

PhD

3<sup>rd</sup>  
CYCLE

FCUP  
YEAR

U.PORTO

Quantum circuit design, discovery and  
optimization from theory to practice




Pedro Miguel  
Miranda Queiroz da Cruz

FC



# Quantum circuit design, discovery and optimization from theory to practice

Pedro Miguel Miranda Queiroz da Cruz  
Doctoral Program in Computer Science  
Department of Computer Science  
Faculty of Sciences of the University of Porto  
2024



# Quantum circuit design, discovery and optimization from theory to practice

**Pedro Miguel Miranda Queiroz da Cruz**

Thesis carried out as part of the Doctoral Program in  
Computer Science  
Department of Computer Science  
2024

**Supervisor**

Prof. Dr. Joaquín Fernández-Rossier, Group Leader,  
International Iberian Nanotechnology Laboratory

**Supervisor**

Prof. Dr. João Pedro Pedroso, Associate Professor,  
University of Porto

*To those who search truth with unwavering integrity.*

# Acknowledgements

Pursuing a PhD was never about obtaining a degree but about the opportunity to conduct research, develop new ideas, and contribute to expanding the boundaries of knowledge. This endeavor would not have been possible without the financial support from Fundação para a Ciência e a Tecnologia through the PhD scholarship grant No. SFRH/BD/150708/2020. This support granted me the freedom to advance research at the International Iberian Nanotechnology Laboratory (INL), in Portugal, and the Institute of Photonic Sciences (ICFO), in Spain, in addition to my home University of Porto, where I benefited from expertise in quantum information theory, computer science and condensed matter physics.

Over these years, I had the honor of collaborating with exceptional researchers across these inspiring institutions. I am grateful to my collaborators and academic supervisors for their availability and guidance, as well as for granting me an extraordinary level of independence throughout this PhD. I extend my sincere gratitude to João Pedro Pedroso, for continuously supporting my interest in deepening my expertise in computer science, building upon my background in physics; to Joaquín Fernández-Rossier, whose early encouragement and scientific guidance in quantum simulation created an exciting environment at INL and opened opportunities that shaped my trajectory; and to Antonio Acín, for welcoming me into his vibrant research group at ICFO where I had an amazing time and made many friends for three years. Additionally, my frequent discussions with fellow PhD student Bruno Murta have made for a highly rewarding and enjoyable collaboration over the course of these years. My work with each of them culminated in the four research projects presented in the chapters of this dissertation.

I cannot briefly mention everyone with whom I had stimulating discussions about the topics covered in this thesis or engaged in tangential collaborations, but I would like to acknowledge Ernesto Galvão, Donato Farina, Victoria Wright, Hippolyte Dourdent, Enky Oudot, Rui Soares Barbosa, Korbinian Kottmann, Luke Mortimer, Elisa Bäumer, Máté Farkas, Borja Requena and Louise Stepanov for enriching this journey.

In addition, I acknowledge the use of IBM Quantum processors to obtain the experimental results presented in this dissertation, both through their public cloud service and through the IBM Quantum Network access tier. The latter provided direct access to reserved execution time on their processors, which was essential for achieving the precision required for one of the projects.

A PhD can be, at times, a solitary journey with plenty of setbacks, but once in a while, it gives you the blissful joy of discovery. Throughout this journey, I am deeply grateful to my parents, my sister, and Linda for their love and permanent support.

# Resumo

**Título:** Concepção, descoberta e otimização de circuitos quânticos da teoria à prática.

Um objetivo de longa data na investigação em computação quântica é não apenas desenvolver novos algoritmos quânticos, mas também transitar de construções teóricas para implementações práticas, alcançando assim uma vantagem computacional quântica tangível sobre os métodos clássicos. A concretização deste objetivo exige avanços coordenados tanto no desenvolvimento teórico como na implementação experimental, cada um enfrentando desafios únicos, mas interligados. Esta dissertação aborda sistematicamente estes desafios através de contribuições que abrangem avanços teóricos fundamentais, metodologias práticas e extensas validações experimentais, todas unificadas no modelo de circuitos quânticos em quatro projetos interconectados.

Em primeiro lugar, dada a necessidade persistente de novos exemplos conceituais de algoritmos e protocolos quânticos, é introduzido um protocolo quântico para resolver o chamado *Puzzle Lógico Mais Difícil de Sempre*. Este resultado fornece uma demonstração rigorosa e formalmente provada de vantagem quântica para além do paradigma da computação clássica, destacando a sua aplicabilidade mais ampla no raciocínio e dedução em interações entre agentes. Em segundo lugar, esta tese avança nos testes fundamentais da mecânica quântica ao realizar experiências de complementaridade em processadores quânticos supercondutores com uma precisão sem precedentes. Uma análise estatística rigorosa expõe o impacto dos erros coerentes em portas de dois qubits na fronteira quântico-clássica, fornecendo *insights* que informam o desenvolvimento de um protocolo de mitigação de erros numa fase posterior. Em terceiro lugar, é introduzida uma metodologia abrangente para a otimização de circuitos quânticos. Baseado em ZX-calculus e em técnicas de inteligência artificial, este método reduz significativamente a contagem de portas de dois qubits em circuitos gerais, igualando ou superando os métodos de otimização de última geração, ao mesmo tempo que reduz substancialmente os custos computacionais. Para além de melhorar o tempo de execução e o desempenho prático em dispositivos quânticos com ruído, esta metodologia também possibilita uma abordagem sistemática para a descoberta e generalização de algoritmos quânticos através da análise das estruturas de portas de emaranhamento nos circuitos. Em

quarto lugar, este método de otimização e classificação de circuitos é aplicado a duas primitivas fundamentais da computação quântica — especificamente as portas de Fredkin e Toffoli — produzindo diversas decomposições logicamente equivalentes para diferentes topologias de conectividade. Um protocolo de execução aleatorizada de circuitos é proposto e validado empiricamente, explorando estas decomposições equivalentes com diferentes estruturas de portas de emaranhamento para mitigar eficazmente erros coerentes sem custos adicionais de hardware.

Ao integrar a concepção de algoritmos quânticos, a validação experimental e a otimização de circuitos, esta dissertação avança um paradigma que liga a computação quântica teórica à sua realização prática. As técnicas propostas não só melhoram o desempenho de circuitos quânticos independentes e dependentes de hardware, mitigando erros, mas também fornecem um quadro sistemático para a descoberta de novos algoritmos quânticos, estabelecendo as bases para subsequente generalização e aplicabilidade mais abrangente. Este trabalho destaca o papel emergente do *co-design* na computação quântica, complementando o desenvolvimento tradicional de algoritmos com métodos assistidos por máquinas. À medida que o hardware quântico evolui, estas metodologias serão determinantes na busca por uma computação quântica escalável e tolerante a falhas e, em última instância, na conquista de uma vantagem computacional quântica prática.

**Palavras-chave:** Computação quântica, Vantagem quântica, Otimização de circuitos quânticos, Mitigação de erros quânticos, Inteligência artificial, Descoberta assistida por máquina.

# Abstract

**Title:** Quantum circuit design, discovery and optimization from theory to practice.

A long-standing goal in quantum computing research is not only to develop novel quantum algorithms but also to transition them from theoretical constructs to practical implementations, thereby achieving a tangible quantum computational advantage over classical methods. Realizing this goal demands coordinated advancements in both theoretical development and experimental implementation, each confronting unique yet interconnected challenges. This dissertation systematically addresses these challenges through contributions spanning fundamental theoretical insights, practical methodologies, and extensive experimental validations, all unified within the quantum circuit model in four interconnected projects.

First, given the persistent need for novel conceptual examples of quantum algorithms and protocols, a quantum protocol is introduced to solve the so-called *Hardest Logic Puzzle Ever*. This result provides a rigorous, formally provable demonstration of quantum advantage beyond the classical computing paradigm, highlighting its broader applicability to reasoning and deduction in agent-based interactions. Second, this thesis advances foundational tests of quantum mechanics by conducting complementarity experiments on superconducting quantum processors with unprecedented precision. A rigorous statistical analysis exposes the impact of coherent errors in two-qubit gates on the quantum-classical boundary, providing insights that inform the development of an error mitigation protocol at a later stage. Third, a comprehensive framework for quantum circuit optimization is introduced. Built upon the ZX-calculus and artificial intelligence techniques, this approach significantly reduces the two-qubit gate count in general circuits, matching or outperforming state-of-the-art optimization methods while incurring substantially lower computational costs. Beyond improving execution time and practical performance on noisy quantum devices, this methodology also enables a systematic framework for discovering and generalizing quantum algorithms by analyzing entangling-gate circuit structures. Fourth, this circuit optimization and classification framework is applied to two fundamental quantum primitives — specifically Fredkin and Toffoli gates — yielding diverse logically equivalent decompositions across arbitrary connectivity topologies. A randomized circuit execution protocol is proposed and empirically validated, leveraging these equivalent decompositions with

distinct entangling-gate structures to effectively mitigate coherent errors without additional hardware costs.

By integrating quantum algorithm design, experimental validation, and circuit optimization, this dissertation advances a paradigm that bridges theoretical quantum computing and its practical realization. The proposed techniques not only refine both hardware-agnostic and hardware-aware quantum circuit performance while mitigating errors, but also provide a systematic approach for discovering new quantum algorithms, laying the groundwork for further generalization and broader applicability. This work highlights the emerging role of co-design in quantum computing, complementing traditional algorithm development with machine-assisted discovery. As quantum hardware advances, these methodologies will be instrumental in the pursuit of scalable, fault-tolerant quantum computing and, ultimately, practical quantum computational advantage.

**Keywords:** Quantum computing, Quantum advantage, Quantum circuit optimization, Quantum error mitigation, Artificial intelligence, Machine-assisted discovery.

# Contents

|                                                                               |             |
|-------------------------------------------------------------------------------|-------------|
| <b>Contents</b>                                                               | <b>vi</b>   |
| <b>Nomenclature</b>                                                           | <b>x</b>    |
| <b>List of Figures</b>                                                        | <b>xiii</b> |
| <b>List of Tables</b>                                                         | <b>xv</b>   |
| <b>List of Algorithms</b>                                                     | <b>xvi</b>  |
| <b>1 Introduction</b>                                                         | <b>1</b>    |
| 1.1 Framework and motivation . . . . .                                        | 1           |
| 1.1.1 Theoretical advances in quantum algorithms . . . . .                    | 1           |
| 1.1.2 Experimental challenges and recent achievements . . . . .               | 5           |
| 1.1.3 Optimizing circuits, discovering algorithms . . . . .                   | 10          |
| 1.2 Dissertation outline . . . . .                                            | 12          |
| 1.2.1 Overview . . . . .                                                      | 12          |
| 1.2.2 Structure . . . . .                                                     | 13          |
| <b>Quantum computing: from theory to practice</b>                             | <b>16</b>   |
| <b>2 Quantum advantage in the solution of the “Hardest Logic Puzzle Ever”</b> | <b>17</b>   |
| 2.1 Introduction . . . . .                                                    | 17          |
| 2.2 Formalizing the HLPE . . . . .                                            | 25          |
| 2.3 Classical communication of a binary value . . . . .                       | 29          |
| 2.3.1 Assumed innate modes of communication . . . . .                         | 29          |
| 2.3.2 Minimizing physical resources and shared assumptions . . . . .          | 31          |

|          |                                                                                 |           |
|----------|---------------------------------------------------------------------------------|-----------|
| 2.3.3    | Realizing classical randomness . . . . .                                        | 36        |
| 2.4      | Classical-communication-based solution to the HLPE . . . . .                    | 38        |
| 2.4.1    | Revisiting Boolos' exact solution . . . . .                                     | 38        |
| 2.4.2    | Probabilistic solutions and success probabilities . . . . .                     | 39        |
| 2.5      | Quantum communication in the HLPE . . . . .                                     | 45        |
| 2.6      | A one of a kind solution to the HLPE in two questions . . . . .                 | 48        |
| 2.7      | A generalized quantum implementation . . . . .                                  | 53        |
| 2.7.1    | Relaxing the computational basis . . . . .                                      | 58        |
| 2.7.2    | Minimum-error state discrimination strategy . . . . .                           | 60        |
| 2.8      | Communication subject to noise . . . . .                                        | 65        |
| 2.8.1    | Protocol performance under equivalent bit-flip errors . . . . .                 | 67        |
| 2.8.2    | Classical communication with a repetition code . . . . .                        | 71        |
| 2.8.3    | Quantum communication with a repetition code . . . . .                          | 73        |
| 2.9      | Discussion and Conclusion . . . . .                                             | 77        |
| 2.10     | Coda: Extensions and Artificial Intelligence . . . . .                          | 81        |
| 2.11     | Appendix . . . . .                                                              | 85        |
| 2.11.1   | Minimum-error discrimination via semidefinite programming . . . . .             | 85        |
| 2.11.2   | Quantum circuit decomposition for the square-root measurement . . . . .         | 86        |
| <b>3</b> | <b>Testing complementarity on a transmon quantum processor</b>                  | <b>90</b> |
| 3.1      | Introduction . . . . .                                                          | 90        |
| 3.2      | Quantum circuits for which-path detection . . . . .                             | 93        |
| 3.2.1    | The interferometer . . . . .                                                    | 93        |
| 3.2.2    | The detector . . . . .                                                          | 94        |
| 3.2.3    | Interference–detection tradeoff . . . . .                                       | 97        |
| 3.2.4    | Quantum erasure . . . . .                                                       | 98        |
| 3.3      | Experimental results . . . . .                                                  | 99        |
| 3.3.1    | Methods . . . . .                                                               | 100       |
| 3.3.2    | Single-qubit interference circuit . . . . .                                     | 102       |
| 3.3.3    | Optimal path detection and the visibility-distinguishability tradeoff . . . . . | 104       |
| 3.3.4    | Quantum Eraser . . . . .                                                        | 107       |
| 3.4      | Accounting for coherent errors . . . . .                                        | 110       |
| 3.4.1    | A biased-CNOT gate model . . . . .                                              | 111       |

|       |                                                                                |     |
|-------|--------------------------------------------------------------------------------|-----|
| 3.5   | Refinement and statistical validation of the circuit models . . . . .          | 115 |
| 3.5.1 | Impact of the BCNOT on exclusive particlelike or wavelike statistics . . . . . | 118 |
| 3.6   | Discussion and conclusions . . . . .                                           | 120 |
| 3.7   | Appendix . . . . .                                                             | 122 |
| 3.7.1 | Conditional statistics in discrete bivariate sampling . . . . .                | 122 |
| 3.7.2 | Model parameter values and other details of model evaluation . . . . .         | 125 |

## **Circuit discovery and optimization 128**

### **4 Logically-equivalent quantum circuit discovery and optimization 129**

|       |                                                                          |     |
|-------|--------------------------------------------------------------------------|-----|
| 4.1   | Introduction . . . . .                                                   | 129 |
| 4.1.1 | Quantum circuit optimization . . . . .                                   | 131 |
| 4.1.2 | Motivation . . . . .                                                     | 137 |
| 4.2   | Entangling-gate topologies in multi-qubit circuits . . . . .             | 138 |
| 4.3   | Adaptive circuit optimization via dual representation feedback . . . . . | 142 |
| 4.3.1 | Best-first search as a metaheuristic . . . . .                           | 145 |
| 4.3.2 | Technical implementation . . . . .                                       | 149 |
| 4.4   | A practical example in quantum simulation . . . . .                      | 151 |
| 4.5   | Towards scalable algorithmic gate discovery . . . . .                    | 155 |
| 4.6   | Conclusions and future work . . . . .                                    | 158 |
| 4.7   | Appendix . . . . .                                                       | 163 |
| 4.7.1 | Quantum circuit decomposition for the VBS simulation state . . . . .     | 163 |

### **5 Optimized Fredkin and Toffoli gate decompositions for coherent error mitigation via equivalent circuit averaging 165**

|       |                                                    |     |
|-------|----------------------------------------------------|-----|
| 5.1   | Introduction . . . . .                             | 165 |
| 5.2   | Decompositions for adjacent qubits . . . . .       | 169 |
| 5.3   | Decompositions for non-adjacent qubits . . . . .   | 173 |
| 5.3.1 | CNOT-SWAP rerouting . . . . .                      | 173 |
| 5.3.2 | Long-range CNOT and Toffoli gates . . . . .        | 175 |
| 5.3.3 | Fredkin gate . . . . .                             | 178 |
| 5.4   | Equivalent circuit averaging . . . . .             | 180 |
| 5.4.1 | Numerical analysis using the BCNOT model . . . . . | 182 |

|                                   |                                                                       |            |
|-----------------------------------|-----------------------------------------------------------------------|------------|
| 5.4.2                             | Experimental validation on superconducting hardware . . . . .         | 185        |
| 5.5                               | Conclusion . . . . .                                                  | 188        |
| 5.6                               | Appendix . . . . .                                                    | 190        |
| 5.6.1                             | Simplifying a controlled gate via symmetric decomposition . . . . .   | 190        |
| 5.6.2                             | Reassigning the target qubit in a Toffoli gate . . . . .              | 190        |
| 5.6.3                             | Gate diagonality condition in a qubit's computational basis . . . . . | 192        |
| 5.6.4                             | Proof of the diamond distance bound for ECA . . . . .                 | 192        |
| <b>Conclusion</b>                 |                                                                       | <b>195</b> |
| <b>Publications by the author</b> |                                                                       | <b>199</b> |
| <b>Bibliography</b>               |                                                                       | <b>200</b> |

# Nomenclature

## Abbreviations

|        |                                         |
|--------|-----------------------------------------|
| AI     | Artificial Intelligence                 |
| AKLT   | Affleck-Kennedy-Lieb-Tasaki             |
| BCNOT  | Biased controlled-NOT                   |
| BFS    | Best-First Search                       |
| BQP    | Bounded-error Quantum Polynomial time   |
| CHSH   | Clauser-Horne-Shimony-Holt inequality   |
| CNOT   | Controlled-NOT                          |
| CR     | Cross-Resonance                         |
| CSD    | Cosine-Sine Decomposition               |
| ECA    | Equivalent Circuit Averaging            |
| EGS    | Entangling-Gate Structure               |
| GQSP   | Generalized Quantum Signal Processing   |
| HHL    | Harrow-Hassidim-Lloyd algorithm         |
| HLPE   | Hardest Logic Puzzle Ever               |
| i.i.d. | independent and identically distributed |
| LBM    | Language-Boolean Map                    |
| LCU    | Linear Combination of Unitaries         |

|      |                                            |
|------|--------------------------------------------|
| MCTS | Monte-Carlo Tree Search                    |
| MEM  | Measurement Error Mitigation               |
| MSE  | Mean Squared Error                         |
| NISQ | Noisy Intermediate-Scale Quantum           |
| POVM | Positive Operator-Valued Measure           |
| QAOA | Quantum Approximate Optimization Algorithm |
| QEC  | Quantum Error Correction                   |
| QEM  | Quantum Error Mitigation                   |
| QFT  | Quantum Fourier Transform                  |
| QITE | Quantum Imaginary Time Evolution           |
| QKD  | Quantum Key Distribution                   |
| QMA  | Quantum Merlin-Arthur                      |
| QPE  | Quantum Phase Estimation                   |
| QSD  | Quantum Shannon Decomposition              |
| QSP  | Quantum Signal Processing                  |
| QSVT | Quantum Singular Value Transformation      |
| RL   | Reinforcement Learning                     |
| SBM  | State-Boolean Map                          |
| SCB  | Shared Computational Basis                 |
| SCE  | Single-Circuit Execution                   |
| SDP  | Semidefinite programming                   |
| SMT  | Satisfiability Modulo Theories             |
| SQG  | Single-Qubit Gate                          |

|      |                                 |
|------|---------------------------------|
| SQGE | Single-Qubit Gate Error         |
| SRF  | Shared Reference Frame          |
| TQG  | Two-Qubit Gate                  |
| VBS  | Valence-Bond-Solid              |
| VQE  | Variational Quantum Eigensolver |

# List of Figures

|     |                                                                                                                                       |     |
|-----|---------------------------------------------------------------------------------------------------------------------------------------|-----|
| 2.1 | State-based communication scheme between Alice and a god. . . . .                                                                     | 32  |
| 2.2 | Operations-based communication scheme between Alice and a god. . . . .                                                                | 35  |
| 2.3 | Classical two-question minimum-error discrimination strategy to solve the HLPE. . . . .                                               | 42  |
| 2.4 | Superdense-coding-based questioning protocol. . . . .                                                                                 | 50  |
| 2.5 | Two approaches to generalize the computational basis. . . . .                                                                         | 55  |
| 2.6 | Communication with an extraterrestrial god. . . . .                                                                                   | 59  |
| 2.7 | Quantum circuit decomposition for the square-root measurement. . . . .                                                                | 63  |
| 2.8 | Success probability for all the HLPE solution protocols under a noise channel with bit-flip<br>error probability $\epsilon$ . . . . . | 69  |
| 2.9 | Quantum communication of three bit types via a repetition code. . . . .                                                               | 76  |
| 3.1 | Which-path detection circuit. . . . .                                                                                                 | 94  |
| 3.2 | Physical implementation of the which-path circuits. . . . .                                                                           | 101 |
| 3.3 | Experimental results for the interference circuit. . . . .                                                                            | 103 |
| 3.4 | Results of the visibility and distinguishability tradeoff experiments. . . . .                                                        | 105 |
| 3.5 | Residual errors from experimental model fitting of the visibility-distinguishability tradeoff. . . . .                                | 106 |
| 3.6 | Experimental results for the quantum eraser circuit. . . . .                                                                          | 108 |
| 3.7 | Theoretical visibility and distinguishability trade-off under a biased-CNOT model. . . . .                                            | 120 |
| 4.1 | Entangling-gate structure calculation. . . . .                                                                                        | 139 |
| 4.2 | Exploration path in two best-first search example runs. . . . .                                                                       | 153 |
| 4.3 | Statistics for the second best-first search example run. . . . .                                                                      | 154 |
| 5.1 | Circuit decompositions of Fredkin and Toffoli gates with full qubit connectivity. . . . .                                             | 170 |
| 5.2 | Circuit decompositions of Toffoli and Fredkin gates under linear qubit connectivity. . . . .                                          | 171 |
| 5.3 | [Control and target-NOT qubit relocation via cnot-swapping . . . . .                                                                  | 175 |
| 5.4 | Long-range CNOT gate implementation. . . . .                                                                                          | 177 |

|     |                                                                                        |     |
|-----|----------------------------------------------------------------------------------------|-----|
| 5.5 | Long-range Toffoli and Fredkin gate implementations under linear connectivity. . . . . | 178 |
| 5.6 | Fredkin and Toffoli gate implementations via ECA under the <b>BCNOT</b> model. . . . . | 184 |
| 5.7 | Experimental evaluation of ECA performance. . . . .                                    | 187 |
| 5.8 | Quantum circuit to control a symmetric $n$ -qubit gate. . . . .                        | 191 |
| 5.9 | Reassigning the Target Qubit in a Toffoli Gate. . . . .                                | 191 |

# List of Tables

|     |                                                                                                                                                                |     |
|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 2.1 | Truth table of the “ <i>if and only if</i> ” (iff) logical connective. . . . .                                                                                 | 34  |
| 2.2 | Truth table and received state for a truthful answer encoded in a two-state system without a shared state-boolean map. . . . .                                 | 34  |
| 2.3 | Truth table and received state transformation for a truthful answer encoded in a two-state system without a shared operation-boolean map. . . . .              | 36  |
| 2.4 | Truth table and received state for an answer to $\mathcal{Q}_{10}$ from either <code>True</code> or <code>False</code> in both possible OMB scenarios. . . . . | 41  |
| 2.5 | Truth table and received state for an answer to $\mathcal{Q}_9$ from either <code>True</code> or <code>False</code> in both possible OMB scenarios. . . . .    | 42  |
| 2.6 | Classical unambiguous discrimination strategy to solve the HLPE in two questions. . . .                                                                        | 44  |
| 3.1 | Model performance evaluation measures. . . . .                                                                                                                 | 117 |
| 3.2 | Optimal parameters obtained for different fitted models. . . . .                                                                                               | 127 |
| 4.1 | CNOT counts for decomposing a general $n$ -qubit unitary using various methods. . . . .                                                                        | 130 |
| 4.2 | Heuristic function portfolio for ZX-diagram simplification . . . . .                                                                                           | 150 |
| 4.3 | CNOT count comparison across different gate decompositions and compilers. . . . .                                                                              | 157 |
| 5.1 | CNOT counts and number of equivalent circuits for Fredkin and Toffoli gates across five different connectivity scenarios. . . . .                              | 172 |
| 5.2 | CNOT counts obtained by different methods for Fredkin and Toffoli gate circuits. . . .                                                                         | 172 |

# List of Algorithms

|   |                                                                                               |     |
|---|-----------------------------------------------------------------------------------------------|-----|
| 1 | Quantum Circuit Normalization . . . . .                                                       | 141 |
| 2 | Entangling-gate structure classification . . . . .                                            | 142 |
| 3 | Metaheuristic best-first search for circuit optimization via dual representation feedback . . | 148 |
| 4 | Entangling-Gate Structure-based Equivalent Circuit Averaging (ECA) protocol . . . . .         | 181 |

## List of source codes

|   |                                                                              |     |
|---|------------------------------------------------------------------------------|-----|
| 1 | Minimum-error state discrimination via SDP. . . . .                          | 86  |
| 2 | OpenQASM circuit instruction for the square-root measurement. . . . .        | 89  |
| 3 | OpenQASM circuit instruction decomposing controlled- $\mathcal{V}$ . . . . . | 164 |

# Chapter 1

## Introduction

### 1.1 Framework and motivation

#### 1.1.1 Theoretical advances in quantum algorithms

The evolution of quantum algorithms is marked by a series of theoretical breakthroughs that have expanded our understanding of computation. The earliest motivations for quantum computation were grounded in the idea of simulating quantum mechanical systems more efficiently than classical computers could [1, 2]. Shor’s algorithm soon followed [3, 4], demonstrating a dramatic exponential speedup for integer factorization, a problem that had resisted efficient classical solutions for centuries. Grover’s search algorithm [5] was another major milestone, offering a quadratic speedup for unstructured search tasks over classical exhaustive search, with applications to a wide range of computational problems.

These initial breakthroughs spurred interest in identifying new problems where quantum resources could be harnessed effectively. Anchored on these three core pillars — quantum simulation, Shor’s algorithm, and Grover’s search — the broader landscape of quantum algorithms soon grew remarkably diverse. Crucially, however, most subsequent developments can be traced back to one or more of these paradigms, building upon their conceptual and technical frameworks. Algorithms leveraging the Quantum Fourier Transform (QFT) [6] as their computational engine, for instance, encompass Quantum Phase Estimation (QPE) [7] — a foundational subroutine for eigenvalue extraction — Shor’s techniques for number-theoretic problems like factorization and discrete logarithms, as well as extensions to non-periodic tasks such as the Bernstein-Vazirani problem.

The second category generalizes Grover’s amplitude amplification into a versatile subroutine for accelerating unstructured search variants [8, 9], while discrete-time quantum walk search exploits similar

principles of constructive interference [10]. In the third family, Hamiltonian simulation encodes computational problems into controlled dynamical evolution. This includes trotterization [11, 12, 13], adiabatic algorithms [14], continuous-time quantum walks [15, 16], linear combinations of unitaries (LCU) [17, 18], and the Harrow-Hassidim-Lloyd (HHL) algorithm [19], which, besides employing the QFT, maps matrix inversion to simulating an engineered Hamiltonian. In a related framework, quantum signal processing (QSP) [20] leverages polynomial approximations to efficiently synthesize transformations of Hamiltonians. For years, this tripartite division — simulation, QFT, and amplitude manipulation — appeared irreducible. Yet recent work has revealed a *grand unification*: these seemingly distinct approaches are now understood as specializations of a single meta-algorithm, the Quantum Singular Value Transformation (QSVT) [21, 22].

The QSVT framework generalizes a large class of quantum algorithms by recasting them as structured transformations of singular values. It starts from a unitary operator that embeds a matrix of interest via block encoding, a technique advanced by methods like qubitization [23], which constructs optimal encodings for Hamiltonian simulation. Then, it interleaves applications of this operator with carefully chosen rotations, effectively applying a polynomial through controlled transformations to the singular values of the embedded matrix. This procedure can implement matrix inversion, matrix exponentiation, and other transformations through a single powerful construct, subsuming many previously distinct methods. This consolidation, and the optimality of QSVT in terms of query complexity, not only simplifies algorithm design but also clarifies the origins of *quantum advantage*: whether via spectral manipulation (Shor), coherent amplification (Grover), or dynamical evolution (simulation), speedups emerge from tailored interference over singular value spectra.

The exploration of quantum machine learning further diversified the field. The HHL algorithm had gained prominence by producing a quantum state encoding the solution to  $Ax = b$  with logarithmic dependence on the system's dimension, contingent on assumptions such as data sparsity and efficient input preparation. Many data-processing and machine-learning tasks involve linear algebra, leading to initial speculation about broad quantum advantages in a number of publications. However, subsequent *dequantization* results [24, 25] exposed the inadequate assumptions in this quantum rush and demonstrated that several quantum-inspired<sup>1</sup> techniques can be efficiently simulated classically under realistic data-access assumptions, reducing the usefulness of their natively quantum implementations. This discovery redefined expectations about quantum advantages, limiting the hope of exponential speedups to specific settings, such as structured or sparse data cases [26, 27]. On the other hand, while claims

---

<sup>1</sup>A quantum-inspired algorithm is a classical algorithm that emulates or approximates a quantum algorithm.

of exponential speedups in classical machine learning remain highly nuanced and conditional, a more promising direction appears to be processing data originating from quantum-mechanical processes [28].

The QSVT framework played a crucial role in clarifying the structural features that distinguish quantum from classical methods. Subsequent work has demonstrated that the QSVT can be efficiently dequantized for low-rank matrices [29] and that, for sparse matrices, while certain quantum chemistry problems can be solved classically with constant precision, attaining inverse-polynomial precision remains BQP-complete — suggesting that the quantum advantage in chemistry may primarily arise from the enhanced precision achievable in the quantum setting [30]. Despite these theoretical advances, the practical implementation of QSVT is hindered by several challenges, including the high circuit depth required for accurate polynomial approximations, the complexity of computing the necessary phase factors, the substantial number of ancilla qubits needed for block encoding, and its sensitivity to noise in multi-qubit operations. Optimizing these components has become an active research area, with subsequent innovations such as generalized quantum signal processing (GQSP), among others [31, 32, 33].

Notwithstanding the unifying role of the QSVT framework, not all theoretical advances in quantum algorithm development have been subsumed under it. In certain instances, problem-specific adaptations can lead to algorithmic approaches that improve upon direct, generic implementations of QSVT. For instance, recent work has shown that carefully tailored mappings from classical dynamical systems to quantum evolution can yield exponential quantum speedups for specific tasks, such as the simulation of large networks of coupled classical oscillators [34]. Moreover, several quantum procedures lie outside the formalism of QSVT. For instance, the measurement-based shadow tomography protocol [35], while not strictly a quantum algorithm, has become an important tool for efficiently and scalably analyzing quantum states and operations when exhaustive characterization is required. At a more foundational level, recent complexity-theoretic work has further shown that quantum advantage need not rely on the presence of exploitable algebraic or spectral structure at all, establishing separations between classical and quantum computation for certain search and sampling problems even relative to uniformly random oracles [36].

Beyond rigorously characterized algorithms, the field has also embraced heuristic approaches designed for near-term, noisy hardware, commonly referred to as Noisy Intermediate-Scale Quantum (NISQ) hardware [37]. Hybrid variational quantum algorithms [38], most notably the Variational Quantum Eigensolver (VQE) [39, 40] and the Quantum Approximate Optimization Algorithm (QAOA) [41], have emerged primarily in response to the constraints imposed by noise, as well as the limited circuit depth and width. These methods leverage classical optimization to refine quantum circuit *ansätze*, mitigating hardware imperfections while still harnessing entanglement. Despite challenges such as selecting an appropriately

expressive ansatz, optimizing noisy high-dimensional functions prone to barren plateaus, and managing the increasing number of required measurements [42, 43], the potential applicability of these methods beyond the NISQ regime and into the fault-tolerant era remains an open question, particularly given their ability to exploit the high-dimensional state space of quantum processors.

Similarly, heuristic methods based on quantum annealing [44, 45] offer a more practical approach to adiabatic quantum computing by relaxing the requirements of the adiabatic theorem [46] and permitting diabatic transitions while operating as an open system. Quantum annealing is naturally suited for analog implementations, whereas its digitized variants on gate-based quantum computers exhibit distinct trade-offs. Other heuristic strategies, such as Quantum Imaginary Time Evolution (QITE) [47], similarly approximate ground states under less restrictive conditions. In contrast to heuristics, alternative non-deterministic approaches were also developed for NISQ processors. Although such algorithms may not succeed on every run and can require multiple repetitions, they guarantee that, when post-selection is successful, the desired state is produced exactly [48, 49]. Quantum Phase Estimation has also been adapted for near-term quantum devices by reducing circuit depth and ancilla qubit requirements in exchange for increased measurement repetitions [50]. This method leverages information from the entire output distribution of phases to improve estimation accuracy and can be readily applied to both Hamiltonian eigenstates and general state inputs. Additionally, it provides resilience to local noise during phase estimation, a feature that has been crucial for its validation on physical quantum processors contrary to other quantum phase estimation proposals only studied with theoretical noise models [50].

These heuristics and non-deterministic methods coexist with problem-specific quantum models such as boson sampling, which has provided an alternative route for exploring quantum advantage, particularly within complexity theory [51]. By exploiting the statistical properties of non-interacting photons in linear-optical networks, boson sampling aimed to provide a definitive example of a quantum system outperforming classical computers in specialized sampling tasks. In parallel, advances in fault-tolerant architectures and studies of complexity classes like QMA (Quantum Merlin-Arthur) [52] have further refined the theoretical roadmap.

Together, these milestones reflect a field maturing from fragmented breakthroughs toward integrative principles and applications. From Shor and Grover's pioneering algorithms to QSVT's unifying lens and the pragmatic rise of heuristics, the evolution of quantum algorithms [53, 54] underscores a deepening dialogue between abstract theory and physical implementation, an interplay that continues to define the frontier of quantum computation and situates the theoretical and practical contributions in this thesis.

### 1.1.2 Experimental challenges and recent achievements

Despite these rich theoretical advancements, practical realization of the promised super-polynomial speedups beyond classical capabilities remains highly challenging. NISQ devices operate with a limited number of qubits and face scalability constraints, requiring a delicate balance between coherence times, gate fidelities, and qubit connectivity [55]. Both near-term and future quantum hardware must contend with two major obstacles: *decoherence*, which limits the time window within which quantum circuits remain operational before the accumulation of incoherent errors, and *coherent errors*, which often arise from imperfect control operations. The relative difficulty of overcoming each of these limitations remains uncertain due to an inherent trade-off between increasing circuit depth — which exacerbates decoherence effects — and introducing parallel operations to reduce it — which can amplify coherent noise.

A wide range of technologies is currently under development for implementing quantum processors, including superconducting circuits [56], trapped-ion systems [57], ultracold neutral atom arrays [58], photonic processors [59], electron spin-based platforms [60, 61], and exploratory attempts to develop elusive topological qubits [62]. Notably, while industry leaders and startups alike are pursuing these diverse implementations, two-qubit gate operations remain a critical performance benchmark, as their cumulative errors can significantly impact overall computational accuracy. Reported error rates for these entangling gates can vary considerably. Superconducting processors such as IBM’s Heron R2 and Google’s Willow achieve two-qubit gate infidelities on the order of  $\sim 10^{-3}$  [63, 64], similar to Quera’s neutral atom arrays [65], IonQ’s Forte ion trap system [66] and PsiQuantum’s manufacturable platform for photonic quantum computing [67], while Quantinuum’s H1-1 [68] and Oxford Ionics’ ion trap processors [69] have reported two-qubit gate infidelities reaching  $\sim 10^{-4}$ .

While physical error rates are a key metric, platform performance comparisons are inherently non-trivial due to additional differences in connectivity, operational speed, and coherence characteristics [70]. Trapped-ion systems provide all-to-all connectivity within a single trap, simplifying multiqubit gates and error correction, whereas superconducting circuits typically include only nearest-neighbour couplings, requiring additional SWAP operations for long-range interactions. Superconducting qubits, however, operate orders of magnitude faster than ions, enabling higher sampling rates despite shorter coherence times. Photonic platforms operate in the fundamentally different regime of measurement-based quantum computing, which operates through the continuous generation and consumption of short-lived entangled resource states rather than through the preservation of long-lived physical qubits [71]. In this paradigm, logical information is propagated via adaptive measurements and classical feedforward, allowing quantum computation to be extended in principle to arbitrary durations without persistent coherence of individual

qubits. Although optical processes support fast operations, two-qubit interactions are typically probabilistic and incur substantial resource overhead. As a result, trade-offs among sampling rate, connectivity, coherence, and computational model critically determine effective performance beyond raw error rates.

While ongoing efforts to scale qubit counts and improve gate fidelities dominate hardware development, current quantum processors already enable experimental tests of quantum theory with unprecedented control and precision. For example, some experiments have sought to validate entanglement on these processors through Mermin inequalities [72], probe the consistency of quantum histories [73], and test the quantumness of correlations in variational algorithms [74, 75]. The ability to prepare, manipulate, and measure quantum states of a single system at millikelvin temperatures, coupled with rapid, deterministic measurement sequences, provides a unique window into microscopic statistical properties, from complementarity principles [76] to Sorkin's and Peres' tests [77]. However, not all studies have been without critique: methodological rigor, particularly in quantifying statistical errors and environmental noise, has sometimes been a concern in experiments claiming to test quantum foundations.

The pursuit of quantum computational advantage<sup>2</sup> has driven remarkable experimental progress in NISQ devices. At its core, the objective is to challenge the Extended Church-Turing Thesis (ECT), which posits that all physically realizable computational processes can be efficiently simulated by a classical Turing machine, at most incurring polynomial overhead [78]. Since quantum mechanics allows for fundamentally different computational primitives, it provides a possible counterexample to the ECT in principle. The goal of quantum supremacy is to demonstrate such a violation in practice: namely, to establish that certain quantum computations can outperform the best-known or conjectured classical algorithms for specific, even if contrived, problems. While these experiments do not yet aim to construct full-fledged, fault-tolerant, universal quantum computers, they seek to exhibit quantum speedups in a robust and statistically verifiable manner.

A pivotal milestone in this pursuit was achieved by Google's Sycamore 53-qubit superconducting processor, which demonstrated a sampling task in 200 seconds that was initially estimated to require 10000 years on classical supercomputers [79]. This experiment employed random circuit sampling with cross-entropy benchmarking, establishing the principle that quantum systems could outperform classical counterparts for specific computational tasks. Although subsequent analyses indicated that more optimized classical algorithms could reduce the computational gap, these developments highlighted a crucial trend: as the number and quality of qubits increase, classical simulations face exponentially growing demands on both memory and processing resources.

---

<sup>2</sup>Also called quantum supremacy.

Shortly thereafter, a 56-qubit processor extended superconducting quantum advantage by implementing deeper random circuits [80], while Gaussian boson sampling with 76 [81] and 255 [82] detected photons demonstrated photonic quantum advantage. These experiments pushed system sizes well beyond previous benchmarks, showing that photonic and superconducting platforms alike could generate statistically verifiable output distributions at scales that rendered naive classical methods infeasible.

As the boundary between quantum and classical performance continues to shift due to algorithmic innovations, the role of Quantum Error Mitigation (QEM) in near-term devices becomes increasingly critical. A recent experiment demonstrated increased quantum utility using a 127-qubit superconducting processor to run quantum circuits 60 layers deep, with approximately 2800 two-qubit gates, to solve 2D Ising model simulations that are intractable by naive classical methods [83]. By combining QEM with problem-specific hardware optimizations, they obtained accurate results suggesting a viable path toward practical quantum advantage before fault-tolerance. Together, these achievements suggest that error reduction techniques might provide *quantum utility* on specific computational tasks in the near future [84].

While QEM [85, 86, 87] does not eliminate errors as Quantum Error Correction (QEC) aims to do, it provides an essential tool for extracting reliable computational results from noisy quantum hardware. QEM addresses systematic deviations in expectation values arising from imperfect quantum computations. Ideally, a quantum circuit would generate an output state  $\rho_0$ , yielding an expectation value  $\text{Tr}[O\rho_0]$  for an observable  $O$ , but real hardware instead produces a noisy state  $\rho$ , introducing systematic deviations. QEM strategies work by manipulating measurement outcomes from multiple circuit executions to construct an estimator  $\hat{O}$  that more closely approximates the ideal expectation value. A key performance measure for such estimators is the *mean square error* (MSE), defined as  $\text{MSE}[\hat{O}] = \mathbb{E}[(\hat{O} - \text{Tr}[O\rho_0])^2]$ , which decomposes into bias and variance components:

$$\text{MSE}[\hat{O}] = \text{Bias}[\hat{O}]^2 + \text{Var}[\hat{O}], \quad (1.1)$$

where  $\text{Bias}[\hat{O}] = \mathbb{E}[\hat{O}] - \text{Tr}[O\rho_0]$ , and  $\text{Var}[\hat{O}] = \mathbb{E}[\hat{O}^2] - \mathbb{E}[\hat{O}]^2$ . While statistical noise diminishes with increased circuit repetitions, residual systematic errors persist as fundamental limitations targeted by QEM methods. However, by diminishing the bias, the error-mitigated estimator will also usually increase variance. As a result, the resource cost of error mitigation is quantified through the *sampling overhead*, given as the ratio of the number of shots  $S$  required to achieve a target shot noise level  $\nu$  between the mitigated and the noisy estimators [88]. For a noisy estimator  $\hat{O}_\rho$  with variance  $\text{Var}[\hat{O}_\rho]$  and its error-

mitigated counterpart  $\hat{O}_{\text{em}}$  with variance  $\text{Var}[\hat{O}_{\text{em}}]$ , this overhead is expressed as

$$C_{\text{em}} = \frac{S_V(\hat{O}_{\text{em}})}{S_V(\hat{O}_\rho)} = \frac{\text{Var}[\hat{O}_{\text{em}}]}{\text{Var}[\hat{O}_\rho]}. \quad (1.2)$$

This ratio reflects the increased sampling required to maintain precision when mitigating errors, a direct consequence of bias reduction. In fact, it has been shown that  $C_{\text{em}}$  scales exponentially with the total number of gates in a circuit, each with an error probability  $\varepsilon$  [89, 90, 91]. However,  $\varepsilon$  appears in the exponent of this dependency on system size, making QEM effective for circuit sizes of  $\mathcal{O}(\frac{1}{\varepsilon})$  [92]. Therefore, as gate errors continue to decrease, improving NISQ device performance and enabling the QEC era, QEM is expected to remain relevant as part of a broader trade-off between mitigating errors via additional measurements and correcting them with extra qubits<sup>3</sup> [93]. Effective QEM protocols must optimize bias suppression against variance amplification while preventing an explosive increase in resource demands—a critical balance in the pursuit of practical quantum advantage.

Despite providing a means to mitigate errors in expectation values without modifying hardware, the effectiveness of QEM techniques depends on the nature of the underlying noise. Many strategies are well-suited for addressing stochastic errors, which arise from random fluctuations and can be statistically suppressed with sufficient averaging. However, a distinct challenge emerges when dealing with *coherent errors*—systematic deviations in quantum operations that accumulate predictably and do not average out over repeated measurements. These can lead to interference effects that bias computational results, making their mitigation particularly critical. In fact, their worst-case error rate scales with the square root of the average error rate, leading to a rapid degradation in circuit fidelity if left uncorrected [94, 95]. Consequently, effective QEM techniques must not only suppress stochastic noise but also prevent the accumulation of coherent errors during circuit execution.

To address this issue, several methodologies have been proposed that modify or augment the set of single-qubit gates applied in a given circuit. Examples of include dynamical decoupling [96], which inserts idle-time pulses to counteract decoherence, composite pulse sequences with arbitrarily high accuracy [97], and various randomization-based strategies such as Pauli twirling [98], Pauli frame randomization [99], and randomized compiling [100]. Another class of techniques involves synthesizing approximate unitaries that differ slightly from the target but collectively average towards the desired transformation [101, 102]. Equivalent Circuit Averaging (ECA) [103] follows a similar rationale: it relies on executing multiple logically equivalent circuits and aggregates their measurement statistics to convert systematic

---

<sup>3</sup>It is worth noting that similar trade-offs—either adding auxiliary qubits or reducing them by compensating with increased sampling—also appear in various quantum protocols. For example, see Fig. 3 in Ref. [50].

errors into effectively stochastic noise.

In addition to modifying the quantum circuit itself to suppress or randomize coherent errors, an alternative strategy relies on *post hoc* analysis of measurement outcomes to mitigate errors after circuit execution. These methods are particularly effective for addressing stochastic noise, where the errors can be more effectively estimated or extrapolated. Examples include probabilistic error cancellation [104, 86], zero-noise extrapolation [86], and tensor-network error mitigation [105, 93]. While some of these techniques can also mitigate coherent errors under certain conditions, their effectiveness or sampling overhead generally worsens if the coherent errors create strong correlations or complex error structures.

While quantum error mitigation techniques have enabled significant progress in reducing errors on near-term devices, the pursuit of scalable quantum computation through QEC continued. By encoding logical qubits into entangled states of multiple physical qubits, QEC aims to suppress errors below a critical threshold where logical error rates decrease as code distance increases [106, 107]. Recent experimental breakthroughs have demonstrated the feasibility of this approach, marking a major step toward fault-tolerant quantum computation.

In particular, the Google Quantum AI team reported a breakthrough in scaling a surface code logical qubit from distance-3 (17 physical qubits) to distance-5 (49 qubits) in superconducting circuits, showing for the first time that increasing the code distance can indeed lower the overall logical error rate from 3.028% to 2.914% [108]. Building on this, a subsequent experiment introduced an enhanced superconducting processor (Willow) implementing distance-7 logical qubits with 101 physical qubits, demonstrating exponential error suppression through improved qubit coherence and advanced real-time decoding via a Google Deepmind-developed machine learning algorithm [109]. This achieved a 2.4-fold error rate reduction per added qubit, surpassing the critical threshold for scalable error correction [63].

Parallel advances in neutral-atom architectures achieved scalable logical processor operation through dynamic reconfiguration. A 48-logical-qubit processor was demonstrated using 228 physical qubits in reconfigurable Rydberg atom arrays, executing error-corrected algorithms with real-time parity checks [110]. This system enabled simultaneous operation of multiple logical gates across the largest logical qubit ensemble to date, showcasing the advantage of atomic repositioning for maintaining entanglement during computation. Both achievements highlight not only the viability of error-corrected qubits but also the promise of multiple architectures for implementing large-scale fault-tolerant quantum computers.

These experimental milestones illustrate the steady progression of quantum hardware from small, noisy prototypes toward increasingly capable platforms. Foundational tests of quantum mechanics, demonstrations of quantum advantage, and the development of error-mitigation and fault-tolerant schemes alike

highlight both the potential and the constraints of practical quantum computation. This thesis situates its contributions within this evolving experimental landscape, acknowledging platform limitations to guide the design of new methods and algorithms, while also performing extensive experimental validations.

### 1.1.3 Optimizing circuits, discovering algorithms

The development of entirely new and conceptually distinct quantum algorithms remains relatively rare. This scarcity can be largely attributed to our limited intuitive grasp of quantum phenomena, which stands in stark contrast to our extensive experience in the classical world. The primary obstacle in developing new quantum algorithms, therefore, is arguably not the structure of quantum mechanics itself but rather our difficulty in conceptualizing profoundly non-classical manipulations of information — an essential prerequisite for designing algorithms that outperform classical methods. True breakthroughs may require cross-pollination with unexpected fields to reframe problems in quantum-native terms.

Quantum circuits are at the heart of quantum algorithms. In the quantum logic-gate model [111], they directly encode the exact sequence of operations that define the computational processes. As a result, the development of classical procedures to optimize them, or machine intelligence tools capable of uncovering new structures within them, can open up a path to the discovery of new quantum algorithms, unveil deeper physical principles, and foster the development of new mathematical frameworks.

A well-established result in quantum circuit synthesis [112] shows that, for a universal gate library consisting of CNOT and single-qubit gates, the minimal number of CNOT gates required in a circuit to implement *any*  $n$ -qubit unitary is

$$N_{\text{theo}}(n) = \left\lceil \frac{1}{4} (4^n - 3n - 1) \right\rceil. \quad (1.3)$$

This provides a theoretical lower bound on the number of CNOT gates required to parameterize the *full* set of  $n$ -qubit unitaries. At the same time, it establishes a constructive upper limit on the worst-case gate count when implementing a specific unitary, ensuring that a circuit with at most  $N_{\text{theo}}$  CNOT gates always exists. However, many structured unitaries — such as diagonal, Clifford, or permutation unitaries, among others — can be realized with significantly fewer two-qubit gates, making this result only a first reference for circuit optimization strategies targeting specific applications.

For efficient algorithms and effective circuit discovery, the number of CNOT gates required to implement a given  $n$ -qubit unitary corresponding to a specific circuit primitive should ideally be much lower than the worst-case theoretical bound,  $N_{\text{theo}}(n)$ . In fact, due to inherently physical or algorithmically imposed

constraints, unitaries of interest can fortunately often be decomposed with sub-exponential numbers of two-qubit gates. Furthermore, when generalizing an algorithm across varying  $n$ , one is not just concerned with one  $n$ -qubit instance, but with a family of related unitaries that must scale efficiently with  $n$ .

To find one of these efficient circuits from scratch in an automated manner, search and optimization tools can be employed to generate shorter gate sequences that implement the desired  $n$ -qubit unitary [113, 114]. For the more relevant case of discovering algorithms requiring a scalable family of unitaries — an example being the SWAP-test with two multi-qubit state inputs — this approach extends to a two-step strategy. First, efficient decompositions are derived for a few small- $n$  instances of the unitary family. Then, through further analysis, these circuits are connected and extrapolated to rigorously generalize them into a closed-form or symbolic circuit template for systems of arbitrary size [115].

While the first approach focuses solely on decomposing an exact unitary operator, the second introduces greater generality by tackling broader unitary templates. Crucially, this requires developing new tools capable of identifying underlying circuit structures and extracting general templates. Ideally, this would be achieved through automated procedures — extending beyond previous manual approaches [115]. Introducing continuous parameters into the algorithm can further complicate this process.

Quantum algorithms, by their nature, are intimately tied to fundamental physical laws, so insights gained from circuit optimization and their complexity analysis contribute to a deeper understanding of quantum mechanics itself [116]. For instance, identifying more efficient quantum gates or alternative representations of entangled states could provide new perspectives on the nature of quantum interactions [117, 118]. Similarly, the discovery of more compact or symmetrically structured circuits may reveal underlying symmetries in quantum systems that were previously unrecognized [119].

From a mathematical standpoint, quantum circuits operate within a formal framework intersecting linear algebra, group theory, and category theory [120]. Optimizing them to find novel or more efficient decompositions of complex unitary operations into fundamental gates can inspire new mathematical techniques, particularly in areas such as tensor networks [121] and graph theory [122]. Additionally, machine learning methods for analyzing quantum circuits may evolve into broader algorithmic paradigms extending beyond quantum computing itself.

Quantum circuit optimization research has been primarily driven by practical motivations. By reducing circuit complexity and adapting decompositions to constrained graph topologies [123], circuit optimization shortens execution times and mitigates NISQ-era imperfections such as decoherence and noise. As quantum hardware advances, optimizing circuits remains crucial for improving practical performance. At the same time, new theoretical avenues for algorithmic development may emerge from circuit optimization

methods as well. This thesis contributes to this evolving paradigm, developing methods that both enhance practical circuit efficiency and inform the discovery of new quantum algorithms.

## **1.2 Dissertation outline**

### **1.2.1 Overview**

This dissertation advances the pursuit of practical quantum advantage by systematically addressing theoretical, experimental, and engineering challenges in quantum computing. While much of the existing research often focuses on isolated aspects of quantum technologies, this dissertation bridges these domains through the framework of the circuit model of quantum computing, unlocking new pathways toward both theoretical insights and practical implementation.

This work comprises different phases of design, optimization, and autonomous discovery of both general and applied quantum circuits, organized around four interconnected projects. First, a concrete problem is analyzed, and quantum protocols are derived to solve it by strategically leveraging quantum phenomena to provably outperform classical counterparts, providing a motivating instance of quantum advantage. The second project then shifts focus to the experimental testing of a fundamental quantum tenet on a superconducting quantum processor, analyzing results in detail to characterize the quantum features and limitations of the platform and assess the implications for quantum algorithms. The third project introduces a competitive framework for the automated discovery and optimization of quantum circuits, classifying their internal structures not only to minimize computational complexity but also to facilitate the discovery of algorithmic generalizations. Finally, these theoretical tools are applied to optimize two essential quantum primitives for a wide range of quantum algorithms and implement an error mitigation protocol validated on the same noisy platforms characterized earlier.

Different complementary perspectives converge to consolidate this dissertation. Theoretical proposals motivate experimental validation, while empirical insights in turn refine and redirect theoretical frameworks, thus spanning the full arc of the scientific method. Similarly, the interplay between conceptual abstractions and applications is a constant theme: foundational quantum principles inform the design of hardware-efficient circuits, and engineering constraints inspire new theoretical tools for circuit discovery and optimization. A commitment to practical constraints is maintained throughout this work, with qubit connectivity limitations addressed alongside a systematic treatment of noise, from its physical origins to mitigation strategies and its potential as a quantum advantage resource. No less significant is the integration of complementary approaches to developing novel quantum computing protocols by combin-

ing theory-grounded circuit design with classical optimization and artificial intelligence techniques. These methods not only enhance quantum circuit performance but also facilitate the automated discovery of new circuit structures, paving the way for autonomous quantum circuit design to complement traditional hand-crafted approaches. Multidisciplinary views also find convergence in this dissertation, drawing insights from fundamental physics and philosophy to computer science and engineering, thereby contributing to the exploration of new frontiers and the sustained ambition of scientific inquiry [124, 125].

### 1.2.2 Structure

This dissertation is organized into two main sections. The first part, *Quantum Computing: from theory to practice*, comprises two chapters and explores foundational aspects of quantum mechanics with far-reaching implications for information processing. This part begins with a fully theoretical chapter and proceeds with a second chapter that presents both theoretical and experimental results.

**Chapter § 2** The scarcity of novel examples of quantum advantage underscores the importance of beginning this dissertation by presenting a problem where a clear instance of fundamental quantum advantage has been identified. In this chapter, solution protocols are developed for an interactive logical problem — *The Hardest Logic Puzzle Ever*— where a solver engages in logical interactions with three unknown agents to infer knowledge from minimal data amid randomness and linguistic ambiguity. Although not a quantum algorithm in the strictest sense, the presented solutions demonstrate that quantum resources provide a computational advantage over classical approaches while offering novel insights into the empirical nature of logic. Since this problem was originally formulated within the philosophical logic community and the aim here is to take it from its abstract form into a physical instantiation, our exploration will carefully develop from foundational principles, introducing quantum mechanical concepts as they arise. This chapter presents original, unpublished work for which a publication is currently in preparation [126].

**Chapter § 3** After establishing how Quantum Physics gives rise to novel phenomena with no classical counterpart and demonstrating their surprising implications for information processing in the previous chapter, we turn to an experimental investigation of one such fundamental feature: the complementarity principle. While many of the distinctive aspects of quantum mechanics have been known for nearly a century, the newest generation of quantum computing technologies enables us to test its foundations with unprecedented precision, uncovering both fundamental insights and practical implications for technological

applications. Specifically, we introduce quantum circuits to implement symmetric two-way interferometers coupled to a which-path detector to test the visibility-distinguishability tradeoff using minimum-error state discrimination. Experiments are performed on a superconducting quantum processor, and results are exhaustively analyzed to detect any slight deviations from theoretical expectations. In doing so, this chapter is not merely about replicating experiments that trace back to the origins of quantum mechanics and have been performed countless times using a novel physical platform. Instead, it focuses on executing them with unprecedented precision, analyzing the data in extraordinary detail, and exploring its implications for quantum technologies. Notably, two-qubit gate errors play a fundamental role in explaining the experimental data, revealing limitations in reaching the full particlelike and wavelike statistics postulated by Bohr’s strong formulation of complementarity. These findings provide valuable insights for quantum circuit design and error mitigation strategies. This chapter incorporates work previously published in Ref. [76].

The second part of this dissertation focuses on *Circuit discovery and optimization*, an essential avenue for improving the practical performance of quantum computing and facilitate the discovery of novel quantum algorithms by leveraging increasingly abstract layers of machine intelligence. As in the first part, this section begins with a fully theoretical chapter before moving on to a second chapter that presents both theoretical and experimental results.

**Chapter § 4** While the first two chapters examined the implications of intrinsically quantum features in specific computational protocols, this chapter shifts focus to general quantum circuits. It presents a method for optimizing general quantum circuits across arbitrary qubit connectivity topologies while simultaneously classifying their entangling-gate structures. The proposed framework leverages a combination of previously developed ZX-diagram techniques and circuit-level optimization heuristics, performing a metaheuristic search over circuit instances categorized via a novel entangling-gate structure classification scheme. This classification guides the search through tailored policies, efficiently minimizing both the CNOT gate count and circuit depth while generating a large variety of classified equivalent circuits. Example executions are provided. Beyond its direct applicability in improving quantum circuit complexity and adapting to hardware constraints—essential for practical quantum computing—this method also offers novel tools to explore generalizations across multi-qubit gate families, driving progress in automated quantum algorithm discovery. For this reason, it is primarily designed for permanent circuit discovery rather than for optimizing runtime efficiency. This chapter also presents unpublished research, with a corresponding publication currently in preparation [127].

**Chapter § 5** In this final chapter, the method developed in the previous one is applied to the optimization of Fredkin and Toffoli gate decompositions on both fully connected and limited-connectivity architectures. These two fundamental quantum primitives have widespread applications across several quantum algorithms, making their efficient decomposition highly valuable in practice. In addition to optimizing these circuits for the case of adjacent qubit decompositions, the case of non-adjacent qubits is also addressed. A novel qubit reordering scheme is proposed, which saves one CNOT per SWAP and enables more efficient rerouting of long-range Fredkin and Toffoli gates, with broader applications to other quantum computing primitives. Since a diverse set of logically equivalent decompositions with similar circuit complexities is provided by the method, we also build on insights from the experimental evidence obtained in previous chapters and leverage them to propose a coherent error mitigation scheme based on equivalent circuit averaging. Experimental validation for this error mitigation procedure is then sought and confirmed through a SWAP-test experiment on a superconducting quantum processor. The results of this chapter have been previously published in Refs. [128, 129].

Finally, in § Conclusion, the dissertation concludes with a brief discussion summarizing the key findings. A list of publications associated with the work presented in this dissertation is provided in § Publications by the author.

# **Quantum computing: from theory to practice**

## Chapter 2

# Quantum advantage in the solution of the “Hardest Logic Puzzle Ever”

“One could caricature Quantum Information Processing as the science of turning Quantum conundrums into potentially useful applications.”

---

Gisin *et al*, 2002 [[130](#)]

### 2.1 Introduction

Logic lies at the foundations of mathematics, the physical sciences, and the systematic approach of science in building explanatory models of the natural world, based upon observation and experimental validation. The history and development of logic began long ago with Aristotle’s profoundly influential works collectively known as the *Organon*, meaning *instrument* in Greek. In the years since, numerous figures have made key contributions to the field, including, but not limited to, Euclid, Leibniz, Boole, Frege, Cantor, Peano, Hilbert, Russell, Whitehead, Wittgenstein, Gödel, Turing, Church, Rosser, Tarski, and Kripke.

Despite the millennia-long development of Aristotelian logic, the history of recreational logic is much more recent, with Lewis Carroll being the unavoidable name in its origins [[131](#)]. In his book *The Game of Logic*, published in 1886, he presents several puzzles that apply the basic principles of Aristotelian categorical syllogisms in the form of a game designed for a non-academic audience [[132](#)]. A *syllogism* is a form of deductive reasoning where the objective is to draw a valid conclusion from two given premises by examining the categorical relationships between subjects and predicates. In his subsequent book, *Symbolic Logic*, Carroll presents more complicated puzzles with more than two categorical premises [[133](#)].

A noteworthy example of a more recent puzzle similar to Carroll's, is the so-called *Zebra Puzzle*, whose earliest appearance dates back to 1962 and has since gained widespread popularity [134]. This and other similar problems can be thought of as an extension of the traditional categorical syllogism with multiple interdependent premises imposing constraints that must be simultaneously satisfied in order to deduce a non-contradictory set of relationships that complete the puzzle. In fact, this riddle has also been used to evaluate algorithms for constraint satisfaction problems and logic programming [135, 136].

Logic riddles of a different kind, with characters who either tell the truth or lie, have been recorded since at least 1931, when philosopher Nelson Goodman anonymously published a “daily brain-teaser” on the front page of the *Boston Post* [137]. Raymond Smullyan later took this form of puzzle to a high level of sophistication, creating several and far more intricate riddles involving multiple speakers who are either liars or truth-tellers. These figures, whom he respectively referred to as *knaves* and *knight*s, along with others who sometimes lie and other times tell the truth (*normals*), became well-known through his numerous works on recreational logic, including [138, 139, 140], among others.

The history of logic puzzles reflects the history of logic itself. Just as Lewis Carroll used puzzles to introduce fundamental principles of Aristotelian logic, Raymond Smullyan employed his knights and knaves riddles to create complex problems in propositional logic, involving combinations of entire propositions through logical connectives such as conjunction, disjunction, negation, implication, and equivalence. In this art, Smullyan even elucidated highly complex concepts in mathematical logic, which extends beyond Aristotelian and propositional logic with greater expressiveness, explaining concepts such as self-reference and Gödel's groundbreaking incompleteness theorems using knights and knaves [141, 142].

By the end of the century, philosopher and logician George Boolos had published the so-called *Hardest Logic Puzzle Ever* (HLPE), first in the *La Repubblica* newspaper, in 1992, and later in a 1996 article in *The Harvard Review of Philosophy* [143]. Rooted in the tradition of logic puzzles involving knights, knaves and normals, here is the puzzle as Boolos formulated it in natural language:

(HLPE<sub>sem</sub>) “Three gods A, B, and C are called, in some order, *True*, *False*, and *Random*. *True* always speaks truly, *False* always speaks falsely, but whether *Random* speaks truly or falsely is a completely random matter. Your task is to determine the identities of A, B, and C by asking three YES-NO questions; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language in which the words for YES and NO are ‘da’ and ‘ja’, in some order. You do not know which word means which.” [143]

The allure of the HLPE lied in its claim to be the hardest of its kind. Boolos credited Smullyan for devising the initial version of the puzzle and computer scientist John McCarthy for the final twist of not

knowing which of “da” or “ja” mean YES. He also provided additional clarifications, here paraphrased:

- (B1) more than one question can be directed to the same god,
- (B2) questions are permitted to depend on previous answers,
- (B3)  $\mathfrak{R}\text{andom}$  decides whether to speak truly or falsely with equal probability before evaluating the answer to the question he was addressed, and
- (B4)  $\mathfrak{R}\text{andom}$  always answers “da” or “ja” when asked any YES-NO question.

The subscript “*sem*” in  $(\text{HLPE}_{\text{sem}})$  stands for “*semantic*,” and its significance will become apparent later in our discussion, as its introduction became necessary only in subsequent analysis of the HLPE.

Alongside presenting the puzzle, Boolos provided a complete solution comprised of three straightforward, *truth-functional* and *non-subjunctive* logical inquiries to the gods making decisive use of the biconditional logical connective<sup>1</sup>. The procedure involved posing the first question to A, “Does ‘da’ mean YES iff (you are  $\mathfrak{T}\text{rue}$  iff B is  $\mathfrak{R}\text{andom}$ )?” . If the response to this query was “da,” the second question would then be posed to C, or else to B, which read: “Does ‘da’ mean YES iff Rome is in Italy?” . Lastly, the third question would be directed to the same god who was asked the second question: “Does ‘da’ mean YES iff A is  $\mathfrak{R}\text{andom}$ )?” .

While being the most challenging of its kind is certainly a questionable assertion, the  $(\text{HLPE}_{\text{sem}})$  became more than a recreational puzzle and attracted great interest also in academia, particularly in the philosophy and logic communities. Five years after Boolos’ publication in the *Harvard Review of Philosophy*, Tim Roberts published a different solution to the puzzle [144]. The distinctive characteristic of his three questions lies in the incorporation of a query  $\mathcal{Q}'$  within the actual question  $\mathcal{Q}$  taking the form of  $\mathcal{Q}$ : “If I asked you  $\mathcal{Q}'$  would you say ‘da’?” .

This strategy requires the questioned god to first compute the truthful answer to  $\mathcal{Q}'$ , act over the result with its logical personality<sup>2</sup>, and finally translate the result into his “da”/“ja” language. This intermediate result then serves as input into the evaluation of an answer to the actual question  $\mathcal{Q}$ , which proceeds similarly to the evaluation of  $\mathcal{Q}'$ . Roberts’ solution, therefore, asks the addressee to simulate his own answer to a given question, gaining an advantage by requiring him to act twice with his personality operator during the functional evaluation of the query. In contrast, the solution Boolos offered does not use nested questions, and the functional evaluation of a god’s logical character is performed only once.

An enigma posed in natural language can create irresistible appeal, but it also risks leaving ambiguities

---

<sup>1</sup>The biconditional operator, denoted as  $\leftrightarrow$  or iff, which stands for “if and only if,” indicates that two propositions  $P$  and  $Q$  are logically equivalent. This means  $P \leftrightarrow Q$  evaluates to TRUE exactly when both  $P$  and  $Q$  are TRUE simultaneously or both are FALSE simultaneously.

<sup>2</sup>By *logical personality*, or *character*, we refer to either keeping the truthful Boolean value unaltered, in case the god is replying truthfully, or negating it, if he is replying falsely. This has also been referred to as the god’s *nature* in previous literature.

in its formulation. In 2008, Rabern and Rabern took advantage of that [145]. First, they noted that Boolos' intention with condition (B3) was unclear, as taken literally, it implies that *Random* replies to some questions in a predictable rather than random manner. For instance, if a question is posed to which both *True* and *False* provide the same answer, *Random*'s reply is deterministic if, by (B3) we assume that his random decision is just about acting as a liar or a truth teller. This insight led Rabern and Rabern to conclude that the (HLPE<sub>sem</sub>) could be solved using a trivial questioning strategy. However, their observation could not enable them to solve the puzzle in less than three binary output questions.

As a result, they secondly considered relaxing the requirement for a binary answer output by including a third option of “no-answer”. Their exploit was to assume the gods could represent and evaluate self-referential questions<sup>3</sup>, hence designing one that is unsatisfiable in some cases and from which the only outcome deemed acceptable is for the god not to answer. By determining whether the response is “da”, “ja” or “no-answer”, they contributed a two-question solution to the (HLPE<sub>sem</sub>), with Boolos' (B3) condition intact. It is important to emphasize, however, that such a solution is only possible because we have three potential outputs per question, which can accommodate nine symbol combinations capable of discriminating among all six possible permutations of the gods.

Rabern and Rabern then proposed amending condition (B3) to make the (HLPE<sub>sem</sub>) harder in the face of their own self-referential question exploit. Their suggestion was to replace condition (B3) with

(B3\*) *Random* replies “da” or “ja” with equal probability, independent of the question received.<sup>4</sup>

A revised formulation of the HLPE, now supplemented by conditions (B1), (B2), (B3\*), and (B4), therefore became necessary to ensure logical accuracy and preserve the hardness spirit of original puzzle:

(HLPE<sub>syn</sub><sup>2</sup>) “Three gods A, B, and C are called, in some order, *True*, *False*, and *Random*. *True* always speaks truly, *False* always speaks falsely, but whether *Random* answers ‘da’ or ‘ja’ is a completely random matter. Your task is to determine the identities of A, B, and C by asking three YES-NO questions; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language in which the words for YES and NO are ‘da’ and ‘ja’, in some order. You do not know which word means which.” [145]

The subscript “syn” in (HLPE<sub>syn</sub><sup>2</sup>) stands for “syntactic,” with the superscript number 2 indicating that *Random* is modeled as a binary random variable.

It is straightforward to check that the solutions provided by Boolos and Roberts remain valid with the amended puzzle (HLPE<sub>syn</sub><sup>2</sup>), as resilience to a completely random answer from *Random* had already been

<sup>3</sup>A self-referential question is a logical expression that refers to itself, e.g. “Are you going to answer ‘ja’ to this question?”.

<sup>4</sup>Paraphrased.

incorporated in their strategies. Additionally, Rabern and Rabern also provided a solution to the revised (HLPE<sup>2</sup><sub>syn</sub>), again resorting to self-referential questions and assuming the possibility of a “no-answer” outcome from *True* and *False*<sup>5</sup>. Their method can resolve the identity of the three gods in no more than three questions while also being able to complete the solution within just the first two questions in two out of the six possible arrangements of the gods. Therefore, they offer a strategy capable of solving the problem with certainty with probability  $1/3 \approx 33.3\%$  in the first two questions, with the complete solution guaranteed after three questions in case it does not succeed in the first two. Nevertheless, it is again important to belabor the obvious and remember that the two-question solution only becomes possible by assuming three possible answers from *True* and *False* gods, namely “da”, “ja”, and “no-answer”.

In 2010, Uzquiano achieved what Rabern and Rabern had proposed but did not: he provided two complete solutions to the (HLPE<sup>2</sup><sub>syn</sub>) in just two questions while continuing to rely on the same three-valued answer possibilities they had assumed for *True* and *False* [146]. First, without resorting to self-reference, the loophole exploited to obtain the enabling “no-answer” output from these two gods was based on the observation that they should not be capable of anticipating *Random*’s truly random replies. Because of that, if either *True* or *False* were asked what one of their neighboring gods would answer to a certain question, they would need to remain silent if this god was actually *Random*, given the impossibility of guaranteeing their answer would be logically valid. With this assumption about *True* and *False*’s responses when ignorant, a two-question solution became possible.

However, Boolos arguably implied with the *god* designation that this party should be thought of as omniscient. Therefore, Uzquiano considered next the scenario in which *True* and *False* have an oracular ability to predict *Random*’s answers before they occur, making it unacceptable for them to avoid answering the questions upon which the trick relied. Addressing this case, he was also able to achieve a two-question solution to (HLPE<sup>2</sup><sub>syn</sub>) by making use of self-reference once again.

Continuing the tradition of escalating difficulty, Uzquiano concluded by extending the (HLPE<sup>2</sup><sub>syn</sub>) to grant *Random* not only the possibility of replying “da” and “ja” but also a “no-answer”, with each outcome having an equal probability of occurrence. This placed the possible answers of the three gods on equal footing. Here is Uzquiano’s HLPE extension proposal:

(HLPE<sup>3</sup><sub>syn</sub>) “Three gods A, B, and C are called, in some order, *True*, *False*, and *Random*. *True* always speaks truly, *False* always speaks falsely, but whether *Random* speaks truly or falsely or whether *Random* speaks at all is a completely random matter. Your task is to determine the identities of A, B, and C by asking three YES-NO questions; each question must be put

---

<sup>5</sup>Remember that, according to (B3\*), *Random* always replies either “da” or “ja”.

to exactly one god. The gods understand English, but will answer all questions in their own language in which the words for YES and NO are ‘da’ and ‘ja’, in some order. You do not know which word means which.” [146]

To this formulation, modeling  $\mathfrak{R}andom$  as a ternary instead of a binary random variable, conditions (B1) and (B2) still apply, but both (B3\*) and (B4) are replaced by:

(B3\*\*) Whether  $\mathfrak{R}andom$  answers “da” or “ja” or whether  $\mathfrak{R}andom$  answers at all should be thought of as depending on the toss of a fair three-sided dice hidden in his brain: if the dice comes down 1, he doesn’t answer at all; if the dice comes down 2, he answers “da”; if 3, “ja”.

thereby removing the asymmetry in the answers of  $\mathfrak{R}andom$  from those of  $\mathfrak{T}rue$  and  $\mathfrak{F}alse$ . Uzquiano did not attempt to solve this puzzle.

In 2011, Wheeler and Barahona showed that  $(HLPE^3_{syn})$  eliminates the source of advantage being exploited by both the Rabern brothers and Uzquiano in their two-question solutions to  $(HLPE_{sem})$  and  $(HLPE^2_{syn})$ , respectively [147]. They used a simple information-theoretic argument to prove that three-value answer outcomes, “da”, “ja”, and “no-answer” for the three gods, could not make  $(HLPE^3_{syn})$  solvable in fewer than three questions, thereby presenting a solution in precisely three questions. In keeping with tradition, they also introduced their own variation of the puzzle, further departing from the original one with a new character they called  $\mathfrak{D}evious$ , for which a three-question solution was also provided drawing on modal logic structures.

At this point, one could criticize the informal manner in which the *Hardest Logic Puzzle Ever* was presented by Boolos for inviting multiple interpretations, leading to a range of solutions and modified versions that challenge each other in poorly defined terms. After all, how can we evaluate and compare different solutions on equal footing when they play by different rules? While this objection is not without merit, it is also worth considering that, inadvertently or not, it was precisely this original non-formal formulation that fostered such a wealth of insights in approaching the puzzle from different angles. Had the problem been more strictly defined in terms of permissible questions and answers, it is possible that we would not have gained as many perspectives into it as already seen.

The next chapter in this saga was also written in 2011 by Wintein, who introduced the nomenclature that we have been using to designate the different versions of the HLPE [148]. Wintein took account of the analysis by Wheeler and Barahona but did not lose heart, instead pointing out its incompleteness. He contended that, just as the Rabern brothers and Uzquiano, Wheeler and Barahona overlooked the possibility that there could be other answers to self-referential YES-NO questions besides “da”, “ja,” and

“no-answer.” He showed that a solution to (HLPE<sup>3</sup><sub>syn</sub>) was still possible in just two questions by adopting four acceptable answers for  $\mathcal{T}\text{rue}$  and  $\mathcal{F}\text{alse}$ .

To examine the issue raised by Wintein, let us first consider the self-referential question  $\lambda$  : “Is it the case that your answer to  $\lambda$  is NO?” , which in earlier versions of the HLPE would be an example of an undecidable question to which  $\mathcal{T}\text{rue}$  could not respond either YES or NO without contradicting his logical personality. Although questions of this type had been addressed before, Wintein drew attention to a different category of questions that had not been studied in previous solutions. For instance, when asked the question  $\tau$  : “Is it the case that your answer to  $\tau$  is YES?” ,  $\mathcal{T}\text{rue}$ ’s response would be valid by providing either YES or NO. This suggested that, while some self-referential questions are undecidable and have no solution, as they cannot be answered with neither YES nor NO, other questions have both solutions, meaning that they can arbitrarily be answered with both YES or NO.

Identifying this issue, Wintein considered that  $\mathcal{T}\text{rue}$  would reply “neither” when faced with an undecidable question, and “both” if faced with a question to which both YES and NO would be valid answers. Speaking falsely, on the other hand,  $\mathcal{F}\text{alse}$  would operate by responding “neither” when the truthful answer would be “both” and vice versa. Importantly, Wintein considered both gods would provide those two additional answer possibilities in English, while providing the YES and NO in their own “da”/“ja” language.

Since an answer repertoire with four options for  $\mathcal{T}\text{rue}$  and  $\mathcal{F}\text{alse}$  again suggested a harder variation of the HLPE in which  $\mathcal{R}\text{andom}$  would be modeled as a four-valued random variable, Wintein readily embraced this HLPE<sup>4</sup><sub>syn</sub> generalization. He proceeded to find a solution to HLPE<sup>4</sup><sub>syn</sub> in two questions, demonstrating that it could also solve (HLPE<sup>3</sup><sub>syn</sub>). Wintein went to great lengths to detail how his HLPE<sup>4</sup><sub>syn</sub> version adheres to Boolos exposition of the HLPE. Moreover, given that all prior solutions to the HLPE had been presented informally in natural language, he also extensively analyzed formal representations of his and all those solutions using Kripkean fixed point theories of truth.

There still remained one conspicuous issue, however: why do the gods have their own words for YES and NO but respond with “neither” and “both” in English? A more plausible interpretation of the HLPE would consider the gods also replying with “neither” and “both” in their own language, using the words “huh” and “duh” in some order. This particular version of the puzzle, named HLPE<sup>4\*</sup><sub>syn</sub>, was left as a challenge at the end of Wintein’s article.

All solutions to the HLPE up until 2012 assumed that the gods knew each other’s identities and were therefore able to respond when questioned about them. However, that year saw Liu and Wang challenging this assumption in their analysis of the puzzle [149]. In an extensive publication, they developed a

logical framework based on public announcement logic<sup>6</sup> that could specify agent *types*—i.e., their logical personality—in communication scenarios. Their framework accommodated dynamic epistemic models able to capture incremental knowledge acquired after each question in the HLPE.

Liu and Wang demonstrated the ability of their formal system to analyze and solve various of Smullyan’s knights and knaves puzzles, showing how different epistemic assumptions about the agents’ knowledge of each other can influence the solvability of the puzzles. Their paper formalized the (HLPE<sup>2</sup><sub>syn</sub>) while restricting it to only consider questions that are always answerable by “da” or “ja,” and subsequently verified Rabern and Rabern’s three-question solution. They then relaxed the implicit epistemic assumption that the gods know each other’s identities, and importantly proved that a version of the HLPE where the gods do not know each other’s identities is unsolvable under their formalization.

Also in 2012, a preprint that has not yet been formally verified and peer-reviewed [152, 153] has studied an additional modification to the puzzle more informally. In this publication, Novozhilov argues that it appears unnatural for the puzzle solver to know that the gods respond with the words “da” and “ja” without knowing their meaning. His interesting proposition is then to study variations of the (HLPE<sub>sem</sub>) and the (HLPE<sup>2</sup><sub>syn</sub>) where knowledge of the symbols used for communication is removed from the enunciation. In other words, we are only told that the gods answer YES or NO questions in an unknown language. This twist prevents us from referring directly to the words replied by the gods when formulating our questions, as previous solution strategies have done. In this publication, Novozhilov devises a three-question solution to his variation of the (HLPE<sub>sem</sub>), which also assumes that if you ask a question that the god cannot answer, he will answer NO. He also provides a demonstration that a solution to his variation of the (HLPE<sup>2</sup><sub>syn</sub>) is not possible in three questions with only two-valued outputs, but possible in two if self-referential questions are allowed.

You might think that philosophers would have it all figured out by now, but in 2017, Carnielli proposed yet another variant of the HLPE grounded in non-classical logic<sup>7</sup> [154]. He recognized Wintein’s interesting solution to the HLPE<sup>4</sup><sub>syn</sub> using expanded, non-binary answers and advanced theories of truth, but aimed to explicitly widen the logic behind the puzzle agents to encompass more advanced forms of reasoning and

---

<sup>6</sup>Public Announcement Logic (PAL) [150, 151] is a formal framework used to model the evolution of knowledge and belief in systems with multiple agents. PAL begins by considering all possible worlds that represent different combinations of facts compatible with the knowledge of the agents. When a public announcement is made, the agents update their knowledge by eliminating the worlds where the announcement is false, hence narrowing down the set of possible worlds and improving their understanding of the actual world.

<sup>7</sup>Classical logic, which is the logic at the foundations of mathematics, is based on principles like the *law of excluded middle*—which states that every proposition is either TRUE or FALSE, there is no intermediate value—and the *law of non-contradiction*—which states that a proposition cannot be both TRUE and FALSE at the same time. Non-classical logic refers to any logical system that challenges or extends one or more of these principles to reason in scenarios classical logic cannot, such as uncertainty or paradoxes. Examples include modal logic, intuitionistic logic, and fuzzy logic.

resolve the problems that previous solutions had been running into.

Carnielli considered that all puzzle agents would reason using the three-value LF11 paraconsistent logic, which can accommodate both *consistent* and *non-consistent* questions. The former can be answered definitively in a YES or NO way, whereas the latter are inherently doubtful and can only be answered as “indeterminate.” Carnielli reformulated the HLPE, making it similar to Uzquiano’s (HLPE<sub>syn</sub><sup>3</sup>) with the difference that all three gods would reply either “da,” “ja,” or “ta,” corresponding to the truth values “yes,” “no,” and “indeterminate,” in an unknown order. He solved this puzzle in three questions, conjecturing that this version cannot be solved in fewer questions. Rather than building on previous versions of the HLPE, Carnielli explicitly sought to inspire new forms of logic puzzles based on non-standard logics.

As we have seen, what began as a somewhat provocative puzzle has evolved into a significant topic of study within the academic and technical fields of logic, largely due to its implications for reasoning about truth, randomness, and language under conditions of minimal information. The *Hardest Logic Puzzle Ever*, though seemingly innocuous and tracing back to the tradition of recreational problem-solving, has proven to be a valuable model for advanced research in formal logic, linguistic structures, epistemology, and the role of randomness across these domains.

In this work, we aim to incorporate ontological reality into the HLPE. We will study the fundamental physical realization of the puzzle as an interaction between agents and analyse the consequences for its solution. As we will soon demonstrate with novel solution strategies exploiting the implementation of logical processes, physical reality can be tied with logical reasoning to uncover remarkable advantages to the solvability of the puzzle.

Our exploration offers important insights into the empirical nature of logic which, given the depth of understanding already achieved in the extensive body of literature on the HLPE, may come as a surprise to the readers unfamiliar with the implications of quantum theory.

## 2.2 Formalizing the HLPE

With the exception of the solutions to the HLPE proposed by Boolos [143], Roberts [144], and Carnielli [154], other attempts to develop more advantageous solutions by Rabern and Rabern [145], Uzquiano [146], Wheeler and Barahona [147], and Wintein [148] have all exploited the use of sophisticated self-referential questions. This approach enabled them to justify an answer output produced from more than two available values, unlocking the possibility to solve the riddle in less than three questions.

While these solutions are certainly interesting in their own right, from a practical standpoint and assuming deception is not part of the game, one might question the purpose of receiving instructions to pose a binary question if there could be more than two possible answers. When presented with a choice between two options, not only it seems natural but also necessary for establishing a well-defined system of communication to assume that the gods should select only one of them, rather than fabricating alternatives such as “silence,” “neither,” or “both” that are not even declared upfront. We will not enter into a philosophical discussion on this matter here; instead, we will operate under *classical logic* and the assumption that the gods are obliged to reply *exclusively binary answers* to binary questions.

In the next section, we will explain why this assumption arises as a fundamental requirement for communication under minimal physical resources and shared assumptions rather than as an arbitrary design decision. Before that, identifying a key consequence of this condition is necessary: constraining answers to a binary value requires the use of *valid questions*. Hence, we enforce the following condition:

(B0) Only *valid questions* can be directed to the gods, which are the subset of all binary questions that only admit YES or NO answers.

As a result, solutions relying on undecidable questions, such as some of those presented in the previous section, cannot be employed under this restriction<sup>8</sup>. Furthermore, by limiting the questions that can be asked, we are effectively reducing our interrogative power and making the puzzle harder to solve.

In addition to *i*) enabling clear communication between the parties by ensuring predictability and removing ambiguity, and *ii*) making the puzzle harder by reducing the number of questions that can be asked, (B0) also *iii*) makes the communication more memory efficient. Transmitting three binary values requires less physical resources than transmitting two ternary or quaternary values. To see why, consider that, while the first can discriminate  $2^3 = 8$  possibilities with the 3 bits, the second can discriminate between  $3^2 = 9$  possibilities, carrying therefore  $\log_2 9 = 3.17$  bits, and the third discriminates among  $4^2 = 16$  possibilities using  $\log_2 16 = 4$  bits. As a result, solutions based on three binary answers also outperform those based on two ternary- or quaternary-valued answers in terms of communication efficiency.

To help us discuss the problem, we will frame it as an interaction between two agent categories: the *solver*, Alice<sup>9</sup>, and the *gods*, *True* ( $\mathcal{T}$ ), *False* ( $\mathcal{F}$ ), and *Random* ( $\mathcal{R}$ ). The designation of *god* serves merely to emphasize the ability of these physical agents to correctly compute truthful answers to any

<sup>8</sup>If more detail is desired, let us assume it is the responsibility of the interrogator to anticipate when a question is potentially doubtful and non-valid before addressing it. And if it does not, the absence of a reply will signify that the directed question was deemed non-valid if there could be at least one scenario in its evaluation by any of the gods in which it would produce a contradictory output. The game is terminated immediately.

<sup>9</sup>Here we follow a certain tradition initiated in [155].

question, and there is nothing else qualitatively different about it.

Formally, once Alice asks a valid question  $\mathcal{Q}$ , it is translated by the recipient god, who is assigned to position A, B, or C, into a logical expression  $f_{\mathcal{Q}} : \{0, 1\}^n \rightarrow \{0, 1\}$  that evaluates and combines  $n$  input propositions  $\{p_i\}_{i=1}^n$ , where the domain  $\{0, 1\}^n$  represents all their possible truth assignments and the codomain  $\{0, 1\}$  represents the output truth value. The positive Boolean value will be labeled either as 1, YES, or TRUE; the complementary value is labeled 0, NO, or FALSE. As usual, logical connectives such as conjunction ( $\wedge$ ), disjunction ( $\vee$ ), negation ( $\neg$ ), implication ( $\rightarrow$ ), and biimplication ( $\leftrightarrow$ ), are used to combine propositions.

After computing the *truthful* value  $x_1 = f_{\mathcal{Q}}(\{p_i\}_{i=1}^n)$ ,  $\mathcal{T}\text{rue}$ ,  $\mathcal{F}\text{alse}$ , or  $\mathcal{R}\text{andom}$  act over it according to their characteristic *logical personality operator*  $U \in \{T, F, R\}$ , which either maintains this value, negates it, or uniformly randomizes it, respectively. Explicitly, an output value  $x_2$  is obtained from one of

$$T(x_1) = x_1, \quad (2.1)$$

$$F(x_1) = x_1 \oplus 1, \quad (2.2)$$

$$R(x_1) = x_1 \oplus x', \text{ where } x' \sim \text{unif}\{0, 1\}, \quad (2.3)$$

where  $x_1$  is the truthful Boolean input,  $\oplus$  represents addition modulo 2 and  $\text{unif}\{0, 1\}$  denotes a uniform binary probability distribution. The resulting  $x_2$  is then mapped into the final  $x_3$  output according to the *language function*

$$L_j(x_2) = x_2 \oplus j, \quad (2.4)$$

where  $j \in \{0, 1\}$  denotes one of the two *language-boolean map* possibilities:

$$\text{LBM}_0 = \{(da, 1), (ja, 0)\} \quad (2.5)$$

$$\text{LBM}_1 = \{(da, 0), (ja, 1)\}. \quad (2.6)$$

In words, either “da” means YES and “ja” means NO, as in Eq. (2.5), or vice-versa, as in Eq. (2.6). Importantly, note that the final output,  $x_3$ , is always a binary value, where made the arbitrary choice to encode the word “da” in label 1<sup>10</sup>.

Alice’s task, as challenged by Boolos, is to identify the unknown  $U$ , and therefore the correct assignment of the gods to their positions A, B, or C, by asking three YES/NO questions without knowledge of the actual value of  $j$ . The questions we will use to solve the HLPE do not involve self-reference, just as

---

<sup>10</sup>We will come back to this point in the next section.

Boolos' and Roberts' questions did not. Therefore, this formalism suffices to represent all the questions we will encounter and solve the puzzle.

To illustrate the framework, let us consider question  $\mathcal{Q}$ : “Does  $1 + 1$  equal to 2 and you are  $\mathfrak{R}andom$ ?” is addressed to  $\mathfrak{F}alse$  and “da” means YES. In this example,  $f_{\mathcal{Q}}(p_1, p_2) = p_1 \wedge p_2$  represents the question; after evaluating the truth values of its two propositions  $p_1 = 1$  and  $p_2 = 0$ , the truthful answer becomes  $x_1 = f_{\mathcal{Q}}(1, 0) = 1 \wedge 0 = 0$ . Then,  $x_2 = F(x_1) = F(0) = 1$  and  $x_3 = L_{j=0}(x_2) = x_2 \oplus j = 1 \oplus 0 = 1$ , meaning the answer is “da”.

The HLPE version we are considering with this formalization enunciates exactly as  $(HLPE_{syn}^2)$ , but strickly constrains all three gods to answer with either “da” or “ja”. It should be understood as supplemented by conditions (B0), (B1), (B2), and (B3\*). It is clear that a solution to our formulation will also be a solution to the  $(HLPE_{syn}^2)$  without (B0) explored by previous authors, as the former is a restriction of the latter. Both the solutions of Boolos and Roberts can be represented in our formalism and solve our restricted  $(HLPE_{syn}^2)$  formulation<sup>11</sup>.

A brief inspection of the problem shows that we have 12 possible configuration scenarios we could be in, as the six possible permutations of the gods with respect to positions A, B and C,

$$\pi_1 = (\mathfrak{T}, \mathfrak{F}, \mathfrak{R}) \quad (2.7)$$

$$\pi_2 = (\mathfrak{T}, \mathfrak{R}, \mathfrak{F}) \quad (2.8)$$

$$\pi_3 = (\mathfrak{F}, \mathfrak{T}, \mathfrak{R}) \quad (2.9)$$

$$\pi_4 = (\mathfrak{F}, \mathfrak{R}, \mathfrak{T}) \quad (2.10)$$

$$\pi_5 = (\mathfrak{R}, \mathfrak{T}, \mathfrak{F}) \quad (2.11)$$

$$\pi_6 = (\mathfrak{R}, \mathfrak{F}, \mathfrak{T}) \quad (2.12)$$

are multiplied by the two possible language-boolean maps. Since obtaining a single bit of information per question, amounting to a total of three bits, will only allow for the discrimination of  $2^3 = 8$  distinct options, it is evident that with only three questions it will be impossible to determine the language of the gods if we aim to first establish their position assignment. Nevertheless, transmitting three bits suffices, as language determination is not a requirement. In fact, existing three-question solutions, such as those proposed by Boolos and Roberts, determine position while leaving the meanings of “da” and “ja” undetermined.

In the following section, we will look at the realization of the HLPE from a fundamental physics per-

---

<sup>11</sup>To actually formalize  $(HLPE_{sem})$ , however, Eq. (2.3) would need to be modified to randomly select wether to apply the  $T$  or  $F$  operator, in Eq. (2.1) and Eq. (2.2), respectively.

spective. This endeavor requires translating the mathematical model into physical processes. Despite the effort to perform this translation in full generality, adapting this abstract model to physical reality inevitably requires explicit assumptions about the communication medium and implementation details, which were ignored in previous analyses. In the next section, we will discuss these assumptions and work towards developing a setup that minimizes them as much as possible.

## 2.3 Classical communication of a binary value

In the previous section, we have formalized the HLPE to strictly adhere to the rules of the puzzle while explicitly establishing well-defined logical personality operators for the gods. Here, we begin exploring the primary objective of this work: *to give the HLPE an ontological framework grounded in physical reality*. Since the scenario involves the interaction between two parties, we will focus on understanding the role of communication and its implications for solution strategies. To this end, the previously introduced formalization serves as the foundation for a seamless translation of the problem setup into a fundamental physical implementation.

Let us recall the basic communication scenario between the two involved parties: Alice and one of the gods. Alice asks the god a valid question; the god needs to reply a simple TRUE or FALSE Boolean value, encoded as “da” or “ja”. We will ignore how Alice communicates her question, assuming only that she gets the message accross<sup>12</sup>. Our purpose is to understand how the minimal amount of information in the answer can be physically communicated. In the rest of this section, we will consider different ways for the god to communicate the answer classically, aiming to use the minimal physical resources and make the fewest shared assumptions possible.

### 2.3.1 Assumed innate modes of communication

Posing a question and receiving an answer verbally, or through other modes of communication presumed to be natural and innate between the agents, appears to be how Boolos intended the puzzle to be interpreted. This is a common approach when the primary focus is on the analysis of argumentative logic. In fact, all the presented solutions to the HLPE assume seamless interaction between the two players and ignore the details of the communication system [143, 144, 145, 146, 147, 148, 149, 152, 154].

However, framing the problem in terms of verbal questions and answers incurs some important implicit

---

<sup>12</sup>Our assumption is based on the HLPE’s assertion that the gods understand Alice’s language, suggesting that comprehension of the posed question is not problematic.

assumptions. First, Alice and the god begin their interaction aware of the medium and symbols used for communication<sup>13</sup>. Second, the god must be capable of performing complex physical manipulations to transduce the information into modulated air pressure waves following a systematic encoding protocol. Third, Alice’s innate auditory system enables her to sense and decode the information contained in the incoming sound waves. Fourth, if Alice and the god both speak the same language, she will be able to instantly understand the answer when processing the auditory signal, as the communicated symbol directly maps to the Boolean value representing the answer.

We could instead consider the answer is shared by means of any other natural or engineered mode of communication. In any case, generalized versions of these assumptions would still be required; informally:

1. Both Alice and the god know the physical medium and the syntax used for communication (e.g., the symbols and the waiting times),
2. The god must manipulate the physical system to encode the symbols,
3. Alice must sense her end of the system to decode the information, and
4. Alice must be able to attribute meaning to the decoded symbols in a language she understands.

The important point is that, in any communication framework, both parties must know how to interact with a shared external medium that serves as the information carrier. This interaction allows abstracting the information into physical states, which can then be manipulated and transmitted over the communication channel. The process of encoding and decoding depends on protocols that can be previously shared or learned through experience, mapping the information to and from the physical states in a mutually intelligible way. Without this dual capacity for manipulating and interpreting a common medium, meaningful information exchange cannot occur.

Therefore, there are implicit assumptions ignored by the solutions of the HLPE that do not consider the communication problem. Consciously or not, both parties have some understanding of how the message is encoded in the physical medium used for communication and are therefore not entirely independent.

One of the interesting aspects of the HLPE is that it explicitly addresses the fourth assumption above. Boolos challenges us to solve the puzzle without knowledge of the semantic decoding protocol, demonstrating it is possible if the questions are carefully formulated. However, none of the previous HLPE formulations specifies the physical substrate for communication. As we more formally examine different communication models, it will become clear why it is also important to recognize the other implicit assumptions in verbal solutions of the HLPE.

---

<sup>13</sup>As established by Boolos’ formulation, Alice already knows the answer of the god will consist in either “da” or “ja”.

### 2.3.2 Minimizing physical resources and shared assumptions

The foundation of a rigorous theoretical framework for the quantification, storage, and communication of information is widely attributed to Claude Shannon in a seminal 1948 article [156]. Shannon's work was motivated by engineering considerations, and in his conceptualization of information transmission, he consequently postulated that both the sender and receiver share a common understanding of the communication system, including the encoding and decoding schemes.

Here, we set out to minimize these assumptions and reduce the physical resources required for this communication system to a fundamental physics level. We will limit ourselves to classical physics to start with. Since the answer to be transmitted consists of only one bit of information, we abstract away the details of the communication medium and consider the simplest physical model that could be used to encode the message: a *classical two-state system*.

#### Encoding on classical two-state system states

A *classical two-state system* is a fundamental physical system governed by classical laws, such as Newtonian mechanics or electrodynamics, that can exist in one of two distinct *effective* states at any given time. The *phase space* of the system is a continuous space where both its generalized coordinates and momenta can take any real value within their respective ranges, with each point in this space representing the state of the system at a moment in time [157, 158]. The effective two-state discreteness of the system arises from constraints or conservation laws that confine its steady-state evolution to one of two *periodic trajectories* in the phase space, or to *fixed-points* in the case of no movement.<sup>14</sup> Each effective state, henceforth simply referred to as *state*,<sup>15</sup> thus has its own phase portrait, with well-defined and fundamentally deterministic<sup>16</sup> operations establishing the transient evolutions between the states.

Properly defined, these states are mutually exclusive and exhaustive, meaning the system can only occupy one of the states at any given time, and the two states cover all possible configurations of the system. Enforcing such a system for communication ensures reliability and predictability in information transmission. Importantly, the two states are distinguishable and can be described based on their attributes; more fundamentally, states of a conservative system can be described based on their phase space portrait or, equivalently, their energy, which is fixed for a given contour line in the phase space.

---

<sup>14</sup>As an example, imagine a pendulum oscillating with one of two possible energy levels in a system where energy is conserved.

<sup>15</sup>For the purposes of information processing, it is often desirable to preserve information over time in a stable steady state that facilitates transmission. Hence, we reserve the term "*state*" specifically for these effective states in the phase space.

<sup>16</sup>At a *fundamental* level, *classical* physics is entirely deterministic: given the initial conditions of the system and the governing laws of motion, its future behavior can be precisely determined.




Figure 2.1: **State-based communication scheme between Alice and a god.** Alice (left) poses her question via a specified communication channel. The binary answer is provided by the god (right) using a two-state system. The parties make mutual assumptions to facilitate reliable communication. In the main text, the minimization of physical resources and shared assumptions is discussed.

For familiarity, we will illustrate our discussion by resorting to the example of a *coin*. The two faces of the coin and their relative orientation to the parties, can configure two states, denoted as *heads* ( $H$ ) and *tails* ( $T$ ). This minimally expressive system can be used as the medium of communication in the HLPE by mapping each of the Boolean values onto each of the states of the system. We begin by considering this simple encoding strategy, represented in Fig. 2.1. Five assumptions shall be made explicit:

Both the sender and the receiver have a mutually intelligible description of the two states that the system can realize, labelled  $H$  and  $T$  for convenience. (A<sub>1</sub>)

The sender can read the state of the system. (A<sub>2</sub>)

The sender can perform two state-transition operations required to modify the steady state of the system : 1)  $H \rightarrow T$  , and 2)  $T \rightarrow H$ . (A<sub>3</sub>)

The receiver can read the state of the system. (A<sub>4</sub>)

Both parties agree beforehand on a *state-boolean map* (SBM), identifying which state corresponds to which Boolean value. (A<sub>5</sub>)

Assumption A<sub>1</sub> ensures that when asking a question or providing an answer, both parties correspond  $H$  and  $T$  to the same faces of the coin, allowing Alice to ask questions referring to these states that are correctly understood by the god as well as to correctly identify the state the god intends to communicate back. The labels  $H$  and  $T$  are used for convenience and do not need to be agreed upon in advance as the states are distinguishable and can be explicitly described based on their physical attributes when the question is posed<sup>17</sup>. What A<sub>1</sub> requires more fundamentally is that the description of the states of the system is known from the outset, since otherwise it is not possible to refer to this description in the question — at least not in the first one.

<sup>17</sup>For instance, Alice can ask “Does ‘the state with such and such phase space portrait or energy’ mean TRUE?”

At first glance, it may seem unnecessary to include  $A_1$  as a requirement, given that the system has only two states available for encoding. These states — and, by extension, the syntax — might appear to be self-evidently constrained by the system. However, it is important to recognize that the definition of the encoding states is not uniquely determined and must be agreed upon in advance. For instance, a coin only has two distinct states after we have chosen a specific direction along which to observe it. Similarly, while it may not be immediately apparent, assumption  $A_1$  is also provided in the formulation of the HLPE: when Alice is informed the answer will consist in either “da” or “ja”, she learns the symbols to be used for communication. Furthermore, when the answer is communicated verbally, the mapping from the transmitted signal to the syntactic symbol is implicit and we may not be consciously aware of it. With  $A_1$ , this assumption is made explicit.

As for  $A_2$ , it ensures the sender can read the state of the system in order to know how to manipulate it to encode the answer following  $A_3$ . The receiver is only required the ability to read the state of the system to recover the answer — assumption  $A_4$ . It is important to note that, with  $A_5$ , the LBM, i.e. Eqs. (2.5) and (2.6), is equivalently transformed into a SBM. Answers are no longer given in terms of “da” and “ja” but, equivalently,  $H$  and  $T$ . The communicated symbols, or syntax, are intelligible to both parties, but their semantics can be encoded using one of two possible SBMs. Communicating in the presence of an unknown language, therefore, corresponds to being unaware of which SBM the other party employs.

To illustrate how this framework would work with an example, suppose Alice asks her question  $\mathcal{X}$ : “Does  $1 + 1$  equal to 2?” and receives the coin with  $T$  facing upwards. Beforehand, they agreed that the answer will be encoded in the face-up of the coin (assumption  $A_1$ ) and that the state-boolean map is  $\text{SBM} = \{(H, 0), (T, 1)\}$  (assumption  $A_5$ ). As a result, Alice understands the god’s answer is TRUE.

## Dropping a shared state-boolean map

Can we drop assumption  $A_5$  and still be able to communicate with full understanding using different languages? This is indeed possible by carefully formulating the question. Let us present the method employed by Boolos in his solution to the HLPE [143], adapted to this communication scenario. Suppose Alice is attempting to determine the answer to a question labeled  $\mathcal{X}$ . First, consider the question

$$\text{Does } T \text{ mean TRUE?} \quad (2_1)$$

We can construct the actual question to be addressed to assess whether the truth value of  $\mathcal{X}$  is equivalent to the one from  $2_1$ . In propositional logic, equivalence is computed using the “if and only if” (iff) logical connective, with the truth table given in Table 2.1. Therefore, the addressed question is

$$\mathcal{X} \text{ iff } 2_1? \quad (2_2)$$

| $\mathcal{P}$ | $\mathcal{Q}$ | $\mathcal{P} \leftrightarrow \mathcal{Q}$ |
|---------------|---------------|-------------------------------------------|
| 0             | 0             | 1                                         |
| 0             | 1             | 0                                         |
| 1             | 0             | 0                                         |
| 1             | 1             | 1                                         |

Table 2.1: Truth table of the “if and only if” (iff) logical connective between propositional variables  $\mathcal{P}$  and  $\mathcal{Q}$ .

| $\mathcal{X}$ | $\mathcal{Q}_1$ | $\mathcal{Q}_2$ | Coin state received |
|---------------|-----------------|-----------------|---------------------|
| 0             | 0               | 1               | $H$                 |
| 0             | 1               | 0               | $H$                 |
| 1             | 0               | 0               | $T$                 |
| 1             | 1               | 1               | $T$                 |

Table 2.2: Truth table and received state for a truthful answer to question  $\mathcal{Q}_2$  encoded in a two-state system without a shared state-boolean map.

When Alice poses question  $\mathcal{Q}_2$  instead of  $\mathcal{X}$ , she awaits the response without making any assumptions regarding the state-boolean map the god might employ. The evaluation of the truthful answer will depend upon the specific SBM utilized by the god. The analysis of all potential responses and the consequent coin state received as an answer can be easily performed and is presented in Table 2.2. It is evident that Alice can retrieve the Boolean value corresponding to the answer to the question she is interested in, question  $\mathcal{X}$ , from the received coin state. The coin state will be  $H$  when the answer to  $\mathcal{X}$  is FALSE and  $T$  when this answer is TRUE.

The only requirement from the sender is that he agrees to communicate by means of a two-state system and to encode the Boolean values into its two possible states, with the exact mapping unknown. For this purpose, we assumed  $A_1$ ,  $A_2$ , and  $A_3$ ; however, it is also conceivable that the sender possesses two distinct two-state systems, each in a different state, and selects which one to transmit based on the required response. If that were the case, assumption  $A_3$  would not be mandatory, but this would be offset by the communication system requiring twice the physical resources; that is, the sender would need to possess two two-state systems instead of only one if the objective was to avoid manipulating their state.

In summary, this framework allows operating the communication system while minimizing shared assumptions to only  $A_1$ ,  $A_2$ ,  $A_3$ , and  $A_4$ .

### Providing the system for hidden state encoding

Consider now the scenario in which the two-state system is provided by Alice to the god together with the question (Fig. 2.2). Since she initiates the interaction by transmitting the question and she is the final recipient of the information, it is appropriate for her to also send the coin that the god will use to encode the answer. This protocol eliminates the need for the god to store the physical substrate for encoding the answer, requiring only that he operates over it upon its reception. As we will now demonstrate, this also allows for a further reduction in the shared assumptions by eliminating  $A_1$  and  $A_2$ .




Figure 2.2: **Operations-based communication scheme between Alice and a god.** Alice (left) poses her question via a specified communication channel along with a hidden state two-state system. The binary answer is provided by the god (right) by either bit-flipping or not the system.

Being a two-state system, acting over it can either preserve its original state or alter it to the other available state. Therefore, if a symmetric operation can be designed to accomplish both the transitions  $H \rightarrow T$  and  $T \rightarrow H$  independently of the system's initial state, the sender is not required the ability to read the received state. The Boolean values no longer need to be encoded in the *state* of the system, but rather by the *operation* that transforms between system states. In this case, only Assumptions  $A_3$  and  $A_4$  need to be retained. However, assumption  $A_3$  has to be tightened to explicitly ensure that a single operation performs both state transitions:

The sender has access to a bit-flip operation that changes the internal state of the system, independent of its initial state. ( $A_3'$ )

It becomes necessary only to reformulate the addressed question in a way that enables the extraction of the answer to  $\mathcal{X}$ , the question of interest, without assuming any specific *operation-boolean map* (OBM). This process follows a similar approach to the one previously described. We first formulate the question

Does bit-flip mean TRUE?, ( $\mathcal{Q}_3$ )

making the addressed question

$\mathcal{X}$  iff  $\mathcal{Q}_3$ ? ( $\mathcal{Q}_4$ )

Alice hides the coin inside a box and sends it to the god along with question  $\mathcal{Q}_4$ , while maintaining a record of the initial state of the coin that only she can see. The god will either flip the box or not, without learning the state of the coin. The possible answers and the operation they produce on the coin states are detailed in Table 2.3. As we can see, Alice can retrieve the Boolean answer by reading the received state and comparing it to the state she initially sent. The state of the coin will remain *unaltered* if the answer to  $\mathcal{X}$  is FALSE and *bit-flipped* if this answer is TRUE. The only requirement from the sender now is that he agrees to communicate using a two-state system and to encode the two possible Boolean values by either maintaining or altering its state, with the exact semantic mapping of the OBM remaining unknown.

| $\mathcal{X}$ | $\mathcal{Q}_3$ | $\mathcal{Q}_4$ | Coin state operation |
|---------------|-----------------|-----------------|----------------------|
| 0             | 0               | 1               | identity             |
| 0             | 1               | 0               | identity             |
| 1             | 0               | 0               | bit-flip             |
| 1             | 1               | 1               | bit-flip             |

Table 2.3: Truth table and received state transformation for a truthful answer to question  $\mathcal{Q}_4$  encoded in the state transformation of a two-state system without a shared operation-boolean map.

A description of the states of the system is no longer necessary since Alice no longer needs to address them explicitly in her questions, avoiding the syntactic problem that required assumption  $A_1$ . To be precise, both parties still share knowledge of the syntax, as it consists of either an `identity` or a `bit-flip` operation. However, the interpretation that the `identity` operation leaves the state of the system unchanged and the `bit-flip` changes it can be communicated unambiguously during the question and does not rely on detailed knowledge of the system. Designing the `bit-flip` operation does depend on system knowledge, hence assumption  $A_1$  could still be implicitly present within  $A_{3'}$ . Nevertheless, the design can be shared from one party to the other, effectively eliminating the requirement for both to have a complete description of the system.<sup>18</sup>

By requiring only a two-state system provided by the final recipient of the information, together with Assumptions  $A_{3'}$  and  $A_4$ , this is the classical communication setup that requires minimal physical resources and shared assumptions.

### 2.3.3 Realizing classical randomness

Another interesting feature of the HLPE that connects it to an important aspect of physical reality is *randomness*, which is assumed to play a critical role in the problem. Hence, a physical realization for randomness — specifically, the implementation of the logical personality operator employed by  $\mathfrak{R}andom$  — must also be addressed. While an in-depth discussion of randomness in the physical world is beyond the scope of this work, it is necessary to outline the essential aspects of this topic.

As we have already mentioned in passing before, classical physics is an entirely deterministic theory at a fundamental level. When we are given the initial conditions of any classical system in terms of its

<sup>18</sup>Although these considerations about the syntax might seem peripheral or overly meticulous, it is worth noting that a version of the HLPE has also been informally studied where syntactic knowledge of the language of the gods is altered or eliminated [152]. As we discussed in § 2.1, Novozhilov investigated how to solve the puzzle assuming only that the gods reply TRUE or FALSE in some language for which both syntax and semantics are *a priori* unknown, and communication is assumed to be innately perceived. In the communication framework here presented, the syntax is effortlessly fixed by the physical system, as its minimal structure leaves no room for alternative choices.

generalized coordinates and applied forces, its future dynamical behavior is already determined and can, *in principle*, be exactly predicted by solving the governing equations of motion.

Classical randomness, therefore, cannot stem from fundamental processes but rather from the absence of perfect knowledge about the system's initial conditions. This ignorance becomes significant in practical situations where exceedingly large numbers of interacting components must be considered in order to accurately model a complex system and predict its behavior. An example is the outcome of rolling a die, where unknown or untrackable initial conditions make the result effectively unpredictable, and our best approach to describing the system is with an effective classical probabilistic process. As a consequence, randomness in classical physics is generally understood as epistemic rather than intrinsic to the nature of reality; that is, it arises from incomplete knowledge about a system's full description rather than from fundamental indeterminism.<sup>19</sup>

In the HLPE, particularly in the (HLPE<sup>2</sup><sub>syn</sub>) formulation,  $\mathfrak{R}andom$  is postulated to act in a completely unpredictable manner. Since randomness has no ontological status in classical physics, it must be acknowledged as an explicit assumption in our physical implementation of the HLPE:

The sender has access to a rand-flip operation that can be used to sample a random Boolean value prior to encoding the message onto the communication system. (A<sub>6</sub>)

The realization of this operation requires an epistemic source, with numerous examples in the physical world related to unpredictability in complex systems that can function as sampling machines. Assuming perfect transmission of the message, the way to introduce randomness into our classical communication scenario is to generate it at the source as controlled by  $\mathfrak{R}andom$ , in agreement with the problem formulation described in § 2.2. It is worth noting that later in § 2.8, we will also examine the effects of epistemic randomness during transmission.

Finally, in contrast to classical physics, randomness has a fully ontological status in quantum physics. We will also explore these implications further in § 2.5. Before that, the next section consolidates the building blocks we have established thus far to demonstrate how the HLPE is solved classically.

---

<sup>19</sup>A short digression is in order. The discussion about randomness in classical physics can be a bit more nuanced due to *chaos*, involving the practical unpredictability of deterministic systems that are highly sensitive to the initial conditions [159, 160]. There is also the more esoteric issue of the possibility of non-uniqueness of solutions for Newton's equations in some carefully crafted systems with fine-tuned artificial potentials, such as in Norton's dome [161]. These can only be rescued by supplementing classical Newtonian mechanics with the additional Lipschitz continuity condition to guarantee uniqueness of solutions to the ordinary differential equations of motion. However, we also know that Newtonian mechanics, while a highly useful approximation in many common situations at low speeds and weak gravitational fields, is ultimately superseded and extended by *general relativity*, which offers a more complete description of gravity. In general relativity, where gravity is described through the geometric curvature of spacetime, determinism is preserved as long as spacetime remains globally hyperbolic and free from singularities, thus avoiding the specific indeterminism presented with Norton's dome [162].

## 2.4 Classical-communication-based solution to the HLPE

### 2.4.1 Revisiting Boolos' exact solution

In the original “*The hardest logic puzzle ever*” publication by Boolos [143], he also presents a three-question solution to (HLPE<sub>sem</sub>). The strategy and the questions Boolos employs are able to solve the problem under conditions (B1), (B2), (B3), and (B4), but can equally be applied to solve (HLPE<sub>syn</sub><sup>2</sup>) under its conditions (B1), (B2), (B3\*), and (B4). The questions Boolos uses are, in fact, *valid* under our condition (B0), and can therefore also be used to solve our restricted (HLPE<sub>syn</sub><sup>2</sup>) as formalized in § 2.2, which assumes conditions (B0), (B1), (B2), and (B3\*)<sup>20</sup>.

Here, we briefly revisit Boolos' solution to illustrate how the questions he employed can be recast in the language of communicating the answer using a classical two-state system presented in the previous section, under Assumptions A<sub>3'</sub>, A<sub>4</sub>, and A<sub>6</sub>. We will not expose a case-by-case analysis to exhaustively examine all twelve scenarios and prove that this strategy solves the problem deterministically. We invite the reader to consult Boolos' publication for that analysis. Rather, we will simply identify the steps taken in the solution strategy and summarize the knowledge gained after each answer, as this will be useful later. The solution proceeds in three steps, in which Alice asks the following questions:

- **Question 1**, addressed to the god in position A:

Does a bit-flip mean YES iif (A is *True* iif B is *Random*)? (Q<sub>5</sub>)

If A answers with a bit-flip operation, Alice is left with possible permutations  $\{\pi_2, \pi_4, \pi_5, \pi_6\}$ . On the other hand, if A answers with the identity operation, Alice learns the possible permutations left are  $\{\pi_1, \pi_3, \pi_5, \pi_6\}$ .

- **Question 2**, addressed to the god in position C if the answer to Q<sub>5</sub> was a bit-flip, or to the god in position B in case it was the identity operation:<sup>21</sup>

Does a bit-flip mean YES iif 1 + 1 is equal to 2? (Q<sub>6</sub>)

Now, Alice is carefully directing the question to a god that she knows cannot be *Random*. If C answers with an identity or a bit-flip operation, she learns that the final permutation is one of  $\{\pi_2, \pi_5\}$  or one of  $\{\pi_4, \pi_6\}$ , respectively. In the case the question has been addressed to B,

<sup>20</sup>Redundantly, also condition (B4).

<sup>21</sup>Boolos' original question uses “Rome is in Italy” in place of “1+1 is equal to 2”.

if he answers with an identity or a bit-flip operation, the final permutation is contained in either  $\{\pi_1, \pi_6\}$  or  $\{\pi_3, \pi_5\}$ , respectively.

- **Question 3**, addressed to the same god who was asked [26](#):

Does a bit-flip mean YES iif A is  $\mathfrak{R}andom$ ? (27)

The answer reveals the final permutation among only two possibilities left.

What we want to emphasize in this discussion is the gradual reduction of the solution space as each successive question is answered. First, the bit obtained from the first answer reduces the possible permutations of the gods from six to four. Then, the second bit further narrows the solution space from four to two possibilities. Finally, the third answer reveals the correct permutation among the two options left. Additionally, note that the structure of this strategy only determines the recipient of the second and third questions depending on the outcome of the first one, and notably, one of the gods remains unqueried throughout. This progression is also the same in the solution by Roberts [\[144\]](#).

## 2.4.2 Probabilistic solutions and success probabilities

We already discussed that it is possible to solve the HLPE<sup>22</sup> using three valid questions following the strategies of Boolos [\[143\]](#) or Roberts [\[144\]](#). This solution is deterministic and exact, therefore, the probability of successfully solving the problem in three valid questions can be denoted as follows, where the subscript refers to the number of allowed questions:

$$P_{\text{succ},3} = 1. \tag{2.13}$$

Given that there are six possible permutations in which the gods could distribute themselves over the three available positions A, B, and C, the probability of correctly guessing the position of each god by random chance without asking any questions is

$$P_{\text{succ},0}^{\text{ME}} = \frac{1}{6}, \tag{2.14}$$

where we use the superscript ME to denote *minimum-error* discrimination, anticipating the language of quantum strategies. A minimum-error discrimination probability is the probability of correctly guessing the

---

<sup>22</sup>One more time, we refer to our constrained ( $\text{HLPE}_{\text{syn}}^2$ ) version under conditions (B0), (B1), (B2), (B3\*).

identity of the gods with the strategy that minimizes the error, which, for an equiprobable distribution, is simply the probability of selecting the correct permutation at random.

To obtain the minimum-error probability of guessing the correct permutation of the gods after just one question, let us consider the discussion in the previous section. As we have seen, in both Boolos and Roberts solutions, the first question reduces the solution space from six to four possible permutations, which can be clearly identified based on the value of the received answer. With the knowledge of the four remaining possible permutations in which the gods can be distributed, Alice can bet in one of them with a 25% chance of guessing it correctly. As a result, following the solution strategies from Boolos or Roberts, the probability of solving the HLPE by random chance after just one question is

$$P_{\text{succ},1}^{\text{ME}} = \frac{1}{4}. \quad (2.15)$$

This probability of success is relatively low, which is expected given the limited amount of information that can be extracted with a single question. The same result can also be derived through an alternative probabilistic reasoning strategy. Consider Alice poses a question to the god in position A, which could be any of the three gods with an equal probability of  $1/3$ . Since  $\mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}$  responds randomly with either bit value, Alice would prefer to craft her question in a way that guarantees the same answer from the other two gods. This could be, for instance, question  $\mathcal{Q}_{10}$  below

$$\text{Are you } \mathfrak{F}\mathfrak{a}\mathfrak{l}\mathfrak{s}\mathfrak{e}? \quad (\mathcal{Q}_8)$$

$$\text{Does the identity mean TRUE?} \quad (\mathcal{Q}_9)$$

$$\mathcal{Q}_8 \text{ iff } \mathcal{Q}_9? \quad (\mathcal{Q}_{10})$$

which would yield a bit-flip answer from both  $\mathfrak{T}\mathfrak{r}\mathfrak{u}\mathfrak{e}$  and  $\mathfrak{F}\mathfrak{a}\mathfrak{l}\mathfrak{s}\mathfrak{e}$  as the exhaustive analysis in Table 2.4 demonstrates. Consequently, if Alice receives an identity answer, she can infer with certainty that it came from  $\mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}$ , although he will also answer bit-flip half of the time. With this in mind, she can formulate the following betting strategy, for which the success probability is computed next.

Suppose first that Alice receives the identity as an answer. In this case, only permutations  $\pi_5$  and  $\pi_6$ , where  $\mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}$  is in position A, could have produced the result. She can thus place her bet on one of these permutations, with a  $1/2$  chance of success. However,  $\mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}$  is in position A with only  $1/3$  chance, and even then, he answers identity only half the time. Therefore, the overall probability for Alice to correctly identify one of these permutations is  $(1/3)(1/2)(1/2) = 1/12$ .

| God   | $\mathcal{Q}_8$ | $\mathcal{Q}_9$ | $\mathcal{Q}_{10}$ | $U(\mathcal{Q}_{10})$ | Answer   |
|-------|-----------------|-----------------|--------------------|-----------------------|----------|
| True  | 0               | 1               | 0                  | 0                     | bit-flip |
| True  | 0               | 0               | 1                  | 1                     | bit-flip |
| False | 1               | 1               | 1                  | 0                     | bit-flip |
| False | 1               | 0               | 0                  | 1                     | bit-flip |

Table 2.4: Truth table and received state for an answer to  $\mathcal{Q}_{10}$  from either `True` or `False` in both possible OMB scenarios. For each possible combination of the OMB (included as the answer to  $\mathcal{Q}_9$ ) and the identity of the god (in the first column), the Boolean values corresponding to each question are shown. The calculation proceeds sequentially from left to right in each row. The value resulting from the action of the logical personality operator of the addressed god is then computed, and the final answer is derived based on this outcome.

Alternatively, if Alice receives the `bit-flip` reply, the situation becomes more ambiguous. Both `True` and `False` will always respond with `bit-flip`, while `Random` does so with a 50% probability. Alice knows, however, that the permutation more likely to have produced this answer is one of  $\{\pi_1, \pi_2, \pi_3, \pi_4\}$ , since these permutations always result in the answer `bit-flip`, whereas  $\pi_5$  and  $\pi_6$  yield this answer only half of the time. Therefore, Alice randomly selects one permutation from  $\{\pi_1, \pi_2, \pi_3, \pi_4\}$  to place her bet, yielding a success probability of  $(2/3)(1/4) = 1/6$ .

Thus, combining both scenarios—success when receiving the `identity` as an answer and success when receiving the `bit-flip` as an answer—the overall probability that Alice correctly identifies the permutation after just one question, using this simple strategy, is identical to that obtained by the protocol where the bet is placed after asking Boolos’ or Roberts’ first questions and deducing the permutations left, that is  $P_{\text{succ},1}^{\text{ME}} = (1/3)(1/2)(1/2) + (2/3)(1/4) = 1/4$ . This result can be easily confirmed through numerical simulation.

More interestingly, we now turn our attention to the success probabilities when two questions are allowed. In particular, we will analyze two distinct strategies: *minimum-error discrimination* and *unambiguous discrimination*.

### Minimum-error discrimination

Because in both the solutions of Boolos [143] and Roberts [144] we are left with two candidate permutations after the second question, if Alice guesses at random at that point, she has a chance of selecting the correct permutation equal to

$$P_{\text{succ},2}^{\text{ME}} = \frac{1}{2}. \quad (2.16)$$

| God   | $\mathcal{Q}_9$ | $U(\mathcal{Q}_9)$ | Answer   |
|-------|-----------------|--------------------|----------|
| True  | 1               | 1                  | identity |
| True  | 0               | 0                  | identity |
| False | 1               | 0                  | bit-flip |
| False | 0               | 1                  | bit-flip |

Table 2.5: Truth table and received state for an answer to  $\mathcal{Q}_9$  from either  $\mathcal{T}\text{rue}$  or  $\mathcal{F}\text{alse}$  in both possible OMB scenarios.




Figure 2.3: **Classical two-question minimum-error discrimination strategy to solve the HLPE.**

Traversing the tree from its root, addressed gods for each configuration are denoted by each node. Arrows denote the selection paths, with the specific prior probabilities denoted with  $p(\cdot)$ . At each node, the probability of successfully identifying the addressed god with the described protocol is denoted by  $P_{\text{succ}}$ . The overall success probability in determining the identity of all the gods using this strategy is  $P_{\text{succ}}^{\text{ME}} = 1/2$ , and it can straightforwardly be confirmed with a numerical simulation.

However, when she places her bet, she does not know if she got it right until the gods reveal the result of the bet to her. She will know, however, the position of at least one of the gods: depending on the answers she receives along the way with Boolos strategy, she will know the positions of either the  $\mathcal{T}\text{rue}$  or the  $\mathcal{F}\text{alse}$  god. This is because the candidate permutations will either be  $\{\pi_3, \pi_5\}$  or  $\{\pi_1, \pi_6\}$  if the second question was addressed to  $B$ , or  $\{\pi_4, \pi_6\}$  or  $\{\pi_2, \pi_5\}$  if the second question was addressed to  $C$ .

It is again important to emphasize that this success probability can also be achieved through a very simple probabilistic reasoning protocol. In this protocol, Alice asks two sequential questions, where she expects a response with the `identity` operation from  $\mathcal{T}\text{rue}$  and with the `bit-flip` operation from  $\mathcal{F}\text{alse}$ , with  $\mathcal{R}\text{andom}$  returning either response with equal probability. This question could be, for instance,  $\mathcal{Q}_9$ , as the simple analysis in Table 2.5 shows.

When she receives `identity` as an answer, she assumes the queried god to be  $\mathcal{T}\text{rue}$ , unless this is the second question and she has already received `identity` in response to the first one, in which case

she attributes the answer to **Random**. Similarly, if she receives a bit-flip as an answer, she should assume the god to be **False**, unless it is her second bit-flip response, in which case she assumes that the god is **Random**.

To compute the probability of correctly identifying the permutation using this protocol, one must consider all possible outcomes across the six potential permutation scenarios. The probability tree diagram illustrated in Fig. 2.3 demonstrates this calculation. As shown, the overall probability of success sums to 50%, which is identical to the probability obtained by betting on one of the two remaining permutations after employing the first two questions in both Boolos' and Roberts' solutions.

### Unambiguous discrimination

Suppose Alice wishes to place a bet only if she is certain she can make the correct choice after using the two questions. This is, in fact, achievable. To demonstrate how, consider the first question in Boolos' solution, which helps her identify a god who is definitely not **Random**. This god must be, therefore, either **True** or **False**. Let us first assume that the god identified as *not Random* occupies position B. After this identification, the remaining possible permutations are  $\{\pi_1, \pi_3, \pi_5, \pi_6\}$ .

As a result, to identify the correct permutation unambiguously, Alice must place a bet on one of these remaining permutations and ask her second question in a manner that gives her a chance to uniquely determine it from the answer value. If she is lucky, this strategy will allow her to identify the correct permutation. However, she faces a probability of failure in three out of four cases. Consequently, the probability of success in this scenario is

$$P_{\text{succ},2}^{\text{UN}} = \frac{1}{4}, \quad (2.17)$$

where the superscript UN stands for *unambiguous discrimination*.

To achieve this, the following strategy can be employed. After the first question has excluded the other two permutations, Alice should ask question  $\mathcal{Q}_{13}$  below to the god in position B:

$$\text{Are you True and C is Random?} \quad (\mathcal{Q}_{11})$$

$$\text{Does bit-flip mean TRUE?} \quad (\mathcal{Q}_{12})$$

$$\mathcal{Q}_{11} \text{ iff } \mathcal{Q}_{12} ? \quad (\mathcal{Q}_{13})$$

The logical analysis of the received answer in all possible scenarios is presented in Table 2.6(a). As shown, if Alice receives **identity** as the answer, she can confidently deduce that the hidden permutation of the gods is  $\pi_5$  and place her bet with the certainty that it will solve the HLPE in two questions. However,

|         | (A,B,C)                                      | $\mathcal{Q}_{11}$ | $\mathcal{Q}_{12}$ | $\mathcal{Q}_{13}$ | $U(\mathcal{Q}_{13})$ | Reply |
|---------|----------------------------------------------|--------------------|--------------------|--------------------|-----------------------|-------|
| $\pi_1$ | $(\mathfrak{T}, \mathfrak{F}, \mathfrak{R})$ | 0                  | 1                  | 0                  | 1                     | bf    |
|         |                                              |                    | 0                  | 1                  | 0                     | bf    |
| $\pi_3$ | $(\mathfrak{F}, \mathfrak{T}, \mathfrak{R})$ | 1                  | 1                  | 1                  | 1                     | bf    |
|         |                                              |                    | 0                  | 0                  | 0                     | bf    |
| $\pi_5$ | $(\mathfrak{R}, \mathfrak{T}, \mathfrak{F})$ | 0                  | 1                  | 0                  | 0                     | id    |
|         |                                              |                    | 0                  | 1                  | 1                     | id    |
| $\pi_6$ | $(\mathfrak{R}, \mathfrak{F}, \mathfrak{T})$ | 0                  | 1                  | 0                  | 1                     | bf    |
|         |                                              |                    | 0                  | 1                  | 0                     | bf    |

(a)

|         | (A,B,C)                                      | $\mathcal{Q}_{14}$ | $\mathcal{Q}_{12}$ | $\mathcal{Q}_{15}$ | $U(\mathcal{Q}_{15})$ | Reply |
|---------|----------------------------------------------|--------------------|--------------------|--------------------|-----------------------|-------|
| $\pi_2$ | $(\mathfrak{T}, \mathfrak{R}, \mathfrak{F})$ | 0                  | 1                  | 0                  | 1                     | bf    |
|         |                                              |                    | 0                  | 1                  | 0                     | bf    |
| $\pi_4$ | $(\mathfrak{F}, \mathfrak{R}, \mathfrak{T})$ | 1                  | 1                  | 1                  | 1                     | bf    |
|         |                                              |                    | 0                  | 0                  | 0                     | bf    |
| $\pi_5$ | $(\mathfrak{R}, \mathfrak{T}, \mathfrak{F})$ | 0                  | 1                  | 0                  | 1                     | bf    |
|         |                                              |                    | 0                  | 1                  | 0                     | bf    |
| $\pi_6$ | $(\mathfrak{R}, \mathfrak{F}, \mathfrak{T})$ | 0                  | 1                  | 0                  | 0                     | id    |
|         |                                              |                    | 0                  | 1                  | 1                     | id    |

(b)

Table 2.6: Exhaustive case-by-case logical analysis of the classical strategy to unambiguously discriminate the correct permutation of the gods and solve the HLPE in two questions with a 25% probability of success. For each permutation scenario, the Boolean values corresponding to each question are shown. The calculation proceeds sequentially from left to right in each row. The value resulting from the action of the logical personality operator of the addressed god is then computed, and the final answer is obtained based on this outcome.

when she receives the alternative answer — which occurs three out of four times — she can only randomly choose between one of the three remaining permutations, giving her a  $1/3$  chance of correctly identifying it. Despite this, this strategy may still be preferable to the protocol discussed in the previous minimum-error discrimination section. Here, Alice benefits from a “best of both worlds approach”: she has 25% chance of unambiguously determining the correct permutation, while still maintaining an overall success probability of 50%. This is evident by summing the probabilities of unambiguous discrimination and a correct guess when unambiguous discrimination fails, yielding a total probability of success equal to  $P_{\text{succ}}^{\text{ME}} = (1/4) + (3/4)(1/3) = 1/2$ .

Finally, since we initially assumed that the god identified as *not*  $\mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}$  occupied position B, we still need to address the case where Alice deduces from the first question that this position is, in fact, C. The analysis follows a similar pattern. In this situation, she proceeds by asking the following question,  $\mathcal{Q}_{15}$ , to the god at position C:

$$\text{Are you } \mathfrak{T}\mathfrak{r}\mathfrak{u}\mathfrak{e} \text{ and B is } \mathfrak{R}\mathfrak{a}\mathfrak{n}\mathfrak{d}\mathfrak{o}\mathfrak{m}? \quad (\mathcal{Q}_{14})$$

$$\mathcal{Q}_{14} \text{ iff } \mathcal{Q}_{12} ? \quad (\mathcal{Q}_{15})$$

The comprehensive logical analysis of the possible answers is presented in Table 2.6(b), leading to the same probability of success as calculated previously, following the straightforward adaptation of the unambiguous bet from  $\pi_5$  to  $\pi_6$ .

In the following sections, we will explore a fundamentally different medium of communication between Alice and the gods, namely quantum communication, examining whether the minimum-error and

unambiguous discrimination probabilities in Eqs. (2.16) and (2.17) can be improved to gain an advantage.

## 2.5 Quantum communication in the HLPE

As discussed in § 2.3.1, both the original HLPE solution due to Boolos [143] and all subsequent analyzes [144, 145, 146, 147, 148, 149, 152, 154] left the physical substrate for the communication rounds between the two parties unspecified, defaulting to classical communication. However, what happens if the god flips a box containing a hidden system inside it that is not a *classical* but a *quantum* coin? In this section, we scrutinize the implicit assumption of classicality in previous solutions by exploring the use of the simplest quantum mechanical system as a communication medium: a *two-level system*, or *qubit*. Just like a classical two-state system, a quantum two-level system can also be used for communication while adhering to the rules of the HLPE formulation in § 2.2. However, it is governed by fundamentally different physical laws, with no counterpart in classical physics. This offers a new perspective on the communication model, the implications of which will be explored in the rest of this chapter.

Qubits can be physically implemented using either fermions or bosons<sup>23</sup>. The minimum spin a fermion can have is  $\frac{1}{2}$ , as exemplified by the electron, making its two perfectly distinguishable spin states — spin-up and spin-down — an ideal platform for representing the qubit. In the case of spin 1 bosons, while massive spin-1 particles possess three spin projections ( $m = -1, 0, +1$ ), for massless spin-1 particles, such as photons, the  $m = 0$  spin projection is not a physical degree of freedom and only two physical spin (helicity) states are left. These two states also span a two-dimensional Hilbert space that can be used to encode a qubit, commonly represented in practice by horizontal and vertical polarization states, which form a convenient basis. Such bosonic qubits are frequently considered for quantum communication and cryptographic protocols, particularly using quantum optical circuits. But how do qubits differ from bits?

Classical systems live in a phase space where states are mutually exclusive, meaning the system definitely occupies only one state at any time. This property conveniently allows for encoding the Boolean values in different states, following the conventions discussed in § 2.3.2. In contrast, the state space of quantum mechanical systems is a complex vector space, specifically a Hilbert space. Quantum states are represented by vectors which can be linearly added to create new states, reflecting the classically-absent *superposition principle*. In the case of the qubit, the Hilbert space is a 2-dimensional complex vector space

---

<sup>23</sup>All particles in nature are either fermions or bosons. This distinction is based on their intrinsic spin, which is a half-integer ( $\frac{1}{2}, \frac{3}{2}, \dots$ ) for fermions and an integer ( $0, 1, 2, \dots$ ) for bosons.

$\mathcal{H} = \mathbb{C}^2$ , and its general *pure state*, or *wavefunction*, is mathematically described as

$$|\psi\rangle = c_0|0\rangle + c_1|1\rangle, \quad |\psi\rangle \in \mathcal{H} \quad (2.18)$$

where the constants  $c_0, c_1 \in \mathbb{C}$  obey the normalization condition  $|c_0|^2 + |c_1|^2 = 1$ . The states  $\{|0\rangle, |1\rangle\}$  form an orthonormal *basis* for  $\mathcal{H}$ , and, as we will explain, correspond to the *observable* states of some physically measurable quantity.

Another fundamental feature of quantum systems, which profoundly distinguishes the qubit from the bit, concerns their observability and predictability. Although the time evolution in both quantum and classical systems is deterministic, the continuous Hilbert space in which quantum states are defined remains fundamentally inaccessible. Observing the qubit requires a *measurement* operation, which non-deterministically leads to the *collapse of the wavefunction* into one of the two possible observable states  $|x\rangle \in \{|0\rangle, |1\rangle\}$ . As postulated by the *Born rule* [163, 164, 165, 166], the probability of measuring each of these states and obtaining the corresponding encoded Boolean value  $x \in \{0, 1\}$  is given by

$$P(x) \equiv |\langle x|\psi\rangle|^2 = |c_x|^2, \quad (2.19)$$

or, equivalently,

$$P(x) \equiv \text{Tr}\{\rho_\psi \Pi_x\}, \quad (2.20)$$

where we use the formalism based on the *density operator*  $\rho_\psi = |\psi\rangle\langle\psi|$  and the *projective measurement* operator  $\Pi_x = |x\rangle\langle x|$ . This ontological nature of randomness in quantum mechanics is in stark contrast with classical physics, in which randomness is an epistemic issue, as discussed in § 2.3.3.

Because of the Born rule, global phases in the wavefunction have no physical significance as they cannot be observed. That is, two quantum states  $|\psi\rangle$  and  $e^{i\theta}|\psi\rangle$  actually represent the same state because the measurement probabilities derived from them yield the same values. More importantly, different quantum states can result in the same measurement probabilities on a given basis. For instance, both  $|\psi_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$  and  $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$  yield  $P(0) = P(1) = \frac{1}{2}$ , rendering them indistinguishable under this measurement, a property that we will come to discuss again below.

Before that, however, the question of how to define the measurement operation that establishes the  $\{|0\rangle, |1\rangle\}$  basis states remains. We first observe that encoding the Boolean values in so-called *orthogonal* states is indeed required for the qubit to be able to supersede the classical bit and process information following classical logic. This ensures that when the qubit is prepared in one of the eigenstates of the measured operator, say state  $|0\rangle$ , its measurement will never return the complementary

state, i.e.,  $P(1) \equiv |\langle 1|0\rangle|^2 = 0$ , thereby giving the qubit a mutually exclusive and deterministic nature necessary for classical logic. To define these basis states, we can, for instance, take  $|0\rangle$  and  $|1\rangle$  to represent the eigenstates of the Pauli-Z operator (see Eq. 2.22) corresponding to the  $+1$  and  $-1$  eigenvalues, respectively. However, the measurement operation can also be performed along a different basis — such as a different spatial coordinate axis — since observables on the qubit correspond to Hermitian operators  $M \in \mathcal{H}$  satisfying  $M = M^\dagger$ . The set of Hermitian  $2 \times 2$  matrices forms a four-dimensional real vector space, for which the identity and the three Pauli matrices constitute a convenient basis. As a result, any qubit observable can be uniquely expressed as a linear combination

$$M = c_0 \mathbb{1} + c_1 X + c_2 Y + c_3 Z, \quad (2.21)$$

where  $c_j \in \mathbb{R}$  and the Pauli matrices are represented as usual,

$$\mathbb{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (2.22)$$

For effective communication between the two parties, a *computational basis* needs to be established in which to define the qubit's states and operations, which involves deciding the exact form of  $M$ .

The definition of the computational basis for a qubit can be superficially compared to the parameters that establish the two states in a classical system. For instance, in the coin-in-a-box analogy, flipping its state or measuring its orientation also requires defining the relevant spatial direction, in a direct analogy to selecting a measurement axis for the qubit. Similarly, in the conservative pendulum example mentioned in § 2.3.2, encoding binary values into two effective discrete states, defined by closed trajectories in phase space, requires selecting the energy difference between them. This is analogous to the way voltage levels are agreed upon in digital circuits to represent classical binary states.

Having defined a computational basis, we can return to the earlier observation that the Born rule implies all quantum states with the same projection onto it yield identical probabilities for the measured states. In particular, the probability of measuring either state can be made  $1/2$  by splitting the wave function equally between the basis, as in the examples of  $|\psi_1\rangle$  and  $|\psi_2\rangle$  above. This feature enables the construction of elementary operations that transform the qubit state to yield a uniformly random truth value upon measurement, providing an ontological realization of  $\mathfrak{R}andom$ 's logical personality operator in Eq. 2.3. In contrast, with a classical two-state system,  $\mathfrak{R}andom$ 's probabilistic bit-flip must be induced

through an abstract epistemic construct rather than a fundamental physical process.<sup>24</sup>

In other words, with quantum communication assumed for the realization of (HLPE<sub>syn</sub><sup>2</sup>),  $\mathfrak{R}\text{andom}$  does not need to prepare one of the orthogonal states of the qubit using classical randomness before sending it. Instead, he can prepare the qubit in a basis that is mutually unbiased with respect to the one Alice is expecting, sending it in a state that is inherently random upon measurement from her perspective. At readout, the result is probabilistically equivalent to that generated by a random classical coin. However, the source of this randomness is now entirely different and more fundamental given the inherently quantum mechanical nature of reality [167].

Given the above, in order to maintain parity with the classical implementation of the HLPE, the assumptions required for a quantum-communication-based solution are simply quantum analogues of the previously considered  $A_{3'}$  and  $A_4$  and  $A_6$ :

The sender has access to a *quantum* bit-flip operation that maps any previously shared computational basis state to its orthogonal state. ( $A_{3''}$ )

The receiver knows how to interact with the qubit system to read its state. ( $A_{4'}$ )

The sender has access to a *quantum* rand-flip operation that maps any previously shared computational basis state to a state that is maximally unbiased with respect to the computational basis. ( $A_{6'}$ )

The next section introduces a remarkable solution to the puzzle in just two questions under a particular implementation of each of these operators.

## 2.6 A one of a kind solution to the HLPE in two questions

Having established the correspondence between classical and quantum implementations of communication between the involved parties in the HLPE, we are now positioned to provide an explicit form for the identity, bit-flip and rand-flip operators, consistent with Assumptions  $A_{3''}$ ,  $A_{4'}$ , and  $A_{6'}$ . We begin by considering the following quantum oracles:

$$\mathcal{O}_{\text{id}} = \mathbb{1}, \tag{2.23}$$

$$\mathcal{O}_{\text{bf}} = X, \tag{2.24}$$

$$\mathcal{O}_{\text{rd}} = \frac{Y + Z}{\sqrt{2}}. \tag{2.25}$$

---

<sup>24</sup>There is also the possibility that the bit-flip value is sampled from an ontological quantum source controlled by  $\mathfrak{R}\text{andom}$  and transmitted through a classical system. However, this scenario would not be classified as quantum communication.

The implicit assumption in the design of these operators is that both parties encode logical values in the basis states of the Pauli-Z operator. When these operators act on a basis state  $|x\rangle$ , which encodes the Boolean value, and this state is subsequently measured, it is straightforward to confirm that the resulting measurement probabilities *indistinguishably* correspond to those of the equivalent classical logic operators, a crucial point in establishing the correspondence with the classical scenario:

$$|\langle x | \mathcal{O}_{\text{id}} | x \rangle|^2 = 1, \quad (2.26)$$

$$|\langle x | \mathcal{O}_{\text{bf}} | x \rangle|^2 = 0, \quad (2.27)$$

$$|\langle x | \mathcal{O}_{\text{rd}} | x \rangle|^2 = 1/2. \quad (2.28)$$

However, the underlying implementation of these logical propositions is quantum mechanical in nature, and with the right knowledge, it can be exploited. Quantum theory presents distinctive features from classical theory beyond those already discussed in the previous section. When multiple qubit states are combined, they can be represented through the tensor product, which enlarges the Hilbert space to describe the composite system, accounting for all possible combinations of the individual states. More importantly, qubits can also become *entangled* when their states develop correlations that cannot be described independently of one another, challenging classical intuitions about separability and locality.

Armed with this knowledge, Alice uses an entangled qubit pair to eavesdrop on the behavior of the god with or without his awareness. The strategy she implements is based on the superdense coding protocol (SDC) [168]. It begins with the preparation of a maximally entangled pair of qubits in the Bell state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \quad (2.29)$$

To prepare  $|\Phi^+\rangle$ , Alice can take two unentangled qubits initialized in a  $|0\rangle \otimes |0\rangle$  product state and apply an Hadamard gate on the first qubit followed by a CNOT gate controlled by the first qubit with a target on the second, as shown in the first part of the circuit in Fig. 2.4 and represented below

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad (2.30)$$




Figure 2.4: **Superdense-coding-based questioning protocol.** Extracts the identity of a queried god in a single question by eavesdropping on the oracle employed to encode the answer in the communication qubit.  $\mathcal{O}$  represents one of  $\mathcal{O}_{\text{id}}$ ,  $\mathcal{O}_{\text{bf}}$  or  $\mathcal{O}_{\text{rd}}$  operators. The protocol enables a strategy to solve the  $(\text{HLPE}_{\text{syn}}^2)$  exactly in just two questions.

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (2.31)$$

After that, she simply poses  $\mathcal{Q}_9$  to the god in position  $A$  as her first question, while also sending the first qubit in the Bell state pair she prepared for him to encode the answer according to the agreed protocol. Following the procedure we described in § 2.2, the god then evaluates the truthful answer to the question, processes it with his logical personality operator, and encodes the final evaluation of  $\mathcal{Q}_9$  in the shared qubit following his OBM language encoding before sending it back.

This question is directed to one of three possible gods, and there are two possible OBMs, resulting in six possible scenarios in which its answer will be computed. However, the possible outcomes are simpler.  $\mathfrak{R}\text{andom}$  will always answer in the same way to any question, acting with the  $\mathcal{O}_{\text{rd}}$  oracle on the input state. Both  $\mathfrak{T}\text{rue}$  and  $\mathfrak{F}\text{alse}$  can, in general, act using either the  $\mathcal{O}_{\text{id}}$  or the  $\mathcal{O}_{\text{bf}}$  oracle; however, in response to this question, they answer differently from each other while providing the same answer regardless of the language. These outcomes for Question  $\mathcal{Q}_9$  were already presented in the exhaustive step-by-step logical analysis shown in Table 2.5.

The key insight is that Alice can distinguish these oracles quantum mechanically using the SDC protocol. Since each of the three gods responds with a different oracle, identifying the employed oracle enables Alice to map it directly to the corresponding god, thus identifying it with a single question. To understand

how the SDC protocol can distinguish each oracle, we analyze the transformation of the joint qubit state by first computing its output after the queried god applies the encoding oracle  $\mathcal{O}$ . Depending on whether  $\mathcal{O}$  is  $\mathcal{O}_{\text{id}}$ ,  $\mathcal{O}_{\text{bf}}$ , or  $\mathcal{O}_{\text{rd}}$ , Alice can obtain from  $(\mathcal{O} \otimes \mathbb{1})|\Phi^+\rangle$  one of the

$$|\Psi_{\mathfrak{T}}\rangle = |\Phi^+\rangle, \quad (2.32)$$

$$|\Psi_{\mathfrak{F}}\rangle = |\Psi^+\rangle, \quad (2.33)$$

$$|\Psi_{\mathfrak{R}}\rangle = \frac{1}{\sqrt{2}}(|\Phi^-\rangle - i|\Psi^-\rangle), \quad (2.34)$$

states, respectively, which map directly to the identities of **True**, **False**, and **Random**. The first qubit in the tensor product is the one acted upon by the god while the second is the one that Alice retains. The other three Bell states are given by

$$|\Phi^-\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \quad (2.35)$$

$$|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), \quad (2.36)$$

$$|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle). \quad (2.37)$$

Finally, Alice receives the communication qubit back and, together with the qubit she kept on her side, she can locally perform a joint measurement to identify the sender. The measurement can be understood in terms of either the Bell basis or the computational basis. In the computational basis, Alice first applies a CNOT gate with the communication qubit as the control and the second qubit as the target, followed by a Hadamard gate on the communication qubit. She then measures both qubits in the computational basis (i.e., using the Pauli-Z eigenstates), which yields the following outcomes:

- A measurement result of  $|00\rangle$  occurs with probability 1 if the qubit was sent by **True**;
- A measurement result of  $|01\rangle$  occurs with probability 1 if the qubit was sent by **False**;
- Measurement results of  $|10\rangle$  or  $|11\rangle$  occur with  $\frac{1}{2}$  probability if the qubit was sent by **Random**.

As a result, Alice can deterministically identify the queried god with a single question. The advantage arises from the fact that, although a qubit can inherently store only one bit of information, the protocol enables her to retrieve two classical bits by using only one qubit of communication entangled with an ancillary qubit for information extraction. With the identity of the god in position A determined, Alice

repeats the exact same protocol to identify **B**, ultimately solving the HLPE exactly in two questions by deducing **C** through exclusion. The probability of success is evidently 1, beating any result presented in § 2.4.2 for two question solutions.

It is essential to conclude this section by delving deeper into the remarkable features of the solution Alice has accomplished. To start with, this solution employs logical queries to the gods that are entirely consistent with classical logic. Specifically, it relies on straightforward, non-subjunctive and non-contrary-to-fact<sup>25</sup> questions that are purely truth-functional — meaning that the truth value of the addressed propositions is determined by the values of their components and the logical operators that connect them. It follows, therefore, while complying with the law of the excluded middle — highlighted by Boolos as crucial to his solution — which, in turn, implies that not all binary questions are allowed and Alice restricts herself to the set of valid questions only, as defined in (B0). This surprising result would have seemed impossible had we assumed only classical communication was allowed.

The second observation that cannot be left unnoted is that information extraction about the identity of the gods is achieved even if each of them has his own language independently of his peers, that is, a different OBM. The third key insight is that the solution is achieved without requiring the gods to have any knowledge of their own identities, since they are not asked about the identity of each other or even of themselves. The formalization studied by Liu and Wang [149] assumes not only condition (B3\*) but also restricts questions to those that can always be answered with “da” or “ja”, as required by our condition (B0), which aligns precisely with the HLPE formulation examined in this work. Thus, although Liu and Wang proved, within their framework, the non-existence of solutions for versions of the HLPE where the gods do not know each other’s identities, we show — perhaps surprisingly — that it is not only possible to solve this problem without requiring the gods to know each other’s identities but also without them being aware of their own identities, besides achieving a solution in only two questions.

Finally, it is remarkable that a solution incorporating all these features is possible while also breaking the classical communication limit of three questions. Arguably, this strategy became possible by assuming specific physical forms for the operators implementing identity, bit-flip and randomization transformations to encode the answer in the communication qubit, as detailed in Eqs. (2.23 – 2.25). As we will see in the next section, these forms, satisfying Assumptions  $A_{3''}$ ,  $A_{4'}$ , and  $A_{6'}$ , are not unique. However, it is impor-

---

<sup>25</sup>Subjunctive or counterfactual questions often lead to problematic situations in logic, as they involve reasoning about hypothetical scenarios that might be poorly defined or non-factual and require additional assumptions rather than allowing straightforward answers grounded in direct reality. As discussed in § 2.1, in several of the previous solutions the gods are asked to predict how their peers would respond to a particular question, for instance. These approaches exploit the potential for third-answer values arising from logical inconsistencies or the impossibility of the gods to predict those answers, and might require adding explicit assumptions about giving an oracular ability for the gods to foresee hypothetical events.

tant to recognize that the bits recovered when employing these operations replicate classical logic and the puzzle’s setup in a way that remains indistinguishable from a classical realization. Furthermore, even if Alice has knowledge of the operations in the classical scenario, it would be impossible for her to identify which one was applied without disturbing the system and risking detection, unlike in the quantum context. For instance, if we postulate that the `bit-flip` operation, in a classical implementation, consists of a  $180^\circ$  rotation around the  $X$  axis — in an analogy with the quantum operator in Eq. 2.24 — this knowledge does not enable Alice to devise a protocol to determine whether `bit-flip` or `rand-flip` was applied upon receiving the coin. In contrast, the quantum solution allows eavesdropping on the applied operation without direct interaction by tapping into the fabric of quantum superposition and entanglement — purely quantum resources<sup>26</sup>.

It is already enlightening to observe that a solution to the HLPE is possible in just two questions. This example highlights a logic problem where we can definitively prove that an advantageous solution is attainable by employing physical principles, and that such a solution is impossible within the realm of purely abstract logical relationships. Nevertheless, in the remaining sections, we will introduce further complications to this scenario by removing the critical assumption of knowing the exact form of the quantum operations used by the gods to encode their answers as well as consider the unavoidable presence of noise in real communication systems, investigating whether a quantum advantage can still be achieved.

## 2.7 A generalized quantum implementation

Although we demonstrated in the previous section that there is an instance where a quantum implementation of the HLPE has an exact solution in two questions, outperforming any possible classical solution, this result depends critically on defining the specific operations performed by the gods. The chosen operations make superdense coding feasible, but for its effective application, Alice must also have prior knowledge of them rather than treating the operations as a black boxes. Her ability to identify the operators in Eqs. (2.23 – 2.25) is precisely what allows her to solve the challenge efficiently.

However, these are not the only operators that could have been chosen to satisfy Assumptions  $A_3'$ ,

---

<sup>26</sup>Since we are targeting the  $(HLPE^2_{syn})$  formulation, which became the standard after the publication by Rabern and Rabern, `Random` is considered to be a fully random agent, as formulated in § 2.2. Therefore, implementing it with the operator in Eq. (2.25) is perfectly compatible with the formulation. Interestingly, however, if `Random` were actually deciding first whether to act like `True` or like `False`, using one of Eq. (2.23) or Eq. (2.24) operators to encode the answer after deciding which identity to assume at random, this protocol would not solve the problem in two questions anymore. Therefore, paradoxically, Boolos’ original problem, as formulated in  $(HLPE_{sem})$  and shown to be solvable in three “trivial” questions in the Rabern and Rabern publication, would resist this quantum-communication-based solution strategy built on top of the SDC protocol.

$A_4$ , and  $A_6$ . For instance, an alternative choice such as

$$\mathcal{P}_{\text{id}} = Z, \quad (2.38)$$

$$\mathcal{P}_{\text{bf}} = Y, \quad (2.39)$$

$$\mathcal{P}_{\text{rd}} = \frac{1+Z}{\sqrt{2}}, \quad (2.40)$$

would yield equivalent logical propositions when applied to the computational basis states  $|x\rangle$  as before:

$$|\langle x | \mathcal{P}_{\text{id}} | x \rangle|^2 = |\langle x | \mathcal{O}_{\text{id}} | x \rangle|^2 = 1, \quad (2.41)$$

$$|\langle x | \mathcal{P}_{\text{bf}} | x \rangle|^2 = |\langle x | \mathcal{O}_{\text{bf}} | x \rangle|^2 = 0, \quad (2.42)$$

$$|\langle x | \mathcal{P}_{\text{rd}} | x \rangle|^2 = |\langle x | \mathcal{O}_{\text{rd}} | x \rangle|^2 = 1/2. \quad (2.43)$$

Yet, employing  $\mathcal{P}_{\text{id}}$ ,  $\mathcal{P}_{\text{bf}}$ , and  $\mathcal{P}_{\text{rd}}$  renders the previous solution protocol unfeasible, as  $\mathcal{P}_{\text{rd}}$  is constructed with a superposition that includes  $\mathcal{P}_{\text{id}}$ , creating ambiguity and the potential for these operators to be misidentified as one another by the protocol.

In fact, as discussed in § 2.5, the same classical logic operations can be implemented by different quantum operators that return output states with the same projection onto the computational basis. From a classical perspective, all the states with an equal projection onto the computational basis are equivalent, as they yield the same measurement probabilities for the represented Boolean values [169].

This observation raises an issue. In classical-communication-based solutions to the HLPE, specifying the physical implementation of the identity and bit-flip operations is unnecessary because eavesdropping on them without somehow observing or interacting with the transmitted bit at the moment of the operation is fundamentally unfeasible from a physics standpoint. Different physical implementations of these operations are still possible; however, the only relevant physical object for logical reasoning is the output state, which encodes the Boolean value of ultimate interest.

In contrast, for the earlier quantum communication protocol based on superdense coding, carefully selecting specific implementations of the operators is essential to enable an effective eavesdropping strategy, which becomes possible remotely, without direct interaction with the qubit or observation of the operation applied by the god. Therefore, a more generalized quantum implementation of the HLPE would address this limitation and preserve a solution advantage even when different quantum operators enacting equivalent classical logic are selected to realize identity, bit-flip, and rand-flip operations.

To more closely align classical and quantum realizations of the HLPE while minimizing additional




Figure 2.5: **Two approaches to generalize the computational basis.** **a**, The computational basis is predetermined and shared by Alice and the god, becoming a shared resource. Any state in the equatorial plane of the Bloch sphere can represent the random state since all such states project equally onto the computational basis, yielding identical measurement probabilities. **b**, The computational basis is relaxed to a less strict common assumption that it may consist of any pair of orthogonal states lying in the  $XZ$  plane of the Bloch sphere. Only two states can encode the random answer, obtained by rotating either basis state by  $\frac{\pi}{2}$  or  $-\frac{\pi}{2}$  radians around the  $Y$ -axis.

assumptions, a more general implementation is necessary. We identify two primary generalization approaches, each suited to a different scenario or application:

### 1. Fixing the computational basis

This approach involves the mutual agreement between Alice and the gods to align the qubit's measurement axis along a fixed direction—typically, and without loss of generality, the  $Z$ -axis, as has been the convention thus far. Fixing this axis defines the  $|0\rangle$  and  $|1\rangle$  basis states, along with an equatorial plane on the Bloch sphere containing all equivalent random states, as illustrated in Fig. 2.5a. Different bit-flip and rand-flip operations can be implemented to transition between these states by tracing different paths through Hilbert space. A robust communication approach would account for all quantum operations that achieve a final state in any of the equivalent classes, regardless of the path taken.

This diversity of compatible operations possible in a fixed-basis framework has a clear analogy in our classical coin example. A coin aligned along the  $Z$ -axis can be flipped by rotating it  $180^\circ$  around any axis in the  $XY$ -plane. It may also undergo multiple rotations, as long as it ultimately flips. Similarly, two effective states of a conservative classical pendulum can be transitioned into each other from different trajectories connecting them in the phase space.

This generalized approach provides flexibility in reaching the desired final qubit state without imposing unnecessary operational constraints that, in any case, do not interfere with the implemented

logic. The possibility of preserving quantum advantage within this framework could be explored.

## 2. Relaxing the computational basis

The previous framework removes shared assumptions about exact quantum operations, but only after fixing a shared computational basis (SCB). As we first argued in § 2.3.2, fixing a SCB is an assumption equivalently present in the classical version of the HLPE, where it appears when the syntax is fixed by informing Alice that each god’s answer will be one of “da” or “ja”, or when defining the two-state system’s effective states, or the logical operations built on top of them.

However, while this approach may seem like a reasonable generalization towards an ultimate solution, fixing a SCB is not physically without cost or limited to a single choice. In one approach, fixing the SCB requires first agreeing on a shared reference frame (SRF) to ensure that both parties align on the directional axes for qubit measurements. It is not a trivial matter to establish this common coordinate system at a distance. In fact, it is an example of a type of information processing task that cannot be achieved with *fungible* information<sup>27</sup>. That is, there is no way to abstractly describe a direction in space using a sequence of classical bits without resorting to external common references. It requires exchanging physical resources — *nonfungible* information — along with previously shared protocols to ensure proper encoding and decoding [170]. Specifically,  $N$  entangled qubits must be transmitted to communicate a single direction [171, 172] or a complete Cartesian frame [173] encoded intrinsically in their polarization states by the emitter and measured optimally by the receiver. A reference frame can also be communicated by transmitting a system with an infinite-dimensional Hilbert space such as an hydrogen atom instead [174], or simply one qubit if both parties previously share an entangled two-qubit state that is consumed to establish the SRF [175]. Even in noiseless channels, these procedures do not achieve perfect fidelity.

On the other hand, there are also protocols for both classical and quantum communication that do not require a SRF to be established first and, importantly, can do so with perfect fidelity. These involve encoding rotation-invariant *logical* qubit states into entangled states of multiple *physical* qubits [176, 170]. However, even in this case, communication is only possible by encoding and decoding the logical qubits following a shared protocol that agrees on the labeling of each physical qubit within the transmitted system. This represents yet another form of SRF that persists [177].

Furthermore, even if one establishes a SRF, an agreement is still necessary on which direction within it to use for the computational basis, just as we previously considered agreement on the  $Z$

---

<sup>27</sup> *Fungible* (also called *speakeable*) information is data that can be fully replicated regardless of the encoding method. In contrast, *nonfungible* (or *unspeakable*) information is intrinsically tied to the unique physical characteristics of its carrier.

direction. Therefore, returning to the beginning, a simpler solution to limit shared protocols when establishing the SCB, at the expense of increasing physical resources, would be to directly fix a common direction, rather than establishing a full-fledged SRF [171, 172]. This would still require mutual understanding between both parties, as an entangled state of  $N$  qubits must be prepared by one party and measured collectively by the other based on knowledge of the encoding. However, this optimal measurement would be independent of  $N$  and fully symmetric under any permutation of the qubits. Nonetheless, this solution cannot transmit information about the directional axis for the SCB with perfect fidelity. As a result, the established SCB would be slightly biased, potentially leading to the detection of some qubits, which are later transmitted as  $|0\rangle$ , as  $|1\rangle$ .

As we can see, shared physical resources and mutual assumptions are generally required to establish a SCB. Both should thus be minimized for a true advantage. Like most communication protocols, those applied in the absence of a SRF typically assume that both parties are not strangers to each other and can agree on the protocols to use. In the present framework, we avoid this assumption and consider the gods and their conception of the world to be as unknown to Alice as possible. While she may have learned some basic facts about them — such as their ability to think logically and use a qubit for communication — these should be the bare minimum necessary for communication. At the same time, we are still constrained by the limits of physical reality. We focus on the use of these *physical constraints* and *minimal mutual inference* to establish a SCB.

A preliminary hypothesis to attempt this could involve eliminating the need to define a computational basis altogether, instead considering an alternative in which the god assumes only that an unknown single-qubit state has been transmitted. He would then be tasked with either applying *identity*, *bit-flip*, or *rand-flip* depending on the question and his logical personality. However, an obvious complication arises: it is not physically possible to realize an operation that can orthogonalize an arbitrary single-qubit state because it would be antiunitary [178].

Since both parties know this basic quantum mechanics fact, to allow the possibility of negating a more generalized single-qubit state rather than a fixed one, they understand that only one viable solution remains: the transmitted state  $|\psi\rangle$  must lie in a specific great circle plane of the Bloch sphere. Consider then  $|\psi\rangle$  restricted to have real amplitudes, which confines it to the intersection of the Bloch sphere with the  $XZ$  plane, as represented in Fig. 2.5b. For this subset of states, a universal-NOT gate is indeed unitary and can be physically implemented by a  $Y$ -axis rotation of  $\pm\pi$  rad. Similarly, the *rand-flip* operation can be achieved through  $Y$ -axis rotations of  $\pm\frac{\pi}{2}$  rad. This generalization applies equally in classical physics. Our example coin cannot be flipped without

knowledge of its directional orientation, although this need not be fully characterized. It suffices to know that it points radially outward within the  $XZ$  plane, allowing a rotation about the  $Y$ -axis to act as a `bit-flip` operation. This example mirrors the treatment of the qubit.

This framework relaxes the SCB from a fixed direction in the Bloch sphere to an arbitrary direction within a plane passing through its center, requiring only that both parties establish a common axis around which the rotational operations are applied. To accomplish that agreement, the crucial observation comes from recognizing that, in many communication scenarios, an intrinsic directional reference exists between the parties: the geodesic connecting them in space. In other words, when Alice and the god point at each other, a direction is naturally established.

Both generalization scenarios could, in principle, be explored independently, depending on the application at hand. For the remainder of this section, we will focus exclusively on the second scenario, which further minimizes the physical resources and syntactic requirement for communication. In doing so, it also constrains the encoding operations that can be applied on the communication qubit.

### 2.7.1 Relaxing the computational basis

To ensure compatibility with the previously motivated scenario, the operations applied to the communication qubit must be capable of preserving, negating, or randomizing the classical readout of any arbitrary single-qubit state  $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$ , where  $c_0, c_1 \in \mathbb{R}$ . The readout is considered within a measurement basis defined by  $|\psi\rangle$  and its orthogonal counterpart  $|\psi^\perp\rangle = c_1|0\rangle - c_0|1\rangle$ . The required oracles thus implement the necessary classical logic operations based on measurement outcomes in this specific basis. To meet this criteria, the unitary operators are defined as

$$\mathcal{O}'_{\text{id}} = R_y(0), \quad (2.44)$$

$$\mathcal{O}'_{\text{bf}} = R_y(\pi), \quad (2.45)$$

$$\mathcal{O}'_{\text{rd},\pm} = R_y\left(\pm\frac{\pi}{2}\right), \quad (2.46)$$

where irrelevant global phases are ignored and the  $Y$  rotation is defined as usual

$$R_y(\theta) = e^{-i\theta Y/2}. \quad (2.47)$$

To ensure generality, we include two distinct oracles for the `rand-flip` operation. Importantly, each of these operators can identically realize the intended classical logic, and this set is uniquely determined




Figure 2.6: **Communication with an extraterrestrial god.** A photon travels from Alice (left) to the god (right) and back along the geodesic connecting both parties. One of the three available operations in Eqs. (2.44 – 2.46) can be applied to this communication qubit, possibly achieved by placing a birefringent crystal in the optical path of the photon. *True* and *False* have access to two of these polarization rotators, while *Random* always applies only the one that can randomize the measurement outcome of the received quantum state. The reference frames of both parties need not match; it is only required that one direction is common, here set to the  $Y$ -axis, which is effortlessly established by the line connecting them.

by the required conditions. In this framework, Boolean values are encoded by  $|\psi\rangle$  and  $|\psi^\perp\rangle$ , and it is only required that Alice and the gods align on the  $Y$ -axis to perform the necessary  $R_y(\theta)$  rotations.

To provide further motivation for this scenario and a concrete realization supporting a natural alignment of the  $Y$  direction between the parties, consider the diagram in Fig. 2.6. In this hypothetical scenario, Alice and the god need to communicate over vast distances in space. A natural method of establishing communication would be for Alice to send a photon, which the god receives and then redirects back after applying the appropriate operation.

No shared coordinate system has been established between them. However, the geodesic path traversed by the photon can install a clear common direction, which both parties readily identify upon locating each other. The orthogonal plane manifold to this geodesic trajectory forms a natural space to establish a basis of real amplitude states using the horizontal,  $|\leftrightarrow\rangle$ , and vertical,  $|\updownarrow\rangle$ , polarization states of the photon. This basis coincides with the  $X$  and  $Z$  axes of each party's frame, and the quantum state of the photon is independently described by both of them as

$$|\psi\rangle = c_0 |\leftrightarrow\rangle + c_1 |\updownarrow\rangle. \quad (2.48)$$

The constraint that  $c_0, c_1 \in \mathbb{R}$  restricts the photon to a linearly polarized state, in which the answer is encoded with the operations in Eqs. (2.44 – 2.46) that rotate its polarization plane. These operations can be enacted on the photon without altering its polarization type—i.e., without inducing circular or elliptical polarization—by using an optical rotator based on the Faraday effect or by placing an appropriate

birefringent crystal in the photon's optical path.

This implementation of  $(\text{HLPE}_{\text{syn}}^2)$  with the supplemental conditions (B0), (B1), (B2), and (B3\*)<sup>28</sup> is solvable exactly in three questions with classical protocols such as those in Refs. [143, 144], just like any other valid realization. However, contrary to the implementation considering operations in Eqs. (2.23 – 2.25), a complete and deterministic solution in just two questions with this implementation could not be found. In the following section, we present a probabilistic solution strategy in two questions.

## 2.7.2 Minimum-error state discrimination strategy

We would now like to find a probabilistic two-question solution strategy to solve this epistemically relaxed implementation of the HLPE with minimal assumptions about the computational basis. Let us first consider the state that would be returned by applying each of the oracles in Eqs. (2.44 – 2.46) to a hidden input state  $|\psi_{\text{in}}\rangle = |0\rangle$  prepared and sent by Alice, where the definition follows her basis convention. Straightforward algebra shows that there are four possible output states, which are divided into three groups depending on the specific oracle applied to  $|\psi_{\text{in}}\rangle$ , namely  $\mathcal{O}'_{\text{id}}$ ,  $\mathcal{O}'_{\text{bf}}$ ,  $\mathcal{O}'_{\text{rd},+}$  or  $\mathcal{O}'_{\text{rd},-}$ , given by

$$|\psi^{\text{id}}\rangle = |0\rangle, \quad (2.49)$$

$$|\psi^{\text{bf}}\rangle = |1\rangle, \quad (2.50)$$

$$|\psi_{\pm}^{\text{rd}}\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}. \quad (2.51)$$

To obtain these outputs, Alice must ask a question to which she knows **True** replies with  $\mathcal{O}'_{\text{id}}$  and **False** with  $\mathcal{O}'_{\text{bf}}$ . As already shown, this question could be  $\mathcal{Q}_9$ . **Random** invariably applies one of its  $\mathcal{O}'_{\text{rd},\pm}$  oracles to a received state given any question.

The two-question solution strategy that we will consider consists in guessing the permutation of the gods,  $\pi_i, i \in \{1, \dots, 6\}$ , by optimally distinguishing the states returned from two of them using a minimum-error discrimination measurement. The protocol starts with Alice addressing question  $\mathcal{Q}_9$  independently to the gods in positions A and B. When she receives both qubits back with their encoded replies, she simultaneously measures them in an attempt to identify the recovered two-qubit product state. This state maps exactly to a single  $\pi_i$ , allowing her to place a bet on it.

We begin by considering the minimal implementation of the generalized oracles, in which there is only one unitary available to **Random**:  $\mathcal{O}'_{\text{rd},+}$ . The possible output states when applying each oracle operator to  $|\psi_{\text{in}}\rangle$  are  $|\psi^{\text{id}}\rangle$ ,  $|\psi^{\text{bf}}\rangle$ , and  $|\psi_+^{\text{rd}}\rangle$ . The resulting states Alice needs to discriminate, obtained

<sup>28</sup>Redundantly, also condition (B4).

with equal prior probability of  $p_i = 1/6$ , are therefore

$$|\phi_1\rangle = |\psi^{\text{id}}\rangle |\psi^{\text{bf}}\rangle = |01\rangle, \quad (2.52)$$

$$|\phi_2\rangle = |\psi^{\text{id}}\rangle |\psi_+^{\text{rd}}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |01\rangle), \quad (2.53)$$

$$|\phi_3\rangle = |\psi^{\text{bf}}\rangle |\psi^{\text{id}}\rangle = |10\rangle, \quad (2.54)$$

$$|\phi_4\rangle = |\psi^{\text{bf}}\rangle |\psi_+^{\text{rd}}\rangle = \frac{1}{\sqrt{2}}(|10\rangle + |11\rangle), \quad (2.55)$$

$$|\phi_5\rangle = |\psi_+^{\text{rd}}\rangle |\psi^{\text{id}}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |10\rangle), \quad (2.56)$$

$$|\phi_6\rangle = |\psi_+^{\text{rd}}\rangle |\psi^{\text{bf}}\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |11\rangle). \quad (2.57)$$

Given this set of  $N = 6$  states  $\{\rho_i\}_{i=1}^N$ , where  $\rho_i = |\phi_i\rangle \langle \phi_i|$ , we look for probability operators  $\{E_i\}_{i=1}^N$  such that, when Alice measures the output associated with  $E_i$ , she identifies the received state as  $\phi_i$ . Since these states are not orthogonal to one another, error-free discrimination is not possible. Therefore, the goal of a minimum-error state discrimination measurement is to find the probability operators that minimize the total error probability. Equivalently, the goal is to maximize the sum of the probabilities of successfully identifying each state,

$$P_{\text{succ}} = \sum_{i=1}^N p_i \text{Tr}(E_i \rho_i). \quad (2.58)$$

The set of probability operators  $\{E_i\}$  configures a positive operator-valued measure (POVM) [179], which satisfies two conditions:

$$E_i \succeq 0, \quad (2.59)$$

$$\sum_{i=1}^N E_i = \mathbb{1}. \quad (2.60)$$

The first condition is *positivity*, which requires each POVM element  $E_i$  to be a positive semidefinite operator. This means that  $E_i$  satisfies  $\langle \psi | E_i | \psi \rangle \geq 0$  for any state  $|\psi\rangle$  in the Hilbert space, from which it follows that it is Hermitian and all its eigenvalues are non-negative. This property ensures that  $\text{Tr}(E_i \rho_i) \geq 0$  can be consistently interpreted as a valid probability associated with each possible measurement outcome. The second condition is *completeness*, which asserts that the operators provide a resolution of the identity; that is, their sum must encompass the entire probability space represented by the identity operator.

A POVM is a generalization of a projection-valued measure (PVM), necessary to describe the effect on a subsystem of a projective measurement performed on a larger system. Providing a POVM demonstrates

that such a measurement can be represented, but its physical realization must be constructed with suitable PVMs acting on the entire system, as importantly demonstrated by Naimark's dilation theorem [180].

What is the optimal POVM for minimum-error discrimination of these  $\{\rho_i\}_{i=1}^N$  states? Before we address this question directly, let's consider the *square-root* measurement, also called *pretty good* measurement [181, 182, 179]. The POVM it specifies consists of the following Hermitian operators

$$E_i = \frac{1}{6} \rho^{-1/2} \rho_i \rho^{-1/2}, \quad (2.61)$$

where

$$\rho = \frac{1}{6} \sum_{i=1}^6 \rho_i \quad (2.62)$$

is the so-called *a priori* density operator. Here,  $\rho^{-1/2}$  denotes the inverse of the positive square root of  $\rho$ . In the present case, the states  $\{\rho_i\}$  collectively span the two-qubit Hilbert space, so that  $\rho$  has full rank and strictly positive eigenvalues. As a result,  $\rho$  is positive definite and admits a well-defined inverse square root<sup>29</sup>  $\rho^{-1/2}$ . It can be checked that these POVM elements are positive semidefinite and sum to the identity, i.e. they satisfy Eq. (2.59) and Eq. (2.60). By computing the probability that Alice successfully identifies the received states using this POVM, we obtain

$$P_{\text{succ},2} = \sum_{i=1}^6 \frac{1}{6} \text{Tr}(E_i \rho_i) = \frac{1}{240} \left( 59 + 10\sqrt{2} + 2\sqrt{5} (6 + 7\sqrt{2}) \right) \approx 0.601028 \dots \quad (2.63)$$

This probability is indeed greater than the classical result found in § 2.4.2 by over 10%, demonstrating that quantum advantage in two questions remains achievable within this generalized oracle framework. Furthermore, contrary to the superdense-coding-based strategy, this is a simpler procedure that does not require entangling the qubits over large distances in space.

Eq. 2.63 establishes the possibility of a solution. To physically implement this POVM, however, one must determine how auxiliary qubits should be entangled with the system and construct the precise operations to apply to them. This procedure involves finding an isometry  $V$  from 2 to 5 qubits, i.e., a complex matrix of dimension  $2^5 \times 2^2$  such that  $V^\dagger \cdot V = \mathbb{1}_{2^2 \times 2^2}$ , implementing the POVM by using 3 auxiliary qubits to accommodate the 6 necessary measurement outcomes, which are mapped onto the set  $\{|000\rangle, |001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle\}$ , with  $\text{Tr}(E_1 \rho_1) = \text{Tr}(|000\rangle\langle 000| \otimes \mathbb{1}_{4 \times 4} V \rho_1 V^\dagger)$ ,  $\text{Tr}(E_2 \rho_2) = \text{Tr}(|001\rangle\langle 001| \otimes \mathbb{1}_{4 \times 4} V \rho_2 V^\dagger)$ ,  $\text{Tr}(E_3 \rho_3) = \text{Tr}(|010\rangle\langle 010| \otimes \mathbb{1}_{4 \times 4} V \rho_3 V^\dagger)$ , etc., for all density operators  $\{\rho_i\}$ . After applying the isometry, the measurement probabilities of the auxiliary

<sup>29</sup>Explicitly, if  $\rho = \sum_k \lambda_k |v_k\rangle\langle v_k|$  is the spectral decomposition of  $\rho$  with  $\lambda_k > 0$ , then  $\rho^{-1/2} = \sum_k \lambda_k^{-1/2} |v_k\rangle\langle v_k|$ .




Figure 2.7: **Quantum circuit decomposition for the square-root measurement.** The circuit is decomposed in the  $\{\text{CNOT}, R_z(\theta), R_x(\theta)\}$  gate basis, where  $R_z(\theta) = e^{-i\theta Z/2}$  (green) and  $R_x(\theta) = e^{-i\theta X/2}$  (red). Qubits are labeled from 0 to 4, ordered from bottom to top. Qubits  $q_1$  and  $q_0$  correspond to the first and second qubits of the states  $|\phi_i\rangle$ , which are associated with gods located at positions A and B, respectively. The success probability described in Eq. (2.63) is determined by measurements on qubits  $q_4$ ,  $q_3$ , and  $q_2$ . The basis states  $\{|i\rangle_3\}_{i=1}^6$  in the  $|q_4 q_3 q_2\rangle$  tensor product order map to the detection outcomes for the set of states  $\{|\phi_i\rangle\}_{i=1}^6$  in the same order. The auxiliary qubit states  $|110\rangle$  and  $|111\rangle$  never occur. The circuit comprises 103 CNOT gates and a total depth of 200 layers. The exact circuit instructions are provided in Appendix 2.11.2 in OpenQASM format to ensure reproducibility.

qubits in the computational basis for every state will match those of measuring the original state using the POVM. Finally, having obtained  $V$ , it can be translated into an explicit quantum circuit via unitary decomposition [183, 184, 185], which, in turn, can be further optimized with the method developed in § 4. The resulting circuit we obtained comprises 103 CNOT gates and 200 layers of depth, as displayed in Fig. 2.7.

It is important to note, however, that this is not yet the optimal POVM. In other words, Alice's probability of success can still be improved through a more effective POVM design. This can be confirmed by evaluating Helstrom and Holevo optimality conditions<sup>30</sup>, the first of which is

$$\sum_{i=1}^6 \frac{1}{6} \rho_i E_i - \frac{\rho_j}{6} \succeq 0, \quad \forall j \quad (2.64)$$

and is not satisfied for all  $j$ , and the second of which is

$$E_j \left( \frac{1}{6} \rho_j - \frac{1}{6} \rho_k \right) E_k = 0, \quad \forall j, k \quad (2.65)$$

and is also not satisfied for all pairs of  $j$  and  $k$ .

Therefore, we now return to the question we posed earlier to address how Alice can maximize her success probability in minimum-error discrimination. To determine the optimal POVM, a numerical approach was employed. Using a semidefinite program (SDP) formulation, it is possible to define the objective function in Eq. (2.58) and maximize the success probability for general positive semidefinite operators  $E_i$ ,

<sup>30</sup>Satisfying these conditions guarantees minimization of the error. For a proof, see Ref. [179].

subject to the constraints in Eqs. (2.59) and (2.60). This computation was carried out in Python with the CVXPY library. As detailed in Appendix 2.11.1, the resulting probability of success is only marginally increased further to

$$P_{\text{succ},2}^{\text{ME}} \approx 0.601271, \quad (2.66)$$

which implies that the square-root measurement is already close to optimal.

Finally, we address the fully generalized scenario in which multiple oracles can be applied to implement the same classical logic. The rationale is as follows: since the logic of the HLPE is classical, it could be argued that  $\mathfrak{Random}$  retains some freedom to apply to the communication qubit any of the compatible quantum oracles presented in Eq. (2.46). In the absence of further knowledge regarding this choice, Alice might reasonably assume that  $\mathfrak{Random}$  selects one of the oracles at random, with equal probability. This generalization increases the number of possible combinations of oracles that produce the output two-qubit state from 6, as presented in Eqs. (2.52 – 2.57), to 10, since for any permutation of A and B where  $\mathfrak{Random}$  appears, the number of oracle combinations increases to 2. These combinations can be grouped into six permutation classes, categorized according to the gods that originated them.

An optimal POVM for minimum-error discrimination in this scenario can readily be obtained using a different SDP. In this case, the objective function is modified to maximize

$$P_{\text{succ}} = \sum_{i=1}^6 \sum_{j \in \{j\}_i} \frac{1}{6n_i} \text{Tr}(E_i \rho_{i,j}), \quad (2.67)$$

where  $j$  indexes the oracle combinations within each of the six permutations of the gods' identities,  $n_i = |\{j\}_i|$  is the number of oracle combinations per permutation group,  $\rho_{i,j} = |\phi_{i,j}\rangle \langle \phi_{i,j}|$ , and  $\{E_i\}$  are the POVM elements to be optimized, satisfying Eqs (2.59) and (2.60). As shown, even though there is a higher diversity of possible states that Alice can receive, we maintain only six POVM elements. These elements are designed to detect multiple possible output states, accounting for degeneracies among different oracles associated with the same gods. However, the success probability achieved does not exceed 50% in this epistemically relaxed scenario, a result entirely equivalent to the classical probability of solving the HLPE in two questions with the minimum-error discrimination strategy that culminated in Eq. (2.16).

To try to improve upon this result, we conducted additional theoretical explorations. In an external optimization cycle, we incorporated two additional features in the search for higher  $P_{\text{succ}}$  besides the POVM measurement optimization via SDP: *i*) the preparation of a general two-qubit initial state, wherein the two qubits sent to two gods are entangled rather than being in the simple product state considered earlier, and *ii*) the inclusion of a fixed general two-qubit operation applied to the joint state after the

oracle applied to the first qubit and before the oracle applied to the second one. The 6 real parameters defining the initial state and the 15 real parameters of the general unitary operation were optimized in an external loop using the `basinhopping` global optimization method from `scipy`'s library, with the POVM measurement on the returned state also being optimized for each choice of these 21 additional parameters along the search. We tested this procedure for both approaches described earlier: (1) with a single oracle fixed per god and (2) with both redundant oracles available for `Random`. The results were identical in both cases, yielding  $P_{\text{succ}} \approx 0.601271$  and  $P_{\text{succ}} \approx 0.5$ , respectively.

If a quantum-advantageous solution to this fully generalized version of the HLPE including all the redundant oracles exists, it appears to require a more sophisticated quantum algorithm. Such an algorithm must operate without prior knowledge of the exact version of the oracle employed by each god, discriminate the unknown permutation in a single shot using only two oracle queries, and function without access to the oracle or its controlled versions.

## 2.8 Communication subject to noise

Even theoretical frameworks grounded in physical principles often require refinement to address the complexities of practical implementation. While conceptually foundational and capturing the essential phenomena, initial simplifications may overlook critical details necessary for practical realization, such as unanticipated environmental interactions, imperfections in control, and unavoidable sources of noise.

To develop intuition for the types of imperfections that arise in realistic implementations, it is instructive to consider concrete physical transmission scenarios before introducing an abstract noise model. Free-space optical communication provides a particularly transparent setting: it interpolates between idealized vacuum propagation and realistic noisy channels, while minimizing material interactions that may obscure the fundamental mechanisms responsible for loss and decoherence. Consider again the example of a photon propagating through space. While over sufficiently short distances its transmission path may be effectively treated as vacuum, practical communication systems typically employ beams comprising many identically prepared photons in order to mitigate losses and improve the signal-to-noise ratio.

The most significant milestones in space-based quantum communication to date have been achieved by China's Micius satellite, launched in 2016. This mission demonstrated two-photon entanglement distribution and the violation of the CHSH inequality between two ground stations separated by 1200 km [186], as well as satellite-to-ground quantum key distribution (QKD) [187] and satellite-mediated QKD between terrestrial locations separated by 7600 km [188]. Orbiting at an altitude of approximately 500 km, *Mi-*

*cius* relied on the polarization degree of freedom of photons to implement these protocols, similar to the framework discussed in § 2.7.

As communication distances increase beyond Earth orbit into interplanetary regimes, additional sources of degradation become increasingly relevant. Pointing errors, beam divergence, and diffraction must be carefully controlled in order to suppress channel loss [189]. At even larger, interstellar distances, photons may encounter further challenges, including absorption and scattering by interstellar gas and dust [190], magneto-optic effects in regions of high plasma density, interactions with cosmic radiation, and gravitational influences [191, 192, 193, 194, 195]. Such interactions can modify a photon's energy, polarization, or temporal profile, entangle it with external degrees of freedom, or destroy its quantum state entirely. While these effects are extreme, they exemplify the same fundamental process of decoherence and information loss that requires careful strategies to avoid losing the quantum advantage of the protocol.

To date, no demonstrations of quantum communication have been performed in deep space. Nevertheless, classical deep-space communication provides a useful point of comparison. Transmitting a single classical bit to a distant spacecraft requires sophisticated error correction and signal processing techniques to compensate for extreme attenuation. A striking example is provided by Voyager 1<sup>31</sup>, which continues to transmit scientific data to Earth from interstellar space [199, 200, 201].

More generally, and independently of implementation, whether the chosen generalization approach in § 2.7 involves fixing or relaxing the computational basis, control errors and noise are invariably present in both the transmission system itself and any raw communication channel<sup>32</sup>. The result is gradual decoherence and loss of information, progressively driving quantum states toward classicality [202].

Such practical considerations must therefore be incorporated when analyzing the feasibility of communication in the HLPE. Introducing a full-fledged error-correction scheme would require additional agreements between Alice and the gods regarding protocol structure, thereby deviating from the minimal-assumption philosophy adopted thus far. A simpler and more natural approach to communicate with minimal effort is to employ a repetition strategy, wherein both parties implicitly recognize that, due to noise, a single qubit cannot reliably convey information and instead transmit a collection of  $N$  identically prepared qubits. The central task then becomes determining a sufficiently large value of  $N$  to guarantee statistical distinguishability of the received states, given an estimated single-qubit error probability.

---

<sup>31</sup>Launched in 1977, Voyager 1 and Voyager 2 crossed the heliopause and entered the interstellar medium in 2012 and 2019, respectively [196]. They remain the only probes from Earth confirmed to be operating in interstellar space. Informal estimates suggest that  $\sim 10^{22}$  photons are emitted per bit by Voyager 1, with only a few hundred photons per bit being detected at Earth, assuming realistic conditions [197]. Advances toward one-photon-per-bit classical communication have also been proposed [198].

<sup>32</sup>By “raw,” we refer to a communication channel that transmits information without any additional abstract layer of encoding or error correction to protect the transmitted state.

As will be shown, this repetition-based strategy — motivated purely by the necessity of overcoming noise in the simplest possible manner — not only restores reliable communication but also enables a two-question solution to the HLPE with success probability arbitrarily close to 1 in  $N$ . This result holds within both generalized frameworks introduced in § 2.7 and, in the relaxed computational basis case, applies to the fully general scenario encompassing both oracles in Eq. 2.46.

Before detailing this repetition-based scheme explicitly, it is instructive to analyze how noise degrades the performance of the previously presented protocols. To this end, we will consider classical and quantum channels with an equivalent symmetric bit-flip error rate, which will also serve as the basis for later studying the repetition-based strategy.<sup>33</sup>

### 2.8.1 Protocol performance under equivalent bit-flip errors

To explore the impact of noise on the solution protocols discussed thus far, we begin with a classical bit-flip error model. Transmitted through this simple channel, a bit  $x$  has an equal probability  $\varepsilon < 1/2$  of being flipped to  $\bar{x} = x \oplus 1$ , regardless of whether its original state is 0 or 1. Using this error model, we can numerically simulate the operation of Boolos' three-question solution protocol and the classical two-question minimum-error discrimination solution protocol described in § 2.4.2 and Fig. 2.3.

We begin by analyzing the scenario in which errors occur exclusively during transmission in one direction, from the gods to Alice, following the gods' application of their respective encoding operations on the received two-state system. Later, we will explain how this analysis can be easily extended to account for errors occurring in both transmission directions.

In the case of Boolos' protocol, one might initially expect the success probability to be  $(1 - \varepsilon)^3$ , since each bit received per question can be corrupted with probability  $\varepsilon$ , thereby potentially compromising all subsequent questions and the final guess. However, because 3 bits can encode  $2^3 = 8$  states, and there are only 6 possible permutations of interest, certain bitstrings are resilient to error, resulting in a slightly higher success probability. For instance, with  $\varepsilon = 20\%$ , the simulation yields  $P_{\text{succ}} = 55.47 \pm 0.03\%$ , whereas the formula yields only 51.2%. The performance of Boolos' solution under different bit-flip error probabilities ranging from 0 to  $1/2$ , is shown in Fig. 2.8. For the probabilistic two-question protocol, the results of the simulation under the same bit-flip error channel are also presented in Fig. 2.8.

In the quantum case, a qubit state can experience not only a bit-flip error but also a phase-flip error, or a combination of both. These errors are represented by the Pauli matrices  $X$ ,  $Z$ , and  $Y$ , respectively.

---

<sup>33</sup>While the conceptual difference enabling quantum advantage in the HLPE is independent of the specific noise model, adopting a concrete error model is necessary for quantitative comparison.

To compare the distinct consequences of classical and quantum errors in the context of the HLPE, we consider a simple depolarizing channel where each of these errors occurs with equal probability<sup>34</sup>.

The single-qubit depolarizing channel,  $\mathcal{D}_\varepsilon$ , introduces errors by applying the Pauli operators  $X$ ,  $Y$ , and  $Z$  with equal probability  $\varepsilon/2$ , while leaving the quantum state unchanged with a probability of  $1 - 3\varepsilon/2$  [203]. It is represented using the Kraus operators  $\{K_i\}$ , given by

$$K_0 = \sqrt{1 - 3\varepsilon/2} \mathbb{1}, \quad (2.68)$$

$$K_1 = \sqrt{\varepsilon/2} X, \quad (2.69)$$

$$K_2 = \sqrt{\varepsilon/2} Y, \quad (2.70)$$

$$K_3 = \sqrt{\varepsilon/2} Z, \quad (2.71)$$

where  $\mathbb{1}$  is the identity operator. When this channel acts on an initial single-qubit density matrix state  $\rho$ , the resulting state is given by:

$$\mathcal{D}_\varepsilon(\rho) = \sum_{i=0}^3 K_i \rho K_i^\dagger = \left(1 - \frac{3\varepsilon}{2}\right) \rho + \frac{\varepsilon}{2} (X\rho X + Y\rho Y + Z\rho Z). \quad (2.72)$$

It is straightforward to verify that this channel, acting on the  $|\psi^{\text{id}}\rangle$  or  $|\psi^{\text{bf}}\rangle$  states, Eq. (2.49) and (2.50) respectively, will output these input states with probability  $1 - \varepsilon$  and the orthogonal states with probability  $\varepsilon$ . This reflects a bit-flip probability equal to the classical communication case, facilitating direct comparison. Simultaneously, the channel will leave the measurement statistics of the random states  $|\psi_\pm^{\text{rd}}\rangle$ , in Eq. (2.51), unaltered, meaning that these states will remain completely random.

To simulate the effect of the depolarizing channel on the communication qubit in the case of the superdense-coding-based protocol, we apply the channel to this qubit only, since the second one always remains in Alice's possession. Mathematically, the channel is applied on the two-qubit state  $\rho$  produced after the god's encoding operation, and before Alice's Bell-basis measurement (see Fig. 2.4), and constructed by tensoring each Kraus operator in Eq. (2.72) with the identity, i.e. replacing  $K_i$  with  $K_i \otimes \mathbb{1}$ .

Since the superdense-coding-based protocol operates sequentially by identifying one god at a time, a bit-flip error on the first qubit can lead to misidentification of the gods addressed in first or second place. For example, if both measurements yield the classical outcome of  $|00\rangle$ , Alice can deduce that one god must be **True** and the other **Random**. This is because a bit-flip on the first qubit cannot transform **False**'s encoded state  $|01\rangle$  into  $|00\rangle$ , while **Random**'s potential state  $|10\rangle$  can be corrupted to  $|00\rangle$ .

<sup>34</sup>A more general model could assign different probabilities to bit-flip, phase-flip, and combined errors.




Figure 2.8: **Success probability for four HLPE solution protocols simulated under a noise channel with bit-flip error probability  $\varepsilon$  on the transmission path from the gods to Alice.** The results for classical protocols are represented by blue markers, while quantum results are shown in red. The performance of the exact three-question solution protocol from Boolos (■) is compared with the exact two-question quantum solution protocol based on superdense coding (▲) presented in § 2.6. The latter consistently outperforms the former across all values of the bit-flip error probability parameter  $\varepsilon$ , ranging from 0 to  $1/2$ . For the generalized oracles implementing the two-question quantum solution protocol, with success probability given by Eq. (2.66) in the absence of noise (▼), we observe the success probability remains superior to the classical two-question solution probability achieved by the probabilistic protocol shown in Fig. 2.3 (●). In all four cases, when  $\varepsilon = 1/2$ , the success probability decays to  $1/6$ , corresponding to the probability of randomly guessing the correct permutation.

To mitigate noise effects, Alice supplements the previously derived protocol (see § 2.6) with the following rules: (1) if  $|00\rangle$  or  $|01\rangle$  is measured a second time, the second god is identified as *Random*; (2) if  $|10\rangle$  or  $|11\rangle$  is measured a second time, the second god is identified as *True* or *False*, respectively.

As Fig. 2.8 clearly shows, the probability of success for the superdense-coding-based protocol outperforms that of Boolos' protocol across all values of  $\varepsilon$ . Consequently, in addition to the distinctive advantages of the quantum solution strategy highlighted previously—the most remarkable being its ability to solve the puzzle with just two questions—our protocol also importantly demonstrates greater resilience to bit-flip errors compared to the classical one.

Lastly, we simulate the effect of depolarizing noise on the solution protocol that employs the generalized oracles  $\mathcal{O}'_{\text{id}}$ ,  $\mathcal{O}'_{\text{bf}}$ , and  $\mathcal{O}'_{\text{rd},+}$  presented in § 2.7.2. To obtain the resulting two-qubit density matrix  $\rho'$  received by Alice, the separable two-qubit state  $\rho$ , prepared after the two sampled gods apply their respective encoding operations on the communication qubits, is passed through a two-qubit depolarizing channel. This channel is constructed by independently applying the single-qubit depolarizing channel to

each qubit. The Kraus operators for the two-qubit channel are obtained by taking the tensor product of the Kraus operators for each qubit,

$$K_{ij} = K_i \otimes K_j, \quad (2.73)$$

where  $i, j \in \{0, 1, 2, 3\}$ . These 16 operators are then used to calculate the output density matrix following

$$\rho' = \sum_{i=0}^3 \sum_{j=0}^3 K_{ij} \rho K_{ij}^\dagger. \quad (2.74)$$

For the final mixed state  $\rho'$  received by Alice, the optimal POVM is determined again using an SPD formulation. The resulting  $P_{\text{succ}}(\varepsilon)$  is also plotted in Fig. 2.8. As shown, this probability consistently outperforms that of the classical two-question solution protocol, starting from  $\varepsilon = 0$ , where the quantum advantage was already demonstrated in § 2.7.2, and extending up to the maximally random error scenario  $\varepsilon = 1/2$ , where all protocols converge to the random-guess success probability of  $1/6$ .

To conclude, we revisit the earlier observation that this analysis accounts solely for errors occurring during transmission from the gods to Alice. Using straightforward algebra and the Pauli commutation relations, it can be checked that all the oracles in Eqs. (2.44 – 2.46) commute with  $\mathcal{D}_\varepsilon$ . Specifically,  $U \mathcal{D}_\varepsilon(\rho) U^\dagger = \mathcal{D}_\varepsilon(U \rho U^\dagger)$ , where  $U$  represents any of the oracles and  $\rho$  denotes an arbitrary single-qubit state. As a result, the effect of errors occurring only during transmission from Alice to the gods is identical to that of the same error process occurring only during transmission from the gods to Alice.

Moreover, in the more realistic scenario where errors occur in both transmission directions, the total effect can be determined by applying the depolarizing channel defined in Eq. (2.72) consecutively. It is well known that the composition of a depolarizing channel  $\mathcal{D}_\varepsilon$  with itself results in another depolarizing channel with an updated depolarizing probability. According to Eq. (2.72), parameterized to facilitate comparison with the classical bit-flip probability  $\varepsilon$ , this composition yields  $\mathcal{D}_\varepsilon \circ \mathcal{D}_\varepsilon = \mathcal{D}_{\varepsilon'}$ , where  $\varepsilon' = 2(\varepsilon - \varepsilon^2)$ <sup>35</sup>.

Therefore, for theoretical analysis of the output mixed state when the communication qubit is subjected to the depolarizing channel  $\mathcal{D}_\varepsilon$  both before and after the god's encoding operations, we can equivalently consider that the qubit arrives at the god without error. After he applies his encoding operation, the resulting state is transmitted back through a depolarizing channel  $\mathcal{D}_{\varepsilon'}$ . This yields a qualitatively similar effect to the analysis previously performed (see Fig. 2.8), only with a faster decay in success probability.

<sup>35</sup>Note that the equivalence of this combined channel with the classical case is straightforward to derive. When a bit undergoes the same error process twice, each with a bit-flip probability  $\varepsilon$ , the bit is flipped if the first process leaves it unchanged and the second flips it, or vice-versa. Thus, the total bit-flip error probability is  $\varepsilon' = (1 - \varepsilon)\varepsilon + \varepsilon(1 - \varepsilon) = 2(\varepsilon - \varepsilon^2)$ .

## 2.8.2 Classical communication with a repetition code

As shown in Fig. 2.8, for an arbitrary number of game repetitions  $N$ , the average success probability of all protocols, whether classical or quantum, decreases with increasing  $\varepsilon$ . However, a more efficient method is possible for transmitting the necessary information over a noisy channel to solve an instance of the HLPE with reduced communication overhead. Here, we consider the classical case.

Specifically, the players can encode the answer to each single question into  $N$  bits, rather than repeating  $N$  instances of the game and averaging the success probabilities. This technique, consisting of a simple repetition code, is easy to justify and implement without requiring complex assumptions.<sup>36</sup> As demonstrated in this section, a repetition code allows the game to recover the noiseless success probability exponentially fast with increasing  $N$ .

Here is the setup: Alice transmits  $N$  copies of the information carrier prepared in her own basis, which are operated on equally by the god, one by one, before being transmitted back and received by Alice. She then decodes the received bitstring  $\mathbf{x}$  using a simple *threshold-based majority rule* to retrieve the bit that the god intended to encode.

Let  $M$  be the number of received bits that match the sent bit  $x$ . Assuming the samples are independent and identically distributed (i.i.d.),  $M \sim \text{Binomial}(N, 1 - \varepsilon)$  follows a binomial distribution with parameters  $N \in \mathbb{N}$ ,  $\varepsilon \in [0, 1]$ , and probability mass function given by

$$\Pr(M; N, \varepsilon) = \binom{N}{M} (1 - \varepsilon)^M \varepsilon^{N-M}, \quad (2.75)$$

where  $\binom{N}{M}$  is the binomial coefficient and  $1 - \varepsilon$  is the probability of successful transmission.

To correctly infer  $x$ , a threshold  $T > N/2$  of the  $N$  bits, to be decided, must match  $x$ . The *reliability* of the inference is defined as the probability that the number of bits communicated without error is greater

---

<sup>36</sup>While mutual assumptions remain necessary, as in the game repetition protocol, they can plausibly emerge in communication scenarios where both parties possess limited knowledge of each other, which we have been aiming to motivate. As previously noted, similar repetition protocols are frequently used in quantum communication experiments, such as those conducted with the Micius satellite, where millions of polarization-entangled photon pairs are emitted per second with a certain fidelity, only a fraction of which are detected by the receiver.

than or equal to  $T$ . In other words, it is the probability of correctly inferring  $x$  when  $x$  was actually sent:

$$R(T; N, \epsilon) \equiv \Pr(M \geq T) \quad (2.76)$$

$$= \sum_{M=T}^N \binom{N}{M} (1 - \epsilon)^M \epsilon^{N-M} \quad (2.77)$$

$$= (1 - \epsilon)^T \sum_{M=T}^N \binom{M-1}{T-1} \epsilon^{M-T}. \quad (2.78)$$

It is apparent that the reliability increases with  $N$  and decreases with  $\epsilon$  or  $T$ , when each of these parameters is varied independently while all others are held fixed. A detailed derivation of the final step in this expression is provided in Ref. [204]. For practical values of these variables, obtaining this probability requires numerical computation. For instance, with  $N = 10$  physical bits per encoded bit,  $\epsilon = 0.1$  for a bit-flip error, and requiring  $T = 8$  for bit decoding, the message will be successfully transmitted with a probability  $R(8; 10, 0.1) \approx 93.0\%$ .

However, there is a small but non-zero probability that only  $N - T$  bits or less will match  $x$  after transmission, causing the received bitstring to appear as if the encoded bit is  $\bar{x}$  when decoded using the threshold-based majority voting rule. This event—decoding  $\bar{x}$  when the encoded bit was  $x$ —is classified as an error, occurring with probability

$$P_{\text{err}} = \Pr(M \leq N - T). \quad (2.79)$$

In addition, there is a third possible outcome. If the maximum number of bits in the received bitstring that are equal to each other, denoted as  $M_{\text{maj}}$ <sup>37</sup>, does not meet the threshold  $T$ , the result is classified as *ambiguous* and the bitstring is rejected. The probability of this outcome is given by

$$P_{\text{amb}} = \Pr(N - T < M < T), \quad (2.80)$$

such that we have  $R(T; N, \epsilon) + P_{\text{err}} + P_{\text{amb}} = 1$ .

Continuing with the example considered previously, where  $N = 10$ ,  $T = 8$ , and  $\epsilon = 0.1$ , we compute  $P_{\text{err}} \approx 0.00004\%$  and  $P_{\text{amb}} \approx 7\%$ . Note that the error probability is made very small due to the choice of  $T$ ; lowering the threshold would increase the reliability  $R$  but also lead to a higher error rate. Overall, with this choice of  $T$ , Alice would face a combined probability of  $P_{\text{err}} + P_{\text{amb}} \approx 7\%$  of failing to decode the message correctly.

---

<sup>37</sup>Note that  $\lceil N/2 \rceil \leq M_{\text{maj}} \leq N$ .

In practice, both parties might agree on the values of  $N$  and  $T$  for a given  $\epsilon$ , which is estimated by previously studying the channel, in order to achieve the desired confidence margin that the decoded bit value corresponds to the intended one, while minimizing errors. Alternatively, Alice may unilaterally choose the value of  $N$ , with the gods implicitly understanding that they should act identically on all carriers.

To summarize the protocol, *i*) Alice and the god begin by deciding the target reliability  $R$  for communication and, based on a previously estimated bit-flip error probability  $\epsilon$  for the channel, they calculate the required number of  $N$  qubits and the threshold  $T$  to be used; *ii*) Alice then counts the number of 0's and 1's in the received  $N$ -bit string  $\mathbf{x}$ , which encodes the bit  $x$ ; *iii*) if the most frequent bit in  $\mathbf{x}$  occurs  $T$  or more times, she identifies  $x$  with that value; and *iv*) if not, she declares the message ambiguous. Finally, if ambiguity is not an option, they just set  $T = \lfloor N/2 \rfloor + 1$ .

This setup is sufficient to communicate the messages in the classical scenario of the HLPE. `True`, `False` and `Random` will always communicate a deterministic system state, which will either be bit-flipped or not relative to the state Alice sent them. The message can then be decoded by employing a  $T$ -out-of- $N$  threshold rule.

### 2.8.3 Quantum communication with a repetition code

Now we analyze the equivalent implementation of a repetition code in the generalized quantum communication-based scenario of the HLPE. Maintaining the parallel between classical and quantum implementations is essential for our goal of comparison. In both cases, the key requirement for handling noise with a repetition code is that all  $N$  transmitted bits, or qubits, are prepared in the same physical state. The idea is that if a fraction of them is corrupted, majority voting can still recover the message. This observation is trivial when the transmitted system is read out deterministically, as occurs when `True` or `False` reply. However, a difference arises with `Random`'s nature.

Classically, `Random` must decide uniformly at random whether to apply a bit-flip operation to the received coins, maintaining the same decision for all  $N$  coins. He must do this instead of deciding the operation uniformly at random for each coin to conceal his random identity. Quantum mechanically, the protocol is equivalent, with `Random` applying the same operation to all received qubits. However, when Alice measures the  $N$  qubits upon receiving them back, she retrieves a random bit from each quantum state. This allows her to immediately identify `Random`. In the classical case, she can only recover a deterministic bit — subject to statistical confidence correction — which appears indistinguishable from the bits received from `True` or `False` and does not reveal the sender's identity.

The fundamental difference in the quantum case is that randomness is recovered at readout, unlike

in classical mechanics. This enables a two-question solution to the HLPE with a success probability exponentially close to 1 with increasing  $N$  in the presence of noise mitigated through a repetition code. This result holds within both generalized frameworks introduced in § 2.7 and, in the relaxed computational basis case, applies to the fully general scenario encompassing both oracles in Eq. 2.46.

To quantify this, we next analyze how the number of copies of the transmitted states,  $N$ , and the threshold value,  $T$ , impact the statistical confidence in distinguishing the transmitted state. Since one of the states is completely random, this separation is slightly less resource-efficient compared to the classical scenario. However, it can still be accomplished.

### Decoupling three different messages

Consider again the case where the bit-flip probability satisfies  $\varepsilon < 1/2$ , as derived from the quantum mechanical properties of the depolarizing channel discussed earlier. Suppose Alice receives an arbitrary bitstring  $\mathbf{x}$ , encoding an unknown bit. To analyze it, we again let  $M$  denote the number of received bits, out of a total of  $N$ , that match the sent bit. The probability mass function of  $M$ , conditioned on the encoded bit being either  $x$  or  $\bar{x}$ , was previously established in Eq. (2.75).

Now, consider the additional case where a fully random bit  $u$  might be encoded into the bitstring. Here,  $u$  is drawn from the uniform distribution  $u \sim \text{Unif}\{x, \bar{x}\}$ , representing an equal probability of encoding  $x$  or  $\bar{x}$  at the readout stage of each sent qubit. Although, in this case, no deterministic bit is encoded in the bitstring, the randomness can still be modeled by assuming that  $M \sim \text{Binomial}(N, 1/2)$ , corresponding to a symmetric binomial distribution, with probability mass function given by

$$\Pr'(M; N, 1/2) = \binom{N}{M} \frac{1}{2^N}, \quad (2.81)$$

where the factor  $1/2^N$  reflects the equal probability of each possible configuration of  $N$  bits.

We aim to determine a suitable value for the threshold  $T$  that is sufficiently far from  $N/2$  to confidently distinguish not only between the two complementary deterministic bits but also from the completely random one. Just as before, the protocol operates by counting  $M_{\text{maj}}$ , the number of occurrences of the most frequent readout value. The encoded bit is classified as 0 (conversely, 1) if the most frequent readout is 0 (conversely, 1) and  $M_{\text{maj}} \geq T$ , or as  $u$  if  $M_{\text{maj}} < T$ .

To compute the *reliability*—i.e., the average probability of successful transmission—we sum the conditional reliabilities for detecting each type of bit, weighted their respective prior probabilities of occurrence.

This expression can be written as:

$$R(T; N, \varepsilon, p_u) = p_x R(T; N, \varepsilon | x) + p_{\bar{x}} R(T; N, \varepsilon | \bar{x}) + p_u R(T; N, \varepsilon | u) \quad (2.82)$$

$$= (1 - p_u) \Pr(M \geq T) + p_u \Pr'(N - T < M < T) \quad (2.83)$$

$$= (1 - p_u)(1 - \varepsilon)^T \sum_{M=T}^N \binom{M-1}{T-1} \varepsilon^{M-T} + p_u \sum_{M=N-T+1}^{T-1} \binom{N}{M} \frac{1}{2^N}, \quad (2.84)$$

where the expression is simplified using  $p_x + p_{\bar{x}} + p_u = 1$ . As evident, in the limit  $p_u \rightarrow 0$ , the reliability  $R(T; N, \varepsilon, p_u)$  converges to the value found in the classical communication scenario. In that case, the bitstring encodes only one deterministic bit, recovering the expression given in Eq. (2.78).

There is also a possibility that, for a given type of bit sent, the sampled value  $M$  does not fall within the desired bounds defined by the chosen threshold  $T$ . The overall probability of error is given by

$$\begin{aligned} P_{\text{err}} &= (p_x + p_{\bar{x}}) \Pr(M < T) + p_u 2 \Pr'(M \geq T) \\ &= (1 - p_u) \left[ \sum_{M=0}^{T-1} \binom{N}{M} (1 - \varepsilon)^M \varepsilon^{N-M} \right] + 2 p_u \sum_{M=T}^N \binom{N}{M} \frac{1}{2^N}. \end{aligned} \quad (2.85)$$

It is important to note that ambiguity errors, as computed in the classical communication case, do not arise in this framework. Unlike in the classical scenario, when the threshold  $T$  is not met, Alice always classifies the received bitstring as a random bit. As a result, in the limit  $p_u \rightarrow 0$ , the probability of error  $P_{\text{err}}$  reduces to the sum of the errors given in Eqs. (2.79) and (2.80). Furthermore, it is evident that with appropriate choices of  $N$  and  $T$ ,  $P_{\text{err}}$  can be made to vanish asymptotically. However, for fixed  $N$ ,  $p_u$ , and  $\varepsilon$ , it is also possible to compute the optimal value of  $T$  that minimizes  $P_{\text{err}}$ .

Let us now illustrate this scenario through a few specific examples. First, we extend the example provided in the previous section: consider  $N = 10$  physical bits per encoded bit, with a bit-flip error probability of  $\varepsilon = 0.1$  for each individual bit. For decoding, we again set the threshold  $T = 8$ , as was done in the classical communication scenario. This threshold is, in fact, optimal in the quantum communication scenario if we assume a prior probability  $p_u = \frac{1}{3}$ . The resulting distributions for all three types of bits are plotted in Fig. 2.9a. Under these parameters, the reliability is computed as  $R(8; 10, 0.1, 1/3) \approx 91.7\%$ , while the probability of error is  $P_{\text{err}} \approx 8.3\% = 1 - R$ .

This contrasts with the results from the equivalent classical example discussed in § 2.8.2, where the reliability and the combined error and ambiguity probability were found to be approximately 93.0% and 7.0%, respectively. However, it is important to note that in this quantum communication scenario, Alice can directly determine the identity of the god with whom she is communicating in a single question with the




Figure 2.9: **Quantum communication of three bit types via a repetition code.** **a**, Binomial distributions for  $M_x$ , the number of received bits in state  $x$ , conditioned on the sent bit being  $x$  (blue),  $\bar{x}$  (red), or a random bit  $u$  (dark grey), for a bitstring of length  $N = 10$ . Notice the overlapping regions of the different distributions. **b**, The same distributions for  $N = 30$ , showing an improved distinguishability, which is confirmed by selecting the optimal threshold  $T_{opt} = 22$  as described in the main text. **c**, Scaling behavior of the minimum error probability  $P_{err}^{min}$  as a function of  $N$ , highlighting exponential decay for various bit-flip error probabilities  $\epsilon$  with fixed prior  $p_u = 1/3$ .

given reliability. This direct identification was not achievable in the classical communication framework.

To increase the reliability, the number of repetitions can be increased. For instance, at  $N = 30$ , the optimal threshold value is  $T_{opt} = 22$  and the reliability is computed as  $R(22, 30; 0.1, \frac{1}{3}) \approx 99.3\%$ , with the error probability reduces to  $P_{err} \approx 0.7\%$  (see Fig. 2.9b).

Due to the central limit theorem, the binomial distribution converges to a normal distribution in the large- $N$  limit. In this regime, the error probability becomes negligibly small, and the distributions become nearly perfectly distinguishable, highlighting the effectiveness of increasing the number of repetitions for improving communication reliability. Figure 2.9(c) illustrates the scaling behavior of  $P_{err}^{min}$  as a function of  $N$ , where  $P_{err}^{min}$  represents the error probability at the optimal threshold  $T_{opt}$  for a fixed prior probability  $p_u = \frac{1}{3}$ . Evidently,  $P_{err}^{min}$  decays exponentially as  $\mathcal{O}(a^{-N})$ , where the base  $a > 1$  depends on the bit-flip error probability  $\epsilon$ . The scaling behavior is provided for different values of the bit-flip parameter  $\epsilon$ .

As a final remark, it is important to note that the error model considered here may not fully capture all the complexities of realistic quantum systems. A more accurate description of qubit-environment interactions can lead to more complex noise models, such as an amplitude damping channel. The amplitude damping channel is particularly relevant for modeling quantum systems where energy loss is significant. For example, it introduces an asymmetry in the probabilities of state transitions, with a higher likelihood of decay from  $|1\rangle$  to  $|0\rangle$  than the reverse transition. Such scenarios would require a recomputation of the decoding strategy to account for the directional nature of the noise.

Nevertheless, the fundamental advantage of communication using qubits, as opposed to classical

bits, would remain intact. This advantage arises from the inherently distinct nature of the random state. While noise is typically regarded as an obstacle that must be mitigated for effective communication, in the context of the HLPE under a repetition code, we demonstrate noise is the key enabler of quantum advantage, allowing for a higher success rate than is achievable through classical means.

This protocol explores the interaction between ontological randomness and epistemic randomness. Ontological randomness is an intrinsic feature of quantum mechanics, manifesting during qubit readout, as described by the Copenhagen interpretation and other non-deterministic quantum frameworks through the Born rule in Eq. (2.19). Epistemic randomness, in contrast, arises from incomplete knowledge of the state of the system during transmission and is modeled using noise channels, a mathematical framework to transform pure quantum states into a probabilistic ensemble, or classical mixture, of quantum states. In this context, leveraging the latter in the presence of the former provides a distinct advantage in solving the  $(\text{HLPE}_{\text{syn}}^2)$ .

## 2.9 Discussion and Conclusion

The *Hardest Logic Puzzle Ever* tests our ability to employ deductive reasoning to obtain knowledge from the absolute minimum amount of information by strategically posing logical questions to three unfamiliar agents who communicate in an unknown language and may exhibit random behavior [143].

The riddle is rooted in the tradition of logic puzzles involving “knights” (who always tell the truth) and “knaves” (who always lie), popularized by Raymond Smullyan in his works on recreational logic, and contains a number of interesting features. From a logical standpoint it involves the interaction with three distinct logical personas: one agent who always preserves truth values, another who systematically negates them, and a third who randomizes them. The inclusion of this random agent is particularly noteworthy from a physical perspective, as it raises questions about the nature and interpretation of randomness — a topic of significant interest in physics and philosophy. In addition, the added difficulty of communicating without a shared language introduces further complexity, presenting challenges with implications for communication theory and the study of epistemic barriers.

Since its inception, the HLPE has attracted significant academic attention and received different analyses. With only a couple of exceptions [144, 149], the focus of most studies has primarily been on extending the range of possible answers provided by the “gods” or altering the behavior of  $\mathfrak{R}and\mathfrak{O}m$ . The aim has been to explore logical validity within different interpretative frameworks and exploit loopholes that might enable solutions in two questions under additional assumptions, such as the possibility of addressing self-

referential or undecidable questions, or reason following non-classical logics [145, 146, 147, 148, 154].

Although never explicitly mentioned, all prior analyses of the HLPE have implicitly assumed innate classical communication between the puzzle’s agents. In this work, we have analyzed the communication problem and its implications for solution strategies in both the classical and the quantum communication scenarios, thereby contributing an ontological dimension grounded in the physical realizability of the puzzle for the first time. We analyze the riddle entirely within a classical logic framework, constraining the acceptable strategies to use only valid YES-NO questions. As a result, difficulty is also increased by restricting the set of acceptable questions. Specifically, our formalization and physical realization correspond precisely to (HLPE<sub>syn</sub><sup>2</sup>) supplemented by conditions (B0), (B1), (B2), (B3\*), and (B4).

We began by revisiting Boolos’ classical solution strategy in the exactly solvable case — requiring three binary-answer questions — before proceeding with a new analysis of classical probabilistic solutions involving zero, one, and two questions. Particular emphasis was placed on solution strategies for both minimum-error and unambiguous discrimination in only two questions. The corresponding success probabilities were found to be 50% and 25%, respectively.

Having established these baselines for novel solutions to outperform, we proceeded to analyze quantum-communication-based implementations designed to precisely replicate the classical-communication-based solutions when the allowed interactions were restricted to classical physics. This approach extended the scope of the HLPE to the quantum domain without changing its fundamental nature.

We focused our analysis on two quantum implementations. On both, the quantum operators used to implement the identity, negation, and randomization logic render the bits recovered from the communication qubit indistinguishable from those obtained by performing an equivalent classical operation instead. This reproduces classical logic and the puzzle’s setup in perfect agreement with the original formulation. Furthermore, all of the solution strategies we discovered use straightforward, truth-functional, non-subjunctive and non-contrary-to-fact logical queries to the gods.

In our first implementation, we postulated a particular qubit encoding using quantum operators that allow their perfect distinguishability via the superdense coding protocol. The solution we presented solves the HLPE in just two binary-answer questions. Additionally, it offers novel advantages: the questions can be formulated in advance, independent of previous answers; each of the gods can employ their own distinct encoding language, independent of its peers; and the gods need not be aware of each other’s identities or even their own, something deemed impossible before [149]. If this solution comes as a surprise, it is because, while making use of argumentative logic, it hinges on exploiting physical principles. Nonetheless, it crucially demonstrates that an exact solution to the puzzle can be found in just two questions using only

YES-NO answers, while avoiding the issues of self-reference and undecidability.

In the aftermath of this implementation, we highlight the fact that different physical operations — whether classical or quantum — can implement the same classical logic. While this observation yields no advantage in a classical-communication based context, it can be exploited quantum mechanically, as we demonstrate with our solution. Recognizing that, our second implementation explores opportunities for provable quantum advantage in a more general framework minimizing assumptions about both parties.

Specifically, we considered the natural constraint imposed by the geodesic path traversed by photons exchanged between two parties communicating in free space, in order to relax mutual assumptions about the computational basis used for communication. This framework, in turn, restricted the set of permissible operators for implementing the required classical logic, thereby facilitating a more natural agreement between the communicating parties. While an exact solution to the HLPE in just two questions could not be obtained in this framework, as in the previous one, a probabilistic advantage was still identified. Specifically, a minimum-error discrimination strategy applied to a non-degenerate subset of these operators yielded a success probability of  $\approx 60\%$ , contrasting with the  $50\%$  achieved classically.

We then made our analysis more realistic by incorporating the unavoidable effect of noise on our protocol. This introduces epistemic randomness alongside ontological randomness, allowing for a clear distinction between their respective roles in this HLPE realization. Using a classical error model for symmetric bit-flip errors and an equivalent quantum noise model implemented via a depolarizing channel with the same bit-flip error probability, we compared the performance of all the previously presented protocols. We observed that quantum solutions outperform classical ones for all values of the bit-flip error probability, both in the exact protocols — Boolos versus superdense-coding-based — and in the probabilistic ones — the classical minimum-error discrimination strategy versus the generalized quantum protocol.

More importantly, we demonstrated that extending the communication protocol with a repetition code — necessary for mitigating noise and achieving a negligible error rate, regardless of implementation — enables the recovery of an asymptotically-exact solution to the HLPE in just two questions. This holds with arbitrarily high success probability, even in the fully generalized scenario that considers all the four possible quantum operators implementing the corresponding classical logic.

To summarize, we demonstrate provable quantum supremacy in a specific interaction-based logical puzzle. Notably, this advantage is not a moving target that could be challenged by classical protocols in the future; it is a definitive edge of a quantum solution over classical ones. Furthermore, it does not arise from engineering artifices or spying mechanisms, but from the operational principles of the physical laws themselves. The definitive advantage of this protocol bears similarity to results such as Deutsch's problem

in quantum computing [205], communication protocols like quantum teleportation [206], and quantum game-theoretic strategies [207], where quantum solutions outperform classical approaches. However, in our case, the computation is interactive, involving minimal communication between parties to solve the problem, rather than just a purely computational or communication framework.

As we advance in developing quantum information processing technologies, identifying quantum algorithms and protocols with provable quantum advantage over classical counterparts remains an important priority, as such examples are still relatively few. The proposed quantum protocols to defeat the HLPE in two questions with minimal assumptions and stricter constraints, though their immediate applications may not be readily apparent, nonetheless provide another concrete example of a problem where quantum advantage can be rigorously demonstrated.

From a more practical perspective, the protocol we propose based on relaxing the computational basis shares parallels with blind delegated quantum computing, where computation is entrusted to a remote quantum processor without revealing its quantum states or the specific nature of the computation. It could also serve as a novel method for benchmarking quantum communication technologies, particularly space-based ones. One could even envision potential applications in unforeseen or contrived space-based communication scenarios, where the vast distances complicate and delay interactions. In such cases, the ability to quickly identify the logical type of an agent with a minimal number of queries could facilitate real-time decision-making and offer practical value.

From a philosophical point of view, the debate over whether logic has an empirical basis often focuses on propositions concerning shared characteristics among sets of physical objects [208]. When these objects are quantum systems, such characteristics correspond to subspaces of their Hilbert space [169]. In contrast, this work draws insights from the physical mechanisms underlying the implementation of classical logic operations. The protocols we introduce demonstrate how optimal knowledge extraction connects these mechanisms to logical reasoning, thereby illuminating the relationship between abstraction and material constraints. In particular, we show that employing quantum mechanics to implement classical logic operations and communication allows for the exploitation of novel resources, such as superposition and entanglement, to achieve more advantageous solution protocols. By bridging formal logic with ontological constructs and empirical reality, we demonstrate that the physical nature of logic is not merely a vacuous observation but has direct implications for how physical agents can represent, process, and reason with both abstract and physical sources of information.

Furthermore, quantum computing also challenges traditional notions of mathematical proof and computation by performing tasks that classical computers cannot and making classical verification impossible.

Since a quantum verification would not produce a physical record, accepting such proofs requires trusting the physical process itself, rather than an explicit, verifiable document [209]. In our case, rather than computing proofs with quantum resources, Alice reasons classically while understanding how classical logic is instantiated through quantum operations, thereby rendering otherwise unprovable statements provable.

To conclude, logic puzzles provide a playground to demonstrate deductive reasoning, logical consistency and various logical paradoxes. They offer a valuable means of exploring the limits of logical reasoning and the nature of truth. Beginning with simple truth-teller and liar puzzles, the genre evolved to incorporate more complex scenarios like the HLPE. This work offers new insights into the nature of reasoning by incorporating physical implications into the genre.

For example, after analyzing the HLPE through the lens of quantum mechanics, and recognizing that Smullyan worked on similar puzzles to illustrate Gödel’s incompleteness theorems, a question comes to mind: Could we use one of Smullyan’s puzzles to perform similar physics-based analyses and arrive at new embodiments of Gödel’s theorem? Specifically, could such analyses reveal intersections between physical principles and unprovable statements, thus grounding abstract logical limits in physical reality?

This chapter ends with a short extra section where we add to the tradition of proposing variations to the HLPE. We see them as a useful playground, not only to challenge human thinkers but also to provide the AI community with novel logic problems to evaluate and improve AI reasoning capabilities.

## 2.10 Coda: Extensions and Artificial Intelligence

Is the HLPE truly the *most challenging* logic puzzle ever devised? While this question might certainly spark debate, what remains undeniable is the rich tradition of proposing variations to the HLPE, a practice that began immediately following the publication of Boolos’s seminal paper [143].

Roberts proposed two variations of the puzzle [144]: in the first, two of the gods are of type **R**andom, while the third is either **T**ru<sup>e</sup> or **F**alse; in the second, two gods are either **T**ru<sup>e</sup> or **F**alse, with the third being **R**andom. Rabern and Rabern modified Boolos’ puzzle by modeling **R**andom as a binary random variable [145] and analyzing it as we described. Uzquiano proposed an extension to this model by treating **R**andom as a ternary random variable [146]. Wheeler and Barahona introduced a variation where **F**alse is replaced by so-called **D**ec<sup>i</sup>ous, who lies when certain he can do so but otherwise responds as a ternary random variable to paradoxical questions [147]. Finally, Wintein [148] proposed a version of the HLPE where all gods have a four-option answer set: “huh”, “duh”, “da”, and “ja”.

With the exception of Robert’s proposals, which involve reconfigurations of the same set of god types,

subsequent variations primarily focus on expanding the possible range of responses the gods can provide or altering the behavior of *Random*. These modifications aim to explore logical validity under different interpretive frameworks and complexity conditions.

However, the question of whether the original HLPE or any of its variations truly qualifies as the hardest logic puzzle ever devised is likely not the most important one. Equally compelling is the prospect of extending the problem to introduce new conceptual layers, which could offer novel insights, interdisciplinary connections, or applications. Therefore, we provide two extensions to the HLPE, which can be combined together or with previous versions to create additional variations. Our goal is not to complicate the problem without purpose, but to expand its conceptual scope.

The first extension involves an adversary that introduces *double encoding*. It builds on top of (HLPE<sup>2,DE</sup><sub>syn</sub>) with the addition of a final language translation. It enunciates as follows, where boldface is used to indicate our specific additions or modifications:

(HLPE<sup>2,DE</sup><sub>syn</sub>) Three gods A, B, and C are called, in some order, *True*, *False*, and *Random*. *True* always speaks truly, *False* always speaks falsely, but whether *Random* answers “da” or “ja” is a completely random matter. Your task is to determine the identities of A, B, and C by asking **only valid YES-NO questions**; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language in which the words for YES and NO are “da” and “ja”, in some order. You do not know which word means which. **A demon secretly intercepts the words from the gods and may or may not flip them, doing so consistently before sending them to you. How few questions suffice to reveal their identities?**

To this version, conditions (B0), (B1), (B2), (B3\*), and (B4) still apply. In addition, we clarify the operational character of the demon with the following condition:

(B5) Once established, the demon’s behaviour — flipping the answers or not — is invariable, although unknown to both the solver and the gods.

Note that Boolo’s artifice, as demonstrated with questions  $\mathcal{Q}_1$  and  $\mathcal{Q}_2$  in Table 2.2, is no longer effective because the outcome in the “state received” column of the table can subsequently be preserved or flipped by the demon. Since the gods are unaware of the demon’s encoding map and cannot be inquired

about it, there are now four possible encodings being employed:

$$\text{LBM}_1 = \{(1, da, da), (0, ja, ja)\}, \quad (2.86)$$

$$\text{LBM}_2 = \{(1, da, ja), (0, ja, da)\}, \quad (2.87)$$

$$\text{LBM}_3 = \{(0, da, da), (1, ja, ja)\}, \quad (2.88)$$

$$\text{LBM}_4 = \{(0, da, ja), (1, ja, da)\}, \quad (2.89)$$

where  $\text{LBM}_i$  is a bijective mapping. Each triplet  $(a, b, c)$  indicates that  $a$  maps to  $b$  and  $c$ , respectively, where  $a$  represents the output Boolean value from a god's evaluation of the question,  $b$  is the god's reply, and  $c$  is the final reply received by the solver.

The second puzzle variation we propose introduces *contextuality* in the problem, where the logical character of one of the gods is determined by the value of previous answers.

( $\text{HLPE}_{syn}^{2,CO}$ ) **Four** gods A, B, C, **and** D are called, in some order, *True*, *False*, *Random*, and *Context*. *True* always speaks truly, *False* always speaks falsely, but whether *Random* answers “da” or “ja” is a completely random matter. *Context* **answers “da” when an odd number of previous answers were “da”, and “ja” otherwise.** Your task is to determine the identities of A, B, C, **and** D by asking **only valid YES-NO questions**; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language in which the words for YES and NO are “da” and “ja”, in some order. You do not know which word means which. **How few questions suffice to reveal their identities?**

As for ( $\text{HLPE}_{syn}^{2,DE}$ ), conditions (B0), (B1), (B2), (B3\*), and (B4) still apply on this version. As is apparent, we now need to identify the correct permutation among  $4! = 24$  possibilities. Considering only classical solutions, it is evident that at least 5 questions will be required, since  $2^4 < 24$ . Therefore, at least one of the gods will inevitably be addressed twice. However, without further analysis, it remains uncertain whether asking only five questions suffices to solve this version.

In this puzzle, we introduce a novel god whose behavior depends on the history of previous answers. This adds a layer of complexity in which the answers are interdependent, and the behavior of neighboring gods cannot be ignored in the answers of at least one of them. The HLPE, as initially presented in ( $\text{HLPE}_{sem}$ ), and later amended in ( $\text{HLPE}_{syn}^2$ ) and tightened in this work with (B0), can be solved by posing questions that do not require the gods to listen to the answers of their peers. In contrast, this version inexorably links the gods' responses.

As a result, not only Alice but also one the gods are forced to interact with each other by keeping a memory record of what their peers have replied. A quantum mechanical implementation might potentially offer a means to execute the operation of `Context`, which is conditioned by the responses of his peers, in a coherent manner, which could provide a way to explore different causal narratives in the puzzle.

$(HLPE_{syn}^{2,DE})$  and  $(HLPE_{syn}^{2,CO})$  are introduced informally to preserve the enticing spirit of the original HLPE, serving as inspiration and a starting point for more rigorous formalizations. In tackling them, we propose limiting both questions and answers to the ones of binary type that can be modeled with the transmission of a classical or quantum two-state system, as described here. We will not pursue their solutions here, leaving this endeavor to be explored by future puzzle solvers.

Finally, we also believe these puzzles provide an ideal framework for evaluating the logical reasoning capabilities of artificial intelligence (AI) models. Recent years have witnessed a surge of innovative techniques and breakthroughs in AI. These include deep neural networks, which exploit the full potential of backpropagation; advancements in reinforcement learning and search algorithms, enabling strategic decision-making in complex environments; and generative models, such as generative adversarial networks and diffusion models, which have substantially improved data synthesis. Additionally, the new transformer architecture has revolutionized sequence modeling through attention mechanisms. Transformers enabled parallelized training and achieved groundbreaking results in natural language understanding, while also driving scalability and performance in large language models and multimodal systems.

Building on these advancements, AI is now beginning to exhibit early signs of reasoning abilities, demonstrating its potential to aid in scientific discovery across a variety of domains bordering with the puzzles we present. Notable examples include breakthroughs in games [210, 211], where reinforcement learning has achieved superhuman performance, and in physics [212], where symbolic regression techniques have uncovered interpretable physical laws. Similarly, in computer science [213], AI has been used to discover new algorithms, while in mathematics [214, 215, 216, 217, 218, 219, 220], it has enabled the exploration of novel insights, the formulation of new conjectures, advancements in open problems, and solutions to mathematical problems in international competitions.

Emerging efforts, such as the integration of these techniques into more comprehensive reasoning frameworks, are paving the way for end-to-end models of scientific discovery [221], highlighting the potential for a convergence of methodologies to produce more capable and general-purpose reasoning systems.

Given these impressive advancements and the promising potential for further breakthroughs in the near future, we encourage challenging AI model's reasoning capabilities using  $(HLPE_{syn}^{2,DE})$  and  $(HLPE_{syn}^{2,CO})$  as benchmarks. Will an AI solve them before a human does?

## 2.11 Appendix

### 2.11.1 Minimum-error discrimination via semidefinite programming

Semidefinite programming (SDP) techniques are broadly applicable in the field of quantum information, and in particular also in quantum state discrimination. This is because the mathematical representation of quantum states and measurements using positive semidefinite operators aligns seamlessly with the SDP framework [222].

In minimum-error quantum state discrimination, the objective is to maximise the average probability of successfully identifying the received quantum state in a single-shot measurement. In our problem, we are tasked with distinguishing between the set of quantum states  $\{\rho_i\}$  in Eqs (2.52 – 2.57), each occurring with equal prior probability  $p_i = 1/6$ , where  $i = 1, \dots, 6$ . To achieve this, we use a measurement described by a set of positive semidefinite operators  $\{E_i\}$ , which form a positive operator-valued measure (POVM) satisfying  $\sum_i E_i = \mathbb{1}$ . The optimization problem can then be formulated as follows:

$$\begin{aligned} P_{\text{succ}}^* = \quad & \text{maximize} \quad \sum_i p_i \text{Tr}(E_i \rho_i) \\ & \text{subject to} \quad E_i \succeq 0 \quad \quad \quad i = 1, \dots, 6 \\ & \quad \quad \quad \sum_i E_i = \mathbb{1} \end{aligned} \quad (2.90)$$

To implement this SPD, we use the `cvxpy` v.1.5.2 library [223, 224] with Python v.3.12.4 and `numpy` v.1.26.4 for numerical computations [225]. Although SDP is a convex optimization problem with an exact optimal solution, the computational efficiency and numerical stability of the solution can vary depending on the solver used. In this work, we use the Splitting Conic Solver (SCS), which implements first-order methods that avoid matrix factorizations and are better suited for large or sparse problems. While this particular problem is small in scale, SCS provides an efficient and reliable solution. The source code can be inspected in Listing 2.

---

```
1 import cvxpy as cp
2 import numpy as np
3
4 zero = np.array([[1],[0]], complex)
5 one  = np.array([[0],[1]], complex)
6
7 psi_IDE_p  = zero
8 psi_NOT_p  = one
9 psi_RAN_pp = (zero + one) / np.sqrt(2)
10
11 phis = [np.kron(psi_IDE_p, psi_NOT_p), # phi_1
12         np.kron(psi_IDE_p, psi_RAN_pp), # phi_2
13         np.kron(psi_NOT_p, psi_IDE_p), # phi_3
14         np.kron(psi_NOT_p, psi_RAN_pp), # phi_4
```

```

15         np.kron(psi_RAN_pp, psi_IDE_p), # phi_5
16         np.kron(psi_RAN_pp, psi_NOT_p)] # phi_6
17
18     rhos = [phi @ phi.T.conj() for phi in phis]
19     priors = [1/6] * 6
20
21     # Define the POVM elements as CVXPY variables
22     E = [cp.Variable((4, 4), PSD=True) for _ in range(6)]
23
24     # Objective: Maximize the sum of probabilities of correct identification
25     objective = cp.Maximize(cp.sum([priors[i] * cp.real(cp.trace(rhos[i] @ E[i])) for i in
26         ↪ range(6)]))
27
28     # Constraint: POVM elements sum to the identity matrix
29     constraints = [cp.sum(E) == np.eye(4)]
30
31     # Solve the problem
32     prob = cp.Problem(objective, constraints)
33     prob.solve(solver=cp.SCS, max_iters=1000, eps_abs=1e-10, eps_rel=1e-10)
34     print("The optimal value is:", prob.value)

```

---

Listing 1: Minimum-error state discrimination via SDP.

### 2.11.2 Quantum circuit decomposition for the square-root measurement

To complement the discussion in § 2.7.2, we provide the explicit OpenQASM code for implementing a quantum circuit that realizes the square-root POVM measurement using a five-qubit register, where three qubits are measured in the computational basis. This listing specifies the gate sequence required for direct experimental implementation. The circuit has been optimized for reduced CNOT gate count using the techniques introduced in § 4.

```

1  OPENQASM 3.0;
2  include "stdgates.inc";
3  qubit[5] q;
4  rz(3*pi/2) q[0];
5  rz(2.0156946614431264) q[1];
6  rz(1.8498725781753507) q[2];
7  rx(0.49518884406633923) q[2];
8  rz(-0.49996410580842987) q[3];
9  rx(2.785002852674138) q[3];
10 rx(-1.3228706501132046) q[4];
11 rz(1.2916579247925615) q[4];
12 cx q[0], q[4];
13 rx(-0.7430636574279852) q[4];
14 rz(-2.770413339491739) q[4];
15 cx q[1], q[4];
16 rz(5.912005993081331) q[4];
17 cx q[0], q[4];
18 rz(3.5127719676880456) q[4];
19 cx q[3], q[4];
20 rx(-1.3755650053974744) q[3];
21 rz(pi/2) q[3];
22 rx(pi/2) q[3];
23 cx q[0], q[3];
24 rz(5.7313641848665045) q[3];
25 cx q[1], q[3];
26 rz(0.5518211223130849) q[3];
27 cx q[0], q[3];
28 rz(5.731364184866502) q[3];
29 cx q[2], q[3];
30 rx(2.2259994448827563) q[2];
31 rz(pi/2) q[2];
32 rx(pi/2) q[2];
33 rz(pi/2) q[3];
34 rx(pi/2) q[3];
35 rz(pi/2) q[3];
36 rz(pi/2) q[4];
37 rx(pi/2) q[4];
38 rz(pi/2) q[4];
39 cx q[4], q[2];
40 rz(5.084691385559514) q[2];

```

```

41  cx q[3], q[2];
42  rz(1.563487720376832) q[2];
43  cx q[0], q[2];
44  rx(pi) q[0];
45  rz(pi/4) q[2];
46  cx q[1], q[2];
47  rz(3*pi/4) q[2];
48  rx(0.8592236731771836) q[2];
49  cx q[0], q[2];
50  rz(pi/2) q[0];
51  rx(pi/2) q[0];
52  rz(-7*pi/16) q[0];
53  cx q[1], q[0];
54  rz(31*pi/16) q[0];
55  rz(-2.1492955707682775) q[2];
56  rx(2.135916490376377) q[2];
57  rz(-0.8839454354143501) q[2];
58  cx q[2], q[0];
59  rz(31*pi/16) q[0];
60  cx q[1], q[0];
61  rz(pi/16) q[0];
62  cx q[3], q[0];
63  rz(pi/16) q[0];
64  cx q[1], q[0];
65  rz(31*pi/16) q[0];
66  cx q[2], q[0];
67  rz(31*pi/16) q[0];
68  cx q[1], q[0];
69  rz(17*pi/16) q[0];
70  cx q[4], q[0];
71  rz(15*pi/16) q[0];
72  rx(-2.2917710380486653) q[4];
73  rz(-pi/2) q[4];
74  cx q[1], q[4];
75  cx q[1], q[0];
76  rz(pi/16) q[0];
77  cx q[2], q[0];
78  rz(pi/16) q[0];
79  cx q[1], q[0];
80  rz(31*pi/16) q[0];
81  cx q[3], q[0];
82  rz(31*pi/16) q[0];
83  cx q[1], q[0];
84  rz(17*pi/16) q[0];
85  cx q[2], q[0];
86  rz(9*pi/16) q[0];
87  rx(2.884017450970407) q[2];
88  rz(pi/2) q[2];
89  rx(pi/2) q[2];
90  rx(1.224920102305071) q[3];
91  rz(pi/2) q[3];
92  rx(pi/2) q[3];
93  rz(pi/2) q[4];
94  rx(-0.8498216155411225) q[4];
95  cx q[4], q[3];
96  rz(1.1998220520969358) q[3];
97  cx q[1], q[3];
98  cx q[1], q[0];

99  rz(pi/2) q[0];
100 rx(pi/2) q[3];
101 rz(pi/2) q[3];
102 rx(0.7168504991877915) q[3];
103 cx q[3], q[2];
104 rz(6.229383781412431) q[2];
105 cx q[4], q[2];
106 rz(3.632801421069135) q[2];
107 cx q[3], q[2];
108 rx(-1.211892475095747) q[2];
109 rz(2.663986897675525) q[2];
110 rx(-0.344140602109392) q[2];
111 cx q[1], q[2];
112 rx(7*pi/8) q[1];
113 rz(pi/2) q[1];
114 rz(2.2453918172948892) q[2];
115 rx(1.0491909156523553) q[2];
116 rz(2.1279753303795177) q[2];
117 cx q[1], q[2];
118 rx(-pi/2) q[1];
119 rz(5*pi/8) q[1];
120 rz(pi/2) q[2];
121 rx(pi/2) q[2];
122 rz(-1.2718692285849418) q[2];
123 cx q[3], q[1];
124 rz(pi/8) q[1];
125 cx q[2], q[1];
126 rz(pi/8) q[1];
127 cx q[4], q[1];
128 rz(pi/8) q[1];
129 cx q[2], q[1];
130 rz(pi/8) q[1];
131 cx q[3], q[1];
132 rz(9*pi/8) q[1];
133 cx q[2], q[1];
134 rx(pi/2) q[1];
135 rz(pi/2) q[1];
136 rx(5*pi/8) q[1];
137 cx q[1], q[0];
138 rz(pi/16) q[0];
139 rz(6.161307177024823) q[1];
140 cx q[2], q[0];
141 rz(pi/16) q[0];
142 cx q[1], q[0];
143 rz(pi/16) q[0];
144 cx q[3], q[0];
145 rz(pi/16) q[0];
146 cx q[1], q[0];
147 rz(pi/16) q[0];
148 cx q[2], q[0];
149 rz(pi/16) q[0];
150 cx q[1], q[0];
151 rz(17*pi/16) q[0];
152 cx q[4], q[0];
153 rz(17*pi/16) q[0];
154 cx q[1], q[0];
155 rz(pi/16) q[0];
156 cx q[2], q[0];

```

```

157 rz(pi/16) q[0];
158 cx q[1], q[0];
159 rz(pi/16) q[0];
160 cx q[3], q[0];
161 rz(pi/16) q[0];
162 cx q[1], q[0];
163 rz(pi/16) q[0];
164 cx q[2], q[0];
165 rz(25*pi/16) q[0];
166 cx q[1], q[0];
167 rz(9*pi/16) q[0];
168 rx(1.6568486249446597) q[2];
169 rz(pi/2) q[2];
170 rx(pi/2) q[2];
171 rx(1.7988450702231678) q[3];
172 rz(pi/2) q[3];
173 rx(pi/2) q[3];
174 rx(-1.9943643727758875) q[4];
175 rz(-0.7438254328581934) q[4];
176 cx q[4], q[0];
177 rz(pi/2) q[0];
178 rx(pi/2) q[0];
179 rz(pi/2) q[0];
180 rx(1.021107634203453) q[4];
181 rz(pi/2) q[4];
182 rx(-2.4763237670922873) q[4];
183 cx q[4], q[3];
184 rz(-2.989804419039243) q[3];
185 cx q[0], q[3];
186 rx(2.704981953542296) q[3];
187 rz(1.223672432754145) q[3];
188 rx(-1.4133594294788199) q[3];
189 cx q[3], q[2];
190 rz(0.08605229814976972) q[2];
191 rz(4.3034364521848625) q[3];
192 rz(4.350826685668594) q[4];
193 cx q[4], q[2];
194 rz(0.5042403442438911) q[2];
195 cx q[3], q[2];
196 rz(0.8944513862514877) q[2];
197 cx q[0], q[2];
198 rx(-9*pi/16) q[0];
199 rz(pi/2) q[0];
200 rx(pi/2) q[0];
201 cx q[1], q[0];
202 rz(31*pi/16) q[0];
203 rx(pi/2) q[2];
204 rz(pi/2) q[2];
205 cx q[2], q[0];
206 rz(pi/16) q[0];
207 cx q[1], q[0];
208 rz(pi/16) q[0];
209 cx q[3], q[0];
210 rz(pi/16) q[0];
211 cx q[1], q[0];
212 rz(pi/16) q[0];
213 cx q[2], q[0];
214 rz(31*pi/16) q[0];

```

```

215 cx q[1], q[0];
216 rz(31*pi/16) q[0];
217 cx q[4], q[0];
218 rz(31*pi/16) q[0];
219 cx q[1], q[0];
220 rz(31*pi/16) q[0];
221 cx q[2], q[0];
222 rz(pi/16) q[0];
223 cx q[1], q[0];
224 rz(pi/16) q[0];
225 cx q[3], q[0];
226 rz(pi/16) q[0];
227 cx q[1], q[0];
228 rz(pi/16) q[0];
229 cx q[2], q[0];
230 rz(4.516039439535327) q[0];
231 cx q[1], q[0];
232 rz(25*pi/16) q[0];
233 rx(7*pi/8) q[1];
234 rz(pi/2) q[1];
235 rx(pi/2) q[1];
236 rx(-2.406767079279307) q[2];
237 rz(-2.339446234921606) q[2];
238 rx(1.8814521636677144) q[3];
239 rz(pi/2) q[3];
240 rx(pi/2) q[3];
241 rx(-1.470628905633335) q[4];
242 rz(-pi/2) q[4];
243 cx q[4], q[3];
244 rz(0.3106558368728146) q[3];
245 cx q[2], q[3];
246 rx(1.15209507928315) q[2];
247 rz(1.172805975478288) q[2];
248 cx q[4], q[2];
249 rx(2.425086454813969) q[2];
250 rz(pi/2) q[2];
251 rx(1.057041538340239) q[2];
252 cx q[2], q[3];
253 rx(-0.6791550671400084) q[2];
254 cx q[2], q[1];
255 rz(6.162965310615633) q[1];
256 rz(3*pi/2) q[2];
257 rz(-pi/2) q[3];
258 rx(-pi/2) q[3];
259 cx q[3], q[1];
260 rz(6.241936596560329) q[1];
261 cx q[2], q[1];
262 rz(6.0952401771755715) q[1];
263 cx q[4], q[1];
264 rz(5.309842013778126) q[1];
265 cx q[2], q[1];
266 rz(0.19919128250865697) q[1];
267 cx q[3], q[1];
268 rz(0.12021999656394633) q[1];
269 cx q[2], q[1];
270 rx(pi/2) q[1];
271 rz(pi/2) q[1];
272 rx(5*pi/8) q[1];

```

```

273 cx q[1], q[0];
274 rz(6.262642597457667) q[0];
275 rz(pi/2) q[1];
276 cx q[2], q[0];
277 rz(0.07143718357630909) q[0];
278 cx q[1], q[0];
279 rz(0.33581350401147225) q[0];
280 cx q[3], q[0];
281 rz(5.899634870787048) q[0];
282 cx q[2], q[0];
283 rz(0.7775809849684469) q[0];
284 cx q[1], q[0];
285 rz(6.271902283760495) q[0];
286 cx q[4], q[0];
287 rz(0.011283023419051815) q[0];
288 cx q[1], q[0];
289 rz(6.275368128750695) q[0];

```

```

290 cx q[2], q[0];
291 rz(0.515618882379504) q[0];
292 cx q[3], q[0];
293 rz(1.2349828227834303) q[0];
294 cx q[1], q[0];
295 rz(0.07143718357630227) q[0];
296 cx q[2], q[0];
297 rz(3.143615990706258) q[0];
298 cx q[1], q[0];
299 rz(-5*pi/8) q[0];
300 rx(pi/2) q[0];

```

---

Listing 2: OpenQASM circuit instruction for the square-root measurement, comprising 103 CNOT gates and a total depth of 200 layers.

## Chapter 3

# Testing complementarity on a transmon quantum processor <sup>1</sup>

“‘Well,’ you say, ‘what about Proposition A? Is it true, or is it *not* true, that the electron either goes through hole 1 or it goes through hole 2?’”

---

Richard Feynman, 1965 [226]

### 3.1 Introduction

Complementarity represents one of the cornerstone principles of quantum mechanics, encapsulating the inherent limitations in acquiring simultaneous information about non-commuting observables. The principle first arose in discussions of wave-particle duality but it was rapidly understood as a general feature of quantum systems. This concept is often illustrated using a quantum system capable of taking two distinct alternatives or paths, denoted as  $|0\rangle$  and  $|1\rangle$ . When prepared in a superposition state, such a system exhibits a wavefunction of the form:

$$|\psi\rangle = c_0|0\rangle + c_1e^{i\phi}|1\rangle, \quad (3.1)$$

where the coefficients  $c_0$  and  $c_1$  describe the probability amplitudes associated with each path, and  $\phi$  represents their relative phase.

Interference patterns emerge when the phase  $\phi$  is varied, and the expectation value of a complemen-

---

<sup>1</sup>This chapter is based on joint work published in collaboration with J. Fernández-Rossier in Ref. [76].

tary observable is measured. The visibility of these interference fringes, proportional to  $\text{Re} \{c_0^* c_1\}$ , arises from the simultaneous traversal of both paths by the quantum object—an inherently *wavelike* behavior. However, introducing a mechanism that couples with the system to determine which path it takes—a *particlelike* property—leads to a suppression of interference visibility, illustrating the dual nature of quantum systems.

This nuanced interplay between wavelike and particlelike behaviors was formalized in Bohr’s principle of complementarity [164], which posits that certain pairs of properties in quantum systems—such as position and momentum, or unambiguous particle trajectories and interference fringes—are mutually exclusive in their precise determination. Bohr’s insights drew upon foundational contributions, including De Broglie’s hypothesis of matter waves [227] and Heisenberg’s uncertainty principle [228]. Together, these works established complementarity as an intrinsic feature of quantum theory.

Bohr’s complementarity principle faced early scrutiny through Einstein’s 1927 proposal of a recoiling double-slit interferometer, which sought to falsify the principle by demonstrating the simultaneous observability of which-path information and interference [229]. This was the first example of a which-path thought experiment. Subsequent theoretical analyses and other experimental proposals further underscored the impossibility of simultaneously observing both aspects of quantum behavior [230, 226, 231]. A significant milestone occurred in 1979 when Wootters and Zurek [232] quantified the trade-off between which-path information and interference visibility by analyzing the momentum exchange between a photon and a detector in Einstein’s *gedanken* experiment. Their work extended Bohr’s qualitative dichotomy into a continuous framework.

Beyond the position-momentum uncertainty principle, new proposals emerged with alternative mechanisms enforcing complementarity [233, 234]. Subsequent investigations emphasized the role of entanglement and correlations in complementarity and advanced an information-theoretic treatment inspired by the analysis of Wootters and Zurek [235, 236, 237]. Notably, Jaeger *et al.* [238] and Englert [239] introduced a duality relation for two-dimensional bipartite systems, expressed as:

$$\mathcal{V}^2 + \mathcal{D}^2 \leq 1, \quad (3.2)$$

where  $\mathcal{V}$  quantifies interference visibility and  $\mathcal{D}$  measures distinguishability of path determination.

The duality in Eq. (3.2) captures both the sharp and partial complementary nature of quantum information and was shortly extended to encompass quantum erasure [240], higher-dimensional and multipath systems [241]. A further breakdown of this relation into a triality was derived early in the new century, showing that the *a posteriori* distinguishability of the paths,  $\mathcal{D}$ , can be divided into the *a priori* predictability

and the genuine quantum correlations as quantified by the concurrence. This triality was also generalized to bipartite systems of any dimension [242, 243].

Englert initially contended that the duality relation expressed in Eq. (3.2) was distinct from Heisenberg's uncertainty principle. However, subsequent developments have unveiled a new entropic formulation of the uncertainty principle, demonstrating its equivalence to the duality relation in the context of multipath interferometers [244, 245]. Additionally, alternative wave-particle duality relations employing various informational measures have been proposed, extending their applicability to Hilbert spaces of arbitrary dimensions [246, 247].

Wave-particle duality continues to be explored through ingenious which-path experiments across diverse platforms, including atom-photon systems [248, 249], photon-photon systems [250, 251, 252, 253, 254, 255], electron-electron systems [256, 257, 258], atomic nuclei within molecules [259, 260, 261], superconducting circuits [262, 263, 264], and optical interferometers, where the which-path information is recorded in additional degrees of freedom of the interfering object itself rather than stored in a different quantum object [265, 266]. Such experiments reveal the versatility of complementarity as a principle governing quantum behavior across various physical systems.

With the advent of transmon-based quantum processors over the past decade, complementarity tests have become feasible on novel platforms offering unprecedented levels of control. These systems enable repeated probing of the same quantum system at very low temperatures in rapid measurement sequences with high precision, facilitating deeper exploration of microscopic statistics and the fundamental principles of quantum mechanics. Recent studies have used multipath interferometers to examine visibility and distinguishability measures using unambiguous state discrimination measurements [267]. The above-mentioned triality relation was tested with a tomography-based two-qubit circuit in [268], while in [269] a tomographic approach was also used to test different complementarity relations obtained from density matrix properties of three-qubit random states. Transmon-based quantum processors have also been employed in testing other fundamental phenomena, such as Bell and Mermin inequalities [72, 270].

In this chapter, we focus on implementing two classes of which-path experiments on a transmon-based quantum processor. The first involves the controlled transfer of which-path information to a detector subsystem, testing the duality relation of Eq. (3.2) via minimum-error state discrimination. The second explores quantum erasure, demonstrating interference recovery. Both experiments are implemented using two-qubit quantum circuits, where one qubit acts as the interferometer and the other as the detector, with results interpreted in fine statistical detail to accurately identify the underlying physical models generating them and their fundamental repercussions.

The subsequent sections are organized as follows. § 3.2 introduces the theoretical framework for implementing which-path detection and erasure on gate-based quantum computers. § 3.3 describes the experimental methodology and presents the results. § 3.5 reassesses the experimental findings in light of deviations from theoretical predictions, and § 3.6 concludes with a discussion of the broader implications of our results.

## 3.2 Quantum circuits for which-path detection

### 3.2.1 The interferometer

We begin our exploration with the quantum circuit representation of the Mach-Zehnder interferometer (MZI) [7], depicted in Fig. 3.1a. It requires a single qubit, denoted as  $q_i$ . First, the Hadamard gate creates an equal superposition of the computational basis states,  $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ , analogous to the role of the first beam splitter in the MZI, which divides the photon into two potential paths. Second, a phase gate  $R_\phi$  introduces a relative phase shift, producing the state  $|\psi\rangle$  in Eq. (3.1) with coefficients  $c_0 = c_1 = \frac{1}{\sqrt{2}}$ . Finally, the second Hadamard gate, mimicking the functionality of the MZI's second beam splitter, enables the measurement of the expectation value of the Pauli  $X$  operator in the prepared state by performing a readout in the computational basis—the Pauli  $Z$  basis. This can be seen by considering the post-Hadamard state  $|\psi'\rangle = H|\psi\rangle$ , which satisfies

$$\langle\psi'|Z|\psi'\rangle = \langle\psi|H^\dagger Z H|\psi\rangle = \langle\psi|X|\psi\rangle = \cos\phi. \quad (3.3)$$

By measuring  $|\psi'\rangle$  in the computational basis, we retrieve this expectation value from the difference in the relative frequencies of observing the states  $|0\rangle$  and  $|1\rangle$ :

$$\langle\psi|X|\psi\rangle = P(|0\rangle) - P(|1\rangle), \quad (3.4)$$

where  $P(|0\rangle)$  and  $P(|1\rangle)$  denote the probabilities of measuring the qubit in the states  $|0\rangle$  and  $|1\rangle$ , respectively<sup>2</sup>. Using this single-qubit circuit, interference patterns can be observed as a function of  $\phi$ , enabling verification of both the coherence of the superposition and the applied phase shift. Nevertheless, this configuration provides no information regarding the specific path taken by the qubit—whether it was in

<sup>2</sup>To align the computational basis notation with the operator-based formalism of quantum measurements, we define  $P(+1) = P(|0\rangle)$  and  $P(-1) = P(|1\rangle)$ , where  $P(+1)$  and  $P(-1)$  refer to the probabilities of measuring the eigenvalues  $+1$  and  $-1$  of the Pauli- $Z$  operator. This allows direct computation of expectation values, e.g.,  $\langle Z \rangle = (+1)P(+1) + (-1)P(-1)$ .




Figure 3.1: **Which-path detection circuit.** **a**, Mach-Zehnder interferometer implemented as an equivalent quantum circuit. **b**, An extended circuit introducing a second qubit in the setup,  $q_d$ , acting as a which-path detector coupled to the interferometer qubit  $q_i$ . The readout basis are defined via  $P_i^\dagger$  and  $P_d^\dagger$  prior to the computational basis measurement. For  $P_i^\dagger = H$ , maximum amplitude interference is observed on the first qubit when  $\alpha = 0$ , while varying  $\alpha$  gradually diminishes interference on  $q_i$ , eventually depleting it at  $\alpha = \pi$ , regardless of the readout state of  $q_d$ . The optimal basis  $P_d^\dagger = O_\alpha^\dagger$  optimizes retrieval of path information from the detector's readout. Using  $P_d^\dagger = \mathbb{1}$  erases path information, facilitating recovery of high-visibility interference patterns on  $q_i$ . Adapted from [76]. © 2021 American Physical Society.

state  $|0\rangle$  or  $|1\rangle$  – prior to measurement in the  $X$  basis.

### 3.2.2 The detector

To implement a which-path detector, we augment the circuit with a second qubit as shown in Fig. 3.1b. This detector qubit,  $q_d$ , is prepared in an equal superposition of the computational basis before a controlled phase gate applies a relative phase  $\alpha$  to the combined system state when both qubits are in state  $|1\rangle$ <sup>3</sup>. The resulting two-qubit wavefunction, expressed in the basis  $|xy\rangle = |x\rangle_i \otimes |y\rangle_d$ , is given by

$$|\Psi\rangle = \frac{1}{\sqrt{2}} \left( |0\rangle \otimes |\delta_0\rangle + e^{i\phi} |1\rangle \otimes |\delta_1\rangle \right), \quad (3.5)$$

where  $|\delta_x\rangle$ , for  $x \in \{0, 1\}$ , represents the state in which the detector qubit is left for each possible state  $|x\rangle$ , or *path*, taken by the interferometer qubit:

$$\begin{aligned} |\delta_0\rangle &= \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle), \\ |\delta_1\rangle &= \frac{1}{\sqrt{2}} (|0\rangle + e^{i\alpha} |1\rangle). \end{aligned} \quad (3.6)$$

The purpose of using the detector qubit is to perform a single measurement on it to infer the path taken by the interferometer qubit following their interaction. To achieve this, a mapping between the

<sup>3</sup>That is,  $CR(\alpha)$  acts over  $|00\rangle$ ,  $|01\rangle$ ,  $|10\rangle$  as the identity, and for the last basis state it yields  $CR(\alpha)|11\rangle = e^{i\alpha}|11\rangle$ .

measurement basis of the detector and the possible paths of the interferometer must be established. Moreover, the measurement basis should be selected to maximize the probability of successful inference. Since the prior probability of preparing the  $|0\rangle$  and  $|1\rangle$  states of the interferometer, and hence the  $\rho_0 = |\delta_0\rangle\langle\delta_0|$  and  $\rho_1 = |\delta_1\rangle\langle\delta_1|$  states of the detector, is the same, we need to look for positive semidefinite probability operators  $E_0$  and  $E_1 = \mathbb{1} - E_0$  that maximize the success probability

$$P_{\text{succ}} = \frac{1}{2}\text{Tr}(E_0\rho_0) + \frac{1}{2}\text{Tr}(E_1\rho_1) \quad (3.7)$$

$$= \frac{1}{2} + \frac{1}{2}\text{Tr}[E_0(\rho_0 - \rho_1)]. \quad (3.8)$$

The first term is constant, thus, maximizing this probability reduces to maximizing the second term.

Since  $\rho_0$  and  $\rho_1$  are Hermitian, positive semidefinite density matrices with  $\text{Tr}(\rho_0) = \text{Tr}(\rho_1) = 1$ , their difference is traceless, i.e.  $\text{Tr}(\Delta) = \text{Tr}(\rho_0 - \rho_1) = 1 - 1 = 0$ . This implies that the eigenvalues  $\{\lambda_-, \lambda_+\}$  of  $\Delta$  are real and must sum to zero, ensuring they have opposite signs. Taking advantage of the diagonalization of  $\Delta$  in its eigenbasis,  $\Delta = \sum_{i \in \{-, +\}} \lambda_i |\Lambda_i\rangle\langle\Lambda_i|$ , where  $|\Lambda_i\rangle$  is the eigenvector corresponding to  $\lambda_i$ , we use the basis invariance of the trace to reduce the second term above to the expectation value of  $E_0$  with respect to this eigenbasis:

$$\text{Tr}[E_0(\rho_0 - \rho_1)] = \lambda_+ \langle\Lambda_+|E_0|\Lambda_+\rangle + \lambda_- \langle\Lambda_-|E_0|\Lambda_-\rangle. \quad (3.9)$$

Therefore, to maximize the success probability,  $E_0$  should project onto the subspace corresponding to the positive eigenvalue [271], that is, the optimal  $E_0$  is simply  $E_0 = |\Lambda_+\rangle\langle\Lambda_+|$  and the maximum success probability becomes

$$P_{\text{succ}}^{\text{max}} = \frac{1}{2} + \frac{1}{2}\lambda_+. \quad (3.10)$$

For single-qubit states, the trace norm  $\|\rho_0 - \rho_1\|_1$  is defined as  $\|\rho_0 - \rho_1\|_1 = |\lambda_+| + |\lambda_-| = 2|\lambda_+|$  because  $\lambda_- = -\lambda_+$ . Thus  $\lambda_+ = \frac{1}{2}\|\rho_0 - \rho_1\|_1$  and the maximum success probability becomes

$$P_{\text{succ}}^{\text{max}} = \frac{1}{2}(1 + \mathcal{D}), \quad (3.11)$$

where we introduced the definition of the *distinguishability*  $\mathcal{D}$  in terms of the trace norm as

$$\mathcal{D} \equiv \frac{1}{2}\|\rho_0 - \rho_1\|_1 \quad (3.12)$$

$$= \sqrt{1 - |\langle\delta_0|\delta_1\rangle|^2} = \sqrt{\sin^2 \frac{\alpha}{2}}, \quad (3.13)$$

where we used

$$\langle \delta_0 | \delta_1 \rangle = \frac{1}{2} (1 + e^{i\alpha}). \quad (3.14)$$

With  $0 \leq \mathcal{D} \leq 1$ , the distinguishability of the single-qubit states  $\rho_0$  and  $\rho_1$  is quantified.  $\mathcal{D} = 0$ , occurring when  $\alpha = 0$ , corresponds to completely indistinguishable states ( $\rho_0 = \rho_1$ ).  $\mathcal{D} = 1$ , for  $\alpha = \pi$ , corresponds to perfectly distinguishable states ( $\rho_0$  and  $\rho_1$  have orthogonal supports). This result demonstrates that the distinguishability  $\mathcal{D}$  directly determines the maximum achievable probability of successfully inferring the state, providing an operational and geometric interpretation of  $\mathcal{D}$ .

To implement the projective measurement characterized by  $E_0$  and  $E_1$  in our circuit, we need to perform a basis transformation that allows the measurement to be carried out in the computational basis. This requires identifying a matrix  $O_\alpha$  that maps the computational basis states  $\{|0\rangle, |1\rangle\}$  to the eigenstates  $\{|\Lambda_+\rangle, |\Lambda_-\rangle\}$ . To construct  $O_\alpha$ , we take the eigenstates  $|\Lambda_+\rangle$  and  $|\Lambda_-\rangle$  in the computational basis and use them as columns of the matrix:

$$O_\alpha = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -ie^{\frac{i\alpha}{2}} & ie^{\frac{i\alpha}{2}} \end{pmatrix}. \quad (3.15)$$

The matrix  $O_\alpha$  transforms the computational basis into the eigenbasis of  $\Delta$ , i.e.,  $|\Lambda_+\rangle = O_\alpha |0\rangle$  and  $|\Lambda_-\rangle = O_\alpha |1\rangle$ . Its Hermitian conjugate  $O_\alpha^\dagger$  performs the reverse transformation, mapping the eigenbasis of  $\Delta$  back to the computational basis, i.e.,  $|0\rangle = O_\alpha^\dagger |\Lambda_+\rangle$  and  $|1\rangle = O_\alpha^\dagger |\Lambda_-\rangle$ . To perform the measurement, the quantum state  $|\delta_x\rangle$  is first transformed into the eigenbasis of  $\Delta$  using  $O_\alpha^\dagger$ . The probability of obtaining the outcome associated with  $E_0$  and  $E_1$  is then:

$$P(|0\rangle) = \langle \delta_x | E_0 | \delta_x \rangle = |\langle \Lambda_+ | \delta_x \rangle|^2 = \left| \langle 0 | O_\alpha^\dagger | \delta_x \rangle \right|^2. \quad (3.16)$$

$$P(|1\rangle) = \langle \delta_x | E_1 | \delta_x \rangle = |\langle \Lambda_- | \delta_x \rangle|^2 = \left| \langle 1 | O_\alpha^\dagger | \delta_x \rangle \right|^2. \quad (3.17)$$

Thus, by setting  $P_d^\dagger = O_\alpha^\dagger$  in the circuit depicted in Fig. 3.1b, the reduced density matrix for  $q_d$  takes the form  $\rho'_x = O_\alpha |\delta_x\rangle \langle \delta_x| O_\alpha^\dagger$ . When  $q_d$  is measured in the computational basis, it outputs the state associated with the path  $|x\rangle$  on  $q_i$  with the highest probability of success, given by

$$P_{\text{succ}}^{\text{max}} = \frac{1}{2} \text{Tr}(|0\rangle \langle 0| \rho'_0) + \frac{1}{2} \text{Tr}(|1\rangle \langle 1| \rho'_1) = \frac{1}{2} (1 + \mathcal{D}). \quad (3.18)$$

To empirically determine  $\mathcal{D}$ , the circuit in Fig. 3.1b can be executed with  $P_i^\dagger = \mathbb{1}$  and  $P_d^\dagger = O_\alpha^\dagger$ , to confirm the success frequency of the detector by measuring the interferometer directly in the computa-

tional basis. In this setup,  $\mathcal{D}$  can be inferred from:

$$P_{\text{succ}}^{\text{max}} = \frac{1}{2}P(|0\rangle_{\text{d}}|0\rangle_{\text{i}}) + \frac{1}{2}P(|1\rangle_{\text{d}}|1\rangle_{\text{i}}), \quad (3.19)$$

where the conditional probability is defined as  $P(|y\rangle_{\text{d}}|x\rangle_{\text{i}}) = P(|xy\rangle) / \sum_y P(|xy\rangle)$ , with  $P(|xy\rangle)$  representing the probability of collapsing to the state  $|xy\rangle$  at the output measurement of the circuit.

### 3.2.3 Interference–detection tradeoff

When the circuit in Fig. 3.1b is executed with  $P_{\text{i}}^{\dagger} = H$ , the interference pattern observed in the readouts of the interferometer is determined from the expectation value

$$\langle X \rangle \equiv \text{Tr}[\sigma_{\text{qi}} X] \quad (3.20)$$

$$= \frac{\cos \phi + \cos(\phi + \alpha)}{2}, \quad (3.21)$$

where  $\sigma_{\text{qi}}$  is the reduced density matrix associated with  $q_{\text{i}}$  prior to the application of the basis change operation. This reduced density matrix is derived by tracing out the degrees of freedom of  $q_{\text{d}}$  from the joint two-qubit state defined in Eq. (3.5), such that

$$\sigma_{\text{qi}} \equiv \text{Tr}_{\text{qd}}[|\Psi\rangle\langle\Psi|] \quad (3.22)$$

$$= \sum_{j=0}^1 (\mathbb{1} \otimes \langle j|) |\Psi\rangle\langle\Psi| (\mathbb{1} \otimes |j\rangle). \quad (3.23)$$

From Eq. (3.21), it becomes evident that the interference contrast as a function of  $\phi$  is determined by the overlap between the potential states of the detector qubit. For maximum overlap, at  $\alpha = 0$ , the interference visibility is highest, and no which-path information is retained. Conversely, when  $\alpha = \pi$  the states are perfectly distinguishable since the overlap vanishes entirely, leading to the complete disappearance of interference. At this point, the detector encodes maximal which-path information, enabling perfect path determination on the first qubit. Interestingly, it is precisely at the condition of maximal entanglement between the qubits, a distinctly quantum phenomenon, that the system exhibits purely classical behavior with particlelike statistics, characterized by a complete absence of interference, regardless of whether the detector is measured.

To quantify how varying  $\alpha$  from zero to  $\pi$  gradually vanishes the amplitude of the interference pattern,

we define the interference visibility below and compute it from the result in Eq. (3.21):

$$\mathcal{V} \equiv \frac{\langle X \rangle_{\alpha}^{\max}(\phi) - \langle X \rangle_{\alpha}^{\min}(\phi)}{2 + \langle X \rangle_{\alpha}^{\max}(\phi) + \langle X \rangle_{\alpha}^{\min}(\phi)} \quad (3.24)$$

$$= \sqrt{\cos^2\left(\frac{\alpha}{2}\right)}. \quad (3.25)$$

Using the expression for overlap given in Eq. (3.14), we see that the visibility is directly linked to

$$\mathcal{V} = |\langle \delta_0 | \delta_1 \rangle|, \quad (3.26)$$

and consequently satisfies

$$\mathcal{V}^2 + \mathcal{D}^2 = 1. \quad (3.27)$$

This relation quantifies the principle of complementarity between visibility and distinguishability: as which-path information becomes more distinct, interference visibility disappears [238, 239].

### 3.2.4 Quantum erasure

As detailed earlier, the act of encoding which-path information onto the detector qubit fundamentally disrupts the interference pattern observed in the interferometer. This disruption occurs regardless of whether the detector qubit is directly measured, as the mere entanglement between the two qubits encodes path distinguishability. However, under certain conditions, the lost interference can be recovered by employing a quantum erasure protocol to irreversibly destroy which-path information.

Consider the circuit in Fig. 3.1b, with  $P_d^{\dagger} = \mathbb{1}$  and  $P_i^{\dagger} = H$ . In this configuration, the detector qubit is measured in the computational basis, collapsing its state into either  $|0\rangle$  or  $|1\rangle$ . This measurement eliminates the which-path information stored in the detector qubit because it projects the two possible detector states with equal probability into the measurement basis, effectively erasing any correlation that could distinguish the paths. As a result, the interference pattern on the first qubit can re-emerge, a phenomenon known as quantum erasure [272, 249, 273].

To analyze this process, we start from the two-qubit state just before the detector qubit is measured. The output state of the system can be written as:

$$|\Psi_{\text{out}}\rangle = \frac{1 + e^{i\phi}}{2\sqrt{2}} |00\rangle + \frac{1 + e^{i\phi+i\alpha}}{2\sqrt{2}} |01\rangle + \frac{1 - e^{i\phi}}{2\sqrt{2}} |10\rangle + \frac{1 - e^{i\phi+i\alpha}}{2\sqrt{2}} |11\rangle. \quad (3.28)$$

The probabilities of observing each computational basis state,  $|xy\rangle$ , where  $x$  and  $y$  denote the states of the

interferometer and detector qubits respectively, are computed with  $P(|xy\rangle) = |\langle xy | \Psi_{\text{out}} \rangle|^2$ , yielding:

$$\begin{aligned} P(|00\rangle) &= \frac{1}{4}(1 + \cos \phi), \\ P(|01\rangle) &= \frac{1}{4}(1 + \cos(\phi + \alpha)), \\ P(|10\rangle) &= \frac{1}{4}(1 - \cos \phi), \\ P(|11\rangle) &= \frac{1}{4}(1 - \cos(\phi + \alpha)). \end{aligned} \quad (3.29)$$

Conditioning on the measurement outcome of the detector qubit, the state of the interferometer qubit collapses from a statistical mixture to a specific superposition, with the probabilities of each path determined by the corresponding detector outcome. By analyzing the expectation value of the interferometer, broken down guided by the detector, two full-visibility ( $\mathcal{V} = 1$ ) interference patterns emerge, with a dependence on the phase  $\phi$  shifted by  $\alpha$ :

$$\begin{aligned} \langle X_0 \rangle &\equiv P(|0\rangle_i | 0\rangle_d) - P(|1\rangle_i | 0\rangle_d) = \frac{P(|00\rangle) - P(|10\rangle)}{P(|00\rangle) + P(|10\rangle)} \\ &= \cos \phi. \end{aligned} \quad (3.30)$$

$$\begin{aligned} \langle X_1 \rangle &\equiv P(|0\rangle_i | 1\rangle_d) - P(|1\rangle_i | 1\rangle_d) = \frac{P(|01\rangle) - P(|11\rangle)}{P(|01\rangle) + P(|11\rangle)} \\ &= \cos(\phi + \alpha). \end{aligned} \quad (3.31)$$

Clearly, aggregating the measurements from these two interference patterns, ignoring therefore any guidance of the detector, recovers the interference pattern in Eq. (3.21), that is  $\frac{1}{2}\langle X_0 \rangle + \frac{1}{2}\langle X_1 \rangle = \langle X \rangle$ .

It is crucial to observe that applying the same decomposition procedure using  $P_d^\dagger = O_\alpha^\dagger$  does not yield interference patterns with full visibility. Such measurement inherently disrupts the phase coherence required for full interference, underscoring the fundamental tradeoff between wavelike statistics and path information gained through measurement. Specifically, with the optimal measurement, the interference patterns for  $\langle X_0 \rangle$  and  $\langle X_1 \rangle$  are both equal to the aggregated pattern  $\langle X \rangle$  in Eq. (3.21).

### 3.3 Experimental results

We now present the methodology and results of experimentally testing the circuits described in the previous section. This was performed using the IBM Quantum platform, which provides access to gate-

based quantum processors made from fixed-frequency transmon qubits coupled via the cross-resonance gate [274]. Making use of the Qiskit open-source software framework [275], we designed and remotely executed experiments on two qubits from *ibmq\_toronto*, a 27-qubit *Falcon r4* processor with a heavy-hexagon lattice architecture.

The transmon qubit architecture comprises two superconducting islands linked through a pair of Josephson junctions (JJs). Key dynamical variables of the JJs are the electron imbalance and the phase difference of the superconducting order parameter, effectively forming conjugate variables. The two lowest energy eigenstates of the transmon Hamiltonian define the qubit states  $|0\rangle$  and  $|1\rangle$ . Control and readout are realized through capacitive coupling to electromagnetic modes in a transmission line resonator. For detailed explanations of the device's principles, we refer the reader to foundational studies [276, 277, 278] and comprehensive reviews [279, 280].

### 3.3.1 Methods

The experimental process for separately testing each of the circuits in Fig. 3.1 involved two primary steps. The first step consisted of estimating the relevant observables for  $C$  configurations of the circuit's variables, using  $S$  shots per configuration. With a repetition rate of  $500 \mu s$  per shot, the total runtime of an experiment is  $500CS \mu s$ . All these circuits, comprising an experiment, were divided into batches of 896. Additionally, each batch included four error mitigation circuits [281, 282, 283], except for the experiment involving the circuit in Fig. 3.1a. For each batch, every circuit was executed once sequentially before repeating the process for the remaining shots. Batches were processed consecutively until all necessary executions were complete. Variables  $\phi$  and  $\alpha$  fix the exact circuit to be executed within each circuit *family*; therefore, the order in which this parameter space is probed is defined by the order in which each fixed circuit is executed; for reasons that will be explained below, to facilitate the subsequent step of the methodology, this space was sampled uniformly at random.

In the second step, the  $C$  measured values were used to determine a global experimental model  $f(\phi, \alpha)$  by minimizing the chi-squared statistic:

$$\chi^2 = \sum_{j=1}^C \left( \frac{z_j - \hat{z}_j}{\sigma_j} \right)^2, \quad (3.32)$$

where  $\hat{z}_j = f((\phi, \alpha)_j)$ , and  $\sigma_j$  is the standard error for each estimate  $\hat{z}_j$ . Although measurement error mitigation was applied later during exploratory analysis of the experimental data to identify the primary sources of error, it is not included in the final analysis of results; all the empirical models presented below




Figure 3.2: **Physical implementation of the which-path circuits.** This circuit adapts Fig. 3.1b to the native gateset of the device. The final gates  $G_i$  and  $G_d$  are fixed depending on the readout basis of each family of circuits, following the correspondence:  $P_i^\dagger = H \Rightarrow G_i = U_2(0, \pi)$ ,  $P_i^\dagger = \mathbb{1} \Rightarrow G_i = \mathbb{1}$ ,  $P_d^\dagger = O_\alpha^\dagger \Rightarrow G_d = U_2(0, 3\pi/2)$ , and  $P_d^\dagger = \mathbb{1} \Rightarrow G_d = U_1(\alpha/2)$ . Measurement of both qubits occurs simultaneously. Reproduced from [76]. © 2021 American Physical Society.

were fitted to the unmitigated raw data.

To construct the model  $f(\phi, \alpha)$  for data fitting, the theoretical prediction  $g(\phi, \alpha)$  from ideal unitary circuit execution was modified to include additional parameters. First, calibration errors in the physical gates were addressed by introducing parameters  $\Theta_1$  and  $\Theta_2$ :

$$\begin{aligned}\phi &\rightarrow \phi + \Theta_1, \\ \alpha &\rightarrow \alpha + \Theta_2.\end{aligned}\tag{3.33}$$

Second, to account for nonunitary dynamics induced by environmental noise, the pure-state expression was scaled by  $\eta$ , and a bias parameter  $\varepsilon$  was added. The resulting model is:

$$f(\phi, \alpha) = \eta g(\phi, \alpha) + \varepsilon.\tag{3.34}$$

After fitting the model, the residuals  $z_j - \hat{z}_j$  were inspected for autocorrelation in the circuit execution order. Autocorrelation, if present, is attributed to drift in device calibration or coherence during the experiment. This identification is enabled by random sampling of  $(\phi, \alpha)$ , which eliminates potential correlations due to using an incorrect theoretical model to fit the data. To detect autocorrelation, the nonparametric Wald-Wolfowitz runs test [284] was applied to the sequence of residuals, divided into positive and negative deviations from the mean. Rejection of the null hypothesis indicates non-independence of errors, that is, evidence for autocorrelation. Residual distributions were also evaluated in the  $(\phi, \alpha)$  space for patterns suggesting model inadequacy.

All circuits in an experiment were executed on the same pair of qubits during a single calibration cycle of the machine to ensure consistency of calibration errors across the variable space. The approach of random sampling the  $(\phi, \alpha)$  parameters with each circuit execution provided an additional advantage: even if calibration shifted before all variable values were sampled, interrupting the experiment, the model could still be fit to the data already obtained, uniformly distributed in the parameter space. In addition, this

protocol mitigates the impact of small coherence drifts by enabling uniform outlier filtering in the  $(\phi, \alpha)$  space.

At the time of the experiment, the hardware supported controlled-NOT (CNOT) gates and three parameterized single-qubit gates. Thus, the circuits in Fig. 3.1 were transpiled into the native operations of the device as shown in Fig. 3.2. The used native single-qubit gates were:

$$\begin{aligned} U_1(y) &= \begin{pmatrix} 1 & 0 \\ 0 & e^{iy} \end{pmatrix}, \\ U_2(x, y) &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{iy} \\ e^{ix} & e^{i(x+y)} \end{pmatrix}. \end{aligned} \quad (3.35)$$

Experimental simplifications based on specific  $\phi$  and  $\alpha$  values were avoided to ensure uniformity in the number of operations and execution time across all circuits, maintaining consistency of the complete experimental dataset with the experimental model.

### 3.3.2 Single-qubit interference circuit

We begin with the experimental testing of the configuration illustrated in Fig. 3.1a. To observe the presence of an interference pattern with this simple single-qubit control circuit, we acquired  $S = 8192$  measurement shots for each of  $C = 401$  uniformly spaced values of  $\phi \in [0, 2\pi]$ . The corresponding experimental results are depicted in Fig. 3.3. At each value of  $\phi_j$ , for  $j = 1, \dots, C$ , a measurement  $X_j \in \{-1, 1\}$  was recorded, and the relative frequencies  $\bar{P}(1)$  and  $\bar{P}(-1)$  of the two possible outcomes were computed. The expectation value  $\langle X \rangle$  could then be straightforwardly estimated using the sample mean

$$\bar{X} = \bar{P}(|0\rangle) - \bar{P}(|1\rangle), \quad (3.36)$$

with the standard error

$$\sigma_{\bar{X}} = s / \sqrt{S}. \quad (3.37)$$

Here, the sample variance  $s^2$  is given by

$$s^2 = \frac{1}{S-1} \sum_{j=1}^S (X_j - \bar{X})^2. \quad (3.38)$$

Using Eq. (3.3), the data was modeled by the expression  $\bar{X} = \eta \cos(\phi + \Theta_1) + \varepsilon$ . With  $S$  and  $C$  set as described, the fitting process yielded a close match between the data and the model, characterized by




Figure 3.3: **Experimental results for the interference circuit in Fig. 3.1a.** The computed  $\bar{X}$  values are presented with nearly negligible  $\pm 3\sigma_{\bar{X}}$  confidence intervals, demonstrating an excellent agreement with the experimental model derived from Eq. (3.34). This model closely resembles the predictions of the pure state model in Eq. (3.3). Reproduced from [76]. © 2021 American Physical Society.

a reduced chi-squared<sup>4</sup> of  $\chi^2_\nu \approx 1.04$  for  $\text{dof} = 398$ <sup>5</sup>, where

$$\chi^2_\nu = \chi^2 / \text{dof}. \quad (3.39)$$

The fitted parameters and their one-standard-deviation uncertainties were determined as follows:

$$\begin{aligned} \eta &= 0.98581 \pm 0.00024, \\ \varepsilon &= 0.00519 \pm 0.00022, \\ \Theta_1 &= -0.03219 \pm 0.00060. \end{aligned} \quad (3.40)$$

The residuals were uniformly distributed over the entire range of  $\phi$ , as verified by a runs test. With a  $p$ -value of 0.61, the test did not reject the null hypothesis at a 0.05 significance level, indicating the stability of the system against quantum fluctuations throughout the approximately 30 minutes required to perform the experiment.

Figure 3.3 compares the experimental results with the theoretical prediction in Eq. (3.3). The parameter  $\eta$  indicates a slight reduction in amplitude, likely attributable to readout errors or to the interaction of

<sup>4</sup>If  $\chi^2_\nu \approx 1$ , the model fits well, with residuals consistent with uncertainties. If  $\chi^2_\nu \gg 1$ , the model may underfit or uncertainties are underestimated. If  $\chi^2_\nu \ll 1$ , the model may overfit or uncertainties are overestimated.

<sup>5</sup>The degrees of freedom (dof) in this context are determined by the number of data points minus the number of parameters estimated in the model,  $\text{dof} = 401 - 3 = 398$ . This reflects the number of independent data points remaining after accounting for the parameters estimated in the model.

the system with the environment due to its open dynamics. The phase offset  $\Theta_1$  reflects systematic calibration error. The asymmetry quantified by  $\varepsilon$  aligns with phenomena such as energy dissipation, causing relaxation of the qubit from  $|1\rangle$  to  $|0\rangle$ , and unequal error rates during state readout.

### 3.3.3 Optimal path detection and the visibility-distinguishability tradeoff

This experimental study required executing two distinct types of circuits. The first family employed  $P_1^\dagger = H$  to estimate  $\langle X \rangle$  on  $q_i$ , which allowed for the computation of the visibility  $\mathcal{V}$ . The second family determined the distinguishability  $\mathcal{D}$  using  $P_1^\dagger = \mathbb{1}$ . Both types of circuits included  $P_d^\dagger = O_\alpha^\dagger$  and were implemented for  $S = 8192$  shots over  $C = 10201$  evenly spaced combinations of  $(\phi, \alpha)$  on a  $101 \times 101$  square grid, with  $\phi, \alpha \in [0, 2\pi]$ . This setup resulted in a total of 20402 circuits, divided into 23 batches, which took approximately 23 hours to complete. The sequence of circuits executed within each batch was ordered to alternate between each circuit family, shot by shot, executing shots at the same  $(\phi, \alpha)$  for every two circuits, while sampling this parameter space uniformly at random.

The data analysis was carried out separately for each family of circuits. For the first family,  $\bar{X}$  and  $\sigma_{\bar{X}}$  were computed on the first qubit, following same procedure in Eqs. (3.36 – 3.38). The experimental data were modeled with a four-parameter function approximating Eq. (3.21), achieving a fit with  $\chi_V^2 \approx 2.19$ . The resulting parameters, with their one standard deviation uncertainties, are

$$\begin{aligned} \eta &= 0.93039 \pm 0.00017, & \Theta_1 &= -0.03704 \pm 0.00026, \\ \varepsilon &= 0.03700 \pm 0.00009, & \Theta_2 &= 0.00512 \pm 0.00037. \end{aligned} \quad (3.41)$$

No significant deviations in calibration are observed, as indicated by the small values of  $\Theta_1$  and  $\Theta_2$ . However, the drop of the amplitude  $\eta$  and the bias  $\varepsilon$  show slight increases compared to the one-qubit experiment.

The second family of circuits was used to estimate distinguishability  $\mathcal{D}$  via the relative frequencies of successfully identifying the path on  $q_i$ . Following Eqs. (3.18) and (3.19),  $\mathcal{D}$  was calculated as

$$\mathcal{D} = \frac{\bar{P}(|00\rangle)}{\bar{P}(|00\rangle) + \bar{P}(|01\rangle)} + \frac{\bar{P}(|11\rangle)}{\bar{P}(|10\rangle) + \bar{P}(|11\rangle)} - 1, \quad (3.42)$$

with the associated standard error given by (see Appendix 3.7.1)

$$\sigma_{\mathcal{D}} = \sqrt{\frac{\bar{P}(|00\rangle) \bar{P}(|01\rangle)}{S(\bar{P}(|00\rangle) + \bar{P}(|01\rangle))^3} + \frac{\bar{P}(|10\rangle) \bar{P}(|11\rangle)}{S(\bar{P}(|10\rangle) + \bar{P}(|11\rangle))^3}}. \quad (3.43)$$




Figure 3.4: **Results of the visibility and distinguishability tradeoff experiments.** Two families of circuit based on Fig. 3.1b were executed, each with  $S = 8192$  shots across  $C = 10201$  sampled pairs of  $\alpha, \phi \in [0, 2\pi]$ . The first family corresponds to the interference experiments using  $P_i^\dagger = H$  and  $P_d^\dagger = O_\alpha^\dagger$ , with the top-right panel depicting the observed  $\bar{X}$  on  $q_i$  for all circuits. A subset of data for six values of  $\alpha$  is shown in the left panel, where experimental points include  $\bar{X} \pm 2\sigma_{\bar{X}}$  and are compared to the model fit (dashed lines) parameterized in Eq. (3.41). The bottom right panel presents visibility, distinguishability, and their squared sum. The distinguishability data originates from the second family of circuits, executed with  $P_i^\dagger = \mathbb{1}$  and  $P_d^\dagger = O_\alpha^\dagger$  for the same values of  $S$  and  $C$ , with  $\mathcal{D}$  calculated via Eqs. (3.18 – 3.19). The plotted distinguishability points, in blue, represent averaged values of  $\mathcal{D}$  over all circuits for each  $\alpha$ , which, in theory, are independent of  $\phi$ . Reproduced from [76]. © 2021 American Physical Society.

This distinguishability values, were modeled with a three-parameter fit based on Eq. (3.13), which resulted in  $\chi^2_\nu \approx 1.34$  and the following parameter estimates:

$$\begin{aligned} \eta &= 0.90216 \pm 0.00027, & \Theta_2 &= -0.00224 \pm 0.00029, \\ \varepsilon &= 0.00253 \pm 0.00023. \end{aligned} \tag{3.44}$$

Although  $\eta$  exhibits a significant reduction, consistent with additional random noise, the bias  $\varepsilon$  in this case lacks a straightforward explanation. Calibration shifts remain minimal.

In the bottom right panel of Fig. 3.4, experimental data are plotted for  $\mathcal{V}$  and  $\mathcal{D}$ . While the fitted visibility curve – obtained from the model with parameters in Eq. (3.41) – approaches zero by design, the raw experimental data reveal a minimum value of  $\mathcal{V}_{\min} = 0.030 \pm 0.004$  at  $\alpha = \pi$ , and a maximum




Figure 3.5: **Residual errors from the second experiment.** The left panel displays  $z_j - \hat{z}_j$  for the  $\bar{X}$  model associated with the first circuit family, while the right panel shows these residuals for the  $\bar{\mathcal{D}}$  model from the second circuit family. The colorbars indicate the minimum and maximum recorded values. Reproduced from [76]. © 2021 American Physical Society.

of  $\mathcal{V}_{\max} = 0.917 \pm 0.004$  at  $\alpha = 1.94\pi$ , where the standard deviation error is obtained with 10000 bootstrap samples. The distinguishability data points represent the averaged values of  $\bar{\mathcal{D}}$  over all circuits for each  $\alpha$ . Notably, the relation in Eq. (3.27) is modified to  $\mathcal{V}^2 + \mathcal{D}^2 < 1$  due to information loss arising from random noise mixed with the final pure state distributions, which lowers the amplitudes of Eq. (3.13) and Eq. (3.25). The largest part of such noise likely arises from readout errors, given how shallow this circuit is, reducing the role of decoherence in the loss of quantum information. This is further corroborated by the observation that applying the error mitigation protocol restores  $\eta \approx 1$  in both models.

The runs test performed on the residual errors in the order of execution yielded  $p$ -values close to zero, rejecting the null hypothesis at the 0.05 significance level and confirming some amount of autocorrelation in the data. This suggests there was some instability in the quantum processor during the experiment. Despite this, both the model in Eq. (3.41) and in Eq. (3.44) achieve strong fits, with  $R^2 \approx 0.99926$  and

0.99888, and residual standard errors<sup>6</sup>

$$\text{RSE} \equiv \sqrt{\frac{1}{\text{dof}} \sum_{j=1}^C (z_j - \hat{z}_j)^2}, \quad (3.45)$$

of 0.01274 and 0.00946, respectively. More significantly, residual error maps, that is  $z_j - \hat{z}_j$  as a function of  $(\phi, \alpha)$ , for both circuit families are presented in Fig. 3.5. These show structured patterns for the first family of circuits and less pronounced, but still visible, deviations for the second. Appropriate fits would produce a random distribution of residual errors within this space, but these systematic deviations from theoretical predictions suggest missing dependencies on  $(\phi, \alpha)$  in the experimental models. These anomalies are also evidenced by  $\chi_V^2 \approx 2.19$  and  $\chi_V^2 \approx 1.34$ , for the first and second models, respectively. A refined theoretical framework for the entangling operation, discussed in Sec. 3.5, addresses these issues.

### 3.3.4 Quantum Eraser

We now turn to the third experiment, employing the circuit in Fig. 3.1b with  $P_i^\dagger = H$  and  $P_d^\dagger = \mathbb{1}$ . Similar to the prior section, we used  $S = 8192$  shots per circuit and evaluated  $C = 10201$  pairs of  $(\phi, \alpha)$  parameter values sampled on a regular  $101 \times 101$  square grid within  $\phi, \alpha \in [0, 2\pi]$ . The execution of these 10201 circuits was distributed over 12 batches, requiring approximately 12 hours to complete.

The objective here is to demonstrate that measuring  $q_d$  in a projection that does not extract the accumulated which-path information nullifies path discrimination, thereby restoring full interference visibility on  $q_i$ . To achieve this, we splitted the readouts on the interferometer into two groups based on the output ket  $|y\rangle \in \{|0\rangle, |1\rangle\}$  of the detector. This allowed us to estimate  $\langle X_y \rangle$  as defined in Eqs. (3.30) and (3.31) from

$$\overline{X}_y = \frac{\overline{P}(|0y\rangle) - \overline{P}(|1y\rangle)}{\overline{P}(|0y\rangle) + \overline{P}(|1y\rangle)}. \quad (3.46)$$

The corresponding standard error,  $\sigma_{\overline{X}_y}$ , was calculated using the same procedure from Eqs. (3.36 – 3.38), after segmenting the outcomes of the first qubit according to the observed states of the second qubit. It

---

<sup>6</sup>The Residual Standard Error (RSE) quantifies the variability of observed data around a fitted regression model by summarizing the average magnitude of residuals, expressed in the same units as the dependent variable. A smaller RSE generally indicates a better fit. However, while the RSE measures the overall “lack of fit,” it does not detect systematic patterns in residuals relative to explanatory variables or fitted values. Identifying such patterns requires complementary diagnostic tools, such as residual plots or tests for autocorrelation and heteroscedasticity.




Figure 3.6: **Experimental results for the quantum eraser circuit.** Expectation values of the  $X$  operator on the first qubit are measured conditioned on the detector outcome:  $\bar{X}_0 \pm 2\sigma_{\bar{X}_0}$  (top row) and  $\bar{X}_1 \pm 2\sigma_{\bar{X}_1}$  (bottom row). Right panels depict results across all  $(\phi, \alpha)$  values, while left panels focus on six specific  $\alpha$ . Reproduced from [76]. © 2021 American Physical Society.

is straightforward to demonstrate that this result can be expressed as:

$$\sigma_{\bar{X}_y} = \sqrt{\frac{4 \bar{P}(|0y\rangle) \bar{P}(|1y\rangle)}{(S-1) (\bar{P}(|0y\rangle) + \bar{P}(|1y\rangle))^3}}, \quad (3.47)$$

where, for large  $S$ , the denominator can be simplified with  $(S-1) \approx S$ . For example, with  $S = 8192$ , such approximation equals the exact standard error of the sample to within a factor of approximately 0.999939, i.e.,  $\tilde{\sigma}_{\bar{X}_y} \approx 0.999939 \sigma_{\bar{X}_y}$ .

The results are presented in Fig. 3.6. The top and bottom rows show  $\bar{X}_0$  and  $\bar{X}_1$ , respectively. Theoretically, both are expected to exhibit interference patterns with full visibility, and the data qualitatively align with this prediction. To quantitatively analyze the results, we fitted  $\bar{X}_0$  to the experimental model

derived from Eq. (3.30), yielding

$$\begin{aligned}\eta &= 0.91740 \pm 0.00011, & \Theta_1 &= -0.03118 \pm 0.00020, \\ \varepsilon &= 0.03361 \pm 0.00009, & & \end{aligned} \quad (3.48)$$

with  $\chi^2_v \approx 8.55$  and  $\text{RSE} = 0.03083$ . For  $\overline{X_1}$ , we fitted the data using the model from Eq. (3.31), obtaining  $\chi^2_v \approx 2.49$ ,  $\text{RSE} = 0.01885$ , and the following parameters:

$$\begin{aligned}\eta &= 0.93765 \pm 0.00011, & \Theta_1 &= -0.00565 \pm 0.00010, \\ \varepsilon &= 0.03379 \pm 0.00009, & \Theta_2 &= -0.00565 \pm 0.00010. \end{aligned} \quad (3.49)$$

The precision with which these parameters are determined is noteworthy in both cases. Additionally, runs tests on the ordered sequence of residuals for each fit separately failed to reject the null hypothesis of randomness at a 0.05 significance level, indicating stable device performance throughout this experiment.

However, the  $\chi^2_v$  values deviate significantly from unity, signaling potential issues. Inspecting the residual plots as a function of  $(\phi, \alpha)$  for both  $\overline{X_0}$  and  $\overline{X_1}$  exhibits again systematic patterns, most evident for  $\overline{X_0}$ . This missing functional dependence can be seen, for instance, in the inset plot on the top-left panel of Fig. 3.6. Notably, a weak dependence on  $\alpha$  emerges, with interference visibility slightly diminishing as  $\alpha$  approaches  $\pi$ . This suggests that some error in the controlled- $R_\alpha$  gate may cause minor leakage of which-path information to the detector qubit, partially undermining interference visibility on  $q_i$  due to the inability of the computational basis readout to erase this information. A detailed analysis of this issue is deferred to Sec. 3.5.

To evaluate the performance of the detector in distinguishing paths through the interferometer, we performed an additional independent experiment<sup>7</sup>. We considered the computational basis readout for both qubits, i.e. setting  $P_i^\dagger = P_d^\dagger = \mathbb{1}$  in Fig. 3.1b, which were implemented with  $G_i = \mathbb{1}$  and  $G_d = U_1(\frac{\alpha}{2})$  in Fig. 3.2. For the same grid of  $C = 10201$  pairs of  $(\phi, \alpha)$  values and  $S = 8192$  shots, we measured distinguishability for each circuit via Eqs. (3.42) and (3.43). When the detector qubit is readout directly on the states from Eq. (3.6), theoretical prediction in Eq. (3.19) dictates that  $p_{\text{succ}}$  is expected to be 0.5, and the *measured* distinguishability [240] should be

$$\mathcal{D}_m = 0, \quad (3.50)$$

---

<sup>7</sup>The fourth experiment.

independent of  $(\phi, \alpha)$ . This contrasts with the results in the previous section where the optimal measurement basis aligns measured and theoretical distinguishability  $\mathcal{D}$  in Eq. (3.12).

We then fitted the experimental data to a simple model derived from the theoretical expectation that  $\mathcal{D}_m = 0$ , which involved only a constant parameter  $\varepsilon$ , according to Eq. (3.34). This yielded  $\varepsilon = -0.00370 \pm 0.00011$ ,  $\chi^2_\nu \approx 2.06$ , and  $\text{RSE} = 0.01581$ . The analysis of the spacial correlations of the residuals again revealed smooth patterns across  $(\phi, \alpha)$ , confirming incomplete erasure of quantum information and a failure to produce entirely wavelike behavior. The result is robust, as the runs tests on the residuals in chronological order yielded a  $p$ -value of 0.96, failing to reject the null hypothesis at a significance level of 0.05 and indicating no significant environmental disturbances during the experiment.

### 3.4 Accounting for coherent errors

As observed in the earlier two-qubit experiments, discrepancies between the experimental results and the theoretical predictions presented in § 3.2 were consistently observed. These deviations could not be fully accounted for through measurement error mitigation, suggesting they originate from systematic factors unrelated with the readout operation. To refine our theoretical models and better align them with experimental data, we now consider the influence of hardware-level gate imperfections on the expected output of the circuits.

Revisiting the physical circuit from Fig. 3.2, we note that all single-qubit gates, except those parameterized by  $\phi$  and  $\alpha$ , have fixed input parameters. Depending on the chosen configurations of  $G_i$  and  $G_d$ , these circuits contain both  $U_1$  and  $U_2$  gates applied to the bottom qubit, but only  $U_2$  gates acting on the top qubit. To model these systematic single-qubit gate errors (SQGEs), we introduce bias parameters into the rotation angles. Specifically, we modify the  $U_2$  gate acting on the interferometer qubit as follows:

$$U_2^{(\text{qi})}(x, y) \equiv U_2(x + \theta_1, y + \theta_2), \quad (3.51)$$

and introduce independent bias parameters for the gates acting on the detector qubit:

$$U_2^{(\text{qd})}(x, y) \equiv U_2(x + \theta_3, y + \theta_4), \quad (3.52)$$

$$U_1^{(\text{qd})}(y) \equiv U_1(y + \theta_5). \quad (3.53)$$

By incorporating these modified gate models, with five additional bias parameters, into the circuits of Fig. 3.2, we recalculate the expressions for output quantities of interest. These include Eq. (3.13) — derived

from Eq. (3.18) and Eq. (3.19) — as well as Eqs. (3.21), (3.30), (3.31), and (3.50). The resulting expressions are:

$$\mathcal{D} = \sqrt{\cos^2(\theta_3 + \theta_4) \sin^2\left(\frac{\alpha}{2} - \theta_5\right)}, \quad (3.54)$$

$$\langle X \rangle = \frac{1}{2} [\cos(\alpha + \phi + \theta_1 + \theta_2 - \theta_5) + \cos(\phi + \theta_1 + \theta_2 + \theta_5)], \quad (3.55)$$

$$\langle X_0 \rangle = \cos(\phi + \theta_1 + \theta_2 + \theta_5), \quad (3.56)$$

$$\langle X_1 \rangle = \cos(\alpha + \phi + \theta_1 + \theta_2 - \theta_5), \quad (3.57)$$

$$\mathcal{D}_m = 0. \quad (3.58)$$

These modified models provide a theoretical underpinning for the empirically introduced parameters  $\Theta_{1,2}$  in Eq. (3.33). Errors modeled by  $\theta_3$  and  $\theta_4$  reduce the maximum achievable value of  $\mathcal{D}$ , while biases due to  $\theta_1$ ,  $\theta_2$ , and  $\theta_5$  only shift the  $\langle X \rangle$  interference pattern without reducing the visibility  $\mathcal{V}$ . However, these SQGEs fail to explain experimental observations such as the  $\alpha$ -dependence of the  $\langle X_0 \rangle$  interference contrast and the variability of  $\mathcal{D}_m$  across different  $\phi$  and  $\alpha$  values.

To address these unresolved discrepancies, we note that while single-qubit gates on state-of-the-art quantum hardware achieve error rates as low as  $10^{-4}$  [274], two-qubit gates remain a major source of error. At the time of these experiments, the CNOT gate, implemented via the cross-resonance (CR) gate combined with single-qubit rotations, exhibits error rates in the range of  $10^{-2}$  to  $10^{-3}$ . Efforts to improve CR gate fidelity are an active area of research, which include designing optimized pulse sequences to suppress coherent error terms and generate only a  $ZX$  interaction. For example, an echoed CR sequence with target rotary pulses has recently demonstrated improved performance by minimizing undesired interactions in the effective Hamiltonian [285]. Nevertheless, residual two-qubit error terms still add a systematic bias to the interaction. The CNOT gates used in our experiments rely on this implementation.

To incorporate these systematic errors, we next introduce a biased model for the CNOT gate. This model conveniently parameterizes the coupling strengths of unwanted error terms relative to the intended  $ZX$  interaction with dimensionless bias ratios  $\beta_j$  and can be used to analyse the effect of these biased CNOT gates on arbitrary quantum circuits.

### 3.4.1 A biased-CNOT gate model

To construct a model for the biased CNOT gate representing the physical CNOT operation implemented in the quantum processor, we restrict our analysis to the two-qubit subspace of the effective cross-

resonance (CR) pulse Hamiltonian. Any entanglement with spectator qubits is neglected. The basis for this model is the  $ZX$  interaction Hamiltonian, whose propagator is defined as  $U_{ZX}(\theta) = \exp(-i\frac{\theta}{2}ZX)$ . When  $\theta = \pi/2$ , this propagator becomes:

$$U_{ZX}\left(\frac{\pi}{2}\right) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i & 0 & 0 \\ -i & 1 & 0 & 0 \\ 0 & 0 & 1 & i \\ 0 & 0 & i & 1 \end{pmatrix} \equiv U_{ZX;\frac{\pi}{2}}, \quad (3.59)$$

which can maximally entangle a product state such as  $(|00\rangle + |10\rangle)/\sqrt{2}$  and is locally equivalent to a CNOT gate. Although the exact single-qubit rotations required to complete the CNOT gate from  $U_{ZX;\frac{\pi}{2}}$  may be unknown to the user of the IBM Quantum platform, these can be encapsulated in a two-qubit unitary  $M$ , derived by solving  $\text{CNOT} = U_{ZX;\frac{\pi}{2}} \cdot M$ . This yields:

$$M = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i & 0 & 0 \\ i & 1 & 0 & 0 \\ 0 & 0 & -i & 1 \\ 0 & 0 & 1 & -i \end{pmatrix}, \quad (3.60)$$

which commutes with  $U_{ZX;\frac{\pi}{2}}$ , i.e.,  $[M, U_{ZX;\frac{\pi}{2}}] = 0$ . This operator can be decomposed as:

$$M = S^\dagger \otimes R_x\left(-\frac{\pi}{2}\right), \quad (3.61)$$

where  $S^\dagger$  and  $R_x(-\pi/2)$  single-qubit gates are applied to the control and target qubits, respectively, and the  $X$  rotation is defined as  $R_x(\theta) = \exp(-i\theta X/2)$ .

The CR gate, used to realize an effective  $ZX$  coupling in the transmon platform, achieves this by applying a microwave drive on the control qubit at the resonance frequency of the target qubit. This interaction induces state-dependent rotations on the target qubit, which have been previously modeled by an effective Hamiltonian in the Pauli basis [286, 287]. Advanced pulse sequences, such as echoed CR combined with target rotary pulses, are designed to minimize unwanted error terms, improving fidelity [285]. The resulting effective interaction Hamiltonian can be written as:

$$H_{\text{eff}} = \tilde{v}_{ZX} \frac{ZX}{2} + H_{\text{err1}} + H_{\text{err2}}, \quad (3.62)$$

where the primary term  $\tilde{v}_{ZX}$  corresponds to the desired  $ZX$  interaction, and the error contributions are:

$$H_{\text{err}_1} = \tilde{v}_{IY} \frac{IY}{2} + \tilde{v}_{IZ} \frac{IZ}{2}, \quad (3.63)$$

$$H_{\text{err}_2} = \tilde{v}_{IX} \frac{IX}{2} + \tilde{v}_{ZY} \frac{ZY}{2} + \tilde{v}_{ZZ} \frac{ZZ}{2}. \quad (3.64)$$

Here,  $H_{\text{err}_1}$  accounts for error terms from the standard CR sequence in the two-qubit subspace, while  $H_{\text{err}_2}$  represents crosstalk errors caused by the CR drive inadvertently affecting the target qubit.

If the error terms were negligible, the perfect entangling operator  $U_{ZX; \frac{\pi}{2}}$  would emerge for a pulse duration  $\tau = \pi/(2\tilde{v}_{ZX})$  in  $\exp(-i\tau H_{\text{eff}})$ . However, the presence of errors introduces systematic deviations, leading to a modified operator

$$U_{\text{eff}; \frac{\pi}{2}}(\mathbf{B}), \quad (3.65)$$

parameterized by bias ratios  $\mathcal{B} = \{\beta_1, \beta_2, \beta_3, \beta_4, \beta_5\}$ , defined as:

$$\beta_1 = \frac{\tilde{v}_{IY}}{\tilde{v}_{ZX}}, \quad \beta_2 = \frac{\tilde{v}_{IZ}}{\tilde{v}_{ZX}}, \quad \beta_3 = \frac{\tilde{v}_{IX}}{\tilde{v}_{ZX}}, \quad \beta_4 = \frac{\tilde{v}_{ZY}}{\tilde{v}_{ZX}}, \quad \beta_5 = \frac{\tilde{v}_{ZZ}}{\tilde{v}_{ZX}}. \quad (3.66)$$

The elements of this  $4 \times 4$  matrix can be compactly encoded as:

$$\begin{aligned} \langle c't' | U_{\text{eff}; \frac{\pi}{2}}(\mathcal{B}) | ct \rangle = & (c + c' - 1)^2 \left[ \cos\left(\gamma_c \frac{\pi}{4}\right) (t + t' - 1)^2 \right. \\ & + \frac{\sin(\gamma_c \pi/4)}{\gamma_c} \left( i(t + t' - 1)(\beta_2 + (-1)^c \beta_5) \right. \\ & \left. \left. - (t + t' - 2tt') \left( i(\beta_3 + (-1)^c) + (-1)^{t'} (\beta_1 + (-1)^c \beta_4) \right) \right) \right]. \end{aligned} \quad (3.67)$$

where  $c, t, c', t' \in \{0, 1\}$  and

$$\gamma_c = \sqrt{(\beta_1 + (-1)^c \beta_4)^2 + (\beta_2 + (-1)^c \beta_5)^2 + (\beta_3 + (-1)^c)^2}. \quad (3.68)$$

When the biases vanish, i.e.,  $\beta_j \rightarrow 0 \forall j$ , the operator reduces to the ideal case:

$$\lim_{\beta_j \rightarrow 0 \forall j} U_{\text{eff}; \frac{\pi}{2}}(\mathcal{B}) = U_{ZX; \frac{\pi}{2}}. \quad (3.69)$$

To approximate the physical realization of the CNOT gate, we now apply the same local operation  $M$  to the effective propagator. However, since now  $[M, U_{\text{eff}; \frac{\pi}{2}}(\mathbf{B})] \neq 0$ , the operator ordering affects the

resulting matrix. Therefore, we define the biased variant of the CNOT as:

$$\text{BCNOT}_5 \equiv U_{\text{eff}; \frac{\pi}{2}}(\mathbf{B}) \cdot M. \quad (3.70)$$

Explicitly, the matrix elements of this  $U(4)$  operator are encoded as:

$$\begin{aligned} \langle c't' | \text{BCNOT}_5 | ct \rangle = & \frac{i(c' - 1) - cc'(1 + i) + ic}{\sqrt{2}} \left[ \cos(\gamma_c \pi/4) (t(1 + i)(2t' - 1) - (1 + i)t' + i) \right. \\ & + \frac{\sin(\gamma_c \pi/4)}{\gamma_c} \left( \beta_1 + \beta_2 + i\beta_3 + (-1)^c (\beta_4 + \beta_5 + i) \right. \\ & + t(1 + i) (-\beta_1 + i(\beta_2 + \beta_3(2t' - 1) + (-1)^c (i\beta_4 + \beta_5 + 2t' - 1))) \\ & \left. \left. + t'(i - 1) (\beta_1 + i\beta_2 - \beta_3 + (-1)^c (\beta_4 + i\beta_5 - 1)) \right) \right]. \end{aligned} \quad (3.71)$$

In the limit of no bias,  $\beta_j \rightarrow 0 \forall j$ , both the  $\text{BCNOT}_5$  and a  $\text{BCNOT}_2$ —when only  $H_{\text{err}_1}$  is considered and  $H_{\text{err}_2} = 0$ —operators converge to the ideal CNOT gate, which is a perfect entangler. The presence of bias, however, impairs the ability of these gates to become perfect entanglers.

An alternative operator ordering,  $\text{BCNOT}'_5 \equiv M \cdot U_{\text{eff}; \frac{\pi}{2}}(\mathbf{B})$ , leads to a different parametrization that can be connected to  $\text{BCNOT}_5$  via  $\{\beta_1 \mapsto \beta_2, \beta_2 \mapsto -\beta_1, \beta_3 \mapsto \beta_3, \beta_4 \mapsto \beta_5, \beta_5 \mapsto -\beta_4\}$  substitutions. This equivalence ensures that either definition can model experimental data, with differences only in parameter interpretation. This is sufficient when the objective is to validate that the entire set of biases effectively models the experimental data, without the need for a detailed characterization of each individual bias.

Importantly, the parametrizations of differently constructed BCNOT gates remain connected in the other two possibilities of constructing them by combining the base two-qubit entangler with the phase and  $X$  rotation gates in different orderings:

$$\text{BCNOT}''_5 \equiv \left( S^\dagger \otimes \mathbb{1} \right) \cdot U_{\text{eff}; \frac{\pi}{2}}(\mathcal{B}) \cdot \left( \mathbb{1} \otimes R_x \left( -\frac{\pi}{2} \right) \right), \quad (3.72)$$

$$\text{BCNOT}'''_5 \equiv \left( \mathbb{1} \otimes R_x \left( -\frac{\pi}{2} \right) \right) \cdot U_{\text{eff}; \frac{\pi}{2}}(\mathcal{B}) \cdot \left( S^\dagger \otimes \mathbb{1} \right). \quad (3.73)$$

While  $\text{BCNOT}''_5 = \text{BCNOT}_5$ , the map  $\text{BCNOT}_5 \mapsto \text{BCNOT}'''_5$  can be established via the same parameter transformations shown before for  $\text{BCNOT}'_5$ . This consistency allows a unified analysis of bias across all gate constructions, simplifying the experimental interpretation. Moreover, the transformations between  $\text{BCNOT}_5$ ,  $\text{BCNOT}'_5$ , and  $\text{BCNOT}'''_5$  maintain the separation of the  $\beta_j$  parameters associated with the distinct error Hamiltonians described in Eqs. (3.63) and (3.64). This property is particularly ad-

vantageous when one of the error terms is excluded from the analysis, as it ensures that the gates remain interconnected and can be used to globally model experimental data without requiring knowledge of the exact placement of single-qubit gates in the pulse sequence.

### 3.5 Refinement and statistical validation of the circuit models

In this section, we extend the theoretical framework introduced earlier by generalizing Eqs. (3.54 – 3.58) to incorporate alternative implementations of the CNOT operation. Specifically, we replace the ideal CNOT in the circuit of Fig. 3.2 with the BCNOT<sub>2</sub> and BCNOT<sub>5</sub> gates, combined with the single-qubit gate corrections given in Eqs. (3.51 – 3.53). The resulting models, enhanced with the scaling and shifting parameters  $\eta$  and  $\varepsilon$  as described in Eq. (3.34), are then employed to analyze experimental data and infer the physical processes at play in the quantum device.

While this approach enables a comprehensive representation, the derived expressions for these models often become prohibitively complex, rendering their direct analytical evaluation and interpretation challenging. To address this, a truncated series expansion for the BCNOT gates, assuming small  $\beta_j$  parameters, could provide a simplified approximation. However, to ensure that higher-order terms are not neglected, we instead rely on the exact numerical computation of the circuit outputs to fit all the results.

We proceed by fitting experimental data for all previously analyzed quantities, including  $\mathcal{D}$ ,  $\langle X \rangle$ ,  $\langle X_0 \rangle$ ,  $\langle X_1 \rangle$ , and  $\mathcal{D}_m$ , against theoretical predictions obtained from these extended models. The coherent models, incorporating empirical adjustments via  $\eta$  and  $\varepsilon$ , are evaluated for their ability to accurately capture systematic deviations observed in the experiments (§ 3.3).

To assess the performance of the models objectively, we employ the statistical measures in Eq. (3.39) and Eq. (3.45), i.e., the reduced chi-squared ( $\chi^2_\nu$ ) and the residual standard error (RSE). Crucially, we mitigate overfitting by using tenfold cross-validation, partitioning the  $(\phi, \alpha)$  space into uniform random subsets. This methodology enables a robust comparison of model adequacy across splits, ensuring that training and testing are conducted on disjoint data sets to maximize accuracy while preventing bias towards more complex and parametrized models. The results of these analyses, summarized in Table 3.1, rank the models based on their predictive accuracy and the structure of their residuals.

From Table 3.1, we first observe that models incorporating the BCNOT<sub>5</sub> gate along with single-qubit gate errors result in a more accurate representation of the experimental data. While BCNOT<sub>2</sub> gates sometimes yield comparable results, BCNOT<sub>5</sub> provides more consistent and physically realistic parameter estimates across all models, as detailed in Appendix 3.7.2. However, the extent of the “goodness

of fit” improvement is not uniform across all the models and depends on the specific observable being considered. This observation can be explained by the fact that different quantities have varying degrees of sensitivity to the types of coherent errors introduced by the quantum gates. This phenomenon is akin to the behavior observed in Eqs. (3.54 – 3.58), where the measured quantities react differently to the perturbations depending on the circuit involved.

Next, let us consider the results for the observable  $\langle X \rangle$ . The data in Table 3.1 demonstrates that the models that account for SQGEs and incorporate both  $\text{BCNOT}_2$  and  $\text{BCNOT}_5$  gates provide the best fits. Interestingly, the noticeable pattern of deviation observed in the left panel of Fig. 3.5 begins to vanish when using the  $\text{BCNOT}_2$  gates, as indicated by the residuals from any of the ten random splits in the cross-validation process. Despite this, when assessed solely through the goodness-of-fit metrics, there is no clear advantage of one model over the other between  $\text{BCNOT}_2$  and  $\text{BCNOT}_5$  for this particular case. This absence of a decisive improvement from  $\text{BCNOT}_5$  could suggest a potential overparameterization, where the added complexity of the  $\text{BCNOT}_5$  gates does not contribute significantly to the overall fit quality. However, a closer examination of the parameters listed in Table 3.2 reveals that the  $\text{BCNOT}_5$  model generally produces more stable and realistic parameter values.

For the conditional interference patterns  $\langle X_0 \rangle$  and  $\langle X_1 \rangle$ , the performance metrics clearly indicate that incorporating SQGEs and utilizing the  $\text{BCNOT}_5$  gate leads to the best fit. The improvement in the goodness of fit is particularly pronounced for  $\langle X_0 \rangle$ , where the  $\text{BCNOT}_5$  model offers a much better representation of the data compared to the other models. To visualize this, the region shown in the inset of the top-left panel of Fig. 3.6 has been plotted using the fitted equations, revealing that the initial single sinusoidal curve is replaced by multiple sinusoids that fit the data points for each value of  $\alpha$ .

Nevertheless, while the residuals of the fits for both  $\langle X_0 \rangle$  and  $\langle X_1 \rangle$  show marked improvement, they still exhibit noticeable patterns, albeit with smaller amplitudes. Notably, high-frequency oscillations are clearly visible in the residuals, which suggests the presence of additional effects that are not captured by the current models. For instance, in the bottom-left panel of Fig. 3.6, these subtle oscillations in the data can be seen around the fitted model curves, and remain even after fitting with the improved models.

Let us now consider the observations for  $\mathcal{D}$  and  $\mathcal{D}_m$ . For both of these quantities, the model choice does not appear to have a significant impact on the fit quality. Both scores remain relatively unchanged regardless of whether  $\text{BCNOT}_2$  or  $\text{BCNOT}_5$  gates are used, although a slight improvement is observed for  $\mathcal{D}_m$  when including the SQGEs and using either of the  $\text{BCNOT}$  gates. To be fair, however, the fit quality for the simplest model of  $\mathcal{D}$  is already the best in Table 3.1 to begin with, as indicated by its already low scores which suggest that little further improvement is possible for this quantity.

|                       |            | No SQGEs | SQGEs   |                    |                    |
|-----------------------|------------|----------|---------|--------------------|--------------------|
|                       |            | CNOT     | CNOT    | BCNOT <sub>2</sub> | BCNOT <sub>5</sub> |
| $\mathcal{D}$         | $\chi^2_v$ | 1.34692  | 1.34766 | 1.34993            | 1.35376            |
|                       | RSE        | 0.00950  | 0.00950 | 0.00950            | 0.00952            |
| $\langle X \rangle$   | $\chi^2_v$ | 5.78430  | 2.20479 | <b>1.81458</b>     | <b>1.81619</b>     |
|                       | RSE        | 0.02062  | 0.01278 | <b>0.01161</b>     | <b>0.01161</b>     |
| $\langle X_0 \rangle$ | $\chi^2_v$ | 11.07644 | 8.61654 | 7.50310            | <b>2.71485</b>     |
|                       | RSE        | 0.03670  | 0.03094 | 0.02775            | <b>0.02016</b>     |
| $\langle X_1 \rangle$ | $\chi^2_v$ | 2.83404  | 2.51185 | 2.48586            | <b>2.42447</b>     |
|                       | RSE        | 0.01997  | 0.01891 | 0.01873            | <b>0.01865</b>     |
| $\mathcal{D}_m$       | $\chi^2_v$ | 2.05986  | 2.07046 | <b>2.02330</b>     | <b>2.02358</b>     |
|                       | RSE        | 0.01582  | 0.01586 | <b>0.01568</b>     | <b>0.01568</b>     |

Table 3.1: **Model performance evaluation measures.** Performance metrics averaged across the fits derived from various single- and two-qubit gate models, assessed using tenfold cross-validation. The best-performing models are highlighted in bold. Corresponding parameter values are provided in Table 3.2.

That said, an intriguing feature emerges when examining the experimental data for  $\mathcal{D}$  and  $\mathcal{D}_m$ : a continuous dependence on  $\phi$  is observed. This is particularly surprising since there is no explicit  $\phi$  dependence in the long closed-form expressions for these quantities obtained using computer algebra software for the models incorporating SQGEs and the BCNOT<sub>5</sub> gate. This unexpected  $\phi$ -dependence in the data again suggests that there are additional effects at play, which are not accounted for by any the current models, and which could provide further insight into the underlying physics of the platform.

To further investigate these unexpected patterns that, while subtle, remain clearly unaccounted for in these models, it would be essential to consider error sources not explicitly included in the analysis. We suggest two potential explanations for these observations which are not addressed by our BCNOT models, themselves based on the CR gate analysis in [285]. One possibility is the influence of couplings to neighboring *spectator* qubits, while another is the effect of crosstalk caused by the rotary tone of the CR gate affecting the control qubit. These factors, though typically small, could manifest as subtle modulations in the measured quantities, particularly in observables like  $\mathcal{D}$ , which aggregate measurements from multiple qubits. While these explanations are plausible, resolving this question definitively would require further investigation, which is particularly relevant not only due to the primary aim of this study to explore foundational aspects of quantum mechanics but also because certain subtle effects may have been overlooked in the statistical analyses of prior explorations.

In spite of that, cross-validation, the results in Table 3.1, and fewer instances of problematic parameter optimization compared to other models as discussed in Appendix 3.7.2, underscore the benefits of

incorporating both single-qubit gate errors and two-qubit gate errors in modeling experimental data. While simpler models can provide adequate fits in certain cases, the  $\text{BCNOT}_5$  model consistently demonstrates superior performance, particularly for observables more sensitive to coherent errors. This consistency highlights the robustness of the  $\text{BCNOT}_5$  model and supports its use in applications requiring precise characterization of experimental data.

Furthermore, the experimental inability to fully realize the anticipated extremes of purely particlelike or purely wavelike statistical behavior can be attributed to the impact of the biases in the  $\text{CNOT}$  operation, as will be demonstrated in the next section.

### 3.5.1 Impact of the $\text{BCNOT}$ on exclusive particlelike or wavelike statistics

From Eq. (3.55), it is evident that single-qubit gate errors do not preclude the expectation value  $\langle X \rangle$  from attaining  $\mathcal{V} = 0$  or  $\mathcal{V} = 1$ . Those errors leave the fringe contrast of Eq. (3.25) unaffected, adding only a phase shift. Similarly, SQGEs alone do not compromise the ideal interference visibility  $\mathcal{V} = 1$  in  $\langle X_0 \rangle$  or  $\langle X_1 \rangle$ , as seen from Eqs. (3.56) and (3.57), which would also undergo only a phase shift. Furthermore, the distinguishability  $\mathcal{D}_m$  in Eq. (3.58) remains invariant under SQGEs. The only case in which SQGEs can limit the achievability of extreme-case statistics is for the distinguishability  $\mathcal{D}$ , where Eq. (3.54) may fail to reach  $\mathcal{D} = 1$ , thereby precluding fully particlelike behavior. In this specific case, however, it becomes experimentally indistinguishable whether the observed deviation arises from SQGEs or from incoherent noise characterized by the  $\eta$  parameter.

To investigate the unique limitations imposed by biases in the  $\text{CNOT}$  gate we analyze the behavior of quantum circuits incorporating a biased  $\text{CNOT}$ , as defined in Eq. (3.71). Unlike single-qubit gate errors, which do not fundamentally obstruct the realization of fully particlelike or wavelike statistics, the presence of biases in two-qubit gates can significantly restrict the ability to observe these extreme cases. To see this, we recompute the output probabilities of all previously studied circuits by replacing the ideal  $\text{CNOT}$  gate in Fig. 3.2 with the biased  $\text{BCNOT}_5$ , while keeping the single-qubit operations as defined in Eq. (3.35). From these recomputed probabilities, we derive the corresponding visibility, distinguishability, and related quantities to assess the impact of the  $\text{CNOT}$  biases on the observed measurement statistics.

Before delving into exact solutions, it is instructive to examine the first-order truncated series expansion for small bias parameters  $\beta_j \approx 0$  ( $j = 1, \dots, 5$ ). The five resulting models, derived from four distinct circuits, are as follows:

$$\mathcal{D} \approx \sin\left(\frac{\alpha}{2}\right) - (\beta_1 + \beta_2 + \beta_4 + \beta_5) \cos\left(\frac{\alpha}{2}\right), \quad (3.74)$$

$$\langle X \rangle \approx \frac{1 + \beta_2 - \beta_4}{2} \cos \phi + \frac{1 - \beta_2 + \beta_4}{2} \cos(\alpha + \phi) - \frac{\beta_1 + \beta_2 + \beta_4 + \beta_5}{2} (\sin \phi - \sin(\alpha + \phi)), \quad (3.75)$$

$$\langle X_0 \rangle \approx \cos(\phi) + \sin(\phi) \left( (\beta_1 - \beta_5) \sin\left(\frac{\alpha}{2}\right) - \beta_1 - \beta_2 - \beta_4 - \beta_5 \right), \quad (3.76)$$

$$\langle X_1 \rangle \approx \cos(\alpha + \phi) + \sin(\alpha + \phi) \left( (\beta_1 - \beta_5) \sin\left(\frac{\alpha}{2}\right) + \beta_1 + \beta_2 + \beta_4 + \beta_5 \right), \quad (3.77)$$

$$\mathcal{D}_m \approx -\beta_1 + \beta_5 + (\beta_2 - \beta_4) \cos\left(\frac{\alpha}{2}\right) - \frac{\pi}{2} \beta_3 \sin\left(\frac{\alpha}{2}\right). \quad (3.78)$$

These approximations reveal significant overparametrization in the models. For instance,  $\mathcal{D}$ ,  $\langle X \rangle$ ,  $\langle X_0 \rangle$ , and  $\langle X_1 \rangle$  could be adequately fitted by considering only nonzero values for  $\beta_1$  and  $\beta_2$ . While some limitations in measurement statistics are evident in these approximations — e.g., Eq. (3.75) excludes  $\mathcal{V} = 0$  — a complete understanding of extreme values for  $\mathcal{D}$  and  $\mathcal{V}$  requires exact computations.

Exact closed-form solutions for these circuits yield expressions too complex for practical use. Thus, we employ a numerical approach to compute all circuit outputs and derived quantities. These results confirm that the  $\text{BCNOT}_5$  biases prevent full particlelike and wavelike statistics to be observed across all tested circuits. The severity of this limitation depends on the relative magnitudes of the bias parameters  $\beta_j$ , with nonzero values for even just one bias inducing observable effects. The single exception is  $\mathcal{V}_{\langle X \rangle}$ , which remains unaffected by  $\beta_3$  alone, incapable of distorting the function or prevent the extremal values.

Figure 3.7 illustrates these limitations in the visibility and distinguishability obtained for all the different circuits by sampling arbitrary values for the biases as an example. In the top row, the visibility of  $\langle X \rangle$  and the corresponding distinguishability are shown. Even minor biases prevent the interferometer qubit from achieving  $\mathcal{V} = 0$  or  $\mathcal{V} = 1$ , and the detector qubit cannot reach  $\mathcal{D} = 1$ . The impact of this effect becomes increasingly evident as the bias parameters grow larger in the panels at the right. It is important to note that  $\mathcal{D}$  no longer represents the true distinguishability but rather the perceived distinguishability derived from measurements in the  $O_\alpha^\dagger$  basis.

The bottom row depicts the visibility of conditional interference patterns  $\langle X_0 \rangle$  and  $\langle X_1 \rangle$ , which ideally should reach unity. However, these patterns exhibit visibility values below one for most cases, with only two  $\alpha$  values within the  $4\pi$  period reaching  $\mathcal{V} \approx 1$ . Furthermore, the reduced visibility of  $\langle X_0 \rangle$  in the range from  $\alpha = 0$  to  $\alpha = \pi$  follows the observations in Fig. 3.6. Additionally,  $\mathcal{D}_m$  confirms that measurement in the  $Z$  basis fails to entirely eliminate two-qubit correlations or enable complete wavelike behavior. The observed  $4\pi$  periodicity arises from the  $2\pi$  range of input variables for the circuits.

These findings align with experimental results, where  $\mathcal{V} = 0$  could not be observed in the optimal which-path detection experiments (full particlelike statistics), and  $\alpha$ -independent constant high visibility or  $\mathcal{D}_m = 0$  was not achieved in quantum eraser experiments (full wavelike statistics). Thus, even mi-




Figure 3.7: **Theoretical visibility and distinguishability trade-off under a biased-CNOT model.** The visibility and distinguishability values were obtained from the exact numerical computation of  $\mathcal{D}$ ,  $\langle X \rangle$ ,  $\langle X_0 \rangle$ ,  $\langle X_1 \rangle$ , and  $\mathcal{D}_m$  with the various circuits implemented via the decomposition in Fig. 3.2, where the CNOT gate is replaced by a BCNOT<sub>5</sub> featuring illustrative bias parameters. Across all panels, the bias values are set to  $\beta_1 = 0.06$ ,  $\beta_3 = -0.05$ ,  $\beta_4 = -0.07$ , and  $\beta_5 = 0.06$ . In the left-hand panels,  $\beta_2 = 0.09$ , while in the right-hand panels,  $\beta_2$  is significantly increased to 0.3 to highlight the pronounced impact of this parameter. Reproduced from [76]. © 2021 American Physical Society.

nor systematic biases in two-qubit gates can impose fundamental constraints on realizing Bohr’s strong complementarity principle, and require careful experimental analysis to detect their impact.

### 3.6 Discussion and conclusions

In this chapter, we examined interferometer-detector complementarity and wave-particle duality using a superconducting transmon qubit platform. To this end, we proposed a versatile two-qubit quantum circuit capable of implementing which-path detection with adjustable sensitivity and readout basis. This design ensures compatibility with any gate-based quantum computing platform equipped with a universal gate set, regardless of the specific hardware.

We explored four distinct configurations of the readout basis, corresponding to four parameterized circuits, each analyzed theoretically. The first two configurations employed the detector’s qubit measurement basis that minimized the average probability of detection error, enabling the extraction of maximal which-path information. This approach optimized the probability of accurately determining the path through the interferometer, independently of its readout basis. By measuring the interferometer qubit in both the  $X$

and  $Z$  bases and combining the results, we tested the duality relation in Eq. (3.2) across a range of visibility and distinguishability values governed by the detector sensitivity parameter,  $\alpha$ .

The remaining two configurations implemented the quantum eraser setup, wherein the detector qubit is measured directly in the computational basis. In this case, theoretical expectations suggest that the measured distinguishability should vanish for all values of  $\alpha$ , as a result of erasing which-path information at the detector readout. This allows the recovery of two interference patterns with full visibility, a phenomenon absent in the optimal readout scenario. By sorting measurement results on the interferometer qubit conditioned on the detector's output, we demonstrated this recovery mechanism theoretically.

The circuits were executed on an IBM Quantum system accessed remotely. At a macroscopic level, the experimental results aligned well with theoretical predictions, as expected due to the relatively low depth of the circuits. However, at finer granularity, we identified small but systematic deviations from theoretical expectations.

Notably, the  $X$ -basis measurements of the interferometer qubit exhibited a persistent *waviness* across all  $\alpha$  values, preventing the realization of entirely particlelike statistics. Similarly, complete wavelike behavior was not attained for any  $\alpha$  in these circuits as well as in the quantum eraser setup, where it should hold for all  $\alpha$ . Residual distinguishability and  $\alpha$ -dependent visibility of the conditional interference patterns were observed for the quantum eraser, indicative of residual particlelike characteristics. Interestingly, this implies the potential to achieve a guessing probability exceeding 50% in the which-path guessing game, even within the quantum eraser framework, by strategically analyzing the circuit and adjusting the strategy before placing a bet.

We attributed these discrepancies between theory and experiment to coherent errors in both single-qubit and, predominantly, CNOT gates. Incorporating these errors into theoretical models of these gates—based on the cross-resonance gate in the case of the CNOT implementation—the alignment between theory and experimental data was improved. This analysis underscored the challenges of achieving purely wavelike or particlelike statistics, given the requirement of precisely controlled and unbiased two-qubit gates. This might seem like an overly meticulous observation, but it is important to recognize that such systematic errors, although small, are statistically resolvable and therefore carry fundamental significance.

These findings raise important questions about whether such limitations are intrinsic to our implementation or if they manifest similarly across other platforms. For instance, quantum circuits implementing a delayed-choice quantum eraser via bipartite entanglement [288] and via Scully-Drühl-type experiments [289] were realized after the publication of the results presented in this chapter on both a similar transmon

processor from IBM Quantum and a trapped ion processor from IonQ. In both cases, the predicted recovery of interference patterns was achieved. However, significant systematic errors were also observed, underscoring the persistent challenges posed by the limitations of current quantum hardware.

Furthermore, the question of how unbiased, maximally entangling gates or circuits can be realized within the decoherence times of particular quantum platforms is crucial for confidently testing Bohr's strong complementarity principle, beyond the weaker inequality relation in Eq. (3.2). Achieving this would require circuits capable of resolving minute systematic deviations with high precision in the measured observables across the entire variable space.

While improving gate calibration is one approach to achieve the statistical validation of this principle, alternative strategies include customizing circuits [290, 291, 292]. However, such modifications increase circuit depth, thereby amplifying susceptibility to incoherent noise, which hides small coherent errors. To focus on fundamentally testing quantum mechanics with high accuracy, our approach deliberately avoids relying on mixtures of different circuit configurations or qubit pairs [100], which exchange coherent errors for incoherent ones and obscure quantum phenomena beneath classical mixtures. In doing so, our findings also demonstrate that which-path experiments can serve as a simple yet important benchmark for assessing quantum hardware performance.

Historically, which-path *gedanken* experiments were conceived as theoretical exercises to illustrate the unfamiliar principles of quantum mechanics, requiring highly specialized equipment for experimental verification, even though many of the earlier experiments remained within the chaotic statistical regime. With the emergence of cloud-based quantum computers, our work demonstrates a practical means of probing this fundamental quantum principle within the circuit model with a highly controlled individual quantum system.

## 3.7 Appendix

### 3.7.1 Conditional statistics in discrete bivariate sampling

Consider a bivariate distribution of two discrete random variables  $Z_1, Z_2 \in \{0, 1\}$ . The joint probability mass function, denoted as  $P(xy) \equiv P(Z_1 = x, Z_2 = y)$ , and the marginal probability distributions for each

variable are summarized in the following table:

|           | $Z_2 = 0$    | $Z_2 = 1$    | Total        |
|-----------|--------------|--------------|--------------|
| $Z_1 = 0$ | $P(00)$      | $P(01)$      | $P_{Z_1}(0)$ |
| $Z_1 = 1$ | $P(10)$      | $P(11)$      | $P_{Z_1}(1)$ |
| Total     | $P_{Z_2}(0)$ | $P_{Z_2}(1)$ | 1            |

The *conditional expectation* of  $Z_2$  given  $Z_1 = x$  is defined as:

$$\mathbb{E}[Z_2 \mid Z_1 = x] = \sum_{y \in \{0,1\}} y \cdot P(Z_2 = y \mid Z_1 = x) \quad (3.79)$$

$$= P(Z_2 = 1 \mid Z_1 = x) \quad (3.80)$$

$$= \frac{P(x1)}{P(x0) + P(x1)}. \quad (3.81)$$

For a random sample of size  $S$ , the sample estimator of this conditional expectation is estimated as:

$$\bar{\mathbb{E}}[Z_2 \mid Z_1 = 0] = \frac{\bar{P}(x1)}{\bar{P}(x0) + \bar{P}(x1)}, \quad (3.82)$$

where  $\bar{P}(xy)$  denotes the relative frequency of the event  $(Z_1 = x, Z_2 = y)$  in the sample.

The *conditional variance*  $\text{Var}[Z_2 \mid Z_1 = x]$  quantifies the variability of  $Z_2$  within the subgroup where  $Z_1 = x$ . Since  $Z_2$  is Bernoulli-distributed within each conditional subgroup, the conditional variance is:

$$\text{Var}[Z_2 \mid Z_1 = x] = P(Z_2 = 0 \mid Z_1 = x) \cdot P(Z_2 = 1 \mid Z_1 = x) \quad (3.83)$$

$$= \frac{P(x0) \cdot P(x1)}{(P(x0) + P(x1))^2}. \quad (3.84)$$

When a sample of size  $S$  is drawn, the variance of the sample estimator  $\bar{\mathbb{E}}[Z_2 \mid Z_1 = 0]$  is computed by dividing the conditional variance by the number of observations that fall in the conditional subgroup, that is

$$\text{Var}[\bar{\mathbb{E}}[Z_2 \mid Z_1 = x]] = \frac{\text{Var}[Z_2 \mid Z_1 = x]}{S(P(x0) + P(x1))} \quad (3.85)$$

$$= \frac{P(x0) \cdot P(x1)}{S(P(x0) + P(x1))^3}. \quad (3.86)$$

In practice, we estimate probabilities using relative frequencies and apply the finite sample correction

factor:

$$\overline{\text{Var}}[\mathbb{E}[Z_2 \mid Z_1 = x]] = \frac{\bar{P}(x0) \cdot \bar{P}(x1)}{(S - 1) \cdot (\bar{P}(x0) + \bar{P}(x1))^3}. \quad (3.87)$$

The correction factor  $S - 1$  accounts for the fact that sample variances tend to underestimate population variances when computed using finite samples. This adjustment ensures that the estimator of the variance remains unbiased. It becomes negligible as  $S \rightarrow \infty$ , but it is essential for smaller samples. The standard deviation is then obtained as the square root of the variance.

Now consider the expression for the *distinguishability*  $\mathcal{D}$ , derived from Eqs. (3.18) and (3.19) as:

$$\mathcal{D} = \frac{P(00)}{P(00) + P(01)} + \frac{P(11)}{P(10) + P(11)} - 1. \quad (3.88)$$

To estimate  $\mathcal{D}$  from a finite sample of size  $S$ , we use the corresponding sample-based estimator:

$$\overline{\mathcal{D}} \equiv \mathbb{E} \left[ \frac{P(00)}{P(00) + P(01)} + \frac{P(11)}{P(10) + P(11)} - 1 \right] \quad (3.89)$$

$$= \frac{\bar{P}(00)}{\bar{P}(00) + \bar{P}(01)} + \frac{\bar{P}(11)}{\bar{P}(10) + \bar{P}(11)} - 1 \quad (3.90)$$

The variance of the sample estimator  $\overline{\mathcal{D}}$  is computed as the sum of the variances of its components  $\bar{X}$  and  $\bar{Y}$ :

$$\text{Var}[\overline{\mathcal{D}}] = \text{Var}[\bar{X}] + \text{Var}[\bar{Y}], \quad (3.91)$$

where:

$$\bar{X} \equiv \frac{\bar{P}(00)}{\bar{P}(00) + \bar{P}(01)}, \quad \bar{Y} \equiv \frac{\bar{P}(11)}{\bar{P}(10) + \bar{P}(11)}. \quad (3.92)$$

The covariance term is omitted because  $X$  and  $Y$  are derived from disjoint subsets of the probability space. Specifically,  $X$  corresponds to outcomes conditional on  $Z_1 = 0$ , while  $Y$  corresponds to outcomes conditional on  $Z_1 = 1$ . Since these subsets are mutually exclusive, no co-occurring observations exist, leading to  $\text{Cov}[\bar{X}, \bar{Y}] = 0$ .

The variances  $\text{Var}[\bar{X}]$  and  $\text{Var}[\bar{Y}]$  follow directly from Eq. (3.86), therefore

$$\text{Var}[\overline{\mathcal{D}}] = \frac{\bar{P}(00) \cdot \bar{P}(01)}{S(\bar{P}(00) + \bar{P}(01))^3} + \frac{\bar{P}(11) \cdot \bar{P}(10)}{S(\bar{P}(10) + \bar{P}(11))^3}. \quad (3.93)$$

For finite samples, the variance is adjusted as described earlier using a correction factor,  $S \rightarrow S - 1$ ,

to ensure an unbiased estimate:

$$\overline{\text{Var}}[\mathcal{D}] = \frac{\bar{P}(00) \cdot \bar{P}(01)}{(S-1) \cdot (\bar{P}(00) + \bar{P}(01))^3} + \frac{\bar{P}(11) \cdot \bar{P}(10)}{(S-1) \cdot (\bar{P}(10) + \bar{P}(11))^3}. \quad (3.94)$$

This correction becomes negligible as  $S \rightarrow \infty$ .

In summary,  $\mathcal{D}$  provides a quantitative measure of the distinguishability of two conditional distributions, and its variance estimator ensures accurate quantification of uncertainty, even for finite samples.

### 3.7.2 Model parameter values and other details of model evaluation

As outlined in § 3.5, we consider various combinations of ideal and noisy models for the physical gates depicted in Fig. 3.2 to construct four distinct models for the output probabilities of each circuit used in our experiments. The experimental data obtained subsequently guides the fitting of these models by enabling parameter optimization under the constraints  $|\beta_j| \leq 0.5$ ,  $|\theta_k| \leq \pi/10$ ,  $\eta \in [0.7, 1]$ , and  $|\varepsilon| \leq 0.5$ . Both this fitting procedure, along with the evaluation of model performance, is conducted using tenfold cross-validation. The mean parameter values obtained from the ten splits of the cross-validation procedure are used as final estimates, with their standard deviations reported as precision errors.

Focusing on the first column of Table 3.2, which corresponds to models employing the ideal CNOT gate without single-qubit gate errors, we observe that these models yield some of the smallest precision margins. This is expected, as the limited parameterization in these models restricts their flexibility, enabling them to align closely with the data without significant optimization freedom. However, Table 3.1 reveals that these models perform poorly in terms of goodness-of-fit metrics.

In the subsequent columns of Table 3.2, representing more complex models, some optimal parameter values are found at the predefined bounds. For the reasons described in § 3.5, numerical calculations serve as the reference for fitting. However, in the case of the models in the second column, closed-form solutions exist, as given by Eqs. (3.54 – 3.58). Examination of the table reveals that parameters with unstable values (underlined) are absent from these expressions. As a result, it becomes clear why the fitting procedure could not optimize these nonexistent parameters, rendering them irrelevant and explaining these pathologies.

In the third and fourth columns, two overlapping factors contribute to the instability of some parameter values. First, as for the previous column, certain error parameters—though present in the gate operations—may not influence the observables measured in these specific circuits. The numerical calculations include all the error parameters for generality, making it difficult to determine which ones could be omitted

for specific circuits without adding further complexity to the analysis.

Second, as highlighted in § 3.5, the data indicate residual coherent effects unaccounted for even by the more complex models. These models, with their increased degrees of freedom, attempt to fit data that partially deviate from the theoretical assumptions, resulting in some parameter values that may lack direct physical significance. Despite these challenges, models incorporating SQGEs and the  $\text{BCNOT}_5$  gate generally achieve the best balance of goodness-of-fit metrics and parameter stability.

It is important to note that the  $\beta_j$  and  $\theta_k$  parameters listed in different rows of Table 3.2 do not necessarily correspond, as they derive from experiments conducted under varying conditions, including different qubit assignment and calibrations, and device configurations. Nevertheless, for models associated with  $\langle X \rangle$  and  $\mathcal{D}$ , or  $\langle X_0 \rangle$  and  $\langle X_1 \rangle$ , where data were collected simultaneously, certain  $\beta_j$  and  $\theta_k$  values in the final column seem comparable, while others do not. This inconsistency aligns with earlier observations that some physical effects remain unmodeled, allowing models the flexibility to arrive at differing optimal parameter values. Additionally, overparameterization may lead to multiple parameter configurations yielding similar goodness-of-fit results, further contributing to observed differences. As explained in this chapter, however, the goal of these experiments is not to precisely estimate each parameter, but rather to use them collectively to model the data.

In summary, our models successfully capture the primary physics underlying the gate operations and provide an interpretable framework for analyzing the data. Nonetheless, minor gaps persist, as evidenced by non-uniformity of the residuals, incomplete fit quality, and the occasional instability of certain parameters.

| No SQGEs              |               |            |               | SQGEs                           |                    |                                 |                    |                                 |                   |
|-----------------------|---------------|------------|---------------|---------------------------------|--------------------|---------------------------------|--------------------|---------------------------------|-------------------|
| CNOT                  |               |            |               | CNOT                            | BCNOT <sub>2</sub> |                                 | BCNOT <sub>5</sub> |                                 |                   |
| $\mathcal{D}$         | $\eta$        | 0.90206    | $\pm 0.00011$ | 0.91506                         | $\pm 0.00453$      | 0.92386                         | $\pm 0.00352$      | 0.92707                         | $\pm 0.00238$     |
|                       | $\varepsilon$ | 0.00261    | $\pm 0.00009$ | 0.00261                         | $\pm 0.00011$      | 0.00240                         | $\pm 0.00013$      | 0.00006                         | $\pm 0.00147$     |
|                       | $\theta_1$    | —          | —             | <u>0.01839 <math>\pi</math></u> | $\pm 0.06088 \pi$  | 0.00471 $\pi$                   | $\pm 0.07067 \pi$  | $-0.01176 \pi$                  | $\pm 0.05021 \pi$ |
|                       | $\theta_2$    | —          | —             | $-0.06065 \pi$                  | $\pm 0.07914 \pi$  | $-0.04070 \pi$                  | $\pm 0.09111 \pi$  | $-0.06065 \pi$                  | $\pm 0.07914 \pi$ |
|                       | $\theta_3$    | —          | —             | 0.02689 $\pi$                   | $\pm 0.00462 \pi$  | 0.02131 $\pi$                   | $\pm 0.00151 \pi$  | 0.01876 $\pi$                   | $\pm 0.00183 \pi$ |
|                       | $\theta_4$    | —          | —             | 0.02604 $\pi$                   | $\pm 0.00430 \pi$  | 0.02161 $\pi$                   | $\pm 0.00174 \pi$  | 0.01876 $\pi$                   | $\pm 0.00183 \pi$ |
|                       | $\theta_5$    | —          | —             | 0.00037 $\pi$                   | $\pm 0.00002 \pi$  | $-0.02377 \pi$                  | $\pm 0.00405 \pi$  | $-0.02842 \pi$                  | $\pm 0.00646 \pi$ |
|                       | $\beta_1$     | —          | —             | —                               | —                  | 0.03751                         | $\pm 0.00629$      | 0.01340                         | $\pm 0.00759$     |
|                       | $\beta_2$     | —          | —             | —                               | —                  | 0.03800                         | $\pm 0.00628$      | 0.02078                         | $\pm 0.00539$     |
|                       | $\beta_3$     | —          | —             | —                               | —                  | —                               | —                  | 0.06064                         | $\pm 0.00880$     |
|                       | $\beta_4$     | —          | —             | —                               | —                  | —                               | —                  | 0.01920                         | $\pm 0.00587$     |
|                       | $\beta_5$     | —          | —             | —                               | —                  | —                               | —                  | 0.03955                         | $\pm 0.01229$     |
| $\langle X \rangle$   | $\eta$        | 0.92975    | $\pm 0.00020$ | 0.93039                         | $\pm 0.00009$      | 0.94163                         | $\pm 0.00077$      | 0.94054                         | $\pm 0.00170$     |
|                       | $\varepsilon$ | 0.03707    | $\pm 0.00013$ | 0.03700                         | $\pm 0.00005$      | 0.03706                         | $\pm 0.00005$      | 0.03706                         | $\pm 0.00005$     |
|                       | $\theta_1$    | —          | —             | $-0.00525 \pi$                  | $\pm 0.00330 \pi$  | $-0.01759 \pi$                  | $\pm 0.00056 \pi$  | $-0.00769 \pi$                  | $\pm 0.00142 \pi$ |
|                       | $\theta_2$    | —          | —             | $-0.00573 \pi$                  | $\pm 0.00330 \pi$  | $-0.01707 \pi$                  | $\pm 0.00099 \pi$  | $-0.00767 \pi$                  | $\pm 0.00142 \pi$ |
|                       | $\theta_3$    | —          | —             | <u>0.05396 <math>\pi</math></u> | $\pm 0.07693 \pi$  | <u>0.10000 <math>\pi</math></u> | $\pm 0.00000 \pi$  | 0.00961 $\pi$                   | $\pm 0.01214 \pi$ |
|                       | $\theta_4$    | —          | —             | <u>0.02049 <math>\pi</math></u> | $\pm 0.06694 \pi$  | $-0.01281 \pi$                  | $\pm 0.07203 \pi$  | $-0.00520 \pi$                  | $\pm 0.08215 \pi$ |
|                       | $\theta_5$    | —          | —             | $-0.00082 \pi$                  | $\pm 0.00004 \pi$  | $-0.05063 \pi$                  | $\pm 0.00196 \pi$  | $-0.00676 \pi$                  | $\pm 0.00396 \pi$ |
|                       | $\beta_1$     | —          | —             | —                               | —                  | 0.17064                         | $\pm 0.00658$      | 0.10534                         | $\pm 0.02135$     |
|                       | $\beta_2$     | —          | —             | —                               | —                  | $-0.01194$                      | $\pm 0.00029$      | $-0.03753$                      | $\pm 0.01307$     |
|                       | $\beta_3$     | —          | —             | —                               | —                  | —                               | —                  | 0.07033                         | $\pm 0.07038$     |
|                       | $\beta_4$     | —          | —             | —                               | —                  | —                               | —                  | $-0.01431$                      | $\pm 0.01384$     |
|                       | $\beta_5$     | —          | —             | —                               | —                  | —                               | —                  | $-0.03618$                      | $\pm 0.01304$     |
| $\langle X_0 \rangle$ | $\eta$        | 0.91745    | $\pm 0.00015$ | 0.91740                         | $\pm 0.00018$      | 0.91900                         | $\pm 0.00011$      | 0.94262                         | $\pm 0.00014$     |
|                       | $\varepsilon$ | 0.03394    | $\pm 0.00008$ | 0.03360                         | $\pm 0.00007$      | 0.03356                         | $\pm 0.00007$      | 0.03237                         | $\pm 0.00005$     |
|                       | $\theta_1$    | —          | —             | $-0.00331 \pi$                  | $\pm 0.00002 \pi$  | 0.05805 $\pi$                   | $\pm 0.00009 \pi$  | 0.00326 $\pi$                   | $\pm 0.00171 \pi$ |
|                       | $\theta_2$    | —          | —             | $-0.00331 \pi$                  | $\pm 0.00002 \pi$  | 0.05805 $\pi$                   | $\pm 0.00010 \pi$  | 0.00327 $\pi$                   | $\pm 0.00172 \pi$ |
|                       | $\theta_3$    | —          | —             | <u>0.00645 <math>\pi</math></u> | $\pm 0.09506 \pi$  | <u>0.10000 <math>\pi</math></u> | $\pm 0.00000 \pi$  | $-0.01545 \pi$                  | $\pm 0.02125 \pi$ |
|                       | $\theta_4$    | —          | —             | <u>0.01769 <math>\pi</math></u> | $\pm 0.07071 \pi$  | $-0.01414 \pi$                  | $\pm 0.08456 \pi$  | <u>0.00289 <math>\pi</math></u> | $\pm 0.09722 \pi$ |
|                       | $\theta_5$    | —          | —             | $-0.00331 \pi$                  | $\pm 0.00002 \pi$  | $-0.10000 \pi$                  | $\pm 0.00000 \pi$  | 0.00512 $\pi$                   | $\pm 0.00443 \pi$ |
|                       | $\beta_1$     | —          | —             | —                               | —                  | $-0.02477$                      | $\pm 0.00063$      | $-0.00394$                      | $\pm 0.02631$     |
|                       | $\beta_2$     | —          | —             | —                               | —                  | $-0.05963$                      | $\pm 0.00029$      | $-0.02698$                      | $\pm 0.02227$     |
|                       | $\beta_3$     | —          | —             | —                               | —                  | —                               | —                  | 0.23815                         | $\pm 0.00155$     |
|                       | $\beta_4$     | —          | —             | —                               | —                  | —                               | —                  | $-0.06506$                      | $\pm 0.03454$     |
|                       | $\beta_5$     | —          | —             | —                               | —                  | —                               | —                  | 0.02333                         | $\pm 0.03759$     |
| $\langle X_1 \rangle$ | $\eta$        | 0.93767    | $\pm 0.00006$ | 0.93765                         | $\pm 0.00005$      | 0.93826                         | $\pm 0.00009$      | 0.94136                         | $\pm 0.00007$     |
|                       | $\varepsilon$ | 0.03391    | $\pm 0.00004$ | 0.03379                         | $\pm 0.00006$      | 0.03382                         | $\pm 0.00005$      | 0.03379                         | $\pm 0.00006$     |
|                       | $\theta_1$    | —          | —             | $-0.00137 \pi$                  | $\pm 0.00047 \pi$  | 0.05483 $\pi$                   | $\pm 0.00050 \pi$  | 0.00477 $\pi$                   | $\pm 0.00164 \pi$ |
|                       | $\theta_2$    | —          | —             | $-0.00116 \pi$                  | $\pm 0.00030 \pi$  | 0.05485 $\pi$                   | $\pm 0.00047 \pi$  | 0.00477 $\pi$                   | $\pm 0.00164 \pi$ |
|                       | $\theta_3$    | —          | —             | $-0.02542 \pi$                  | $\pm 0.06461 \pi$  | $-0.03627 \pi$                  | $\pm 0.00287 \pi$  | 0.01919 $\pi$                   | $\pm 0.00582 \pi$ |
|                       | $\theta_4$    | —          | —             | <u>0.00000 <math>\pi</math></u> | $\pm 0.09982 \pi$  | $-0.03084 \pi$                  | $\pm 0.07377 \pi$  | 0.04070 $\pi$                   | $\pm 0.09111 \pi$ |
|                       | $\theta_5$    | —          | —             | 0.00106 $\pi$                   | $\pm 0.00030 \pi$  | <u>0.10000 <math>\pi</math></u> | $\pm 0.00000 \pi$  | $-0.00225 \pi$                  | $\pm 0.00134 \pi$ |
|                       | $\beta_1$     | —          | —             | —                               | —                  | $-0.00272$                      | $\pm 0.00027$      | $-0.01068$                      | $\pm 0.00611$     |
|                       | $\beta_2$     | —          | —             | —                               | —                  | 0.04713                         | $\pm 0.00304$      | 0.00853                         | $\pm 0.00825$     |
|                       | $\beta_3$     | —          | —             | —                               | —                  | —                               | —                  | 0.06933                         | $\pm 0.00161$     |
|                       | $\beta_4$     | —          | —             | —                               | —                  | —                               | —                  | 0.07775                         | $\pm 0.01049$     |
|                       | $\beta_5$     | —          | —             | —                               | —                  | —                               | —                  | $-0.02379$                      | $\pm 0.00541$     |
| $\mathcal{D}_m$       | $\eta$        | —          | —             | <u>1.00000</u>                  | $\pm 0.00000$      | 0.99999                         | $\pm 0.00002$      | 0.98819                         | $\pm 0.03324$     |
|                       | $\varepsilon$ | $-0.00370$ | $\pm 0.00004$ | $-0.00370$                      | $\pm 0.00007$      | 0.01448                         | $\pm 0.00224$      | $-0.32232$                      | $\pm 0.01093$     |
|                       | $\theta_1$    | —          | —             | 0.00425 $\pi$                   | $\pm 0.07391 \pi$  | $-0.01039 \pi$                  | $\pm 0.07781 \pi$  | $-0.00380 \pi$                  | $\pm 0.06965 \pi$ |
|                       | $\theta_2$    | —          | —             | $-0.02309 \pi$                  | $\pm 0.06652 \pi$  | $-0.00139 \pi$                  | $\pm 0.08569 \pi$  | <u>0.00209 <math>\pi</math></u> | $\pm 0.07158 \pi$ |
|                       | $\theta_3$    | —          | —             | 0.03880 $\pi$                   | $\pm 0.07545 \pi$  | <u>0.09986 <math>\pi</math></u> | $\pm 0.00032 \pi$  | <u>0.06215 <math>\pi</math></u> | $\pm 0.04653 \pi$ |
|                       | $\theta_4$    | —          | —             | $-0.00000 \pi$                  | $\pm 0.00000 \pi$  | $-0.00591 \pi$                  | $\pm 0.06771 \pi$  | $-0.01688 \pi$                  | $\pm 0.06911 \pi$ |
|                       | $\theta_5$    | —          | —             | $-0.03308 \pi$                  | $\pm 0.07253 \pi$  | $-0.03750 \pi$                  | $\pm 0.08945 \pi$  | $-0.09324 \pi$                  | $\pm 0.01471 \pi$ |
|                       | $\beta_1$     | —          | —             | —                               | —                  | $-0.01573$                      | $\pm 0.00234$      | 0.12234                         | $\pm 0.04425$     |
|                       | $\beta_2$     | —          | —             | —                               | —                  | $-0.00364$                      | $\pm 0.00140$      | $-0.37500$                      | $\pm 0.04456$     |
|                       | $\beta_3$     | —          | —             | —                               | —                  | —                               | —                  | 0.05512                         | $\pm 0.01691$     |
|                       | $\beta_4$     | —          | —             | —                               | —                  | —                               | —                  | $-0.10399$                      | $\pm 0.04620$     |
|                       | $\beta_5$     | —          | —             | —                               | —                  | —                               | —                  | $-0.39596$                      | $\pm 0.05204$     |

Table 3.2: **Optimal model parameters obtained for different fitted models.** Using various single- and two-qubit gate models to simulate the circuits, final parameters were obtained from tenfold cross-validation. Underlined values indicate unstable cases where the fitting returned one of the imposed variable bounds; a situation occurring in only one of the splits when corresponding precision errors are null.

## **Circuit discovery and optimization**

# Chapter 4

## Logically-equivalent quantum circuit discovery and optimization

“It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience.”

---

Albert Einstein, 1933 [293]

### 4.1 Introduction

Circuit synthesis is a fundamental procedure in quantum computing, as every  $2^n \times 2^n$  unitary operator designed for a given computational task on  $n$  qubits must be decomposed into the operational gate set of the device that will execute it. At its simplest, such a decomposition can be constructed entirely from single-qubit and two-qubit gates, provided the employed gate set is *universal*. When this gate set is discrete and finite, it has long been established that any multiqubit unitary can be efficiently *approximated* using a circuit whose size grows polylogarithmically with the inverse of the approximation error [294]. For an *exact* implementation of the target unitary, it is also well known that a gate library consisting of parameterized single-qubit gates and at least one entangling two-qubit gate is sufficient [111].

Early results established exact synthesis methods for two-qubit circuits [295], which were followed by generalized techniques for arbitrary  $n$ -qubit unitaries. However, as  $n$  increases, achieving an exact decomposition requires a growing number of two-qubit gates to account for the increasing independent degrees

| Number of qubits               | 2        | 3         | 4          | 5          | $n$                                                             |
|--------------------------------|----------|-----------|------------|------------|-----------------------------------------------------------------|
| Theoretical lower bounds [112] | 3        | 14        | 61         | 252        | $\left\lceil \frac{1}{4} (4^n - 3n - 1) \right\rceil$           |
| QR-optimized [183]             | 8        | 62        | 344        | 1642       | $2 \cdot 4^n - (2n + 3) \cdot 2^n + 2n$                         |
| Givens rotations [299, 300]    | 4        | 64        | 536        | 4156       | $\approx 8.4 \cdot 4^n$                                         |
| Recursive CSD [301]            | 14       | 92        | 504        | 2544       | $\frac{1}{2} \cdot n \cdot 4^n - \frac{1}{2} \cdot 2^n$         |
| Recursive CSD-optimized [297]  | 4        | 26        | 118        | 494        | $\frac{1}{2} \cdot 4^n - \frac{1}{2} \cdot 2^n - 2$             |
| QSD [183]                      | 6        | 36        | 168        | 720        | $\frac{3}{4} \cdot 4^n - \frac{3}{2} \cdot 2^n$                 |
| <b>QSD-optimized [183]</b>     | <b>3</b> | <b>20</b> | <b>100</b> | <b>444</b> | $\frac{23}{48} \cdot 4^n - \frac{3}{2} \cdot 2^n + \frac{4}{3}$ |

Table 4.1: CNOT **counts for decomposing a general  $n$ -qubit unitary using various methods.** The columns present count values for  $n$  ranging from 2 to 5, along with the asymptotic scaling with  $n$ .

of freedom in the circuit. The first algorithm to decompose an arbitrary multiqubit unitary was based on QR factorization and produced a circuit comprising  $\mathcal{O}(n^3 4^n)$  CNOT gates [111, 296], rendering it infeasible even for moderately sized systems. Subsequent research on recursive cosine-sine decomposition (CSD) greatly improved this result [297]. Around the same time, Shende *et al.* introduced recursive Quantum Shannon Decomposition (QSD), a framework that uses multiplexors as building blocks and guaranteed a circuit with no more than  $\frac{23}{48} \cdot 4^n - \frac{3}{2} \cdot 2^n + \frac{4}{3}$  CNOT gates for general unitaries, improving over all previous methods while remaining within a factor of two of the theoretical minimum [183]. Table 4.1 summarizes the two-qubit gate complexity of the main *exact* synthesis methods for general unitaries.<sup>1</sup>

Crucially, all exact synthesis methods incur an exponential cost as  $n$  increases, making them computationally prohibitive for large systems. More fundamentally, this exponential growth in gate count is not merely a limitation of current techniques but reflects an inherent lower bound [112] on the number of CNOT gates required to decompose an arbitrary  $n$ -qubit unitary given by

$$N_{\text{theo}}(n) = \left\lceil \frac{1}{4} (4^n - 3n - 1) \right\rceil. \quad (4.1)$$

This is essentially a parameter-counting argument: the number of independent parameters that can be introduced by circuits with fewer than this many CNOTs is insufficient to cover the full  $4^n$ -dimensional manifold (up to global phase) of  $n$ -qubit unitaries.

Nonetheless, when focusing on specific classes of unitaries with particular structures, it is often possible to significantly reduce the complexity of the circuits synthesized by these general methods. For instance, specialized results provide efficient circuits with subexponential numbers of gates for linear re-

<sup>1</sup>Approximate circuit synthesis methods also exist but have limitations in applicability and are not included in this category. Unlike exact decomposition methods, which construct unitary decompositions algebraically and guarantee exact correctness, approximate decomposition techniques, such as the recent sequential optimization (SO) method [298], lack such formal guarantees. The SO method is a heuristic numerical approach that approximates the decomposition of general unitaries up to an error that is bounded in a suitable operator norm.

versible circuits [302] — which contain only CNOT gates — and Clifford circuits [303] — which consist of CNOT, Hadamard, and Phase ( $S$ ) gates — among others.

More importantly, when targeting individual unitaries corresponding to specific algorithms or applications, one can frequently obtain efficient circuit decompositions. This led to the emergence of a vast literature on *quantum circuit optimization* methods. While synthesis algorithms construct a circuit from scratch, circuit optimization usually<sup>2</sup> starts with an existing circuit, obtained either by the methods above or derived from a quantum algorithm, and iteratively simplifies it without changing its functionality. Over the past two decades, a broad spectrum of optimization techniques has been introduced, ranging from simple local gate cancellation rules to complex global transformations.

### 4.1.1 Quantum circuit optimization

For small circuits with restricted discrete gate sets, one can brute-force optimal solutions. This approach has been applied to Clifford [304] and CNOT +  $T$  [305] circuits using SAT solvers. Additionally, a meet-in-the-middle algorithm with a quadratic speedup over naive brute force has been implemented for Clifford+ $T$  circuits [113]. While limited in scalability due to their exponential worst-case complexity, these methods have successfully found optimal circuits in certain cases, providing a gold standard for benchmarking and guiding more scalable optimization strategies.

For more general qubit numbers, a simple starting strategy for circuit optimization is to exploit local algebraic identities to cancel or merge gates. When gates commute, they can be re-ordered to allow merging local sequences — e.g. successive rotations about the same axis — or to cancel inverse gate sequences through  $UU^\dagger = \mathbb{1}$ . Gate commutation also allows reordering the circuit for parallel execution, reducing depth. Similar to these principles is the systematic use of templates — pre-identified small subcircuits that realize the identity, or another simple equivalence, and can be inserted or removed to shorten a circuit [306]. The idea is analogous to peephole optimization in reversible circuits [307], which replaces small subcircuits with more efficient equivalents. While peephole optimization ensures local optimality within small subcircuits, template matching targets broader simplifications that may not always be locally optimal. These pattern-based techniques are intuitive, but implementing them at scale requires careful tracking of operations. Due to their simplicity and speed, they remain widely used in many compilers.

Between purely local and fully global methods lies the approach of re-synthesizing portions of the circuit. The idea is to identify a subcircuit, like a block of gates or a set of qubits, and replace that entire

---

<sup>2</sup>For small qubit numbers, some optimization methods can search exhaustively over possible decompositions without requiring an initial circuit, relying only on the target unitary.

sub-unit with a new, optimized subcircuit that has the same effect. An example is a recently proposed symbolic peephole method for CNOT gate minimization in Clifford circuits, which projects the circuit onto a small subset of qubits, optimally recompiles it while tracking interactions with its context, and then reinserts it back [308]. This approach achieved near-optimal results for up to six qubits while scaling to 64 qubits with meaningful gate reductions, demonstrating that even when local optimizations plateau, smarter re-synthesis can yield significantly shorter circuits.

A more global optimization approach is the algebraic phase polynomial method, which optimizes circuits with controlled-phase operations, often represented as CNOT +  $R_z$  rotation networks, including Clifford+ $T$  circuits. This method encodes phase interactions as a polynomial over  $\mathbb{Z}_2$ , enabling structured transformations that optimize the circuit while preserving its functional behavior. While initially developed to minimize  $T$  gates [309], this framework has also been applied to directly reduce CNOT gate count [310], yielding an average 23% CNOT reduction on standard benchmarks with 5 to 40 qubits. These optimizations are particularly effective when combined with Clifford simplifications, as they reorganize phase dependencies in a way that reduces entangling operations, demonstrating the power of algebraic representations in circuit optimization.

While phase polynomial techniques optimize Clifford+ $T$  circuits effectively, more general circuits with continuous parameters require different strategies. A key contribution introduced a scalable heuristic-based optimization framework for large-scale circuits while preserving connectivity structure [114]. By handling continuous rotation parameters symbolically, this method significantly reduced two-qubit gate counts across benchmark algorithms with minimal computational overhead, striking a balance between scalability and efficiency for practical quantum computing.

In the last decade, algebraic and graph-based optimization methods have advanced significantly. Chief among these is the ZX-calculus, a graphical language for representing arbitrary linear maps between qubits, based on tensor networks and category theory. It expresses quantum states and operations using a monoidal string diagram notation comprising two primary node types:  $Z$ -spiders and  $X$ -spiders, which correspond to different qubit bases. These nodes are connected by edges representing qubits, forming a graph-like tensor network representation of a quantum process defined by a  $2^k \times 2^j$  complex matrix, where  $j$  is the number of input edges and  $k$  is the number of output edges. The diagrams can be transformed using a set of rewrite rules, such as spider fusion (merging adjacent nodes of the same color), color change (interchanging  $Z$ - and  $X$ -spiders via Hadamard gates), bialgebra laws (capturing the interactions between  $Z$ - and  $X$ -spiders), the Hopf law (eliminating self-loop), identity removal (simplifying trivial wires and nodes), and Euler decomposition (expressing arbitrary single-qubit rotations as phased spiders), as

well as graph-theoretic rules like local complementation, pivoting, and phase gadget manipulation. These well-developed graphical rewrite rules preserve the underlying linear map while enabling simplification, verification, and optimization of complex algebraic transformations that are often less transparent in matrix form. A comprehensive review of this formalism and its numerous applications is beyond the scope for this dissertation. Excellent pedagogical introductions can be found elsewhere [311, 312].

$ZX$ -calculus has become a powerful approach for circuit optimization, with particularly rapid advancements and practical implementations in recent years. At its core, the technique involves converting a quantum circuit into an equivalent  $ZX$ -diagram, applying graph-theoretic transformations to simplify its structure, and then extracting an optimized circuit back [122]. Unlike conventional algebraic techniques, which manipulate gate sequences directly, this graphical representation can capture global structures. This enables cancellations and simplifications that do not rely on gate adjacency in the original circuit and are often obscured in standard circuit form. For instance, entanglement between qubits can sometimes be eliminated entirely through graphical transformations, even when the corresponding gates are spatially separated. This ability to exploit structural symmetries and non-local effects gives  $ZX$ -calculus a unique advantage over conventional optimization approaches.

Many of the early  $ZX$ -based circuit optimization works focused on circuits with significant Clifford+ $T$  structure, demonstrating that  $ZX$ -diagram simplifications could match or outperform traditional  $T$ -count reduction techniques, while also often — but not always — reducing Clifford gate overhead, particularly in two-qubit gates such as CNOTs [122, 313]. The development of software tools like PyZX has made this practical, and subsequent works integrated  $ZX$ -based rewrites into compiler toolchains like tket [314].

While early  $ZX$ -based methods occasionally reduced the CNOT count as a byproduct of Clifford simplification, directly optimizing CNOT gates posed unique challenges. The graph-like structure of  $ZX$ -diagrams does not inherently encode a unique circuit decomposition, and conventional extraction algorithms, based on so-called *generalized flow* properties, often reintroduced redundant two-qubit gates [315]. This bottleneck was mitigated in subsequent work by integrating  $ZX$ -diagram simplification with extraction-aware rewrite strategies that preserved the systematic and efficient Gaussian elimination-based circuit extraction techniques developed before. Heuristic methods prioritizing CNOT minimization enabled the achievement of two-qubit gate counts lower than those attainable through purely algebraic techniques, even for circuits not dominated by Clifford operations [316]. Parallel work has also addressed the problem of extracting circuits suitable for architectures with limited connectivity by introducing an adaptation of Gaussian elimination-based circuit extraction that employs Steiner-tree routing heuristics to efficiently synthesize CNOT-only circuits [317]. While clearly outlining potential avenues for generalizing

their technique to broader classes of circuits, explicit implementations and detailed analyses of such extensions were left open for future research. Subsequent work further reduced CNOT counts, albeit at the expense of introducing final qubit permutations, thus limiting its applicability primarily to full-circuit compilation rather than modular gate simplification for circuit integration [318].

Following these works, further progress has been made by introducing a stronger flow condition called *causal flow*, whose explicit preservation during ZX-diagram transformations guarantees straightforward circuit extraction and reliable reduction of two-qubit gates [319]. Optimized circuits following this approach outperformed the previous ZX-only heuristic methods from [316] on the same benchmarking dataset, though it did not consistently surpass state-of-the-art non-ZX techniques from [114] in every instance.

Another major development in the past decade is the application of artificial intelligence techniques to quantum circuit optimization. The first reinforcement learning (RL)-based approach for optimizing general quantum circuits was designed specifically to minimize circuit depth and total gate count while adapting to specific hardware architectures [320]. By training on randomly generated 12-qubit circuits, the method achieved a depth reduction of approximately 27% and a gate count reduction of 15%, outperforming simulated annealing in both optimization quality and computational efficiency. The approach ensured unitary preservation by construction, as all applied transformations were logically equivalent. Due to its convolutional neural network architecture, the trained RL agent generalized effectively to significantly larger circuits — up to 50 qubits — without requiring further training, providing a scalable and adaptive alternative to manually tuned optimization strategies. However, these advantages came at the cost of substantial upfront computational resources — approximately one week of CPU-based training — highlighting the critical dependency on suitable datasets for training robust optimization agents.

VOQC [321] was the first fully verified quantum circuit optimizer introduced for general quantum circuits, ensuring that all applied transformations are mathematically proven to preserve the original unitary. Unlike RL-based or heuristic approaches, VOQC leverages formal verification using the Coq proof assistant, eliminating the possibility of incorrect optimizations—an issue that had occurred in earlier circuit compilers. The optimizer reduces total gate count, CNOT count, and  $T$ -gate count, and has demonstrated competitive performance in benchmarks on a standard 28-circuit dataset. However, VOQC requires manual specification and verification of optimization rules, making it labor-intensive and limiting its ability to discover new optimization strategies automatically.

Quartz [322] also targets general quantum circuit optimization, introducing a method combining automated symbolic theorem proving with cost-based search. In contrast to VOQC's manual rule verification and strict pattern matching, Quartz automatically discovers circuit transformations verified through the Z3

Satisfiability Modulo Theories (SMT) solver. Its major innovation lies in handling parameterized gates and generalizing optimizations across diverse quantum hardware architectures. Quartz systematically constructs equivalent circuit classes by iteratively identifying alternative subcircuit representations, applies equivalence rules, and prunes redundant operations to minimize gate count through heuristic-driven cost-based backtracking search, indirectly optimizing CNOT and  $T$  gates. It matches or surpasses hand-tuned optimizers, yet its transformation discovery demands substantial computational resources (128-core CPU cluster), a one-time but considerable cost. While efficient for known circuit transformations, adapting Quartz to unseen structures may require costly retraining, limiting its practical flexibility.

Building directly on Quartz, QUESO [323] introduced a scalable approach to automatically synthesize rewrite rules tailored specifically to quantum architectures and general circuits. Replacing Quartz’s SMT-based symbolic verification with Polynomial Identity Testing (PIT) significantly improved efficiency, achieving synthesis of rewrite rules in approximately 1.2 minutes on a standard CPU. The introduction of a Polynomial Identity Filter (PIF) prevented quadratic explosion during rule synthesis, further enhancing scalability. While the probabilistic nature of PIT means QUESO only guarantees correctness with high probability rather than absolute certainty, it nevertheless achieved state-of-the-art optimization performance, efficiently applying synthesized rewrite rules through beam search and matching or outperforming existing optimizers, including Quartz, on standard benchmarks.

Expanding beyond these rule-based and symbolic approaches, Quarl [324] advanced RL-based optimization by introducing a novel hierarchical action space and a graph neural network representation of quantum circuits. This structure allows Quarl to balance local optimization decisions with global circuit-wide improvements, effectively addressing the complexity of quantum circuit optimization. Unlike Quartz or QUESO, Quarl learns long optimization strategies directly from circuit data, breaking down the optimization process into a two-step hierarchical action space: first gate selection, then transformation selection, significantly simplifying the learning task compared to previous single-step RL methods, even when some intermediate steps momentarily increase circuit cost. Quarl ensures optimized circuits preserve the original unitary through symbolic equivalence verification, inherited from Quartz, and complements this with explicit matrix equivalence checks as a robust fallback mechanism to catch edge cases that symbolic verification might fail to prove. Benchmarking using extensive GPU resources — multiple NVIDIA A100 GPUs on the Perlmutter supercomputer at NERSC — demonstrated Quarl’s superior performance, with gate-count reductions of up to 36.6% compared to existing methods. Nonetheless, Quarl’s significant GPU-intensive training requirements may pose practical limitations, particularly when retraining is needed to adapt to entirely new circuit structures or quantum architectures.

Additionally, around the same time as Quarl, another RL-based optimization method leveraging the  $ZX$ -calculus was introduced specifically for Clifford+ $T$  circuits [325]. The approach trained an RL agent on small circuits to iteratively select simplification rules applicable to a given  $ZX$ -diagram, restricting its choices to a subset of transformations that preserve the diagram's gflow. The primary optimization targets were total and two-qubit gate counts in the extracted circuits. Unlike previous methods [114, 316, 319, 321, 322, 323, 324], this approach was benchmarked not on standard circuits such as adders or Toffoli gates but instead on randomly generated Clifford+ $T$  circuits containing up to 80 qubits and 2100 gates, where it outperformed the best previous  $ZX$ -only circuit optimizer [319]. Nonetheless, as typical in previous RL-based optimization methods, training incurred a substantial computational cost, especially for large-scale circuits, thus posing practical limitations for real-time or large-scale optimization tasks. Furthermore, the method does not generalize to circuits involving continuous parameters, such as arbitrary rotational gates.

Also specifically designed for Clifford+ $T$  circuits and using RL, but focusing on  $T$ -gate reduction instead of total gate or CNOT count, and unable to deal with continuous rotation parameters, the AlphaTensor-Quantum method was recently introduced, leveraging domain-specific knowledge through gadgets to further reduce  $T$ -count compared to previous tensor-decomposition-based methods [326]. In parallel, another specialized RL-based circuit optimization method demonstrated “near-optimal” synthesis and routing for linear function, Clifford, and permutation circuits up to 9, 11, and 65 qubits respectively [327]. Finally, RL has also been explored in stochastic or noisy environments, autonomously identifying quantum circuits and error correction codes tailored to specific hardware and noise models [328, 329]. These approaches typically remain context-dependent, adapted to particular hardware constraints and noise profiles, functioning as heuristic optimization techniques rather than methods for discovering universally optimal quantum circuits that can be reused across diverse platforms.

In summary, nearly two decades after the introduction of quantum Shannon decomposition, quantum circuit compilation and optimization has matured significantly, evolving into an active and highly competitive research arena [330]. The rise of algebraic, symbolic, and graph-based methods such as the  $ZX$ -calculus has substantially advanced our capability to compress circuits beyond what traditional algebraic synthesis methods allowed. More recently, artificial intelligence techniques, particularly reinforcement learning, have further transformed optimization strategies, enabling adaptive algorithms that generalize beyond manually crafted rules and heuristics. Coupled with the development of industrial-strength compilers, these innovations have resulted in substantial reductions in circuit complexity, driven largely by the urgent practical need to execute quantum computations reliably on current noisy quantum hardware. De-

spite these remarkable strides, quantum circuit optimization remains an inherently challenging problem, balancing rigorous guarantees of correctness, scalability, flexibility, and computational efficiency. Addressing these challenges not only advances theoretical understanding but also directly contributes to enabling practical, reliable quantum computing at scale.

### 4.1.2 Motivation

In this chapter, we present a novel circuit optimization methodology applicable to quantum circuits that implement arbitrary unitary transformations and are defined over general qubit topologies. The proposed approach builds on both  $ZX$ -calculus and circuit-level optimization by integrating them within a unified higher-level framework that incorporates optimization, feedback, and search procedures. We focus primarily on reducing the number of two-qubit gates, although the methodology can readily be adapted to minimize alternative cost functions. Prioritizing the reduction of two-qubit gates emphasizes the identification of minimal circuit structures connecting qubits and operations. Beyond reducing circuit complexity and often also algorithmic execution times, this cost metric also directly targets the predominant source of error in contemporary quantum processors, as discussed in § 3.

This emphasis on two-qubit gate minimization is further motivated by recent developments in fault-tolerant quantum computing. While  $T$ -gate count reduction has attracted significant attention over the past decade, due to the traditionally high overhead of magic-state distillation protocols [331, 309, 332, 313, 326], this cost model has recently been challenged [333]. In particular, magic state cultivation [334] has demonstrated that  $T$ -gates can now be implemented at costs approaching that of a single CNOT gate. This shift renews interest in optimizing CNOT count and circuit depth — particularly in architectures with limited connectivity — where routing and qubit storage overhead can become dominant. Therefore, while two-qubit gate minimization remains critical for near-term devices, it is also expected to persist as a key optimization target in large-scale fault-tolerant architectures, where non-Clifford gates no longer constitute the primary resource bottleneck.

Furthermore, beyond its role as a circuit optimization tool, the methodology is also designed to support theoretical insight by enabling the identification of structural patterns within optimized circuits. In favorable instances, such patterns may serve as the basis for extrapolating more abstract algorithmic constructions. While not guaranteed in general, this feature distinguishes the approach from prior methods, as it facilitates a principled pathway from low-level optimization to higher-level algorithmic discovery.

The remainder of this chapter is organized as follows. § 4.2 introduces a normalization procedure that rewrites general circuits into a canonical form composed of alternating layers of single- and two-qubit

gates, which is then used to classify the circuit's entangling-gate structure. § 4.3 presents our metaheuristic optimization routine, based on  $ZX$ -diagram simplification combined with circuit-level transformations and feedback. In addition to minimizing circuit complexity, this approach generates multiple equivalent decompositions with broader applicability across other contexts. § 4.4 illustrates an application to quantum simulation, providing a detailed examination of the inner workings of the method. § 4.5 discusses potential pathways for algorithmic generalization from optimized circuit instances. Finally, § 4.6 summarizes the main findings and outlines future research directions building upon the developed techniques.

## 4.2 Entangling-gate topologies in multi-qubit circuits

Base two-qubit gates, such as the CNOT and CZ gates, play a crucial role in circuit synthesis by connecting qubits and regulating the propagation of entanglement and information throughout the system. The strategic placement of these gates within the circuit directly shapes the interference landscape and impacts its ability to execute complex operations efficiently. Furthermore, the two-qubit gate skeleton of the circuit not only establishes the interconnections between the causal past and future of all the instructions of the algorithm but also defines the connectivity topology of its associated  $ZX$ -diagram representation.

The connectivity topology of the  $ZX$ -diagram linked to a circuit encodes fundamental algebraic symmetries, making it a natural framework for efficient circuit optimization through diagrammatic reasoning [335]. Our optimization methodology, to be detailed in § 4.3, begins by converting an input circuit into its  $ZX$ -graph representation, upon which we apply a suite of heuristic optimization techniques before converting the optimized graph back into a circuit. Therefore, it is useful to identify different circuits whose  $ZX$ -graph retains a similar<sup>3</sup> connectivity structure after this initial conversion, even if their sequences of operations differ, for instance, due to reordering of commuting gates.

To achieve this, we introduce a classification methodology that categorizes circuits with the same number of two-qubit gates according to their *entangling-gate structure* (EGS). The EGS is defined for any multi-qubit circuit as a *categorical* label that encodes the positions of the two-qubit gates applied throughout the execution sequence when the circuit is transcribed into a *normalized* form. Computing the EGS, therefore, requires first applying a *normalization* procedure. This normalization takes a circuit previously expressed in the  $\{R_x(\theta), R_y(\theta), R_z(\theta), CZ\}$  universal gate set basis and segments it into layers of single- and two-qubit gates, where all two-qubit gate layers are composed of CZ gates in a sorted order.

This choice of base two-qubit gate for the normalized form is motivated by their symmetry and com-

---

<sup>3</sup>As will be explained shortly, our classification is not invariant under qubit permutation, whereas a truly graph-based equivalence would be.




Figure 4.1: **Entangling-gate structure calculation.** **a**, Illustration of an example sequence of twenty-two successive instructions in the input circuit. The allowed gate set consists of  $\{R_x(\theta), R_y(\theta), R_z(\theta), CZ\}$  gates. **b**, Normalized circuit representation, segmented into layers of single- and two-qubit gates. **c**, Extraction of the circuit ID string based on its entangling-gate structure. Digits encode the qubits on which CZ gates act, while dashes separate the indices of gates in different layers. The inclusion of a dot between digits is necessary to prevent ambiguity between single- and multi-digit indices, particularly when the circuit involves more than ten qubits.

mutativity. Because CZ gates are symmetric—i.e., their matrix representation is invariant under the permutation of the two qubits they act upon—and mutually commuting, they can be reordered within a two-qubit gate layer without affecting the overall operation. These properties allow us to standardize the two-qubit gate layers by sorting the CZ gates according to the qubit indices where they are applied. This canonical ordering is instrumental for comparing different circuits and identifying equivalent entangling-gate structures that might not be immediately apparent from a superficial inspection.

## Circuit normalization

The normalization procedure transforms an arbitrary circuit, defined over the restricted but universal gate set  $\{R_x(\theta), R_y(\theta), R_z(\theta), CZ\}$ , into a canonical form (see Fig. 4.1b). This canonical form is particularly useful for equivalence comparisons and subsequent analysis, as it guarantees a consistent temporal ordering of operations, clear layer separation, and a unique ordering of two-qubit interactions.

Consider a general circuit described as a *list*<sup>4</sup>  $I$  of single-qubit and two-qubit gate instructions. Algorithm 1 outlines the normalization procedure at a high-level, abstracting away implementation details while retaining its core logical structure. It comprises three phases.

<sup>4</sup>Formally, an ordered sequence of gate instructions.

The first phase (line 1) advances  $R_z$  gates by processing the full instruction sequence and promoting them to the earliest feasible timestep. This is achieved by iteratively swapping each  $R_z$  gate with its preceding instruction, provided the swap does not introduce a conflict. A conflict arises if the preceding instruction is an  $R_x$  or  $R_y$  gate acting on the same qubit, in which case the swap is not performed.

The second phase (line 10) partitions the instruction sequence into distinct sections of SQGs and TQGs, ensuring that single-qubit gates are scheduled before two-qubit gates whenever possible. To enforce this ordering, the algorithm employs a locking mechanism that flag qubits upon scheduling an operation, preventing improper reordering.

The final phase (line 32) establishes a canonical ordering of two-qubit gates within their own sections. The collection of CZ gates is sorted lexicographically: first by the index of the first qubit, and in cases where these indices coincide, by the second qubit index. This guarantees a consistent and unambiguous representation of two-qubit interactions within each layer.

This procedure enforces a standardized temporal ordering of operations and clear layer separation, transforming any circuit in the given gate set into a unique, well-defined form. The resulting normalized form enables the classification of entangling-gate structures and also allows detailed structural pattern comparisons between different circuits.

## Identifier attribution

Once the circuit is expressed in its normalized form, its entangling gate structure (EGS) label is determined by identifying the indices of the qubit pairs on which CZ gates act. These indices are listed in ascending order according to their appearance in the standardized circuit instruction sequence. In this process, single-qubit gates are disregarded, ensuring that two normalized circuits differing only by local transformations along the circuit execution, as encoded by their single-qubit gates, are classified under the same EGS (see Fig. 4.1c).

Algorithm 2 provides a detailed procedure for computing the EGS from the normalized circuit representation. In applications where circuit categorization is the primary objective, a hash function can be applied to the label to shorten its length for convenience. However, when analyzing the circuit structure itself or comparing different circuits, preserving the numerical indices of the qubits provides essential information for identifying structural patterns.

While alternative circuit classification methods could conceivably be developed — such as employing numerical quantifiers instead of categorical labels — this approach serves a dual purpose. It not only groups *structurally equivalent* circuits under the same category but also reveals *qubit connectivity information*.

---

**Algorithm 1:** Quantum Circuit Normalization

---

**Input:**

$I$ : Instructions list for a quantum circuit  $Q$  expressed with the  $\{R_x(\theta), R_y(\theta), R_z(\theta), CZ\}$  gate set

**Output:**

$I'$ : Normalized instructions list associated with  $Q$  sectioned into SQGs and sorted TQGs

```
1 // Phase 1: Pull  $R_z$  gates to the earliest permissible timestep;
2 for  $i \leftarrow 0$  to  $|I| - 1$  do
3   while  $i \geq 1$  do
4     if  $I[i]$  is an  $R_z$  gate then
5        $q \leftarrow$  qubit index on which  $I[i]$  acts;
6        $q' \leftarrow$  qubit index on which  $I[i - 1]$  acts;
7       if  $I[i - 1]$  is a CZ gate or ( $I[i - 1]$  is a single-qubit gate acting on  $q' \neq q$ ) then
8         Swap  $I[i]$  and  $I[i - 1]$  in  $I$ ;
9      $i \leftarrow i - 1$ ;
10 // Phase 2: Sectioning;
11  $I' \leftarrow []$ ;
12 while  $|I| > 1$  do
13    $L_{SQ}, L_{TQ} \leftarrow [], []$   $\triangleright$  Empty lists to store instructions for SQG and TQG sections
14    $K_{SQ}, K_{TQ} \leftarrow \emptyset, \emptyset$   $\triangleright$  Empty sets to lock qubits from receiving any more SQGs or TQGs
15    $i \leftarrow 0$ ;
16   while  $i < |I|$  do
17     if  $I[i]$  is a single-qubit gate then
18        $q \leftarrow$  qubit index on which  $I[i]$  acts;
19       if  $q \notin K_{SQ}$  then
20         Append  $I[i]$  to  $L_{SQ}$  and remove it from  $I$ ;
21          $i \leftarrow i - 1$ ;
22       else
23         Add  $q$  to  $K_{TQ}$ ;
24     else if  $I[i]$  is a two-qubit gate then
25        $(q_1, q_2) \leftarrow$  qubit indices on which  $I[i]$  acts;
26       if  $q_1, q_2 \notin K_{TQ}$  then
27         Append  $I[i]$  to  $L_{TQ}$  and remove it from  $I$ ;
28          $i \leftarrow i - 1$ ;
29       Add  $q_1, q_2$  to  $K_{SQ}$ ;
30      $i \leftarrow i + 1$ ;
31   Append the elements of  $L_{SQ}$  and  $L_{TQ}$  to  $I'$ , in this order;
32 // Phase 3: Sorting CZ Gates;
33 foreach TQG section in  $I'$  (identified by inspecting neighboring gates) do
34   Extract all CZ gates in the section;
35   Sort the extracted CZ gates lexicographically by qubit indices;
36   Replace the original CZ gates in the section with the sorted sequence;
37 return normalized instruction list  $I'$ ;
```

---

---

**Algorithm 2:** Entangling-gate structure classification

---

**Input:**  $I$ : Normalized instructions list of a quantum circuit  
**Output:**  $id$ : String with ordered qubit indices of CZ gates separated by "." and layers by "-"  
1  $idx \leftarrow []$   $\triangleright$  Empty list to store qubit indices from each TQG layer  
2 **for**  $i \leftarrow 0$  **to**  $|I| - 1$  **do**  
3     **if**  $i = 0$  **then**  
4         **if**  $I[i]$  is a CZ gate **then**  
5              $s \leftarrow$  string with the ordered qubit indices on which  $I[i]$  acts separated by ".";  
6             Append  $s$  to  $idx$ ;  
7     **else**  
8         **if**  $I[i]$  is a CZ **and**  $I[i - 1]$  is not a CZ **then**  
9              $s \leftarrow$  string with the ordered qubit indices on which  $I[i]$  acts separated by ".";  
10             Append  $s$  to  $idx$ ;  
11         **else if**  $I[i]$  is a CZ **and**  $I[i - 1]$  is a CZ **then**  
12              $s \leftarrow$  string with the ordered qubit indices on which  $I[i]$  acts separated by ".";  
13             Concatenate "." followed by  $s$  to the last element of  $idx$ , i.e.,  $idx[-1]$ ;  
14  $id \leftarrow$  Join the elements of  $idx$  using "-" as the separator;  
15 **return**  $id$ ;

---

This feature can be leveraged for further structural analysis, enabling the detection of regularities and general templates that may inform broader algorithmic insights.

Another key advantage of this method is that it enables EGS classification solely from the classical description of the circuit, without requiring explicit representation of the unitary at any stage. This is particularly important for comparing circuits when the number of qubits employed by the algorithm becomes considerably large.

Notably, the EGS classification derived from the normalized circuit form differs from one based solely on ZX-graph structural equivalence, as it is not invariant under qubit permutations, whereas a graph-based analyses would be. This distinction highlights the specificity of this approach in preserving the exact unitary representation and qubit ordering of the circuit, an often desirable feature, rather than abstracting away such details through graphical equivalences.

As we will explain in the next section, this EGS classification is essential to our optimization methodology as it functions collaboratively with heuristics to guide the search.

### 4.3 Adaptive circuit optim. via dual representation feedback

The ZX-calculus provides a powerful framework for the algebraic manipulation of quantum circuits at a high level of abstraction by transforming them into ZX-diagrams and applying various diagrammatic

rewrite simplification rules. Once the diagram is simplified, the final step is to extract an optimized circuit from it. However, optimal circuit extraction remains a significant bottleneck in this workflow.

While  $ZX$ -diagrams can be significantly simplified through the various rules of the calculus, the conventional process of converting the final diagram back into a circuit is non-trivial and often suboptimal in CNOT count [315]. The full simplification of  $ZX$ -diagrams — minimizing nodes and edges — tends, in principle, to translate into a circuit with fewer CNOTs, as the number of edges correlates with multi-qubit interactions. However, the critical issue is that this only works if the optimized diagram still allows for efficient circuit extraction, which in turn depends on its topological properties, most notably the preservation of generalized flow (*gflow*). *Gflow* is a structural property that ensures deterministic measurement-based quantum computation and enables a polynomial-time circuit extraction procedure [315].

When  $ZX$ -diagram simplifications are restricted to *gflow*-preserving rewrites, circuit extraction remains systematic and efficient, following a causal sequence of graph operations that guarantee correct unitary reconstruction. This ensures that  $T$ -count reductions are reliably achieved. However, *gflow*-preserving transformations may prevent further reductions in CNOT count, as certain *gflow*-preserving rewrites do not fully simplify the diagram. If the diagram could have been further simplified without *gflow* constraints, a circuit with fewer CNOT gates might have been obtained, but at the cost of making the extraction procedure inefficient. Conversely, aggressive  $ZX$ -diagram simplifications that do not preserve *gflow* can render circuit extraction exponentially costly or even infeasible, with results indicating that optimal extraction can be  $\#P$ -hard in general cases [336]. For this reason, practical optimization workflows have been prioritizing *gflow*-preserving transformations to maintain computational efficiency at the cost of forgoing potential additional reductions in two-qubit gate count.

To address this challenge, heuristic methods have been developed to improve  $ZX$ -based circuit optimization more robustly, with the goal of reducing CNOT or  $T$  gate usage. A common approach is to operate directly on  $ZX$ -diagrams, simplifying them in a way that facilitates more efficient extraction. Some methods employ graph-theoretic heuristics to guide diagram simplifications while preserving extractability via Gaussian elimination-based methods [122]. Others focus on local transformations within the  $ZX$ -diagram, leveraging structural changes that significantly impact the complexity of the extracted circuit [337]. Another approach specifically targets Hadamard wire reduction, optimizing the structure of the diagram before extraction to indirectly minimize two-qubit gate count [316, 338]. In contrast, other methods operate at the level of the extracted circuit, optimizing its structure by modifying CNOT gates through simulated annealing while preserving the  $ZX$  phase circuit representation [339].

In an alternative approach [319], causal flow (*cflow*) — a structural property stronger than *gflow* — is

strictly preserved, which ensures a direct and trivial extraction of the circuit. This enables exact and transparent tracking of how each diagrammatic transformation affects the two-qubit gate count. In particular, by leveraging a generalization of the neighbour unfusion rule, this structured method can systematically achieve two-qubit gate reductions while maintaining efficient circuit extraction.

All these methods and heuristics—including genetic algorithms, simulated annealing, greedy optimization, and random search, among others—have demonstrated success in various settings, but each exhibits inherent biases and limitations. For instance, genetic algorithms promote diversity in search space exploration but may struggle with fine-tuned local improvements. Greedy methods converge quickly but risk getting trapped in suboptimal plateaus. Simulated annealing balances exploration and exploitation but its tunable randomness may fail in high-complexity settings. As a result, even with well-tuned parameters, these heuristics often reach local optima under their optimization landscape and saturate in performance after a certain runtime, with additional computational effort yielding diminishing returns.

However, a key empirical observation in our work is that when a single heuristic stagnates, an alternative one may still find improvements that the first failed to uncover. More importantly, we find that iteratively feeding some of the extracted circuits back into a  $ZX$ -diagram and re-extracting them after a new round of simplification, possibly using a different heuristic, can result in further complexity reductions significantly beyond what a single heuristic can achieve given additional runtime. This insight forms the foundation of our quantum circuit optimization via dual representation feedback, where circuits are alternately simplified in both the  $ZX$ -diagram and circuit representations through a metaheuristic search and optimization procedure driven by adaptive feedback.

Essential to this approach is to employ a *portfolio of heuristics*, each implementing a distinct strategy for  $ZX$ -diagram simplification prior to circuit extraction. At each iteration, a given circuit is transformed into its  $ZX$  representation and processed independently by a collection of  $N$  heuristics, each producing a distinct simplified diagram. These  $ZX$ -diagrams, corresponding to alternative heuristic pathways, are then separately extracted back into  $N$  circuits, which all undergo the same circuit-level optimization pass. This results in a final set of candidate solutions per iteration, ready to be considered for feedback.

A significant challenge in implementing such a feedback mechanism is the diversity of extracted circuits. Given a simplified  $ZX$ -diagram, there exist multiple valid quantum circuits that represent the same unitary transformation. This diversity creates both an opportunity and a challenge: it provides a rich space for exploration when their structures differ significantly, but it also introduces redundancy, as many extracted circuits may be structurally equivalent under local transformations. To address this, we employ the entangling-gate structure (EGS) classification method, as introduced in the previous section, to sys-

tematically identify and remove redundant circuit structures before feeding extracted circuits back into the optimization process. Crucially, this method prunes the search space while preserving essential solution diversity, enabling a focused and effective strategy.

Once structurally equivalent circuits are filtered out, the remaining ones must be selected for feedback. As a result, selecting which circuits to reintegrate into the feedback loop becomes a discrete search problem, where the challenge lies in effectively navigating the space of extracted candidates. A naïve approach would be to select circuits at random, but this would be suboptimal, as it fails to leverage implicit structural information in the search space. While the space of extracted circuits under feedback is not strictly well-defined, empirical evidence suggests that certain circuit structures are more conducive to further simplification, while others tend to lead to stagnation. An effective feedback-driven method must therefore employ an adaptive exploration policy that balances exploitation, where circuits that have previously demonstrated optimization potential are prioritized, and exploration, where structurally diverse candidates are tested to avoid premature convergence.

#### 4.3.1 Best-first search as a metaheuristic

Unlike conventional search problems with well-defined, deterministic state transitions [340], the space of extracted circuits under iterative heuristic feedback is *partially* stochastic and lacks a hierarchical state transition structure. The first issue arises because each round of circuit optimization employs a combination of deterministic and stochastic heuristics. Deterministic heuristics always yield the same output for a given input, while stochastic ones can generate a diverse range of circuit structures, introducing variability into the search trajectory. This effectively increases the branching factor at each search node, making the exploration space significantly more expansive.

The second issue occurs because the search process may rediscover previously encountered solutions. This indicates that the search unfolds in a directed graph where nodes can have multiple predecessors, rather than in a tree with a unique predecessor-successor relationship. These transitions introduce computational inefficiencies, but we mitigate their effect by storing each structurally equivalent circuit only once. This prevents redundant re-expansion of previously visited nodes, eliminating search loops and effectively restoring a minimal exploration tree in the space of *discovered* circuit dependencies.

Despite the challenge of navigating this space, empirical observations suggest that certain circuit structures consistently lead to better optimizations when fed back into the *ZX* simplification-extraction loop. This indicates that the search space is not entirely random but instead exhibits an implicit structure, where some circuit classes are more amenable to simplification than others. However, due to the heuristic

nature of transformations, the iterative process may occasionally cause a temporary increase in the cost function before ultimately converging to a better solution—an observation also noted and emphasized from a different perspective in previous work [324].

To explore solutions efficiently, we introduce a best-first search (BFS) [341] metaheuristic that prioritizes circuits according to a predefined cost function. Formally, let  $\mathcal{C}$  denote the space of extracted quantum circuits, where each circuit  $C \in \mathcal{C}$  is characterized by its CNOT count, circuit depth, and entangling-gate structure (EGS). Given an initial circuit  $C_0$ , our iterative optimization process seeks to find an optimal circuit  $C^* \in \mathcal{C}$  that minimizes a lexicographic cost function defined as:

$$f(C) = (N_{\text{CNOT}}(C), \text{depth}(C)), \quad (4.2)$$

where the optimization problem is expressed as

$$C^* = \underset{C \in \mathcal{C}}{\text{lexmin}} f(C). \quad (4.3)$$

Here, the lexicographic ordering prioritizes minimizing the number of CNOT gates first, and among circuits with the same CNOT count, the secondary objective is to minimize circuit depth.

The main data structure maintained by BFS is a *priority queue*  $Q$ , where circuit candidates for further exploration are dynamically sorted according to the lexicographic cost function ordering, prioritizing the most promising candidates found so far. The search proceeds iteratively as follows: *(i)* Initialize  $Q$  with the seed circuit  $C_0$ , or a collection of circuits if available. *(ii)* Extract from  $Q$  the circuit

$$C_t = \underset{C \in Q}{\arg \text{lexmin}} f(C) \quad (4.4)$$

that lexicographically minimizes the cost function. *(iii)* Convert  $C_t$  into a ZX-diagram, apply  $N$  heuristic-based transformations, and extract a new set of circuits, each of which is further optimized using circuit-based heuristics to obtain the output set of circuits in the current iteration:

$$\mathcal{N}(C_t) = \{C_t^1, C_t^2, \dots, C_t^N\}. \quad (4.5)$$

*(iv)* Filter out redundant circuits using Algorithm 2 to check for structural equivalence using their EGS identifier:

$$\mathcal{N}_{\text{filtered}}(C_t) = \{C_t^i \mid C_t^i \notin \mathcal{H}\}, \quad (4.6)$$

where  $\mathcal{H}$  is the history of previously encountered structurally equivalent circuits. Together, steps (iii) and (iv) induce an implicit search tree over  $\mathcal{C}$ . (v) Update the priority queue with the newly discovered circuits:

$$Q \leftarrow (Q \cup \mathcal{N}_{\text{filtered}}(C_t)) \setminus \{C_t\}; \quad (4.7)$$

(vi) Repeat (ii-v) until a termination criterion is met, such as convergence or resource constraints.

Checking circuit equivalence in step (iii) is currently performed by directly and explicitly computing the corresponding unitary matrices, up to a global phase. While this approach ensures correctness—subject to floating-point precision—it scales poorly with the number of qubits  $n$ . When this becomes a limiting factor for wider circuits, replacing this verification step with more scalable methods, such as those based on the ZX-calculus or formal symbolic techniques, may offer a more efficient alternative. However, for use cases involving the application of this optimization framework to primitive circuit discovery and algorithmic generalization from small- $n$  instances, direct unitary comparison may remain sufficient.

While standard BFS is a greedy strategy that primarily prioritizes *exploitation* by focusing on local optimization, due to the inherent stochasticity of heuristic-driven circuit extraction, BFS indirectly benefits from *exploration*, as different heuristics may produce structurally distinct circuits even when applied to the same input. This allows BFS to escape local minima. To further encourage structured exploration, we also introduce a  $k$ -beam extension to BFS, wherein the search simultaneously expands the top  $k$  lowest-cost circuits at each iteration rather than only the single best candidate. This results in an update rule:

$$Q \leftarrow (Q \cup \bigcup_{j=1}^k \mathcal{N}_{\text{filtered}}(C_j)) \setminus \{C_1, C_2, \dots, C_k\}. \quad (4.8)$$

This modification introduces a trade-off between focused exploitation and broader search horizons: while standard BFS continually refines the most promising circuit at each step,  $k$ -beam BFS maintains a diverse pool of candidate circuits, thereby mitigating the risk of premature convergence. The degree of exploration is controlled by the parameter  $k$ , with larger values promoting structural diversity in the search trajectory and smaller values emphasizing depth-first refinement within locally optimal regions. By adjusting  $k$ , the search strategy can be tailored to balance exploration and exploitation, ensuring more efficient traversal of the circuit optimization landscape.

Algorithm 3 presents the pseudocode for this BFS metaheuristic search and optimization strategy, including the  $k$ -beam search variant. Unlike circuit optimization approaches focused on runtime compilation, this particular implementation stores all circuits on disk during execution rather than retaining only the most optimized one. This is done to allow users to select relevant circuits post-execution for further

---

**Algorithm 3:** Metaheuristic best-first search for circuit optimization via dual representation feedback

---

**Input:** A set of seed circuits  $\mathcal{S}$ , heuristics portfolio  $\mathcal{H}$ , maximum iterations  $T$ , beam width  $k$

**Output:** Circuits stored on disk in files labeled with (CNOT count, depth, EGS signature)

```
1 Initialize a priority queue  $Q$   $\triangleright$  Implemented as a min-heap for best performance
2 foreach  $C \in \mathcal{S}$  do
3   Push  $(f(C), C)$  onto  $Q$ ;
4 Initialize visited set  $\mathcal{V} \leftarrow \{\text{EGS}(C) : C \in \mathcal{S}\}$   $\triangleright$  Stores EGS signatures to detect duplicates
5 foreach  $C \in \mathcal{S}$  do
6   Write  $C$  to disk with a filename encoding  $(\text{CNOT}(C), \text{depth}(C), \text{EGS}(C))$ ;
7 for  $t \leftarrow 1$  to  $T$  do
8   if  $Q$  is empty then
9     stop  $\triangleright$  No more circuits to explore; all circuits stored on disk
10     $\{C_1, \dots, C_m\} \leftarrow$  dequeue up to the top  $k$  circuits from  $Q$   $\triangleright m \leq k$ , if  $|Q| < k$ 
11    foreach  $C_j \in \{C_1, \dots, C_m\}$  do
12      Generate  $N$  new candidate circuits  $\mathcal{N}(C_j)$  using each of the  $N$  heuristics in  $\mathcal{H}$ ;
13      foreach  $C' \in \mathcal{N}(C_j)$  do
14        if  $\text{EGS}(C') \notin \mathcal{V}$  then
15           $\mathcal{V} \leftarrow \mathcal{V} \cup \{\text{EGS}(C')\}$ ;
16          Push  $(f(C'), C')$  onto  $Q$ ;
17          Write  $C'$  to disk with a filename encoding  $(\text{CNOT}(C'), \text{depth}(C'), \text{EGS}(C'))$ ;
18 stop  $\triangleright$  All discovered circuits have been saved to disk
```

---

analyses, highlighting our focus on permanent circuit discovery. Additionally, multiple decompositions with minimal CNOT counts but varying EGS can find use in coherent error mitigation protocols. For instance, a transpiler can deliberately choose to execute decompositions that minimize CNOT gates between the noisiest qubit pairs, among other strategies discussed in § 5.

In summary, this adaptive circuit optimization metaheuristic leverages an iterative feedback mechanism between ZX-diagram simplification, circuit extraction, and circuit-level optimization, systematically guided by  $k$ -beam best-first search. The framework naturally balances exploration and exploitation, ensuring both fast progression towards optimized solutions and sustained global search. By combining multiple heuristics within a unified framework, it is guaranteed to perform at least as well as the best heuristic in the portfolio. Beyond this worst-case guarantee, the interaction among diverse heuristics can induce synergistic effects that enable the method to escape local minima that may trap individual strategies, thereby substantially exceeding the performance of any single heuristic in a wide range of scenarios. The parallelizable nature of the search allows independent circuit expansions across multiple heuristics and node candidates, making the method highly scalable. Finally, unlike reinforcement learning-based approaches, this method does not require compute-intensive training, making it efficient and adaptable for a wide range

of circuit optimization tasks.

### 4.3.2 Technical implementation

The methods described in this chapter were implemented in Python, using a combination of open-source libraries. The resulting software framework, currently under development and not yet publicly released, is referred to as QCMotherLode<sup>5</sup>. In particular, Qiskit [342] and PyZX [343] provide core functionality for manipulating quantum circuits and ZX-diagrams, respectively. The OpenQASM language<sup>6</sup> [344] is used to represent and store all circuits generated during optimization, ensuring compatibility with standard quantum computing workflows. The procedures for normalizing circuits (Algorithm 1) and extracting their EGS (Algorithm 2) are also implemented on top of the OpenQASM representation.

Initial circuits to seed the optimization process can be generated using unitary synthesis functions in libraries such as Qiskit [342], BQSKit [345], or tket [346, 314], among others. In some cases, generating multiple decompositions of a given unitary can improve search space exploration by increasing circuit diversity.

The implementation provides access to a set of ZX-calculus-based simplification heuristics sourced from the PyZX repository and two other forks of this repository [347, 347] based on the works in Refs. [316, 319], comprising the ten methods listed in Table 4.2. In addition, the converted ZX-diagram is first passed through the `teleport_reduce()` method before applying the heuristics. These heuristics serve as the building blocks for constructing the portfolio of heuristics  $\mathcal{H}$  used during ZX-diagram simplification. The portfolio is user-defined, allowing for flexible selection of heuristics, including repeated instances of the same heuristic with different parameter settings to generate diverse solutions at each iteration. However, empirical observations suggest that parameter variation within a single heuristic has minimal impact on performance, while the feedback mechanism itself plays a decisive role in continuously improving optimization results. Although the current implementation primarily uses heuristics from these three sources, it is designed to be easily extensible, enabling the integration of new heuristic tools into subsequent versions of the portfolio as they become available in future works.

Once extracted, all circuits undergo a standardized circuit-level optimization pass in Qiskit before feedback. When qubit connectivity constraints are imposed, this optimization pass may introduce randomness due to methods such as the SABRE routing algorithm [123, 350], which employs Monte Carlo

---

<sup>5</sup>The name QCMotherLode reflects the intended role of the framework as a tool for unearthing useful quantum circuits for storage and subsequent reuse, in analogy with a rich source of valuable material.

<sup>6</sup>OpenQASM (Open Quantum Assembly Language) is a hardware-agnostic language for specifying quantum circuits through a simple, sequential list of operations—including gate instructions, measurements, and control directives.

| Codeword | Heuristic ZX-diagram function                     | Deterministic |
|----------|---------------------------------------------------|---------------|
| BO       | <code>basic_optimization</code> [343, 348]        | yes           |
| FR       | <code>full_reduce</code> [343, 348]               | yes           |
| GA       | <code>GeneticOptimizer</code> [343, 348]          | no            |
| GS       | <code>greedy_simp</code> [316, 347]               | yes           |
| GSN      | <code>greedy_simp_neighbors</code> [316, 347]     | yes           |
| RS       | <code>random_simp</code> [316, 347]               | no            |
| RSN      | <code>random_simp_neighbors</code> [316, 347]     | no            |
| SAS      | <code>simulated_annealing_simp</code> [316, 347]  | no            |
| SASN     | <code>sim_anneal_simp_neighbors</code> [316, 347] | no            |
| F2QS     | <code>flow_2Q_simp</code> [319, 349]              | yes           |

Table 4.2: **Heuristic function portfolio for ZX-diagram simplification.** The table lists the codes, corresponding Python function implementations, and their classification as deterministic or stochastic, refers to the possible EGS of the circuits they return.

techniques that can yield different results in different runs. Feeding back the resulting circuit in the next iteration can refocus the ZX-diagram simplification on the constrained graph topology. However, circuit extraction from the ZX representation in this situation must also incorporate the routing methods available within PyZX to ensure compliance with the required topology in both steps of the optimization [318].

A key design choice is that circuits are stored on disk rather than retained in memory. This reflects the primary focus of the method on circuit discovery rather than runtime compilation. For this reason, although circuits generated with the same EGS during the search are not reintroduced into the priority queue, they are preserved in the filesystem to enable post-execution analysis of their differing single-qubit gate parameterizations. This is particularly relevant for: (i) Investigating single-qubit gate angles to enable precise synthesis into alternative gatesets (e.g., Clifford+ $T$ ); (ii) Facilitating data-driven discovery of circuit structures through subsequent artificial intelligence techniques.

While the implementation is currently design for circuit discovery, it can be adapted for real-time compilation as well, though further evaluation would be required to assess its runtime performance. Similarly, while optimization is performed based on the lexicographic cost function defined in Eq. (4.2), the framework readily supports substitution with alternative cost functions, such as  $T$ -gate minimization.

Designed with modularity and adaptability in mind, the codebase aims to enable the seamless integration and removal of methods from prior research while allowing for the incorporation of future advancements in quantum compilation techniques. By providing a flexible environment for conducting comprehensive metaheuristic searches, it allows users to combine heuristic strategies into customized optimization pipelines tailored to specific circuit synthesis tasks.

## 4.4 A practical example in quantum simulation

The preparation of Valence-Bond-Solid (VBS) states, a class of entangled quantum many-body states, is a crucial task in quantum simulations of strongly correlated spin systems, as demonstrated by the methods recently introduced in Ref. [49]. A well-known example of a VBS state is the exact ground state of the Affleck-Kennedy-Lieb-Tasaki (AKLT) model. In a VBS state, localized spins interact via exchange couplings to form spin-singlet pairs,  $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$ . On a honeycomb lattice, each site hosts three spin-1/2 degrees of freedom, and each spin is coupled to a different neighboring site via a single valence bond. To recover the physical spin-3/2 Hilbert space, the three-qubit space at each lattice site must be projected onto a permutation-symmetric subspace. This requires implementing the non-unitary symmetrization operator  $\mathcal{S}$ , which enforces permutation symmetry among the three underlying spin-1/2 components — a key challenge in practical realizations.

On quantum hardware,  $\mathcal{S}$  can be implemented probabilistically using ancilla-assisted post-selection via a Hadamard test procedure, which requires compiling a controlled- $e^{-i\pi\mathcal{S}}$  unitary. This unitary acts on its three target qubits as an  $8 \times 8$  matrix given by

$$e^{-i\pi\mathcal{S}} = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & -\frac{2}{3} & 0 & -\frac{2}{3} & 0 & 0 & 0 \\ 0 & -\frac{2}{3} & \frac{1}{3} & 0 & -\frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & -\frac{2}{3} & -\frac{2}{3} & 0 \\ 0 & -\frac{2}{3} & -\frac{2}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{2}{3} & 0 & \frac{1}{3} & -\frac{2}{3} & 0 \\ 0 & 0 & 0 & -\frac{2}{3} & 0 & -\frac{2}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix} \equiv e^{i\pi\mathcal{V}}, \quad (4.9)$$

as given in Eq. (19) of Ref. [49].

In the broader VBS-state preparation procedure, this method enables a shallow-circuit approach. To achieve this, the controlled version of  $\mathcal{V}$  must be implemented. However, since the resulting controlled- $\mathcal{V}$  matrix is a nontrivial  $16 \times 16$  operation, its decomposition can easily become a bottleneck in terms of gate complexity. Even using the optimized quantum Shannon decomposition (QSD-optimized), the circuit requires 63 CNOT gates, which is already prohibitively deep for many NISQ devices. Therefore, obtaining a circuit decomposition that is both accurate and gate-efficient is essential for reducing the overall cost of preparing spin-3/2 VBS states on near-term devices.

In what follows, we demonstrate how the method described in this chapter can optimize the decomposition of the controlled- $\mathcal{V}$  gate. We assume full connectivity among all qubit pairs. As mentioned, the best deterministic method available at the time of Ref. [49] was QSD-optimized, yielding a circuit with 63 CNOT gates. However, other methods exist across different software packages that, despite producing less efficient results, can serve as seed circuits that provide diverse starting points for our metaheuristic search. One such method generated a circuit with 82 CNOT gates, which was subsequently refined using our optimization procedure as described below. In fact, this seed circuit led to the best decomposition obtained by our metaheuristic algorithm. Henceforth, when referring to the seed circuit, we refer to this particular one, and we will describe the optimization process that reduced its CNOT count from 82 to an impressive 22 gates.<sup>7</sup>

We performed two optimization runs starting from this seed circuit to compare performance under different search strategies. The first strategy emphasized broad exploration by applying multiple heuristic methods from our metaheuristic toolbox over a limited number of iterations. The second prioritized a greater number of iterations while restricting the heuristic toolbox to faster-executing methods. In both cases, we used a beam width of  $k = 1$ , resulting in sequential node expansion. It should be noted that different heuristics may occasionally generate circuits with the same entangling-gate structure, in which case they are considered a single unique circuit.

In the first optimization run, we incorporated a comprehensive set of heuristics into the circuit extraction toolbox for ZX-diagrams, namely: FR, GA, GS, GSN, RS, RSN, SAS, and SASN. Additionally, each case incorporated circuit-based simplification passes applied after circuit extraction, as implemented in qiskit's `transpile` function with `optimization_level=3`. The optimization was set to run for a total of 50 iterations. The lowest CNOT count achieved was 26, with a circuit depth of 73, found at iteration 38, at which point the number of distinct entangling-gate circuit IDs discovered had reached 149. By the end of the 50 iterations, the total number of different IDs identified had increased to 185. The exploration tree of this optimization history is depicted in Fig. 4.2a.

For the second optimization run, we restricted the heuristic toolbox to only the faster methods: FR, GS, GSN, RS, and RSN. The same seed circuit was used as the starting point, and the run continued until it was manually aborted at iteration 8666. The lowest CNOT count obtained was 22, achieved at iteration 3030, in a circuit with depth 65, when the number of distinct circuit-IDs found had reached 4964. At iteration 3031, the depth of this circuit was reduced to 61 while maintaining the same CNOT count. By the end, seven distinct entangling-gate circuit IDs were found for the optimized circuit with CNOT count

---

<sup>7</sup>This result was obtained after the publication of Ref. [49], using an improved version of the optimization pipeline.




Figure 4.2: **Exploration path in two best-first search example runs.** The unitary decomposition of controlled- $\mathcal{V}$  is considered. **a**, A run with 50 iterations, reducing the number of CNOT gates from 82 to 26. **b**, A run with 8666 iterations, further reducing the number of CNOT gates from the same initial count of 82 to 22. The path in the tree from the root node to the first node where a circuit with the minimal CNOT count was found is marked in red. This node has 142 ancestor nodes, indicating that 142 iterations were required before this circuit was discovered. Its parent node had a CNOT count of 25. All heuristics in the portfolio were used to obtain the nodes connecting the minimum to the root. All nodes with the same minimal gate count of 22 CNOT gates, added after the first such circuit was found, are also colored in red. In both cases, exactly one node is selected for expansion per iteration, with all child nodes expanded in parallel. Details regarding the methods used and the relevant performance metrics are provided in the main text. Nodes are colored according to their iteration number.

22, providing multiple alternative implementations with the same CNOT cost.

This optimization procedure thus achieved a substantial reduction in circuit complexity. Starting from an initial CNOT count of 82, the final circuit contained only 22 CNOT gates, representing a 73.2% reduction from the seed circuit and a 65.1% reduction relative to the exact synthesis output produced by QSD-optimized. The explicit OpenQASM instruction sequence for one of these optimized decompositions is provided in Listing 3.

By the end of the second run, at the moment of manual termination, the total number of distinct circuit IDs discovered had reached 14074. These circuits were distributed across different CNOT counts, with their histogram at the final iteration shown in Fig. 4.3d. If all single-qubit gate variations were counted as distinct circuits, the number of circuits found would be significantly larger.

Empirical observations suggest that the procedure generates circuits with novel entangling-gate structure at an approximately linear rate with respect to the number of iterations. This behavior is depicted in




Figure 4.3: **Statistics for the second best-first search example run.** As described in the main text and Fig. 4.2, the unitary decomposition of controlled- $\mathcal{V}$  is considered. **a**, The number of novel circuits discovered per iteration, classified by their entangling-gate structure ID, ranges from 0 to 5, as five heuristic methods are employed. **b**, The total number of unique discovered circuits accumulates linearly with the iteration count **c**, The number of unexplored circuits in the queue also exhibits linear growth over time. **d**, A histogram showing the distribution of distinct entangling-gate structures as a function of circuit CNOT count, corresponding to the complete set of results obtained by the last iteration.

Fig. 4.3**b**, which shows the growth of unique entangling-gate circuit IDs over the course of the optimization. Similarly, the queue length, representing still unexplored circuits, also exhibits a linear trend, reaching 5408 by iteration 8666, Fig. 4.3**c**. At each iteration, between 0 and 5 new circuit IDs were identified, with the rate of discovery remaining nearly constant (Fig. 4.3**a**).

Fig. 4.3**d** shows the histogram of the total number of distinct entangling-gate structures for each CNOT-count value accumulated by the final iteration value of 8666. The histogram peaks at 27, exhibiting a smooth decay in both directions, with a sharper decline in the negative direction. Examining whether the evolution of the distribution of distinct circuits — i.e., their entangling-gate structure IDs — across iterations can inform a stopping criterion for the search procedure warrants further study.

This application demonstrates the effectiveness of the presented metaheuristic in systematically reducing the number of two-qubit gates and achieving a more compact implementation of critical building blocks for VBS-state preparation. This illustrates how structured unitaries in quantum many-body physics can be efficiently compiled, overcoming the limitations of generic quantum circuit synthesis software.

## 4.5 Towards scalable algorithmic gate discovery

The most prevalent paradigm for quantum algorithm development begins with its ingenious high-level conceptual formulation, followed by the construction of the corresponding quantum circuits to realize those underlying ideas. Canonical algorithms such as Shor's, Grover's, and quantum phase estimation all emerged from identifying quantum resources — such as the quantum Fourier transform, amplitude amplification, or eigenvalue estimation — that could address specific computational problems before designing the circuits to implement them.

In contrast, a different route for developing quantum algorithms starts from a unitary transformation and seeks to decompose it into a sequence of elementary gates using either exact or approximate quantum circuit synthesis techniques. This *unitary-first* approach is typically not used to *design* algorithms in the abstract, but rather to implement known or desired quantum transformations, often arising from physical or mathematical models. The algorithmic goal is embedded in the unitary operation itself (or a family thereof), and the task becomes one of constructing efficient quantum circuits to realize those operations.

While this approach is less commonly associated with the original discovery of major quantum algorithms, it nonetheless plays a central role in several important settings. A paradigmatic example is Hamiltonian simulation, where the algorithm consists of approximating the unitary evolution  $\exp(-iHt)$  for some Hermitian operator  $H$ . This has led to a number of algorithmic developments: Trotter-Suzuki product formulas, Taylor series methods, quantum signal processing, and qubitization.

Efficient algorithms for implementing many useful multi-qubit quantum operations are often unknown, which underscores the importance of circuit optimization in the standard pipeline for realizing quantum algorithms [330]. However, quantum circuits inherently encode the algorithms they implement. As a result, even in the absence of an explicit algorithmic construction for a given unitary decomposition, it seems possible, in principle, to reverse-engineer the underlying algorithm directly from the circuit. Although computationally demanding and far from trivial, such a reverse-engineering approach could, under favorable conditions, yield novel insights or even groundbreaking general quantum algorithms.

Building on this perspective, the pattern-matching methodology we begin to implement for identifying general quantum algorithms from numerically optimized quantum circuits involves the following steps: (i) Select classes of unitary operations whose scalable decompositions remain unknown or inadequately understood; (ii) Numerically synthesize and thoroughly optimize explicit gate-level circuit implementations for small qubit counts (e.g., 3–10 qubits); (iii) Analyze these implementations to identify recurring structural patterns and gate configurations, while formulating general hypotheses regarding the  $n$ -qubit scalability of the identified decompositions; (iv) Rigorously test these hypotheses through formal meth-

ods and theoretical analysis, aiming to establish provable general quantum algorithms.

Several steps in this pipeline must be substantially well developed for the methodology to succeed. The first critical one is the availability of a robust circuit optimization framework specifically tailored to this task in step (ii). Such a framework should be capable of generating multiple optimized circuit decompositions for a given unitary and organizing them into a standardized canonical form. This canonicalization step is essential for eliminating redundancies and enabling systematic comparison of structural patterns across varying qubit counts from an underlying circuit family.

The QCMotherLode framework is well aligned with these requirements, as it outputs multiple circuit decompositions and assigns an EGS classification label to each of them. While further studies are warranted to evaluate its performance in comparison with alternative quantum circuit optimization frameworks and to assess its behavior under different exploration policies aimed at optimizing the trade-off between exploration and exploitation, we now present evidence suggesting that, even with the current BFS metaheuristic policy, QCMotherLode already proves valuable for circuit discovery and generalization.

Specifically, we evaluated our framework by optimizing circuits derived from three distinct decompositions of a multi-controlled NOT gate under the assumption of full qubit connectivity—also commonly referred to as multi-controlled Toffoli gate. With  $m$  control qubits and one target qubit, these gates are denoted  $C^mX$ ,  $\text{Tof-Barenc}_m$ , and  $\text{Tof-NC}_m$ . The first gate uses only  $m + 1$  qubits, while the last two require  $m - 2$  ancillary qubits in addition to control and target qubits. The difference between the latter two is that  $\text{Tof-Barenc}_m$  permits ancilla qubits to begin in arbitrary states [111], whereas the Nielsen-Chuang implementation  $\text{Tof-NC}_m$  explicitly requires ancilla qubits to be initialized in the computational basis state  $|0\rangle$  [203]. Both of these decompositions ensure ancillary qubits are restored to their original states upon gate completion.

The minimal CNOT gate count obtained by QCMotherLode for each resulting circuit is reported in Table 4.3 and compared against the best results in the literature obtained via alternative methods. The first thing to note is that our method demonstrates state-of-the-art performance in CNOT gate count reduction, matching or outperforming the best-known results across all tested instances, while offering a simpler framework and significantly lower computational cost. In contrast to Quarl, which relies on extensive training of a RL agent on supercomputer-grade hardware and requires tuning at least ten hyperparameters, QCMotherLode operates effectively on a standard personal computer with modest hardware resources, although its performance does scale with the number of CPU cores available for parallelization.

Let us now consider the optimized circuit decompositions we obtained for  $C^mX$ , with  $m = 3, 4, 5$ , which require 14, 30, and 62 CNOT gates, respectively (see Table 4.3). According to Lemma 7.1 in

| Circuit                  | $n$ | Original          | Nam <i>et al.</i> [114] | VOQC <sup>a</sup> [321] | Quartz <sup>a</sup> [322] | Staudacher <i>et al.</i> <sup>b</sup> [316] | Staudacher + Nam <sup>b</sup> [316] | Holker <sup>b</sup> [319] | QUESO <sup>b</sup> [323] | Quarl w/o R.M. <sup>c</sup> [324] | Quarl w/ R.M. <sup>c</sup> [324] | Tket 2.0.0 <sup>d</sup> [346] | Qiskit 2.0.0 <sup>e</sup> [342] | QCMotherLode 0.1.0 |
|--------------------------|-----|-------------------|-------------------------|-------------------------|---------------------------|---------------------------------------------|-------------------------------------|---------------------------|--------------------------|-----------------------------------|----------------------------------|-------------------------------|---------------------------------|--------------------|
| $C^3X$                   | 4   | 14 <sup>g</sup>   | –                       | –                       | –                         | –                                           | –                                   | –                         | –                        | –                                 | –                                | 14                            | 14                              | 14                 |
| $C^4X$                   | 5   | 36 <sup>g</sup>   | –                       | –                       | –                         | –                                           | –                                   | –                         | –                        | –                                 | –                                | 36                            | 36                              | 30                 |
| $C^5X$                   | 6   | 92 <sup>g</sup>   | –                       | –                       | –                         | –                                           | –                                   | –                         | –                        | –                                 | –                                | 92                            | 92                              | 62                 |
| $C^{10}X$                | 11  | 3068 <sup>g</sup> | –                       | –                       | –                         | –                                           | –                                   | –                         | –                        | –                                 | –                                | 3068                          | 3068                            | ?                  |
| Tof-NC <sub>3</sub>      | 5   | 18 <sup>f</sup>   | 14                      | 16                      | 14                        | 15                                          | 14                                  | –                         | 14                       | 12                                | 12                               | 18                            | 18                              | 12                 |
| Tof-NC <sub>4</sub>      | 7   | 30 <sup>f</sup>   | 22                      | 26                      | 22                        | 24                                          | 22                                  | 24                        | 22                       | 18                                | 18                               | 30                            | 30                              | 18                 |
| Tof-NC <sub>5</sub>      | 9   | 42 <sup>f</sup>   | 30                      | 36                      | 30                        | 33                                          | 30                                  | 33                        | 30                       | 26                                | 25                               | 42                            | 42                              | 24                 |
| Tof-NC <sub>10</sub>     | 19  | 102 <sup>f</sup>  | 70                      | 86                      | 70                        | 78                                          | 70                                  | 78                        | 70                       | 70                                | 65                               | 102                           | 102                             | ?                  |
| Tof-Barenco <sub>3</sub> | 5   | 24 <sup>f</sup>   | 18                      | 20                      | 16                        | 21                                          | 18                                  | –                         | 16                       | 14                                | 13                               | 24                            | 24                              | 13                 |
| Tof-Barenco <sub>4</sub> | 7   | 48 <sup>f</sup>   | 34                      | 38                      | 30                        | 40                                          | 34                                  | 37                        | 30                       | 24                                | 26                               | 48                            | 48                              | 24                 |
| Tof-Barenco <sub>5</sub> | 9   | 72 <sup>f</sup>   | 50                      | 56                      | 44                        | 57                                          | 48                                  | 55                        | 44                       | 42                                | 38                               | 72                            | 72                              | 38                 |

<sup>a</sup> Best reported result as evaluated in Ref. [323].

<sup>b</sup> Best reported result.

<sup>c</sup> Values represent averages across multiple optimization runs using their best-performing gate set. Standard deviations were also provided but are omitted here for simplicity. However, the minimal values obtained were not explicitly disclosed. While it is conceivable that some optimizations might have achieved lower values than the reported averages, without detailed knowledge of the underlying distribution, it is not possible to infer this for certain. Specifically, if the distribution of results is asymmetric (skewed), values below the reported averages may not have been observed at all.

<sup>d</sup> Obtained by combining all of the following optimization passes: `SynthesiseTket`, `FullPeepholeOptimise`, `CliffordSimp`, `KAKDecomposition`, `CommuteThroughMultis`, `RemoveRedundancies`, `AutoRebase`, `CommuteThroughMultis`, `OptimisePhaseGadgets`, `DecomposeBoxes`.

<sup>e</sup> Obtained using the `transpile` method with `optimization_level=3`.

<sup>f</sup> Sourced from Ref. [114, 351, 309].

<sup>g</sup> An  $(n - 1)$ -controlled NOT gate without auxiliary qubits, generated using the `MCXGrayCode` method in Qiskit 2.0.0.

Table 4.3: **CNOT count comparison across different gate decompositions and compilers.** The minimal number of CNOT gates achieved for each circuit and optimization method, as reported in the most favorable instance found in the literature, is listed – except for one method, where the average over multiple optimization runs was reported instead. The circuits labeled Tof-NC <sub>$m$</sub>  and Tof-Barenco <sub>$m$</sub>  refer to two distinct  $m$ -controlled Toffoli gates with  $m - 2$  ancillary qubits, both commonly used for benchmarking quantum circuit optimizers. In contrast,  $C^mX$  denotes an exact  $m$ -controlled Toffoli gate implemented without ancillary qubits.  $n$  denotes the number of qubits in the circuit. Green (yellow) cells highlight the best (second-best) result. Columns are sorted chronologically from left to right. QCMotherLode is the only method that outputs multiple logically equivalent and EGS-classified optimized circuits in a single run.

the foundational work of Barenco *et al.* [111], a general  $C^mX$  gate can be decomposed without ancillary qubits using  $2^{m+1} - 3$  two-qubit gates of three distinct types: CNOT, controlled- $V$ , and controlled- $V^\dagger$ , where  $V^2 = X$ . However, an explicit construction or gate count restricted exclusively to CNOT gates was not provided. For small values of  $m$ , specifically  $m = 2$  (the standard Toffoli gate) and  $m = 3$ , optimized constructions employing solely CNOT and single-qubit gates are known, requiring exactly 6 and 14 CNOT gates, respectively. When combined with the optimized circuit decompositions we derived

for  $m = 4$  and  $m = 5$ , the resulting sequence of CNOT gate counts — 6, 14, 30, 62, ... — suggests a scaling pattern of  $2^{m+1} - 2$  for  $m \geq 2$ , which hints that a deeper structural relationship underlying these decompositions has been found by this optimization procedure.

To confirm this hypothesis, a detailed examination of the circuit structures is required, which could yield valuable insights and potentially reveal the underlying algorithm governing these decompositions. Confirmation of this pattern would make generalization straightforward, obviating the need of explicit optimization for larger values of  $m$ . For instance, under this hypothesized relation, the decomposition for a  $C^{10}X$  gate is predicted to require exactly 2046 CNOT gates, indicated by the  symbol in Table 4.3. Although the precise optimal decomposition restricted exclusively to CNOT gates for larger  $m$  remains unresolved in the literature, exponential scaling in gate complexity is expected in the ancilla-free setting.

Despite this exponential complexity rendering these decompositions impractical for large-scale implementations, the significance of this example highlights a broader theoretical possibility: the discovery of general algorithmic structures through systematic circuit optimization, structural evaluation, and extrapolation might become feasible with intensive circuit optimization as provided by QCMotherLode.

Approaches for the implementation of multi-controlled Toffoli gates that employ auxiliary qubits typically achieve linear scaling in the gate count relative to the number of qubits involved. The remaining two constructions presented in the table exemplify this more efficient methodology. In particular, in the case of  $\text{Tof-NC}_m$ , a progression following the sequence 6, 12, 18, 24, ... also seems to be developing, suggesting the possibility of generalizing an algorithmic decomposition containing  $6(m - 1)$  CNOT gates. As a result, the hypothesis in  is that a circuit for  $\text{Tof-NC}_{10}$  might exist comprising 54 CNOTs. In the case of  $\text{Tof-Barenco}_m$ , although the optimal circuits obtained for  $m = 3, \dots, 5$  do not suggest any clear progression in CNOT count, further optimization of these and subsequent circuit instances might still reveal a pattern.

## 4.6 Conclusions and future work

In this chapter, we introduced QCMotherLode, a novel methodology for optimizing general quantum circuits on arbitrary qubit connectivity topologies aiming to minimize their two-qubit gate count. The proposed framework belongs to the class of exact optimization methods—i.e., it preserves the unitary representation of the circuit—and extends both  $ZX$ -calculus and circuit-level optimization by integrating them into a higher-level search and feedback procedure.

At each iteration, an input circuit is first transformed into its  $ZX$ -graph representation, simplified using

$ZX$  rewrite rules, and then independently re-extracted into  $N$  circuits using several recently proposed heuristic procedures collected in a user-defined portfolio. Following extraction, each circuit undergoes an additional circuit-level simplification pass employing state-of-the-art compilation techniques. The resulting circuits are then labeled based on a novel entangling-gate structure classification and each of the  $M \leq N$  unique circuits is added to a priority queue, which is sorted lexicographically by CNOT count and circuit depth for a best-first search procedure. The search iteratively feeds extracted circuits back into the  $ZX$  optimization-extraction pipeline. As a result, an extensive pool of logically equivalent circuit decompositions emerges, allowing this high-level search method to effectively explore the large discrete space of circuit structures while systematically pursuing configurations that minimize the objective function.

This feedback-based methodology consistently outperforms the heuristics on which it is built upon across all tested instances by combining their individual strengths while mitigating their biases. Additionally, it is inherently modular and versatile, enabling new heuristics to be seamlessly integrated into its portfolio as they emerge in future research. With the set of heuristics summarized in Table 4.2 and the normalization procedure described in Algorithm 1, this method has already been used to optimize several quantum circuits throughout this dissertation, including the critical component circuit for the preparation of VBS states in quantum hardware presented in § 4.4; benchmark circuits for multi-controlled Toffoli gates reported in § 4.5 and Table 4.3; the quantum circuit that implements the square-root POVM measurement in § 2.7.2; and multiple Fredkin and Toffoli gate circuits in various connectivity topologies, which will be presented in § 5.

Several features distinguish this framework from previous circuit optimization approaches. First, unlike the heuristics in the modular portfolio presented in Table 4.2, which explore the optimization landscape within the space of  $ZX$ -diagrams, this higher-level dual representation procedure, while leveraging  $ZX$ -based heuristics, directly navigates the space of extracted circuits. By iteratively extracting circuits, normalizing them, converting them back into  $ZX$ -diagrams, and re-optimizing, the method refocuses the optimization on the final circuit space, enabling best-first search to effectively uncover the optimal circuits.

This search procedure operates over a solution transition graph, where each node represents a complete quantum circuit, and edges correspond to transformations induced by heuristic procedures. Unlike traditional state-space search algorithms, where individual states incrementally contribute to a final solution, our framework navigates an evolving landscape of fully realized circuits. Through its feedback mechanism, it selects which solutions to expand further, leveraging properties of previously discovered circuit structures to facilitate convergence toward optimal configurations. To prevent redundant exploration and ensure efficiency, each newly generated decomposition is stored only once in the search space, effectively

restoring a search tree from the underlying dependency graph. This metaheuristic approach balances exploration and exploitation by selectively prioritizing candidates that demonstrate potential for reducing the two-qubit gate count while retaining structurally diverse circuits in the priority queue and leveraging the global reachability features enabled by the combination of multiple heuristics.

Unlike learning-based approaches, our method immediately achieves strong performance without requiring extensive computational pretraining. It also offers clear interpretability, particularly when compared with black-box reinforcement learning agents, whose internal decision-making processes can sometimes be more difficult to analyze. In contrast, the entire optimization trajectory and its intermediate circuits—from the initial input to the final circuit—is accessible and can be examined in terms of rule-based transformations. Since the procedure is largely deterministic, with randomness confined to specific intermediate stages, its behavior remains analyzable and reproducible.

Another key feature of this approach is its ability to generate multiple logically equivalent decompositions tailored to arbitrary coupling architectures, each exhibiting distinct entangling-gate structures. This reveals a significant abundance of logically equivalent circuits near optimal solutions (see Fig. 4.3d), where equivalence extends beyond local gate transformations to more complex multi-qubit transformations. For certain practical applications, circuits close to optimal may be sufficient to achieve good experimental results, particularly given the computational efficiency and speed of identifying these circuits using this simpler, non-learning-based optimization pipeline.

It is also worth noting that, by outputting many logically equivalent circuits with different CNOT counts and entangling-gate structures in a single run, QCMotherLode enables subsequent analysis of the stability of these circuits to random fluctuations in their input parameters. For instance, it might turn out that highly optimized circuits are more sensitive to small fluctuations in their single-qubit gate parameters, which can arise from hardware miscalibrations or qubit drift, making slightly less optimal circuits preferable in practice. While detailed exploration of these stability considerations is reserved for future work, our method provides a robust mechanism to generate a rich variety of near-optimal circuits, facilitating comprehensive investigation of such questions.

More importantly, this capability to generate multiple circuits provides insight into which ones correspond to fundamentally different classes of paths through the Hilbert space connecting input and output states. This feature holds significant potential for advancing automated quantum algorithm development by revealing the two-qubit gate structure of the circuits and facilitating the identification of recurring decomposition patterns across different qubit numbers within the same gate family. In addition, this variety of circuits can be leveraged to construct more resilient channels for coherent error mitigation by actively

selecting decompositions better suited to the noise profiles of physical quantum processors.

In fact, rather than focusing on runtime compilation, our method emphasizes the permanent discovery of circuit primitives. This distinction is significant because optimal decompositions of common gates, once identified, can be reused across multiple algorithms, thereby reducing execution time and error rates. Moreover, discovering novel algorithms for decomposing general gates over varying  $n$  qubits by generalizing structural patterns identified from small- $n$  decompositions can provide valuable new theoretical insights into quantum computing and simulation.

To systematically explore and leverage these structural patterns, we introduce tools such as the normalized circuit form, which enables the precise capture and comparison of connectivity structures across different circuits. This allows for the identification of commonalities across multiple qubit numbers and supports the discovery of general decomposition strategies. As an example, in the next chapter, we present a circuit rerouting method derived by optimizing the decomposition of the long-range CNOT gate for a small number of qubits and identifying, through extrapolation, a general structural pattern for arbitrary  $n$ . Even when more general gates involve continuous parameters, it remains plausible to derive useful structures and effective algorithms by leveraging the concepts presented in this chapter.

By systematically uncovering structural patterns in quantum circuits, this method not only facilitates optimization but also opens new avenues for circuit design. These insights naturally suggest extensions that could further refine the optimization process, motivating several directions for future research.

One promising direction is enhancing BFS with more adaptive backtracking and diversification strategies beyond  $k$ -beam BFS, which might further improve search efficiency. Potential strategies include a diversity-aware selection policy that prioritizes structurally distinct circuits over strictly lowest-cost candidates — requiring some novel quantitative measure of circuit dissimilarity, possibly based on geometric phases — thereby promoting broader exploration, or stochastic backtracking, which would reintroduce previously encountered circuits when no improvement is observed, helping to prevent stagnation. Additionally, leveraging the structure of the search tree by prioritizing prolific branches — those that consistently yield strong optimizations — could further focus exploration on the most promising regions.

We have begun exploring more advanced metaheuristic search strategies, particularly Monte Carlo Tree Search (MCTS), for navigating the discrete space of extracted circuits. Preliminary tests suggest that the current MCTS-based feedback procedure does not yet outperform the BFS-based approach. One possible explanation is the apparent abundance of logically equivalent circuits near optimal solutions (see Fig. 4.3d), which may allow certain heuristics to transition between equally promising branches of the search tree with sufficiently high probability, reducing the necessity of a structured hierarchical search.

However, these results remain inconclusive and are not yet suitable for formal presentation or for drawing reliable conclusions. Further analysis is required to determine the optimal tuning of MCTS parameters to balance exploration and exploitation, as well as to develop more sophisticated mechanisms for parallelizing the search and adapting the execution frequency of different heuristics in the portfolio to mitigate runtime imbalances. These refinements, which would strengthen the implicit hyperheuristic layer of the method by making heuristic application more dynamic and adaptive, are left as future work.

Arguably, to properly assess how our method compares with other circuit optimization approaches developed in recent years, further benchmarking would be beneficial in addition to the results gathered in Table 4.3. Many prior studies have benchmarked their methods using a standard suite of quantum arithmetic and reversible computing circuits [309, 114, 313, 316, 319]—including multi-controlled Toffoli gates, adders, multiplexers, Quantum Fourier Transform circuits, and circuits for modular and finite field arithmetic—or random circuit instances [325]. These evaluations typically focus on circuit complexity metrics such as gate count, two-qubit gate count, and circuit depth, though some also assess runtime performance. While the latter metric is primarily relevant for runtime compilation on specific platforms, it also provides insight into the practicality of these methods for permanent gate discovery, especially when computational costs or execution times vary significantly.

This dissertation has prioritized applying the developed method to discover circuits with immediate practical application rather than undertaking an extensive benchmarking study. Through our example applications, we have demonstrated its strong performance in optimization tasks beyond the Clifford+ $T$  gate set. Nevertheless, future work could incorporate large-scale numerical simulations to benchmark performance across these established metrics to better situate its effectiveness within existing optimization frameworks. In addition to reporting the best gate counts in common benchmarking circuits, this could include optimizations of random circuits for different numbers of qubits and varying computational time budgets, to empirically derive scaling laws that characterize performance in approximating the lower bound in Eq. (4.1). Additionally, testing different variations of the search strategies discussed earlier could provide deeper insight into which approach performs best across different circuit classes and optimization regimes.

In conclusion, the evolution and increasing sophistication of quantum circuit optimization as a research domain have introduced powerful methods and applications that extend far beyond optimizing circuits for hardware execution. In addition to improving execution times and minimizing errors on real hardware, our framework also enables novel error mitigation protocols and can complement theoretical algorithm development by facilitating machine-assisted extrapolation and discovery of new circuit primitives.

Since the invention of general-purpose computers, computer science has continuously demonstrated

a prolific growth driven significantly by the availability of hardware capable of rapidly testing new theoretical ideas, thereby accelerating conceptual iteration and exploration. Similarly, an optimization- and AI-based approach to quantum circuit implementation may well stimulate a new generation of quantum algorithms, unlocking solutions to computational challenges that remain beyond our current reach.

## 4.7 Appendix

### 4.7.1 Quantum circuit decomposition for the VBS simulation state

To complement the discussion in § 4.4, we provide the explicit OpenQASM code for the optimized decomposition of the controlled- $\mathcal{V}$  unitary. This listing details a gate sequence that achieves a substantial reduction in circuit depth and CNOT count, facilitating efficient implementation on near-term quantum devices.

```

1 OPENQASM 3.0;
2 include "stdgates.inc";
3 qubit[4] q;
4 cx q[3],q[0];
5 rx(-0.7639951343964286) q[0];
6 rz(1.1158382641639877) q[0];
7 rx(-2.0008195041933377) q[0];
8 cx q[3],q[1];
9 rx(3.064914452243082) q[1];
10 rz(1.6249100052412424) q[1];
11 rx(0.5468353521006657) q[1];
12 cx q[0],q[1];
13 rx(-2.6275132905411143) q[0];
14 rz(3.8335923997282695) q[0];
15 rz(-1.5752308984270087) q[1];
16 rx(1.6372168351540282) q[1];
17 rz(3.074877162154234) q[1];
18 cx q[0],q[1];
19 rx(0.30179380092061003) q[0];
20 rz(-2.8252844891168785) q[0];
21 rz(1.5707963267948974) q[1];
22 rx(1.570796326794896) q[1];
23 rz(-0.2859769674201251) q[1];
24 cx q[3],q[2];
25 rz(1.5707963267948954) q[2];
26 cx q[2],q[0];
27 rx(1.5707963267948966) q[0];
28 rz(-1.0471975511965979) q[0];
29 cx q[1],q[0];
30 rz(2.0943951023931953) q[0];
31 rx(1.7723200724795767) q[1];
32 cx q[2],q[0];
33 rx(1.5707963267948948) q[0];
34 rz(1.7780609457008603) q[0];
35 rx(-0.34325397615222286) q[0];
36 cx q[0],q[1];
37 rx(-2.526112944919406) q[0];
38 rz(-2.501257315317108) q[1];
39 rx(1.5840083469745991) q[1];
40 rz(3.1017511553135417) q[1];
41 cx q[0],q[1];
42 rx(1.5707963267948977) q[0];
43 rz(1.570796326794878) q[1];
44 rx(1.5707963267948977) q[1];
45 rz(-1.8346320608114404) q[1];
46 rx(-2.5261129449197046) q[2];
47 rz(1.5707963267948966) q[2];
48 rx(1.5707963267948966) q[2];
49 cx q[0],q[2];
50 rz(0.615479708670396) q[2];
51 cx q[3],q[1];
52 rx(1.7241550901022693) q[1];
53 rz(1.5707963267948968) q[1];
54 rx(1.5707963267948966) q[1];
55 cx q[3],q[2];
56 rz(5.66770559850919) q[2];
57 cx q[0],q[2];
58 rx(-2.7854644635105736) q[0];
59 rz(1.1751182385118137) q[0];
60 rx(-0.14240663324553715) q[0];
61 rz(4.096909271714302) q[2];
62 cx q[1],q[2];
63 rx(3.0578418259035454) q[1];
64 rz(1.5707963267948968) q[1];
65 rx(-0.8554899220033789) q[1];
66 cx q[0],q[1];
67 rx(-2.828279040316661) q[0];
68 rz(4.675298299831765) q[0];
69 rz(-1.6992233688580625) q[1];

```

```

70 rx(1.6289719608041442) q[1];
71 rz(3.0754267475024264) q[1];
72 cx q[0],q[1];
73 rx(0.11509147631954965) q[0];
74 rz(-1.6815725948794231) q[0];
75 cx q[0],q[2];
76 rx(-2.3561944901923453) q[0];
77 rz(1.466494821223268) q[0];
78 rx(0.06656816377582306) q[1];
79 rz(2.31138373166121) q[1];
80 cx q[0],q[1];
81 rx(1.4670575760178082) q[0];
82 rz(-2.3507746781660375) q[0];
83 rx(-0.3498006371684532) q[1];
84 rz(1.5707963267948966) q[2];
85 rx(1.5707963267948966) q[2];
86 cx q[2],q[0];

```

```

87 rz(-0.4525050988063064) q[0];
88 rx(-1.636781585996145) q[0];
89 cx q[0],q[1];
90 rx(-3.0752142681300185) q[0];
91 rz(-1.5752308984270087) q[1];
92 rx(1.6372168351540282) q[1];
93 rz(3.074877162154234) q[1];
94 cx q[0],q[1];
95 rx(0.06656816377584053) q[1];
96 rz(0.06407282357673472) q[1];

```

---

Listing 3: OpenQASM circuit instruction decomposing the controlled- $\mathcal{V}$  unitary from Ref. [49] assuming full qubit connectivity and employing 22 CNOT gates.

# Chapter 5

## Optimized Fredkin and Toffoli gate decompositions for coherent error mitigation via equivalent circuit averaging <sup>1</sup>

“Computation — whether by man or by machine — is a physical activity, and is ultimately governed by physical principles.”

---

Edward Fredkin & Tommaso Toffoli, 1982 [[352](#)]

### 5.1 Introduction

*Conservative logic* is a classical computational model in which operations mirror fundamental physics principles, such as the reversibility of dynamical laws and the conservation of additive quantities, such as energy or particle number. Unlike traditional computational models, which often disregard these physical constraints, conservative logic emphasizes both *reversibility* and *conservativity*, providing a framework that is better aligned with natural principles and well-suited for its physical implementation. Within this model, reversibility requires logic circuits to realize invertible Boolean functions, ensuring that every input maps uniquely to a corresponding output without information loss. In addition, *conservativity* enforces the output bitstrings to be a permutation of the input bits, thereby preserving their Hamming weight.

Designed within this framework, the Fredkin gate — also referred to as controlled-SWAP — and the Toffoli gate — also referred to as controlled-controlled-NOT — are two notable examples of three-input,

---

<sup>1</sup>This chapter is based on joint work published in collaboration with Bruno Murta in Refs. [[129](#), [128](#)].

three-output logic gates introduced by Fredkin and Toffoli in their seminal work [352]. The Fredkin gate operates by swapping two target bits if and only if the control bit, left unchanged, is found in state 1<sup>2</sup>. In Boolean notation, where juxtaposition  $pq$  represents the logical conjunction ( $p \wedge q$ ),  $\bar{p}$  denotes the logical negation of  $p$ , and  $\oplus$  denotes the logical exclusive OR (XOR), we may write this operation as

$$(c, t_1, t_2) \mapsto (c, ct_2 \oplus \bar{c}t_1, ct_1 \oplus \bar{c}t_2). \quad (5.1)$$

In contrast, the Toffoli gate flips its single target bit if the input of both control bits is set to 1:

$$(c_1, c_2, t) \mapsto (c_1, c_2, t \oplus c_1c_2). \quad (5.2)$$

Although the Fredkin gate is conservative and the Toffoli gate is not, both of these logic gates serve as universal primitives for reversible classical logic. Specifically, any invertible Boolean function can be constructed entirely out of Fredkin or Toffoli gates, provided the freedom to introduce ancillary bits that may be initialized in either the 0 or 1 state as needed.

Because their reversible nature is compatible with the unitarity required for quantum operations in closed systems, these gates were seamlessly adapted to the quantum domain. Beyond their capability to operate classical logic circuits using a quantum substrate, they also enabled the extension of Boolean-like operations to quantum data by acting on qubits in superposition states. Their significance was further underscored by the realization that the Toffoli gate, when combined with the Hadamard gate, forms a universal gate set for quantum computation [353, 354]. Additionally, since the Toffoli gate can be synthesized using a combination of Fredkin gates and auxiliary qubits—at least one of which must be initialized in the state  $|1\rangle$ —a gate set comprising Fredkin, Hadamard, and  $X$  gates is likewise universal for quantum computation, provided fiducial states are prepared in  $|0\rangle$  and sufficient ancillary qubits are available.

Both gates soon became fundamental components in quantum computation due to their versatility and wide applicability. The Toffoli gate serves as the cornerstone for its generalizations to multi-qubit systems [355, 203, 356], which are widely employed in quantum arithmetic operations and the design of oracles for quantum algorithms [357, 3, 5, 358]. Toffoli gates also found application in quantum error correction schemes [359], linear combination of unitaries [17], and quantum singular value transformations [360]. A recent variation, termed the iToffoli gate, modifies the behavior of the Toffoli by acting with  $iX$  instead of  $X$  in the target qubit when the control qubits are in state  $|0\rangle$ , and has been applied in proposals for

---

<sup>2</sup>Fredkin and Toffoli's original publication defines the swap action triggered by the control bit in state 0. However, following quantum information conventions, we adopt the trigger state to be  $|1\rangle$  in the quantum versions of these gates.

calculating molecular response properties in the frequency domain [361, 362].

The Fredkin gate is equally significant, underpinning numerous quantum algorithms and protocols. It forms the basis of the SWAP test, a standard tool for evaluating the fidelity between two quantum states with performing tomography on each of them separately [363, 364]. Beyond this, the Fredkin gate has been applied in diverse areas, such as quantum state preparation [365, 366, 49], estimating density operator functionals [367], and constructing quantum switches [368, 369, 370]. Additional applications include optimal quantum cloning [371], stabilization of quantum computations [363], and the preparation of states in the Hamiltonian eigenbasis, often in conjunction with the Toffoli gate [372]. Furthermore, it has been employed in calculating Bargmann invariants, a task relevant to quantum geometric and topological analyses [373, 374].

Finally, both the Fredkin and Toffoli gates have demonstrated utility in quantum routines optimized for noisy intermediate-scale quantum (NISQ) devices. These gates are integral to various algorithms designed specifically for near-term quantum hardware, where limited coherence times and gate fidelities require careful optimization of resources [49, 374, 375, 376, 377, 48].

Given their extensive list of applications, it is easy to understand the significant attention that has been devoted to the challenge of implementing Fredkin and Toffoli gates on digital quantum computers. Numerous strategies have been proposed for their native implementation on specific quantum platforms such as trapped ions [378, 379], superconducting circuits [380, 359, 381, 361, 382], and optical systems [383, 384, 385, 386, 387]. In most cases, however, Fredkin and Toffoli gates are not directly implemented; instead, they are synthesized from easier-to-realize single- and two-qubit gates that form the basis elements of universal gate sets, with which all other gates can be expressed.

In this chapter, a similar high-level and hardware-agnostic methodology is adopted, where the Fredkin and Toffoli gates are represented in terms of common single- and two-qubit operations, using the CNOT gate as the reference two-qubit operation. In contrast to some previous works [113, 388, 389], which have prioritized minimizing the use of non-Clifford operations, such as  $T$  gates, to facilitate fault-tolerant implementations [390], in what follows we pursue decompositions minimizing the number of CNOT gates while also accounting for qubit connectivity constraints — an approach particularly important for NISQ hardware [391]. Besides, while it is possible to further reduce the CNOT count with decompositions up to a relative phase factor [355, 389, 392] or by making use of qutris in place of some qubits [393, 394, 395, 396], we focus exclusively on qubits, aiming to implement Fredkin and Toffoli gates with exact CSWAP and CCNOT operators with the following representation, up to a global phase:

$$\text{CSWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} = \begin{array}{c} \bullet \\ \text{---} \\ \times \\ \text{---} \\ \times \\ \text{---} \end{array} \quad (5.3)$$

$$\text{CCNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} = \begin{array}{c} \bullet \\ \text{---} \\ \bullet \\ \text{---} \\ \oplus \\ \text{---} \end{array} \quad (5.4)$$

Here, we adopt the standard lexicographic ordering for the computational basis  $|q_1\rangle \otimes |q_2\rangle \otimes |q_3\rangle$ , where the leftmost qubit corresponds to the top wire in circuit diagrams and the rightmost qubit corresponds to the bottom wire.

Moreover, beyond simply optimizing CSWAP and CCNOT decompositions, a key objective of this chapter is to address coherent errors in their implementation on real devices. The proposed methodology for mitigating these errors places particular emphasis on two-qubit gates, given their disproportionately high impact on the accuracy of algorithms executed on real quantum platforms, as identified in § 3. In particular, the circuit optimization strategy introduced in § 4 is employed to generate multiple distinct decompositions of the Fredkin and Toffoli gates, each featuring a unique two-qubit gate structure. By combining these logically equivalent but structurally different circuits, systematic control errors can be effectively mitigated, as will be explained in detail.

The structure of this chapter is organized as follows. In § 5.2, we discuss the minimization of the CNOT count in the decomposition of Fredkin and Toffoli gates for three qubits that are adjacently connected, considering both all-to-all and linear qubit connectivity. § 5.3 addresses the scenario where the three qubits involved in the operations are not directly connected. Specifically, a method is developed to temporarily reposition the qubits to adjacent locations and then return them to their original positions, achieving a reduction of one CNOT gate for each SWAP operation. § 5.4 leverages the multiple circuit representations of the Fredkin and Toffoli gates to mitigate coherent errors through an equivalent circuit

averaging procedure; we analyze its performance with both computational simulations and experimental validations on a transmon quantum processor. Finally, § 5.5 provides a summary of all the findings.

## 5.2 Decompositions for adjacent qubits

The possibility of decomposing the Toffoli gate using only five two-qubit gates has been rigorously demonstrated in previous studies [397, 398]. Nevertheless, the gate sets natively implemented in most quantum computing platforms usually include only a single type of two-qubit gate, such as the CNOT. For this reason, the practical realization of the Toffoli gate typically requires a quantum circuit decomposition with a minimum of six two-qubit gates. The standard decomposition, shown within the solid blue box in Fig. 5.1(c), is a textbook example [203] of a Toffoli gate implementation optimized for CNOT count. This decomposition serves as our basis for deriving similarly shallow implementations of both the Toffoli and Fredkin gates constrained to all possible qubit connectivity graphs.

For the Fredkin gate, its standard quantum circuit [397] is derived by adapting the established decomposition of the SWAP gate into three CNOT gates that is known at least since the work in [399] (Fig. 5.1(a)). To make this decomposition controlled by a third qubit, a simple approach would attach the control qubit to each CNOT gate. However, owing to the symmetric structure of the SWAP decomposition, only the middle CNOT gate needs to be controlled, as illustrated in Fig. 5.1(b) and detailed in § 5.6.1. Incorporating the standard Toffoli decomposition mentioned before, a Fredkin gate circuit with 8 CNOT gates and a depth of 14 layers is obtained, as depicted in Fig. 5.1(c). By simplifying the subcircuit highlighted by the red dashed box and adjusting the single-qubit gates at the end of the circuit, one CNOT gate and one layer of single-qubit gates can be eliminated, reducing the circuit to 7 CNOT gates and depth 13 (Fig. 5.1(d)). To date, this circuit achieves the Fredkin gate decomposition structure with minimal CNOT count and depth.

The quantum circuits displayed in Fig. 5.1 are designed under the assumption of complete qubit connectivity, enabling direct implementation of CNOT gates between any qubit pair. However, many real-world quantum processors, such as those employing superconducting circuits [56] or silicon quantum dots [400], often impose significant connectivity limitations. In these systems, qubits that are physically distant require their information to be relayed through intermediary SWAP gates [401], which inevitably increases the overall circuit depth. As a result, designing shallower decompositions that circumvent these SWAP networks while accommodating restricted connectivity is vital for optimizing the performance of quantum devices with these constraints. This is already a challenge for three-qubit gates since the majority




Figure 5.1: **Circuit decompositions of Fredkin and Toffoli gates with full qubit connectivity.** **a**, Standard SWAP gate decomposition using 3 CNOT gates [399]. **b**, Fredkin gate decomposition derived by integrating a Toffoli gate between two CNOT gates, adapted from the circuit in (a) and controlled following the results in § 5.6.1. **c**, Incorporation of the textbook Toffoli gate decomposition [203], highlighted by the solid blue box, into the circuit in (b), yielding a Fredkin gate implementation with 8 CNOT gates and a circuit depth of 14. **d**, The circuit in (c) is simplified twice. First, by replacing the subcircuit enclosed by the red dashed box, reducing the CNOT count by one. Second, a layer of single-qubit gates is removed at the end by appropriately commuting them with a CNOT. As a result, the Fredkin gate can be implemented using 7 CNOT gates and a depth of 13. The single-qubit Hadamard ( $H$ ), phase ( $S$ ), and  $\pi/8$  ( $T$ ) gates follow their conventional definitions [203], with  $\sqrt{X} = HSH$ . Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

of current quantum processors lacks fully connected qubit triplets.

To generate alternative but logically equivalent circuit implementations for the Toffoli and Fredkin gates under varying connectivity constraints, the method in § 4 was then employed. In the case of all-to-all connectivity, we began with the circuits with optimal CNOT-count shown in Fig. 5.1, creating equivalent decompositions that preserved the number of CNOT gates while modifying their arrangements. For linear connectivity, the starting circuits for the search procedure were derived from the same decompositions in Fig. 5.1 but modified by incorporating additional SWAP gates. These gates, decomposed into three CNOTs, were employed to resolve the lack of direct connectivity between the qubits involved in the implementation of 2 CNOTs. This straightforward approach, necessary to ensure the input to the method is a constrained circuit, adds an excessive number of CNOT gates. Nevertheless, the optimization and discovery method presented in § 4 significantly reduced it, obtaining several equivalent circuits with different entangling-gate structure.

Since our goal is not only to minimize CNOT count but also to maximize the number of logically equivalent circuit decompositions with distinct entangling-gate structure, additional circuit variations were generated from the search results by leveraging inherent symmetries of the Toffoli and Fredkin gates. Key symmetries include the invariance of the Toffoli gate under qubit permutation when represented




Figure 5.2: **Circuit decompositions of Toffoli and Fredkin gates under linear qubit connectivity.** Just like for full qubit connectivity, these circuits were optimized for CNOT-count using the methods in § 4. **a**, Toffoli gate. **b**, Fredkin gate with the control qubit positioned at one end of the three-qubit register. **c**, Fredkin gate with the control qubit located in the middle of the three-qubit register. The single-qubit gates Hadamard ( $H$ ), phase ( $S$ ),  $\pi/8$  ( $T$ ), Pauli-X ( $X$ ), and Pauli-Z ( $Z$ ) adhere to usual definitions [203], with  $\sqrt{X} = HSH$  and  $\sqrt{X}^\dagger$  representing its inverse. The circuit in (a) assumes the target qubit of the Toffoli gate occupies the central position; however, this target qubit can be reassigned by repositioning the two Hadamard gates in the desired target qubit (see § 5.6.2). Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

as a controlled-controlled-Z (CCZ) gate through Hadamard transformations, as detailed in § 5.6.2, the symmetry of the Fredkin gate under target-qubit exchange, and the self-inverse property of both gates. In scenarios constrained by linear connectivity, transformations that introduced CNOT gates between indirectly connected qubits were excluded. This systematic approach expanded the repertoire of valid, logically-equivalent circuit implementations, the usefulness of which will be discussed in § 5.4.

Table 5.1 provides a comprehensive summary of CNOT counts and the variety of logically-equivalent circuits generated for each connectivity configuration and odd qubit placement—target qubit for the Toffoli, control qubit for the Fredkin. In fully connected architectures, the placement of the odd qubit is inconsequential. In the case of the Toffoli gate, even in systems with linear connectivity, the location of the target qubit in the circuit does not affect the efficiency of its decomposition. This is, again, because the target qubit can be reassigned by repositioning the Hadamard gates at the circuit ends, a property that arises from the CCZ gate invariance under qubit reordering (see § 5.6.2). Figure 5.2 illustrates representative circuits achieving minimal CNOT counts in linear connectivity configurations. The full set of circuit instructions was published in Ref. [128] in OpenQASM file format.

The optimization technique here presented produced circuits for the Toffoli and Fredkin gates in five different scenarios of qubit connectivity and odd qubit placement that achieve the lowest CNOT counts

| Gate    | Connectivity | Odd qubit placement | CNOT count | No. equiv. circ. | No. equiv. circ. with +1 CNOT |
|---------|--------------|---------------------|------------|------------------|-------------------------------|
| Toffoli | All-to-all   | Anywhere            | 6          | 48               | —                             |
|         | Linear       |                     | 8          | 18               | 54                            |
| Fredkin | All-to-all   | Anywhere            | 7          | 40               | —                             |
|         | Linear       | Ends                | 8          | 8                | 22                            |
|         |              | Center              | 10         | 2                | 69                            |

Table 5.1: **Summary of CNOT counts and number of logically equivalent circuits for Fredkin and Toffoli gates across five different connectivity scenarios.** Varying qubit connectivity and placement of the odd qubit — control qubit for the Fredkin gate and target qubit for the Toffoli gate — are considered. Alongside the circuits with optimized CNOT number, additional circuits were also obtained under linear connectivity that contain one extra CNOT gate compared to the optimal decompositions, which could be advantageous for applications like equivalent circuit averaging or other error mitigation strategies (see § 5.4). These circuits were made publicly accessible in the OpenQASM file format [128].

|                        | Toffoli<br>all-to-all | Toffoli<br>linear | Fredkin<br>all-to-all | Fredkin<br>linear (ends) | Fredkin<br>linear (center) |
|------------------------|-----------------------|-------------------|-----------------------|--------------------------|----------------------------|
| Here                   | 6                     | 8                 | 7                     | 8                        | 10                         |
| BQSKit [402]           | 6                     | 8                 | 7                     | 8                        | 10                         |
| CPFlow [403]           | 6                     | 8                 | 7                     | 8                        | 11                         |
| Qiskit [404]           | 6                     | 10                | 7                     | 11                       | 17                         |
| Duckering et al. [377] | 6                     | 8                 | -                     | -                        | -                          |
| Liu et al. [405]       | 6                     | 8                 | -                     | -                        | -                          |

Table 5.2: **CNOT counts obtained by different methods for Fredkin and Toffoli gate circuits.** Just like in Table 5.1, five scenarios involving different qubit connectivity constraints and placements of the odd qubit are considered. The methods used for benchmarking include BQSKit [402], CPFlow [403], and Qiskit [404]. Additionally, the lowest CNOT counts previously reported in the literature [377, 405] for the Toffoli gate are included for comparison; however, no equivalent results for the Fredkin gate were found. Beyond achieving the lowest CNOT counts in all scenarios considered, the circuits produced by our approach offer further advantages: they feature exact single-qubit-gate parameters expressed as fractions of  $\pi$ , ensuring precision, and have been precomputed and stored, allowing instant retrieval and eliminating the need for recomputing their unitary decomposition when applied to algorithms.

reported to date. Table 5.2 compares these results with alternative methods. Results in the literature primarily address Toffoli decompositions under all-to-all and linear connectivity [377, 405], with analogous data for Fredkin gates being scarce. Using the BQSKit [402] and CPFlow [403] software packages, the lowest CNOT counts reported in this work have also been achieved for most of these decompositions, the only exception being the Fredkin gate under linear qubit connectivity with the control qubit positioned at the center. However, our approach provides three key advantages. First, the circuits produced in this study are decomposed in the  $\{\text{CNOT}, R_z(\theta), R_x(\theta)\}$  basis<sup>3</sup> and all  $\theta$  parameters expressed as exact

<sup>3</sup>Where  $R_x(\theta) = \exp(-i\theta X/2)$ , and  $R_z(\theta) = \exp(-i\theta Z/2)$

fractions of  $\pi$ , ensuring both numerical accuracy and simplified decompositions in fault-tolerant systems. Second, multiple logically equivalent decompositions are provided, enhancing future flexibility in the global circuit optimization procedures where these gates can be integrated, besides opening up new possibilities such as described in § 5.4. Third, the full set of logically equivalent circuits that were produced is freely accessible online, enabling efficient storage and retrieval as needed [128].

To elaborate, the availability of diverse equivalent circuit decompositions enables numerous strategies to mitigate the limitations of near-term quantum hardware. For instance, distinct gate decompositions can be selected to better integrate with surrounding circuit primitives [405]. Additionally, if specific qubit pairs exhibit high CNOT error rates, a decomposition minimizing interactions between those qubits can be employed to improve overall fidelity. Most critically, equivalent circuit averaging [103, 102, 101] can significantly reduce coherent error effects. A detailed examination of this application is presented in § 5.4. Before that, the next section focuses on the implementation of Fredkin and Toffoli gates in configurations where the qubits are not directly connected.

## 5.3 Decompositions for non-adjacent qubits

This section explores the realization of Fredkin and Toffoli gates when the three involved qubits are not directly connected. To implement them in this scenario, intermediate qubits along the path between active qubits are leveraged to construct the required global unitary operations over longer ranges. Instead of employing a direct decomposition into basic gates, we propose the *cnot-swapping* technique, which efficiently repositions the qubits prior to and following the execution of the three-qubit circuits illustrated in Fig. 5.2. The general applicability of this method is first demonstrated for moving qubits where the gate matrix is diagonal in the computational basis, encompassing the cases of control qubits that need to be relocated in controlled operations. Subsequently, we detail its use for repositioning target qubits in multi-controlled-NOT operations. The long-range CNOT and Toffoli gates emerge as natural applications of these principles. Finally, the method is extended to the Fredkin gate.

### 5.3.1 CNOT-SWAP rerouting

We begin by introducing the main primitive to be used for rerouting qubits in Fredkin and Toffoli gates—the CNOT-SWAP gate. Let us examine its action on two classical bits:

$$\begin{array}{c} \text{---} \times \text{---} \\ | \\ \oplus \end{array} = \begin{array}{c} |i_1\rangle \text{---} \bullet \text{---} \oplus \text{---} |i_2\rangle \\ |i_2\rangle \text{---} \oplus \text{---} \bullet \text{---} |i_1 \oplus i_2\rangle, \end{array}$$

where  $i_1, i_2 \in \{0, 1\}$ . As can be seen, the CNOT-SWAP gate applies a NOT operation on the second bit controlled by the state of the first one, while simultaneously replacing the state of the first qubit with that of the second one. Its action, therefore, resembles that of a CNOT on the second qubit and that of a SWAP on the first one, as represented by the shorthand diagram. At the end of the operation, one of the qubit states is preserved in its original form (*clean*), while the other one carries computational results (*dirty*). Applied to a general two-qubit state  $|\psi\rangle = \sum_{i_1, i_2} a_{i_1 i_2} |i_1\rangle \otimes |i_2\rangle$ , the resulting state is

$$\begin{aligned} |\psi'\rangle &= \text{CNOT-SWAP}_{1,2} |\psi\rangle \\ &= \sum_{i_1, i_2} a_{i_1 i_2} |i_2\rangle \otimes |i_1 \oplus i_2\rangle, \end{aligned} \quad (5.5)$$

showing that the operation permutes the computational basis without altering their amplitudes. Represented in the computational basis, the CNOT-SWAP is written as

$$\text{CNOT-SWAP}_{1,2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \quad (5.6)$$

When we wish to implement a two-qubit gate  $V$  between qubits that are not physically adjacent, the standard approach involves three steps. First, the qubits must be relocated to adjacent positions so that the gate can be applied locally. This is typically done using a network of SWAP gates, which exchange the states of neighboring qubits, effectively moving the qubits across intermediate positions. Once the two qubits are adjacent, the gate  $V$  is applied. Finally, to preserve the logical structure of the circuit, the qubits are returned to their original positions by reversing the sequence of SWAP gates. This ensures that the overall computation behaves as though the qubits had remained in their original locations throughout.

However, this process can be optimized if the gate  $V$  has a special property: being diagonal in the computational basis of the qubit being moved (see § 5.6.3 for an exact definition). In such cases, the SWAP gates in the relocation sequence can be replaced with CNOT-SWAP gates. For each unit of a forward and reverse movement of a qubit, two fewer CNOT gates are required when using CNOT-SWAP gates instead of SWAP gates, as depicted in Fig. 5.3(b). This reduction becomes significant in circuits with many qubit movements.

The advantage of this approach is that the qubit being moved remains in its well-defined *clean* state throughout the process. Even though the intermediate qubits along the path of movement are left *dirty* after




Figure 5.3: **Control and target-NOT qubit relocation via cnot-swapping.** **a**, Shorthand representation of the CNOT-SWAP gate, which is functionally equivalent to the *fermionic* SWAP [406] and *i*SWAP [407] gates up to single-qubit transformations. Previously, the CNOT-SWAP was also referred to as the “double-CNOT” and recognized as a maximally non-local operator [408]. **b**, One-hop relocation of the control-qubit in a general controlled gate, achieved with two CNOT-SWAP gates. **c**, One-hop relocation of the target-qubit in a multi-controlled-NOT gate using two CNOT-SWAP gates oriented in the reverse direction compared to the scheme in (b). Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

the first passage, they are restored to their original states after the CNOT-SWAP sequence is reversed, ensuring that the computation is not disrupted. The key insight is that  $V$  preserves the computational basis states of the moving qubit, ensuring that any garbage generated on the dirty qubits remains directly tied to specific basis states of the moving qubit (see Eq. 5.5). This relationship allows for effectively cleaning the dirty qubits when the initial network is reversed, by disentangling them from the moving qubit.

This method can be applied more broadly to relocate a qubit in any  $n$ -qubit gate that does not change its amplitudes, apart from introducing a phase factor. In practical terms, the diagonality condition of the moving qubit can first be verified algebraically, as explained in § 5.6.3, to determine whether *cnot-swapping* can then be applied in place of standard *swapping*. By leveraging the special structure of such gates, this CNOT-SWAP technique provides an efficient alternative to the standard SWAP-based approach, reducing both the gate count and the overall depth of the circuit.

### 5.3.2 Long-range CNOT and Toffoli gates

An important application of cnot-swapping lies in rerouting control qubits in quantum gates, an essential operation for long-range CNOT and Toffoli gates, as depicted in Fig. 5.3(b). Consider an initial state  $|i_1\rangle \otimes |i_2\rangle$  of the two topmost qubits. Applying  $\text{CNOT-SWAP}_{2,1}$  modifies this state to  $|i_1 \oplus i_2\rangle \otimes |i_1\rangle$ . As a result, any subsequent controlled operation on the bottom  $n + 1$  qubits will now depend on  $|i_1\rangle$  as the control, also maintaining this state after its operations. Finally, the top two-qubit state is reverted to

its original form,  $|i_1\rangle \otimes |i_1 \oplus i_1 \oplus i_2\rangle = |i_1\rangle \otimes |i_2\rangle$ , by applying another CNOT-SWAP in the reverse direction. This process naturally generalizes to longer control qubit movements via iterative application.

Surprisingly, the CNOT-SWAP gate also enables the relocation of target qubits in multi-controlled-NOT (MCX) operations, as illustrated in Fig. 5.3(c). Unlike the control-qubit case, the moving target-qubit serves as the dirty qubit in the CNOT-SWAP operation, leaving the idle qubits it traverses clean. This reversal of roles explains the opposite orientation of the CNOT-SWAP gates with respect to the direction of movement of the relocated qubit in Figs. 5.3(b) and (c). Specifically, for the example in Fig. 5.3(c), an initial state  $|i_{n+1}\rangle \otimes |i_{n+2}\rangle$  of the bottom two qubits is transformed into  $|i_{n+1} \oplus i_{n+2}\rangle \otimes |i_{n+1}\rangle$  by the first CNOT-SWAP. After the MCX gate acts, the state evolves to  $|(\bigoplus_{j=1}^n i_j) \oplus (i_{n+1} \oplus i_{n+2})\rangle \otimes |i_{n+2}\rangle$ , and the final CNOT-SWAP <sub>$n+1, n+2$</sub>  operates on it to produce the intended result given by  $|i_{n+2}\rangle \otimes |(\bigoplus_{j=1}^n i_j) \oplus i_{n+1} \oplus (i_{n+2} \oplus i_{n+2})\rangle = |i_{n+2}\rangle \otimes |(\bigoplus_{j=1}^n i_j) \oplus i_{n+1}\rangle$ .

When constructing CNOT-SWAP networks to extend the movement of control and target qubits, it is possible to optimize the relocation circuit further by rearranging CNOT gates. Specifically, the final CNOT in one CNOT-SWAP can be interchanged with the initial CNOT of the next CNOT-SWAP in the sequence. As the common qubit in these subsequent CNOT gates is either a control or a target to both, they can be permuted, facilitating concurrent execution of some of the CNOTs along the network path and reducing circuit depth even further. This optimization is illustrated in Fig. 5.4(c).

The simplest case of rerouting the control qubit in a CNOT gate corresponds to the *long-range* CNOT gate, which connects two non-adjacent qubits. The CNOT-SWAP methodology applied to the long-range CNOT yields a decomposition with the minimum CNOT count and circuit depth in the literature, minimized in this order. Compared to traditional approaches, this method offers significant savings. For instance, the basic SWAP-based approach requires  $6n + 1$  CNOT gates and a best-case depth of approximately  $3n^4$ , where  $n \geq 1$  is the number of intermediary qubits. Shende *et al.* [183] introduced an alternative to this basic SWAP-based approach, efficiently reducing the CNOT gate count to exactly  $4n$  but increasing circuit depth to  $4n$  as well. Subsequently, Kutin *et al.* [409] developed a circuit design that decreased the circuit depth to  $\sim n$ , while incurring only a marginal increase of one additional CNOT gate compared to the construction by Shende *et al.*

The long-range CNOT decomposition derived using the CNOT-SWAP approach, illustrated in Fig. 5.4, achieves a circuit depth<sup>5</sup> of  $\sim n$ , as the one reported by Kutin *et al.*, while preserving the minimal CNOT count of  $4n$  obtained by Shende *et al.* We have confirmed the optimality of this decomposition for

<sup>4</sup>Achieved when the control and target qubits of the long-range CNOT are both moved simultaneously towards one another.

<sup>5</sup>Specifically, the decomposition in Fig. 5.4 has a depth of  $n + 7$  for  $n \geq 4$ , which is just one CNOT layer higher than the  $n + 6$  depth guaranteed by Kutin *et al.* [409] for arbitrary  $n$  (see Theorem 2.1).




Figure 5.4: **Long-range CNOT gate implementation.** **a**, Example of a  $3 \times 3$  square qubit lattice with only nearest-neighbor connectivity. The five qubits involved in the long-range CNOT operation are highlighted, where active qubits (in blue) are linked through idle qubits (in red) along a Manhattan-distance-minimizing path. **b**, Circuit decomposition of the long-range CNOT between qubits 1 and 9, achieved by symmetrically moving both towards each other to reduce circuit depth. **c**, Generalized construction of a long-range CNOT gate, optimized first for minimal CNOT count and subsequently for depth. The section within the blue box represents the optimal decomposition for a CNOT acting across a single idle qubit. For longer distances between control and target qubits, networks of CNOT-SWAP gates are introduced on both sides. The red box highlights the commutation of adjacent CNOT gates from subsequent CNOT-SWAP gates, demonstrating an arrangement that allows each rerouting layer to accommodate two to four concurrent CNOT executions. Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

cases where  $n = 1$  and  $n = 2$  idle qubits separate the two active qubits. These instances were minimized first with respect to the CNOT count and second with respect to CNOT depth, with exhaustive searches performed using a gate set restricted to nearest-neighbor CNOT gates.

With a view to implementing the Toffoli gate on three non-adjacent qubits, treating every CNOT in the decomposition as a long-range CNOT may not be the most efficient approach. Instead, the same cnot-swapping principles can be leveraged to reposition the qubits and realize the Toffoli gate locally using circuits that assume linear qubit connectivity introduced in § 5.2. Specifically, the movement of both the control and target qubits in the Toffoli gate can be performed analogously to the relocation of the control and target qubits in the long-range CNOT. For clarity, Fig. 5.5(a) presents an explicit example of a CNOT-SWAP-based decomposition for the Toffoli gate. A review of the literature revealed no prior documentation of this type of decomposition.

To conclude, in the case of both the long-range CNOT and Toffoli gates, the SWAP operations in the




Figure 5.5: **Long-range Toffoli and Fredkin gate implementations under linear connectivity.** **a**, Toffoli gate decomposition in which each SWAP gate has been replaced with a CNOT-SWAP, eliminating one CNOT for every replacement. **b**, For the Fredkin gate, rerouting its target qubits using CNOT-SWAP networks is feasible only when both targets are moved together to traverse the same set of idle qubits, as elaborated in the main text. The depicted approach, which moves the control qubit and both targets in parallel, is thus designed to minimize circuit depth instead of CNOT count. An alternative strategy, that only relocates the control qubit via CNOT-SWAP networks would minimize the overall CNOT count, albeit at the expense of increased circuit depth. The depth optimization achieved through this method can be better understood by analyzing the CNOT-SWAP decomposition at the level of individual CNOT gates and leveraging the gate permutation technique illustrated in Fig. 5.4(c). Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

rerouting layers — both preceding and following the local gate — can be substituted with CNOT-SWAP gates. This modification results in the CNOT gate count within the rerouting network dropping from  $6n$  to  $4n$ , while the overall depth of the circuit is even more significantly decreased from  $\sim 6n$  to  $\sim n$ . Here,  $n$  represents the total number of qubit hops required to bring together the three (two) qubits of the Toffoli (long-range CNOT).


### 5.3.3 Fredkin gate

For the Fredkin gate, it is clear that the control qubit can be relocated using a network of CNOT-SWAP gates analogously to the control qubits in long-range CNOT and Toffoli operations. However, the direct application of such networks to relocate the target qubits is not evident from the previous discussion, as the SWAP operation induced by the Fredkin gate is neither diagonal in the computational basis nor equivalent to a NOT operation. Nevertheless, a valid approach involves relocating the target qubits collectively via the CNOT-SWAP network, as demonstrated in Fig. 5.5(b).

Suppose the two target qubits of the long-range Fredkin gate are located in adjacent positions. By relocating them in tandem towards the control, they encounter the same set of idle qubits. This symmetry ensures that any unintended computation introduced into those idle qubits by the first CNOT-SWAP

network can be removed by the second network: even though the Fredkin gate swaps the target qubits partway through the procedure, because the basis states of both targets can still be propagated backwards, they can still uncompute the garbage at the end.

To illustrate this, consider the setup depicted in Fig. 5.5(b). Below, we break it down to analyze the effect of the network on the computational basis states before and after the Fredkin operation:



Let  $\{|i_n\rangle\}_{n=1}^4$  represent the computational basis states of the idle qubits, and the control qubit  $|c\rangle$  be initialized in state  $|1\rangle$ , since  $|x\rangle = |0\rangle$  is trivial. As can be seen, the Fredkin gate, applied locally, exchanges the states  $|t_1\rangle$  and  $|t_2\rangle$ . Once this swap occurs, both target qubits again pass through the same idle qubits they affected initially, so the undesired operations imprinted in those idle qubits are reversed by the second CNOT-SWAP network.

On the other hand, if the two target qubits are not adjacent to begin with, the CNOT-SWAP gate cannot universally replace the SWAP gate. Nevertheless, one can strategically reposition first the more distant target with SWAP gates that bring it closer to the other target, and then move the two target qubits together towards the control through CNOT-SWAP gates. Simultaneously, the control qubit would also be rerouted toward the target qubits through a similar network of CNOT-SWAP operations, allowing for a parallelized approach to optimize circuit depth.

In summary, when only the control qubit is non-adjacent to the target qubits, relocating it alone via cnot-swapping minimizes the CNOT count in implementing the long-range Fredkin gate. However, if the primary objective is to reduce circuit depth, this strategy results in a depth scaling of  $\sim 2n$  within the rerouting network, where  $n$  is the number of idle qubits. By contrast, concurrently repositioning all three qubits to an intermediate location achieves a depth scaling of  $\sim n$ . More precisely, when at least two idle qubits separate the control from both targets (i.e.,  $n \geq 2$ ), moving only the control qubit results in a depth of  $2n + 4$  in the rerouting network, whereas relocating all three qubits concurrently from the same initial configuration yields a depth of  $n + \frac{15+(-1)^n}{2}$  for  $n \geq 6$ , when the network's circuit pattern has fully emerged. These scalings can be easily verified through circuit sketches. The latter method involves moving

the control qubit by one position if  $n = 1$ , or by  $n + 1 - \lfloor n/2 \rfloor$  positions for  $n > 1$ , while simultaneously relocating the target qubits in tandem in the opposite direction to connect them directly to the control. This simultaneous rerouting of all three qubits can be advantageous in NISQ hardware when computation is primarily limited by decoherence rather than gate errors, as well as in fault-tolerant architectures, where it helps reduce overall execution time.

## 5.4 Equivalent circuit averaging

Building scalable quantum computers capable of executing long quantum circuits with high fidelity remains a fundamental challenge due to the presence of both decoherence and coherent noise. While the former sets a limit on circuit depth, the latter — particularly control errors, § 3.4.1 — can lead to systematic deviations that accumulate over time. As discussed in § 1.1.2, several techniques have been developed to mitigate these errors, including randomization and error suppression strategies. In this section, we focus on an approach known as Equivalent Circuit Averaging (ECA) [103, 102, 101], which leverages the execution of multiple logically equivalent circuits to transform coherent errors into stochastic noise. Following the randomization obtained with ECA, further noise suppression can be achieved through increased sampling, thereby improving the robustness of quantum computations.

Leveraging the diverse set of optimized decompositions of Fredkin and Toffoli gates presented in § 5.2, we introduce a structured ECA protocol to mitigate coherent errors in algorithms that rely on these gates. The core idea is to construct multiple logically equivalent circuits with distinct entangling-gate structures, distribute measurement shots among them, execute each instance, and finally aggregate the measurement statistics to suppress coherent noise. These steps are detailed in Algorithm 4.

To implement this protocol, we first generate  $M$  distinct but logically equivalent circuit instances, denoted as  $C_1, C_2, \dots, C_M$ . Each of these circuits is derived from the original circuit  $C$  by replacing every Fredkin or Toffoli gate that it contains with a randomly selected decomposition from our precomputed sets of Fredkin  $\mathcal{F}$  or Toffoli  $\mathcal{T}$  decompositions — summarized in Table 5.1 and available in Ref. [128]. In the presence of systematic control errors in the native gates of the quantum processor, this ensures that each instance  $C_i$  applies a slightly different unitary transformation  $C_i$  from the target unitary  $T$ , effectively diversifying the error landscape. Once the circuits are constructed, we distribute the total number of available measurement shots,  $S$ , equally among them, so that each instance is executed  $s = S/M$  times. This uniform distribution ensures that no single circuit implementation dominates the statistical outcomes.

Each circuit  $C_i$  is then executed on the quantum processor, and the measurement results  $R_i$  are

---

**Algorithm 4:** Entangling-Gate Structure-based Equivalent Circuit Averaging (ECA) protocol

---

**Data:**  $\mathcal{F}$ : Set of logically equivalent Fredkin gate decompositions with distinct EGS

$\mathcal{T}$ : Set of logically equivalent Toffoli gate decompositions with distinct EGS

**Input:**  $C$ : Quantum circuit containing Fredkin and Toffoli gates

$M$ : Number of circuit instances

$S$ : Number of measurement shots (with  $S \geq M$ )

**Output:**  $R$ : Measurement statistics dataset with reduced coherent errors

```
1 // Step 1: Construct  $M$  equivalent circuits;
2 for  $i \leftarrow 1$  to  $M$  do
3   Generate circuit  $C_i$  by replacing each Fredkin and Toffoli gate in  $C$  with a randomly selected
   decomposition from  $\mathcal{F}$  and  $\mathcal{T}$ , respectively;
4 // Step 2: Distribute the  $S$  measurement shots;
5 Assign  $s \leftarrow S/M$  shots to each circuit  $C_i$ ;
6 // Step 3: Execute circuits;
7 for  $i \leftarrow 1$  to  $M$  do
8   Execute  $C_i$  on the quantum processor with  $s$  shots;
9   Collect measurement results  $R_i$ ;
10 // Step 4: Collect measurement results;
11 return Aggregated measurement counts per computational basis state into a single data set  $R$ ;
```

---

recorded. Since different instances will exhibit slightly different coherent noise characteristics, this randomized execution prevents systematic errors from accumulating in a single direction. Finally, we compute the final relative frequencies of each outcome by aggregating over all measurement results. This process effectively transforms coherent errors into stochastic noise, reducing bias and improving computational robustness. Formally, the protocol is modeled as a combination of  $M$  unitaries, i.e., a *uniformly-mixed-unitary* channel, that transforms the input state  $\rho$  following

$$\mathcal{E}(\rho) = \frac{1}{M} \sum_{i=1}^M C_i \rho C_i^\dagger, \quad (5.7)$$

where each  $C_i$  represents the unitary transformation implemented by an individual circuit instance.

The motivation for this entangling-gate structure-based ECA strategy lies in the insight discussed in § 3, where we observed not only that the predominant source of control errors in contemporary quantum processors arises from two-qubit gate implementations, but also that these errors introduce two-qubit correlations that single-qubit gates cannot remove. Existing techniques such as Pauli twirling focus on mitigating errors by randomizing local operations to achieve global randomization. By contrast, an ensemble of circuits that directly modifies the two-qubit gate structure of the circuits, while not as flexible as Pauli twirling, offers an orthogonal and hardware-aware strategy which may serve as a complementary error-mitigation layer of randomization to Pauli twirling. The procedure uses the method detailed in § 4

to maximize the structural diversity among circuits. By distributing systematic errors across multiple logically equivalent circuits, the ECA strategy improves the reliability of quantum computations and offers a practical error-mitigation solution for near-term quantum devices.

### 5.4.1 Numerical analysis using the BCNOT model

Coherent errors, being repeatable, can in principle be suppressed through recalibration or compensation techniques. Although progress has been made in reducing their impact at the implementation level of quantum operations, achieving significant suppression remains a formidable challenge. This difficulty arises from the need for extreme precision engineering and extensive physical characterization. A major obstacle in the latter is the lack of scalable protocols that can rapidly and comprehensively characterize coherent error dynamics across all qubits and native gate implementations.

As a result, without explicit knowledge of the specific error mechanisms within a quantum device, determining *a priori* which of the  $M$  logically equivalent circuits is more robust to error is practically infeasible. However, it can be shown that the channel  $\mathcal{E}$  produced by ECA with Eq. (5.7) is guaranteed to achieve a better or equal approximation, in the worst case, to the target unitary than an average individual  $C_i$ , as quantified by the inequality

$$d_{\diamond}(\mathcal{E}, \mathcal{T}) \leq \langle d_{\diamond}(\mathcal{C}_i, \mathcal{T}) \rangle = \frac{1}{M} \sum_{i=1}^M d_{\diamond}(\mathcal{C}_i, \mathcal{T}). \quad (5.8)$$

Here, the quantum channels corresponding to the target unitary  $T$  and the unitary of the equivalent circuit,  $C_i$ , are denoted as  $\mathcal{T}(\rho) = T\rho T^\dagger$  and  $\mathcal{C}_i(\rho) = C_i\rho C_i^\dagger$ , respectively, and  $d_{\diamond}$  represents the *diamond distance*. A proof is provided in § 5.6.4. The diamond distance<sup>6</sup> is defined between two completely positive trace-preserving maps  $\mathcal{M}$  and  $\mathcal{M}'$  as

$$\begin{aligned} d_{\diamond}(\mathcal{M}, \mathcal{M}') &\equiv \|\mathcal{M} - \mathcal{M}'\|_{\diamond} \\ &= \sup_{\rho} \|(\mathcal{M} \otimes \mathcal{I})\rho - (\mathcal{M}' \otimes \mathcal{I})\rho\|_1, \end{aligned} \quad (5.9)$$

where  $\|\cdot\|_1$  is the trace norm, defined for an operator  $X$  as  $\|X\|_1 = \text{Tr}\sqrt{X^\dagger X}$ ,  $\mathcal{I}$  denotes the identity map on an ancillary system with the same dimension  $n$  as  $\mathcal{M}$  and  $\mathcal{M}'$ , and the supremum is taken over all  $n^2$ -dimensional density matrices  $\rho$ , including those entangled with the ancillary system [410, 411]. Geometrically, the diamond distance, constrained by  $0 \leq d_{\diamond} \leq 2$ , measures the maximum trace distance

<sup>6</sup>The diamond \*norm\* — often referred to as the *completely bounded trace norm* — underlies the definition of the diamond distance when applied to the difference of two quantum channels.

between the possible output states produced by the channels. Put differently, it quantifies the worst-case distinguishability between the output states of the two maps under any input state.

While Eq. (5.8) is *universal* for any set of channels  $\{\mathcal{C}_i\}_{i=1}^M$ , regardless of the error model, its practical utility depends on the structure of the errors. For example, if errors in  $\mathcal{C}_i$  are coherent and aligned (e.g., systematic over-rotations), the bound may be tight. However, if the errors are incoherent or randomized, destructive interference in  $\mathcal{E}$  can make  $d_\diamond(\mathcal{E}, \mathcal{T})$  significantly smaller than the average  $\langle d_\diamond(\mathcal{C}_i, \mathcal{T}) \rangle$ .

Thus, numerical simulations with a concrete coherent-error model are an important first step to assess whether ECA provides practical advantages over a Single Circuit Execution (SCE), as well as the tightness of the bound for specific error processes and the circuit decompositions gathered in  $\{\mathcal{C}_i\}_{i=1}^M$ . For example, if  $\mathcal{C}_i$  are Haar-randomly perturbed versions of  $\mathcal{T}$ , a rigorous analysis (e.g., via matrix concentration inequalities) shows that the diamond norm of the channel obtained by averaging over the circuits,  $d_\diamond(\mathcal{E}, \mathcal{T})$ , can decrease rapidly with  $M$ , even though the expected diamond norm of each individual circuit instance,  $\langle d_\diamond(\mathcal{C}_i, \mathcal{T}) \rangle$  clearly remains constant. The overall effect is that the probability of large deviations in the error—the worst-case error—is quickly reduced compared to that of a single instance, even though the average error per circuit does not itself decrease.

Since *in silico* quantum circuit simulators that incorporate noise typically do not account for coherent errors, the BCNOT model introduced in § 3.4.1 and Ref. [76] is particularly useful. We now employ this model to simulate the performance of our ECA protocol on a circuit with a single Fredkin or Toffoli gate making use of the gate decompositions in Table 5.1.

To analyze the impact of these errors, we modify the equivalent circuits for the Fredkin and Toffoli gates by systematically replacing each CNOT with a corresponding BCNOT gate. We assume distinct error parameters for each unique qubit pair but maintain these values fixed across multiple instances of the same two-qubit gate acting in the same qubit pair along the circuit, effectively modeling systematic miscalibrations. To assign these bias parameters  $\{\beta_j\}_{j=1}^5$ , we sample each one from a uniform distribution within the range  $[-\beta_{\max}, \beta_{\max}]$ . Once these values are assigned, the unitary representations of all the noisy circuits are constructed by replacing each CNOT with its corresponding BCNOT. The unitary representation of a given BCNOT between two qubits in a given direction is fixed for all the circuits. The diamond distance between these unitary transformations and the target operation is computed, and the average is taken across all equivalent circuits. Additionally, the diamond distance of the mixed-unitary channel generated through ECA (Eq. (5.7)) is evaluated. These calculations are implemented using the Qutip open-source library [412], employing a simplified semi-definite programming approach [413]. This process is repeated for  $B$  independent trials with different sets of sampled bias parameters at each




Figure 5.6: **Fredkin and Toffoli gate implementations via ECA under the BCNOT model.** The analysis makes use of the set of equivalent circuits summarized in Table 5.1 and considers a coherent-noise model in which every CNOT operation is replaced by a BCNOT (see § 3.4.1). The fidelity of the approximated unitary is assessed using the diamond distance  $d_{\diamond}$ , which is plotted as a function of the maximum bias ratio,  $\beta_{\max}$ , allowed in sampling the noise model. The numerical simulation methodology is detailed in the main text. For each of the 20 values of  $\beta_{\max}$ , an additional  $B = 20$  distinct BCNOT models were generated. Furthermore, the diamond distance between a single BCNOT and an ideal CNOT is plotted for reference. Data points correspond to averages and shaded regions indicate standard deviations. The results confirm that ECA significantly reduces  $d_{\diamond}$  for both Fredkin and Toffoli gates in comparison to individual circuit implementations. The consistently lower value of  $d_{\diamond}$  for the Toffoli gate, both with and without ECA, results from the Fredkin circuit requiring an additional CNOT gate, rendering it more susceptible to coherent errors. Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

fixed value of  $\beta_{\max}$ , to generalize over noise models under a given  $\beta_{\max}$ . Statistical averages along with standard deviations are computed with the resulting  $B$  values of  $d_{\diamond}(\mathcal{E}, \mathcal{T})|_{\beta_{\max}}$  and  $\langle d_{\diamond}(\mathcal{C}_i, \mathcal{T}) \rangle|_{\beta_{\max}}$ . Additionally, these simulations are iterated for different  $\beta_{\max}$  values ranging from 0 to 0.5.

The numerical results presented in Fig. 5.6 were obtained using  $B = 20$  distinct BCNOT models generated for each of 20 values of  $\beta_{\max}$ . All-to-all qubit connectivity was assumed, resulting in the use of 40 and 48 equivalent circuits to construct the ECA channels of the Fredkin and Toffoli gates, respectively (see Table 5.1). The figure illustrates that, for both the Fredkin (red) and Toffoli (blue) gates, the diamond distance relative to the uniformly-mixed-unitary channel derived through the ECA procedure,  $d_{\diamond}(\mathcal{E}, \mathcal{T})$ , is significantly smaller compared to the mean diamond distance relative to an individual equivalent circuit  $\langle d_{\diamond}(\mathcal{C}_i, \mathcal{T}) \rangle$ . As a reference, the black line represents the diamond distance between a single BCNOT

gate and the ideal CNOT, i.e.,  $d_{\diamond}(\text{BCNOT}, \text{CNOT})$  averaged over the noise models at each point. As expected, the Fredkin and Toffoli circuits accumulated more errors than a single BCNOT due to their higher gate count, leading to higher diamond distance values. The Toffoli circuits exhibit a consistently lower diamond distance than the Fredkin circuits. This asymmetry is most likely explained by the fact that, with 6 CNOT gates, the Toffoli gate has one less CNOT in its decomposition relative to the Fredkin.

### 5.4.2 Experimental validation on superconducting hardware

The previous section studied circuits with a single Fredkin or Toffoli gate using *in silico* simulations. We now report on an experimental study using an IBM Quantum processor to evaluate the performance of ECA on a real device. For this purpose, the SWAP test [363, 364] was selected as the experimental testbed. The SWAP test is a well-established quantum algorithm designed to estimate the fidelity  $F = |\langle \psi_1 | \psi_2 \rangle|^2$  between two given quantum states  $|\psi_1\rangle$  and  $|\psi_2\rangle$  without requiring full state tomography. The execution of the SWAP test on single-qubit states requires three qubits: one for preparation of each state and an auxiliary qubit for measurement in the computational basis to extract the fidelity estimation as  $\hat{F} \equiv \langle Z \rangle = P(0) - P(1)$ . This protocol primarily consists of two Hadamard gates and a single Fredkin gate, which was implemented using the aforementioned equivalent and optimized gate decompositions.

Given the connectivity constraints of the processor, we specifically tested the CSWAP gate decompositions designed for linear connectivity, with the control qubit positioned at the edge. We then gathered equivalent circuits with eight distinct entangling gate structures, as summarized in Table 5.1. It is important to note that, although not included in Table 5.1, the method developed in § 4 generated numerous additional equivalent circuits differing only in single-qubit gates. These circuits were not initially counted, as priority was given to variations in two-qubit gate structures due to the greater impact of two-qubit gate errors over single-qubit ones. Moreover, techniques such as randomized compiling [100] are already effective in introducing diversity in single-qubit gates when needed.

Nonetheless, since these circuits were already available, we leveraged them to enhance the diversity of our pool of equivalent decompositions. As a result, a total of 404 circuits with the same minimal CNOT count were transpiled into the native gate set of IBM Quantum,  $\mathcal{G}_{\text{IBMQ}} = \{\text{CNOT}, R_z, \sqrt{X}, X\}$ , and subsequently sorted by depth. Since circuit depths varied widely, ranging from 28 to 43, a cutoff depth of 31 was imposed to ensure a balanced representation of all eight entangling gate structures. In the end, only 40 CSWAP circuits were selected, including 2 with depth 28, 7 with depth 29, 6 with depth 30, and 25 with depth 31. The distribution of circuits among the eight entangling structures was as follows: 6, 3, 5, 3, 7, 1, 9, and 6.

The experimental procedure began with generating 200 pairs of Haar random single-qubit states by applying a Haar random unitary to a fixed initial state, ensuring that all chosen pairs satisfied  $F > 0.01$ . Fidelity estimation for each pair of states was performed using two approaches: SCE and ECA. Importantly, each experiment executed both protocols on the same pair of states, with experimental conditions varying between different pairs. Each protocol was allocated  $S = 980,000$  measurement shots. In the SCE method, a single decomposition was randomly selected from the 40 equivalent Fredkin implementations, and the **SWAP** test circuit was constructed by incorporating the corresponding state preparation sequence [183]. The selected circuit was then executed  $S$  times to estimate  $\hat{F}$ . For ECA, each of the eight entangling gate structures was assigned an equal fraction  $s = S/8$  of shots. The circuits for each structure were sampled from the respective Fredkin decompositions, and after incorporating the same state preparation circuit as before and executing the full algorithm for  $s$  shots, the combined measurement outcomes were used to compute  $\hat{F}$ .

To ensure a fair comparison, all circuits associated with a given pair of quantum states were executed within the same *job* on `ibmq_lima`, a Falcon r4T processor with Quantum Volume 8, freely accessible in the IBM Quantum cloud. Additionally, two extra circuits per job were used to calibrate the measurement error mitigation (MEM) protocol [414]. Each job sequentially cycled through all the circuits, executing one shot at a time before starting each subsequent round of shots and reaching the final shot count number of 20,000. With 98 circuits per job — 49 copies of one of the circuits for the SCE protocol, and 49 logically equivalent circuits for the ECA protocol — we then reached  $49 \times 20,000 = 980,000$  shots per protocol. With the randomly selected 200 pairs of states and an approximate execution time of 9 minutes per job, the overall runtime for all independent experiments was around 30 hours, distributed over two days<sup>7</sup>.

Post-experimental analysis revealed some expected deviations between the observed fidelity, denoted as  $\hat{F}$ , and the theoretical expectation,  $F$ . In both protocols, when  $F \approx 0$ , the measured fidelity  $\hat{F}$  tends to slightly exceed the theoretical value due to the added challenge in further reducing the error by random noise processes during the execution. Conversely, in cases where  $F \approx 1$ , any introduced error systematically reduces  $\hat{F}$ , highlighting the sensitivity of fidelity estimation in this extreme regimes. A more detailed investigation is provided in Fig. 5.7, which illustrates the relative accuracy error in fidelity estimation, defined as  $\varepsilon = |\hat{F} - F|/F$ , both with and without the ECA protocol. The figure demonstrates the substantial improvements achieved with ECA, which not only suppresses the magnitude of the errors but

---

<sup>7</sup>It is instructive to mention that `ibmq_lima` may have been recalibrated during this period. However, all executions assessing the performance of the **SWAP** test for a single pair of random states, under both protocols, were carefully prepared to be completed within a single, uninterrupted job lasting approximately 9 minutes. Therefore, analyzing the results of all experiments collectively remains valid, even if recalibration occurred between jobs. The primary goal of this experiment is to determine whether ECA exhibits improved performance compared to SCE within each job, independently, rather than comparing results from different jobs.




Figure 5.7: **Experimental evaluation of ECA performance.** The results from the SWAP test applied to 200 pairs of Haar-random single-qubit states of varying fidelities and executed on IBM Corp.'s `ibmq_lima` quantum processor are shown for two execution protocols: SCE and ECA. ECA is assessed by comparing its performance against SEC using the relative accuracy error metric  $\varepsilon$ . The introduction of ECA leads to a substantial decrease in the average  $\varepsilon$  value (indicated by horizontal colored lines) from 0.28 to 0.17, achieved by aggregating shot statistics from multiple circuit decompositions. Additionally, ECA also noticeably lowers the standard deviation of errors, reducing it from 0.56 to 0.25. The marginal distribution histograms are also presented. Experiments were uniformly conducted across the fidelity range  $F \in [0.01, 1]$ . Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

also mitigates their variability. To quantitatively support this observation, we computed both the mean ( $\bar{\varepsilon}$ ) and standard deviation ( $\sigma_{\varepsilon}$ ) of all relative errors. The ECA methodology significantly improves error suppression in comparison to the SCE approach, reducing the mean relative error from 28% to 17%, while also lowering the standard deviation from 56% to 25%.

Although Fig. 5.7 focuses exclusively on the impact of ECA, it is worth noting that this protocol can be effectively combined with additional error mitigation techniques, such as measurement error mitigation (MEM) [414, 415] or the more directly related Pauli twirling method. While not explicitly depicted in the figure, we conducted an auxiliary comparison between ECA and MEM [414] to further validate our approach. Specifically, applying MEM to the SCE protocol results in only a marginal improvement, reducing  $\bar{\varepsilon}$  from 28% to 27%, with an increase in  $\sigma_{\varepsilon}$  to 62%. In contrast, the ECA approach yields a significantly

more pronounced improvement, reducing the mean relative error to 17%, as discussed in the previous paragraph. Moreover, when ECA is combined with MEM, performance slightly improves, yielding an average relative error of  $\bar{\epsilon} = 16\%$  and a correspondingly reduced standard deviation of  $\sigma_{\epsilon} = 16\%$ .

Beyond integrating MEM with ECA, it is also worthwhile to explore the potential advantages of pairing ECA with techniques specifically designed to mitigate incoherent errors. Since these errors are generally stochastic and uncorrelated with circuit decomposition, their magnitude should remain comparable across circuits of similar depth. As a result, ECA alone is not expected to directly influence incoherent errors. Therefore, employing a complementary error mitigation strategy targeting incoherent noise could provide further improvements in fidelity.

## 5.5 Conclusion

Fredkin and Toffoli gates are fundamental to quantum computing, highlighting the imperative of efficiently decomposing these three-qubit operations using only CNOT and single-qubit gates. This chapter introduced several distinct circuit decompositions for these gates, achieving the optimal CNOT count for all of them in the literature. This optimization is particularly significant for near-term quantum processors, where minimizing entangling gate overhead is crucial. Our approach, leveraging the ZX-calculus-based optimization method presented in § 4, produced substantial CNOT count reduction particularly under qubit connectivity constraints. Due to the extensive computational resources required to generate the full set of equivalent circuit decompositions, these circuits have been precomputed and stored for future use.

Beyond the case in which the three active qubits of the Toffoli and Fredkin gates are adjacent, we have also investigated scenarios in which they are separated by idle qubits in hardware subject to connectivity limitations. Notably, we introduced an enhanced qubit rerouting strategy that replaces the conventional SWAP gate with a CNOT-SWAP operation. While this transformation does not execute a complete SWAP, as it leaves one qubit in a corrupted state, it requires only two CNOT gates instead of the three needed for a standard SWAP. By employing this CNOT-SWAP-based technique, the active qubits can be efficiently repositioned adjacent to one another, enabling the optimized local decompositions of the Fredkin and Toffoli gates to be applied. At the end, CNOT-SWAP networks return the active qubits back to their starting positions while restoring the states of idle qubits. As a result, both the total CNOT count and circuit depth were significantly reduced.

The utility of the CNOT-SWAP operation extends beyond the decomposition of Fredkin and Toffoli gates. More generally, replacing standard SWAP gates with CNOT-SWAP operations enables an overall

reduction in CNOT usage for a variety of scenarios, including the rerouting of control qubits in multi-controlled operations and target qubits in multi-controlled-NOT gates. This approach is also applicable to any multi-qubit gate for which the matrix representation remains diagonal in the computational basis of the moving qubit. One particularly notable application of this *cnot-swapping* technique is the efficient decomposition of a long-range CNOT, where the control and target qubits are not directly connected. With CNOT-SWAP networks, one can achieve the optimal CNOT count for cases where either one or two idle qubits separate the active qubits, as verified through an exhaustive search restricted to circuits composed exclusively of CNOT gates. For higher number of qubits,  $4n$  CNOT gates are required to implement a long-range CNOT for arbitrary  $n$  number of idle qubits, with circuit depth scaling as  $\sim n$ .

The availability of multiple logically equivalent circuits with different CNOT structures for implementing Fredkin and Toffoli gates provides new opportunities for global circuit optimization and advanced error mitigation. In this work, we explored equivalent circuit averaging, a technique in which measurement results from multiple logically equivalent circuits are aggregated instead of repetitively executing a single optimized circuit. This approach mitigates the impact of systematic coherent errors. By employing the realistic noise model introduced in § 3.4.1, it was demonstrated that the application of ECA produces an effective noisy channel that better approximates the ideal Fredkin and Toffoli unitaries, as quantified via the diamond distance. Furthermore, to validate ECA on real quantum hardware, we conducted an experiment on an IBM Quantum processor, estimating the fidelity between pairs of single-qubit states using a SWAP test. The implementation of ECA led to a reduction in both the average relative error and its variance when compared to conventional single-circuit execution. This technique also shows potential to be combined with other mitigation methods, such as standard measurement error mitigation, Pauli twirling or more sophisticated tensor-network error mitigation [105, 93] to further improve performance.

The decompositions presented in this work should be broadly applicable both to noisy and fault-tolerant quantum computing platforms. They are expected to be particularly beneficial for solid-state architectures such as superconducting circuits and silicon quantum dots, where connectivity constraints pose significant challenges. However, even systems with all-to-all connectivity, such as those based on trapped ions and cold atoms, may benefit from these optimized decompositions, particularly in contexts such as equivalent circuit averaging for coherent noise mitigation and general circuit optimization. While our analysis in Section 5.4 focused on superconducting qubit implementations, future studies exploring alternative hardware architectures are encouraged to assess the general applicability of our findings.

One promising direction for future work could be in exploring specialized hardware calibrations that exploit the availability of multiple gate decompositions. Instead of compiling all circuits at once, one might

adaptively select from a pool of Fredkin and Toffoli decompositions based on real-time noise statistics, achieving dynamic error avoidance. Such an adaptive scheme could further boost performance beyond uniform random averaging.

## 5.6 Appendix

### 5.6.1 Simplifying a controlled gate via symmetric decomposition

Consider an  $n$ -qubit gate  $U$  with a symmetric decomposition  $U = V^\dagger W V$ , where  $V$  and  $W$  are also  $n$ -qubit gates. We aim to construct a controlled version of  $U$ , denoted as

$$CU = \begin{pmatrix} \mathbb{1} & 0 \\ 0 & U \end{pmatrix} \quad (5.10)$$

and depicted in the left-hand side of the equality in Fig. 5.8(c). We will demonstrate that the controls on  $V$  and  $V^\dagger$ , as assumed in the straightforward approach illustrated in Fig. 5.8(b), can be omitted. Controlling only to the central gate  $W$ , as shown in Fig. 5.8(c), is enough. To show that, we start from the right-hand side of Fig. 5.8(c) and simplify it as follows:

$$(\mathbb{1} \otimes V^\dagger)CW(\mathbb{1} \otimes V) = \begin{pmatrix} V^\dagger & 0 \\ 0 & V^\dagger \end{pmatrix} \begin{pmatrix} \mathbb{1} & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} V & 0 \\ 0 & V \end{pmatrix} \quad (5.11)$$

$$= \begin{pmatrix} V^\dagger V & 0 \\ 0 & V^\dagger W V \end{pmatrix} \quad (5.12)$$

$$= \begin{pmatrix} \mathbb{1} & 0 \\ 0 & U \end{pmatrix} \quad (5.13)$$

Where, in the last step, we used the unitarity of  $V$ , i.e.,  $V^\dagger V = \mathbb{1}$ , and replaced the symmetric decomposition  $U = V^\dagger W V$ . This demonstrates that controlling  $V$  and  $V^\dagger$  is unnecessary, as controlling only  $W$  suffices to achieve the desired controlled operation on  $U$ .

### 5.6.2 Reassigning the target qubit in a Toffoli gate

The target qubit of the Toffoli gate can be reassigned as a control qubit, and vice versa, by applying Hadamard gates before and after the gate operation at the respective qubits. Specifically, placing a




Figure 5.8: **Quantum circuit to control a symmetric  $n$ -qubit gate.** **a**, Circuit representation of a  $n$ -qubit gate with symmetric decomposition  $U = V^\dagger W V$ . **b**, Straightforward approach to control  $U$  by adding controls to every gate. **c**, Alternative approach where it suffices to control the central gate  $W$ . Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

Hadamard gate on both the original target qubit and one of the control qubits achieves this transformation, as illustrated in Fig. 5.9(a). This property follows directly from the well-established identity  $HXH = Z$  and the inherent symmetry of the controlled-controlled- $Z$  (CCZ) gate, which remains unchanged under any permutation of the three qubits involved in its operation.




Figure 5.9: **Reassigning the Target Qubit in a Toffoli Gate.** **a**, A method to interchange the target and control qubits of a Toffoli gate using Hadamard gates applied before and after the operation at the designated qubits. **b**, A two-qubit equivalent transformation demonstrating the reversal of a CNOT gate direction. **c**, Generalization to a multi-controlled-Toffoli gate, where  $n = m_1 + m_2 + m_3 + 1$  is the number of control qubits. Reproduced from [129]. © 2024 Author(s). Licensed under a Creative Commons Attribution (CC BY) license <http://creativecommons.org/licenses/by/4.0/>.

This transformation is a direct extension of a well-known two-qubit property, depicted in Fig. 5.9(b), in which the control and target qubits of a CNOT gate are swapped by applying Hadamard gates to both qubits before and after the operation. More generally, as shown in Fig. 5.9(c), this technique can be applied to an arbitrary number of control qubits. A multi-controlled-Toffoli (MCX) gate can be transformed into a multi-controlled- $Z$  (MCZ) gate by applying Hadamard gates on the target qubit. Subsequently, by placing another set of Hadamard gates at a different qubit, the MCZ operation can be converted back into an MCX gate with a new target qubit.

It is important to note that this transformation holds exclusively for operations with a single target qubit. That is, applying multiple pairs of Hadamard gates at different qubits in a MCZ gate does not produce a multi-target multi-controlled-Toffoli gate. Instead, each transformation must be applied individually,

ensuring that only one target qubit is modified at a time.

### 5.6.3 Gate diagonality condition in a qubit's computational basis

Consider a two-qubit Hilbert space spanned by the computational basis states  $\{|ab\rangle \mid a, b \in \{0, 1\}\}$ , where  $a$  and  $b$  denote the states of qubit 1 and qubit 2, respectively. A two-qubit operator  $V$  is said to be *diagonal* with respect to qubit 1 if it does not couple basis states that differ solely in the state of qubit 1. In Dirac notation, this condition can be expressed as:

$$\langle 1b'|V|0b\rangle = 0 \quad \text{and} \quad \langle 0b'|V|1b\rangle = 0 \quad \forall b, b' \in \{0, 1\}. \quad (5.14)$$

In other words, the matrix elements of  $V$  connecting  $|0b\rangle$  with  $|1b'\rangle$  must vanish for all  $b$  and  $b'$ . This ensures that, for any fixed  $b$ , the subspaces spanned by  $\{|0b\rangle\}$  and  $\{|1b\rangle\}$  remain invariant under  $V$ . As a result, in the  $4 \times 4$  matrix representation of  $V$ , ordered according to the computational basis  $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ , the operator takes on a block-diagonal form:

$$V = \begin{pmatrix} V_0 & 0 \\ 0 & V_1 \end{pmatrix}, \quad (5.15)$$

where  $V_0$  and  $V_1$  are  $2 \times 2$  submatrices corresponding to the subspaces where qubit 1 is in the states  $|0\rangle$  and  $|1\rangle$ , respectively. Similarly,  $V$  is diagonal with respect to qubit 2 if it does not couple basis states that differ only in the state of qubit 2. This condition can be stated formally as:

$$\langle a1|V|a'0\rangle = 0 \quad \text{and} \quad \langle a0|V|a'1\rangle = 0 \quad \forall a, a' \in \{0, 1\}. \quad (5.16)$$

In this case, matrix elements connecting  $|a0\rangle$  and  $|a'1\rangle$  must vanish for all  $a$  and  $a'$ . If the computational basis is reordered such that all states with qubit 2 in the state  $|0\rangle$  appear before those with qubit 2 in the state  $|1\rangle$ , the matrix representation of  $V$  acquires a block-diagonal structure.

### 5.6.4 Proof of the diamond distance bound for ECA

We consider a set of  $M$  quantum circuits, each described by a unitary  $C_i$ . The corresponding quantum channels are given by  $\mathcal{C}_i(\rho) = C_i \rho C_i^\dagger$ . Since the circuits  $\{C_i\}$  are logically equivalent, each one is expected to approximate the target unitary  $T$ , which defines the target channel  $\mathcal{T}(\rho) = T \rho T^\dagger$ . The effective channel produced by ECA is a uniform mixture of the unitary channels corresponding to each

circuit, given in Eq. (5.7) as  $\mathcal{E}(\rho) = \frac{1}{M} \sum_{i=1}^M C_i \rho C_i^\dagger = \frac{1}{M} \sum_{i=1}^M \mathcal{C}_i(\rho)$ . Our goal is to show that the worst-case error, as measured by the diamond distance, satisfies the bound in Eq. (5.8), repeated below for convenience

$$d_\diamond(\mathcal{E}, \mathcal{T}) \leq \frac{1}{M} \sum_{i=1}^M d_\diamond(\mathcal{C}_i, \mathcal{T}), \quad (5.17)$$

which ensures that the diamond distance between the ensemble-averaged channel  $\mathcal{E}$  and the target unitary channel  $\mathcal{T}$  is at most the arithmetic mean of the individual distances  $d_\diamond(\mathcal{C}_i, \mathcal{T})$ .

To prove Eq. (5.17), we begin by writing the difference between the channels:

$$\mathcal{E} - \mathcal{T} = \frac{1}{M} \sum_{i=1}^M (\mathcal{C}_i - \mathcal{T}). \quad (5.18)$$

Taking the diamond norm on both sides and applying the convexity property of the norm, we obtain

$$\|\mathcal{E} - \mathcal{T}\|_\diamond \leq \frac{1}{M} \sum_{i=1}^M \|\mathcal{C}_i - \mathcal{T}\|_\diamond. \quad (5.19)$$

This inequality follows from a fundamental property of the diamond norm: for any set of completely positive trace-preserving maps  $\mathcal{A}_i$  and nonnegative coefficients  $\alpha_i$  summing to 1,

$$\left\| \sum_i \alpha_i \mathcal{A}_i \right\|_\diamond \leq \sum_i \alpha_i \|\mathcal{A}_i\|_\diamond. \quad (5.20)$$

This result is a direct consequence of the triangle inequality and subadditivity of the trace norm under quantum operations. Since the diamond distance is defined as  $d_\diamond(\mathcal{M}, \mathcal{M}') \equiv \|\mathcal{M} - \mathcal{M}'\|_\diamond$ , applying this to Eq. (5.19) directly proves the desired bound in Eq. (5.17).



# Conclusion

Computer science is undergoing a historical development as practical quantum computers begin to emerge as a complementary paradigm to classical computation. Unlike the advances of previous decades built upon continuous improvements to a single paradigm — classical, transistor-based machines following Moore’s Law — quantum computing introduces a fundamentally distinct mathematical framework. By leveraging superposition, entanglement, and interference, it opens new computational frontiers that transcend the capabilities of classical systems. Realizing the full potential of this paradigm requires progress on multiple fronts. Theoretical advancements are needed to develop new algorithms and protocols that achieve reliable quantum advantages over the best-known classical counterparts across various applications. At the same time, experimental innovations are crucial to realize those algorithms on real hardware, overcoming practical challenges such as decoherence, gate errors, and limited qubit counts.

This dissertation contributed to this overarching effort by addressing interconnected theoretical and experimental challenges in quantum computing within the framework of the quantum circuit model. By exploring fundamental quantum principles and classical optimization techniques, this thesis presented new results on quantum protocols, algorithmic techniques, and hardware-oriented developments, bridging conceptual insights with practical implementations.

One of the key objectives in quantum computing research is to establish clear instances of quantum advantage, demonstrating scenarios where quantum resources outperform classical counterparts. Despite the growing repertoire of quantum algorithms, novel examples of fundamental quantum advantage remain scarce, often constrained to highly structured problems or specific computational paradigms. This thesis introduced a quantum protocol for solving the *Hardest Logic Puzzle Ever* (HLPE), an interactive logical problem requiring a solver to deduce knowledge from minimal data amid randomness and linguistic ambiguity (§ 2). By leveraging quantum communication features, the proposed protocol provides a solution that classical strategies cannot replicate, showcasing a concrete instance of fundamental quantum advantage in a non-traditional computational setting. While the puzzle itself is a specialized problem, the HLPE quantum protocol goes beyond its immediate application and exemplifies the broader potential of

quantum procedures to tackle reasoning and deduction tasks by harnessing uniquely quantum effects.

The exploration of fundamental quantum principles is essential not only for deepening our understanding of quantum mechanics but also for assessing the quantum processing power of existing hardware. The experimental study of quantum complementarity presented in this dissertation provided an empirical test of quantum mechanics on superconducting quantum processors (§ 3). By implementing two-path interferometry experiments coupled to a which-path detector, the visibility-distinguishability tradeoff was tested with unprecedented precision. More significantly, the results revealed that even small coherent errors in two-qubit gate operations fundamentally constrain the extent to which a processor exhibits fully quantum or classical behavior. These findings have direct implications for quantum computing, highlighting the necessity of minimizing coherent two-qubit gate errors to preserve the quantum nature of the computational processes. Furthermore, they provided valuable insights into the physical limitations of current quantum processors, informing subsequent work in circuit optimization and error mitigation.

Optimizing quantum circuits is a crucial step toward practical quantum computing, as circuit depth and gate complexity directly impact performance on noisy devices. This thesis introduced a general framework for quantum circuit optimization that robustly reduces two-qubit gate count (§ 4). By leveraging entangling-gate structure classifications, search techniques, and heuristics-based circuit and ZX-diagram simplifications, this approach systematically generated optimized quantum circuits across arbitrary connectivity topologies. This method was applied across projects to optimize circuits used for the implementation of a square-root measurement POVM, primitives for the preparation of VBS states, and Fredkin and Toffoli gates. Besides improvements over previous methods, as demonstrated for the tested use cases, the circuit optimization method herein presented also offers a structured approach for algorithm discovery. By analyzing the underlying entangling structures of circuits, the method enables a deeper understanding of how unitary operations can be efficiently decomposed, providing a pathway for identifying new quantum algorithms through artificial intelligence methods built upon it.

A key theme of this work is the role of co-design of quantum circuits and algorithms—the integration of classical optimization techniques and artificial intelligence tools to enhance quantum computation. The methodology developed in this thesis applies classical resources to systematically refine quantum circuits by searching for optimal circuit structures, tuning parameters, and even learning from data to suggest model improvements. Optimized circuits have also facilitated the identification of general algorithmic structures, leading to the proposal of the *cnot-swapping* methodology, while the introduction of the entangling-gate structure classification aims to further advance this research direction. This interplay between classical and quantum paradigms was central to this research, demonstrating that classical

computing remains an indispensable tool in quantum algorithm design.

The insights gained from the complementarity experiments and circuit optimization directly informed the development of optimized Fredkin and Toffoli gate decompositions (§ 5). These multi-qubit gates serve as fundamental building blocks in many quantum algorithms, and their efficient realization contributes to the broader goal of reducing the resource overhead associated with quantum computations. The optimized decompositions presented in this dissertation reduce the gate count required to implement these operations while also improving their routing on constrained qubit topologies, thereby enhancing their feasibility for both near- and long-term quantum devices. Furthermore, by providing a library of equivalent circuits for these gates, this work enables tailoring quantum circuits to specific error profiles in physical hardware, complementing existing error mitigation techniques such as Pauli twirling by addressing coherent errors through equivalent circuit averaging. This approach exploits multiple logically equivalent decompositions of the same unitary transformation carrying different entangling-gate structures, distributing errors in a manner that improves computational accuracy without requiring additional hardware overhead.

These error mitigation strategies are particularly well-suited for current quantum processors, where coherent errors in two-qubit gates remain a dominant source of computational inaccuracies. By systematically varying the entangling-gate structure of equivalent circuits, the method effectively neutralizes systematic biases in gate execution, offering a practical means to improve algorithm fidelity. Experimental validation of this technique on superconducting quantum hardware confirmed its effectiveness, demonstrating that strategic circuit design can significantly enhance computational reliability.

A defining aspect of this thesis was the extensive experimental validation of theoretical proposals on real quantum processors. The implementation of quantum circuits on IBM hardware required careful consideration of noise, hardware constraints, and calibration procedures. One key lesson was the importance of understanding how jobs are executed on this quantum cloud service at a low level, including details such as qubit calibration schedules, shot sequences, queueing systems, and the effects of concurrent users on a device. To ensure data integrity, rigorous experimental protocols were established, carefully managing execution conditions to reflect the true performance of quantum circuits rather than artifacts of external disturbances. For example, multiple runs of the same *which-path* quantum circuits were often necessary to gather sufficient statistics, yet fluctuations in hardware performance over time posed a challenge when using shared public resources. To mitigate this, the experiments in this case strategically employed reserved quantum processors for extended periods to ensure stable calibrations and predictable execution conditions. The ability to manage these hardware-level constraints was crucial for obtaining reliable data, with post-execution analysis confirming the consistency and robustness of the results.

Beyond individual experiments, this research underscored the importance of hardware awareness in quantum computing. Achieving optimal performance on near-term quantum devices requires an intimate understanding of how qubits are controlled, how noise manifests at different operational layers, and how job execution pipelines affect results. Engineering diligence played a critical role in this process, as effective experimental execution required custom low-level code, careful hardware-specific optimizations, and continuous monitoring of device performance to adapt to any drifts. These efforts ensured that the theoretical contributions were validated in a realistic setting, with experimental findings backed by well-managed and carefully controlled data collection. The hands-on experience gained through this work demonstrated that experimental expertise is indispensable for harnessing the full potential of existing quantum processors, bridging the gap between theoretical quantum computing and practical implementation.

The contributions of this dissertation provide concrete examples of how hybrid quantum-classical methodologies and algorithmic co-design can be effectively applied, illustrating their role in broader developments within quantum computing. As hardware capabilities improve, the integration of classical and quantum computational resources will become increasingly sophisticated, with AI-driven approaches playing a central role in optimizing quantum circuits and algorithm discovery. Future research will likely continue refining hybrid strategies, leveraging machine learning to autonomously explore quantum circuit spaces, adaptively optimizing execution protocols, and error mitigation strategies.

A key lesson from this research is that quantum algorithm design cannot be divorced from hardware considerations. The pursuit of quantum advantage will require continued refinements in circuit synthesis, error mitigation, and co-design principles, ensuring that quantum algorithms are tailored to the strengths and constraints of the hardware that executes them. Moving forward, interdisciplinary collaborations will be essential in steering quantum computing from an experimental discipline to a mature technological paradigm. The advances presented in this dissertation contribute to this ongoing effort, providing both foundational insights and practical tools to navigate the next stages of quantum computing's evolution.

# Publications by the author

## Presented in this thesis

1. M. Q. Cruz, P. & Fernández-Rossier, J. “Testing complementarity on a transmon quantum processor”. *Phys. Rev. A* **104**, 032223 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevA.104.032223>
2. M. Q. Cruz, P. & Murta, B. “Shallow unitary decompositions of quantum Fredkin and Toffoli gates for connectivity-aware equivalent circuit averaging”. *APL Quantum* **1**, 016105 (2024). URL <https://doi.org/10.1063/5.0187026>
3. M. Q. Cruz, P. & Murta, B. “Fredkin and Toffoli quantum gate decompositions in diverse connectivity topologies”. *Zenodo* (2023). URL <https://doi.org/10.5281/zenodo.10047422>
4. M. Q. Cruz, P. & Acín, A. “Quantum advantage in the *Hardest Logic Puzzle Ever*”. (*in preparation*)
5. M. Q. Cruz, P. & Pedroso, J. P. “QCMotherLode: Autonomous discovery and optimization of logically-equivalent quantum circuits”. (*in preparation*)

## Other publications

1. Murta, B., M. Q. Cruz, P. & Fernández-Rossier, J. “Preparing valence-bond-solid states on noisy intermediate-scale quantum computers”. *Phys. Rev. Research* **5**, 013190 (2023). URL <https://doi.org/10.1103/PhysRevResearch.5.013190>
2. M. Q. Cruz, P., Catarina, G., Gautier, R. & Fernández-Rossier, J. “Optimizing quantum phase estimation for the simulation of hamiltonian eigenstates”. *Quantum Science and Technology* **5**, 044005 (2020). URL <https://dx.doi.org/10.1088/2058-9565/abaa2c>

# Bibliography

- [1] Feynman, R. P. Simulating physics with computers. *International Journal of Theoretical Physics* **21**, 467–488 (1982). URL <https://doi.org/10.1007/BF02650179>.
- [2] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996). URL <https://www.science.org/doi/abs/10.1126/science.273.5278.1073>.
- [3] Shor, P. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134 (1994). URL <https://www.doi.org/10.1109/SFCS.1994.365700>.
- [4] Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review* **41**, 303–332 (1999). URL <https://doi.org/10.1137/S0036144598347011>.
- [5] Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, 212–219 (Association for Computing Machinery, New York, NY, USA, 1996). URL <https://doi.org/10.1145/237814.237866>.
- [6] Coppersmith, D. An approximate Fourier transform useful in quantum factoring. Research report RC 19642, IBM T.J. Watson Research Division, Yorktown Heights, New York (USA) (1994).
- [7] Cleve, R., Ekert, A., Macchiavello, C. & Mosca, M. Quantum algorithms revisited. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **454**, 339–354 (1998). URL <https://doi.org/10.1098/rspa.1998.0164>.
- [8] Brassard, G., Høyer, P., Mosca, M. & Tapp, A. Quantum amplitude amplification and estimation (2002). URL <http://dx.doi.org/10.1090/conm/305/05215>.
- [9] Yoder, T. J., Low, G. H. & Chuang, I. L. Fixed-point quantum search with an optimal number of queries. *Phys. Rev. Lett.* **113**, 210501 (2014). URL <https://link.aps.org/doi/10.1103/PhysRevLett.113.210501>.
- [10] Aharonov, Y., Davidovich, L. & Zagury, N. Quantum random walks. *Phys. Rev. A* **48**, 1687–1690 (1993). URL <https://link.aps.org/doi/10.1103/PhysRevA.48.1687>.
- [11] Trotter, H. F. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society* **10**, 545–551 (1959). URL <https://doi.org/10.1090/S0002-9939-1959-0108732-6>.
- [12] Suzuki, M. Generalized trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics* **51**, 183–190 (1976). URL <https://doi.org/10.1007/BF01609348>.

- [13] Childs, A. M., Su, Y., Tran, M. C., Wiebe, N. & Zhu, S. Theory of trotter error with commutator scaling. *Phys. Rev. X* **11**, 011020 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevX.11.011020>.
- [14] Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution (2000). URL <https://arxiv.org/abs/quant-ph/0001106>. quant-ph/0001106.
- [15] Farhi, E. & Gutmann, S. Quantum computation and decision trees. *Phys. Rev. A* **58**, 915–928 (1998). URL <https://link.aps.org/doi/10.1103/PhysRevA.58.915>.
- [16] Childs, A. M. Universal computation by quantum walk. *Phys. Rev. Lett.* **102**, 180501 (2009). URL <https://link.aps.org/doi/10.1103/PhysRevLett.102.180501>.
- [17] Childs, A. M. & Wiebe, N. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Information and Computation* **12** (2012). URL <http://doi.org/10.26421/qic12.11-12>.
- [18] Berry, D. W., Childs, A. M., Cleve, R., Kothari, R. & Somma, R. D. Simulating hamiltonian dynamics with a truncated taylor series. *Phys. Rev. Lett.* **114**, 090502 (2015). URL <https://link.aps.org/doi/10.1103/PhysRevLett.114.090502>.
- [19] Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009). URL <https://link.aps.org/doi/10.1103/PhysRevLett.103.150502>.
- [20] Low, G. H. & Chuang, I. L. Optimal hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.* **118**, 010501 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.010501>.
- [21] Gilyén, A., Su, Y., Low, G. H. & Wiebe, N. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, 193–204 (Association for Computing Machinery, New York, NY, USA, 2019). URL <https://doi.org/10.1145/3313276.3316366>.
- [22] Martyn, J. M., Rossi, Z. M., Tan, A. K. & Chuang, I. L. Grand unification of quantum algorithms. *PRX Quantum* **2**, 040203 (2021). URL <https://link.aps.org/doi/10.1103/PRXQuantum.2.040203>.
- [23] Low, G. H. & Chuang, I. L. Hamiltonian Simulation by Qubitization. *Quantum* **3**, 163 (2019). URL <https://doi.org/10.22331/q-2019-07-12-163>.
- [24] Tang, E. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, 217–228 (Association for Computing Machinery, New York, NY, USA, 2019). URL <https://doi.org/10.1145/3313276.3316310>.
- [25] Tang, E. Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Phys. Rev. Lett.* **127**, 060503 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.060503>.
- [26] Wiebe, N., Braun, D. & Lloyd, S. Quantum algorithm for data fitting. *Phys. Rev. Lett.* **109**, 050505 (2012). URL <https://link.aps.org/doi/10.1103/PhysRevLett.109.050505>.

- [27] Clader, B. D., Jacobs, B. C. & Sprouse, C. R. Preconditioned quantum linear system algorithm. *Phys. Rev. Lett.* **110**, 250504 (2013). URL <https://link.aps.org/doi/10.1103/PhysRevLett.110.250504>.
- [28] Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L. & Coles, P. J. Challenges and opportunities in quantum machine learning. *Nature Computational Science* **2**, 567–576 (2022). URL <https://doi.org/10.1038/s43588-022-00311-3>.
- [29] Chia, N.-H. *et al.* Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *J. ACM* **69** (2022). URL <https://doi.org/10.1145/3549524>.
- [30] Gharibian, S. & Le Gall, F. Dequantizing the quantum singular value transformation: Hardness and applications to quantum chemistry and the quantum PCP conjecture. *SIAM Journal on Computing* **52**, 1009–1038 (2023). URL <https://doi.org/10.1137/22M1513721>.
- [31] Motlagh, D. & Wiebe, N. Generalized quantum signal processing. *PRX Quantum* **5**, 020368 (2024). URL <https://link.aps.org/doi/10.1103/PRXQuantum.5.020368>.
- [32] Dong, Y., Meng, X., Whaley, K. B. & Lin, L. Efficient phase-factor evaluation in quantum signal processing. *Phys. Rev. A* **103**, 042419 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevA.103.042419>.
- [33] Tosta, A., de Lima Silva, T., Camilo, G. & Aolita, L. Randomized semi-quantum matrix processing. *npj Quantum Information* **10**, 93 (2024). URL <https://doi.org/10.1038/s41534-024-00883-0>.
- [34] Babbush, R., Berry, D. W., Kothari, R., Somma, R. D. & Wiebe, N. Exponential quantum speedup in simulating coupled classical oscillators. *Phys. Rev. X* **13**, 041041 (2023). URL <https://link.aps.org/doi/10.1103/PhysRevX.13.041041>.
- [35] Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nature Physics* **16**, 1050–1057 (2020). URL <https://doi.org/10.1038/s41567-020-0932-7>.
- [36] Yamakawa, T. & Zhandry, M. Verifiable quantum advantage without structure. *J. ACM* **71** (2024). URL <https://doi.org/10.1145/3658665>.
- [37] Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018). URL <https://doi.org/10.22331/q-2018-08-06-79>.
- [38] Cerezo, M. *et al.* Variational quantum algorithms. *Nature Reviews Physics* **3**, 625–644 (2021). URL <https://doi.org/10.1038/s42254-021-00348-9>.
- [39] Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* **5**, 4213 (2014). URL <https://doi.org/10.1038/ncomms5213>.
- [40] Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications* **10**, 3007 (2019). URL <https://doi.org/10.1038/s41467-019-10988-2>.
- [41] Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028* (2014). URL <https://arxiv.org/abs/1411.4028>.

- [42] Tilly, J. *et al.* The variational quantum eigensolver: A review of methods and best practices. *Physics Reports* **986**, 1–128 (2022). URL <https://doi.org/10.1016/j.physrep.2022.08.003>.
- [43] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nature Communications* **9**, 4812 (2018). URL <https://doi.org/10.1038/s41467-018-07090-4>.
- [44] Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse ising model. *Phys. Rev. E* **58**, 5355–5363 (1998). URL <https://link.aps.org/doi/10.1103/PhysRevE.58.5355>.
- [45] Rajak, A., Suzuki, S., Dutta, A. & Chakrabarti, B. K. Quantum annealing: an overview. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **381**, 20210417 (2023). URL <https://www.doi.org/10.1098/rsta.2021.0417>.
- [46] Born, M. & Fock, V. Beweis des adiabatensatzes. *Zeitschrift für Physik* **51**, 165–180 (1928). URL <https://doi.org/10.1007/BF01343193>.
- [47] Motta, M. *et al.* Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics* **16**, 205–210 (2020). URL <https://doi.org/10.1038/s41567-019-0704-4>.
- [48] Murta, B. & Fernández-Rossier, J. Gutzwiller wave function on a digital quantum computer. *Phys. Rev. B* **103**, L241113 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevB.103.L241113>.
- [49] Murta, B., Cruz, P. M. Q. & Fernández-Rossier, J. Preparing valence-bond-solid states on noisy intermediate-scale quantum computers. *Phys. Rev. Res.* **5**, 013190 (2023). URL <https://link.aps.org/doi/10.1103/PhysRevResearch.5.013190>.
- [50] Cruz, P. M. Q., Catarina, G., Gautier, R. & Fernández-Rossier, J. Optimizing quantum phase estimation for the simulation of hamiltonian eigenstates. *Quantum Science and Technology* **5**, 044005 (2020). URL <https://dx.doi.org/10.1088/2058-9565/abaa2c>.
- [51] Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, 333–342 (Association for Computing Machinery, New York, NY, USA, 2011). URL <https://doi.org/10.1145/1993636.1993682>.
- [52] Watrous, J. Succinct quantum proofs for properties of finite groups. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, 537–546 (2000).
- [53] Dalzell, A. M. *et al.* *Quantum Algorithms: A Survey of Applications and End-to-end Complexities* (Cambridge University Press, 2025). URL <https://doi.org/10.1017/9781009639651>.
- [54] Arnault, P., Arrighi, P., Herbert, S., Kasnetsi, E. & Li, T. A typology of quantum algorithms. *arXiv preprint arXiv:2407.05178* (2024). URL <https://arxiv.org/abs/2407.05178>.
- [55] Cheng, B. *et al.* Noisy intermediate-scale quantum computers. *Frontiers of Physics* **18**, 21308 (2023). URL <https://doi.org/10.1007/s11467-022-1249-z>.
- [56] Blais, A., Grimsmo, A. L., Girvin, S. M. & Wallraff, A. Circuit quantum electrodynamics. *Rev. Mod. Phys.* **93**, 025005 (2021). URL <https://link.aps.org/doi/10.1103/RevModPhys.93.025005>.

- [57] Bruzewicz, C. D., Chiaverini, J., McConnell, R. & Sage, J. M. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews* **6**, 021314 (2019). URL <https://doi.org/10.1063/1.5088164>.
- [58] Henriët, L. *et al.* Quantum computing with neutral atoms. *Quantum* **4**, 327 (2020). URL <https://doi.org/10.22331/q-2020-09-21-327>.
- [59] Slussarenko, S. & Pryde, G. J. Photonic quantum information processing: A concise review. *Applied Physics Reviews* **6**, 041303 (2019). URL <https://doi.org/10.1063/1.5115814>.
- [60] Burkard, G., Ladd, T. D., Pan, A., Nichol, J. M. & Petta, J. R. Semiconductor spin qubits. *Rev. Mod. Phys.* **95**, 025003 (2023). URL <https://link.aps.org/doi/10.1103/RevModPhys.95.025003>.
- [61] Baydin, A. *et al.* Carbon nanotube devices for quantum technology. *Materials* **15** (2022). URL <https://www.mdpi.com/1996-1944/15/4/1535>.
- [62] Aghaee, M. *et al.* Interferometric single-shot parity measurement in InAs–Al hybrid devices. *Nature* **638**, 651–655 (2025). URL <https://doi.org/10.1038/s41586-024-08445-2>.
- [63] Acharya, R. *et al.* Quantum error correction below the surface code threshold. *Nature* (2024). URL <https://doi.org/10.1038/s41586-024-08449-y>.
- [64] Quantum processing units (2024). URL <https://quantum.ibm.com/services/resources>. Accessed on Jan 1, 2025.
- [65] Evered, S. J. *et al.* High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature* **622**, 268–272 (2023). URL <https://doi.org/10.1038/s41586-023-06481-y>.
- [66] Chen, J.-S. *et al.* Benchmarking a trapped-ion quantum computer with 30 qubits. *Quantum* **8**, 1516 (2024). URL <https://doi.org/10.22331/q-2024-11-07-1516>.
- [67] Alexander, K. *et al.* A manufacturable platform for photonic quantum computing. *Nature* **641**, 876–883 (2025). URL <https://doi.org/10.1038/s41586-025-08820-7>.
- [68] Mayer, K. *et al.* Benchmarking logical three-qubit quantum fourier transform encoded in the steane code on a trapped-ion quantum computer (2024). URL <https://arxiv.org/abs/2404.08616>.
- [69] Löschnauer, C. M. *et al.* Scalable, high-fidelity all-electronic control of trapped-ion qubits. *arXiv preprint arXiv:2407.07694* (2024). URL <https://arxiv.org/abs/2407.07694>.
- [70] Proctor, T., Young, K., Baczewski, A. D. & Blume-Kohout, R. Benchmarking quantum computers. *Nature Reviews Physics* **7**, 105–118 (2025). URL <https://doi.org/10.1038/s42254-024-00796-z>.
- [71] Bartolucci, S. *et al.* Fusion-based quantum computation. *Nature Communications* **14**, 912 (2023). URL <https://doi.org/10.1038/s41467-023-36493-1>.
- [72] Alsina, D. & Latorre, J. I. Experimental test of Mermin inequalities on a five-qubit quantum computer. *Phys. Rev. A* **94**, 012314 (2016). URL <https://link.aps.org/doi/10.1103/PhysRevA.94.012314>.

- [73] Arrasmith, A., Cincio, L., Sornborger, A. T., Zurek, W. H. & Coles, P. J. Variational consistent histories as a hybrid algorithm for quantum foundations. *Nature Communications* **10**, 3438 (2019). URL <https://doi.org/10.1038/s41467-019-11417-0>.
- [74] Kirby, W. M. & Love, P. J. Contextuality test of the nonclassicality of variational quantum eigensolvers. *Phys. Rev. Lett.* **123**, 200501 (2019). URL <https://link.aps.org/doi/10.1103/PhysRevLett.123.200501>.
- [75] Kirby, W. M., Tranter, A. & Love, P. J. Contextual Subspace Variational Quantum Eigensolver. *Quantum* **5**, 456 (2021). URL <https://doi.org/10.22331/q-2021-05-14-456>.
- [76] Cruz, P. M. Q. & Fernández-Rossier, J. Testing complementarity on a transmon quantum processor. *Phys. Rev. A* **104**, 032223 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevA.104.032223>.
- [77] Sadana, S., Maccone, L. & Sinha, U. Testing quantum foundations with quantum computers. *Phys. Rev. Res.* **4**, L022001 (2022). URL <https://link.aps.org/doi/10.1103/PhysRevResearch.4.L022001>.
- [78] Aaronson, S. & Chen, L. Complexity-theoretic foundations of quantum supremacy experiments. *arXiv preprint arXiv:1612.05903* (2016). URL <https://arxiv.org/abs/1612.05903>.
- [79] Arute, F. *et al.* Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019). URL <https://doi.org/10.1038/s41586-019-1666-5>.
- [80] Wu, Y. *et al.* Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.* **127**, 180501 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.180501>.
- [81] Zhong, H.-S. *et al.* Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020). URL <https://www.science.org/doi/abs/10.1126/science.abe8770>.
- [82] Deng, Y.-H. *et al.* Gaussian boson sampling with pseudo-photon-number-resolving detectors and quantum computational advantage. *Phys. Rev. Lett.* **131**, 150601 (2023). URL <https://link.aps.org/doi/10.1103/PhysRevLett.131.150601>.
- [83] Kim, Y. *et al.* Evidence for the utility of quantum computing before fault tolerance. *Nature* **618**, 500–505 (2023). URL <https://doi.org/10.1038/s41586-023-06096-3>.
- [84] Abbas, A. *et al.* Challenges and opportunities in quantum optimization. *Nature Reviews Physics* **6**, 718–735 (2024). URL <https://doi.org/10.1038/s42254-024-00770-9>.
- [85] Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevX.7.021050>.
- [86] Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.180509>.
- [87] Endo, S., Benjamin, S. C. & Li, Y. Practical quantum error mitigation for near-future applications. *Phys. Rev. X* **8**, 031027 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevX.8.031027>.

- [88] Cai, Z. *et al.* Quantum error mitigation. *Rev. Mod. Phys.* **95**, 045005 (2023). URL <https://link.aps.org/doi/10.1103/RevModPhys.95.045005>.
- [89] Tsubouchi, K., Sagawa, T. & Yoshioka, N. Universal cost bound of quantum error mitigation based on quantum estimation theory. *Phys. Rev. Lett.* **131**, 210601 (2023). URL <https://link.aps.org/doi/10.1103/PhysRevLett.131.210601>.
- [90] Takagi, R., Tajima, H. & Gu, M. Universal sampling lower bounds for quantum error mitigation. *Phys. Rev. Lett.* **131**, 210602 (2023). URL <https://link.aps.org/doi/10.1103/PhysRevLett.131.210602>.
- [91] Quek, Y., França, D. S., Khatri, S., Meyer, J. J. & Eisert, J. Exponentially tighter bounds on limitations of quantum error mitigation. *Nature Physics* **20**, 1648–1658 (2024). URL <https://doi.org/10.1038/s41567-024-02536-7>.
- [92] Zimborás, Z. *et al.* Myths around quantum computation before full fault tolerance: What no-go theorems rule out and what they don't. *arXiv preprint arXiv:2501.05694* (2025). URL <https://arxiv.org/abs/2501.05694>.
- [93] Filippov, S. N., Maniscalco, S. & García-Pérez, G. Scalability of quantum error mitigation techniques: from utility to advantage. *arXiv preprint arXiv:2403.13542* (2024). URL <https://arxiv.org/abs/2403.13542>.
- [94] Iverson, J. K. & Preskill, J. Coherence in logical quantum channels. *New J. Phys.* **22**, 073066 (2020). URL <https://dx.doi.org/10.1088/1367-2630/ab8e5c>.
- [95] Cai, Z., Xu, X. & Benjamin, S. C. Mitigating coherent noise using Pauli conjugation. *npj Quantum Inf.* **6**, 17 (2020). URL <https://doi.org/10.1038/s41534-019-0233-0>.
- [96] Tripathi, V. *et al.* Suppression of crosstalk in superconducting qubits using dynamical decoupling. *Phys. Rev. Appl.* **18**, 024068 (2022). URL <https://link.aps.org/doi/10.1103/PhysRevApplied.18.024068>.
- [97] Brown, K. R., Harrow, A. W. & Chuang, I. L. Arbitrarily accurate composite pulse sequences. *Phys. Rev. A* **70**, 052318 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevA.70.052318>.
- [98] Geller, M. R. & Zhou, Z. Efficient error models for fault-tolerant architectures and the Pauli twirling approximation. *Phys. Rev. A* **88**, 012314 (2013). URL <https://link.aps.org/doi/10.1103/PhysRevA.88.012314>.
- [99] Ware, M. *et al.* Experimental Pauli-frame randomization on a superconducting qubit. *Phys. Rev. A* **103**, 042604 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevA.103.042604>.
- [100] Wallman, J. J. & Emerson, J. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325 (2016). URL <https://link.aps.org/doi/10.1103/PhysRevA.94.052325>.
- [101] Hastings, M. B. Turning gate synthesis errors into incoherent errors. *Quantum Info. Comput.* **17**, 488–494 (2017).

- [102] Campbell, E. Shorter gate sequences for quantum computing by mixing unitaries. *Phys. Rev. A* **95**, 042306 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevA.95.042306>.
- [103] Hashim, A. *et al.* Optimized swap networks with equivalent circuit averaging for qaoa. *Phys. Rev. Res.* **4**, 033028 (2022). URL <https://link.aps.org/doi/10.1103/PhysRevResearch.4.033028>.
- [104] van den Berg, E., Mineev, Z. K., Kandala, A. & Temme, K. Probabilistic error cancellation with sparse pauli–lindblad models on noisy quantum processors. *Nature Physics* **19**, 1116–1121 (2023). URL <https://doi.org/10.1038/s41567-023-02042-2>.
- [105] Filippov, S., Leahy, M., Rossi, M. A. & García-Pérez, G. Scalable tensor-network error mitigation for near-term quantum computing. *arXiv preprint arXiv:2307.11740* (2023). URL <https://arxiv.org/abs/2307.11740>.
- [106] Devitt, S. J., Munro, W. J. & Nemoto, K. Quantum error correction for beginners. *Rep. Prog. Phys.* **76**, 076001 (2013). URL <https://dx.doi.org/10.1088/0034-4885/76/7/076001>.
- [107] Gottesman, D. *Surviving as a Quantum Computer in a Classical World* (2024). URL <https://www.cs.umd.edu/class/spring2024/cmsc858G/>. Accessed: 2024-05-07.
- [108] Acharya, R. *et al.* Suppressing quantum errors by scaling a surface code logical qubit. *Nature* **614**, 676–681 (2023). URL <https://doi.org/10.1038/s41586-022-05434-1>.
- [109] Bausch, J. *et al.* Learning high-accuracy error decoding for quantum processors. *Nature* **635**, 834–840 (2024). URL <https://doi.org/10.1038/s41586-024-08148-8>.
- [110] Bluvstein, D. *et al.* Logical quantum processor based on reconfigurable atom arrays. *Nature* **626**, 58–65 (2024). URL <https://doi.org/10.1038/s41586-023-06927-3>.
- [111] Barenco, A. *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467 (1995). URL <https://link.aps.org/doi/10.1103/PhysRevA.52.3457>.
- [112] Shende, V. V., Markov, I. L. & Bullock, S. S. Minimal universal two-qubit controlled-not-based circuits. *Phys. Rev. A* **69**, 062321 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevA.69.062321>.
- [113] Amy, M., Maslov, D., Mosca, M. & Roetteler, M. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **32**, 818–830 (2013). URL <https://www.doi.org/10.1109/TCAD.2013.2244643>.
- [114] Nam, Y., Ross, N. J., Su, Y., Childs, A. M. & Maslov, D. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information* **4**, 23 (2018). URL <https://doi.org/10.1038/s41534-018-0072-4>.
- [115] Cincio, L., Subaşı, Y., Sornborger, A. T. & Coles, P. J. Learning the quantum algorithm for state overlap. *New Journal of Physics* **20**, 113022 (2018). URL <https://dx.doi.org/10.1088/1367-2630/aae94a>.
- [116] Brown, A. R. & Susskind, L. Second law of quantum complexity. *Phys. Rev. D* **97**, 086015 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevD.97.086015>.

- [117] Yordanov, Y. S., Arvidsson-Shukur, D. R. M. & Barnes, C. H. W. Efficient quantum circuits for quantum computational chemistry. *Phys. Rev. A* **102**, 062612 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevA.102.062612>.
- [118] Tacchino, F., Chiesa, A., Carretta, S. & Gerace, D. Quantum computers as universal quantum simulators: State-of-the-art and perspectives. *Advanced Quantum Technologies* **3**, 1900052 (2020). URL <https://doi.org/10.1002/qute.201900052>.
- [119] van den Berg, E. & Temme, K. Circuit optimization of Hamiltonian simulation by simultaneous diagonalization of Pauli clusters. *Quantum* **4**, 322 (2020). URL <https://doi.org/10.22331/q-2020-09-12-322>.
- [120] Abramsky, S. & Coecke, B. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.*, 415–425 (2004). URL <https://www.doi.org/10.1109/LICS.2004.1319636>.
- [121] Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* **91**, 147902 (2003). URL <https://link.aps.org/doi/10.1103/PhysRevLett.91.147902>.
- [122] Duncan, R., Kissinger, A., Perdrix, S. & van de Wetering, J. Graph-theoretic simplification of quantum circuits with the ZX-calculus. *Quantum* **4**, 279 (2020). URL <https://doi.org/10.22331/q-2020-06-04-279>.
- [123] Li, G., Ding, Y. & Xie, Y. Tackling the qubit mapping problem for NISQ-era quantum devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19*, 1001–1014 (Association for Computing Machinery, New York, NY, USA, 2019). URL <https://doi.org/10.1145/3297858.3304023>.
- [124] Park, M., Leahey, E. & Funk, R. J. Papers and patents are becoming less disruptive over time. *Nature* **613**, 138–144 (2023). URL <https://doi.org/10.1038/s41586-022-05543-x>.
- [125] Koltun, V. & Hafner, D. The h-index is no longer an effective correlate of scientific reputation. *PLoS One* **16**, e0253397 (2021). URL <https://doi.org/10.1371/journal.pone.0253397>.
- [126] M. Q. Cruz, P. & Acín, A. Quantum advantage in the “hardest logic puzzle ever”. *[in preparation]* (2024).
- [127] M. Q. Cruz, P. & Pedroso, J. P. QCMotherLode: Autonomous discovery and optimization of logically-equivalent quantum circuits. *[in preparation]* (2025).
- [128] M. Q. Cruz, P. & Murta, B. Fredkin and Toffoli quantum gate decompositions in diverse connectivity topologies (2023). URL <https://zenodo.org/records/10047422>.
- [129] M. Q. Cruz, P. & Murta, B. Shallow unitary decompositions of quantum fredkin and toffoli gates for connectivity-aware equivalent circuit averaging. *APL Quantum* **1**, 016105 (2024). URL <https://doi.org/10.1063/5.0187026>.
- [130] Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Rev. Mod. Phys.* **74**, 145–195 (2002). URL <https://link.aps.org/doi/10.1103/RevModPhys.74.145>.
- [131] Rosenhouse, J. *Games for Your Mind: The History and Future of Logic Puzzles* (Princeton University Press, Princeton, NJ, 2020).

- [132] Carroll, L. *The Game of Logic* (Macmillan and Co., London, 1886).
- [133] Carroll, L. *Symbolic Logic* (Macmillan and Co., London, 1896).
- [134] A problem for the logical: Who owns the zebra? *Life International magazine* 95 (1962). Reprinted in various puzzle collections.
- [135] Prosser, P. Hybrid algorithms for the constraint satisfaction problem. *Computational intelligence* **9**, 268–299 (1993).
- [136] Salavati, S., Hajjarzadeh, S. & Mazloom, M. An optimized method for solving zebra puzzle. In *2009 Second International Conference on Computer and Electrical Engineering*, vol. 1, 448–451 (2009).
- [137] Goodman, N. *Problems and Projects* (The Bobbs-Merrill Company, Inc., Indianapolis, 1972).
- [138] Smullyan, R. M. *What Is the Name of This Book? - The riddle of Dracula and other logical puzzles* (Prentice-Hall, Englewood Cliffs, New Jersey, 1978).
- [139] Smullyan, R. M. *The Lady or the Tiger? and Other Logic Puzzles including a mathematical novel that features Gödel's great discovery* (Alfred A. Knopf, Inc., New York, 1982).
- [140] Smullyan, R. M. *To Mock a Mockingbird - And Other Logic Puzzles Including an Amazing Adventure in Combinatory Logic* (Alfred A. Knopf, Inc., New York, 1985).
- [141] Smullyan, R. M. *Forever Undecided: A Puzzle Guide to Gödel* (Alfred A. Knopf, Inc., New York, 1987).
- [142] Smullyan, R. M. Self-referential systems. In *Gödel's Incompleteness Theorems*, 116–135 (Oxford University Press, Oxford, UK, 1992).
- [143] Boolos, G. The hardest logic puzzle ever. *The Harvard Review of Philosophy* **6**, 62–65 (1996). URL <https://doi.org/10.5840/harvardreview1996615>.
- [144] Roberts, T. S. Some thoughts about the hardest logic puzzle ever. *Journal of Philosophical Logic* **30**, 609–612 (2001). URL <https://doi.org/10.1023/A:1013344220298>.
- [145] Rabern, B. & Rabern, L. A simple solution to the hardest logic puzzle ever. *Analysis* **68**, 105–112 (2008). URL <https://doi.org/10.1093/analys/68.2.105>.
- [146] Uzquiano, G. How to solve the hardest logic puzzle ever in two questions. *Analysis* **70**, 39–44 (2010).
- [147] Wheeler, G. & Barahona, P. Why the hardest logic puzzle ever cannot be solved in less than three questions. *Journal of Philosophical Logic* **41**, 493–503 (2012). URL <https://doi.org/10.1007/s10992-011-9181-7>.
- [148] Wintein, S. On the behavior of true and false. *Minds and Machines* **22**, 1–24 (2012). URL <https://doi.org/10.1007/s11023-011-9260-0>.
- [149] Liu, F. & Wang, Y. Reasoning about agent types and the hardest logic puzzle ever. *Minds and Machines* **23**, 123–161 (2013). URL <https://doi.org/10.1007/s11023-012-9287-x>.
- [150] Plaza, J. Logics of public communications. *Synthese* **158**, 165–179 (2007). URL <https://doi.org/10.1007/s11229-007-9168-7>.

- [151] Gerbrandy, J. & Groeneveld, W. Reasoning about information change. *Journal of Logic, Language and Information* **6**, 147–169 (1997). URL <https://doi.org/10.1023/A:1008222603071>.
- [152] Novozhilov, N. The hardest logic puzzle ever becomes even tougher. *arXiv:1206.1926 [math.LO]* (2012). URL <https://arxiv.org/abs/1206.1926>.
- [153] Webb, R. God, what a problem. *New Scientist* **2896**, 50–52 (2012). URL <https://www.newscientist.com/issue/2896/>.
- [154] Carnielli, W. *Making The ‘Hardest Logic Puzzle Ever’ a Bit Harder*, 181–190 (Springer International Publishing, Cham, 2017). URL [https://doi.org/10.1007/978-3-319-68732-2\\_11](https://doi.org/10.1007/978-3-319-68732-2_11).
- [155] Rivest, R. L., Shamir, A. & Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* **21**, 120–126 (1978). URL <https://doi.org/10.1145/359340.359342>.
- [156] Shannon, C. E. A mathematical theory of communication. *The Bell System Technical Journal* **27**, 379–423 (1948).
- [157] Landau, L. D. & Lifshitz, E. M. *Mechanics*, vol. 1 of *Course of Theoretical Physics* (Butterworth-Heinemann, Oxford, 1976), 3rd edn.
- [158] Goldstein, H., Poole, C. & Safko, J. *Classical Mechanics* (Addison-Wesley, San Francisco, 2002), 3rd edn.
- [159] Lorenz, E. N. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences* **20**, 130 – 141 (1963). URL [https://journals.ametsoc.org/view/journals/atsc/20/2/1520-0469\\_1963\\_020\\_0130\\_dnf\\_2\\_0\\_co\\_2.xml](https://journals.ametsoc.org/view/journals/atsc/20/2/1520-0469_1963_020_0130_dnf_2_0_co_2.xml).
- [160] Lorenz, E. N. The predictability of a flow which possesses many scales of motion. *Tellus* **21**, 289–307 (1969). URL <https://doi.org/10.3402/tellusa.v21i3.10086>.
- [161] Norton, J. D. Causation as folk science. *Philosopher’s Imprint* **3**, 1–22 (2003).
- [162] Wald, R. M. *General Relativity* (University of Chicago Press, Chicago, 1984).
- [163] Born, M. Zur quantenmechanik der stoßvorgänge. *Zeitschrift für Physik* **37**, 863–867 (1926). URL <https://doi.org/10.1007/BF01397477>.
- [164] Bohr, N. The quantum postulate and the recent development of atomic theory. *Nature* **121**, 580–590 (1928).
- [165] von Neumann, J. *Mathematische Grundlagen der Quantenmechanik* (Springer, Berlin, 1932). English translation: *Mathematical Foundations of Quantum Mechanics*, Princeton University Press, 1955.
- [166] Gleason, A. M. Measures on the closed subspaces of a hilbert space. *Journal of Mathematics and Mechanics* **6**, 885–893 (1957). URL <https://doi.org/10.1512/iumj.1957.6.56050>.
- [167] Bera, M. N., Acín, A., Kuś, M., Mitchell, M. W. & Lewenstein, M. Randomness in quantum mechanics: philosophy, physics and technology. *Reports on Progress in Physics* **80**, 124001 (2017). URL <https://dx.doi.org/10.1088/1361-6633/aa8731>.

- [168] Bennett, C. H. & Wiesner, S. J. Communication via one- and two-particle operators on einstein-podolsky-rosen states. *Phys. Rev. Lett.* **69**, 2881–2884 (1992). URL <https://link.aps.org/doi/10.1103/PhysRevLett.69.2881>.
- [169] Birkhoff, G. & Neumann, J. V. The logic of quantum mechanics. *Annals of Mathematics* **37**, 823–843 (1936). URL <http://www.jstor.org/stable/1968621>.
- [170] Bartlett, S. D., Rudolph, T. & Spekkens, R. W. Reference frames, superselection rules, and quantum information. *Rev. Mod. Phys.* **79**, 555–609 (2007). URL <https://link.aps.org/doi/10.1103/RevModPhys.79.555>.
- [171] Peres, A. & Scudo, P. F. Entangled quantum states as direction indicators. *Phys. Rev. Lett.* **86**, 4160–4162 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevLett.86.4160>.
- [172] Bagan, E., Baig, M., Brey, A., Muñoz Tapia, R. & Tarrach, R. Optimal encoding and decoding of a spin direction. *Phys. Rev. A* **63**, 052309 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevA.63.052309>.
- [173] Bagan, E., Baig, M. & Muñoz Tapia, R. Aligning reference frames with quantum states. *Phys. Rev. Lett.* **87**, 257903 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevLett.87.257903>.
- [174] Peres, A. & Scudo, P. F. Transmission of a cartesian frame by a quantum system. *Phys. Rev. Lett.* **87**, 167901 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevLett.87.167901>.
- [175] Acín, A., Jané, E. & Vidal, G. Optimal estimation of quantum dynamics. *Phys. Rev. A* **64**, 050302 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevA.64.050302>.
- [176] Bartlett, S. D., Rudolph, T. & Spekkens, R. W. Classical and quantum communication without a shared reference frame. *Phys. Rev. Lett.* **91**, 027901 (2003). URL <https://link.aps.org/doi/10.1103/PhysRevLett.91.027901>.
- [177] Bartlett, S. D., Rudolph, T. & Spekkens, R. W. Decoherence-full subsystems and the cryptographic power of a private shared reference frame. *Phys. Rev. A* **70**, 032307 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevA.70.032307>.
- [178] Bužek, V., Hillery, M. & Werner, R. F. Optimal manipulations with qubits: Universal-not gate. *Phys. Rev. A* **60**, R2626–R2629 (1999). URL <https://link.aps.org/doi/10.1103/PhysRevA.60.R2626>.
- [179] Barnett, S. *Quantum Information* (Oxford University Press, Great Clarendon St, Oxford OX2 6DP, United Kingdom, 2009).
- [180] Peres, A. *Quantum Theory: Concepts and Methods*, vol. 72 of *Fundamental Theories of Physics* (Kluwer Academic Publishers, 2002). URL <https://link.springer.com/book/10.1007/0-306-47120-5>.
- [181] Belavkin, V. P. Optimal multiple quantum statistical hypothesis testing. *Stochastics* **1**, 315–345 (1975). URL <https://doi.org/10.1080/17442507508833114>.

- [182] Hausladen, P. & Wootters, W. K. A ‘pretty good’ measurement for distinguishing quantum states. *Journal of Modern Optics* **41**, 2385–2390 (1994). URL <https://doi.org/10.1080/09500349414552221>.
- [183] Shende, V. V., Bullock, S. S. & Markov, I. L. Synthesis of quantum-logic circuits. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **25**, 1000–1010 (2006).
- [184] Iten, R., Colbeck, R., Kukuljan, I., Home, J. & Christandl, M. Quantum circuits for isometries. *Phys. Rev. A* **93**, 032318 (2016). URL <https://link.aps.org/doi/10.1103/PhysRevA.93.032318>.
- [185] Malveti, E., Iten, R. & Colbeck, R. Quantum Circuits for Sparse Isometries. *Quantum* **5**, 412 (2021). URL <https://doi.org/10.22331/q-2021-03-15-412>.
- [186] Yin, J. *et al.* Satellite-based entanglement distribution over 1200 kilometers. *Science* **356**, 1140–1144 (2017). URL <https://www.science.org/doi/abs/10.1126/science.aan3211>.
- [187] Liao, S.-K. *et al.* Satellite-to-ground quantum key distribution. *Nature* **549**, 43–47 (2017). URL <https://doi.org/10.1038/nature23655>.
- [188] Liao, S.-K. *et al.* Satellite-relayed intercontinental quantum network. *Phys. Rev. Lett.* **120**, 030501 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.030501>.
- [189] Sidhu, J. S. *et al.* Advances in space quantum communications **2**, 182–217 (2021). URL <https://doi.org/10.1049/qtc2.12015>.
- [190] Whittet, D. C. B. *Dust in the Galactic Environment (Third Edition)*. 2514-3433 (IOP Publishing, 2022). URL <https://dx.doi.org/10.1088/2514-3433/ac7204>.
- [191] Dehnen, H. Gravitational faraday-effect. *International Journal of Theoretical Physics* **7**, 467–474 (1973). URL <https://doi.org/10.1007/BF00713248>.
- [192] Bruschi, D. E., Ralph, T. C., Fuentes, I., Jennewein, T. & Razavi, M. Spacetime effects on satellite-based quantum communications. *Phys. Rev. D* **90**, 045041 (2014). URL <https://link.aps.org/doi/10.1103/PhysRevD.90.045041>.
- [193] Bruschi, D. E., Datta, A., Ursin, R., Ralph, T. C. & Fuentes, I. Quantum estimation of the schwarzschild spacetime parameters of the earth. *Phys. Rev. D* **90**, 124001 (2014). URL <https://link.aps.org/doi/10.1103/PhysRevD.90.124001>.
- [194] Jonsson, R. H., Aruquipa, D. Q., Casals, M., Kempf, A. & Martin-Martinez, E. Communication through quantum fields near a black hole. *Phys. Rev. D* **101**, 125005 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevD.101.125005>.
- [195] Exirifard, Q., Culf, E. & Karimi, E. Towards communication in a curved spacetime geometry. *Communications Physics* **4**, 171 (2021). URL <https://doi.org/10.1038/s42005-021-00671-8>.
- [196] Strauss, R. D. T. Voyager 2 enters interstellar space. *Nature Astronomy* **3**, 963–964 (2019). URL <https://doi.org/10.1038/s41550-019-0942-5>.

- [197] Community, S. E. How many photons are received per bit transmitted from Voyager 1? (2024). URL <https://physics.stackexchange.com/questions/816698/how-many-photons-are-received-per-bit-transmitted-from-voyager-1>. Accessed: 2024-11-19.
- [198] Kakarla, R., Schröder, J. & Andrekson, P. A. One photon-per-bit receiver using near-noiseless phase-sensitive amplification. *Light: Science & Applications* **9**, 153 (2020). URL <https://doi.org/10.1038/s41377-020-00389-2>.
- [199] Burlaga, L. F. *et al.* A large magnetic hump in the VLISM observed by Voyager 1 in 2020–2022. *The Astrophysical Journal* **953**, 135 (2023). URL <https://dx.doi.org/10.3847/1538-4357/acd6eb>.
- [200] Kurth, W. S. Voyager 1 electron densities in the very local interstellar medium to beyond 160 au. *The Astrophysical Journal Letters* **963**, L6 (2024). URL <https://dx.doi.org/10.3847/2041-8213/ad2617>.
- [201] Dialynas, K. *et al.* Anisotropies of 40–139 keV ions measured beyond the termination shock and in the very local interstellar medium. *The Astrophysical Journal* **974**, 174 (2024). URL <https://dx.doi.org/10.3847/1538-4357/ad7601>.
- [202] Schlosshauer, M. Quantum decoherence. *Physics Reports* **831**, 1–57 (2019). URL <https://doi.org/10.1016/j.physrep.2019.10.001>. Quantum decoherence.
- [203] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).
- [204] Kuo, W. & Zuo, M. J. *Optimal Reliability Modeling: Principles and Applications* (John Wiley & Sons, Hoboken, New Jersey, 2003).
- [205] Deutsch, D. & Jozsa, R. Rapid solution of problems by quantum computation. *Proc. R. Soc. Lon. A* **439**, 553–558 (1992).
- [206] Bennett, C. H. *et al.* Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.* **70**, 1895–1899 (1993). URL <https://link.aps.org/doi/10.1103/PhysRevLett.70.1895>.
- [207] Eisert, J., Wilkens, M. & Lewenstein, M. Quantum games and quantum strategies. *Phys. Rev. Lett.* **83**, 3077–3080 (1999). URL <https://link.aps.org/doi/10.1103/PhysRevLett.83.3077>.
- [208] Putnam, H. *Is Logic Empirical?*, 216–241 (Springer Netherlands, Dordrecht, 1969). URL [https://doi.org/10.1007/978-94-010-3381-7\\_5](https://doi.org/10.1007/978-94-010-3381-7_5).
- [209] Aaronson, S. *Quantum Computing Since Democritus* (Cambridge University Press, 2013).
- [210] Silver, D. *et al.* Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016). URL <https://doi.org/10.1038/nature16961>.
- [211] Schrittwieser, J. *et al.* Mastering atari, go, chess and shogi by planning with a learned model. *Nature* **588**, 604–609 (2020). URL <https://doi.org/10.1038/s41586-020-03051-4>.

- [212] Udrescu, S.-M. & Tegmark, M. Ai feynman: A physics-inspired method for symbolic regression. *Science Advances* **6**, eaay2631 (2020). URL <https://www.science.org/doi/abs/10.1126/sciadv.aay2631>.
- [213] Fawzi, A. *et al.* Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **610**, 47–53 (2022). URL <https://doi.org/10.1038/s41586-022-05172-4>.
- [214] Davies, A. *et al.* Advancing mathematics by guiding human intuition with AI. *Nature* **600**, 70–74 (2021). URL <https://doi.org/10.1038/s41586-021-04086-x>.
- [215] Trinh, T. H., Wu, Y., Le, Q. V., He, H. & Luong, T. Solving olympiad geometry without human demonstrations. *Nature* **625**, 476–482 (2024). URL <https://doi.org/10.1038/s41586-023-06747-5>.
- [216] Romera-Paredes, B. *et al.* Mathematical discoveries from program search with large language models. *Nature* **625**, 468–475 (2024). URL <https://doi.org/10.1038/s41586-023-06924-6>.
- [217] Shao, Z. *et al.* DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024). URL <https://arxiv.org/abs/2402.03300>.
- [218] Lu, P. *et al.* MathVista: Evaluating mathematical reasoning of foundation models in visual contexts. In *International Conference on Learning Representations (ICLR)* (2024). URL <https://arxiv.org/abs/2310.02255>.
- [219] Zhang, D., Huang, X., Zhou, D., Li, Y. & Ouyang, W. Accessing gpt-4 level mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *arXiv preprint arXiv:2406.07394* (2024). URL <https://arxiv.org/abs/2406.07394>.
- [220] Chervonyi, Y. *et al.* Gold-medalist performance in solving olympiad geometry with alphageometry2. *arXiv preprint arXiv:2502.03544* (2025). URL <https://arxiv.org/abs/2502.03544>.
- [221] Lu, C. *et al.* The AI Scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292* (2024). URL <https://arxiv.org/abs/2408.06292>.
- [222] Skrzypczyk, P. & Cavalcanti, D. *Semidefinite Programming in Quantum Information Science*. 2053-2563 (IOP Publishing, 2023). URL <https://dx.doi.org/10.1088/978-0-7503-3343-6>.
- [223] Diamond, S. & Boyd, S. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* **17**, 1–5 (2016).
- [224] Agrawal, A., Verschueren, R., Diamond, S. & Boyd, S. A rewriting system for convex optimization problems. *Journal of Control and Decision* **5**, 42–60 (2018).
- [225] Harris, C. R. *et al.* Array programming with numpy. *Nature* **585**, 357–362 (2020). URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [226] Feynman, R., Leighton, R. & Sands, M. *The Feynman Lectures on Physics, Vol. III* (Addison-Wesley, 1965).
- [227] De Broglie, L. *Recherches sur la théorie des quanta*. Ph.D. thesis, Migration-université en cours d'affectation (1924).

- [228] Heisenberg, W. Über den anschaulichen inhalt der quantentheoretischen kinematik und mechanik. *Zeitschrift für Physik* **43**, 172–198 (1927).
- [229] Bohr, N. *Albert Einstein: Philosopher-Scientist*, vol. VII, chap. 7 - Discussion with Einstein on Epistemological Problems in Atomic Physics, 201–241 (Library of Living Philosophers, Evanston, Illinois, 1949).
- [230] Heisenberg, W. *The Physical Principles of the Quantum Theory* (University of Chicago Press, 1930).
- [231] Wheeler, J. A. The “past” and the “delayed-choice” double-slit experiment. In Marlow, A. (ed.) *Mathematical Foundations of Quantum Theory*, 9 – 48 (Academic Press, 1978).
- [232] Wootters, W. K. & Zurek, W. H. Complementarity in the double-slit experiment: Quantum nonseparability and a quantitative statement of bohr’s principle. *Physical Review D* **19**, 473 (1979).
- [233] Scully, M. O., Englert, B.-G. & Schwinger, J. Spin coherence and humpty-dumpty. iii. the effects of observation. *Phys. Rev. A* **40**, 1775–1784 (1989).
- [234] Scully, M. O. & Walther, H. Quantum optical test of observation and complementarity in quantum mechanics. *Phys. Rev. A* **39**, 5229–5236 (1989).
- [235] Deutsch, D. Uncertainty in quantum measurements. *Phys. Rev. Lett.* **50**, 631–633 (1983).
- [236] Mittelstaedt, P., Prieur, A. & Schieder, R. Unsharp particle-wave duality in a photon split-beam experiment. *Foundations of Physics* **17**, 891–903 (1987).
- [237] Greenberger, D. M. & Yasin, A. Simultaneous wave and particle knowledge in a neutron interferometer. *Physics Letters A* **128**, 391–394 (1988).
- [238] Jaeger, G., Shimony, A. & Vaidman, L. Two interferometric complementarities. *Physical Review A* **51**, 54 (1995).
- [239] Englert, B.-G. Fringe visibility and which-way information: An inequality. *Phys. Rev. Lett.* **77**, 2154 (1996).
- [240] Björk, G. & Karlsson, A. Complementarity and quantum erasure in welcher weg experiments. *Phys. Rev. A* **58**, 3477–3483 (1998).
- [241] Dürr, S. Quantitative wave-particle duality in multibeam interferometers. *Phys. Rev. A* **64**, 042113 (2001).
- [242] Jakob, M. & Bergou, J. A. Complementarity and entanglement in bipartite qudit systems. *Phys. Rev. A* **76**, 052107 (2007).
- [243] Jakob, M. & Bergou, J. A. Quantitative complementarity relations in bipartite systems: Entanglement as a physical reality. *Optics Communications* **283**, 827 – 830 (2010).
- [244] Coles, P. J., Kaniewski, J. & Wehner, S. Equivalence of wave–particle duality to entropic uncertainty. *Nature communications* **5**, 5814 (2014).
- [245] Coles, P. J. Entropic framework for wave-particle duality in multipath interferometers. *Phys. Rev. A* **93**, 062111 (2016).
- [246] Angelo, R. M. & Ribeiro, A. D. Wave–particle duality: An information-based approach. *Foundations of Physics* **45**, 1407–1420 (2015).

- [247] Bagan, E., Bergou, J. A., Cottrell, S. S. & Hillery, M. Relations between coherence and path information. *Phys. Rev. Lett.* **116**, 160406 (2016).
- [248] Dürr, S., Nonn, T. & Rempe, G. Origin of quantum-mechanical complementarity probed by a ‘which-way’ experiment in an atom interferometer. *Nature* **395**, 33–37 (1998).
- [249] Bertet, P. *et al.* A complementarity experiment with an interferometer at the quantum–classical boundary. *Nature* **411**, 166 (2001).
- [250] Herzog, T. J., Kwiat, P. G., Weinfurter, H. & Zeilinger, A. Complementarity and the quantum eraser. *Phys. Rev. Lett.* **75**, 3034–3037 (1995).
- [251] Kim, Y.-H., Yu, R., Kulik, S. P., Shih, Y. & Scully, M. O. Delayed “choice” quantum eraser. *Phys. Rev. Lett.* **84**, 1–5 (2000).
- [252] Walborn, S. P., Terra Cunha, M. O., Pádua, S. & Monken, C. H. Double-slit quantum eraser. *Phys. Rev. A* **65**, 033818 (2002).
- [253] Neves, L. *et al.* Control of quantum interference in the quantum eraser. *New Journal of Physics* **11**, 073035 (2009).
- [254] Peruzzo, A., Shadbolt, P., Brunner, N., Popescu, S. & O’Brien, J. L. A quantum delayed-choice experiment. *Science* **338**, 634–637 (2012).
- [255] Kaiser, F., Coudreau, T., Milman, P., Ostrowsky, D. B. & Tanzilli, S. Entanglement-enabled delayed-choice experiment. *Science* **338**, 637–640 (2012).
- [256] Buks, E., Schuster, R., Heiblum, M., Mahalu, D. & Umansky, V. Dephasing in electron interference by a ‘which-path’ detector. *Nature* **391**, 871 (1998).
- [257] Sprinzak, D., Buks, E., Heiblum, M. & Shtrikman, H. Controlled dephasing of electrons via a phase sensitive detector. *Phys. Rev. Lett.* **84**, 5820–5823 (2000).
- [258] Neder, I., Heiblum, M., Mahalu, D. & Umansky, V. Entanglement, dephasing, and phase recovery via cross-correlation measurements of electrons. *Phys. Rev. Lett.* **98**, 036803 (2007).
- [259] Teklemariam, G., Fortunato, E. M., Pravia, M. A., Havel, T. F. & Cory, D. G. Nmr analog of the quantum disentanglement eraser. *Phys. Rev. Lett.* **86**, 5845–5849 (2001).
- [260] Teklemariam, G. *et al.* Quantum erasers and probing classifications of entanglement via nuclear magnetic resonance. *Phys. Rev. A* **66**, 012309 (2002).
- [261] Peng, X. *et al.* An interferometric complementarity experiment in a bulk nuclear magnetic resonance ensemble. *Journal of Physics A: Mathematical and General* **36**, 2555–2563 (2003).
- [262] Zheng, S.-B. *et al.* Quantum delayed-choice experiment with a beam splitter in a quantum superposition. *Phys. Rev. Lett.* **115**, 260403 (2015).
- [263] Liu, K. *et al.* A twofold quantum delayed-choice experiment in a superconducting circuit. *Science Advances* **3**, e1603159 (2017).
- [264] Bienfait, A. *et al.* Quantum erasure using entangled surface acoustic phonons. *Phys. Rev. X* **10**, 021055 (2020).

- [265] Kwiat, P. G., Steinberg, A. M. & Chiao, R. Y. Observation of a “quantum eraser”: A revival of coherence in a two-photon interference experiment. *Phys. Rev. A* **45**, 7729–7739 (1992).
- [266] Schwindt, P. D. D., Kwiat, P. G. & Englert, B.-G. Quantitative wave-particle duality and nonerasing quantum erasure. *Phys. Rev. A* **60**, 4285–4290 (1999).
- [267] Amico, M. & Dittel, C. Simulation of wave-particle duality in multipath interferometers on a quantum computer. *Phys. Rev. A* **102**, 032605 (2020).
- [268] Schwaller, N., Dupertuis, M.-A. & Javerzac-Galy, C. Evidence of the entanglement constraint on wave-particle duality using the IBM Q quantum computer. *Phys. Rev. A* **103**, 022409 (2021).
- [269] Pozzobom, M. B., Basso, M. L. W. & Maziero, J. Experimental tests of the density matrix’s property-based complementarity relations. *Phys. Rev. A* **103**, 022212 (2021).
- [270] García-Martín, D. & Sierra, G. Five experimental tests on the 5-qubit IBM quantum computer. *Journal of Applied Mathematics and Physics* **6**, 1460–1475 (2018).
- [271] Helstrom, C. W. *Quantum detection and estimation theory* (Academic press, 1976).
- [272] Scully, M. O., Englert, B.-G. & Walther, H. Quantum optical tests of complementarity. *Nature* **351**, 111 (1991).
- [273] Ionicioiu, R. & Terno, D. R. Proposal for a quantum delayed-choice experiment. *Phys. Rev. Lett.* **107**, 230406 (2011).
- [274] IBM Quantum. <https://quantum-computing.ibm.com> (Access: 2025-01-01).
- [275] Qiskit contributors. Qiskit: An open-source framework for quantum computing (2023).
- [276] Blais, A. *et al.* Quantum-information processing with circuit quantum electrodynamics. *Physical Review A* **75**, 032329 (2007).
- [277] Koch, J. *et al.* Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A* **76**, 042319 (2007).
- [278] Majer, J. *et al.* Coupling superconducting qubits via a cavity bus. *Nature* **449**, 443–447 (2007).
- [279] Krantz, P. *et al.* A quantum engineer’s guide to superconducting qubits. *Applied Physics Reviews* **6**, 021318 (2019).
- [280] Blais, A., Grimsmo, A. L., Girvin, S. M. & Wallraff, A. Circuit quantum electrodynamics. *Rev. Mod. Phys.* **93**, 025005 (2021).
- [281] Chen, Y., Farahzad, M., Yoo, S. & Wei, T.-C. Detector tomography on IBM quantum computers and mitigation of an imperfect measurement. *Phys. Rev. A* **100**, 052315 (2019).
- [282] Maciejewski, F. B., Zimborás, Z. & Oszmaniec, M. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum* **4**, 257 (2020).
- [283] Asfaw, A. *et al.* Learn quantum computation using qiskit. <http://community.qiskit.org/textbook> (2020).
- [284] Wald, A. & Wolfowitz, J. On a test whether two samples are from the same population. *The Annals of Mathematical Statistics* **11**, 147–162 (1940).

- [285] Sundaresan, N. *et al.* Reducing unitary and spectator errors in cross resonance with optimized rotary echoes. *PRX Quantum* **1**, 020318 (2020).
- [286] Sheldon, S., Magesan, E., Chow, J. M. & Gambetta, J. M. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Phys. Rev. A* **93**, 060302 (2016). URL <https://link.aps.org/doi/10.1103/PhysRevA.93.060302>.
- [287] Malekakhlagh, M., Magesan, E. & McKay, D. C. First-principles analysis of cross-resonance gate operation. *Phys. Rev. A* **102**, 042605 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevA.102.042605>.
- [288] Chiou, D.-W. & Hsu, H.-C. Complementarity relations of a delayed-choice quantum eraser in a quantum circuit. *Quantum Information Processing* **23**, 18 (2024). URL <https://doi.org/10.1007/s11128-023-04214-8>.
- [289] Chen, B.-H., Chiou, D.-W. & Hsu, H.-C. Demonstration of scully-drühl-type quantum erasers on quantum computers. *arXiv preprint arXiv:2409.08053* (2024).
- [290] Bremner, M. J. *et al.* Practical scheme for quantum computation with any two-qubit entangling gate. *Phys. Rev. Lett.* **89**, 247902 (2002).
- [291] Zhang, J., Vala, J., Sastry, S. & Whaley, K. B. Exact two-qubit universal quantum circuit. *Phys. Rev. Lett.* **91**, 027903 (2003).
- [292] Dawson, C. M. & Nielsen, M. A. The solovay-kitaev algorithm. *Quantum Info. Comput.* **6**, 81–95 (2006).
- [293] Einstein, A. B. On the method of theoretical physics. *Philosophy of Science* **1**, 163–169 (1934). URL <https://doi.org/10.1086/286316>.
- [294] Dawson, C. M. & Nielsen, M. A. The solovay-kitaev algorithm. *Quantum Info. Comput.* **6**, 81–95 (2006). URL <https://dl.acm.org/doi/10.5555/2011679.2011685>.
- [295] Vidal, G. & Dawson, C. M. Universal quantum circuit for two-qubit transformations with three controlled-not gates. *Phys. Rev. A* **69**, 010301 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevA.69.010301>.
- [296] Reducing quantum computations to elementary unitary operations. *Computing in Science & Engineering* **3**, 27–32 (2001).
- [297] Möttönen, M. & Vartiainen, J. J. Decompositions of general quantum gates. In Shannon, S. (ed.) *Trends in Quantum Computing Research*, chap. 7, 149–160 (Nova Science Publishers, Inc., New York, 2006).
- [298] Rakyta, P. & Zimborás, Z. Approaching the theoretical limit in quantum gate decomposition. *Quantum* **6**, 710 (2022). URL <https://doi.org/10.22331/q-2022-05-11-710>.
- [299] Vartiainen, J. J., Möttönen, M. & Salomaa, M. M. Efficient decomposition of quantum gates. *Phys. Rev. Lett.* **92**, 177902 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevLett.92.177902>.
- [300] Krol, A. M., Sarkar, A., Ashraf, I., Al-Ars, Z. & Bertels, K. Efficient decomposition of unitary matrices in quantum circuit compilers. *Applied Sciences* **12** (2022). URL <https://www.doi.org/10.3390/app12020759>.

- [301] Tucci, R. R. A rudimentary quantum compiler(2cnd ed.). *arXiv preprint arXiv:9902062* (1999). URL <https://arxiv.org/abs/quant-ph/9902062>.
- [302] Patel, K. N., Markov, I. L. & Hayes, J. P. Optimal synthesis of linear reversible circuits. *Quantum Info. Comput.* **8**, 282–294 (2008). URL <https://dl.acm.org/doi/10.5555/2011763.2011767>.
- [303] Aaronson, S. & Gottesman, D. Improved simulation of stabilizer circuits. *Phys. Rev. A* **70**, 052328 (2004). URL <https://link.aps.org/doi/10.1103/PhysRevA.70.052328>.
- [304] Schneider, S., Burgholzer, L. & Wille, R. A sat encoding for optimal clifford circuit synthesis. In *Proceedings of the 28th Asia and South Pacific Design Automation Conference, ASPDAC '23*, 190–195 (Association for Computing Machinery, New York, NY, USA, 2023). URL <https://doi.org/10.1145/3566097.3567929>.
- [305] Meuli, G., Soeken, M. & De Micheli, G. SAT-based {cnot, t} quantum circuit synthesis. In Kari, J. & Ulidowski, I. (eds.) *Reversible Computation*, 175–188 (Springer International Publishing, Cham, 2018). URL [https://link.springer.com/chapter/10.1007/978-3-319-99498-7\\_12](https://link.springer.com/chapter/10.1007/978-3-319-99498-7_12).
- [306] Maslov, D., Young, C., Miller, D. & Dueck, G. Quantum circuit simplification using templates. In *Design, Automation and Test in Europe*, 1208–1213 Vol. 2 (2005). URL <https://doi.org/10.1109/DATE.2005.249>.
- [307] Prasad, A. K., Shende, V. V., Markov, I. L., Hayes, J. P. & Patel, K. N. Data structures and algorithms for simplifying reversible circuits. *J. Emerg. Technol. Comput. Syst.* **2**, 277–293 (2006). URL <https://doi.org/10.1145/1216396.1216399>.
- [308] Bravyi, S., Shaydulin, R., Hu, S. & Maslov, D. Clifford Circuit Optimization with Templates and Symbolic Pauli Gates. *Quantum* **5**, 580 (2021). URL <https://doi.org/10.22331/q-2021-11-16-580>.
- [309] Amy, M., Maslov, D. & Mosca, M. Polynomial-time T-depth optimization of clifford+T circuits via matroid partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **33**, 1476–1489 (2014). URL <https://doi.org/10.1109/TCAD.2014.2341953>.
- [310] Amy, M., Azimzadeh, P. & Mosca, M. On the controlled-not complexity of controlled-not–phase circuits. *Quantum Science and Technology* **4**, 015002 (2018). URL <https://dx.doi.org/10.1088/2058-9565/aad8ca>.
- [311] Coecke, B. & Kissinger, A. *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning* (Cambridge University Press, 2017). URL <https://doi.org/10.1017/9781316219317>.
- [312] van de Wetering, J. ZX-calculus for the working quantum computer scientist. *arXiv preprint arXiv:2012.13966* (2020). URL <https://arxiv.org/abs/2012.13966>.
- [313] Kissinger, A. & van de Wetering, J. Reducing the number of non-clifford gates in quantum circuits. *Phys. Rev. A* **102**, 022406 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevA.102.022406>.
- [314] Sivarajah, S. *et al.* t|ket>: a retargetable compiler for nisq devices. *Quantum Science and Technology* **6**, 014003 (2020). URL <https://dx.doi.org/10.1088/2058-9565/ab8e92>.

- [315] Backens, M., Miller-Bakewell, H., de Felice, G., Lobski, L. & van de Wetering, J. There and back again: A circuit extraction tale. *Quantum* **5**, 421 (2021). URL <https://doi.org/10.22331/q-2021-03-25-421>.
- [316] Staudacher, K., Guggemos, T., Grundner-Culemann, S. & Gehrke, W. Reducing 2-qubit gate count for ZX-calculus based quantum circuit optimization. In *Quantum Physics and Logic (QPL) 2022*, vol. 394 of *Electronic Proceedings in Theoretical Computer Science (EPTCS)*, 29–45 (2023). URL <https://doi.org/10.4204/EPTCS.394.3>.
- [317] Kissinger, A. & de Griend, A. M.-v. CNOT circuit extraction for topologically-constrained quantum memories. *Quantum Information and Computation* **20**, 581–596 (2020). URL <https://doi.org/10.26421/QIC20.7-8-4>.
- [318] Meijer-van de Griend, A. & Li, S. M. Dynamic qubit routing with cnot circuit synthesis for quantum compilation. In *Quantum Physics and Logic (QPL) 2022*, vol. 394 of *Electronic Proceedings in Theoretical Computer Science (EPTCS)*, 363–399 (2023). URL <https://doi.org/10.4204/EPTCS.394.18>.
- [319] Holker, C. Causal flow preserving optimisation of quantum circuits in the ZX-calculus. *arXiv preprint 2312.02793* (2023). URL <https://arxiv.org/abs/2312.02793>.
- [320] Fösel, T., Niu, M. Y., Marquardt, F. & Li, L. Quantum circuit optimization with deep reinforcement learning. *arXiv preprint arXiv:2103.07585* (2021). URL <https://arxiv.org/abs/2103.07585>.
- [321] Hietala, K., Rand, R., Hung, S.-H., Wu, X. & Hicks, M. A verified optimizer for quantum circuits. *Proc. ACM Program. Lang.* **5** (2021). URL <https://doi.org/10.1145/3434318>.
- [322] Xu, M. *et al.* Quartz: superoptimization of quantum circuits. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2022*, 625–640 (Association for Computing Machinery, New York, NY, USA, 2022). URL <https://doi.org/10.1145/3519939.3523433>.
- [323] Xu, A., Molavi, A., Pick, L., Tannu, S. & Albarghouthi, A. Synthesizing quantum-circuit optimizers. *Proc. ACM Program. Lang.* **7** (2023). URL <https://doi.org/10.1145/3591254>.
- [324] Li, Z. *et al.* Quarl: A learning-based quantum circuit optimizer. *Proc. ACM Program. Lang.* **8** (2024). URL <https://doi.org/10.1145/3649831>.
- [325] Riu, J., Nogué, J., Vilaplana, G., Garcia-Saez, A. & Estarellas, M. P. Reinforcement learning based quantum circuit optimization via zx-calculus. *arXiv preprint arXiv:2312.11597* (2024).
- [326] Ruiz, F. J. R. *et al.* Quantum circuit optimization with alphasensor. *arXiv preprint arXiv:2402.14396* (2024). URL <https://arxiv.org/abs/2402.14396>.
- [327] Kremer, D. *et al.* Practical and efficient quantum circuit synthesis and transpiling with reinforcement learning. *arXiv preprint arXiv:2405.13196* (2024). URL <https://arxiv.org/abs/2405.13196>.
- [328] Fösel, T., Tighineanu, P., Weiss, T. & Marquardt, F. Reinforcement learning with neural networks for quantum feedback. *Phys. Rev. X* **8**, 031084 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevX.8.031084>.

- [329] Olle, J., Zen, R., Puviani, M. & Marquardt, F. Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent. *npj Quantum Information* **10**, 126 (2024). URL <https://doi.org/10.1038/s41534-024-00920-y>.
- [330] Ge, Y. *et al.* Quantum circuit synthesis and compilation optimization: Overview and prospects. *arXiv preprint arXiv:2407.00736* (2024). URL <https://arxiv.org/abs/2407.00736>.
- [331] O’Gorman, J. & Campbell, E. T. Quantum computation with realistic magic-state factories. *Phys. Rev. A* **95**, 032338 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevA.95.032338>.
- [332] Heyfron, L. E. & Campbell, E. T. An efficient quantum compiler that reduces t count. *Quantum Science and Technology* **4**, 015004 (2018). URL <https://dx.doi.org/10.1088/2058-9565/aad604>.
- [333] Litinski, D. Magic State Distillation: Not as Costly as You Think. *Quantum* **3**, 205 (2019). URL <https://doi.org/10.22331/q-2019-12-02-205>.
- [334] Gidney, C., Shu, N. & Jones, C. Magic state cultivation: growing t states as cheap as cnot gates. *arXiv preprint arXiv:2409.17595* (2024). URL <https://arxiv.org/abs/2409.17595>.
- [335] Coecke, B. & Duncan, R. Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics* **13**, 043016 (2011). URL <https://dx.doi.org/10.1088/1367-2630/13/4/043016>.
- [336] de Beaudrap, N., Kissinger, A. & van de Wetering, J. Circuit Extraction for ZX-Diagrams Can Be #P-Hard. In Bojańczyk, M., Merelli, E. & Woodruff, D. P. (eds.) *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, vol. 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 119:1–119:19 (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022). URL <https://www.doi.org/10.4230/LIPIcs.ICALP.2022.119>.
- [337] Krueger, R. *Vanishing 2-Qubit Gates with Non-Simplification ZX-Rules*. Msc dissertation, University of Oxford (2022). URL <https://doi.org/10.48550/arXiv.2209.06874>.
- [338] Staudacher, K. *Optimization Approaches for Quantum Circuits using ZX-calculus*. Master’s thesis, Ludwig-Maximilians-Universität München (2021).
- [339] Gogioso, S. & Yeung, R. Annealing optimisation of mixed zx phase circuits. In *Quantum Physics and Logic (QPL) 2022*, vol. 394 of *Electronic Proceedings in Theoretical Computer Science (EPTCS)*, 415–431 (2023). URL <https://doi.org/10.4204/EPTCS.394.20>.
- [340] Russell, S. & Norvig, P. *Artificial Intelligence: A Modern Approach* (Pearson Education, Inc., Hoboken, NJ, 2021), 4th edn.
- [341] Pearl, J. *Heuristics: Intelligent Search Strategies for Computer Problem Solving* (Addison-Wesley, 1984), first edn.
- [342] Qiskit: an open-source SDK for working with quantum computers at the level of extended quantum circuits, operators, and primitives. v2.0.0. GitHub repository. URL <https://github.com/Qiskit/qiskit>.

- [343] Kissinger, A. & van de Wetering, J. PyZX: Large scale automated diagrammatic reasoning. In Coecke, B. & Leifer, M. (eds.) *Proceedings 16th International Conference on Quantum Physics and Logic*, Chapman University, Orange, CA, USA., 10-14 June 2019, vol. 318 of *Electronic Proceedings in Theoretical Computer Science*, 229–241 (Open Publishing Association, 2020).
- [344] Cross, A. *et al.* OpenQASM 3: A Broader and Deeper Quantum Assembly Language. *ACM Transactions on Quantum Computing* **3** (2022). URL <https://doi.org/10.1145/3505636>.
- [345] Younis, E., Iancu, C. C., Lavrijsen, W., Davis, M. & Smith, E. Berkeley quantum synthesis toolkit (bqskit) v1. [Computer Software] <https://doi.org/10.11578/dc.20210603.2> (2021). URL <https://doi.org/10.11578/dc.20210603.2>.
- [346] tket - source code for the TKET quantum compiler, python bindings and utilities. v2.0.0. GitHub repository (2025). URL <https://github.com/CQCL/tket>.
- [347] PyZX Heuristixz. GitHub repository (2021). URL <https://github.com/mnm-team/pyzx-heuristics>.
- [348] PyZX - Python library for quantum circuit rewriting and optimisation using the ZX-calculus. v0.9.0 2025-01-30. GitHub repository (2025). URL <https://github.com/zxcalc/pyzx>.
- [349] Holker, C. PyZX Fork. GitHub repository (2025). URL <https://github.com/calumholker/pyzx>. Accessed: 2025-01-30.
- [350] Zou, H., Treinish, M., Hartman, K., Ivrii, A. & Lishman, J. LightSABRE: A lightweight and enhanced SABRE algorithm. *arXiv preprint arXiv:2409.08368* (2024). URL <https://arxiv.org/abs/2409.08368>.
- [351] optimizer. 2017-06-01. GitHub repository (2025). URL <https://github.com/njross/optimizer>.
- [352] Fredkin, E. & Toffoli, T. Conservative logic. *Int. J. Theor. Phys.* **21**, 219–253 (1982). URL <https://doi.org/10.1007/BF01857727>.
- [353] Shi, Y. Both toffoli and controlled-not need little help to do universal quantum computing **3**, 84–92 (2003).
- [354] Aharonov, D. A simple proof that Toffoli and Hadamard are quantum universal. *arXiv preprint quant-ph/0301040* (2003). URL <https://arxiv.org/abs/quant-ph/0301040>.
- [355] Barenco, A. *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467 (1995). URL <https://link.aps.org/doi/10.1103/PhysRevA.52.3457>.
- [356] Shende, V. V. & Markov, I. L. On the CNOT-Cost of Toffoli Gates. *Quantum Info. Comput.* **9**, 461–486 (2009).
- [357] Vedral, V., Barenco, A. & Ekert, A. Quantum networks for elementary arithmetic operations. *Phys. Rev. A* **54**, 147–153 (1996). URL <https://link.aps.org/doi/10.1103/PhysRevA.54.147>.
- [358] Gilliam, A., Pistoia, M. & Gonciulea, C. Canonical Construction of Quantum Oracles. *arXiv preprint arXiv:2006.10656* (2020). URL <https://arxiv.org/abs/2006.10656>.

- [359] Reed, M. D. *et al.* Realization of three-qubit quantum error correction with superconducting circuits. *Nature* **482**, 382–385 (2012). URL <https://doi.org/10.1038/nature10786>.
- [360] Gilyén, A. *Quantum Singular Value Transformation & Its Algorithmic Applications*. Ph.D. thesis, Universiteit van Amsterdam (2019). URL <https://eprints.illc.uva.nl/id/eprint/2165/>.
- [361] Kim, Y. *et al.* High-fidelity three-qubit iToffoli gate for fixed-frequency superconducting qubits. *Nat. Phys.* **18**, 783–788 (2022). URL <https://doi.org/10.1038/s41567-022-01590-3>.
- [362] Sun, S.-N. *et al.* Quantum Computation of Frequency-Domain Molecular Response Properties Using a Three-Qubit iToffoli Gate. *arXiv preprint arXiv:2302.04271v1* (2023). URL <https://arxiv.org/abs/2302.04271v1>.
- [363] Barenco, A. *et al.* Stabilization of quantum computations by symmetrization. *SIAM J. Comput.* **26**, 1541–1557 (1997). URL <https://doi.org/10.1137/S0097539796302452>.
- [364] Buhrman, H., Cleve, R., Watrous, J. & de Wolf, R. Quantum fingerprinting. *Phys. Rev. Lett.* **87**, 167902 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevLett.87.167902>.
- [365] Ozaydin, F. *et al.* Fusing multiple  $W$  states simultaneously with a Fredkin gate. *Phys. Rev. A* **89**, 042311 (2014). URL <https://link.aps.org/doi/10.1103/PhysRevA.89.042311>.
- [366] Araujo, I. F., Park, D. K., Petruccione, F. & da Silva, A. J. A divide-and-conquer algorithm for quantum state preparation. *Scientific Reports* **11**, 6329 (2021). URL <https://doi.org/10.1038/s41598-021-85474-1>.
- [367] Ekert, A. K. *et al.* Direct estimations of linear and nonlinear functionals of a quantum state. *Phys. Rev. Lett.* **88**, 217901 (2002). URL <https://link.aps.org/doi/10.1103/PhysRevLett.88.217901>.
- [368] Chiribella, G., D’Ariano, G. M., Perinotti, P. & Valiron, B. Quantum computations without definite causal structure. *Phys. Rev. A* **88**, 022318 (2013). URL <https://link.aps.org/doi/10.1103/PhysRevA.88.022318>.
- [369] Araújo, M., Guérin, P. A. & Baumeler, A. Quantum computation with indefinite causal structures. *Phys. Rev. A* **96**, 052315 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevA.96.052315>.
- [370] Castro-Ruiz, E., Giacomini, F. & Brukner, C. Dynamics of quantum causal structures. *Phys. Rev. X* **8**, 011047 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevX.8.011047>.
- [371] Hofmann, H. F. How weak values emerge in joint measurements on cloned quantum systems. *Phys. Rev. Lett.* **109**, 020408 (2012). URL <https://link.aps.org/doi/10.1103/PhysRevLett.109.020408>.
- [372] Babbush, R. *et al.* Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity. *Phys. Rev. X* **8**, 041015 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevX.8.041015>.
- [373] Oszmaniec, M., Brod, D. J. & Galvão, E. F. Measuring relational information between quantum states, and applications. *New Journal of Physics* **26**, 013053 (2024). URL <https://dx.doi.org/10.1088/1367-2630/ad1a27>.

- [374] Wagner, R. *et al.* Quantum circuits for measuring weak values, kirkwood–dirac quasiprobability distributions, and state spectra. *Quantum Science and Technology* **9**, 015030 (2024). URL <https://dx.doi.org/10.1088/2058-9565/ad124c>.
- [375] Tacchino, F., Macchiavello, C., Gerace, D. & Bajoni, D. An artificial neuron implemented on an actual quantum processor. *npj Quantum Inf.* **5**, 26 (2019). URL <https://doi.org/10.1038/s41534-019-0140-4>.
- [376] Satoh, T., Ohkura, Y. & Van Meter, R. Subdivided phase oracle for nisq search algorithms. *IEEE Transactions on Quantum Engineering* **1**, 1–15 (2020).
- [377] Duckering, C., Baker, J. M., Litteken, A. & Chong, F. T. Orchestrated Trios: Compiling for Efficient Communication in Quantum Programs with 3-Qubit Gates. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '21*, 375–385 (Association for Computing Machinery, New York, NY, USA, 2021). URL <https://doi.org/10.1145/3445814.3446718>.
- [378] Ivanov, S. S., Ivanov, P. A. & Vitanov, N. V. Efficient construction of three- and four-qubit quantum gates by global entangling gates. *Phys. Rev. A* **91**, 032311 (2015). URL <https://link.aps.org/doi/10.1103/PhysRevA.91.032311>.
- [379] Rasmussen, S. E., Groenland, K., Gerritsma, R., Schoutens, K. & Zinner, N. T. Single-step implementation of high-fidelity  $n$ -bit Toffoli gates. *Phys. Rev. A* **101**, 022308 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevA.101.022308>.
- [380] Fedorov, A., Steffen, L., Baur, M., da Silva, M. P. & Wallraff, A. Implementation of a Toffoli gate with superconducting circuits. *Nature* **481**, 170–172 (2012). URL <https://doi.org/10.1038/nature10713>.
- [381] Zahedinejad, E., Ghosh, J. & Sanders, B. C. High-fidelity single-shot toffoli gate via quantum control. *Phys. Rev. Lett.* **114**, 200502 (2015). URL <https://link.aps.org/doi/10.1103/PhysRevLett.114.200502>.
- [382] Bowman, M. A., Gokhale, P., Larson, J., Liu, J. & Suchara, M. Hardware-conscious optimization of the quantum toffoli gate. *ACM Transactions on Quantum Computing* **4** (2023). URL <https://doi.org/10.1145/3609229>.
- [383] Milburn, G. J. Quantum optical Fredkin gate. *Phys. Rev. Lett.* **62**, 2124–2127 (1989). URL <https://link.aps.org/doi/10.1103/PhysRevLett.62.2124>.
- [384] Fiurášek, J. Linear-optics quantum Toffoli and Fredkin gates. *Phys. Rev. A* **73**, 062313 (2006). URL <https://link.aps.org/doi/10.1103/PhysRevA.73.062313>.
- [385] Gong, Y.-X., Guo, G.-C. & Ralph, T. C. Methods for a linear optical quantum Fredkin gate. *Phys. Rev. A* **78**, 012305 (2008). URL <https://link.aps.org/doi/10.1103/PhysRevA.78.012305>.
- [386] Patel, R. B., Ho, J., Ferreyrol, F., Ralph, T. C. & Pryde, G. J. A quantum Fredkin gate. *Sci. Adv.* **2**, e1501531 (2016). URL <https://www.science.org/doi/abs/10.1126/sciadv.1501531>.
- [387] Li, Y. *et al.* Quantum Fredkin and Toffoli gates on a versatile programmable silicon photonic chip. *npj Quantum Inf.* **8**, 112 (2022). URL <https://doi.org/10.1038/s41534-022-00627-y>.

- [388] Jones, C. Low-overhead constructions for the fault-tolerant Toffoli gate. *Phys. Rev. A* **87**, 022328 (2013). URL <https://link.aps.org/doi/10.1103/PhysRevA.87.022328>.
- [389] Selinger, P. Quantum circuits of  $T$ -depth one. *Phys. Rev. A* **87**, 042302 (2013). URL <https://link.aps.org/doi/10.1103/PhysRevA.87.042302>.
- [390] Preskill, J. Fault-tolerant quantum computation. In *Introduction to quantum computation and information*, 213–269 (World Scientific, 1998).
- [391] Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018). URL <https://doi.org/10.22331/q-2018-08-06-79>.
- [392] Maslov, D. Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization. *Phys. Rev. A* **93**, 022311 (2016). URL <https://link.aps.org/doi/10.1103/PhysRevA.93.022311>.
- [393] Ralph, T. C., Resch, K. J. & Gilchrist, A. Efficient toffoli gates using qudits. *Phys. Rev. A* **75**, 022313 (2007). URL <https://link.aps.org/doi/10.1103/PhysRevA.75.022313>.
- [394] Bækkegaard, T. *et al.* Realization of efficient quantum gates with a superconducting qubit-qutrit circuit. *Sci. Rep.* **9**, 13389 (2019). URL <https://doi.org/10.1038/s41598-019-49657-1>.
- [395] Gokhale, P. *et al.* Asymptotic improvements to quantum circuits via qutrits. In *Proceedings of the 46th International Symposium on Computer Architecture*, ISCA '19, 554–566 (Association for Computing Machinery, New York, NY, USA, 2019). URL <https://doi.org/10.1145/3307650.3322253>.
- [396] Liu, W.-Q. & Wei, H.-R. Optimal synthesis of the Fredkin gate in a multilevel system. *New J. Phys.* **22**, 063026 (2020). URL <https://doi.org/10.1088/1367-2630/ab8e13>.
- [397] Smolin, J. A. & DiVincenzo, D. P. Five two-bit quantum gates are sufficient to implement the quantum Fredkin gate. *Phys. Rev. A* **53**, 2855–2856 (1996). URL <https://link.aps.org/doi/10.1103/PhysRevA.53.2855>.
- [398] Yu, N. & Ying, M. Optimal simulation of deutsch gates and the fredkin gate. *Phys. Rev. A* **91**, 032302 (2015). URL <https://link.aps.org/doi/10.1103/PhysRevA.91.032302>.
- [399] Feynman, R. P. Quantum mechanical computers. *Foundations of Physics* **16**, 507–531 (1986). URL <https://doi.org/10.1007/BF01886518>.
- [400] Zwanenburg, F. A. *et al.* Silicon quantum electronics. *Rev. Mod. Phys.* **85**, 961–1019 (2013). URL <https://link.aps.org/doi/10.1103/RevModPhys.85.961>.
- [401] O’Gorman, B., Huggins, W. J., Rieffel, E. G. & Whaley, K. B. Generalized swap networks for near-term quantum computing. *arXiv:1905.05118* (2019). URL <https://arxiv.org/abs/1905.05118>.
- [402] Younis, E. *et al.* Berkeley Quantum Synthesis Toolkit (BQSKit) v1 (2021). URL <https://www.osti.gov/biblio/1785933>.
- [403] Nemkov, N. A., Kiktenko, E. O., Luchnikov, I. A. & Fedorov, A. K. Efficient variational synthesis of quantum circuits with coherent multi-start optimization. *Quantum* **7**, 993 (2023). URL <https://doi.org/10.22331/q-2023-05-04-993>.

- [404] Qiskit Documentation: transpile (Qiskit). URL <https://qiskit.org/documentation/stubs/qiskit.compiler.transpile.html>. Last accessed on 20-07-2023.
- [405] Liu, J. *et al.* Qcontext: Context-aware decomposition for quantum gates. In *2023 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5 (2023).
- [406] Verstraete, F., Cirac, J. I. & Latorre, J. I. Quantum circuits for strongly correlated quantum systems. *Phys. Rev. A* **79**, 032316 (2009). URL <https://link.aps.org/doi/10.1103/PhysRevA.79.032316>.
- [407] Williams, C. P. *Quantum Gates*, 51–122 (Springer London, London, 2011). URL [https://doi.org/10.1007/978-1-84628-887-6\\_2](https://doi.org/10.1007/978-1-84628-887-6_2).
- [408] Collins, D., Linden, N. & Popescu, S. Nonlocal content of quantum operations. *Phys. Rev. A* **64**, 032302 (2001). URL <https://link.aps.org/doi/10.1103/PhysRevA.64.032302>.
- [409] Kutin, S. A., Moulton, D. P. & Smithline, L. Computation at a distance. *Chic. J. Theor. Comput. Sci.* **2007** (2007).
- [410] Wilde, M. M. *Quantum Information Theory* (Cambridge University Press, Cambridge, 2013). URL <https://doi.org/10.1017/CB09781139525343>.
- [411] Watrous, J. *The Theory of Quantum Information* (Cambridge University Press, Cambridge, UK, 2018). URL <https://doi.org/10.1017/9781316848142>.
- [412] Johansson, J., Nation, P. & Nori, F. QuTiP: An open-source Python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.* **183**, 1760–1772 (2012). URL <https://www.sciencedirect.com/science/article/pii/S0010465512000835>.
- [413] Watrous, J. Simpler semidefinite programs for completely bounded norms. *Chicago Journal of Theoretical Computer Science* **2013**.
- [414] Qiskit Documentation: Measurement Error Mitigation (Qiskit). URL [https://qiskit.org/documentation/stable/0.26/tutorials/noise/3\\_measurement\\_error\\_mitigation.html](https://qiskit.org/documentation/stable/0.26/tutorials/noise/3_measurement_error_mitigation.html). Last accessed on 18-09-2023.
- [415] Nation, P. D., Kang, H., Sundaresan, N. & Gambetta, J. M. Scalable mitigation of measurement errors on quantum computers. *PRX Quantum* **2**, 040326 (2021). URL <https://link.aps.org/doi/10.1103/PRXQuantum.2.040326>.