Representation, Characterization, and Mitigation of Noise in Quantum Processors

by

Junan Lin

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics (Quantum Information)

Waterloo, Ontario, Canada, 2023

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Gerardo Adesso Professor, School of Mathematical Sciences, University of Nottingham
Supervisor(s):	Raymond Laflamme Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal Member:	Kazi Rajibul Islam Assistant Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal-External Member:	Christopher Wilson Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
Other Member(s):	Joel Wallman Adjunct Assistant Professor, Dept. of Physics and Astronomy, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

NOTE: Supervision and contribution from Prof. Raymond Laflamme should be assumed in all of the following. Unless otherwise stated, all authors are involved in reading and improving the final manuscript for each work. Contributors to the work with non-overlapping contributions with Junan Lin are not mentioned.

1. Chapter 2

    - Junan Lin (J.L.) encountered the gauge ambiguity when first studying the SPAM characterization problem. Joel Wallman, Brandon Buonacorsi, and J.L. conceived the original ideas for the paper.
    - J.L. developed the examples based on supervision from Joel Wallman.
    - J.L. produced the figures, ran the simulation for the MVE experiment and generated the plot, and prepared the manuscript. J.L. and Brandon Buonacorsi maintained the manuscript together.
    - This chapter is an improved version of a partially finished work in J.L.'s masters thesis.

2. Chapter 3

    - J.L. designed the 1-qubit and N-qubit experiments based on supervision from Joel Wallman.
    - J.L. derived the bounds under noisy gates, performed all the experiments on the IBMQ platform, performed data analyses, and generated the plots.
    - J.L. and Ian Hincks derived the estimators for the variance of estimated parameters.
    - J.L. wrote and maintained the manuscript.

3. Chapter 4

    - Raymond Laflamme, Tal Mor, and J.L. conceived the idea of using measurement-based algorithmic cooling to characterize SPAM errors.
    - J.L. proved the technical results, including the final state error, bounds on the final error with noisy measurements, and the bound on the expected number of trials. J.L. also generated the plots.
    - J.L. wrote and maintained the manuscript.

4. Chapter 5

    - Raymond Laflamme, Nayeli Azucena Rodríguez-Briones, and J.L. conceived the idea of analyzing thermodynamics of AC.
    - J.L. and Nayeli Azucena Rodríguez-Briones performed extensive studies on AC protocols together, including examining the state evolution, cooling limits, entropy production, mutual information creation, etc. Some materials are still under development.
    - Nayeli Azucena Rodríguez-Briones discovered the unified cooling limits and target evolution.
    - J.L. grouped the protocols under the coherent cooling framework.
    - J.L. and Nayeli Azucena Rodríguez-Briones performed the transition matrix calculations and generated various plots. J.L. produced the plots currently shown in this thesis.
    - J.L. constructed the Landauer Ratio as an efficiency measure, proved Theorem 5.5.1, and realized its relation with the CoP.
    - J.L. developed the energetically improved protocols over the original ones.
    - J.L. and Nayeli Azucena Rodríguez-Briones wrote and maintained the manuscript.

5. Chapter 6

    - Bei Zeng, Ningping Cao, and J.L. first conceived the ideas for a paper contrasting QEM and QEC.
    - J.L. performed the reviews and developed the examples under classical and quantum settings. Bei Zeng and J.L. developed the idea of demonstrating QEM's limitations by drawing connections with entanglement purification protocols.
    - Ningping Cao developed the Drazin inverse as a quasi-inverse, constructed the examples, and proved Theorem 6.3.5.
    - Ningping Cao developed pictorial illustrations (Figs. 6.5 and 6.6), proposed and proved propositions on when QEM can improve experimental results, and constructed examples showing the consequences of a noise model mismatch (Fig. 6.7).
    - J.L. and Ningping Cao wrote and maintained the manuscript.

The work presented contains material from the following publications and preprints:

- Chapter 2: Junan Lin, Brandon Buonacorsi, Raymond Laflamme, Joel J Wallman, "On the freedom in representing quantum operations", *New Journal of Physics, 21*(2), 023006 (2019)

- Chapter 3: Junan Lin, Joel J Wallman, Ian Hincks, Raymond Laflamme, "Independent state and measurement characterization for quantum computers", *Physical Review Research*, 3 (3), 033285 (2021)

- Chapter 4: Raymond Laflamme, Junan Lin, Tal Mor, "Algorithmic cooling for resolving state preparation and measurement errors in quantum computing", *Physical Review A 106 (1), 012439* (2022)

- Chapter 5: Junan Lin, Nayeli Azucena Rodríguez-Briones, Eduardo Martín-Martínez, Raymond Laflamme,"Cooling algorithms and Quantum Thermodynamics", *in preparation* (2023)

- Chapter 6: Ningping Cao, Junan Lin, David Kribs, Yiu-Tung Poon, Bei Zeng, Raymond Laflamme, "NISQ: Error correction, mitigation, and noise simulation", *arXiv:2111.02345* (2021)

Abstract

Quantum computers have the potential to outperform classical computers on several families of important problems, and have a great potential to revolutionize our understanding of computational models. However, the presence of noise deteriorates the output quality from near-term quantum computers and may even offset their advantage over classical computers. Studies on noise in these near-term quantum devices has thus become an important field of research during the past years. This thesis addresses several topics related to this subject including representing, quantifying, and mitigating noise in quantum processors.

To study noise in quantum processors, it is first necessary to ask how noise can be accurately represented. This is the subject of Chapter 2. The conventional way is to use a gate-set, which include mathematical objects assigned to each component of a quantum processor, and compare individual gate-set elements to their ideal images. Here, we present some clarifications on this approach, pointing out that a gauge freedom exists in this representation. We demonstrate with experimentally relevant examples that there exists equally valid descriptions of the same experiment which distribute errors differently among objects in a gate-set, leading to different error rates. This leads us to rethink about the operational meaning to figures of merit for individual gate-set elements. We propose an alternative operational figure of merit for a gate-set, the mean variation error, and develop a protocol for measuring this figure. We performed numerical simulations for the mean variation error, illustrating how it suggests a potential issue with conventional randomized benchmarking approaches.

Next, we study the problem of whether there exist sufficient assumptions under which the gauge ambiguity can be removed, allowing one to obtain error rates of individual gate-set elements in a more conventional manner. We focus on the subset of errors including state preparation and measurement (SPAM) errors, both subject to a gauge ambiguity issue. In Chapter 3, we provide a sufficient assumption that allows a separate SPAM error characterization, and propose a protocol that achieves this in the case of ideal quantum gates. In reality where quantum gates are imperfect, we derived bounds on the estimated SPAM error rates, based on gate error measures which can be estimated independently of SPAM processes. We tested the protocol on a publicly available quantum processor and demonstrated its validity by comparing our results with simulations.

In Chapter 4, we present another protocol capable of separately characterizing SPAM errors, based on a different principle of algorithmic cooling (AC). We propose an alternative AC method called measurement-based algorithmic cooling (MBAC), which assumes the ability to perform (potentially imperfect) projective measurements on individual qubits and

is available on various modern quantum computing platforms. Cooling reduces the error on initial states while keeping the measurement operations untouched, thereby breaking the gauge symmetry between the two. We demonstrate that MBAC can significantly reduce state preparation error under realistic assumptions, with a small overhead that can be upper bounded by measurable quantities. Thus, our results can be a valuable tool not only for benchmarking near-term quantum processors, but also for improving the quality of state preparation processes in an algorithmic manner.

The capability of AC for improving initial state quality has inspired us to perform a parallel study on the thermodynamic cost of AC protocols. The motivation is that since cooling a subset of qubits may result in finite energy increase in its environment, applying them in certain platforms that are temperature-sensitive could induce a negative impact on the overall stability. Meanwhile, previous studies on AC have largely focused on subjects like cooling limits, without paying attention to their thermodynamics. Understanding the thermodynamic cost of AC is of both theoretical and practical interest. These results are presented in Chapter 5. After reviewing their procedure, cooling limits, and target state evolution of various AC protocols, we propose two efficiency measures based on the amount of work required, or the amount of heat released. We show how these measures are related to each other and how they can be computed for a given protocol. We then compare the previously studied protocols using both measures, providing suggestions on which ones to use when these protocols are to be carried out experimentally. We also propose improved protocols that are energetically more favorable over the original proposals.

Finally, in Chapter 6, we present a study on a different family of methods aiming at reducing effective noise level in near-term hardware called quantum error mitigation (QEM). The principle behind various QEM approaches is to mimic outputs from the ideal circuit one wants to implement using noisy hardware. These methods recently became popular because many near-term hybrid quantum-classical algorithms only involve relatively shallow depth circuits and limited types of local measurements, implying a manageable cost of performing data processing to alleviate the effect of noise. Using some intuitions built upon classical and quantum communication scenarios, we clarify some fundamental distinctions between quantum error correction (QEC) and QEM. We then discuss the implications of noise invertibility for QEM, and give an explicit construction called *Drazin-inverse* for non-invertible noise, which is trace preserving while the commonly-used Moore-Penrose pseudoinverse may not be. Finally, we study the consequences of having an imperfect knowledge about the noise, and derive conditions when noise can be reduced using QEM.

Acknowledgements

First and foremost, I would like to thank my supervisor, Prof. Raymond Laflamme, for all his guidance and support throughout my time at IQC. Choosing to finish my PhD in Ray's group has simply been the best decision that I have made in my life so far. I feel particularly privileged to have become Ray's student, from whom I learned so many invaluable lessons on how to be a good researcher. I am thankful for all the trust he has given me to explore different research topics, as well as the opportunities he provided me with to present at international conferences since my early stages. Ray, thank you for being supportive in every aspect of my life as a PhD student, I really enjoyed every single moment in the group.

I would like to thank the members of my PhD supervisory committee: Joel J. Wallman, Christopher Wilson, and Christine Muschik. I appreciate the support and input from every one of them during our regular committee meetings. Special thanks to Joel for providing guidance during my early stages and attending my defence as an examiner, and extra thanks to Chris for being on both my qualifying exam and my defence committee. I have learned a lot about noise characterization from Joel, and I have gained knowledge about superconducting qubits thanks to Chris. In addition, I would like to thank Richard Cleve for being on my qualifying exam committee, and Kazi Rajibul Islam for being my thesis defence examiner. I also thank Gerardo Adesso for coming all the way from Nottingham to attend my thesis defence as the external examiner, and giving me many good questions about my thesis.

I have met many wonderful colleagues in Ray's group, without whom my PhD journey would be incomplete. Xian Ma, Dawei Lu, Keren Li, Guanru Feng, Daniel Park, Franklin Cho, and Tomas Jochym-O'Connor have helped me a lot to integrate into Ray's group when I first joined. Stefanie Beale and I shared a very similar trajectory and has always been a great friend to chat about everything. Christopher Chamberland has been a great officemate as well as someone to look up to when it comes to doing research. Hemant Katiyar and Brandon Buonacorsi have both been knowledgeable friends/colleagues, who also provided valuable career suggestions at later stages of my PhD. Nayeli Azucena Rodriguez Briones has not only been a great friend but also a great collaborator when it comes to algorithmic cooling. Shayan Majidi's enthusiasm for science and life has been constantly inspiring me to be not only a better researcher, but also a better person. I was impressed by Michael Chen's positive attitudes and endless novel ideas during his time here. Many thanks to Annie Ray for constantly challenging me and others in the group with insightful comments, and for sharing the love for badminton. I have always enjoyed conversations with Esha Swaroop about research as well as music, reading, and life. Ningping Cao has

been a talented collaborator with keen mathematical insights, and a great friend to chat on random things with. In addition to the above, I appreciate the presence of all previous and current members of Ray's group, with whom I had wonderful interactions regardless of our overlap. These include but are not limited to: Pooya Ronagh, Aleksander Kubica, Dongsheng Wang, John Peterson, Tian Lan, Michael Vasmer, Anirban Ch Narayan Chowdhury, Matthew Graydon, Danny Paulson, Maria Julia Maristany, Natalie Parham, Matthew Duschenes, Maria Christina Rodriguez, Arsalan Motamedi, Alev Orfi, Emiliia Dyrenkova, and Daniel Honciuc. I would also like to give special thanks to Theerapat Tansuwannont for being a wonderful friend, classmate, and colleague. Ink, thank you for all the great discussions about science and life, and introducing me to bird-watching.

The physics department and the Institute for Quantum Computing have provided me with much help during my time at Waterloo. I would like to thank Steve Weiss in particular for all his amazing IT support and introducing Linux to me. I also acknowledge financial support from NSERC which allowed me to focus fully on my research without having to worry about money.

Finally, I would like to thank my family for their endless love and support. My parents Bin Lin and Xiaoru Cheng have always had my back throughout my life, no matter what I am pursuing and where I decide to go. My partner and best friend Junjie Yin has shared the same PhD life as mine, and I am so glad that we went through this journey together. Thank you for shaping me into a better person with your endless love.

Dedication

This is dedicated to my family.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

The field of quantum computation have witnessed an explosive development over the past two decades. It has transformed from being a niche field studied by a handful of professors in universities and national labs, to a large territory shared by countless research labs and companies around the globe, with a projected billion-dollar market size. Considering their potential to outperform classical computers on certain types of important problems [1, 2, 3, 4], the hype on quantum computation has resulted in a race among almost all major tech companies towards building a useful, large-scale quantum computer.

Meanwhile, it is important to always keep in mind that there is still a long way to go in order for quantum computing to become practically useful. The major obstacle that impedes their usefulness is the presence of noise in quantum processors. The current stage is accurately described by the phrase Noisy, Intermediate-Scale Quantum (NISQ) [5], which corresponds to having processors with the size of roughly 50 to a few hundred qubits, and with lowest error rates per gate falling in the rough range between 10^{-2} to 10^{-4}. Studying the noise behavior in quantum hardware is a crucially important task for the progress of quantum computation in the near term, due to the following reasons. While there exists works which aim to demonstrate the advantage of NISQ devices over classical computers on certain tasks [6, 7], others have argued that such advantages may disappear due to the noisy nature of the output [8]. Furthermore, noisier hardware imply a higher overhead on the number of qubits required, with the implementation of quantum error correcting codes to decrease effective error rates in the hardware. Meanwhile, knowing the form of noise makes it possible to tailor error correcting codes to reduce the overhead if certain characteristics are present in the noise [9], and also allow experimentalists to locate and, consequently, reduce noise in hardware. In this thesis, we examine several topics including representing, quantifying, and reducing noise in quantum processors.

Chapter 2: Gauge and representing noise in quantum processors

Chapter 2 is devoted to the representation of noise. Knowing how to characterize one's control over a quantum system is of great importance in quantum information processing. The subject of characterizing quantum states and processes has been an important subfield in quantum information for decades. The earliest proposal for benchmarking a quantum system is by using quantum state tomography (QST), where given many copies of a system prepared identically, one can measure different components of it and reconstruct the state. This can be dated back to as early as 1957, when Fano [10] mentioned the idea of using multiple copies to measure non-commuting observables. In 1989, Vogel and Risken [11] first described a procedure to determine the Wigner function using homodyne detection in combination with an inverse Radon transform [12]. Since then, various methods have been proposed to refine and improve QST, initially in the field of quantum optics and gradually extended to other systems. Earliest approaches works by linear inversion of an observation matrix [13, 14], but many other reconstruction methods such as using maximum likelihood estimation [15], Bayesian estimation [16], or recent machine-learning-based algorithms [17] have later been proposed, just to name a few.

The same concept was later applied to reconstruct descriptions of the measurement apparatus using well-characterized states, under the name quantum detector (or measurement) tomography [18, 19, 20]. This is easily understood from the dual role between states and measurements under the Born rule. Similarly, with sufficiently many well-characterized states and measurements, one can reconstruct an arbitrary quantum process by applying quantum process tomography (QPT) [15, 21, 16]. While these seem to suggest a possibility of completely characterizing a given quantum computer, a conceptual loophole exists, since characterizing the state requires known measurement apparatus and vice versa, while characterizing the process requires both known states and measurements. Such assumptions are, fundamentally, unrealistic.

A more formal way to express this conceptual difficulty was later coined a "gauge freedom" in representing quantum operations [22, 23, 24], by similarity to gauge theories in physics where there exists unobservable degrees of freedom in the model describing a given physical phenomenon. Similar to a change of reference frame, this implies that there exists different descriptions of the same system. The most notable attempt to deal with this is gate-set tomography (GST), which avoids making any assumption about the uncharacterized system by self-consistently inferring all gate-set elements from experimentally estimated probabilities [25, 22]. This leads to difficulty in interpreting many conventional quality measures, such as the fidelity or trace distance [26], which depend on the particular representation for quantum operations. Despite the broad conceptual importance of

representing quantum operations, the impact of the gauge freedom has only occasionally been analyzed, and primarily in the context of gate-dependent noise in randomized benchmarking [27, 28, 29]. In Chapter 2, we clarify how this gauge freedom affects experimental descriptions and demonstrate some of its implications for interpreting experimentally reconstructed representations of quantum objects. We first illustrate with experimentally relevant examples of how the gauge may appear in a characterization experiment. We then mathematically define what gauge transformations are, and illustrate how it leads to an ambiguity in the representation of quantum errors. We propose a gauge-invariant measure which describes the holistic performance of a quantum processor, and discuss its implications for benchmarking protocols.

Chapter 3: resolving SPAM errors with noise propagation

While there exists conceptual difficulty in fully characterizing all parts of a quantum computer simultaneously, a parallel research direction is to extract partial information about certain parts of a quantum computer, in a manner that is independent of the other parts. The most successful and widely-used method along these lines is randomized benchmarking (RB) and its variants [30, 31, 32, 33, 34, 35, 36], which returns a quality measure (the "RB number") for a quantum processes, or quantum gates, independently of state preparation and measurement. The central idea is that the measurement fidelity would decrease exponentially under a wide range of error models [29, 37] with an increasing number of random gate sequences. These provide an efficient way to obtain a gauge-invariant quality measure for quantum gates. Having been tested on various platforms [38, 39, 40], RB is now among the standard characterization procedures in modern quantum processors.

Compared to the efforts on characterizing quantum gates, much fewer have studied the problem of characterizing state preparation and measurement (SPAM) errors. This may be largely due to the consensus that since SPAM errors do not increase with circuit depth, they play a relatively insignificant role to the overall quality of quantum computation. Meanwhile, SPAM errors nowadays can be on the same order as (and sometimes surpass) gate errors in some current QPUs. For example, the combined SPAM error in current superconducting transmon qubit systems has been reported to range from 0.8% to 2% [41], while 1- and 2-qubit gates may achieve fidelities over 99.9% and 99% respectively [42]. The requirement to repeatedly prepare qubits in well-defined initial states and perform syndrome measurements in quantum error correcting codes also puts SPAM operations on the same level of importance as gate operations. This demands alternative approaches that can provide more trustworthy estimates of SPAM error.

It is also not immediately clear how SPAM errors can be separately obtained in a gauge-invariant manner. While gate errors can be amplified by repeatedly applying more gates with fixed SPAM operators, the converse appears more conceptually difficult since neither the initial state nor the final measurement can be "repeated" in a straightforward manner. Therefore, one may ask under what assumptions can SPAM errors be separately obtained in a gauge-invariant manner, and what procedure one can use to achieve this task? This is the subject of Chapter 3, where we answer these questions. In particular, we discuss the principle of "noise propagation" and how it can be useful in separating SPAM errors. We propose a protocol to separately estimate SPAM error rates for single- and multi-qubit processors, based on this principle. We then show how gate errors can be accounted for in a self-consistent manner, and demonstrate it in an experiment using IBMQ cloud processors as well as corresponding simulations. Our results demonstrate the possibility of separate SPAM error characterization in the regime where SPAM error is more prominent, and can be helpful for further improving the quality of current NISQ devices.

Chapter 4: resolving SPAM errors with algorithmic cooling

Next, in Chapter 4, we provide a second solution to the SPAM characterization problem by means of algorithmic cooling (AC). AC refers to a family of methods that aim to algorithmically purify a target system. These algorithms re-distribute the entropy among subsystems, so that the state of a certain target systems after AC is more pure, as if they were put in a lower temperature than the initial temperature of the original thermal state. This gives rise to the term cooling in the name. A more detailed review of the history and different types of AC will be given in Section 4.1. The simple idea here is that since purifying the initial state breaks the gauge symmetry between state and measurement, this provides another possibility to separate SPAM errors. However, conventional AC approaches are limited by their performance and are difficult to use for the task of characterizing SPAM errors. Here, we introduced measurement as a resource in AC and propose a simple protocol called "measurement-based AC" (MBAC), under the assumption that (potentially imperfect) measurements can be applied on individual qubits. The simple behavior of post-cooling states allows a straightforward procedure to extract SPAM error rates. We study the noisy measurement situation and derive an upper bound for the expected overhead for using this method in practice. The two methods presented in Chapter 3 and Chapter 4 provide new insights into the problem of SPAM characterization, and is valuable to validating QPUs. Moreover, they complement the many existing protocols that measure gate errors.

Chapter 5: thermodynamic studies of AC

Chapter 4 provides compelling evidence that AC can be used not only in noise characterization, but also in noise reduction. This motivates our further investigation into using AC as an algorithmic refrigerator to produce higher quality qubits. The next chapter, Chapter 5, is devoted to studying the thermodynamic cost of using AC to reduce state preparation error. Why is thermodynamics of concern in the first place? Classically, it has long been recognized that erasing information is accompanied by an inevitable energy increase in the environment, a statement typically referred to as Landauer's principle [43]. Moreover, variants of this principle holds in the quantum case too, and since AC effectively drives the state closer towards the ground state, it will also result in an increase of the environment's energy. Meanwhile, many current QPUs operate under low temperature in the milli-kelvin regime, at which small amounts of additional heat release may lead to deterioration in QPU performance. This motivates our study on analyzing AC from a thermodynamic perspective, a path which has not been extensively explored in the literature. First, we give a more in-depth review of common AC protocols, and formally classify them under the same category of "coherent cooling" (defined later). We apply the transfer matrix formalism to consistently calculate their performance metrics, including their cooling limits and target state evolution. Then, we propose two generic efficiency metrics which are applicable to all coherent cooling protocols, namely, the coefficient of performance K and the Landauer Ratio R_L, and derive an explicit formula to compute one from the other. We evaluate the performance of selected HBAC protocols against both metrics, analyzing their relative strengths. Finally, we propose improved versions of HBAC protocols with better thermodynamic performances, which are capable of cooling the target to the same temperature as before while using lower work inputs.

Chapter 6: views on quantum error mitigation

The last chapter, Chapter 6, includes some results on quantum error mitigation (QEM), which refers to a family of numerical approaches to lower the effective noise level through post-processing. Conventionally, the main tool for protecting the processor from noise has been quantum error correction (QEC). QEC protocols are designed to allow a user to detect, and eventually correct, errors that happen during a quantum computation. While many approaches for QEC have been developed, few have been tested on real quantum processors due to the significant demands on the hardware. First, QEC generally encodes quantum information into a much larger Hilbert space, which requires the hardware size to be large as well. Second, quantum operations (gates) on a processor must be below

a certain threshold value for QEC to successfully reduce the effective error, instead of introducing more errors. Meeting both requirements is generally difficult on most state-of-the-art devices available today.

Recently, QEM has become a fast-growing research area due to their immediate applicability on NISQ devices. These protocols generally aim for decreasing the effective noise level, while circumventing the two aforementioned obstacles of QEC. The idea is that if one has some knowledge about the noise processes in a particular hardware, then one should be able to utilize that knowledge to reduce (part of) the effect of that noise. Importantly, it is more desirable to have protocols that introduce no or very little additional hardware overhead in order to improve the computation accuracy. Numerous protocols have been developed during the past few years [44, 45, 46, 47, 48] that fall into this category and we give a further review in Section 6.1.

The parallel development of both fields naturally leads to the question: under what circumstances should one apply QEC over QEM, and vice versa? To answer this question, it is useful to clearly illustrate the differences between QEC and QEM. In this chapter, we examine the relation between QEC and QEM from a high-level perspective, and discuss several implications when using QEM in practice. From a communication viewpoint, we first examine the fundamental differences between QEC and QEM. Because QEM involves effectively applying the inverse noise map to an output probability vector, we examine the special case of non-invertible noise, which may lead to a failure of QEM. We propose a construction called Drazin-inverse for non-invertible noises and demonstrated advantages over the conventional pseudoinverse in certain cases. We then study the effect of imperfect knowledge about the real noise channel, and derive a sufficient condition on when QEM can lead to an improvement in terms of the uncertainty about the actual noise.

Summary

In summary, this thesis has addressed the important field of studying noise in near-term quantum computers, which is crucial for the development and improvement of these devices. Through the examination of noise representation and quantification, the characterization of state preparation and measurement errors, the development of algorithmic cooling protocols and the investigation of quantum error mitigation approaches, this thesis has contributed to advancing the understanding and management of noise in quantum processors. The protocols and methods proposed in this thesis have the potential to be useful for benchmarking near-term quantum processors and improving the quality of state preparation processes. Additionally, this thesis provides insights into the theoretical and

practical aspects of algorithmic cooling and quantum error mitigation. Overall, this thesis has demonstrated the significance and relevance of noise research in the development of quantum computing.

Chapter 2

Representing errors in quantum processors

Quantum mechanics allows us to assign representations that describe the state of quantum objects and processes, and many figures of merit have been developed to evaluate them based on their representations [26]. Interestingly, this representation admits some degrees of freedom, commonly known as the gauge freedom [22, 23, 24], which causes the representation to be non-unique. Since most conventional quality measures are gauge-dependent, changing the representation (or reference frame) will cause the quality measures themselves to change, even if they describe the same underlying hardware. In this chapter, we clarify how this gauge freedom affects experimental descriptions and demonstrate some of its implications for interpreting experimentally reconstructed representations of quantum objects. In Section 2.2 we demonstrate with an experimentally motivated example that the gauge freedom makes assigning errors to individual operations ambiguous. In Section 2.3 we give definitions for gauges and gauge transformations, as well as their role in representing quantum operations as mathematical objects. In Section 2.4 we discuss the implications brought by this gauge freedom, in particular addressing why many figures of merit such as the diamond norm distance between a measured gate and a target do not have a concrete operational meaning. We also mention some common practices in tomography that are related to this problem. Lastly, in Section 2.5, we define and motivate the mean variation error (MVE), a gauge-invariant figure of merit for gate-sets. We propose a protocol to experimentally measure the MVE and demonstrate its behaviour relative to randomized benchmarking through numerical simulations.

2.1 Pauli-Liouville representation

Before illustrating the issue with representing quantum errors due to the gauge freedom, we will first give an introduction the Pauli-Liouville representation formalism. This representation, along with variants of it such as the superoperator or the transition matrix formalisms, will be used multiple times later throughout this thesis. The central idea is to

1. Represent general quantum state as vectors,
2. Represent general quantum evolutions as matrices,
3. Construct the representations in such a way that the actions of quantum evolutions on quantum states are written as matrix multiplications.

For this chapter, we will focus on system of n qubits where quantum states can be represented as $2^n \times 2^n$ Hermitian operators, while the same formalism can be generalized to arbitrary dimension systems by using, for example, the generalized Gell-Mann matrices [49] instead of Pauli matrices. It is commonly known that the set of (tensor products of) normalized Pauli matrices, which we denote as $\mathbf{P}_n$, form an orthonormal basis for all $2^n \times 2^n$ Hermitian operators. Every element P in $\mathbf{P}_n$ is of the form

$$P = \bigotimes_k \left(\frac{\sigma_k}{\sqrt{2}} \right) \tag{2.1}$$

and each σ_k is a member from the single-qubit Pauli group, $\mathbf{P}_1 \coloneqq \{\sigma_0 = I,\ \sigma_1 = X,\ \sigma_2 = Y,\ \sigma_3 = Z\}$. The orthonormality is defined with respect to the Hilbert-Schmidt inner product

$$\langle P_i,\ P_j \rangle_{HS} \coloneqq \mathrm{Tr}\left[P_i^\dagger P_j \right] = \delta_{ij} \tag{2.2}$$

for all $P_i, P_j \in \mathbf{P}_n$. Any 2^n by 2^n Hermitian matrix can be represented as a real linear combination of Pauli basis matrices. Writing these inner products as components of a vector will define a representation in the space of $2^{2n} \times 1$ real vectors, which is isomorphic to the set of $2^n \times 2^n$ real matrices.

For every 2^n by 2^n Hermitian matrix ρ, we define its Pauli-Liouville representation as follows:

$$|\rho\rangle\rangle \coloneqq \sum_i \mathrm{Tr}[\rho P_i] |i\rangle\rangle \tag{2.3}$$

and define an element σ in the dual space (e.g., representing a measurement operator) as

$$\langle\langle\sigma| \coloneqq \sum_i (\mathrm{Tr}[P_i\sigma])^* \langle\langle i| \tag{2.4}$$

where $|i\rangle\rangle$ and $\langle\langle i|$ are standard computational (column and row) basis vectors with 1 in the i-th entry and 0 elsewhere, and * denotes complex conjugation for a scalar. We see that the Hilbert-Schmidt inner product is now transformed into an Euclidean inner product:

$$\begin{aligned}
\langle\langle\sigma|\rho\rangle\rangle &= \sum_{i,j} \mathrm{Tr}[\rho P_i](\mathrm{Tr}[P_j\sigma])^* \langle\langle j|i\rangle\rangle \\
&= \sum_{i,j} \mathrm{Tr}[\rho P_i](\mathrm{Tr}[P_j\sigma])^* \,\mathrm{Tr}[P_iP_j] \\
&= \mathrm{Tr}\left[\sum_i \mathrm{Tr}[\rho P_i] P_i \sum_j (\mathrm{Tr}[P_j\sigma])^* P_j\right] \\
&= \mathrm{Tr}\big[\sigma^\dagger \rho\big]
\end{aligned} \tag{2.5}$$

where the following identity is used:

$$\begin{aligned}
\sum_i \mathrm{Tr}[P_iA]^* P_i &= \sum_i \left(\mathrm{Tr}\Big[P_i^\dagger A\Big]\right)^* P_i \\
&= \sum_i \mathrm{Tr}\Big[\overline{(A^\dagger P_i)^\dagger}\Big] P_i \\
&= \sum_i \mathrm{Tr}\big[(A^\dagger P_i)^T\big] P_i \\
&= \sum_i \mathrm{Tr}\big[A^\dagger P_i\big] P_i = A^\dagger
\end{aligned} \tag{2.6}$$

where the overhead bar denotes element-wise conjugation for a matrix.

Now, define the Pauli-Liouville representation of a (linear) map $\mathcal{G}$ as $\mathcal{A}_\mathcal{G}$, which has components

$$(\mathcal{A}_\mathcal{G})_{ij} \coloneqq \mathrm{Tr}[P_i\mathcal{G}(P_j)] \tag{2.7}$$

then the post-state of $\mathcal{G}$ acting on a state ρ, written all in the Pauli basis, can be shown

to be equal to a matrix multiplication:

$$\begin{aligned}
|\mathcal{G}(\rho)\rangle\!\rangle &= \sum_i \mathrm{Tr}[\mathcal{G}(\rho)P_i]|i\rangle\!\rangle \\
&= \sum_i \mathrm{Tr}\left[\mathcal{G}(\sum_j \mathrm{Tr}[\rho P_j]P_j)P_i\right]|i\rangle\!\rangle \\
&= \sum_{ij} \mathrm{Tr}[\rho P_j]\,\mathrm{Tr}[\mathcal{G}(P_j)P_i]|i\rangle\!\rangle \\
&= \sum_{ij} (\mathcal{A}_\mathcal{G})_{ij}(|\rho\rangle\!\rangle)_j|i\rangle\!\rangle \\
&= \mathcal{A}_\mathcal{G}|\rho\rangle\!\rangle
\end{aligned} \tag{2.8}$$

Thus, series of gates are conveniently expressed as matrix multiplications (from the left) in the Pauli-Liouville representation.

2.2 Assigning errors to operations

We now illustrate the gauge freedom with a simple, experimentally relevant example, namely, amplitude damping. We will do so mostly in the gate-set language. A gate-set is a mathematical description of the possible actions executable in an experiment, typically consisting of models for initial states ($\mathbb{S}$), gate operations ($\mathbb{G}$), and measurements ($\mathbb{M}$). If an experimentalist with an ideal quantum system could initialize a qubit in the state $|0\rangle$, apply an arbitrary unitary gate, and measure the expectation value of Z, then their control can be represented by the gate-set

$$\Phi = \left\{ \mathbb{S}_\Phi = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},\ \mathbb{G}_\Phi = \mathrm{SU}(2),\ \mathbb{M}_\Phi = \mathrm{Z} \right\}. \tag{2.9}$$

Now suppose that the experimentalist prepares a mixed initial Z state with polarization ϵ_1 and performs a measurement with signal-to-noise ratio ϵ_2. Suppose further that before each gate is applied, the system undergoes amplitude damping with strength γ but that the target Hamiltonian is still implemented perfectly. The Pauli-Liouville representation (see Section 2.1 or e.g., [50]) of the noisy gate-set is then

$$\Theta = \left\{ \mathbb{S}_\Theta = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ \epsilon_1 \end{pmatrix},\ \mathbb{G}_\Theta = \{\mathcal{U}\mathcal{A}_\gamma : U \in \mathrm{SU}(2)\},\ \mathbb{M}_\Theta = \sqrt{2}\begin{pmatrix} 0 & 0 & 0 & \epsilon_2 \end{pmatrix} \right\} \tag{2.10}$$

with $\sqrt{2}$ and $\frac{1}{\sqrt{2}}$ being normalization factors, and $\mathcal{U} = U\rho U^\dagger$ denotes the unitary channel acting via conjugation and

$$\mathcal{A}_\gamma = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-\gamma} & 0 & 0 \\ 0 & 0 & \sqrt{1-\gamma} & 0 \\ \gamma & 0 & 0 & 1-\gamma \end{pmatrix}. \tag{2.11}$$

The expectation value of an operator M given an input state ρ is the vector inner product between the Pauli-Liouville representations of the state and measurement operators,

$$\text{prob} = \langle\langle M|\rho\rangle\rangle. \tag{2.12}$$

If m gates $\mathcal{G}_1, \ldots, \mathcal{G}_m \in \mathbb{G}$ are applied to the state in chronological order before the measurement takes place, the expectation value becomes

$$\text{prob} = \langle\langle M|\mathcal{G}_{m:1}|\rho\rangle\rangle \tag{2.13}$$

where we use the shorthand notation

$$\mathcal{G}_{b:a} \coloneqq \begin{cases} \mathcal{G}_b\mathcal{G}_{b-1}...\mathcal{G}_a & \text{if } b \geq a \\ \mathcal{I} & \text{otherwise.} \end{cases} \tag{2.14}$$

The above probabilities are preserved under the family of gate-set transformations

$$|\rho\rangle\rangle \to B|\rho\rangle\rangle,\ \langle\langle M| \to \langle\langle M|B^{-1},\ \mathbb{G}_\Phi \to B\mathbb{G}_\Phi B^{-1} \tag{2.15}$$

for some invertible matrix B. Because these probabilities are the only experimentally accessible quantities, the same experimental results can be predicted equally well by these two gate-sets. This is the gauge freedom inherent in mathematically representing quantum experiments, in analogy with concepts in thermodynamics and electromagnetism [51], with B being called a gauge transformation matrix. The analogy arises from the fact that changing the gauge does not result in observable effects in an experiment, just as changing the electromagnetic gauge would not result in any difference in the measurable electric or magnetic fields.

Generally, a gate-set is considered valid if all quantum states can be represented as density matrices, measurements as expectation values of Hermitian operators, and quantum gates as completely-positive, trace-preserving (CPTP) maps as these conditions ensure that

probabilities for arbitrary experiments are positive. Gauge transformations do not generally preserve these *canonical constraints*, although the resulting gate-set is nevertheless an equally valid mathematical description of the same experiment.

We now present a simple, physically motivated, gauge transformation that yields a new gate-set which suggests a different physical interpretation of the experimental system. Applying the gauge transformation matrix

$$B = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & q & 0 & 0 \\ 0 & 0 & q & 0 \\ 0 & 0 & 0 & q \end{pmatrix} \tag{2.16}$$

for any $q \in [-1, 1]$ to the noisy gate-set in Eq. (2.10) yields the equivalent gate-set

$$\Theta_q = \left\{ \mathbb{S}_{\Theta_q} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ q\epsilon_1 \end{pmatrix},\ \mathbb{G}_{\Theta_q} = \{\mathcal{U}\mathcal{A}_{\gamma,q} : U \in \mathrm{SU}(2)\},\ \mathbb{M}_{\Theta_q} = \sqrt{2} \begin{pmatrix} 0 & 0 & 0 & \frac{\epsilon_2}{q} \end{pmatrix} \right\}, \tag{2.17}$$

where

$$\mathcal{A}_{\gamma,q} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-\gamma} & 0 & 0 \\ 0 & 0 & \sqrt{1-\gamma} & 0 \\ q\gamma & 0 & 0 & 1-\gamma \end{pmatrix} \tag{2.18}$$

and we have used the fact that $\mathcal{U}$ commutes with B for any $U \in \mathrm{SU}(2)$. The gauge transformation results in equivalent statistics but suggests a different noise model, namely, relaxation to a mixed state rather than a pure state (corresponding to a different effective temperature). As long as $|q| \in [|\epsilon_2|, 1]$, the states, measurements, and transformations all satisfy the canonical constraints for gate-set elements.

Note that this gauge freedom does not change the average gate fidelity as $\mathrm{Tr}[\mathcal{A}_{\gamma,q}]$ is independent of q [52, Eq. 2.5]. However, the diamond distance from the identity depends on q, with $\|\mathcal{A}_{\gamma,1} - \mathcal{I}\|_\diamond \approx 2\|\mathcal{A}_{\gamma,0} - \mathcal{I}\|_\diamond$ for $\gamma \in [0, 1]$ [53]. Moreover, this example illustrates that noise can be artificially reassigned to different objects, as the state in Eq. (2.10) is closer to pure than the one in Eq. (2.17). Note that the range of gauge transformations is constrained by ϵ_2 and so cannot significantly change the effective temperature for systems with high quality readout. We could have added larger errors by considering a non-unital (e.g., $\mathcal{A}_{\gamma,q}B$) or unitary gauge transformation at the cost of making the errors gate-dependent

and consequently giving a more complicated example. We did not do this as our intent is to clarify that the gauge freedom is more than a basis mis-match [37] and is distinct from the issue of gate-dependent noise [27, 29, 37]. In particular, we note that the full effect of the gauge freedom for states and measurements is unknown.

2.3 Gauge and Representation of Quantum States

We have seen that under realistic circumstances, the same experiment can be described by distinct gate-sets that suggest different physical noise models due to a gauge freedom. In this section we illustrate how representations of quantum states are related to the concepts of gauges and gauge transformations. For clarity we focus on the representation for quantum states, but similar arguments can be made about gate and measurement operations.

From the point of view of scientific realism, the apparatus (e.g., a qubit) has a physical existence and properties (which may be relative to the environment) independent of our representation. We describe the abstract state of this physical object as a *noumenal state* following the terminology in [54], denoted as $\mathbb{N}$ in Figure Fig. 2.1. Here we slightly change their definition to include in $\mathbb{N}$ both physically allowed (denoted as $\mathbb{P}$) and forbidden states, such that the set $\mathbb{N}$ contains both states that the system can be in, and ones that it cannot be in based on the physics. Quantum mechanics allows us to assign to each noumenal state a mathematical *representation* which is an element of a Hilbert space $\mathcal{H}^d$: for example, one can associate the system with a matrix that summarizes its properties, and the set of all $d \times d$ matrices is called $\mathbb{R}$ in the same figure. Such an association is what we call a *gauge* Γ, which is a bijective map from $\mathbb{N}$ to $\mathbb{R}$: the bijectivity of the map should be clear from our inclusion of physically-forbidden states in $\mathbb{N}$, which allows assigning "some state" to every $d \times d$ matrix. Different choices of Γ thus correspond to different mathematical descriptions of the noumenal states.

The common formulation of quantum mechanics says that every state of a quantum object can be described by a density operator [55], which is a part of the canonical constraints defined in Section 2.2. This means that there exists a *canonical gauge*

$$\Gamma_1 : \ N \to R, \ \Gamma_1(\mathbb{P}) = \mathbb{D}_d \tag{2.19}$$

where $\mathbb{D}_d$ is the set of $d \times d$ density operators. In fact, there exists a family of canonical gauges that are all related to Γ_1 through unitary gauge transformations which preserves the shape of $\mathbb{D}_d$. Satisfying the canonical constraint implies that we should work in one of

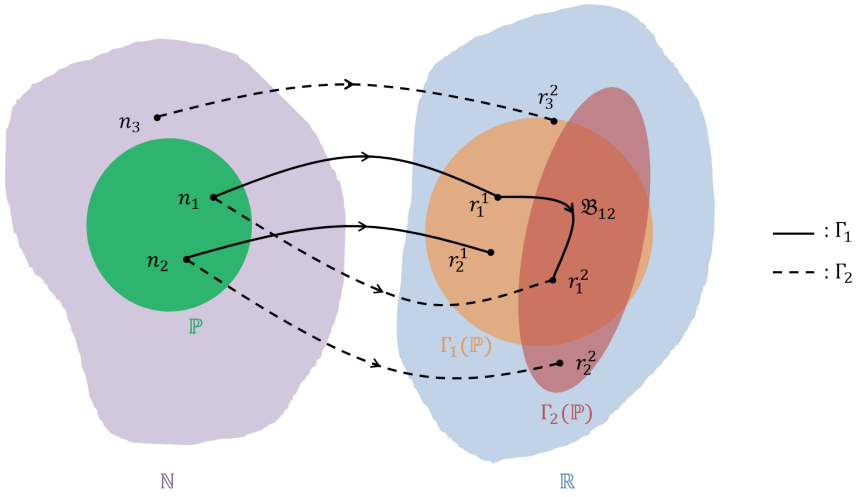

Figure 2.1: Diagram illustrating the distinction between the set of noumenal states $\mathbb{N}$ and their mathematical representations $\mathbb{R}$. For a single qubit, $\mathbb{R}$ is the set of 2×2 matrices. A gauge is a bijective map $\Gamma : N \to R$. Let $\mathbb{P}$ be the set of all noumenal states that can possibly be prepared physically. In general, $\mathbb{P}$ is unknown, but $\Gamma(\mathbb{P})$ is assumed to satisfy the canonical constraints, that is, to be the set of density matrices. In our example, $\Gamma_1(\mathbb{P})$ is the set of density operators in the corresponding Hilbert space and $\mathfrak{B}_{12}$ is a gauge transformation from Γ_1 to Γ_2. Whether an object in $\mathbb{R}$ directly corresponds to objects in $\mathbb{P}$ or not depends on the particular gauge under which the object is represented. For example, $r_2^1 = \Gamma_1(n_2)$ is a density operator while $r_2^2 = \Gamma_2(n_2)$ is not, despite both being the image of the same physical state.

these canonical gauges. Now, consider another gauge Γ_2 which can be converted from Γ_1 with a *gauge transformation* $\mathfrak{B}_{12}$, defined by

$$\mathfrak{B}_{12} \coloneqq \Gamma_2(\Gamma_1^{-1}), \ \mathfrak{B}_{12}(r_*^1) = r_*^2 \tag{2.20}$$

where the r's are members in $\mathbb{R}$ and the superscript denotes the gauge in which they are represented. In the light of Eq. (2.15), this transformation can be represented in Pauli-Liouville representation as

$$|\mathfrak{B}_{12}(\rho)\rangle\!\rangle \coloneqq B_{12}|\rho\rangle\!\rangle, \ \langle\!\langle \mathfrak{B}_{12}(M)| \coloneqq \langle\!\langle M|B_{12}^{-1}, \ \mathbb{G}_{\mathfrak{B}_{12}(\Phi)} \coloneqq B_{12}\mathbb{G}_\Phi B_{12}^{-1} \tag{2.21}$$

As a subset of $\mathbb{R}$, $\Gamma_1(\mathbb{P})$ is generally not invariant under an arbitrary gauge transformation: consider a general trace-preserving transformation given by the following transformation matrix

$$B_{12} = \begin{pmatrix} 1 & 0 \\ \vec{x} & y \end{pmatrix} \tag{2.22}$$

where $\vec{x}$ is a (d^2-1) by 1 real vector and y is a (d^2-1) by (d^2-1) real matrix: the image of this affine transformation of $\Gamma_1(\mathbb{P})$ is a different subset of $\mathbb{R}$. Such a gauge is perfectly valid in principle, provided that *all* the gates and measurement operators are transformed according to Eq. (2.21) as well, even though $\Gamma_2(\mathbb{P})$ is no longer the set of density operators.

The existence of a non-canonical gauge implies, for example, that a physical state may or may not be represented by a density operator: as illustrated in Fig. 2.1, $r_1^2 \in \Gamma_1(\mathbb{P})$ whereas $r_2^2 \notin \Gamma_1(\mathbb{P})$. Similarly, a density operator in a non-canonical gauge does not necessarily correspond to a physical state, as $r_3^2 \in \Gamma_1(\mathbb{P})$ but $n_3 \notin \mathbb{P}$. One example for the state n_3 is a qubit state represented as $\frac{1}{2}(I + \frac{10}{9}\sigma_z)$ in a canonical gauge. It is not positive semidefinite and thus lies outside I_1, representing an abstract state the qubit cannot be in. Now, using $B_{12} = B$ from Eq. (2.16) with $q = \frac{9}{10}$, the image of n_3 under Γ_2 becomes $\frac{1}{2}(I+\sigma_z)$, which *is* a density operator, but only as a consequence of this non-canonical gauge. We conclude that if the gauge is unknown, the mathematical representation does not imply the noumenal state is physically possible. Representations satisfying the canonical constraints are easier to work with, so it is often implicitly assumed that all gate-set elements (obtained from a tomography experiment, for example) are expressed in a canonical gauge. However, this assumption can only be verified by performing perfect experiments, which are axiomatically the operations specified by the canonical constraints (up to a unitary change of basis).

2.4 Operational interpretations of figures of merit

The existence of this gauge freedom has direct implications for figures of merit used to evaluate quantum operations. The main problem is that there is no way to know whether an experimentally-determined gate-set element is expressed in a canonical gauge. We have already seen in Section 2.2 that by changing the gauge, the states can appear as having different expressions; the same holds true for gates and measurement operators.

From quantum information theory, we have successfully attached some operational meanings to various distance metrics: an important example is the interpretation for the diamond norm distance between two channels $\mathcal{A}$ and $\mathcal{B}$ as the maximum distinguishability between output states under a fixed input [56]. Mathematically,

$$\frac{1}{2}\|\mathcal{A}-\mathcal{B}\|_\diamond = \max_{M\in\Gamma(\mathbb{M}),\rho\in\Gamma(\mathbb{P})} \langle\langle M|(\mathcal{A}-\mathcal{B})\otimes I)|\rho\rangle\rangle \tag{2.23}$$

where $\mathbb{P}$ and $\mathbb{M}$ are the set of physically possible states and measurements respectively. This operational meaning is gauge invariant, provided one consistently transforms $\mathcal{A}$, $\mathcal{B}$,

$\Gamma(\mathbb{P})$, and $\Gamma(\mathbb{M})$. However, when $\mathcal{A}$ is an experimentally reconstructed gate and $\mathcal{B}$ is its ideal target, $\Gamma(\mathbb{P})$ and $\Gamma(\mathbb{M})$ are unknown and so the above maximization that leads to its operational meaning cannot be performed. To obtain concrete numbers, people calculate

$$\frac{1}{2}\|\mathcal{A}-\mathcal{B}\|_{\diamond} = \max_{M\in\mu(\mathbb{D}_d),\rho\in\mathbb{D}_d} \langle\langle M|(\mathcal{A}-\mathcal{B})\otimes I)|\rho\rangle\rangle \tag{2.24}$$

where $\mu(\mathbb{D}_d)$ is the set of all POVMs. However, this assumes that the reconstructed $\mathcal{A}$ and the ideal target $\mathcal{B}$ are expressed in a canonical gauge. While $\mathcal{B}$ is an ideal gate, its representation may not be unitary in a non-canonical (and unknown) gauge. Other works have reported that the quantity $\frac{1}{2}\|\mathcal{A}-\mathcal{B}\|_{\diamond}$ can be changed by changing the gauge and used this to minimize reported error rates [23], however, such changes are obtained by implicitly changing the set of physically allowed states and measurements. Note that even in one special case of interest where $\mathcal{B}=\mathcal{I}$, which is gauge invariant, $\mathcal{A}$ is still reconstructed in an unknown gauge.

We briefly discuss several common practices related to this gauge freedom in quantum tomography. First, a common statistical method used in tomographic reconstructions is Maximum Likelihood Estimation (MLE), which takes the estimated gate-set to be the one that maximizes the likelihood function of obtaining the experimental data, while restricting the gate-set elements to satisfy the canonical physicality constraints [57, 58]. However, all gauge-equivalent gate-sets are equally likely to produce the data by definition. In the process of optimization, one will find that the likelihood function profile has the same value wherever two points are related by a gauge-transformation, and the actual output is largely a matter of the optimization algorithm and the initial parameters [59].

Second, the process known as "gauge optimization" is commonly adopted in GST experiments whereby the gauge transformation matrix B is varied to minimize the distance from the target gate-set according to a (non-gauge-invariant) weighted distance measure [23]. Such optimization undermines a common use of tomography, namely assessing the performance of a system against some external threshold (e.g., a fault-tolerance threshold) because this optimized gauge is just as arbitrary as any other gauge. Specifically, one still cannot know whether the resultant gate-set is a faithful representation of the apparatus, in particular, whether the states and measurements that satisfy the canonical constraints are actually the images of the set of physically possible states and measurements respectively.

Furthermore, the optimization can change the relative size of errors from different components of the device, leading to misunderstandings about their relative quality. This misidentification of error was recently observed experimentally in a trapped ion processor: in particular, fig. 4(h) in [60] demonstrates that the gauge optimization procedure may

effectively cancel some gate errors that were purposely added, resulting in a smaller "diamond norm distance" than expected, which implies unrealistically good quantum gates. In that paper, and in other systems where, for example, a basis change in the classical software is used to achieve certain "virtual" gates [61, 62], experimentalists have *a priori* information about which operations are better. This is also true for state and measurement operations in some systems: for example, in an NMR spectrometer there is a well-defined Z direction set by the external magnetic field, which defines the initial qubit states with which all other operations are calibrated. But incorporating such information into the optimization procedure is a challenging task, and the resulting representation will only be as reliable as the prior information.

2.5 A Gauge-Invariant Measure for Gate-Sets

The gauge freedom prevents one from using conventional distance measures to faithfully evaluate the quality of individual quantum operations. Note that our discussion is carried out in the absence of any additional errors such as finite-counting, and in a real experiment the situation becomes even more complicated. Fundamentally, this problem is due to the limited information that can be gained from experimental probabilities. A gauge-transformation re-assigns state, gate, and measurement "errors" by adjusting their relative appearance in different representations, while keeping the experimental measurables unchanged, although some degrees of freedom can be fixed by convention (e.g., that the state preparation is diagonal in the Z eigenbasis).

We now propose a gauge-invariant figure of merit for a *gate-set*. To the best of our knowledge, this is the first fully gauge-invariant measure, addressing a problem raised in Ref. [59]. Let Φ denote the gate-set $\{\mathbb{S}, \mathbb{G}, \mathbb{M}\}$ and C denote a particular experiment with input state $\rho \in \mathbb{S}$, measurement $M \in \mathbb{M}$, and a set of m gates $\mathcal{G}_1 ... \mathcal{G}_m$ each selected from $\mathbb{G}$. The only observable property of an experiment C is the probability distribution over outcomes. We can quantify the error of the experiment by the total variation distance between the observed and ideal distributions over outcomes,

$$\delta d(C, \tilde{C}) \coloneqq \frac{1}{2} \sum_i \left| \mathrm{Tr}\left[\tilde{M}_i^\dagger \tilde{G}_{m:1}(\tilde{\rho})\right] - \mathrm{Tr}\left[M_i^\dagger G_{m:1}(\rho)\right] \right| \tag{2.25}$$

where the tilde represents real versions of the operations, and the subscript i denotes different measurement operators such that $\sum_i M_i = I$. For a particular experiment, this number only depends on outcome probabilities, and therefore is gauge-invariant by definition. The total variation distance is a metric between two probability distributions (for

further motivation and background on this metric, see chapter 4.1 of [63]). Denoting the set of all experiments with m gates by $\mathbb{A}_m$, we further define the *Mean Variation Error* (MVE) over $\mathbb{A}_m$ with the underlying gate-set Φ as

$$x(\Phi, m) \coloneqq \frac{1}{|\mathbb{A}_m|} \sum_{C \in \mathbb{A}_m} [\delta d(C, \tilde{C})] \tag{2.26}$$

Note that MVE is a special case of the *mean variation distance* (MVD) between two gate-sets Φ_1 and Φ_2 with the same number of elements in $\mathbb{S}$, $\mathbb{G}$ and $\mathbb{M}$, where Φ_2 is chosen to be an ideal version of Φ_1. The MVD is a metric between two gate-sets, up to a gauge transformation, if the images of state preparation and measurement processes are tomographically complete. Also note that the size of $\mathbb{A}_m$ is given by $|\mathbb{A}_m| = |\mathbb{S}||\mathbb{G}|^m|\mathbb{M}|$.

The MVE quantifies on average how well the apparatus performs a random experiment from the gate-set. In the case where the measurement is a projective measurement in the basis of the initial state (i.e., $\rho = M_i$ for some i) and the gate sequence is self-inverting[1], $\delta d(C, \tilde{C})$ can be simplified as

$$\begin{aligned} \delta d(C, \tilde{C}) &= \frac{1}{2} \left(\left| \mathrm{Tr}\left[\tilde{M}_i^\dagger \tilde{G}_{m:1}(\tilde{\rho}) \right] - 1 \right| + \sum_{j \neq i} \left| \mathrm{Tr}\left[(I - \tilde{M}_j)^\dagger \tilde{G}_{m:1}(\tilde{\rho}) \right] - 0 \right| \right) \\ &= 1 - \mathrm{Tr}\left[\tilde{M}_i^\dagger \tilde{G}_{m:1}(\tilde{\rho}) \right] \end{aligned} \tag{2.27}$$

whose average over $\mathbb{A}_m$ is just 1 minus the "survival probability" plotted in a conventional randomized benchmarking experiment. When $\mathbb{G}$ is a unitary 2-design, the MVE restricted to self-inverting gate sequences is well-approximated by a linear relation to first order in the average error rate [29, 27].

However, for generic gate sequences, the MVE behaves differently depending on the underlying error model. This behavior provides additional information about the underlying error mechanism compared to a conventional randomized benchmarking experiment [64]. To illustrate this, we simulated random circuits of varying length m sampled from the gate-set $\Phi = \{\mathbb{S} = |0\rangle\langle 0|, \mathbb{G} = Cl_1, \mathbb{M} = |0\rangle\langle 0|\}$ (with Cl_1 denoting the set of 1-qubit Clifford gates), where erroneous gates are represented as $\tilde{\mathcal{G}} = \mathcal{E}\mathcal{G}$ for a fixed error channel $\mathcal{E}$. We simulated two types of random circuits: circuits from the entire set of possible experiments allowed by the gate-set, and circuits restricted to self-inverting gate sequences. In both simulations, the state and measurements are assumed to be error-free. The results

[1] A self-inverting circuit is one where all the gates performed in the circuit compose to the identity operation in the absence of gate errors.

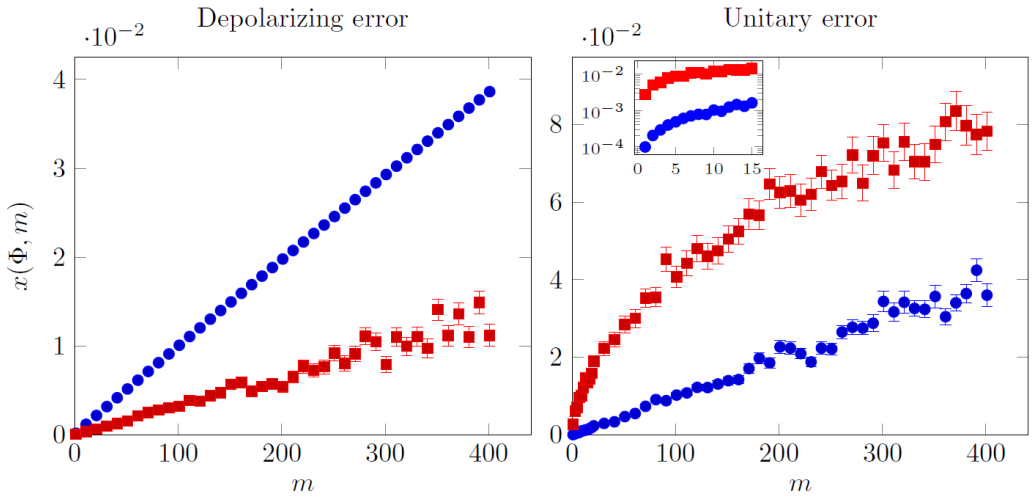

Figure 2.2: Simulated mean variation error from Eq. (2.26) under two error models for a gate-set with $\rho = M_1 = |0\rangle\langle 0|$, and $\mathbb{G} = Cl_1$ being the 1-qubit Clifford group. The depolarizing error channel is $\mathcal{E}_D(\rho, r) = (1-2r)\rho + rI$, whereas the unitary error is $\mathcal{E}_U(\rho, \theta) = e^{-i\theta Z}$ with $\theta = \arccos\left(\sqrt{1-3r/2}\right)$, such that the error channel on every gate has an averaged infidelity of $r = 10^{-4}$. Blue circles indicate self-inverting (identity) circuits whereas red squares indicate random circuits. Each point is generated from averaging 200 random circuits with length m; error bars are standard error in the mean and data show significant spread for unitary error. MVE may have different behaviors under different error types (m or $\sqrt{m}$) for a random circuit, as compared to the linear behavior for a self-inverting circuit. The inset in the second plot is a zoom-in view for small m, showing the significant underestimation of MVE by restricting to self-inverting circuits.

are shown in Fig. 2.2. When the error is a depolarizing channel, the MVE scales linearly with the gate sequence length m for both random and self-inverting circuits, with the slope for random circuits being $\sim 1/3$ the slope for self-inverting circuits. This is because when the state is transformed onto the xy-plane of the Bloch sphere right before measurement (which happens about 2/3 of the time), the depolarizing channel does not affect the outcome probability of a z-axis measurement, resulting in an MVE of 0 for those circuit sequences. Additionally, there is no statistical error present for the self-inverting circuit under this error model because all circuits of the same sequence length have exactly the same overall error, as the error channel commutes with all the gates in Cl_1. In contrast, for a gate-independent unitary error, the scaling remains linear for the self-inverting cir-

cuits but exhibits a $\sqrt{m}$ scaling for generic circuits. This occurs because when the state system is in the xy-plane before measurement, each error contributes a random sign to the probability of each outcome, whereas when the system is on the z axis each error has to contribute a systematic sign [64]. As shown in Fig. 2.2, this implies that restricting to self-inverting circuits can underestimate the MVE by over an order of magnitude in the small-m regime, which is relevant for near-term quantum computer applications.

Unlike other distance measures where an improvement in quality can be caused by a bias in choosing a gauge, a decrease in MVE is unequivocally an improvement due to its gauge-invariance and because, by definition, the output probability distribution gets closer to the ideal distribution. Furthermore, the MVE captures the relevant behavior for generic circuits, rather than just self-inverting circuits which, by design, perform no useful computation.

A protocol for estimating the MVE of a gate-set $\Phi = \{\mathbb{S}, \mathbb{G}, \mathbb{M}\}$ is as follows:

1. Select N_m random experiments $C \in \mathbb{A}_m$, for some N_m large enough to accurately approximate the average.
2. Repeat each experiment C K_m times to estimate $\langle \tilde{M}_i, \tilde{G}_{m:1}(\tilde{\rho})\rangle$ for each C.
3. Compute the ideal probabilities $\langle M_i, G_{m:1}(\rho)\rangle$ for each observed outcome of $\tilde{C}$.
4. Calculate $\delta d(C, \tilde{C})$ for each experiment C, average over them to estimate $x(\Phi, m)$.
5. Repeat step 1–4 for different values of m to measure the scaling behaviour of MVE.

Note that if $\mathbb{G}$ is a unitary 2-design and the states and measurements are chosen appropriately, the applied operations are identical to those used to estimate the unitarity [65]. The primary difference in the protocol is that it is more scalable, more general, and has different post-processing.

The scalability of the above protocol is affected by the number of experiments N_m, the number of repetitions for each experiment K_m, and the complexity of calculating the ideal probabilities. The number of experiments determine the accuracy of the MVE, and can be estimated using Hoeffding's inequality independently of the number of qubits [66]. The complexity in the protocol is determined by the complexity of calculating probabilities and by the number of repetitions required to estimate $\delta d(C, \tilde{C})$ to a fixed precision. The number of repetitions required to estimate $\delta d(C, \tilde{C})$ to a fixed precision is polynomial in the number of outcomes [67]. To efficiently characterize multi-qubit gate-sets (where the number of raw outcomes grows exponentially with the number of qubits), we can coarse-grain the

measurements over sets of outcomes. The computational complexity of calculating each probability will depend on the gate-set in question. The ideal probabilities can be efficiently computed if $\mathbb{G}$ is the N-qubit Clifford group [68, 69]. For gate-sets containing only one- and two-qubit gates and product states and measurements, the MVE can be computed for small values of m. However, for a generic gate-set, each probability will be hard to compute. Of course, for small systems with a few qubits, this procedure can nonetheless be performed quickly on a classical computer.

An experimentalist can perform a feedback loop whereby they update the control parameters, rerun the MVE evaluation experiment (potentially for some fixed value of m) and compare to the previous result to see if the error has decreased. Protocols that use feedback from experimental outcomes to improve control over quantum devices have been proposed before, such as in [70] where control parameters were optimized by maximizing the randomized benchmarking survival probability for a fixed sequence length. Optimizing the MVE instead of the randomized benchmarking survival probability corresponds to minimizing the effect of errors on generic quantum circuits, rather than minimizing the effect of errors on self-inverting circuits. As demonstrated in Fig. 2.2, errors in self-inverting circuits may be substantially smaller than those in generic circuits because such circuits suppress coherent errors and implement a form of randomized dynamical decoupling [71].

2.6 Summary

In this chapter, we demonstrated how a gauge degree of freedom arises when representing a physical apparatus with a gate-set including a finite number of initial states, gates, and measurements, in the experimentally-relevant case where only measured probabilities are available. We discussed the difficulty in interpreting conventional quality metrics for individual gate-set elements, when the quantum processor was characterized in a non-gauge-invariant manner. Next, an alternative operational quality measure for a gate-set called mean variation error was proposed, along with a protocol to measure the MVE. Numerical simulations showed that MVE has different scaling for random vs. self-inverting circuits under certain noise models, suggesting that the restriction to self-inverting circuits in conventional randomized benchmarking procedures may lead to an underestimation of error. We also briefly sketched a method to improve experimental control that is not prone to the gauge ambiguity.

Chapter 3

SPAM error characterization through noise propagation

We have seen that a gauge freedom prevents one to obtain operationally meaningful quality metrics for individual gate-set elements, with no additional assumptions about the form of these elements. While gauge-independent measures can be constructed, few have appeared in the literature and none has been widely used by now. Therefore, from a characterization point of view, it is equally interesting to study the following parallel question: under what assumption can one eliminate the gauge freedom to obtain a unique gauge-dependent measure, and how can one estimate such measures experimentally?

In this chapter we study this problem, with a particular focus on obtaining unique gauge-dependent measures for state preparation and measurement (SPAM) processes. We look at the problem of SPAM error characterization from a new perspective, by demonstrating a sufficient condition for eliminating the gauge ambiguity, providing a protocol to estimate SPAM error rates, while taking gate imperfections into account. After quickly re-illustrating the problem caused by the gauge freedom, we show that distinguishing state preparation error from measurement error is possible in Section 3.1, if one assumes that there exists one qubit whose initial state is uncorrelated with other qubits in a processor. We provide a simple protocol from which the SPAM operators can be separately determined in Section 3.2, in the case where quantum gates are ideal, by propagating the error from these other qubits to the initially uncorrelated qubit. We then derive in Section 3.3 upper and lower bounds on the estimated parameters in case of non-ideal quantum gates, based on an error metric that can be estimated independently of SPAM, resolving the self-consistency problem. We performed our protocol on a publicly available 5-qubit quantum processor and obtained consistent results with a simulation.

3.1 SPAM characterization and gauge ambiguity

The ideal operations on a QPU generally include initializing the qubits in a state described by a density operator ρ, applying an arbitrary sequence of unitary gates, and making a final measurement described by a k-outcome POVM $M = \{M_1 \dots M_k\}$. We will assume here that the state which the QPU can be initialized to is unique. The implementation of each of these operations is imperfect due to a variety of noise processes. Noisy implementations of operations are denoted with an overset $\sim$ so that, for example, $\tilde{\rho}$ is the noisy implementation of ρ. To avoid overcrowding the text, we do not put additional $\sim$'s on *parameters* describing any operator: their meaning can usually be understood from context, and additional special notations on parameters will be defined prior to being used.

Denoting the N-qubit Pauli basis as $\mathbb{P}^N = \{I, X, Y, Z\}^{\otimes N}$, we can uniquely write

$$\begin{aligned} \tilde{\rho} &= \sum_{P \in \mathbb{P}^N} 2^{-N} \operatorname{Tr}\left[P^\dagger \tilde{\rho}\right] P \coloneqq \sum_{P \in \mathbb{P}^N} 2^{-N} s_P P, \\ \tilde{M}_i &= \sum_{P \in \mathbb{P}^N} 2^{-N} \operatorname{Tr}\left[P^\dagger \tilde{M}_i\right] P \coloneqq \sum_{P \in \mathbb{P}^N} 2^{-N} m_{P,i} P. \end{aligned} \tag{3.1}$$

In the Pauli-Liouville representation, $\tilde{\rho}$ is represented as a $4^N \times 1$ vector $|\tilde{\rho}\rangle\!\rangle$ with components $2^{-N/2} s_P$ [1], and similarly $\tilde{M}_i$ as $|\tilde{M}_i\rangle\!\rangle$. A linear map $\mathcal{G}$ is represented by a $4^N \times 4^N$ matrix $\Phi_{\mathcal{G}}$ with elements

$$(\Phi_{\mathcal{G}})_{P,Q} = 2^{-N} \operatorname{Tr}[P\ \mathcal{G}(Q)]. \tag{3.2}$$

Φ is called the Pauli transfer matrix, or PTM. In this picture, the result of mapping $\tilde{\mathcal{G}}$ to a state $\tilde{\rho}$ is given by a matrix multiplication: $|\tilde{\mathcal{G}}(\tilde{\rho})\rangle\!\rangle = \Phi_{\tilde{\mathcal{G}}} |\tilde{\rho}\rangle\!\rangle$. The probability $p(\tilde{\rho}, \tilde{\mathcal{G}}, \tilde{M})$ of an outcome corresponding to a POVM element $\tilde{M}$ given an input state $|\tilde{\mathcal{G}}(\tilde{\rho})\rangle\!\rangle$ can be computed by the inner product via $p(\tilde{\rho}, \tilde{\mathcal{G}}, \tilde{M}_i) = \langle\!\langle \tilde{M}_i | \Phi_{\tilde{\mathcal{G}}} | \tilde{\rho} \rangle\!\rangle$.

We now give an operational definition of what "SPAM errors" and "SPAM error rates" mean. Consider the experiment where one prepares the initial state and performs a measurement. The ideal and actual probabilities of obtaining outcome i are $\langle\!\langle M_i | \rho \rangle\!\rangle$ and $\langle\!\langle \tilde{M}_i | \tilde{\rho} \rangle\!\rangle$ respectively. We thus define the SPAM error to be the difference between these probabilities for this experiment, that is, as the *vector* $\boldsymbol{\delta}_{\text{SPAM}}$ with components

$$\boldsymbol{\delta}_{\text{SPAM},i}(\tilde{\rho}, \tilde{M}) \coloneqq \langle\!\langle M_i | \rho \rangle\!\rangle - \langle\!\langle \tilde{M}_i | \tilde{\rho} \rangle\!\rangle. \tag{3.3}$$

Next, the state preparation (SP) error vector $\boldsymbol{\delta}_{\text{SP}}$ and the measurement (M) error vector $\boldsymbol{\delta}_{\text{M}}$ are defined by $\boldsymbol{\delta}_{\text{SP},i} = \boldsymbol{\delta}_{\text{SPAM},i}(\tilde{\rho}, M)$ and $\boldsymbol{\delta}_{\text{M},i} = \boldsymbol{\delta}_{\text{SPAM},i}(\rho, \tilde{M})$, i.e., the SPAM error

[1] $2^{-N/2}$ serves as a normalization factor.

vector with the measurement/state preparation operators replaced by their ideal versions, respectively.

For a single qubit with a 2-outcome measurement, we can always write $\boldsymbol{\delta}_{\text{SPAM}} \coloneqq (1-\delta, \delta)^T$. In the usual case where $\rho = M_0 = |0\rangle\langle 0|$, δ corresponds to the probability of returning an outcome 1 when measuring ρ, which we will refer to as the SPAM *error rate*. Then, $\boldsymbol{\delta}_{\text{SP}}$ and $\boldsymbol{\delta}_{\text{M}}$ are also each characterized by a single parameter, which we will refer to as the SP-error rate δ_{SP}, and the M-error rate δ_{M}, respectively.

Ideally, one would like to obtain a full, unique description of $\tilde{\rho}$, $\tilde{M}$, and all possible control operations $\tilde{\mathcal{G}}_j$. We have already seen in the previous chapter that this is impossible due to a gauge freedom [72, 24, 73]: these operators are hidden and we can only infer their values from probabilities based on the Born rule. The choice of $(\tilde{\rho}, \tilde{\mathcal{G}}_j, \tilde{M})$ given a list of $p(\tilde{\rho}, \tilde{\mathcal{G}}_j, \tilde{M})$ is non-unique and are related by a gauge transformation

$$|\tilde{\rho}\rangle\!\rangle \to B|\tilde{\rho}\rangle\!\rangle, \ \langle\!\langle \tilde{M}_i| \to \langle\!\langle \tilde{M}_i|B^{-1}, \ \Phi_{\tilde{\mathcal{G}}_j} \to B\Phi_{\tilde{\mathcal{G}}_j}B^{-1}, \tag{3.4}$$

where B is an invertible matrix. This preserves all outcome probabilities when applied to all elements simultaneously, making the transformed set equally valid as the original set. On the other hand, most quality metrics for *individual components* (such as δ_{SP}, δ_{M}, or gate error rates) are *not* gauge-invariant. Since separate components in a QPU often require individual calibration in reality, having non-unique metrics is problematic because it becomes unclear whether an operation has improved (e.g., due to a change in control parameter) or not.

Next, we consider the weaker question of SPAM characterization, which boils down to estimating s_P and $m_{P,i}$. Previous studies on quantum state and detector tomography showed that $\tilde{\rho}$ or $\tilde{M}$ can be determined if the other is fully known. If we assume both to be in the most general form (satisfying only the physicality constraints that $\tilde{\rho}$ is a density matrix and $\tilde{M}$ is a POVM), but allow an arbitrary set of known, unitary gates $\mathcal{G}_j$, can we learn either $\tilde{\rho}$ or $\tilde{M}$? Interestingly, the answer is still negative. In particular, since a unitary gate $\mathcal{G}_j$ is trace-preserving and unital, it can be parametrized by

$$\Phi_{\mathcal{G}_j} = \begin{pmatrix} 1 & 0 \\ 0 & \phi_j \end{pmatrix} \tag{3.5}$$

where ϕ_j is a block matrix with components $\phi_{P,Q,j}$. All outcome probabilities are thus in the form

$$p(\tilde{\rho}, \mathcal{G}_j, \tilde{M}_i) = 2^{-N}(m_{I^{\otimes N},i} + \sum_{P,Q \in \mathbb{P}^N \setminus \{I^{\otimes N}\}} \phi_{P,Q,j} s_Q m_{P,i}) \tag{3.6}$$

where we followed the definitions in Eq. (3.1) [2]. Among all such equations which can be

[2] note that $s_{I^{\otimes N}} = 1$ by the unit trace constraint.

constructed, s_Q and $m_{P,i}$ always appear in a product form and cannot be separately solved for, assuming that ϕ only consists of constants. A gauge transformation (named "blame gauge" in [24]) of the form

$$s_Q \to x s_Q, \; m_{P,i} \to m_{P,i}/x \tag{3.7}$$

in the second term for some real number x will keep the equations unaltered [3]. While this transformation also needs to maintain the physicality constraints on $\tilde{\rho}$ and $\tilde{M}$, it is valid for most experimentally relevant cases [73]. Therefore, in addition to assuming ideal gates, one needs further assumptions about the structure of SPAM elements, and needs to design an effective operation that breaks this symmetry. Below, we will state a sufficient assumption, and present a protocol that achieves this by engineering Φ to depend upon the SPAM coefficients.

3.2 Protocol assuming ideal gates

To develop a straightforward protocol, we engineer simplified effective SPAM operators based on an averaging technique in [74], by removing undesired components in $\tilde{\rho}$ and $\tilde{M}$. From now on, the qubit or system of qubits whose SPAM operators we would like to know will be called the target qubit (system), and we will use the subscript t to indicate parameters of the target qubit (system). For now we assume all quantum gates to be ideal, and will relax this later. Consider a single qubit initialized to $\tilde{\rho}$ and has a 2-outcome POVM $\tilde{M} = \{\tilde{M}_0, \tilde{M}_1 = I - \tilde{M}_0\}$, parametrized by

$$|\tilde{\rho}\rangle\!\rangle = \frac{(1, s_X, s_Y, s_Z)^T}{\sqrt{2}}, \langle\!\langle \tilde{M}_0| = \frac{(m_I, m_X, m_Y, m_Z)}{\sqrt{2}} \tag{3.8}$$

We will assume that the ideal state and measurement are $\rho = M_0 = |0\rangle\langle 0|$, corresponding to $s_X = s_Y = m_X = m_Y = 0$, and $s_Z = m_Z = 1$ in Eq. (3.8). In reality these parameters deviate from the ideal, but we can use the following technique to eliminate some undesired components. By linearity of quantum operations and probabilities, applying two Pauli gates from the set $\{I, Z\}$ immediately after state preparation and before measurement, and averaging over outputs from all possible circuits would set $s_X = s_Y = m_X = m_Y = 0$ (this is similar to the phase cycling technique commonly used in NMR spectroscopy to suppress spurious signals [75]). To fix m_I, we apply an additional gate from $\{I, X\}$ immediately before the measurement, and relabel the outcome when we apply an X so that the

[3] Importantly, it also keeps the gate intact, because the matrix B here commutes with $\Phi_{\mathcal{G}_j}$ for all unitary gates $\mathcal{G}_j$

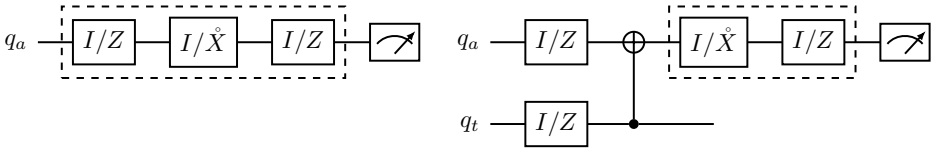

Figure 3.1: Circuits for determining the coefficient $s_{Z,t}$ on the target qubit q_t, assuming ideal gates. The one on the left/right gives α_a and β_t, respectively. Combinations of SPAM averaging gates are run as separate experiments (the outcome is flipped classically when the first M-averaging gate is X, indicated by an overhead circle). Adjacent single-qubit gates (grouped by the dashed box) are logically compiled to a single gate when running the circuits.

outcome 0 corresponds to the POVM element $\tilde{M}_1$ and *vice versa*. We label this with an overhead circle in Fig. 3.1 and Fig. 3.2. Averaging the results from these two circuits (with an I or X averaging gate) effectively sets $m_I = 1$. Combining the above, the problem is now reduced to finding $s_{Z,t}$ and $m_{Z,t}$ on the target qubit q_t.

A few words regarding SPAM averaging shall take place before we proceed. While the SPAM operators after averaging deviate from the original ones, they do represent the ones that actually enter the circuit, if SPAM averaging is consistently applied in all future circuits. Since our protocol estimates exactly the parameters in this averaged model, they will predict the correct experimental outputs for future circuits as well.

Now, we provide a simple protocol that estimates $s_{Z,t}$ and $m_{Z,t}$. We assume that there exists an ancillary qubit q_a which can be prepared and measured *independently*, i.e. is described by two independent but unknown coefficients $s_{Z,a}$ and $m_{Z,a}$ (after applying the same SPAM averaging). If we apply a CNOT gate controlled on q_t and targeted on q_a (see Fig. 3.1), which we will call $\mathcal{C}_{t,a}$, the PTM on q_a can be calculated to be a diagonal matrix

$$\Phi = \mathrm{diag}(1, 1, s_{Z,t}, s_{Z,t}), \tag{3.9}$$

which depends upon $s_{Z,t}$ as desired. In other words, the entangling CNOT gate *propagates SP parameters* of q_t to q_a. Because measurement on q_a is unaffected by the gate, we can learn $s_{Z,t}$ as follows: define α and β as the expectation values $\langle \tilde{M}_0 - \tilde{M}_1 \rangle_a$ on q_a in the *absence* and *presence* of the CNOT gate, respectively. A direct calculation (see Eq. (3.8) and Eq. (3.9)) shows

$$\alpha_a = s_{Z,a} m_{Z,a}, \;\; \beta_t = s_{Z,t} s_{Z,a} m_{Z,a}, \tag{3.10}$$

which gives $s_{Z,t} = \beta_t / \alpha_a$. $m_{Z,t}$ can then be determined by a separate experiment that

measures α_t on q_t with

$$m_{Z,t} = \alpha_t / s_{Z,t}. \tag{3.11}$$

Note that the subscript for α refers to the qubit *being measured*, while for β it refers to the qubit whose *parameter is being propagated*. The SP- and M-error rates can then be computed as

$$\delta_{\mathrm{SP},t} = (1 - s_{Z,t})/2, \ \delta_{\mathrm{M},t} = (1 - m_{Z,t})/2. \tag{3.12}$$

We can generalize this idea to measuring parameters of an N-qubit system, as long as we assume that there exists *one* additional ancilla q_a that can be prepared and measured independently. This is thus a sufficient condition for breaking the gauge symmetry. Here, α_a is estimated using the same circuit as the one on the left of Fig. 3.1. To estimate $\beta_{P,t}$ where P labels the Pauli components of the unknown initial state, the CNOT in the 1-qubit case is generalized to $\mathcal{U}_P = (\mathcal{H} \otimes \mathcal{I})\mathcal{C}_P(\mathcal{H} \otimes \mathcal{I})$, where $\mathcal{C}_P$ corresponds to the unitary $|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes P$, $P = P_1 \otimes \cdots \otimes P_N$, $P_i \in \{I, Z\}$, and $\mathcal{H}$ corresponds to a Hadamard gate on q_a. $\mathcal{C}_P$ is controlled on q_a and targeted on the N qubit system (see Fig. 3.2). As shown in Appendix A.1, the effect on q_a is identical to Eq. (3.9) with $s_{Z,t}$ replaced by $s_{P,t}$. Repeating for all possible P's will fully determine ρ_t, allowing one to determine the POVMs by detector tomography. This is summarized in Algorithm 1. Note that the number of circuits to average over grows exponentially with the size of q_t. For large systems, one should instead randomly sample from the space of all SPAM averaging gates. Additionally, the controlled-P gate requires $\mathcal{O}(N)$ controlled-Z gates and has depth $\mathcal{O}(N)$ for the worst case, but can be achieved by 2 all-to-all Mølmer-Sørensen gates [76, 77].

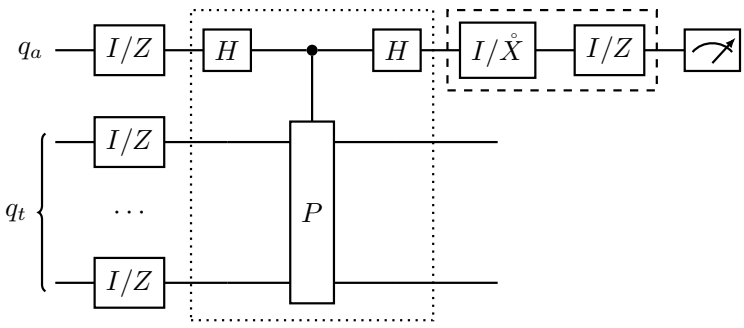

Figure 3.2: The circuit for estimating $\beta_{P,t}$ of an N-qubit system q_t, assuming ideal gates. The dotted box indicate the propagating cycle $\mathcal{U}_P$.

We would also like to point out that small modifications to Algorithm 1 would allow one to obtain *all* components of the SPAM operators in principle. For example, in the 1-qubit

case, one could average over $\{I, X\}$ instead of $\{I, Z\}$ to obtain $|\tilde{\rho}\rangle\!\rangle \sim (1, s_{X,t}, 0, 0)^T$. Then performing a $Y(-\pi/2)$ rotation would result in $|\tilde{\rho}\rangle\!\rangle \sim (1, 0, 0, s_{X,t})^T$, so that s_X can then be determined in exactly the same way as $s_{Z,t}$ before. But this is unnecessary if SPAM averaging is consistently applied for all future circuits, as we have discussed previously.

Algorithm 1 Estimating the SP- and M-operators of an N-qubit system, assuming ideal gates

1: Choose an ancillary qubit q_a
2: Measure q_a (see left of Fig. 3.1) and record the result α_a
3: **for** each P in $\{I, Z\}^{\otimes N}$ **do**
4: Apply the circuits in Fig. 3.2 with gate P, measure q_a and record the result $\beta_{P,t}$
5: Calculate $s_{P,t} = \beta_{P,t}/\alpha_a$
6: **end for**
7: Apply detector tomography on q_t to estimate $m_{P,j}$ for each POVM element $\tilde{M}_j$, where $j = 1...k$ for a k-outcome POVM.

3.3 Protocol with imperfect gates

We now take into account gate imperfections. The effects of gate errors must be treated in a way that does not rely on any prior information on SPAM, since they are assumed unknown. This prohibits using protocols like process tomography to extract the full effect of the gate in question, and substitute to replace Φ in Eq. (3.9). Protocols that estimate gate error strengths independently of SPAM offer a solution to the problem. Here we utilize the recently proposed cycle benchmarking (CB) [78] procedure. CB estimates the process infidelity of a composite cycle (consisting of a round of the original gates $\tilde{\mathcal{G}}$ composed with a round of "dressing" gates $\tilde{\mathcal{D}}$), averaged over all Pauli dressing gates, namely

$$r_{\mathrm{CB}}(\tilde{\mathcal{G}}, \mathcal{G}) := \sum_{\mathcal{D} \in \{\mathcal{I}, \mathcal{X}, \mathcal{Y}, \mathcal{Z}\}^{\otimes N}} 4^{-N} r_p(\tilde{\mathcal{G}}\tilde{\mathcal{D}}, \mathcal{G}\mathcal{D}), \tag{3.13}$$

where the process infidelity is

$$r_p(\tilde{\mathcal{G}}, \mathcal{G}) := 1 - 4^{-N} \sum_{P \in \mathbb{P}^N} 2^{-N} \operatorname{Tr}\Big[\mathcal{G}(P)\tilde{\mathcal{G}}(P)\Big]. \tag{3.14}$$

The figure r_{CB} is relevant when a quantum computation task is used in conjunction with a noise-tailoring procedure called randomized compiling (RC) [74]. Here, random twirling

gates are inserted into the original circuit, such that the logical circuit is preserved. Uniformly averaging over all twirling gates turns the error of a composite cycle into a Stochastic Pauli channel $\mathcal{P}$: $\mathcal{P}(\rho) = \sum_{P\in\mathbb{P}^N} c_P P\rho P^\dagger$, where c_P's form a probability distribution. The error rate of $\mathcal{P}$ is then precisely characterized by $r_{\mathrm{CB}}(\tilde{\mathcal{G}}_j, \mathcal{G}_j)$ [78]. This twirling is exact under the standard assumption that errors on twirling gates are gate-independent (which we assume throughout), and has a relatively small correction when gate dependence is present [74]. For simplicity we also make the standard assumption that 1-qubit gates have a 1-qubit error channel, however we conjecture that this can be relaxed.

$\delta_{\mathrm{SP},t}$, lower	$\delta_{\mathrm{SP},t}$, upper	$\delta_{\mathrm{M},t}$, lower	$\delta_{\mathrm{M},t}$, upper
$\frac{1}{2} - \frac{\beta_t + 2r_{t,a}}{2\alpha_a}$	$\frac{1}{2} - \frac{\beta_t - 2r_{t,a}}{2\alpha_a}$	$\frac{1}{2} - \frac{\alpha_t\alpha_a}{2\beta_t - 4r_{t,a}}$	$\frac{1}{2} - \frac{\alpha_t\alpha_a}{2\beta_t + 4r_{t,a}}$

Table 3.1: Upper and lower bounds for 1-qubit SPAM error rates (Eq. (3.12)) on a target qubit q_t. α and β are defined in Eq. (3.10). $r_{t,a}$ is shorthand for $r_{\mathrm{CB}}(\tilde{\mathcal{C}}_{t,a}, \mathcal{C}_{t,a})$.

Let's now denote the parameters that *would have been obtained with an ideal propagating cycle* with a superscript ic (i.e., ideal cycle). These are the actual parameters describing our unknown SPAM operators, which are not affected by the imperfect gates. On the other hand, the ones that are actually obtained in experiments will be denoted as normal letters. We will show in Appendix A.2 that $\beta^{\mathrm{ic}}_{P,t}$ can be bounded using the measured $\beta_{P,t}$ and $r_{\mathrm{CB}}(\tilde{\mathcal{U}}_P, \mathcal{U}_P)$ as:

$$\beta^{\mathrm{ic}}_{P,t} \in [\beta_{P,t} - 2r_{\mathrm{CB}}(\tilde{\mathcal{U}}_P, \mathcal{U}_P), \beta_{P,t} + 2r_{\mathrm{CB}}(\tilde{\mathcal{U}}_P, \mathcal{U}_P)], \tag{3.15}$$

which holds independently of the dimension of q_t. Since $s_{P,t} = \beta_{P,t}/\alpha_a$, and because α_a does not involve gates with unknown effects, we see that

$$s^{\mathrm{ic}}_{P,t} \in [\frac{\beta_{P,t} - 2r_{\mathrm{CB}}(\tilde{\mathcal{U}}_P, \mathcal{U}_P)}{\alpha_a}, \frac{\beta_{P,t} + 2r_{\mathrm{CB}}(\tilde{\mathcal{U}}_P, \mathcal{U}_P)}{\alpha_a}]. \tag{3.16}$$

Repeating for all values of P would give a bound on each parameter $s^{\mathrm{ic}}_{P,t}$ of the estimated initial state $\tilde{\rho}$. Recall from our previous definition that the i-th component of $\boldsymbol{\delta}_{\mathrm{SP}}$ is given by

$$\boldsymbol{\delta}_{\mathrm{SP},i} = \langle\langle M_i|(|\rho\rangle\rangle - |\tilde{\rho}\rangle\rangle) = \sum_P m_{P,i}(s^{\mathrm{ideal}}_P - s_P), \tag{3.17}$$

where we used the superscript "ideal" to represent the ideal parameters of ρ. This is a linear function of s_P, whose bounds are given by Eq. (3.16). Therefore the upper and lower

bounds for $\boldsymbol{\delta}_{\mathrm{SP},i}$ can simply be obtained by optimizing each term in the sum, resulting in

$$\begin{aligned}\boldsymbol{\delta}_{\mathrm{SP},i,\ \mathrm{lower}} &= \sum_P m_{P,i} s_P^{\mathrm{ideal}} - \sum_P m_{P,i} s_{\mathrm{sgn}(m_{P,i})},\\ \boldsymbol{\delta}_{\mathrm{SP},i,\ \mathrm{upper}} &= \sum_P m_{P,i} s_P^{\mathrm{ideal}} - \sum_P m_{P,i} s_{-\mathrm{sgn}(m_{P,i})},\end{aligned} \tag{3.18}$$

where sgn is the sign function, and we have used the shorthand s_- and s_+ to represent the lower and upper bounds in Eq. (3.16).

The bounds for measurement parameters is more complicated, since one would need to perform measurement tomography based on the learned initial state, and different tomography approaches will lead to different bounds. But, the principles behind all approaches will be similar. Here we demonstrate with the simplest case of a linear inversion (LI) tomography. In LI detector tomography, one prepares an informationally-complete set of initial states $\tilde{\rho}_1 \dots \tilde{\rho}_{4^N}$, which we assume have all been characterized using our procedure. For qubit measurements in the computational basis, the unknown POVM elements will correspond to the outcomes $|0\rangle^{\otimes N} \dots |1\rangle^{\otimes N}$, so there are a total of 2^N of them. Arrange the column vectors $|\tilde{\rho}_j\rangle\!\rangle$ into a $4^N \times 4^N$ matrix S. Arrange the vectorized POVM elements $|\tilde{M}_i\rangle\!\rangle$ into a $2^N \times 4^N$ matrix R. One then measures each basis state $\tilde{\rho}_i$ and record the data matrix with components $D_{i,j} = \langle\!\langle \tilde{M}_i | \tilde{\rho}_j \rangle\!\rangle$. This gives the matrix relation:

$$R \cdot S = D \tag{3.19}$$

which can be inverted as $R = D \cdot S^{-1}$ to solve for the unknown matrix R. In the absence of gate error and measurement shot noise, this results in a noiseless reconstruction of the POVM elements $\tilde{M}_i$. In the presence of gate errors when measuring the states $\tilde{\rho}_j$, we have learned from Eq. (3.18) that each component of the vector $|\tilde{\rho}_j\rangle\!\rangle$ is bounded. This uncertainty translates into uncertainties in $\tilde{M}_i$ through the matrix inverse, S^{-1}. In this case, each component of the resulted $\langle\!\langle \tilde{M}_i|$ is a (potentially highly nonlinear) function of the components of $|\tilde{\rho}_j\rangle\!\rangle$. Nonetheless, max and min values of $\langle\!\langle \tilde{M}_i|_k$ is guaranteed by the extreme value theorem, and can be found using numerical programs such as scipy. From this, bounds on components of $\boldsymbol{\delta}_M$ can be derived in the same way as what we did for $\boldsymbol{\delta}_{SP}$.

The situation becomes particularly simple for 1 qubit with a 2-outcome measurement, along with SPAM averaging. In this case there is only one unknown parameter s_Z for $\tilde{\rho}$ and another one m_Z for $\tilde{M}_0$ ($\tilde{M}_1$ is fixed by $\tilde{M}_0$). The bound for s_Z is given directly in Eq. (3.16). Since m_Z and s_Z are inversely proportional (see Eq. (3.11)), the maximum of s_Z gives the minimum of m_Z, and vice versa. We then use Eq. (3.12) to convert to bounds on the error rates δ_{SP} and δ_{M}. These are summarized in Table 3.1. Intuitively, a

smaller gate error corresponds to a narrower range, and the region restores the previous point estimate (Eq. (3.10)) in the limit of perfect gates.

We incorporate gate error effects into Algorithm 1 by proposing a simple procedure to separately estimate the single qubit SP- and M-errors in a QPU. For each qubit i, label it as the "target" (t) and find an "ancilla" (a) such that a CNOT gate $\mathcal{C}_{t,a}$ is allowed by the QPU's connectivity. We then run the 1-qubit protocol to estimate the SPAM parameters in Table 3.1. Repeatedly identifying each of the N qubits as the "target" gives a single-qubit SPAM characterization of the full device. This is summarized in Algorithm 2. While providing experimentally relevant single-qubit error rates for the full system, the protocol has an overhead that only scales linearly with the system size, making it a practical tool for many scenarios.

Algorithm 2 Estimating single qubit SP- and M-error rates on an N-qubit device

1: **for** each of the N qubits **do**
2: Label it as q_t; choose an ancilla q_a where $\mathcal{C}_{t,a}$ is allowed
3: Run the 1-qubit protocol to estimate α_a, α_t, and β_t
4: Use cycle benchmarking to estimate $r_{\mathrm{CB}}(\tilde{\mathcal{C}}_{t,a}, \mathcal{C}_{t,a})$
5: Compute a regional estimate on $\delta_{\mathrm{SP},t}$ and $\delta_{\mathrm{M},t}$ on q_t according to Table 3.1.
6: **end for**

3.4 Experimental results

We performed the above protocol on a publicly available 5-qubit QPU (ibmq-santiago [79]) to estimate δ_{SP} and δ_{M} on each qubit. For each target qubit, we chose the ancilla to be the one connected to it with the lowest error CNOT. The presented experiment was performed on the ibmq-santiago machine on Jan 28th, 2021. The specifications for each different type of experiment are summarized as follows:

1. Each α was obtained by averaging exhaustively over the 4 possible cases, corresponding to the cases where the compiled SPAM randomizing gates belong to $\{I, X, Y, Z\}$. Each circuit is sampled with 8192 measurement shots.

2. Each β is estimated by averaging over 60 randomly compiled circuits (with a total of 256). Each circuit is sampled with $k = 1024$ measurement shots. For practical convenience, we performed sampling with replacement because the estimate precision is sufficiently high; if higher precision is desired, one may switch to sampling without replacement or even exhaustive sampling as in estimating α.

3. The infidelity $r_{t,a}$ for each CNOT was estimated using cycle benchmarking by repeating the CNOT cycle $\{4, 84\}$ times, and averaging over all 16 Pauli decay strings, each by sampling 30 random circuits with 128 shots. The specific choice of CNOT gates used in the experiment were: $\mathcal{C}_{0,1},\ \mathcal{C}_{1,0},\ \mathcal{C}_{2,1},\ \mathcal{C}_{3,2},\ \mathcal{C}_{4,3}$.

We used the TrueQ [80] software to generate circuits, submit to the IBM-Q server, and perform data analysis. The range of possible error rates (i.e., between the upper and lower bounds in Table 3.1) is shown in shaded regions for each estimated δ_{SP} or δ_{M}. The 95% confidence intervals (CIs) of the upper and lower bounds are shown as error bars; their derivations are slightly involved and can be found in Appendix A.3. Any region below 0 is discarded due to the physicality constraint that error rates are positive by definition (Eq. (3.12)).

Since IBM-Q does not provide separate SP/M error rates for us to compare with, we simulated the same circuits to gauge the validity of our experiment. For each qubit, we individually add a state preparation error to it by replacing the ideal initial state $|0\rangle$ with a density matrix $\rho = \mathrm{diag}(1-\epsilon_{\mathrm{SP}}, \epsilon_{\mathrm{SP}})$. We add a measurement error by classically flipping the outcome (symmetrically, from 0 to 1 and from 1 to 0) with probability ϵ_{M}. Gate errors are simulated using a simple T_1, T_2 relaxation model: for each clock cycle in the circuit, we apply a noise process to each qubit defined by the following Choi matrix

$$C = \begin{pmatrix} 1 & 0 & 0 & e^{-t/T_2} \\ 0 & e^{-t/T_1} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ e^{-t/T_2} & 0 & 0 & 1 - e^{-t/T_1} \end{pmatrix}. \tag{3.20}$$

We modeled noisy gates by assuming a simple $T_1 + T_2$ relaxation model, and using the relaxation time and gate time obtained from the provider. The T_1 and T_2 relaxation times are obtained from the provider and are tabulated in Table 3.2. Single-qubit gates all have the same (35.6ns) gate time, according to the provider. The gate time of the CNOT gates used in the experiments are tabulated in Table 3.3.

Qubit	0	1	2	3	4
T_1 (μs)	75.9	134.667	120.21	137.32	100.68
T_2 (μs)	140.18	96.34	87.25	94.1	133.12

Table 3.2: T_1 and T_2 for individual qubits on IBMQ-santiago.

By manually adjusting the relative magnitudes between $\delta_{\mathrm{SP},i}$ and $\delta_{\mathrm{M},i}$ (indicated by the green stars) while fixing the total SPAM error to the measured value, we obtained a similar

Gate	$\mathcal{C}_{0,1}$	$\mathcal{C}_{1,0}$	$\mathcal{C}_{2,1}$	$\mathcal{C}_{3,2}$	$\mathcal{C}_{4,3}$
Gate time (ns)	526.2	561.7	568.9	412.4	341.3

Table 3.3: Gate time of each CNOT gate used in the experiment.

behavior compared to the data, strengthening our claim about noise on the physical device (see the bottom half of Fig. 3.3). Our result shows that state preparation contributes more to the total SPAM error on qubit 4, while not conclusively on the other qubits. Better distinguishability can be achieved if higher quality gates are available, as shown by the darker shaded regions in the bottom of Fig. 3.3, which represent the same bounds (error bars omitted) in Table 3.1 with all gate times reduced to 1/5 of their original values.

3.5 Summary

In this chapter, we proposed a method to characterize state preparation and measurement errors independently on a QPU. In the case where quantum gates are ideal, our method returns the exact state preparation and measurement errors, resolving the gauge ambiguity issue raised in the previous chapter. In the case where quantum gates are imperfect, by utilizing randomized compiling and cycle benchmarking techniques, we derived upper and lower bounds for the estimated SPAM errors in terms of gate error rates that can be measured independently of SPAM. We demonstrated our protocol on a publicly available QPU and observed consistent results between the data and a computer simulation. We believe this protocol can be a valuable tool for benchmarking near-term quantum devices, in complement to the existing protocols that estimate errors on quantum gates.

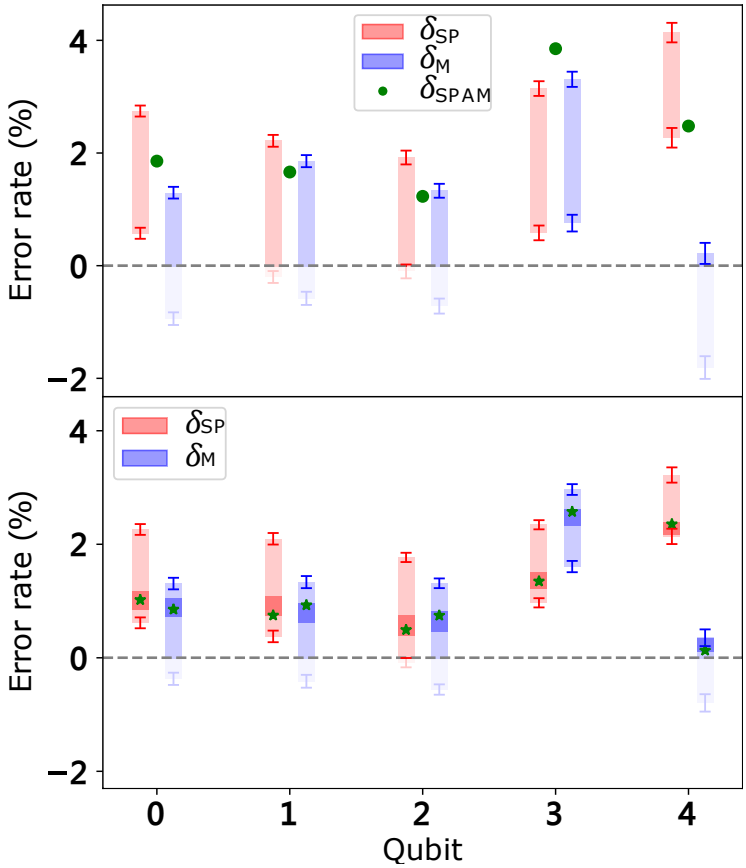

Figure 3.3: Top: estimated single qubit state preparation (in red) and measurement (in blue) error rates on ibmq-santiago QPU. Shaded regions represent the range of error rates consistent with the measured gate error, and error bars are 95% CIs for the endpoints. Green dots indicate measured total SPAM error δ_{SPAM}. All estimates are cut-off below 0%. Bottom: a simulation assuming T_1, T_2 relaxation gate errors on 2-qubit gates and ideal 1-qubit gates, using the device's specifications. Green stars mark the magnitudes of SP- and M-errors used in the simulation, which combine to δ_{SPAM} on each qubit. Darker regions mark the same bounds when all gate times are reduced to 1/5 of their original values.

Chapter 4

SPAM error characterization through algorithmic cooling

In the previous chapter, we showed that the SPAM gauge ambiguity can be eliminated by assuming the existence of one independent ancillary qubit, plus the application of untiary gates. An alternative possibility to eliminate the gauge ambiguity is by applying non-unital operations. For example, the 1-qubit amplitude damping channel is non-unital and thus is not described by Eq. (3.5). Indeed, if one can fully relax the qubit to $|0\rangle$, then all errors coming from measuring the state comes from the measurement, providing a method to directly estimate ϵ_{M}. This can be seen as a novel application of algorithmic cooling (AC) techniques which aim to reduce state preparation errors, and is the focus of this chapter.

In this chapter, we show that an extended form of AC, which we call measurement-based algorithmic cooling (MBAC), can used to distinguish state preparation and measurement errors. Moreover, it can also be used to reduce the initial state error by construction. In Section 4.1 we review the history of AC. In Section 4.2 we give a review of conventional AC protocols, which motivate the new MBAC protocols. In Section 4.3 we describe k-qubit MBAC assuming ideal measurements on the ancillary qubits, and compare it with the optimal reversible AC scheme. In Section 4.4 we describe a procedure to separately characterize SPAM errors using MBAC. In Section 4.5 we relax the assumption made in Section 4.3 and analyze MBAC when the ancillary qubits have finite measurement errors, and derive a lower bound for its performance. In Section 4.6 we study the number of trials needed to cool down a target qubit by a desired multiplicative factor, and illustrate the practical usefulness of MBAC despite its probabilistic nature.

4.1 Backgrounds

The original studies of AC were largely motivated by seeking ways to increase the effective polarization of spins in a magnetic field. Ole W. Sørensen was the first to observe [81] that certain unitary dynamics can lead to partial polarization transfer from a subset of spins to the others in nuclear magnetic resonance (NMR) spectroscopy. He then derived the so-called "universal bound on spin dynamics" [82], giving the maximum possible polarization boost in such scenarios. Later, Schulman and Vazirani [83] studied a more abstract version of this problem, and derived the maximum fraction of pure spins that can be extracted from a given initial temperature. They considered a simplified model which ignored relaxation of the spins, and based their studies purely on entropy transfer theories. This is usually referred to in the literature as the "reversible scheme", because it only rely on reversible, unitary operations on the system which is considered as closed.

Later on, a new family of AC protocols were proposed under the name "heat bath algorithmic cooling" (HBAC) [84, 85, 86]. The assumption is slightly different than in reversible AC: instead of requiring all spins to have infinitely long relaxation times, it was assumed that a subset of "computational spins" have a much longer relaxation time than the rest, so-called "reset spins". These assumptions are more realistic than the fully-closed one in the reversible setting. HBAC performs alternate rounds of "entropy compression - reset" subroutines, where entropy is transferred from the computational spins to the reset spins during the compression stage, and eventually released into a thermal bath during the reset stage. It has been shown that by considering the system to be semi-open, one is able to achieve further entropy reduction in the computational spins compared to the reversible scheme. Moreover, HBAC protocols usually involves simpler controls, and therefore has been experimentally tested in different platforms [87, 88, 89, 90, 91]. More recently, AC algorithms utilizing additional resources such as qubit-environment interactions [92] or internal interactions [93] have been proposed. There have also been efforts on generalizing the reset stage and find optimal thermalization operations within the set of all "thermal operations" [94] (defined later in Chapter 5).

Meanwhile, there has been few efforts in studying the role of measurement for the task of cooling. There may be historical reasons behind this lack of attention, presumably rooted in how measurements are perceived in earlier days of quantum computing. In ensemble quantum processors, the type of operations permitted is typically "global" or "coarse", in the sense that operations cannot be applied onto individual qubits, but only to the same qubits in each of the many parallel quantum computers. Consequently, one is limited to weak ensemble measurements that estimate expectation values of certain operators.

On the other hand, quantum processors nowadays are less restrictive: if projective

measurements are allowed, a trivial way to cool a qubit is to measure it. If the outcome is 0, the system is in the $|0\rangle$ state immediately after the measurement, corresponding to a zero temperature. If the outcome is 1, then one can perform a Pauli X gate to flip it to $|0\rangle$, achieving the same result. There are two major issues with this approach: first, in some existing gate-based quantum computing architectures, the measurement process only takes place at the final stage of the computation. No computation can proceed on the measured qubit and a new round of computation must be initiated. In other words, the measurement "destroys" the system of interest, so this naive cooling method is not so useful for doing further computations. Second, this method works only if measurement errors are absent, and there is a perfect correspondence between the measurement result and the post-measurement state; if this is not true, we cannot safely deduce the post-measurement state. In this chapter, we will study how one may utilize measurement to perform cooling, leading to measurement-based AC (MBAC), while circumventing these issues.

4.2 Brief recap on conventional AC

We first describe AC protocols, notations and goals, in a way that leads more naturally to MBAC. Below, we will denote the Pauli matrices X, Y, Z by $\sigma_x, \sigma_y, \sigma_z$, and the 2×2 identity matrix by σ_I. The ideal state preparation step should initialize a single qubit to the $|0\rangle$ state. We assume that due to imperfect state preparation processes, a bit-flip (equivalently, σ_x) error occurs with probability δ, so that the actual initial state is described by the following mixed state density matrix:

$$\rho = (1-\delta)\,|0\rangle\langle 0| + \delta\,|1\rangle\langle 1|\,. \tag{4.1}$$

This represents a state preparation error on the quantum processor. We will assume that $0 \leq \delta < 1/2$, where $\delta = 0$ corresponds to the pure state $|0\rangle$ and $\delta \to 1/2$ corresponds to the completely mixed state.

Algorithmic cooling procedures are designed to reduce δ on the target qubit towards 0. The term "cooling" comes from viewing ρ as a thermal state,

$$\rho = \frac{1}{Z} e^{-\beta \hat{H}} = \frac{1}{Z} \begin{pmatrix} e^{\frac{\hbar\omega}{2k_B T}} & 0 \\ 0 & e^{-\frac{\hbar\omega}{2k_B T}} \end{pmatrix} \tag{4.2}$$

where $\hat{H} = -\frac{1}{2}\hbar\omega\sigma_z$ is the bare qubit Hamiltonian, k_B is the Boltzmann constant, T is an effective temperature, and Z is the partition function so that ρ is normalized. From

Eq. (4.1) and Eq. (4.2) we can identify $\delta = (e^{\frac{\hbar\omega}{k_B T}} + 1)^{-1}$. Solving for T gives $k_B T = \hbar\omega / \log\left(\frac{1}{\delta} - 1\right)$, so that T is closer to 0 when δ is closer to 0, for a fixed ω. Therefore, reducing its effective temperature increases its probability of being in the ground state.

We first review a primitive type of AC scheme using a minimal example with $m = 3$ qubits [85], which inspires our measurement-based protocol. We will review the reversible AC later in Section 4.3 when we compare MBAC with conventional AC.

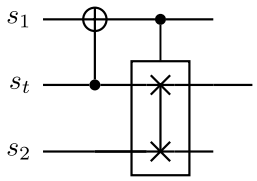

Figure 4.1: The circuit for 3-qubit BCS aiming to cool down the target spin s_t. Within the box is a SWAP gate, which is controlled by spin s_1.

Let us now assume that there is a target spin to be cooled called s_t, and two ancillary spins s_1 and s_2. Starting with the system in a state $\rho_1 \otimes \rho_t \otimes \rho_2$ where all ρ's are described by Eq. (4.1) with the same value of δ, we can apply a basic compression subroutine (BCS) to reduce δ on s_t (see Fig. 4.1). The BCS involves essentially two steps: first, a CNOT operation controlled by s_t and targeted on s_1; second, a controlled-SWAP (CSWAP) operation controlled by s_1, and targeted on s_t and s_2. To see why this protocol works, first consider the evolution s_1 and s_t after the CNOT gate. We replace the noisy states of s_t and s_1 by assuming an error model is applied onto pure qubit states. The error model is such that a σ_x channel occurs independently on each qubit after the preparation. Referring to Fig. 4.2, this corresponds to the following four possibilities where U equals to

1. $\sigma_I \otimes \sigma_I,\ p = (1-\delta)^2$;
2. $\sigma_I \otimes \sigma_x,\ p = \delta(1-\delta)$;
3. $\sigma_x \otimes \sigma_I,\ p = \delta(1-\delta)$;
4. $\sigma_x \otimes \sigma_x,\ p = \delta^2$;

with the corresponding probabilities listed after each case. In the above the left error-operator is applied onto s_1 and the right operator onto s_t. From Fig. 4.2 and the mapping rule of Pauli operators under the CNOT gate, the four cases can be written equivalently where U' equals to

1. $\sigma_I \otimes \sigma_I,\ p = (1-\delta)^2$;
2. $\sigma_x \otimes \sigma_x,\ p = \delta(1-\delta)$;
3. $\sigma_x \otimes \sigma_I,\ p = \delta(1-\delta)$;
4. $\sigma_I \otimes \sigma_x,\ p = \delta^2$.

Since s_1 and s_t both start from $|0\rangle$, the CNOT does nothing and can be removed. Thus the output states are

1. $|0\rangle \otimes |0\rangle\,,\ p = (1-\delta)^2$;
2. $|1\rangle \otimes |1\rangle\,,\ p = \delta(1-\delta)$;
3. $|1\rangle \otimes |0\rangle\,,\ p = \delta(1-\delta)$;
4. $|0\rangle \otimes |1\rangle\,,\ p = \delta^2$.

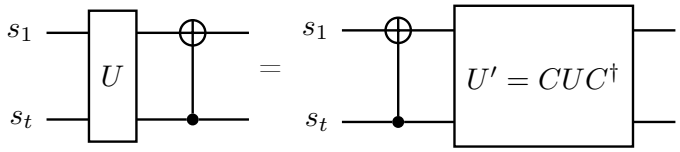

Figure 4.2: Two equivalent circuits where the order of two unitary gates are exchanged, and the second gate on the RHS is replaced by the original gate conjugated by the first gate. C stands for the CNOT gate in this case.

It is now clear that s_1 and s_t become correlated, since the probability of them being in the same state (case 1 and 2) is higher than being in different states (case 3 and 4), for all $0 \leq \delta < 1/2$. When s_1 is in $|0\rangle$, s_t is more likely to be in $|0\rangle$ (case 1) than in $|1\rangle$ (case 4), so we keep this purified spin s_t in the second step of Fig. 4.1. When s_1 is in $|1\rangle$, s_t is equally likely to be in $|0\rangle$ or $|1\rangle$ (case 2 and 3), correspond to the completely mixed state and an effective temperature $T \to \infty$. In this case s_t has been heated up, so we can restore its original state by swapping s_t and s_2 in the second step of Fig. 4.1. Overall, the density matrix of s_t becomes closer to $|0\rangle$ at the output end.

We can calculate the exact reduced state of s_t at the output, by taking an average over the aforementioned 4 cases (after applying the CSWAP step). Denote the probability of

being in $|1\rangle$ for s_t after a BCS round as δ_t'. For cases 1 and 4, s_1 is in $|0\rangle$ so no SWAP gate is applied, and the probability of s_t being in $|1\rangle$ is 0 and 1, respectively. For cases 2 and 3, s_t is swapped with s_2, so the probability of s_t being in $|1\rangle$ is restored to δ in both cases. Averaging over these 4 cases yields

$$\delta_t' = \delta^2 + 2\delta^2(1-\delta) = 3\delta^2 - 2\delta^3 \leq \delta, \ \forall\ 0 \leq \delta < \frac{1}{2}. \tag{4.3}$$

For small δ, it is reduced to order $\mathcal{O}(\delta^2)$. In the more general case where the initial δ's on each spin can be different, we have

$$\delta_t' = \delta_1\delta_t + \delta_2\delta_t + \delta_1\delta_2 - 2\delta_1\delta_t\delta_2. \tag{4.4}$$

If we further assume that the ancillas s_1 and s_2 can relax much faster to their original states than s_t, so that we can effectively repeat this BCS round (where s_t now has error δ_t'), then s_t can be further purified. In the limit of performing infinitely many rounds, s_t arrives at the steady state, whose error can be calculated by setting $\delta_t' = \delta_t = \delta_t^\infty$ in Eq. (4.4), which gives

$$\delta_t^\infty = \frac{\delta_1\delta_2}{1 - \delta_1 - \delta_2 + 2\delta_1\delta_2}. \tag{4.5}$$

4.3 MBAC

Based on the previous analysis, one can see that the second CSWAP gate (and consequently, s_2 as well) is not needed if we can learn the state of s_1 after the first CNOT gate. If s_1 is in $|0\rangle$, we know that s_t has a higher probability of being in $|0\rangle$. If we discard the cases where s_1 is in $|1\rangle$ and only keep the ones where it is in $|0\rangle$, we can reduce the error δ on the target qubit. The working principle behind this method is similar to an error *detecting* code, where the effective noise level is reduced by accepting the cases where no error occurs, and discarding those with an error occurring.

Referring to the circuit in Fig. 4.3, measuring $|0\rangle$ on s_1 updates the state of s_t to

$$\rho_t' = \frac{\langle 0|_1 \, \tau \, |0\rangle_1}{\langle 0| \, \mathrm{Tr}_t[\tau] \, |0\rangle} \tag{4.6}$$

where, throughout this chapter, we'll use τ to denote the state of the full system right before the measurement. Starting from two states with initial error rates δ_1 and δ_t, the final error rate on s_t upon measuring $|0\rangle$ on s_1 becomes

$$\delta_t' = \frac{\delta_1\delta_t}{1 - \delta_1 - \delta_t + 2\delta_1\delta_t}. \tag{4.7}$$

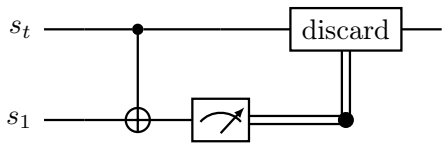

Figure 4.3: A circuit for 2-qubit measurement-based algorithmic cooling which increases the bias of the target qubit s_t. The double-line notation means "controlled on classical outcome", i.e., s_t is kept for further computations when the measurement outcome is 0 and discarded if the outcome is 1.

Then, taking $\delta_t = \delta_1 = \delta$ again for simplicity, we see that in the small-δ limit, the error is also reduced to $\mathcal{O}(\delta^2)$ to leading order. Importantly, the target qubit s_t is not being measured, which resolves the first issue we raised in Section 4.1. Assuming that the ancillary qubit s_1 can maintain sufficiently long coherence, we can store them in the quantum computer and measure them at the end of the computation, along with other qubits. We then post-select only those measurement outcomes where s_1 measures to 0. Interestingly, comparing with Eq. (4.5), we see that applying the above protocol once achieves the same polarization on s_t as applying infinitely many rounds of BCS, if initially $\delta_t = \delta_2$.

The above 2-qubit protocol, which we'll call MBAC-2, forms a basis for analyzing expanded versions of MBAC. For conventional AC protocols there are two major ways to expand them, either by using more ancillary qubits, or by repeating the protocol for multiple rounds. Since we have assumed that the ancillary qubits cannot be reused once they have been measured, MBAC protocols cannot be expanded by repeating it for multiple rounds. However, it is feasible to use more ancillary qubits to achieve better cooling. Specifically, imagine expanding s_1 to $k-1$ qubits in Fig. 4.3. We apply $k-1$ CNOT gates controlled by s_t and targeted on $s_i,\ i = 1, ..., k-1$, and measure all $s_1, ..., s_{k-1}$ at the end. The target s_t is kept only if *all* measurement outcomes are 0, and is discarded otherwise. We call the above protocol MBAC-k, which serves as an expanded version of MBAC-2.

To analyze MBAC-k, observe that it is equivalent to repeating $k-1$ "rounds" of MBAC-2, if the ancilla in MBAC-2 is allowed to return to its original state after being measured and can be reused again. The evolution of noise in a single "round" is already given in Eq. (4.7), so the evolution for MBAC-k can be recursively calculated from Eq. (4.7). Furthermore, by assuming again $\delta_t = \delta_1 = \delta$ initially and observing the first few terms, we can verify the analytic solution for this series, given by

$$\delta_t[k] = \frac{\delta^k}{\delta^k + (1-\delta)^k}, \tag{4.8}$$

where from now on we use $\delta_t[k]$ to denote "SP error on s_t after applying MBAC-k". In particular, $\delta_t[1]$ will be used to denote the initial SP error on s_t. The solution in Eq. (4.8) can be readily verified by plugging it back into Eq. (4.7). From this solution one can see how δ_t decreases approximately exponentially in k, especially in the small-δ limit.

To make a fair comparison between MBAC and conventional AC, we now briefly review the reversible scheme first proposed by Schulman and Vazirani [83]. The main step in reversible AC is called entropy compression, where through a unitary map $\mathcal{U}$, the entropy is extracted from some subset of qubits and transferred to another subset of qubits. If we constrain the system to start and end in diagonal states, and the goal is to cool down only one qubit, then it can be shown that the optimal unitary $\mathcal{U}$ is to perform a descending sort on the diagonal elements of the full system's density matrix [95, 96]. We thus consider this to be an upper bound on the performance of Schulman-Vazirani type of cooling, and will call this scheme SV-k if it uses a total of k qubits to cool one qubit. In Fig. 4.4, we compared MBAC-k (in circles) and SV-k (in squares) starting from two initial noise levels, $\delta = 0.1$ (in blue) and $\delta = 0.45$ (in red). The advantage of allowing projective measurement into the task of cooling, compared to the optimal reversible scheme, is clearly visible.

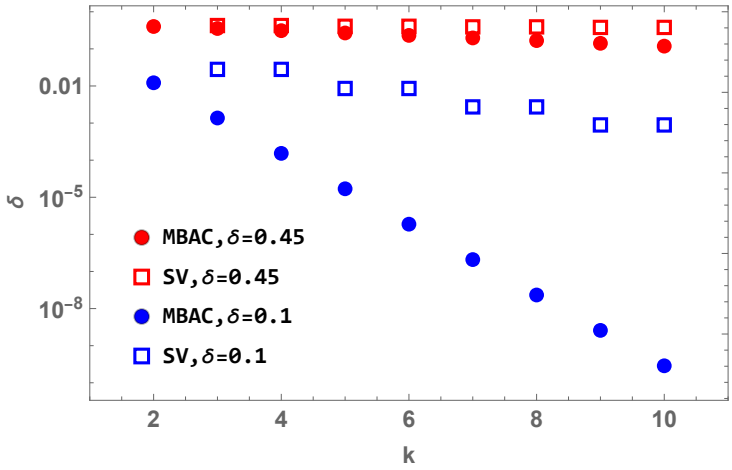

Figure 4.4: Simulated evolution of δ_t between MBAC-k and SV-k. The initial error (red: $\delta = 0.45$; blue: $\delta = 0.1$) are assumed to be the same on all qubits. The effect of decoherence is assumed negligible and gates are assumed to be ideal.

4.4 SPAM error characterization

In Section 4.2, we have already seen how the parameter δ can be interpreted as a measure of state preparation (SP) error. We will now formally define what SP and M errors are, then demonstrate how MBAC can be used to characterize SPAM errors. Our definition here will closely follow the one given in [97]. The ideal SPAM operators are denoted by a density operator ρ and a 2-outcome POVM $M = \{M_0, M_1\}$, which satisfy the physicality constraints that ρ and M_i are positive-semidefinite Hermitian operators, and $M_0 + M_1 = \sigma_I$. We assume that ideally, $\rho = M_0 = |0\rangle\langle 0|$. The SPAM error is defined as the probability of obtaining an incorrect outcome when measuring the initial state: that is,

$$\delta_{\text{SPAM}} \coloneqq 1 - \text{Tr}[\rho M_0]. \tag{4.9}$$

We now define the SP error, denoted as δ_{SP}, to be equal to δ_{SPAM} with an ideal measurement operator. Similarly, the M error (denoted as δ_{M}) is defined as δ_{SPAM} with an ideal input state. While the total SPAM error δ_{SPAM} is a measurable quantity, δ_{SP} and δ_{M} are not, when both state preparation and measurement errors are present. However, one may design algorithms that allow one to isolate the contributions to the total error from state preparation or measurement processes.

We now demonstrate how one can separately estimate δ_{SP} and δ_{M}, based on the concept of AC. For simplicity, we first assume that both the imperfect state and measurement operators are of the same form

$$\rho = \begin{pmatrix} 1-\delta_{\text{SP}} & 0 \\ 0 & \delta_{\text{SP}} \end{pmatrix}, \; M_0 = \begin{pmatrix} 1-\delta_{\text{M}} & 0 \\ 0 & \delta_{\text{M}} \end{pmatrix}, \tag{4.10}$$

so that the total SPAM error is given (from Eq. (4.9)) by

$$\delta_{\text{SPAM}} = \delta_{\text{SP}} + \delta_{\text{M}} - 2\delta_{\text{SP}}\delta_{\text{M}}. \tag{4.11}$$

As before, we will also assume that $\delta_{\text{SP}}, \delta_{\text{M}} \in [0, 0.5)$. Our goal is to separately estimate $\delta_{\text{SP},t}$ and $\delta_{\text{M},t}$ on s_t. For now we also assume, for simplicity, that measurement operations on all *ancillary* qubits are *ideal*, and the only noisy measurement is the one on s_t. This allows us to directly apply the previous calculations. This assumption will be relaxed in the next section.

Since $\delta_{\text{SPAM},t}$ can be obtained from directly measuring s_t, the problem of separately characterizing SPAM is then reduced to estimating either $\delta_{\text{SP},t}$ or $\delta_{\text{M},t}$: from the symmetry

between the two, once either is known, the other can also be calculated from the second equation.

It is now possible to intuitively see how AC can be used to resolve SPAM errors. We saw in Eq. (4.8) and Fig. 4.4 that MBAC can quickly reduce the error δ_t on s_t close to 0. Now imagine two separate experiments where in the first one, we measure s_t directly and obtain $\delta_{\text{SPAM},t}$. In the second one, we first apply multiple rounds of MBAC until the final bias on the target is sufficiently close to 1 (we'll later show how this can be determined), then measure the target qubit. Since the measurement operation is independent of the qubit state, $\delta_{\text{M},t}$ is directly obtained from the measurement result, since the input state is now ideal. From here, $\delta_{\text{SP},t}$ can be easily calculated from Eq. (4.11). We have thus separately estimated SPAM errors by first eliminating the SP error, determining the measurement error, then inferring the SP error from the total $\delta_{\text{SPAM},t}$.

We have so far focused on the case of diagonal state and measurement operators. To justify this, below we describe an averaging technique to convert arbitrary 1-qubit SPAM elements to this simpler case. Begin by noting that we can generally write

$$\begin{aligned} \rho &= \frac{1}{2}(\sigma_I + s_x\sigma_x + s_y\sigma_y + s_z\sigma_z) \\ M_0 &= \frac{1}{2}(m_i\sigma_I + m_x\sigma_x + m_y\sigma_y + m_z\sigma_z) \end{aligned} \tag{4.12}$$

where the s's and m's are unknown parameters ($s_i = 1$ because $\text{Tr}[\rho] = 1$). Assuming ideal quantum gates, we can obtain an effective initial state with $s_x = s_y = 0$ for an arbitrary 1-qubit circuit as follows. We perform two separate experiments, where in the first we apply the original circuit, and in the second we apply a σ_z gate immediately after the state preparation, then carry out the same circuit. Due to linearity of quantum operations, the average of measurement outcomes from the two experiments is then equivalent to one where $\rho_{\text{eff}} = \frac{1}{2}(\rho + \sigma_z\rho\sigma_z^\dagger) = \frac{1}{2}(\sigma_I + s_z\sigma_z)$. Similarly, we can also make $m_x = m_y = 0$ by averaging the results from the original experiment and one where a σ_z gate is applied immediately before the measurement. To set $m_i = 1$, we can average the original circuit with one where a σ_x is applied immediately before the measurement, and the outcomes 0 and 1 are relabelled (so that outcome 0 corresponds to the POVM element M_1 and *vice versa*). This reduces the SPAM operators to the ones described by Eq. (4.10).

4.5 MBAC with measurement errors

We now deal with the second issue raised in Section 4.1, and study the performance of MBAC when the measurement error is non-zero on all *ancillary* qubits as well, thereby

relaxing the assumption made in the previous section. We will show that MBAC still performs well if δ_{SPAM} on each ancilla is not very large, thereby relaxing the most crucial assumption in SPAM characterization. Again, in this section we will assume that quantum gates are ideal, which means that the SPAM averaging processes are also ideal, and the state and measurement operators can each be described by a single parameter.

The problem setup is as follows. We assume that each qubit i in a quantum computer has an independent state preparation error $\delta_{\text{SP},i}$, and measurement error $\delta_{\text{M},i}$, where i labels the qubit. The target qubit is labelled by t as usual. The goal is again to learn $\delta_{\text{SP},t}$ and $\delta_{\text{M},t}$. But as we show later, similar arguments can be used to learn all $\delta_{\text{SP},i}$ and $\delta_{\text{M},i}$ if desired.

Consider performing one successful round of the MBAC-2 protocol with two qubits s_t and s_1. With an imperfect measurement, we generalize the case of projective measurement in Eq. (4.6) to a POVM measurement, so that performing one successful round of MBAC updates the state of s_t to [56]

$$\rho_t' = \frac{\text{Tr}_1[\tau(I \otimes M_0)]}{\text{Tr}[\text{Tr}_t[\tau]M_0]}, \tag{4.13}$$

where, again, τ denotes the state of the full system immediately before measurement. Using Eq. (4.13), following again the circuit in Fig. 4.3 and the parametrization in Eq. (4.10), we calculate the SP-error on s_t after one round of MBAC to be

$$\begin{aligned} \delta_{\text{SP},t}' &= \delta_{\text{SP},t} \frac{2(\delta_{\text{SP},1} + \delta_{\text{M},1} - 2\delta_{\text{SP},1}\delta_{\text{M},1})}{1 + (1 - 2\delta_{\text{SP},1})(1 - 2\delta_{\text{SP},t})(1 - 2\delta_{\text{M},1})} \\ &= \delta_{\text{SP},t} \frac{2\delta_{\text{SPAM},1}}{1 + (1 - 2\delta_{\text{SP},1})(1 - 2\delta_{\text{SP},t})(1 - 2\delta_{\text{M},1})} \end{aligned} \tag{4.14}$$

where the second equality comes from Eq. (4.11). The ratio

$$\frac{\delta_{\text{SP},t}}{\delta_{\text{SP},t}'} = \frac{1 + (1 - 2\delta_{\text{SP},1})(1 - 2\delta_{\text{SP},t})(1 - 2\delta_{\text{M},1})}{2\delta_{\text{SPAM},1}} \tag{4.15}$$

is a measure of the improvement of SP-error on s_t after one round of MBAC, which is better when larger. For example, a ratio of 100 implies that $\delta_{\text{SP},t}$ has been reduced by a factor of 100. Intuitively, the improvement should be more significant when there is less error on s_1: indeed, if $\delta_{\text{SP},1} = \delta_{\text{M},1} = 0$ in Eq. (4.14), then $\delta_{\text{SP},t}' = 0$ and s_t will always be projected to $|0\rangle\langle 0|$ when the measurement outputs 0 on s_1. In the more general case where SPAM error on s_1 is present, we observe that the numerator on the RHS of Eq. (4.15) is always ≥ 1, in the relevant region where $\delta_{\text{SP},1}, \delta_{\text{SP},t}, \delta_{\text{M},1} \in [0, 1/2)$. Therefore,

$$\frac{\delta_{\text{SP},t}}{\delta_{\text{SP},t}'} \geq \frac{1}{2\delta_{\text{SPAM},1}}. \tag{4.16}$$

Furthermore, in the limit where all the $\delta_{\mathrm{SP},1}$, $\delta_{\mathrm{SP},t}$, $\delta_{\mathrm{M},1} \ll 1$, the bound in Eq. (4.16) can be improved by approximately a factor of 2 to simply $1/\delta_{\mathrm{SPAM},1}$.

Next, recall from Section 4.3 that MBAC-k is equivalent to repeating $k-1$ rounds of MBAC-2. Using mathematical induction, we see that by applying one successful run of MBAC-k, the final SP-error on s_t is upper bounded by

$$\delta_{\mathrm{SP},t}[k] \leq \delta_{\mathrm{SP},t} \prod_{i=1}^{k-1} (2\delta_{\mathrm{SPAM},i}). \tag{4.17}$$

Recall that $\delta_{\mathrm{SPAM},i}$ is a measurable quantity obtained by measuring the initial state on ancilla s_i. The product simply corresponds to the probability of directly measuring all ancillary qubits $s_1 \dots s_{k-1}$ after they are prepared, and getting the output 1 on all qubits. Therefore, given $\delta_{\mathrm{SPAM},i}$ on each qubit (which can be obtained before the experiment, during the calibration step), Eq. (4.17) guarantees that $\delta_{\mathrm{SP},t}$ from the output of a successful run of MBAC-k is at least reduced by $\prod_{i=1}^{k-1}(2\delta_{\mathrm{SPAM},i})$. This shows that finite measurement error on the ancillary qubits do not pose a fundamental limitation to the cooling power of MBAC. In particular, as long as each $\delta_{\mathrm{SPAM},i} < 1/2$ (which we will assume from now on), the output state is guaranteed to be more pure than the input. Moreover, if both $\delta_{\mathrm{SP},t}$ and all $\delta_{\mathrm{SPAM},i}$'s are upper bounded by a constant δ, then $\delta_{\mathrm{SP},t}[k]$ is simply upper bounded by δ^k.

4.6 Number of trials needed in MBAC-k

Next, we study whether the probabilistic nature of MBAC hinders its usefulness in practice. We will see that it remains practically useful for a wide range of experimentally relevant SPAM error rates. To illustrate the problem, first ignore any measurement error and consider a run of MBAC-2. The post-measurement state upon measuring 1 (which imply a failed run) can be computed by changing all 0's to 1's in Eq. (4.6), leading to

$$\delta'_{t,\mathrm{fail}} = \frac{\delta_t(1-\delta_1)}{\delta_t + \delta_1 - 2\delta_t\delta_1} = \frac{1}{2} - \frac{\delta_1 - \delta_t}{2(\delta_t + \delta_1 - 2\delta_t\delta_1)}. \tag{4.18}$$

One sees that, if $\delta_1 = \delta_t$ initially, then obtaining measurement outcome 1 on s_1 will heat up the state to a completely mixed one. The probability of failure is

$$p_{\mathrm{fail}} = \delta_1 + \delta_t - 2\delta_1\delta_t. \tag{4.19}$$

For MBAC-k, since we need all the $k-1$ measurements to succeed, the rate of success decreases exponentially with k, and it becomes increasingly difficult to get a successful run.

Fortunately, recall from Eq. (4.17) that the target polarization also improves exponentially fast with increasing number of ancillas. Therefore, the intuition is that one does not need a large k value to achieve significant cooling, which in turn will not have a vanishingly small success probability. Below we'll make this intuition more concrete. Specifically, we ask the following question: if we would like to reduce $\delta_{\mathrm{SP},t}$ by a factor of r, i.e., we want $\delta_{\mathrm{SP},t}/\delta_{\mathrm{SP},t}[k] = r$, how many runs are needed to achieve this cooling ratio? To answer this we will start from Eq. (4.17), and first derive a relation between k and r. Note that because each $\delta_{\mathrm{SPAM},i}$ can be different, the most general expression will involve all $\delta_{\mathrm{SPAM},i}$'s. In order to obtain an expression in k, we now make the assumption that all $k-1$ ancillary qubits have the same δ_{SPAM}, which we will denote as $\delta_{\mathrm{SPAM},a}$ where a stands for the word "ancillary". The case where the $\delta_{\mathrm{SPAM},i}$'s are different can be bounded similarly by setting $\delta_{\mathrm{SPAM},a}$ to the highest among all $\delta_{\mathrm{SPAM},i}$.

With this assumption, we obtain the following inequality between k and r:

$$r = \frac{\delta_{\mathrm{SP},t}}{\delta_{\mathrm{SP},t}[k]} \geq (2\delta_{\mathrm{SPAM},a})^{-(k-1)}. \tag{4.20}$$

Taking the logarithm of both sides and rearranging the terms (note that $\log(2\delta_{\mathrm{SPAM},a}) < 0$) results in

$$k - 1 \leq \frac{\log(r)}{-\log(2\delta_{\mathrm{SPAM},a})}. \tag{4.21}$$

This allows the experimentalist to determine the total number of ancillary qubits needed in order to cool to the desired noise level, based on their hardware specifications (i.e., the $\delta_{\mathrm{SPAM},a}$ on their hardware). Importantly, this upper bound scales logarithmically with r.

Next, we compute the expectation value of the total number of runs needed before having a successful run, in order to achieve a cooling ratio r. We show in Appendix B.1 that the expected number of runs is upper bounded by a function,

$$N_{\mathrm{upper}}(r) = (r)^{\frac{\log(A)}{\log(B)}} \tag{4.22}$$

where

$$\begin{aligned} A &= (1 - \delta_{\mathrm{SP},t}[1] - \delta_{\mathrm{SP,a}} + 2\delta_{\mathrm{SP},t}[1]\delta_{\mathrm{SP,a}})(1 - \delta_{\mathrm{M},a}) \\ B &= 2\delta_{\mathrm{SPAM},a}. \end{aligned} \tag{4.23}$$

We see that Eq. (4.22) scales polynomially in r. Note that A is simply the success probability (i.e., measuring 0) of doing the first run of MBAC-2, so the exponent is a measurable quantity. Thus, the upper bound can be calculated given a target r.

To understand the behavior of $N_{\text{upper}}(r)$ more concretely, we can further simplify Eq. (4.22) by assuming again $\delta_{\text{SP},t}[1] = \delta_{\text{SP,a}} \coloneqq \delta_{\text{SP}}$, i.e., both the target and the ancillas have the same initial error δ_{SP}. The exponent then becomes (recall δ_{SPAM} from Eq. (4.11))

$$\frac{\log(A)}{\log(B)} = \frac{\log((1 - 2\delta_{\text{SP}} + 2\delta_{\text{SP}}^2)(1 - \delta_{\text{M},a}))}{\log(2(\delta_{\text{SP}} + \delta_{\text{M},a} - 2\delta_{\text{SP}}\delta_{\text{M},a}))}. \tag{4.24}$$

Shown in Fig. 4.5 are plots of $N_{\text{upper}}(r)$ in Eq. (4.22) as a function of δ_{SP}, after making the simplifying assumption in Eq. (4.24), for a few chosen values of r and $\delta_{\text{M},a}$. As the initial SP error rate approaches 0.5 (the theoretical maximum), MBAC fails since the measurement is simply giving random outputs. In this case, $N_{\text{upper}}(r)$ diverges as expected. On the other hand, for reasonably low values of δ_{SP}, we see that the expected number of runs before achieving a successful one is rather low. For example, for $\delta_{\text{SP}} = 0.1$, $\delta_{\text{M},a} = 0$, and $r = 1000$, we expect to obtain a successful run in about 2 trials. For $\delta_{\text{SP}} = 0.1$, $\delta_{\text{M},a} = 0.1$, and $r = 1000$, the expected number of trials is approximately 3. Note that this corresponds to a case of 10% SP error rate plus 10% measurement error rate on the ancillas, combining to almost 20% of total SPAM error rate. Many modern quantum computing platforms [98, 41, 99] can now achieve SPAM error rates below this level. In these cases, MBAC will be a useful and easy tool to significantly improve the quality of state preparation in quantum computers.

4.7 Summary

In this chapter, we introduced a new variant of AC protocol based on the ability to perform imperfect measurements on individual qubits, which we call measurement-based AC (MBAC). Using this method, we developed a novel and simple way to separately characterize state preparation and measurement errors, by eliminating the former using MBAC and directly obtaining the latter. Our approach is applicable to many current quantum computing platforms, and significantly outperforms the optimal reversible AC protocol in the absence of measurement errors. Moreover, we showed that its cooling power retains the exponential scaling when measurement errors are present. Despite the probabilistic nature of MBAC, we have demonstrated its practical usefulness under realistic error rates for current quantum processors. We believe our method can be a helpful tool for benchmarking and improving current NISQ-era quantum computers, and can provide further insights to measurement as a resource for cooling.

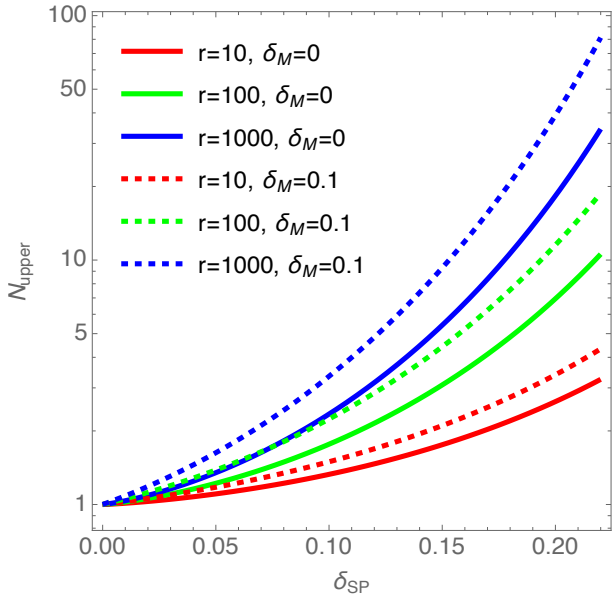

Figure 4.5: Upper bound N_{upper} on expected number of runs required for different values of r (the improvement ratio defined in Eq. (4.20)), versus the initial SP error δ_{SP} (assumed to be the same on the target and all ancillary qubits). Solid and dashed lines represent the cases of $\delta_{\text{M},a} = 0$ and $\delta_{\text{M},a} = 0.1$, respectively.

Chapter 5

Thermodynamic analyses of algorithmic cooling

We have seen how algorithmic cooling can be used to produce qubits in highly pure initial states from the previous chapter. This motivates us to further investigate the performance of AC as algorithmic refrigerators. In the past, many AC protocols have been proposed and tested, each requiring slightly different resources. These were mostly analyzed from a quantum information point of view. Recently, there is an increasing trend to study the thermodynamic performance of AC protocols. For example, the work by Soldati et al. [100] looked at thermodynamics of the 3-qubit version of the PPA algorithm (will be introduced later). The goal of this chapter is to expand upon this topic, establishing a more general framework for studying thermodynamics of AC for analyzing a wide family of protocols. In particular, we will examine different AC protocols belonging to the family of coherent cooling protocols, by first reviewing their procedure, cooling limits, and target state evolution in Section 5.2 and Section 5.3. In doing so, we use the transition matrix formalism to derive the quantities we need, showing that a unified view can be established among these protocols. We then study the thermodynamic aspects of each protocol. In Section 5.4 and Section 5.5, we propose two efficiency measures based on the amount of work required, or the amount of heat released. We show how each measure can be computed for a given protocol, and compare the previously studied protocols using both measures, providing suggestions on which ones to use when these protocols are to be carried out experimentally. Finally, in Section 5.6 we propose improved protocols that are energetically more favorable over the original proposals.

5.1 Notation

In this chapter, we will denote the dimension of a subsystem on the superscript with an angled bracket. We denote the identity operator by the symbol $\mathbb{1}$. A subscript on an operator implies that it is a reduced operator: for example ρ_1 refers to the density operator on subsystem 1. A subscript with a square bracket implies that it is a full-system operator, where the (lower dimensional) operator acts nontrivially on the subsystem in the bracket: for example $\sigma_{x,[1]}$ is a tensor product of σ_x on qubit 1 and $\mathbb{1}$ on other subsystems. Qubit numbering will start from 1 in this chapter.

Given a system with Hamiltonian H, the thermal (or, Gibbs) state at inverse temperature β is defined as

$$\rho^{th}(\beta, H) = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]} \tag{5.1}$$

and will be denoted by a superscript "th".

In this chapter we are mostly working with states that are diagonal in the energy eigenbasis. In particular, any diagonal single qubit state can be characterized by a single parameter, with the different choices which have appeared in the literature include but are not limited to:

1. The ground ($|0\rangle$) state population, p
2. The excited ($|1\rangle$) state population δ
3. The polarization ϵ, defined as the difference between ground and excited state populations
4. The effective temperature βE, where E denotes the energy gap between $|0\rangle$ and $|1\rangle$, by considering the qubit as a thermal state at inverse temperature β

Unless otherwise stated, these symbols will be reserved for the corresponding quantities in this chapter. For convenience, we list the transformations between these quantities in Table 5.1.

For a system of N independent qubits, the stationary (noninteracting) Hamiltonian is given by

$$H_{\mathrm{tot}} = \sum_{i=1}^{N} h_{[i]}(E_i) \tag{5.2}$$

	p	δ	ϵ	βE
p	-	$p = 1 - \delta$	$p = \frac{1+\epsilon}{2}$	$p = \frac{1}{1+e^{-\beta E}}$
δ	$\delta = 1 - p$	-	$\delta = \frac{1-\epsilon}{2}$	$\delta = \frac{1}{1+e^{\beta E}}$
ϵ	$\epsilon = 2p - 1$	$\epsilon = 1 - 2\delta$	-	$\epsilon = \tanh\frac{\beta E}{2}$
βE	$\beta E = \log\frac{p}{1-p}$	$\beta E = \log\frac{1-\delta}{\delta}$	$\beta E = \log\frac{1+\epsilon}{1-\epsilon}$	-

Table 5.1: Conversion table between parameters describing single-qubit states.

where $h(E) = E\,|1\rangle\langle 1|$ is the single-qubit Hamiltonian for the i-th qubit with energy gap E. The value of E can differ from qubit to qubit. Note that this is a shifted version of the Z-type Hamiltonian, $h' = -\frac{E}{2}\sigma_z$ where σ_z is the Pauli-Z operator. We will mostly choose h over h' for computational convenience.

5.2 Review of AC schemes: protocols and framework

We introduced the original ideas on AC from Sørensen and Schulman/Vazirani (SV) in Section 4.1. Sørensen first observed the phenomenon where the polarization of a subset of spins can be boosted at the expense of reduced polarization of nearby spins, during a unitary evolution. Schulman and Vazirani were the first to give explicit ways to implement the redistribution of entropy in the context of quantum information [101, 102]. They considered performing reversible operations to a chain of qubits, to create a separation of cold and warm regions within the chain. Utilizing results from classical data compression, they suggested that starting with n spins with low initial polarization ϵ, one can reversibly re-distribute the entropy within the spins and produce $cn\epsilon^2$ useful qubits that are close to the pure state $|0\rangle$ where c is some constant. In this chapter, we will focus on newer protocols derived from these original proposals, as introduced below.

HBAC and PPA-HBAC

An improved AC protocol over the original SV-cooling appeared in 2007, bearing the name heat-bath algorithmic cooling (HBAC) [103]. In addition to unitary compressions within the string, the authors assumed the ability to interact with the system with an external heat bath. This allows the entropy to be released outside the system, thereby significantly reducing the lowest achievable polarization using the same number of initial spins compared to SV-cooling. Here, we denote the spins to be cooled as the target spins, and those used as thermal machines as the computational spins. The protocol consists of repeated rounds, each consisting of two subroutines, where the first step compresses entropy from target spins into computational spins (similar to SV-cooling), and the second step is thermalization of selected computational spins with higher-than-bath temperature. Because different spins have different magnetic moments, which correspond to different relaxation rates, it is possible to design the system so that the target spins remain relatively unchanged while the computational spins relax to the bath temperature. This eventually allows the target to be cooled to lower temperatures compared to SV-cooling.

Among all possible unitaries, the one that achieves optimal extraction of entropy from the target was later studied under the name partner-pairing algorithm [104], or PPA for short. The HBAC protocol that uses the PPA scheme is thus called PPA-HBAC, or PPA for short. This compression unitary is later defined in Section 5.3.4. Small instances of PPA has been tested experimentally to show its ability to improve target polarization [87].

NOE-HBAC and SRG-HBAC

After the discovery of PPA, more works have appeared which aimed at achieving higher final polarization than PPA, using the same number of qubits. Two notable examples are nuclear Overhauser-effect HBAC (NOE-HBAC) and state-relaxation HBAC (SR-HBAC) protocols, both appearing in 2017 [92]. In addition to the ability of interacting with a heat bath, both protocols require the possibility of a state-relaxation step. More specifically, this refers to the ability to "thermalize" populations of selected energy levels in the composite spin system. This was inspired from the Nuclear Overhauser Effect, where the presence of cross-relaxation between two nearby spins allows the polarization of one spin to increase when the other spin is saturated (to the completely mixed state). Both NOE-HBAC and SR-HBAC can cool a target qubit using only one computational qubit, unlike the PPA which requires at least 2, and the SR-HBAC achieves higher final polarization than NOE-HBAC. The SR step is denoted by Γ_2 in the circuits, Fig. 5.1b and Fig. 5.1c. This shows

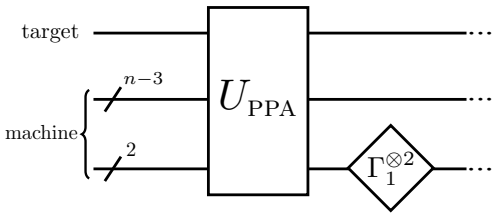

(a) Round of PPA

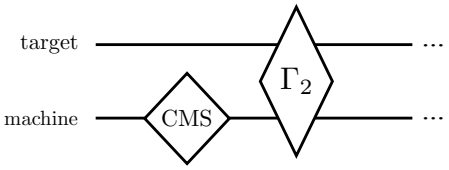

(b) Round of NOE

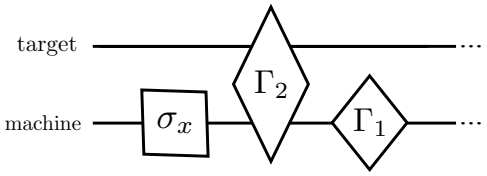

(c) Round of SRG2

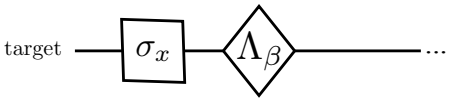

(d) Round of xHBAC

Figure 5.1: Circuit diagrams for the rounds of the HBAC protocols. The diamond shape indicates non-unitary gates, while the box shape represents unitary gates.

that a more refined relaxation pathway for individual energy levels can be a useful tool for cooling.

xHBAC

More recently, further generalizations of the SR-based HBAC has been proposed, most notably in the name extended HBAC (xHBAC) [94]. By realizing that the SR step belongs to a wider family of quantum operations called "dephasing thermalization", the authors showed that even further cooling can be achieved if one can perform arbitrary dephasing thermalizations, by optimizing over this family. In particular, they showed that a single qubit can be cooled towards polarization 1 (corresponding to zero temperature) exponentially fast, using the optimal thermalization operation, denoted by Λ_β in Fig. 5.1d. This seeming violation of the second law of thermodynamics can be resolved by realizing that performing Λ_β requires detailed interaction with an infinite dimensional ancillary system, which has a diverging control complexity as a resource [105].

5.2.1 Coherent control AC

The AC protocols described above all fall under the category of "coherent control" cooling in the literature [105], defined as follows. The setup consists of three components, namely, a *target* system to be cooled (denoted by a subscript t in this chapter), a thermal *machine* (denoted by a subscript m), and an outside environment acting as a heat *bath* (denoted by a subscript b). Collectively, we will refer to the total of target and machine as simply the *system*. Following a conventional treatment in thermodynamics, we will take the bath component to consist of the rest of the universe outside of the system, so that the target, machine, and bath combined is taken to be closed.

Coherent control cooling protocols are composed of two basic subroutines, namely, a unitary *control* step and a *thermalization* step. The meaning of these two subroutines will be made exact in this chapter through the following definitions.

Definition 5.2.1. A control subroutine refers to a map achieved by a unitary $U_{t,m}$ that acts on the target and the machine subsystems.

In the control step, one applies a unitary between the target system and the machine. The goal of this step is typically (but not necessarily) to transfer energy from the target into the machine. Importantly, this unitary can be an arbitrary one and need not be energy-preserving, i.e., there can be a net work input when applying $U_{t,m}$.

Definition 5.2.2. A thermalization subroutine refers to a map $\mathcal{E}$ that is a *thermal operation* (TO) acting on the target and machine subsystems, that is, one which can be written as

$$\mathcal{E}(\rho_{t,m}) = \mathrm{Tr}_b[V_{t,m,b}(\rho_{t,m} \otimes \rho^{th}(\beta_b, H_b))V_{t,m,b}^\dagger] \tag{5.3}$$

where the joint unitary $V_{t,m,b}$ acting on the target, machine, and bath satisfies $[V_{t,m,b}, H_{[t]} + H_{[m]} + H_{[b]}] = 0$.

In the thermalization step, one allows (part of) the system to exchange energy with the heat bath. Usually the net result is a transfer of energy from the system into the bath. When a protocol is repeated for multiple rounds, we will assume that the bath is sufficiently large and always start from a thermal state, and ignore correlation effects that might build up during the relaxation.

The following holds true regarding the transition matrix of any TO map:

Corollary 5.2.1. *The transition matrix $G_{\mathcal{E}}$ of a TO map $\mathcal{E}$ is Gibbs-stochastic.*

Here, Gibbs-stochastic means that it is a stochastic matrix (as demonstrated before) that preserves a Gibbs state input, i.e., $G_{\mathcal{E}}|\rho^{th}\rangle\rangle = |\rho^{th}\rangle\rangle$. This can be seen by inserting the definition of ρ^{th} into Definition 5.2.2 and noticing that $[V_{t,m,b}, H_{[t]} + H_{[m]} + H_{[b]}] = 0$.

Crucially, the thermalization subroutine is one which can be realized without any external work input. For example, the act of resetting a qubit to the thermal state at the bath temperature (i.e., a complete thermalization) falls under the TO category. This makes intuitive sense, since one can imagine achieving this by putting a system in contact with a heat bath for a long period of time, which indeed does not require any work input. However, here we have used a broader definition for the thermalization subroutine than just complete thermalization. These more general TO operations will appear as constituents of the cooling protocols introduced later.

Definition 5.2.3. A coherent control AC is a sequential procedure where each step is either a control subroutine, or a thermalization subroutine.

The specific form of control and thermalization operations differ from one protocol to another, but all the coherent control AC protocols involve only these two types of operations. In the following sections, we will review in further detail the structure of each protocol studied throughout this chapter. The circuit view of all the protocols being studied in this chapter are collectively shown in Fig. 5.1.

5.3 Performance of HBAC protocols

The standard approach for evaluating the effectiveness of HBAC protocols is to examine their cooling limits, i.e., what's the lowest achievable temperature of the target qubit depending on the available resources, such as machine size and allowed operations. One can analyze the performance of a given protocol by monitoring the evolution of the target system. In this case, one can disregard the evolution of the machine, and focus solely on maximizing cooling of the target. Here, we present the state evolution as well as cooling limits of the different HBAC protocols using the *transition matrix* formalism, which is naturally suited for this task. This section starts with a brief introduction to the transition matrix formalism. We then apply this formalism to each HBAC protocol presented in this chapter to gain a unified view of their cooling behavior. Finally, we present the unified cooling limits in Section 5.3.6, and the unified target state evolution in Section 5.3.7.

5.3.1 Transition matrix formalism

We first give a brief review of the transition matrix formalism used in this chapter. A similar superoperator description of quantum objects has already been used multiple times throughout this thesis (e.g., in Section 2.2); meanwhile, the treatment in this section is slightly different than that in previous chapters and requires further clarifications. The state of a general quantum system is described by a density operator, which can be represented in vector form by choosing some orthonormal basis σ_k. This is typically represented by a supervector $|\rho\rangle\!\rangle$, with components $|\rho\rangle\!\rangle_k = \langle\sigma_k, \rho\rangle$ where σ_k ranges over all the basis elements. Given a CPTP map $\mathcal{E}$, if we construct a matrix $G_\mathcal{E}$ with the i, j-th component being $\langle\sigma_i, \mathcal{E}(\sigma_j)\rangle$ where σ_i, σ_j ranges over the basis elements, then this matrix entails all the information of $\mathcal{E}$. This matrix is called the transition matrix for the process $\mathcal{E}$, and it acts on the state $|\rho\rangle\!\rangle$ through matrix multiplication:

$$|\mathcal{E}(\rho)\rangle\!\rangle = G_\mathcal{E}|\rho\rangle\!\rangle \tag{5.4}$$

by a simple linearity argument.

For the HBAC protocols under consideration in this chapter, the density matrices are diagonal with respect to the energy eigenbasis. Furthermore, all quantum processes under consideration map diagonal operators to diagonal operators. It can be straightforwardly seen that in this case, there is a natural vector representation for the operators by simply taking the diagonal elements of a N-qubit state as an $2^N \times 1$ vector. The orthonormal bases are the elementary operators which are diagonal in the energy basis:

$$\sigma_k = |k\rangle\langle k|$$

where k denotes the energy eigenbases. For example, the state

$$\rho = \begin{pmatrix} \frac{1}{3} & 0 \\ 0 & \frac{2}{3} \end{pmatrix}$$

is represented by the vector $(\frac{1}{3}, \frac{2}{3})$. We will slightly abuse the notation here by denoting vector representations resulting from this mapping by $|\cdot\rangle\!\rangle$ again, such that the standard basis vectors are written as

$$|\sigma_k\rangle\!\rangle \coloneqq |k\rangle\!\rangle \tag{5.5}$$

and the vector representing a diagonal operator ρ is $|\rho\rangle\!\rangle = \sum_k \langle\!\langle k|\rho\rangle\!\rangle |k\rangle\!\rangle$. The Hilbert-Schmidt inner product becomes an inner vector product as

$$\langle A, B\rangle = \mathrm{Tr}\big[A^\dagger B\big] = \langle\!\langle A|B\rangle\!\rangle \tag{5.6}$$

for A, B both diagonal, where $\langle\langle A|$ is the conjugate transpose of $|A\rangle\rangle$. Matrices representing transformations will be denoted by G. It can be easily verified that the basis vectors are orthonormal, i.e., $\langle\langle i|j\rangle\rangle = \delta_{ij}$. Therefore, the transition matrix for a map Φ has components

$$(G_\Phi)_{i,j} = \langle\langle \sigma_i|\Phi(\sigma_j)\rangle\rangle \tag{5.7}$$

In the literature, these are typically referred to as population transition matrices, since they entail the change in population of each energy level during a transformation.

In this chapter we often encounter expectation values of the form $\mathrm{Tr}[AB]$, which equals to $\mathrm{Tr}\left[A^\dagger B\right]$ when A is Hermitian. For this reason we will sometimes write $\mathrm{Tr}[AB] = \langle\langle A|B\rangle\rangle$ when it is clear that A is Hermitian.

The restriction to diagonal states requires that $G_\mathcal{E}$ maps probability vectors to probability vectors, making $G_\mathcal{E}$ a left-stochastic matrix. This property does not hold for the fully general superoperator formalism used in previous chapters. The matrix $G_\mathcal{E}$ fully describes the effect of an operation and is uniquely determined by the cooling protocol. Note that the dimensions of $G_\mathcal{E}$ may be unequal for maps that change the dimension of $|\rho\rangle\rangle$.

Knowing the exact form of G for each step enables a direct simulation of the cooling process through direct matrix multiplication. For certain HBAC protocols that consist of repeated rounds of the same operation, the computation is greatly simplified by first diagonalizing G as $G = TDT^{-1}$ where D is diagonal (assuming that G is diagonalizable), then raising D to the desired power n. Taking the limit of $n \to \infty$ gives the asymptotic final state of the system, from which quantities such as the asymptotic polarization can be easily computed. In the section below, we will utilize this technique to unify the different HBAC protocols. One shall see that the transition matrix formalism provides a consistent way of simulating HBAC protocols, especially when the protocol involves repeating the same operations over multiple rounds. We now apply this on the protocols of interest to derive several important characterizations, including the polarization evolution and the asymptotic state. Later in Section 5.4, we will again apply it to obtain thermodynamic quantities like work, energy flow, and efficiency measures.

5.3.2 The 2-qubit NOE

In the 2-qubit NOE HBAC (NOE2), the target system is the first qubit and the machine is the second qubit, as shown in Fig. 5.1b. It involves two steps: the first one is a randomization of the second qubit to the completely mixed state (CMS). In practice, this can be achieved through a random rotation, and averaging the results from multiple experiments. For this reason, CMS will be treated as effectively a probabilistic mixture of

unitary processes. The second step is the Γ_2 relaxation, defined by the following Kraus operators:

$$\begin{aligned} A_1 &= \sqrt{p_2}\,|00\rangle\langle 00| \\ A_2 &= \sqrt{p_2}\,|00\rangle\langle 11| \\ A_3 &= \sqrt{1-p_2}\,|11\rangle\langle 11| \\ A_4 &= \sqrt{1-p_2}\,|11\rangle\langle 00| \\ A_5 &= |01\rangle\langle 01| \\ A_6 &= |10\rangle\langle 10|\,, \end{aligned} \tag{5.8}$$

where p_2 is defined by

$$p_2 \coloneqq \frac{1}{1+e^{-2\beta_b E}}. \tag{5.9}$$

Thus, p_2 is the ground state population of a thermal qubit at bath temperature with energy gap $2E$, which can also be interpreted as the relative ground state population between the states $|00\rangle$ and $|11\rangle$ in thermal equilibrium, if both qubits have an energy gap of E. In comparison, the relaxation of a single qubit alone with the bath achieves a thermal equilibrium between the states $|0\rangle$ and $|1\rangle$, and will be denoted by Γ_1.

One can verify that the above Kraus operators correspond to a dephasing map, which removes all off-diagonal elements. In the reduced space of diagonal states, Γ_2 can be written as the following transition matrix:

$$G_{\Gamma_2} = \begin{pmatrix} p_2 & 0 & 0 & p_2 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1-p_2 & 0 & 0 & 1-p_2 \end{pmatrix}. \tag{5.10}$$

One can see that Γ_2 preserves the Gibbs state $\rho^{th}_{t,m}(\beta_b, H_{t,m})$ and acts non-trivially on only two energy levels, making it fall into the category of 2-level Gibbs stochastic matrices. It is known that every 2-level Gibbs stochastic transition matrix corresponds to a thermal operation [106]. Therefore, it belongs to the thermalization subroutine category by our definition. The same argument would trivially apply to Γ_1, leading to the corollary below.

Corollary 5.3.1. *Both Γ_1 and Γ_2 maps are thermal operations.*

We will use the transition matrix to directly calculate the evolution of SRG2. The detailed calculations are shown in Appendix C.4 for all protocols under consideration. Here, we will state the final results, then group them into a unified formula in Section 5.3.6

and Section 5.3.7. Starting from a target qubit with ground state population p_t, its final ground state population after n rounds will become

$$p_t(n) = \frac{x}{1+x} - 2^{-n}(\frac{x}{1+x} - p_t) \tag{5.11}$$

where $x = \frac{p_2}{1-p_2}$. The asymptotic state is

$$|\rho_{\mathrm{NOE},t}(\infty)\rangle\rangle = \frac{1}{1+x}\begin{pmatrix} x \\ 1 \end{pmatrix} = \begin{pmatrix} p_2 \\ 1-p_2 \end{pmatrix}. \tag{5.12}$$

Since $p_2 > p_b$, this is a purified state compared to the original bath-temperature state.

5.3.3 The 2-qubit SRG2

In SRΓ_2-HBAC (SRG2), the target system is the first qubit and the machine is the second qubit, as shown in Fig. 5.1c. The first step is to apply a Pauli σ_x gate to the second qubit. The next step is the Γ_2 relaxation on both qubits, and the third step is a Γ_1 relaxation on the second qubit.

Starting with a target qubit characterized by p_t, the ground state population after n rounds of cooling is

$$p_t(n) = \frac{w}{w+v} - u^n(\frac{w}{w+v} - p_t) \tag{5.13}$$

where $w = p_2 p_b$, $v = 1 - p_2 - p_b + p_2 p_b$, $u = 1 - v - w$. The asymptotic state is

$$|\rho_{\mathrm{SRG2},t}(\infty)\rangle\rangle = \frac{1}{w+v}\begin{pmatrix} w \\ v \end{pmatrix}. \tag{5.14}$$

It is possible to extend the SRG2 to larger systems with more qubits [92], by using a generalized relaxation process Γ_N on N qubits between the states $|0\ldots0\rangle$ and $|1\ldots1\rangle$. However, the original proposal uses the SRG2 as a subroutine and requires running it many times to reach the asymptotic state. As will be seen later, this has an infinite energy cost and is unrealistic to perform in the original form. In practice, one can truncate up to a finite number of repetitions, resulting in a modified protocol. Here, we will only state a known result on the asymptotic achievable polarization for SRGN as analyzed in the original protocol in Section 5.3.6.

5.3.4 PPA

In PPA-HBAC, the target system is the first qubit represented by the most significant bit in a binary representation, and the machine is the rest of the qubits. The compression unitary U_{PPA} corresponds to a decreasing SORT of the diagonal elements in the input state. This SORT is with respect to an ascending, lexicographical basis where 000 comes before 001, which comes before 010, etc. After the sorting unitary, the population in the 0 subspace of the first qubit has increased, implying a decrease of entropy. The entropy compression results in a separation of cold and warm qubits, where the target defines the "cold end".

The second step is a Γ_1 thermalization of qubits at the warm end of the string. Here we relax the last two qubits, which always have higher temperatures than the bath after applying U_{PPA}. It is also assumed that this step occurs at a much higher rate than the thermalization of the other qubits, so that the final state of those are effectively unchanged.

The difficulty in analyzing PPA lies in calculating the compression step. The smallest, 3-qubit variant of PPA has the property that starting from the all-thermal state at the same temperature, all the compression unitaries are identically equal to [107]

$$U_{\text{PPA3}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{5.15}$$

Starting with a target qubit characterized by p_t, the target's ground state population at the end of the n-th round is

$$p_t(n) = \frac{z}{1+z} - y^n\left(\frac{z}{1+z} - p_t\right). \tag{5.16}$$

where $y = 2p_b(1-p_b)$ and $z = p_b^2/(1-p_b)^2$. The asymptotic state is

$$|\rho_{\text{PPA3},t}(\infty)\rangle\rangle = \frac{1}{1+z}\begin{pmatrix} z \\ 1 \end{pmatrix}. \tag{5.17}$$

For more than 3 qubits, U_{PPA} generally changes from round to round. So far, no analytical expression has been obtained for the optimal compression unitary during a

general round of $N > 3$ PPA. There exist alternative compression unitaries that are fixed for each cooling round and can achieve the same asymptotic temperature as the optimal one, with a slower convergence rate. Here, we will rely on numerical simulations that performs the optimal PPA in larger systems.

5.3.5 The 1-qubit xHBAC

The 1-qubit xHBAC protocol involves only a single qubit as the target, as shown in Fig. 5.1d, although realizing the β-swap operation involves interaction with a well-controlled bosonic mode. It consists of a 1-qubit Pauli σ_x gate, followed by a β-swap operation Λ_β, defined in [94] as the dephasing operation with the following Kraus operators:

$$\begin{aligned} A_1 &= |0\rangle\langle 1| \\ A_2 &= \sqrt{e^{-\beta_b E}}\,|1\rangle\langle 0| \\ A_3 &= \sqrt{1-e^{-\beta_b E}}\,|0\rangle\langle 0|\,. \end{aligned} \tag{5.18}$$

The transition matrix describing the β-swap is

$$G_\beta = \begin{pmatrix} 1 - Ee^{-\beta_b} & 1 \\ e^{-E\beta_b} & 0 \end{pmatrix}. \tag{5.19}$$

We calculate the ground state population after n rounds of cooling to be

$$p_{\text{xHBAC1},t}(n) = 1 - s^n(1 - p_t) \tag{5.20}$$

where $s = e^{-\beta_b E}$. The asymptotic state is

$$|\rho_{\text{xHBAC1},t}(\infty)\rangle\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \tag{5.21}$$

i.e., the pure state $|0\rangle$.

5.3.6 Unified cooling limits

The asymptotic polarizations for the different algorithmic cooling protocols discussed above are captured by the general expression:

$$\epsilon_\infty(\epsilon_b, x) := \frac{(1+\epsilon_b)^x - (1-\epsilon_b)^x}{(1+\epsilon_b)^x + (1-\epsilon_b)^x} \tag{5.22}$$

$$= \tanh\left[x * \operatorname{arctanh}(\epsilon_b)\right]. \tag{5.23}$$

where x depends on the implemented protocol, as shown below. Equivalently, since $1+\epsilon = 2p$, and $1-\epsilon = 2\delta$, in the cooling limit, the ground state population of the target qubit corresponds to

$$p_{t,\infty}(p_b, x) = \frac{p_b^x}{p_b^x + \delta_b^x}. \tag{5.24}$$

The corresponding temperature of the target qubit in the asymptotic limit is given by

$$T_{t,\infty} = \frac{T_b}{x}, \tag{5.25}$$

where T_b is the temperature of the thermal bath.

- **PPAN:** $x = 2^{N-2}$, for a total of N qubits using one reset qubit. In general, $x = Md$, for PPA implemented using an auxiliary system of dimension d (also known as scratch system) and M reset qubits.
- **NOE2:** $x = 2$, for two qubits. Note that the cooling limits of the NOE protocol on 2 qubits is the same as the cooling limit of the PPA-HBAC on 3 qubits.
- **SRGN:** $x = 2^N - 1$, for a total of N qubits. Specifically, for the above SRG2, $x = 3$.
- **xHBAC:** $x = \infty$, giving $\epsilon_\infty^{\text{xHBAC}} \to 1$, for all values of ϵ_b.

5.3.7 Unified polarization evolution

Under the implementation of different AC protocols, the target qubit evolves, increasing its polarization exponentially with the number of rounds at different convergence speeds. We observe that the general expression of the polarization as a function of the number of rounds k and convergence rate r is given by

$$\epsilon_k\left(r, \epsilon_\infty, \epsilon_b\right) = \epsilon_\infty - r^k\left(\epsilon_\infty - \epsilon_b\right), \tag{5.26}$$

where ϵ_b is the polarization of the bath and the parameters r and ϵ_∞ depend on the protocol used as follows:

- **PPA3:** $r = \dfrac{1-\epsilon_b^2}{2}$, with $\epsilon_\infty = \dfrac{2\epsilon_b}{1+\epsilon_b^2}$, for three qubits.
- **NOE2:** $r = \dfrac{1}{2}$ with $\epsilon_\infty = \dfrac{2\epsilon_b}{1+\epsilon_b^2}$, for two qubits.

- **SRG2:** $r = \frac{1}{2}\left(\frac{1-\epsilon_b^2}{1+\epsilon_b^2}\right)$ with $\epsilon_\infty = \tanh\left[3 * \text{arctanh}(\epsilon_b)\right] = \epsilon_b \frac{3+\epsilon_b^2}{1+3\epsilon_b^2}$, for two qubits.
- **xHBAC1:** $r = e^{-\beta_b E} = e^{-2\text{arctanh}(\epsilon_b)}$ with $\epsilon_\infty = 1$.

This unified view reveals an interesting fact that *all* AC protocols under consideration here cools down the target qubit exponentially fast towards the (different) asymptotic states, in terms of the number of cooling rounds.

5.4 Cooling efficiency of HBAC protocols

Efficiency analyses for classical cooling protocols were of both theoretical and practical interest. These analyses both provide an upper limit for the achievable efficiency and suggest aspects in practical protocols which can be optimized. Conventionally, the focus for developing quantum cooling protocols have largely been to improve the cooling limit instead of efficiency. Meanwhile, as the scale of quantum processors gets larger, efficiency considerations will start to play a more significant role. In the remaining sections of this chapter, we analyze two different notions of efficiencies in quantum thermodynamics, and compare different HBAC protocols in terms of these efficiencies. We then propose a few modified HBAC protocols which have better energy efficiency compared to the original ones we previously analyzed.

A note on notation before going into the discussions: in this section we are frequently working with differences, such as the change in energy of a given system. Unless otherwise stated, we will always use the Greek letter Δ to denote the net change of a state function for a process. Our convention for the net change is "final - initial", i.e., $\Delta A = A_f - A_i$ where A_f is the final value of some quantity A after the process, and A_i is the initial value of A before the process.

5.4.1 Work cost and the coefficient of performance

The first and primary figure of merit of interest concerns the amount of work input during a protocol, and is inspired by the classical coefficient of performance (CoP). A classical cooling engine operates between a cold bath and a hot bath, pumping energy out of the cold bath and transferring it to the hot bath. For any operation performed by the engine, the CoP provides a measure of how effective it is in removing heat from the cold bath,

giving the ratio of the desired effect (the heat removed from the target) to the required input (work input to extract that energy). Specifically, the CoP, denoted by K, is given by

$$K \coloneqq \frac{-\Delta E_t}{W}, \tag{5.27}$$

where $-\Delta E_t$ is the total amount of heat removed from the target cold bath during the operation, and W is the total work input for the operation. Note that certain operations do not result in a net change of energy in the cold bath. For example, thermalizing the engine with the hot bath does not result in a net change in the cold bath (and moreover, does not require any work input either), so calculating the CoP for that process alone is not meaningful. For this reason, we will restrict to processes that result in a positive energy decrease in the target. As the definition of "operation" in the above varies, one can study the CoP of a single step during a protocol, or that of a complete cooling process.

We can analogously define the same quantum CoP for a quantum cooling procedure. HBAC protocols considered here all aim to cool down a finite quantum system instead of an infinite cold bath. The expression will be identical to Eq. (5.27), with $-\Delta E_t$ now denoting the energy decrease in the target system (which can now be interpreted as a finite cold bath), and W denoting the work input. Again, we will restrict to processes with $-\Delta E_t > 0$. In particular, two processes that are the most relevant for HBAC is either a single round, or the full cooling procedure up to a given round. We will denote the CoP for the n-th round by a lower-case $k(n)$, and the cumulative CoP *up to* the n-th round by an upper-case $K(n)$.

The next question is how k or K can be computed given a protocol. It is easy to recognize that since energy is a state function, the numerator is always given by

$$-\Delta E_t = E_{t,i} - E_{t,f}. \tag{5.28}$$

By further assuming that the Hamiltonians H_t, H_m both return to their initial values at the end of the unitary, the energy changes are simplified to

$$-\Delta E_t = \mathrm{Tr}[H_t(\rho_{t,i} - \rho_{t,f})], \tag{5.29}$$

where the subscripts i and f stand for initial and final states of the target.

On the other hand, the definition for W is slightly more subtle. An infinitesimal change in the internal energy can be split into two components, which may be identified as infinitesimal heat and work as [108]

$$dU = d\,\mathrm{Tr}[\rho H] = \mathrm{Tr}[d\rho\ H] + \mathrm{Tr}[\rho\ dH] \coloneqq \delta Q + \delta W, \tag{5.30}$$

where δ is used to denote path-dependent infinitesimal quantities. For a process where the evolution $\rho(t)$ and $H(t)$ from time 0 to τ is known, the mean heat and work during the process can be calculated as

$$\begin{aligned} Q &\coloneqq \int_0^\tau dt \, \mathrm{Tr}[\dot{\rho}(t) H(t)] \\ W &\coloneqq \int_0^\tau dt \, \mathrm{Tr}\Big[\rho(t) \dot{H}(t)\Big] \end{aligned} \tag{5.31}$$

with the first law being $\Delta U = \int dU = Q + W$.

In general, it can be nontrivial to find Q and W for a general quantum process. Here, we face a simpler situation where the analysis only needs to be performed for coherent control AC protocols. First, for the control subroutine, the target and machine undergo a unitary evolution, so the joint system evolves under the Schrodinger equation, $\dot{\rho_{t,m}}(t) = -\frac{i}{\hbar}[H_{t,m}(t), \rho_{t,m}(t)]$. This implies that Q vanishes:

$$Q = -\frac{i}{\hbar} \int_0^\tau dt \, \mathrm{Tr}[[H_{t,m}(t), \rho_{t,m}(t)] H_{t,m}(t)] = 0 \tag{5.32}$$

since trace is invariant under cyclic permutation. Therefore, for the control subroutine, the amount of work input is $W = \mathrm{Tr}[\rho(\tau)H(\tau) - \rho(0)H(0)]$, i.e., the total energy change of the full system correspond to work input. This agrees with the intuition that, since the system is isolated from the bath in the control step, the amount of heat released into the bath is 0. We thus have the following expression for the control subroutine:

$$W_{\mathrm{c}} = \Delta E_{t,\mathrm{c}} + \Delta E_{m,\mathrm{c}} \tag{5.33}$$

where the subscript c stands for control.

For the second part of the round, for the thermalization subroutine, we have established in Section 5.2.1 that the transformation on the system *can be* achieved by *some* energy-preserving unitary $V_{[t,m,b]}$ between the system and the bath. Thus, the thermalization step can be performed without work input in principle. Of course, it is possible to construct another unitary $V'_{[t,m,b]}$ that achieves the same transformation on the system, but with a nonzero work input. Therefore, the exact form of the thermalization unitary $V_{[t,m,b]}$ is non-unique and may be unknown in practice. So, instead, we will consider the lower bound situation for the work calculation, where the thermalization on the system *is* achieved by an energy-preserving unitary, so that the work input is indeed zero for the thermalization part

$$W_{\mathrm{th}} = 0 \tag{5.34}$$

where the subscript th stands for thermalization. This gives a recipe to compute the total lower bound work for a coherent cooling procedure consisting of any combination of control and thermalization subroutines as

$$W = \sum_{\alpha} W_{\mathrm{c},\alpha} \tag{5.35}$$

where α denotes different control steps of the protocol. As a result, the CoP for coherent control cooling protocols can be expressed purely in terms of (changes in) state functions. The CoP computed under this assumption will serve as an upper bound of the actual CoP in practice, due to the aforementioned reasons.

5.4.2 Comparing the CoP of the cooling protocols

We utilize the transition matrix formalism introduced in Section 5.3.1 to compute the efficiencies for different protocols in a systematic manner. This time we monitor the evolution of both the target system and the machine – thus, involving matrices of larger size. As before, we will only present the final results here, while the detailed derivations are given in Appendix C.4.

1. PPA3

   In PPA3, the first step is entropy compression, which is a global unitary on all 3 qubits. The second step is thermalization of both ancillary qubits with a bath at temperature β_b. We identify the energy change during the first step to be the total work input for each round of the protocol. For the n-th round, this is given by

   $$w_{\mathrm{PPA3}}(n) = E(p_b - \frac{1}{2})y^n, \ n \geq 1 \tag{5.36}$$

   with the same y defined previously. Using the target state evolution in Eq. (5.16), the energy decrease for the target during the n-th round is

   $$-\Delta e_{\mathrm{PPA3},t}(n) = E(p_b - \frac{1}{2})y^n, \ n \geq 1. \tag{5.37}$$

   We see that the CoP is

   $$k_{\mathrm{PPA3}}(n) = 1, \ n \geq 1 \tag{5.38}$$

   and the cumulative CoP is

   $$K_{\mathrm{PPA3}}(n) = 1, \ n \geq 1. \tag{5.39}$$

We note that the same CoP result has been presented in a recent work by Soldati et al. [100]. For PPA on more than 3 qubits, we use numerical simulation to compute the CoP for each round.

2. NOE2

In NOE2, the first step is a driven operation to the CMS state on qubit 2, and the second step is a 2-qubit Γ_2 relaxation. We identify the energy change during the CMS operation to be the work input, because it is achieved by applying a random rotation unitary to qubit 2 and averaging over the results in practice. The (average) work input during the n-th round is

$$w_{\mathrm{NOE2}}(n) = \begin{cases} E(p_2 - p_b)2^{1-n}, & n > 1 \\ E(p_b - \frac{1}{2}), & n = 1 \end{cases} \tag{5.40}$$

and the energy decrease for the target during the n-th round is

$$-\Delta e_{\mathrm{NOE2},t}(n) = E(p_2 - p_b)2^{-n} \tag{5.41}$$

for all $n \geq 1$. Using these, we calculate the CoP to be

$$k_{\mathrm{NOE2}}(n) = \frac{-\Delta E_t(n)}{W_{\mathrm{NOE}}(n)} = \begin{cases} \frac{p_2 - p_b}{2p_b - 1}, & n = 1 \\ \frac{1}{2}, & n > 1 \end{cases} \tag{5.42}$$

(note that p_2 and p_b are both functions of the bath temperature). The cumulative CoP up to round n is

$$K_{\mathrm{NOE2}}(n) = \frac{-\Delta E_t(n)}{W_{\mathrm{NOE}}(n)} = \frac{(p_2 - p_b)(1 - 2^{-n})}{(p_2 - \frac{1}{2}) - 2^{1-n}(p_2 - p_b)}. \tag{5.43}$$

3. SRG2

In the SRG2 protocol the first step is a σ_x gate on qubit 2, the second step is a Γ_2 relaxation on both qubits, and the third step is a Γ_1 relaxation on qubit 2. We identify the energy change during the first step (σ_x gate) to be the work input, given by

$$w_{\mathrm{SRG2}}(n) = E(2p_b - 1) \tag{5.44}$$

for all $n \geq 1$. In fact, this can be trivially obtained by noticing that the σ_x gate always flips the same thermal state at bath temperature in each round. This explains our

comment before that the total energy cost from applying the SRG2 many times is unbounded. Meanwhile, the energy decrease for the target during the n-th round is

$$-\Delta e_{\mathrm{SRG2},t}(n) = Ep_b(1-p_b)(2p_2-1)u^{n-1} \tag{5.45}$$

for all $n \geq 1$. The CoP of the n-th round is

$$k_{\mathrm{SRG2}}(n) = \frac{p_b(1-p_b)(2p_2-1)}{2p_b-1}(p_2+p_b-2p_2p_b)^{n-1} \tag{5.46}$$

for all $n \geq 1$. The cumulative CoP up to round n is

$$K_{\mathrm{SRG2}}(n) = \frac{p_b(1-p_b)(2p_2-1)}{2p_b-1}\frac{1-u^n}{n(1-u)}. \tag{5.47}$$

4. xHBAC1

In 1-qubit xHBAC, one first applies a σ_x gate, followed by a β-swap relaxation Λ_β. The work input corresponds to the energy change during the σ_x gate, given by

$$w_{\mathrm{xHBAC1}}(n) = E(1-2s^{n-1}(1-p_b)). \tag{5.48}$$

The energy decrease for the target during the n-th round is

$$-\Delta e_{\mathrm{xHBAC1},t}(n) = E(1-p_b)(1-s)s^{n-1} \tag{5.49}$$

for all $n \geq 1$. The CoP of the n-th round is given by

$$k_{\mathrm{xHBAC1}}(n) = \frac{-\Delta E_t(n)}{W_{\mathrm{xHBAC1}}(n)} = \frac{(1-p_b)(1-s)}{1-2s^{n-1}(1-p_b)}s^{n-1} \tag{5.50}$$

for all $n \geq 1$, where $s = e^{-E\beta_b}$. The cumulative CoP up to round n is

$$K_{\mathrm{xHBAC1}}(n) = \frac{(1-p_b)(1-s)(1-s^n)}{n(1-s)-2(1-p_b)(1-s^n)} \tag{5.51}$$

We first discuss the per-round CoP, k, of these protocols. The PPA3 exhibits the simplest behavior with a constant $k_{\mathrm{PPA3}} = 1$. The NOE has an almost constant behavior, with $k_{\mathrm{NOE}} = 1/2$ for all rounds except the first one. Protocols like these are considered to be the most energetically effective, since as the target is cooled, the work input required at each round also decreases proportionally with the decrease in target energy per round. Meanwhile, the SRG2 and xHBAC1 both have k's that decreases exponentially with the

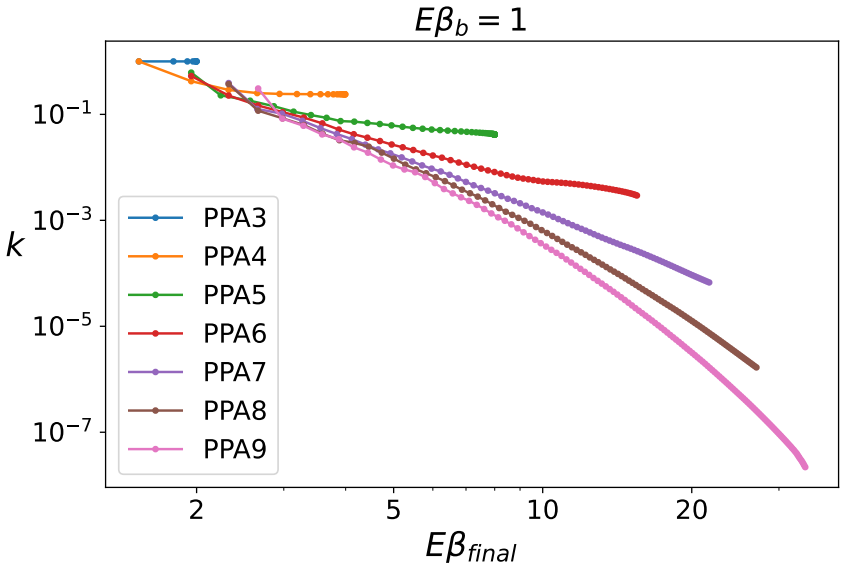

Figure 5.2: A plot of per-round CoP, k_{PPA3}, plotted against the final inverse temperature $E\beta_{\text{final}}$ for selected PPA protocols.

round n. This means that as the qubit is cooled, the amount of energy extracted from the target becomes exponentially small compared to the work input for that round. Therefore, protocols like these becomes more energetically inefficient for deeper cooling.

Next we turn to PPAN with $N \geq 3$, where we numerically study their efficiencies. This is plotted in Fig. 5.2. Each point shows the per-round CoP on the y-axis, versus the scaled final temperature ($\beta_{\text{final}}E$) of the target qubit on the x-axis, so that we go from left to right as the target cools down. The first point (on the leftmost) in each curve represents the cumulative CoP after the first round of cooling, and one follows the line towards the right as cooling proceeds further. Note that except the PPA3 which has a constant k, all larger instances of PPA have decreasing k's as cooling progresses. This shows that larger PPA protocols also becomes inefficient at extracting energy from the target qubit as it is being cooled down.

Although the per-round CoP decreases for most of the protocols considered here, the magnitude of the work input may also decrease as cooling proceeds (while the heat removed from target decreases even faster), and will contribute less and less to the total amount of work inputted. Therefore, the more practically meaningful figure of merit for determining the total energy efficiency is K, the cumulative CoP. We plot the cumulative CoP for selected protocols in Fig. 5.3. From top to bottom are the results when starting from a

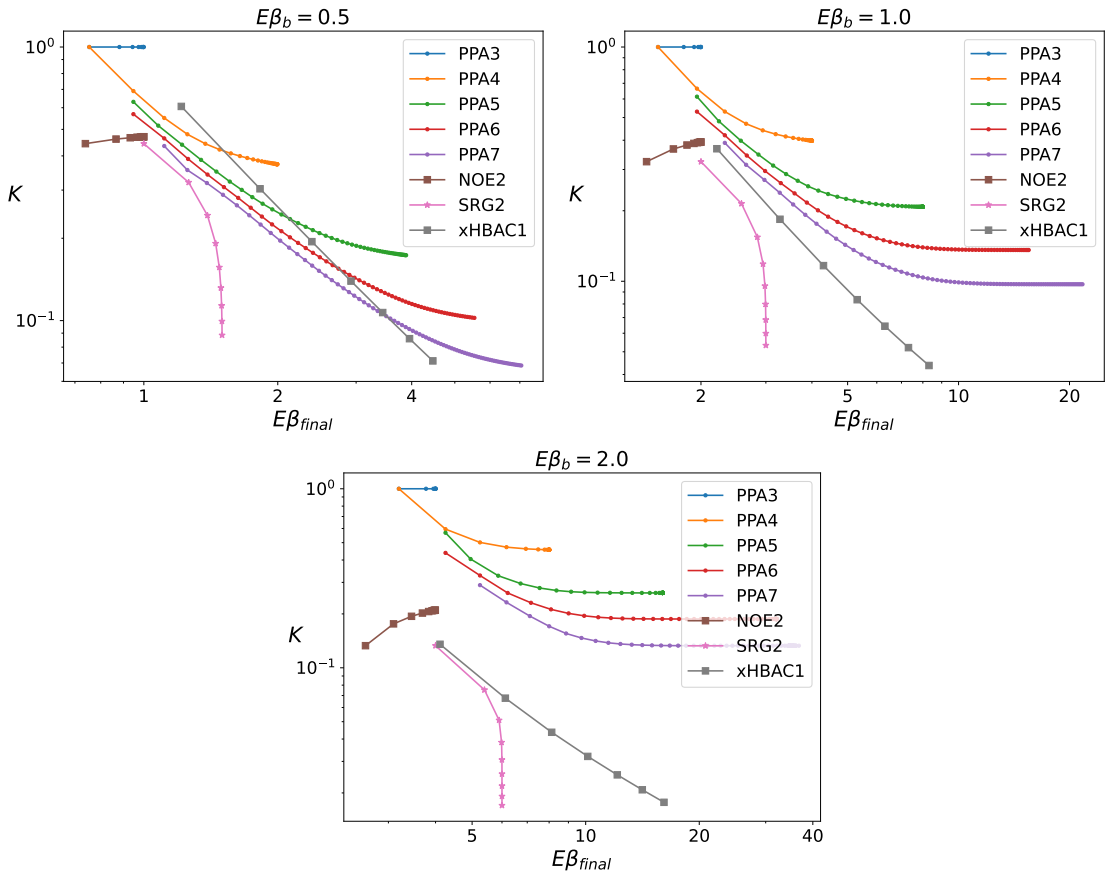

Figure 5.3: Comparison of cumulative CoP between HBAC protocols under different initial bath temperatures, plotted against the final inverse temperature $E\beta_{\text{final}}$.

high ($\beta_b E = 0.5$), intermediate ($\beta_b E = 1$), and low ($\beta_b E = 2$) bath temperature (scaled by the qubit's energy gap E), where β_b is the bath inverse temperature and E represents the qubit energy gap[1]. Recall from Section 5.3 that the xHBAC has a cooling limit of 0 (the lowest possible), whereas all other protocols under consideration here have finite cooling limits. Here, a higher CoP implies better energy efficiency in extracting a certain total amount of heat from the target qubit.

Several important observations are listed as follows:

1. The CoP for both the xHBAC1 and the SRG2 goes towards 0 as cooling proceeds,

[1]As a reference, for a transmon qubit with a transition frequency of 5GHz at temperature T_b =20mK, $\beta_b E \approx 12$ which corresponds to a very low temperature.

whereas the other protocols have a finite CoP as we approach equilibrium;

2. Both tending towards $K = 0$ as the number of cooling rounds increase, the xHBAC1 has an equilibrium temperature of 0 (corresponding to $E\beta_{\text{final}} = \infty$) whereas SRG2 has a finite equilibrium temperature;

3. The CoP of PPA protocols plateau as we approach equilibrium, and PPA protocols with more qubits are capable of cooling to lower temperatures at the cost of having a lower final cumulative CoP;

4. Both the PPA3 and NOE2 achieve the same final target temperature, but the PPA3 has a higher CoP than the NOE2;

5. The relative magnitude of K between different protocols vary as the bath temperature changes.

The diverging behavior of both the xHBAC1 and the SRG2 can be easily understood from the construction of their protocols, which involves applying σ_x gates at every round of cooling. Meanwhile, the optimized relaxation scheme of xHBAC1 allows it to achieve final temperature 0 exponentially fast, which SRG2 cannot do. On the other hand, the PPA appears to be a milder family of protocols, which cools to a finite temperature at a finite energy cost. We can also see that the constant $K = 1$ performance of PPA3 is a special case rather than a general behavior for all PPA protocols.

The temperature dependence of HBAC is also interesting. It appears that as temperature lowers (going from top to bottom in Fig. 5.3), the PPA becomes more advantageous energetically compared to the other protocols. Conversely, protocols like the xHBAC1 becomes more favorable when the system starts at higher initial temperatures, because it may cool to the same temperature as the PPA's in less number of steps and with less total energy input, under certain regimes. Indeed, comparing the energy cost of protocols in Fig. 5.3 can be done by drawing a vertical line at a certain final target temperature which we wish to achieve. For example, we can see that at the final temperature achieved by the first round of xHBAC1 cooling, the point representing xHBAC1 lies above the ones for the PPA protocols which are capable of cooling to a similar temperature. The second round of xHBAC1 lies below PPA4, but above PPA5 and PPA6. On the other hand, the line for xHBAC1 lies below all PPA protocols under consideration here, for lower initial temperatures $\beta_b E = 1$ and $\beta_b E = 2$. In summary, selecting the best protocol in practice requires one take into consideration both time and energy cost.

5.5 Heat production and the Landauer Ratio

In addition to the work required, another equally interesting aspect is to look at the heat released during a cooling procedure. This is particularly relevant for certain quantum processing platforms that are sensitive to the environmental temperature, such as superconducting qubits. In these platforms, one must minimize heating up of the environment while performing information processing. Therefore, protocols that produces less heat will be more desirable to be used in practice, motivating our study on heat production.

5.5.1 Minimal heat production

Landauer [43] was the first to realize that classical information erasure (more precisely, logically irreversible operations) must come with heat released to the environment. This phenomenon has been referred to as Landauer's Principle (LP) in later studies. The original arguments can be summarized as follows. Consider a single spin which acts as a memory. The initial state of the spin is taken to be unknown; hence statistically it is equally probable to be in either spin up or spin down. Invoking the Boltzmann definition of statistical entropy, $S = k\log(W)$ where W is the number of possible microstates a system can be in, one sees that if one is to perform a *reset* operation on the spins, then W is reduced to half of its initial value. Therefore, the entropy of the memory has decreased by $k\log(2)$. Since the entropy of a closed system (memory plus environment) cannot decrease, one can then infer that the entropy of the environment must have increased by at least $k\log(2)$.

It is expected that a similar argument should hold when the system of interest is of quantum nature. The analogy goes that unitary evolutions are reversible and does not change the entropy, therefore is analogous to reversible operations in classical thermodynamics. Meanwhile, a non-unitary transformation may result in a finite entropy change, and a negative entropy change is analogous to classical information removal. Therefore, the goal of formulating a quantum Landauer's Principle has largely been focusing on bounds of heat production in terms of entropy change in the target system. In particular, note that the goal of algorithmic cooling is to drive the target system towards a fixed state (the ground state in our examined cases), and necessarily accompanies an entropy change. Below we will review two inequalities of this type, then propose an efficiency measure for algorithmic cooling based on entropy considerations.

In the first scenario, there exists a target system of interest initially in an arbitrary state ρ_t, and a bath initially in a thermal state ρ_b^{th} with inverse temperature β, jointly in a tensor product state $\rho_{t,b} = \rho_t \otimes \rho_b$. The full state of system plus bath is thermodynamically

closed. The system and bath then undergo a joint unitary transformation denoted by $U_{t,b}$, taking them to a final state $\rho'_{t,b}$. The work by Reeb and Wolf [109] showed that (see Appendix C.1 for a derivation)

$$\beta\Delta E_b = -\Delta S_t + I(t' : b') + D(\rho'_b||\rho_b) \tag{5.52}$$

where $\Delta E_b = \mathrm{Tr}[(\rho'_b - \rho_b)H_b]$ is the energy increase of the bath, $-\Delta S_t = -(S(\rho'_t) - S(\rho_t))$ is the entropy decrease of the system, $I(t' : b')$ is the quantum mutual information between target and bath after the evolution, and $D(\rho'_b||\rho_b)$ is the quantum relative entropy between initial and final states of the bath,

$$\begin{aligned} I(t' : b') &= S(\mathrm{Tr}_b[\rho'_{t,b}]) + S(\mathrm{Tr}_t[\rho'_{t,b}]) - S(\rho'_{t,b}) \\ D(\rho'_b||\rho_b) &= -S(\mathrm{Tr}_t[\rho'_{t,b}]) - \mathrm{Tr}\big[\mathrm{Tr}_t[\rho'_{t,b}]\log\rho_b\big] \end{aligned} \tag{5.53}$$

Due to the non-negativity of I and D, this result implies that the energy increase in the bath (after scaled by the inverse temperature) will be lower bounded by the entropy decrease in the system.

In the second scenario, consider a single target system undergoing a general CPTP evolution $\mathcal{E}_t$. When the system is open, $\mathcal{E}_t$ may be a non-unitary evolution. In general we may extend this CPTP map to be a unitary map on a larger Hilbert space, but since the extension is non-unique, the exact evolution on the ancillary system is also non-unique. Interestingly, a Landauer-type inequality can be derived for a subset of such processes, by only examining the change in the system alone. For a general CPTP evolution on the target system where $\rho'_t = \mathcal{E}_t(\rho_t)$, one can derive (see Appendix C.2) the following equality:

$$-\beta\Delta E_t = -\Delta S_t + (D(\rho_t||\rho_t^{th}) - D(\rho'_t||\rho_t^{th})) \tag{5.54}$$

where, as before, $-\beta\Delta E_t = -\mathrm{Tr}[(\rho'_t - \rho_t)H_t]$ is the negative energy increase in the system, $-\Delta S_t = S(\rho_t) - S(\rho'_t)$ is the entropy decrease in the system, and ρ_t^{th} is the system's thermal state.

Interpreting quantum relative entropy as a distance measure, one sees that for processes where the final state is closer to ρ_t^{th} than the initial state, the difference term in the bracket on the right is positive. For example, this will hold true for a thermalization process towards ρ_t^{th}. Therefore, for processes where the state is becoming closer to the thermal state, Eq. (5.55) states that the decrease in energy (scaled by the inverse bath temperature) will be larger than the decrease in entropy for the system. We therefore refer to Eq. (5.54) as the LP for thermalization processes.

To relate to energy change in the bath, a common assumption (e.g., in collision models [110]) is that the actual joint evolution of the system and bath is energy-preserving.

This implies that $\beta\Delta E_b = -\beta\Delta E_t$, so that the energy *increase* in the bath can now be related to the entropy decrease in the system as

$$\beta\Delta E_b = -\Delta S_t + (D(\rho_t||\rho_t^{th}) - D(\rho_t'||\rho_t^{th})). \tag{5.55}$$

We have seen that for both scenarios considered above, one can establish Landauer-type inequalities which lower bounds the amount of heat increase in the bath by the amount of entropy decrease in the system. Since the goal of HBAC is exactly reducing entropy of the system, it is well motivated to define an "entropic efficiency" for a given protocol by considering the energy increase *outside* the target system. A slight complication for the coherent control cooling schemes is that outside the target, there exists both a machine and a heat bath where the energy can be directed to. While the bath is assumed to be large such that it remains in a thermal state at fixed temperature, the machine may not be thermal after the first round of cooling, preventing a direct application of Eq. (5.52) for the "control" step. Nonetheless, the following quantity

$$R_L \coloneqq \frac{\beta_b(\Delta E_b + \Delta E_m)}{-\Delta S_t} \tag{5.56}$$

where β_b is the inverse temperature of the bath, is well-defined for any nontrivial cooling procedure with $\Delta S_t < 0$. We call this the *Landauer Ratio* (LR) and propose it as a measure of entropic efficiency, by observing that for the two settings discussed previously, this ratio is lower bounded by 1, and a lower value implies more entropy extracted from the target system per amount of energy released. Like the CoP, we can study the LR for a particular round where the difference (Δ) quantities above refer to the change for a round, or the cumulative LR where the difference quantities refer to the cumulative change. We will denote the prior by r_L, and reserve R_L as the cumulative LR, noting that the two quantities coincide for a single round of cooling.

Eq. (5.56) is again an operational definition, like the CoP, because it is a ratio of cost over desired effect. What we have in mind is the coherent cooling scheme with some predetermined target system, machine, and heat bath. Interestingly, if we change the definition of what a target is, the LR may reveal further insights about the evolution of a cooling protocol. A particularly insightful example is when we focus on a single unitary process, e.g., the compression step in PPA. Specifically, we know that the purpose of compression is to concentrate entropy on the two relaxation qubits, which release those entropy to the bath during the thermalization step. How efficient is this concentration operation? We first note the following theorem which is proved in Appendix C.3.

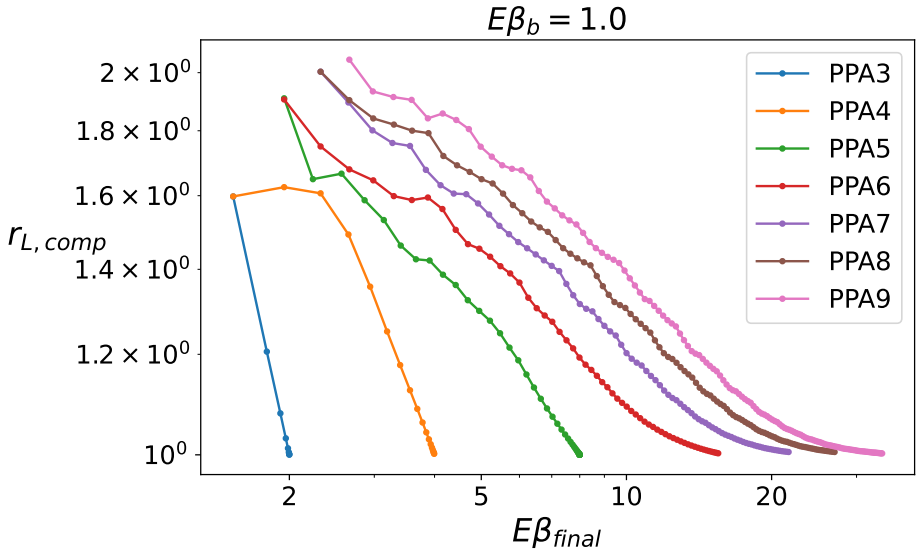

Figure 5.4: A plot of per-step LR, showing the ratio of the energy increase in the last two reset qubits over the entropy reduction in the target plus computational qubits, for PPA protocols. This illustrates the asymptotic optimality of PPA for cooling the computational qubits.

Theorem 5.5.1. *Consider a finite dimensional system with initial state* $\rho_{xy} = \rho_x \otimes \rho_y^{th}$ *where* y *has an inverse temperature* β*. Let* $T(\rho, \sigma) \coloneqq \frac{1}{2}\|\rho - \sigma\|_1$ *denote the trace distance between* ρ *and* σ*. Then for a unitary process* $\mathcal{U}(\rho_{xy}) = \rho'_{xy}$ *where* $T(\rho'_x, \rho_x) \leq \frac{1}{2e}$ *(where* e *is the natural log base) and* $-\Delta S_x > 0$*,*

$$\frac{\beta \Delta E_y}{-\Delta S_x} \geq 1 + \frac{\gamma_y^2 + \lambda_{xy}^2}{\epsilon_x \log\left(\frac{d_x}{2\epsilon_x}\right)} \tag{5.57}$$

where $\epsilon_x = T(\rho'_x, \rho_x)$*,* $\gamma_y = T(\rho'_y, \rho_y)$*,* $\lambda_{xy} = T(\rho'_{xy}, \rho'_x \otimes \rho'_y)$*, and* d_x *is the dimension of subsystem* x*.*

Clearly, this is exactly the setting described prior to Eq. (5.52), where a bipartite unitary process occurs from an initially tensor product state, with the second subsystem in the thermal state. Therefore we expect the ratio on the LHS to be greater than or equal to 1. Meanwhile, this theorem provides some insights into when the LHS may approach the theoretical minimum of 1. Specifically, this has to do with how "drastic" the unitary is, how "evenly distributed" the change is between subsystems x and y, and how much

mutual information has been created. We observe that the ratio would approach 1 for a unitary that satisfy the following:

1. it only results in a marginally small variation in the joint state ρ_{xy} (when measured by the trace distance T), so that all of ϵ_x, γ_y, and λ_{xy} are small;

2. The change in subsystem x, ϵ_x, is less than or on the same order as the change in subsystem y, γ_y, and λ_{xy} which is proportional to the square root of the mutual information created by the unitary.

One could see that if the above are true, then the numerator in the extra term on the RHS goes quadratic in the small quantities, while the denominator only goes as $\epsilon \log(\epsilon)$, so the ratio would go towards 0. Of course, this is only a rough intuition, and we shall now see how this lower bound is asymptotically achieved in the PPA protocol. In addition, we note here that a tighter but more sophisticated bound has been derived in [109]. The significance of our result is to give an intuition for when the lower bound of 1 may be saturated, in terms of the change and correlation quantities as a result of the unitary.

To relate back to PPA, notice that immediately before the compression step of each round, the first $n-2$ qubits are in a tensor product state with the last two qubits, which have just been refreshed to the bath temperature. Therefore, we may take the first $n-2$ qubits to be subsystem x and the last two qubits to be subsystem y, so that Eq. (5.57) is satisfied. Comparing with Eq. (5.56), one sees that the LHS of Eq. (5.57) is precisely R_L where $\Delta E_b = 0$ (since the system is closed during the unitary), if we identify the target system to be the first $n-2$ qubits, and the machine to be the last 2 qubits. We will refer to the RL with this identification as $r_{L,\text{comp}}$ since it is a per-round LR where the new "target" system consists of all *computational* qubits (a historical term referring to all except the relaxation qubits). What this represents is the ability of PPA to extract entropy from the computational qubits, and concentrate them into the two relaxation qubits. A ratio close to 1 implies maximum entropy extracted from the computational qubits, per energy increase in the relaxation qubits, by the compression unitary. We plot this quantity $r_{L,\text{comp}}$ in Fig. 5.4, for different PPA protocols. As can be seen from the plot, all the protocols indeed show a consistent behavior where the $r_{L,\text{comp}}$ goes towards 1 as cooling proceeds towards equilibrium. Therefore, the PPA compression subroutine is asymptotically optimal in terms of extracting entropy from the computational qubits.

5.5.2 Relating to the coefficient of performance

Having operationally defined the Landauer Ratio, we then ask the same question of whether it can be computed given a protocol. The difficulty lies in computing ΔE_b for a general process, since the bath degree of freedom is typically not known in full detail, whereas the protocol only dictates how the system (target and machine) evolves. Meanwhile, we will see that for coherent cooling protocols, the LR can be calculated in principle. For the control subroutines, the joint system of target plus machine is closed, and undergoes a unitary transformation. The quantity ΔE_b is thus 0 for these processes. Meanwhile, for the thermalization subroutines, we again assume that it occurs with a joint energy-preserving unitary on the system and bath, so that $\Delta E_b = -(\Delta E_t + \Delta E_m)$. Therefore, for coherent control protocols,

$$\Delta E_b = -(\Delta E_{t,\alpha} + \Delta E_{m,\alpha}) \tag{5.58}$$

where α denotes all thermalization steps during the protocol.

It is instructive to see how the LR relates to the CoP defined in the previous section. Since the full system of (target, machine, bath) is closed, energy conservation implies that

$$\Delta E_t + \Delta E_m + \Delta E_b = W. \tag{5.59}$$

Replacing W in the definition of the CoP gives

$$K = \frac{-\Delta E_t}{W} = \frac{-\Delta E_t}{\Delta E_t + \Delta E_m + \Delta E_b}, \tag{5.60}$$

and rearranging the terms gives

$$\frac{\Delta E_b + \Delta E_m}{\Delta E_t} = -\left(\frac{1}{K} + 1\right). \tag{5.61}$$

Recalling the definition of R_L, we have

$$R_L = \frac{\beta_b(\Delta E_b + \Delta E_m)}{-\Delta S_t} = \frac{-\beta_b \Delta E_t}{-\Delta S_t}\left(\frac{1}{K} + 1\right). \tag{5.62}$$

This above equation shows how K and R_L are related to each other. For a given cooling procedure with known initial and final states, knowing K immediately gives R_L, and vice versa. Moreover, it can be seen that for a given process with fixed initial and final states, a higher CoP corresponds to a lower LR, confirming our previous intuition about what "more efficient" means under both measures. An equivalent way of saying this is that a higher work input necessarily accompanies more heat released into the (machine and) bath.

When the target system is a qubit with energy gap E, the cumulative LR up to the n-th round is simplified to

$$R_L(n) = \frac{\beta_b E(p_t(n) - p_t)}{H(p_t) - H(p_t(n))}\left(\frac{1}{K(n)} + 1\right). \tag{5.63}$$

In Fig. 5.5, we show a plot of the cumulative LR for the HBAC protocols under consideration, starting from different initial temperatures. Comparing with Fig. 5.3 at the same bath temperature, we may verify our result in Eq. (5.62) that protocols with a lower CoP necessarily have a higher LR, so both are meaningful figures of merit for efficiency. At $E\beta_b = 1$ and above, the PPA protocols under consideration releases less amount of heat into the environment when cooling to the same final target temperature. Meanwhile, at a higher temperature of $E\beta_b = 0.5$, there are certain regions in which using the xHBAC1 is more efficient than the PPAs, when both can reach the same final temperature.

Interestingly, note that both the PPA3 and NOE have LR's that decrease with further cooling. This can be understood from Eq. (5.62) that, for K being fixed, the LR is proportional to the ratio of energy change over entropy change in the target system. When the target is a single qubit, both changes are functions of a single parameter p. Since the rate of decrease in energy (linear in p) is faster than the rate of decrease in entropy (log in p) as $p \to 1$, the ratio will grow smaller as p increases. Therefore, for protocols like the PPA3 and NOE which have constant or near-constant K, the change in LR is dominated by this effect and decreases with the cooling progress.

5.6 Designing protocols with better energy efficiency

We noted previously that conventionally, analyses on HBAC have mostly focused on the cooling speed and asymptotic achievable temperature. This drives searches for faster protocols or ones which can cool to lower temperatures, but largely ignores the amount of energy cost. Having taken energy aspects into consideration for HBAC protocols, we now study possible methods to improve their energy performance. To be more concrete, we will first look at the compression step in PPA, and show how it can be improved energetically. Later, we will see how this idea leads to more efficient variants of the PPA and xHBAC protocols.

Consider the same setup as the PPA which consists of a target system with X energy levels and a machine with Y energy levels, with the possibility of having degeneracies. For simplicity, we assume that both the target and the ancillas are diagonal in the energy

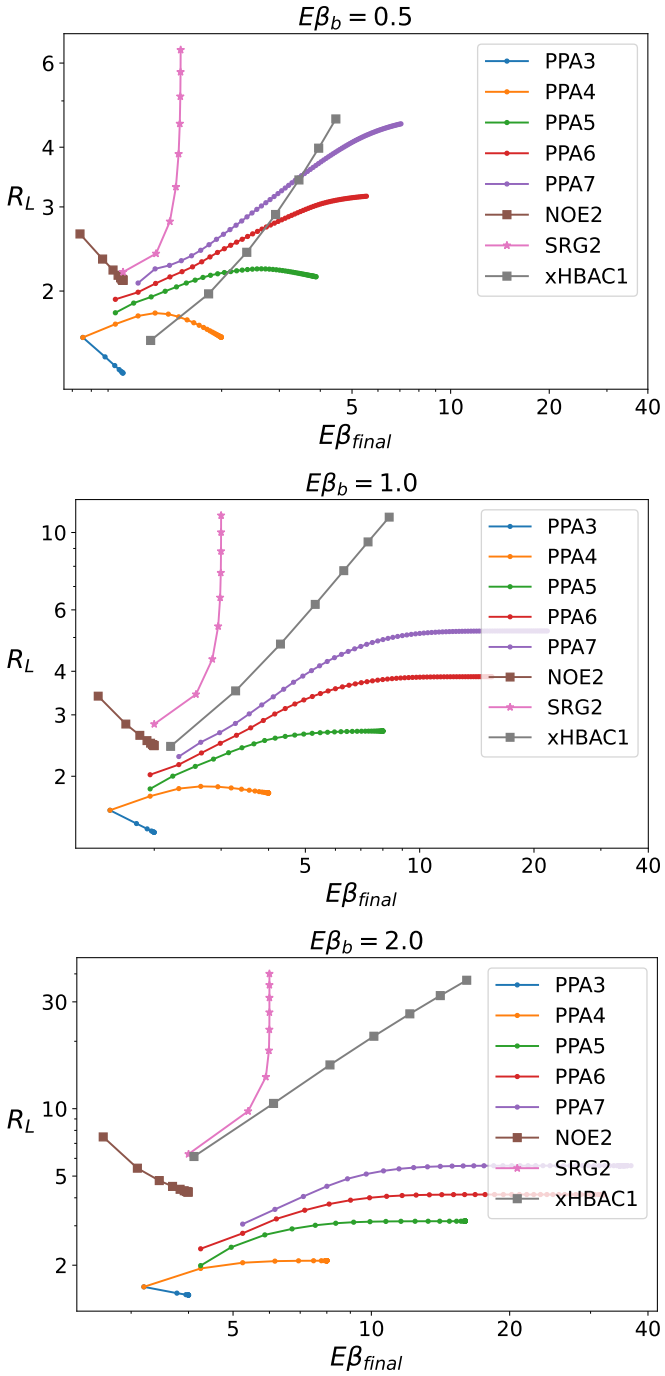

Figure 5.5: Comparison of cumulative LR between HBAC protocols under different initial bath temperatures, plotted against the final inverse temperature $E\beta_{\text{final}}$.

eigenbasis. In such systems, a unitary corresponds to permutations of the diagonal elements (up to a global phase), so it is sufficient to consider the simpler problem of optimizing over all possible permutations. We ask the question: among the unitaries that achieve maximal energy reduction in the target, which is the most energy-efficient one? Note that here we do not try to optimize the CoP, and instead, only optimize within the set of unitaries that achieve the most cooling possible given the initial system.

5.6.1 Improved PPA

We first examine the original PPA entropy compression, which performs a descending SORT to the diagonal elements with respect to the lexicographic basis order. For simplicity, we will assume that all qubits have the same energy gap. Recalling that the target system in PPA is a single qubit, so $X = 2$. Therefore, the descending SORT always achieve the maximal energy reduction in the target, since it maximally populates the lowest energy ground state. Therefore, one only needs to look at the energy cost of this operation and check whether there exists a more efficient one, whose existence we will show below.

For PPA3, the states in the $|0\rangle$ subspace of the first qubit include

$$000,\ 001,\ 010,\ 011,$$

and similarly for the $|1\rangle$ subspace. We see that these states are already ordered in terms of energy assuming the same energy gap between all qubits, so a descending SORT would result in the lowest energy configuration within each subspace of the target.

In contrast, for $n = 4$ qubits, the states in the $|0\rangle$ subspace of the first qubit are

$$0000,\ 0001,\ 0010,\ 0011,\ 0100,\ 0101,\ 0110,\ 0111$$

and similarly for the $|1\rangle$ subspace. These states are not fully ordered in terms of energy: specifically the state $|0011\rangle$ comes before $|0100\rangle$, but has a higher energy. Consider the following new ordering

$$0000,\ 0001,\ 0010,\ 0100,\ 0011,\ 0101,\ 0110,\ 0111$$

which is now ordered in ascending energy within the target $|0\rangle$ subspace. Compared with the original compression, a descending SORT with respect to this new basis order achieves the same cooling in the target, but uses potentially less energy since the low-energy state 0100 is now more populated than the high-energy state 0011.

We now go back to the more general setting, where we wish to reduce the energy in an X-level target system as much as possible with a single unitary. The available ancillary system is a machine with Y energy levels, such that the full system has a total of XY levels. Following the previous intuition, we propose the following compression procedure:

1. Label the target subspaces in non-decreasing order, such that $E_{t,1} \leq \cdots \leq E_{t,X}$;
2. Label the machine subspaces in non-decreasing order, such that $E_{m,1} \leq \cdots \leq E_{m,Y}$. This results in an ordering for the full system as

$$|1,1\rangle, \ldots, |1,Y\rangle, |2,1\rangle, \ldots, |2,Y\rangle, \ldots, |X,Y\rangle$$

   where the first index labels the target and the second labels the machine;
3. Perform a descending SORT with respect to this basis ordering, so that the population profile after the unitary satisfies:

$$p_{1,1} \geq \cdots \geq p_{1,Y} \geq p_{2,1} \geq \cdots \geq p_{X,Y}.$$

The above procedure ensures maximal energy reduction in the target, while requiring the least amount of work input, in a single unitary step. It is interesting to see how this affects the performance of PPA, when combined with non-unitary relaxation steps in the original protocol. A round in our new protocol consists of: 1) a control subroutine using the above compression procedure, and 2) a thermalization subroutine where all qubits higher than the bath temperature are relaxed to the bath temperature T_b. Interestingly, starting from an all-thermal state at T_b, we find that all qubits except the target becomes warmer than the bath, and are therefore all relaxed after the new compression. According to [107], relaxing more qubits implies a lower achievable temperature of the target. In particular, we show in Appendix C.5 that the asymptotic achievable temperature using this new procedure is $T_b/(n-1)$ with a total of n qubits, compared to $T_b/2^{n-2}$ in the original PPA. Consequently, a fair CoP comparison is between the new scheme using $2^{n-2}+1$ qubits and the original one using n qubits, such that both achieve the same asymptotic temperature. This is shown in Fig. 5.6 where we compared the first three instances, and see indeed that the improved protocol has a higher cumulative CoP than the original one.

We end with a brief discussion on the above results. Compared with the original scheme, the new scheme creates a more "averaged" compression, such that the temperatures of qubits other than the target are more evenly distributed. This "softer" compression uses less energy, but also creates less temperature gradient among the machine bits. Thus, this

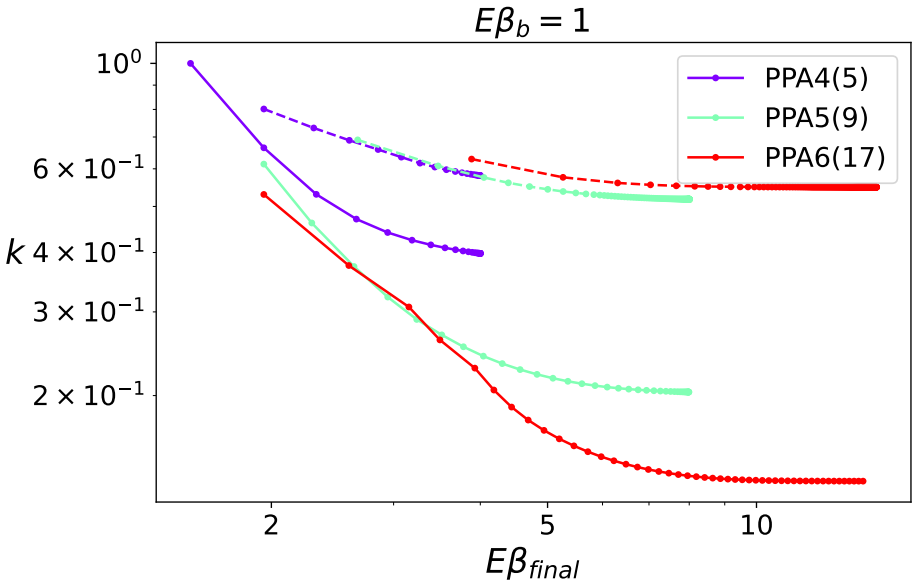

Figure 5.6: Comparison of CoP between the original (solid lines) and energetically improved (dashed lines) PPA protocols. The improved PPA uses a larger number of qubits (indicated in the brackets in the plot legend) compared with the original PPA in order to cool the target to the same asymptotic temperature, but uses less energy in total to achieve so.

change the structure of the asymptotic state of the full system: while that of the original PPA has the following temperature profile

$$\frac{T_b}{2^{n-2}},\ \frac{T_b}{2^{n-3}},\ \ldots,\ \frac{T_b}{2},\ T_b,\ T_b$$

for each qubit, the new PPA has

$$\frac{T_b}{n-1},\ T_b,\ T_b,\ \ldots, T_b.$$

One can see this as a way to trade spatial cost with energetic cost: in order to create the temperature gradient in the original PPA, it is necessary to cool the first $n-2$ qubits to temperatures lower than the bath. The new PPA avoids such unnecessary cooling, by using a larger total number of qubits but with much lower energy cost.

5.6.2 Improved xHBAC

Next, we see how the same principle can be used to obtain an improved realization of the xHBAC1 protocol, compared with the original proposal which uses a bosonic bath interacting with the qubit to implement the Λ_β operation [94]. While the original proposal can already cool down the target qubit exponentially fast, as can be seen from Fig. 5.3, it is not thermodynamically optimal in terms of energy consumption, since the total amount of work input is unbounded as more cooling rounds are performed. In fact, it is well known that using an ancillary system with an unbounded spectrum, a target system can be cooled asymptotically to zero temperature using unitary control. Using this same ancillary system, we can explicitly construct the asymptotic final state of the combined system and calculate the work cost.

Again, the target qubit and the harmonic oscillator start in a tensor product state, each being in a thermal state at the same temperature T_b. This corresponds to a situation with $X = 2$ and $Y = \infty$. We will consider the "resonant" scenario where $E_t = E_a \coloneqq E$, since this is the necessary condition for Λ_β to be a thermal operation [94]. Then, the thermal population of a state $|m, n\rangle$ only depends on the "total quantum number" $m + n$: for example, the population of $|1, 0\rangle$ and that of $|0, 1\rangle$ will be identical. Moreover, the ordering for the full system will be given by

$$|0, 0\rangle\,,\ |0, 1\rangle\,,\ |0, 2\rangle\,,\ |0, 3\rangle\,,\ \ldots$$

where the first index labels the qubit and the second labels the bosonic mode. Thus, the improved protocol puts the largest diagonal elements in the joint density matrix into the $|0\rangle$ subspace of the qubit, in decreasing order of magnitude with the increasing energy levels of the harmonic oscillator $|n\rangle$. Due to the infinite energy levels present in the harmonic oscillator, the final state would have all the significant components in the qubit $|0\rangle$ subspace, which can be formally seen by first considering a truncated Y-level machine, and taking the limit $Y \to \infty$.

Therefore, after the optimized compression, the final state would have the $|0, 0\rangle$ population unchanged, then the original $|0, i\rangle$ population "copied twice" into $|0, 2i - 1\rangle$ and

$|0, 2i\rangle$, for $i = 1 \ldots \infty$. The energy of this final state is given by

$$\begin{aligned} E_{\text{fin}} &= \sum_{i=1}^{\infty} \frac{1}{Z_a} p_0 e^{-i\beta E}(2i + 2i - 1)E \\ &= \sum_{i=1}^{\infty}(1 - e^{-\beta E})(\frac{1}{1 + e^{-\beta E}})e^{-i\beta E}(4i - 1)E \\ &= \frac{1 + 3e^{\beta E}}{e^{2\beta E} - 1}E \end{aligned} \tag{5.64}$$

where the ground state $|0, 0\rangle$ has been set to have energy 0. The energy of the initial state is simply the energy of the qubit plus that of the harmonic oscillator,

$$E_{\text{ini}} = \delta_b E + \sum_{i=0}^{\infty} \frac{1}{Z_a} e^{-i\beta E} iE = \frac{2e^{\beta E}}{e^{2\beta E} - 1}E \tag{5.65}$$

The work input is

$$W = E_{\text{fin}} - E_{\text{ini}} = \frac{1}{e^{\beta E} - 1}E. \tag{5.66}$$

Since the final energy of the qubit is 0, the energy decrease in the qubit is the initial energy of the qubit, $\delta_b E$. The CoP is thus

$$K = \frac{-\Delta E_t}{W} = \frac{1 - e^{-\beta E}}{1 + e^{-\beta E}} = \tanh \frac{\beta E}{2}. \tag{5.67}$$

We see that the optimal CoP is bounded between 0 and 1, depending on the bath temperature. Compared with the exponentially decreasing CoP (see Fig. 5.3) using the original proposal, this result suggests that there is much room for improvement in the energy efficiency when designing practical cooling protocols using a bosonic bath.

5.7 Summary

In this chapter, we investigated the thermodynamic properties of algorithmic cooling within a comprehensive framework of coherent cooling, encompassing a broad range of distinct protocols. By employing the transition matrix formalism, we identified a consistent behavior among these protocols concerning their cooling limits and target state evolution. Meanwhile, their thermodynamic performance displayed markedly diverse characteristics.

To compare the protocols, we proposed two key metrics: the coefficient of performance and the Landauer ratio, each highlighting a different aspect of a protocol. Furthermore, we demonstrated that the PPA becomes asymptotically efficient in terms of the Landauer ratio. Finally, we proposed improved versions of the PPA and xHBAC with enhanced energetic efficiency, inspired by our thermodynamic analysis. Our findings reinforce the connection between algorithmic cooling protocols and classical heat engines, thereby establishing a solid theoretical foundation for future research in the field of quantum thermodynamics.

Chapter 6

Views on quantum error mitigation

We saw in the previous chapter that AC could serve as a useful tool for physically improving the quality of initial state preparation in quantum computers. In this chapter, we discuss another near-term approach for reducing errors in quantum computation called quantum error mitigation (QEM), which refers to a family of protocols that reduce effective error rates through data post-processing. The emergence of QEM is partially driven by the difficulty to perform full scale quantum error correction in NISQ devices, as well as the fact that many near-term quantum-classical hybrid algorithms only use limited types of local measurements. Here, we examine the task of QEM from several perspectives. After reviewing some common QEM protocols in Section 6.1, in Section 6.2 we illustrate the fundamental distinctions between QEC and QEM from the perspective of classical and quantum communication. In Section 6.3 we study implications of noise invertibilities in QEM, and illustrate how non-invertible noise, which is largely omitted in the literature, may arise in an experiment. We propose a construction called Drazin-inverse for non-invertible noises, and prove that compared to a conventional choice of pseudoinverse (the Moore-Penrose pseudoinverse), the Drazin-inverse has the advantage of being trace preserving.In Section 6.4 we study the consequences due to imperfect knowledge about actual noise, and give a sufficient condition for when an optimal QEM can improve the expectation value of any observable.

6.1 Review of QEM protocols

The goal of QEM protocols is usually to recover the ideal expectation value of some observable A from the output state ρ_{out}, namely $\langle A\rangle = \text{Tr}[A\rho_{\text{out}}]$. These can be further classified

into two categories [111, 112]. The first one, which we call error-based QEM, requires prior knowledge on the form of noise occurring in a quantum processor. The second one does not require the particular form of noise being known, and will be called error-agnostic QEM. This section aims to review some of the commonly-used protocols, with a focus on error-based QEM. Our treatment will be based on the review article by Zhang et al. [112] with some modifications.

6.1.1 Extrapolation Methods

Extrapolation-based QEM is among the earliest proposed protocols that fall under the QEM category [44, 113]. The intuition is that while it is in general difficult to reduce the physical error rate on a processor, *increasing* the error rate in a controlled manner might be possible on some systems. Once sufficiently many observations of the expectation value under different noise strengths have been obtained, it could be possible to infer the noiseless value from the noisy ones. Consider a stochastic noisy physical process of the form

$$\mathcal{E} = (1-\epsilon)\mathcal{I} + \epsilon\mathcal{N} \tag{6.1}$$

where $\mathcal{I}$ is the identity map, $\mathcal{N}$ is the noise map, and ϵ is a small error rate. For a target quantum circuit with depth n, the ideal and noisy output states can be written as

$$\begin{aligned} \rho_{\text{out}}^{\text{ideal}} &= \mathcal{U}_n \circ \cdots \circ \mathcal{U}_1(\rho_{\text{in}}), \\ \rho_{\text{out}}^{\text{exp}} &= \mathcal{N} \circ \mathcal{U}_n \circ \cdots \circ \mathcal{N} \circ \mathcal{U}_1(\rho_{\text{in}}), \end{aligned} \tag{6.2}$$

where we have made a simplifying assumption that the noise is Markovian and gate-independent. The expectation value can be expanded in a series as

$$\langle A\rangle(\epsilon) = \langle A\rangle(0) + \sum_{m=0}^{k} A_m \epsilon^m + \mathcal{O}(\epsilon^{m+1}), \tag{6.3}$$

where A_m are expansion coefficients, and $\langle A\rangle(\epsilon)$ denotes the expectation value corresponding to an error rate of ϵ. One can then perform a series of experiments with different noise strength $\{\lambda_i\epsilon\}$, where $i = 0, ..., k$ and $\lambda_i > 0$, and collect the expectation values $\langle A\rangle(\lambda_i\epsilon)$. Applying the Richardson extrapolation method, one obtains the parameters r_i to infer the noiseless value up to precision $\mathcal{O}(\epsilon^{k+1})$:

$$\langle A\rangle_{\text{EM}} = \sum_{i=0}^{k} r_i \langle A\rangle(\lambda_i\epsilon) = \langle A\rangle(0) + \mathcal{O}(\epsilon^{k+1}). \tag{6.4}$$

Besides the original approach using Richardson extrapolation, other extrapolation methods exist. For example, the exponential or linear extrapolation method used by Endo et al. [45] uses two data points only. We note that extrapolation methods are accurate up to order $\mathcal{O}(\epsilon^{k+1})$, compared with quasiprobability method below which can be fully accurate in principle.

6.1.2 Quasiprobability Methods

QEM using quasiprobability sampling was also proposed by Temme et al. in the same paper where Richardson extrapolation QEM was introduced [44], and was later generalized by Endo et al. in [45]. Consider a target unitary gate $\mathcal{U}$ and its noisy version $\mathcal{U}_{\text{noisy}} = \mathcal{N}\mathcal{U}$, i.e., the ideal gate followed by a noise process. One can always find a CPTP map $\mathcal{N}$ that satisfies the above equation, since the inverse of $\mathcal{U}$ is always CPTP, and the composition of two CPTP maps is CPTP. The assumption is that there exists a complete set of (noisy) operations available to the experimentalist, denoted by

$$\{\mathcal{G}_1, \ldots, \mathcal{G}_K\} \tag{6.5}$$

which has been pre-characterized. This set is complete in the sense that they form a basis for the inverse noise channel $\mathcal{U}_{\text{noisy}}^{-1}$,

$$\mathcal{N}^{-1} = \sum_i a_i \mathcal{G}_i = \tau \sum_i \frac{|a_i|}{\tau} \text{sgn}(a_i) \mathcal{G}_i \tag{6.6}$$

where $\tau = \sum_i |a_i|$. The coefficients $p_i \coloneqq \frac{a_i}{\tau}$ now form a probability distribution, so one can append additional gates $\mathcal{G}_i$ with probability p_i after the noisy circuit to reverse the effects of $\mathcal{N}$. For an observable A, we can express its ideal value by

$$\langle A \rangle_{\text{EM}} = \text{Tr}\big[\mathcal{N}^{-1}\mathcal{N}\mathcal{U}(\rho)A\big] = \tau \sum_i p_i \text{sgn}(a_i) \langle A \rangle_i \tag{6.7}$$

where $\langle A \rangle_i = \text{Tr}[\mathcal{G}_i \mathcal{N}\mathcal{U}(\rho)A]$. We see that by adding in $\mathcal{G}_i$ at the end of the circuit with probability p_i, and keeping track of each $\text{sgn}(a_i)$, one obtains the ideal expectation value. Note that $\sum_i a_i = 1$ due to the trace-preserving constraint. Since some of the a_i's are negative, $\tau \geq 1$, so the $\langle A \rangle_{\text{EM}}$ has a variance that is approximately τ^2 times larger than that measuring with the ideal circuit.

It is interesting to point out that the original proposal for quasiprobability QEM aimed not to append $\mathcal{N}^{-1}$ at the end, but to directly simulate the *ideal* unitary $\mathcal{U}$. It may

appear that this scenario falls under the error-agnostic category. However, since it is still required that all noisy gates be characterized precisely, knowledge of noise in the real system is necessary. Therefore the method is still implicitly error-based. The effects of noise characterization comes in nontrivially through the gate decomposition procedure. In practice, it is sometimes more favorable to use the $\mathcal{N}^{-1}$ approach introduced here. This may happen when, for example, a decomposition for the near-identity map $\mathcal{N}$ has a much smaller cost τ than that for the target unitary $\mathcal{U}$.

6.1.3 Readout Error Mitigation

Readout error mitigation [47, 114] aims to reduce the effect of noisy measurement operations during the readout step of quantum computation. A general K-outcome measurement process can be described by a set of positive operator-valued measure (POVM) elements,

$$\{M_1, \ldots, M_K\}, \tag{6.8}$$

such that the probability of obtaining outcome j given an input state ρ is $p_j = \mathrm{Tr}[M_j \rho]$. We may arrange the outcome probabilities with an ideal measurement apparatus as a vector, $P_{\text{ideal}} = (p_1, \ldots, p_K)^T$, and similarly for the actual outcome probability vector P_{noisy}. One can see that there exists a transformation T between the two such that

$$P_{\text{noisy}} = T \cdot P_{\text{ideal}}. \tag{6.9}$$

The matrix T can be learned through experimentally measuring different input states, and its number of parameters may be less than the size of T if one makes further assumptions about the form of noise [115]. Once T is learned, the ideal probability vector can be obtained by $P_{\text{ideal}} = T^{-1} \cdot P_{\text{noisy}}$. In order to ensure invertibility of T, it is common to parametrize T in such a way that it is invertible by construction, when the parameters fall within a certain range.

6.1.4 Error-Agnostic QEM

We have so far focused on error-based QEM protocols. Below we briefly go over some other protocols that do not require explicit knowledge about the form of error, and are therefore error-agnostic. These approaches generally utilize certain structures of the problem that is known to the user prior to conducting any experiment. Prominent examples include virtual distillation (VD), symmetry verification (SV), and N-representability. In VD [48, 116, 117],

one obtains effective error-mitigated expectation values of an observable A as $\langle A \rangle_{\text{EM}} = \text{Tr}\left[A\rho_{\text{VD}}^{(m)}\right]$ where $\rho_{\text{VD}}^{(m)} = \rho^m / \text{Tr}[\rho^m]$, through entangling operations between m copies of the output state ρ. The purity of this effective state increases with m, thereby eliminating the effects of stochastic errors.

In SV [118, 46], one examines certain symmetry properties of the output state from a quantum program, and performs post-selection to remove outputs that disobey these symmetries. For example, in variational quantum eigensolver (VQE) algorithms, one prepares an ansatz input state which is then updated variationally, in order to approach the true ground state of a target Hamiltonian. There are often certain symmetries that are obeyed by the true ground state, such as the total number of electrons, as well as the number of spin-up spin-down electrons. One can thus measure these symmetries in the output state, and discard results that violate any symmetry tested.

The idea of N-representability QEM [119] is similar to that of SV. Its name originated from the N-representability conditions in quantum chemistry, which stand for a set of necessary conditions that must be satisfied by the reduced (marginal) density matrices. Due to the presence of noise, states measured in experiments might violate some of these conditions, so one may reduce the effects of these noise by projecting the experimental state onto the closest state satisfying the N-representability conditions.

6.2 Error Correction vs. Mitigation: a Communication Viewpoint

One can see from the previous review that many protocols with different nature are currently summarized under the name "quantum error mitigation". In this section we will provide a high-level comparison between QEM and QEC, with a focus on error-based QEM. As previously shown, these are procedures that utilize knowledge about the noise occurring during computation and *actively* try to eliminate their effects, making them more similar to *active* QEC. On the other hand, error-agnostic QEM are, in our opinion, closer to passive error correction. Nonetheless, insightful comparison between error-agnostic QEM and QEC can be an interesting future research direction and may facilitate integrations between QEC and QEM in general.

The development of QEC has historically stemmed from previous studies of classical error correction (CEC). Therefore, before discussing QEC or QEM, it is instructive to first look at their classical counterparts. We will start our discussion by considering a

communication task, which is essentially a trivial computation with the target unitary being the identity, but with two spatially separated parties being the input and output ends. This illustrates some simple yet important distinction between EC and EM. We then move on to computation and discuss the usefulness of QEM in quantum computation, using our intuition from the communication case.

6.2.1 Classical Communication

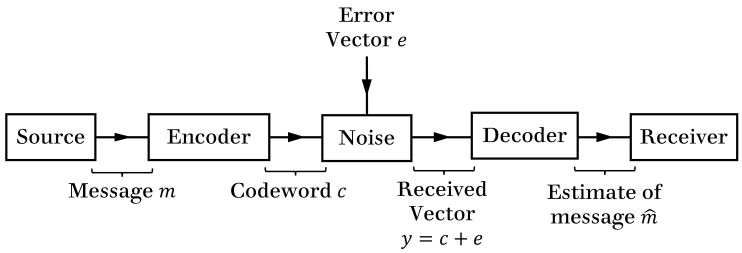

Figure 6.1: An illustrative diagram for noisy classical communication.

In a classical communication task, a sender Alice would like to transmit a k-bit string to Bob. An example of the string would look like

$$s = 1101010100010011010101001 1...$$

Alice and Bob share a classical communication channel C which is subject to noise. She does so by sending each letter of the text through C, involving a total of k uses of C. For simplicity, assume that the noise is described by a binary symmetric noise channel with strength p, denoted as BSC_p. This channel preserves the sent bit with probability $1-p$, and flips it (symmetrically from 0 to 1 and from 1 to 0) with probability p. The matrix representation of the noise is given by

$$\mathrm{BSC}_p \to \begin{pmatrix} 1-p & p \\ p & 1-p \end{pmatrix} \coloneqq N_{\mathrm{BSC},p} \tag{6.10}$$

where both the input and output basis are ordered as $\{0,1\}$. If Alice sends a bit 0, it will have a distribution

$$N_{\mathrm{BSC},p} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1-p \\ p \end{pmatrix}, \tag{6.11}$$

at the output end.

An illustrative diagram for classical error correcting codes against this noise is given in Fig. 6.1. The simplest example is the 3-bit repetition code, defined by the encoding

$$0 \to 000, \ 1 \to 111, \tag{6.12}$$

i.e., each bit is repeatedly encoded 3 times. The number of uses of the channel has now increased 3 times to $3k$. The decoding is done by performing a majority vote on the received bits, so that at the receiving end, Bob again obtains a bit string of length k. Assuming that BSC_p acts independently on each bit, the probability of error is reduced from p to $3p^2(1-p) = \mathcal{O}(p^2)$ by this code.

Next we ask the question: Can Bob improve the quality of communication, given some knowledge about the noise? From the form of noise in Eq. (6.10), the best possible knowledge Bob could have is the precise value of p. Suppose in addition that $p < 1/2$. If Bob receives a bit 1, then he only knows that Alice more likely sent a 1 than a 0, so the best deterministic procedure is simply to keep the bit. Applying this argument to all received bits, we see that the best Bob can do is to simply keep all received bits intact. Similarly, if $p > 1/2$, then Bob's best action is to flip all the received bits. Clearly, this does not increase Bob's information on Alice's message, as measured by the classical mutual information. In fact, what we have shown is a special case of the classical data-processing inequality, which states that no post-processing can increase the mutual information between Alice and Bob.

What Bob *can* do is to recover the *distribution* of Alice's input for sufficiently large k. Specifically, he can apply the inverse map of $N_{\mathrm{BSC},p}$ on his received distribution, resulting in

$$N_{\mathrm{BSC},p}^{-1} N_{\mathrm{BSC},p} v_A = v_A, \tag{6.13}$$

where $v_A = (p_{0,a}, p_{1,a})^T$ is Alice's input distribution. However, Bob cannot further use this restored distribution to recover Alice's *message*. The best he can do is to use the restored distribution to randomly generate a new k-bit string, during which Alice's message is completely destroyed.

6.2.2 Quantum Communication

In quantum communication, Alice and Bob communicates by sharing a quantum channel Q capable of transmitting quantum particles. Alternatively, one may also model the situation as having a central source sending out (potentially entangled) particles to Alice and Bob. If this source is located in Alice's lab, then the channel from source to Alice is ideal and

(a) (b)

Figure 6.2: (a) Quantum communication model where Alice sends a state ρ_{in} to Bob through a noisy channel $\mathcal{N}$. (b) An attempt to reduce the effect of $\mathcal{N}$ using QEC, through an encoding operation $\mathcal{E}$ and a recovery operation $\mathcal{R}$.

the one to Bob is described by a noisy channel $\mathcal{N}$, as shown in Fig. 6.2a. Here, they are interested in either encoding classical information in the particles and use Q to transmit a classical message, or directly sharing quantum particles which may have been prepared in some special states.

Quantum Communication of Classical Information

In the first scenario, the classical information may be extracted at Bob's end by measuring the received state. For each use of the channel, we will name the state that Alice sent ρ_{in} and the one coming out of Bob's end $\rho_{\text{out}} = \mathcal{N}(\rho_{\text{in}})$. Without loss of generality, assume that the information is encoded in the expectation value of some observable O. Due to the nature of quantum measurements, every possible outcome occurs randomly with probabilities determined by the Born rule. Therefore, Alice must send multiple copies of the state in order for Bob to extract the expectation value, $\text{Tr}[O\rho_{\text{out}}]$. This is to be contrasted with classical communication, where sending 1 bit of information only involves using C once, in the limit of an ideal channel. One thus sees that, sending classical information using a quantum channel inevitably results in Bob measuring a distribution of possible outcomes.

If Bob knows the exact form of $\mathcal{N}$, then he may reconstruct the matrix ρ_{out} using sufficiently many copies of received states, assuming that the size of ρ_{out} is not too large. He may then apply $\mathcal{N}^{-1}$ to perfectly restore ρ_{in}, and consequently use this classical copy to infer the expectation value of any observable. This is in direct analogy to restoring the distribution of input characters using CEM described in the previous section.

Alternatively, Alice and Bob may also use QEC to fight against this noise. A quantum error correcting code (QECC) is defined by an encoding scheme, which is a completely positive and trace preserving (CPTP) map $\mathcal{E} : \mathbb{C}_{2^k} \to \mathbb{C}_{2^n}$ (where $n > k$), and then decode at Bob's end using another CPTP map $\mathcal{R} : \mathbb{C}_{2^n} \to \mathbb{C}_{2^k}$. Here $\mathbb{C}_{2^n}$ denotes the complex Euclidean space with dimension 2^n. This involves a total of nk uses of the quantum

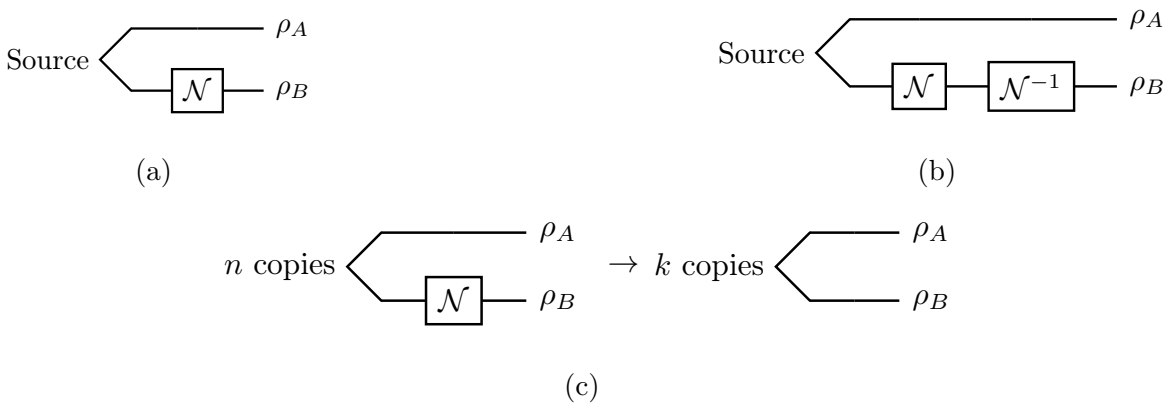

Figure 6.3: (a) Figure where Alice and Bob receives an EPR pair Φ^+ from a source, and a noise $\mathcal{N}$ occur for Bob's channel. (b) A QEM approach. (c) A QEC approach.

channel. Assuming that $\mathcal{N}$ is correctable by QECC [120], the output state ρ_{out} will be equal to the input ρ_{in}, so directly measuring this error-corrected state will give Bob the correct expectation value for any observable.

One may now recognize the similarity between QEM and CEM for reconstructing the input *distribution*. Indeed, a density matrix *is* probabilistic description of outcomes of any possible measurement on a quantum system. In this scenario, it is sufficient for Bob to reconstruct ρ_{in} as a mathematical object in order to eliminate the effect of $\mathcal{N}$, since this determines the outcome of *any measurement* Bob can possibly make on the system.

Quantum Communication of Quantum Information

In the second scenario, the goal is to send particles encoding quantum information from Alice to Bob. The physical states of the particles represent the quantum information encoded, which must be kept "alive" for such a task. This scenario is relevant when, for example, ρ_{in} is the output from a computation done by Alice, which needs to be transferred to Bob for further processing. Clearly, the QEC approach shown in Fig. 6.2b is capable of eliminating the effect of $\mathcal{N}$ if it falls within the set of correctable errors, because what Bob gets at the end is an actual quantum object. In contrast, one cannot use the "error mitigation" approach described in the previous section, where Bob applies $\mathcal{N}^{-1}$ onto a *classical image* of ρ_{out}, since this only corrects Bob's description about ρ_{out} rather than the physical state itself.

Another case is when Alice and Bob tries to establish shared entanglement to achieve certain tasks, such as winning a nonlocal game. This is relevant because entanglement is among the most commonly accepted benchmarks for quantum information, and was speculated to have a similar role as classical information [121]. As shown in Fig. 6.3, Alice prepares k copies of maximally entangled Bell pairs

$$\Phi^+ = (|00\rangle_{AB} + |11\rangle_{AB})/\sqrt{2} \tag{6.14}$$

and sends half of each pair to Bob. Again, the noisy channel $\mathcal{N}$ to Bob potentially reduces the entanglement shared between Alice and Bob. Historically, this task is achieved by a family of procedures called entanglement purification protocols (EPPs) [121], illustrated in Fig. 6.3c. In EPP, Alice and Bob needs to start from $n > k$ copies of the noisy Bell state, and obtain k pairs at the end which are closer to the pure state Φ^+. Profoundly, a sub-class of EPP protocols called one-way EPP (or 1-EPP), where only communication from Alice to Bob is allowed, have a one-to-one correspondence with QECCs [121].

On the contrary, one sees that the QEM approach described in the previous section cannot be used to purify entanglement. In particular, all entanglement between Alice and Bob would be destroyed due to the measurement. In fact, analogous to the classical case where Bob recovers Alice's input distribution and generate a random k-bit string, here Bob knows in advance that he will ideally get a maximally mixed state $I/2$; so the above protocol is simply equivalent to Bob generating $(I/2)^k$ locally, and discarding all qubits received from the source!

The above example illustrates a fundamental distinction between QEM, which is only capable of restoring the classical image of a quantum system, and QEC, which is capable of restoring the quantum object *itself*, along with all possible non-classical resources possessed by that object. It is instructive to recall again the case where one needs to preserve classical information (see Section 6.2.1), where we have also argued that CEC is helpful for such a task, while CEM is not. Furthermore, recall in Section 6.2.2 we argued that recovering density matrices in QEM is analogous to recovering classical distributions in CEM. These complete our comparisons between EM and EC, which are summarized in Table 6.1.

6.3 Noise Invertibility and the Drazin-Inverse

Having illustrated the fundamental distinction between error-based QEM and QEC under the communication setting, it is then natural to ask under what circumstances are these two methods useful. For QEC, the set of errors (and linear combinations within) are

	EC	EM
Classical	Classical information being transmitted	Classical distribution of possible outputs
Quantum	Physical quantum objects being transmitted	Density matrices describing the physical quantum objects

Table 6.1: Summary of what EC and EM could recover at the output end, under classical and quantum communication settings.

correctable if they obey the Knill-Laflamme conditions [122, 120] for the particular code being used. While there exists no parallel, quantitative results for QEM, the form of some error-based QEM protocols suggests that the inverse noise channel plays an important role, since it is what these protocols aim to apply to the output (directly in the case of readout QEM, and indirectly in the case of quasiprobability QEM, for instance). In this section, we discuss the implications of noise invertibility on error-based QEM, and propose an alternative quasi-inverse construction in the case of non-invertible noise.

A quantum noise process is, on the physical level, described by a completely-positive (CP) and trace-preserving (TP), or CPTP map, $\mathcal{N}$. Below we first define a matrix representation for quantum states and maps. In this chapter, we denote the space of linear operators mapping Hilbert space H_A to H_B as $L(H_A, H_B)$, or $L(H_A)$ in short if $H_A = H_B$. Let $T(H_A, H_B)$ be the space of linear maps from $L(H_A)$ to $L(H_B)$. Let e_i be the standard basis of H_i with a 1 at position i and 0 elsewhere. Let $E_{a,b}$ be the standard basis of $L(H_A, H_B)$ with a 1 at position (a, b) and 0 elsewhere.

Definition 6.3.1. (Vectorization of linear operators.) The vec mapping $\mathbf{v}(\cdot) : L(H_A, H_B) \to H_B \otimes H_A$ is the unique mapping that satisfies $\mathbf{v}(E_{a,b}) = e_b \otimes e_a$.

Next we define two representations for quantum maps.

Definition 6.3.2. (Choi representation.) The Choi representation of a map $\mathcal{M} \in T(H_A, H_B)$ is defined by $C(\mathcal{M}) = \sum_{a,b} E_{a,b} \otimes \mathcal{M}(E_{a,b})$.

Definition 6.3.3. (Natural representation.) The natural (or equivalently, superoperator) representation of a map $\mathcal{M} \in T(H_A, H_B)$ is defined by the unique linear operator $\mathbf{v}(\mathcal{M}) \in L(H_A \otimes H_A, H_B \otimes H_B)$ that satisfies $\mathbf{v}(\mathcal{M})\mathbf{v}(A) = \mathbf{v}(\mathcal{M}(A))$ for all $A \in L(H_A)$.

In the natural representation, the channel $\mathcal{N}$ acting on a quantum state ρ can be written as the superoperator $\mathbf{v}(\mathcal{N})$ multiplying the vector representation $\mathbf{v}(\rho)$ of the quantum

state ρ [56]. The vector representation $\mathbf{v}(\rho)$ of ρ inherits its ordering from the superoperator, hence we abuse the notation $\mathbf{v}(\cdot)$ for vector representations of quantum states and observables (which are often written as double kets $|\rho\rangle\rangle$ in other literature).

The following theorem directly comes from representation theory of linear maps.

Theorem 6.3.1. *The quantum channel $\mathcal{N}$ is invertible iff $\mathbf{v}(\mathcal{N})$ is an invertible matrix.*

Below we give an example where the inverse $\mathcal{N}^{-1}$ of a CPTP map $\mathcal{N}$ is non-CP.

Example 6.3.1. Let the Choi representation of a quantum channel $\mathcal{N}$ be

$$C(\mathcal{N}) = \left(\begin{array}{cc|cc} \frac{3}{4} & 0 & -\frac{i}{8} & \frac{1}{2}+\frac{i}{8} \\ 0 & \frac{1}{4} & -\frac{i}{8} & \frac{i}{8} \\ \hline \frac{i}{8} & \frac{i}{8} & \frac{1}{4} & 0 \\ \frac{1}{2}-\frac{i}{8} & -\frac{i}{8} & 0 & \frac{3}{4} \end{array}\right).$$

The superoperator is

$$\mathbf{v}(\mathcal{N}) = \begin{pmatrix} \frac{3}{4} & \frac{i}{8} & -\frac{i}{8} & \frac{1}{4} \\ 0 & \frac{1}{2}-\frac{i}{8} & -\frac{i}{8} & 0 \\ 0 & \frac{i}{8} & \frac{1}{2}+\frac{i}{8} & 0 \\ \frac{1}{4} & -\frac{i}{8} & \frac{i}{8} & \frac{3}{4} \end{pmatrix}.$$

Therefore, the inverse of $\mathbf{v}(\mathcal{N})$ is

$$\mathbf{v}(\mathcal{N}^{-1}) = \begin{pmatrix} \frac{3}{2} & \frac{1}{4}-\frac{i}{2} & \frac{1}{4}+\frac{i}{2} & -\frac{1}{2} \\ 0 & 2+\frac{i}{2} & \frac{i}{2} & 0 \\ 0 & -\frac{i}{2} & 2-\frac{i}{2} & 0 \\ -\frac{1}{2} & -\frac{1}{4}+\frac{i}{2} & -\frac{1}{4}-\frac{i}{2} & \frac{3}{2} \end{pmatrix}$$

with a Choi representation

$$C(\mathcal{N}^{-1}) = \left(\begin{array}{cc|cc} \frac{3}{2} & 0 & \frac{1}{4}+\frac{i}{2} & 2-\frac{i}{2} \\ 0 & -\frac{1}{2} & \frac{i}{2} & -\frac{1}{4}-\frac{i}{2} \\ \hline \frac{1}{4}-\frac{i}{2} & -\frac{i}{2} & -\frac{1}{2} & 0 \\ 2+\frac{i}{2} & -\frac{1}{4}+\frac{i}{2} & 0 & \frac{3}{2} \end{array}\right).$$

The Choi representation $C(\mathcal{N}^{-1})$ has negative eigenvalues. Therefore, $\mathcal{N}^{-1}$ is a Hermitian preserving (HP) and trace preserving (TP) map, but not CP.

There are three distinct possibilities regarding noise invertibility. The first is that $\mathcal{N}$ is invertible, and $\mathcal{N}^{-1}$ is CPTP. In this case, the inverse $\mathcal{N}^{-1}$ is unique, and is Hermitian preserving (HP) and trace preserving (TP) [123]. Note that the set of CPTP maps is a subset of HPTP maps. For $\mathcal{N}^{-1}$ to be CPTP, the channel $\mathcal{N}$ has to be a unitary channel, or can be seen as an unitary channel acting on the input state along with an ancilla prepared to a fixed state [124]. If the dimensions of input and output space are the same, the channel has a CPTP inverse *iff* the channel is an unitary channel [125, 124]. This is a relatively easy scenario, because the quantum information can be coherently restored by physically applying $\mathcal{N}^{-1}$ to the output state.

The second possibility is that $\mathcal{N}$ is invertible, but $\mathcal{N}^{-1}$ is not CPTP. A condition for when this will happen is later given in Proposition 6.3.2. Many experimentally relevant noise models, such as the phase damping channel and the depolarizing channel, fall under this category. Since $\mathcal{N}^{-1}$ is not a physically realizable operation, it cannot be experimentally implemented on the target system, so our above method to restore quantum information without redundancy fails. Using QEM procedures, one can still recover the classical information in principle, by first extracting the classical output density matrix through measurements, and numerically apply the inverse map $\mathcal{N}^{-1}$. But the process of measurement will inevitably disturb the system being measured, and destroy any entanglement it possibly has with other systems.

The third possibility is that $\mathcal{N}$ is non-invertible. First, we motivate where a non-invertible noise channel may originate from. Any CPTP map can be dilated to a unitary channel in a larger Hilbert space, and unitary channels are all invertible. What happens while tracing out the environment? For a physical action to be written as a CPTP map, one needs to assume that the system and environment are separable in the beginning, which means that the input state has to have the form $\rho \otimes \sigma$. If the unitary U is also separable, tracing out the environment will also result in a unitary evolution on the system. The correlation between system and environment causes the evolution of the system to become a general CPTP map. The non-invertibility of a channel signals strong non-locality, i.e., there is information in the whole system (system and environment) that is entirely invisible in the local system.

An example is the CNOT gate in Fig. 6.4. Let the initial state ρ_{AB} be $\rho_A \otimes |0\rangle\langle 0|$. Considering the controlling qubit B as the environment, the reduced channel $\mathcal{N}_A$ on the system qubit A is non-invertible. The density matrix of qubit A after the first CNOT gate is

$$\rho_{A'} = \mathcal{N}_A(\rho_A) = \mathrm{tr}_B[U_{CNOT}(\rho_A \otimes |0\rangle\langle 0|)U_{CNOT}^\dagger].$$

Some information (more precisely, certain off-diagonal entries) in the 2-qubit density matrix

$\rho_{A'B'}$ does not reflect in the density matrix $\rho_{A'}$ of the system qubit (in fact, this nonlocal information cannot be seen locally in $\rho_{B'}$ either). However, after the second CNOT gate, the whole system AB backtracks to the original state. That is to say

$$\begin{aligned}\rho_A &= \rho_{A''} \\ &= \mathrm{tr}_B[U_{CNOT}^\dagger U_{CNOT}(\rho_A \otimes |0\rangle\langle 0|)U_{CNOT}U_{CNOT}^\dagger] \\ &\neq \mathcal{N}_A^{-1}\mathcal{N}_A(\rho_A).\end{aligned}$$

This is because the second CNOT gate utilizes nonlocal information that is not available locally for qubit A. Although this information loss through the null space of $\mathcal{N}_A$ cannot be recovered without collecting information from the environment, the local rotation and contraction caused by the noise channel can be restored.

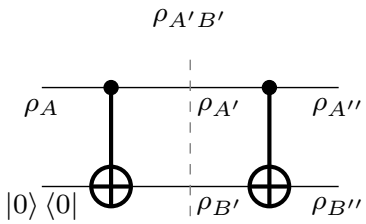

Figure 6.4: The equivalent channel $\mathcal{N}_\mathcal{A}$ on system A is non-invertible. However, the second CNOT gate brings the whole device back to the original state, i.e. $\rho_{A''B''} = \rho_{A''} \otimes \rho_{B''} = \rho_A \otimes |0\rangle\langle 0|$. The locally unseen information in $\rho_{A'B'}$ flow back to system A after the second CNOT.

Second, how can we deal with these non-invertible noise channels? It is known that the superoperator $\mathbf{v}(\mathcal{N}^{-1})$ of the inverse channel $\mathcal{N}^{-1}$ equals to the inverse $\mathbf{v}(\mathcal{N})^{-1}$ of the superoperator $\mathbf{v}(\mathcal{N})$. However, if the channel $\mathcal{N}$ is not invertible, the generalized inverse of $\mathbf{v}(\mathcal{N})$ is not unique. A commonly used generalized inverse is the Moore-Penrose inverse [126, 127], but in Example 6.3.2 we show that the Moore-Penrose inverse of a CPTP map is not necessarily TP. In the following, we provide a construction of inverse-like channel $\mathcal{N}^+$.

Let the dimension of input and output space be d. Take the Jordan decomposition of the superoperator of $\mathcal{N}$,

$$\mathbf{v}(\mathcal{N}) = Q \cdot J \cdot Q^{-1} \tag{6.15}$$

where $J = \oplus_i J_{\lambda_i}$ is the Jordan normal form, J_{λ_i} is a Jordan block corresponding to the eigenvalue λ_i, and Q is a invertible matrix contains the generalized eigenvectors of $\mathbf{v}(\mathcal{N})$.

If $\mathbf{v}(\mathcal{N})$ is diagonalizable, the Jordan normal form $J = \mathrm{diag}[\lambda_1, \cdots, \lambda_{d^2}]$ is the diagonal matrix contains eigenvalues λ_i of $\mathbf{v}(\mathcal{N})$.

We take the inverse-like channel $\mathcal{N}^+$ to be

$$\mathbf{v}(\mathcal{N}^+) = Q \cdot J' \cdot Q^{-1}. \tag{6.16}$$

If $\mathbf{v}(\mathcal{N})$ is diagonalizable, J' is the diagonal matrix that leaves the 0's in J untouched and take the reciprocal of the rest elements in J. If $\mathbf{v}(\mathcal{N})$ is defective (meaning that it contains blocks of zero matrices), we can construct each Jordan block in the following way: a k by k Jordan block J_{λ_i} of λ_i ($\lambda_i \neq 0$) in J is

$$J_{\lambda_i} = \begin{pmatrix} \lambda_i & 1 & & \\ & \lambda_i & \ddots & \\ & & \ddots & 1 \\ & & & \lambda_i \end{pmatrix},$$

let the corresponding block J'_{λ_i} in J' be the inverse of J_{λ_i}

$$J'_{\lambda_i} \coloneqq J_{\lambda_i}^{-1} = \begin{pmatrix} \frac{1}{\lambda_i} & -\frac{1}{\lambda_i^2} & \cdots & (-1)^{k+1}\frac{1}{\lambda_i^k} \\ & \frac{1}{\lambda_i} & \cdots & (-1)^{k}\frac{1}{\lambda_i^{k-1}} \\ & & \ddots & \vdots \\ & & \frac{1}{\lambda_i} & -\frac{1}{\lambda_i^2} \\ & & & \frac{1}{\lambda_i} \end{pmatrix}.$$

For a k by k Jordan block of diagonal zero ($\lambda_i = 0$), which is the nilpotent matrix N, we can set the corresponding block in J' as a zero matrix 0_k . Since N is not invertible, letting the block be 0_k will have the same result as setting it as N^{k-1}. There is a certain freedom in the choice of this block.

Note that, for invertible channels, $\mathcal{N}^+$ described above provides the inverse $\mathcal{N}^{-1}$ of the channel ($\mathcal{N}^+ = \mathcal{N}^{-1}$). For non-invertible channels, this construction Eq. (6.16) does not satisfy the condition of generalized inverse ($\mathcal{N} \circ \mathcal{N}^+ \circ \mathcal{N} \neq \mathcal{N}$ when the dimension of the nilpotent Jordan block is greater than one). We will call $\mathcal{N}^+$ the *Drazin-inverse* since it is the same construction as the Drazin inverse in matrix analysis [128].

The resulting composed map $\mathbf{v}(\mathcal{N})\mathbf{v}(\mathcal{N}^+) = QJ''Q^{-1}$, where J'' is a diagonal matrix with only 0's and 1's on its main diagonal. When the noise channel $\mathcal{N}$ already only contains 1 and 0 in its spectrum, the Drazin-inverse is itself, and does not recovery more information. In fact, any generalized inverse would not improve the outcome in this case.

The following proposition tells us another condition for a quantum channel to have a non-CP (Drazin-) inverse.

Proposition 6.3.2. *If a non-zero eigenvalue λ of a quantum channel $\mathcal{N}$ has modulus less than 1 ($|\lambda| < 1$), then the inverse (or Drazin-inverse) channel $\mathcal{N}^+$ is not completely positive.*

Proof. $\mathcal{N}$ is a CPTP map, therefore its spectral radius is one [129], i.e. $|J_{ii}| \leq 1$ for any main diagonal element J_{ii} in J. Since $\mathcal{N}$ has eigenvalues less than 1, there exists $|J_{jj}| < 1$ for some $j \in \{1, \cdots, d^2\}$. As defined above, $|J'_{jj}| > 1$, i.e. the spectral radius of $\mathcal{N}^+$ is greater than one. Therefore, $\mathcal{N}^+$ is not complete positive. □

A non-unitary quantum channel cannot be fully recovered by another channel since CPTP maps cannot resolve contractions in the subspace corresponding to $|\lambda_i| < 1$. To deal with contractions, the CP property has to be broken. That is to say, restricting the generalized inverse $\mathcal{N}^g$ to be CPTP has to scarifies the quality of recovery. However, while HPTP maps cannot be directly implemented in a physical system, Steinspring dilation theory guarantees CPTP maps can be extended to unitary channels and hence executed in physical devices. This inspires several beautiful works on finding CPTP inverses that optimize the average fidelity of the recovery [130, 131, 132]. On the other hand, HPTP maps can be decomposed as a linear combination of CPTP maps, and so can be implemented in a physical device. The physical implementability of HPTP maps is defined and discussed in [123, 133]. Also note that the spectrum of a quantum channel can be defined independently from its representations. In this section, we mainly work with superoperators (natural representation), but the Proposition 6.3.2 still holds in other representations (e.g. the Pauli representation).

Unlike the Choi representation, the natural representation does not directly show a lot of critical properties of quantum channels, such as CP, TP, or HP. However, we found that the eigen-structure of the superoperator is essential for its property. Lemma 6.3.3 and Lemma 6.3.4 provide insights into why Moore-Penrose inverse is not TP in certain cases. Then, we prove that the Drazin-inverse for a TP map is always TP in Theorem 6.3.5.

Denote the trace operation in the vector representation $\mathbf{v}(A)$ of a d by d matrix A as $\mathrm{sTr}[\cdot]$, where $\mathrm{sTr}[\mathbf{v}(A)] \coloneqq \mathrm{Tr}(A)$.

Lemma 6.3.3. *If a linear map $\mathcal{N} : M_d \to M_d$ is trace preserving, the eigenvectors v and generalized eigenvectors v^g of eigenvalue $\lambda \neq 1$ of the superoperator $\mathbf{v}(\mathcal{N})$ is trace zero, i.e. $sTr[v] = sTr[v^g] = 0$.*

Proof. For an eigenvector v of $\mathbf{v}(\mathcal{N})$, we have $\mathbf{v}(\mathcal{N})v = \lambda v$. Since $\mathcal{N}$ is trace preserving, $\mathrm{sTr}[v] = \mathrm{sTr}[\lambda v]$. And the eigenvalue $\lambda \neq 1$, we have $\mathrm{sTr}[v] = 0$

For a k by k Jordan block of eigenvalue λ^g, where $k > 1$, denote the first generalized eigenvector as v^{g_1}, we have

$$[\mathbf{v}(\mathcal{N}) - \lambda^g I]v^{g_1} = v, \tag{6.17}$$

where v is the eigenvector corresponding to λ^g.

Taking the trace on both sizes, $\mathrm{sTr}[(\mathbf{v}(\mathcal{N}) - \lambda^g I)v^{g_1}] = \mathrm{sTr}[v]$, the left hand side is $\mathrm{sTr}[v^{g_1} - \lambda^g v^{g_1}] = (1-\lambda^g)\mathrm{sTr}[v^{g_1}]$, and the right hand side is zero from the argument above. Since $\lambda^g \neq 1$, $\mathrm{sTr}[v^{g_1}] = 0$. By deduction, all v^{g_i} are trace zero for $i \in \{1, \cdots, k-1\}$. □

Lemma 6.3.4. *For a trace persevering linear map $\mathcal{N} : M_d \to M_d$, if there is a k by k ($k > 1$) defective Jordan Block of eigenvalue $\lambda = 1$ in $\mathbf{v}(\mathcal{N})$, the eigenvector v and first $k-2$ generalized eigenvector v^{g_i} has to be trace zero, i.e. $sTr[v] = sTr[v^{g_i}] = 0$ for $i \in \{1, \cdots, k-2\}$.*

Proof. Assume that $\mathrm{sTr}[v] \neq 0$. The first generalized eigenvector v^g satisfy that $[\mathbf{v}(\mathcal{N}) - I]v^{g_1} = v$. Taking trace on both size, the left hand side equals to zero, and the right hand side does not equal to zero, leading to a contradiction. The same argument holds for the rest of the generalized eigenvectors except the last one. □

From Lemma 6.3.3 and Lemma 6.3.4, we know that all eigenvectors v_λ for $\lambda \neq 1$ of a TP map has to be traceless. When $\lambda = 1$, if its algebraic multiplicity equals to its geometry multiplicity, $\mathrm{sTr}[\mathbf{v}(\mathcal{N})v_\lambda] = \mathrm{sTr}[v_\lambda]$ (i.e. the trace of v_λ will not be changed under the action of $\mathbf{v}(\mathcal{N})$); if the algebraic multiplicity does not equal to the geometry multiplicity, the eigenvectors and generalized eigenvectors is traceless except for the last generalized eigenvector. This tells us that the eigen-structure of the superoperator $\mathbf{v}(\mathcal{N})$ is crucial for $\mathcal{N}$ to be TP. The way that we construct the Drazin-inverse $\mathcal{N}^+$ largely preserves the eigen-structure, while the Moore-Penrose inverse $\mathcal{N}^p$ focuses more on the singular value structure. It hints that $\mathcal{N}^+$ should be TP and $\mathcal{N}^p$ may not.

Theorem 6.3.5. *The Drazin-inverse $\mathcal{N}^+$ of a trace preserving map $\mathcal{N}$ is also trace preserving.*

To prove that $\mathcal{N}^+$ is trace preserving, we need to prove

$$\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)v_\lambda\right] = \mathrm{sTr}\left[v_\lambda\right],$$

for every eigenvectors and generlized eigenvectors v_λ of $\mathbf{v}(\mathcal{N})$ in Q. From the construction of $\mathcal{N}^+$, we almost get trace preserving for free. The full proof can be found in Appendix D.1. Moreover, it is easy to see from the proof that the composed map $\mathcal{N}^+ \circ \mathcal{N}$ is also trace preserving.

Example 6.3.2. Here we give an example where the Moore-Penrose inverse $\mathcal{N}^p$ of a CPTP map is not TP, while the Drazin-inverse $\mathcal{N}^+$ is TP. Consider a noise channel $\mathcal{N}$ whose Choi representation is given by

$$C(\mathcal{N}) = \frac{1}{20}\left(\begin{array}{cc|cc} 8 & 0 & 1 & 6 \\ 0 & 12 & 2 & -1 \\ \hline 1 & 2 & 8 & 0 \\ 6 & -1 & 0 & 12 \end{array}\right), \tag{6.18}$$

and its superoperator is

$$\mathbf{v}(\mathcal{N}) = \frac{1}{20}\begin{pmatrix} 8 & 1 & 1 & 8 \\ 0 & 6 & 2 & 0 \\ 0 & 2 & 6 & 0 \\ 12 & -1 & -1 & 12 \end{pmatrix}$$

The Jordan normal form is given by $J = \mathrm{diag}(0, 1, \frac{2}{5}, \frac{1}{5})$, and its inverse is $J' = \mathrm{diag}(0, 1, \frac{5}{2}, 5)$. The superoperator of Drazin-inverse $\mathcal{N}^+$ is

$$\mathbf{v}(\mathcal{N}^+) = \begin{pmatrix} \frac{2}{5} & \frac{5}{16} & \frac{5}{16} & \frac{2}{5} \\ 0 & \frac{15}{4} & -\frac{5}{4} & 0 \\ 0 & -\frac{5}{4} & \frac{15}{4} & 0 \\ \frac{3}{5} & -\frac{5}{16} & -\frac{5}{16} & \frac{3}{5} \end{pmatrix}$$

The Choi representation of $\mathcal{N}^+$ is

$$C(\mathcal{N}^+) = \left(\begin{array}{cc|cc} \frac{2}{5} & 0 & \frac{5}{16} & \frac{15}{4} \\ 0 & \frac{3}{5} & -\frac{5}{4} & -\frac{5}{16} \\ \hline \frac{5}{16} & -\frac{5}{4} & \frac{2}{5} & 0 \\ \frac{15}{4} & -\frac{5}{16} & 0 & \frac{3}{5} \end{array}\right)$$

The Choi representation has negative eigenvalues. Therefore, the channel $\mathcal{N}^+$ is trace preserving (partial trace of $C(\mathcal{N}^+)$ is identity), Hermitian preserving ($C(\mathcal{N}^+)$ is Hermitian), but not complete positive.

The Moore-Penrose inverse of $\mathbf{v}(\mathcal{N})$ is

$$\mathbf{v}(\mathcal{N}^p) = \begin{pmatrix} \frac{115}{294} & \frac{10}{441} & \frac{10}{441} & \frac{505}{882} \\ \frac{50}{147} & \frac{3245}{882} & -\frac{1165}{882} & -\frac{100}{441} \\ \frac{50}{147} & -\frac{1165}{882} & \frac{3245}{882} & -\frac{100}{441} \\ \frac{115}{294} & \frac{10}{441} & \frac{10}{441} & \frac{505}{882} \end{pmatrix},$$

and its Choi representation is

$$C(\mathcal{N}^p) = \left(\begin{array}{cc|cc} \frac{115}{294} & \frac{50}{147} & \frac{10}{441} & \frac{3245}{882} \\ \frac{50}{147} & \frac{115}{294} & -\frac{1165}{882} & \frac{10}{441} \\ \hline \frac{10}{441} & -\frac{1165}{882} & \frac{505}{882} & -\frac{100}{441} \\ \frac{3245}{882} & \frac{10}{441} & -\frac{100}{441} & \frac{505}{882} \end{array}\right),$$

which is Hermitian preserving but not trace preserving.

To our knowledge, the exploration of generalized non-CPTP inverses for non-invertible quantum channels has been limited in existing literature. The new understanding of the natural representation opens the possibility of studying the structures and properties of these maps from mathematical interests. This research direction has the potential to serve as a valuable guide for implementing these maps in error mitigation techniques and other aspects of quantum information science.

6.4 QEM in Quantum Computation

In the previous sections we discussed how the nature of noise determines whether it is theoretically possible to fully recover the quantum and/or classical information. We explored the possibility of non-invertible noise and constructed a Drazin-inverse under such a case. In this section, we will study the effects of recovery operations when performing QEM in quantum computation. The task of quantum computation may be viewed as a modified version of communication, where Alice and Bob are no longer spatially separated, but the channel from Alice (the input) to Bob (the output) becomes a nontrivial unitary $\mathcal{U}$.

Using the intuition built from the communication setting, there can again be two situations depending on the goal. In the first case, one is interested in obtaining the expectation value of some observable $\mathrm{Tr}[O\rho_{\mathrm{out}}]$ from the output state. We can call this "quantum computation of classical information", inspired by the usage of a similar term when discussing about quantum communication. Accurately obtaining $\mathrm{Tr}[O\rho_{\mathrm{out}}]$ for some observable O is

the goal of many quantum algorithms, such as the Harrow-Hassidim-Lloyd algorithm [134] for solving linear systems and the variational quantum eigensolver algorithm [135] for estimating ground/excited state energies of molecules. From our discussion before, it is now clear that both QEC and QEM can be useful for these types of problems. Indeed, to the authors' knowledge, all preexisting QEM protocols have been developed towards solving these problems.

In the second case, the goal is to output a particular quantum state, which needs to be kept in coherence and perhaps sent to another party later. This is particularly relevant in, for example, the problem of distributed quantum computing [136, 137] or active quantum memories [138, 139]. This is thus the case of "quantum computation of quantum information". Here, the idea of QEM cannot be used (at least locally at the output end), to reduce noise effects for these types of tasks, and only QEC is useful.

As mentioned previously, quantum computation has an additional layer of complexity that comes from the composition of gates. In communication, it is known prior to sending particles that the ideal "gate" is the identity, so applying the inverse noise map would directly yield the ideal input distribution. In computation, typically only a condensed description of the target unitary in the form of a quantum circuit is available prior to an experiment. The full form of the target unitary map (and consequently, that of the noise process) is typically not known explicitly. Instead, one usually have knowledge on components (e.g., the form of one- or two-qubit gates) in the circuit as well as the noise on these components. Thus, QEM protocols reviewed in Section 6.1 were developed to implement this noise inversion procedure more efficiently in practice, for "quantum computation of classical information" tasks. In the following we study this scenario in more detail, with a focus on the effects of imperfect knowledge on noise in Section 6.4.2.

6.4.1 QEM in Multi-layer Quantum Computation

Consider again a noisy quantum circuit with depth n, where each layer can be represented by a unitary map $\mathcal{U}_i$ with $i = 1, ..., n$. The ideal output would be

$$\rho_{\text{out}}^{\text{ideal}} = \mathcal{U}_n \circ \cdots \circ \mathcal{U}_1(\rho_{\text{in}}).$$

In practice the gates $\mathcal{U}_i$ are implemented imperfectly. Making the standard Markovian assumption on the noise, each imperfect $\mathcal{U}_i$ can be decomposed as $\mathcal{N}_i\mathcal{U}_i$, where each $\mathcal{N}_i$ is a CPTP map and can be distinct for different i. We thus have

$$\rho_{\text{out}}^{\text{exp}} = \mathcal{N}_n \circ \mathcal{U}_n \circ \cdots \circ \mathcal{N}_1 \circ \mathcal{U}_1(\rho_{\text{in}}) \tag{6.19}$$

where ρ_{in} is the input quantum state, ρ_{out} is the quantum state came out of the noisy circuits, $\mathcal{U}_i$ are the desired operations, and $\mathcal{N}_i$ are the noise channels corresponding to gate $\mathcal{U}_i$.

To perform QEM, one first tries to learn (part or all of) the noise models, then recover the ideal gates through either physical or numerical means. Thus, if we wish to analyze the performance of the *best possible* QEM strategy, we may wish that all $\mathcal{N}_i$'s are known exactly. In reality, these $\mathcal{N}_i$'s are obtained from experiments either during the calibration stage or as part of the QEM process, which necessarily involves inaccuracies when being reconstructed. Denote the experimentally characterized noise models $\tilde{\mathcal{N}}_i$, and let $\tilde{\mathcal{N}}_i^{-1}$ denote the inverse of $\tilde{\mathcal{N}}_i$. In this section we will consider channels with the same input and output dimensions.

First, consider the case where $\mathcal{N}_i^{-1}$ exists and is CPTP for all i. Recall that this is true *iff* $\mathcal{N}_i$ is a unitary channel when the input and output dimensions equal. Then in principle one can insert an additional gate implementing $\mathcal{N}_i^{-1}$ after each $\mathcal{U}_i$ to fully invert the noise effect [123]. In reality, the experimentally obtained noise models are $\tilde{\mathcal{N}}_i$. Thus, the output from this method will be

$$\rho_{\text{EM}} = \tilde{\mathcal{N}}_n^{-1} \circ \mathcal{N}_n \circ \mathcal{U}_n \circ \cdots \circ \tilde{\mathcal{N}}_1^{-1} \circ \mathcal{N}_1 \circ \mathcal{U}_1(\rho_{\text{in}}). \tag{6.20}$$

Naturally, there are two main sources of additional errors. First, the experimentally learned noise model $\tilde{\mathcal{N}}_i$ is not always equal to $\mathcal{N}_i$, so $\tilde{\mathcal{N}}_i \circ \mathcal{N}_i$ is not necessarily equal to the identity. Second, even if $\mathcal{N}_i$ can be learned ideally, physically implementing $\mathcal{N}_i^{-1}$ will also not be ideal and can introduce extra errors.

Next, consider the case where $\mathcal{N}_i^{-1}$ exists but is not CPTP. In this case it is impossible to physically restore the ideal output state. However, we can still perform the inverse numerically to recover the output density matrix. This can be thought of as first numerically inverting all the channels in Eq. (6.19), then applying the ideal gates in the original order. Specifically, we define the ideal "reversal" channel $\mathcal{R}$ to be the one that maps $\rho_{\text{out}}^{\text{exp}}$ to ρ_{in}, constructed as

$$\mathcal{R} \coloneqq \mathcal{U}_1^\dagger \circ \mathcal{N}_1^{-1} \circ \cdots \circ \mathcal{U}_n^\dagger \circ \mathcal{N}_n^{-1}. \tag{6.21}$$

Correspondingly, replacing $\mathcal{N}_i$ in the above by $\tilde{\mathcal{N}}_i$ gives the realistic reversal channel,

$$\tilde{\mathcal{R}} \coloneqq \mathcal{U}_1^\dagger \circ \tilde{\mathcal{N}}_1^{-1} \circ \cdots \circ \mathcal{U}_n^\dagger \circ \tilde{\mathcal{N}}_n^{-1}, \tag{6.22}$$

which represents the experimentalist's best knowledge about the ideal reversal channel $\mathcal{R}$. We thus have

$$\rho_{\text{EM}} = \mathcal{U}_{n\cdots 1} \circ \tilde{\mathcal{R}}(\rho_{\text{out}}^{\text{exp}}) \tag{6.23}$$

where the shorthand $\mathcal{U}_{n\cdots 1} := \mathcal{U}_n \circ \cdots \circ \mathcal{U}_1$ is used for the ideal circuit sequence. The composition of such channels first maps the experimental output state $\rho_{\text{out}}^{\text{exp}}$ back to the input state ρ_{in}, then perform the ideal operations $\mathcal{U}_{n\cdots 1}$; this is illustrated by the blue arrows in Fig. 6.5. The numerical inverse method does not involve implementing physical gates, but still require that noise processes are accurately characterized. As reviewed in Section 6.1, we can consider Eq. (6.23) to be the desired output of QEM protocols.

A naive numerical implementation of the channel inverse requires simulating the quantum circuit $\mathcal{U}_{n\cdots 1}$, which is naturally expensive. Generally speaking, the computational complexity for computing Eq. (6.23) can be higher than classically simulating the ideal circuit, even without including the cost of characterizing noise channels. Therefore, we regard Eq. (6.23) as a theoretical tool for upper bounding the performance of QEM, rather than a specific procedure. An alternative use of this expression is to contrast the error mitigated results with the ideal ones, which measures one's knowledge about the noise in the device. Recalling from Chapter 2, this measure is gauge invariant, yet the exact operational meaning of it remains unclear.

Finally, we mention briefly that if only an approximate version of the ideal output is wanted, it may be sufficient to apply one effective recovery map $\mathcal{N}_{\text{eff}}^{-1}$ to the noisy output state, in the hope that it will eliminate most of the noise effects. The mitigated output from this approximate method is given by

$$\rho_{\text{EM}} = \mathcal{N}_{\text{eff}}^{-1}(\rho_{\text{out}}^{\text{exp}}). \tag{6.24}$$

The effective recover map $\mathcal{N}_{\text{eff}}^{-1}$ may again be applied either physically or numerically. Typically, $\mathcal{N}_{\text{eff}}^{-1}$ contains a few tunable parameters which can be experimentally optimized to achieve the best noise-mitigating performance on some test runs. Methods that fall into this category include decoherence compensation in NMR experiments, and depolarizing-model-based EM [140]. Recent work also considered continuous inversion through the Petz recovery map [141]. Although it leads to resource savings in practice, this general approach can be problematic in some instances due to channel mismatching. We provide one such example in Example 6.4.1.

6.4.2 QEM with Imperfect Knowledge on Noise

Next we study the effects of imperfectly characterized noise channels on the performance of QEM. It is generally acknowledged that characterizing noise models in a quantum system is highly resource demanding [142]. In many current error mitigation protocols, the noise channel is assumed to have a simple form [143], such as being a depolarizing channel with

a few parameters. It is then natural to ask the question of how incorrectly characterized noise channels $\{\mathcal{N}_i\}$ would affect the mitigation outcome. As mentioned before, we will assume all $\{\mathcal{N}_i\}$'s to be invertible in this subsection.

Fig. 6.5 shows an illustrative diagram showing the relationship between different objects under this discussion. One sees that while the ideal circuits and the experimental operations are CPTP maps, the channels $\mathcal{R}$ and $\tilde{\mathcal{R}}$ are not necessarily CPTP anymore. The gap $\tilde{\mathcal{N}}_i - \mathcal{N}_i$ between the estimations $\tilde{\mathcal{N}}_i$ and the actual channels $\mathcal{N}_i$ upper bounds the result of EM, independent of how the inverses are achieved. This gap only affects the difference between $\mathcal{R}$ and $\tilde{\mathcal{R}}$.

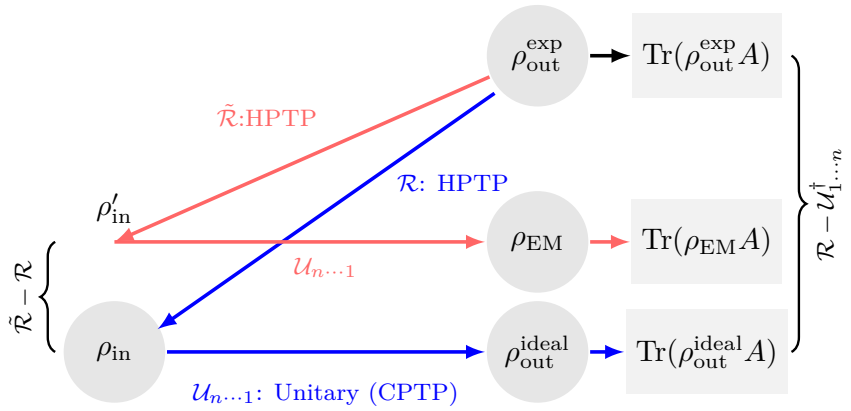

Figure 6.5: The schematic diagram of maps. The blue arrows indicate the map $\mathcal{U}_{n\cdots 1} \circ \mathcal{R}$ for ideal error mitigation, and the red arrows indicate the map $\mathcal{U}_{n\cdots 1} \circ \tilde{\mathcal{R}}$ for error mitigation with imperfect noise characterization. The error between actual noise channels $\mathcal{N}_i$ and estimations $\tilde{\mathcal{N}}_i$ cause the difference between $\mathcal{R}$ and $\tilde{\mathcal{R}}$, which leads to a deviation in the mitigated result.

From the perspective of output states, the goal of EM is to bring the output states closer to the ideal. In terms of state fidelity, this is to ensure that

$$F(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}}) > F(\rho_{\text{out}}^{\text{exp}}, \rho_{\text{out}}^{\text{ideal}}), \tag{6.25}$$

where $F(\rho_1, \rho_2) \coloneqq \text{tr}\big(\sqrt{\sqrt{\rho_1}\rho_2\sqrt{\rho_1}}\big)$ is the fidelity between ρ_1 and ρ_2.

If the actual noise channels $\{\mathcal{N}_i\}$ are invertible and the noise characterization is perfect ($\tilde{\mathcal{N}}_i = \mathcal{N}_i$), theoretically the errors can be perfectly mitigated, with Eq. (6.25) naturally satisfied. Realistically, $\tilde{\mathcal{N}}_i \neq \mathcal{N}_i$, which opens the gap between ideal output states $\rho_{\text{out}}^{\text{ideal}}$ and error mitigated state ρ^{EM}. We next answer the question of how much effect will an imperfect characterization of $\mathcal{N}$ has on the mitigated output fidelity.

Let $\Delta\mathcal{N}_i \coloneqq \tilde{\mathcal{N}}_i - \mathcal{N}_i$ and $\Delta\mathcal{N}_i^{-1} \coloneqq \tilde{\mathcal{N}}_i^{-1} - \mathcal{N}_i^{-1}$[1]. Fig. 6.5 shows that the errors $\{\Delta\mathcal{N}_i^{-1}\}$ only affect $\mathcal{R}$ and $\tilde{\mathcal{R}}$ in the error mitigation maps. The difference between ρ_{EM} and $\rho_{\text{out}}^{\text{ideal}}$ is

$$\begin{aligned}\rho_{\text{EM}} - \rho_{\text{out}}^{\text{ideal}} &= \mathcal{U}_{n\cdots 1} \circ \tilde{\mathcal{R}}(\rho_{\text{out}}^{\text{exp}}) - \mathcal{U}_{1\cdots n}(\rho_{\text{in}}) \\ &= \mathcal{U}_{n\cdots 1} \circ \left[\tilde{\mathcal{R}} - \mathcal{R}\right](\rho_{\text{out}}^{\text{exp}}) \qquad (6.26) \\ &\coloneqq \mathcal{U}_{n\cdots 1} \circ \Delta\mathcal{N}(\rho_{\text{out}}^{\text{exp}}). \qquad (6.27)\end{aligned}$$

In the middle bracket in Eq. (6.26), the errors $\{\Delta\mathcal{N}_i^{-1}\}$ scramble in the layers of unitaries $\mathcal{U}_i^\dagger$. Denote the first order estimation of $\Delta\mathcal{N}$ to be $\Delta\mathcal{N}^{(1)}$, where each term in $\Delta\mathcal{N}^{(1)}$ only contain one of $\Delta\mathcal{N}_i^{-1}$ (see Eq. (D.4) in Appendix D.2 for the explicit expression). The first order error between states is $\Delta\rho_{\text{EM}} \coloneqq \mathcal{U}_{1\cdots n} \circ \Delta\mathcal{N}^{(1)}(\rho_{\text{out}}^{\text{exp}})$. We then define $F(\rho_{\text{EM}}, \rho_{\text{EM}} + \Delta\rho_{\text{EM}})$ to be the first order estimation $F^{(1)}(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}})$ of the state fidelity $F(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}})$. The following proposition gives a bound on this quantity.

Proposition 6.4.1. *The first order estimation of fidelity between ρ_{EM} and ρ_{out}^{ideal} is*

$$\begin{aligned}&\left(1 - \frac{1}{2}\sqrt{d}C_{exp}\left\|\mathbf{v}(\Delta\mathcal{N}^{(1)})\right\|\right)^2 \leq F^{(1)}(\rho_{EM}, \rho_{out}^{ideal}) \\ &\leq 1 - \frac{1}{4}\left(l_U \cdot \left\|\mathbf{v}(\Delta\mathcal{N}^{(1)})\mathbf{v}(\rho_{out}^{exp})\right\|\right)^2, \qquad (6.28)\end{aligned}$$

where $C_{exp} \coloneqq \|\mathbf{v}(\mathcal{U}_{n\cdots 1})\| \cdot \|\mathbf{v}(\rho_{out}^{exp})\|$ is an experiment-related constant, and

$$l_U \coloneqq \inf_{\|x\|=1} \|\mathbf{v}(\mathcal{U}_{n\cdots 1})x\|$$

is the lower Lipschitz constant of the ideal operations $\mathcal{U}_{n\cdots 1}$. The norm $\|\cdot\|$ is 2-norm for vectors and is the induced matrix norm for matrices.

We see that $F^{(1)}(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}})$ is bounded by $\Delta\mathcal{N}^{(1)}$ and experimental constants (including norms of the ideal circuits and the Frobenius norm of the experimental outcome state $\rho_{\text{out}}^{\text{exp}}$). Therefore, by bounding the errors $\{\Delta\mathcal{N}_i^{-1}\}$ in channel estimation, one can constrain the fidelity by using Eq. (6.28). In fact, this result can be understood easily from the left-hand side of Fig. 6.5 – closing the gap between $\tilde{\mathcal{R}}$ and $\mathcal{R}$ can bring ρ_{in}' and ρ_{in} closer, therefore bounding the fidelity afterwards. Further details can be found in Appendix D.2.

[1]Note that $\Delta\mathcal{N}_i$ and $\Delta\mathcal{N}_i^{-1}$ are related by $\Delta\mathcal{N}_i\tilde{\mathcal{N}}_i^{-1} + \mathcal{N}_i\Delta\mathcal{N}_i^{-1} = 0$.

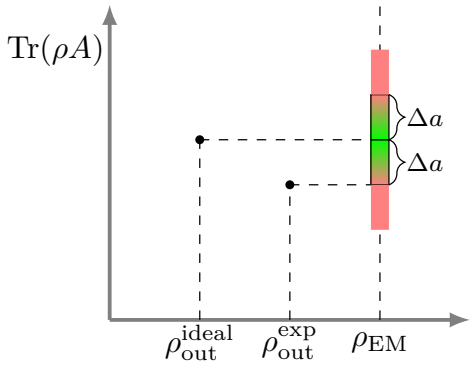

Figure 6.6: A schematic diagram for improving the expectation value $\mathrm{Tr}(\rho A)$ using QEM. Δa denotes the difference (in absolute value) between the ideal output, $\rho_{\mathrm{out}}^{\mathrm{ideal}}$, and the actual output $\rho_{\mathrm{out}}^{\mathrm{exp}}$. The goal of QEM is to achieve a mitigated output state ρ_{EM}, such that $\mathrm{Tr}(\rho_{\mathrm{EM}} A)$ in the green zone for an observable A of interest.

If the task realized by the given circuit only concerns the expectation value of a set of observables $\{A_i\}$, then the goal of QEM can be simplified as recovering the ideal expectation value, $\mathrm{Tr}\big(\rho_{\mathrm{out}}^{\mathrm{ideal}} A_i\big)$. As shown in Fig. 6.6, one would like the error mitigated result to be closer to the ideal than the one directly coming from experiments. Since one cannot perfectly characterize the noise models $\mathcal{N}_i$, it is desirable to know the condition which guarantees $\mathrm{Tr}(\rho_{\mathrm{EM}} A)$ to land in the green region. We show in Appendix D.3 that the following is a sufficient condition for such a goal.

Proposition 6.4.2. *If Eq. (6.29) is satisfied, Quantum Error Mitigation has the ability to improve the expectation value of any observable A for any circuit $\mathcal{U}_{n\cdots 1}$:*

$$\|\mathbf{v}(\Delta\mathcal{N})\| \leq l_{\textit{ideal-exp}}, \tag{6.29}$$

where $l_{\textit{ideal-exp}} := \inf_{\|x\|=1} \left\|\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})x\right\|$ *is the lower Lipschitz constant of* $\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})$.

In the above result, the channels $\mathcal{R}$ and $\mathcal{U}_{1\cdots n}^{\dagger}$ maps $\rho_{\mathrm{out}}^{\mathrm{exp}}$ and $\rho_{\mathrm{out}}^{\mathrm{ideal}}$ back to ρ_{in} respectively. The condition Eq. (6.29), in general, is requiring $\Delta\mathcal{N}$ to be smaller than $\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger}$. This can again be observed from the brackets in Fig. 6.5. Since this proposition is for any observables and any circuit, it will also work for quantum state fidelity.

Note that Eq. (6.29) is a stringent requirement. If $\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})$ has a nontrivial null space, then it will force the noise channel estimation $\tilde{\mathcal{N}}_i$ to be perfect, i.e. $\tilde{\mathcal{N}}_i = \mathcal{N}_i$ for $\forall i \in \{1, \cdots, n\}$. This is because we do not make extra assumptions on circuits and noises

while deriving this sufficient condition. In general, having additional information about the circuit and noise channels can loosen this requirement.

Roughly speaking, for an error mitigation protocol to be able to improve the experimental outcome, one's knowledge of the noise channels needs to be more accurate than the experimental accuracy. Since the unitary $\mathcal{U}_{n\cdots 1}$ is an isometry under Frobenius norm ($\|M\|_F \coloneqq \sqrt{\sum_{ij} |M_{ij}|^2}$), the difference between error mitigated state ρ_{EM} and ideal output state $\rho_{\mathrm{out}}^{\mathrm{ideal}}$ equals to $\Delta_\rho \coloneqq \|\rho_{\mathrm{EM}} - \rho_{\mathrm{out}}^{\mathrm{ideal}}\|_F = \|\Delta\mathcal{N}(\rho_{\mathrm{out}}^{\mathrm{exp}})\|_F = \|\tilde{\mathcal{R}}(\rho_{\mathrm{out}}^{\mathrm{exp}}) - \rho_{\mathrm{in}}\|_F$. The error Δ of error mitigation is bounded by Δ_ρ:

$$\begin{aligned}\Delta &= \mathrm{tr}\left[A(\rho_{\mathrm{EM}} - \rho_{\mathrm{out}}^{\mathrm{ideal}})\right] \\ &\leq \|A\|_F \cdot \|\rho_{\mathrm{EM}} - \rho_{\mathrm{out}}^{\mathrm{ideal}}\|_F \\ &= \|A\|_F \cdot \Delta_\rho. \end{aligned} \tag{6.30}$$

Assume that the input state ρ_{in} is the all-zero state $|0,\cdots,0\rangle\langle 0,\cdots,0|$, i.e. ρ_{in} is a very sparse matrix with 1 as the first entry and 0 elsewhere. The noisy $\tilde{\mathcal{R}}$ can be efficiency computed using tensor network methods, and $\rho_{\mathrm{out}}^{\mathrm{ideal}}$ can be constructed by classical shadow tomography [144]. Thus, the upper bound Eq. (6.30) of value Δ can be efficiently computed.

Finally, we discuss potential consequences of incorrect assumptions about the actual noise. Normally, certain noise models are assumed while identifying device noise, which leads to savings in parameters and resources in characterization. However, the distance between the actual noise $\mathcal{N}$ in the system and the model assumed will not be arbitrarily close, which opens a gap between the ideal and error mitigated outcomes. In particular, if the error model is overly simplified, it can cause problems on EM performance.

We consider the following simple example of a depth-1, single qubit quantum channel, where the actual noise $\mathcal{N}$ is a Pauli Channel, but a depolarizing channel is assumed when mitigating error.

Example 6.4.1. Suppose one believes that the noise in the system is mainly depolarizing, and tries to use the depolarizing channel $\mathcal{D}$ to approximate the actual noise. After optimizing the parameters in $\mathcal{D}$, the inverse $\mathcal{D}^{-1}$ of the estimated $\mathcal{D}$ is used to recover information (i.e. $\mathcal{D}^{-1} \circ \mathcal{N}(\rho)$).

The Kraus operators of $\mathcal{N}$ are $\{\sqrt{p_1}I, \sqrt{p_2}X, \sqrt{p_3}Y, \sqrt{(1-p_1-p_2-p_3)}Z\}$, and those for $\mathcal{D}$ are $\{\sqrt{1-\frac{3\lambda}{4}}I, \sqrt{\frac{\lambda}{4}}X, \sqrt{\frac{\lambda}{4}}Y, \sqrt{\frac{\lambda}{4}}Z\}$. For a given set of $\{p_1, p_2, p_3\}$, the optimal λ to minimize $\|\mathcal{N} - \mathcal{D}\|_?$ varies according to different representations and different choices of norm, $\|\cdot\|_?$. The symmetry on the parameters in $\mathcal{D}$ makes it impossible to perfectly capture the noise $\mathcal{N}$ for p_i's that do not have such a symmetry.

Note that the two vectors

$$\vec{n} \coloneqq (\sqrt{p_1}, \sqrt{p_2}, \sqrt{p_3}, \sqrt{(1-p_1-p_2-p_3)})$$
$$\vec{d} \coloneqq (\sqrt{1-\frac{3\lambda}{4}}, \sqrt{\frac{\lambda}{4}}, \sqrt{\frac{\lambda}{4}}, \sqrt{\frac{\lambda}{4}})$$

are also representations for $\mathcal{N}$ and $\mathcal{D}$ respectively. Since $\vec{n}$ and $\vec{d}$ are normalized, minimizing the distance between $\mathcal{N}$ and $\mathcal{D}$ is equivalent to maximizing $\vec{n} \cdot \vec{d}$. i.e.

$$\max_{\lambda \in [0,1]} \left\{ \sqrt{p_1(1-\frac{3\lambda}{4})} + [\sqrt{p_2} + \sqrt{p_3} + \sqrt{(1-p_1-p_2-p_3)}] \sqrt{\frac{\lambda}{4}} \right\}.$$

When $p_1 = \frac{1}{2}$ and $p_2 = p_3 = 0$, a channel will have a phase flip error with probability $\frac{1}{2}$ and will stay unchanged with probability $\frac{1}{2}$, corresponding to an optimal $\lambda_{\max}$ value of $\frac{1}{3}$. This $\lambda_{\max}$ bounds the distance between $\mathcal{N}$ and $\mathcal{D}$ from above for this metric. Assume one fits the parameter λ from experiments, and obtains the estimation that $\lambda = \frac{1}{3}$, the channel $\{\sqrt{\frac{3}{4}}I, \sqrt{\frac{1}{12}}X, \sqrt{\frac{1}{12}}Y, \sqrt{\frac{1}{12}}Z\}$ will be believed to be $\tilde{\mathcal{N}}$. Then $\tilde{\mathcal{N}}^{-1} = \mathcal{D}^{-1}$ will be used to perform error mitigation. In Fig. 6.7, we can see that while the actual channel $\mathcal{N}$ preserves the expectation value of Z, the mitigated results are actually worse due to the incorrect assumptions on noise model (see the blue triangles in Fig. 6.7). Also note that, since $\mathcal{D}^{-1}$ is non-CP, the outputs $\mathcal{D}^{-1} \circ \mathcal{N}(\rho)$ are not valid quantum states anymore. In this case the fidelity function is not bounded below 1, thus is no longer a valid metric. We give further details in Appendix D.4.

While the above is a rather extreme example of channel mismatching, the message in this example is alerting, because it illustrates how a misunderstanding of noise can lead to failures when mitigating errors. Although we can lower bound the fidelity of the error mitigated state ρ_{EM} and $\rho_{\text{out}}^{\text{ideal}}$ from Proposition 6.4.1, mitigating errors to improve the results still imply a competition between the experimental and noise characterization accuracies (Proposition 6.4.2 and Fig. 6.5). In order to improve experimental readout using QEM, the increasing accuracy of the experiments demands better knowledge of device noise, which will translate into a high sampling cost during noise characterization. The additional overhead due to the need for accurately characterizing noise channels is typically ignored when estimating the cost of QEM [145], but should also be taken into account for future analyses on QEM protocols.

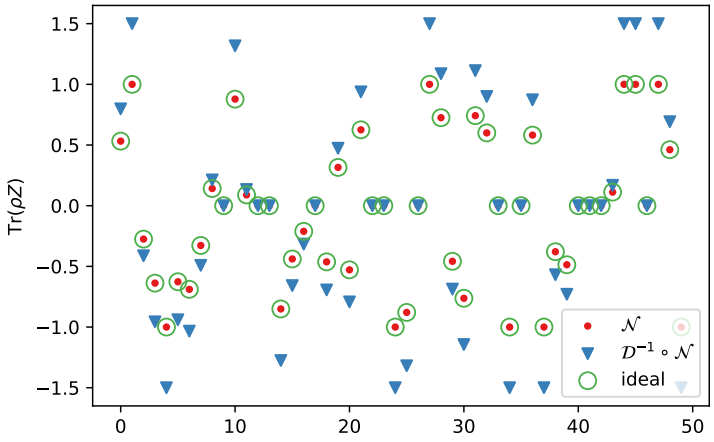

Figure 6.7: Expectation value of Z for 50 randomly generated states, under the noise models in Example 6.4.1. The x-axis is the dummy label for tested states. While this noise does not affect the Z expectation, the incorrectly applied "recovery" operation $\mathcal{D}^{-1}$ leads to incorrect (and sometimes unphysical) outcomes.

6.5 Summary

In this chapter, we have examined several aspects of quantum error mitigation. Using intuitions from classical and quantum communication, we clarify that QEM is fundamentally different from QEC because each method has distinct goals, namely, QEC preserves the physical objects themselves while QEM restores the mathematical descriptions (i.e., density matrices) of the physical objects. This motivates that inverses of the noise channels play an important role when evaluating the ultimate performance of QEM protocols. For the case of invertible but non-unitary noise where the inverse cannot be physically implemented, we show that one may decompose the (HPTP) inverse as multiple implementable CPTP channels, and recombine the results to obtain the inverse using post-processing. For the case of non-invertible channels, we explicitly provide a generalized-inverse construction called Drazin-inverse, and prove that the Drazin-inverse of any channel is trace preserving, while another commonly used Moore-Penrose generalized-inverse may not be.

When the noise channels are invertible, the improvement from EM protocols is constrained by our knowledge about the noise in the device of interest. The gap between $\{\tilde{\mathcal{N}}_i\}$ and $\{\mathcal{N}_i\}$ can bound the fidelity between the ideal state and the error mitigated state. A

sufficient condition is derived for guaranteeing an improvement, which is motivated from the idea that the accuracy of noise characterization needs to be higher than that of the experiment. Therefore, more accurate experiments require higher overheads during noise characterization. We derive a first order approximation to the overall error in multi-layer circuits, which can be used to estimate the minimal cost associated with learning the noise before applying EM procedures. We also show that if the mismatch between assumed and real noise channels is too large, error mitigation may yield unphysical outcomes and fail. Overall, our analyses demonstrate that the complications and subtleties when implementing QEM demand more in-depth studies on several topics, such as optimal decomposition of inverse maps, and more realistic estimates on the cost of accurate QEM.

Chapter 7

Conclusion and Outlooks

The presence of noise in quantum processors is a critical issue that affects the output quality of these devices. In this thesis, we have presented several results addressing this important and challenging subject of studying noise in near-term quantum processors, which is crucial for the development and improvement of these devices. We first studied how noise can be accurately represented in a device, by studying a gauge freedom in the representation of quantum operations. After illustrating the relevance of gauge using realistic examples, we propose the mean variation error (MVE) as an operational figure of merit, along with a protocol to measure the MVE. Our numerical study on the MVE suggests a potential underestimation of error in conventional randomized benchmarking experiments. Moreover, it shows that MVE carries useful information regarding the nature of noise occurring in the hardware.

We then study under what assumptions can the gauge ambiguity be removed, and how conventional error rates can be obtained. Our focus has been on state preparation and measurement (SPAM) errors. We have shown that one can unambiguously distinguish state preparation errors from measurement errors, in the presence of an ancillary qubit that is independent of the system of interest. We have also derived bounds on the estimated SPAM error rates based on gate error measures that can be estimated independently of SPAM processes, which has been cross-validated between data from quantum hardware and classical simulations.

An alternative way to remove gauge ambiguity between SPAM errors is by reducing the state preparation noise independently from the measurement noise. We explore this direction by developing a simple algorithmic cooling (AC) protocol which utilizes imperfect measurements. We compute how the overhead using our method can be upper bounded

by measurable quantities, and show that it can efficiently reduce state preparation noise under experimentally feasible regimes. This offers a new pathway to benchmark as well as algorithmically improve existing quantum processors.

Since using AC as a state-purification procedure comes with an inevitable thermodynamic cost, we then studied the cost of various AC protocols. We proposed two related, operationally meaningful measures, and benchmarked the efficiencies of various protocols under the framework of coherent cooling. We studied the cost of various AC protocols using simulations, providing a guideline for the advantageous regime of each protocol. Our focus on thermodynamic costs of AC has also led to interesting discoveries of new protocols that are capable of cooling to the same final temperature with a lower thermodynamic cost.

In addition to algorithmic cooling, the use of numerical methods to reduce effective error rates in noisy hardware has recently emerged to form a field named quantum error mitigation. We contrasted the family of QEM techniques with quantum error correction, pointing out fundamental differences between the two from a communication perspective. We then proposed a new construction of pseudo-inverse to be used in QEM when the noise is non-invertible, and proved important properties of this construction including its trace-preserving nature. We also studied the problem of under what condition can QEM provide definitive improvements to the final results, when one's knowledge about the noise is imperfect. Our results suggest that more work is needed for both applying and estimating the cost of accurate QEM protocols.

Outlook

There are many possible future research directions derived from this thesis that could be worth further investigating. First, the topic of gauge-invariant quality metrics is worth a closer look. For example, when defining the mean variation error in Eq. (2.26), there may exist more reasonable practices than assigning equal weights to all experiments with the same length, such as adaptively giving more weights to certain circuits for a given computational task. This is particularly relevant when a small subset of the operations have much higher error rates than the rest, such that for a task that relies heavily on these elements, the MVE will not correctly reflect the true performance of the device. Alternatively, it is possible to define a different measure based on the variational distance, but using tailored circuits instead of fully random ones, in order to amplify certain types of errors in the device.

Second, regarding the work of distinguishing SPAM errors using noise propagation, we have yet to optimize the protocol to improve the measurement overhead. For large target

systems, we expect the real density matrix to be close to a sparse matrix, with many vanishingly small elements. In such cases it may be beneficial to integrate compressed-sensing techniques [146] to more efficiently estimate the whole density matrix. Another possibility to improve the scaling is by placing further assumptions on the form of the initial state. For example, it is typically reasonable to assume that components corresponding to low-weight errors on the ideal state are much larger than those corresponding to high-weight errors. Therefore, sampling only the major components could significantly improve the scaling of the protocol, while sacrificing only slightly on the accuracy.

On the other hand, there may be interesting theoretical relationships between our assumption-based approach with the assumption-free approach in gate-set tomography (GST). It is currently unclear whether there exists a set of additional assumptions (i.e., constraints) to the general GST protocol that, upon taking into consideration, will fully eliminate the gauge freedom. Our results points towards a positive answer to this question, and further studies along these lines could lead to a reduction in the GST overhead when the underlying physical system satisfies these assumptions.

Third, similar arguments can be applied to improve the SPAM-characterization protocol based on AC. Using the same principles as in noise propagation, it should be possible take gate errors into account in the final estimates using the AC-based approach [147, 97]. Another natural extension for AC-based protocols is to explore whether it is still possible to reduce state preparation error rates in the presence of correlated initial state errors, demonstrating a fault-tolerance capability of AC. We speculate that this should hold true in the limit of weak correlated error. In addition, it can be very interesting to integrate the concept of MBAC into other algorithmic cooling protocols (see, e.g., the reviews [148, 149]), which goes beyond the BCS subroutine mentioned in Section 4.3, to develop new protocols with potentially better performances.

Forth, thermodynamic studies of AC could be continued as a standalone field that further extends studies on classical heat engines to the quantum regime. On the one hand, the coherent cooling paradigm may be extended to include a wider family of protocols, including those involving measurements such as MBAC as suggested by Chapter 4. Integrating measurement operations into the current quantum resource theory framework could lead to interesting discoveries such as placing a minimum thermodynamic cost on quantum measurements. On the other hand, experimentally measuring the actual amount of energy flow during an AC protocol will provide a better understanding on how far we are from the ideal coherent paradigm, revealing potential sources of energy loss in the experimental design. An experiment along both lines has very recently been demonstrated in a superconducting transmon system [150] and it could be interesting to see if similar results hold in other platforms.

Finally, our exploration on QEM provides insights from the point of view of superoperators, opening up new possibilities for applying generalized matrix inverse theory to tackle this problem. Properties of different generalized inverses and when to use them require further investigation. For practical QEM, more rigorous studies on the impact of imperfect characterization on the noise channel is needed for each protocol. This will provide not only more accurate error bars for QEM-assisted algorithms, but also more realistic overhead estimates in order to achieve a target accuracy. Practical ways to implement QEC in conjunction with QEM, as well as alternative error-reducing procedures such as AC, is likely to remain an active research area in the near future. Overall, it is fair to say that more needs to be done in representing, characterizing, and mitigating noise in quantum processors, and progresses in these fields will be essential stepping stones towards building large-scale, universal quantum computers.

References

[1] D. Deutsch and R. Jozsa, "Rapid Solution of Problems by Quantum Computation," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 439, pp. 553–558, dec 1992.

[2] L. K. Grover, "A fast quantum mechanical algorithm for database search," *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp. 212–219, 1996.

[3] D. R. Simon, "On the Power of Quantum Computation," *SIAM Journal on Computing*, vol. 26, pp. 1474–1483, oct 1997.

[4] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM Review*, vol. 41, pp. 303–332, jan 1999.

[5] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, aug 2018.

[6] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, pp. 505–510, oct 2019.

[7] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, "Quantum computational advantage using photons," *Science*, vol. 370, pp. 1460–1463, dec 2020.

[8] F. Pan, K. Chen, and P. Zhang, "Solving the Sampling Problem of the Sycamore Quantum Circuits," *Physical Review Letters*, vol. 129, p. 090502, aug 2022.

[9] Q. Xu, N. Mannucci, A. Seif, A. Kubica, S. T. Flammia, and L. Jiang, "Tailored XZZX codes for biased noise," *Physical Review Research*, vol. 5, p. 013035, jan 2023.

[10] U. Fano, "Description of states in quantum mechanics by density matrix and operator techniques," *Reviews of modern physics*, vol. 29, no. 1, p. 74, 1957.

[11] K. Vogel and H. Risken, "Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase," *Physical Review A*, vol. 40, no. 5, p. 2847, 1989.

[12] G. M. D'Ariano, M. G. Paris, and M. F. Sacchi, "Quantum tomography," *Advances in imaging and electron physics*, vol. 128, pp. 206–309, 2003.

[13] G. M. D'Ariano, C. Macchiavello, and M. G. A. Paris, "Detection of the density matrix through optical homodyne tomography without filtered back projection," *Phys. Rev. A*, vol. 50, pp. 4298–4302, Nov 1994.

[14] S. Wallentowitz and W. Vogel, "Unbalanced homodyning for quantum state measurements," *Phys. Rev. A*, vol. 53, pp. 4528–4533, Jun 1996.

[15] Z. Hradil, "Quantum-state estimation," *Physical Review A*, vol. 55, no. 3, pp. R1561–R1564, 1997.

[16] R. Blume-Kohout, "Optimal, reliable estimation of quantum states," *New Journal of Physics*, vol. 12, p. 043034, apr 2010.

[17] Y. Quek, S. Fort, and H. K. Ng, "Adaptive quantum state tomography with neural networks," *npj Quantum Information*, vol. 7, p. 105, jun 2021.

[18] J. Fiurášek, "Maximum-likelihood estimation of quantum measurement," *Physical Review A*, vol. 64, p. 024102, jul 2001.

[19] D. Mogilevtsev, A. Ignatenko, A. Maloshtan, B. Stoklasa, J. Rehacek, and Z. Hradil, "Data pattern tomography: Reconstruction with an unknown apparatus," *New Journal of Physics*, vol. 15, p. 025038, feb 2013.

[20] A. C. Keith, C. H. Baldwin, S. Glancy, and E. Knill, "Joint quantum-state and measurement tomography with incomplete measurements," *Physical Review A*, vol. 98, p. 042318, oct 2018.

[21] A. I. Lvovsky and M. G. Raymer, "Continuous-variable optical quantum-state tomography," *Reviews of Modern Physics*, vol. 81, no. 1, pp. 299–332, 2009.

[22] R. Blume-Kohout, J. K. Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz, "Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit," *arXiv preprint arXiv:1310.4492*, 2013.

[23] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, "Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography," *Nature Communications*, vol. 8, p. 14485, feb 2017.

[24] C. Jackson and S. J. van Enk, "Detecting correlated errors in state-preparation-and-measurement tomography," *Physical Review A*, vol. 92, p. 042312, oct 2015.

[25] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, "Self-consistent quantum process tomography," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 87, no. 6, p. 62119, 2013.

[26] A. Gilchrist, N. K. Langford, and M. A. Nielsen, "Distance measures to compare real and ideal quantum processes," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 71, p. 062310, jun 2005.

[27] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout, "What randomized benchmarking actually measures," *arXiv preprint arXiv:1702.01853*, 2017.

[28] Ł. Rudnicki, Z. Puchała, and K. Zyczkowski, "Gauge invariant information concerning quantum channels," *Quantum*, vol. 2, p. 60, apr 2018.

[29] J. J. Wallman, "Randomized benchmarking with gate-dependent noise," *Quantum*, vol. 2, p. 47, jan 2018.

[30] J. Emerson, R. Alicki, and K. Życzkowski, "Scalable noise estimation with random unitary operators," *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 7, pp. S347–S352, oct 2005.

[31] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, "Randomized benchmarking of quantum gates," *Physical Review A*, vol. 77, p. 012307, jan 2008.

[32] E. Magesan, J. M. Gambetta, and J. Emerson, "Scalable and Robust Randomized Benchmarking of Quantum Processes," *Physical Review Letters*, vol. 106, p. 180504, may 2011.

[33] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, "Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking," *Physical Review Letters*, vol. 109, p. 080505, aug 2012.

[34] A. Carignan-Dugas, J. J. Wallman, and J. Emerson, "Characterizing universal gate sets via dihedral benchmarking," *Physical Review A*, vol. 92, p. 060302(R), dec 2015.

[35] S. Sheldon, L. S. Bishop, E. Magesan, S. Filipp, J. M. Chow, and J. M. Gambetta, "Characterizing errors on qubit operations via iterative randomized benchmarking," *Physical Review A*, vol. 93, p. 012301, jan 2016.

[36] A. W. Cross, E. Magesan, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, "Scalable randomised benchmarking of non-Clifford gates," *npj Quantum Information*, vol. 2, p. 16012, nov 2016.

[37] A. Carignan-Dugas, K. Boone, J. J. Wallman, and J. Emerson, "From randomized benchmarking experiments to gate-set circuit fidelity: how to interpret randomized benchmarking decay parameters," *New Journal of Physics*, vol. 20, p. 092001, sep 2018.

[38] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, "Randomized benchmarking of quantum gates," *Physical Review A*, vol. 77, p. 012307, jan 2008.

[39] J. M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, J. A. Smolin, J. M. Chow, C. A. Ryan, C. Rigetti, S. Poletto, T. A. Ohki, M. B. Ketchen, and M. Steffen, "Characterization of Addressability by Simultaneous Randomized Benchmarking," *Physical Review Letters*, vol. 109, p. 240504, dec 2012.

[40] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland, "Randomized Benchmarking of Multiqubit Gates," *Physical Review Letters*, vol. 108, p. 260503, jun 2012.

[41] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, "Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits," *Physical Review Applied*, vol. 7, p. 054020, may 2017.

[42] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, "Superconducting quantum circuits at the surface code threshold for fault tolerance," *Nature*, vol. 508, pp. 500–503, apr 2014.

[43] R. Landauer, "Irreversibility and Heat Generation in the Computing Process," *IBM J. Res. Dev.*, vol. 5, no. 3, pp. 183–191, 1961.

[44] K. Temme, S. Bravyi, and J. M. Gambetta, "Error Mitigation for Short-Depth Quantum Circuits," *Physical Review Letters*, vol. 119, p. 180509, nov 2017.

[45] S. Endo, S. C. Benjamin, and Y. Li, "Practical Quantum Error Mitigation for Near-Future Applications," *Physical Review X*, vol. 8, p. 031027, jul 2018.

[46] S. McArdle, X. Yuan, and S. Benjamin, "Error-Mitigated Digital Quantum Simulation," *Physical Review Letters*, vol. 122, p. 180501, may 2019.

[47] F. B. Maciejewski, Z. Zimborás, and M. Oszmaniec, "Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography," *Quantum*, vol. 4, p. 257, apr 2020.

[48] B. Koczor, "Exponential Error Suppression for Near-Term Quantum Devices," *Physical Review X*, vol. 11, p. 031057, sep 2021.

[49] R. A. Bertlmann and P. Krammer, "Bloch vectors for qudits," *Journal of Physics A: Mathematical and Theoretical*, vol. 41, p. 235303, jun 2008.

[50] D. Greenbaum, "Introduction to quantum gate set tomography," *arXiv preprint arXiv:1509.02921*, 2015.

[51] C. Jackson and S. van Enk, "Nonholonomic tomography. I. The Born rule as a connection between experiments," *Physical Review A*, vol. 95, p. 052327, may 2017.

[52] S. Kimmel, M. P. da Silva, C. A. Ryan, B. R. Johnson, and T. Ohki, "Robust Extraction of Tomographic Information via Randomized Benchmarking," *Physical Review X*, vol. 4, p. 011050, mar 2014.

[53] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, "Comparing Experiments to the Fault-Tolerance Threshold," *Physical Review Letters*, vol. 117, p. 170502, oct 2016.

[54] G. Brassard and P. Raymond-Robichaud, "The equivalence of local-realistic and no-signalling theories," *arXiv preprint arXiv:1710.01380*, 2017.

[55] L. Hardy, "Quantum theory from five reasonable axioms," *arXiv preprint quant-ph/0101012*, 2001.

[56] J. Watrous, *The Theory of Quantum Information*. Cambridge University Press, apr 2018.

[57] J. Medford, J. Beil, J. M. Taylor, S. D. Bartlett, A. C. Doherty, E. I. Rashba, D. P. DiVincenzo, H. Lu, A. C. Gossard, and C. M. Marcus, "Self-consistent measurement and state tomography of an exchange-only spin qubit," *Nature Nanotechnology*, vol. 8, pp. 654–659, sep 2013.

[58] G. Brida, L. Ciavarella, I. P. Degiovanni, M. Genovese, L. Lolli, M. G. Mingolla, F. Piacentini, M. Rajteri, E. Taralli, and M. G. A. Paris, "Quantum characterization of superconducting photon counters," *New Journal of Physics*, vol. 14, p. 085001, aug 2012.

[59] R. J. Blume-Kohout, J. K. Gamble, E. Nielsen, P. L. W. Maunz, T. L. Scholten, and K. M. Rudinger, "Turbocharging Quantum Tomography," tech. rep., Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA (United States), jan 2015.

[60] S. Mavadia, C. L. Edmunds, C. Hempel, H. Ball, F. Roy, T. M. Stace, and M. J. Biercuk, "Experimental quantum verification in the presence of temporally correlated noise," *npj Quantum Information*, vol. 4, p. 7, dec 2018.

[61] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient Z gates for quantum computing," *Physical Review A*, vol. 96, p. 022330, aug 2017.

[62] E. Knill, R. Laflamme, R. Martinez, and C.-H. Tseng, "An algorithmic benchmark for quantum information processing," *Nature*, vol. 404, pp. 368–370, mar 2000.

[63] D. A. Levin, Y. Peres, E. L. Wilmer, J. Propp, and D. B. Wilson, *Markov chains and mixing times*. Providence, Rhode Island: American Mathematics Society, 2017.

[64] J. J. Wallman, "Bounding experimental quantum error rates relative to fault-tolerant thresholds," *arXiv preprint arXiv:1511.00727*, 2015.

[65] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, "Estimating the coherence of noise," *New Journal of Physics*, vol. 17, p. 113020, nov 2015.

[66] W. Hoeffding, "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, vol. 58, pp. 13–30, mar 1963.

[67] S.-O. Chan, I. Diakonikolas, P. Valiant, and G. Valiant, "Optimal Algorithms for Testing Closeness of Discrete Distributions," in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, (Philadelphia, PA), pp. 1193–1203, Society for Industrial and Applied Mathematics, jan 2014.

[68] D. Gottesman, "The Heisenberg Representation of Quantum Computers," *arXiv preprint arXiv:quant-ph/9807006*, 1998.

[69] M. V. den Nest, "Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond," *arXiv preprint arXiv:0811.0898*, 2008.

[70] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis, "Optimal Quantum Control Using Randomized Benchmarking," *Physical Review Letters*, vol. 112, p. 240504, jun 2014.

[71] L. Viola and E. Knill, "Random Decoupling Schemes for Quantum Dynamical Control and Error Suppression," *Physical Review Letters*, vol. 94, p. 060502, feb 2005.

[72] R. Blume-Kohout, J. K. Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz, "Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit," *arXiv preprint arXiv:1310.4492*, oct 2013.

[73] J. Lin, B. Buonacorsi, R. Laflamme, and J. J. Wallman, "On the freedom in representing quantum operations," *New Journal of Physics*, vol. 21, p. 023006, feb 2019.

[74] J. J. Wallman and J. Emerson, "Noise tailoring for scalable quantum computation via randomized compiling," *Physical Review A*, vol. 94, p. 052325, nov 2016.

[75] M. H. Levitt, *Spin dynamics: Basics of Nuclear Magnetic Resonance*. Wiley, 2 ed., 2001.

[76] A. Sørensen and K. Mølmer, "Quantum Computation with Ions in Thermal Motion," *Physical Review Letters*, vol. 82, pp. 1971–1974, mar 1999.

[77] A. Sørensen and K. Mølmer, "Entanglement and quantum computation with ions in thermal motion," *Physical Review A*, vol. 62, p. 022311, jul 2000.

[78] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, "Characterizing large-scale quantum computers via cycle benchmarking," *Nature Communications*, vol. 10, p. 5347, dec 2019.

[79] IBM Q team, "IBM Q 5 Santiago backend specification V1.3.6," 2019.

[80] S. J. Beale, A. Carignan-Dugas, D. Dahlen, J. Emerson, I. Hincks, A. Jain, D. Hufnagel, E. Ospadov, J. Saunders, J. J. Wallman, and A. Winick, "True-Q software," jun 2020.

[81] O. W. Sørensen, "Polarization transfer experiments in high-resolution nmr spectroscopy," *Progress in Nuclear Magnetic Resonance Spectroscopy*, vol. 21, no. 6, pp. 503–569, 1989.

[82] O. W. Sørensen, "A universal bound on spin dynamics," *Journal of Magnetic Resonance (1969)*, vol. 86, no. 2, pp. 435–440, 1990.

[83] L. J. Schulman and U. V. Vazirani, "Molecular scale heat engines and scalable quantum computation," in *Proceedings of the thirty-first annual ACM symposium on Theory of computing - STOC '99*, (New York, New York, USA), pp. 322–329, ACM Press, 1999.

[84] P. O. Boykin, T. Mor, V. Roychowdhury, F. Vatan, and R. Vrijen, "Algorithmic cooling and scalable NMR quantum computers," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, pp. 3388–3393, mar 2002.

[85] J. M. Fernandez, S. Lloyd, T. Mor, and V. Roychowdhury, "Algorithmic Cooling of Spins: A Practicable Method for Increasing Polarization," *International Journal of Quantum Information*, vol. 02, pp. 461–477, dec 2004.

[86] Y. Elias, T. Mor, and Y. Weinstein, "Semioptimal practicable algorithmic cooling," *Physical Review A*, vol. 83, p. 042340, apr 2011.

[87] J. Baugh, O. Moussa, C. A. Ryan, A. Nayak, and R. Laflamme, "Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance," *Nature*, vol. 438, pp. 470–473, nov 2005.

[88] C. Ryan, O. Moussa, J. Baugh, and R. Laflamme, "Spin Based Heat Engine: Demonstration of Multiple Rounds of Algorithmic Cooling," *Physical Review Letters*, vol. 100, p. 140501, apr 2008.

[89] Y. Elias, H. Gilboa, T. Mor, and Y. Weinstein, "Heat-bath cooling of spins in two amino acids," *Chemical Physics Letters*, vol. 517, pp. 126–131, dec 2011.

[90] G. Brassard, Y. Elias, J. M. Fernandez, H. Gilboa, J. A. Jones, T. Mor, Y. Weinstein, and L. Xiao, "Experimental heat-bath cooling of spins," *The European Physical Journal Plus*, vol. 129, p. 266, dec 2014.

[91] D. K. Park, G. Feng, R. Rahimi, S. Labruyère, T. Shibata, S. Nakazawa, K. Sato, T. Takui, R. Laflamme, and J. Baugh, "Hyperfine spin qubits in irradiated malonic acid: heat-bath algorithmic cooling," *Quantum Information Processing*, vol. 14, pp. 2435–2461, jul 2015.

[92] N. A. Rodríguez-Briones, J. Li, X. Peng, T. Mor, Y. Weinstein, and R. Laflamme, "Heat-bath algorithmic cooling with correlated qubit-environment interactions," *New Journal of Physics*, vol. 19, p. 113047, nov 2017.

[93] N. A. Rodríguez-Briones, E. Martín-Martínez, A. Kempf, and R. Laflamme, "Correlation-Enhanced Algorithmic Cooling," *Physical Review Letters*, vol. 119, p. 050502, aug 2017.

[94] Á. M. Alhambra, M. Lostaglio, and C. Perry, "Heat-Bath Algorithmic Cooling with optimal thermalization strategies," *Quantum*, vol. 3, p. 188, sep 2019.

[95] N. A. Rodríguez-Briones and R. Laflamme, "Achievable Polarization for Heat-Bath Algorithmic Cooling," *Physical Review Letters*, vol. 116, p. 170501, apr 2016.

[96] L. J. Schulman, T. Mor, and Y. Weinstein, "Physical Limits of Heat-Bath Algorithmic Cooling," *Physical Review Letters*, vol. 94, p. 120501, apr 2005.

[97] J. Lin, J. J. Wallman, I. Hincks, and R. Laflamme, "Independent state and measurement characterization for quantum computers," *Physical Review Research*, vol. 3, p. 033285, sep 2021.

[98] P. J. Shadbolt, M. R. Verde, A. Peruzzo, A. Politi, A. Laing, M. Lobino, J. C. F. Matthews, M. G. Thompson, and J. L. O'Brien, "Generating, manipulating and measuring entanglement and mixture with a reconfigurable photonic circuit," *Nature Photonics*, vol. 6, pp. 45–49, jan 2012.

[99] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pisenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe, and J. Kim, "Benchmarking an 11-qubit quantum computer," *Nature Communications*, vol. 10, p. 5464, dec 2019.

[100] R. R. Soldati, D. B. R. Dasari, J. Wrachtrup, and E. Lutz, "Thermodynamics of a Minimal Algorithmic Cooling Refrigerator," *Physical Review Letters*, vol. 129, p. 030601, jul 2022.

[101] L. J. Schulman and U. Vazirani, "Scalable NMR Quantum Computation," *arXiv preprint arXiv:quant-ph/9804060*, 1998.

[102] L. J. Schulman and U. V. Vazirani, "Molecular scale heat engines and scalable quantum computation," in *Proceedings of the thirty-first annual ACM symposium on Theory of computing - STOC '99*, (New York, New York, USA), pp. 322–329, ACM Press, 1999.

[103] P. O. Boykin, T. Mor, V. Roychowdhury, F. Vatan, and R. Vrijen, "Algorithmic cooling and scalable NMR quantum computers," *Proc. Natl. Acad. Sci.*, vol. 99, pp. 3388–3393, mar 2002.

[104] L. J. Schulman, T. Mor, and Y. Weinstein, "Physical Limits of Heat-Bath Algorithmic Cooling," *SIAM J. Comput.*, vol. 36, pp. 1729–1747, jan 2007.

[105] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, M. P. E. Lock, G. Vitagliano, F. C. Binder, T. Debarba, E. Schwarzhans, F. Clivaz, and M. Huber, "Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?," *arXiv preprint arXiv:2106.05151*, 2021.

[106] M. Lostaglio, Á. M. Alhambra, and C. Perry, "Elementary Thermal Operations," *Quantum*, vol. 2, p. 52, feb 2018.

[107] N. A. Rodríguez-Briones and R. Laflamme, "Achievable Polarization for Heat-Bath Algorithmic Cooling," *Phys. Rev. Lett.*, vol. 116, p. 170501, apr 2016.

[108] J. Anders and V. Giovannetti, "Thermodynamics of discrete quantum processes," *New J. Phys.*, vol. 15, no. 3, p. 033022, 2013.

[109] D. Reeb and M. M. Wolf, "An improved Landauer principle with finite-size corrections," *New Journal of Physics*, vol. 16, p. 103011, 2014.

[110] S. Lorenzo, R. McCloskey, F. Ciccarello, M. Paternostro, and G. M. Palma, "Landauer's Principle in Multipartite Open Quantum System Dynamics," *Physical Review Letters*, vol. 115, p. 120403, sep 2015.

[111] N. Yoshioka, H. Hakoshima, Y. Matsuzaki, Y. Tokunaga, Y. Suzuki, and S. Endo, "Generalized quantum subspace expansion," *arXiv preprint arXiv:2107.02611*, 2021.

[112] Y. Zhang and X. Yuan, "Quantum Error Mitigation: A Review," *Journal of Computer Research and Development*, vol. 58, p. 1843, sep 2021.

[113] Y. Li and S. C. Benjamin, "Efficient Variational Quantum Simulator Incorporating Active Error Minimization," *Physical Review X*, vol. 7, p. 021050, jun 2017.

[114] Y. Chen, M. Farahzad, S. Yoo, and T.-C. Wei, "Detector tomography on IBM quantum computers and mitigation of an imperfect measurement," *Phys. Rev. A*, vol. 100, p. 052315, nov 2019.

[115] S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay, and J. M. Gambetta, "Mitigating measurement errors in multiqubit experiments," *Physical Review A*, vol. 103, p. 042605, apr 2021.

[116] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, "Virtual Distillation for Quantum Error Mitigation," *Physical Review X*, vol. 11, p. 041036, nov 2021.

[117] M. Huo and Y. Li, "Dual-state purification for practical quantum error mitigation," *Physical Review A*, vol. 105, p. 022427, feb 2022.

[118] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, "Low-cost error mitigation by symmetry verification," *Physical Review A*, vol. 98, p. 062339, dec 2018.

[119] N. C. Rubin, R. Babbush, and J. McClean, "Application of fermionic marginal constraints to hybrid quantum algorithms," *New J. Phys.*, vol. 20, p. 053020, may 2018.

[120] E. Knill, R. Laflamme, and L. Viola, "Theory of Quantum Error Correction for General Noise," *Physical Review Letters*, vol. 84, pp. 2525–2528, mar 2000.

[121] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, "Mixed-state entanglement and quantum error correction," *Physical Review A*, vol. 54, pp. 3824–3851, nov 1996.

[122] E. Knill and R. Laflamme, "Theory of quantum error-correcting codes," *Physical Review A*, vol. 55, pp. 900–911, feb 1997.

[123] J. Jiang, K. Wang, and X. Wang, "Physical implementability of linear maps and its application in error mitigation," *Quantum*, vol. 5, p. 600, 2021.

[124] A. Nayak and P. Sen, "Invertible quantum operations and perfect encryption of quantum states," *arXiv preprint quant-ph/0605041*, 2006.

[125] J. Preskill, "Lecture notes for physics 229: Quantum information and computation," *California Institute of Technology*, vol. 16, 1998.

[126] E. H. Moore, "On the reciprocal of the general algebraic matrix," *Bull. Am. Math. Soc.*, vol. 26, pp. 394–395, 1920.

[127] R. Penrose, "A generalized inverse for matrices," *Mathematical proceedings of the Cambridge philosophical society*, vol. 51, pp. 406–413, jul 1955.

[128] M. Drazin, "Pseudo-inverses in associative rings and semigroups," *The American mathematical monthly*, vol. 65, no. 7, pp. 506–514, 1958.

[129] M. M. Wolf and D. Perez-Garcia, "The inverse eigenvalue problem for quantum channels," *arXiv preprint arXiv:1005.4545*, 2010.

[130] V. Karimipour, F. Benatti, and R. Floreanini, "Quasi-inversion of qubit channels," *Physical Review A*, vol. 101, no. 3, p. 032109, 2020.

[131] F. Shahbeigi, K. Sadri, M. Moradi, K. Życzkowski, and V. Karimipour, "Quasi-inversion of quantum and classical channels in finite dimensions," *Journal of Physics A: Mathematical and Theoretical*, vol. 54, no. 34, p. 345301, 2021.

[132] E. Aurell, J. Zakrzewski, and K. Życzkowski, "Time reversals of irreversible quantum maps," *Journal of Physics A: Mathematical and Theoretical*, vol. 48, no. 38, p. 38FT01, 2015.

[133] B. Regula, R. Takagi, and M. Gu, "Operational applications of the diamond norm and related measures in quantifying the non-physicality of quantum maps," *Quantum*, vol. 5, p. 522, 2021.

[134] A. W. Harrow, A. Hassidim, and S. Lloyd, "Quantum Algorithm for Linear Systems of Equations," *Physical Review Letters*, vol. 103, p. 150502, oct 2009.

[135] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature Communications*, vol. 5, p. 4213, sep 2014.

[136] H. J. Kimble, "The quantum internet," *Nature*, vol. 453, pp. 1023–1030, jun 2008.

[137] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead," *Science*, vol. 362, oct 2018.

[138] D. Bacon, "Operator quantum error-correcting subsystems for self-correcting quantum memories," *Physical Review A*, vol. 73, p. 012340, jan 2006.

[139] B. M. Terhal, "Quantum error correction for quantum memories," *Reviews of Modern Physics*, vol. 87, pp. 307–346, apr 2015.

[140] J. Vovrosh, K. E. Khosla, S. Greenaway, C. Self, M. Kim, and J. Knolle, "Efficient mitigation of depolarizing errors in quantum simulations," *arXiv preprint arXiv:2101.01690*, 2021.

[141] H. Kwon, R. Mukherjee, and M. Kim, "Reversing open quantum dynamics via continuous petz recovery map," *arXiv preprint arXiv:2104.03360*, 2021.

[142] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, "Gate set tomography," *Quantum*, vol. 5, p. 557, 2021.

[143] J. Vovrosh, K. E. Khosla, S. Greenaway, C. Self, M. S. Kim, and J. Knolle, "Simple mitigation of global depolarizing errors in quantum simulations," *Physical Review E*, vol. 104, Sep 2021.

[144] H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements," *Nature Physics*, vol. 16, no. 10, pp. 1050–1057, 2020.

[145] R. Takagi, S. Endo, S. Minagawa, and M. Gu, "Fundamental limits of quantum error mitigation," *arXiv preprint arXiv:2109.04457*, 2021.

[146] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, "Quantum state tomography via compressed sensing," *Phys. Rev. Lett.*, vol. 105, p. 150401, Oct 2010.

[147] M. Ben-Or, D. Gottesman, and A. Hassidim, "Quantum Refrigerator," *arXiv preprint arXiv:1301.1995*, 2013.

[148] G. Brassard, Y. Elias, T. Mor, and Y. Weinstein, "Prospects and limitations of algorithmic cooling," *The European Physical Journal Plus*, vol. 129, p. 258, nov 2014.

[149] D. K. Park, N. A. Rodriguez-Briones, G. Feng, R. Rahimi, J. Baugh, and R. Laflamme, "Heat Bath Algorithmic Cooling with Spins: Review and Prospects," in *Biological Magnetic Resonance*, vol. 31, pp. 227–255, Springer, New York, NY, 2016.

[150] X. Linpeng, L. Bresque, M. Maffei, A. N. Jordan, A. Auffèves, and K. W. Murch, "Energetic Cost of Measurements Using Quantum, Coherent, and Thermal Light," *Physical Review Letters*, vol. 128, p. 220506, jun 2022.

[151] M. A. Nielsen and I. L. Chuang, "Quantum noise and quantum operations," in *Quantum Computation and Quantum Information*, pp. 353–398, Cambridge: Cambridge University Press, 2010.

[152] S. K. Thompson, *Sampling*. Wiley Series in Probability and Statistics, Hoboken, NJ, USA: John Wiley & Sons, Inc., feb 2012.

[153] E. Witten, "A mini-introduction to information theory," *La Rivista del Nuovo Cimento*, vol. 43, pp. 187–227, apr 2020.

[154] K. M. R. Audenaert and J. Eisert, "Continuity bounds on the quantum relative entropy," *Journal of Mathematical Physics*, vol. 46, p. 102104, oct 2005.

[155] M. Fannes, "A continuity property of the entropy density for spin lattice systems," *Communications in Mathematical Physics*, vol. 31, pp. 291–294, dec 1973.

[156] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, jun 2012.

APPENDICES

Appendix A

Additional information for Chapter 3

A.1 Calculating the effect of the n-qubit entangling cycle in Fig. 3.2

Here we show that the entangling cycle has the effect of "propagating" the desired component from the target qubits to the ancillary qubit. Specifically, first note that due to the commutation relations between Pauli operators, the only non-zero components of an N-qubit state after averaging over $\{I, Z\}$ on each qubit are tensor products of I and Z: that is,

$$\rho_t = \frac{1}{2^N} \sum_R s_R R, \;\; R \in \{I, Z\}^{\otimes N}. \tag{A.1}$$

To calculate the effective PTM on q_a, we first note that $[P, Q] = 0$ for all $P, Q \in \{I, Z\}^{\otimes N}$, and that all elements of $\{I, Z\}^{\otimes N}$ are involutory (meaning that they square to the identity). First consider the effect of the entangling gate only. We will consider a particular gate $U = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes T$, where T also belongs to the group of $\{I, Z\}^{\otimes N}$. The PTM on q_a is given by

$$(\Phi_{\mathcal{G}})_{P,Q} = \frac{1}{2^N} \mathrm{Tr}[P \; \mathcal{G}(Q)] \tag{A.2}$$

where the map $\mathcal{G}$ is the effect on q_a by first attaching a state q_t, applying the controlled-P gate, and tracing out q_t. We can write the term $\mathcal{G}(Q)$ as

$$\begin{aligned}\mathcal{G}(Q) &= \mathrm{Tr}_t[U(Q\otimes\rho_t)U^\dagger] \\ &= \mathrm{Tr}_t[\frac{1}{2}(|0\rangle\langle 0|\otimes I + |1\rangle\langle 1|\otimes T)(\sum_R s_R\, Q\otimes R)(|0\rangle\langle 0|\otimes I + |1\rangle\langle 1|\otimes T^\dagger)]\end{aligned} \tag{A.3}$$

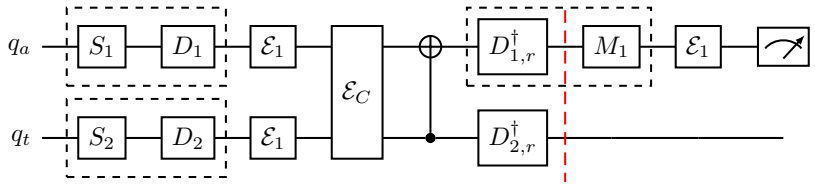

Figure A.1: An expanded version of the two-qubit propagating circuits that are actually carried out. The restoring gates are given by $\mathcal{D}_r^\dagger = \mathcal{C}\mathcal{D}^\dagger\mathcal{C}^\dagger$. Note that SPAM averaging gates are compiled with adjacent dressing gates from randomized compiling (denoted by dashed boxes) and are implemented as one gate, hence there is only one noise channel $\mathcal{E}_1^{\otimes 2}$. The red dashed line indicates the place where we make the comparison (see text).

Consider two separate cases:

1. $Q = I$ or $Q = Z$. Then Q is diagonal and the only non-zero elements are $Q_{00} \coloneqq \langle 0|\, Q\, |0\rangle$ and $Q_{11} \coloneqq \langle 1|\, Q\, |1\rangle$. Therefore we can simplify Eq. (A.3) as

$$\begin{aligned}&\mathrm{Tr}_t[\frac{1}{2^N}\sum_R s_R(|0\rangle\langle 0|\, Q_{00}\otimes R + |1\rangle\langle 1|\, Q_{11}\otimes TRT^\dagger)] \\ &= \mathrm{Tr}_t[\frac{1}{2^N}\sum_R s_R(|0\rangle\langle 0|\, Q_{00}\otimes R + |1\rangle\langle 1|\, Q_{11}\otimes R)] \\ &= \mathrm{Tr}_t[Q\otimes(\frac{1}{2^N}\sum_R s_R R)] \\ &= Q\end{aligned} \tag{A.4}$$

   where we used the fact that T commutes with R for all $T, R \in \{I, Z\}^{\otimes N}$ in the first step, that Q is diagonal in the second step, and that density operators have unit trace in the third step. Therefore,

$$(\Phi_{\mathcal{G}})_{P,Q} = \frac{1}{2}\,\mathrm{Tr}[P\; Q] = \delta_{PQ}. \tag{A.5}$$

2. $Q = X$ or $Q = Y$. Then Q has only off-diagonal elements $Q_{01} \coloneqq \langle 0| Q |1\rangle$ and $Q_{10} \coloneqq \langle 1| Q |0\rangle$. So we can write Eq. (A.3) as

$$\mathrm{Tr}_t[\frac{1}{2^N} \sum_R s_R(|0\rangle\langle 1| Q_{01} \otimes RT^\dagger + |1\rangle\langle 0| Q_{10} \otimes TR)] \tag{A.6}$$

Here we use the fact that all elements of $\{I, Z\}^{\otimes N}$ are both Hermitian and involutory, which leads to the relation

$$RT^\dagger = TR = I^{\otimes N} \text{ iff } T = R. \tag{A.7}$$

Since the only element of $\{I, Z\}^{\otimes N}$ with a nonzero trace is $I^{\otimes N}$, these would be the only remaining terms when the partial trace is performed. Keeping only these terms, we further simplify Eq. (A.6) as

$$\mathrm{Tr}_t[\frac{1}{2^N} \sum_R s_{R=T}(|0\rangle\langle 1| Q_{01} + |1\rangle\langle 0| Q_{10}) \otimes I^{\otimes N}] \tag{A.8}$$

The term in the parenthesis is just Q since $Q = X$ or $Q = Y$ has off-diagonal elements only. Thus we get $\mathcal{G}(Q) = s_T Q$ and

$$(\Phi_\mathcal{G})_{P,Q} = \frac{1}{2} \mathrm{Tr}[P\ s_T Q] = s_T \delta_{PQ}. \tag{A.9}$$

From the above we get the form of $\Phi_\mathcal{G}$ as:

$$\Phi_\mathcal{G} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & s_T & 0 & 0 \\ 0 & 0 & s_T & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{A.10}$$

Finally, the PTM of the a 1-qubit Hadamard gate is

$$\Phi_H = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{A.11}$$

The effect of the full cycle is given by the matrix product

$$\Phi_H \Phi_\mathcal{G} \Phi_H = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & s_T & 0 \\ 0 & 0 & 0 & s_T \end{pmatrix} \tag{A.12}$$

which is identical in form to Eq. (3.9) in the main text.

A.2 Proof of Equation (3.15)

In this section we prove Eq. (3.15) in the main text. First, we define the operator 1-norm $\|A\|_1 \coloneqq \mathrm{Tr}\left[\sqrt{A^\dagger A}\right]$, the induced superoperator norm $\|\mathcal{G}\|_{1\to 1} \coloneqq \max\{\|\mathcal{G}(A)\|_1 : \|A\|_1 \leq 1\}$ and the diamond norm $\|\mathcal{G}\|_\diamond \coloneqq \|\mathcal{G}\otimes\mathcal{I}_d\|_{1\to 1}$ for quantum channels [26, 56]. Note that $\|\rho\|_1 = 1$ for any density matrix ρ. We would like to compare the distance between the final states of the ancillary qubit q_a immediately before the measurement (see Fig. A.1), under two cases: the imaginary case where the propagation cycle is ideal (denoted as ρ_a^{ic}, "ideal cycle"), and the actual case with imperfect gates (denoted as ρ_a).

Below we denote a round of CNOT gate as $\mathcal{C}$, and the channel represented by the round of dressing gate $D_1 \otimes D_2$ as $\mathcal{D}$. Under the assumption of gate-independent error on the dressing gates, we'll denote this (single-qubit) error channel as $\mathcal{E}_1$ so that $\tilde{\mathcal{D}} = \mathcal{E}_1\mathcal{D}$. The error on the CNOT can be a general channel denoted by $\tilde{\mathcal{C}} = \mathcal{C}\mathcal{E}_C$. The dressing gates are randomly sampled from the 2-qubit Pauli channels $\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}$. Recall that ρ_a denotes the state of the ancilla qubit immediately before measurement. Further, let's denote the full system immediately before measurement as ρ_{full}. We then have the following chain of inequalities:

$$\begin{aligned}\left\|\rho_{a,\mathcal{D}} - \rho^{\mathrm{ic}}_{a,\mathcal{D}}\right\|_1 &= \left\|\mathrm{Tr}_2[\mathcal{D}_r^\dagger(\rho_{\mathrm{full}} - \rho^{\mathrm{ic}}_{\mathrm{full}})]\right\|_1 \\ &\leq \left\|\mathcal{D}_r^\dagger(\rho_{\mathrm{full}} - \rho^{\mathrm{ic}}_{\mathrm{full}})\right\|_1 \\ &\leq \left\|\rho_{\mathrm{full}} - \rho^{\mathrm{ic}}_{\mathrm{full}}\right\|_1 \\ &\leq \left\|\tilde{\mathcal{C}}\tilde{\mathcal{D}} - \mathcal{C}\mathcal{D}\right\|_\diamond\end{aligned} \tag{A.13}$$

for each choice of dressing gates $\mathcal{D}$. The first inequality is because partial tracing does not increase the trace distance [151]. The second inequality is because trace distance is non-increasing upon action of any CPTP map. The third inequality is from the definition of the diamond norm.

The diamond norm is related to the process infidelity r_p (Eq. (3.14) in the main text) by [74]

$$2r_p(\tilde{\mathcal{C}}\tilde{\mathcal{D}}, \mathcal{C}\mathcal{D}) \leq \left\|\tilde{\mathcal{C}}\tilde{\mathcal{D}} - \mathcal{C}\mathcal{D}\right\|_\diamond \leq 2d\sqrt{r_p(\tilde{\mathcal{C}}\tilde{\mathcal{D}}, \mathcal{C}\mathcal{D})} \tag{A.14}$$

for d-dimensional channels. From our assumption on the error model, $\tilde{\mathcal{C}}\tilde{\mathcal{D}} = \mathcal{C}\mathcal{E}_C\mathcal{E}_1^{\otimes 2}\mathcal{D}$. Note that for any channel $\mathcal{E}$ and unitary processes $\mathcal{U}$ and $\mathcal{V}$,

$$r_p(\mathcal{E},\mathcal{U}) = r_p(\mathcal{U}^\dagger\mathcal{E},\mathcal{I}),\ \ \|\mathcal{U}\mathcal{E}\mathcal{V}\|_\diamond = \|\mathcal{E}\|_\diamond \tag{A.15}$$

and since both $\|\cdot\|_\diamond$ and r_p are linear functions in their arguments,

$$\begin{aligned} r_{\mathrm{CB}}(\tilde{\mathcal{C}},\mathcal{C}) &= 4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} r_p(\tilde{\mathcal{C}}\tilde{\mathcal{D}},\mathcal{C}\mathcal{D}) \\ &= 4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} r_p(\mathcal{D}^\dagger\mathcal{C}^\dagger\tilde{\mathcal{C}}\tilde{\mathcal{D}},\mathcal{I}) \\ &= 4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} r_p(\mathcal{D}^\dagger\mathcal{E}_C\mathcal{E}_1^{\otimes 2}\mathcal{D},\mathcal{I}) \\ &= r_p(\mathcal{P},\mathcal{I}) \end{aligned} \tag{A.16}$$

and

$$4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} \left\|\tilde{\mathcal{C}}\tilde{\mathcal{D}}-\mathcal{C}\mathcal{D}\right\|_\diamond = \|\mathcal{P}-\mathcal{I}\|_\diamond \tag{A.17}$$

where $\mathcal{P}$ is the twirled error channel mentioned in the main text. After averaging over all Pauli dressing gates (i.e., the circuit is randomly compiled),

$$\begin{aligned} \left\|\rho_{a,\mathrm{RC}}-\rho^{\mathrm{ic}}_{a,\mathrm{RC}}\right\|_1 &\coloneqq 4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} \left\|\rho_{a,\mathcal{D}}-\rho^{\mathrm{ic}}_{a,\mathcal{D}}\right\|_1 \\ &\le 4^{-2} \sum_{\mathcal{D}\in\{\mathcal{I},\mathcal{X},\mathcal{Y},\mathcal{Z}\}^{\otimes 2}} \left\|\tilde{\mathcal{C}}\tilde{\mathcal{D}}-\mathcal{C}\mathcal{D}\right\|_\diamond \\ &= \|\mathcal{P}-\mathcal{I}\|_\diamond \\ &= 2r_p(\mathcal{P},\mathcal{I}) \end{aligned} \tag{A.18}$$

where the last equality is because the lower bound of Eq. (A.14) is saturated for a Pauli noise channel. Combining with Eq. (A.16), we finally have

$$\left\|\rho_{a,\mathrm{RC}}-\rho^{\mathrm{ic}}_{a,\mathrm{RC}}\right\|_1 \le 2r_{\mathrm{CB}}(\tilde{\mathcal{C}},\mathcal{C}). \tag{A.19}$$

On the other hand, $\left\|\rho_{a,\mathrm{RC}}-\rho^{\mathrm{ic}}_{a,\mathrm{RC}}\right\|_1$ is related to β and β^{ic} through the Holevo-Helstrom theorem for distinguishing quantum states [56]. We quote theorem 3.4 in [56] as the following lemma:

Lemma A.2.1. *Let ρ_1, ρ_2 be density operators. Let $\lambda \in [0,1]$. For an arbitrary two-outcome POVM measurement described by elements $\{M_0, M_1\}$, it holds that*

$$\lambda\langle M_0,\rho_0\rangle + (1-\lambda)\langle M_1,\rho_1\rangle \le \frac{1}{2}(1+\|\lambda\rho_0-(1-\lambda)\rho_1\|_1). \tag{A.20}$$

From the definition of β, we can rewrite it as

$$\begin{aligned}\beta_{\mathcal{D}} &= 2\langle M_0, \rho_{a,\mathcal{D}}\rangle - 1 = 1 - 2\langle M_1, \rho_{a,\mathcal{D}}\rangle \\ \beta^{\mathrm{ic}}_{\mathcal{D}} &= 2\langle M_0, \rho^{\mathrm{ic}}_{a,\mathcal{D}}\rangle - 1 = 1 - 2\langle M_1, \rho^{\mathrm{ic}}_{a,\mathcal{D}}\rangle\end{aligned} \tag{A.21}$$

for each particular $\mathcal{D}$. Since the measurement on q_a is unchanged by the propagation cycle, we can apply the above lemma with $\lambda = \frac{1}{2}$ twice: first, using the first definition of $\beta_{\mathcal{D}}$ and the second definition of $\beta^{\mathrm{ic}}_{\mathcal{D}}$,

$$\beta_{\mathcal{D}} - \beta^{\mathrm{ic}}_{\mathcal{D}} = 2(\langle M_0, \rho_{a,\mathcal{D}}\rangle + \langle M_1, \rho^{\mathrm{ic}}_{a,\mathcal{D}}\rangle - 1) \leq \left\|\rho^{\mathrm{ic}}_{a,\mathcal{D}} - \rho_{a,\mathcal{D}}\right\|_1 \tag{A.22}$$

and next, using the second definition of $\beta_{\mathcal{D}}$ and the first definition of $\beta^{\mathrm{ic}}_{\mathcal{D}}$,

$$\beta^{\mathrm{ic}}_{\mathcal{D}} - \beta_{\mathcal{D}} = 2(\langle M_0, \rho^{\mathrm{ic}}_{a,\mathcal{D}}\rangle + \langle M_1, \rho_{a,\mathcal{D}}\rangle - 1) \leq \left\|\rho_{a,\mathcal{D}} - \rho^{\mathrm{ic}}_{a,\mathcal{D}}\right\|_1 \tag{A.23}$$

Combining the above two equations we get:

$$-\left\|\rho^{\mathrm{ic}}_{a,\mathcal{D}} - \rho_{a,\mathcal{D}}\right\|_1 \leq \beta^{\mathrm{ic}}_{\mathcal{D}} - \beta_{\mathcal{D}} \leq \left\|\rho^{\mathrm{ic}}_{a,\mathcal{D}} - \rho_{a,\mathcal{D}}\right\|_1. \tag{A.24}$$

Thus, combining with Eq. (A.19) and averaging over all $\mathcal{D}$'s, we obtain the desired result

$$\left|\beta^{\mathrm{ic}}_{\mathrm{RC}} - \beta_{\mathrm{RC}}\right| \leq 2r_{\mathrm{CB}}(\tilde{\mathcal{C}}, \mathcal{C}). \tag{A.25}$$

The proof can be trivially extended to the case where a multi-qubit propagation cycle, $(\mathcal{H}\otimes\mathcal{I})\mathcal{C}_P(\mathcal{H}\otimes\mathcal{I})$, is used in place of the CNOT gate, hence Eq. (3.15) in the main text.

A.3 Calculating uncertainties in estimated parameters

In this section, we derive the expressions for uncertainties in our experiments. The directly measured quantities in our scheme are the expectation values α and β, as well as the infidelity r_{CB} measured by cycle benchmarking. Here, we will focus on the uncertainties on α and β. The one for r_{CB} is based on the same ideas but involves more technical details, and we refer to the Supplementary Information of the original paper [78] for the exact expressions. In this section only, we'll denote the estimated values of the random variables with an overhead tilde, such that the measured value of α is written as $\tilde{\alpha}$. We will first derive estimators for the expectation value (denoted with $\mathbb{E}$) and variance (denoted with

$\mathbb{V}$) of a desired quantity in the general case, then apply it to the case of α and β. Finally we use standard error propagation to obtain the uncertainty on the upper and lower bounds.

The quantity of interest which we try to estimate can generally be described by the following average value:

$$\lambda = \frac{1}{N}\sum_{i=1}^{N} p_i, \tag{A.26}$$

where the value of N depends on the context. In our case, λ can be α or β, so the p_i's are expectation values of single qubit measurements. There are two things to be noted about estimating this quantity: first, N can be very large in general, so that it is sometimes not possible to exhaustively sample all p_i's. Second, each p_i cannot be measured perfectly because of finite sampling error. From now on, we'll assume that we sample n out of the N elements with or without replacement. For a particular n-element sample s, the value of each sampled element $\tilde{p}_{s_i}$ is a random variable which is denoted with a hat. Furthermore we assume that $\mathbb{E}[\tilde{p}_i] = p_i$, and that the variance $\mathbb{V}[\tilde{p}_i] \coloneqq \sigma_i$ exists and can be estimated using an unbiased estimator $\tilde{\sigma}_i$.

We now derive the estimators of interest. By a simple counting argument (see, for example, section 2.6 of [152]) and the law of total expectation, it is easy to see that

$$\mathbb{E}\left[\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i}\right] = \mathbb{E}_s\left[\frac{1}{n}\sum_{s_i}^{n}\mathbb{E}[\tilde{p}_{s_i}]\right] = \frac{1}{N}\sum_{i} p_i, \tag{A.27}$$

where the standard abuse of notation of denoting the sample with subscript (and summation) is used. This holds true whether we're sampling with or without replacement. Therefore the estimator

$$\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i} \tag{A.28}$$

is an unbiased estimator for the population mean. Next, we can use the law of total variance to compute the variance of this estimator. Note that in the absence of noisy measurements ($\tilde{p}_i = p_i$) the problem reduces to estimating the variance of the sample mean, which has a well-known formula (see section 2.5 and 2.6 of [152]) when the sampling is done without replacement:

$$\mathbb{V}_s\left[\frac{1}{n}\sum_{s_i}^{n} p_{s_i}\right] = \left(1 - \frac{n}{N}\right)\frac{\sigma^2}{n} \tag{A.29}$$

or with replacement:

$$\mathbb{V}_s\left[\frac{1}{n}\sum_{s_i}^{n} p_{s_i}\right] = \frac{N-1}{N}\frac{\sigma^2}{n} \tag{A.30}$$

where σ^2 is the population variance of p_i:

$$\sigma^2 = \frac{1}{N-1}\sum_{i=1}^{N}(p_i - \bar{p})^2,\ \bar{p} = \sum_{i=1}^{N}\frac{p_i}{N}. \tag{A.31}$$

By the law of total variance we can extend to account for noisy measurements: for sampling without replacement we have

$$\mathbb{V}\left[\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i}\right] \tag{A.32}$$

$$=\mathbb{V}_s\mathbb{E}\left[\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i}\middle| s\right] + \mathbb{E}_s\mathbb{V}\left[\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i}\middle| s\right] \tag{A.33}$$

$$=\mathbb{V}_s\left[\frac{1}{n}\sum_{s_i}^{n}p_{s_i}\middle| s\right] + \mathbb{E}_s\left[\frac{1}{n^2}\sum_{s_i}^{n}\sigma_i^2\right] \tag{A.34}$$

$$=\left(1-\frac{n}{N}\right)\frac{\sigma^2}{n} + \frac{1}{n}\sum_{i=1}^{N}\frac{\sigma_i^2}{n} \tag{A.35}$$

$$=\left(1-\frac{n}{N}\right)\frac{\sigma^2}{n} + \frac{1}{nN}\sum_{i=1}^{N}\sigma_i^2, \tag{A.36}$$

and similarly for with replacement,

$$\mathbb{V}\left[\frac{1}{n}\sum_{s_i}^{n}\tilde{p}_{s_i}\right] = \left(1-\frac{1}{N}\right)\frac{\sigma^2}{n} + \frac{1}{nN}\sum_{i=1}^{N}\sigma_i^2. \tag{A.37}$$

An important implication from the above expression is that, for situations where each σ_i is small, or where N is very large, the above expression depends mostly on the spread of the quantities over the set of values (i.e., σ^2) and only very weakly on N. A practical example that aligns with our protocol is where N grows exponentially in the number of qubits, and σ_i decreases as the square root of measurement "shots". This ensures that randomly sampling from a large population is scalable in practice.

We then need an estimator for σ^2 and $\frac{1}{N}\sum_{i=1}^{N}\sigma_i^2$. For the second quantity it's simply $\frac{1}{n}\sum_{i=1}^{n}\tilde{\sigma}_{s_i}^2$, which are variances of each element in the chosen sample. For the first quantity, it can be shown that the sample variance corrected by the average of $\sigma_{s_i}^2$'s gives an unbiased estimator for σ^2: i.e.,

$$\mathbb{E}\left[\tilde{s}^2 - \frac{1}{n}\sum_{i}^{n}\tilde{\sigma}_{s_i}^2\right] = \sigma^2. \tag{A.38}$$

where $\tilde{s}^2$ is the variance for the chosen sample. To see this, note that σ^2 is equal to

$$\sigma^2 = \frac{1}{N-1}(p_i^2 - \frac{2}{N}\sum_j^N p_i p_j + \frac{1}{N^2}\sum_{jj'} p_j p_{j'}) \tag{A.39}$$

$$= \frac{1}{N-1}(\sum_i^N p_i^2 - \frac{1}{N}\sum_{ij} p_i p_j) \tag{A.40}$$

$$= \frac{1}{N-1}(\sum_i^N p_i^2 - \frac{1}{N}(\sum_i^N p_i^2 + 2\sum_{i\neq j}^N p_i p_j)) \tag{A.41}$$

$$= \frac{1}{N}\sum_i^N p_i^2 - \frac{2}{N(N-1)}\sum_{i\neq j}^N p_i p_j. \tag{A.42}$$

Meanwhile, the expectation value of $\tilde{s}^2$ is

$$\mathbb{E}\left[\tilde{s}^2\right] = \mathbb{E}\left[\frac{\sum_{i=1}^n (\tilde{p}_{s_i} - \frac{1}{n}\sum_{j=1}^n \tilde{p}_{s_j})^2}{n-1}\right] \tag{A.43}$$

$$= \frac{1}{n-1}\mathbb{E}\left[\sum_i^n \tilde{p}_{s_i}^2 - \frac{1}{n}\sum_{i,j}^n \tilde{p}_{s_i}\tilde{p}_{s_j}\right]. \tag{A.44}$$

By a counting argument, the first term evaluates to

$$\mathbb{E}\left[\sum_i^n \tilde{p}_{s_i}^2\right] = \frac{n}{N}\sum_i^N \mathbb{E}\left[\tilde{p}_i^2\right] = \frac{n}{N}\sum_i^N p_i^2 + \sigma_i^2. \tag{A.45}$$

The second term evaluates to

$$\mathbb{E}\left[\frac{1}{n}\sum_{i,j}^n \tilde{p}_{s_i}\tilde{p}_{s_j}\right] = \mathbb{E}\left[\frac{1}{n}(\sum_{i=j}^n + \sum_{i\neq j}^n)\tilde{p}_{s_i}\tilde{p}_{s_j}\right] \tag{A.46}$$

$$= \frac{1}{N}\sum_i^N \mathbb{E}\left[\tilde{p}_i^2\right] + \frac{2(n-1)}{N(N-1)}\sum_{i\neq j}^N \mathbb{E}\left[\tilde{p}_i\tilde{p}_j\right] \tag{A.47}$$

$$= \frac{1}{N}\sum_i^N (p_i^2 + \sigma_i^2) + \frac{2(n-1)}{N(N-1)}\sum_{i\neq j}^N p_i p_j \tag{A.48}$$

where the last equality is given by independence of p_i's.

Combining Eqs. (A.44), (A.45) and (A.48) and compare with Eq. (A.42), one sees that the corrected estimator (Eq. (A.38)) is unbiased. We can then use this in Eq. (A.36) and simplify to get our unbiased estimator for the variance for sampling without replacement:

$$(\frac{1}{n}-\frac{1}{N})(\tilde{s}^2-\frac{1}{n}\sum_i^n\tilde{\sigma}_{s_i}^2)+\frac{1}{n^2}\sum_i^n\tilde{\sigma}_{s_i}^2 \tag{A.49}$$

$$=(\frac{1}{n}-\frac{1}{N})\tilde{s}^2-(\frac{1}{n^2}-\frac{1}{nN})\sum_i^n\tilde{\sigma}_{s_i}^2+\frac{1}{n^2}\sum_i^n\tilde{\sigma}_{s_i}^2 \tag{A.50}$$

$$=(\frac{1}{n}-\frac{1}{N})\tilde{s}^2+\frac{1}{nN}\sum_i^n\tilde{\sigma}_{s_i}^2, \tag{A.51}$$

and sampling with replacement:

$$(\frac{1}{n}-\frac{1}{nN})(\tilde{s}^2-\frac{1}{n}\sum_i^n\tilde{\sigma}_{s_i}^2)+\frac{1}{n^2}\sum_i^n\tilde{\sigma}_{s_i}^2 \tag{A.52}$$

$$=(\frac{1}{n}-\frac{1}{nN})\tilde{s}^2+\frac{1}{n^2N}\sum_i^n\tilde{\sigma}_{s_i}^2. \tag{A.53}$$

Eq. (A.28) and Eq. (A.51) (or Eq. (A.53)) allow us to write down mean and variance estimators for any quantity in the form of Eq. (A.26). We then use standard (linear-approximated) error propagation to estimate uncertainties on the parameters of interest, i.e., upper and lower bound on the error rates ϵ_{SP} and ϵ_{M} from Table 1 in the main text. Each bound is an independent estimate and is a function of the four parameters: α_t, α_a, β_t, and $r_{t,a}$. The uncertainty on individual parameters are independent of each other, so its covariance matrix Σ^p is diagonal (where p stands for "parameter"). The first order approximation to the covariance matrix of the bounds is given by $\Sigma^b = J\Sigma^p J^T$ where J is the Jacobian. We then take the diagonal elements of Σ^b to be the uncertainties of the bounds.

Finally we mention how each $\sigma_{s_i}^2$ is estimated for our experiment. Since $\tilde{p}_i$ equals to a binomial variable divided by the sample size k, it has mean p_i and variance $\frac{p_i(1-p_i)}{k}$. Using again the relation $\mathbb{E}[\tilde{p}^2] = p^2 + \sigma^2$, it can be verified that an unbiased estimator for the variance is $\frac{\tilde{p}_i(1-\tilde{p}_i)}{k-1}$. Because $\tilde{\alpha}_i = 2\tilde{p}_i - 1$, $\sigma_{\alpha,i}^2 = 4\sigma_{p,i}^2$. We can express the estimator in terms of $\tilde{\alpha}_i$ as

$$\tilde{\sigma}_{\alpha,i}^2 = \frac{(1+\tilde{\alpha}_i)(1-\tilde{\alpha}_i)}{k-1}. \tag{A.54}$$

The estimator for β is identical, except changing the $\tilde{\alpha}_i$ to $\tilde{\beta}_i$ in the above expression.

Appendix B

Additional information for Chapter 4

B.1 Calculating expected number of runs before a successful one

Here we calculate the expected number of failed MBAC-k runs before having a successful one. The scenario is that one will continue running the experiment, until a successful round of MBAC-k occurs. Assume that one would like to reduce $\delta_{\mathrm{SP},t}$ by a factor of r (e.g., $r = 100$ or $r = 1000$ can be set by the experimentalist), by using MBAC-k. According to Eq. (14) in the main text, the final SP error is exponentially suppressed in terms of successful cooling rounds.

We again think about MBAC-k as an $(k-1)$-step protocol, where each step corresponds to measuring an ancilla and tracing it out after the measurement. The probability of measuring 0 on the i-th step is

$$p_{0,i} = (1 - \delta_{\mathrm{SP},t}[i] - \delta_{\mathrm{SP,a}} + 2\delta_{\mathrm{SP},t}[i]\delta_{\mathrm{SP,a}})(1 - \delta_{\mathrm{M},a}). \tag{B.1}$$

The measurements are independent for each round, so the probability of measuring all 0's is given by the product of $p_{0,i}$ from each round. Since $p_{0,i}$ improves as i increases, we can lower bound all of them by $p_{0,1}$, i.e., the probability of getting 0 in the first run of MBAC-2.

Thus, we see that the expected total number of runs is upper bounded by a simpler case, where n independent Bernoulli trials are conducted in series, each having success probability $p_{0,1}$. The probability of having all trials successful, which we will call the success probability, is $(p_{0,1})^{k-1}$. Note that for a Bernoulli trial that has success probability

p, the expected number of tests n can be calculated to be

$$\begin{aligned}\mathbb{E}(n) =& \sum_i n_i \cdot p_i(n_i) \\ =& 1 \cdot p + 2p(1-p) + 3p(1-p)^2 + \dots \\ =& p \sum_{n=0}^{\infty} (n+1)(1-p)^n \\ =& \frac{p}{p^2} = \frac{1}{p}\end{aligned} \tag{B.2}$$

as expected. Thus, if we denote the expected number of tests *in reality* by N, we have

$$N \leq \frac{1}{(p_{0,1})^{k-1}}, \tag{B.3}$$

where $p_{0,1}$ is defined in Eq. (B.1). Next, recall that Eq. (4.21) gives an upper bound on the required number of ancillary qubits, $k-1$. Because $p_{0,1} \in (0,1)$, $p_{0,1}^x$ decreases with x for $x \geq 0$, so that $(1/p_{0,1})^x$ increases with x for $x \geq 0$. Combining this with Eq. (B.3), we have

$$N \leq \left(\frac{1}{p_{0,1}}\right)^{k-1} \leq \left(\frac{1}{p_{0,1}}\right)^{\frac{\log(r)}{-\log\left(2\delta_{\text{SPAM},a}\right)}} := A^{\frac{\log(r)}{\log(B)}} \tag{B.4}$$

where

$$A = p_{0,1}, \;\; B = 2\delta_{\text{SPAM},a}. \tag{B.5}$$

We now define the expression on the RHS of Eq. (B.4) as an upper bound on N, i.e.,

$$N_{\text{upper}}(r) := A^{\frac{\log(r)}{\log(B)}}. \tag{B.6}$$

Note that this upper bound is only a function of r and some measurable quantities that are specific to the machine. Taking the log on both sides of Eq. (B.6) gives

$$\log(N_{\text{upper}}(r)) = \frac{\log(r)}{\log(B)} \log(A) = \log(r) \frac{\log(A)}{\log(B)} \tag{B.7}$$

and taking the exponential on both sides again gives

$$N_{\text{upper}}(r) = (r)^{\frac{\log(A)}{\log(B)}} \tag{B.8}$$

which is polynomial in r, as claimed.

Appendix C

Additional information for Chapter 5

C.1 Derivation for Landauer's Principle in the unitary scenario, Eq. (5.52)

This is a more concise version of the proof in [109]. We first prove the case where β is finite. Due to the invariance of the von Neumann entropy under unitary transformations, $S(\rho'_{t,b}) = S(\rho_{t,b}) = S(\rho_t) + S(\rho_b)$. Expand the mutual information term as

$$\begin{aligned} I(t' : b') &= S(\rho'_t) + S(\rho'_b) - S(\rho'_{t,b}) \\ &= S(\rho'_t) + S(\rho'_b) - S(\rho_{t,b}) \\ &= S(\rho'_t) + S(\rho'_b) - S(\rho_t) - S(\rho_b) \\ &= \Delta S_t + S(\rho'_b) - S(\rho_b). \end{aligned} \tag{C.1}$$

Also

$$\begin{aligned} D(\rho'_b||\rho_b) &= \mathrm{Tr}[\rho'_b \log \rho'_b] - \mathrm{Tr}[\rho'_b \log \rho_b] \\ &= -S(\rho'_b) - \mathrm{Tr}[\rho'_b \log \rho_b] \end{aligned} \tag{C.2}$$

Therefore

$$\begin{aligned} -\Delta S_t + I(t' : b') + D(\rho'_b||\rho_b) &= -S(\rho_b) - \mathrm{Tr}[\rho'_b \log \rho_b] \\ &= \mathrm{Tr}[(\rho_b - \rho'_b) \log \rho_b]. \end{aligned} \tag{C.3}$$

The bath is initially in a thermal state $\rho_b = e^{-\beta_b H_b} / \mathrm{Tr}\left[e^{-\beta_b H_b}\right]$, so

$$\log \rho_b = -(\beta_b H_b + \mathbb{1} \log \mathrm{Tr}\left[e^{-\beta_b H_b}\right]). \tag{C.4}$$

Plug this back into one equation above gives

$$-\Delta S_t + I(t':b') + D(\rho'_b||\rho_b) = \beta_b \operatorname{Tr}[(\rho_b - \rho'_b)H_b] \tag{C.5}$$

since $\rho_b - \rho'_b$ is traceless. We can identify this to be the energy change in the bath, ΔE_b.

For the case where $\beta = \infty$, $\rho_b = P_g/\dim(P_g)$ is the normalized projector onto the ground state space of H_b. Thus $\Delta E_b \geq 0$ and there are two possibilities:

1. $\Delta E_b = 0$, then ρ'_b is also supported only in the ground state space of H_b. The key is to realize that $\log \rho_b = \log P_g - \log \dim(P_g)$, and that $\log P_g$ is the 0 operator on the support of P_g. Therefore

$$\begin{aligned} -\Delta S_t + I(t':b') + D(\rho'_b||\rho_b) &= -S(\rho_b) - \operatorname{Tr}[\rho'_b \log \rho_b] \\ &= \operatorname{Tr}[(\rho_b - \rho'_b)\log \rho_b] \\ &= \operatorname{Tr}[(\rho_b - \rho'_b)\log P_g] - \operatorname{Tr}[(\rho_b - \rho'_b)\log \dim(P_g))] \\ &= 0 \end{aligned} \tag{C.6}$$

   where the first term is 0 because both ρ_b andρ'_b are supported only on the ground state space of H_b, and the second term is 0 because $\rho_b - \rho'_b$ is traceless. Thus, both sides of Eq. (C.5) are 0.

2. $\Delta E_b > 0$, then ρ'_b is supported outside the ground state of H_b. Thus $D(\rho'_b||\rho_b) = \infty$, and both sides of Eq. (C.5) are ∞.

C.2 Derivation for Landauer's Principle in the general CPTP scenario, Eq. (5.54)

The derivation can be most easily seen by expressing the difference in relative entropy using its definition,

$$\begin{aligned} &D(\rho_t||\rho_t^{th}) - D(\rho'_t||\rho_t^{th}) \\ &= -S(\rho'_t) - \operatorname{Tr}\left[\rho'_t \log\left(\rho_t^{th}\right)\right] + S(\rho_t) + \operatorname{Tr}\left[\rho_t \log\left(\rho_t^{th}\right)\right] \\ &= (S(\rho_t) - S(\rho'_t)) + \operatorname{Tr}\left[(\rho_t - \rho'_t)\log\left(\rho_t^{th}\right)\right] \\ &= -\Delta S_t + \operatorname{Tr}\left[(\rho_t - \rho'_t)\log\left(\rho_t^{th}\right)\right] \end{aligned} \tag{C.7}$$

By definition, $\rho_t^{th} = e^{-\beta H_t} / \operatorname{Tr}\left[e^{-\beta H_t}\right]$, so replacing this within the log() gives

$$\begin{aligned} \ldots &= -\Delta S_t + \operatorname{Tr}\left[(\rho_t - \rho_t')(-\beta H_t - \mathbb{1} \log \operatorname{Tr}\left[e^{-\beta H_t}\right])\right] \\ &= -\Delta S_t - \operatorname{Tr}[(\rho_t - \rho_t')\beta H_t] - \log \operatorname{Tr}\left[e^{-\beta H_t}\right] \operatorname{Tr}[(\rho_t - \rho_t')] \end{aligned} \tag{C.8}$$

The last term is 0 since all physical processes are trace-preserving. Defining $\Delta E_t \coloneqq \operatorname{Tr}[(\rho_t' - \rho_t)H_t]$ and $\Delta S_t \coloneqq S(\rho_t') - S(\rho_t)$, we arrive at

$$-\beta \Delta E_t = -\Delta S_t + (D(\rho_t || \rho_t^{th}) - D(\rho_t' || \rho_t^{th})). \tag{C.9}$$

C.3 Proof of Theorem 5.5.1

Consider the unitary evolution $\mathcal{U}(\rho_{ab}) = \rho_{ab}'$ which results in a negative entropy change on subsystem a, $-\Delta S_a > 0$. We can rearrange Eq. (5.52) and divide by $-\Delta S_a$ to obtain an expression for the Landauer Ratio for a round of unitary operation:

$$R_L = 1 + \frac{I(a' : b') + D(\rho_b' || \rho_b)}{-\Delta S_a}. \tag{C.10}$$

First note that the mutual information term $I(a' : b')$ can be rewritten as (see, e.g., Eq. (3.55) in [153])

$$I(a' : b') = D(\rho_{ab}' || \rho_a' \otimes \rho_b') \tag{C.11}$$

A quadratic lower bound on the relative entropy is [154]

$$D(\rho || \sigma) \geq 2(T(\rho, \sigma))^2 \tag{C.12}$$

On the other hand, an upper bound on the entropy change is given by the Fannes' inequality [155, 156]:

$$|S(\rho) - S(\sigma)| \leq 2T(\rho, \sigma) \log\left(\frac{d}{2T(\rho, \sigma)}\right) \tag{C.13}$$

for $T(\rho_a', \rho_a) \leq 1/2e$. Directly plugging in the above into Eq. (C.10) gives the desired bound on R_L.

C.4 Detailed evolution of HBACs using transition matrices

C.4.1 Calculation of PPA-HBAC

A round of PPA3 consists of applying the 3-qubit entropy compression U_{PPA3} (given in Eq. (5.15)), and thermalizing both machine qubits to a given bath temperature. The thermalization can be effectively calculated as tracing out qubit 2 and 3, and replacing with two thermal qubits at the bath temperature. The transition matrices representing these two operations are

$$G_{tr23} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} \vec{1}_{1\times 4} & \vec{1}_{0\times 4} \\ \vec{0}_{1\times 4} & \vec{1}_{1\times 4} \end{pmatrix}, \tag{C.14}$$

$$G_{\otimes \rho_b^{\otimes 2}} = \begin{pmatrix} p_b^2 & 0 \\ p_b(1-p_b) & 0 \\ (1-p_b)p_b & 0 \\ (1-p_b)^2 & 0 \\ 0 & p_b^2 \\ 0 & p_b(1-p_b) \\ 0 & (1-p_b)p_b \\ 0 & (1-p_b)^2 \end{pmatrix} := \begin{pmatrix} \vec{s}_{b,2} & \vec{0}_{4\times 1} \\ \vec{0}_{4\times 1} & \vec{s}_{b,2} \end{pmatrix}, \tag{C.15}$$

where we defined

$$\vec{s}_{b,2} = \begin{pmatrix} p_b^2 \\ p_b(1-p_b) \\ (1-p_b)p_b \\ (1-p_b)^2 \end{pmatrix} \tag{C.16}$$

to be the vector representing the state of two independent qubits both at the bath temperature (p_b is the probability of being in 0 for a thermal bath qubit at bath temperature). The full transition matrix is given by

$$\begin{aligned} G_{\text{PPA3}}^{\text{relax}} &= G_{\otimes \rho_b^{\otimes 2}} G_{tr23} \\ &= \begin{pmatrix} \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} \\ \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} \end{pmatrix}. \end{aligned} \tag{C.17}$$

Combining with the transition matrix for the compression step,

$$G_{\mathrm{PPA3}}^{\mathrm{comp}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.18}$$

the transition matrix for a complete round of PPA3 is

$$\begin{aligned} G_{\mathrm{PPA3}} &= G_{\mathrm{PPA3}}^{\mathrm{relax}} G_{\mathrm{PPA3}}^{\mathrm{comp}} \\ &= \begin{pmatrix} \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{0}_{4\times 1} & \vec{s}_{b,2} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} \\ \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{0}_{4\times 1} & \vec{s}_{b,2} & \vec{0}_{4\times 1} & \vec{s}_{b,2} & \vec{s}_{b,2} & \vec{s}_{b,2} \end{pmatrix}. \end{aligned} \tag{C.19}$$

Thus, the post-cooling state after repeated rounds of PPA3 can be computed by simply raising G_{PPA3} to the desired power, and multiplying it to the initial state vector.

Focusing on the target qubit alone, the evolution in a given round is first introducing two qubits at the bath temperature, apply the compression, then tracing out the two ancillary qubits. The reduced transition matrix on the target qubit is given by

$$\begin{aligned} G_{\mathrm{PPA3},t} &= G_{tr23} G_{\mathrm{PPA3}}^{\mathrm{comp}} G_{\otimes \rho_b^{\otimes 2}} \\ &= \begin{pmatrix} 2p_b - p_b^2 & p_b^2 \\ (1-p_b)^2 & 1 - p_b^2 \end{pmatrix}. \end{aligned} \tag{C.20}$$

This provides a simpler view on how the target evolves during cooling. Diagonalizing this matrix gives

$$G_{\mathrm{PPA3},t} = \begin{pmatrix} z & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & y \end{pmatrix} \frac{1}{1+z} \begin{pmatrix} 1 & 1 \\ -1 & z \end{pmatrix} \tag{C.21}$$

where $y \coloneqq 2p_b(1-p_b)$ and $z \coloneqq p_b^2/(1-p_b)^2$. We can raise it to the n-th power by

$$(G_{\mathrm{PPA3},t})^n = \begin{pmatrix} z & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & y^n \end{pmatrix} \frac{1}{1+z} \begin{pmatrix} 1 & 1 \\ -1 & z \end{pmatrix} \tag{C.22}$$

Therefore, the state of the target after n rounds starting from the initial state $|\rho_{\mathrm{PPA3},t}(0)\rangle\rangle = (p_t, 1-p_t)^T$ is

$$\begin{aligned} |\rho_{\mathrm{PPA3},t}(n)\rangle\rangle &= (G_{\mathrm{PPA3},t})^n |\rho_{\mathrm{PPA3},t}(0)\rangle\rangle \\ &= \frac{1}{1+z} \begin{pmatrix} z - y^n(z - (z+1)p_t) \\ 1 + y^n(z - (z+1)p_t) \end{pmatrix}. \end{aligned} \tag{C.23}$$

for $n \geq 1$. Since $y \in [0, 1/2]$, it is easy to obtain the asymptotic state as

$$|\rho_{\mathrm{PPA3},t}(\infty)\rangle\!\rangle = \lim_{n\to\infty} |\rho_{\mathrm{PPA3},t}(n)\rangle\!\rangle = \frac{1}{1+z}\begin{pmatrix} z \\ 1 \end{pmatrix}. \tag{C.24}$$

We now switch back to using the full transition matrix to compute thermodynamic quantities for calculating cooling efficiencies. The work input for the n-th round is the energy increase for the full system during the compression step,

$$\begin{aligned} w_{\mathrm{PPA3}}(n) &= \langle\!\langle H_{\mathrm{PPA3}}|(G^{\mathrm{comp}}_{\mathrm{PPA3}} - \mathbb{1})|\rho_{\mathrm{PPA3}}(n-1)\rangle\!\rangle \\ &= \langle\!\langle H_{\mathrm{PPA3}}|(G^{\mathrm{comp}}_{\mathrm{PPA3}} - \mathbb{1})(G_{\mathrm{PPA3}})^{n-1}|\rho_{\mathrm{PPA3}}(0)\rangle\!\rangle. \end{aligned} \tag{C.25}$$

The vectorized Hamiltonian is

$$\langle\!\langle H_{\mathrm{PPA3}}| = \sum_{i=1}^{3} \langle\!\langle h_{[i]}| = E\begin{pmatrix} 0 & 1 & 1 & 2 & 1 & 2 & 2 & 3 \end{pmatrix} \tag{C.26}$$

where we assumed the same energy gap E for each qubit. The vectorized initial state is

$$|\rho_{\mathrm{PPA3}}(0)\rangle\!\rangle = |\bigotimes_{i=1}^{3} \rho_i\rangle\!\rangle = \begin{pmatrix} p_b^3 \\ p_b^2(1-p_b) \\ p_b^2(1-p_b) \\ p_b(1-p_b)^2 \\ p_b^2(1-p_b) \\ p_b(1-p_b)^2 \\ p_b(1-p_b)^2 \\ (1-p_b)^3 \end{pmatrix} \tag{C.27}$$

where we assumed that all qubits start from the same (bath) temperature for simplicity. Directly calculating Eq. (C.25) using a computational software gives

$$w_{\mathrm{PPA3}}(n) = E(p_b - \frac{1}{2})y^n, \ n \geq 1 \tag{C.28}$$

Meanwhile, the energy decrease for the target during the n-th round can be computed from Eq. (C.23):

$$\begin{aligned} -\Delta e_t(n) &= \langle\!\langle H_1|(|\rho_{\mathrm{PPA3},t}(n-1)\rangle\!\rangle - |\rho_{\mathrm{PPA3},t}(n)\rangle\!\rangle) \\ &= E(p_b - \frac{1}{2})y^n, \ n \geq 1 \end{aligned} \tag{C.29}$$

It can be seen from here that both the per-round and cumulative CoP are simply

$$k_{\mathrm{PPA3}}(n) = K_{\mathrm{PPA3}}(n) = 1. \tag{C.30}$$

C.4.2 Calculation of NOE-HBAC

A round of NOE-HBAC consists of first applying the CMS operation on the second qubit, then applying the Γ_2 relaxation. The transition matrices representing these two operations are

$$G_{\text{CMS}} = \frac{1}{2}\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}, \tag{C.31}$$

$$G_{\Gamma_2} = \begin{pmatrix} p_2 & 0 & 0 & p_2 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1-p_2 & 0 & 0 & 1-p_2 \end{pmatrix}, \tag{C.32}$$

where p_2 is defined in Eq. (5.9). A round of NOE-HBAC is represented by the transition matrix

$$\begin{aligned} G_{\text{NOE}} &= G_{\Gamma_2} G_{\text{CMS}} \\ &= \frac{1}{2}\begin{pmatrix} p_2 & p_2 & p_2 & p_2 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1-p_2 & 1-p_2 & 1-p_2 & 1-p_2 \end{pmatrix}. \end{aligned} \tag{C.33}$$

The transition matrix representing the evolution of the target qubit is

$$\begin{aligned} G_{\text{NOE},t} &= G_{tr2} G_{\Gamma_2} G_{\otimes \rho_{\text{CMS}}} \\ &= \frac{1}{2}\begin{pmatrix} p_2+1 & p_2 \\ 1-p_2 & 2-p_2 \end{pmatrix}. \end{aligned} \tag{C.34}$$

Let $x \coloneqq \frac{p_2}{1-p_2}$; then $G_{\text{NOE},t}$ can be diagonalized as

$$G_{\text{NOE},t} = \begin{pmatrix} x & -1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{2} \end{pmatrix}\frac{1}{1+x}\begin{pmatrix} 1 & 1 \\ -1 & x \end{pmatrix} \tag{C.35}$$

Raising to the nth power:

$$G^n_{\text{NOE},t} = \frac{1}{1+x}\begin{pmatrix} x+(\frac{1}{2})^n & x(1-(\frac{1}{2})^n) \\ 1-(\frac{1}{2})^n & 1+(\frac{1}{2})^n x \end{pmatrix}. \tag{C.36}$$

The state of the target after n rounds starting from the initial state $|\rho_{\mathrm{NOE},t}(0)\rangle\!\rangle = (p_t, 1-p_t)^T$ is

$$\begin{aligned} |\rho_{\mathrm{NOE},t}(n)\rangle\!\rangle &= (G_{\mathrm{NOE},t})^n |\rho_{\mathrm{NOE},t}(0)\rangle\!\rangle \\ &= \frac{1}{1+x}\begin{pmatrix} x + 2^{-n}(p_t - x + p_t x) \\ 1 - 2^{-n}(p_t - x + p_t x) \end{pmatrix} \end{aligned} \tag{C.37}$$

for $n \geq 1$. The asymptotic state is

$$|\rho_{\mathrm{NOE},t}(\infty)\rangle\!\rangle = \lim_{n\to\infty} |\rho_{\mathrm{NOE},t}(n)\rangle\!\rangle = \frac{1}{1+x}\begin{pmatrix} x \\ 1 \end{pmatrix}. \tag{C.38}$$

The work input for the n-th round is the energy increase for the full system during the CMS step,

$$\begin{aligned} w_{\mathrm{NOE}}(n) &= \langle\!\langle H_{\mathrm{NOE}} | (G_{\mathrm{CMS}} - \mathbb{1}) | \rho_{\mathrm{NOE}}(n-1) \rangle\!\rangle \\ &= \langle\!\langle H_{\mathrm{NOE}} | (G_{\mathrm{CMS}} - \mathbb{1})(G_{\mathrm{NOE}})^{n-1} | \rho_{\mathrm{NOE}}(0) \rangle\!\rangle. \end{aligned} \tag{C.39}$$

The vectorized Hamiltonian is

$$\langle\!\langle H_{\mathrm{NOE}}| = \sum_{i=1}^{2} \langle\!\langle h_{[i]}| = E\begin{pmatrix} 0 & 1 & 1 & 2 \end{pmatrix} \tag{C.40}$$

where we assumed the same energy gap E for each qubit. The vectorized initial state is

$$|\rho_{\mathrm{NOE}}(0)\rangle\!\rangle = |\bigotimes_{i=1}^{2} \rho_i\rangle\!\rangle = \begin{pmatrix} p_b^2 \\ p_b(1-p_b) \\ p_b(1-p_b) \\ (1-p_b)^2 \end{pmatrix} \tag{C.41}$$

where we assumed that all qubits start from the same (bath) temperature for simplicity. Directly calculating Eq. (C.39) using a computational software gives

$$w_{\mathrm{NOE}}(n) = \begin{cases} E(p_2 - p_b)2^{1-n}, & n > 1 \\ E(p_b - \frac{1}{2}), & n = 1 \end{cases} \tag{C.42}$$

The cumulative work up to round n is

$$W_{\mathrm{NOE}}(n) = \sum_{i=1}^{n} w_{\mathrm{NOE}}(i) = E((p_2 - \frac{1}{2}) - 2^{1-n}(p_2 - p_b)). \tag{C.43}$$

Meanwhile, the energy decrease for the target during the n-th round can be computed from Eq. (C.37):

$$\begin{aligned} -\Delta e_t(n) &= \langle\langle H_1 | (|\rho_{\mathrm{NOE},t}(n-1)\rangle\rangle - |\rho_{\mathrm{NOE},t}(n)\rangle\rangle) \\ &= E(p_2 - p_b)2^{-n} \end{aligned} \tag{C.44}$$

for all $n \geq 1$. The total energy decrease up to round n is

$$-\Delta E_t(n) = \sum_{i=1}^{n} -\Delta e_t(i) = (p_2 - p_b)(1 - 2^{-n})E \tag{C.45}$$

The COP for the n-th round is

$$k_{\mathrm{NOE}}(n) = \frac{-\Delta e_t(n)}{w_{\mathrm{NOE}}(n)} = \begin{cases} \frac{1}{2}, & n > 1 \\ \frac{p_2 - p_b}{2p_b - 1}, & n = 1 \end{cases} \tag{C.46}$$

and the cumulative CoP up to round n is

$$K_{\mathrm{NOE}}(n) = \frac{-\Delta E_t(n)}{W_{\mathrm{NOE}}(n)} = \frac{(p_2 - p_b)(1 - 2^{-n})}{(p_2 - \frac{1}{2}) - 2^{1-n}(p_2 - p_b)}. \tag{C.47}$$

C.4.3 Calculation of SRG2-HBAC

A round of SRG2-HBAC consists of first applying the σ_x gate on the second qubit, then applying the Γ_2 relaxation, and finally thermalizing the second qubit. The transition matrices are

$$G_{\sigma_x,[2]} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{C.48}$$

$$G_{\Gamma_1,[2]} = \begin{pmatrix} p_b & p_b & 0 & 0 \\ 1 - p_b & 1 - p_b & 0 & 0 \\ 0 & 0 & p_b & p_b \\ 0 & 0 & 1 - p_b & 1 - p_b \end{pmatrix}, \tag{C.49}$$

and G_{Γ_2} as defined as above. A round of SRG2-HBAC is represented by the transition matrix

$$\begin{aligned} G_{\mathrm{SRG2}} &= G_{\Gamma_1,[2]} G_{\Gamma_2} G_{\sigma_x,[2]} \\ &= \begin{pmatrix} p_b & p_b p_2 & p_b p_2 & 0 \\ 1 - p_b & (1 - p_b)p_2 & (1 - p_b)p_2 & 0 \\ 0 & p_b(1 - p_2) & p_b(1 - p_2) & p_b \\ 0 & (1 - p_b)(1 - p_2) & (1 - p_b)(1 - p_2) & 1 - p_b \end{pmatrix}. \end{aligned} \tag{C.50}$$

The transition matrix representing the evolution of the target qubit is

$$\begin{aligned} G_{\mathrm{SRG2},t} &= G_{tr2} G_{\Gamma_2} G_{\sigma_x,[2]} G_{\otimes\rho_b} \\ &= \begin{pmatrix} p_b + (1-p_b)p_2 & p_b p_2 \\ (1-p_b)(1-p_2) & 1-p_b p_2 \end{pmatrix}. \end{aligned} \tag{C.51}$$

Let $w \coloneqq p_2 p_b$, $v = 1 - p_2 - p_b + p_2 p_b$, $u = 1 - v - w$, then $G_{\mathrm{SRG2},t}$ can be diagonalized as

$$G_{\mathrm{SRG2},t} = \begin{pmatrix} \frac{w}{v} & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & u \end{pmatrix} \frac{v}{w+v} \begin{pmatrix} 1 & 1 \\ -1 & \frac{w}{v} \end{pmatrix} \tag{C.52}$$

Raising to the power of n:

$$G^n_{\mathrm{SRG2},t} = \frac{1}{w+v} \begin{pmatrix} w + v u^n & w - w u^n \\ v - v u^n & v + w u^n \end{pmatrix}. \tag{C.53}$$

The state of the target after n rounds starting from the initial state $|\rho_{\mathrm{SRG2},t}(0)\rangle\!\rangle = (p_t, 1-p_t)^T$ is

$$\begin{aligned} |\rho_{\mathrm{SRG2},t}(n)\rangle\!\rangle &= (G_{\mathrm{SRG2},t})^n |\rho_{\mathrm{SRG2},t}(0)\rangle\!\rangle \\ &= \frac{1}{w+v} \begin{pmatrix} w + u^n(-w + p_t(v+w)) \\ v - u^n(-w + p_t(v+w)) \end{pmatrix} \end{aligned} \tag{C.54}$$

for $n \geq 1$. The asymptotic state is

$$|\rho_{\mathrm{SRG2},t}(\infty)\rangle\!\rangle = \lim_{n\to\infty} |\rho_{\mathrm{SRG2},t}(n)\rangle\!\rangle = \frac{1}{w+v} \begin{pmatrix} w \\ v \end{pmatrix} \tag{C.55}$$

since $|u| < 1$.

The work input for the n-th round is the energy increase for the full system during the σ_x step,

$$\begin{aligned} W_{\mathrm{SRG2}}(n) &= \langle\!\langle H_{\mathrm{SRG2}} | (G_{\sigma_x,[2]} - \mathbb{1}) | \rho_{\mathrm{SRG2}}(n-1) \rangle\!\rangle \\ &= \langle\!\langle H_{\mathrm{SRG2}} | (G_{\sigma_x,[2]} - \mathbb{1}) (G_{\mathrm{SRG2}})^{n-1} | \rho_{\mathrm{SRG2}}(0) \rangle\!\rangle. \end{aligned} \tag{C.56}$$

The vectorized Hamiltonian is again

$$\langle\!\langle H_{\mathrm{SRG2}}| = \sum_{i=1}^{2} \langle\!\langle h_{[i]}| = E \begin{pmatrix} 0 & 1 & 1 & 2 \end{pmatrix} \tag{C.57}$$

where we assumed the same energy gap E for each qubit. The vectorized initial state is

$$|\rho_{\mathrm{SRG2}}(0)\rangle\!\rangle = |\bigotimes_{i=1}^{2} \rho_i\rangle\!\rangle = \begin{pmatrix} p_b^2 \\ p_b(1-p_b) \\ p_b(1-p_b) \\ (1-p_b)^2 \end{pmatrix} \tag{C.58}$$

where we assumed that all qubits start from the same (bath) temperature for simplicity. Directly calculating Eq. (C.56) using a computational software gives

$$w_{\mathrm{SRG2}}(n) = E(2p_b - 1) \tag{C.59}$$

for all $n \geq 1$. The cumulative work up to round n is

$$W_{\mathrm{SRG2}}(n) = nE(2p_b - 1) \tag{C.60}$$

Meanwhile, the energy decrease for the target during the n-th round can be computed from Eq. (C.54):

$$\begin{aligned} -\Delta e_{\mathrm{SRG2},t}(n) &= \langle\langle H_1 | (|\rho_{\mathrm{SRG2},t}(n-1)\rangle\rangle - |\rho_{\mathrm{SRG2},t}(n)\rangle\rangle) \\ &= Ep_b(1-p_b)(2p_2-1)u^{n-1} \end{aligned} \tag{C.61}$$

for all $n \geq 1$. The total energy decrease in the target is

$$-\Delta E_{\mathrm{SRG2},t}(n) = Ep_b(1-p_b)(2p_2-1)\frac{1-u^n}{1-u} \tag{C.62}$$

The COP for the n-th round is

$$k_{\mathrm{SRG2}}(n) = \frac{-\Delta e_{\mathrm{SRG2},t}(n)}{w_{\mathrm{SRG2}}(n)} = \frac{p_b(1-p_b)(2p_2-1)}{2p_b-1}u^{n-1} \tag{C.63}$$

for all rounds $n \geq 1$. The cumulative CoP up to round n is

$$\begin{aligned} K_{\mathrm{SRG2}}(n) &= \frac{-\Delta E_{\mathrm{SRG2},t}(n)}{W_{\mathrm{SRG2}}(n)} \\ &= \frac{p_b(1-p_b)(2p_2-1)}{2p_b-1}\frac{1-u^n}{n(1-u)}. \end{aligned} \tag{C.64}$$

C.4.4 Calculation of xHBAC

A round of 1-qubit xHBAC consists of first applying the σ_x gate on the target, then applying the β-swap operation. The transition matrix for β-swap is

$$G_\beta = \begin{pmatrix} 1 - Ee^{-\beta_b} & 1 \\ e^{-E\beta_b} & 0 \end{pmatrix} \tag{C.65}$$

where E is the energy gap of the qubit and β_b is the bath inverse temperature. The total G-matrix is given by

$$G_{\text{xHBAC1}} = G_\beta G_{\sigma_x} = \begin{pmatrix} 1 & 1 - Ee^{-\beta_b} \\ 0 & e^{-E\beta_b} \end{pmatrix}. \tag{C.66}$$

To diagonalize G_{xHBAC1}, let $s = e^{-E\beta_b}$, we obtain

$$G_{\text{xHBAC1}} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & s \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \tag{C.67}$$

The n-th power of G_{xHBAC1} is given by

$$G^n_{\text{xHBAC1}} = \begin{pmatrix} 1 & 1 - s^n \\ 0 & s^n \end{pmatrix} \tag{C.68}$$

The state of the target after n rounds starting from the initial state $|\rho_{\text{xHBAC1},t}(0)\rangle\!\rangle = (p_t, 1-p_t)^T$ is

$$\begin{aligned} |\rho_{\text{xHBAC1},t}(n)\rangle\!\rangle &= (G_{\text{xHBAC1},t})^n |\rho_{\text{xHBAC1},t}(0)\rangle\!\rangle \\ &= \begin{pmatrix} 1 - s^n(1-p_b) \\ s^n(1-p_b) \end{pmatrix} \end{aligned} \tag{C.69}$$

for $n \geq 1$. The asymptotic state is

$$|\rho_{\text{xHBAC1},t}(\infty)\rangle\!\rangle = \lim_{n\to\infty} |\rho_{\text{xHBAC1},t}(n)\rangle\!\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \tag{C.70}$$

i.e., the pure state $|0\rangle$.

The work input for the n-th round is the energy increase for the full system during the σ_x step,

$$\begin{aligned} w_{\text{xHBAC1}}(n) &= \langle\!\langle H_{\text{xHBAC1}} | (G_{\sigma_x} - \mathbb{1}) | \rho_{\text{xHBAC1}}(n-1) \rangle\!\rangle \\ &= \langle\!\langle H_{\text{xHBAC1}} | (G_{\sigma_x} - \mathbb{1})(G_{\text{xHBAC1}})^{n-1} | \rho_{\text{xHBAC1}}(0) \rangle\!\rangle \\ &= E(1 - 2s^{n-1}(1-p_b)). \end{aligned} \tag{C.71}$$

The total work input is

$$W_{\text{xHBAC1}}(n) = E\left(n - 2(1-p_b)\frac{1-s^n}{1-s}\right). \tag{C.72}$$

The energy decrease for the target during the n-th round is

$$\begin{aligned} -\Delta e_{\text{xHBAC1},t}(n) &= \langle\!\langle H_1 | (|\rho_{\text{xHBAC1}}(n-1)\rangle\!\rangle - |\rho_{\text{xHBAC1}}(n)\rangle\!\rangle) \\ &= E(1-p_b)(1-s)s^{n-1} \end{aligned} \tag{C.73}$$

for all $n \geq 1$. The total energy decrease up to round n is

$$-\Delta E_{\mathrm{xHBAC1},t}(n) = E(1-p_b)(1-s^n). \tag{C.74}$$

The COP for the n-th round is

$$k_{\mathrm{xHBAC1}}(n) = \frac{-\Delta e_t(n)}{w_{\mathrm{xHBAC1}}(n)} = \frac{(1-p_b)(1-s)}{1-2s^{n-1}(1-p_b)} s^{n-1} \tag{C.75}$$

for all rounds $n \geq 1$. The cumulative COP up to round n is

$$K_{\mathrm{xHBAC1}}(n) = \frac{(1-p_b)(1-s)(1-s^n)}{n(1-s)-2(1-p_b)(1-s^n)} \tag{C.76}$$

C.5 Proof of the equilibrium state of the improved PPA

In the improved PPA, the goal is to minimize the work input during the compression step, while keeping the same final temperature of the target qubit as in PPA. This is achieved by placing the diagonal elements in the $|0\rangle_t$ subspace of the target qubit in the lowest energy levels in decreasing order, and similarly for the $|1\rangle_t$ subspace.

In every cycle, the system will start from an initial tensor product state where the first qubit is in the state with diagonal elements $(1-\delta_t, \delta_t)$, and all other qubits have diagonal elements $(1-\delta_b, \delta_b)$, where $\delta_b = \langle 1| \rho_b^{th} |1\rangle$. The full state has 2^n diagonal elements, which can be grouped by the Hamming weight of their corresponding state. For example, the population corresponding to the state $|00101\rangle$, which has weight 2, is given by $(1-\delta_t)(1-\delta_b)^2\delta_b^2$. In general, for a total weight of i, the diagonal elements are

$$\begin{aligned} \binom{n-1}{i} &: \ \delta_0(i) = (1-\delta_t)(1-\delta_b)^{n-i-1}\delta_b^i, \ i = 0 \ldots (n-1) \\ \binom{n-1}{i-1} &: \ \delta_1(i) = \delta_t(1-\delta_b)^{n-i}\delta_b^{i-1}, \ i = 1 \ldots n \end{aligned} \tag{C.77}$$

where the binomial coefficients denote how many identical terms there are in total, $\delta_0(i)$ corresponds to what's originally in the $|0\rangle_t$ subspace and the $\delta_1(i)$ in the $|1\rangle_t$ subspace of the target qubit.

Consider the situation where the target starts from a state with a δ_t satisfying the following relation:

$$\frac{1-\delta_t^*}{\delta_t^*} = \left(\frac{1-\delta_b}{\delta_b}\right)^{n-1}. \tag{C.78}$$

We can now rewrite $\delta_0(i)$ and $\delta_1(i)$ as

$$\begin{aligned}\{\delta_0(i), \delta_1(i)\} &= \left\{\left(\frac{1-\delta_b}{\delta_b}\right)^{n-1}(1-\delta_b)^{n-i-1}\delta_b^i\delta_t^*, \delta_t^*(1-\delta_b)^{n-i}\delta_b^{i-1}\right\} \\ &= \delta_t^*(1-\delta_b)^{n-1}\left\{\left(\frac{1-\delta_b}{\delta_b}\right)^{n-i-1}, \left(\frac{1-\delta_b}{\delta_b}\right)^{-i+1}\right\}\end{aligned} \tag{C.79}$$

where we have factored out the common term $\delta_t^*(1-\delta_b)^{n-1}$. The fraction $(1-\delta_b)/\delta_b$ is the ratio of $|0\rangle$ over $|1\rangle$ populations for a qubit at the bath temperature and is always ≥ 1 for all $\delta_b \in (0, 1/2]$. Therefore, each term increases with an increasing exponent. Specifically, since both exponents decrease with i, the *smallest* term from the $|0\rangle_t$ subspace is achieved at $i = n-1$, and the *largest* term from the $|1\rangle_t$ subspace is achieved at $i = 1$. Since

$$n - (n-1) - 1 = -(1) - 1 = 0, \tag{C.80}$$

these two terms are in fact equal to each other. This implies that all the largest 2^{n-1} terms are already in the $|0\rangle_t$ subspace, so the temperature of the target qubit will remain the same after the compression step. And since the target's initial temperature is the only variable during the compression, we conclude that this state is a stable fixed point for the E-PPA protocol.

Finally, we see that the equilibrium population satisfying Eq. (C.78) has an inverse temperature

$$\begin{aligned}\beta E &= \log\left(\left(\frac{1-\delta_b}{\delta_b}\right)^{n-1}\right) = (n-1)\log\left(\frac{1-\delta_b}{\delta_b}\right) \\ &= (n-1)\beta_b E\end{aligned} \tag{C.81}$$

where β_b is the bath's inverse temperature. Therefore, the equilibrium temperature is $T_b/(n-1)$.

Appendix D

Additional information for Chapter 6

D.1 Proof that the Drazin-inverse of a TP map is TP

Here we provide the proof for Theorem 6.3.5.

Proof. Let the Jordan block of eigenvalue λ in J be J_λ, where J is defined in Eq. (6.15).

The inverse of a k by k Jordan block J_λ ($\lambda \neq 0$) of $\mathbf{v}(\mathcal{N})$ is

$$\begin{aligned} J_\lambda^{-1} &= (\lambda I + N)^{-1} = \lambda^{-1}(I - \lambda^{-1}N + \cdots \lambda^{-(k-1)}N^{(k-1)}) \\ &= \lambda^{-1}\left(\sum_{i=0}^{k-1}(-\lambda^{-1}N)^i\right) =: J'_\lambda \end{aligned}$$

where N is the k by k nilpotent matrix, λ is the eigenvalue.

By the construction (Eq. (6.16)), we have

$$\mathbf{v}(\mathcal{N}^+)Q = QJ', \tag{D.1}$$

where Q contains the eigenvectors and generalized eigenvectors of $\mathbf{v}(\mathcal{N})$. For the particular block that we concern, the corresponding eigenvector and generalized eigenvector is

$$Q = \begin{pmatrix} \cdots & v^{g_0} & v^{g_1} & \cdots & v^{g_{k-1}} & \cdots \end{pmatrix}$$

Let $e_i, i \in \{0, \cdots, k-1\}$ be the standard basis vectors for this block. Acting e_j on both sides of Eq. (D.1), the left hand side is

$$\mathbf{v}(\mathcal{N}^+)Qe_j = \mathbf{v}(\mathcal{N}^+)v^{g_j},$$

and the right hand side is

$$QJ'e_j = Q\lambda^{-1}\left(\sum_{i=0}^{k-1}(-\lambda)^{-i}N^i e_j\right).$$

It is easy to show that $N^i e_j = e_{j-i}$ for $j \geq i$, and $N^i e_j = 0 \cdot e_j$ for $j < i$. Therefore

$$QJ'e_j = Q\lambda^{-1}\left(\sum_{i=0}^{j}(-\lambda)^i e_{j-i}\right) = \lambda^{-1}\left(\sum_{i=0}^{j}(-\lambda)^i v^{g_{j-i}}\right).$$

Thus

$$\mathbf{v}(\mathcal{N}^+)v^{g_j} = \lambda^{-1}\left(\sum_{i=0}^{j}(-\lambda)^i v^{g_{j-i}}\right). \tag{D.2}$$

For λ not equals to 1 and 0, taking trace on both sides of Eq. (D.2),

$$\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)v^{g_j}\right] = \mathrm{sTr}\left[\lambda^{-1}\left(\sum_{i=0}^{j}(-\lambda)^i v^{g_{j-i}}\right)\right].$$

From Lemma 6.3.3, we know that $\mathrm{sTr}\left[v^{g_j}\right] = 0$, so the right hand side is also 0. That is, $\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)v^{g_j}\right] = \mathrm{sTr}\left[v^{g_j}\right]$ holds for every $j \in \{0, \cdots, k-1\}$.

When $\lambda = 1$, according to Lemma 6.3.4, we have the same results except for $j = k-1$. Now we check $j = k-1$ case for $\lambda = 1$,

$$\mathbf{v}(\mathcal{N}^+)v^{g_{k-1}} = \left(\sum_{i=0}^{k-1}(-1)^i v^{g_{k-1-i}}\right).$$

Taking trace on both sides,

$$\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)v^{g_{k-1}}\right] = \mathrm{sTr}\left[\left(\sum_{i=0}^{k-1}(-1)^i v^{g_{k-1-i}}\right)\right]$$

From Lemma 6.3.4, we know that all the eigenvector and generalized eigenvectors have trace zero except for the $(k-1)$th one. We have

$$\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)v^{g_{k-1}}\right] = \mathrm{sTr}\left[v^{g_{k-1}}\right]$$

Finally, when $\lambda = 0$, $J'_\lambda = 0_k$, where 0_k is the k by k zero matrix. Thus, $\mathbf{v}(\mathcal{N}^+)v^{g_j} = 0 \cdot e_j$. From Lemma 6.3.3, the trace of both sides are zero.

Now we have proved that the trace of all columns (eigenvectors and generalized eigenvectors) in Q are unchanged under the action of $\mathbf{v}(\mathcal{N}^+)$.

Since Q is invertible, any $\mathbf{v}(\rho)$ can be expanded by columns $v_i^{\lambda_j}$ in Q. And we have

$$\mathrm{sTr}\left[\mathbf{v}(\mathcal{N}^+)(\sum_{ij} a_{ij} v_i^{\lambda_j})\right] = \mathrm{sTr}\left[\sum_{ij} a_{ij}\mathbf{v}(\mathcal{N}^+)(v_i^{\lambda_j})\right] = \mathrm{sTr}\left[\sum_{ij} a_{ij} v_i^{\lambda_j}\right],$$

Hence, $\mathcal{N}^+$ is trace preserving.

□

D.2 The Effect of Imperfect Knowledge about Noise Channels on Fidelity

From the main text, we know that

$$\rho_{\mathrm{EM}} = \mathcal{U}_{n\cdots 1} \circ \tilde{\mathcal{R}}(\rho_{\mathrm{out}}^{\mathrm{exp}}),$$
$$\rho_{\mathrm{out}}^{\mathrm{ideal}} = \mathcal{U}_{n\cdots 1} \circ \mathcal{R}(\rho_{\mathrm{out}}^{\mathrm{exp}}),$$

where $\mathcal{U}_{n\cdots 1} \coloneqq \mathcal{U}_n \circ \cdots \circ \mathcal{U}_1$ is the ideal set of circuits (See Fig. 6.5).

Imperfect knowledge about $\mathcal{N}_i$ leads to imperfect inverse $\tilde{N}_i^{-1}$. Let $\tilde{N}_i^{-1} = \mathcal{N}_i^{-1} + \Delta\mathcal{N}_i^{-1}$, where $\mathcal{N}_i^{-1}$ is the perfect inverse of $\mathcal{N}_i$.

Let $\Delta\rho_{\mathrm{EM}} \coloneqq \rho_{\mathrm{EM}} - \rho_{\mathrm{out}}^{\mathrm{ideal}}$, then

$$\Delta\rho_{\mathrm{EM}} = \mathcal{U}_{n\cdots 1} \circ \left[\tilde{\mathcal{R}} - \mathcal{R}\right](\rho_{\mathrm{out}}^{\mathrm{exp}}) = \mathcal{U}_{n\cdots 1} \circ \Delta\mathcal{N}(\rho_{\mathrm{out}}^{\mathrm{exp}}). \tag{D.3}$$

If we only consider the first order approximation $\Delta\mathcal{N}^{(1)}$ of $\Delta\mathcal{N}$, the first order correction term $\Delta\rho_{\mathrm{EM}}^{(1)}$ would be

$$\Delta\rho_{\mathrm{EM}}^{(1)} = \mathcal{U}_{n\cdots 1} \circ \Delta\mathcal{N}^{(1)} \circ \mathcal{U}_{\mathrm{exp}}(\rho_{\mathrm{in}}) \tag{D.4}$$
$$= \mathcal{U}_{n\cdots 1} \circ \left(\sum_{i=1}^{n} \mathcal{U}_1^\dagger \circ \tilde{\mathcal{N}}_1^{-1} \cdots \circ \mathcal{U}_i^\dagger \circ \Delta\tilde{\mathcal{N}}_i^{-1} \circ \cdots \mathcal{U}_n^\dagger \circ \tilde{\mathcal{N}}_n^{-1}\right) \circ \mathcal{U}_{\mathrm{exp}}(\rho_{\mathrm{in}}), \tag{D.5}$$

where $\mathcal{U}_{\text{exp}} \coloneqq \mathcal{N}_n \circ \mathcal{U}_n \circ \cdots \circ \mathcal{N}_1 \circ \mathcal{U}_1$ is the actual experimental operator.

Then $\rho_{\text{EM}} \approx \rho_{\text{out}}^{\text{ideal}} + \Delta\rho_{\text{EM}}^{(1)}$. The first order approximation of $F(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}})$, the fidelity between ρ_{EM} and $\rho_{\text{out}}^{\text{ideal}}$, is

$$F(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}}) \approx F^{(1)}(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}}) \coloneqq F(\rho_{\text{out}}^{\text{ideal}} + \Delta\rho_{\text{EM}}^{(1)}, \rho_{\text{out}}^{\text{ideal}}) = \|\sqrt{\rho_{\text{out}}^{\text{ideal}} + \Delta\rho_{\text{EM}}^{(1)}}\sqrt{\rho_{\text{out}}^{\text{ideal}}}\|_{\text{tr}}.$$

By the Fuchs–van de Graaf inequalities,

$$[1 - D(\Delta\rho_{\text{EM}}^{(1)})]^2 \leq F^{(1)}(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}}) \leq 1 - D^2(\Delta\rho_{\text{EM}}^{(1)}), \tag{D.6}$$

where $D(\cdot) \coloneqq \frac{1}{2}\|\cdot\|_{\text{tr}}$ is the trace distance, and $\|\cdot\|_{\text{tr}}$ is the trace norm. It is also known that $\|A\|_F \leq \|A\|_{\text{tr}} \leq \sqrt{r}\|A\|_F$, where $\|A\|_F$ is the Frobenius norm which equals to $\|\mathbf{v}(A)\|$. The norm $\|\cdot\|$ is the 2-norm. Since

$$\|\Delta\rho_{\text{EM}}^{(1)}\|_F = \|\mathbf{v}(\Delta\rho_{\text{EM}}^{(1)})\| = \|\mathbf{v}(\mathcal{U}_{n\cdots 1})\mathbf{v}(\Delta\mathcal{N}^{(1)})\mathbf{v}(\mathcal{U}_{\text{exp}})\mathbf{v}(\rho_{\text{in}})\| = \|\mathbf{v}(\mathcal{U}_{n\cdots 1})\mathbf{v}(\Delta\mathcal{N}^{(1)})\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\|,$$

therefore we can bound $\|\Delta\rho_{\text{EM}}^{(1)}\|_F$ by

$$l_U \cdot \|\mathbf{v}(\Delta\mathcal{N}^{(1)})\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\| \leq \|\Delta\rho_{\text{EM}}^{(1)}\|_F \leq \|\mathbf{v}(\mathcal{U}_{n\cdots 1})\| \cdot \|\mathbf{v}(\Delta\mathcal{N}^{(1)})\| \cdot \|\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\| \tag{D.7}$$

where $l_U \coloneqq \inf_{\|x\|=1} \|\mathbf{v}(\mathcal{U}_{n\cdots 1})x\|$ is the lower Lipschitz constant of the superoperator of the ideal circuits. Notice that, on the right hand side of Eq. (D.7), $\|\mathbf{v}(\mathcal{U}_{n\cdots 1})\|$ and $\|\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\|$ are known for a given experiment. Denote $(\|\mathbf{v}(\mathcal{U}_{n\cdots 1})\| \cdot \|\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\|)$ as C_{exp}. From Eq. (D.6), we know the fidelity between the mitigated state and the ideal state is bounded by

$$\left(1 - \frac{1}{2}\sqrt{d}C_{\text{exp}}\|\mathbf{v}(\Delta\mathcal{N}^{(1)})\|\right)^2 \leq F^{(1)}(\rho_{\text{EM}}, \rho_{\text{out}}^{\text{ideal}}) \leq 1 - \frac{1}{4}\left(l_U \cdot \|\mathbf{v}(\Delta\mathcal{N}^{(1)})\mathbf{v}(\rho_{\text{out}}^{\text{exp}})\|\right)^2.$$

In addition, the norm of $\mathbf{v}(\Delta\mathcal{N}^{(1)})$ satisfies that

$$\left\|\mathbf{v}(\Delta\mathcal{N}^{(1)})\right\| \leq \prod_{k=1}^{n}\left\|\mathbf{v}(\mathcal{U}_k^\dagger)\right\| \cdot \sum_{i=1}^{n}\left[\left\|\mathbf{v}(\Delta\mathcal{N}_i^{-1})\right\| \prod_{\substack{j\in\{1,\cdots,n\}\\ j\neq i}}\left\|\mathbf{v}(\tilde{\mathcal{N}}_j^{-1})\right\|\right], \tag{D.8}$$

where $\mathcal{U}_k^\dagger$ and $\tilde{\mathcal{N}}_j^{-1}$ are known for a given EM tasks. The error on each inverse $\Delta\mathcal{N}_i^{-1}$ can be exposed at the lower bound of the fidelity. And by counting the sampling cost of getting $\Delta\mathcal{N}^{-1}$, one can bound the fidelity from the sampling cost.

D.3 A Sufficient Condition on Improving Expectation Values

The goal of error mitigation on the expectation value of an observables A is

$$|\operatorname{Tr}(\rho_{\mathrm{EM}}A) - \operatorname{Tr}\big(\rho_{\mathrm{out}}^{\mathrm{ideal}}A\big)| \leq |\operatorname{Tr}\big(\rho_{\mathrm{out}}^{\mathrm{ideal}}A\big) - \operatorname{Tr}(\rho_{\mathrm{out}}^{\mathrm{exp}}A)|. \tag{D.9}$$

The left hand side of Eq. (D.9) is

$$\begin{aligned}\left|\operatorname{Tr}\left[(\rho_{\mathrm{EM}} - \rho_{\mathrm{out}}^{\mathrm{ideal}})A\right]\right| &= |\operatorname{Tr}(\Delta\rho_{\mathrm{EM}}A)| = |\operatorname{Tr}(\mathcal{U}_{n\cdots 1}\circ\Delta\mathcal{N}(\rho_{\mathrm{out}}^{\mathrm{exp}})\cdot A)| \\ &= \left|\left\langle \mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger}), \mathbf{v}(\Delta\mathcal{N})\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\rangle\right|.\end{aligned} \tag{D.10}$$

The right hand side of Eq. (D.9) equals to

$$\left|\operatorname{Tr}\left[(\rho_{\mathrm{out}}^{\mathrm{ideal}} - \rho_{\mathrm{out}}^{\mathrm{exp}})A\right]\right| = |\operatorname{Tr}[(\mathcal{U}_{n\cdots 1}\circ\mathcal{R} - \mathcal{I})(\rho_{\mathrm{out}}^{\mathrm{exp}})\cdot A]| \tag{D.11}$$

$$\begin{aligned}&= \left|\operatorname{Tr}\left[\mathcal{U}_{n\cdots 1}\circ[\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger}](\rho_{\mathrm{out}}^{\mathrm{exp}})\cdot A\right]\right| \\ &= \left|\left\langle \mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger}), \mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\rangle\right|\end{aligned} \tag{D.12}$$

It is difficult to draw conclusions directly from Eq. (D.10) and Eq. (D.11) since $\mathbf{v}(\Delta\mathcal{N})$ and $\mathbf{v}(\mathcal{U}^{\dagger}\mathcal{N}_{1\cdots n}^{-1} - \mathcal{U}_{1\cdots n}^{\dagger})$ can be arbitrary. However,

$$\left|\left\langle \mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger}), \mathbf{v}(\Delta\mathcal{N})\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\rangle\right| \leq \left\|\mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger})\right\| \left\|\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\| \left\|\mathbf{v}(\Delta\mathcal{N})\right\|,$$

$$\left|\left\langle \mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger}), \mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\rangle\right| \geq \left\|\mathbf{v}(\mathcal{U}_{n\cdots 1}^{\dagger})\mathbf{v}(A^{\dagger})\right\| \left\|\mathbf{v}(\rho_{\mathrm{out}}^{\mathrm{exp}})\right\| \inf_{\|x\|=1}\left\|\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})x\right\|.$$

Therefore, if

$$\left\|\mathbf{v}(\Delta\mathcal{N})\right\| \leq \inf_{\|x\|=1}\left\|\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})x\right\|,$$

then Eq. (D.9) is guaranteed. This means that the EM process will improve the expectation value for any observable A and any desired circuit $\mathcal{U}_{n\cdots 1}$ when the above is satisfied. This is a stringent requirement, since if $\mathbf{v}(\mathcal{R} - \mathcal{U}_{1\cdots n}^{\dagger})$ has a nontrivial null space, then the right hand side equals to 0. This implies $\tilde{\mathcal{N}}_i = \mathcal{N}_i$ for $\forall i \in \{1, \cdots, n\}$, i.e., the estimation on all noise channels $\tilde{\mathcal{N}}_i$ must be perfect.

D.4 Examples of Noise Channel Mismatch

The Kraus representation of $\mathcal{N}$ and $\mathcal{D}$ are

$$\mathcal{N}: \left\{\sqrt{p_1}I, \sqrt{p_2}X, \sqrt{p_3}Y, \sqrt{(1-p_1-p_2-p_3)}Z\right\};$$

$$\mathcal{D}: \left\{\sqrt{1-\frac{3\lambda}{4}}I, \sqrt{\frac{\lambda}{4}}X, \sqrt{\frac{\lambda}{4}}Y, \sqrt{\frac{\lambda}{4}}Z\right\}.$$

For a given set of $\{p_1, p_2, p_3\}$, what is the optimal λ to minimize $\|\mathcal{N}-\mathcal{D}\|$ of a chosen norm $\|\cdot\|$? One approach is to write down a matrix representation of $\mathcal{N}$ and $\mathcal{D}$, then solve λ by minimizing $\|N-D\|$ for a particular choice of norm. For different representations and/or norms, the optimization outcome could be different. The optimal λ will bound the distance $\|\mathcal{N}-\mathcal{D}\|$ from below for any possible experimental implementation for this particular norm $\|\cdot\|$.

As mentioned in the main text, the two vectors, $\vec{n} \coloneqq (\sqrt{p_1}, \sqrt{p_2}, \sqrt{p_3}, \sqrt{(1-p_1-p_2-p_3)})$ and $\vec{d} \coloneqq \left(\sqrt{1-\frac{3\lambda}{4}}, \sqrt{\frac{\lambda}{4}}, \sqrt{\frac{\lambda}{4}}, \sqrt{\frac{\lambda}{4}}\right)$, are also representations for $\mathcal{N}$ and $\mathcal{D}$ respectively. Since $\vec{n}$ and $\vec{d}$ are normalized, minimizing the distance between $\mathcal{N}$ and $\mathcal{D}$ is equivalent to maximizing $\vec{n}\cdot\vec{d}$. i.e.

$$\max_{\lambda\in[0,1]} \left\{\sqrt{p_1(1-\frac{3\lambda}{4})} + [\sqrt{p_2}+\sqrt{p_3}+\sqrt{(1-p_1-p_2-p_3)}]\sqrt{\frac{\lambda}{4}}\right\}.$$

This can be solved by taking the derivative of the expression, and setting it to zero. The result is

$$\lambda_{max} = \frac{[\sqrt{p_2}+\sqrt{p_3}+\sqrt{(1-p_1-p_2-p_3)}]^2 p_1}{\frac{9}{4}p_1^2 + \frac{3}{4}p_1[\sqrt{p_2}+\sqrt{p_3}+\sqrt{(1-p_1-p_2-p_3)}]^2}, \text{ or } \lambda=1, \text{ or } \lambda=0. \tag{D.13}$$

The superoperators of $\mathcal{N}$ and $\mathcal{D}$ are

$$\begin{aligned}\mathbf{v}(\mathcal{N}) = p_1 \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} + p_2 \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} + p_3 \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \\ + (1 - p_1 - p_2 - p_3) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},\end{aligned} \tag{D.14}$$

$$\begin{aligned}\mathbf{v}(\mathcal{D}) = (1 - \frac{3\lambda}{4}) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} + \frac{\lambda}{4} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} + \frac{\lambda}{4} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \\ + \frac{\lambda}{4} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.\end{aligned} \tag{D.15}$$

Even with the optimal λ in Eq. (D.13), when p_2, p_3 and $1 - p_2 - p_3$ are not equal to each other, the distance between $\mathcal{N}$ and $\mathcal{D}$ is not zero.

The following are two examples of different sets of $\{p_i\}$.

1. When $p_1 = p_3 = 0$ and $p_2 = 1$, the optimal $\lambda_{\max}$ is 1. Therefore

$$\mathbf{v}(\mathcal{N}) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \ \mathbf{v}(\mathcal{D}) = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

   In this case, the estimated $\mathcal{D}$ is non-invertible while $\mathcal{N}$ is invertible. Any generalized inverse $\mathcal{D}$ will definitely worsen the outcomes.

2. When $p_1 = \frac{1}{2}$ and $p_2 = p_3 = 0$, we have $\lambda_{\max} = \frac{1}{3}$ according to Eq. (D.13).

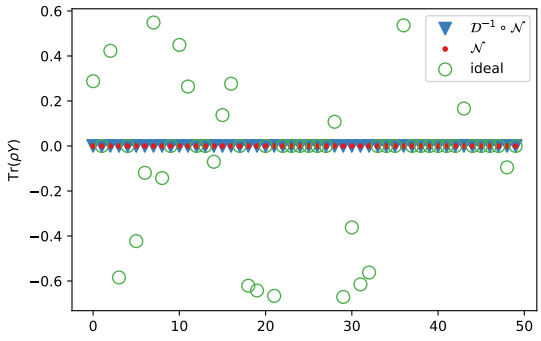

(a) Expectation values of Pauli Y for 50 randomly generated states.

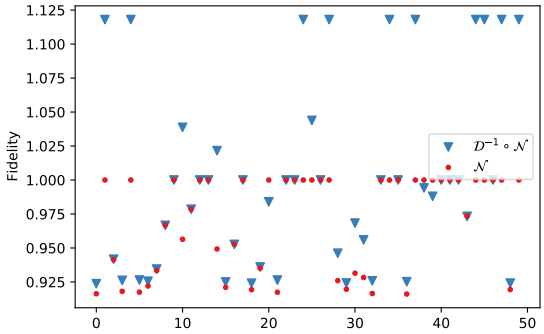

(b) Fidelities of 50 randomly generated states.

Figure D.1: Effects of applying the mismatched noise channel, $\mathcal{D}^{-1}$, to 50 randomly generated noisy output states. The x-axis is a dummy label for the tested states. Because the channel $\mathcal{D}^{-1} \circ \mathcal{N}$ is not physical (not CP), the "mitigated" outputs $\mathcal{D}^{-1} \circ \mathcal{N}(\rho)$ are not valid quantum states. In this case the fidelity is no longer a good metric for distinguishing two "states".

$$\mathbf{v}(\mathcal{N}) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{v}(\mathcal{D}) = \frac{1}{6} \begin{pmatrix} 5 & 0 & 0 & 1 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 1 & 0 & 0 & 5 \end{pmatrix}$$

The inverse of $\mathcal{D}$ is

$$\mathbf{v}(\mathcal{D}^{-1}) = \frac{1}{4} \begin{pmatrix} 5 & 0 & 0 & -1 \\ 0 & 6 & 0 & 0 \\ 0 & 0 & 6 & 0 \\ -1 & 0 & 0 & 5 \end{pmatrix}.$$

Therefore,

$$\mathbf{v}(\mathcal{N}^{\text{re}}) \coloneqq \mathbf{v}(\mathcal{D}^{-1})\mathbf{v}(\mathcal{N}) = \frac{1}{4} \begin{pmatrix} 5 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 5 \end{pmatrix}$$

This resulting channel $\mathcal{N}^{\text{re}}$ has eigenvalues $\{\frac{3}{2}, 1, 0, 0\}$, which will worsen the outcome. In Fig. D.1, we tested 50 randomly generated quantum state ρ for this example.

Fig. D.1a shows the expectation value of Pauli Y for these 50 states. Since the expectation value $\mathrm{Tr}(Y\rho)$ is erased by the noise channel $\mathcal{N}$, $\mathcal{D}^{-1}$ has no effect on improving $\mathrm{Tr}(Y\rho)$. In comparison, for $\mathrm{Tr}(Z\rho)$ in Fig. 6.7, the channel $\mathcal{D}^{-1}$ has made the outcome worse. Fig. D.1b shows the fidelities $F(\mathcal{N}(\rho),\rho)$ and $F(\mathcal{N}^{\mathrm{re}}(\rho),\rho)$. Since $\mathcal{D}^{-1}$ is non-CP, the outputs $\mathcal{D}^{-1}\circ\mathcal{N}(\rho)$ are not valid quantum states anymore. The fidelity function does not always smaller than 1, thus is no longer a good metric. This explains why the recovery $\mathcal{D}^{-1}$ does not improve any expectation value but seems to have higher fidelities.