8:1–8:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.CCC.2022.8.

14 James Cook and Ian Mertz. Tree evaluation is in space $O(\log n \cdot \log\log n)$. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 1268–1278, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649664.

15 Samir Datta, Chetan Gupta, Rahul Jain, Vimal Raj Sharma, and Raghunath Tewari. Randomized and symmetric catalytic computation. *Electron. Colloquium Comput. Complex.*, TR20-024, 2020. URL: https://eccc.weizmann.ac.il/report/2020/024.

16 Christopher M. Dawson and Michael A. Nielsen. The solovay-kitaev algorithm. *Quantum Info. Comput.*, 6(1):81–95, January 2006.

17 David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

18 Yfke Dulek. Catalytic space: on reversibility and multiple-access randomness. 2015.

19 Bill Fefferman and Zachary Remscrim. Eliminating intermediate measurements in space-bounded quantum computation. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1343–1356, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3406325.3451051.

20 Marten Folkertsma, Ian Mertz, Florian Speelman, and Quinten Tupker. Fully characterizing lossy catalytic computation. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, January 7-10, 2025, Columbia University, New York, NY, USA*, volume 325 of *LIPIcs*, pages 50:1–50:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.ITCS.2025.50.

21 Uma Girish, Ran Raz, and Wei Zhan. Quantum logspace algorithm for powering matrices with bounded norm. *arXiv preprint arXiv:2006.04880*, 2020.

22 Uma Girish, Ran Raz, and Wei Zhan. Quantum logspace computations are verifiable. In *2024 Symposium on Simplicity in Algorithms (SOSA)*, pages 144–150. SIAM, 2024.

23 Chetan Gupta, Rahul Jain, Vimal Raj Sharma, and Raghunath Tewari. Lossy catalytic computation. *Computing Research Repository (CoRR)*, abs/2408.14670, 2024.

24 Dustin G. Mixon (https://mathoverflow.net/users/29873/dustin-g mixon). How many non-orthogonal vectors fit into a complex vector space? MathOverflow. URL:https://mathoverflow.net/q/458508 (version: 2023-11-16). arXiv:https://mathoverflow.net/q/458508.

25 Neil Immerman. Nondeterministic space is closed under complementation. *SIAM Journal on computing*, 17(5):935–938, 1988.

26 Richard Jozsa, Barbara Kraus, Akimasa Miyake, and John Watrous. Matchgate and space-bounded quantum computations are equivalent. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2115):809–830, 2010.

27 A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, December 1997. doi:10.1070/rm1997v052n06abeh002155.

28 E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672–5675, December 1998. doi:10.1103/physrevlett.81.5672.

29 Michal Koucký, Ian Mertz, Ted Pyne, and Sasha Sami. Collapsing catalytic classes. *Electronic Colloquium on Computational Complexity (ECCC)*, TR25-018, 2025. URL: https://eccc.weizmann.ac.il/report/2025/018.

30 Patryk Lipka-Bartosik, Henrik Wilming, and Nelly HY Ng. Catalysis in quantum information theory. *Reviews of Modern Physics*, 96(2):025005, 2024.

31 Ian Mertz. Reusing space: Techniques and open problems. *Bulletin of the EATCS (B.EATCS)*, 141:57–106, 2023.

32 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2010. doi:10.1017/CBO9780511976667.

33 Harumichi Nishimura and Masanao Ozawa. Computational complexity of uniform quantum circuit families and quantum turing machines. *Theor. Comput. Sci.*, 276(1–2):147–181, April 2002. doi:10.1016/S0304-3975(01)00111-6.

34 Harumichi Nishimura and Masanao Ozawa. Perfect computational equivalence between quantum turing machines and finitely generated uniform quantum circuit families. *Quantum Information Processing*, 8(1):13–24, January 2009. doi:10.1007/s11128-008-0091-8.

35 Tetsuro Nishino. Mathematical models of quantum computation. *New Generation Computing*, 20(4):317–337, December 2002. doi:10.1007/bf03037370.

36 Aaron Potechin. A note on amortized branching program complexity. In Ryan O'Donnell, editor, *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, volume 79 of *LIPIcs*, pages 4:1–4:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICS.CCC.2017.4.

37 Edward Pyne, Nathan S. Sheffield, and William Wang. Catalytic communication. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, January 7-10, 2025, Columbia University, New York, NY, USA*, volume 325 of *LIPIcs*, pages 79:1–79:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.ITCS.2025.79.

38 Walter J Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.

39 Dan Shepherd. Computation with unitaries and one pure qubit, 2006. arXiv:quant-ph/0608132.

40 Peter W. Shor and Stephen P. Jordan. Estimating jones polynomials is a complete problem for one clean qubit, 2008. arXiv:0707.2831.

41 Róbert Szelepcsényi. The method of forced enumeration for nondeterministic automata. *Acta informatica*, 26:279–284, 1988.

42 Amnon Ta-Shma. Inverting well conditioned matrices in quantum logspace. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 881–890, 2013.

43 John Watrous. Quantum algorithms for solvable groups. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 60–67, 2001.

44 John Harrison Watrous. *Space-bounded quantum computation.* The University of Wisconsin-Madison, 1998.

45 Ryan Williams. Simulating time in square-root space. *Electron. Colloquium Comput. Complex.*, TR25-017, 2025. URL: https://eccc.weizmann.ac.il/report/2025/017/.

46 Andrew Chi-Chih Yao. Quantum circuit complexity. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 352–361. IEEE Computer Society, 1993. doi:10.1109/SFCS.1993.366852.

A Simulation of TC^1

In this section we show that QCL can simulate Boolean threshold circuits. As in the classical world, the ability to simulate TC^1 is also a reason to believe that catalytic logspace is strictly more powerful than logspace. This follows from the fact that $\mathsf{QL} = \mathsf{PL}$ [44], which is itself contained in TC^1:

▶ **Lemma 43.** $\mathsf{QL} \subseteq \mathsf{TC}^1$

Since TC^1 can compute powerful functions such as determinant, this containment is largely believed to be strict. Thus Theorem 3 gives us a candidate class of problems for separating QL from QCL.

A.1 Reversibility and obliviousness

In [10] the authors showed that TC^1 can be simulated by *transparent register programs*, which themselves are computable in CL; thus our goal is to extend the CL simulation of transparent programs to QCL. More broadly, we show that *reversible, oblivious, time-bounded* CL is enough to simulate transparent programs, and such a model is structured enough that, while we cannot show that all of CL is in QCL, we can at least prove the containment for this small fragment.

We first make the following definitions which we use for our simulations. We begin by recalling a result of Dulek [18] which shows that catalytic Turing machines can be made *reversible* (see c.f. [12] for a proof)

▶ **Theorem 44.** *For every catalytic machine M with space s and catalytic space c, there exist catalytic machines $M_{\rightarrow}$, $M_{\leftarrow}$ with space $s+1$ and catalytic space c such that for any pair of configurations (τ_1, v_1), (τ_2, v_2) of $M_{\rightarrow}$ and $M_{\leftarrow}$, if $M_{\rightarrow}$ transitions from (τ_1, v_1) to (τ_2, v_2) on input x, then $M_{\leftarrow}$ transitions from (τ_2, v_2) to (τ_1, v_1) on input x.*

We will also need to consider *oblivious* machines, i.e. ones where the tape head movement is solely a function of the input length $|x|$ and does not depend at all on the content of the catalytic tape c. While any Turing machine can be made oblivious, it requires relaxing the definition of obliviousness to not forcing the machine to halt at the same time on every input; we simply require that every machine that continues to run carries out its execution in an oblivious manner. We will bar this restriction in this section.

▶ **Definition 45.** *We say that a CL machine is totally oblivious if the following holds. Let t, q, h be special registers on the free work tape, all initialized to 0, representing the time, state, and tape heads of the machine. At each point in time our machine consist of one mega-step: for every setting of t, q, h there is a fixed transformation, computable in logspace, which the machine applies to the catalytic tape and to q, h, and a mega-step consists of applying each of these operations, conditioned on the values of t, q, h on the free work tape, in order. At the end of every mega-step we increment t, and our machine halts iff t reaches a predetermined step T.*

Totally oblivious machines are ones that in essence apply the same bundle of transformations at every time step, with the information about which one to to actually apply being written on the free work tape, and the halting behavior being determined only by the clock.

Such machines are clearly in poly-time bounded CL (see c.f. [12] for a discussion of this class), since the clock must fit on the free work tape. This causes issues when we seek total obliviousness in tandem with reversibility; in general it is not known, and is highly unlikely, that a polynomially time-bounded Turing machine can be made reversible while remaining polynomially time-bounded.

However, there is an important class of algorithms which is both reversible and totally oblivious: *clean register programs*. For our purposes we will use a very restricted version of clean register programs (see c.f. [31] for a discussion).

▶ **Definition 46.** *A register program $\mathcal{P}$ is a list of instructions $P_1 \dots P_t$ where each P_i either has the form $R_j \mathrel{+}= x_k$ for some input variable x_k or has the form $R_j \mathrel{+}= q_i(R_1 \dots R_m)$ for some polynomial q_i. A register program cleanly computes a value v if for any initial values $\tau_1 \dots \tau_m$, the net result of running $\mathcal{P}$ on the registers $R_1 \dots R_m$, where each R_j is initialized to the value τ_j, is that $R_1 = \tau_1 + v$ and $R_j = \tau_j$ for all $j \neq 1$.*

If we think of these registers as being written on the catalytic tape, it is clear that clean register programs are totally oblivious, as the instruction at every moment in time is based only on the timestep. This is nearly immediate, although we note a few minor complications here. We need to preprocess the catalytic tape to ensure our registers have values over the same ring as our register program; for example, if we represent numbers mod p using $\lceil \log p \rceil$ bits, some initial values will exceed p. This can be handled obliviously by observing that for either τ or $\overline{\tau}$, half the registers are already correct, and so we take one full pass over τ to keep a count of which case we are in, store this as a bit b (1 iff we need to flip τ), and XOR τ with b at the beginning and end of the computation. We subsequently ignore all blocks which are initialized to improper values; when we go to operate on register R_j, say, as we obliviously pass over the whole catalytic tape we will count how many *valid* registers we have seen, and act only when we see the counter reach j.

Besides being totally oblivious, however, such programs are also *reversible*, as every step of the form $R_j \mathrel{+}= c$ can be inverted by a step of the form $R_j - = c$. Thus such programs appear highly constrained in terms of what they can and cannot achieve. Nevertheless, such programs are sufficient to compute TC^1.

▶ **Lemma 47** ([10]). *Let L be a language in TC^1. Then L can be decided by a clean register program, and, hence, by a totally oblivious reversible* CL *machine.*

A.2 Simulation by QCL machines

We now show that reversibility plus total obliviousness is sufficient for simulation by QCL.

▶ **Lemma 48.** *Let L be a language which can be computed be a totally oblivious reversible* CL *machine. Then $L \subseteq$* QCL*.*

Proof. Let M be a totally oblivious reversible CL machine. We will treat our quantum catalytic tape as a superposition over classical catalytic tapes, i.e. a superposition over computational basis states. It is thus sufficient to show that the operation of machine M can be simulated by a fixed quantum circuit containing Toffoli gates, as such a circuit will correctly operate on each of our catalytic basis states in each branch of the superposition.

By total obliviousness, every step that M takes is a fixed transformation conditioned on the value of t, q, and h; since we additionally know that such a step is reversible, it must be isomorphic to a Toffoli gate applied to a fixed position of the catalytic tape conditioned on some fixed mask applied to t, q, and h, and furthermore each transformation can be computed by our logspace controlling machine. Since these operations are fixed for each timestep, we can move t to our space controlling machine and have it construct a circuit, comprised of Toffoli gates on q, h, and the catalytic tape, of polynomial length. ◀

This is sufficient to prove our main result for this section:

Proof of Theorem 3. Combine Lemma 47 with Lemma 48. ◀

B Simulating catalytic space in DQC$_1$

Lastly we will discuss the relationship between catalytic computing and a pre-existing yet closely related quantum model, namely the one clean qubit setting. We will introduce the model and then prove that it can simulate unitary QCL. In the full version of the paper, we further show that classical CL is also contained in the one clean qubit model.

B.1 One clean qubit model

In the one-clean qubit model, first introduced by Knill and Laflamme [28], a quantum machine is given a single input qubit initialized in the zero state and n qubits initialized in the maximally mixed state. We will formalize the definition of this computational model:

▶ **Definition 49** (One clean qubit). *Let $\{Q_x\}_x$ be a log-space uniform family of unitary quantum circuits. The* one clean qubit model *is a model of computation in which Q_x is applied to the $n+1$-qubit input state*

$$\rho = |0\rangle\langle 0| \otimes \frac{I_n}{2^n},$$

where $n = |x|$ and I_n operator is the identity on n qubits. After execution of Q_x the first qubit is measured, giving output probabilities:

$$p_0 = 2^{-n}\operatorname{Tr}\big[(|0\rangle\langle 0| \otimes I)Q_x(|0\rangle\langle 0| \otimes I)Q_x^\dagger\big],$$
$$p_1 = 1 - p_0$$

▶ Remark 50. Two points stand out in this definition. First, note that Q_x are unitary circuits, and hence do not allow intermediate measurements; such measurements would allow for resetting the qubits initialized in the maximally mixed state, making the model significantly stronger. Second, in this paper we consider log-space uniform families of unitary circuits, rather than the more common deterministic polynomial-time uniform families, in order to align more closely with the QCL model that we study.

The one-clean qubit model is a probabilistic model of computation, and hence we typically talk about computing a function $f(x)$ in terms of success probability for computing $f(x)$ being bounded away from 1/2. The exact bound on the error probability does not matter; while we often use 2/3 in defining e.g. BQP, even a 1/poly(n) gap is sufficient as there we can employ standard error-correction to boost our success, namely by running the algorithm multiple times. However, this is not known to be possible in the one-clean qubit model, as such a machine can only reliably run once.

▶ **Definition 51** ([28, 39]). DQC$_1$ *is the set of all languages $L = (L_{yes}, L_{no}) \subset \{0,1\}^* \times \{0,1\}^*$ for which there exists a one-clean qubit machine M and a polynomial $q(n)$ that on input $x \in L$ of length $n = |x|$,*

- *if $x \in L_{yes}$ then the output probability $p_1 \geq \frac{1}{2} + \frac{1}{q(n)}$*
- *if $x \in L_{no}$ then the output probability $p_0 \geq \frac{1}{2} + \frac{1}{q(n)}$*

On the other hand, somewhat surprisingly the one-clean qubit model is robust to the number of clean qubits allowed, up to a logarithmic number:

▶ **Lemma 52** ([40]). *$\mathsf{DQC}_k = \mathsf{DQC}_1$ for $k = \mathcal{O}(\log(n))$, where DQC_k means having access to k clean qubits instead of one.*

B.2 Containment of unitary QCL in DQC$_1$

We now move on to establishing a formal connection between QCL and DQC$_1$. A QCL machine is allowed to apply intermediate measurements to its quantum tape as well as its catalytic tape, which is not possible in DQC$_1$; however, if we restrict the QCL machine to not make any intermediate measurements we can show that such a machine can in fact be simulated by the one-clean qubit model.

▶ **Definition 53** (Q$_\mathsf{U}$CL)**.** *A Q$_\mathsf{U}$CL machine is a QCL machine in which the quantum circuit is unitary. In the final step of the unitary the Q$_\mathsf{U}$CL machine measures the first qubit, which then gives the outcome of the calculation. Similarly we define BQ$_\mathsf{U}$CL to be BQCL with the unitary restriction.*

Using this definition we can give the following proof of containment:

Proof of Theorem 4. Let C be a log-space uniform BQ$_\mathsf{U}$CL quantum channel. Since C is unitary up until the last measurement step, it preserves all possible density matrices from the catalytic tape, and in particular it preserves the maximally mixed state I_n. Let U be the unitary part of C. The action of U on the work-tape and the catalytic tape, with the catalytic tape initialized in I_n, is:

$$U\left|0\right\rangle\left\langle 0\right|_w \otimes \frac{I_n}{2^n}U^\dagger = (\sqrt{p_0}\left|0\right\rangle\left\langle 0\right|_{w_0}\left|\psi_0\right\rangle\left\langle\psi_0\right|_w + \sqrt{p_1}\left|1\right\rangle\left\langle 1\right|_{w_0}\left|\psi_1\right\rangle\left\langle\psi_1|_w\right.) \otimes \frac{I_n}{2^n}$$

with $|p_1| \geq 2/3$ in a "yes" instance and $|p_0| \geq 2/3$ in a "no" instance. Note that this calculation is of the exact form of a $\log(n)$-clean qubit machine and that the output probabilities are a constant bounded away from $1/2$; hence this problem is in DQC$_\mathsf{k}$, and by Lemma 52 is therefore in DQC$_1$ ◀

B.3 Containment of CL in DQC$_1$

We aim to show that CL $\subseteq$ DQC$_1$. The idea is that CL, as per Theorem 44, can always be made reversible. While as discussed before we cannot maintain reversibility and total obliviousness, a CL machine can also always be made "almost oblivious" while maintaining reversibility; the tape head movements are independent of the input, but the machine does not know when to halt. Instead, after any given amount of time, we know that the machine has halted on a fraction $1/\mathsf{poly}(n)$ of possible initial catalytic states. Since the DQC$_1$ model can be interpreted as sampling from a uniform distribution of computational basis states, this shows the probability of finding the correct output is $1/2 + 1/\mathsf{poly}(n)$, which is sufficient for the proof.

▶ **Definition 54.** *A* non-halting reversible oblivious *catalytic Turing machine is a reversible oblivious catalytic Turing machine that need not halt absolutely. In particular, for every input x and initial catalytic state c there exists a time $t(x,c)$ where the correct output has been written to the output tape and the catalytic tape has been reset to its initial state. In addition, the output state has an additional binary cell that indicates whether or not the output has been determined yet, or is still "unknown" by the machine.*

▶ **Definition 55.** *We say a reversible oblivious catalytic Turing machine* halts with polynomial success probability *if there exists polynomials p, q such that for any valid input x to a promise problem, after time $p(|x|)$ the output tape of the catalytic Turing machine contains the correct output to the problem on a fraction of at least $1/q(|x|)$ when the initial catalytic tapes are taken uniformily at random. After time $p(|x|)$, the output tape of the catalytic Turing machine never contains the wrong answer, but it may leave the output undetermined.*

We show that any CL machine can be transformed into a reversible oblivious catalytic Turing machine that halts with polynomial success probability. We defer the proof of this fact to the full version of the paper.

▶ **Lemma 56.** *Any catalytic Turing machine M that has a logarithmic clean space and polynomial size catalytic tape can be turned into a non-halting oblivious reversible catalytic Turing machine M^o with a logarithmic clean tape and polynomial catalytic tape.*

We call the machine formed this way M^o for oblivious M. Since the catalytic and clean tape are no more than polynomial length, this procedure adds at most a polynomial factor to the runtime. However, since the runtime of M may be super-polynomial and an oblivious machine has the same runtime for all inputs x of the same length and catalytic tapes c, the machine does not have enough clean space to keep a clock to know whether or not it has terminated. This means we cannot assume it to be halting. However, we can show that it is halting with sufficient probability (we again defer this proof to the full version of the paper):

▶ **Lemma 57.** *For any language L in CL that is recognized by a catalytic Turing machine M, there exists a reversible oblivious catalytic Turing machine N that halts with polynomial probability that also recognizes L. Furthermore, N also uses $O(\log|x|)$ clean space and polynomial catalytic space.*

This completes all technical components necessary to show that $\mathsf{CL} \subseteq \mathsf{DQC}_1$.

Proof of Theorem 5. The maximally mixed state of DQC_1 can be interpreted as uniformly randomly sampling computational basis states. If we take these basis states to be the catalytic tape and use the fact that DQC_1 is unchanged if we allow a logarithmic number of clean qubits, then we can run the machine N from Lemma 57 by using unitary gates instead of reversible, oblivious operations. When we measure the output bit at the end, we get either an indeterminate state or the correct output with certainty. If we get an indeterminate state, we output a random bit and thus output the correct answer with probability $1/2$. If not, then we output the correct answer, which occurs with probability at least $1/\mathsf{poly}(n)$. ◀

The Rotation-Invariant Hamiltonian Problem Is $\mathrm{QMA}_{\mathrm{EXP}}$-Complete

Jon Nelson
Joint Center for Quantum Information and Computer Science (QuICS),
Department of Computer Science, University of Maryland, College Park, MD, USA

Daniel Gottesman
Joint Center for Quantum Information and Computer Science (QuICS),
Department of Computer Science, University of Maryland, College Park, MD, USA

Abstract

In this work we study a variant of the local Hamiltonian problem where we restrict to Hamiltonians that live on a lattice and are invariant under translations and rotations of the lattice. In the one-dimensional case this problem is known to be $\mathrm{QMA}_{\mathrm{EXP}}$-complete. On the other hand, if we fix the lattice length then in the high-dimensional limit the ground state becomes unentangled due to arguments from mean-field theory. We take steps towards understanding this complexity spectrum by studying a problem that is intermediate between these two extremes. Namely, we consider the regime where the lattice dimension is arbitrary but fixed and the lattice length is scaled. We prove that this rotation-invariant Hamiltonian problem is $\mathrm{QMA}_{\mathrm{EXP}}$-complete answering an open question of [6]. This characterizes a broad parameter range in which these rotation-invariant Hamiltonians have high computational complexity.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Hamiltonian complexity, Local Hamiltonian problem, Monogamy of entanglement

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.12

Funding *Jon Nelson*: supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2236417.

1 Introduction

In order to understand the behavior of a quantum many-body system, it is crucial to study its Hamiltonian. The Hamiltonian operator not only governs the system's dynamics through the Schrödinger equation but also encodes its low-energy states and energy spectrum. It is thus important to understand which Hamiltonians are tractable to analyze. In this work, we study the computational complexity of estimating the ground-state energy of a Hamiltonian with only short-range interactions, which is known as the local Hamiltonian problem.

Often the goal is to show that a specific variant of this problem is QMA-complete, which implies that for certain Hamiltonians not even a quantum computer can be expected to find its ground state energy. On the one hand, this is a negative result for being able to calculate ground state energies. On the other hand, these results lead to constructions of highly complex quantum systems that are interesting objects of study in their own right.

Kitaev initiated the study of local Hamiltonian problems in his landmark result proving that this problem in its most general form is QMA-complete [7]. However, Kitaev's result only applies for a worst-case family of Hamiltonians, which are not physically natural. In order to study the complexity of more physically relevant cases, subsequent work has extended Kitaev's result to apply under additional constraints that capture what it means to be "natural". This has included restricting the local dimension of each particle as well as

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 12; pp. 12:1–12:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

constraining the geometry of the interactions. For example, [10] extends Kitaev's result to qubits on a 2D lattice while [2] further restricts to particles on a line with constant local dimension.

Additional follow-up work has also emphasized symmetry constraints. This is motivated by the observation that many systems in nature are highly symmetric. For instance, the laws of gravity, electromagnetism, etc., do not change depending on where you are or how you are oriented; thus, these laws are translation and rotation-invariant. For translation-invariant one-dimensional spin chains, [6] showed that this local Hamiltonian problem is $\mathrm{QMA_{EXP}}$-complete.

Although much work is now known about the complexity of translation-invariant systems [6, 4, 3, 8, 11], there have been very few results for the rotation-invariant case. In fact, it was posed as an open question of [6] whether their results can be extended to rotation-invariant Hamiltonians in higher dimensions. In this work, we solve this question and hope similar techniques can be used to lift other translation-invariant results to the rotation-invariant case.

Translation invariance with reflection symmetry and rotation invariance coincide in 1D and so the main challenge is to extend [6] to Hamiltonians on higher dimensional lattices. It is trivial to extend their result to higher dimensional translation-invariant lattices simply by ignoring all but one dimension. However, this breaks the rotation symmetry, which requires that the Hamiltonian terms act identically in all directions. In this case, the key challenge is to handle the increasingly high degree of interaction without increasing the number of parameters in the Hamiltonian. This presents issues, for example, when attempting to encode computation into the Hamiltonian's ground state, which is an essential step for proving hardness. Controlling this computation requires the ability to track time, which can be accomplished in the 1D setting by moving a clock pointer along the spin chain [6]. However, this same idea cannot be used in higher dimensions since the paths can branch in many directions throughout the lattice. In order to pick out a specific time direction, we must engineer a family of Hamiltonians that spontaneously breaks the rotation symmetry.

Technical difficulties aside, it may seem intuitive that increasing the lattice dimension only makes the local Hamiltonian problem more difficult, and so one might assume that the complexity for higher dimensional cases follows from the one-dimensional case. However, due to the rotation symmetry, the increase in lattice dimension does not correspond to more Hamiltonian parameters, and so it is unclear how the complexity actually compares. In fact, standard condensed matter arguments imply that increasing the dimension can instead make the problem easier. This follows from the observation that for higher-dimensional lattices, mean-field theory (which uses a product-state ansatz to approximate the ground state) becomes more and more accurate [12]. In the quantum setting, this can be explained by an effect called monogamy of entanglement [14], which states that a particle cannot be highly entangled with many other particles. Thus, for high lattice dimension, each particle has many neighbors and so on average they must be nearly unentangled. Due to this effect, the product state becomes a good approximation of the ground state, suggesting that this problem could now be more tractable than the lower dimensional cases.

This has been formalized in [5] which shows that for lattice dimension r there is a product state that approximates the ground-state energy by an average error of $O(r^{-1/3})$ per term of the Hamiltonian. Furthermore, [9] rigorously show that in the limit as $r \to \infty$ the ground state is exactly a product state when the Hamiltonian is translation and rotation invariant. These results suggest that if the lattice dimension is high enough, the problem loses its quantum hardness since the low-energy states become unentangled. Another result that captures this phenomenon is [1], who show that a commuting version of the local Hamiltonian problem becomes easier as the interaction graphs become more expanding, which intuitively corresponds to more interaction.

In this work, we consider a lattice dimension that is in an intermediate regime between one-dimensional spin chains, which are hard and spin chains with $r \to \infty$, which are easy. In particular, we consider an arbitrary but fixed lattice dimension and show that this rotation-invariant Hamiltonian problem is quantumly hard as you scale the lattice length.

1.1 Results

We informally describe the rotation-invariant Hamiltonian problem as follows.

▶ **Definition 1** (Rotation-invariant Hamiltonian problem (Informal)). *Consider the Hamiltonian where a single two-body term is applied to each neighboring pair of qudits on an r-dimensional lattice of side length n. Is the ground-state energy below a or greater than b?*

Our main result is that the rotation-invariant Hamiltonian problem is $\mathrm{QMA}_{\mathrm{EXP}}$-complete, where $\mathrm{QMA}_{\mathrm{EXP}}$ is the same as QMA except the witness and verification circuit are allowed to be exponentially large in the input size. The reason we consider $\mathrm{QMA}_{\mathrm{EXP}}$ rather than QMA is that the input size to our problem is actually very small in comparison to the size of the Hamiltonian. To see this, notice that our Hamiltonian can be completely described by 1) the two-body term, 2) the dimension r, and 3) the lattice length n. The first two of these only require a constant number of bits to specify, while n requires $\log n$ bits to specify. Therefore, the Hamiltonian description length is exponentially smaller than the total number of qudits. However, we still would like to allow an "efficient" algorithm to run in polynomial time with respect to the number of qudits, which in turn is exponential in the input size. To accommodate this technicality, we must prove quantum hardness even when the quantum computer is allowed an exponential amount of computation time.

To prove $\mathrm{QMA}_{\mathrm{EXP}}$-completeness, we use the standard method of reducing an instance x of an arbitrary $\mathrm{QMA}_{\mathrm{EXP}}$ problem to an instance $R(x)$ of the rotation-invariant Hamiltonian problem. It turns out that in this reduction only n depends on the original problem instance x, so we take everything else (such as the two-body term and the lattice dimension) to be parameters rather than inputs to the problem. This is described in the more technical definition of the problem in Section 2. This differs from the standard QMA-completeness result, where the Hamiltonian terms themselves are given as input. We argue that this is a more natural setting since often one is studying a particular Hamiltonian, and so it is more suitable to consider the hardness of a given Hamiltonian for increasingly large system sizes. A desirable feature of our reduction is that the Hamiltonian we construct has no dependence on the system size or lattice dimension.

The $\mathrm{QMA}_{\mathrm{EXP}}$-completeness of this problem has a number of interesting implications. First, it suggests that not even a quantum computer can find the ground-state energy of certain rotation-invariant Hamiltonians. Since nature can be viewed as a quantum computer this means that the system itself cannot find its own ground state either, suggesting the emergence of spin glass behavior at low temperatures. Next, our result implies (assuming $\mathrm{QCMA}_{\mathrm{EXP}} \neq \mathrm{QMA}_{\mathrm{EXP}}$) that the ground state of these Hamiltonians cannot have an efficient classical description and thus cannot be well approximated by a product state. In fact, our result directly implies a lower bound of $\Omega(n^{-r}r^{-1})$ for how close the average ground-state energy per term can be approximated by a product state. This complements the result in [5] by providing a corresponding lower bound for the product-state approximation error.

1.2 Techniques

Our main approach is to carve out one-dimensional chains within our higher dimensional lattice so that we can apply [6]'s 1D construction to these chains. To do this, we take inspiration from the closely related classical problem of tiling. In this problem, imagine fitting together square tiles to cover an entire floor, where we are given a penalty for placing certain color tiles next to each other. The task is now to design a set of tiling rules where the least penalized configuration is a pattern of stripes. If such a set of tiling rules exists, we can use a portion of each qudit's Hilbert space to represent a tile and encode the tiling rules in our two-body Hamiltonian term. This enforces that the tiling of the ground state will have this striped pattern. Our construction then proceeds by applying the 1D construction onto neighboring qudits with same-colored tiles, which are now effectively spin chains.

Unfortunately, such a set of tiling rules does not actually exist, and so we will have to modify this classical technique to incorporate some quantum phenomena. To see this, consider the 3D case with periodic boundary conditions. No matter what set of rules are given, the optimal tiling will always take the following form. Start by tiling the first column of the first 2D slice with the optimal 1D configuration. Then, for each subsequent column in this slice, tile it by offsetting this sequence by exactly one. Finally for each subsequent 2D slice, tile by shifting the entire previous 2D configuration by one. With some thought, one can convince themselves that this is the correct tiling. The issue is that this configuration cuts the 3D lattice into diagonal planes which is not the desired 1D structure. Additionally, this argument also shows that these rotation-invariant tiling problems are in P whereas their translation-invariant (but not rotation-invariant) counterparts are NEXP-complete for dimension 2 and higher [6]. This further shows how the additional symmetry constraints can potentially simplify the complexity.

It is possible to still achieve our tiling goal by combining it with some purely quantum effects. In particular, we introduce a new technique that uses the monogamy of entanglement to enforce an effective 1D geometry. This is performed by first appending two qubits to the Hilbert space of each particle. The key idea is to enforce that same-colored neighbors share an EPR pair among their qubits. Since each site only has two qubits, it can only share an EPR pair with two neighbors by the monogamy of entanglement. Thus, it can only have two same-colored neighbors. This is already close to our goal, since now our lattice must be colored by disjoint same-colored loops. It remains to make sure that these loops do not have any turns but instead cut straight across the lattice. This can be handled by imposing some further classical tiling constraints, which we discuss in more detail in Section 3. We hope that this EPR pair technique can also find use in other Hamiltonian complexity problems that benefit from embedding lower dimensional geometries into higher ones.

One last technicality to resolve is that [6]'s 1D Hamiltonian is frustrated and has an energy of at least 1/2. This results in an overall energy of $n^{r-1}/2$ when this Hamiltonian is embedded into n^{r-1} 1D lines in the lattice. In order to balance this energy penalty with the rest of the Hamiltonian terms, it is necessary to normalize this contribution with a coefficient that depends on n. However, such a system size dependence is unnatural and is preferably avoided. To fix this, we first embed a 2D translation-invariant Hamiltonian into the lattice by encoding stripes in two different directions as opposed to just one. In this case, the same result can be achieved as in the 1D case except the construction can now be made to be nearly frustration-free, removing the need for a system-size dependent normalization term.

1.3 Outline

We begin by describing our notation and briefly state the technical version of our result in Section 2. Next, we define the Hamiltonian construction in Section 3. Finally, the proof of our main theorem is presented in Section 4 where it is shown that our construction satisfies both completeness and soundness.

2 Notation and technical result

Define $\Lambda_r(n) := \mathbb{Z}^r/n\mathbb{Z}^r$ to be a periodic lattice. In other words, the lattice is r-dimensional, where each dimension has length n and each site is denoted by integer coordinates. For a lattice point $u \in \Lambda_r(n)$ we denote the ith coordinate as u_i. To define distance between points in the lattice while respecting the periodic boundary conditions, we use the Lee metric, which is defined as follows:

▶ **Definition 2** (Lee metric). *Let $x, y \in \Lambda_r(n)$.*

$$d(x,y) = \sum_{i=1}^{r} \min(|x_i - y_i|, n - |x_i - y_i|)$$

The set of nearest-neighbor pairs is defined as $E_{\Lambda_r(n)} = \{\{x,y\} : x, y \in \Lambda_r(n), d(x,y) = 1\}$. If h is a 2-local Hamiltonian term and $u, v \in \Lambda_r(n)$ then $h^{u,v}$ denotes the Hamiltonian term h applied to sites u and v. For an operator H, we denote the lowest eigenvalue of H by $E_0(H)$.

▶ **Definition 3** (QMA$_{\text{EXP}}$). *A language L is in QMA_{EXP} if there exists a quantum verifier V such that on input x, V has runtime $O(2^{|x|^k})$ for some k. In addition, if $x \in L$ then there exists a state $|\psi\rangle \in \mathbb{C}^{O(2^{|x|^k})}$ such that $V(x, |\psi\rangle)$ accepts with probability at least $2/3$. If $x \notin L$, then $V(x, |\psi\rangle)$ accepts with probability at most $1/3$.*

With this notation in hand, the formal definition of the rotation-invariant Hamiltonian problem can be stated as follows:

▶ **Definition 4** (r-DIM-RIH (Rotationally-Invariant Hamiltonian)).

Problem Parameter: *The geometric dimension of the lattice r. A permutation-invariant Hermitian operator h. Two polynomials p and q.*

Input: *Integer n specified in binary.*

Promise: *Let $N = |\Lambda_r(n)| = n^r$. Consider the Hamiltonian $H = \sum_{\{u,v\} \in E_{\Lambda_r(n)}} h^{u,v}$. The ground state energy of H is either at most $p(n)$ or at least $p(n) + 1/q(n)$.*

Output: *Determine whether the ground-state energy of H is at most $p(n)$ or at least $p(n) + 1/q(n)$.*

In particular, notice that h does not depend on the system size n or the lattice dimension r in this definition. Our main result is the following theorem.

▶ **Theorem 5.** *r-DIM-RIH is QMA_{EXP}-complete for $q(n) = 1$.*

3 Hamiltonian construction

Our strategy will be first to embed a two-dimensional translation-invariant Hamiltonian without reflection symmetry into our *rotationally*-invariant Hamiltonian. In the case where our lattice has dimension r, we will break up the lattice into n^{r-2} 2D slices and embed the Hamiltonian into each of these slices. Then we can utilize the extra parameters of the 2D translation-invariant Hamiltonian to embed a one-dimensional hard Hamiltonian that is nearly frustration-free.

To accomplish this, it is crucial to have a mechanism to break the rotation symmetry by selecting a direction. We will do this by embedding two sets of directed stripes to indicate the two directions of the 2D grid.

3.1 Embedding directed stripes

In our construction, we will attach the following Hilbert spaces to each site in the lattice:

$$\mathcal{H}_{T_1} \otimes \mathcal{H}_{T_2}$$

Our Hamiltonian will act diagonally in these subspaces and, therefore, it will only enforce classical constraints with respect to a given set of basis states, which we denote by the sets T_1 and T_2, respectively. We refer to these basis states as "tiles" that we can assign to each site.

We define $T_1 := \{\text{red}, \text{yellow}, \text{blue}\}$ and $T_2 := \{0, 1, 2\}$, which associate a color and number with each tile, respectively. It can be enforced that two tiles t_1 and t_2 cannot be placed next to each other by including the term $|t_1, t_2\rangle\langle t_1, t_2|$ in the Hamiltonian. In this way, we incorporate the following rules for which tiles are allowed to be placed next to each other:

1. If two neighboring tiles have the same color then they must have different numbers.
2. If two neighboring tiles have different colors then they must have the same number.

To write down the Hamiltonian terms associated with these rules more explicitly, let V represent the set of illegal neighboring tiles. Then we include the following Hamiltonian term:

$$h_{\text{tile}} = 8 \sum_{(s,t)\in V} |s, t\rangle\langle s, t|$$

The energy cost of 8 is carefully chosen to balance out other competing terms introduced later.

3.1.1 EPR projections

Next, we would like to enforce that the qubits of same-colored neighbors form EPR pairs with each other. We can do this by attaching the following two additional Hilbert spaces to each site:

$$\mathcal{H}_{\sigma_1} \otimes \mathcal{H}_{\sigma_2}$$

Since each site only has two qubits it can only form two EPR pairs and therefore can only have two same-colored neighbors. Thus, this accomplishes our goal of forcing the same-colored tiles to form one-dimensional chains.

Given two same-colored neighbors u and v, it remains to determine which of their qubits must form EPR pairs. To do this, we first define directed edges between each basis state of T_2 such that $0 \to 1$, $1 \to 2$ and $2 \to 0$. Due to the constraints in the previous section, u and v must both have different numbers. Without loss of generality, if the directed edge between these numbers points towards v's tiles then we enforce that u's σ_2 qubit forms an EPR pair with v's σ_1 qubit. We denote this EPR pair state as $|\Phi^+\rangle_{u_{\sigma_2}v_{\sigma_1}} = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)_{u_{\sigma_2}v_{\sigma_1}}$. More formally we add the following term acting on sites u and v:

$$A^{u,v} = 16 \sum_{i\in T_2} |i, i+1 \bmod 3\rangle\langle i, i+1 \bmod 3| \otimes \left(\frac{\mathbb{I} - |\Phi^+\rangle\langle\Phi^+|}{2}\right)_{u_{\sigma_2}v_{\sigma_1}}$$

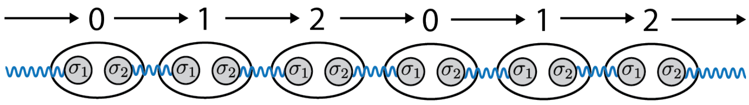

Figure 1 When a chain of sites is numbered sequentially around the cycle $\mathbb{Z}_3$, each qubit is matched with exactly one other qubit to form an EPR pair and so the EPR constraint can easily be satisfied.

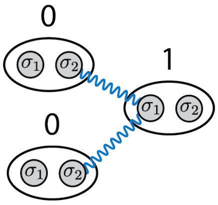

Figure 2 When there are three consecutive same-color neighbors that are not numbered monotonically around $\mathbb{Z}_3$ (for instance $0, 1, 0$) then two different qubits are matched with the same qubit to form an EPR pair. Due to the monogamy of entanglement this constraint cannot be satisfied and incurs an energy penalty.

Our convention throughout is to separate sites of the lattice by commas within the braket notation and to separate subspaces within each site by tensor product symbols. Notice that if u and v have different numbers then this implies they also have the same color and so this does not need to be additionally conditioned on in this Hamiltonian term. In order to preserve rotation invariance, we add this term for both orderings of the particles u and v:

$$h_{\mathrm{EPR}} = A^{u,v} + A^{v,u} \tag{1}$$

In addition to enforcing that contiguous regions of same-colored sites form 1D chains, these EPR constraints also require that each same-colored chain is numbered as the periodic sequence: $0, 1, 2, 0, 1, 2, \ldots$ either in the forwards or backwards direction (see Figure 1). This is because there can never be a site where the directed edges incident to it are both pointing away or both pointing towards it. In either case, this requires one of that site's qubits to be in two different EPR pairs, which is not allowed by the monogamy of entanglement. This scenario is depicted in Figure 2. This sequential numbering, is very helpful because it defines a direction to each chain. Notice that such a numbering is only possible when n is a multiple of 3, so we will later incorporate this restriction into our hardness reduction.

3.1.2 1D chain boundary conditions

Given the current Hamiltonian terms, the same-colored sites can form chains with either open or periodic boundary conditions (lines or loops). It will be convenient later that the boundary conditions are fixed and so we add another term to enforce periodic boundary conditions. In particular, we define the following term acting on sites u and v:

$$h_{\mathrm{loop}} = 2 \sum_{c,d \in T_1 \;|c \neq d} |c, d\rangle\langle c, d| \,,$$

which penalizes neighboring sites with different colors. This adds a penalty of $2r - 2$ for every site in the middle of the chain. This is because each site has $2r$ neighbors, where all but exactly two are colored differently. This results in a penalty of $2(2r - 2)/2$ where we have divided by two to fix double-counting. Using similar reasoning, a penalty of $2r - 1$ is

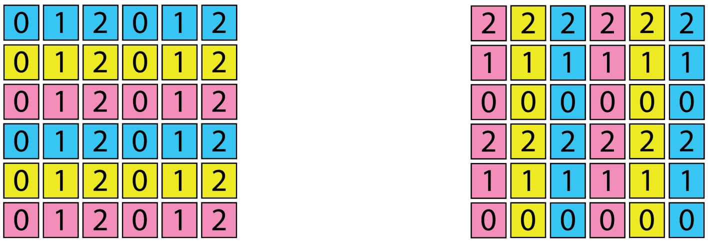

(a) First copy of tiles. **(b)** Second copy of tiles.

Figure 3 An example of how a 2D lattice can be tiled to optimize the Hamiltonian terms. Specifically, each copy forms a striped pattern where each stripe is numbered in a cyclic sequence. In addition, the two copies of tiles must have stripes pointing in different directions. Finally, the rows/columns that do not hold stripes must be numbered with the same number. Now a translation-invariant Hamiltonian can be simulated by using these tilings as a guideline for which sites are above, below, left, or right of each other. As our convention, we take the first copy to denote the horizontal direction and the second copy to denote the vertical direction.

incurred for every site at the endpoint of an open chain. Assuming that the classical tiling and EPR constraints are satisfied, this term is optimized when all 1D chains form loops so that each site always neighbors two other sites of the same color. The scaling of 2 is chosen so that the energy savings of coloring neighboring sites the same will never outweigh the energy penalty of violating the EPR constraints. This tradeoff will be worked out in detail in Section 4.2.

3.2 Adding the second dimension

Now that we have embedded stripes in one direction of the lattice, we must repeat this process to embed stripes in another direction. To do this, we can simply make a copy of each site's Hilbert space and apply the same Hamiltonian terms to the copy. It remains to ensure that both copies do not have stripes oriented in the same direction. To do this, it is sufficient to simply disallow two neighboring tiles from having matching colors on both copies. In other words, we add the following term where we let $\mathcal{H}_{T_1}$ denote the T_1 subspace of the first copy and $\mathcal{H}'_{T_1}$ denote that of the second copy.

$$h_{\text{copy}} = \Big(\sum_{c\in T_1} |c,c\rangle\langle c,c|\Big)_{\mathcal{H}_{T_1}} \otimes \Big(\sum_{d\in T_1} |d,d\rangle\langle d,d|\Big)_{\mathcal{H}'_{T_1}}$$

3.3 Embedding the translation-invariant Hamiltonian

An arbitrary two-dimensional translation-invariant Hamiltonian H_{TI} can now be embedded into our rotation-invariant Hamiltonian by using the directed stripes as guidelines. Our convention will be to let the first copy of each site represent the horizontal stripes and the second copy represent the vertical stripes. These tiling patterns are depicted in Figure 3. Now, the horizontal Hamiltonian term is applied only when the first copy tiles have the same color (i.e., different numbers). In addition, the orientation of which site is on the left and which is on the right will be decided by the directed edge in between the two tile's numbers where each arrow points from left to right. The vertical Hamiltonian term is applied similarly with respect to the second copy of each site's tiles.

To make this more concrete, we first attach to each site of our lattice the Hilbert space of a site in H_{TI} which we denote by $\mathcal{H}_{\mathrm{2D}}$. Denote the horizontal 2-body term of H_{TI} by h_{TI}. To incorporate this into our construction we add the following term:

$$h_{\mathrm{RI}} = \sum_{i \in T_2} |i, i+1 \bmod 3\rangle\langle i, i+1 \bmod 3| \otimes h_{\mathrm{TI}} \tag{2}$$

$$+ \sum_{i \in T_2} |i, i-1 \bmod 3\rangle\langle i, i-1 \bmod 3| \otimes S h_{\mathrm{TI}} S \tag{3}$$

where S is the swap operator on the two sites and the term acts on the first copy of $\mathcal{H}_{T_2}$. Note that while h_{TI} does not have reflection symmetry, h_{RI} does. The vertical term is implemented in the same way but acts on the second copy. We denote the vertical term by v_{RI}.

It remains to now describe the 2D translation-invariant Hamiltonian that encodes the computational hardness. We will combine techniques from [6] in order to encode the desired ground state energy in the yes and no cases. The details are deferred to Section A but the key result is outlined below.

▶ **Theorem 6.** *Let L be a QMA_{EXP}-complete language. There exists an efficiently computable function $f : \{0,1\}^* \to \mathbb{Z}$ and 2-local positive semidefinite Hamiltonian terms h_{TI} and v_{TI} with the following properties:*

1. *$f(x)$ is a multiple of 3 and $f(x)/3$ is prime. Furthermore, $\log f(x) = O(\mathrm{poly}(|x|))$ and f is computable in $O(\mathrm{poly}(|x|))$ time.*
2. *Let $f(x) = n \geq n_0$ for some constant n_0 and a given problem instance $x \in \{0,1\}^*$. For an $n \times n$ 2D lattice with periodic boundary conditions, let E_h be the set of ordered pairs of horizontal neighbors and let E_v be the set of ordered pairs of vertical neighbors. Consider the Hamiltonian $H_{TI} = \sum_{(u,w) \in E_h} h_{TI}^{u,w} + \sum_{(x,y) \in E_v} v_{TI}^{x,y}$.*
   - **a.** *If $x \in L$ then $E_0(H_{TI}) \leq O(n^{-k})$ for an arbitrarily large constant k.*
   - **b.** *If $x \notin L$ then $E_0(H_{TI}) \geq \Omega(1/n^3)$*

Using the prime symbol to denote terms acting on the second copy of a site's Hilbert space, we can now write the entire 2-body Hamiltonian term of our construction as follows:

$$h = h_{\mathrm{tile}} + h_{\mathrm{EPR}} + h_{\mathrm{loop}} + h'_{\mathrm{tile}} + h'_{\mathrm{EPR}} + h'_{\mathrm{loop}} + h_{\mathrm{copy}} + h_{\mathrm{RI}} + v_{\mathrm{RI}} \tag{4}$$

We can now define the reduction from any $\mathrm{QMA}_{\mathrm{EXP}}$ problem instance to an r-DIM-RIH problem instance.

▶ **Definition 7** (Reduction from $\mathsf{QMA}_{\mathsf{EXP}}$ to r-DIM-RIH)**.** *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Define the function $R(L,x,r) = \sum_{\{u,v\} \in E_{\Lambda_r(f(x))}} h^{u,v}$ where $f : \{0,1\}^* \to \mathbb{Z}$ is constructed as in Theorem 6 and h as in Eq. 4.*

Next, the detailed analysis of this construction is provided.

4 Analysis

In this section we prove the main theorem that the rotation-invariant Hamiltonian problem is $\mathrm{QMA}_{\mathrm{EXP}}$-complete.

4.1 Completeness

We begin by first considering the case where $x \in L$. In this case, we present a ground state that achieves energy below $p(n) = 4n^r(r-1) + 1/g(n)$ for an arbitrarily large polynomial $g(n)$.

▶ **Lemma 8.** *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \in L$, then $E_0(H) \leq 4n^r(r-1) + 1/g(n)$ for any polynomial $g(n)$.*

Proof. Consider the following tiling, which generalizes the tiling of the 2D lattice depicted in Figure 6. We first construct a 3-coloring of the $(r-1)$-dimensional lattice. This always exists because an $(r-1)$-dimensional lattice can be decomposed as a Cartesian product of cycle graphs. Since each cycle graph is 3-colorable their Cartesian product is also 3-colorable by a result by Sabidussi [13]. To color a site in the full r-dimensional lattice we simply drop the 1st coordinate and assign the color from the 3-coloring of the remaining $r-1$ coordinates. This enforces that any two neighboring sites that have the same 1st coordinate are colored by different colors and any two neighbors that only differ in the 1st coordinate are tiled by the same colors. This has the effect of coloring 1D chains of sites that travel in a straight line along the 1st coordinate dimension.

For this coloring, it is easy to satisfy the numbering constraints. This can be accomplished by tiling all particles u with the number u_1 mod 3. This simultaneously tiles all 1D chains with the ordering $0, 1, 2, 0, 1, 2, \ldots$ which eventually wraps around since n is a multiple of 3. In addition, all neighboring tiles with different color tiles are tiled with the same number since this only occurs when the two particles have the same 1st coordinate.

The EPR constraint can easily be satisfied since the particles have been tiled as disjoint 1D chains with the appropriate numbering. In addition, all chains are loops and so the constraint on having periodic boundary conditions is also satisfied. This ensures that only a penalty of $n^r(2r-2) = 2n^r(r-1)$ is introduced by the h_{loop} term. We can repeat this for the second copy of tiles but now directing the 1D chains along the second coordinate. This introduces another penalty of $2n^r(r-1)$.

With this choice of tiles the lattice has effectively been broken up into n^{r-2} 2D slices where we can now apply the 2D translation-invariant Hamiltonian construction. By Theorem 6, since $x \in L$ we have that each $2D$ slice contributes an energy of at most $O(n^{-k})$. The total energy is thus $O(n^{r-2-k}) = O(n^{-k'})$ where we have chosen $k = r - 2 + k'$. This results in a final energy upper bound of $4n^r(r-1) + O(n^{-k'})$. ◀

4.2 Soundness

In this section we will prove the following lemma:

▶ **Lemma 9.** *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L, x, r)$. If $x \notin L$, then $E_0(H) \geq 4n^r(r-1) + 1$.*

Our general strategy will be to first lower bound the energy of any state with a given classical tiling. We call these "tile states" and define them as follows:

▶ **Definition 10** (Tile state). *A* ***tile state*** *is a state $|\psi_{c,c'}\rangle = |c\rangle \otimes |c'\rangle \otimes |\phi\rangle$ where $c, c' \in T_1^N \times T_2^N$ and $|\phi\rangle \in (\mathcal{H}_{\sigma_1} \otimes \mathcal{H}_{\sigma_2} \otimes \mathcal{H}'_{\sigma_1} \otimes \mathcal{H}'_{\sigma_2} \otimes \mathcal{H}_{2D})^{\otimes N}$. The notation $|\psi_c\rangle$ is used whenever only one of the tilings is relevant.*

Lower bounding the energy of an arbitrary state follows straightforwardly, since each of the Hamiltonian terms is diagonal with respect to the tile Hilbert spaces.

We start by first establishing the fact that a qubit cannot be in two EPR pairs at once.

▶ **Fact 11.** *Consider a Hamiltonian on three qubits u, v and w defined by* $(\frac{\mathbb{I}-|\Phi^+\rangle\langle\Phi^+|}{2})_{u,v} + (\frac{\mathbb{I}-|\Phi^+\rangle\langle\Phi^+|}{2})_{v,w}$. *By direct computation this has a minimum eigenvalue of* $1/4$.

We next focus on one copy of the tilings at a time and define the following notation.

▶ **Definition 12.** *For a given site u, let n_u be the number of u's neighbors that are tiled with the same color as u.*

With these in hand, we can now lower bound the energy of any tile state in terms of the number of same-colored neighbors of each site.

▷ Claim 13.

$$\langle\psi_c| \sum_{\{u,v\}\in E_{\Lambda_r(n)}} (h_{\text{tile}} + h_{\text{EPR}} + h_{\text{loop}})^{u,v} |\psi_c\rangle \geq 2n^r r - \sum_{u\in\Lambda_r(n)} n_u + 4 \sum_{u\in\Lambda_r(n)} \lfloor n_u/3 \rfloor \tag{5}$$

Proof. Recall that h_{loop} gives an energy penalty of 2 for neighboring particles that are tiled by opposite colors. Every particle has $2r$ neighbors and so if every particle is tiled with a different color than all of its neighbors then the total penalty due to this Hamiltonian term would be $2\frac{n^r(2r)}{2} = 2n^r r$. The total number of neighboring pairs with the same color tiling is $\sum_{u\in\Lambda_r(n)} \frac{n_u}{2}$. Since each of these pairs saves an energy of 2, the total energy with respect to h_{loop} applied to each edge is $2n^r r - 2\sum_{u\in\Lambda_r(n)} \frac{n_u}{2} = 2n^r r - \sum_{u\in\Lambda_r(n)} n_u$.

Now for every particle u we can group its same-colored neighbors into groups of three until a full group of three cannot be formed. There will be $\lfloor n_u/3 \rfloor$ such groups. For a given group of three neighbors, label the particles v, w, z. If h_{tile} applied to edge $\{u, v\}$, $\{u, w\}$, or $\{u, z\}$ is violated then this incurs a penalty of 4 per particle involved (8 in total for the edge).

If none are violated then this means at least two of the three particles, v, w, z, must be tiled with the same number. This is because h_{tile} enforces that same-colored neighbors of u are tiled with a different number than u and there are only two such numbers. Without loss of generality assume that v and w are tiled with the same number and it is exactly one higher (mod 3) than u's number. Then we have the following bound

$$E_0(h^{u,v}_{\text{EPR}} + h^{u,w}_{\text{EPR}}) \geq E_0(A^{u,v} + A^{u,w}) \tag{6}$$
$$\geq 16E_0((\frac{\mathbb{I}-|\Phi^+\rangle\langle\Phi^+|}{2})_{u_{\sigma_2}v_{\sigma_1}} + (\frac{\mathbb{I}-|\Phi^+\rangle\langle\Phi^+|}{2})_{u_{\sigma_2}w_{\sigma_1}}) \tag{7}$$
$$\geq 4 \qquad \text{By Fact 11} \tag{8}$$

Repeating this argument for each site u in the lattice does not double count energy penalties since we only used the $A^{u,v}$ term of h_{EPR}, and so when considering v we would use the $A^{v,u}$ term instead. Therefore, either h_{tile} or h_{EPR} must be violated and either way a penalty of at least 4 is added to the energy. This argument can be repeated for each group of three same-colored neighbors and so this contributes at least $4\sum_{u\in\Lambda_r(n)} \lfloor n_u/3 \rfloor$ to the energy. ◁

It follows immediately from Equation (5) that each particle must have exactly two same colored-neighbors. Otherwise this induces an energy penalty of at least one which is enough to imply the bound in Lemma 9. This is helpful since we have now shown that the tiling pattern forms loops. It still remains to show that these loops are straight. Before moving on to this, we first make the above statements rigorous as follows:

▶ **Definition 14.** *A classical tiling $c \in T_1^N \times T_2^N$ is* ***looped*** *if* $n_u = 2\ \forall u \in \Lambda_r(n)$

▷ Claim 15. If c or c' is not looped, then $\langle\psi_{c,c'}| H |\psi_{c,c'}\rangle \geq 4n^r(r-1)+1$.

Proof. Without loss of generality assume that c is not looped and that c' is any tiling. The inequality in Eq. 5 is minimized when $n_u = 2$ for all $u \in \Lambda_r(n)$. This is because $-n_u + 4\lfloor n_u/3 \rfloor = -2$ for $n_u = 2$ and $-n_u + 4\lfloor n_u/3 \rfloor \geq -1$ for $n_u \neq 2$ where n_u is a nonnegative integer. The right-hand side of Eq. 5 is equal to $2n^r r - 2n^r = 2n^r(r-1)$ when $\forall u\ n_u = 2$. Therefore, when $n_u \neq 2$ for some $u \in \Lambda_r(n)$, $\langle\psi_c| \sum_{\{u,v\}\in E_{\Lambda_r(n)}} (h_{\text{tile}} + h_{\text{EPR}} + h_{\text{loop}})^{u,v} |\psi_c\rangle \geq 2n^r(r-1)+1$. The terms $h'_{\text{tile}}, h'_{\text{EPR}}$, and h'_{loop} add a penalty of at least $2n^r(r-1)$ as we have just argued. Finally, each term in H is positive semidefinite and so the energy with respect to the entire Hamiltonian is lower bounded by the energy with respect to a subset of the terms. The claim then follows. ◁

We now must show that these loops are in fact straight and do not contain any turns.

▶ **Definition 16.** *Let $g : \Lambda_r(n)^2 \to \mathbb{N}$ be a function that outputs the number of coordinates that differ between two lattice sites.*

▶ **Definition 17.** *A classical tiling $c \in T_1^N \times T_2^N$ has a* ***turn*** *if there exists three particles u, v, and w where the following is true: u and w are neighbors of v (i.e. $d(u,v) = d(v,w) = 1$), u, v, and w are all tiled by the same color, and $g(u,w) = 2$.*

We can now penalize turns by using the tiling rule that different color neighbors must have the same number. This is because if there is a turn then there exists some neighboring site outside of the loop that borders two sites right where the turn occurs. This causes a penalty because both sites in the loop must have the same number as the site outside of the loop, which contradicts the sequential numbering of the loop. This is argued in more detail below.

▷ Claim 18. If c or c' is looped but has a turn, then

$$\langle\psi_{c,c'}| H |\psi_{c,c'}\rangle \geq 4n^r(r-1)+4 \tag{9}$$

Proof. Once again without loss of generality, assume that c is looped and has a turn and that c' is any tiling. If c has a turn then there exists u, v, and w that are all tiled by the same color and $g(u,w) = 2$. Without loss of generality, let $u_1 = 0$, $u_2 = 0$, $v_1 = 0$, $v_2 = 1$, $w_1 = 1$, $w_2 = 1$. This simplification is for clarity but note that the following argument works for any coordinates such that $d(u,v) = d(v,w) = 1$ and $g(u,w) = 2$. Consider a fourth site z at coordinates $z_1 = 1$, $z_2 = 0$. Note that $d(z,u) = d(z,w) = 1$ and so it neighbors u and w. This situation is depicted in Figure 4. If z is tiled by the same color as u, v, and w then it would form a loop of length four. This violates h_{tile} since the loop is not a multiple of three and so the numbers can not sequentially wrap around $\mathbb{Z}_3$ (see Figure 4a). If z is tiled with a different color than u, v, and w then it must be tiled with the same number as u and w since h_{tile} enforces that different colored neighbors must have the same number. However, as we have argued previously, a site cannot have two same-colored neighbors that are tiled with the same number as each other since this causes the EPR constraint to be violated (see Figure 4b). In either case, there is an energy penalty of at least 4. In addition, the h_{loop} and h'_{loop} terms together add a penalty of $4n^r(r-1)$. Finally, noting the positive semidefiniteness of each Hamiltonian term concludes the proof. ◁

So far we have given a sufficient lower bound for any tile states that do not consist of only straight loops. It will be helpful to make a few more observations on the structure of these 1D loops.

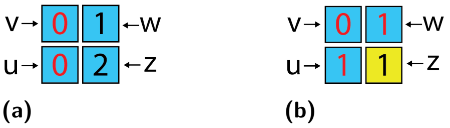

Figure 4 Two examples of how energy penalties can arise when there is a turn in the loop. Here, the loop consists partially of u, v, and w which contains a turn since u and w differ in more than one coordinate. In (a), z is colored the same as the rest but this results in a penalty since a loop of size 4 cannot be numbered cyclically around $\mathbb{Z}_3$. This results in the illegal configuration of u and v having the same color and the same number. In (b), z is colored differently but this leads to part of the loop being numbered as $1, 0, 1$, which is also illegal as depicted in Figure 2.

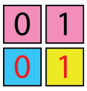

Figure 5 This configuration always arises if the tiling is not uniformly directed since otherwise each loop is always pointed in the same direction. This causes a rule violation since the blue and yellow tiles must have the same number but are forced to hold different numbers due to the red tiles.

▶ **Definition 19.** *A classical tiling $c \in T_1^N \times T_2^N$ is* ***uniformly directed*** *if it is looped without turns and each 1D loop is oriented along the same dimension.*

▷ Claim 20. If c or c' is looped without turns but not uniformly directed then

$$\langle \psi_{c,c'} | \, H \, | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 8 \tag{10}$$

Proof. If the tiling is not uniformly directed then the following situation will necessarily arise, which is illustrated in Figure 5. Without loss of generality, consider a square of sites within the lattice such that the top two sites are the same color and thus a part of the same 1D chain, but the bottom two sites have two different colors from the top and from each other. Since different color tiles must have the same numberings, then all four squares would be required to have the same number. However, since the top two have the same color, they are required to have different numbers and so it is impossible to assign numbers that do not incur any penalties. ◁

One final property we will need is that each 1D loop is numbered consistently.

▶ **Definition 21.** *A classical tiling $c \in T_1^N \times T_2^N$ is* ***numbered consistently*** *if it is uniformly directed towards a given dimension d and every site with the same dth coordinate value has the same number.*

▷ Claim 22. If c or c' is uniformly directed, but not numbered consistently then

$$\langle \psi_{c,c'} | \, H \, | \psi_{c,c'} \rangle \geq 4n^r(r-1) + 8 \tag{11}$$

Proof. Consider the $r - 1$ dimensional sublattice of sites that each have the same dth coordinate value. Each pair of neighbors in this lattice must have different colors. Otherwise, the coloring would not be uniformly directed towards the dth dimension since this means there is some 1D loop that points in a different direction. Since each pair of numbers has different colors, they all must have the same number, i.e. the tiles must be numbered consistently. Otherwise, a penalty of at least 8 is incurred. ◁

All that remains is to now lower bound the ground-state energy when both tilings are numbered consistently.

▷ Claim 23. Let c and c' be numbered consistently. Let L be a language in $\mathrm{QMA}_{\mathrm{EXP}}$ and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L,x,r)$ If $x \notin L$, then $\langle \psi_{c,c'} | \, H \, |\psi_{c,c'}\rangle \geq 4n^r(r-1) + \Omega(n^{r-5})$.

Proof. First, we must handle the case where c and c' both have 1D chains pointing in the same direction. This would incur a penalty of $n^r r$ from the h_{copy} term alone, which would clearly imply the desired lower bound. Next, we focus on the case where they point in different directions. We let the 1D chains in the first tiling represent the horizontal rows of each 2D slice and those of the second tiling represent vertical rows. In addition, we let the order of the numberings define the left, right, up and down directions of the slices. In this way we can embed n^{r-2} 2D translation-invariant Hamiltonians. Since $x \notin L$, by Theorem 6, these terms contribute an energy penalty of $\Omega(n^{r-5})$. ◁

To complete the proof it remains to deal with the non-tile states but these can easily be handled since the Hamiltonian is diagonal in the tile Hilbert spaces.

▶ **Lemma 24** (Restatement of Lemma 9). *Let L be a language in QMA_{EXP} and let $x \in \{0,1\}^*$ be a problem instance. Let $H = R(L,x,r)$. If $x \notin L$, then $E_0(H) \geq 4n^r(r-1) + 1$.*

Proof. First we can write an arbitrary state as a superposition of tile states: $|\zeta\rangle = \sum_i \alpha_i \, |c_i\rangle \, |\phi_i\rangle$ where $c_i \in T_1^{2N} \times T_2^{2N}$. Note that $\langle c_i| \, \langle \phi_i| \, H \, |c_j\rangle \, |\phi_j\rangle = 0$ for $c_i \neq c_j$. This is because all terms are diagonal on the Hilbert space of the classical tiles. Therefore, we have $\langle \zeta | \, H \, |\zeta\rangle = \sum_i |\alpha_i|^2 \, \langle c_i| \, \langle \phi_i| \, H \, |c_i\rangle \, |\phi_i\rangle$. This is an affine combination of tile state energies which we have already lower bounded by $4n^r(r-1)+1$. Thus, the energy itself is also bounded from below by $4n^r(r-1)+1$. ◀

5 Open boundary conditions

So far we have only considered the case where the Hamiltonian has periodic boundary conditions. It turns out that the same construction also works for open boundary conditions. Our method of embedding directed stripes still works in this case except now instead of closed loops the stripes form spin chains with open boundary conditions. This leaves the sites at the ends of the chain with one unpaired qubit that can still form an EPR pair with another site; however, this would require introducing a turn. The energy penalty for having a turn outweighs the energy bonus of having one more same-colored neighbor and so it is optimal to leave the qubit unpaired. Thus, the same construction can once again be used to embed a 2D translation-invariant non-reflection-invariant Hamiltonian with the exception that this Hamiltonian now has open boundary conditions. To complete the result, it remains to show the equivalent of Theorem 6 in the case of open boundary conditions. The proof of this statement is shown in Section A.2.

6 Conclusion

In this work, we have resolved the complexity for rotation-invariant Hamiltonians with constant lattice dimension, but it still remains interesting to better understand the complexity at even higher lattice dimensions. For instance, we know that as $r \to \infty$ the ground state becomes a product state, but how fast does it converge? Consider the problem where the lattice length is now fixed, and the lattice dimension is given as input. Is there a small enough promise gap for which this problem is quantumly hard? Another direction to consider is to study more general permutation symmetries. In some sense, the rotation-invariant Hamiltonian problem interpolates between systems with comparatively low symmetry in the one-dimensional case to highly symmetric systems as the lattice dimension increases. It would be an interesting question to generalize this interpolation and probe whether there exists a complexity phase transition with respect to some symmetry parameter.

References

1 Dorit Aharonov and Lior Eldar. Commuting local hamiltonians on expanders, locally testable quantum codes, and the qpcp conjecture, 2013. arXiv:1301.3407.

2 Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Communications in Mathematical Physics*, 287(1):41–65, January 2009. doi:10.1007/s00220-008-0710-3.

3 Johannes Bausch, Toby S. Cubitt, and James D. Watson. Uncomputability of phase diagrams. *Nature Communications*, 12(1), January 2021. doi:10.1038/s41467-020-20504-6.

4 Johannes Bausch and Stephen Piddock. The complexity of translationally invariant low-dimensional spin lattices in 3d. *Journal of Mathematical Physics*, 58(11), November 2017. doi:10.1063/1.5011338.

5 Fernando G.S.L. Brandao and Aram W. Harrow. Product-state approximations to quantum ground states. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 871–880, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2488608.2488719.

6 Daniel Gottesman and Sandy Irani. The quantum and classical complexity of translationally invariant tiling and hamiltonian problems. *Theory of Computing*, 9(2):31–116, 2013. doi:10.4086/toc.2013.v009a002.

7 A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, USA, 2002.

8 Tamara Kohler and Toby Cubitt. Translationally Invariant Universal Classical Hamiltonians. *Journal of Statistical Physics*, 176(1):228–261, July 2019. doi:10.1007/s10955-019-02295-3.

9 Christina V. Kraus, Maciej Lewenstein, and J. Ignacio Cirac. Ground states of fermionic lattice hamiltonians with permutation symmetry. *Physical Review A*, 88(2), August 2013. doi:10.1103/physreva.88.022335.

10 Roberto Oliveira and Barbara Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum information & computation*, 8, May 2005. doi:10.26421/QIC8.10-2.

11 Stephen Piddock and Johannes Bausch. Universal translationally-invariant hamiltonians, 2020. arXiv:2001.08050.

12 Ch. Rickwardt, P. Nielaba, and K. Binder. A finite size scaling study of the five-dimensional ising model. *Annalen der Physik*, 506(6):483–493, 1994. doi:10.1002/andp.19945060606.

13 Gert Sabidussi. Graphs with given group and given graph-theoretical properties. *Canadian Journal of Mathematics*, 9:515–525, 1957. doi:10.4153/CJM-1957-060-7.

14 B. M. Terhal. Is entanglement monogamous? *IBM Journal of Research and Development*, 48(1):71–78, January 2004. doi:10.1147/rd.481.0071.

A Constructing the 2D translation-invariant Hamiltonian

A.1 Periodic boundary conditions

In this section we prove the following theorem:

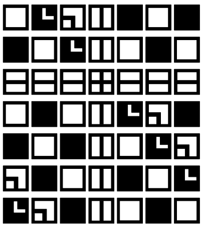

Figure 6 An example of an allowed configuration for a given set of translation-invariant tiling rules on a 2D grid defined in [6]. The key feature is that when the side length is a prime number, exactly one row can contain the ⊟ and ⊞ tiles and one site within this row can contain the ⊞ tile. The above image is directly reproduced from [6] under CC-BY 3.0.

▶ **Theorem 25** (Restatement of Theorem 6). *Let L be a QMA$_{EXP}$-complete language. There exists an efficiently computable function $f : \{0,1\}^* \to \mathbb{Z}$ and 2-local positive semidefinite Hamiltonian terms h_{TI} and v_{TI} with the following properties:*

1. *$f(x)$ is a multiple of 3 and $f(x)/3$ is prime. Furthermore, $\log f(x) = O(\mathrm{poly}(|x|))$ and f is computable in $O(\mathrm{poly}(|x|))$ time.*
2. *Let $f(x) = n \geq n_0$ for some constant n_0 and a given problem instance $x \in \{0,1\}^*$. For an $n \times n$ 2D lattice with periodic boundary conditions, let E_h be the set of ordered pairs of horizontal neighbors and let E_v be the set of ordered pairs of vertical neighbors. Consider the Hamiltonian $H_{TI} = \sum_{(u,w)\in E_h} h_{TI}^{u,w} + \sum_{(x,y)\in E_v} v_{TI}^{x,y}$.*
   a. *If $x \in L$ then $E_0(H_{TI}) \leq O(n^{-k})$ for an arbitrarily large constant k.*
   b. *If $x \notin L$ then $E_0(H_{TI}) \geq \Omega(1/n^3)$*

Much of the proof will directly utilize techniques from [6]. When needed, a brief summary of these ideas is given but we direct the interested reader to [6] for the full details. The main idea of the construction is to use the 2D translation-invariant tiles to embed a 1D chain with a designated starting tile. This allows us to avoid the 1/2 additive penalty that is required in the 1D construction when there is no designated starting tile (see section 6 "The quantum cycle" of [6]). Fortunately, the construction in tables 5 and 6 of section 4.1 of [6] accomplish exactly this. This construction is quite complicated and so we only present the end result. The last piece to handle is that this construction requires the grid length to be prime while our tiling rules require it to be a multiple of three. This can easily be remedied by simulating each tile in the 2D construction with a 3×3 grid of nine tiles. We explain each of these steps in more detail below.

We start by describing how to construct the function f referenced in the theorem statement. [6] show that given x there is a randomized algorithm a running in expected time $O(\mathrm{poly}\,|x|)$ to find a prime number p such that the 1/3 most significant digits represent x and $\log p = O(\mathrm{poly}(|x|))$. Using this result, we simply define f as $f := 3 \cdot a(x)$.

Next, to construct the 2D translation-invariant Hamiltonian, we start by using the below result.

▶ **Lemma 26** (see section 4.1 of [6]). *There exists a set of translation-invariant, non-reflection-invariant tiling rules involving a constant number of total tile types on an $n \times n$ 2D grid with periodic boundary conditions such that when n is prime exactly one row must contain tiles of the form ⊟ and ⊞ and no other row can contain either of these tiles. In addition, exactly one site in this row must contain ⊞ and all other sites in this row must contain ⊟. This configuration is depicted in Figure 6.*

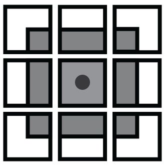

Figure 7 This is the only allowed configuration for the subtiles associated with each 3×3 grid. Each such grid represents a tile in the original tiling system.

Since this construction only works for prime n but our lattice length must be a multiple of three, it is necessary to replace each tile with a 3×3 grid of tiles that serve the same function as the original. First, for each tile in the original tiling, a corresponding center tile is defined as ▣. Then, the following tiles and rules are added. ⬓ must be placed above ▣. Similarly, ⬒ must be placed below ▣, ◨ placed to the left of it and ◧ placed to the right of it. Now to fix the corners in place we enforce that ◲ must be placed to the left of ⬓, ◱ placed to the right of ⬓, ◳ placed to the left of ⬒ and ◰ placed to the right of ⬒. In other words, the center tile must always be surrounded by the border tiles. The reciprocal of each of these rules is also included. This enforces that the border tiles must always be accompanied by the center tile in the appropriate location. This results in Figure 7 being the only allowed configuration. The last thing to resolve is to ensure that these 3×3 grids are aligned with each other. To do this, we must regulate which border tiles can be placed next to those of a different 3×3 grid. We enforce that only ◧ type tiles are allowed to the left of ◨ type tiles. We incorporate the same rule for the other sides as well: only ◨ type tiles are allowed to the right of ◧ type tiles, only ⬒ type tiles are allowed above ⬓ type tiles and only ⬓ type tiles are allowed below ⬒ type tiles.

Now we can incorporate an original rule between tiles by applying it to their corresponding border tiles. This will exactly simulate the original but with each site replaced by a 3×3 grid. This results in an $n \times n$ grid where $n = 3p$ and p is a prime number.

With this tiling in hand, the 1D construction can now be embedded into the 3×3 grids associated with the ⊟ and ⊞ tiles. In particular, only the middle row of the 3×3 grids associated with each ⊟ and ⊞ tile is used for the chain (i.e. only the tiles ◨, ▣, and ◧). In addition, the ▣ tile associated with the ⊞ tile is used to mark the left endpoint of the 1D construction.

It remains to construct a 1D translation-invariant Hamiltonian on a $f(x)$-length spin chain with the desired ground-state energy properties. To accomplish this [6] is directly used and is briefly summarized here. The main idea is to use the Hamiltonian terms to simulate a quantum Turing machine where each site in the length-$f(x)$ spin chain represents a different cell of the Turing machine tape except for two sites which are used to mark the boundaries. The goal of the first part of the Turing machine is to infer x from the length of the tape and write it on the tape. The second part of the Turing machine is quantum and uses x as the input along with a quantum witness to execute the $\mathrm{QMA}_{\mathrm{EXP}}$ verification algorithm for L.

We now discuss the first part of the Turing machine, which we denote as M_{BC}. M_{BC} is a purely classical Turing machine implementing a binary counter. By incrementing a clock pointer from one side of the tape to the other, we can ensure that M_{BC} is run for exactly $f(x) - 3$ steps. Recall that $f(x) = 3p$ where $1/3$ of the most significant digits of p is x. Therefore, by using a sufficiently slow binary counter, it is possible to ensure that x is always written on the tape at the end of the $f(x) - 3$ steps.

Now that x is on the tape, we can run the quantum verification algorithm on x along with an arbitrary quantum witness. The verifier is also allowed a total of $f(x) - 3$ timesteps. Notice this is more strict than $\mathrm{QMA}_{\mathrm{EXP}}$, which allows the verifier $2^{\mathrm{poly}\,|x|}$. $\mathrm{QMA}_{\mathrm{EXP}}$ can be

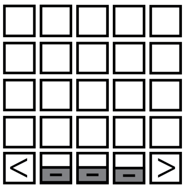

Figure 8 The only allowed configuration for the given set of translation-invariant tiling rules on a 2D grid with open boundary conditions.

reduced to this case by a standard padding argument where x is padded by zeros to have length poly$(|x|)$. Finally, an energy penalty is applied if the verifier does not accept. If the verifier accepts with probability $1-\epsilon$ in the $x \in L$ case then the ground-state energy will be upper bounded by ϵ/n^2. In order to drive $\epsilon \leq O(n^{-k})$ for an arbitrarily large constant, it is possible to use witness amplification. This incurs a $O(k \log n)$ overhead in the verifier's runtime [7], which can easily be accommodated by padding. In our construction, we would like to set $k = O(r)$ where r is the dimension, but would prefer not to have the verification algorithm depend on r. Therefore, we can instead set $k = \log n$ where $\log n \geq O(r)$ for some $n \geq n_0$ since r is a parameter of the problem and does not scale with n. Importantly, even though the number of rounds of witness amplification depends on n, our Hamiltonian term still does not depend on n since n is deduced from the length of the lattice and then given as input to the verification algorithm. For this construction [6] also show that in the $x \notin L$ case the ground-state energy is lower bounded by $\Omega(1/n^3)$.

Finally, each term in the construction is of one of two forms called Type 1 terms: $|ab\rangle\langle ab|$ and Type II terms:

$$\frac{1}{2}(|ab\rangle\langle ab| + |cd\rangle\langle cd| - |ab\rangle\langle cd| - |cd\rangle\langle ab|). \tag{12}$$

Both types are positive semidefinite and so the overall Hamiltonian term is also positive semidefinite.

A.2 Open boundary conditions

An equivalent theorem to Theorem 6 is also true in the case of open boundary conditions. The construction is also inspired by [6]. First, we define the tiles $\boxed{<}$, ⊟, $\boxed{>}$, and □. The idea is then to introduce the following tiling rules: no tile is allowed to the left of or below $\boxed{<}$, no tile is allowed below ⊟ and no tile is allowed to the right of or below $\boxed{>}$. Additionally, the only tile that is allowed to the right of $\boxed{<}$ is ⊟ and the only tile that is allowed to the left of $\boxed{>}$ is also ⊟. Finally, □ is not allowed to the left or right of ⊟ and the only tile allowed above $\boxed{<}$, ⊟ and $\boxed{>}$ is □. This results in the configuration depicted in Figure 8. The 1D Hamiltonian can then be embedded into the bottom row of the 2D grid where the $\boxed{<}$ and $\boxed{>}$ tiles denote the endpoints.