Quantum Nescimus

Improving the characterization of quantum systems from limited information

Robin Harper

2018

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

School of Physics

The University of Sydney

Abstract

We are currently approaching the point where quantum systems with 15 or more qubits will be controllable with high levels of coherence over long timescales. One of the fundamental problems that has been identified is that, as the number of qubits increases to these levels, there is currently no clear way to use efficiently the information that can be obtained from such a system to make diagnostic inferences and to enable improvements in the underlying quantum gates. Even with systems of only a few bits the exponential scaling in resources required by techniques such as quantum tomography or gate-set tomography will render these techniques impractical. Randomized benchmarking (RB) is a technique that will scale in a practical way with these increased system sizes. Although RB provides only a partial characterization of the quantum system, recent advances in the protocol and the interpretation of the results of such experiments confirm the information obtained as helpful in improving the control and verification of such processes. This thesis examines and extends the techniques of RB including practical analysis of systems affected by low frequency noise, extending techniques to allow the anisotropy of noise to be isolated, and showing how additional gates required for universal computation can be added to the protocol and thus benchmarked. Finally, it begins to explore the use of machine learning to aid in the ability to characterize, verify and validate noise in such systems, demonstrating by way of example how machine learning can be used to explore the edge between quantum non-locality and realism.

Author attribution statements.

The papers included in Part IV have been published as follows:

- Chapter 6: *Estimating the fidelity of T gates using standard interleaved randomized benchmarking*, Quantum Science and Technology, Volume 2, Number 1 — 9 March 2017 [1]. In this paper I was lead author having been responsible for the initial ideas contained in the paper and the production of the first and subsequent drafts. Steven Flammia provided comments and guidance in the construction of the paper. Joel Wallman and David Gross provided helpful conversations developing the idea of analyzing the anisotropy of the noise.

- Chapter 7: *Explaining quantum correlations through evolution of causal models*, Phy. Rev. A 95, 042120 — 17 April 2017 [2]. In this paper I was joint lead author, responsible for the theoretical parts and computational aspects (the genetic algorithm) of the paper. Robert Chapman, the other lead author, was responsible for conducting the experiment that supplied data to the algorithm. Chris Ferrie and Chris Granade were both significantly involved in the drafting of the paper and the initial exploratory coding of the experiment. Richard Kueng was instrumental in the theoretical side of the paper, providing the initial drafts of the theoretical bounds reported. Steve Flammia provided invaluable assistance and guidance, commenting on drafts and also providing some of the theoretical bounds.

- Chapter 8: *Nonexponential fidelity decay in randomized benchmarking with low-frequency noise*, Phys. Rev. A 92, 022326 — 11 August 2015 [3]. In this paper Stephen Bartlett, Steve Flammia and I provided theoretical support to the experimental results reported in the paper. The experimental team were the lead authors of the paper. Whilst the theoretical ideas related to the experimental design in the paper were developed in group discussions, I was responsible for first drafts of the relevant sections and providing the statistical analysis relating to the experiments, the introduction of alternative models and the calculation and write up of the Akaike bounds used to discriminate between the models.

- Chapter 9: *Estimating the coherence of noise*, New Journal of Physics, Volume 17 — November 2015 [4]. In this paper Joel Wallman was the lead author. I was involved in all relevant discussions and meetings relating to the drafting of the paper and the development of the theories as well as being responsible for many of the numerics and the production of numerical examples illustrating the principles behind the paper.

The first part of this thesis forms a review of the concepts and papers relating to the matters contained in this thesis, attempting to place a context around the work included in Part IV. Further original work contained within the literature review, but not in the published papers, includes:

- Subsection 3.1.4, where I provide a way to visualize Wallman's analysis ([5]) of the effect of gate dependent errors in a single qubit model.

- Subsection 3.2.5, where I begin the process of analyzing how [5] might impact on randomized benchmarking's interleaved protocol.
- Subsection 3.3.4, subsection 3.3.5 and subsection 3.3.6, where I set out previously unpublished work relating to the statistics surrounding the analysis of the results derived from Randomized Benchmarking experiments. This includes analysis relating to variations of experimental design (this is touched on in the paper reproduced in chapter 8). Most of the ideas have resulted from discussions between Joel Wallman, Chris Ferrie, Steven Flammia, Chris Granade and myself and will form the basis of a paper currently being written.

Acknowledgements

I would like to thank my supervisor Steven Flammia and co-supervisor Stephen Bartlett. My route to quantum information theory was not an orthodox one and I am grateful for their leap of faith in agreeing to accept me as a PhD candidate and for their continuous help, support, advice and encouragement during my candidature.

I am also grateful for the help, discussions, explanations and patience of those I have worked with, including Chris Granade, Chris Ferrie, Richard Kueng, and Joel Wallman and all the postgrads in the department, but especially Chris Chubb, Jacob Bridgeman, and Sam Roberts who always appeared to be happy to explain the finer points of some aspect of theory that I might not have quite fully grasped.

Finally I would like to thank my family. My wife Kate and children Imogen, Harriet and Hamish who tolerated with good humour my excited enthusiasm for all things science, bore with good grace my tendency to over-explain and provided the love and support that allowed me to complete this thesis. The title of this thesis came from Harriet's love of latin and loosely translated means "what a lot we don't know", which seems entirely appropriate.

Certificate

This is to certify that to the best of my knowledge, save as expressly acknowledged, the content of this thesis is my own work. This thesis has not previously been submitted for a degree at this or any other university.

Signature:

Name: Robin Harper

Contents

Part I

Introduction

Chapter 1

What are Quantum Computers?

— and why do we care?

> "A single particle reflects the light of creativity."
>
> Randomly generated fictional Deepak Chopra quote courtesy of *wisdomofchopra.com*.

A quantum computer is an embodiment of the desire to use the mathematics describing the quantum mechanical nature of reality in a way that will allow us to compute and solve certain problems that will never, in a practical sense, be able to be solved by a classical binary computer. It is not the case that the classical computer cannot compute such problems, rather, the limited resources available make such calculations practically impossible. Some problems scale in a way that is exponential in the size of the inputs and quickly, with such problems, even if they are of a very modest size the (classical) computing resources of the universe would be insufficient to solve them.

Richard Feynman pointed out that one such problem is the simulation of even a modest quantum mechanical system: "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly its a wonderful problem, because it doesnt look so easy" [6].

Since then the theory surrounding quantum computers has matured, and so have the potential uses of one — if it could be built. Shor's algorithm [7] shows how a quantum computer could be used to factorize large numbers efficiently, opening up a range of problems otherwise unsolvable by any presently known method, including ones on which modern day encryption relies. Quantum cryptography shows how quantum computers can be used to better secure the information they would otherwise make vulnerable. Grover's algorithm [8] opens the door to improved search techniques, and so on. Whilst the instances where a quantum computer would be exponentially more able to solve problems than a classical one are limited, they are tantalising instances. Moreover, the ability to be better able to simulate the quantum mechanical nature of reality will be of

immense help to every other area of science.

As will be discussed later, one of the main differences between a quantum bit (a *qubit*), and a classical bit is the ability for the qubit to be in a superposition, effectively 'partly' on and 'partly' off at the same time. This superposition is delicate and easily destroyed. A loosely similar claim might have once been made for analog computers — where each analog bit could store (to an arbitrary precision) a number. Other than in very specialised applications analog computers were made obsolete by digital computers, mainly because of the lack of ability to control the precision of the computer in the presence of noise, which meant that processes could not be reliably repeated with exact equivalence. The lack of a fixed state is a liability when carrying out such computations. Whilst the 'power' of a quantum computer comes from many aspects of the nature of quantum mechanics (such as the wave-like behaviour of the quantum mechanical processes), it is reasonable to ask what gives us confidence that a quantum computer will overcome the noise limitations that so hamper classical analog computers.

The key to answering this question is the fact that quantum states can be error corrected even when the gates and measurements implementing the error correction are themselves faulty. This theory uses the existence of error correcting codes (such as the Shor code), stabilizer states (e.g. [9]) and the results of the threshold theorem (briefly described in section 4.1.3) to show that, providing the error rates and error correlations can be reduced below a certain threshold level, noise is a surmountable barrier in the quest to build a universal quantum computer.

This thesis does not look at the error correction of a system (see chapter 10 of [10] for an introduction to this subject), rather it looks at the current thinking of how we might determine whether the control of the system is sufficient to bring the error rates below the required threshold levels.

1.1 What are we trying to measure?

In any quantum information system initial states are prepared, quantum gates applied, and measurements made. Given that this is a physical process the results of such experiments are likely to differ from those that might be expected by theory. Assuming the theory is correct, the inconsistency between the results of the process and the theoretical results are due to errors in the system, typically classified as noise. As previously mentioned, the threshold theorem assures us that if we engineer our system so that these errors are below a certain rate and that there is a sufficiently low correlation between errors then we can reliably correct the errors as they occur and commence with large scale quantum computation. The exact level of fidelity required depends, amongst other things, on the underlying system, the error correction employed, and the amount of parallelism inherent in the error correction regime. Work by Kitaev [11] and many others give us aspirational failure rate goals that people are on the cusp of achieving.

The rest of this thesis explores how we can measure the failure rates of the gates and processes underlying our quantum information system.

1.2 What are the difficulties?

The most direct method of characterizing a quantum state is quantum state tomography [12]. Quantum state tomography is a method of estimating a quantum state by creating the state multiple times, measuring it sufficiently many times and in sufficiently many different ways that the information gained can be used to estimate the state. Because of the nature of quantum mechanics, this is primarily a probabilistic methodology allowing, at best, with increased measurements an increasing confidence level that the estimated state is likely to be the one that produced the observed data. Subsequent to its introduction, better methods of fitting the data such as maximum likelihood estimator [13], Bayesian mean estimation [14] and adaptive Bayesian techniques [15] have been proposed. Other papers have concentrated, for example, on probing pure states [16] where techniques such as compressed sensing can be used to reduce the number of measurements and measurement settings required.

However in all these cases the number of measurements required, which scale exponentially with the number of qubits, means that tomography will not scale past the several qubit level.

Knowing a state, however, is not sufficient to know how the quantum system will work. For that, we need to know the noise characteristics or the noise profile of the system. Quantum process tomography attempts to do that by probing the system with quantum states, to see how it affects the states; that is, it attempts to characterize the dynamics of the system [17]. For our purposes, the important dynamic is the noise of the system, which we can think of as a quantum channel (these terms are all made more precise in chapter 2). For each state with which we probe the system we need to perform quantum tomography on the state that comes out, again leading to an exponential increase in the data we need to gather. In addition, the inability to separate our measurement and preparation errors from our control errors makes it difficult to know where in the system errors are occurring.

Finally, as the control of quantum systems improves, with fidelities exceeding 99% being reported (for example, [18, 19]) the fact we are trying to measure a low probability event makes it increasingly difficult to observe the failure and so the number of measurements required to obtain useful information about the error increases.

One protocol designed to overcome some of these difficulties is gate set tomography (*GST*). GST is described in [20], with a detailed description of its application to a particular experiment to be found in [21]. In brief it has grown out of quantum process tomography using sequences containing informationally complete *fiducials* to probe the system. These sequences are increased in length, allowing small errors to be magnified

and measured. All this information, of course, comes at a cost. As Greenbaum notes in [20], GST is more demanding than quantum process tomography in the number of experiments, the scaling over qubits and the post processing. Whilst experiments have been done with one qubit, and recently 2-qubit GST experiments have been proposed, after that we will likely have to look elsewhere.

It is, of course, clear that we will never be able to fully characterize a complex quantum system of more than a few qubits. If we could fully characterize it we could simulate it and the ability to simulate such a system is exactly one of the reasons we desire to build quantum computation devices. It is not just a question of how little we know, but also a question of how little we can ever know. What we do need to know is how close the (multi) qubit system is to being compliant with the error correcting threshold theorem. That is one of the things that randomized benchmarking attempts to enable.

1.3 How does randomized benchmarking attempt to deal with these difficulties?

The randomized benchmarking protocol [22, 23, 24, 25] is described in detail in Part II. In essence, it deals with the problem of small error rates by using long sequences of quantum gates, with the aim of amplifying such errors. The protocol provides a scalable method for doing this.

By varying the number of gates used, and noting that the increasing failure of the system to return to the desired state fits an exponential decay curve, the error rate associated with the control of the gates is isolated from the error rates relating to state preparation and measurement. This renders the protocol robust to such preparation and measurement errors. Finally, because we do not seek to know everything about the system, but just a 'figure of merit', the system is scalable and can be used to extract this figure of merit from multi-qubit systems. What is this figure of merit? As will be discussed later, we now know that randomized benchmarking returns the loss of fidelity in the system as a result of the average noise between idealized gates — this is what is meant by the average fidelity.

Unfortunately, whilst randomized benchmarking returns a gauge free* figure of merit, the average fidelity is not exactly the figure we need to satisfy the threshold theorems. To see why this is the case it is only necessary to appreciate that an average error rate, on its own, does not convey information about how extreme some of the errors might be — only what they average out to.

The diamond distance is a metric (discussed in section 4.1.3) that might, loosely, be thought of as measuring the worst-case error. This metric is one that does meet the criteria required by the threshold theorems but there are a number of problems. One

*see section 3.1.3 for a discussion of this term.

is that as finding the diamond distance involves a search problem, there cannot be an efficient direct experimental protocol to measure it. Another is that, in the absence of further information about the noise, the diamond distance can scale with the square root of the average *infidelity* of the system. Accordingly, if we are trying to ensure a diamond distance error of $O(10^{-4})$, this means — without further information — we would need gates with an average infidelity of less than $O(10^{-8})$.

Fortunately it appears that there is a further metric we can obtain, in a robust scalable way, using a protocol very similar to randomized benchmarking, that will allow us to determine the appropriate scaling for the diamond distance. This is discussed in more detail in subsection 2.3.1 and in the paper that introduces it (attached as chapter 9). Together these two pieces of information tell us a great deal about the noise afflicting the system in question, allowing us to characterize better the quantum system from this limited, but knowable, information.

Chapter 2

Quantum Information, States and Fidelity

> Hilbert space is big. Really big. You just won't believe how vastly, hugely, mindbogglingly big it is. I mean, you may think it's a long way down the road to the chemist, but that's just peanuts to Hilbert space.
>
> With apologies to Douglas Adams and Carl Caves.

This chapter aims to introduce the concepts and notation used in the rest of this thesis. Whilst many of the concepts will already be known to those familiar with the field (especially the first section), it may be helpful to glance through the relevant sections to ensure familiarity with the notation that will be employed.

This chapter is not designed to serve as an introduction to quantum information theory in general, for that the reader is referred to textbooks such as Nielsen and Chuang's *Quantum Computation and Quantum Information* [10] or an even more gentle introduction such as [26]. Rather this chapter outlines the aspects of quantum information theory that are pertinent to randomized benchmarking and, more generally, quantum characterization, verification and validation ($QCVV$). Some familiarity with linear algebra will be assumed, such as complex numbers, vectors, matrices and matrix multiplication.

2.1 How do we represent quantum information?

2.1.1 Kets/Bras, the quantum state

Mathematically quantum mechanical systems can be represented as existing on complex Hilbert spaces[†]. A Hilbert space is an abstract vector space that possesses the structure

[†]Sticking with the theme of the introductory quotation ... if you are unfamiliar with the concept of Hilbert spaces ... Don't Panic. If you just think of it as something for vectors (as directed lines) to exist in that will probably suffice for this introduction.

of an inner product (discussed in subsection 2.1.2) and allows length and angle to be measured. Typically we can represent states of a quantum mechanical system as normalized complex vectors on such a complex Hilbert space. For the remainder of this thesis we will almost invariably be dealing with finite-dimensional complex Hilbert spaces (denoted as $\mathcal{H}$). If we have a d-dimensional complex Hilbert space this will consist of a vector space of $\mathbb{C}^d$ over $\mathbb{C}$ and the quantum states in this system will consist of vectors that will be represented using the notation below and can be thought of as $d \times 1$ matrices (i.e. d rows and 1 column).

Applying the above to the computation of information in such a quantum mechanical system, we can imagine a two state system (a *qubit*), in loose analogue to a classical binary computation (where a bit represents something that can be either off (0) or on (1)). Here we might imagine a quantum mechanical system with two states, such as the spin-up, spin-down state of an electron or two different energy levels of a real or artificial atom. For a single qubit we can represent these states as orthogonal vectors in our complex Hilbert space (for a single qubit, this will be a vector field $\mathbb{C}^2$ over $\mathbb{C}$), thus:

$$\text{Zero state } = |0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} ; \text{ One state } = |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} . \tag{2.1}$$

Where $|0\rangle$ and $|1\rangle$ represent the states in the *bra-ket* or *Dirac* notation originally introduced by Paul Dirac in 1939. Although a vector space can be spanned by a set of arbitrary, linearly independent, vectors, we will typically choose as a basis the basis outlined above, often referred to as the *computational basis*.

2.1.2 Pure states

Unlike bits in a classical digital computer, which are either 1 or 0, a qubit can be a 'partly' on and 'partly' off at the same time, that is it can be in a superposition. Pure states $|\psi\rangle$ (as opposed to mixed states, see subsection 2.1.6) will be represented by unit length vectors in our Hilbert space. Using the computational basis state set out above, for a single qubit we have:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \, . \tag{2.2}$$

To enforce the unit length we need to turn to the inner product of the Hilbert space. Here the inner product on $\mathbb{C}^n$ is defined as:

$$\langle (w_1, ..., w_n), (z_1, ..., z_n) \rangle = w_1^* z_1 + ... + w_n^* z_n \, , \tag{2.3}$$

where w_1^* represents the complex conjugate of w_1.

Therefore we enforce the unit length by requiring the inner product of a pure state with itself to be equal to 1 and so, if we write $|\psi\rangle$ as (α, β):

$$\langle |\psi\rangle, |\psi\rangle \rangle = \langle (\alpha, \beta), (\alpha, \beta) \rangle = \alpha^* \alpha + \beta^* \beta = |\alpha|^2 + |\beta|^2 = 1 \, . \tag{2.4}$$

The notation here can be considerably simplified by using the Dirac notation introduced earlier. In this case we denote the inner product between two states $|\psi\rangle$ and $|\phi\rangle$ as $\langle\phi||\psi\rangle$ or even more concisely as $\langle\phi|\psi\rangle$. $\langle\phi|$ is the *bra* of Dirac's bra-ket notation and represents the *dual vector* of the vector represented by $|\phi\rangle$. Using our single qubit example, then looking at $|\psi\rangle$, we have $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, which can be represented by vectors as follows:

$$|\psi\rangle = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} . \tag{2.5}$$

Since $|\psi\rangle$ is represented by a 2 x 1 vector, the dual vector, $\langle\psi|$, is represented by a 1 x 2 vector, namely:

$$\langle\psi| = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix} , \tag{2.6}$$

and the inner product with itself $\langle\psi|\psi\rangle$ is therefore:

$$\langle\psi|\psi\rangle = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \alpha^*\alpha + \beta^*\beta = |\alpha|^2 + |\beta|^2 , \tag{2.7}$$

as before.

The expansion of these ideas to multiple qubits is dealt with through *tensor products*. Tensor products are a way of combing vector spaces to form larger vector spaces. Just as in a classical system where the amount of information that can be stored increases with the number of bits, the size of the Hilbert space required to represent the quantum mechanical state increases with the number of qubits.

For a binary classical system with two bits, instead of two values (on, off), we now have 4 values (on on, on off, off on, off off). In our two qubit system, in our computational basis, we would now have four basis vectors (if you like, representing each of the classical values) spanning our increased vectors space, each of which can be 'partly' on:

$$|\psi\rangle = \alpha \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} + \gamma \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} + \delta \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \\ \gamma \\ \delta \end{pmatrix} , \tag{2.8}$$

where, as before, $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$.

The tensor product (or, as it is called, where we are working in matrix representations, the Kronecker product) allows us to construct these larger Hilbert spaces. Assuming we have two Hilbert spaces V and W, the tensor product of these spaces is written as $V \otimes W$. If we have a collection of vectors $\{|i\rangle\}$ that are an orthonormal basis for V and a set $\{|j\rangle\}$ that form an orthonormal basis for W then $\{|i\rangle \otimes |j\rangle\}$ will form an orthonormal basis for $V \otimes W$. To see how this works we can look at the definition of the Kronecker product for matrices. If A is an m x n matrix and B is a p x q matrix, then $A \otimes B$ is a

$mp \times nq$ matrix constructed as follows:

$$A \otimes B = \begin{pmatrix} A_{1,1}B & A_{1,2}B & ... & A_{1,n}B \\ A_{2,1}B & A_{2,2}B & ... & A_{2,n}B \\ \vdots & \vdots & \vdots & \vdots \\ A_{m,1}B & A_{m,2}B & ... & A_{m,n}B \end{pmatrix}. \tag{2.9}$$

For a two qubit system, working in the computational basis we have a basis for V (qubit 1) of $\{\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}\}$ and W (qubit 2) also $\{\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}\}$.

The Hilbert space of the two qubit system will then have basis vectors as follows:

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 * \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ 0 * \begin{pmatrix} 1 \\ 0 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 * \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ 0 * \begin{pmatrix} 0 \\ 1 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 * \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ 1 * \begin{pmatrix} 1 \\ 0 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 * \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ 1 * \begin{pmatrix} 0 \\ 1 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \tag{2.10}$$

as before.

Tensor products can be taken over matrices as well as vectors, using the same method.

The following are important examples of the properties of the tensor product:

1. For a scalar c,
$$c(|\psi\rangle \otimes |\phi\rangle) = (c|\psi\rangle) \otimes |\phi\rangle = |\psi\rangle \otimes (c|\phi\rangle). \tag{2.11}$$

2. For arbitrary $|\psi_1\rangle$ and $|\psi_2\rangle$
$$(|\psi_1\rangle + |\psi_2\rangle) \otimes |\phi\rangle = |\psi_1\rangle \otimes |\phi\rangle + |\psi_2\rangle \otimes |\phi\rangle \tag{2.12}$$
$$|\phi\rangle \otimes (|\psi_1\rangle + |\psi_2\rangle) = |\phi\rangle \otimes |\psi_1\rangle + |\phi\rangle \otimes |\psi_2\rangle. \tag{2.13}$$

3. For arbitrary operators (A and B) and arbitrary states
$$(A \otimes B)(|\psi_1\rangle|\psi_2\rangle\langle\phi_1|\langle\phi_2|) = (A|\psi_1\rangle\langle\phi_1|) \otimes (B|\psi_2\rangle\langle\phi_2|). \tag{2.14}$$

4. Using linearity to extend to all elements in the relevant Hilbert spaces

$$(A \otimes B)(\sum_{i,j} |\psi_i\rangle|\psi_i\rangle\langle\phi_j|\langle\phi_j|) = \sum_{i,j}(A|\psi_i\rangle\langle\phi_j|) \otimes (B|\psi_i\rangle\langle\phi_j|)\,. \tag{2.15}$$

2.1.3 The Born rule

One question that remains is that if our model of a quantum state can be 'partly' on and 'partly' off, how do we represent measuring it, given that for the quantum mechanical system we are trying to measure such measurement will always give us that it is 'on' or 'off', 'up' or 'down' etc. For this we use the Born rule, first stated by Max Born in 1926. The Born rule is a postulate of quantum mechanics. The rule states that if we have an observable A (more on what this means later), then the measured result will be one of the eigenvalues of A. If the system is in a state $|\psi\rangle$ and measured with the observable A, then the probability that an eigenvalue λ_i of an eigenvector $|e_i\rangle$ of A being found is:

$$P(a = \lambda_i|\,|\psi\rangle) = |\langle e_i|\psi\rangle|^2\,. \tag{2.16}$$

Using our ongoing example, in the computational basis we have two eigenvectors $|0\rangle, |1\rangle$. With an appropriate measurement the chance of finding our state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ in the $|0\rangle$ state is:

$$|\langle 0|\psi\rangle|^2 = |\langle 0|(\alpha|0\rangle + \beta|1\rangle)|^2 = |\alpha\langle 0|0\rangle + \beta\langle 0|1\rangle|^2 = |\alpha|^2\,. \tag{2.17}$$

or using our matrix representation

$$|\langle 0|\psi\rangle|^2 = |(\begin{array}{cc} 1 & 0 \end{array})\begin{pmatrix} \alpha \\ \beta \end{pmatrix}|^2 = |\alpha|^2\,. \tag{2.18}$$

Measurement will be discussed more in subsection 2.1.5, but from above it can be seen that the Born rule only gives you a probability of *observing* the quantum state in one state (or the other). Thought of this way, a state corresponds to a linear functional that maps observables onto real numbers, being the probabilities that each eignevalue of A will be found, where A here is the observable. A single measurement only provides limited information as to what the state was (and then changes the state). If it is possible to prepare the state multiple times, then multiple measurements can be taken and the statistical likelihoods of various αs and βs can be calculated. This is, however, essentially a statistical exercise. The inability to reduce this to anything other than a probability is linked to the concept of *quantum indeterminacy* and *quantum randomness*. The discord between this view and Einstein's local realism is a fascinating area of quantum foundations with implications in diverse areas of quantum theory. Some of these issues are explored in more detail in the paper published in chapter 7 where we use a genetic algorithm to probe such matters.

2.1.4 Operators

The next stage is to look at how we can represent the evolution of a closed quantum system, that is, how can we transform our qubit. In the classical case this is analogous to asking what type of logical operations can be performed on our bits. The formalism of quantum mechanics used in this thesis is the unitary formulation of quantum dynamics. This postulates that the evolution of a system can be described by unitary transformations. (See, e.g., [10] for a description of the one-to-one correspondence between the Hamiltonians of Schrödinger's equations and such unitary operators.) That is, we can describe the transformation of our state $|\psi\rangle$ to a different state $|\psi'\rangle$ by way of unitary *operator* U. We can represent this as:

$$U|\psi\rangle = |\psi'\rangle\,. \tag{2.19}$$

The fact that the operator is unitary (i.e. $U^\dagger U = UU^\dagger = I$, where $U^\dagger$ is the complex conjugate transpose of U) embodies the reversibility concept of quantum mechanics.

$$\begin{aligned} & U|\psi\rangle = |\psi'\rangle \\ \implies\ & U^\dagger U|\psi\rangle = U^\dagger|\psi'\rangle \\ \implies\ & |\psi\rangle = U^\dagger|\psi'\rangle\,. \end{aligned}$$

More formally we can describe the quantum operator as a unitary linear operator that takes a state in the Hilbert space to another state in the Hilbert space i.e. $U : \mathcal{H} \to \mathcal{H}$. We will return to this later.

In our matrix representation, for a single qubit, a quantum operator would be represented by a 2x2 matrix. For instance, the Pauli X operator (discussed later), would be represented as

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

and acts as follows:

$$X|\psi\rangle = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \beta \\ \alpha \end{pmatrix} \tag{2.20}$$

Thus the Pauli X operator has 'flipped' the amplitudes of finding the state in the $|0\rangle$ or the $|1\rangle$ state. That is why it is commonly known as the *bit-flip* operator.

2.1.5 Measurement

The above formalism for the evolution of a quantum system, assumed the system was closed, that is it didn't interact with the rest of the world. This is obviously not the

case where we perform a measurement on the system. Performing a measurement, to find out what is going on in the system, is an interaction with the system and means it is not necessarily subject to unitary evolution.

A quantum measurement can be described by a set of operators, in which there is one operator for each possible outcome and the outcome's probability is given in terms of the corresponding operator by:

$$p(m) = \langle\psi|M_m^\dagger M_m|\psi\rangle\,. \tag{2.21}$$

Clearly, as the measurement operators cover all possible measurement outcomes then, when we measure a state, we expect one of the these outcomes. Accordingly, the probabilities of each of the different possible measurement outcomes should sum to one. This means that:

$$\sum_m p(m) = 1 \tag{2.22}$$

$$\implies \sum_m \langle\psi|M_m^\dagger M_m|\psi\rangle = 1\,. \tag{2.23}$$

This further implies that the following must hold (the *completeness equation*)

$$\sum_m M_m^\dagger M_m = I\,. \tag{2.24}$$

To illustrate the above we use our single qubit example and imagine we wish to measure in the computational basis. We are expecting two measurement outcomes, either it is $|0\rangle$ or it is $|1\rangle$. Using our matrix representation we can construct the two measurement operators by taking the outer product (the tensor product of the state and its conjugate) of the two relevant states.

$$M_0 = |0\rangle\langle 0| = |0\rangle \otimes \langle 0| = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \tag{2.25}$$

$$M_1 = |1\rangle\langle 1| = |1\rangle \otimes \langle 1| = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \tag{2.26}$$

We can see that this satisfies the completeness equation:

$$\sum_{m=0}^{1} M_m^\dagger M_m = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I\,. \tag{2.27}$$

Looking at the probability of measuring our generic ψ in the $|0\rangle$ state, the measurement is given by M_0 and the probability is calculated as follows:

$$p(0) = \langle\psi|M_0^\dagger M_0|\psi\rangle = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \alpha^*\alpha = |\alpha|^2\,, \tag{2.28}$$

in accordance with the Born rule.

Quantum measurement affects the state of the system measured, the state of the system after measurement is:

$$\frac{M_m|\psi\rangle}{\sqrt{\langle\psi|M_m^\dagger M_m|\psi\rangle}}\,. \tag{2.29}$$

In this thesis measurement will be the final part of the process, and so I do not further explore how measurement affects the state, leaving me free to use a measurement formalism known as POVM (positive-operator valued measure). Essentially it should be noted that $M_m^\dagger M_m$ is a positive operator. Letting $P_m = M_m^\dagger M_m$, we have a set of measurements where:

$$\sum_m P_m = I \text{ and } p(m) = \langle\psi|P_m|\psi\rangle\,. \tag{2.30}$$

Projective measurements are an important subset of the general measurements described above which in addition to satisfying the completeness equation also satisfy the requirement that they are orthogonal projectors, that is for any two of the operators

$$P_i P_j = \delta_{i,j} P_j\,, \tag{2.31}$$

where $\delta_{i,j}$ is just the Dirac delta function, that is it is 1 if $i = j$ and 0 otherwise.

2.1.6 Density Operators

Whilst we now have the mechanisms to deal with a quantum state that is known, we have yet to introduce a notation that allows us to easily deal with a state that is actually an ensemble of possible states or that represents a subsystem of a larger entangled state. A single state, known as a *pure state*, can be represented by a ket, such as $|\psi\rangle$. To represent a system where we don't completely know the state, but may be in one of a number of states $|\psi_i\rangle$ each with probability p_i, we write:

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|\,. \tag{2.32}$$

It is worth noting that the density matrix doesn't represent a system in superposition (like $|\psi\rangle$) but rather is an ensemble of such systems, using entirely classical probabilities. Such a system is called a *mixed state*. An important property of such a set of density matrices is that they are convex, that is if ρ_i (for some i from $1...n$) are density operators, then the appropriately weighted mixture $\sum_{i=1}^n p_i\rho_i$ is also a density operator. Interestingly, there is no unique decomposition of a density operator, that is — whilst

every density operator can be decomposed into a sum of mixed states — there are infinitely many such decompositions. This is not through a lack of knowledge or ignorance of certain parameters, but rather reflects an important property of quantum mechanics, namely that there is no way to distinguish between preparations that cause a particular state.

A pure state, which we have previously represented as $|\psi\rangle$, can be written as a density matrix as $|\psi\rangle\langle\psi|$. Where we are using POVM elements to measure, then we can calculate the probabilities associated with such a measurement from the the trace of the product of the density operator and the relevant measurement. To illustrate, if we assume an arbitrary POVM P_0, then by using the cyclic property of the trace (Tr(ABC)=Tr(BCA)=Tr(CAB)), and remembering that $\langle\psi|P_0|\psi\rangle$ is a number (so the trace of it is just that number), we see:

$$\begin{aligned} p(P_0) &= \langle\psi|P_0|\psi\rangle && (2.33)\\ &= \mathrm{Tr}(\langle\psi|P_0|\psi\rangle) && (2.34)\\ &= \mathrm{Tr}(P_0|\psi\rangle\langle\psi|)\,. && (2.35) \end{aligned}$$

This extends naturally where we have a mixed state $\rho = \sum_i p_i|\psi_i\rangle\langle\psi_i|$, as follows:

$$\begin{aligned} p(P_0) &= \mathrm{Tr}(P_0\rho) && (2.36)\\ &= \mathrm{Tr}(P_0 \sum_i p_i|\psi_i\rangle\langle\psi_i|) && (2.37)\\ &= \sum_i p_i \mathrm{Tr}(P_0|\psi_i\rangle\langle\psi_i|) && (2.38)\\ &= \sum_i p_i\langle\psi_i|P_0|\psi_i\rangle\,. && (2.39) \end{aligned}$$

That is, the probability of measuring P_0 with the ensemble represented by ρ is the sum of the probabilities of the mixed state being in state $|\psi_i\rangle$ times the probability of measuring P_0 if it were in state $|\psi_i\rangle$.

Unitary evolution is equally straightforward. Again if we start with the pure state $|\psi\rangle$ represented as density matrix $\rho = |\psi\rangle\langle\psi|$, then a unitary operator U would evolve $|\psi\rangle$ as $U|\psi\rangle = |\psi'\rangle$. We would expect the density operator representing $|\psi'\rangle$ to equal $|\psi'\rangle\langle\psi'|$ and we know that $(U|\psi\rangle)^\dagger = \langle\psi|U^\dagger$. Therefore evolution of ρ by U is given as: $U\rho U^\dagger$, since

$$\begin{aligned} U\rho U^\dagger &= U|\psi\rangle\langle\psi|U^\dagger && (2.40)\\ &= |\psi'\rangle(U|\psi\rangle)^\dagger && (2.41)\\ &= |\psi'\rangle(|\psi'\rangle)^\dagger && (2.42)\\ &= |\psi'\rangle\langle\psi'|\,, && (2.43) \end{aligned}$$

as required.

With the linear nature of the operators we can see the same holds true for mixed states, the probabilities being unaffected, each component $U|\psi_i\rangle\langle\psi_i|U^\dagger \rightarrow |\psi_i'\rangle\langle\psi_i'|$.

Finally, it is worth noting that the trace of density operator ρ will equal 1, $\mathrm{Tr}(\rho^2) \leq 1$, with equality if and only if ρ represents a pure state, and the sum of all pure state projectors $|\psi\rangle\langle\psi|$ will give a density matrix (ρ) satisfying $\mathrm{Tr}(\rho) = 1$, where $\rho \geq 0$.

2.1.7 Entangled States and Partial Traces

We have discussed how composite quantum systems (e.g. more than one qubit) can be modeled by using the tensor product. The simplest type of combined system is a composite system where the states are all product states, that is a tensor product of individual states. However it is possible to create states that do not simply factorize as a tensor product of simpler states, that is for some observables $\langle A \otimes B\rangle \neq \langle A \otimes I\rangle\langle I \otimes B\rangle$.

Examples of this are the Bell states over 2 qubits. To illustrate — the $|\Phi^+\rangle$ state can be written as:

$$|\Phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}\,. \tag{2.44}$$

This state has no decomposition into two single qubit states, that is, it cannot be written as $|\psi_1\rangle \otimes |\psi_2\rangle$. The phenomena of entangled states were the subject of the EPR paradox after a 1935 paper by Einstein, Podolosky and Rosen [27]. Examples of using entanglement for such things as super-dense coding, violating the Bell inequality and quantum teleportation can all be found in [10]. These issues and their interaction with local realism are the subject matter of the published paper attached as chapter 7.

Of relevance to this section, however, are the tools we can use to deal with such composite systems. The main one described here is the reduced density operator.

If we have a composite system, with two distinct and disjoint subsystems, whose state is described by a density operator $\rho^{\mathcal{H}_1\mathcal{H}_2}$, then it is possible to throw away or discard one of the subsystems. This physical operation is represented mathematically by the partial trace: throwing away the subsystem represented by $\mathcal{H}_2$ corresponds to tracing it out, i.e. performing a *partial trace* over the factor space $\mathcal{H}_2$ to obtain a density matrix $\rho^{\mathcal{H}_1}$ supported only on $\mathcal{H}_1 E$. This is defined as:

$$\rho^{\mathcal{H}_1} \equiv \mathrm{Tr}_{\mathcal{H}_2}(\rho^{\mathcal{H}_1\mathcal{H}_2})\,. \tag{2.45}$$

The partial trace itself is defined as:

$$\mathrm{Tr}_{\mathcal{H}_2}\left(|\psi_1\rangle\langle\psi_2| \otimes |\phi_1\rangle\langle\phi_2|\right) \equiv |\psi_1\rangle\langle\psi_2| \mathrm{Tr}\left(|\phi_1\rangle\langle\phi_2|\right) \ ,$$

where in the first equivalence ψ_1 and ψ_2 are two vectors in $\mathcal{H}_1$ and ϕ_1, ϕ_2 in $\mathcal{H}_2$.

Using our two-qubit example, we can calculate partial traces by calculating the smaller matrix, through multiplication by rectangular matrices.

For example suppose we wished to calculate $\mathrm{Tr}_A(\rho_{AB})$, we need

$$\mathrm{Tr}_A(\rho_{AB}) = \sum_i \left(\langle i| \otimes \mathbb{1}\right) \rho_{AB} \left(|i\rangle \otimes \mathbb{1}\right) \tag{2.46}$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \rho_{AB} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \rho_{AB} \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \tag{2.47}$$

Similarly $\mathrm{Tr}_B(\rho_{AB}) = \sum_i \left(\mathbb{1}\otimes\langle i|\right) \rho_{AB} \left(\mathbb{1}\otimes|i\rangle\right)$.

2.1.8 Paulis and Cliffords

The last part of this overview section discusses the Pauli operators and the Clifford operators, both of which are used extensively in later sections and the papers attached to this thesis.

Pauli Operators

Using the matrix representation we have been building up, the Pauli operators are defined as follows:

$$\sigma_I \equiv \sigma_0 \equiv I \equiv \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \sigma_X \equiv \sigma_1 \equiv X \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

$$\sigma_Y \equiv \sigma_2 \equiv Y \equiv \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_Z \equiv \sigma_3 \equiv Z \equiv \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} .$$

There are a number of reasons why these operators will prove useful. The first is to note that they form a basis over the vector space of operators, that is any unitary operator can be written as a linear combination of the Pauli operators. They are Hermitian and unitary (i.e. $XX^\dagger = XX = YY^\dagger = YY = ZZ^\dagger = ZZ = II^\dagger = I$). They anti-commute $XZ = -ZX$, $XY = -YX$ and $YZ = -YZ$ and further $XZ = iY$. Only I has trace (the other Paulis have 0 trace).

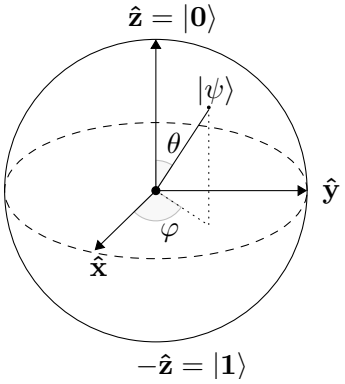

Figure 2.1: Illustration of the Bloch Sphere.

The usefulness of these operators can be more easily demonstrated by introducing a technique that can help visualize quantum operations, namely the Bloch Sphere.

Recalling that a general pure state $|\psi\rangle$ can be written as a combination of the computational basis vectors, e.g. $\alpha|0\rangle + \beta|1\rangle$, with the constraint that $|\alpha|^2 + |\beta|^2 = 1$. Now note that we can also write it as $|\psi\rangle = e^{i\gamma}\left(\cos(\frac{\theta}{2}) + e^{i\varphi}\sin(\frac{\theta}{2})\right)$, where we can ignore the phase factor out the front, as it has no observable effects. Writing it this way the parameters θ and φ can be used to define a point on a unit three-dimensional sphere. This sphere is known as the Bloch sphere, illustrated in figure 2.1.

Thinking of a state as a point on the sphere, the X, Y and Z Pauli operators then represent rotations of π radians around the X,Y and Z axes respectively. It can be seen that all the relations specified above hold true (e.g. each rotation is its own inverse, a rotation about the Z axis, followed by a rotation around the X axis is the equivalent of a rotation about the Y axis). In fact the Paulis together from a basis for all 2×2 matrices in $\mathbb{C}$. For any M which is a 2×2 matrix we can write $M = \frac{1}{2}(x_0\mathbb{1} + \vec{x} \cdot \vec{\sigma})$ where x_0 is the trace of M and $\vec{x} \cdot \vec{\sigma}$ is given by

$$\vec{x} \cdot \vec{\sigma} = \sum_{i=1}^{3} x_i \sigma_i \,. \tag{2.48}$$

Every pure state lies on the surface of the sphere. As previously discussed, a mixed state represented by the density operator ρ can be thought of as an ensemble of pure states, weighted by the probability of the system being in each of the respective pure states. If we imagine a mixed state consisting of a combination of two pure states, then we can imagine that ρ would appear in the Bloch sphere as a point on a line drawn between the two pure states, the position on the line depending on the probabilities assigned to each of the two states. Consequently, ρ would be below the surface of the sphere. The line joining ρ to the center of the sphere would thus be less than one (since the Bloch sphere is a unit sphere), thus we can see in this model the radius to the point represented by

ρ, (which equals $\mathrm{Tr}(\rho^2)$) is < 1 unless ρ is on the surface of the sphere, i.e. is a pure state.

The Bloch sphere, however, doesn't generalise in a way that aids comprehension for more than one qubit.

For more than one qubit the Pauli Group consists of the individual Paulis on each qubit tensored together in all possible combinations (so there would be 16 combinations for 2 qubits and 4^3 i.e. 64 for 3 qubits).

Clifford Operators

The next important group of operators is the Clifford operators. The Clifford operators are defined as the set of operators that conjugate the Pauli operators to themselves, ignoring global phase. More formally:

$$C \equiv U \,|UPU^\dagger \subset P\,, \tag{2.49}$$

where P is the set of Pauli Operators. Technically this is the second level of the *Clifford hierarchy* [28], although, if not otherwise specified a reference to the Cliffords will mean a reference to this second level of the hierarchy. The Cliffords form a finite group, albeit one that grows quickly as the number of qubits increase. Where I refer to the group formed by the Cliffords, I will use $\mathfrak{C}$ to represent the group of these operators and $|\mathfrak{C}|$ to represent the number of elements in the group. $\mathfrak{C}^{\otimes n}$ represents the group over n qubits. The size of the group is [29]:

$$|\mathfrak{C}^{\otimes n}| = 2^{n^2+2n} \prod_{j=1}^{n} \left(4^j - 1\right)\,, \tag{2.50}$$

which starts at 24 for one qubit, is 11,520 for 2 qubits and shortly after that becomes quite unwieldy.

For one qubit the Clifford group can can be generated as products of just two distinct generator gates. For more than one qubit a third gate (being a two qubit gate), needs to be introduced. The generating group then consists of the one qubit gates acting on each of the qubits and the two qubit gate, acting on each pair of qubits. Although the number of generators needed to create each Clifford group gate will vary, ref. [30] gives an efficient algorithm for decomposing any n-qubit Clifford (in fact, any n-qubit stabilizer circuit) into no more than $O(\frac{n^2}{\log n})$ generating gates.

If it is only required to generate a unitary 2-design (subsection 2.3.2) (using Clifford gates) it is possible to use a subset of the Clifford group. In [31] it is shown that such a subset (being an ensemble of Clifford group elements that mix the Paulis, followed by a random Pauli) forms a unitary 2-design. They then show how, for n qubits, the action of the Clifford group on the Pauli group elements can be associated with elements of the Galois Field of size 2^n ($\mathrm{GF}(2^n)$). Given this they then provide a construction of such gates that uses a maximum of $O(n \log n \log\log n)$ one and two qubit Clifford gates. It

should be noted that their construction does, however, require the use of ancillas and assumes that the extended Riemann Hypothesis is true.

Two one qubit gates that can generate the single qubit Clifford group are the Hadamard and Phase gates, given by

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}. \tag{2.51}$$

The two qubit gate is the CNOT gate, which applies an X gate to a target qubit, if the control qubit is one (otherwise it does nothing). In matrix form (in the computational basis) for two qubits this looks as follows:

$$\text{CNOT}_{1,2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \tag{2.52}$$

More generally, to construct a different CNOT gate using this representation note that $|0\rangle\langle 0|$ picks out the state where a bit is zero, and $|1\rangle\langle 1|$ where it is 1. The above matrix can be formed as follows $\text{CNOT}_{1,2} = |0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes X$ which shows us how to generalise for an arbitrary number of qubits. If the control (say on qubit c) is before the target (on qubit t) then for an n-qubit system $\text{CNOT}_{c,t}$ becomes:

$$\text{CNOT}_{c,t} = \mathbb{1}^{\otimes(1..c-1)} \otimes ((|0\rangle\langle 0| \otimes \mathbb{1}^{\otimes(c..t)}) + |1\rangle\langle 1| \otimes (\mathbb{1}^{\otimes(c..t-1)} \otimes X)) \otimes \mathbb{1}^{\otimes(t+1..n)}, \tag{2.53}$$

where an expression of the form $\mathbb{1}^{\otimes(1..i)}$ represents the identity tensored with itself i times. The reverse of this CNOT (where the target is before the control) is created from the following (and similarly expanded) $\text{CNOT}_{2,1} = \mathbb{1} \otimes |0\rangle\langle 0| + X \otimes |1\rangle\langle 1|$.

Together the Pauli group and the Clifford group play a crucial role in quantum computation, for many reasons, including:

- They provide a natural class of circuits that can be simulated efficiently on a classical computer [32].
- Most error correcting codes make extensive use of Clifford circuits [9].
- They have the property of satisfying the technical condition of being a unitary 2-design [24], (discussed in subsection 2.3.1) which means they are used in protocols designed to measure quantum systems (e.g. [22, 33])

2.2 Channels

So far we have discussed the representation of quantum states, density operators, measurements and the unitary operators that allow us to describe evolution of the quantum

system over time. The next section discusses quantum channels; it looks at how we can describe what happens to a quantum system, containing quantum information, if it is transmitted through a communication channel. There are many reasons to study this, the most pertinent to this thesis being that one can view the noise processes inherent in any quantum operation as being the equivalent of transmitting the system through a noise channel. Since noise is inherent in any system, and it is exactly the impact of this noise we wish to measure in QCVV, then we need a way to describe it so we can measure it.

2.2.1 Stuff you can alter stuff with

Formally we can think of a quantum channel as a linear map over the relevant vector space, and indeed I use the two terms, map and channel, interchangeably. Unlike an operator (which is a bounded linear operator), a channel is the space of completely positive linear operator maps. So, whilst an operator describes an operation that changes a state $|\psi\rangle \rightarrow |\psi'\rangle$ a channel is a map of operators $U \rightarrow U'$.

Intuitively one can see that noise might turn an X-gate into an imperfect X-gate, a Y-gate into a (differently) imperfect Y-gate and so on. Such a noise process can be modeled by such a linear map (that is one from operators to operators). Moreover the noise map will describe a physical system (one that can be created in reality) provided some specific criteria are met. Initially we will be dealing with systems where the quantum state does not just vanish. In this case we would expect our map (or channel) to be trace preserving as that represents the total probabilities of finding the system in one of our measurable states. That is, if we have a noise channel $\mathcal{E}$ then for all density operators ρ we would expect $\mathrm{Tr}(\mathcal{E}(\rho))=\mathrm{Tr}(\rho)$.

Of course, in some systems it is possible for the state to leave the computational basis states, i.e. escape from the system. If, for instance, the qubit is encoded as occupying the ground state and first excited state of an ion then if the ion was to enter a second or higher state we would have a loss from the computational basis; such a map would not be trace preserving. However, for most of this thesis we will be dealing with trace preserving (TP) channels. As well as being trace preserving, we need a map that maps positive operators to positive operators and not a physically impossible operator (e.g. a non-positive operator). Clearly a physically possible map must map physically feasible operators to physically possible operators. Moreover we still require this to be the case, even where the map is just part of a bigger system. To achieve this we require our map to *completely positive* (CP). In general complete positivity and being trace-preserving ($CPTP$) are both necessary and sufficient conditions for a noise map to represent a valid quantum operation [34].

There are many ways to represent such a CPTP map such as using a Stinespring dilation, a linear superoperator (Liouville) representation, by use of Kraus operators (also known as the operator-sum representation) or through the Choi-Jamiolkowski isomorphism or

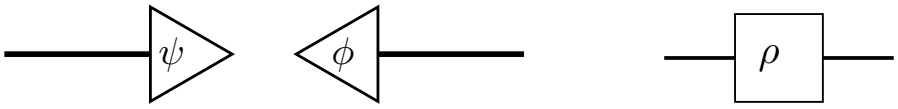

(a) vector - a ket $|\psi\rangle$ (b) vector - a bra $\langle\phi|$ (c) density operator - ρ

Figure 2.2: Some examples of the graphical calculus used.

the related χ representation. Whilst we will mainly concentrate on a variation of the Liouville superoperator each of the other representations have their uses and so I will briefly describe each of them below.

2.2.2 Graphical calculus - an interlude

A graphical calculus is a calculus of graphs, typically designed to aid defining and reasoning about relations. It provides a method of representing relations together with graph transformation rules that allow the graphs to be transformed i.e. reasoned about. In the current context there are two popular approaches based on tensor networks and category theory. Here I will briefly introduce a form of tensor networks described in detail by Wood et al. in [35]. For a more detailed introduction to tensor networks in general, with a specific view to quantum systems [36] is recommended. Whilst such a graphical calculus can be used to rigorously prove transformations between diagrams it can also help to provide a more intuitive insight to the various representations of channels. It is for the latter reason I will use such a calculus here and therefore the descriptions and the transformations might not be as detailed as would be required for a proof. The interested reader is referred to [35].

Unlike typical tensor networks we impose a left-to-right order on the diagrams. Each wire represents a specific size in our system, being the dimension of the Hilbert space of the system in question. States (vectors) and the dual-vectors (the bras introduced earlier) are represented as triangles. Linear operators (including density operators) are represented as boxes. An index corresponds to the open indexes (on the equivalent matrix representation) so higher order tensors are represented with increasingly more wires. For instance in figure 2.2(a) the ket $|\psi\rangle$ has a single wire leading into it representing the open index of the $d \times 1$ matrix representation of ψ, whereas in (c) the density operator ρ, a $d \times d$ matrix has a left and right wire, representing each of the two open indexes.

Unlike normal tensor diagrams, orientation here is important, with bending the wires being the equivalent of a transpose (not the dagger — so we have to be careful with complex conjugates). These various rules are illustrated in figure 2.3. The final manipulations we will cover here are taking the trace (joining two wires together), multiplication and the tensor product, these are illustrated in figure 2.4. Anything else required will be introduced as needed.

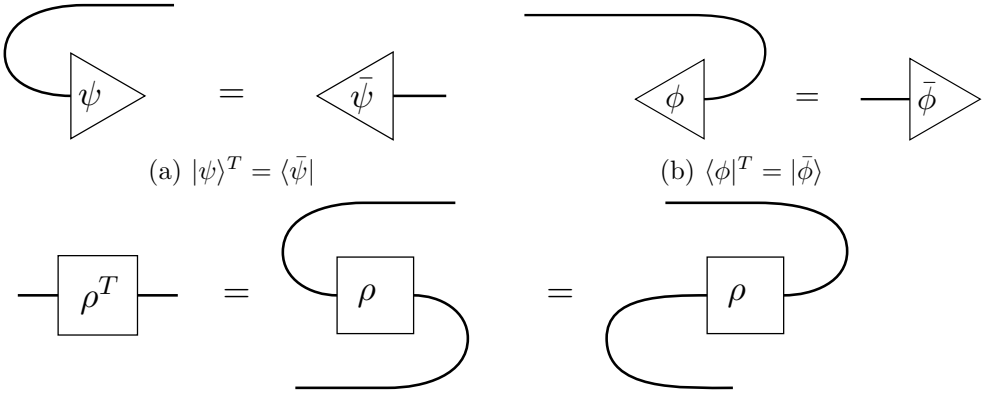

(a) $|\psi\rangle^T = \langle\bar{\psi}|$ (b) $\langle\phi|^T = |\bar{\phi}\rangle$

(c) Transposition of an operator

Figure 2.3: Transposing and dual spaces.

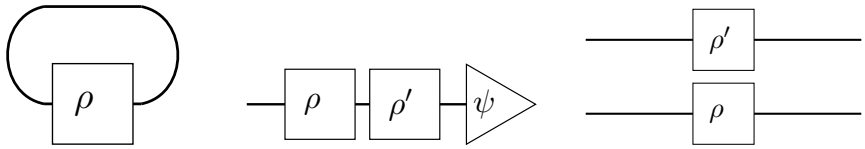

(a) $\mathrm{Tr}(\rho)$ (b) $\rho(\rho'(|\psi\rangle))$ (c) $\rho' \otimes \rho$

Figure 2.4: Tracing, multiplication and tensor product.

2.2.3 Stinespring Dilation

The first representation of CPTP maps is probably the most physically intuitive, although perhaps not as useful for the purposes of QCVV as the ones to be discussed later. It is an application of the Stinespring dilation theorem [37]. In essence Stinespring's theorem states that any CPTP channel can be built up from taking the original system, tensoring it with a second system in a second state, applying unitary transformations and then reducing back to the original subsystem. In particular this representation comes with a bound on the size of the ancillary system.

It goes like this, let $\mathcal{E} : S(\mathcal{H}) \to S(\mathcal{H})$ be our CPTP map between operators on a finite dimensional Hilbert space $\mathcal{H}$. Then there exists a Hilbert space $\mathcal{H}'$ and a unitary operation (U) acting on $\mathcal{H} \otimes \mathcal{H}'$, such that:

$$\mathcal{E}(\rho) = \mathrm{Tr}_{\mathcal{H}'}\left(U\left(\rho \otimes |0\rangle\langle 0|\right)U^\dagger\right), \tag{2.54}$$

where $|0\rangle\langle 0|$ is an arbitrary state in $\mathcal{H}'$. The advantage of this description is that more complicated CPTP maps can now be described by unitary operations (albeit on a larger Hilbert space), a process typically described as *purification*.

It turns out that we need only consider the ancillary system as being in a Hilbert space

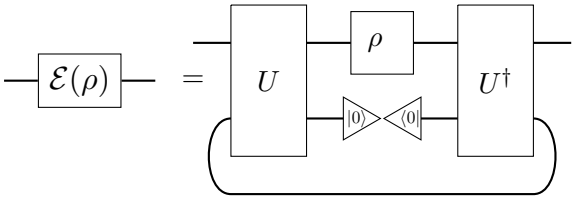

Figure 2.5: Stinespring dilation.

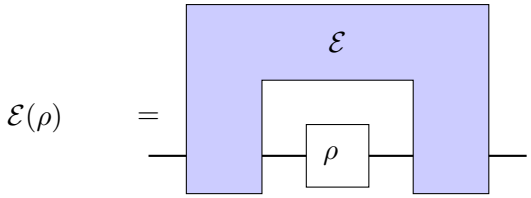

Figure 2.6: Generalised diagram of channel that maps a density operator to a new density operator, note the similarity to figure 2.5. The Stinespring dilation can be thought of as a refinement of the colored $\mathcal{E}$ block.

of no more than d^2 dimensions if our original system is of dimension d (see, for instance chapter 8 of [10]).

Using the graphical calculus shown earlier we can write our CPTP channel ($\mathcal{E}$) as shown in figure 2.5.

2.2.4 Superoperator/Liouville representation

The main representation used in this thesis is a variation of the superoperator/Liouville representation, being the Liouville representation in the Pauli basis.

However, to start with we will look at the representation in the standard (computational) basis.

The representation starts from the viewpoint that we can represent linear maps as matrices. To do this we need to represent a space of complex $d \times d$ matrices (being the operators) as vectors ($d^2 \times 1$ vectors). As a vector space $\mathcal{M}_{d,d'}(\mathbb{C})$ is isomorphic to $\mathbb{C}^{dd'}$. The Liouville representation is based on the vectorization of the density operator $\rho \to |\rho\rangle\rangle$. Given a set of basis vectors in the operator space then any density operator can be written as a sum over all of the basis vectors of a scalar times each basis vector. Consequently we can expand the density operator as $\rho = \sum_{k \in N_{d^2}} \langle A_k, \rho \rangle A_k$, where here $\{A_1..A_{d^2}\}$ represent orthonormal basis operators spanning the $d \times d$ space. Once we have done this we can identify ρ with a column vector $|\rho\rangle\rangle \in \mathbb{C}^{d^2}$ where the kth entry is the Hilbert-Schmidt inner product with the relevant basis operator i.e. $\langle A_k, \rho \rangle$. In the computational basis this merely involves column stacking the density matrix to form a

vector, e.g. for one qubit, we have:

$$\rho = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \equiv |\rho\rangle\rangle = \begin{pmatrix} a \\ c \\ b \\ d \end{pmatrix} \tag{2.55}$$

We can see why this is so by looking at the at the above ρ in the basis $\{|0\rangle\langle 0|, |1\rangle\langle 0|, |0\rangle\langle 1|, |1\rangle\langle 1|\}$ and taking the appropriate traces:

$$\begin{aligned} &\mathrm{Tr}(\rho|0\rangle\langle 0|) = \mathrm{Tr}\left(\rho \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}\right) = a \\ &\mathrm{Tr}(\rho|1\rangle\langle 0|) = c\,, \mathrm{Tr}(\rho|0\rangle\langle 1|) = b\,, \mathrm{Tr}(\rho|1\rangle\langle 1|) = d\,. \end{aligned} \tag{2.56}$$

The Liouville representation of a channel $\mathcal{E}$ is then the unique matrix $\mathcal{E} \in \mathbb{C}^{d^2 \times d^2}$ such that $\mathcal{E}|\rho\rangle\rangle = |\mathcal{E}(\rho)\rangle\rangle$. It will have entries

$$\mathcal{E}_{kl} = \langle A_k, \mathcal{E}(A_l)\rangle = \langle\langle A_k|\mathcal{E}|A_l\rangle\rangle\,. \tag{2.57}$$

There are a number of interesting and useful consequences that arise out of this representation (see [38] for proofs of relevant results)

- Channels compose by matrix multiplication, that is if we have two channels $\mathcal{E}_1$ and $\mathcal{E}_2$, then the concatenation of the maps $\mathcal{E}'(\rho) = \mathcal{E}_1(\mathcal{E}_2(\rho))$ corresponds to matrix multiplication $\mathcal{E}' = \mathcal{E}_1\mathcal{E}_2$ (where I have slightly abused the notation).
- Unitary operators can be embedded in channels, that is a unitary can be written as a channel which transforms a density operator ρ to the density operator ρ', where $\rho' = U\rho U^\dagger$. Typically such a channel is denoted by writing the unitary operator in calligraphic style, i.e. $\mathcal{U}$. It is constructed as $\mathcal{U}_{i,j} = \langle A_i, UA_jU^\dagger\rangle$, where $\{A_1..A_{d^2}\}$ represents the relevant basis vectors. Importantly it should be noted that applying such a channelized operator is done as $|\rho'\rangle\rangle = \mathcal{U}|\rho\rangle\rangle$. Later on we will see constructs such as $\mathcal{U}\mathcal{E}\mathcal{U}^\dagger$, this are quite different from evolution by a unitary operator — we will find such constructs useful for such things as calculating the fidelity of channels.
- Measurement operators and POVMs can be vectorised in exactly the same manner and applied straightforwardly to the resulting matrices. POVM elements are dual to the states and thus where we represent a state by a column vector, the POVM would be a row vector. An element E of a POVM can be expanded $E = \sum E_j A_j^\dagger$, where $E_j = \langle E, A_j^\dagger\rangle$. E can then be identified with the row vector, with its elements being each E_j where (as before) $\{A_1..A_{d^2}\}$ represent the chosen basis operators. Measurement then becomes $\langle\langle E|\mathcal{E}|\rho\rangle\rangle$.

Some insight can be obtained as to why superoperators have the properties described above by looking at how they are presented in a graphical calculus, but to do so we need to introduce a few more elements into our toolbox of diagrams. The unnormalized maximally entangled Bell State (see Equation 2.44), can be written as:

$$|\Phi^+\rangle = \sum_{i=0}^{d-1} |i\rangle \otimes |i\rangle \,, \tag{2.58}$$

Noting that we represent the sum as an extra index from a ket, in this case joining to the other ket as the sum is over the same index. Then we represent the Bell state in the graphical calculus[†] as a curve as shown:

$|\Phi^+\rangle = \sum_{i=0}^{d-1} |i\rangle \otimes |i\rangle =$ 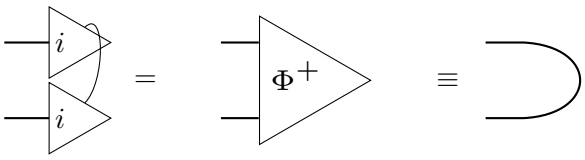

The vectorization we use here is a col-vec mapping (as opposed to row-vec mapping). We can see that the following gives us exactly the vectorization we want:

$$(\mathbb{1} \otimes \rho)|\Phi^+\rangle = |\rho\rangle\rangle \,, \tag{2.59}$$

which is clear when it is realised that the unnormalized Bell state $|\Phi^+\rangle$ is just the vectorised identity.

Putting this all together it can be seen the vectorised version of ρ, is represented as:

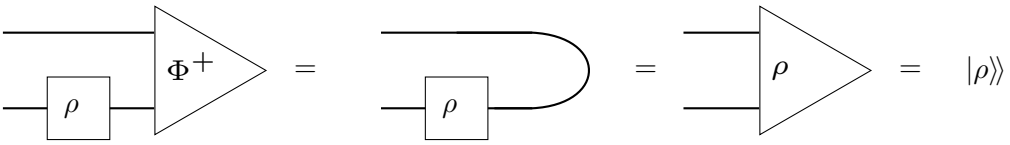

Remembering that we are looking for a representation of a mapping that takes a density operator to another density operator, we can write a generalized diagram of such a channel $\mathcal{E}$ in our graphical calculus as something like figure 2.6.

The graph can then be re-written (using the graphical calculus rules), by bending the right hand leg up and round, as follows (where I have coloured the channel for ease of identification):

[†]Note the convention here differs from that used in [35], where color is used to indicate extra indices being summed over, but is more in line with standard tensor network diagrams such as described by [36].

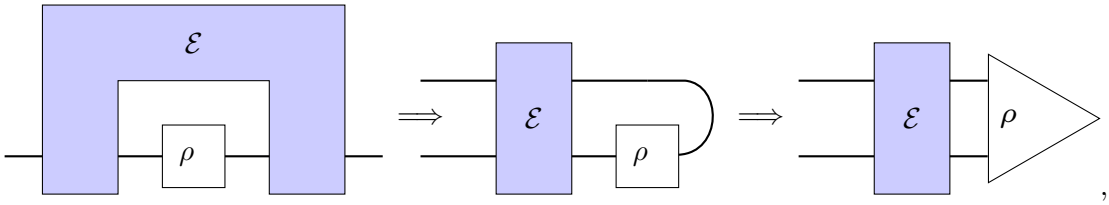

which is exactly what the Liouville representation is. It also becomes easy to see why, in this representation, the channels compose with simple matrix multiplication.

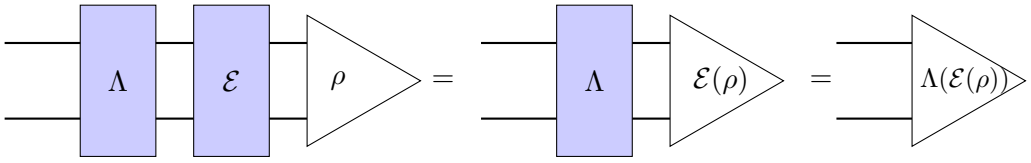

Measurement is equally obvious, with the POVM E being represented as a bra vector:

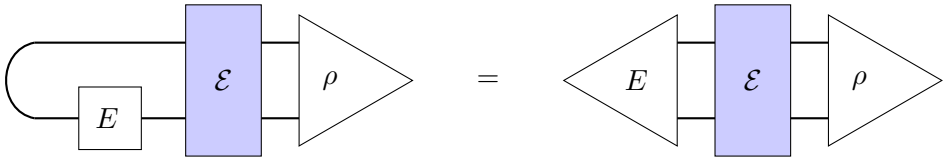

$\mathrm{Tr}((\mathbb{1} \otimes E)\mathcal{E}|\rho\rangle\rangle)$ $\qquad\qquad$ $\langle\langle E|\mathcal{E}|\rho\rangle\rangle$

2.2.5 Kraus Operators

The Kraus operators [39], or as the representation is referred to in [10], the operator-sum representation, is closely related to the Stinespring dilation. With the Stinespring dilation we expanded the size of the system (by tensoring it with a representation of the 'environment') so the quantum channel (that is the linear map) could be represented by unitary evolution on the expanded system. If we let $\{|e_i\rangle\}$ be an orthonormal basis for this finite dimension environment and $\mathbb{1}$ be the identity operator on the principal system (i.e. the one containing ρ), then as shown in [10] we can rewrite the Stinespring dilation (Equation 2.54) using Equation 2.46 as:

$$\mathcal{E}(\rho) = \mathrm{Tr}_{\mathrm{env}} \sum_i \left[\mathbb{1} \otimes \langle e_i|\right] U \left(\rho \otimes |e_0\rangle\langle e_0|\right) U^\dagger \left[\mathbb{1} \otimes |e_i\rangle\right] \tag{2.60}$$

$$= \sum_i \langle e_i|U \left(\rho \otimes |e_0\rangle\langle e_0|\right) U^\dagger |e_i\rangle \tag{2.61}$$

$$= \sum_i K_i \rho K_i^\dagger \,, \tag{2.62}$$

where K_i is defined to be $\langle e_i|U|e_0\rangle$ each such K_i representing operators on the state space of the principal system. Effectively this means that the Kraus operators are read from the left hand columns of the Stinespring unitary.

There are a number of useful implications that arise from the Kraus representation. The first is that a linear map, T, is completely positive iff it admits to a representation in the Kraus form (i.e. Equation 2.62).

The second is that T is trace preserving iff $\sum_j K_j^\dagger K_j = \mathbb{1}$.

The maximum number of Kraus operators needed for a Kraus description of T is equal to the dimensions of the vector spaces the linear map is acting on, for the case where they are equal and the principal dimension is d then no more than d^2 Kraus operators are required.

The graphical calculus representation of this is simple, if not particularly enlightening. A representation of a particular Kraus operator is just that of an operator. Multiple operators can be thought of as another index in the matrix representation, i.e. instead of a $d \times d$ matrix, we have a three-dimensional matrix, say $d \times d \times n$, where n represents the number of $d \times d$ matrices. In [35], this is represented by coloring in the representation. Here, as previously discussed, I prefer the normal tensor network representation of adding another leg onto the tensor in the diagram (a virtual leg). By joining two legs together this simply means we sum over the virtual legs they represent. The Kraus operators are shown as follows:

$$\sum_j K_j^\dagger K_j = \mathbb{1} \qquad \mathcal{E}(\rho) = \sum_j K_j \rho K_j^\dagger$$

Finally the Kraus representation is not unique as there is unitary freedom in choosing the operators. Whilst the Kraus operators are very useful in certain fields and are a particularly useful condition for the construction of completely positive maps, other than using them to help describe the Choi-Jamiolkowski isomorphism (discussed below) we will not use them much in the rest of this thesis.

2.2.6 Choi-Jamiolkowski Isomorphism

The Choi-Jamiolkowski isomorphism represents a bijection between linear maps and linear operators [40], in a similar way that vectorization of an operator represented a bijection between linear operators and vectors, that is $L(\mathcal{H}, \mathcal{H}) \equiv \mathcal{H} \otimes \mathcal{H}$. Although the isomorphism itself doesn't require a basis to be chosen, for our present purposes we will

choose the normal computational basis. Given this we can write the isomorphism as

$$\Lambda = \sum_{i=0,j=0}^{d-1} |i\rangle\langle j| \otimes \mathcal{E}(|i\rangle\langle j|) \tag{2.63}$$

Using the distributive properties of the tensor operation we note this is the same as

$$\Lambda = \sum_{i=0,j=0}^{d-1} (\mathcal{I} \otimes \mathcal{E})|i\rangle|i\rangle\langle j|\langle j| \tag{2.64}$$

$$= (\mathcal{I} \otimes \mathcal{E}) \sum_{i=0}^{d-1} |i\rangle|i\rangle \sum_{j=0}^{d-1} \langle j|\langle j| \tag{2.65}$$

$$= (\mathcal{I} \otimes \mathcal{E})|\Phi^+\rangle\langle\Phi^+|\,, \tag{2.66}$$

where the first identity uses Equation 2.15 (in reverse) and last one uses Equation 2.58.

We can use the graphical calculus used above to show, maybe more intuitively, how this comes about. Here we are going to write $\mathcal{E}(|i\rangle\langle j|)$ as the Kraus operators (see above) acting on the basis states. As before, because we are summing over the various basis states, they will be joined with a third line, representing summing over the virtual leg of the tensor. Finally, as before, the Kraus matrices and error channels will be colored, not representing anything special, but just to aid keeping track of them in the diagrams.

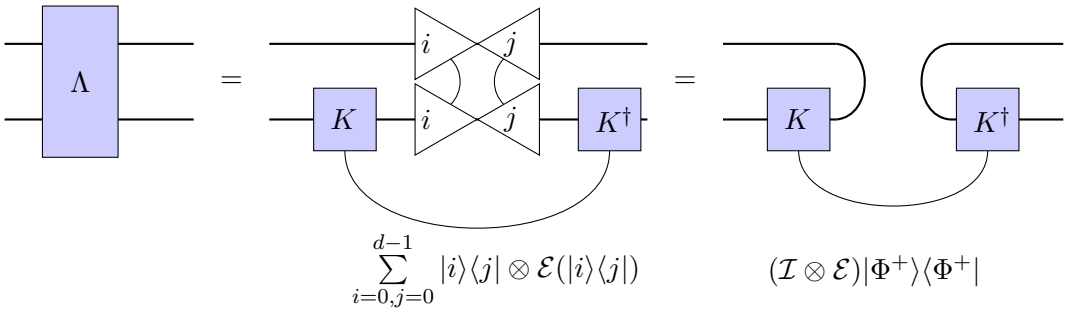

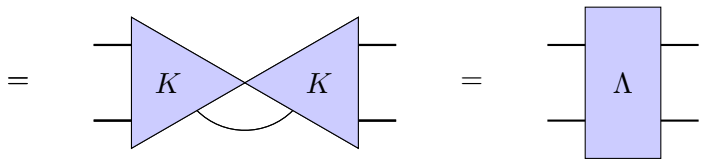

Following the steps outlined in [35] we can also use this graphical calculus and the Kraus operators to show how a quantum operator would change (or evolve under the noise map) using the Choi-matrix representation. That is how to calculate $\mathcal{E}(\rho)$ given Λ.

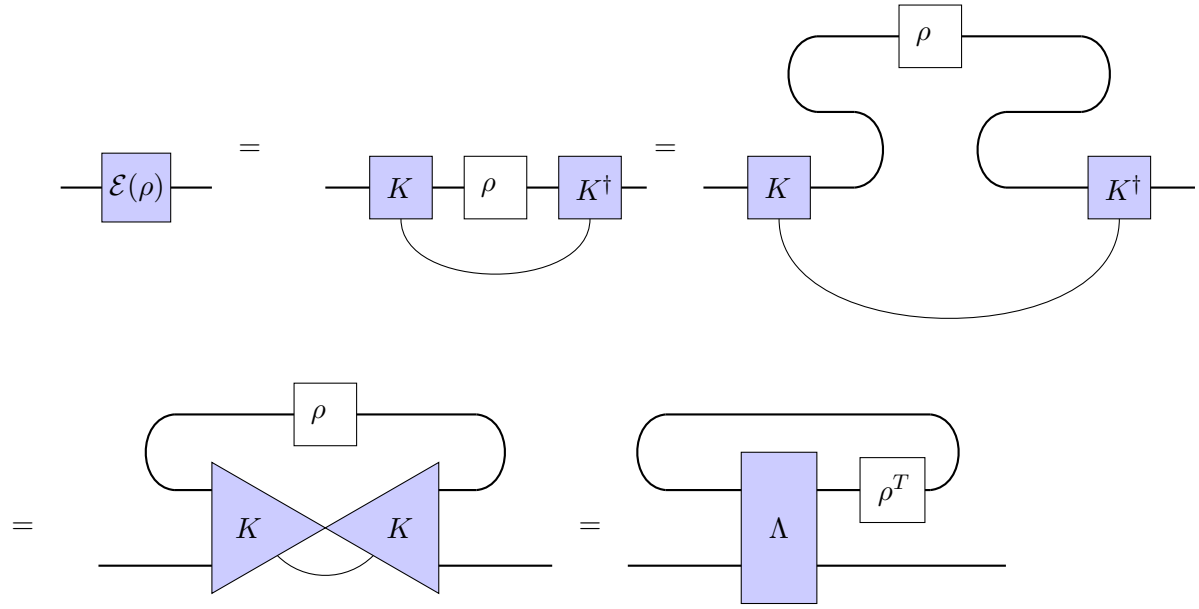

Reading off the picture, we can see that $\mathcal{E}(\rho) = \mathrm{Tr}_{\mathcal{X}}\left(\Lambda(\rho^T \otimes \mathbb{1}_{\mathcal{Y}})\right)$.

The state $J(\mathcal{E}) = \frac{1}{d}\Lambda(\mathcal{E})$ is commonly called the Jamiolkowski state associated with the Choi matrix Λ, which is why the isomorphism is commonly known as the Choi-Jamiolkowski isomorphism, although it should be noted that some authors choose a different normalization.

Related to this is the χ-matrix, or χ-representation.

If we assume that our linear map $\mathcal{E}$ has a Kraus decomposition $\{K_i\}$ and if $\{Q_j\}$ is a basis over the operators (i.e. a basis over $L(\mathcal{H}_1, \mathcal{H}_2)$) then the Kraus operators can be expanded in that basis as follows:

$$\mathcal{E}(\rho) = \sum_i K_i \rho K_i^\dagger = \sum_{i,j} \chi_{i,j} Q_i \rho Q_j^\dagger \,. \tag{2.67}$$

The matrix formed from the entries $\chi_{i,j}$ is known as the χ matrix with respect to that basis.

As is pointed out in [34] the χ matrix is equal to the Jamiolkowski state, written with respect to the basis $\{(Q_i \otimes \mathbb{1})|\Phi^+\rangle\}$. For the interested reader [35] shows how this can be represented using graphical calculus.

Where we are dealing with qubits then a useful basis for the χ matrix is the Pauli operator basis, where by convention we write $P_0 = \mathbb{1}$, with every other P_k being traceless. For P_0 this basis is $(P_0 \otimes \mathbb{1})|\Phi^+\rangle = |\Phi^+\rangle$. The (0,0) element of the χ matrix (written as $\chi_{0,0}$) using the Paulis as a basis is, therefore, $\langle\Phi^+|\Lambda|\Phi^+\rangle$, which is unitary invariant and proportional to the Choi representation of the identity channel. As we will see the

fidelity of the channel $\mathcal{E}$ is linearly related to the $\chi_{0,0}$ element (see subsection 2.4.6). One of the hallmarks of the χ matrix is that there exist experimental techniques for using a form of tomography to directly calculate the individual entries of the matrix [41].

2.2.7 Liouville representation in the Pauli basis

The Liouville representation (detailed earlier in subsection 2.2.4) can also be put into a the Pauli basis, which gives us the Pauli Liouville representation, which is the representation that will be used extensively throughout the rest of this thesis.

As previously mentioned we could choose any orthonormal basis $\mathbb{A} = \{A_0, A_1, \dots A_{d^2-1}\}$ to span the relevant $d \times d$ space. Here we will choose the Pauli operators (which fit this criteria) and as we did with the χ-matrix above, set $A_0 = \mathbb{1}/\sqrt{d}$. Doing this fixes the first operator as the one with trace, all the others being traceless. Where we are dealing with one qubit we will fix the others as the X, Y and Z Pauli operators respectively.

The vectorization of a density operator ρ is now based on the Hilbert-Schmidt inner product of the chosen basis with the operator i.e. $\rho_j = \langle A_j, \rho \rangle$, thus

$$|\rho\rangle\!\rangle = \begin{pmatrix} \rho_0 \\ \vec{\rho} \end{pmatrix}. \tag{2.68}$$

By way of example the 'up' state might be represented by $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$. The density operator is $\begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$. Vectorising this we get

$$\rho_0 = \frac{1}{\sqrt{2}} \mathrm{Tr}\left(\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right) = \frac{1}{\sqrt{2}} \tag{2.69}$$

$$\rho_1 = \frac{1}{\sqrt{2}} \mathrm{Tr}\left(\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right) = 0 \tag{2.70}$$

$$\rho_2 = \frac{1}{\sqrt{2}} \mathrm{Tr}\left(\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right) = 0 \tag{2.71}$$

$$\rho_3 = \frac{1}{\sqrt{2}} \mathrm{Tr}\left(\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right) = \frac{1}{\sqrt{2}} \tag{2.72}$$

$$\implies |\rho\rangle\!\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}. \tag{2.73}$$

The channel itself is a matrix, where each row and column represent a different Pauli

operator. So for example a channel with the same effect as an X gate (written $\mathcal{X}$) would look like:

$$\text{Pauli Liouville - X gate} = \mathcal{X} = \begin{array}{c} \\ \mathbb{1} \\ X \\ Y \\ Z \end{array} \begin{array}{c} \begin{array}{cccc} \mathbb{1} & X & Y & Z \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \end{array}, \tag{2.74}$$

which can be read as a channel that takes the identity to the identity, the X gate to the X gate, the Y gate to the -Y gate and the Z gate to the -Z gate. We can easily see that if we multiply Equation 2.73 with this matrix, we get the result we would expect (i.e. the vectorised version of the 'down' state).

There are a number of interesting facts about this representation:

For any CP channel $\mathcal{E}$, the Pauli Liouville matrix (which will have $d^2 \times d^2$ elements) takes the form [4]:

$$\mathcal{E} = \begin{pmatrix} \mathcal{S}(\mathcal{E}) & \mathcal{E}_{\text{sdl}} \\ \mathcal{E}_{\text{n}} & \mathcal{E}_{\text{u}} \end{pmatrix}, \tag{2.75}$$

where $\mathcal{E}_{\text{sdl}}$ is a $1 \times (d^2 - 1)$ vector, the *state-dependent leakage* block, $\mathcal{E}_{\text{n}}$ is a $(d^2 - 1) \times 1$ vector, the *nonunital* block and $\mathcal{E}_{\text{u}}$ is a $(d^2 - 1) \times (d^2 - 1)$ array, the *unital* block. $\mathcal{S}(\mathcal{E})$ represents the survival probability of the channel (that is the probability that after the channel has been applied that we will still be in the computational basis). Where we have a TP map, then $\mathcal{S}(\mathcal{E}) = 1$ and $\mathcal{E}_{\text{sdl}} = \vec{0}$. The meaning of $\mathcal{E}_{\text{n}}$ and $\mathcal{E}_{\text{u}}$ are explored more later (subsection 4.1.1), although it can be seen that the channel will only be unital if $\vec{\mathcal{E}_{\text{n}}} = 0$ (a unital channel is one where the identity is a fixed point of the channel). If the map were a trace decreasing map then $\mathcal{S}(\mathcal{E})$ will be less than 1.

Since the operator basis is Hermitian, all the entries of the Pauli Liouville matrix are real valued and the representation of the adjoint channel ($\mathcal{E}^\dagger$) is the transpose of the representation of $\mathcal{E}$.

The representation is, technically, a projective representation since any overall global phase is lost (such a global phase is not observable). This has the advantage of making it particularly easy to determine if two states or operators are identical (even if they are subject to a global phase) as, if we project them into this representation, their matrix elements will be identical, even if their global phase differs.

Single qubit trace preserving maps (such as stochastic maps) are also considered in [42] where they note that all trace-preserving maps can be written in the form:

$$\Phi(\rho) = U\Phi_D(V\rho V^\dagger)U^\dagger, \tag{2.76}$$

where U, V are unitary and

$$\Phi_D = \begin{pmatrix} 1 & 0 & 0 & 0 \\ t_1 & \lambda_1 & 0 & 0 \\ t_2 & 0 & \lambda_2 & 0 \\ t_3 & 0 & 0 & \lambda_3 \end{pmatrix}. \tag{2.77}$$

In the special case of unital maps $t_1, t_2, t_3 = 0$ and if such a unital map is completely positive the eigenvalues (which are not necessarily real) will satisfy:

$$|\lambda_1 \pm \lambda_2| \leq |1 \pm \lambda_3|. \tag{2.78}$$

This issue (as well as what gauge invariant information can be extracted from the superoperator describing a single qubit quantum channel) was further examined in [43], where the analysis was extended to encompass single qubit non-unital maps as well as unital maps not in the form of Φ_D.

2.3 Unitary Designs

2.3.1 What is a unitary t-design?

To answer this, we must first understand what the *Haar measure* is, specifically when applied to quantum systems. The Haar measure is the unique group-invariant measure over locally compact topological groups. Restricting it to the case relevant here it is the unique measure over all unitary operations acting on n qubits that is invariant under left or right multiplication by any unitary operation. Relevantly, it can be used to average out properties that are not unitarily invariant over all unitaries.

In a d dimensional quantum system, the Hilbert space is d dimensional and so if we are considering quantum pure states, the natural group is SU(d), which is the special unitary group of dimension d. With a qubit the relevant group is SU(2), that is, using the matrix representation previously discussed, all 2×2 unitary operators. Each of these is a rotation of the Bloch Sphere (figure 2.1). The Haar measure is thus invariant under all rotations of the Bloch sphere — meaning that it is *the* rotationally invariant measure of the Bloch sphere. One way of thinking of this is as a constant density distribution over the surface of the sphere.

A unitary t-design, or quantum t-design, is a distribution over unitary operators that can duplicate the properties of the full Haar measure for all polynomial functions of degree t or less. The definition of unitary t-design in [24] states that for every polynomial $P_{(t,t)}(U)$ of degree at most t in the matrix elements of U and at most t in the complex conjugates

of those matrix elements:

$$\frac{1}{K}\sum_{k=1}^{K} P_{(t,t)}(U_k) = \int_{\mathcal{U}(D)} dU\, P_{(t,t)}(U)\,. \tag{2.79}$$

2.3.2 Unitary 2-Designs

A unitary 2-design is, therefore, a unitary t-design where $t = 2$.

Cleve et al. [31] give an intuitive example of one of the consequences of a unitary 2-design, namely that of *two-query indistinguishability* from Haar-random unitaries.

> Imagine a game where, at the flip of a coin, U is sampled either according to the Haar measure or with respect to the unitary 2-design. A two-query distinguishing procedure can make two queries to U (each being either in the forward direction as U or in the reverse direction as $U^\dagger$) as well as other quantum operations that do not depend on U and then outputs a bit. A unitary 2-design has the property that no two-query distinguishing procedure can distinguish between the Haar-random case and the 2-design case with probability greater than 1/2.

It should be noted that the Clifford gates (discussed in Equation 2.1.8) are a unitary 2-design (indeed they fall just short of a unitary 4-design [44]), a fact we will make use of later.

Twirling

Related (and indeed proven in [31] to be identical) to the definition above is the fact that we can use a unitary 2-design to obtain the same results as twirling through all possible unitaries.

If we have a channel $\mathcal{E}$ acting on a state ρ, then the *twirl* over all unitaries (drawn from the Haar measure) is as follows:

$$\int_{U(d)} (U\mathcal{E}(U^\dagger \rho U)U^\dagger)\, dU\,. \tag{2.80}$$

As will be discussed in subsection 2.4.2 this is how we can measure the average fidelity of a channel. The impact of the unitary 2-design is that it can be used to reduce the number unitaries we need to twirl over. Because the integrand in Equation 2.80 is a 2^{nd}

degree polynomial in $U \otimes U$, we can replace the integral over all unitaries U with a sum over any 2-design and we have:

$$\sum_{j=1}^{K} q_j (U_j \mathcal{E}(U_j^\dagger \rho U_j) U_j^\dagger) = \int_{U(d)} (U \mathcal{E}(U^\dagger \rho U) U^\dagger) \, dU \,, \tag{2.81}$$

where the sum is over the group elements of our unitary 2-design, and q_j the associated probability distribution (weight) relating to the relevant group element.

Identifying a 2-design

So given that we have a use for a unitary 2-design, how do we check that the proposed group of unitaries form one?

In [45], Gross et al. explore these and other related questions. Initially they note that the problem is analogous to the classical problems relating to the structure on spheres in $\mathbb{R}^n$. Adapting the definition of spherical designs to complex vector spaces they note that 'unitary designs aim to identify finite nets of unitaries, which cover the entire group as tightly as possible'.

Using this insight they produce a simple method of checking whether a set of unitaries form a unitary 2-design. They define the *frame potential* of a set of unitaries ($\mathcal{D} = \{U_k\}, k = 1..K$) as:

$$\mathcal{P}(\mathcal{D}) = \frac{1}{K^2} \sum_{U_k, U_{k'}} |\mathrm{Tr} U_k^\dagger U_{k'}|^4 \,. \tag{2.82}$$

Theorem 1 of [45] informs us that $\mathcal{D}$ is a 2-design (or greater) iff $\mathcal{P}(\mathcal{D}) = 2$.

2.3.3 Representation Theory, Schur's lemma, 2-Designs and the Pauli-Liouville Representation

We can use representation theory to gain insight as to the effect of a unitary 2-twirl (or where we are using the Clifford gates as our unitary 2-design, the Clifford twirl).

Representation theory tells us that a finite group (here the elements of the unitary 2-design) $\mathfrak{G}$ can be thought of as a pair (ϕ, V), where V is a vector space (here this may be a real or complex representation space) and ϕ represents a group homomorphism, that maps the group elements to the general linear group over V i.e. $\phi : \mathfrak{G} \to \mathrm{GL(V)}$. Because the mapping is a homomorphism it preserves the group structure; here this simply means that the mapping of the product of two group elements (say g_1 and g_2) is the same as the composition of the mappings of each group element i.e. $\phi(g_1 g_2) = \phi(g_1)\phi(g_2), \forall g \in \mathfrak{G}$. A *faithful* representation is one where the mapping ϕ is injective, that is it maps each

group element to a unique element in GL(V). A representation does not need to be faithful to be a valid representation. For instance, the trivial representation just maps every group element to 1. This still preserves group structure because if $\phi(g_1) = 1$ and $\phi(g_2) = 1$, then $\phi(g_1)\phi(g_2) = 1 = \phi(g_1 g_2)$.

The dimension of a representation is the dimension of V. A *sub-representation* (ϕ_U, U) of the representation (ϕ, V) consists of the homomorphism ϕ_U and a subspace U, where $U \subset V$ and $\phi_U(g)$ preserves U, ϕ_U denoting the restriction of ϕ to the subspace U. An alternative way of expressing this is that $\phi_U(g)u \in U, \forall g \in \mathfrak{G}$ and $u \in U$. From now on I will write a representation (ϕ, V) as simply V where the mapping ϕ can safely be left implicit. A representation is *unitary* if $\phi(g)$ is a unitary operator for all $g \in \mathfrak{G}$.

A representation is *irreducible* or an *irrep* if the only sub-representations are $\emptyset$ and V. All representations of finite groups are semi-simple (*Mashcke's Theorem*), this means they can be composed by the direct sum of irreps (which may occur with some multiplicity). A corollary of this is that if U is a sub-representation of V, the orthogonal complement $U^\perp$ is a sub-representation as well. In this case we have that $V = U \oplus U^\perp$. The final definition we need is of an *intertwining operator*. If we have two representations, say (ϕ, V) and (σ, W) of the same group, here $\mathfrak{G}$. Then an intertwining operator is a linear map $(T : V \to W)$ that has the following property:

$$T(\phi(g)v) = \sigma(g)(T(v)), \forall g \in \mathfrak{G}, v \in V \tag{2.83}$$

The next step is to look at the twirling operator. Given a representation (ϕ, V) for our group $\mathfrak{G}$ and a matrix $A \in \mathrm{GL}(V)$, we define an action (the *twirl*)

$$A^g = \phi(g)\, A\, \phi(g^{-1})\,, \tag{2.84}$$

where g^{-1} is the inverse of g. We can do this for each element of $\mathfrak{G}$, which gives us the uniform average of this action defined as:

$$A^G = \frac{1}{|\mathfrak{G}|} \sum_{g \in \mathfrak{G}} A^g, \tag{2.85}$$

where $|\mathfrak{G}|$ is the number of elements in the group. The important thing to notice here is that A^G commutes with the action of $\mathfrak{G}$ for any representation V (whether reducible or not), that is for all $g_x \in \mathfrak{G}$ and $g \in \mathfrak{G}$:

$$\begin{aligned}
(A^{g_x})^G &= \frac{1}{|\mathfrak{G}|} \sum_{g \in \mathfrak{G}} g(g_x A g_x^{-1}) g^{-1} \\
&= \frac{1}{|\mathfrak{G}|} \sum_{g \in \mathfrak{G}} g_x(g_x^{-1} g g_x) A (g_x^{-1} g^{-1} g_x) g_x^{-1} \\
&= \frac{1}{|\mathfrak{G}|} \sum_{g_y \in \mathfrak{G}} g_x(g_y) A (g_y^{-1}) g_x^{-1} \\
&= (A^G)^{g_x}\,,
\end{aligned}$$

where for notational convenience I have just written $\phi(g)$ as g and we use the properties of the group to rewrite $g_x^{-1} g g_x$ as g_y.

Now we are ready to use Schur's Lemma, for algebraically closed fields. This provides that if V is a finite dimensional irreducible representation and ϕ is an intertwining operator, then $\phi = \lambda \mathbb{1}$ for some scalar λ. In this case the action A^G commutes with all representations and therefore is an intertwining operator.

Next we hunt for irreducible representations.

From subsection 2.2.4 we know that the superoperator representation (Liouville representation) is a homomorphism $U \to \mathcal{U}$. One of the proofs in [45] tells us that, where U is a unitary representation of a unitary 2-design, then the homomorphism has exactly two irreducible representations.

As discussed in subsection 2.2.7, the Pauli Liouville representation uses the Paulis as a basis for the Liouville superoperator and, by convention, we use the identity (the only basis element with trace non-zero) as the first basis vector. It is clear that the identity is a trivial irreducible representation of the group. As discussed above Mashcke's Theorem tells us that the orthogonal complement is another representation of the group and from Gross's paper we know that in this homomorphism (the superoperator representation) a unitary 2-design only has two representations, both of which are irreducible. If we call the second irreducible representation $(\psi, \mathbb{C}^{(d^2-1)})$, then we can write our group as the direct sum $\mathbb{1} \oplus \psi$.

Pulling this all together, we have an action — the twirl — which commutes with any representation, we have two irreducible representations and we have Schur's Lemma that tells us the action of the twirl on an irreducible representation results in a matrix that is scalar. This means that if we twirl any quantum channel represented by a Pauli Liouville superoperator we end up with a matrix that looks like:

$$A^G(\mathcal{E}) = \frac{1}{|\mathfrak{G}|} \sum_{g \in \mathfrak{G}} A^g(\mathcal{E}) = \frac{1}{|\mathfrak{G}|} \sum_{g \in \mathfrak{G}} g(\mathcal{E}) g^{-1} = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda_2 \end{pmatrix}. \tag{2.86}$$

The beauty of this representation is that it now becomes clear that a unitary 2-twirl (the action A^G) will allow us to transform channels into a quantum depolarizing channel. The ramifications of this will be discussed later (subsection 2.4.3).

2.4 Fidelity

2.4.1 What is the fidelity of a state?

This question relates to finding a metric to determine in some sense how close two different quantum states are to each other.

For a single qubit, for instance, one could imagine that there is a natural metric that captures the sense of the distance between two states on the Bloch sphere. This is the trace distance $D(\rho, \sigma) = \frac{1}{2}\mathrm{Tr}(|\rho - \sigma|)$.

The fidelity is slightly different and is not itself a metric (for example the fidelity of two identical states is 1, whereas a *metric* requires it to be zero). It is, however, a useful, natural, concept and can be related to an appropriate metric.

Schumacher [46] provides a definition for the fidelity of pure states as $F(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = |\langle\phi|\psi\rangle|^2$. Which generalizes in an obvious way when one, but only one, of the states is not pure,

$$\mathcal{F}\left(|\psi\rangle\langle\psi|, \rho\right) = \langle\psi|\rho|\psi\rangle\,. \tag{2.87}$$

(This follows from the cyclic properties of the trace.)

Jozsa [47] uses Uhlmann's theorem to generalise this for the case where both states are impure, to obtain the following definition for Fidelity:

$$\mathcal{F}(\rho, \sigma) \equiv \mathrm{Tr}\left(\sqrt{\rho^{\frac{1}{2}}\sigma\rho^{\frac{1}{2}}}\right)^2\,. \tag{2.88}$$

Importantly, as he also shows, it has all the properties one would expect, namely:

- $0 \leq F(\rho_1, \rho_2) \leq 1, F(\rho_1, \rho_2) = 1 \leftrightarrow \rho_1 = \rho_2$.
- $F(\rho_1, \rho_2) = F(\rho_2, \rho_1)$.
- If one of the states is a pure state, $|\psi\rangle\langle\psi|$, the fidelity is the same as Equation 2.87.
- $F(\rho_1, \rho_2)$ is invariant under unitary transforms.
- If $\rho_1, \rho_2, p_1, p_2 \geq 0, p_1 + p_2 = 1$ then $F(\rho, p_1\rho_1 + p_2\rho_2) \geq p_1 F(\rho, \rho_1) + p_2 F(\rho, \rho_2)$.
- $F(\rho_1, \rho_2) \geq \mathrm{Tr}(\rho_1\rho_2)$.
- $F(\rho_1 \otimes \rho_2, \rho_3 \otimes \rho_4) = F(\rho_1, \rho_3)F(\rho_2, \rho_4)$.

2.4.2 What is the fidelity of a channel?

As with states the question relates to finding a way of determining how close one channel is to another. Typically when we talk about the fidelity of a channel we mean how close is that channel to the identity channel. In this sense the average fidelity of a quantum channel described by a CPTP quantum operation is given in [10] as

$$\bar{\mathcal{F}}(\mathcal{E}) = \int d\psi \, \langle\psi|\mathcal{E}(\psi)|\psi\rangle = \int d\psi \, \mathrm{Tr}(|\psi\rangle\langle\psi|\mathcal{E}(\psi)) \,, \tag{2.89}$$

where the integral is over the uniform Haar measure, normalized so that $\int d\psi = 1$.

This can be extended to give us a fidelity estimate as to how closely $\mathcal{E}(\psi)$ approximates a quantum gate, U,

$$\bar{\mathcal{F}}(\mathcal{E}, U) = \int d\psi \, \langle\psi|U^\dagger\mathcal{E}(\psi)U|\psi\rangle = \int d\psi \, \mathrm{Tr}\left(U|\psi\rangle\langle\psi|U^\dagger\mathcal{E}(\psi)\right) . \tag{2.90}$$

One way of viewing this is as the average overlap between (i) all states ψ operated on by U and (ii) all states ψ transformed by $\mathcal{E}$.

As noted in subsection 2.3.1, since $d\psi$ is over the Haar measure and since $\bar{\mathcal{F}}(\mathcal{E})$ is at most of degree 2 in the elements of ψ, then an average over a unitary 2-design gives us exactly the same fidelity as the integral. Consequently, if $\mathfrak{G}$ is a unitary 2-design, with equal group element weighting:

$$\bar{\mathcal{F}}(\mathcal{E}, U) = \int d\psi \, \langle\psi|U^\dagger\mathcal{E}(\psi)U|\psi\rangle = \frac{1}{|\mathfrak{G}|}\sum_{\psi\in\mathfrak{G}} \langle\psi|U^\dagger\mathcal{E}(\psi)U|\psi\rangle \,. \tag{2.91}$$

2.4.3 The depolarizing channel

The depolarizing channel is a type of quantum channel that is particularly relevant to randomized benchmarking. The easiest way to imagine it is as a channel that leaves the state untouched with probability p and that depolarizes it, that is replaces it with the completely mixed state with probability $1 - p$.[‡]

$$\mathcal{E}(\rho) = (1-p)\frac{\mathbb{1}}{d} + p\rho \,. \tag{2.92}$$

For a single qubit the depolarizing channel serves to uniformly shrink the Bloch sphere, contracting it to the middle. Intuitively once can see this by recalling that $\mathbb{1}/2$ is the

[‡]Note conventions differ here, often p (rather than $1-p$) will represent the probability of the channel completely depolarizing the state. I have adopted the convention set out in Equation 2.92 since it better matches the convention used with randomized benchmarking (section 3.1.1).

center of the Bloch sphere. The mixed state caused by the channel will lie on a line from the original state ρ to the middle of the sphere, the position determined by the relative probabilities of the constituent states, i.e. dependent on p. Applied to all possible states, this is the same as causing the sphere to universally shrink.

It can readily be seen that repeated applications of a depolarizing channel will increase the probability that the initial state will be transformed into the completely mixed state. As we will see this is one of the key concepts behind randomized benchmarking and its variants.

2.4.4 Relating average fidelity to a depolarizing channel

Nielsen [33] expanded on work by Horodecki et al. [48] to provide a simple way to determine the average fidelity of a quantum channel.

The first of the two relevant steps in that paper is to show that twirling (see section 2.3.2) doesn't change the fidelity of the channel being twirled. The second is to show that twirling changes the channel into a depolarizing channel. At that point it is simply a matter of calculating the fidelity of the resulting depolarizing channel to determine the fidelity of the channel of interest.

The proof of the first step, where $\mathcal{E}_T$ is the twirled channel and $\mathcal{E}$ is the original channel, is as follows:

We previously had:

$$\mathcal{E}_T \equiv \int dU\, U^\dagger \mathcal{E}(U\rho U^\dagger)U \tag{2.93}$$

$$\bar{\mathcal{F}}(\mathcal{E}) = \int d\psi\, \langle\psi|\mathcal{E}(\psi)|\psi\rangle\,. \tag{2.94}$$

Where ψ is the density operator corresponding to $|\psi\rangle\langle\psi|$ and the integral in Equation 2.94 is over the uniform (Haar) measure $d\psi$ (although as discussed in subsection 2.3.1 the same results are achieved if we use the averaged sum of a unitary 2-design).

Combining the above we see:

$$\bar{\mathcal{F}}(\mathcal{E}_T) = \int d\psi \int dU \langle\psi|U^\dagger \mathcal{E}(U\psi U^\dagger)U|\psi\rangle \tag{2.95}$$

$$= \int dU \int d\psi \langle\psi|(U^\dagger)\mathcal{E}(U\psi U^\dagger)U|\psi\rangle \tag{2.96}$$

$$= \int dU \bar{\mathcal{F}}(\mathcal{E}) = \bar{\mathcal{F}}(\mathcal{E})\,. \tag{2.97}$$

Where the last step was accomplished by a change of variables $|\psi'\rangle \equiv U|\psi\rangle$.

The next stage is to show that $\mathcal{E}_T$ in fact creates a depolarizing channel (and thus *the* depolarizing channel with the same average fidelity as $\mathcal{E}$).

Although this has already been shown, through the implications of the results discussed in subsection 2.3.3, I will reproduce the arguments in [33] as they are nice and concise.

He first notes that for any unitary V

$$V\mathcal{E}_T(\rho)V^\dagger = \int dU\, VU^\dagger \mathcal{E}(U\rho U^\dagger)UV^\dagger = \mathcal{E}_T(V\rho V^\dagger)\,, \tag{2.98}$$

where the last step was obtained by change of variables $W \equiv UV^\dagger$.

Then let P be a one dimensional projector with $Q \equiv (1-P)$. Let V be block diagonal with respect to the spaces projected onto by P and Q, so that $VPV^\dagger = P$ and thus $V\mathcal{E}_T V^\dagger = \mathcal{E}_T(P)$.

The argument continues that it follows $\mathcal{E}_T(P) = \alpha P + \beta Q$ for some α and β, which can be rewritten as $(1-p)\frac{\mathbb{1}}{d} + pP$ for some p. Using Equation 2.98 it is clear that p must be the same for any choice of P and so we have by linearity of $\mathcal{E}_T$:

$$\mathcal{D}_p(\rho) = \mathcal{E}_T(\rho) = (1-p)\frac{\mathbb{1}}{d} + p\rho\,, \tag{2.99}$$

as required.

Finally we relate the depolarizing parameter p to the average fidelity as:

$$\bar{\mathcal{F}}(\mathcal{D}_p) = \frac{(d-1)p+1}{d}\,. \tag{2.100}$$

2.4.5 The depolarizing channel in the Pauli-Liouville basis

Equation 2.86 showed that in the Pauli Liouville basis a twirl by a unitary 2-design would result in a channel that looks like:

$$\begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda_2 \end{pmatrix}. \tag{2.101}$$

As previously discussed for a TP map $\lambda_1 = 1$. As shown above a twirled channel is in fact a depolarizing channel. Informally this is readily apparent from the facts that: (i) in this representation channels compose by matrix multiplication; (ii) λ_2 is less than 1 (or else the channel is the identity); and (iii) repeated applications of this channel will result in a power-up of the diagonal matrix, thus any state will move towards the

	$\bar{\mathcal{F}}$	r	p	$\chi_{0,0}$
$\bar{\mathcal{F}}$	$\bar{\mathcal{F}}$	$1-r$	$\frac{(d-1)p+1}{d}$	$\frac{d\chi_{0,0}+1}{d+1}$
r	$1-\bar{\mathcal{F}}$	r	$\frac{d-1}{d}(1-p)$	$\frac{d}{d+1}(1-\chi_{0,0})$
p	$\frac{d\bar{\mathcal{F}}-1}{d-1}$	$1-\frac{d}{d-1}r$	p	$\frac{d^2\chi_{0,0}-1}{d^2-1}$
$\chi_{0,0}$	$\frac{(d+1)\bar{\mathcal{F}}-1}{d}$	$1-\frac{d+1}{d}r$	$\frac{(d^2-1)p+1}{d^2}$	$\chi_{0,0}$

Table 2.1: Reproduction of a table appearing in [49], setting out the linear relationships between average fidelity $\bar{\mathcal{F}}$, infidelity (r), the depolarizing factor (p) and $\chi_{0,0}$

maximally mixed state (which is $\begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix}^T$ in this basis). Note that λ_2 is equal to the average of the diagonal elements of the $\mathcal{E}_\mathrm{u}$ block (defined in subsection 2.2.7).

More formally, if, as before, we use the normalized Paulis as our orthonormal unit basis of $\mathbb{C}^d$, $\{A_1, \ldots, A_{d^2}\}$, then we can write:

$$\int_{U(d)} dU\, \mathcal{U}^\dagger \mathcal{E} \mathcal{U} = \lambda_1 |A_1\rangle\!\rangle\langle\!\langle A_1| + \lambda_2 (\mathbb{1}_{d^2} - |A_1\rangle\!\rangle\langle\!\langle A_1|)\,, \tag{2.102}$$

as $|A_1\rangle\!\rangle$ carries the trivial representation and therefore $(\mathbb{1}_{d^2} - |A_1\rangle\!\rangle\langle\!\langle A_1|)$ is the other irreducible representation.

The depolarizing parameter p in equation Equation 2.99 can be read directly from this matrix, in that $p = \lambda_2$. Note that as previously mentioned it is a slightly unfortunate convention that in the depolarizing formula Equation 2.99 p and $(1-p)$ are sometimes written so that low p gives high fidelity (which would make it inversely related to the value of λ_2). As previously noted I will refer to p as the depolarizing *parameter* or *factor* in the sense set out above, where $p{=}\lambda_2$ and a high value of p relates to high fidelity. Substituting this into Equation 2.100 we see that:

$$\bar{\mathcal{F}}(\mathcal{E}) = \frac{\lambda_2(d-1)+1}{d}\,. \tag{2.103}$$

2.4.6 Relating fidelity to the χ Matrix

As will be recalled the χ matrix is the Choi matrix, in a particular basis, namely the Pauli operator basis. Using the Pauli basis $\chi_{0,0}$ is the element of the χ matrix corresponding to $\langle\Phi^+|\Lambda|\Phi^+\rangle$, where Λ is the Choi matrix, and $|\Phi^+\rangle$ one half of the Bell state.

Ref. [34] shows how fidelity is related to $\chi_{0,0}$, namely:

$$\bar{\mathcal{F}} = \frac{d\chi_{0,0} + 1}{d + 1} \,. \tag{2.104}$$

Finally we note that all of these methods of calculating fidelity are, of course, linearly related. Ref. [49] sets out the linear relationships between average fidelity, infidelity, the depolarizing factor and $\chi_{0,0}$ in a useful table, which is reproduced here (table 2.1) for convenience of reference.

Historically there has also been work relating quantum gate fidelity to the Choi matrix, see for example [50].

Part II

Randomized Benchmarking

Chapter 3

The Randomized Benchmarking Protocol

> "The most merciful thing in the world, I think, is the inability of the human mind to correlate all its contents. We live on a placid island of ignorance in the midst of black seas of infinity, and it was not meant that we should voyage far."
>
> H.P. Lovecraft

3.1 Basic Protocol

3.1.1 Outline of the protocol

Introduction

Randomized benchmarking is, essentially, an experimental protocol designed with the aim of turning a noise channel into a depolarizing channel with the same average gate fidelity. With a depolarizing channel it is comparatively easy to measure the decay rate of the channel which allows the fidelity of the depolarizing channel and, hence, the original noise channel to be measured.

In a quantum experimental setting we typically identify three points of error. There could be errors in preparing the initial state, errors in the control of the state and finally errors in the measurement of the state. The state preparation and measurement errors are typically referred to as *SPAM* errors. Randomized benchmarking is designed to give a figure of merit relating to the control of the state (i.e. not the SPAM errors) and is designed to do so in a way that is robust to (not affected by) the SPAM errors.

Another key aim of randomized benchmarking is to allow small errors to be measured. One could imagine that if, for instance, a gate was accurate 99.5% of the time then it would require a large amount of data to be gathered on the performance of that

specific gate before there was sufficient data to allow the creation of any reliable or useful statistics relating to the gate. Randomized benchmarking aims to amplify such small errors with a view to providing a method that is scalable and yet also able to quantify such errors.

The basic idea behind the protocol stems from the properties of a depolarizing channel. As previously discussed a depolarizing channel, with a probability related to the depolarizing parameter, will replace a state with the maximally mixed state. With many applications of a depolarizing channel the probability of the state surviving (that is not being replaced with the maximally mixed state) decays exponentially. The statistics gathered from such a channel will be identical (and therefore indistinguishable) from a channel that shifts a starting state towards the maximally mixed state at a rate per application of the channel dependent on its depolarizing parameter (for a single qubit system this is like moving the state from its preparation state towards the middle of the Bloch sphere at a rate dependent on the depolarizing parameter). As discussed in subsection 2.4.3 the density operator concept of an ensemble of states is another way to visualize the effect of a depolarizing channel.

Overview

If one starts in a particular state and then measures for that state (that is, make a measurement with maximum overlap with the initial state you had hoped to prepare), a failure to measure the state would relate only to the SPAM errors. Now if between preparation and measurement one applies a depolarizing channel, say, 10 then 20 then 30 times the chance of a successful measurement will decrease as the state becomes depolarized by the channel. The SPAM errors will (more or less) remain the same, with the *decrease* of successful measurement being attributable solely to the depolarizing channel. By applying successively longer sequences of the depolarizing channel, and noting the decay in the percentage of successful measurements, we can determine the depolarizing parameter of the channel in a way that is independent of the initial SPAM errors. This decrease of the rate of successful measurements (the *decay rate*) will be measurable even for a small depolarizing parameter (since we can just increase the number of times the channel is applied). With the decay rate and thus the depolarizing parameter of the depolarizing channel noted we can calculate the fidelity of the channel.

What we then need is to identify the noise we want to measure and turn it into a depolarizing channel. For that we will use unitary 2-designs — introduced earlier (subsection 2.3.1).

The Protocol

The standard randomized benchmarking protocol is as follows [25].

1. Choose a positive integer m.

2. Choose a random sequence s so that $\vec{G}_s = (G_1, \ldots, G_m) \in \mathfrak{G}^m$ of gates uniformly from $\mathfrak{G}$ and set $G_{m+1} = (G_m \ldots G_1)^\dagger$.

3. Prepare a d-dimensional system in some state ρ, which should be close to an ideal computational basis state.

4. At each time step $t = 1, \ldots, m+1$, apply $\mathcal{G}_t$.

5. Perform a POVM $\{E, I_d - E\}$ for some E, which should be close to a projection in the computational basis and return a 1 if E is observed and 0 otherwise.

6. Repeat steps 3–5 l_m times to obtain an *estimate* of the survival probability for the sequence s so that $\hat{F}_s$ is an estimate of the probability $F_s = \Pr(E|s, \rho)$.

7. Repeat steps 2–6 n times to obtain an *estimate* of the average survival probability over all sequences of length m, $\hat{\mathbb{E}}_s(F_s)$.

8. Repeat steps 1–7 for different values of m and fit to the model

$$\mathbb{E}_s(F_{m,s}) = A_{(\rho,E)}p^m + B_E \tag{3.1}$$

where p is related to the average gate fidelity f by

$$p = \frac{df - 1}{d - 1} \tag{3.2}$$

and $A_{(\rho,E)}$ and B_E are constants that absorb the SPAM error.

The choice of ρ and E are arbitrary, however, the contrast will be proportional to $\mathrm{Tr}(\rho E)$ (which is why they should be close to computational basis states).

Why the protocol should be thus, what the values of m, l_m and n should be and how to analyze the results obtained, is dealt with below.

3.1.2 Gate independent analysis – Magesan

The first full analysis of randomized benchmarking in its current form and for multiple qubits is Magesan, Gambetta and Emerson's paper [51] which expanded on the ground work laid out in [25]. An earlier paper by Knill [23] set out a similar protocol, but without the analysis which will be discussed below.

The protocol was introduced as a way of benchmarking the average fidelity of quantum gates that implement Clifford operators. The Clifford gates have been discussed in Equation 2.1.8, but to recap the Clifford gates themselves can be simulated on classical computers and therefore can be uniformly sampled and their inverses can be found in a scalable manner. They are easily extended to allow for universal computation either through the introduction of one additional gate (such as a $\frac{\pi}{8}$ gate) or through the use

of ancillas and measurements [10]. Finally the Cliffords form the basis of most fault-tolerant schemes and stabilizer codes, so being able to characterize the accuracy with which they can be implemented is an important problem.

The analysis is as follows.

For any unitary operator, such as a quantum gate (C), one can represent the operation of that operator (that is, $\rho' = C\rho C^\dagger$) as a channel, e.g. in the Pauli Liouville basis $|\rho'\rangle\rangle = \mathcal{C}|\rho\rangle\rangle$.

For an imperfect gate, that is one that doesn't quite perform C, we can separate out the noise as a different channel:

$$\tilde{\mathcal{C}} = \mathcal{E}_c\mathcal{C}\,, \tag{3.3}$$

where $\tilde{\mathcal{C}}$ represents a noisy implementation of $\mathcal{C}$ and $\mathcal{E}_c$ represents the noise on the Clifford gate. Note that here I have written the noise on the left hand side of the gate, but I could just as easily have written it as $\mathcal{C}\mathcal{E}'_c$, for a slightly different noise map. In this gate-independent analysis, nothing turns on the convention I have followed.

The *gate-independent* analysis of randomized benchmarking assumes that the noise channel $\mathcal{E}$ does not depend on the actual gate that is attempted, is Markovian (in particular it does not depend on what previous gates have been attempted) and it is not time varying. Therefore the noise model becomes $\tilde{\mathcal{C}} = \mathcal{E}\mathcal{C}\,, \forall \mathcal{C} \in \mathfrak{G}$, where $\mathfrak{G}$ is the unitary 2-design of gates we are going to analyze using randomized benchmarking (in [51] the Cliffords).

Whilst [51] does analyze randomized benchmarking for cases where the noise is not gate-independent, as we will see later this analysis has now been superseded ([5]) so we are going to assume that the error on the gates is gate independent and drop any subscripts on the noise channel $\mathcal{E}$.

With this stated then the purpose of the randomized benchmarking experiment is to characterize the average fidelity of the Clifford gates, by determining the fidelity (to the identity) of the noise channel $\mathcal{E}$, the closer $\mathcal{E}$ is to the identity, the closer each $\tilde{\mathcal{C}}$ is to $\mathcal{C}$.

The protocol for randomized benchmarking involves performing a number of sequences of the form (reading from right to left)

$$\langle\langle E|\tilde{\mathcal{G}}_{m+1}\,\tilde{\mathcal{G}}_m\,\tilde{\mathcal{G}}_{m-1}\ldots\tilde{\mathcal{G}}_2\,\tilde{\mathcal{G}}_1|\rho\rangle\rangle\,. \tag{3.4}$$

From the description of the protocol (section 3.1.1) it should be noted:

$$\mathcal{G}_{m+1} = (\mathcal{G}_m\ldots\mathcal{G}_1)^\dagger\,, \tag{3.5}$$

i.e. the last gate is the inverse of all the gates that have gone before it, meaning that if the gates are perfect the initial state (ρ) will be restored. Since the gates form a closed group (here, the Cliffords) such an inverse gate can be found and will be part of

the group. (See [52] for an example where non-groups are used in measurement based computing.)

We can re-write this with the noise decomposition of Equation 3.3 as:

$$\langle\langle E|\mathcal{E}\mathcal{G}_{m+1}\,\mathcal{E}\mathcal{G}_m\,\mathcal{E}\mathcal{G}_{m-1}\ldots\mathcal{E}\mathcal{G}_2\,\mathcal{E}\mathcal{G}_1|\rho\rangle\rangle \tag{3.6}$$

Now since each of the gates is a perfect gate (followed by the imperfect noise) and because they form a closed group any of these gates can be analyzed as being composed of two (or more) other group gates. We are going to decompose and re-label these gates, effectively performing something akin to a change of variables.

Let $\mathcal{C}_1 = \mathcal{G}_1$, then let $\mathcal{C}_2 = \mathcal{G}_2\,\mathcal{C}_1$, this means that $\mathcal{G}_2 = \mathcal{C}_2\,\mathcal{C}_1^\dagger$. Substituting in so far we have:

$$\langle\langle E|\mathcal{E}\mathcal{G}_{m+1}\,\mathcal{E}\mathcal{G}_m\,\mathcal{E}\mathcal{G}_{m-1}\ldots\mathcal{E}\mathcal{C}_2\,\left[\mathcal{C}_1^\dagger\mathcal{E}\mathcal{C}_1\right]|\rho\rangle\rangle\,, \tag{3.7}$$

where I have added brackets to help clarify the reasons for the transform.

This transform can be chased up the equation, transforming each $\mathcal{G}$ in turn, so that $\mathcal{C}_3 = \mathcal{G}_3\,\mathcal{G}_2\,\mathcal{G}_1 = \mathcal{G}_3\,\mathcal{C}_2$, and accordingly $\mathcal{G}_3$ becomes $\mathcal{C}_3\,\mathcal{C}_2^\dagger$. More generally we have $\mathcal{C}_n = \mathcal{G}_n\,\mathcal{G}_{n-1}\ldots\mathcal{G}_1 = \mathcal{G}_n\,\mathcal{C}_{n-1}$ and therefore $\mathcal{G}_n$ is replaced by $\mathcal{C}_n\,\mathcal{C}_{n-1}^\dagger$. Since by definition (Equation 3.5) we have $\mathcal{G}_{m+1} = (\mathcal{G}_m\ldots\mathcal{G}_1)^\dagger$, then we can see $\mathcal{G}_{m+1} = \mathcal{C}_m^\dagger$, and the rewritten equation becomes:

$$\langle\langle E|\mathcal{E}\left[\mathcal{C}_m^\dagger\mathcal{E}\mathcal{C}_m\right]\left[\mathcal{C}_{m-1}^\dagger\mathcal{E}\mathcal{C}_{m-1}\right]\ldots\left[\mathcal{C}_2^\dagger\mathcal{E}\mathcal{C}_2\right]\left[\mathcal{C}_1^\dagger\mathcal{E}\mathcal{C}_1\right]|\rho\rangle\rangle\,, \tag{3.8}$$

where I have inserted brackets to draw particular groupings to the reader's eye.

One thing to note is that since each of the gates $\mathcal{G}_{1..m}$ were drawn uniformly at random, then each of the changed variables $\mathcal{C}$ are also drawn uniformly at random. Therefore each element of the sequence (where an element here equates to each section in square brackets) is independent of the other. If we look at all possible sequences and average over this ensemble of sequences, this produces independent twirls. Taking

- m as the number of gates in each sequence,
- k as the total number of sequences (which will be $|\mathfrak{C}|^m$),
- $\{\vec{C}\}$ as the collection of gates selected for each of the sequences (each $\vec{C}$ being m gates long)
- $\{C_j\}$ as a collection of single gates, extracted from $\{\vec{C}\}$, the collection having k gates, being the gates in the j^{th} position of each of the the k sequences in $\{\vec{C}\}$; and
- $\mathbb{E}_{\vec{C}}$ being the expectation value of the expression following it, taken over the subscript to the $\mathbb{E}$

We have:

$$\mathbb{E}(\text{survival}) \tag{3.9}$$
$$= \mathbb{E}_{\{\vec{C}\}} \langle\langle E|\mathcal{E}\left[\mathcal{C}_m^\dagger \mathcal{E}\mathcal{C}_m\right]\left[\mathcal{C}_{m-1}^\dagger \mathcal{E}\mathcal{C}_{m-1}\right]\ldots\left[\mathcal{C}_1^\dagger \mathcal{E}\mathcal{C}_1\right]|\rho\rangle\rangle \tag{3.10}$$
$$= \langle\langle E|\mathcal{E}\mathbb{E}_{\{C_m\}}\left[\mathcal{C}_m^\dagger \mathcal{E}\mathcal{C}_m\right]\mathbb{E}_{\{C_{m-1}\}}\left[\mathcal{C}_{m-1}^\dagger \mathcal{E}\mathcal{C}_{m-1}\right]\ldots\mathbb{E}_{\{C_1\}}\left[\mathcal{C}_1^\dagger \mathcal{E}\mathcal{C}_1\right]|\rho\rangle\rangle \tag{3.11}$$
$$= \langle\langle E|\mathcal{E}\frac{1}{k}\sum_{\mathcal{C}\in\{\mathcal{C}_m\}}\left[\mathcal{C}^\dagger \mathcal{E}\mathcal{C}\right]\frac{1}{k}\sum_{\mathcal{C}\in\{\mathcal{C}_{m-1}\}}\left[\mathcal{C}^\dagger \mathcal{E}\mathcal{C}\right]\ldots\frac{1}{k}\sum_{\mathcal{C}\in\{\mathcal{C}_1\}}\left[\mathcal{C}^\dagger \mathcal{E}\mathcal{C}\right]|\rho\rangle\rangle \tag{3.12}$$

If $\mathfrak{C}$ represents the Clifford group of appropriate dimension, from subsection 2.3.3 we know that $|\mathfrak{C}|^{-1}\sum_{\mathcal{C}\in\mathfrak{C}}\left[\mathcal{C}^\dagger \mathcal{E}\mathcal{C}\right]$ is equal to a depolarizing channel with the same fidelity as $\mathcal{E}$, writing this as $\mathcal{E}_d$, then we can write

$$\frac{1}{k}\sum_{\mathcal{C}=\{\mathcal{C}_j\}}\left[\mathcal{C}^\dagger \mathcal{E}\mathcal{C}\right] = \mathcal{E}_d. \tag{3.13}$$

Thus we have:

$$\mathbb{E}(\text{survival}) = \langle\langle E|\mathcal{E}\mathcal{E}_d\mathcal{E}_d\ldots\mathcal{E}_d|\rho\rangle\rangle = \langle\langle E|\mathcal{E}[\mathcal{E}_d]^m|\rho\rangle\rangle \tag{3.14}$$

Finally we can absorb the extraneous $\mathcal{E}$ into what (we assume) is an already slightly faulty measurement (referred to as *absorbing it into the SPAM*), to get:

$$\mathbb{E}(\text{survival}) = \langle\langle E'|[\mathcal{E}_d]^m|\rho\rangle\rangle \tag{3.15}$$

Of course it is not experimentally feasible to iterate over all possible sequences for any m bigger than, say, 2 and so we have to estimate the true mean for any particular m by averaging over a random subset of possible sequences. How many sequences we should average over is discussed shortly.

It might be worth noting that if the noise had been written on the right-hand side of the perfect gate, the analysis would be identical save that the noise would be absorbed into the state preparation rather than the measurement.

As will be immediately apparent from Equation 3.15 as we vary sequence lengths then, assuming the noise is Markovian[§], gate and time independent, the only variable changing the survival rate expectation is the sequence length.

[§]see subsection 3.1.5 for a discussion about what this means.

Since we know that a depolarizing channel (here $\mathcal{E}_d$) can be written as

$$\mathcal{E}_d(\rho) = p\rho + (1-p)\frac{\mathbb{1}}{d}$$

(As previously noted for this thesis I have adopted the form of [51], where high p = high fidelity, rather than the more usual form where low p = high fidelity.) With the Pauli Liouville representation of a depolarizing channel (Equation 2.86) in mind it is easy to see that:

$$\mathcal{E}_d^m(\rho) = p^m\rho + (1-p^m)\frac{\mathbb{1}}{d} \tag{3.16}$$

Equation 3.15 then becomes:

$$\begin{aligned}
\mathbb{E}(\text{survival}) &= \langle\langle E'|[\mathcal{E}_d]^m|\rho\rangle\rangle && (3.17)\\
&= \text{Tr}\left(E'[\mathcal{E}_d^m(\rho)]\right) && (3.18)\\
&= \text{Tr}\left(E'\left(p^m\rho + (1-p^m)\frac{\mathbb{1}}{d}\right)\right) && (3.19)\\
&= \text{Tr}\left(E'\left(p^m(\rho - \frac{\mathbb{1}}{d}) + \frac{\mathbb{1}}{d}\right)\right) && (3.20)\\
&= \text{Tr}\left(E'(\rho - \frac{\mathbb{1}}{d})\right)p^m + \text{Tr}\left(E'\frac{\mathbb{1}}{d}\right) && (3.21)\\
&= Ap^m + B\,, && (3.22)
\end{aligned}$$

Where:

$$A \coloneqq \text{Tr}\left(E'(\rho - \frac{\mathbb{1}}{d})\right)\ B \coloneqq \text{Tr}\left(E'\frac{\mathbb{1}}{d}\right) \tag{3.23}$$

Equation 3.22 is the canonical randomized benchmarking fitting model (sometimes known as the *zeroth order model*), where A and B represent the SPAM parameters. In the Pauli Liouville representation of the depolarizing channel (Equation 2.86) $p = \lambda_2$. B represents the expectation value, given the faulty measurement E' when the channel has completely depolarized, in an ideal system $B = \frac{1}{d}$. A represents the probability of the faulty measurement giving a positive result when measuring the faulty state, with the identity subtracted, without any gates being applied. In an ideal system, where the state preparation and measurement are chosen to give maximal overlap, this would be equal to $\frac{d-1}{d}$. In an ideal single qubit experiment we would expect $A = B = 0.5$.

An experimental implementation would then consist of choosing certain sequence lengths (m) and running the experiment sufficient[¶] times for that number of gates to get an

[¶]what I mean by sufficient is discussed in more detail in section 3.3.

estimate of the survival rate. This is repeated for a number of lengths. The experimental data is then fit to an exponential decay curve of the form of Equation 3.22 to estimate the value of p (A and B as well). As previously discussed p is related to the average fidelity of the noise channel $\mathcal{E}$ as $\tilde{\mathcal{F}} = \frac{(d-1)p+1}{d}$.

3.1.3 Gate dependent errors — A re-analysis

Background

As a matter of practice, when randomized benchmarking was carried out in experiments, the predicted exponential decays were typically found and were used to successfully benchmark a number of single and two-qubit gates on a wide variety of platforms (e.g [23, 53, 54, 55, 56, 57, 58, 54, 59, 19])

However, the underlying assumptions of the zeroth order model, that of gate independent noise, would not in practice be met. The model assumes a unit or indivisible Clifford gate, whereas in practice these gates are constructed from more fundamental operations, typically being rotations around two (or all three) of the qubit axes. Since the Clifford gates were typically made up of between 1 to 3 fundamental operations, the idea that the same noise process would apply to each Clifford gate, was not realistic.

In a numerical investigation of the limits of randomized benchmarking, Epstein et al. [60] tested the benchmarking protocol, using realistically generated gates, against a number of realistic error modes, finding that, even in the case of non gate-independent noise, the benchmarking, in almost all cases, provided better than a factor of 2 estimate of fidelity (see subsection 3.1.5 for a discussion of $1/f$ noise and leakage errors).

As well as the zeroth order model discussed above Magesan provided a first order model that was formulated by perturbatively expanding the underlying sequences and throwing away the higher order terms, however in the cases where the zeroth order model proved not to be accurate, it did not seem that the higher order models helped ([3]).

The issue

Then, from researchers working on gate-set tomography, came an example of specific gate generators together with a specific noise model arose that seemed to indicate that for that example the number returned by randomized benchmarking was several orders of magnitude different the number that the theory predicted [61]. In brief, gate-set tomography attempts to use a variant of quantum process tomography to reconstruct the actual matrix representation of the gates being characterized [62] (see also [20] for a more general introduction). One of the difficulties with trying to reconstruct an actual matrix representing the gates being used is that in any black-box there is a gauge freedom

in the processes. For instance, if we conduct the following experiment:

$$\mathbb{E} = \sum_{i=1..k} \langle\langle E|\mathcal{G}_m \dots \mathcal{G}_1|\rho\rangle\rangle \tag{3.24}$$

where $\mathcal{G}$ represents some series of the gates that we wish to characterize, then if we apply a gauge to the system, i.e. map:

$$|\rho\rangle\rangle \to \mathcal{S}|\rho\rangle\rangle \tag{3.25}$$

$$\mathcal{G} \to \mathcal{S}\mathcal{G}\mathcal{S}^{-1} \tag{3.26}$$

$$\langle\langle E| \to \langle\langle E|\mathcal{S}^{-1}\,, \tag{3.27}$$

it can be seen you receive exactly the same results. It is impossible experimentally to determine (and arguably not meaningful to ask) whether your gates are implementing $\mathcal{G}$ or $\mathcal{S}^{-1}\mathcal{G}\mathcal{S}$ and therefore you can't write down specific canonical numbers for a particular gate in a particular representation. As is stated in [61] "The most obvious reasonable option is to arbitrarily fix a CPTP representation of the perfect gate set and to choose the representation of the imperfect gate set in which the gates are all CPTP and [the average gateset infidelity] is minimal" ... thus fixing the gauge i.e to choose the gauge that keeps reconstructed matrices CPTP but that closest matches the numbers it is expected the gate matrices should contain. Problems, however, remain. For instance, when the reconstructed matrices are used to predict the results of a randomized benchmarking experiment, with predicted values differ from those seen [21]. Similar claims have been made regarding the difficulty of reconciling the results of different, simpler, experiments.

In [43] Rudnicki et al. consider what gauge invariant information can be obtained concerning quantum channels. Looking at a single error channel their conclusion is that all gauge invariant information about the superoperator (describing the channel) is stored in its eigenvalues. Average gate fidelity (to the identity) is, therefore, gauge independent whereas unitarity (see subsection 4.1.1) can only be lower bounded in a gauge independent way. By arguing that any such gauge transformation acts on all gates from a particular gate set, the conclusion they draw is that this analysis is directly applicable to protocols such as GST. Their work confirms the difficulty caused for GST by the gauge degrees of freedom, however, as they note, one possibility might be to use other experiments as alternative ways of fixing the gauge, at the expense of breaking the black-box model which originally motivated GST.

In [61] it was noted that if the generating gates consisted only of gates which could perform a positive $\frac{\pi}{4}$ rotation around the X axis and a positive $\frac{\pi}{4}$ rotation around the Y axis, with the noise being a small amount of coherent Z-axis noise, then if you calculated the fidelity/error of each of the generated Clifford gates (which could involve up to 5 primitive gates) and average the error of each individual gate to get the average gate error, the number that comes back from a randomized benchmarking experiment indicates a much higher fidelity than the former calculation would indicate (the error calculated

from the randomized benchmarking experiment was orders of magnitude smaller than the averaged error of each specific gate). This didn't indicate that there was anything wrong with the theory of randomized benchmarking, the assumptions of the zeroth order model had not been met (although the noise was weak, it was strongly gate correlated). Using the first order model did not improve the match and, on calculation of the bounds indicated by the analysis of the first order model, the bounds turned out to be loose enough to encompass the results seen and therefore not helpful. However, it was widely believed that randomized benchmarking was robust to mild gate dependent noise, so what was happening?

The Bloch sphere picture of a single quantum bit helps visualize what is happening. Intuitively one can see that any continuous application of gates using these generators involves continued traversal of the Bloch sphere (compare to a system that might allow you to apply negative rotations and therefore there could conceivably be sequences of gates that remained on one side of the sphere). The Z-noise in the system anti-commutes with the rotations being applied, this means that the first order terms between, say, a quarter rotation around the X-axis and a 3 quarter rotation around the X-axis would cancel. The noise was effectively refocusing, in that with every traversal of the Bloch sphere through the application of 4 X-generators and/or 4 Y-generators all first order terms of the noise model would cancel. Although (in GST terms) the Clifford gates weren't doing exactly what they were theoretically supposed to do, and thus when you took the fidelity of each Clifford gate to its perfect representation you would see an error (which you could then average), if you look at the system as a whole, the average noise between gates was minimal and on average cancelled.

More precisely [61] defined an AGI as the average gate infidelity, where the gate infidelity is the infidelity of a particular gate, to its ideal counterpart. The average of the gate infidelity of each of these gates (being each of the gates used to invoke the unitary 2-twirl in the randomized benchmarking experiment) is the AGI. They noted that the average of these gate infidelities was gauge dependent and therefore that randomized benchmarking (which because it is an experiment must be gauge independent) couldn't return the average AGI. This prompted the question, asked in [61] 'What is randomized benchmarking measuring'. As we will see this was answered in [5].

The answer

Although [61] gives two methods for attempting to calculate the theoretical result of a randomized benchmarking the analysis by Wallman in [5] gives a more detailed understanding of what is happening and explains why, even in the regimes discussed above, we are getting a useful figure of merit. Wallman's work paves the way for more detailed understanding of interleaved benchmarking (section 3.2) and the unitarity of gates suffering gate-dependent noise (subsection 4.1.1).

Wallman identifies one of the problems with the bounds derived in [51] as the fact that the

gate dependent factor has to be extremely small for the higher factors (that are thrown away in the derivation) not to have an impact, this arises because there is a disconnect between norms and fidelity. His alternative analysis makes most of the perturbative terms vanish and (for good implementations of the gate) the additional decay term is smaller than the average noise term and also decays exponentially, meaning it quickly vanishes. The argument proceeds as follows:

1. Assume our unitary 2-design gates are taken from the group $\mathfrak{G}$. Identify a noise channel for each of them, written as follows $\tilde{\mathcal{G}} = \mathcal{L}\mathcal{G}\mathcal{R}_G$, where $\tilde{\mathcal{G}}$ represents the noisy gate and $\mathcal{L}$ is a noise channel applying to each of the noisy gates, $\mathcal{R}_G$ being a noise channel particular to each of these gates.
2. The average noise between gates, then becomes $|\mathfrak{G}|^{-1}\sum_{\mathcal{G}\in\mathfrak{G}}\mathcal{R}_G\mathcal{L}$, which is a gauge independent noise channel.
3. We can find an $\mathcal{L}$ and $\mathcal{R}$ such that $\tilde{\mathcal{G}} \approx \mathcal{L}\mathcal{G}\mathcal{R}$.
4. The average survival probability over all randomized sequences of length m is $Ap(\mathcal{E})^m + B + \epsilon_m$, where $\mathcal{E} = \mathcal{R}\mathcal{L}$, and
$$p(\mathcal{E}) = \frac{df(\mathcal{E},\mathcal{I}) - 1}{d-1} \tag{3.28}$$
$$|\epsilon_m| \leq \delta_1\delta_2^m\,, \tag{3.29}$$
where δ_1, δ_2 quantify the gate-dependence of the noise
5. δ_2 is small for good implementations of the gates (and thus decays to the point of non-detectability for $m \gtrsim 3$).

Sketch of the proof contained in [5]

A sketch of the proof is as follows. (Note that in the actual proof Wallman deals with the possibility of non trace-preserving maps, which adds a little bit of complexity to the proof. Rather than sketch out this additional complexity I will just state the final results — see [5] for full details.) It should be noted the proof uses characteristics of the Pauli Liouville representation of the noise channels. In the belief that it adds clarity I will move between abstract versions of noise channels and concrete representations in this form without heralding such maneuvers. They should be obvious from the context.

We need to find the following linear maps

$$\mathbb{E}_G(\tilde{\mathcal{G}}\mathcal{L}\mathcal{G}^\dagger) = \mathcal{L}\mathcal{D}_p \tag{3.30a}$$
$$\mathbb{E}_G(\mathcal{G}^\dagger\mathcal{R}\tilde{\mathcal{G}}) = \mathcal{D}_p\mathcal{R} \tag{3.30b}$$
$$\mathbb{E}_G(\mathcal{G}\mathcal{R}\mathcal{L}\mathcal{G}^\dagger) = \mathcal{D}_p\,, \tag{3.30c}$$

where $\mathcal{D}_p$ represents a depolarizing channel, such as Equation 2.86, where $\lambda_1 = 1$ and $\lambda_2 = p$.

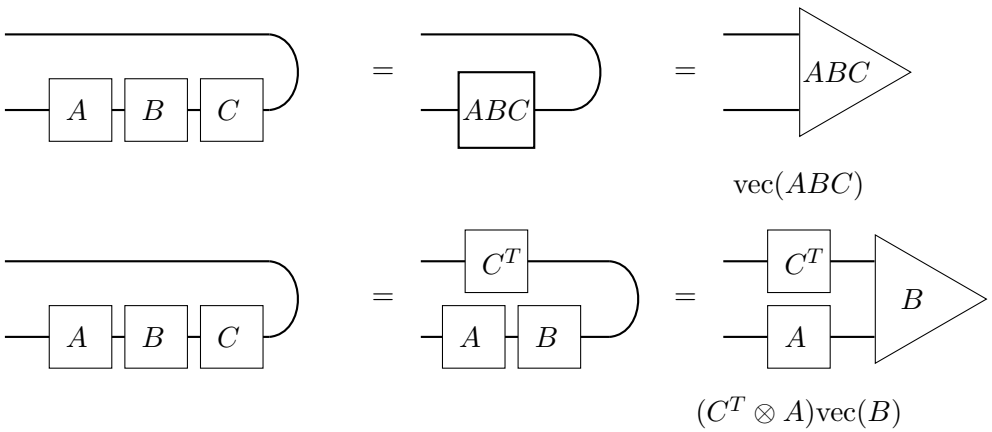

Figure 3.1: Graphical calculus proof of the identity: $\text{vec}(ABC)=(C^T \otimes A)\text{vec}(B)$

Because we are actually interested in the $\mathcal{E}_u$ section of the Pauli Liouville matrix representation of our noise channels we will define $\mathcal{G}_u$ to be the representation of $\mathcal{G}$, with the top left element set to zero, i.e. $\mathcal{G}_u(A) = \mathcal{G} - \text{Tr}(A)/d$. We can similarly split up $\mathcal{L} = |L\rangle\!\rangle\langle\!\langle \hat{I}_d| + \mathcal{L}'$, ($\hat{I}_d$, is the normalized identity matrix). Where we choose $|L\rangle\!\rangle$ such that $\mathcal{L}'|\hat{I}_d\rangle\!\rangle = 0$. (This just means means that $\mathcal{L}'$ has the same matrix representation as $\mathcal{L}$, save that the left most column is all zero). We do a similar process with the $\mathcal{R}$, noting that $\mathcal{R} = |\hat{I}_d\rangle\!\rangle\langle\!\langle R| + \mathcal{R}'$, such that $\langle\!\langle \hat{I}_d|\mathcal{R}' = 0$ that is the top row of $\mathcal{R}'$ is zero. The purpose of this manipulation is solely to confine our interest to the part of the noise matrix that controls the fidelity.

The relevant parts of Equation 3.30a and Equation 3.30b then become:

$$\mathbb{E}_G(\tilde{\mathcal{G}}\mathcal{L}'\mathcal{G}_u^\dagger) = p\mathcal{L}' \tag{3.31a}$$
$$\mathbb{E}_G(\mathcal{G}_u^\dagger\mathcal{R}'\tilde{\mathcal{G}}) = p\mathcal{R}' \tag{3.31b}$$

We then use a variation of Roth's Lemma to note that $\text{vec}(ABC)=(C^T\otimes A)\text{vec}(B)$ (Note that this proof is a one of the canonical proofs used in championing graphical calculus — as such its inclusion proved irresistible see figure 3.1). Since in the Pauli Liouville basis all the matrix elements are real and so the complex conjugate is equal to the transpose, Equation 3.31a and Equation 3.31b then become:

$$\mathbb{E}_G(\mathcal{G}_u \otimes \tilde{\mathcal{G}})\text{vec}(\mathcal{L}') = p\,\text{vec}(\mathcal{L}') \tag{3.32a}$$
$$\mathbb{E}_G(\tilde{\mathcal{G}} \otimes \mathcal{G}_u)^T\text{vec}(\mathcal{R}') = p\,\text{vec}(\mathcal{R}')\,, \tag{3.32b}$$

So p can be set to any eigenvalue of $\mathbb{E}_G(\mathcal{G}_u \otimes \tilde{\mathcal{G}})$, with $\mathcal{L}'$ and $\mathcal{R}'$ being the corresponding eigenvectors. As the identified solutions are only determined up to a normalization constant, and given that $\mathcal{D}_p$ commutes with all $\mathcal{G}$, then we can determine the solutions to the equations listed in 3.30.

So why does this help us? First note that we can write the gate dependent noise as $\Delta_G = \tilde{\mathcal{G}} - \mathcal{LGR}$, or re-arranging we have $\tilde{\mathcal{G}} = \mathcal{LGR} + \Delta_G$.

The average map over all randomized benchmarking sequences $\vec{\mathcal{G}}$ of length m is:

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_1\right] \tag{3.33}$$

We can expand any $\tilde{\mathcal{G}}$ as $\tilde{\mathcal{G}} = \mathcal{LGR} + \Delta_G$.

If we expand $\tilde{\mathcal{G}}_1$, then we have two terms to deal with:

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_2\mathcal{LG}_1\mathcal{R}\right] \tag{3.34}$$

and

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_2\Delta_1\right] \tag{3.35}$$

Noting that $\mathcal{G}_{m+1} = (\mathcal{G}_m\mathcal{G}_{m-1}\dots\mathcal{G}_2\mathcal{G}_1)^\dagger$ (since $\mathcal{G}_{m+1}$ is the inverting gate), we can rearrange as $\mathcal{G}_1 = (\mathcal{G}_{m+1}\mathcal{G}_m\mathcal{G}_{m-1}\dots\mathcal{G}_2)^\dagger$. More generally we can write:

$$\mathcal{G}_{j+1} = (\mathcal{G}_{m+1}\dots\mathcal{G}_{j+2})^\dagger(\mathcal{G}_{j-1}\dots\mathcal{G}_1)^\dagger\mathcal{G}_j^\dagger \tag{3.36}$$

Equation 3.34 then becomes:

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_2\mathcal{L}\left(\mathcal{G}_{m+1}\mathcal{G}_m\mathcal{G}_{m-1}\dots\mathcal{G}_2\right)^\dagger\mathcal{R}\right] \tag{3.37}$$

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_2\mathcal{L}\left(\mathcal{G}_2^\dagger\mathcal{G}_3^\dagger\dots\mathcal{G}_m^\dagger\mathcal{G}_{m+1}^\dagger\right)\mathcal{R}\right] \tag{3.38}$$

Since the gates are chosen independently we can bring the expectation inside the sequence and write the middle term as

$$\mathbb{E}_{\mathcal{G}_2}\left[\tilde{\mathcal{G}}_2\mathcal{L}\mathcal{G}_2^\dagger\right] \tag{3.39}$$

Using Equation 3.30a this becomes $\mathcal{LD}_p$. Since $\mathcal{D}_p$ commutes with the various channels (the Pauli Liouville matrix only has diagonal elements) it can be moved outside the expectation, the process repeated, to give us:

$$\mathcal{L}\left[\mathcal{D}_p\right]^m\mathcal{R}. \tag{3.40}$$

Next we turn to Equation 3.35, by expanding $\tilde{\mathcal{G}}_2$.

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_2\Delta_1\right] \tag{3.41}$$

$$= \mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \left[\mathcal{LG}_2\mathcal{R} + \Delta_2\right]\Delta_1\right] \tag{3.42}$$

$$= \mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_3\left[\mathcal{LG}_2\mathcal{R}\Delta_1\right]\right] + \mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_3\left[\Delta_2\right]\Delta_1\right] \tag{3.43}$$

Taking the left hand side of the equation and using Equation 3.36 and then the fact that $\Delta_1 = \tilde{\mathcal{G}}_1 - \mathcal{L}\mathcal{G}_1\mathcal{R}$ this becomes:

$$\mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_3\mathcal{L}(\mathcal{G}_{m+1}\dots\mathcal{G}_3)^\dagger\left[\mathcal{G}_1^\dagger\mathcal{R}\Delta_1\right]\right] \tag{3.44}$$

$$= \mathbb{E}_{\vec{G}}\left[\tilde{\mathcal{G}}_{m+1}\tilde{\mathcal{G}}_m \dots \tilde{\mathcal{G}}_3\mathcal{L}(\mathcal{G}_{m+1}\dots\mathcal{G}_3)^\dagger\left[\mathcal{G}_1^\dagger\mathcal{R}(\tilde{\mathcal{G}}_1 - \mathcal{L}\mathcal{G}_1\mathcal{R})\right]\right] \tag{3.45}$$

$$\tag{3.46}$$

Bringing the relevant expectation inside the brackets (using the fact that the gates are independent), then by virtue of Equation 3.30b and Equation 3.30c, the relevant part of the expression is:

$$\mathbb{E}_{\mathcal{G}}\left[\mathcal{G}_1^\dagger\mathcal{R}(\tilde{\mathcal{G}}_1 - \mathcal{L}\mathcal{G}_1\mathcal{R})\right] \tag{3.47}$$

$$= \mathbb{E}_{\mathcal{G}}\left[\mathcal{G}_1^\dagger\mathcal{R}\tilde{\mathcal{G}}_1\right] - \mathbb{E}_{\mathcal{G}}\left[\mathcal{G}_1^\dagger\mathcal{R}\mathcal{L}\mathcal{G}_1\mathcal{R}\right] \tag{3.48}$$

$$= 0. \tag{3.49}$$

For the right hand side of Equation 3.43 we repeat the process progressively expanding the next $\tilde{\mathcal{G}}_j$ and eliminating terms, the only non-zero term being:

$$\mathbb{E}_{\vec{\mathcal{G}}}\left[\Delta_{m+1}\Delta_m \dots \Delta_2\Delta_1\right] \tag{3.50}$$

The upshot of this means that every term of the expansion other than the first term (Equation 3.40) and the last term (Equation 3.50) vanishes.

So what have we done? We have shown that there exists (and we can find) a noise map $\mathcal{R}\mathcal{L}$ that represents the average noise between perfect gates. We have shown that if we twirl the noisy gates (using the noisy gates), then we get two decay terms and only two decay terms, namely the average noise map ($\mathcal{R}\mathcal{L}$) and a gate dependent noise map, given by Equation 3.50, which we call ϵ. Like the average noise, ϵ decays exponentially with the number of gates. In [5] Wallman bounds ϵ showing that Δ_G is of order $O(\sqrt{(1-p)})$, that is in the regime of reasonable gates e.g. $p \gtrsim 0.9$, ϵ will be much smaller the p and so after several gates will cease to have a measurable effect.

Finally note that if we provide a gauge transform $\tilde{\mathcal{G}} \to S\tilde{\mathcal{G}}S^{-1}$, then this becomes $\mathcal{L}\mathcal{G}\mathcal{R} \to S\mathcal{L}\mathcal{G}\mathcal{R}S^{-1}$ i.e. $\mathcal{L} \to S\mathcal{L}$ and $\mathcal{R} \to \mathcal{R}S^{-1}$. Under the gauge transform the average inter-gate noise is now $\mathcal{R}\mathcal{L} \to \mathcal{R}S^{-1}S\mathcal{L} = \mathcal{R}\mathcal{L}$, that is, it is unchanged.

As previously mentioned Wallman's full proof is more rigorous and deals with the possibility of trace-decreasing noise. The final equation derived is of the form:

$$\mathbb{E}_{\vec{G}}\mathcal{Q}_{\vec{G}} = Ap^m + Bt^m + \epsilon_m\,, \tag{3.51}$$

where the new terms here are the constants B and t which represent an exponential decay from loss in the system (in terms of the Pauli Liouville matrix of the noise, it is the λ_1 parameter in Equation 2.86).

Summary

So what is it that randomized benchmarking is measuring? Answer: the average noise between idealized gates, which — with the benefit of hindsight — is exactly what you would expect.

3.1.4 Interpretation of the average noise between idealized gates

It is possible to gain some intuition as to what the average noise between gates actually means, although with the caveat that we are dealing with a system where the dimensionality is greater than 3 and therefore it will often confound what we expect.

Whilst I am not aware of a method of reducing the map to a simple 3-dimensional object, it turns out it is possible to visualize 12-elements of the Clifford group, which in themselves are sufficient to constitute a unitary 2-design. We can then look at how a noise map alters this visualization. In [45], the authors offer up the following visualization. They note that up to phases $SU(2) \simeq SO(3)$ so we can think of every qubit unitary as corresponding to a three dimensional rotation. Splitting it into a rotation around each X, Y and Z-axis then a unitary can be plotted in three-dimensional space, with the displacement from the origin along a specific access corresponding to the rotation around the axis. Thus a rotation of π around the X-axis would correspond to the point (π,0,0), where we have assumed an order to the axes of (X,Y,Z). The mapping is, of course, a sphere, rotations of π and $-\pi$ being on the boundary. Whilst in [45] they chose to represent points on the boundary as two points in the diagram, (on the edge of the sphere), here I will plot just one.

Choosing the 12 Cliffords that correspond to such discrete rotations and setting the scale so that $2 = \pi$ we can visualize these Cliffords as shown in figure 3.2.

In the figure we can see the identity gate appears in the middle of the sphere; no rotation is applied. The other 3 Pauli gates are at the edge of the sphere (where I have arbitrarily plotted the π point). The remaining 8 Cliffords (the *non-Paulis*) that form part of this unitary 2-design, correspond to the points $(\pm 1, \pm 1, \pm 1)$, in the scale chosen.

If we apply a uniform z-rotation noise to the unitaries, we would expect the increasing z rotation of each of the non-Paulis, as well as an impact on the other axes of rotation. For the X and Y Paulis, there will be no increase in the z-rotation, but there will be a change in the X/Y rotations. The identity will gain a z-rotation. If we apply increasing small amounts of z-noise and re-plot then we see the results as shown in figure 3.3.

Although the 2D rendition of a 3D plot might make it difficult to see, the changes to the rotations are uniform. It can be imagined how a simple linear map can map the original design on to the noisy design. Because we have a uniform rotation on the gates, the frame potential remains close to 2. The identity can be seen drifting 'up' as the Z-noise is applied.

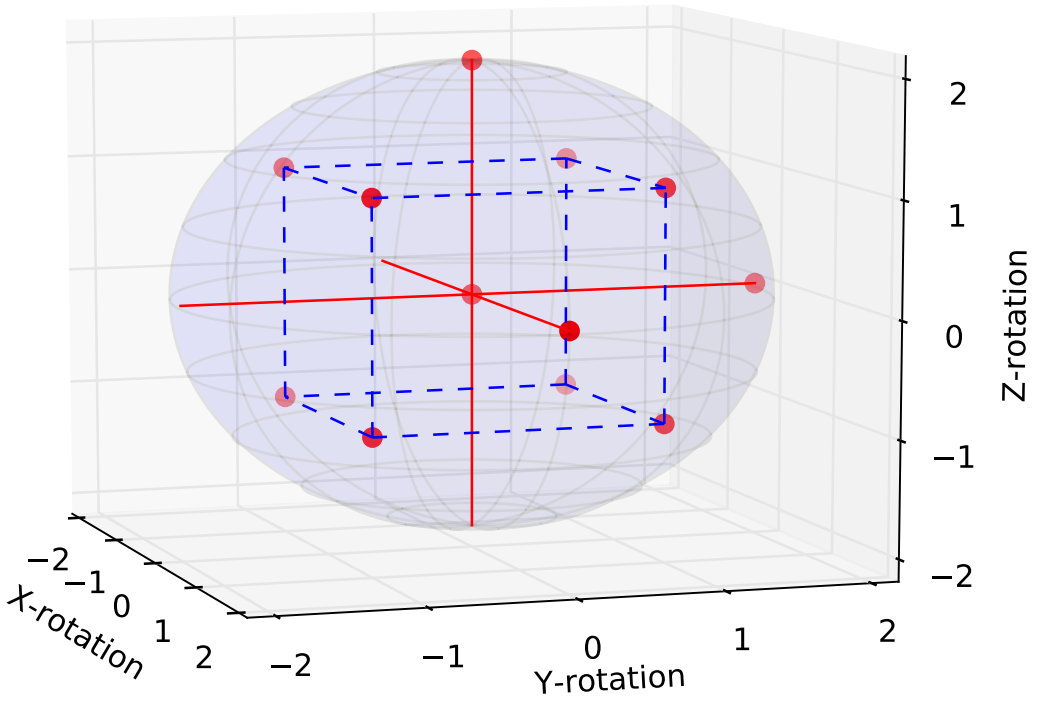

Figure 3.2: Visualisation of 12 rotation Cliffords forming a unitary 2-design. See text for details.

We can contrast this with the picture of what happens when we allow gate dependent noise (figure 3.4. The model used assumes each Clifford is made from positive $\pi/4$ x and y rotations, with a small amount of z-noise applied to each generator (this is the same model discussed in section 3.1.3). In the diagram I have kept the blue drift of increasing amounts of gate independent noise (i.e. the noise shown above) and have also plotted the drift of the Cliffords composed of the gx/gy generators in green.

As can be seen the noise on each of these Cliffords is different, specifically noting that the identity Pauli remains unaffected (to the first order) by the noise — since it is generated by four gx generators, with self cancelling noise. The noise for each of the non-Pauli Cliffords is the reverse of the Clifford on the other face of the sphere. The best linear map from the original (noiseless gates) to these gates (that is the one which minimizes the average difference between the mapped gates and the noise gates), is the one given by Wallman's decomposition, and in this case is close to identity.

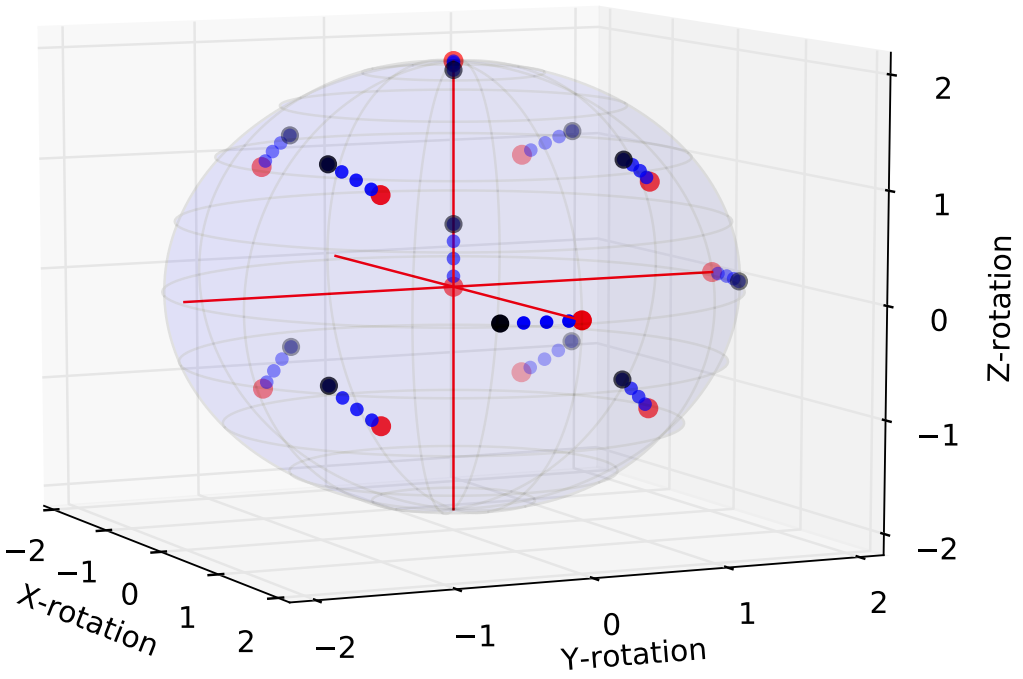

Figure 3.3: Visualisation of 12 rotation Cliffords, under a increasing (but small) z-noise rotation. Since the noise is gate independent the displacement is uniform over the gates.

3.1.5 Non-Markovian noise and time-varying noise

Markovian noise is a stateless or memoryless type of noise, that is the noise is independent of the history of, say, the experiment done. This is quite different from the accumulated error in the system, which will intimately depend on the actions taken to date, rather here we are looking at the noise map, or additional error map, applied at each arbitrary time-step. If we can determine the noise map to apply based only on the current state of the system (which can include information about how long the system has been running) then the process is Markovian.

Non-Markovian noise is, simply, noise which is not Markovian in nature. For example it might be the noise is different when you apply a particular sequence of gates or it might be affected by some process outside the system being considered, where that process cannot be effectively estimated by information currently available in the system. Time-varying noise can be both Markovian (the noise varies in such a way that it is a direct function of the age of the system) and non-Markovian (the noise varies in time but is not determinable solely from the current state and age of the system). However

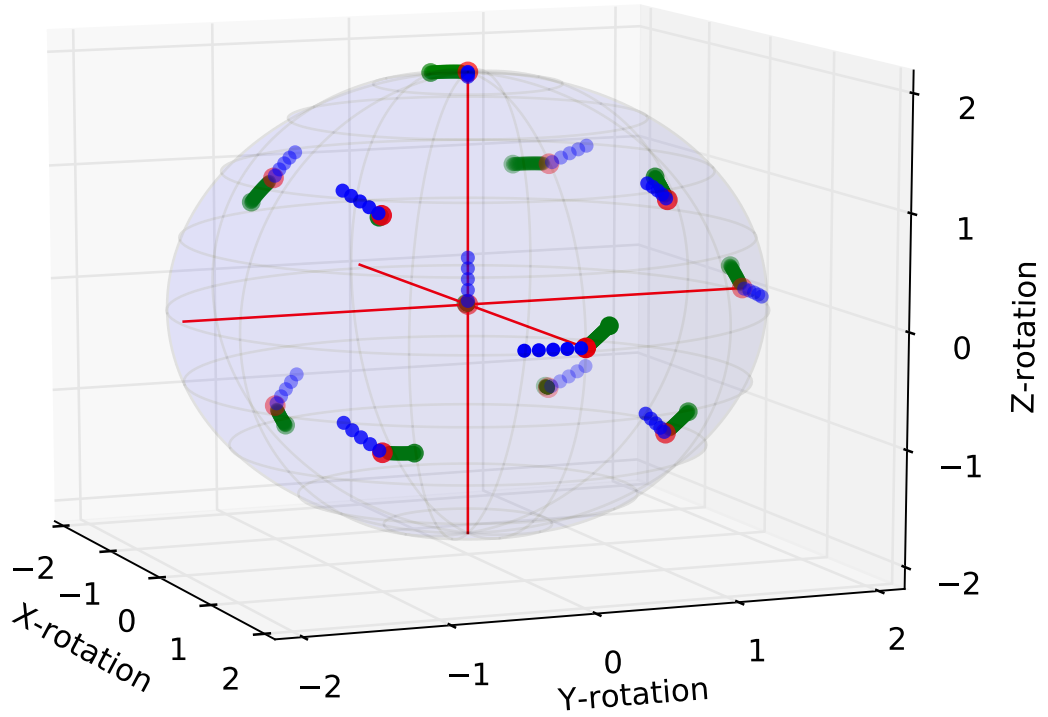

Figure 3.4: Visualisation of 12 rotation Cliffords, under a increasing (but small) gate dependent z-noise rotation. The blue circles show the gate independent noise plotted in figure 3.3, the green the result of gate dependent noise (the same noise affecting the generators as opposed to the gates). The best fit linear map between the displaced green gates is one close to the identity.

both types of time-varying noise violate the assumptions in randomized benchmarking, namely that the noise applied to any particular gate remains constant over the life of the experiment.

The theory of how non-Markovian and time-varying noise impacts on quantum experiments, and here randomized benchmarking experiments, is not yet well developed.

So far, perhaps the best analyzed type of noise is time-dependent or low frequency noise. Initial numerical results were presented in [60] where they found RB estimated the error rate to within a factor of two, the protocol producing a fidelity decay that could be modeled by a composition of correlated depolarizing channels. The degree of correlation appeared to affect the extent to which the averaged decays could be modeled by a simple exponential decay (this manifested as the confidence interval of the randomized benchmarking estimate becoming saturated as the sequence samples reached 100 or greater).

Ref. [60] also looked at the impact of leakage errors (where the computation leaks into levels outside the qubit subspace and back again — a type of non-Markovian error if you limit your system to the qubit subspace). There they find that the benchmarking protocol still estimates error well, albeit with different SPAM parameters as a result of the increased Hilbert space available to the system.

In [38] it is shown that randomized benchmarking can be used to characterize time-dependent noise, provided the gate-dependence is negligible. In [3] (attached as chapter 8) a non-exponential decay curve was noticed in the experimental data. Since the system in question had no loss-rates as such it was hypothesized that the system was being affected by low-frequency noise, effectively meaning that the fidelity of the system was varying from sequence to sequence, which — as shown in the paper — would lead to a non-exponential decay curve in the combined data set. Further analysis showed that the data could be fit to a double exponential decay curve, which operationally speaking, might be thought of as a worst case and best case run. Statistical tools such as the Akaike information criteria were used to provide support that this two-decay model was indeed a viable hypothesis. Finally in [63] extensive analysis was carried out of temporal noise correlations on the outcome of RB protocols. Although there is a wealth of material in the paper one of the main conclusions was that low frequency noise can lead to an overestimation of the fidelity of the channel, specifically where insufficient sequence randomizations are sampled. In extreme cases with a $\sigma = 0.015$ and a confidence interval of $\pm 10\%$, the analysis indicated that at least 443 random sequences should be sampled.

Recently in [64] Gaussian time-correlated noise was studied in relation to its effect on a single qubit. The authors show that the expected sequence fidelity can be linked to a long-range coupled spin-one Ising model. A high effective temperature expansion for such a partition function shows that the decay will vary from the usual exponential decay pattern (seen in the case of uncorrelated noise) to a power law for quasistatic noise. In the case where such time-correlated noise affects the system being benchmarked, blind fitting of an exponential decay to a sequence which is not in fact an exponential decay can give unreliable estimates of average gate error.

3.2 Interleaved Protocol

3.2.1 Background

Although randomized benchmarking gives us information about the average noise channel between all the gates forming a unitary 2-design in the system (typically gates forming the Clifford group), it is always interesting to explore whether more information can be extracted. The interleaved randomized benchmarking protocol was designed as an attempt to be able to characterize the fidelity of the individual gates forming the unitary 2-design. Extensions have been proposed to allow specific or arbitrary additional gates to be characterized. The initial protocol was introduced by Magesan et al. [65]. Although the bounds in that paper were rather loose, the idea of benchmarking individual gates has received a lot of experimental interest (for example, [66, 67, 68, 19, 69]). Kimmel et al. [70] introduced a method to allow randomized benchmarking to perform state tomography, including a variant to allow the fidelity of individual gates to be characterized. The bounds found in [70] were tighter than those in [65] and gave a clearer insight into the problems associated with composed noise channels. Carignan-Dugas et al. [71] and separately Cross et al. [72] introduced a form of benchmarking based, not around a unitary 2-design, but rather dihedral groups, which allows arbitrary gates to be characterized albeit with some additional complexity. (Ref [72] encompassed multiple qubits by considering the CNOT-dihedral group.) One of the contributions of this thesis shows how it is possible to use the fact that a T-gate is part of the third level of the Clifford hierarchy to benchmark the T-gate, without abandoning the benefits of a unitary 2-design [1]. I will explore each of these in some more detail. In discussing the above it should be borne in mind that none of the analysis takes into account Wallman's recent analysis of gate dependent noise ([5]), but rather uses the standard randomized benchmarking assumptions, together with a belief that randomized benchmarking was robust to small gate dependent errors. In the last section I will sketch out some ideas as to how Wallman's analysis might be extended to the interleaved protocols.

3.2.2 How is it designed?

The interleaved protocol is a protocol that appears intuitively simple, but in reality is difficult to analyze. As mentioned above, the first description of using randomized benchmarking appears as an extension of Magesan's previous work [65]. (Previous protocols to measure the fidelity of a single qubit used slightly different methodologies, rendering them more vulnerable to SPAM errors e.g. [73].)

In essence the gate of interest, being one of the gates that comes from the unitary 2-design group used to perform the twirling in randomized benchmarking (say $\mathcal{C}$) is interleaved between an otherwise classic randomized benchmarking sequence, e.g:

$$\langle\langle E|\tilde{\mathcal{G}}_{m+1}\,\tilde{\mathcal{C}}\,\tilde{\mathcal{G}}_m\,\tilde{\mathcal{C}}\,\tilde{\mathcal{G}}_{m-1}\ldots\tilde{\mathcal{G}}_2\,\tilde{\mathcal{C}}\,\tilde{\mathcal{G}}_1|\rho\rangle\rangle\,, \tag{3.52}$$

where, in this case $\tilde{\mathcal{G}}_{m+1}$ inverts the entire sequence (that is, it takes into account the interleaved gate).

With the standard randomized benchmarking assumptions of gate independent noise channels (although, obviously, it is clear that $\mathcal{C}$ is expected to have a different error) and Markovian noise, a sketch of the analysis is as follows:

We can write each of the noisy twirling gates as $\tilde{\mathcal{G}} = \Lambda_G \mathcal{G}$ and the noisy interleaved gate as $\tilde{\mathcal{C}} = \mathcal{C}\Lambda_C$.

A relevant section of the sequence then becomes:

$$\ldots \Lambda_{G_{n+1}} \mathcal{G}_{n+1} \mathcal{C} \Lambda_C \Lambda_{G_n} \mathcal{G}_n \ldots , \tag{3.53}$$

then since $\mathcal{C}$ is assumed to come from the same group as the twirling gates $\mathcal{G}$, we can re-write $\mathcal{G}_{n+1}\mathcal{C}$ as $\mathcal{G}'_{n+1}$ (noting that $\mathcal{G}'_{n+1}$ is still independently random) which gives us

$$\ldots \Lambda_{G_{n+1}} \mathcal{G}'_{n+1} \Lambda_C \Lambda_{G_n} \mathcal{G}_n \ldots . \tag{3.54}$$

Doing this for the whole sequence and then changing variables as discussed with the standard randomized benchmarking model, gives us a twirl of the combined noise $\Lambda_C \Lambda_G$

3.2.3 Underlying assumption – and analysis of error

It will be recalled that the fidelity of a channel, to the identity, is linearly related to the depolarizing factor of the twirled channel (that is, if Λ is the twirled channel, let $p(\Lambda)$ be the depolarizing factor and table 2.1 contains the relevant relations). The aim of the interleaved benchmarking protocol then is to (i) calculate the depolarizing factor of a reference channel using standard randomized benchmarking, (the Λ_G channel); then (ii) to calculate the combined channels as discussed above (the $\Lambda_C \Lambda_G$) channel; and (iii) assert that

$$p(\Lambda_C) \approx \frac{p(\Lambda_C \Lambda_G)}{p(\Lambda_G)} \tag{3.55}$$

This assumption is only the case when one of the composed channels is actually a depolarizing channel (as opposed to something that is twirled into a depolarizing channel). Intuitively this can be seen if one keeps the Pauli Liouville representation in mind.

In [70], Kimmel et al. introduced a different way to use randomized benchmarking, to benchmark not just the average noise channel to identity, but rather to benchmark against different orthogonal maps in order to allow the unital part of trace preserving maps to reconstructed. They also showed how fidelities could be estimated for gates beyond the Clifford group. For the first method Clifford gates were used to create linearly independent Clifford group maps, and the fidelity of these maps to the fixed noise channel in question $\mathcal{E}$ was characterized. Linearity then allows the reconstruction of the

unital part of the noise map. On a practical level the rapid decay of such benchmarking experiments, means that statistically useful results require a lot of data (e.g. [74]), but the protocol is robust. In the second part of the paper, the non-Clifford fidelity estimates, one also requires analysis of channel composition, with a similar assumption to Equation 3.55 and hence they introduce an analysis of these bounds.

Using the complete positivity constraints on the χ-matrix, Kimmel et al. bound the $\chi_{0,0}$ element of the composed matrix (recalling the $\chi_{0,0}$ element is linearly related to the quantities of interest table 2.1) as:

$$\chi^{\mathcal{AB}}_{0,0} = \chi^{\mathcal{A}}_{0,0}\chi^{\mathcal{B}}_{0,0} \pm \left(2\sqrt{(1-\chi^{\mathcal{A}}_{0,0})\chi^{\mathcal{A}}_{0,0}(1-\chi^{\mathcal{B}}_{0,0})\chi^{\mathcal{B}}_{0,0}} + (1-\chi^{\mathcal{A}}_{0,0})(1-\chi^{\mathcal{B}}_{0,0})\right) \tag{3.56}$$

In [49] it is shown that the bounds in Equation 3.56 can be saturated with reasonable noise models, but that the unitarity of the channels (see section 4.1.3) can be used to tighten them.

3.2.4 Where might this be useful

The form of the bounds in Equation 3.56 means that where the group of gates performing the twirl have similar fidelities to the gate being benchmarked then, in the absence of other information, little knowledge of the fidelity of the gate in question can be gleaned from an application of Equation 3.55. On the other hand, as noted in [71] if we are in the regime where the gates performing the twirl are of high fidelity, but the gate that we are interested in might be lower in fidelity, then we can indeed gain insight as to the performance of the lower fidelity gate.

In the case of [71] they were looking at using Dihedral Benchmarking to estimate the fidelity of a T-gate (also known as a $\frac{\pi}{8}$ gate). Although the protocol is useable for a wide variety of gates.

Why a T-gate?

As previously mentioned the Clifford gates are fundamental to many error correcting protocols and so are believed to be fundamental to the implementation of quantum information processors (computers). However, they are not, in themselves, sufficient to allow universal quantum computing [32]. Whilst there is a wide choice in which gate to add, the T-gate is one which, with the Cliffords, allows for universal quantum computing.

Dihedral Benchmarking

Dihedral benchmarking utilizes the properties of the dihedral groups $\mathbb{D}_8$ and $\mathbb{D}_4$ to allow the benchmarking of the T-gate. In [71] the analysis was carried out for a single qubit,

whereas the analysis in [72] included multiple qubit systems. For the ease of discussion I will limit this review to the single qubit scenario. The dihedral groups represent rotations around the Z-axis of the Bloch sphere (although note, the orientation of the Bloch sphere is completely arbitrary), 8 equally spaced rotations for $\mathbb{D}_8$, 4 for $\mathbb{D}_4$, in both cases together with a π rotation around the X-axis. The gates needed to implement a $\mathbb{D}_4$ design can be drawn from the Clifford group, the addition of a T-gate (which represents, confusingly, the $\pi/4$ rotation) completes the gates needed for $\mathbb{D}_8$. The idea being that the difference in the depolarizing factor between implementations of the $\mathbb{D}_4$ and $\mathbb{D}_8$ groups will give us information about the T-gate. D However, unlike the two irreducible representations of a unitary 2-design, the representation of the dihedral group elements in the Pauli Liouville representation decomposes as a direct sum of 3 irreducible representations, being the trivial representation, the faithful representation and the parity representation. The means that even with trace preserving noise any exponential decay will have two parameters. In [71] the protocol includes methods to construct runs that allow the different gathered data-sets to be manipulated to remove, with different manipulations, one of these two parameters from the fitting process. Numerical evidence is also offered to support the estimates of fidelity of the T-gate using this protocol.

Estimating the T gate with standard interleaved randomized benchmarking

In [1] the fact that the Pauli gates conjugate the T-gate to the Clifford group was exploited to allow the characterization of the T-gate using the unitary 2-design twirl of randomized benchmarking. The paper describing this is attached as chapter 6.

3.2.5 Interleaved benchmarking and the gate-dependent analysis

Although there is much work to be done applying Wallman's analysis of gate dependent noise [5] to the interleaved protocol, there are some immediate consequences that fall out when we interleave members of the twirling group. Here I will assume the twirling group is composed of Clifford gates, but nothing turns on this.

With the interleaved protocol we are applying sequences of $\tilde{\mathcal{G}}_j\tilde{\mathcal{C}}$, where $\tilde{\mathcal{C}}$ is the noisy interleaved Clifford, and $\tilde{\mathcal{G}}_j$ is the j^{th} gate from our unitary 2-design. Typically the group formed by $\mathfrak{G}\mathfrak{C} :- \{\mathcal{G}\mathcal{C}\}\,\forall\mathcal{G} \in \mathfrak{G}$, will also be a unitary 2-design, always if $\mathcal{C} \in \mathfrak{G}$.

Thus if we treat each of the combined gates $\tilde{\mathcal{G}}\tilde{\mathcal{C}}$ as a single noisy composite gate $\widetilde{\mathcal{G}\mathcal{C}}$ (there will be as many of them as there are elements of $\mathfrak{G}$) we can treat the gate dependent noise on these combined gates, the same way we did before and write:

$$\widetilde{\mathcal{G}\mathcal{C}} = \mathcal{L}_x\mathcal{G}\mathcal{C}\mathcal{R}_x + \Delta_{\mathcal{G}\mathcal{C}}\,, \tag{3.57}$$

and find the eigenvalues, as per [5] (see Equation 3.30)

$$\mathbb{E}_G(\widetilde{\mathcal{GC}}\mathcal{L}_x(\mathcal{GC})^\dagger) = \mathcal{L}_x\mathcal{D}_{(\mathcal{R}_x\mathcal{L}_x)} \tag{3.58a}$$
$$\mathbb{E}_G((\mathcal{GC})^\dagger\mathcal{R}_x\widetilde{\mathcal{GC}}) = \mathcal{D}_{(\mathcal{R}_x\mathcal{L}_x)}\mathcal{R}_x \tag{3.58b}$$
$$\mathbb{E}_G(\mathcal{GC}\mathcal{R}_x\mathcal{L}_x(\mathcal{GC})^\dagger) = \mathcal{D}_{(\mathcal{R}_x\mathcal{L}_x)}\,, \tag{3.58c}$$

The figure delivered by randomized benchmarking is then the depolarizing factor from twirling the channel $\mathcal{R}_x\mathcal{L}_x$, which represents the averaged noise between perfect versions of the $\mathcal{GC}$ gate sets. Where we have truly gate independent noise, this will just an application of the $\mathcal{RL}$ noise channel, otherwise the imperfect $\tilde{C}$ gate will alter the gate dependent factors, that is $\Delta_\mathcal{G}$s. This will impact the averaged noise channel, depending on whether it brings the gate perturbations closer to something that can be replicated with a linear map (the averaged noise will increase) or to something where the noise is more self-cancelling (the averaged noise will decrease). Both a decrease in average fidelity and an increase in fidelity with the addition of interleaved gates have, in practice, been observed [3].

This insight doesn't yet help us relate the twirled depolarizing factor to the fidelity of a specific group gate, which still remains an active area of research. The following observations can be made.

Define $\mathcal{L}$ and $\mathcal{R}$ as the solutions to Equation 3.30 for the reference benchmarking experiment, and assuming here $\mathcal{C} \in \mathfrak{G}$ so that $\tilde{\mathcal{C}} = \mathcal{LCR}_c = \mathcal{LCR} + \Delta_c$. Then we have:

$$\widetilde{\mathcal{GC}} = \mathcal{L}_x\mathcal{GCR}_x + \Delta_{(GC)} \text{ and also, expanding } \tilde{\mathcal{G}} \tag{3.59}$$
$$\widetilde{\mathcal{GC}} = (\mathcal{LGR} + \Delta_G)\tilde{\mathcal{C}} = \mathcal{LGR}\tilde{\mathcal{C}} + \Delta_G\tilde{\mathcal{C}}. \tag{3.60}$$

Clearly the only gate dependent components are $\Delta_{(GC)}$ in Equation 3.59 and $\Delta_G\tilde{\mathcal{C}}$ in Equation 3.60. Given this we have $\mathbb{E}_G\left(\Delta_{(GC)}\right) = \mathbb{E}_G\left(\Delta_G\tilde{\mathcal{C}}\right)$.

Looking at the expansion of $\tilde{\mathcal{C}}$ above, it can be seen that the gate perturbations from the non-interleaved noisy gates are further shifted by an addition of the gate specific perturbations $\Delta_\mathcal{C}$, confirming the intuition of the discussion above.

We can also use a similar expansion of Equation 3.58a and Equation 3.58b, together

with some re-arrangement to note that:

$$\mathbb{E}_G\left((\mathcal{GC})^\dagger R_x \widetilde{(\mathcal{GC})}\right) = \mathcal{D}_{(R_x L_x)}\mathcal{R}_x \tag{3.61}$$

$$\mathcal{C}^\dagger \mathbb{E}_G\left(\mathcal{G}^\dagger \mathcal{R}_x \tilde{\mathcal{G}}\right)\tilde{\mathcal{C}} = \mathcal{D}_{(R_x L_x)}\mathcal{R}_x \tag{3.62}$$

$$\mathbb{E}_G\left(\mathcal{G}^\dagger \mathcal{R}_x \mathcal{L}\mathcal{G}\mathcal{R}\right) + \mathbb{E}_G(\mathcal{G}^\dagger \mathcal{R}_x \Delta_G) = \mathcal{D}_{(R_x L_x)}\mathcal{C}\mathcal{R}_x\tilde{\mathcal{C}}^\dagger \tag{3.63}$$

$$\mathcal{D}_{(R_x L_x)}\mathcal{C}\mathcal{R}_x\tilde{\mathcal{C}}^\dagger - \mathcal{D}_{(R_x L)}\mathcal{R} = \mathbb{E}_G(\mathcal{G}^\dagger \mathcal{R}_x \Delta_G) \tag{3.64}$$

$$\mathcal{D}_{(R_x L_x)}\mathcal{R}_x - \mathcal{D}_{(R_x L)}\mathcal{C}^\dagger\mathcal{R}\tilde{\mathcal{C}} = \mathcal{C}^\dagger \mathbb{E}_G(\mathcal{G}^\dagger \mathcal{R}_x \Delta_G)\tilde{\mathcal{C}}\,, \tag{3.65}$$

which is showing us that the difference between the $\mathcal{R}_x$ of the interleaved protocol and the $\mathcal{R}$ of the reference protocol is related to the residual noise of the gates (Δ_G) perturbed by the noise in the interleaved Clifford. Similarly with $\mathcal{L}_x$:

$$\mathbb{E}_G\left(\widetilde{(\mathcal{GC})}\mathcal{L}_x(\mathcal{GC})^\dagger\right) = \mathcal{D}_{(R_x L_x)}\mathcal{L}_X \tag{3.66}$$

$$\mathbb{E}_G\left(\mathcal{L}\mathcal{G}\mathcal{R}\tilde{\mathcal{C}}\mathcal{L}_x\mathcal{C}^\dagger\mathcal{G}^\dagger\right) + \mathbb{E}_G\left(\Delta_{\mathcal{G}}\tilde{\mathcal{C}}\mathcal{L}_x\mathcal{C}^\dagger\mathcal{G}^\dagger\right) = \mathcal{D}_{(R_x L_x)}\mathcal{L}_X \tag{3.67}$$

$$\mathcal{D}_{(R_x L_x)}\mathcal{L}_X - \mathcal{D}_{(\mathcal{R}\tilde{\mathcal{C}}\mathcal{L}_x\mathcal{C}^\dagger)}\mathcal{L} = \mathbb{E}_G\left(\Delta_{\mathcal{G}}\tilde{\mathcal{C}}\mathcal{L}_x\mathcal{C}^\dagger\mathcal{G}^\dagger\right)\,. \tag{3.68}$$

Whilst the above analysis confirms that the Wallman analysis holds in the interleaved protocol and goes someway towards showing how the average noise fidelity between random Cliffords is affected by the interleaving of a fixed, noisy, gate, it still lacks the bounds given by Theorem 3 and the bound on the perturbation term given in Theorem 4 of [5]. In addition, with certain assumptions, it might be possible to use Equation 3.65 and Equation 3.68 to provide more insight to the bounds given in Equation 3.56.

3.3 Practicalities of fitting the data

3.3.1 What needs to be fit

So far in the analysis of randomized benchmarking I have assumed that one gathers *enough* data to fit the decay curve with appropriate precision. This is primarily a statistical exercise. Whilst one might hope the application of such statistical models is straightforward, it turns out that working out the variance between sequences and the best way to utilize limited experimental resources is challenging.

The protocol outlined in section 3.1.1 contains two relevant unspecified parameters, l_m, which specifies how many measurements are taken of a specific sequence in order to estimate the survival probability of that sequence and n_m which specifies how many

sequences (of a specific length) should be sampled in order to estimate the average survival probability of all sequences of that length. To give guidance as to what these values should be we need to look at the underlying variance of the sequences of a randomized benchmarking experiment. Related to how much and what data to gather is the problem of the best way to fit the curve and this leads to different forms of experimental design. These issues are explored below.

3.3.2 Underlying variance

To recap, for the purposes of this section, randomized benchmarking involves choosing a gate length, then constructing random sequences of gates, implementing a final gate that inverts all the other gates, and measuring the *survival probability* of the initial state. (A state being said to survive if it was measured as being in its original state.) Clearly for a given gate length m where $m \gg 1$, the number of potential sequences are far too large to sample (growing exponentially in m). The question then becomes how many sequences do we have to sample in order to estimate the average sequence survival rate to within our required precision (what our required precision should be is dealt with later).

Initial bounds proposed in [51] used Höeffding's inequality to estimate the number of sequences to be sampled, suggesting that no more than 7×10^4 samples would be required to measure a 0.99 fidelity gate, with an accuracy (ϵ) of 10^{-3} and a confidence level of $\pm 5\%$. Experimental results and the numeric analysis by [60] suggested that far fewer, in the order of 10–100 random sequences, were actually required. Ref. [60] suggested that this might be because of the fitting process, where multiple sequence lengths were used, rendering the per-sequence bounds of [51] pessimistic.

In [38] tighter bounds are found for the variance between different sequences (*the sequence variance*). Where r represents the infidelity (i.e. 1-Fidelity), then with m representing the number of gates in the sequence length, in the regime where $mr \ll 1$, Wallman and Flammia find the following bounds for qubits

$$\sigma_m^2 \leq m^2 r^2 + \frac{7}{4} m r^2 + O(m^2 r^3)\,, \tag{3.69}$$

where they have used the properties of a unitary 2-design for qubits to tighten the bounds.

3.3.3 Cliffords, almost a 4-Design, impact on variance of sequences

Recently [44] and separately [75] analyze all the irreducible representations of the full Clifford group, determining that it falls just short of a unitary 4-design[‖]. In [76] these discovered properties of the Clifford group are used to:

[‖] In a precise sense, it turns out that the 4th tensor power of the Clifford group affords only one more invariant subspace than the 4th tensor power of the unitary group.

- provide bounds on the sequence variance that are independent of the number of qubits being benchmarked; and
- For a single-qubit version of RB provide bounds on the sequence variance for long sequence lengths.

Single qubit

For a single qubit the bounds found in [76] are, for small sequence lengths:

$$\sigma_m^2 \leq \frac{13}{2} m r^2 + 2\eta m r \tag{3.70}$$

and for larger sequence lengths:

$$\sigma_m^2 \leq \frac{7}{2} r + \eta \tag{3.71}$$

where in both cases η is a prefactor (normally small) that depends entirely on the SPAM error.

Multi-qubit

The bounds presented by [75] for a multi-qubit system are based on two different scenarios. The first, shown here, is where there is minimal SPAM error, in which case they find the following inequality:

$$\mathbb{V}_0^2 \leq 2\left(\frac{1-\chi^m}{1-\chi}\right)\frac{d+1}{d-1}r^2 + \frac{1}{4}\left(\frac{1-f^{2m}}{1-f^2}\right)\left(\frac{d^2-2}{(d-1)^2}\right)r^2\,, \tag{3.72}$$

where χ represents the unitarity of the noise channel (see subsection 2.3.1) and f is the depolarizing factor. The important point to notice is that it scales quadratically in the infidelity and is asymptotically independent of system size d. The paper contains more analysis, including where there are non-negligible amounts of SPAM and the interested reader is referred to the paper.

For present purposes the important take-away is that, especially when the Cliffords are used as the twirling-gates, we have bounds on the variance that only require experimentally feasible sequences to be taken.

3.3.4 Not throwing away your shot — analyzing changing sequence versus multiple measurements per sequence

All the analysis regarding the number of sequences that are needed has, so far, assumed we have an estimate of the survival rate of each sequence. Given the nature of quantum

mechanics this is almost certainly not the case. For each random sequence of a particular gate length m the question then arises do we take many measurements (*shots*) of a sequence to estimate the survival probability for that sequence and then estimate the average survival probability over a number of sequences; or is it better to take a single shot over multiple sequences and average those results. The answer to this question is intrinsically tied up in the consumption and availability of experimental resources.

Using Bayesian analysis techniques [77] shows (without considering experimental resources) that the single-shot limit is not only analytically tractable but provides the most information per shot.

Here I will provide a frequentist analysis of the same question, showing how this analysis can be used to determine shots versus sequences where, for instance, changing sequences consumes more experimental resources (e.g. takes longer) as compared to taking multiple measurements of the same sequence.

There are two relevant sources of statistical fluctuation, the finite measurement statistics (for a given sequence, what is the variance in observed survival probabilities for that sequence — *the in-sequence variance*) and the uncertainty caused because we are sampling over a finite number of sequences (given we can only measure a few sequences, what is the variance in the observed average survival probability over all sampled sequences — *the per-sequence variance*). What we want to do is to combine these two levels of uncertainty into a single estimate of the variance of the average survival probability.

We turn first to the in-sequence variance. For a given sequence s of length m each measurement is a binary-outcome measurement (a 1 or a 0). The statistics of such measurements will be binomially distributed according the survival probability of that particular sequence. Assuming we perform X measurements then for that specific sequence we have will have an estimator of the survival probability $\hat{F}_{(m,s,X)}$, given by a binomial distribution based on the actual survival probability $F_{(m,s)}$. Accordingly we can write:

$$\hat{F}_{(m,s,X)} \sim B(X, F_{m,s}). \tag{3.73}$$

Since each measurement outcome is either 1 or 0, then the estimator $\hat{F}_{(m,s,X)}$ is going to be a rational number of the form x/X where x ranges from $0 \ldots X$, being the number of times a 1 measurement was obtained.

The probability of observing any particular estimator will equal the probability of drawing any particular sequence multiplied by the probability that for that sequence we obtain the estimator in question, summed over all possible sequences.

That is, if S is the set of all possible sequences, and $\Pr(s)$ is the probability of drawing

any particular sequence (which will be $|\mathfrak{G}|^{-m}$ for sequences of length m), we have:

$$\Pr(\hat{F}_{(m,x)} = x/X) = \sum_{s \in S} \Pr(s) \Pr(\hat{F}_{(m,s,X)} = x/X) \tag{3.74}$$

$$= \frac{1}{|\mathfrak{G}|^m} \sum_{s \in S} \left(\begin{array}{c} X \\ x \end{array} \right) (F_{(m,s)})^x (1 - F_{(m,s)})^{(X-x)} \,. \tag{3.75}$$

To work out the variance of such an estimator, let $\mathbb{E}$ be the mean of the number of $x's$ averaged over all the estimators, then to calculate the variance in the number of x's returned we have:

$$\mathbb{V}(\hat{F}_{(m,s,X)}) = \mathbb{E}(\hat{F}^2_{(m,s,X)}) - (\mathbb{E}(\hat{F}_{(m,s,X)}))^2 \tag{3.76}$$

$$= \mathbb{E}(\hat{F}^2_{(m,s,X)}) - (F_{(m,X)})^2 \,. \tag{3.77}$$

From Equation 3.75 it is clear that we can write $x \sim B(X, F_{(m,s)})$ averaged over s, for the distribution of x, this allows us to use the fact that for a Binomial distribution $B(n,p)$ the variance of the distribution, where we average over n trials, is $\frac{1}{n}p(1-p)$. Here $n = X, p = |\mathfrak{G}|^{-m} \sum_s F_{(m,s)}$. We can use this to write:

$$\mathbb{V}(\hat{F}_{(m,s,X)}) = \frac{1}{X}|\mathfrak{G}|^{-m} \sum_{s \in S} F_{(m,s)} - \frac{1}{X}(|\mathfrak{G}|^{-m} \sum_{s \in S} F_{(m,s)})^2 \tag{3.78}$$

$$= \frac{1}{X}F_{(m)} - \frac{1}{X}(F_{(m)})^2 - \frac{1}{X}\left((|\mathfrak{G}|^{-m} \sum_{s \in S} F_{(m,s)})^2 - (F_{(m)})^2 \right) \tag{3.79}$$

$$= \frac{1}{X}(F_{(m)} - (F_{(m)})^2) - \frac{1}{X}\sigma_m^2 \,, \tag{3.80}$$

where σ_m^2 is the per-sequence variation.

This means that we can write:

$$\mathbb{E}[\mathbb{V}(\hat{F}_{(m,X)})|s] = \mathbb{E}_s[\mathbb{V}(\hat{F}_{(m,s,X)}] \tag{3.81}$$

$$= \frac{1}{X}(F_{(m)} - (F_{(m)})^2) - \frac{1}{X}\sigma_m^2 \,. \tag{3.82}$$

The law of total variance gives us:

$$\mathbb{V}(\hat{F}_{(m,X)}) = \mathbb{E}[\mathbb{V}(\hat{F}_{m,X}|s)] + \mathbb{V}(\mathbb{E}[\hat{F}_{(m,X)}|s]) \,. \tag{3.83}$$

Noting that $\mathbb{V}(\mathbb{E}[\hat{F}_{m,X}|s])$ is σ_m^2, then the variance we need then is:

$$\mathbb{V}(\hat{F}_{(m,X)}) = \frac{1}{X}(F_{(m)} - (F_{(m)})^2) - \frac{1}{X}\sigma_m^2 + \sigma_m^2 \tag{3.84}$$

$$= \frac{1}{X}(F_{(m)} - (F_{(m)})^2) + \frac{X-1}{X}\sigma_m^2 \,. \tag{3.85}$$

If we measure K sequences and average them, then the variance of our estimator for F is

$$\mathbb{V}(\hat{F}_{(m,K,X)}) = \frac{F_{(m)}(1-F_{(m)})}{KX} + \frac{X-1}{KX}\sigma_m^2 \,. \tag{3.86}$$

The first point to notice is that assuming and we have noise then, unless we are afflicted by pure depolarising noise, we can safely assume σ_m^2 is going to be positive. Therefore the relevant variance is minimized by setting $X = 1$, i.e. one shot per new sequence. This is the same result as [77]. Some numerics showing the impact of reducing the variance, especially where the number of sequences measured is limited is shown in figure 3.5.

This formula however also allows us to determine the optimum mixture of numbers of sequences and samples per sequence where there is a cost to changing sequences. If we assume a total cost which is the combined cost of changing a sequence (C_s) and the cost of measuring a sequence (C_r). We note that we will change sequences K times, and we will measure KX sequences, then — for a particular sequence length m — we have:

$$C_{(total,m)} = KC_s + KXC_r \tag{3.87}$$

We can write this as $tK + KX = c$ where t is our unit time measurement unit and c is the total time allocated to the measurement of sequences of that length. Rearranging for K and substituting into Equation 3.86 we get

$$\mathbb{V}(\hat{F}_{(m,t,X)}) = \frac{t+X}{c}\left(\sigma^2 + \frac{F_{(m)}(1-F_{(m)}) - \sigma^2}{X}\right) . \tag{3.88}$$

If we differentiate this with respect to X we can get the optimal choice for X, namely X^*

$$\frac{d\mathbb{V}(\hat{F}_{(m,t,X)})}{dX} = 0 \implies X^* = \frac{\sqrt{t}\sqrt{F_{(m)}(1-F_{(m)}) - \sigma^2}}{\sigma} \,, \tag{3.89}$$

which will always be a minimum, as the sign of the second derivative is always non-negative given that $F(1-F) - \sigma^2 \geq 0$.

There are a number of points to note: (i) we are in a bit of a catch-22 position since we need to know the parameters we are supposed to be measuring in order to work out the best way to measure them; and (ii) this formula is sequence length dependent.

It is, however, possible to estimate all the required parameters, which will give a reasonable indication as to the best way to allocate resources.

To show a practical example of the above, assume we are in an experimental set up where it takes 100 times as long to change a sequence as to measure it (certain superconducting

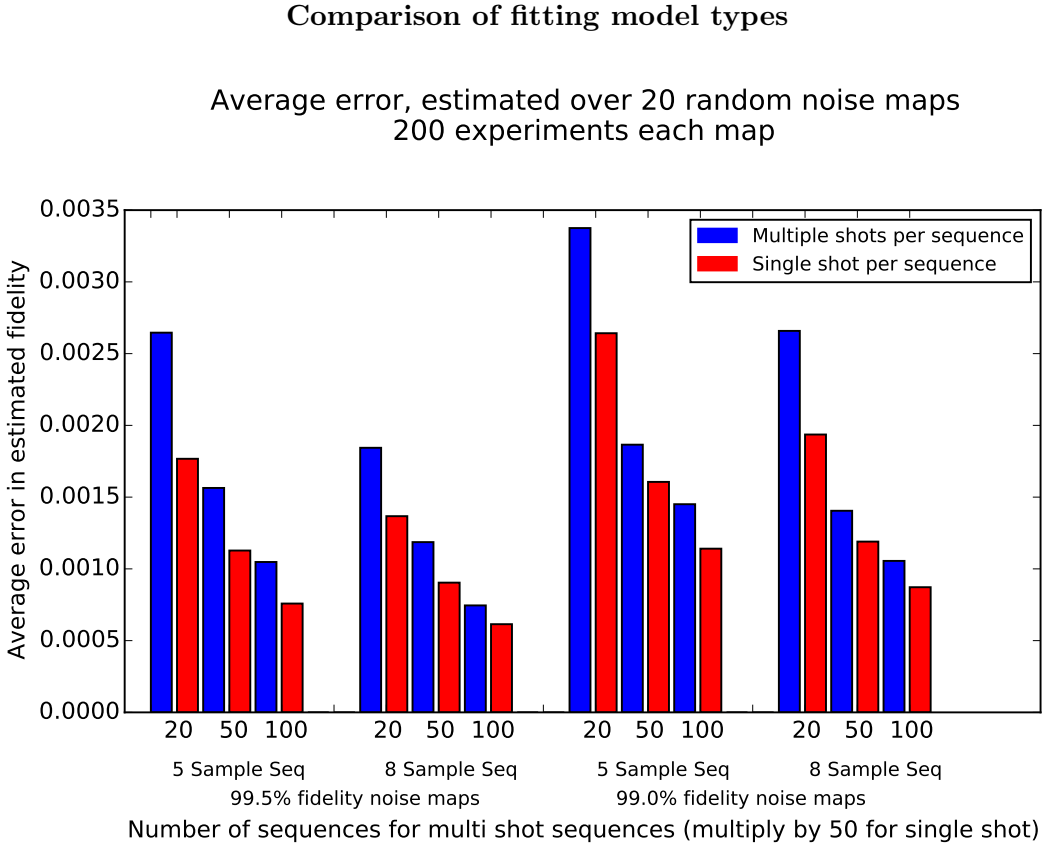

Figure 3.5: Example of the decrease in fitting error where a single shot is used as opposed to multiple measurements of a sequence. The blue bars show the average error in measuring the fidelity where 50 shots per sequence are used, for a varying number of sequence measure points. The red bar shows the average error in fidelity estimation, where the same number of measurements are taken, but each one on a new sequence. The data was averaged over 20 randomly generated noise maps of the required fidelity. For each map the experiment was run 200 times and an estimated fit obtained for each experiment. The average fidelity error is the average of the absolute difference between the estimated value and the actual value for the map in question. A 5 sample sequence commenced measuring at 3 gates, then again at 50 gate increments until gate length 203. An 8 gate sequence incremented by 30 until gate length 213. (Gate lengths 1 and 2 were omitted for reasons discussed in subsection 3.2.5.) Here I have assumed that the cost of taking a new sequence is the same as that of taking a new measurement of a pre-existing sequence. See the text for analysis where this is not the case. Clearly in these numerics there is less error in the fidelity estimation where single shots are measured in 50 times the number of sequences than if each sequence is sampled 50 times. This effect diminished where there are many samples. There were a number of outliers in the first data set that were discarded from the multi-shot data as they occurred in only one of the 20 maps — I believe that these inaccurate estimates may have occurred because the 20 sequences are not sufficient to correctly normalize the distribution in accordance with the CLT for certain noise maps (see text).

Gate length	3	10	30	50	70	90	110	130	150
Estimated F_m	0.9851	0.9522	0.8699	0.8025	0.7474	0.7024	0.6655	0.6354	0.6107
Estimated σ^2	0.0005	0.0016	0.0049	0.0081	0.0114	0.0146	0.0175	0.0175	0.0175
Estimated X^*	54	52	47	43	39	36	34	35	35

Table 3.1: Estimated optimal choice for sequence reuse (X^*) whilst minimising experimental resources. Here I assume that the fidelity of the gates is approximately 99.5% and it takes 100 times as long to re-run a sequence as opposed to changing it.

qubit and spin-qubit experiments fall into a regime such as this). In this case $\sqrt{t} = 10$. Assume we are benchmarking a single qubit and we expect it to have a relatively high fidelity, say we expect:

- $F \cong 0.995$, which means $p \approx 0.99$.
- The sample variance (σ^2) can be estimated using Equation 3.70 and Equation 3.71
- We will assume the two SPAM parameters A, B are 0.5 (little turns on this, assuming we are approximately correct)

Plugging these in we get the values shown in table 3.1.

3.3.5 Curve Fitting, weights and the CLT

From the above it is quite clear that the data returned by randomized benchmarking is heteroscedatic in nature since the variance changes depending on the number of gates in the sequence.

Typically packages used to fit non-linear functions use non-linear least squares regression to fit the curve. Basically this involves fitting the parameters for a specified non-linear function (here $Ap^m + B$) so as to minimize the sum of the squares of the differences between the observed data and the data given by the function to be fit.

The assumptions behind these fitting algorithms typically include constant variance and normal distribution of observed data. Violation of these assumptions can lead to poorly motivated fits, being overly sensitive to data error and/or biased in terms of results and confidence intervals.

As discussed above the variance of the data varies with the gate lengths of the different random sequences, meaning that some form of robust least squares methodology needs to be used. Perhaps the most straight forward way is to weight the observations by the reciprocal of the variance. With single shots, it is clear the observed variance is related to the binomial variance of the fidelity for that sequence length, with multiple shots per sequence the observed variance will be related to the variance in the fidelities of the sequences (which will be smaller). However for weighted least squares the important aspect is the relative weighting of the different data points and in both cases the observed variance should suffice. This was the methodology used in [18].

The second concern is that distribution of survival probabilities over the different sequences is not normally distributed. However, the design of the experiment is such that we are trying to fit the curve to the average survival probability of the sequences and, accordingly, if the central limit theorem (CLT) applies, the averages will be approximately normally distributed. The CLT requires that the samples must be independent and that the sample size must, in some sense, be big enough.

In the single shot regime it is clear that each of the samples are independent but, if we are taking a number of shots/measurements per sequence — each of these measurements will not be independent — although the sequences themselves will be independent. Therefore it is only the data about each sequence (the average survival measured per sequence) that will constitute independent data for determining if the CLT applies. That is, no matter how many times we measure each sequence there will still be a minimum number of sequences we need to measure in order to be able to assume our average sequence fidelity is normally distributed.

As a rule of thumb the CLT requires np and $n(1-p)$ to both be greater than 10. Accordingly near the beginning of the sequence, assuming a survival probability of 0.9, we would need to sample in excess of 100 sequences. Towards the end (where the survival is closer to 0.5) we would need to sample only about 20. Note this requirement is only a rule of thumb as to the minimum number of sequences that must be sampled in order to be able to argue that the CLT applies. For this reason it is important to ensure that the curve fit is to the average survival probabilities for the sequences of a particular length and not to each measured survival probability of each sequence. An example of the application of the CLT is shown in figure 3.6.

3.3.6 Improving the Experimental Design.

Given that we are fitting to three parameters, two of them which are effectively nuisance SPAM parameters, is there a way we can slightly alter the protocol to eliminate, or at least reduce the influence of, one or more of the nuisance parameters?

One method of doing this is to use protocol detailed in [3] (which is attached to this thesis as chapter 8) to eliminate the B parameter — although here I will detail a slight refinement.

In the event we believe we have trace preserving noise and our system has high enough fidelity that we can ignore the gate dependent term (see Equation 3.51)**, then the equation we wish to fit is of the form $\mathbb{E}(\mathcal{F}_m) = Ap^m + B$. If we have a fixed measurement E then B is given by $B = \mathrm{Tr}(E\mathcal{E}\mathbb{1}_d/d)$, where the important point is that B is independent of our choice of starting state ρ.

In this case if we perform randomized benchmarking for two different states ρ and τ, then

**given that there may be a gate dependent term, sequences lengths should avoid $m = 1$ or 2

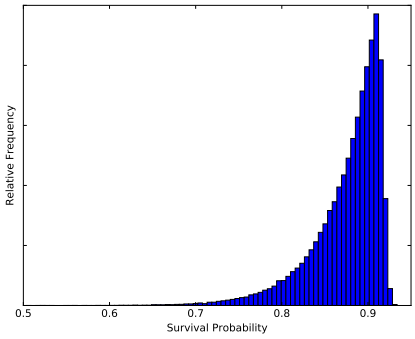

(a) Sequence fidelity distribution for a specific noise map, gate length 100.

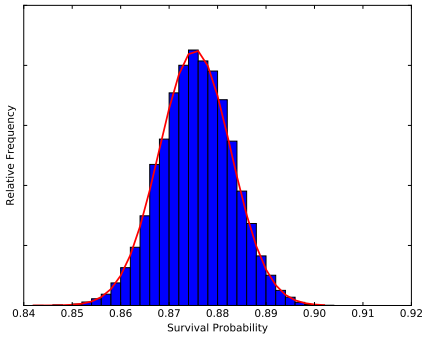

(b) Average sample fidelity distribution for the same noise map, each sample being the average of 1000 single shot sequences. The red line shows a normal distribution with same mean and variance.

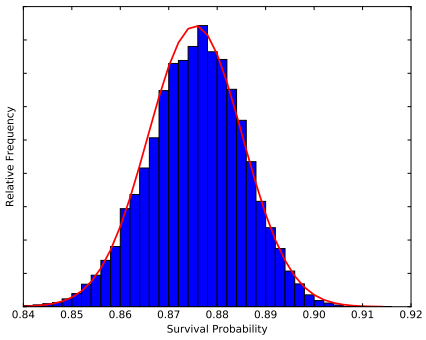

(c) Average sample fidelity distribution for the same noise map, each sample being the average of 40 sequences with 50 shots per sequence. The red line shows a normal distribution with same mean and variance.

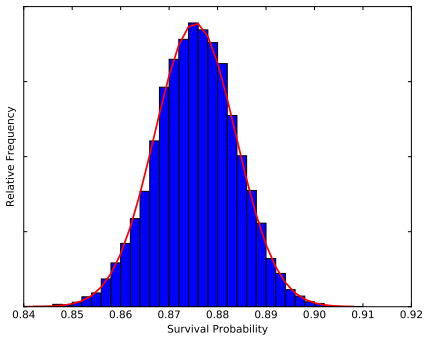

(d) Average sample fidelity distribution for the same noise map, each sample being the average of 100 sequences with 20 shots per sequence. The red line shows a normal distribution with same mean and variance.

Figure 3.6: Subfigure (a) shows the actual sequence distribution for a random high-fidelity noise map. Whilst a typical random noise map might not display such an extreme distribution it has been chosen to make it clear that for this noise map the sequence distribution is not normally distributed. Subfigures (b) through (d) show the effect of the central limit theorem when the sequences are averaged. In sub-figure (b) only a single shot (i.e. a 1 or 0) is recorded for each random sequences, averaged over 1000 sequences. These averaged samples are approximately normally distributed. Subfigures (c) and (d) show similar results where multiple shots are taken per sequence.

for a set number of gates, m we will get estimates of $F_{m,\rho}$ and $F_{m,\tau}$, such that:

$$\mathbb{E}(F_{m,\rho} - F_{m,\tau}) = (A_\rho - A_\tau)p^m \,. \tag{3.90}$$

Using this insight to examine single-qubit randomized benchmarking we see that in the ideal case where $\Lambda\rho = |0\rangle\langle 0|$ and $\Lambda\tau = |1\rangle\langle 1|$ (or any states where the results are orthogonal to each other) then $A_\rho - A_\tau = 1$, which is the maximum possible, since the left hand side of Equation 3.90 is just the difference between two probabilities. A similar result can be achieved if, instead of preparing two different states for one of the benchmarking runs (say, the τ run), we apply an initial gate (which is not inverted by the C_{m+1} gate) transforming $\rho \to \tau$. The SPAM relating to that initial gate will be incorporated into A_τ.

The resulting data set (where one data-set is subtracted from the other) can be fit to a function of the form $\mathbb{E}(F_{\text{diff},m}) = A_d p^m$ (where $A_d = A_\rho - A_\tau$) to achieve an estimate of p. A slightly more rigorous methodology (at the expense of increased fit complexity) is to keep the data-sets separate but to fit the combined data-sets simultaneously to equations of the form:

$$\mathbb{E}(F_\rho) = A_\rho p^m + B \text{ and } \mathbb{E}(F_\tau) = A_\tau p^m + B \,, \tag{3.91}$$

where p and B are held to have the same values over the two simultaneous fittings of the data-sets. This has the advantage of not throwing away any information. In addition we can analyze such a fit by looking at the Fisher information matrix of the equations to determine the influence the parameter B has on our estimation of p.

To analyze the Fisher information matrix, we can recast the equations in Equation 3.91 into the one equation by labelling the samples. Assume we have an indicator parameter i where $i = 1$ if the data belongs to the ρ data-set and $i = 0$ if it is from the τ data-set. In this case we can write:

$$\mathbb{E}(F_m) = ((iA_\rho) - (1-i)A_\tau)\, p^m + B \,. \tag{3.92}$$

This allows us to compare the relevant p/B Fisher information of the combined equation, with twice the Fisher information of the p/B parameters in standard randomized benchmarking (in the latter case we use twice the Fisher information as we need to do twice as many runs to get the combined data set). Such an analysis tells us that using this alternative experimental design strongly unlinks the estimate of p from the parameter B, when compared to the original model.

Numerically we can test this out over random noise maps and over a variety of simulated fidelities. Simulating both methods of experimental design and looking at multiple replications of each experiment we can analyze the variability of the estimate of p. If we repeat each simulated experiment 200 times we get a good idea of the variance of the estimate (in a real experiment, p would only be calculated once). Over a wide range of

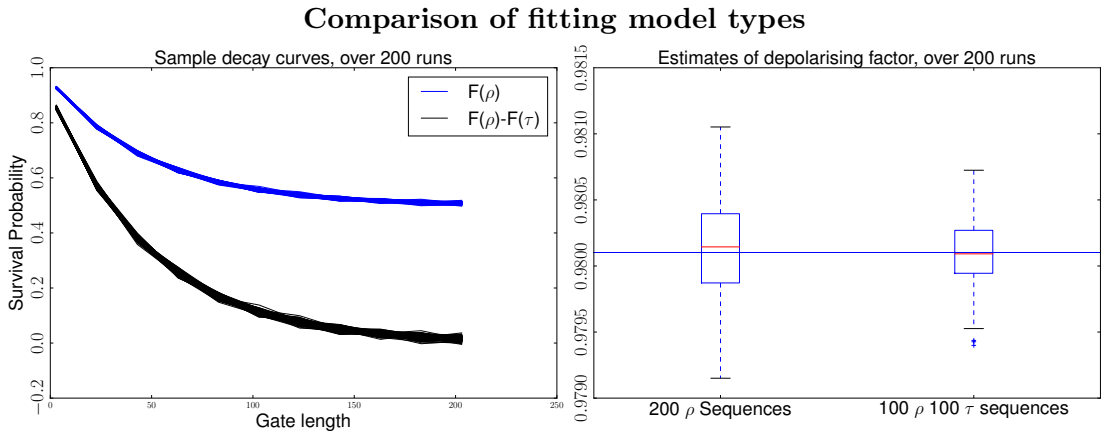

Figure 3.7: Example of the decrease in fitting variability with the alternative design where the input state is varied (here $\rho = |0\rangle\langle 0|$ and $\tau = |1\rangle\langle 1|$). The experiment simulated was to generate a random noisy map, with fidelity = 0.99. Sequences from 3 gates, to 203 gates, every 20 gates were sampled - 200 sequences where the input state was kept as ρ and 100 each where the input state was varied between ρ and τ. The estimated depolarizing factor (p) was calculated using weighted least squares non-linear fit to $Ap^m + B$ for the first case and Ap^m for the second. This was repeated 200 times and the variance in the calculated p plotted in a box-whisker chart. As can be seen the alternative experiment design produces tighter estimates of p. The was true over a wide range of maps and sequence lengths explored.

maps the variability of p was reduced by a factor of between 2 and 4 when this alternative design was used. An alternative way of looking at this is that this translates to a saving of 30%–50% of data to get the same accuracy. There was no noticeable difference between the fidelities calculated using the simpler Ap^m model or the model contained in Equation 3.92. Figure 3.7 shows a typical plot for one of the random maps used.

3.3.7 Bayesian techniques with Randomized Benchmarking

In [77] Granade et al. present a methodology of using Bayesian techniques together with a sequential Monte Carlo (SMC) algorithm to obtain accurate estimates of the randomized benchmarking parameters. In very broad terms, SMC involves a swarm of candidate solutions (each a hypothesis) representing the posterior distribution of a stochastic process (here the random benchmarking process). The data from measured sequence lengths is used to update the hypotheses according to Bayes' rule, this can be done efficiently. They note a number of advantages of this approach including (1) that the SMC can accurately characterize its own performance, obtaining accuracy close the ultimate theoretical Bayesian Cramer-Rao bounds; and (2) their methodology can used to update estimates as the data is generated — potentially allowing the fidelity of a system to be estimated and updated as the measurements are being taken. In theory this could allow real-time fidelity estimation. Real-time fidelity estimation might, for

instance, allow the system to be altered and resulting changes to fidelity monitored or it might allow the observation of fidelity drift over time, for instance if there were time dependent noise factors acting in a time frame of many runs of the experiment. There has not been much work adapting the frequentist models discussed above to allow such updating of the hypothesis and thus the estimate. The numerical part of [77] confirms in simulation the rapid convergence of the algorithm presented to the theoretical bounds.

3.4 Logical Randomized Benchmarking

As the number of qubits increase and researchers start to apply error correcting codes to such qubits to create error corrected *logical qubits*, the question is now being asked as to whether randomized benchmarking can be applied to such logical qubits. This is analysed in [78]. Combes et al. propose a specific model for applying randomized benchmarking in conjunction with an error correction protocol. In particular they note that by applying randomized benchmarking at a logical level it becomes possible to benchmark, not only total error, but also the errors that have been corrected, increasing the amount of diagnostic information available to researchers.

The paper also highlights several important matters relating to logical randomized benchmarking:

- the fidelity of single physical qubits, or even multiple physical qubits, is not an honest predictor of the fidelity that might be achieved by implementing a logical qubit over such physical qubits. In particular the authors provide examples in which the logical fidelity rate would be both underestimated and overestimated if various types of errors exist on the underlying physical system;
- the particular choice of error correcting regime will have an impact on the analysis of the randomized benchmarking protocol. In particular for the authors' choice of protocol they analyse the model to ensure that the underlying assumptions and analysis of RB apply in the logical setting;
- logical randomized benchmarking is more efficient than attempting to recover the same information by running RB and the various RB related protocols that would be required to accurately predict the logical fidelity; and
- logical randomized benchmarking does, however, require additional assumptions compared to those for standard randomized benchmarking. In particular in order to complete the analysis the authors found it necessary to assume that the noise on preparation and measurement could be written as a tensor product of channels supported on the logical and syndrome registers.

Chapter 4

Other uses for the RB Protocol

"It was the day my grandmother exploded."

Iain Banks - The Crow Road.

4.1 Characterization of Noise

4.1.1 Discussion of Unitarity

Unitarity or *purity benchmarking* (as it has subsequently sometimes been referred to) is a method of using the concepts behind the randomized benchmarking protocol to gain some understanding of the coherence of the noise in the system being tested.

Whilst I will leave most of the details of the protocol to be explained in the attached paper (see chapter 9) (published as [4]), I will cover the basics of the definition here, to aid a discussion of subsequent papers and work that has built upon the basic concepts of unitarity.

If we were to analyze the noise processes in a quantum system, one could broadly characterize them as being coherent or incoherent. An example of coherent noise might be rotation errors, incoherent noise might be a depolarizing of the channel. Coherent noise mechanisms can lead to large fluctuations in the apparent fidelity of the gates of the system, admitting sequences of gates where the errors combine. With randomized benchmarking we have an idea of the average fidelity of the system, but there is no indication as to whether errors are close to the average, or if there might be a worst-possible case error that is relatively far from the average. When we are looking to determine if the noise and errors in an experimental implementation are small enough to be consistent with fault tolerant computation, such information is vital. Conversely, although coherent noise can lead to larger worst-case errors, such noise is also likely to be easier to rectify once it has been adequately characterized. In all, being able to characterize the type of noise in the system is useful.

4.1.2 Definition and Protocol

The definition of unitarity is

$$u(\mathcal{E}) = \frac{d}{d-1}\int d\psi \mathrm{Tr}\left[\mathcal{E}'(\psi)^\dagger \mathcal{E}'(\psi)\right], \tag{4.1}$$

where we define $\mathcal{E}'(A) = \mathcal{E}(A) - [\mathrm{Tr}(\mathcal{E}(A))/\sqrt{d}]\mathbb{1}$ for all traceless A and $\mathcal{E}'(\mathbb{1}_d) = 0$. That is it is the average change in the purity of a pure state after applying the noise channel, with the contributions due to the identity component subtracted off.

Using the Pauli Liouville basis, it is relatively easy to see that the relevant part of the channel is the $\mathcal{E}_\mathrm{u}$ block (see Equation 2.75) and the unitarity of the channel $\mathcal{E}$ is simply

$$u(\mathcal{E}) = \frac{1}{d^2-1}\mathrm{Tr}(\mathcal{E}_\mathrm{u}^\dagger \mathcal{E}_\mathrm{u})\,. \tag{4.2}$$

To measure this we need to average over all pure input states, which means that we can use a unitary 2-design to do the averaging (subsection 2.3.1). As with randomized benchmarking we can use the Cliffords, or any other unitary 2-design. Unlike randomized benchmarking we do not need to invert the gates at the end of sequence, which means that our unitary 2-design does not need to be a group nor does the sequence of gates, or their inverse, need to be classically computable.

The protocol then is very similar to the randomized benchmarking protocol. An initial state ρ is prepared, a number of random sequences of our unitary 2-design are then applied, varying numbers of gates to generate a decay curve which we then fit.

The measurement we use for the protocol is the expectation value Q_j of an operator Q and (as proved in [4]) the expected value, for trace preserving noise, of Q_j^2 over all random sequences will be of the form:

$$\mathbb{E}_j[Q_j^2] = A + Bu(\mathcal{E})^{m-1}\,. \tag{4.3}$$

Although the expectation of an operator Q is mathematically equivalent to the probability of a single outcome (such as the type of measurement used in randomized benchmarking), there is a question of increased variability in the observations (requiring a commensurate increase in sequences sampled) if we attempt to reconstruct it by using only one observable (since we don't have an inversion gate, there will be sequences where the overlap between the measurable and the the state is very low, making accurate estimation difficult). However, we can reduce this by averaging over multiple observables, for example averaging over the non-identity Pauli operators. Ideally we would use maximum overlap measurement, which involving the swap operator S, but as shown in [4] we can simulate this by calculating $\langle X\rangle^2 + \langle Y\rangle^2 + \langle Z\rangle^2$. Although this protocol is not scalable given that there are exponentially many n-qubit Paulis, for a single qubit system it is quite practical and indeed has been implemented [79]. An alternative, scalable,

method of realising this is if one can have two copies of the system and a SWAP is performed at the end. In this case there is only one observable to measure.

For a single qubit system, the protocol thus becomes — for each random sequence the experiment is performed 3 times, each different run of that sequence being measured in the X, Y and then the Z basis, with the square of the expectation values summed.

The easiest way to visualize what is going on is to think of the Bloch sphere picture of the single qubit. By measuring $X^2 + Y^2 + Z^2$ of a particular state (after it has been twirled by the sequence) we are simply determining the length of the line from the centre of the Bloch sphere to the twirled state. Averaged by the unitary 2-design this will give us the average purity of the channel i.e. the average distance to the edge of the twirled Bloch sphere — that is the unitarity. In density operator terms this is just the normalized trace of the square of the density matrix, with the identity subtracted i.e. $2\mathrm{Tr}(\rho^2) - 1$

4.1.3 Impact of Unitarity

One of the important uses of unitarity is that it allows us to find bounds that relate the average fidelity to the diamond norm.

The diamond distance

One of the reasons the diamond norm is of interest is because of the fault-tolerance threshold theorem. This theorem gives a (theoretical) guarantee that quantum computers can be built if the noise strengths and correlations afflicting a quantum information process are below a threshold value [11, 80, 81]. Unfortunately the experimentally measurable average fidelity is not the same as the (in)fidelity required to be within the fault-tolerance threshold. The average fidelity does not give us a direct estimate of the worst-possible fidelity — it is, after all, an average — and it is the worst possible case that is required for the fault tolerant guarantees. Ref. [82] emphasizes the distinction providing specific examples where the average fidelity is high, but the worst case error is also large. Average fidelity alone does not provide sufficient guarantee that the fault-tolerant threshold has been met.

The diamond distance has a formal definition as follows:

$$\mathcal{D}(\mathcal{E}, \mathcal{F}) = \frac{1}{2} \max_{\rho} ||\mathcal{I} \otimes \mathcal{F}(\rho) - \mathcal{I} \otimes \mathcal{E}(\rho)||_1 \,, \tag{4.4}$$

where the reason it is tensored with an Identity channel is so that the measure is stable under entangled inputs (contrast with the trace distance which is also analyzed in [83]). This has an operational interpretation as the maximum probability of distinguishing the output of a noisy gate from the ideal output. Whilst there is no efficient way to determine this experimentally (see for instance [21] as to how this might be attempted

with gate set tomography), it can be computed efficiently using methods of semi-definite programming (see [84] for a discussion of this).

Until about 2015, [38] contained the best known bounds for a d-level quantum gate:

$$\frac{d+1}{d}r \leq \mathcal{D} \leq \sqrt{d(d+1)r}\,. \tag{4.5}$$

Importantly and rather worryingly the upper bound scales with the square root of the infidelity. Given that we need to be in the regime where the diamond norm is somewhere in the order of 10^{-2} to 10^{-4}, depending on error correcting schemes and analysis, this would entail achieving an average fidelity in the region of 10^{-4} to 10^{-8}. Little more could be known analytically without more information about the type of noise.

Using the unitarity of the noise

In two papers that came out at roughly the same time ([84] and [83]) it was shown that the unitarity of the noise could be used to further bound the relationship of the average fidelity to the diamond distance.

As previously noted in [4] (a more general, concise proof is also given in [84]) the unitarity and the average fidelity are related as follows (here $r(\mathcal{E})$ is the infidelity, being 1-fidelity):

$$1 \geq u(\mathcal{E}) \geq \left(1 - \frac{dr(\mathcal{E})}{d-1}\right)^2 . \tag{4.6}$$

The bounds found in [83] for a trace preserving quantum channel are:

$$\frac{C}{\sqrt{2}} \leq \mathcal{D} \leq \sqrt{\frac{d^3C^2}{4} + \frac{(d+1)^2 r(\mathcal{E})^2}{2}}\,, \tag{4.7}$$

where:

$$C^2 = \frac{d^2-1}{d^2}(u(\mathcal{E}) - 2p + 1)\,, \tag{4.8}$$

p here being the randomized benchmarking decay constant (i.e. the depolarizing factor).

The immediate impact of this is that we can use the unitarity to work out if we have non-favourable scaling on the diamond distance (compared to the average fidelity).

Since the unitarity can be estimated from benchmarking experiments (see for instance [79]) this may be an experimental way to determine the type of the noise afflicting the system and the corresponding bounds on diamond distance errors.

Unitarity and interleaved randomized benchmarking

As discussed in section 3.2 one of the problems relating to interleaved benchmarking is that for a generalized CPTP channel the bounds on the underlying assumption (that $F(\mathcal{E})F(\Lambda) \approx F(\mathcal{E}\Lambda)$) is generally only true within some relatively loose bounds, unless one of the channels is close to a depolarizing channel.

In [49] they show that the unitarity can be used to improve the bounds discussed in section 3.2.

The define a (positive) coherence angle as

$$\theta(\mathcal{E}) = \arccos\left(\frac{p(\mathcal{E})}{\sqrt{u(\mathcal{E})}}\right). \tag{4.9}$$

Since the unitarity is always less than or equal to 1, this give us the following:

$$\theta(\mathcal{E}) = \begin{cases} 0 & \text{iff } \mathcal{E} \text{ is depolarizing} \\ O(r) & \text{iff } \mathcal{E} \text{ is Pauli noise} \\ \arccos p(\mathcal{E}) = O(\sqrt{r}) & \text{iff } \mathcal{E} \text{ is unitary.} \end{cases} \tag{4.10}$$

To give the following bound for two quantum channels, $\mathcal{X}$ and $\mathcal{Y}$, when the unitarity is known:

$$|p(\mathcal{X}) - \gamma_{xy}\cos[\theta(\mathcal{Y})]| \leq \sin[\theta(\mathcal{Y})\sqrt{1-\lambda_{xy}^2}], \tag{4.11}$$

where

$$\lambda_{xy} = p(\mathcal{X}\mathcal{Y})/\sqrt{u\mathcal{Y}}. \tag{4.12}$$

4.2 Loss rates and leakage

The final additional variation to the randomized benchmarking protocol I shall talk about relates to loss rates and leakage errors.

Loss rates relate to the possibility that the qubit might be lost, which can be a substantial obstacle to many quantum information protocols such as error correction protocols. In [85] a variation to the RB protocol is introduced. It is similar to the unitarity protocol (which it precedes), in that there is no returning gate, rather a unitary 1-design (or higher) is used to twirl the maximally mixed state which is measured by the trivial (identity) measurement. Loss rates then show up as an exponential decay curve or, if the system is also afflicted by trace decreasing noise, a double exponential decay curve.

Leakage errors differ from loss rates as here we are contemplating not the irretrievable loss of the qubit, but rather the leakage of the subspace of interest, into a different subspace where it might continue to evolve coherently and may even re-enter the subspace in question. The was the subject of two relevant papers. In [86], a specific heuristic was presented applying to superconducting qubits and in [87] they analyze how the results of the protocol in [85] can be adapted to deal with the situation of leakage errors.

Part III

Automatic Learning

Chapter 5

Automatic Learning

> "It is well known that a vital ingredient of success is not knowing that what you're attempting can't be done."
>
> Terry Pratchett.

5.1 Brief overview of automatic learning and symbolic regression

Automatic learning or, as it is often known, machine learning is a field of computer science where the algorithm learns from the data without being explicitly programmed. Typically the methodologies used mean the algorithms are classified as *weak problem solvers*. A weak problem solver is a solver that has not been pre-programmed with information about the problem domain, as opposed to a *strong problem solver*, where human insight has been used to pre-program the solver with the information it needs to solve that particular task. A classic example of the latter are chess programs (the problem here being to beat the opponent) that have many hard coded opening and closing strategies. Of the former are the type of image classification neural networks championed by Google and made famous through its release of the DeepDream computer vision program.

Weak problem solvers tend to be more general, better able to adapt when the area being probed is not fully understood and are capable of being used to find patterns and to make connections in areas where expertise is thin on the ground (or non-existent). It is the use of such techniques that have allowed present day inroads in speech recognition, image classification and improved search — to name just a few areas. A specific class of automatic learners, evolutionary programs, are well suited to finding patterns and making relations in a way that sometimes provides insight as to how this is achieved. This can be contrasted with, say, neural networks where little insight is achieved as to why the network behaves as it does. These evolutionary programs are a class of learners that seek to harness algorithms inspired by biological evolution. Although typically classified as weak problem solvers (and thus being adaptable), some domain specific

information can be encoded in the *fitness function* that forms part of the algorithm, although in areas that are not well understood it is normally inadvisable to do so.

In relation to quantum computing there is much interest in how a quantum computer might speed up the learning process of such machines, whether allowing fully connected Boltzmann machines to be trained (as opposed to the restricted Boltzmann machines currently being used) [88, 89] or taking advantage of other quantum features, such as allowing training in noisy environments, where classical methods might fail ([90]).

This section looks, however, at the other end, namely how classical machine learning might help us to better characterize early quantum information systems as we strive to build a quantum computer. There are several ways this can be done.

5.2 Neural networks

One of the main benefits of neural networks to this field is their ability to compress high-dimensional data. The architecture of deep networks, coupled with the correct training regimes allows such networks to automatically extract the features of the data most relevant to explaining the data — allowing patterns that would otherwise be obscured to emerge. It is through this compression of information that such networks might lead to greater understanding of the underlying data.

There have been a number of papers where neural networks have been used to reduce the dimension of the underlying problem and thus make tractable the problem of quantifying the system. By way of example in [91] convolutional neural networks were trained to identify and locate quantum phase transitions in quantum many-Fermion systems and in [92] a restricted Boltzmann machine was used to model many-body quantum states, trained from simple measurements.

5.3 Evolutionary algorithms

Like neural networks, evolutionary algorithms can be thought of as an optimisation problem, although in this case the representation of the problem is decoupled from the optimisation algorithm. There are two main types of evolutionary algorithm pertinent to this thesis. Because evolutionary algorithms *search* for solutions (as opposed to deriving or learning them) they are particularly well suited to domains where there might be a Pareto front of solutions where multiple and contra-indicative criteria have to be satisfied.

5.3.1 Genetic algorithms

In a genetic algorithm the evolutionary algorithm is used to evolve a set of candidate solutions towards better solutions. To utilise the algorithm once needs to be able to represent the solution domain in a suitable way, for example as a string of numbers or as an array of bits. One also needs to be able to create a *fitness function* which will take a potential candidate solution and assess it in some way, such that solutions which are in some way closer to solving the problem receive a better score from the function. This function provides the *fitness gradient* that drives the search process. For instance if you were evolving parameters for a model then, perhaps, the fitness function would be based on how well a candidate model explains the data seen.

In [2] such a technique was used to search for and find the specific correlations that formed a Bayesian directed graph. Despite the complexity of the underlying search landscape, the search was able to find the full Pareto front of relevant distributions.

5.3.2 Symbolic Regression

Genetic programming is a technique for evolving computer programs, originally represented using tree-based structures, to solve problems. Initially introduced by [93] it was popularized and greatly expanded upon by John Koza, the canonical form appearing in his first book *Genetic Programming: On the Programming of Computers by Means of Natural Selection* [94]. Symbolic regression as currently practiced grew out of this form of genetic programming and uses such evolutionary techniques to evolve equations that contain the correct symbols to map input data to the output data set. So instead of, say, teaching a network to map inputs to outputs, symbolic regression attempts to find the underlying equation that will do such mapping.

Extensions of genetic programming, such as grammatical evolution [95], allow the syntax of the expression to be evolved to be more precisely specified, including the ability: (i) to allow the evolution of arbitrary sub-functions [96]; (ii) to specify balanced grammars — enabling greater coverage of the search space [97]; and (iii) to use various techniques such as co-evolution to keep the evolved expression size as small (and therefore as comprehensible) as possible [98].

Such techniques have been applied to physics problems before, for instance in [99] a symbolic regression system was used to derive some of the fundamental equations underlying physical systems such as a double-pendula. It seems likely that we should be able to harness symbolic regression in some aspects of QCVV.

One potential candidate is the recently publicized disconnect between metrics such as the fidelity or diamond distance (to the identity) of an error channel and the ability for a concatenated error correcting code to correct the error. Similar metrics on different error channels mapped to corrected error rates that varied over many orders of magnitude [100]. A currently active area of research is to determine if symbolic regression

might be able to discover a metric based on the noise matrix that might better predict the ability of such an error code to correct errors occurring as a result of such noise. Chapter 7 contains a paper containing a prototype genetic algorithm that we hope to apply to problems such as this in the near future.

Part IV

Papers

Chapter 6

Estimating the fidelity of T gates using standard interleaved randomized benchmarking

Randomized benchmarking (RB) is an important protocol for robustly characterizing the error rates of quantum gates. The technique is typically applied to the Clifford gates since they form a group that satisfies a convenient technical condition of forming a unitary 2-design, in addition to having a tight connection to fault-tolerant quantum computing and an efficient classical simulation. In order to achieve universal quantum computing one must add at least one additional gate such as the T gate (also known as the $\pi/8$ gate). Here we propose and analyze a simple variation of the standard interleaved RB protocol that can accurately estimate the average fidelity of the T gate while retaining the many advantages of a unitary 2-design and the fidelity guarantees that such a design delivers, as well as the efficient classical simulation property of the Clifford group. Our work complements prior methods that have succeeded in estimating T gate fidelities, but only by relaxing the 2-design constraint and using a more complicated data analysis.

IOP Publishing Quantum Sci. Technol. 2 (2017) 015008 https://doi.org/10.1088/2058-9565/aa5f8d

Quantum Science and Technology

RECEIVED
24 August 2016

REVISED
29 January 2017

ACCEPTED FOR PUBLICATION
9 February 2017

PUBLISHED
9 March 2017

PAPER

Estimating the fidelity of T gates using standard interleaved randomized benchmarking

Robin Harper[1] **and Steven T Flammia**

Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, Australia
[1] Author to whom any correspondence should be addressed.

E-mail: robin.harper@sydney.edu.au

Keywords: quantum computing, fidelity of T-gates, interleaved benchmarking, randomized benchmarking

Abstract

Randomized benchmarking is an important protocol for robustly characterizing the error rates of quantum gates. The technique is typically applied to the Clifford gates since they form a group that satisfies a convenient technical condition of forming a unitary 2-design, in addition to having a tight connection to fault-tolerant quantum computing and an efficient classical simulation. In order to achieve universal quantum computing one must add at least one additional gate such as the T gate (also known as the $\pi/8$ gate). Here we propose and analyze a simple variation of the standard interleaved RB protocol that can accurately estimate the average fidelity of the T gate while retaining the many advantages of a unitary 2-design and the fidelity guarantees that such a design delivers, as well as the efficient classical simulation property of the Clifford group. Our work complements prior methods that have succeeded in estimating T gate fidelities, but only by relaxing the 2-design constraint and using a more complicated data analysis.

Randomized benchmarking (RB) [1–4] uses long sequences of quantum gates with the aim of amplifying small errors in the implementation of the gates, providing a scalable method for quantifying these errors. One key advantage of RB over quantum process tomography is that it is robust to errors associated with state preparation and measurement (SPAM) noise, so that it is able to isolate the contribution of the noise due solely to the gates. RB is also substantially more scalable than quantum process tomography, and these crucial advantages have made it an indispensable tool for quantum computing experiments.

The Clifford group plays a central role in the theory of RB, and there are several theoretical motivations for this. First, it provides a natural class of circuits that can be simulated efficiently on a classical computer. Moreover, most fault-tolerant architectures make extensive use of Clifford circuits, so these are precisely the types of gates that will likely appear in implementations. The Clifford group also satisfies the technical condition of being a unitary 2-design [3], meaning that the average over the Clifford group transforms general noise sources into just depolarizing noise. (See section 3 and [1, 5] for more details.) Finally, standard RB estimates the average error in an ensemble of Clifford gates, and the method of interleaved RB [6, 7] enables the estimation of average errors on individual Clifford gates, which yields quite detailed information about the errors in a Clifford circuit given the low cost of the method.

It is important to be able to estimate the fidelity of the physical realization of non-Clifford gates as well, because while Clifford gates play an important role in current approaches to fault-tolerant quantum computing, they do not allow for universal quantum computation. One must add at least one non-Clifford gate to achieve universality, and the most common choice of additional gate is the so-called T gate, or $\pi/8$ gate.

Several protocols in the literature have expanded the notion of RB so that it can efficiently estimate the average error of T gates [8–14]. We will review these methods in the subsequent section.

Our contribution is a method for estimating the average error on T gates that complements these other protocols. We propose and analyze a simple variation of the standard interleaved RB protocol that, under the usual RB assumptions explained below, at least diagonalizes the noise between subsequent T gates. This renders the protocol more robust to coherent errors on the T gate. As a consequence this extension is a method of

measuring the average fidelity of a T gate in a scalable fashion, while preserving the advantageous unitary 2-design property up to second-order corrections and using only elements in the Clifford group (except for the interleaved T gates). In fact, our method works on a slightly more general set of gates that includes T as a special case, as we discuss below. Our approach also provides a simplification with respect to the fit parameters compared to existing methods and, within the assumptions made, retains the strong theoretical guarantees on the error associated with RB [15–18].

1. Prior work

The unitary 2-design property mentioned above is needed to average out generic errors in gates, but recent work by Carignan-Dugas *et al* [10] and Cross *et al* [11] has shown that something nearly as good can be obtained by relaxing this condition while enabling estimation of T-gate error rates. Instead of the Clifford group, they use a dihedral group for single qubits or a CNOT-dihedral group for multiple qubits and demonstrate a practical method to benchmark T gates (and other gates not in the Clifford group) using this approach.

Since these protocols involve averages over gate sets that are not unitary 2-designs, there are several points that distinguish them from standard RB or interleaved RB. The dihedral-type groups are not sufficient to completely average a generic error into a depolarizing channel and require fitting to two exponential decay parameters, although mechanisms exist to allow the extraction of the individual decay parameters by analyzing the differences between the average survival probabilities of differently constructed runs.

In [8], a method of using RB to calculate the fidelity of gates by measuring the overlap between a specific Clifford channel and any target unitary was introduced. This was the first method that demonstrated the idea of using RB to benchmark any unitary and consequently is more flexible than the protocol described in this paper. It alters the RB protocol slightly to change the channels generated to channels representing specific Clifford 'maps' and allows the overlap between these Clifford maps and the target gate to be measured and thus the target unitary to be reconstructed. However, for a single qubit T gate it requires the fidelity of the overlap of three composed channels to the T gate to be estimated and then these overlaps are 'best-fit' to obtain an estimated T gate fidelity. In addition the overlap between the specific Clifford channels and the T gate will not be near unity, increasing the binomial variance of the results—meaning many more samples are required for any specific accuracy.

In [9], two T gates were inserted between each Clifford element, so that the fidelity of a double application of a T gate could be estimated. One can then extract the average error on a single T gate if one is willing to assume that the errors are not correlated and compose in a straightforward manner. This scheme has the advantage of being a unitary-2 design, but compounds the estimation errors that occur when attempting to decompose error sources from the estimation of the composed errors.

Finally, in [12] a protocol is presented that obviates the need to calculate an inversion gate at the end of the sequence by combining RB with the use of Monte Carlo sampling for fidelity estimation [13, 14]. While such a protocol allows for fidelity estimation for arbitrary quantum gates it scales exponentially in the number of qubits.

2. Protocol

RB uses the Clifford group (or any unitary 2-design) to efficiently estimate the average fidelity of a noise map $\mathcal{E}$ associated with this group of operations. Here we are implicitly assuming that this noise map obeys various convenient simplifying assumptions, namely that the noise is time-independent and gate-independent, such that the noise really is characterized by a single map $\mathcal{E}$. Below we will discuss what happens when these simplifying assumptions are relaxed. The average fidelity of $\mathcal{E}$ is defined as

$$\mathcal{F}_{\text{avg}}(\mathcal{E}) \coloneqq \int \mathrm{d}\psi \, \langle\psi|\mathcal{E}(\psi)|\psi\rangle, \tag{1}$$

where the average is over the unitarily invariant Haar measure. The average fidelity of a noise map can be efficiently estimated by twirling the error operator over the Clifford group (or indeed any unitary 2-design). The twirling operation over a group $\mathcal{G}$ is defined as simply performing the group average, yielding

$$\mathcal{E}^{\mathcal{G}} = \frac{1}{|\mathcal{G}|} \sum_{g\in\mathcal{G}} g^{\dagger}\mathcal{E}g, \tag{2}$$

where g is the unitary operation corresponding to a group element. This operation of twirling preserves the average fidelity of the noise map $\mathcal{E}$, and even more, it outputs a pure depolarizing channel that can be used to estimate the average fidelity by doing a simple fit to an exponential decay curve. As alluded to above, the advantages of using the Clifford group in this twirling operation include that Clifford gates are relatively easy to

implement and that the gates form a group so that their operation can be efficiently inverted after a random sequence of such operators. The partial information provided by the average fidelity is useful in practice for improving the implementation of the gates as well as providing a bound on the threshold error rate required for fault-tolerant quantum computing [16, 18, 19].

The set of gates consisting of any Clifford group element C followed by T, a T gate, is a unitary 2-design. This is intuitively clear given that the Clifford group itself forms a unitary 2-design, and it is easily verified by calculating the frame potential of the gate set [20]. Consequently, this TC gate set can be used in some related protocols that require a unitary 2-design, such as measuring unitarity [21] or measuring loss rates [22] (which in fact only requires a weaker condition of a unitary 1-design). However, this set of gates does not form a finite group. In fact its image is dense on the entire unitary group, so there is no scalable method of inverting the sequence of gates at the end of a benchmarking sequence and it becomes difficult to use the gates to measure average fidelity in an accurate and precise way.

We now describe the protocol that allows T gates to be used with the RB protocol while preserving the unitary 2-design property. First, we note that the T gate is a gate in the third level of the Clifford hierarchy [23]. This means that it conjugates members of the Pauli group to the Clifford group (ignoring any overall global phases), namely for every Pauli gate P there exists a Clifford C such that

$$TPT^{\dagger} = C. \tag{3}$$

For an arbitrary number of qubits, we will write $\mathcal{C}$ to represent the normal Clifford gates, $\mathcal{P}$ to represent the Pauli operators and $\mathcal{C}_t$ to represent the subset of Clifford gates formed by conjugation of the elements of $\mathcal{P}$ with a T gate.

In this protocol, gates from the Clifford group (or a subgroup of the Clifford group, as discussed in section 5) form the basis of our unitary-2 design. We will denote this group of Clifford gates by $\mathcal{G}$, and the full gate set we use will be any gate of the form TC for $C \in \mathcal{G}$. In line with standard RB assumptions we assume that an experimental implementation of a gate $U \in \mathcal{G}$ can be written as $\mathcal{U}\mathcal{E}$ where $\mathcal{U}$ denotes the channel corresponding to conjugation by U and $\mathcal{E}$ is a completely positive trace preserving (CPTP) channel independent of U. We also note that the assumption that $\mathcal{E}$ is independent of U can be relaxed without significant effect on the results [4, 22, 24]. Specifically, since the Pauli gates are a subset of the Clifford group, we are assuming the error on the Pauli gates is the same as for the other Clifford gates (subject to the relaxation noted above). With respect to the interleaved gate (here the T gate), it is anticipated that the error channel associated with that gate will be different from those forming $\mathcal{G}$, and we will treat this error separately.

Our T gate interleaved RB protocol is very similar to the usual interleaved RB procedure. We begin by determining the reference fidelity using a slight modification of standard benchmarking. But first some notation: if $\mathcal{G}$ is a group, then $|\mathcal{G}|$ denotes the total number of elements in $\mathcal{G}$. We also say, by an abuse of notation, that $|\mathcal{G}|$ is the set $\{1, \ldots, |\mathcal{G}|\}$. Thus, a statement such as $j \in |\mathcal{G}|$ simply means that j is a label for the jth element in $\mathcal{G}$.

2.1. Determining the reference fidelity

1. For an even integer $m = 2n$, choose a sequence $\mathbf{j} = (j_1, \ldots, j_n)$ for $j_k \in |\mathcal{G}|$ and a sequence $\mathbf{p} = (p_1, \ldots, p_n)$ for $p_k \in |\mathcal{P}|$, both uniformly at random.
2. For each sequence
$$U_{\mathbf{j}}P_{\mathbf{p}} = U_{j_n}P_{p_n}U_{j_{n-1}}P_{p_{n-1}} \ldots U_{j_1}P_{p_1}$$
determine the gate U_{inv} which is $(U_{\mathbf{j}}P_{\mathbf{p}})^{\dagger}$
3. Apply the sequence
$$U_{\text{inv}}U_{j_n}P_{p_n}U_{j_{n-1}}P_{p_{n-1}} \ldots U_{j_1}P_{p_1}$$
to some initial state $\rho \neq \frac{1}{d}\mathbb{1}$ (usually taken to be the pure state $|0\rangle$) and perform a POVM $\{E, \mathbb{1} - E\}$ for some E (usually taken to be $|0\rangle\langle 0|$). The important point being that the ideal measurement aligns, so far as possible, with the ideal initial state to minimize the binomial variance in the results.
4. Repeat steps 1–3 sufficiently many times to estimate the survival probability to a desired precision over the randomized sequences. (See [15, 25] for guidance on choosing a sufficient number of samples.)
5. Repeat steps 1–4 for different values of m and fit the results for the averaged sequence fidelity ($\bar{F}_{\text{ref}}$) to the standard RB model:

$$\bar{F}_{\text{ref}}(m) = A_0 p_{\text{ref}}^m + B_0, \tag{4}$$

where A_0 and B_0 are constants that absorb SPAM errors.

2.2. Determine the interleaved fidelity

1. For an even integer $m = 2n$, choose a sequence $\mathbf{j} = (j_1, \ldots, j_n)$ for $j_k \in |\mathcal{G}|$ and a sequence $\mathbf{p} = (p_1, \ldots, p_n)$ for $p_k \in |\mathcal{P}|$, both uniformly at random.
2. For each sequence
$$U_{\mathbf{j}} TP_{\mathbf{p}} T = U_{j_n} TP_{p_n} TU_{j_{n-1}} TP_{p_{n-1}} \ldots U_{j_1} TP_{p_1} T$$
determine the gate U_{inv} which is $(U_{\mathbf{j}} TP_{\mathbf{p}} T)^\dagger$
3. Apply the sequence
$$U_{\text{inv}} U_{j_n} TP_{p_n} TU_{j_{n-1}} TP_{p_{n-1}} T \ldots U_{j_1} TP_{p_1} T$$
to the initial state ρ used in A.3 and measure as chosen in A.3.
4. Repeat steps 1–3 sufficiently many times to estimate the survival probability to a desired precision over the randomized sequences.
5. Repeat steps 1–4 for different values of m and fit the results to the model
$$\bar{F}_T(m) = A_1 p_T^m + B_1. \tag{5}$$

We note that for the purpose of calculating the inverse gate in step 2 of the interleaved protocol that each P_iT can be replaced by an equivalent sequence TC_i where $C_i \in \mathcal{C}_t$. The sequence therefore is equivalent to

$$U_{j_n} TTC_{i_n t} U_{j_{n-1}} TTC_{i_{n-1} t} \ldots U_{j_1} TTC_{i_1 t},$$

which as the product TT gives a Phase gate, collapses to a Clifford circuit, and the inverse gate can then easily be found.

These procedures allow one to collect estimates of p_{ref} and p_T. The latter quantity conflates the error of the Clifford gates with the interleaved T gates. We now wish to relate p_T to the average fidelity of a single T gate, subject of course to some natural assumptions on the noise. In the next section we demonstrate how the fidelity of the T gate can be estimated as

$$\mathcal{F}(\mathcal{E}_T) = \frac{1}{d+1}\left(d\frac{(d^2-1)p_T + 1}{(d^2-1)p_{\text{ref}} + 1} + 1\right), \tag{6}$$

which we derive in section 3. When estimating the fidelity of the T gate using equation (6) the intrinsic error caused by the assumptions used in interleaved benchmarking, the assumption relating to the noise in the system as well as the uncertainty error caused by stochastic measurements needs to be borne in mind. This is further explored in the next two sections.

As an aside, we see that that our protocol works more generally than on just T gates. In fact, our protocol works for every gate A such that both A^2 and $APA^\dagger$ are elements of the Clifford group for all Pauli operators P. This condition is satisfied for many but not all of the square roots of Clifford gates.

3. Fidelity estimation

The average fidelity of a general noise map is related to a depolarizing channel [1, 5]

$$\mathcal{D}_p(\rho) = p\rho + (1-p)\frac{\mathbb{1}}{d}, \tag{7}$$

with a specific depolarizing parameter p. For a d-dimensional quantum system, the depolarizing parameter is related to the average fidelity by

$$\mathcal{F}_{\text{avg}}(\mathcal{D}_p) = \frac{(d-1)p + 1}{d}. \tag{8}$$

The connection to arbitrary noise maps $\mathcal{E}$ crucially uses the property of a unitary 2-design. Sampling over a unitary 2-design reproduces the second moments of the Haar measure over all unitaries. As shown in [1, 5] this means that if the group $\mathcal{G}$ is a unitary 2-design, we have the following identity

$$\mathcal{E}^{\mathcal{G}} = \mathcal{D}_p, \tag{9}$$

where p is determined by the relations

$$\mathcal{F}_{\text{avg}}(\mathcal{E}) = \mathcal{F}_{\text{avg}}(\mathcal{E}^{\mathcal{G}}) = \mathcal{F}_{\text{avg}}(\mathcal{D}_p), \tag{10}$$

yielding

$$p = \frac{(d\mathcal{F}_{\text{avg}}(\mathcal{E}) - 1)}{d - 1}. \tag{11}$$

This is the fundamental result of RB, and relates the average gate fidelity of the noise $\mathcal{E}$ to the fit parameter p of the benchmarking experiments. Note that the other crucial assumptions for this derivation to hold are that the noise is time- and gate-independent.

Our design differs slightly in that it embeds an average over a unitary 1-design (the Pauli group) inside an average over a unitary 2-design (the Clifford group). Consequently, while the resulting channel is a depolarizing channel, the depolarizing parameter differs from that given by standard RB by a factor related to the anisotropy of the noise. We now quantify the deviation and show that in the regime of interest this difference is not significant.

To show this it is convenient to work in the superoperator representation of quantum channels. The superoperator representation is defined relative to a trace-orthonormal operator basis. Here we will use the Pauli basis representation of a channel $\mathcal{E}$, which chooses the suitably normalized Pauli operators as the trace-orthonormal basis. This representation consists of a matrix of inner products between the each Pauli matrix (P_j) and $\mathcal{E}(P_j)$ (see e.g. [15] for more details). Density operators ρ are represented by column vectors $|\rho)$ whose jth element is $\text{Tr}(P_j^\dagger \rho)/\sqrt{d}$. In particular, it is customary to fix P_1 to be the rescaled identity operator $\mathbb{1}/\sqrt{d}$. In this case any completely positive channel $\mathcal{E}$ can be written in block form as

$$\boldsymbol{\mathcal{E}} = \begin{pmatrix} S(\mathcal{E}) & \boldsymbol{\mathcal{E}}_{\text{sdl}} \\ \boldsymbol{\mathcal{E}}_{\text{n}} & \boldsymbol{\mathcal{E}}_{\text{u}} \end{pmatrix}, \tag{12}$$

where we refer to $\boldsymbol{\mathcal{E}}_{\text{sdl}}$, $\boldsymbol{\mathcal{E}}_{\text{n}}$ and $\boldsymbol{\mathcal{E}}_{\text{u}}$ as the *state-dependent leakage*, *nonunital* and *unital* blocks, respectively (see [21] for more details about this decomposition). If the channel is trace preserving then $\boldsymbol{\mathcal{E}}_{\text{sdl}} = 0$. The unital block ($\boldsymbol{\mathcal{E}}_{\text{u}}$) contains all the information necessary to extract the fidelity of the channel and, in particular, when the channel is twirled by a unitary-2 design, the resultant noise channel will have $\boldsymbol{\mathcal{E}}_{\text{u}}$ diagonalized, with each element being the average of the diagonal elements of the original $\boldsymbol{\mathcal{E}}_{\text{u}}$. By way of illustration in the single-qubit case, a CPTP noise channel will be of the form

$$\boldsymbol{\mathcal{E}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \alpha_1 & \sigma_x & \delta_1 & \delta_2 \\ \alpha_2 & \delta_3 & \sigma_y & \delta_4 \\ \alpha_3 & \delta_5 & \delta_6 & \sigma_z \end{pmatrix}, \tag{13}$$

where all the elements are real. The matrix elements themselves obey certain constraints on account of the requirement for complete positivity; see [26] for an explicit description of these constraints. We will make explicit use of the parameters σ_x, σ_y, and σ_z below.

After averaging over a unitary 2-design, the channel maps to $\mathcal{E}^G = \frac{1}{|G|}\sum_{C\in G}(C^\dagger \mathcal{E} C)$, and the unital block $\boldsymbol{\mathcal{E}}_{\text{u}}$ of the twirled error channel looks like

$$\begin{pmatrix} \frac{1}{3}(\sigma_x + \sigma_y + \sigma_z) & 0 & 0 \\ 0 & \frac{1}{3}(\sigma_x + \sigma_y + \sigma_z) & 0 \\ 0 & 0 & \frac{1}{3}(\sigma_x + \sigma_y + \sigma_z) \end{pmatrix}, \tag{14}$$

where the depolarizing factor measured by RB is $\frac{1}{3}(\sigma_x + \sigma_y + \sigma_z)$. In particular if the twirled noise channel is composed with a twirled noise channel twice then we obtain

$$\frac{1}{|G|^2}\sum_{C\in G}(C^\dagger \mathcal{E} C)\sum_{C\in G}(C^\dagger \mathcal{E} C) = \left(\frac{1}{|G|}\sum_{C\in G}(C^\dagger \mathcal{E} C)\right)^2 \tag{15}$$

and the depolarizing factor is simply squared, giving

$$\left(\tfrac{1}{3}(\sigma_x + \sigma_y + \sigma_z)\right)^2. \tag{16}$$

Recalling that C represents a perfect gate taken from the Clifford group and $\mathcal{E}$ represents the error channel equating to the average error on these gates, this then gives us the depolarizing parameter equating to the average fidelity after the application of two faulty Clifford gates.

In our protocol rather than applying two separate Clifford twirls, in the non-interleaved version, we apply a Clifford gate followed by a Pauli gate. Using the techniques described in [27] we can decompose a random benchmarking sequence as follows:

$$C_{\text{inv}} \ldots \mathcal{E}P_2\mathcal{E}C_2\mathcal{E}P_1\mathcal{E}C_1, \tag{17}$$

where we are using the standard assumptions that $\mathcal{E}$ is the gate error channel for the gates in the Clifford group (including the Paulis), and have absorbed the error on the inverting gate into the SPAM. We can always rewrite C_2 as $\tilde{C}_2 C_1^\dagger P_1^\dagger$, which means that equation (17) becomes:

$$C_{\text{inv}} \ldots \mathcal{E}P_2\mathcal{E}\tilde{C}_2[C_1^\dagger[P_1^\dagger\mathcal{E}P_1]\mathcal{E}C_1]. \tag{18}$$

When the channel is averaged over multiple runs with random choices for each of the Cliffords and Paulis, then for an n-gate sequence of Pauli gates followed by Clifford gates, we obtain

$$C_{\text{inv}}\mathcal{E}[C_1^\dagger[P_1^\dagger\mathcal{E}P_1]^n\mathcal{E}C_1]^n \tag{19}$$

which equates to a depolarizing channel

$$\frac{1}{|G_{\text{C}}|}\sum_{C\in G_{\text{C}}}\left(C^\dagger\mathcal{E}\left(\frac{1}{|G_{\text{P}}|}\sum_{P\in G_{\text{P}}}P^\dagger\mathcal{E}P\right)C\right), \tag{20}$$

where G_{C} represents the Clifford group and G_{P} the Pauli group. In superoperator form in the Pauli basis, the Clifford–Pauli twirled noise of equation (13) has a unital block of the form

$$\begin{pmatrix} \frac{1}{3}(\sigma_x^2+\sigma_y^2+\sigma_z^2) & 0 & 0 \\ 0 & \frac{1}{3}(\sigma_x^2+\sigma_y^2+\sigma_z^2) & 0 \\ 0 & 0 & \frac{1}{3}(\sigma_x^2+\sigma_y^2+\sigma_z^2) \end{pmatrix}, \tag{21}$$

which yields a depolarizing parameter of

$$\tfrac{1}{3}(\sigma_x^2+\sigma_y^2+\sigma_z^2). \tag{22}$$

In summary, two applications of a Clifford twirl yields the square of the average of the σ_j noise parameters (equation (16)), while the Clifford–Pauli twirl yields the average of the square of these parameters (equation (22)).

The difference between equations (22) and (16) can be viewed as the *variance* in the depolarizing parameters. For the purposes of this protocol we are assuming that the experimental Clifford gates are reasonably high fidelity. This means that it is reasonable to assume that the fidelity parameters are individually high, since for small dimensions d the average cannot be high unless each term is reasonably high. In this regime, the variance must also be low, as can be seen from an analysis of the anisotropy in the noise. Let us write the diagonal elements of $\boldsymbol{\mathcal{E}}_{\text{u}}$ as being

$$\sigma_x = \sigma, \quad \sigma_y = \sigma(1+\epsilon_y), \quad \sigma_z = \sigma(1+\epsilon_z), \tag{23}$$

where σ is close to 1 and $|\epsilon_y|$, $|\epsilon_z| \leqslant \epsilon$ are small. Expanding the difference of equations (22) and (16) as a function of the anisotropy of the noise, it can be seen that the variance cancels to $\mathcal{O}(\sigma^2\epsilon^2)$. Accordingly in the regime of reasonably accurate gates, the error introduced by the additional Pauli twirl is not significant in calculating the average fidelity of the gates. Our numerical analysis confirms this (see especially figure 4).

For completeness we note that when T gates are introduced the protocol can be written as

$$C_{\text{inv}}\mathcal{E}\ldots T\mathcal{E}_t\mathcal{E}P_1T\mathcal{E}_t\mathcal{E}C_1, \tag{24}$$

where the error on the T gate has been written as $\mathcal{E}_t$, so the noisy T gate is $T\mathcal{E}_t$. Using the method outlined in [6], this can be rewritten as

$$C_{\text{inv}}[C\mathcal{E}_t\mathcal{E}[(PT)^\dagger\mathcal{E}_t\mathcal{E}(PT)]^nC]^n. \tag{25}$$

Noting that PT (being a perfect Pauli followed by a perfect T gate) is a unitary 1-design, it can be seen that the decay parameter given by the protocol is that of the composed error channel $\mathcal{E}_t\mathcal{E}$ and (within the error noted above) gives us the fidelity of the composed error channel.

Now we wish to extract the specific contributions to the noise from the interleaved T gate. As above, we will denote this noise by $\mathcal{E}_t$. The average gate fidelity of equation (1) of the gates forming the group G, being $\mathcal{F}_{\text{avg}}(\mathcal{E})$ and the average gate fidelity of the the combined $\mathcal{G}T$ gates, being $\mathcal{F}_{\text{avg}}(\mathcal{E}_t\mathcal{E})$ by using the value of p calculated in step 5 of the relevant part of the protocol (being, p_{ref} and p_T respectively) and using equation (11). The average fidelity of the T gate, $\mathcal{F}_{\text{avg}}(\mathcal{E}_t)$ can then be estimated from the approximation $\chi_{00}^{\mathcal{E}_t\mathcal{E}} = \chi_{00}^{\mathcal{E}_t}\chi_{00}^{\mathcal{E}}$, where for qubits

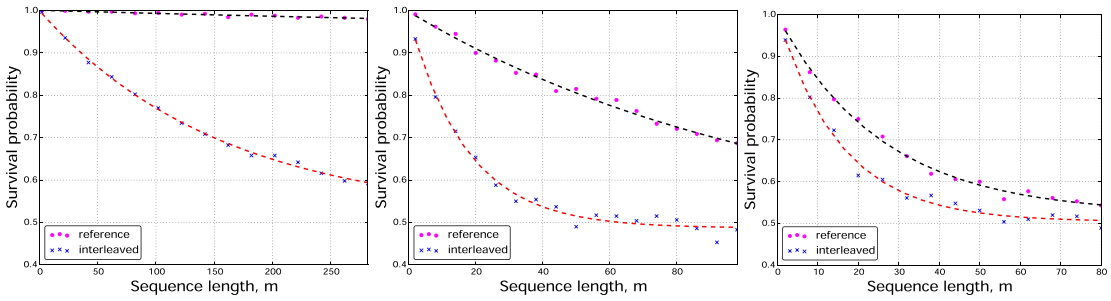

(a) Rotation Errors (b) Depolarizing and Rotation Errors (c) Amplitude Damping channel

Figure 1. In (a) we simulate a typical randomized benchmarking plot where the gates making up the Clifford group are assumed to suffer from a systematic over-rotation around the x-axis of 0.02 radians (i.e. $\exp(-\mathrm{i}X\frac{\theta}{2})$), whereas the T gate suffers an over-rotation of 0.12 radians. The fidelity of the Clifford gates is 99.993%. In this run the estimated fidelity of the T gate is 99.73%, compared with the actual fidelity of 99.76%. In this regime the error bounds given by equation (29) are $\pm 0.08\%$. In (b) we simulate the application of a depolarizing and rotation error channel. The Clifford gates are affected by a depolarizing channel (equation (7)) with $p = 0.01$ with an over-rotation of 0.01 radians whereas the T gates are affected by a depolarizing channel with $p = 0.02$ and an over-rotation of 0.05 radians. The estimated fidelity of the T gates is 97.6% compared to the actual fidelity of 97.2%. Finally in (c) a generalized amplitude damping channel has been simulated (see equation (30)). For the Clifford group $p = 0.995$ and $\gamma = 0.01$. For the T gate $p = 0.99$ and $\gamma = 0.04$. The Estimated fidelity of the T gate is 98.8% and the actual fidelity is 98.7%, for a relative error of about 10^{-3}.

$$\chi_{00}^{\mathcal{E}} = \frac{3}{2}\mathcal{F}_{\mathrm{avg}}(\mathcal{E}) - \frac{1}{2}, \tag{26}$$

where here $\chi_{00}^{\mathcal{E}}$ represents the upper left entry of the χ matrix representation of $\mathcal{E}$ [28].

Finally, to derive equation (6), we use the fact that fidelity of the noise to the identity channel ($\mathcal{F}(\mathcal{E})$), the upper left entry of the χ matrix representation of the noise ($\chi_{00}^{\mathcal{E}}$) and the measured decay parameter ($p_{\mathcal{E}}$) are all linearly related (see [29] for a convenient table showing these relations). Initially we note that we have p_T, which is related to $\chi_{00}^{\mathcal{E}_t\mathcal{E}}$ by the relation

$$\chi_{00}^{\mathcal{E}_t\mathcal{E}} = \frac{(d^2-1)p_T + 1}{d^2}, \tag{27}$$

similarly for $\chi_{00}^{\mathcal{E}}$ and p_{ref}. We estimate $\chi_{00}^{\mathcal{E}_t}$ as $\frac{\chi_{00}^{\mathcal{E}_t\mathcal{E}}}{\chi_{00}^{\mathcal{E}}}$ which simplifies to

$$\chi_{00}^{\mathcal{E}_t} \approx \frac{(d^2-1)p_T + 1}{(d^2-1)p_{\mathrm{ref}} + 1}, \tag{28}$$

then noting that $\mathcal{F} = \frac{d\chi_{00}+1}{d+1}$ we obtain equation (6).

The above approximation is valid to within the bound derived in [8]:

$$\chi_{00}^{\mathcal{E}_t\mathcal{E}} - \chi_{00}^{\mathcal{E}_t}\chi_{00}^{\mathcal{E}} \leqslant 2\sqrt{(1-\chi_{00}^{\mathcal{E}})\chi_{00}^{\mathcal{E}}(1-\chi_{00}^{\mathcal{E}_t})\chi_{00}^{\mathcal{E}_t}} + (1-\chi_{00}^{\mathcal{E}})(1-\chi_{00}^{\mathcal{E}_t}). \tag{29}$$

As noted in [10], this bound is loose in general but tight in the regime where the gates forming G have a much higher fidelity than the T gate, which is fortunately the regime of interest.

4. Numerical simulation

As previously discussed, RB is robust, both in theory and practice, to some level of gate-dependent noise [6, 15]. We now present some numerical simulations illustrating that the robustness in the protocol presented where (as anticipated) the error profile of the T gate is different from the gates in $\mathcal{G}$. For the purpose of these simulations we are assuming a single qubit, where $\mathcal{G}$ is taken to be all 24 single-qubit Cliffords $\mathcal{C}$ and $\mathcal{P}$ is the four single-qubit Paulis (which are contained in $\mathcal{C}$).

Initially, three different noise models were explored. As with all interleaved protocols, care must be taken to reduce the error of each of the estimates, as effectively the protocol requires one estimation to be divided by another. In each simulation, for various gate lengths, a number of random sequences (2000 for figure 1(a) and 1000 for figures 1(b) and (c)) were generated. At the end of each sequence a single measurement (returning 1 or 0) was made. This follows the design outlined in [25]. Curves were fit using standard least square methodologies. Figure 1(a) illustrates the case where the rotational errors on the Clifford gates are different from that of the implementation of the T gate (modeled as over-rotatons on the x-axis i.e. $\exp(-\mathrm{i}X\frac{\theta}{2})$ for some θ—see caption). As previously discussed, the low error rate on the Clifford gates allows an accurate estimation of the T gate even

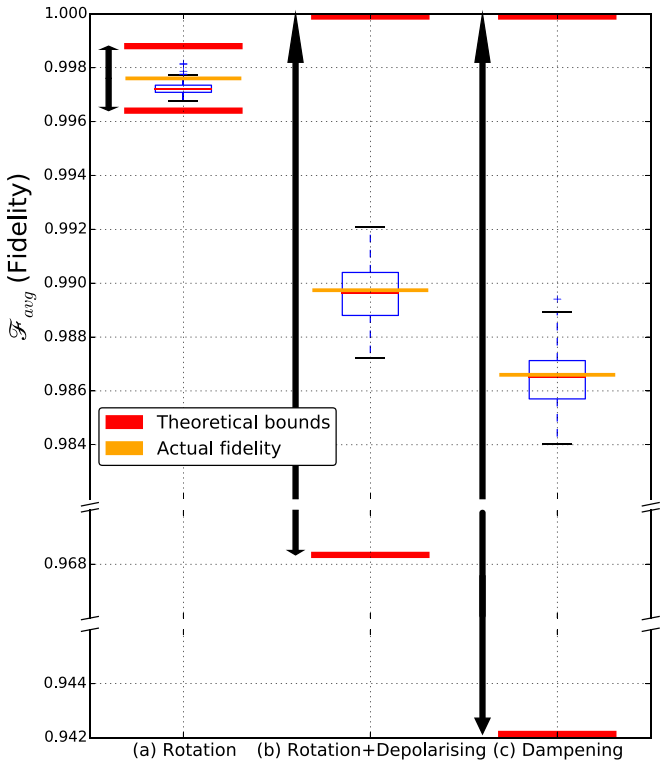

Figure 2. Tukey plots showing estimated fidelities over 100 runs of each of the scenarios presented in figure 1. The box represents the inter-quartile range (IQR), the whiskers 1.5 times the IQR, with outliers plotted separately. The theoretical bounds are shown as thick red lines and the actual fidelity of the T gate as a longer orange line. As can be seen in the regime of figure 1(a) the theoretical bounds of equation (29) are reasonably tight, but in the other cases are a lot looser than the simulated results.

with rotational errors. The first plot in figure 2 shows the variance in the estimation over 100 simulations as a Tukey box plot. As can be seen even in this region, the bounds given by equation (29) are not tight, although it can be seen the protocol does, slightly, underestimate the fidelity of the T gates in this example. The 100 simulations have a median error estimate of 99.72% with a σ of 0.00025 (0.025%), compared with an actual fidelity of 99.76%.

For figure 1(b), we used a depolarization channel equation (7) composed with the x-axis over-rotation channel $(\exp(-\mathrm{i}X\frac{\theta}{2}))$(see caption for details). In this case the fidelity of the Clifford gates is 99.5% and the T gate 99.0%. The estimate in this case was 98.9%. As can be seen from the the second boxplot in figure 2 bounds in equation (29) are not saturated, the standard deviation of the estimated fidelity (over 100 simulations) being 0.001 (0.1%). The final illustrative run figure 1(c) uses an amplitude damping channel parameterized by p and γ with the Kraus operators given by

$$E_0 = \sqrt{p}\begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix} \qquad E_1 = \sqrt{p}\begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}$$
$$E_2 = \sqrt{1-p}\begin{pmatrix} \sqrt{1-\gamma} & 0 \\ 0 & 1 \end{pmatrix} \quad E_3 = \sqrt{1-p}\begin{pmatrix} 0 & 0 \\ \sqrt{\gamma} & 0 \end{pmatrix}. \tag{30}$$

Choosing $p = 0.995$ and $\gamma = 0.01$ for the Clifford gates and $p = 0.99$ and $\gamma = 0.04$ for the T gates, gives us a Clifford gate fidelity of 99.67% and a T gate actual fidelity of 98.7% compared with the estimated fidelity of 98.6%. The standard deviation of the estimate over 100 runs was 0.001 (0.1%).

Figures 3(a) and (b) show the error in the estimate of the T gate over more varied parameters. In figure 3(a), the rotational errors on the Clifford gates and T gate were varied. On the same graph we plotted the plane corresponding to the error bounds detailed in equation (29). As can be seen, the error bounds are not saturated in the regime of higher rotational errors. Figure 3(b) shows a similar treatment for the generalized amplitude dampening channel with varying γ. Again we see the protocol performs better than the limits in equation (29).

Finally, we tested the protocol two different ways against randomly generated unital CTPT error channels, close to Identity. To create such channels we used the parameterization of the unital block $\boldsymbol{E}_{\mathrm{u}}$ (see section 3) introduced in [26]. The block can be represented as $U\Sigma V$, where U and V are unitary and Σ is is diagonal with real entries $(\lambda_1, \lambda_2, \lambda_3)$, constrained such that $|\lambda_1 \pm \lambda_2| \leqslant |1 \pm \lambda_3|$. The unitary matrices were generated by combining three rotations (around the x, y and x axis), where the rotation amount was drawn from a normal

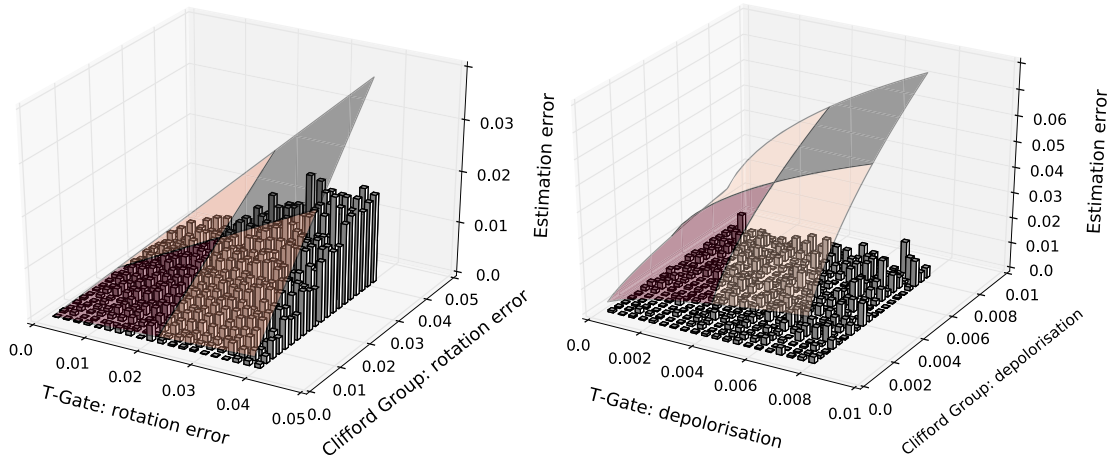

(a) Rotation error (b) Depolarizing noise

Figure 3. (a) Error in the T gate estimation from various numerical simulations where we have varied the over-rotation levels in the Clifford group and in the T gates (rotation (θ) shown as radians around the x-axis i.e. $\exp(-iX\frac{\theta}{2})$). Each point was estimated using 2000 single shot measurements per gate length. The transparent plane above the graph shows the bounds in equation (29). As can be seen the estimate error does not appear to saturate these bounds as the error rate in the Clifford group increases. (b) Error in the T gate estimation for a general amplitude dampening channel, where we have 1000 single shot measurements per gate length. Each bar represents the error in the estimate for a single run of the protocol using the amplitude dampening γ given by the x, y coordinates. The transparent plane above the graph shows the bounds calculated in accordance with equation (29). In both graphs, given that relatively few measurements have been simulated for each of the data points, there is uncertainty in each of the bars because of the stochastic nature of the process. This means that any individual bar height is of limited informational value. However, even when a limited number of measurements are simulated, the graphs are quite clear in confirming the overall trend of decreasing error in the estimation as the fidelity of the Clifford gates and T gates improves and the increased laxity of the theoretical bounds where the noise type is one of low unitarity [21].

distribution with $\mu = 0$ and a standard a deviation (σ) of 0.02 (0.1 for the T gate error channel). The diagonal matrix consisted of values drawn from a normal distribution with $\mu = 0.98$ ($\mu = 0.95$ for the T gate error channel) and $\sigma = 0.01$ ($\sigma = 0.1$ for the T gate). The diagonal elements were re-chosen if they exceeded one or failed to satisfy the requirement stated above. As can be seen from the inset to figure 4 this led to random 'Clifford' error channels that were roughly normally distributed between fidelity ranges from 0.98 to 1.0 and T gate error channels with fidelities from 0.9 to 1.0.

In the first test, each of these maps was then used as an error channel on the Clifford Gates and the T gate, respectively, and (1) the fidelity of the Clifford gates were estimated using a standard RB protocol; (2) the fidelity of the Clifford gates were estimated using the Clifford/Pauli protocol introduced in this paper; and (3) the interleaved protocol discussed above was used to estimate the fidelity of the T gate. In each case sequences of 2–100 gate lengths (increasing by 4) were simulated, 4000 such random sequences for each 'length', with each sequence being run once, returning either a 1 or 0. As can be seen from figure 4 the 'Clifford Pauli' twirl protocol used in this paper returned accurate results in the regime of high-fidelity gates and the T gate estimation is accurate to within the theoretical bounds, even where the noise channel for the T gates is randomly different from the noise channel for the Clifford group.

In the second test, we fixed the error on the Clifford gates (arbitrarily to 0.992 00) by randomly generating noise maps and throwing them away if they did not have the required fidelity (to 5 significant figures) until we had 1000 such noise maps. Similarly with the T gate error (set at 0.985 00). The protocol was then applied to each of these 1000 pairs of noise maps, using different numbers of simulated single shot measurements (being 2000, 4000 and 8000 sequences per gate length). In addition, for the T gate estimates we calculated the fidelity assuming perfect measurements to explore the actual approximation error (i.e. to eliminate stochastic shot variability). As can be seen where we have such relatively high fidelities the bounds in equation (29) are not saturated (with 8000 measurements the worst observed error was 0.003, the bounds are ± 0.033 an order of magnitude worse).

5. Compatibility with small subgroups

In this section we discuss the compatibility of this protocol with the work of Cleve *et al* [30]. In [30], it is shown how a subset of the Clifford group that can be used to construct exact unitary-2 designs with a gate complexity that is nearly linear in the number of qubits (in the simplest case it scales as $\mathcal{O}(n \log n \log\log n)$ assuming the extended Riemann Hypothesis is true). For a given number of qubits n, they note that a collection of Clifford

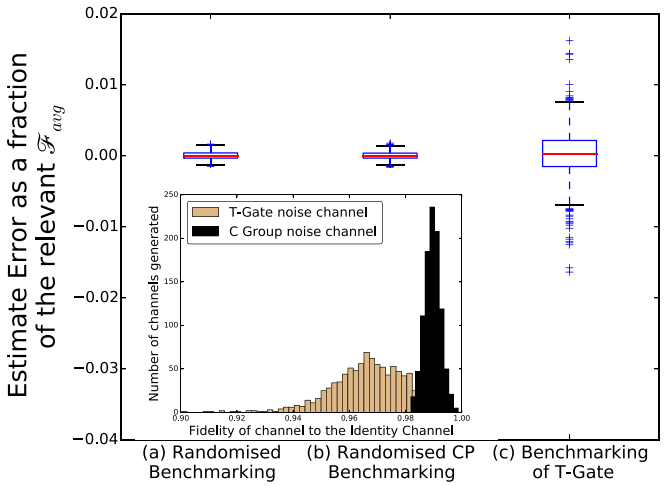

Figure 4. Tukey plots showing estimated fidelities over 1000 randomly generated noise channels, in each case taking 4000 single measurements (one per random sequence) at various gate lengths. The noise channels were generated to be near Identity (see text for details) with the noise channel applied to the Clifford group gates having a fidelity to the Identity channel of between 0.98 and 1. The noise channel applied to the T gates had a fidelity to the Identity of between 0.9 and 1. The inset barchart shows the fidelity distribution of these random noise channels. The Tukey plots show the estimation error as a fraction of the actual fidelity (i.e. $(\mathcal{F}_{\text{estimate}} - \mathcal{F}_{\text{actual}})/\mathcal{F}_{\text{actual}}$). Plot (a) shows the results from a 'standard' randomized benchmarking simulation, plot (b) the reference fidelity generated by the protocol presented in this paper (i.e. a random Clifford followed by a random Pauli). They are indistinguishable in this regime. The error in the estimation of the T gate is shown in plot (c) and can be further reduced by increasing the number of sequences measured, subject to the over-all intrinsic error in the approximation (see for example figure 5(b)).

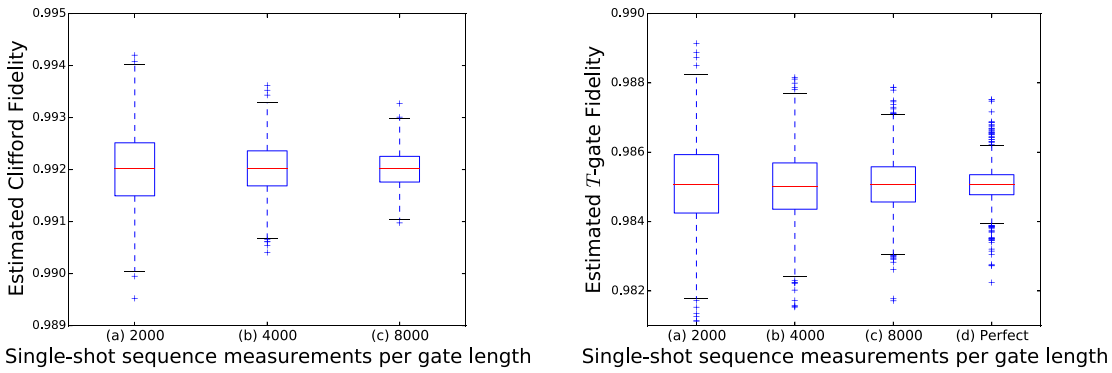

(a) Clifford Error Estimate

(b) T gate Error Estimate

Figure 5. Tukey plots showing estimated fidelity for Clifford gates and T gates where we have randomly generated CPTP noise maps (see text for details) with a fidelity of 0.9920 for the Clifford gates and 0.9850 for the T gates. The x-axis shows the number of simulated measurements (single shot measurements returning a 1 or a 0) for various gate sequence lengths (1–100 incrementing by 6). Plot (a) shows the estimated Clifford gate fidelity and plot (b) the estimated T gate fidelity for various numbers of measurements. The final boxplot in plot (b) shows the results assuming a theoretical 'perfect' measurement and represents the intrinsic approximation error resulting from estimating $\chi_{00}^{\mathcal{E}_T\mathcal{E}}$ as $\chi_{00}^{\mathcal{E}_T}\chi_{00}^{\mathcal{E}}$ in this regime. This error cannot be reduced by making more measurements and varies depending on the nature of the noise and the fidelity of the gates in question. Here for the fidelities stated and assuming random CPTP noise maps the intrinsic error is at most an order of magnitude lower than the the bounds in equation (29) (which give an error bound of ± 0.033).

gates isomorphic to $\mathrm{SL}_2(\mathrm{GF}(2^n))$ (of which there are $2^{3n} - 2^n$ gates) when mixed with the Paulis (2^n), give a unitary-2 design. The mixed gates are a subgroup of Clifford gates, where the total number of gates is $2^{5n} - 2^{3n}$. By way of illustration, the Clifford group has cardinality $|\mathcal{C}_n| = 2^{n^2+2n}\prod_{j=1}^{n}(4^j - 1)$. With two qubits this means that we can reduce the 11 520 Cliffords to a mere 60 gates, twirled by the 16 Pauli gates (giving 960 gates in total). The main contribution of [30] is to show how these gates can be constructed with near-linear gate costs (the number of one- and two-qubit gates). Here we use their results to reduce the variety of gates needed to conduct RB of T gates. Since our changes are only in respect to the additional single-qubit gates, the near-linear time complexity is preserved.

Let $\mathcal{C}_\gamma$ represent the set of gates isomorphic to $\mathrm{SL}_2(\mathrm{GF}(2^n))$. As before $\mathcal{P}$ represents the 2^n Paulis and $\mathcal{C}_t$ represents the Cliffords formed by conjugation of the elements of $\mathcal{P}$ with a T gate. Let S be the phase gate (i.e. the gate formed by two applications of a T gate).

The protocol previously discussed needs to be changed so that the reference run consists of n-sequences of gates drawn randomly from $\mathcal{C}_\gamma\mathcal{P}$. Reference [30] contains the details how to generate $\mathcal{C}_\gamma$, the set of gates $\mathcal{C}_\gamma\mathcal{P}$, is the unitary 2-design we wish to use and it becomes the group $\mathcal{G}$ referred to in the protocol detail in section 2.

The interleaved sequence would be physically implemented using repeated applications of the following gates: $\mathcal{C}_{\gamma'}T\mathcal{C}_tT$, where $\mathcal{C}_{\gamma'}$ is drawn from the gates formed by the sequence $\mathcal{C}_\gamma(S)^\dagger$. That is the $\mathcal{G}$ in the interleaved part of the protocol (section 2.2) consists of the group formed by the gates $\mathcal{C}_\gamma\mathcal{P}(S)^\dagger$ (which are all Clifford gates), rather than the group formed by $\mathcal{C}_\gamma\mathcal{P}$. As is readily apparent the interleaved sequence can be analyzed as

$$\mathcal{C}_{\gamma'}T\mathcal{C}_tT \to \mathcal{C}_\gamma(S)^\dagger TT\mathcal{P} \to \mathcal{C}_\gamma\mathcal{P}, \tag{31}$$

which is an efficiently invertible unitary-2 design as required.

6. Discussion

We provided a simple extension of the RB protocol for interleaved gates that allows the fidelity of T gates to be robustly estimated. Unlike previous proposals for benchmarking the T gate, ours retains the benefits of an exact unitary 2-design that at a minimum at least diagonalizes the noise between the T gates. Our numerical analysis confirms that, subject to the usual RB assumptions of time- and gate-independent noise, the fidelity of T gates can be estimated to a high degree of accuracy. As discussed above, the analyses in [4, 22, 24] will apply, confirming that the gate independence assumption can be dropped without a significant effect on the results. Our simulations confirm that the theoretical bounds for interleaved benchmarking are only tight in the region of high gate fidelity where the gate of interest has a worse fidelity than those forming the rest of the benchmarking group. Outside such regions the protocol appears to still give accurate estimations of the fidelity of the gate in question.

Finally, we note that the use of the protocol consisting of sequences of a random Clifford gate followed by a random Pauli gate creates a depolarizing channel that differs from the usual RB depolarizing channel by the variance of the trace of the diagonal elements $\mathrm{Tr}(P_j\mathcal{E}(P_j))$ where P_j runs over each of the three traceless Paulis. This provides a method of using RB protocols to analyze the anisotropy of the noise, which may provide useful future work.

Acknowledgments

This work was supported by the US Army Research Office grant numbers W911NF-14-1-0098 and W911NF-14-1-0103, and by the Australian Research Council Centre of Excellence for Engineered Quantum Systems CE110001013. STF acknowledges support from an Australian Research Council Future Fellowship FT130101744.

The authors thank David Gross, Joel Wallman and Arnaud Dugas for helpful and illuminating discussions.

References

[1] Emerson J, Alicki R and Życzkowski K 2005 Scalable noise estimation with random unitary operators *J. Opt.* B **7** S347
[2] Knill E, Leibfried D, Reichle R, Britton J, Blakestad R, Jost J D, Langer C, Ozeri R, Seidelin S and Wineland D J 2008 Randomized benchmarking of quantum gates *Phys. Rev.* A **77** 012307
[3] Dankert C, Cleve R, Emerson J and Livine E 2009 Exact and approximate unitary 2-designs and their application to fidelity estimation *Phys. Rev.* A **80** 012304
[4] Magesan E, Gambetta J M and Emerson J 2011 Scalable and Robust randomized benchmarking of quantum processes *Phys. Rev. Lett.* **106** 180504
[5] Nielsen M A 2002 A simple formula for the average gate fidelity of a quantum dynamical operation *Phys. Lett.* A **303** 249
[6] Magesan E *et al* 2012 Efficient measurement of quantum gate error by interleaved randomized benchmarking *Phys. Rev. Lett.* **109** 080505
[7] Gaebler J P *et al* 2012 Randomized benchmarking of multiqubit gates *Phys. Rev. Lett.* **108** 260503
[8] Kimmel S, da Silva M P, Ryan C A, Johnson B R and Ohki T 2014 Robust extraction of tomographic information via randomized benchmarking *Phys. Rev.* X **4** 011050
[9] Barends R *et al* 2014 Superconducting quantum circuits at the surface code threshold for fault tolerance *Nature* **508** 500–3
[10] Carignan-Dugas A, Wallman J J and Emerson J 2015 Characterizing universal gate sets via dihedral benchmarking *Phys. Rev.* A **92** 060302
[11] Cross A W, Magesan E, Bishop L S, Smolin J A and Gambetta J M 2016 Scalable randomised benchmarking of non-Clifford gates *NPJ Quantum Inf.* **2** 16012

[12] Chasseur T, Reich D M, Koch C P and Wilhelm F K 2016 Hybrid benchmarking of arbitrary quantum gates arXiv:1606.03927 [quant-ph]
[13] Flammia S T and Liu Y-K 2011 Direct fidelity estimation from few Pauli measurements *Phys. Rev. Lett.* **106** 230501
[14] da Silva M P, Landon-Cardinal O and Poulin D 2011 Practical characterization of quantum devices without tomography *Phys. Rev. Lett.* **107** 210404
[15] Wallman J J and Flammia S T 2014 Randomized benchmarking with confidence *New J. Phys.* **16** 103032
[16] Kueng R, Long D M, Doherty A C and Flammia S T 2016 Comparing experiments to the fault-tolerance threshold *Phys. Rev. Lett.* **117** 170502
[17] Ball H, Stace T M, Flammia S T and Biercuk M J 2016 Effect of noise correlations on randomized benchmarking *Phys. Rev.* A **93** 022303
[18] Wallman J J 2015 Bounding experimental quantum error rates relative to fault-tolerant thresholds arXiv:1511.00727 [quant-ph]
[19] Sanders Y R, Wallman J J and Sanders B C 2015 Bounding quantum gate error rate based on reported gate fidelity *New J. Phys.* **18** 012002
[20] Gross D, Audenaert K and Eisert J 2007 Evenly distributed unitaries: on the structure of unitary designs *J. Math. Phys.* **48** 052104
[21] Wallman J, Granade C, Harper R and Flammia S T 2015 Estimating the coherence of noise *New J. Phys.* **17** 113020
[22] Wallman J J, Barnhill M and Emerson J 2015 Robust characterization of loss rates *Phys. Rev. Lett.* **115** 060501
[23] Gottesman D and Chuang I L 1999 Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations *Nature* **402** 390–3
[24] Epstein J M, Cross A W, Magesan E and Gambetta J M 2014 Investigating the limits of randomized benchmarking protocols *Phys. Rev.* A **89** 062321
[25] Granade C, Ferrie C and Cory D G 2015 Accelerated randomized benchmarking *New J. Phys.* **17** 013042
[26] King C and Ruskai M B 2001 Minimal entropy of states emerging from noisy quantum channels *IEEE Trans. Inf. Theory* **47** 192–209
[27] Magesan E, Gambetta J M and Emerson J 2012 Characterizing quantum gates via randomized benchmarking *Phys. Rev.* A **85** 042311
[28] Nielsen M A and Chuang I L 2011 *Quantum Computation and Quantum Information: 10th Anniversary Edition* 10th edn (New York: Cambridge University Press)
[29] Dugas A C, Wallman J J and Emerson J 2016 Efficiently characterizing the total error in quantum circuits arXiv:1610.05296 [quant-ph]
[30] Cleve R, Leung D, Liu L and Wang C 2016 Near-linear constructions of exact unitary 2-designs *Quantum Inf. Comput.* **16** 0721–56

Chapter 7

Explaining quantum correlations through evolution of causal models

We propose a framework for the systematic and quantitative generalization of Bell's theorem using causal networks. We first consider the multi-objective optimization problem of matching observed data while minimizing the causal effect of non-local variables and prove an inequality for the optimal region that both strengthens and generalizes Bell's theorem. To solve the optimization problem (rather than simply bound it), we develop a novel genetic algorithm treating as individuals causal networks. By applying our algorithm to a photonic Bell experiment, we demonstrate the trade-off between the quantitative relaxation of one or more local causality assumptions and the ability of data to match quantum correlations.

PHYSICAL REVIEW A **95**, 042120 (2017)

Explaining quantum correlations through evolution of causal models

Robin Harper,[1,2] Robert J. Chapman,[2,3] Christopher Ferrie,[1,2,4] Christopher Granade,[1,2] Richard Kueng,[5] Daniel Naoumenko,[2] Steven T. Flammia,[1,2] and Alberto Peruzzo[2,3]

[1]*Centre for Engineered Quantum Systems, University of Sydney, Sydney, NSW, Australia*
[2]*School of Physics, University of Sydney, Sydney, NSW, Australia*
[3]*Quantum Photonics Laboratory, School of Engineering, RMIT University, Melbourne, Victoria 3000, Australia*
[4]*University of Technology Sydney, Centre for Quantum Software and Information, Ultimo NSW 2007, Australia*
[5]*Institute for Theoretical Physics, University of Cologne, Cologne, Germany*

(Received 29 August 2016; revised manuscript received 29 November 2016; published 17 April 2017)

We propose a framework for the systematic and quantitative generalization of Bell's theorem using causal networks. We first consider the multiobjective optimization problem of matching observed data while minimizing the causal effect of nonlocal variables and prove an inequality for the optimal region that both strengthens and generalizes Bell's theorem. To solve the optimization problem (rather than simply bound it), we develop a genetic algorithm treating as individuals causal networks. By applying our algorithm to a photonic Bell experiment, we demonstrate the trade-off between the quantitative relaxation of one or more local causality assumptions and the ability of data to match quantum correlations.

DOI: 10.1103/PhysRevA.95.042120

I. INTRODUCTION

While it seems conceptually obvious that causality lies at the heart of physics, its exact nature has been the subject of constant debate. The fundamental implications of quantum theory shed new light on this debate. It is thought these implications may lead to new insights into the foundations of quantum theory, and possibly even quantum theories of gravity [1–10].

These realizations have their roots in the Einstein-Podolski-Rosen thought experiment [11] and the fundamental theorems of Bell [12] and of Kochen and Specker [13]. A cornerstone of modern physics, Bell's theorem, rigorously excludes classical concepts of causality. Roughly speaking Bell's theorem states that the following concepts are mutually inconsistent: (1) reality, (2) locality, (3) measurement independence, and (4) quantum mechanics.

In philosophical discussions, typically one rejects (1) or (2), which together are often referred to as *local causality*, though the other options have been considered as well. In studies with an operational bent, however, one often considers relaxations of (2) or (3) which is what we concern ourselves with here. These relaxations have been addressed from different perspectives, but only regarding specific causal influences in isolation [14–23], whereas here we wish to study all possible relaxations of the causal assumptions implied by (2) and (3) simultaneously.

The framework of *causal networks* [24,25] is wildly successful within the field of machine learning and has led some physicists to utilize them to elucidate the tension between causality and Bell's theorem. Recently, Wood and Spekkens have shown that existing principles behind causal discovery algorithms (namely, the absence of fine-tuning) still cannot be reconciled with entanglement induced quantum correlations even if one admits nonlocal models [9]. However, such results only hold for the exact distributions, and would not necessarily apply to experimental data due to measurement noise, or a relaxation of the demand of reproducing exactly the quantum correlations. Clearly, the further away from the quantum correlations one is allowed to stray, the more likely a locally causal model can be found.

Here we propose a framework for systematic and quantitative generalizations of Bell's theorem by using causal networks. The idea, depicted in Fig. 1, is to consider the multiobjective optimization problem of matching the observed data from an experiment while minimizing the causal effect of nonlocal variables. It is in this sense of matching experimental data that we are explaining the quantum correlations. Our first contribution is a rigorous lower bound for this optimization problem, demonstrating a generalization of Bell's theorem. Theorem 1 below establishes that there must exist a tradeoff between the goodness of fit to experimental data and the quantitative amount of causal influence for *any* model.

This theorem rules out a portion of the space allowed by this new framework, but the bounds are not tight. To solve the optimization problem, and hence numerically find the *optimal* bounds, we develop a type of genetic algorithm called a multiobjective evolutionary algorithm (MOEA) to quantify the relaxations necessary to reproduce the data generated by experiments on entangled quantum systems [26–28]. Our genetic algorithm treats as individuals causal networks and we develop genetic operators which represent the evolution of these networks. By applying our algorithm to a photonic Bell experiment, we show that the tradeoff between the quantitative relaxation of one or more local causality assumptions and the ability to match quantum correlations appears linear.

The outline of the paper is as follows. In Secs. II and III we set out the background of the causal models we use and the mathematics required to convert a probability distribution into a fitness function. In Sec. IV we provide analytic bounds on causal influence. In Sec. V we describe the experiment that provided the input to the algorithm. Section VI briefly describes the result of applying the genetic algorithm to the experimental data. Section VII describes the process by which we convert the problem into one that can be explored using evolutionary operators and details the construction of the algorithm. We conclude in Sec. VIII with a discussion.

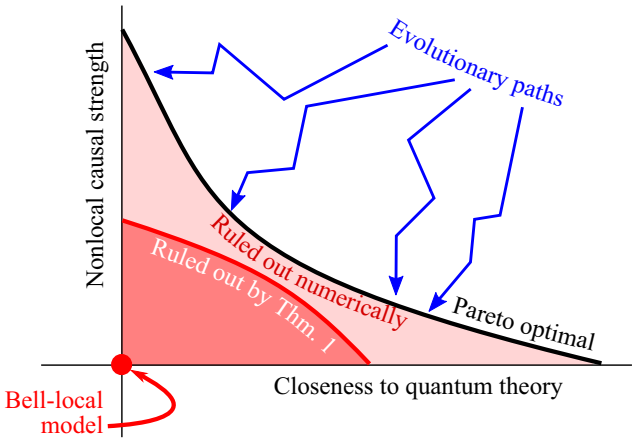

FIG. 1. Sketch of the concept of Pareto optimality for demarcating the boundary between local causality and quantum correlations. In this picture, Bell's theorem rules out the origin only. Our results rule out an entire region of possible models in the presence of relaxations of Bell's assumptions. We rule out this region both rigorously with Theorem 1 and numerically with the evolutionary algorithm that we developed specifically for this task.

II. CAUSAL MODELS FOR BELL EXPERIMENTS

The formalism of causal models allows us to quantify the relaxations necessary to avoid the contradiction in Bell's theorem and, more importantly, explore the trade-offs necessary in minimizing the amount by which the assumptions are violated. Building off the work of Chaves *et al.* [29], we make all this concrete through a quantification of the relaxation of each assumption in the context of causal models. The task of minimizing the amount of the relaxation is a multiobjective optimization problem. Bell's theorem is recast as the statement that all objectives cannot be simultaneously minimized. We explore the trade-offs through the concept of Pareto optimality.

The prototypical "Bell experiment" has two distant parties, often named Alice and Bob. We suppose that Alice and Bob each have devices with binary measurement settings, respectively labeled x and y. Conditioned on these measurement settings, their devices also record binary events, labeled a (Alice) and b (Bob). Suppose it is empirically observed that a and b are correlated. Bell defined a locally causal model of such correlations as follows: there exists a "hidden variable" λ which is the common cause both of a conditioned on x, and of b conditioned on y. We write these random variables as $a \mid x$ and $b \mid y$, respectively. Formally, the general conditional distribution is assumed to satisfy

$$\Pr(a,b|x,y,\lambda) = \Pr(a|x,\lambda)\Pr(b|y,\lambda). \tag{1}$$

Moreover, it is assumed that the choices of settings can be made such that each of x and y can be set independently of the hidden variable λ,

$$\Pr(x,y|\lambda) = \Pr(x|\lambda)\Pr(y|\lambda) = \Pr(x)\Pr(y). \tag{2}$$

Such an assumption is often motivated by the injection of randomness into the measurement settings or the free will of Alice and Bob. Bell's theorem can be stated succinctly as follows: the conditional distributions describing the outcomes of some experiments on quantum systems cannot be factorized as in Eqs. (1) and (2).

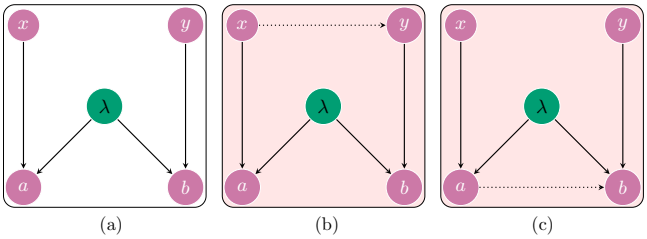

FIG. 2. Causal networks for Bell-type experiments. On the left is the local hidden variable model, which respects the assumptions going into Bell's famous no-go theorem. Such a model cannot account for certain correlations obtained from measuring entangled particles. The graph in the middle contains a causal link between the measurement settings. Such a model exploits the detection loophole and violates measurement independence. Finally, on the right is a superluminal model which contains a causal link between the outcomes of the experiments.

A causal network is a directed acyclic graph with nodes representing random variables and edges denoting causal relationships between variables. The defining feature of such networks is the factorization of joint probabilities. Generally, suppose we have nodes $\{x_0,x_1,\ldots,x_K\}$, each of which represents a random variable in our model. We will assume that each such random variable is discrete, and without loss of generality will assume integer labels $x_i \in \{0,\ldots,\dim x_i - 1\}$ for its possible values. The edges in the causal network of these variables are defined such that

$$\Pr(x_0,x_1,\ldots,x_K) = \prod_{i=0}^{K} \Pr(x_i|\mathrm{pa}_i), \tag{3}$$

where pa_i denotes the parents of node i.

Take, for example, the causal network in Fig. 2(a). In general, we can decompose the joint distribution $\Pr(a,b,x,y,\lambda)$ in terms of conditional distributions as

$$\Pr(a,b,x,y,\lambda) = \Pr(a|b,x,y,\lambda)\Pr(b|x,y,\lambda)\Pr(x|y,\lambda)\Pr(y|\lambda)\Pr(\lambda). \tag{4}$$

Using the causal network to eliminate conditionals, Eq. (3) implies

$$\Pr(a,b,x,y,\lambda) = \Pr(a|x,\lambda)\Pr(b|y,\lambda)\Pr(x)\Pr(y)\Pr(\lambda), \tag{5}$$

which are identical to Bell's assumptions on local hidden variable models. Thus Bell's theorem is equivalent to the statement that certain quantum correlations cannot be realized by the causal network in Fig. 2(a).

III. RELAXING BELL's ASSUMPTIONS

It is known that quantum mechanical correlation arising in a Bell-type experiment *can* however be explained by adding a new causal link to the local hidden variable network [15,18]. Two examples are shown in Fig. 2. In many practical cases, these causal links are not entirely unphysical from the standpoint of respecting relativity and free will, for example, the reason being that experiments do not actually conform to

the exact assumptions Bell made—there are noisy detectors, nonrandom number generation, losses, inability to spacelike separate "Alice" and "Bob," and so on. When this is the case, such causal models are said to be exploiting *loopholes*.

In Fig. 2(b), a causal model that allows correlations between the measurement settings is shown. In the same spirit, we could have had either x or y be causally dependent on λ or another hidden variable. Such models are often called *superdeterministic* and are ruled out by the assumption that Alice and Bob are not colluding and have free will or access to independent randomness. If the experiment only approximately satisfies these assumptions—perhaps due to low detection efficiency—one can still model the data with a local hidden variable said to be exploiting the *detection loophole* [30]. The question of quantifying the amount of independence of the measurement settings necessary has been addressed from multiple perspectives and has practical quantum cryptographic consequences [15–22].

In Fig. 2(c), a causal model which allows correlations between the measurement outcomes is shown. This is, and similar models are, called *nonlocal* and could potentially even allow for superluminal signaling. A quintessential example of a nonlocal model which reproduces the predictions of quantum theory is Bohmian mechanics. Toner and Bacon studied the amount of nonlocality necessary to simulate quantum correlation in the context of classical communication costs [14,23], while Wolf has expressed nonlocality in terms of the compressibility of experimental observations [31].

The current studies, mentioned above, quantifying the relaxations of the causal assumption necessary to replicate quantum correlations are rather disjoint. Recently, Chaves *et al.* placed the question in context of causal networks and found that some measures of these relaxations can be cast as efficiently solvable linear programs [29]. We build on this idea and consider a completely abstract framework amenable to any set of random variables using a single measure of the causal influence of one variable on another. This allows us to consider all possible relaxations simultaneously and thus explore the trade-offs necessary to simulate quantum correlations with hidden variable models.

We will now state our model more technically. For consistency we formulate the problem in the context of the two-party Bell experiment, but we emphasize that this approach generalizes in an obvious way to any set of random variables. A model, M, is specified by a joint distribution

$$\Pr(a,b,x,y,\lambda|M). \tag{6}$$

We label the empirical frequencies $F(a,b,x,y)$ and denote the total variational distance (TVD) of a model to these frequencies by

$$\mathrm{TVD}(M) = \|\Pr(a,b,x,y|M) - F(a,b,x,y)\|_1, \tag{7}$$

where the vector being normed is labeled by (a,b,x,y). Here the 1-norm of a vector x is simply $\|x\|_1 = \sum_i |x_i|$.

The causal influence is defined for a *general* graph as follows:

$$C_{x_i\to x_j}(M) := \max_{x_i,x_i',\mathrm{pa}_j}{}^{*} \|\Pr(x_j|x_i',\mathrm{pa}_j \setminus \mathrm{pa}_j^2, M) - \Pr(x_j|x_i,\mathrm{pa}_j \setminus \mathrm{pa}_j^2, M)\|_1, \tag{8}$$

where $\mathrm{pa}_j \setminus \mathrm{pa}_j^2$ is the set of parents of x_j that are not also grandparents of x_j, and where $\max^*$ indicates that the maximization over x_i, x_i' and pa_j is restricted to feasible assignments. That is, the maximization does not consider assignments outside the support of M. In other words, the causal influence is nonzero when changing x_i leads to a change in x_j. It is quantified by maximizing over latent variables of the target that are not also latent variables of the control.

For example, if we want to minimize the causal influence between two variables a to b in Fig. 2(c) we consider

$$C_{a\to b}(M) := \max_{a,a',y}{}^{*} \|\Pr(b|a,y,M) - \Pr(b|a',y,M)\|_1. \tag{9}$$

We include the conditions $\Pr(a), \Pr(a') \neq 0$ to prevent the causal influence being maximized by an assignment outside the support of the random variable A; the maximization should be taken over all feasible assignments.

Intuitively, this definition represents how distinguishable the different settings of a are when viewed through measurements of b. That is, if a does not causally affect b, then it is not possible for a change in a to be detectable through b alone. We adopt this definition in lieu of the traditional approach of using *interventions*, wherein an external agent imposes a particular value of a while holding all else fixed, effectively cutting out any causal links incident on a other than one originating from the experimentalist themselves. Though some experiments have been performed using intervention to reason about quantum mechanics [32], we cannot intervene on quantum mechanical models in general, such that we must instead maximize over conditions for the experiment, here represented by the maximization over a and y.

The task then is to find a model M which minimizes TVD and $C_{\alpha\to\beta}$ for each $\alpha \to \beta$ ruled out by local causality and measurement independence. If the empirical frequencies contain some causal dependence between two variables, then either the model must also contain such causal dependence or the observed frequencies from the model must be different from the empirical frequencies. Perhaps interestingly, one might be able to "trade" unwanted causal influence between one pair of variables for another, while maintaining the same TVD. Thus the problem of determining "how much" relaxation of Bell's causal assumptions is necessary to match an empirical observed frequency becomes much more interesting and nuanced.

Suppose two models M_1 and M_2 both match the data equally well—i.e., $\mathrm{TVD}(M_1) = \mathrm{TVD}(M_2)$—but M_2 has some unwanted casual influence $a \to b$, say, and M_1 does not—that is, $0 = C_{a\to b}(M_1) < C_{a\to b}(M_2)$. Clearly, M_1 is preferred and we say M_2 is *dominated* by M_1. For many objectives, the situation is more complex but can be handled by the concept of Pareto optimality.

Let $\mathcal{M}$ be the set of all models. Let each model's fitness be represented by the function $f : \mathcal{M} \to \mathbb{R}^n$, where n is the number of objectives. Define the partial order $\prec$ as follows:

$$M \succcurlyeq M' \Leftrightarrow f(M)_k \geqslant f(M')_k, \tag{10}$$

for all $k \in \{0,n-1\}$. If $M \succcurlyeq M'$, we say M' dominates M (or is equivalent to M, if $M' \succcurlyeq M$ holds as well). The set $\mathcal{P} \subset \mathcal{M}$

of *Pareto optimal* models is now defined as follows:

$$\mathcal{P} = \{M \in \mathcal{M} : \{M' \in \mathcal{M} : M \succcurlyeq M', M' \neq M\} = \emptyset\}. \quad (11)$$

This says that a model is Pareto optimal if the set of other models which dominate it is empty. In other words, the Pareto optimal is the set of nondominated models.

IV. ANALYTICAL BOUNDS

In this section, we provide analytical bounds which relate the amount of causal influences exhibited by any model M to its agreement with the empirical frequencies $F(a,b,x,y)$. For the sake of simplicity, we restrict ourselves to analyzing the causal influence between the variables a and b—see Fig. 2(c). However, we emphasize that analogous statements are valid for causal influences between any two variables. For the variables a and b, the empirical frequencies themselves admit a causal influence

$$C_{a\to b}(F) = \max^{*}_{a,a',y} \| F(b|a,y) - F(b|a',y) \|_1, \quad (12)$$

which is defined in complete analogy to (9). To state our theorem, we must define two more quantities. Let $\mathcal{M}_\tau = \mathcal{M}_\tau(F)$ be the set of models having $\mathrm{TVD}(M) \leqslant \tau$ with respect to the empirical frequencies F, and denote by $f^\star = \min_a F(a)$ the minimum empirical marginal frequency.

Theorem 1. For all models $M \in \mathcal{M}_\tau$ and $\tau < 2f^\star$,

$$|C_{a\to b}(F) - C_{a\to b}(M)| \leqslant \frac{2\tau(4f^\star - \tau)}{f^\star(2f^\star - \tau)}. \quad (13)$$

We point out that the bound (13) becomes loose and eventually diverges if the minimum empirical marginal frequency $f^\star$ approaches zero or if the TVD of the class of models becomes too large relative to $f^\star$.

The proof of Theorem 1 can be found in Appendix A.

V. BELL EXPERIMENT AND DATA

As input for the MOEA we use data from a polarization photonic Bell experiment, shown in Fig. 3 [33]. Indistinguishable horizontally polarized ($|H\rangle$) photon pairs are generated via type-1 spontaneous parametric down-conversion. Both polarization qubits are rotated into a diagonal state $\frac{1}{\sqrt{2}}(|H\rangle + |V\rangle)$ by a half wave plate (HWP) with fast axis at $\frac{\pi}{8}$ from vertical, where $|V\rangle$ denotes vertical polarization. A polarization phase rotation is applied to photon 1 by two quarter wave plates (QWPs) and a HWP, while photon 2 has its state optimized by a polarizing beam splitter (PBS). Both photons are collected in polarization maintaining optical fiber (PMF) and are incident on the two input faces of a fiber-coupled PBS, which transmits $|H\rangle$ and reflects $|V\rangle$, preparing Alice's and Bob's qubits. The configuration of the optical fibers results in a σ_x operation applied to Alice's qubit. By measuring in the coincidence basis, we postselect the state

$$\rho = \frac{1+\gamma}{2}|\Phi^+\rangle\langle\Phi^+| + \frac{1-\gamma}{2}|\Phi^-\rangle\langle\Phi^-|, \quad (14)$$

where $|\Phi^\pm\rangle$ are the Bell states $\frac{1}{\sqrt{2}}(|H_A V_B\rangle \pm |V_A H_B\rangle)$ with subscript A (B) corresponding to Alice's (Bob's) qubit. The parameter γ defines the coherence of the state which depends on the overlap of the two photons after the fiber-coupled PBS

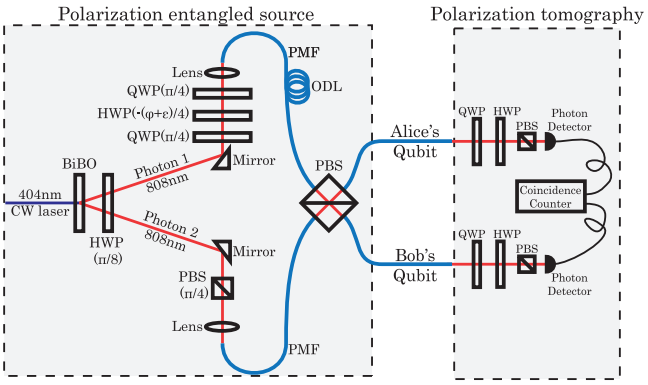

FIG. 3. Photon pairs at 808 nm are emitted via spontaneous parametric down-conversion from a 404 nm pumped bismuth triborate (BiB0) crystal and we prepare polarization enabled photons, labeled "Alice's qubit" and "Bob's qubit." The components used are half wave plates (HWP), quarter wave plates (QWP), polarizing beam splitters (PBS), polarization maintaining fibers (PMF), and an optical delay line (ODL). Once the state is prepared a polarization tomography setup enables projection of each qubit onto any pure polarization state, which is sufficient to perform two-qubit tomography.

and is controlled by the optical delay line (ODL). The state prepared when $\gamma = 1$ is a maximally entangled Bell state $|\Phi^+\rangle$ and when $\gamma = 0$ is an incoherent mixture $\frac{1}{2}(|H_A V_B\rangle\langle H_A V_B| + |V_A H_B\rangle\langle V_A H_B|)$. The polarization tomography setup in Fig. 3 enables projection onto any pure state and can be used for two-qubit state tomography [34]. The photons are detected with silicon avalanche photodiodes and coincidence counts recorded by a timing card.

Our input for the MOEA is a normalized frequency distribution $F(a,b,x,y)$ across binary measurement settings for Alice ($x = \{x_1, x_2\}$) and Bob ($y = \{y_1, y_2\}$), and binary measurement outcomes $a = \{|H_A\rangle, |V_A\rangle\}$ and $b = \{|H_B\rangle, |V_B\rangle\}$, respectively. The measurement settings are controlled by wave plate angles in the tomography and the measurement outcome is the collapse of the state onto one of the four basis states $|H_A H_B\rangle$, $|H_A V_B\rangle$, $|V_A H_B\rangle$, or $|V_A V_B\rangle$. A single measurement is the number of photon pairs recorded for a fixed integration time and can be written as $N^{xy}_{ab} = \mathcal{N}\tau\langle ab|U^{xy}\rho U^{xy\dagger}|ab\rangle$ for measurement settings x,y and measurement outcomes a,b. $\mathcal{N}$ is the total photon flux, τ is the integration time, and U^{xy} is the operation of the wave plates. We calculate $F(a,b,x,y)$ by measuring all combinations of x, y, a, and b, and normalizing by the total number of photon pairs recorded. We note that this experiment is not performed in a loophole-free way, but nonetheless provides us with the quantum correlations we wish to analyze.

Typically, Bell experiments aim to violate the CHSH inequality [35], confirming that quantum mechanical systems cannot be described with local hidden variable models. The CHSH inequality is calculated as

$$|S| \leqslant 2, \quad (15a)$$

$$\text{where} \quad S = \mathbb{E}[x_1 y_1] - \mathbb{E}[x_2 y_1] + \mathbb{E}[x_1 y_2] + \mathbb{E}[x_2 y_2], \quad (15b)$$

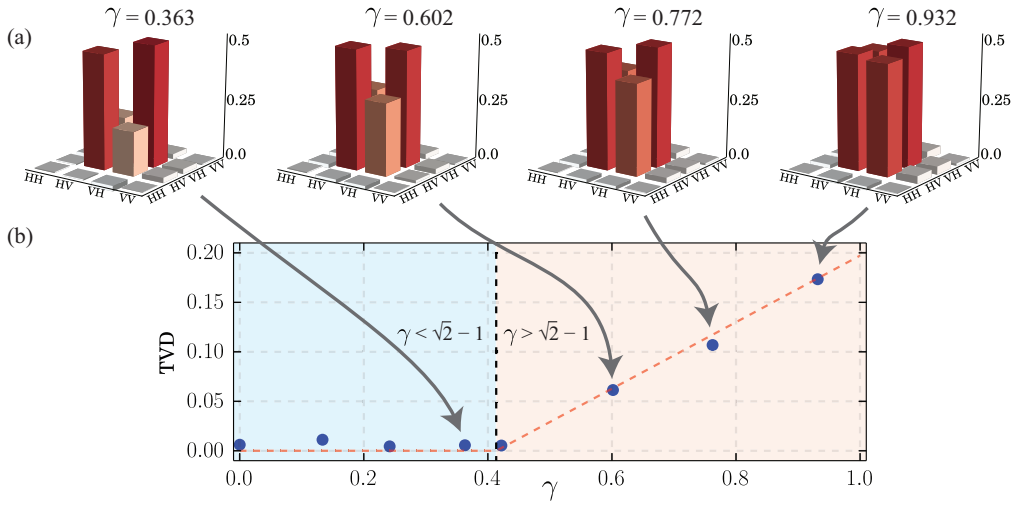

FIG. 4. (a) Real component of experimentally measured density matrices for a range γ values measured via two-qubit state tomography. (b) The results of the MOEA (run as an EA) on a local model see Fig. 2(a) to generate the best achievable TVD for various values of γ from a Bell-like experiment described in Sec. V. As theoretically predicted above a certain threshold ($\gamma = \sqrt{2} - 1$) the local model can no longer explain the measurement results with zero TVD. This threshold corresponds to violating the CHSH inequality in (15). The linear increase in TVD corresponds to the linear increase in S as discussed in Sec. V.

where $\mathbb{E}[xy]$ defines the correlation between Alice's ($x = \{x_1, x_2\}$) and Bob's ($y = \{y_1, y_2\}$) measurements, given as

$$\mathbb{E}[xy] = \frac{N^{xy}_{H_A H_B} - N^{xy}_{H_A V_B} - N^{xy}_{V_A H_B} + N^{xy}_{V_A V_B}}{\mathcal{N}\tau}. \tag{16}$$

While the CHSH inequality holds for systems which respect local causality, a pair of quantum entangled particles can achieve a maximum value of $|S| = 2\sqrt{2}$. By tuning the γ parameter in (14), then for measurement settings fixed to be optimal for the case $\gamma = 1$, we can prepare states that obey the CHSH inequality when $\gamma \leqslant \sqrt{2} - 1$ and states that violate it. In order to achieve the maximum violation of the CHSH inequality, it is necessary to choose specific wave plate angles for x and y. Here, we are not interested in violating the CHSH inequality; however, we can use it to benchmark our results from the MOEA. (See Fig. 4.)

VI. EDGE OF REALITY

Using the experimental data (where $\gamma = 0.984$), we searched for the Pareto optimal models by developing a multiobjective evolutionary algorithm to find the best underlying probability distributions for a causal network. Since this represents a trade-off between a local realistic model and real-world correlations, we call the Pareto optimal surface the "edge" of reality.

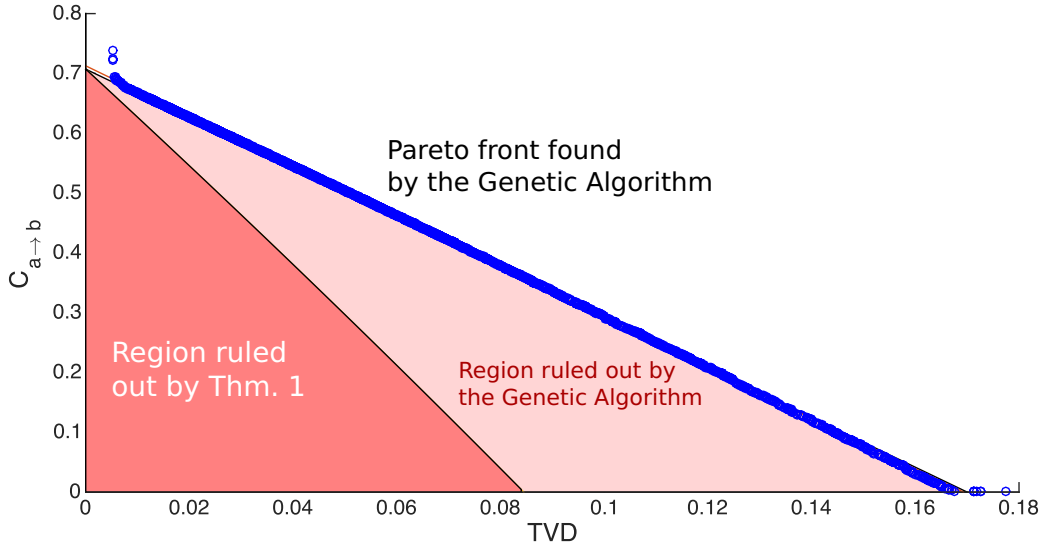

FIG. 5. Pareto front for the causal network in Fig. 2(c) using the data from a photonic Bell experiment. The vertical axis labels the causal influence (9) while the horizontal axis labels the closeness to experimental data (7). The blue circles are the values for the nondominated models found by the evolutionary algorithm. For comparison purposes, the straight line is a linear fit to these data.

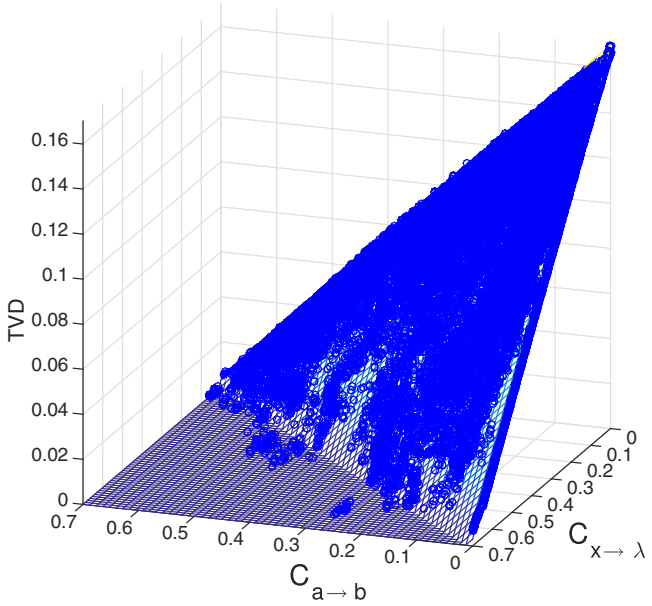

FIG. 6. Pareto front for a local causal network with added $a \to b$ and $x \to \lambda$ edges using the data from a photonic Bell experiment. The vertical axis labels the closeness to experimental data (7). The two horizontal axes label the causal influences (9) for the added edges. The blue circles are the values for the nondominated models found by the evolutionary algorithm. The flat surface is a linear fit to these data.

An *individual* of the population is a probability distribution over the nodes of a given causal network (each such individual is a causal model, M) and its *multiobjective fitness* depends on how close the model can reproduce the experimental data and the amount of causal influences between nonlocally separated variables.

As an initial step, we examined relaxing one casual edge at a time, beginning with a causal influence from a to b—that is, Alice's outcome is allowed to influence Bob's. The Pareto front (the numerical approximation to the Pareto optimal) is shown below in Fig. 5. Like the theoretical bounds (which, while not linear, are nearly so in the considered domain), the front appears to be linear (Pearson's ρ^2 value of 0.997, with bisquare robust fitting). That is, increasing locality violations allows observed (quantum) correlations to be more exactly matched, the trade-off being approximately *linear* in nature. Next, we relax the causal edges $a \to b$ and $x \to \lambda$ simultaneously. The found Pareto front is shown in Fig. 6. Again, we see that the front appears linear ($\rho^2 = 0.9902$). We have also used our algorithm to test other causal networks and found the results to be quantitatively identical to these two cases.

VII. EVOLUTIONARY ALGORITHM

In order to find the Pareto front of solutions, it is necessary to find feasible probability distributions that give rise to the required TVD with the required causal violations(s). There is no known way of doing this analytically. Even in simple single edge causal models the search space is prohibitively large and objective nonconvex. This search space grows rapidly with additional causal edges. Evolutionary algorithms are known methods for finding such Pareto fronts where there is only limited knowledge of the underlying search landscape. We wish to numerically find the Pareto optimal set of models representing Bell experiment data. To do so, we use evolutionary computation [36].

Such algorithms are generally well studied for functions of the form $f : \mathbb{R}^m \to \mathbb{R}^n$. However, here the domain of our objective function f is $\mathcal{M}$, i.e., the probability distributions on the causal network. Consequently, there are implicit constraints on the relative values these distributions can take (for instance, in each node they need to sum to 1) and so we have devised a set of evolutionary operators that allow the probability distribution of an arbitrary causal network to be evolved. With this we combine several evolutionary computation strategies to evolve and explore the Pareto front of a given arbitrary network.

A. Evolutionary algorithm overview

As the cornerstone of our multiobjective evolutionary algorithm (MOEA) we utilize the well-known and well-understood NSGA-II algorithm [37]. Although the NSGA-II algorithm specifies both the generation and selection procedures, we utilize the the DEAP software library [38] which provides the NSGA-II algorithm only for the "select" stage. The method by which we proceed is to use the $(\mu + \lambda)$ algorithm (detailed in Algorithm 1) where we set $\lambda = \mu$ to be the population size. For the purposes of avoiding confusion we note that the $(\mu + \lambda)$ is more properly an algorithm used with a subset of evolutionary algorithms known as evolutionary strategies, and thus is not part of the toolkit of the seperate branch known as genetic algorithms. Consequently, our algorithm is not strictly a genetic algorithm but is an evolutionary algorithm. Although we use an implementation of $(\mu + \lambda)$, by setting $\lambda = \mu$ the algorithm is functionally equivalent to the generation algorithms used in genetic algorithms. In this paper we make no distinction between genetic algorithms and the more general term evolutionary algorithm in the classification of the algorithms used. The overall implementation of the algorithm is thus functionally identical to the original NSGA-II algorithm, save that the selection of parents is random rather than by binary tournament selection.

Consequently this evolutionary algorithm proceeds in *generations*, each of which consists of producing λ offspring from the previous generation's population, then selecting μ individuals from the combination of the previous population and the new offspring to form the new population. As detailed in Algorithm 1 the $(\mu + \lambda)$ algorithm is expressed abstractly in terms of *genetic operators* that create, crossover, evaluate, and select individuals within each population. Thus we form our algorithm by specifying what an individual is, the fitness functions that we use in evaluating individuals, and by providing suitable genetic operators to create "children" causal networks.

B. Representation of individuals

Effectively, our genetic algorithm searches for Pareto optimal models $M \in \mathcal{M}$ by representing M as an assignment of conditional distributions to each node in a causal network

with a fixed structure. Since the random variables at each node are constrained to be discrete, we represent the conditional distributions by *tensors*, such that finding arbitrary joint, marginal, and conditional distributions over subsets of the nodes is then an exercise in standard tensor contractions.

In particular, consider a node x_i with n causal parents $\mathrm{pa}_i = \{x_{i_1}, x_{i_2}, \ldots, x_{i_n}\}$. Then, the distribution $\Pr(x_i|\mathrm{pa}_i) = \Pr(x_i|x_{i_1}, \ldots, x_{i_n})$ is given by the tensor

$$X[j_0, j_1, \ldots, j_n] := \Pr(x_i = j_0 | x_{i_1} = j_1, \ldots, j_n), \qquad (17)$$

where we have used square brackets to indicate indices (similar to C- or Python-style notation).

Algorithm 1 Evolutionary Algorithm

Input: Population sizes $\mu, \lambda \in \mathbb{N}$, crossover and mutation probabilities $p_\times, p_\mu$.
Input: An initial population P_0, a number of generations N_{gen}.
Input: A genetic operator EVALUATE(I) that annotates individuals with their fitness $\vec{f}(I)$.
Input: A genetic operator MUTATE(I) that mutates an individual in-place.
Input: A genetic operator CROSSOVER(I_1, I_2) that crosses over two individuals in-place.
Input: A genetic operator SELECT(P, μ) that selects μ individuals from the population P.
Output: A Pareto front P_* of individuals with respect to the fitness functions implemented by EVALUATE.

▷ In this Algorithm, we follow the DEAP [38] convention of storing an individuals' fitness as *metadata*. This prevents having to reevaluate fitnesses for every comparison.
$P \leftarrow P_0$
$P_* \leftarrow$ KDTREE({}) ▷ Initialize the Pareto front to an empty k-d tree [39].
EVALUATE(P) ▷ Evaluate each individual in the initial population.
for $i_{\mathrm{gen}} \leftarrow 1, \ldots, N_{\mathrm{gen}}$ **do**
  $P_{\mathrm{offspring}} \leftarrow$ {}
  while $|P_{\mathrm{offspring}}| < \lambda$ **do**
    Draw two individuals uniformly at random from P and copy them as I_1 and I_2.
    switch $u \sim \mathrm{Uni}(0,1)$
      case $u \in [0, p_\times)$
        CROSSOVER(I_1, I_2)
      case $u \in [p_\times, p_\times + p_\mu)$
        MUTATE(I_1)
      case $u \in [p_\times + p_\mu, 1]$
        ▷ Leave I_1 and I_2 unmodified.
    end switch
    $P_{\mathrm{offspring}} \leftarrow P_{\mathrm{offspring}} \cup \{I_1\}$
  end while
  EVALUATE($P_{\mathrm{offspring}}$)
  $P \leftarrow$SELECT($P \cup P_{\mathrm{offspring}}, \mu$) ▷ Using the NSGA-II crowding operator, order the individuals and select the next generation from this one and the new offspring.
  for $I \in P$ **do**
    if there does not exist $I' \in P_*$ such that $I' \succeq I$ **then**
  ▷ Average time complexity O($\log |P_*|$) for k-d trees.
      $P_* \leftarrow P_* \cup \{I\}$
    end if
  end for
  if any individuals were added to P_* this generation **then**
      $P_* \leftarrow \{I | I \in P_*$ such that $\forall I' \in P_*, I' \not\preceq I\}$ ▷ Remove dominated individuals from the Pareto front.
      Rebalance P_*.
  end if
end for
return P_*

Algorithm 2 Joint and Conditional Distribution Tensors from Individuals

Input: Individual I, random variables $x_1, \ldots, x_n$, random variables $y_1, \ldots, y_m$.
Output: Tensor $J[i_1, \ldots, i_n, j_i, \ldots, j_m] = \Pr(x_1 = i_1, \ldots, x_n = i_n | y_1 = j_1, \ldots, y_m = j_m)$ for the distribution represented by I.
$X' \leftarrow \{x_1, \ldots, x_n\}$ ▷ X' holds those rvs we must still include.
$F \leftarrow$ {} ▷ F holds those tensor factors we include in the final contraction.
while X' is not empty **do**
  $F \leftarrow F \cup X'$
  $X' \leftarrow \bigcup_{x \in X'} \mathrm{pa}_x \setminus F$ ▷ Add in any parents that we have not already added.
end while
$J \leftarrow$ Einstein sum over of F, holding indices $\{x_1, \ldots, x_n, y_1, \ldots, y_m\}$. ▷ Marginalize over parents not appearing as x or y.
return J

We can contract repeated indices of two such tensors with the tensor at a corresponding third node to perform expectation values. For example, let A be the tensor for $\Pr(a|x,\lambda)$, B be the tensor for $\Pr(b|y,\lambda)$, and Λ be the tensor for $\Pr(\lambda)$ in the model of Fig. 2(c). Then, to find $\Pr(a,b|x,y)$, we compute

$$\Pr(a,b|x,y) = \sum_\lambda A[a,x,\lambda] B[b,y,\lambda] \Lambda[\lambda]. \qquad (18)$$

The general case, allowing for arbitrary numbers of random variables and conditions, is given as Algorithm 2.

C. Fitness functions

Our algorithm uses two different kinds of fitness functions.

(1) The total variational distance (TVD) between the joint distribution computed from an individual and the observed frequencies.

(2) Causal influences along penalized edges, as generalized from the definition given by (9) in Sec. III.

Dealing with each in turn, the TVD is calculated by taking the vector 1-norm between the flattened joint distribution

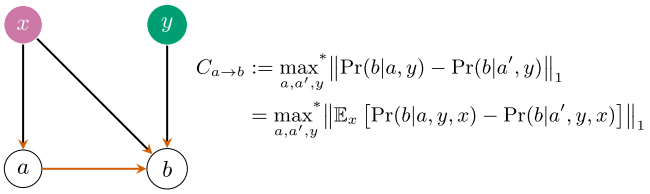

FIG. 7. Example of the causal influence measure $C_{a\to b}(I)$ given by (9) applied to a more complicated graph. The random variable y is conditioned on and maximized over, as it is a parent of b but not of a. By contrast, the variable x is also a parent of a and so it is marginalized over, resulting in the causal influence definition at right.

tensor the observable variables calculated as in Sec. VII B and the flattened observed frequencies,

$$f_{\mathrm{TVD}}(I) = \left\| J^{\flat}_{\mathrm{obs}}(I) - F^{\flat} \right\|_1, \tag{19}$$

where I is an individual with joint distribution tensor $J_{\mathrm{obs}}(I)$ over all observables, F is the tensor of observed frequencies, and where $\flat$ indicates flattening—that is, reduction of an arbitrary-rank tensor to a rank-1 tensor.

As discussed in Sec. III, we adopt a definition of causal influence that allows us to reason even in lieu of interventions. Our definition of the causal influence $C_{a\to b}(I)$ for an individual I proceeds in three steps. First, we maximize over pairs of settings of a to find which are most distinguishable through observations of b alone. We then maximize over the conditions under which these observations are made, represented by maximizing over feasible assignments to the parents of b. Finally, we *marginalize* over those nodes which are also parents of a to prevent "hiding" causal influence; this is illustrated in Fig. 7.

D. Genetic operators

Having defined the mapping from models to individuals, we complete the specification of our algorithm by detailing the various genetic operators which act on these individuals.

a. Creation. In order to create a new individual I, we must specify a new conditional distribution at each node of the causal graph. We do so randomly by assigning a tensor $X_i(I)$ with entries drawn uniformly from $[0,1]$ to each node, then renormalizing to ensure $\sum_{x_i} \Pr(x_i|\mathrm{pa}_i) = 1$. Using the tensor notation defined above,

$$X_i(I)[j_0, j_1, \ldots, j_{|\mathrm{pa}_i|}] = \frac{\tilde{X}_i(I)[j_0, j_1, \ldots, j_{|\mathrm{pa}_i|}]}{\sum_{j'_0} \tilde{X}_i(I)[j'_0, j_1, \ldots, j_{|\mathrm{pa}_i|}]}, \tag{20}$$

where $\tilde{X}_i(I)[j_0, \ldots, j_{|\mathrm{pa}_i|}] \sim \mathrm{Uni}(0,1)$ is the unnormalized tensor of I at X_i.

b. Crossover-mating. Given two individuals I_1 and I_2, we mate them to produce two new individuals I'_1 and I'_2 by swapping the tensors at each node with probability $p_\chi = 0.5$. That is, for each node x_i, the corresponding tensor $X_i(I'_1)$ of I'_1 is given by

$$X_i(I'_1) = \begin{cases} X_i(I_2) & \text{with probability } p_\chi, \\ X_i(I_1) & \text{with probability } 1 - p_\chi. \end{cases} \tag{21}$$

c. Mutation. Given a single individual that has been selected for mutation, we proceed by first picking a node x_μ on the causal graph uniformly at random, with corresponding tensor (assuming n parents) $X_\mu[i_0, \ldots, i_n]$.

One of the conditional events represented by X_μ (that is, a single element of the tensor) is selected at random and the value (and hence the probability assigned to the selected outcome) is randomly increased or decreased by a sample from a zero-mean Gaussian distribution, where the variance is a user supplied parameter. The mutated element of X_μ is then clipped to the interval $[0,1]$, and the relevant tensor index renormalized such that (20) holds.

By way of example, if we had a node a with binary values, which in turn had one parent x also with binary values, then the information pertaining to that node would be stored in a 2×2 tensor $A[a,x]$ corresponding to the probability distribution

$$\Pr(A = a|X = x) = A[a,x] = \begin{array}{c} \\ a=0 \\ a=1 \end{array} \begin{array}{c} \begin{array}{cc} x=0 & x=1 \end{array} \\ \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \end{array}, \tag{22}$$

where α represents the probability of a being zero given x is zero and so on. From this it can be seen that it is necessary that $\alpha + \gamma = \beta + \delta = 1$. One of α, β, γ, or δ would be modified as discussed above, and the remaining values renormalized to ensure that the relevant probabilities continue to sum to 1.

d. Selection. For selection we used an unmodified version of the NSGA-II algorithm [37]. NSGA-II uses a fast sort algorithm to locate the nondominated individuals and then applies a crowding distance sorting algorithm to prefer those individuals that explore different parts of the pareto front. The "best" μ individuals are retained for the next generation.

E. Decomposition of the multiobjective optimization (the island model)

Here we present an enhancement to the basic genetic algorithm discussed above that aids the discovery of the global Pareto front in multidimensional scenarios, where—as is the case here—it is possible to evolve populations to occupy the extremes of any particular front.

As discussed in Sec. VII G it is well known that the NSGA-II crowding becomes less effective with the exponential increase in the size of the front with the number of dimensions. However, in our case we are able to force the population to start at extreme points of the Pareto front by preevolving the population on structurally reduced graphs or with reduced fitness criteria. These populations are able to seed the graph we wish to explore and spread over the front, fleshing it out over multiple runs. This can then be repeated as we increase the dimensions of the fronts. This is not dissimilar to the mechanism used in NeuroEvolution of Augmenting Topologies (NEAT) [40] where populations are evolved on small neural networks prior to allowing additional links to be added. A similar idea of decomposing the objectives is explored in [41]. Effectively, where a multiple dimension Pareto front needs to be explored, different populations are evolved on all permutations of the simpler graphs (on separate

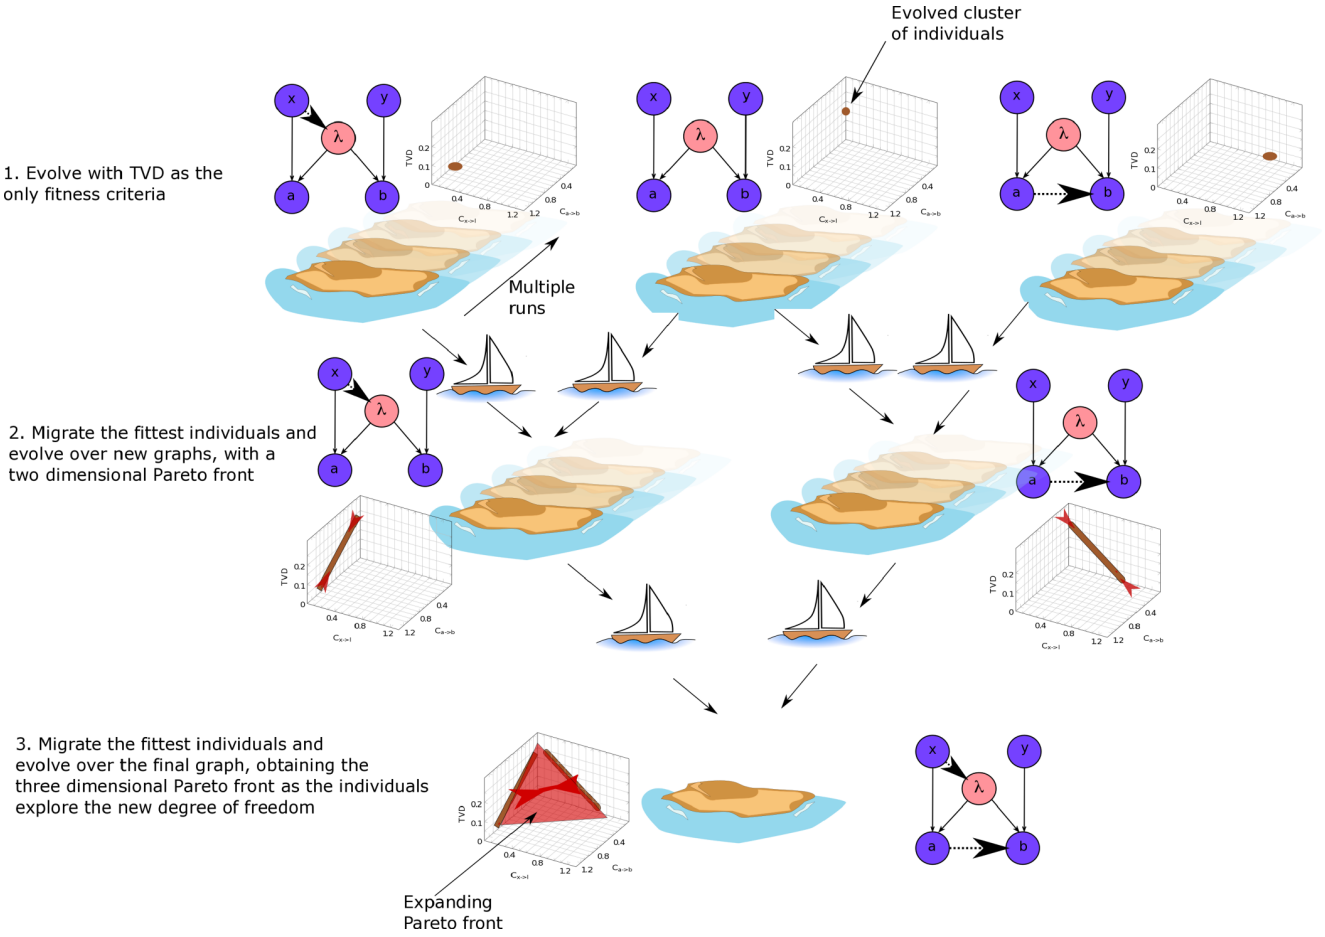

FIG. 8. *Island model* diagram showing the steps in evolving a three-dimensional Pareto front. This allows the edges of the Pareto front to be found by exploring lower-dimensional graphs with lower populations. Multiple runs in each of step 1 and step 2 can be done concurrently.

"islands") before being brought together for evolution over the full graph. This is illustrated in Fig. 8.

This technique allows us to find a three-dimensional Pareto front based on a graph with two causal edges, from $x \to \lambda$ as well as $a \to b$. This was evolved using five runs of the "island model" detailed above. For each run the initial islands had a population of 300 and comprised four runs of 400 generations. The initial runs found populations clustered around the three extremes: (1) min(TVD), hold $C_{x\to\lambda} = 0$, $C_{a\to b} = 0$; (2) min($C_{a\to b}$), hold TVD $= 0$, $C_{x\to\lambda} = 0$; (3) min($C_{x\to\lambda}$), hold TVD $= 0$, $C_{a\to b} = 0$.

The second set of islands take the relevant individuals generated above, reduce them to the best 400 individuals representing the extreme of the Pareto fronts for that island, and transplant them to expanded causal network graphs. In this case there are two second generation islands: one generating the two-dimensional Pareto front for {TVD and $C_{a\to b}$}, with $C_{x\to\lambda}$ held to be zero (i.e., no causal $x \to \lambda$ link); and the second the two-dimensional Pareto front {TVD and $C_{x\to\lambda}$}, with $C_{a\to b}$ held to be zero. These populations are then evolved on the respective causal networks generating two-dimensional Pareto fronts similar to Fig. 10. These populations are placed in an ϵ-dominance archive. (In other words they are only kept if they dominate all previous individuals by at least ϵ, where in this implementation ϵ was $10^{-8} + 10^{-5} \times |\text{value}|$). The entire process so far is repeated several times (in this experiment five times) to ensure we have 2000 suitable individuals in the archive. These individuals are, effectively, clustered on the two-dimensional fronts specified in the second set of islands. This final population is used to generate the three-dimensional Pareto front shown in Fig. 9. The final island had a population of 2000 individuals (extracted initially from the ϵ-dominance archive), evolved for 800 generations. This constituted one run of the island model. The model was run five times, with every individual generated by the model being submitted to (but not necessarily accepted by) the global ϵ-dominance archive.

To illustrate the advantage of using this model, we have also plotted (in red) the best Pareto front found using just the basic algorithm (i.e., evolving only over the full graph). These additional points were collected over eight runs, using a high population (6000) and represented five times the computing power required for the island model. As can be seen the global ϵ-dominance archive for the basic algorithm contained few individuals on the best global Pareto front found by the island model. The Pareto front for the island model (plotted blue) does appear to be a viable candidate for the actual global front, indicating that—for this model—the trade-offs in the different causality violations considered are linear. The front fits a linear plane with a Pearson's ρ^2 value of 0.9902, the nonfitting points being those with extremely low TVD values (the points

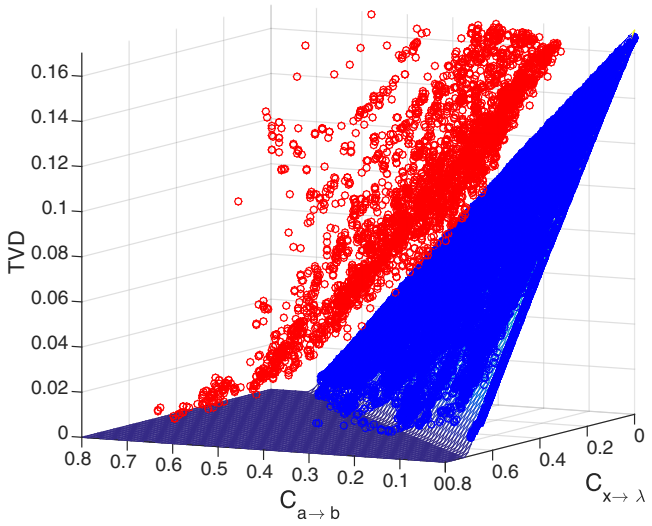

FIG. 9. 3D Pareto front for $a \to b$ and $x \to \lambda$ violations found using the "island model" is shown in blue. Beneath it is a linear mesh which fits the Pareto front with a Pearson's ρ^2 of 0.9902. By comparison the nonoptimal Pareto front found by combining the best individuals from eight runs using the basic algorithm (i.e., no preevolution) on the full graph is shown in red. The front shown in blue took a fraction of the computing time to find compared to the nonoptimal red runs.

which appear on the horizontal part of the mesh). While this still needs to be investigated further we believe it is related to experimental noise which might require increased causal violation to match the noisy data exactly. This is evidence that tradeoffs in multiple causal violations are also linear for such graphs.

F. Previous work and design decisions

Although there has been previous work in using genetic algorithms to explore Bayesian causal networks (e.g., Refs. [42], [43], and [44]), the focus of such works has been to create the network and the links therein. For instance, Ref. [45] uses a multiobjective genetic algorithm (MOGA) to evolve dynamic Bayesian networks. There the multiobjectives explored were the ability of the evolved networks to explain the data, compared with the complexity of the network in question. The MOGA was used instead of, for instance, a minimum description length (MDL) constraint. In all of these cases the network is being used to model something of interest and then, given some observed values, infer the likely causes. The genetic algorithms are used to construct different models which are then trained, with typically the success (or otherwise) of a particular model being its performance on withheld data.

Our work differs because of the way we wish to utilize the Bayesian causal networks; specifically we specify the networks we are interested in, namely those which model a physical view of "reality" with specified local causality violation. Training such networks to replicate observed correlations is of limited interest because successful training results in one specific probability distribution that explains the data. What we are interested in finding are all the relevant probability distributions where the ability to match the observed correlations is contrasted with the strength of the local causality violations. The evolutionary algorithm is used, not to evolve networks, but rather to find these probability distributions given the network. The MOEA is used to guide evolution along these Pareto fronts.

In order to explore the Pareto front some type of MOEA algorithm is required. MOEA on two or three dimensions are relatively well understood. Algorithms to explore large dimensions are still an active area of research (see, for example, Ref. [46]). Since our initial experiments (reported here) would only require causal networks with no penalized edges (a single-value optimization), one penalized edge (a MOEA with a two-dimensional Pareto front), or two penalized edges (a MOEA with a three-dimensional Pareto front), we decided to use the well understood NSGA-II [37]. Although NSGA-II attempts to return the whole of the Pareto front in a single run it was quite clear that the search space (being the required probability assignments for all the nodes in the network) was not smooth, even though the Pareto fronts may be (and, in fact, turned out to be) smooth. Given this an ϵ-dominance archive [47] was maintained and updated through multiple runs. In order to maintain diversity between runs the archive was not used to guide the evolution, but rather served as an updated archive of the best Pareto front found so far. After completing multiple runs, the individuals in the archive thus represent the Pareto front for the entire procedure, rather than for each run taken in isolation.

In Sec. VII E we describe how the ϵ-dominance archive generated from a lower dimension front can be used to seed evolution when a higher-dimensional Pareto front is explored, in a manner not dissimilar to the algorithm presented by Liu *et al.* [41] or utilized by Stanley and Miikkulainen [40].

G. Implementation methodology and details

Our initial runs with the EA (i.e., the MOEA with a single objective, being to minimize the TVD) were used to verify that the EA could match known results. In this case we start with a causal network that reflects Bell's nonlocality assumptions [as shown in Fig. 2(a)] and for various values of γ in (14) use the EA to try and match the experimentally observed joint probability distribution. This is a single-objective EA, with fitness being governed solely by the TVD, i.e., by how closely the observed probability distribution of the model (being the observed joint probability distribution for an individual over the local causal graphs) matches the experimental data. It is known that when γ is less than $\sqrt{2} - 1$ the observed data can be modeled with a local causal network. The TVD values should increase as γ increases to 1 since the empirical distribution no longer factorizes into a locally causal distribution.

An initial population of 300 ($\mu = 300$, $\lambda = 300$) was chosen, with the probability of crossover being 0.1 (and mutation 0.9). The mutation operator used a standard deviation of 0.1 (see Sec. VII D). As is typical for experiments using genetic algorithms no systematic attempt was made to find the "best" parameters for the algorithm. Rather during the course of some initial testing runs, runs with variations of parameters were tried and the parameters of the ones that seemed to find solutions quickest were used. Population sizes reflected those

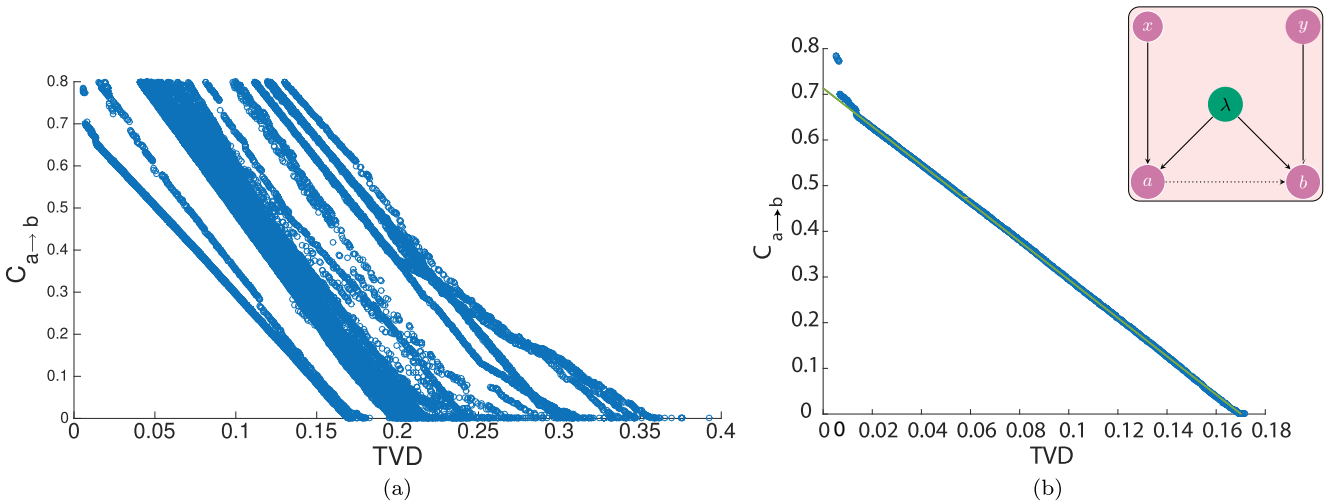

FIG. 10. Results of 40 typical runs of the EA with a causal graph allowing "Alice to Bob" local causality violations see Fig. 2(c). (a) For the purpose of producing this graph, each run had its own ϵ-dominance archive. As can be seen several runs failed to find the correct front at all, indicating that the interplay between hidden nodes and conditioned variables results in a nontrivial search when attempting to match the observed distributions of observable variables. (b) At the conclusion of these 40 runs the ϵ-dominances archive are combined to form the *global* ϵ-dominance archive generating the best estimate of the actual Pareto. The front is well fit by a straight line (Pearson's ρ^2 value of 0.997, with bisquare robust fitting). Exact fitting of the distribution (very low TVD) requires additional causality violation. Experimental noise might be the reason for this. Although around 100 runs were conducted to produce the reported results (Fig. 5), very few additional points were found on the Pareto front.

minimum populations required to avoid runs being trapped early on in local minima. The parameters reported are not reported in a claim of optimality, but rather are reported for the purposes of reproducibility. In any case, once a solution is obtained, its validity does not depend on the means by which it was found.

Figure 4(a) shows experimentally measured density matrices for a range of state γ values. The reduction in coherence is observed as decreasing off-diagonal terms. Figure 4(b) shows the minimum TVD values emerging from 20 runs of the graph for various γ values together with a linear line fitting the data, running from the known y intercept of 0 TVD for $\gamma = (\sqrt{2} - 1)$. As can be seen the EA fits the expected linear results (Pearson's $\rho^2 = 0.9952$).

Having ensured that the algorithm could correctly match the known results on a causal network consistent with Bell's nonlocality assumptions, the next stage is to require a relaxation of local causality to allow the EA (now operating as a MOEA) to match the correlations present in entangled states. As an initial step, we examined relaxing one casual edge at a time, beginning with a causal influence from a to b—that is, Alice's outcome is allowed to influence Bob's. This is illustrated in Fig. 2(c). This becomes a multiobjective problem with a two-dimensional Pareto front, ostensibly well within the capabilities of NSGA-II. The tensor contractions required are not overly complex but with increasing numbers of hidden variables (as a result of additional causal links) each run takes a nontrivial amount of time. As is typical where the search landscape (being the underlying conditional probability distributions) is not smooth (even though the fitnesses such distributions reduce to are smooth) a number of runs failed to converge to any part of the Pareto front, with most runs finding part, but not all of the Pareto front. In Fig. 10(a) we show the individual results of 40 such runs. As can be seen from the figure just under half of the 40 runs had a large percentage of their front nonoptimal, with approximately half the runs being plotted on top of each other on the Pareto front. To generate the final Pareto front each of the individuals in the ϵ-dominance archive from each of the 40 runs are submitted to the *global* ϵ-dominance archive, so that the best estimate of the true Pareto front can emerge, as shown in Fig. 10(b).

Two points arise from these results. The first is that the front appears to be linear, that is increasing locality violations allows observed (quantum) correlations to be more exactly matched, the tradeoff being *linear* in nature. The second arises from the number of failed or only partially successful runs. In particular, we note that while in the majority of runs the MOEA was able to find many points on (or close to) the Pareto front, other runs could be trapped and all runs had difficulty at either extreme of the front. It is clear that the search landscape in general is not smooth—the interplay between the conditioning on the hidden local variable and the other probability distributions allow the MOEA to become trapped in some local minima. The larger front (in this case a two-dimensional line) allowed the population to "slide" away from the edge cases. In addition as observed in [9] it is likely the edge cases represent very specific distributions. Whilst some of these observed difficulties could, in part, be ameliorated by using a larger population and relying on the NSGA-II crowding mechanism to prevent such slippage, as is known this will not be feasible if the front consists of three (or more) dimensions. The front grows exponentially with the number of dimensions, requiring an exponential increase in population size. An alternative MOEA such as NSGA-III may help but each alternative comes with their own difficulties and assumptions. It is, however, possible to use the specifics of the problem space to address these concerns. We know we can evolve the population on a more limited graph (such as

the purely local graph) and force the population to find the lowest TVD with a causal violation of zero (i.e., in the local graph $C_{a\to b} = 0$, since there is no link $a \to b$). This evolved population can then be "transplanted" on to a graph that does have an $a \to b$ link [e.g., Fig. 2(c)]. The other extreme (i.e., lowest $C_{a\to b}$ violation for TVD $= 0$) can be found with a small alteration to the fitness function. To find this point we evolve the population on the graph representing Fig. 2(c) but with a single-objective fitness function, implemented as minimizing the TVD, but where two individuals have the same TVD, the one with the lowest causality violation is preferred. This drives the population towards zero TVD and then minimizes the causality. Even with this the observed correlations were unable to achieve an exact TVD $= 0$, it is speculated this is a result of experimental noise. The ability to generate populations (distributions) that sat at the extreme points of the Pareto front allowed the entire two-dimensional front to be revealed and, as discussed below, can similarly be utilized to reveal the three-dimensional front created by two simultaneously relaxed local causality constraints.

VIII. DISCUSSION

In this work, we have developed a method to allow the study of several relaxations of local hidden variables models simultaneously in a single framework using the tools of casual networks and genetic algorithms.

With further refinement, we hope that our approach can shed light on other scenarios where quantum correlations display richer structure than classical systems would allow. For example, generalizations of the standard Bell scenario to more stations [48,49] and more outcomes [50,51], as well as multiple hidden variables [3,52,53]. In the latter scenario, very little is known since classical correlations are no longer given by linear constraints. Very recently, Chaves has used the framework of causal networks to systematically study such higher-order constraints [54]. Such measured quantities will be particularly useful to our approach as they can be seen as highly relevant coarse grainings of the exponentially growing data space. Such dimension reduction techniques will be crucial for scaling up our numerical algorithm to the analysis of multiparty quantum correlations.

In addition, there is nothing specifically "quantum" about our core numerical methods. Thus our approach should find application outside of the problem of understanding quantum correlations. Recently, Lee and Spekkens have also used inspiration from the causal analysis of quantum correlations to develop causal discovery protocols [55]. Like Lee and Spekkens, we depart from the usual considerations of observed correlation to considering the entire joint probability distribution. Our goals differ, however; whereas the aim of Lee and Spekkens is to find all causal models consistent with data, our goal is to find nondominated models of the plausible correlations. These two approaches are likely to find a harmonious union in the future.

ACKNOWLEDGMENTS

We thank Martin Ringbauer and Andrew White for discussions. This work was supported by the US Army Research Office Grants No. W911NF-14-1-0098 and No. W911NF-14-1-0103, and by the Australian Research Council Centre of Excellence for Engineered Quantum Systems CE110001013. S.T.F. acknowledges support from an Australian Research Council Future Fellowship No. FT130101744. A.P. acknowledges an Australian Research Council Discovery Early Career Researcher Award, Project No. DE140101700 and an RMIT University Vice-Chancellor's Senior Research Fellowship. R.K. acknowledges additional support from the Excellence Initiative of the German Federal and State Governments (Grants ZUK No. 43 and No. 81) and the DFG.

R.H. and R.J.C. contributed equally to this work.

APPENDIX: DERIVATION OF THE INEQUALITY PRESENTED IN SEC. IV

For the sake of being self-contained, let us start this section with reviewing some basic facts about discrete probability distributions and introduce some notation. Throughout this section, we focus on the empirical frequencies $F(a,b,x,y)$ and the probability distribution $\Pr(a,b,x,y|M)$ associated to a fixed model M. Here a,b are the special nodes whose causal relationship is of interest, y will denote the parents that are not grandparents of b, and x is any set of additional random variables which might include hidden variables λ as well as additional measurement outcomes. Therefore, the discussion is completely general and not specific to the models considered, e.g., in the experiment.

If we marginalize these distributions over any variable, say y, we produce new distributions

$$\begin{aligned} F(a,b,x) &= \sum_y F(a,b,x,y), \quad \Pr(a,b,x|M) \\ &= \sum_y \Pr(a,b,x,y|M), \end{aligned} \tag{A1}$$

respectively. As outlined in (A1), we indicate marginalization over any variable, by simply omitting the corresponding variable in the description. Having such a notation at hand, the *product rule* (for discrete probability distribution) assures that as an immediate consequence of the definition of conditional distributions,

$$\begin{aligned} F(a,b,x,y) &= F(a,b,x|y)F(y), \\ \Pr(a,b,x,y|M) &= \Pr(a,b,x|y,M)\Pr(y|M) \end{aligned} \tag{A2}$$

for the variable y. Analogous formulas are true for any combination of the variables present in the distributions [i.e., $\{a,b,x,y\}$ for $F(\cdot)$ and $\{a,b,x,y\}$ for $\Pr(\cdot|M)$].

With these rules and notational concepts at hand, the following statement is an immediate consequence of the triangle inequality.

Lemma 1. Let $F(a,b,x,y)$ and $\Pr(a,b,x,y|M)$ be as above. Then

$$\begin{aligned} \|\Pr(b|M) - F(b)\|_1 &\leqslant \|\Pr(a,b|M) - F(a,b)\|_1 \\ &\leqslant \|\Pr(a,b,x,y|M) - F(a,b,x,y)\|_1 \\ &= \mathrm{TVD}(M). \end{aligned} \tag{A3}$$

This Lemma encapsulates two particular instances of the well-known fact that marginalization contracts the total

variational distance. Since the latter is a measure of how well two probability distributions can be distinguished and marginalization corresponds to ignoring certain variables, Lemma 1 can be intuitively paraphrased as "knowing more doesn't hurt."

Proof of Lemma 1. Inserting the definitions of marginalization and total variational distance yields

$$\begin{aligned}\|\Pr(b|M) - F(b)\|_1 &= \sum_b \left| \sum_a (\Pr(a,b|M) - F(a,b)) \right| \\ &\leqslant \sum_{a,b} |\Pr(a,b|M) - F(a,b)| \\ &= \|\Pr(a,b|M) - F(a,b)\|_1 \end{aligned} \tag{A4}$$

upon employing the triangle inequality. The second inequality can be established in complete analogy. ■

We are now ready to establish the main auxiliary result necessary to establish Theorem 1. It requires the concept of the *harmonic mean* for two variables. For $x_1, x_2 > 0$ the harmonic mean is defined as $H(x_1,x_2) = \frac{2x_1x_2}{x_1+x_2}$.

Lemma 2. Consider two bivariate probability distributions $p(u,v)$ and $q(u,v)$ over finitely many elements labeled by u and v, respectively. Then, the following inequality is valid for any fixed variable v:

$$\begin{aligned}&\|p(u|v) - q(u|v)\|_1 \\ &\quad\leqslant \frac{\sum_u |p(u,v) - q(u,v)| + |p(v) - q(v)|}{H(p(v),q(v))}. \end{aligned} \tag{A5}$$

We point out that this estimate is responsible for introducing the on first sight unfavorable scaling of the bounds (13). However, inequality (A5) is actually tight, making the aforementioned behavior essentially unavoidable. To see this, let $u,v, \in \{0,1\}$ be binary variables and let p be the uniform probability distribution over the four possible joint instances. If one chooses q to be a perfectly correlated bivariate distribution—i.e., $q(0,0) = q(1,1) = 1/2$—it is easy to see that equality is attained in the assertion of Lemma 2.

Proof of Lemma 2. Fix an arbitrary label v. Inverting the product rule allows us to rewrite the left-hand side of (A5) as

$$\|p(u|v) - q(u|v)\|_1 = \left\| \frac{p(u,v)}{p(v)} - \frac{q(u,v)}{q(v)} \right\|_1 = \frac{1}{p(v)q(v)} \sum_u |q(v)p(u,v) - p(v)q(u,v)|. \tag{A6}$$

For $p(v)$ and $q(v)$ we now define

$$\mu := \tfrac{1}{2}(p(v) + q(v)), \quad \delta := \tfrac{1}{2}(p(v) - q(v)),$$

which obey $p(v) = \mu + \delta$ as well as $q(v) = \mu - \delta$ by construction. Inserting these decompositions into (A6) reveals

$$\begin{aligned} p(v)q(v)\|p(u|v) - q(u|v)\|_1 &= \sum_u |(\mu - \delta)p(u,v) - (\mu + \delta)q(u,v)| = \sum_u |\mu(p(u,v) - q(u,v)) - \delta(p(u,v) + q(u,v))| \\ &\leqslant \mu \sum_x |p(u,v) - q(u,v)| + |\delta| \sum_u (p(u,v) + q(u,v)) \\ &= \mu \sum_u |p(u,v) - q(u,v)| + |\delta|(p(v) + q(v)) = \mu \left(\sum_u |p(u,v) - q(u,v)| + 2|\delta| \right), \end{aligned}$$

where we have employed the triangle inequality and the definition of marginalization. Replacing μ and δ with the original expressions then yields

$$\|p(u|v) - q(u|v)\|_1 \leqslant \frac{p(v) + q(v)}{2p(v)q(v)} \left(\sum_u |p(u,v) - q(u,v)| + |p(v) - q(v)| \right).$$

The desired statement then follows from this estimate by identifying the prefactor as $1/H(p(v),q(v))$. ■

We can now show that a bound holds that relates the maximum deviation between the causal influence of any fixed model M and the frequencies F.

Lemma 3. For any fixed model $M \in \mathcal{M}$ and fixed set of empirical frequencies F, let y denote the parents that are not grandparents of the random variable b. Then the following inequality holds:

$$|C_{a\to b}(F) - C_{a\to b}(M)| \leqslant \frac{4\ \mathrm{TVD}(M)}{\min^*_{a,y} H\big(\Pr(a,y|M), F(a,y)\big)}, \tag{A7}$$

where $\min^*$ denotes the minimization over feasible assignments to the variables a, y.

Proof. Choose an arbitrary model $M \in \mathcal{M}$. To ease notation, denote by y the parents that are not grandparents of the variable b. In order to derive the upper bound presented in (13), we start with inserting the definition (9) of $C_{a\to b}(M)$ and observe

$$\begin{aligned} C_{a\to b}(M) &= \max^*_{a,a',y} \|\Pr(b|a,y,M) - \Pr(b|a',y,M)\|_1 \\ &= \max^*_{a,a',y} \|\Pr(b|a,y,M) - F(b|a,y) \\ &\quad - \Pr(b|a',y,M) + F(b|a',y) + F(b|a,y) \\ &\quad - F(b|a',y)\|_1 \\ &\leqslant \max^*_{a,a',y} \|\Pr(b|a,y,M) - F(b|a,y)\|_1 \\ &\quad + \max^*_{a,a',y} \|\Pr(b|a',y,M) - F(b|a',y)\|_1 \end{aligned}$$

$$
\begin{aligned}
&+ \max_{a,a',y}{}^{*} \| F(b|a,y) - F(b|a',y) \|_1 \\
&= 2 \max_{a,y}{}^{*} \| \Pr(b|a,y,M) - F(b|a,y) \|_1 \\
&+ C_{a\to b}(F), \qquad \text{(A8)}
\end{aligned}
$$

where we have identified the last term as the empirical average causal effect defined in (12). As a simple bookkeeping device, let us define $v = (a,y)$ to be the Cartesian product of the random variables a and y. The first term in (A8) can be bounded by invoking Lemma 2 and Lemma 1. Doing so results in

$$
\begin{aligned}
C_{a\to b}(M) - C_{a\to b}(F) &\leqslant 2 \max_{v}{}^{*} \| \Pr(b|v,M) - F(b|v) \|_1 && \text{(A9)} \\
&\leqslant 2 \max_{v}{}^{*} \frac{\sum_b |\Pr(v,b|M) - F(v,b)| + |\Pr(v|M) - F(v)|}{H(\Pr(v|M), F(v))} && \text{(A10)} \\
&\leqslant 4 \max_{v}{}^{*} \frac{\| \Pr(v,b|M) - F(v,b) \|_1}{H(\Pr(v|M), F(v))} && \text{(A11)} \\
&\leqslant 4 \max_{v}{}^{*} \frac{\mathrm{TVD}(M)}{H(\Pr(v|M), F(v))}, && \text{(A12)}
\end{aligned}
$$

which is equivalent to the upper bound presented in (A7). The corresponding lower bound can be derived in a completely analogous fashion by starting off with $C_{a\to b}(F)$ instead of $C_{a\to b}(M)$. ■

This bound is not yet useful because the right-hand side still depends on the unknown model. We seek an inequality that is independent of the model as long as the model has a fixed and sufficiently small value of $\mathrm{TVD}(M)$ with respect to the empirical frequencies. The bound in Sec. IV is a way to avoid this difficulty, and we have now assembled all prerequisites necessary to prove it. We restate the main theorem for completeness.

Theorem 1. Let M denote any model and let $\mathrm{TVD}(M)$, $C_{a\to b}(M)$, and $C_{a\to b}(F)$ be as in (7), (9), and (12), respectively. Denote by $\mathcal{M}_\tau$ the set of models having $\mathrm{TVD}(M) \leqslant \tau$ with respect to the empirical frequencies F, and let $f^\star = \min_{a,y} F(a,y)$, where y denotes all parents of the variable b that are not grandparents of b. Then for all $M \in \mathcal{M}_\tau$ and $\tau < 2f^\star$ we have

$$
|C_{a\to b}(F) - C_{a\to b}(M)| \leqslant \frac{2\tau(4f^\star - \tau)}{f^\star(2f^\star - \tau)}. \qquad \text{(A13)}
$$

Proof of Theorem 1. Again for the sake of bookkeeping we introduce a variable $v = (a,y)$. We begin with the inequality from Lemma 3 and note that we can simply maximize the right-hand side over all $M \in \mathcal{M}_\tau$ to get a universal bound. We have

$$
\begin{aligned}
&\max_{M\in\mathcal{M}_\tau} \frac{4\ \mathrm{TVD}(M)}{\min_v{}^{*} H(\Pr(v|M), F(v))} \\
&\qquad \leqslant \frac{4\ \tau}{\min_{M\in\mathcal{M}_\tau} \min_v{}^{*} H(\Pr(v|M), F(v))}. \qquad \text{(A14)}
\end{aligned}
$$

Therefore, we must establish a lower bound on the denominator. Plugging in the definition of the harmonic mean, a simple calculation confirms that $\partial_x H(x,y) = \frac{2y^2}{(x+y)^2} \geqslant 0$, so the denominator is bounded from below as

$$
\min_{M\in\mathcal{M}_\tau} \min_v{}^{*} \frac{2\ \Pr(v|M) F(v)}{\Pr(v|M) + F(v)} \geqslant \min_{M\in\mathcal{M}_\tau} \min_v{}^{*} \frac{2\ \Pr(v|M) f^\star}{\Pr(v|M) + f^\star}. \qquad \text{(A15)}
$$

Now we relax slightly to allow *all* possible probability distributions (not necessarily ones coming from a causal model M, and denote $\mathcal{P}_\tau$ to be the set of all probability distributions p with $\|p - F\|_1 \leqslant \tau$. Minimizing over a potentially larger set $\mathcal{P}_\tau$ may only decrease the function (or keep its minimum unchanged). We find the denominator is now bounded by

$$
\begin{aligned}
\min_{M\in\mathcal{M}_\tau} \min_v{}^{*} \frac{2\ \Pr(v|M) f^\star}{\Pr(v|M) + f^\star} &\geqslant \min_{p\in\mathcal{P}_\tau} \min_v{}^{*} \frac{2p(v) f^\star}{p(v) + f^\star} \\
&\geqslant \min_{p\in\mathcal{P}_\tau} \frac{2p^\star f^\star}{p^\star + f^\star}. \qquad \text{(A16)}
\end{aligned}
$$

Here in the second inequality we have used the same monotonicity argument for the harmonic mean above (since it is a symmetric function) and replaced the minimum over v with $p^\star = \min_v^{*} p(v)$.

Now we appeal to the monotonicity result of Lemma 1, so that $p \in \mathcal{P}_\tau$ implies that $\|p(v) - F(v)\|_1 \leqslant \tau$. The claim then follows if we can establish the following result:

$$
\min_{p\in\mathcal{P}_\tau} p^\star \geqslant f^\star - \frac{\tau}{2}. \qquad \text{(A17)}
$$

A weaker result, that $\min_{p\in\mathcal{P}_\tau} p^\star \geqslant f^\star - \tau$, is easy to see if we relax the requirement that p is normalized and add the more stringent requirement that $\tau < f^\star$. Begin with the choice $p(a) = F(a)$, and then subtract τ from the smallest component, keeping all other components fixed. This achieves the least value of this relaxed problem. This is a valid solution since the resulting vector is still nonnegative, owing to the constraint $\tau < f^\star$. The slightly tighter result follows from reasserting the constraint that the entries of p must sum to 1, and allows us to weaken the constraint on τ to $\tau < 2f^\star$. With the normalization condition in place, subtracting any deviation of size δ from a component of p must be compensated by adding δ elsewhere in the vector, and this contributes a total of 2δ to the TVD between these differing vectors. The largest such a deviation can be is half of τ, and to minimize our objective function we put this deviation on the smallest component. This component remains positive because of the condition $\tau < 2f^\star$, so this remains a valid probability distribution.

Again by the monotonicity of the harmonic mean, this minimal value can be used to lower bound the denominator.

The final inequality is obtained by plugging in the value of (A17) into the denominator expression and simplifying. ■

We remark that the maximum possible value for $f^\star$ in a Bell experiment where a takes d possible outcomes is $1/d$. Because this inequality is monotonically decreasing with $f^\star$, the bound becomes weaker as the number of outcomes increases, and the requirement that $\tau < 2f^\star$ becomes more demanding.

[1] M. S. Leifer, *Conditional Density Operators and the Subjectivity of Quantum Operations*, AIP Conf. Proc. No. 889 (AIP, Melville, NY, 2007), pp. 172–186.

[2] M. S. Leifer and D. Poulin, Quantum graphical models and belief propagation, Ann. Phys. (NY) **323**, 1899 (2008).

[3] T. Fritz, Beyond Bell's theorem: correlation scenarios, New J. Phys. **14**, 103001 (2012).

[4] J. Fitzsimons, J. Jones, and V. Vedral, Quantum correlations which imply causation, arXiv:1302.2731.

[5] M. S. Leifer and R. W. Spekkens, Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference, Phys. Rev. A **88**, 052130 (2013).

[6] M. S. Leifer and R. W. Spekkens, A Bayesian approach to compatibility, improvement, and pooling of quantum states, J. Phys. A: Math. Theor. **47**, 275301 (2014).

[7] Č. Brukner, Quantum causality, Nat. Phys. **10**, 259 (2014).

[8] E. G. Cavalcanti and R. Lal, On modifications of Reichenbach's principle of common cause in light of Bell's theorem, J. Phys. A: Math. Theor. **47**, 424018 (2014).

[9] C. J. Wood and R. W. Spekkens, The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning, New J. Phys. **17**, 033002 (2015).

[10] K. Ried, M. Agnew, L. Vermeyden, D. Janzing, R. W. Spekkens, and K. J. Resch, A quantum advantage for inferring causal structure, Nat. Phys. **11**, 414 (2015).

[11] A. Einstein, B. Podolsky, and N. Rosen, Can quantum-mechanical description of physical reality be considered complete? Phys. Rev. **47**, 777 (1935).

[12] J. S. Bell, On the Einstein Podolsky Rosen paradox, Physics **1**, 195 (1964).

[13] S. Kochen and E. Specker, The problem of hidden variables in quantum mechanics, Indiana Univ. Math. J. **17**, 59 (1968).

[14] B. F. Toner and D. Bacon, Communication Cost of Simulating Bell Correlations, Phys. Rev. Lett. **91**, 187904 (2003).

[15] J. Barrett and N. Gisin, How Much Measurement Independence is Needed to Demonstrate Nonlocality? Phys. Rev. Lett. **106**, 100406 (2011).

[16] M. J. W. Hall, Complementary contributions of indeterminism and signaling to quantum correlations, Phys. Rev. A **82**, 062117 (2010).

[17] M. J. W. Hall, Local Deterministic Model of Singlet State Correlations Based on Relaxing Measurement Independence, Phys. Rev. Lett. **105**, 250404 (2010).

[18] M. J. W. Hall, Relaxed Bell inequalities and Kochen-Specker theorems, Phys. Rev. A **84**, 022102 (2011).

[19] D. E. Koh, M. J. W. Hall, Setiawan, J. E. Pope, C. Marletto, A. Kay, V. Scarani, and A. Ekert, Effects of Reduced Measurement Independence on Bell-Based Randomness Expansion, Phys. Rev. Lett. **109**, 160404 (2012).

[20] M. Banik, Lack of measurement independence can simulate quantum correlations even when signaling cannot, Phys. Rev. A **88**, 032118 (2013).

[21] L. P. Thinh, L. Sheridan, and V. Scarani, Bell tests with min-entropy sources, Phys. Rev. A **87**, 062121 (2013).

[22] G. Pütz, D. Rosset, T. J. Barnea, Y.-C. Liang, and N. Gisin, Arbitrarily Small Amount of Measurement Independence is Sufficient to Manifest Quantum Nonlocality, Phys. Rev. Lett. **113**, 190402 (2014).

[23] K. Maxwell and E. Chitambar, Bell inequalities with communication assistance, Phys. Rev. A **89**, 042108 (2014).

[24] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2009).

[25] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction and Search* (MIT Press, Cambridge, MA, 2001).

[26] B. Hensen, H. Bernien, A. E. Dreau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellan, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres, Nature (London) **526**, 682 (2015).

[27] M. Giustina, M. A. M. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-A. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, and A. Zeilinger, Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons, Phys. Rev. Lett. **115**, 250401 (2015).

[28] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici, A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam, Strong Loophole-Free Test of Local Realism*, Phys. Rev. Lett. **115**, 250402 (2015).

[29] R. Chaves, R. Kueng, J. B. Brask, and D. Gross, Unifying Framework for Relaxations of the Causal Assumptions in Bell's Theorem, Phys. Rev. Lett. **114**, 140403 (2015).

[30] P. M. Pearle, Hidden-variable example based upon data rejection, Phys. Rev. D **2**, 1418 (1970).

[31] S. Wolf, Nonlocality without counterfactual reasoning, Phys. Rev. A **92**, 052102 (2015).

[32] M. Ringbauer, C. Giarmatzi, R. Chaves, F. Costa, A. G. White, and A. Fedrizzi, Experimental test of nonlocal causality, Sci. Adv. **2**, e1600162 (2016).

[33] R. J. Chapman, M. Santandrea, Z. Huang, G. Corrielli, A. Crespi, M.-H. Yung, R. Osellame, and A. Peruzzo, Experimental perfect state transfer of an entangled photonic qubit, Nat. Commun. **7**, 11339 (2016).

[34] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Measurement of qubits, Phys. Rev. A **64**, 052312 (2001).

[35] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Proposed Experiment to Test Local Hidden-Variable Theories, Phys. Rev. Lett. **23**, 880 (1969).

[36] J. H. Holland, *Adaptation in natural and artificial systems* (University of Michigan Press, Ann Arbor, 1975).

[37] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Trans. Evol. Comput. **6**, 182 (2002).

[38] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné, DEAP: Evolutionary algorithms made easy, J. Machine Learning Res. **13**, 2171 (2012).

[39] J. L. Bentley, Multidimensional binary search trees used for associative searching, Commun. ACM **18**, 509 (1975).

[40] K. O. Stanley and R. Miikkulainen, Evolving neural networks through augmenting topologies, Evol. Comput. **10**, 99 (2002).

[41] H. L. Liu, F. Gu, and Q. Zhang, Decomposition of a multiobjective optimization problem into a number of simple multiobjective subproblems, IEEE Trans. Evol. Comput. **18**, 450 (2014).

[42] A. G. a. C. de Sá and G. L. Pappa, Towards a method for automatically evolving Bayesian network classifiers, in *Proceedings of the 15th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '13 Companion* (ACM, New York, 2013), pp. 1505–1512.

[43] J. Muruzábal and C. Cotta, *Advances in Probabilistic Graphical Models* (Springer, Berlin, 2007), pp. 193–213.

[44] P. Larrañaga, H. Karshenas, C. Bielza, and R. Santana, A review on evolutionary algorithms in Bayesian network learning and inference tasks, Inf. Sci. **233**, 109 (2013).

[45] B. J. Ross and E. Zuviria, Evolving dynamic Bayesian networks with multi-objective genetic algorithms, Appl. Intell. **26**, 13 (2006).

[46] Y. Yuan, H. Xu, and B. Wang, An improved NSGA-III procedure for evolutionary many-objective optimization, in *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation, GECCO '14* (ACM, New York, 2014), pp. 661–668.

[47] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler, Combining convergence and diversity in evolutionary multiobjective optimization, Evol. Comput. **10**, 263 (2002).

[48] R. F. Werner and M. M. Wolf, All-multipartite Bell-correlation inequalities for two dichotomic observables per site, Phys. Rev. A **64**, 032112 (2001).

[49] M. Żukowski, Č. Brukner, W. Laskowski, and M. Wieśniak, Do All Pure Entangled States Violate Bell's Inequalities for Correlation Functions? Phys. Rev. Lett. **88**, 210402 (2002).

[50] D. Collins, N. Gisin, N. Linden, S. Massar, and S. Popescu, Bell Inequalities for Arbitrarily High-Dimensional Systems, Phys. Rev. Lett. **88**, 040404 (2002).

[51] D. Collins and N. Gisin, A relevant two qubit Bell inequality inequivalent to the CHSH inequality, J. Phys. A: Math. Gen. **37**, 1775 (2004).

[52] C. Branciard, N. Gisin, and S. Pironio, Characterizing the Nonlocal Correlations Created via Entanglement Swapping, Phys. Rev. Lett. **104**, 170401 (2010).

[53] J. Henson, R. Lal, and M. F. Pusey, Theory-independent limits on correlations from generalized Bayesian networks, New J. Phys. **16**, 113043 (2014).

[54] R. Chaves, Polynomial Bell Inequalities, Phys. Rev. Lett. **116**, 010402 (2016).

[55] C. M. Lee and R. W. Spekkens, Causal inference via algebraic geometry: Feasibility tests for functional causal structures with two Binary observed variables, Journal of Causal Inference (2017).

Chapter 8

Non-exponential fidelity decay in randomized benchmarking with low-frequency noise

We show that non-exponential fidelity decays in randomized benchmarking experiments on quantum-dot qubits are consistent with numerical simulations that incorporate low-frequency noise and correspond to a control fidelity that varies slowly with time. By expanding standard randomized benchmarking analysis to this experimental regime, we find that such non-exponential decays are better modeled by multiple exponential decay rates, leading to an *instantaneous control fidelity* for isotopically purified silicon metal-oxide-semiconductor quantum-dot qubits which is 98.9% when the low-frequency noise causes large detuning, but can be as high as 99.9% when the qubit is driven on resonance and system calibrations are favourable. These advances in qubit characterization and validation methods underpin the considerable prospects for silicon as a qubit platform for fault-tolerant quantum computation.

Nonexponential fidelity decay in randomized benchmarking with low-frequency noise

M. A. Fogarty,[1] M. Veldhorst,[1] R. Harper,[2] C. H. Yang,[1] S. D. Bartlett,[2] S. T. Flammia,[2] and A. S. Dzurak[1]
[1]*Centre for Quantum Computation and Communication Technology, School of Electrical Engineering and Telecommunications, The University of New South Wales, Sydney, New South Wales 2052, Australia*
[2]*Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, New South Wales 2006, Australia*
(Received 16 February 2015; published 11 August 2015)

We show that nonexponential fidelity decays in randomized benchmarking experiments on quantum-dot qubits are consistent with numerical simulations that incorporate low-frequency noise and correspond to a control fidelity that varies slowly with time. By expanding standard randomized benchmarking analysis to this experimental regime, we find that such nonexponential decays are better modeled by multiple exponential decay rates, leading to an *instantaneous control fidelity* for isotopically purified silicon metal-oxide-semiconductor quantum-dot qubits which is 98.9% when the low-frequency noise causes large detuning but can be as high as 99.9% when the qubit is driven on resonance and system calibrations are favorable. These advances in qubit characterization and validation methods underpin the considerable prospects for silicon as a qubit platform for fault-tolerant quantum computation.

DOI: 10.1103/PhysRevA.92.022326 PACS number(s): 03.67.Lx, 03.67.Ac

I. INTRODUCTION

Randomized benchmarking experiments [1,2] quantify the accuracy of quantum gates by estimating the average decay in control fidelity as a function of the number of operations applied to a qubit. Benchmarking enjoys several advantages over the traditional methods of characterizing gate fidelity that involve quantum process tomography [3,4], namely, that it is insensitive to state preparation and measurement (SPAM) errors and scales efficiently with the system size. As such, benchmarking protocols (see Fig. 1) have become a standard against which different qubit technologies and architectures are compared. Benchmarking experiments have been performed in many different technologies, including trapped ions [1,5,6], superconducting qubits [7–9], nuclear-magnetic-resonance architectures [10], nitrogen-vacancy centers in diamond [11], semiconductor quantum dots in silicon [12], and phosphorous atoms in silicon [13]. Most experiments are fitted using an exponential fidelity decay, which is in line with original theoretical predictions [1,14] and consistent with the assumption of weak correlation between noise on the gates that is important for accurate fidelity estimates.

When the assumptions of randomized benchmarking are violated, there is no guarantee of observing the characteristic exponential decay curves determined by the average fidelity. This has been noted before in NMR experiments due to spatial inhomogeneity across the sample [10] as well as in numerical simulations [15] of $1/f$ noise and leakage to states outside of the computational subspace. Recent experimental results in spin-based silicon metal-oxide-semiconductor (Si-MOS) quantum-dot qubits [12] have also shown nonexponential fidelity decay, and here we directly apply our theoretical modeling to these experiments, but our conclusions are widely applicable.

Here we argue that nonexponential fidelity decay in this semiconductor qubit is indicative of a dephasing-limited decay caused by non-Markovian noise. We first propose a numerical simulation method that incorporates time-dependent effects, primarily a drift in frequency detuning. This detuning drift and other time-dependent low-frequency noise sources lead to decay curves that are effectively integrated over an ensemble of experimental results, each with slightly different "instantaneous" average fidelities, i.e., fidelities that are approximately stable over the course of a single benchmarking run but that drift over the course of the entire sequence of experiments. These simulations show good qualitative agreement with the observed nonexponential decay from the experiments on isotopically purified silicon quantum-dot qubits [12]. We then give a more quantitative analysis that compares two very simple models that both give good fits to the data: the first is a simplified version of the drift model that postulates that each experimental run has one of only two possible average fidelities; the second model attributes the nonexponential decay to fluctuating SPAM errors. Both of these models have only one additional parameter over the standard benchmarking model, but our quantitative likelihood analysis shows that the simplified drift model is much more probable.

The conclusion of this analysis for the Si-MOS quantum-dot qubit is that, while the total average fidelity over a long series of benchmarking runs is 99.6% [12], the instantaneous fidelity can be as high as 99.9% or more when naturally fluctuating environmental noise sources and system calibrations are most favorable. We emphasize that consistent high fidelities such as these may be within reach: further improvements in the system calibration, such as more frequent and accurate estimates of the qubit detuning, could allow these high fidelities directly by exploiting the low-frequency character of the noise. Achieving such high fidelities for single-qubit gate operations gives optimism for exceeding the demanding error thresholds for fault-tolerant quantum computation.

II. BENCHMARKING REVIEW

The standard randomized benchmarking procedure involves subjecting a quantum system to long sequences of randomly sampled Clifford gates followed by an inversion step and a measurement, as depicted in Fig. 1. The unitary operations of the Clifford group G are those that map the set of Pauli operators to itself under conjugation. They are a discrete set of gates that exactly reproduce the uniform average

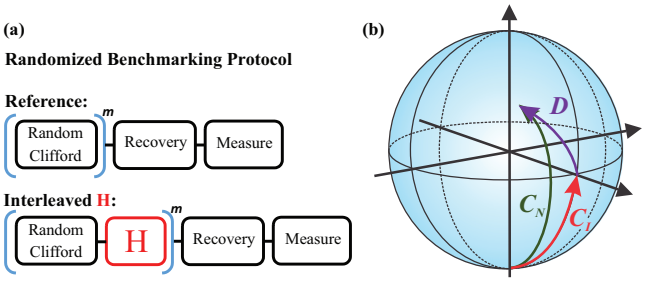

FIG. 1. (Color) (a) Randomized benchmarking consists of applying multiple sequences of random Clifford gates, a final recovery Clifford gate to ensure that each sequence ends with the qubit in an eigenstate, and reading out the qubit state. In interleaved randomized benchmarking, an additional test Clifford gate is inserted in between the random Clifford gates. (b) Bloch sphere representation for the breakdown of a general noisy operation $\mathcal{C}_N$ into an ideal $\mathcal{C}_I$ rotation followed by a noise operation $\mathcal{D}$.

gate fidelity, averaged over the set of all input pure states [16]. An alternate version known as interleaved benchmarking [8] inserts a systematic application of a given gate, such as the H gate shown in Fig. 1. The difference from the reference sequence gives information about the specific average gate fidelity of the given gate, rather than the average fidelity additionally averaged over the ensemble of gates.

Consider a general noise process $\mathcal{D}$, depicted in Fig. 1, which represents the deviation of a noisy Clifford gate $\mathcal{C}_N$ from an ideal unitary Clifford operation $\mathcal{C}_I$:

$$\mathcal{C}_N = \mathcal{D}\mathcal{C}_I\,.$$

We note that the above equation uses the formalism of completely positive maps [17], and the multiplication corresponds to composition of maps. The standard approach to randomized benchmarking makes the assumption that $\mathcal{D}$ does not depend on the choice of $\mathcal{C}_I$ or other details such as time, but our simulations and, of course, real experiments will include such a dependence.

The fundamental result of randomized benchmarking [2] is that for sufficiently well behaved noise the observed fidelities depend on only the average error operation $\mathcal{E}_\mathcal{D}$ averaged over the Clifford group G given by

$$\mathcal{E}_\mathcal{D} = \frac{1}{|\mathsf{G}|}\sum_{\mathcal{C}_I\in\mathsf{G}} \mathcal{C}_I\mathcal{D}\mathcal{C}_I^{-1}\,,$$

as well as any SPAM errors present in the system. Furthermore, standard tools from representation theory reduce this average error operation to one that is nearly independent of $\mathcal{D}$ and is characterized by just a single parameter p. In particular, it is a depolarizing channel $\mathcal{E}$, with $p = p(\mathcal{D})$ being the polarization parameter (i.e., the probability of the information remaining uncorrupted as it passes through the channel). For a d-dimensional quantum system, the action of the depolarizing channel is given by

$$\mathcal{E}(\rho) = p\rho + (1-p)\frac{\mathbb{1}}{d}\,,$$

and the polarization parameter is related to the noisy deviation $\mathcal{D}$ by the average gate fidelity $\bar{\mathcal{F}}_{\rm avg}(\mathcal{D})$ according to [2]

$$\bar{\mathcal{F}}_{\rm avg}(\mathcal{D}) = \int \mathrm{d}\psi\,\langle\psi|\mathcal{D}(|\psi\rangle\langle\psi|)|\psi\rangle = p + \frac{1-p}{d}\,, \qquad (1)$$

where the integral is a uniform average over all pure states.

For a randomized benchmarking sequence composed of $m+1$ total Clifford gates (including the $+1$ for the recovery operation), the average sequence fidelity is given by [2]

$$\bar{F}_m = Ap^m + B\,. \qquad (2)$$

Here the parameters A and B quantify the SPAM errors and are given by [2]

$$A = \mathrm{Tr}[E\mathcal{D}(\rho - \mathbb{1}/d)], \quad B = \mathrm{Tr}[E\mathcal{D}(\mathbb{1}/d)]\,,$$

and ρ and E are the noisy state preparations and measurements implemented instead of the ideal desired states and measurements.

A typical benchmarking experiment proceeds by estimating $\bar{F}_m$ for several values of m and fitting to the model in Eq. (2) to extract the p, A, and B fit parameters and then using Eq. (1) to report an ensemble average of the average gate fidelities $\bar{\mathcal{F}}_{\rm avg}$ of the gates.

This derivation of Eq. (2) assumes certain features about the noise, namely, that it has negligible time and gate dependence and that non-Markovian effects are not present at time scales on the order of the gate time. The limits to the validity of this assumption have been probed before [15,18,19], and in particular, it was noted via numerical simulations by Epstein *et al.* [15] that the exponential model of fidelity decay no longer holds in the presence of $1/f$ noise, resulting in a noise floor to the accuracy of the benchmarking experiment.

III. NONEXPONENTIAL FIDELITY DECAY

A clear deviation from the fidelity decay predicted by Eq. (2) has been observed in a silicon quantum-dot qubit [12]. In order to understand the possible origin of this deviation, we have used the qubit characteristics to numerically simulate randomized benchmarking with a realistic noise model. In the experiment, the qubit is defined by the spin state of a single electron. A magnetic field $B_0 = 1.4$ T is applied to create a Zeeman splitting, and the qubit is operated using electron-spin-resonance (ESR) techniques by applying an ac magnetic field with frequency $\omega_0 = \frac{g\mu_B B_0}{\hbar}$. A Rabi π pulse is realized in $\tau_{op} = 1.6$ μs, and using a Ramsey sequence, the dephasing time $T_2^* = 120$ μs has been obtained with state-preparation and measurement fidelities of 95% and 92%, respectively [12]. In between consecutive pulses, a waiting time $\tau_w = 0.5$ μs has to be incorporated due to the operation of the analog microwave source.

The set of Clifford gates is generated using the set $[\pm X, \pm\frac{1}{2}X, \pm Y, \pm\frac{1}{2}Y]$ that is realized using Rabi pulses and the identity simulated with a waiting time equal to a π pulse.

A nonexponential fidelity decay can be caused by leakage, where population in the two-qubit states is lost to other levels [15,20,21]. For example, in Ref. [15] it has been shown that in the presence of a third level, the sequence fidelity for large m could approach $1/3$ instead of $1/2$ (although this benchmarking protocol used a different gate set than the

standard one). While leakage is an important aspect in multidot qubits, such as singlet-triplet qubits or exchange-only qubits that possess accessible nonqubit spin states, a qubit encoded in a spin-1/2 system is inherently two-dimensional. Higher energy levels of the quantum dot, or valley degeneracies, represent different degrees of freedom, rather than leakage channels. Leakage can occur through loss of the electron, but we note that the qubit system experiences a T_1 time on the order of seconds, and we have measured the absence of tunneling during a sequence. As further evidence of negligible leakage, we note that, experimentally, we observe that the spin-up and -down fractions are symmetric around the half visibility plus offset, as observed in Figs. 4(b) and 4(c) below.

A. Non-Markovian noise in a quantum-dot qubit

Within the experiment, Ramsey sequences have been performed in between benchmarking sequences to recalibrate the resonance frequency of the qubit and to compensate drifts due to, for example, the superconducting magnet. These drifts, in combination with errors in setting the resonance frequency, cumulate in a time dependency within the system and result in an apparent T_2^* for the randomized benchmarking experiment. This decoherence time is also dependent upon the duration of the data acquisition. The non-Markovian noise processes that are expected to determine T_2^* can be modeled as a random walk of the detuning $\Delta\omega$ away from the ideal operation frequency ω_0 over time scales greater than a single run of a random Clifford sequence. In order to simulate an ensemble of results, the $\Delta\omega$ term is selected randomly from a Gaussian distribution of normalized variance:

$$\sigma_{\mathrm{op}} = \frac{\tau_{\mathrm{op}}}{2\pi\sqrt{2\ln(2)}T_2^*}.$$

Using this distribution, we have numerically simulated benchmarking experiments, and the results are shown in Fig. 2. In this simulation we have approximated the time scale of the low-frequency noise to be on the order of a single run, and as such, the detuning is constant over a single trace. However, between each trace the detuning is sampled randomly from a distribution, as shown in Figs. 2(b) and 2(c). The individual traces correspond to a given detuning $\Delta\omega$ and result in the "instantaneous" fidelity of the qubit. While the individual traces are decaying exponentially, the averaged fidelity (thick blue line) obtained from the Gaussian ensemble is clearly nonexponential. We have also included the case of a Lorentzian distribution of detunings (red), resulting in a nonexponential decay as well. In the simulation, the only error source is dephasing, whereas in the experiment, other errors might be present such as pulse errors. Inclusion of such errors will still result in nonexponential decays, provided dephasing is a significant source of error.

As low-frequency drift of the qubit resonance frequency can lead to nonexponential fidelity decay, we hypothesize that some ensemble of experiments with varying decay rates is the correct explanation for the nonexponential behavior of the experimental benchmarking data [12]. To support this hypothesis, we use the Akaike information criterion to show that a simple model allowing for differing fidelity rates better explains the data than an alternative explanation that assumes fluctuating SPAM errors in the standard (zeroth-order) model.

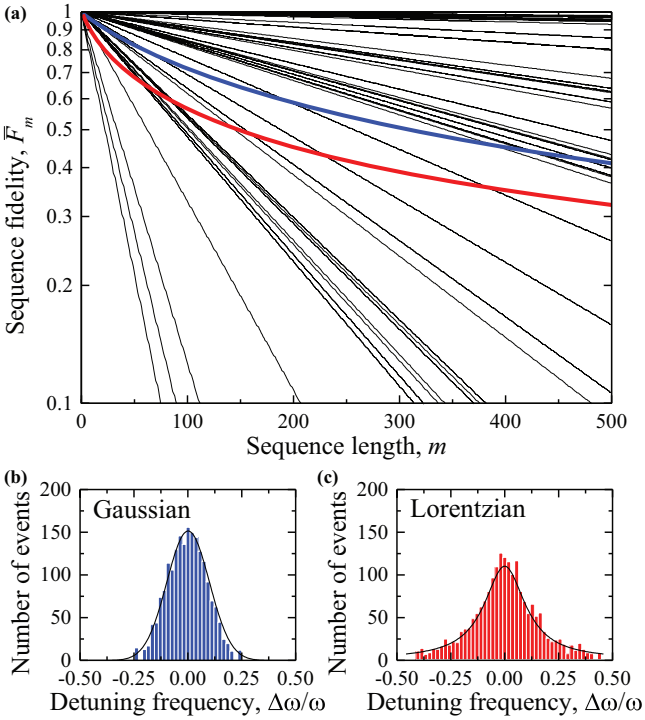

FIG. 2. (Color) (a) Sequence fidelity as a function of sequence length m, with the qubit subject to Gaussian distributed T_2^* associated noise. Each black line represents a fidelity decay for one particular value of detuning $\Delta\omega$ (only 10% of traces shown). The linear decay on the logarithmic scale illustrates that these individual traces are indeed exponential, while the ensemble average (thick blue line) is nonexponential. The thick red line is the ensemble average for a Lorentzian distributed noise. (b) Gaussian distributed detuning frequencies and (c) Lorentzian distributed detuning frequencies associated with individual traces.

B. Eliminating the constant for a single-qubit randomized benchmarking model

The parameters A and B in Eq. (2) are nuisance parameters that do not convey information about the desired control fidelity. Eliminating one of these parameters, in this case B, will further constrain the zeroth-order model and allows deviations to be more clearly identifiable. A further advantage of removing the parameter B is to allow fitting of a linear function on a log-linear plot where deviation from standard assumptions of randomized benchmarking will show clearly as a nonexponential decay trace. In Ref. [12] the randomized benchmarking protocol was modified to eliminate B from the zeroth-order model. We first provide a theoretical justification for this approach that conforms with the standard assumptions of randomized benchmarking. We note that this approach applies only to qubits ($d = 2$) and demonstrate the deviation of the measured data from the expected exponential is highlighted via this method. Recall that the zeroth-order model fits the average fidelity of a gate sequence to a simple formula as

follows [2]:

$$\bar{F}_m^{\uparrow} = A^{\uparrow} p^m + B^{\uparrow}\,, \tag{3}$$

where the qubit is initialized as $|\uparrow\rangle\langle\uparrow|$, the final gate in the random benchmarking sequence is chosen to return the state to $|\uparrow\rangle\langle\uparrow|$, and $\bar{F}_m^{\uparrow}$ is the survival probability of this state. To eliminate the constant $B^{\uparrow}$ from this sequence, we perform a second set of similar randomized sequences, with the difference being that the final $(m+1)$th gate is set to change the state to $|\downarrow\rangle\langle\downarrow|$. For these runs, we consider the survival probability for yielding the measurement outcome $E^{\downarrow}$, where in the ideal case the final state $\rho = E^{\downarrow} = |\downarrow\rangle\langle\downarrow|$. This is the survival probability for each run $\bar{F}_m^{\downarrow}$. Under the same assumptions we have

$$\bar{F}_m^{\downarrow} = A^{\downarrow} p^m + B^{\downarrow}\,. \tag{4}$$

Combining these two equations by defining $\tilde{F}_m \equiv \bar{F}_m^{\uparrow} - (1 - \bar{F}_m^{\downarrow})$, we have

$$\tilde{F}_m = \tilde{A} p^m + (B^{\uparrow} + B^{\downarrow}) - 1\,, \tag{5}$$

where $\tilde{A} = A^{\uparrow} + A^{\downarrow}$.

Recall that $B^{\uparrow} = \mathrm{Tr}[E^{\uparrow}\mathcal{D}(\mathbb{1}/d)]$, where $\mathcal{D}$ is the average noise operator. For the $\bar{F}_m^{\downarrow}$ runs, the derivation is identical, apart from the final change to the $|\downarrow\rangle\langle\downarrow|$ state using a π pulse $\mathcal{X}$, so we have $B^{\downarrow} = \mathrm{Tr}[E^{\downarrow}\mathcal{D}(\mathcal{X}(\mathbb{1}/d))]$. One expects that $\mathcal{X}$ is close to unital, i.e., $\mathcal{X}(\mathbb{1}/d) \simeq \mathbb{1}/d$. Under the assumption that this is true (an assumption that will be respected up to violations no larger than the gate infidelity) and noting that $E^{\uparrow} + E^{\downarrow} = \mathbb{1}$ for qubits ($d = 2$) and that $\mathcal{D}$ is trace preserving, $B^{\downarrow}$ can be reexpressed as follows:

$$\begin{aligned} B^{\downarrow} &= \mathrm{Tr}[E^{\downarrow}\mathcal{D}(\mathbb{1}/2)] = \mathrm{Tr}[(\mathbb{1} - E^{\uparrow})\mathcal{D}(\mathbb{1}/2)] \\ &= \mathrm{Tr}[\mathcal{D}(\mathbb{1}/2)] - \mathrm{Tr}[E^{\uparrow}\mathcal{D}(\mathbb{1}/2)] = 1 - B^{\uparrow}\,. \end{aligned} \tag{6}$$

Therefore by subtracting the average results of the data set $(1 - \bar{F}_m^{\downarrow})$ from the average results of the data set $\bar{F}_m^{\uparrow}$ we can obtain a data set that is distributed according to the model

$$\tilde{F}_m = \tilde{A} p^m \tag{7}$$

under the standard benchmarking assumptions on the noise.

The data from Ref. [12] consist of eight data sets (one reference set and seven interleaved sets). The experiment, in order of operation, comprised measuring 50 single shots which were randomly distributed over $\bar{F}_m^{\uparrow}$ and $\bar{F}_m^{\downarrow}$. The sequence randomization protocol was carried out ten times, first for the reference and then for all seven interleaved sets. For each data set, gate sequences of lengths $m \in \{2,3,5,8,13,21,30,40,50,70,100,150\}$ were measured. This entire process was then repeated 50 times for a total of 2 400 000 single-shot measurements. The fast Ramsey recalibrations were performed at approximately 10-min intervals.

This amount of randomization is at least an order of magnitude more than in previous experiments [1,5,7–11,13]. Each randomized protocol was performed 50 times in order to estimate the survival probability.

C. Analysis of experimental data

To quantify the quality of our fits, we require estimates of the variance in the data of Ref. [12]. The observed variance of the data matched to within 5%–40% of the theoretical upper bounds derived in Ref. [18] when the gate length was shorter than 20 [so that the $m(1 - \bar{\mathcal{F}}_{\mathrm{avg}}) \ll 1$ assumption discussed in that reference was satisfied]. Accordingly, the observed experimental variance was used as a reliable estimate of the actual variance of the distribution. It should be noted that the observed variance actually decreased for gate lengths of 100 or greater. One explanation for this unexpected behavior is that some of the sequences become saturated to something close to a completely mixed state before reaching those sequence lengths.

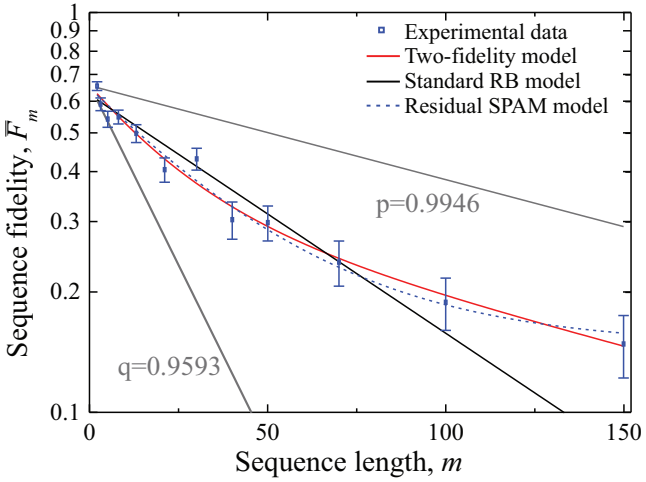

FIG. 3. (Color) Semilog plot of $F_m^{\uparrow} - (1 - F_m^{\downarrow})$ for the reference sequence of randomized benchmarking on a silicon quantum-dot qubit [12]. Both the two-fidelity model and a single-fidelity model including residual SPAM can fit the data, but for the single-fidelity model an unreasonably large SPAM has to be included.

Figure 3 shows the data from the reference data set plotted on a semilog plot. The confidence bounds are 95%, and the data are clearly nonlinear (i.e., the decay is not a simple exponential). A similar deviation from the linear fit was noted in each of the data sets, with the best-fit linear model consistently underestimating $\bar{F}_m$ for $m \geqslant 100$. Reference [2] outlines a higher (first) order fitting for the fidelity decay which includes gate-dependent noise. With the elimination of the parameter B as outlined above, the inclusion of higher orders results in $\tilde{F}_m^1 = \tilde{A} p^m + \tilde{C}(m-1)(q - p^2)p^{m-2}$. The first-order equation did not fit the data significantly better than the zeroth-order equation.

Two other possible explanations are considered. First, it may not be possible to entirely eliminate the constant term (B) due to a violation of one of the assumptions in the above derivation. A second explanation is that low-frequency noise leads to detuning and hence time-dependent errors on the gates in some of the experiments. The first, which we denote the *residual SPAM model*, can be modeled by reverting to a formula of the form $\tilde{F}_m = \tilde{A} p^m + \tilde{B}$, where now $\tilde{B}$ represents residual SPAM errors that were not eliminated under the assumptions that led to the derivation of Eq. (7). We consider the simplest possible model for the second explanation, the *two-fidelity model*, by fitting the fidelity decay to a formula of the form $\tilde{F}_m = \tilde{A} p^m + \tilde{A} q^m$. This represents an attempt to

model the data by simplifying the ensemble of experiments by reducing them to just two different *equally weighted sequence behaviors*: one with a high-fidelity rate (related by the usual measure to p) and one with a lower-fidelity rate (similarly related to q). This model has fewer parameters than the Gaussian or Lorentzian drift models and is much easier to fit. In this interpretation, we have successfully eliminated the B parameter as per Eq. (7), but time variation gives us the two different polarization parameters p and q, with the decay rate for each sequence sampled randomly with equal probability. As can be seen in Fig. 3, both models fit the data substantially better than the simple exponential of the zeroth-order model.

Although the residual SPAM model produces a good fit to the experimental data, it does so with the equivalent of an unusually large SPAM parameter $\tilde{B}$ of around 0.14 corresponding to an individual fitting of $B^{\uparrow} = 0.56$ and $B^{\downarrow} = 0.58$. This represents in the theoretical model a very large bias in the expectation value of the spin-up measurement on the asymptotic value of the sequence fidelity away from the theoretical value of 0.5, which is not observed in the experiment.

To compare the residual SPAM and two-fidelity models quantitatively, it is possible to calculate the log likelihood and Akaike information criterion [22] for the two models. Because we do not have the actual distribution of the test statistic, we make the assumption that the samples contained in the underlying data are independent and the Gaussian distributed limit is appropriate. This assumption is well justified as we have a large number of independent data sets. The distribution $\tilde{F}_m$ can therefore be approximated by a Gaussian distribution with a variance estimated by the observed variance at each gate length. The log likelihood of the observed data, given each of the two models, can then be calculated using standard methods as follows.

The Akaike information criteria used was $2\times$ (number of parameters $-$ loglikehood). The log likelihood is calculated as

$$\sum_{s=\substack{\text{sequence}\\ \text{lengths}}} \left[\ln\left(\frac{1}{\sqrt{2\pi\sigma_s^2}}\right) - \frac{1}{2\sigma_s^2}[\mu_s - x_s]^2 \right], \qquad (8)$$

where σ_s^2 is the variance for sequence s, μ_s is the measured average at that sequence, and x_s is the predicted average.

Table I shows the calculated Akaike information criterion for each of the experimental data sets. As can be seen, the two-fidelity model better explains the data, significantly so on all but one of the data sets. Although such a model is a simplified version of the drift model, the fact that it fits the data well and is physically motivated supports its adoption as the most likely explanation of the nonexponential curve seen in the data.

D. Interpreting the two-fidelity model

Since the two-fidelity model is the quantitatively preferred model, a natural question arises: How should we interpret the model parameters? The obvious interpretation of parameters p and q is as presented in Table II, i.e., that their difference represents the characteristic spread of the actual underlying ensemble of fidelities from which the benchmarking data are sampled. Such an interpretation is natural and compelling; however, it remains an open problem to quantify such a connection more carefully. In particular, it would be interesting to give a direct connection to a more general drift model since these models are easier to interpret physically but much harder to fit and analyze statistically.

By considering the nonexponential decay manifesting as the average over an ensemble of results, the fidelity can be considered to be operating under two regimes, as depicted in Fig. 4(a). First, dominating the observed fidelity decay at low m, there is a rapid decay rate dominated by traces of large detuning $\Delta\omega$. Second, for large m, these traces of large detuning will fast approach the constant B and so will not influence the decay slope in the fidelity; the fidelity decay rate for large m is then governed by long-lived traces of smaller detuning.

In Fig. 4(a), each of the data points is an average over 25 000 experimental repetitions as presented by the two accompanying histograms [Fig. 4(b) for $m = 2$ and Fig. 4(c) for $m = 150$]. Each histogram separately shows the measured probability, averaged over 50 repetitions, for the spin-up and spin-down observables as expected at the end of a noiseless

TABLE I. Akaike information criterion for standard and interleaved randomized benchmarking. The comparison column specifies how many times as probable the $\tilde{A}p^m + \tilde{A}q^m$ model to minimize information loss is compared to the $\tilde{A}p^m + \tilde{B}$ model.

	Akaike information criteria		
Data set	$\tilde{A}p^m + \tilde{B}$	$\tilde{A}p^m + \tilde{A}q^m$	Comparison
Ref.	−16.93	−25.29	65.44
I	−46.19	−57.12	238.10
X	−54.52	−59.99	15.43
X/2	−62.89	−63.79	1.56
−X/2	−57.77	−64.34	26.69
Y	−36.06	−50.43	1317
Y/2	−36.04	−46.39	172.00
−Y/2	−46.37	−63.32	4815

TABLE II. Calculated p and q values for the two-fidelity model. The gate fidelity estimates $\bar{\mathcal{F}}_{\text{avg}}$ reported for the reference run are the high (p) gate fidelity estimate and low (q) gate fidelity, with 95% confidence margins of ±0.06% and ±0.5%, respectively. We further note that calculated p and q values will result in an inaccurate interleaved gate fidelity as given by the process outlined by Ref. [8] due to the low-frequency noise.

Data set	p	q
Ref.	0.995	0.959
I	0.993	0.946
X	0.993	0.952
X/2	0.993	0.947
−X/2	0.991	0.947
Y	0.993	0.964
Y/2	0.991	0.952
−Y/2	0.990	0.911
Data set	$\bar{\mathcal{F}}^p_{\text{avg}}$	$\bar{\mathcal{F}}^q_{\text{avg}}$
Ref.	99.9%	98.9%

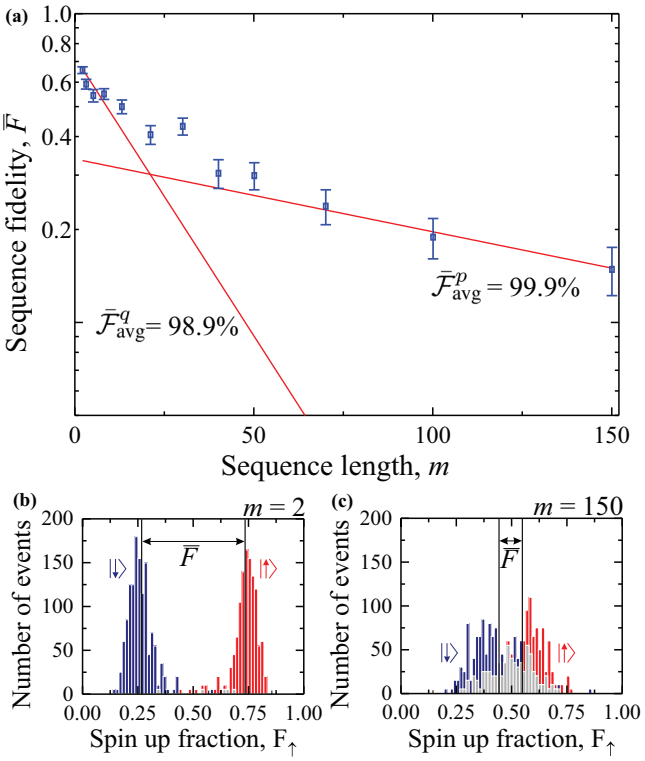

FIG. 4. (Color) (a) Reference sequence of randomized benchmarking on a silicon quantum-dot qubit [12]. The separate fidelities from the two-fidelity model have been plotted to show how the initial decay is dominated by the low q value, whereas the higher value of p is indicative of the average decay in the longer-lived high-fidelity regime. Histogram of spin up $|\uparrow\rangle$ and spin down $|\downarrow\rangle$ corresponding to data point (b) $m = 2$ and (c) $m = 150$. Results with expected spin-up outcome are shown in red, while blue represents data with expected spin-down result. The gray regions illustrate the overlapping areas.

version of the applied random sequence. From Fig. 4 we infer that the instantaneous fidelity approaches a peak fidelity of 99.9% when the microwave frequency is on resonance and that the fidelity can drop to 98.9% in the presence of large detunings. The time-averaged performance will lead to a fidelity in between, consistent with the 99.6% fidelity quoted in Ref. [12].

A similar analysis can also be applied to the interleaved gate sequences to extract average fidelities for the individual gates. See Table II for p and q values for each set. For each of the interleaved gates, we find a high and a low gate fidelity comparable to the two values quoted for the standard benchmarking scheme above, although the numerical instability in calculating the gate fidelities for interleaved benchmarking leads to larger uncertainties in these values, of the order of 1% for 95% confidence margins. We note that a naive application of the method to derive fidelities for interleaved gate sequences as outlined by Ref. [8] will, for certain gates such as the Y gate, yield a fidelity $\bar{\mathcal{F}}_{\text{avg}} > 100\%$. Such results may be an indication of correlated low-frequency noise, where some Clifford gates can echo out noise; such an effect would result in decays which are slower than that of the reference set and break the assumptions of interleaved randomized benchmarking. However, the large uncertainty in our estimated fidelities for interleaved gates due to statistical noise does not allow us to definitively test for such an effect with the current data set.

IV. CONCLUSIONS

We have analyzed the nonexponential decay in randomized benchmarking experiments on Si-MOS quantum-dot qubits and found that the most plausible explanation of this decay is drift in detuning frequencies. Our simulation of temporal integration over a spectrum of time-dependent detuning frequencies qualitatively reproduces the observed fidelity decay of previously conducted experiments [12]. In addition, we have quantitatively ruled out a competing model by showing agreement of a simplified ensemble (the two-fidelity model) that is much more probable. This yields confidence that detuning drift is the correct explanation for the origin of such a nonexponential fidelity decay.

Low-frequency noise leads to a time-varying fidelity that is relatively constant over a given sequence but can vary between sequences. We have therefore defined an "instantaneous fidelity" to characterize the performance of a gate during a single sequence, and we can consider how this instantaneous fidelity varies in time from sequence to sequence. Fitting the randomized benchmarking data with a two-fidelity model demonstrates that silicon MOS quantum-dot qubits can already exhibit an instantaneous control fidelity of 99.9%. This is achieved when the system is correctly calibrated and the microwave frequency is on resonance. However, when the noise causes large detuning, the fidelity drops to 98.9%. We anticipate that the higher fidelity can be achieved consistently (i.e., made time independent) as improvements in the readout fidelity appear feasible and better calibration could be obtained by performing optimized Ramsey protocols to calibrate the resonance frequency for each experiment [23].

These results raise several intriguing questions. The first is to quantitatively link the simple and easy to analyze two-fidelity model to the Gaussian or Lorentzian drift model. Alternatively, directly fitting a drift ensemble to the data would give a better picture of the source of the nonexponential fidelity decay, but this approach risks overfitting and is already difficult for the simple case of Gaussian distributed detuning.

Second, there are several different models discussed in this work, each of which is capable of analyzing different forms of breaches in the standard assumptions of randomized benchmarking. The reduced parameter representation for the fidelity $\tilde{F}$ not only allows for higher accuracy in fitting but is capable of immediately identifying a deviation from the expected result under standard assumptions. Further it is capable of identifying the presence of certain types of noise if $\tilde{F}$ is to asymptote to a nonzero offset.

Finally, there is at least one other natural competing explanation for the nonexponential decay. It might be the case that long benchmarking sequences saturate the exponential decay rates and have slower decay on very long time scales. If this were the case, then fitting to sequences that were "too long" would certainly bias one toward seeing nonexponential decay and reporting fidelities that were higher than warranted by the analysis. Therefore deriving stopping criteria for the

maximum sequence length and deriving tests that rule out this alternate explanation is a further important open question for future work.

ACKNOWLEDGMENTS

We thank C. Ferrie and C. Granade for helpful discussions. The authors acknowledge support from the Australian Research Council (CQC2T - CE110001027 and EQuS - CE11001013), the NSW Node of the Australian National Fabrication Facility, the U.S. Army Research Office (W911NF-13-1-0024, W911NF-14-1-0098, W911NF-14-1-0103 and W911NF-14-1-0133), and iARPA via the MQCO program. M.V. also acknowledges support from the Netherlands Organization for Scientific Research (NWO) through a Rubicon Grant. S.T.F. also acknowledges support from an ARC Future Fellowship (FT130101744).

[1] 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, Phys. Rev. A **77**, 012307 (2008).

[2] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and robust randomized benchmarking of quantum processes, Phys. Rev. Lett. **106**, 180504 (2011).

[3] I. Chuang and M. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, J. Mod. Opt. **44**, 2455 (1997).

[4] J. F. Poyatos, J. I. Cirac, and P. Zoller, Complete characterization of a quantum process: The two-bit quantum gate, Phys. Rev. Lett. **78**, 390 (1997).

[5] J. P. Gaebler *et al.*, Randomized benchmarking of multiqubit gates, Phys. Rev. Lett. **108**, 260503 (2012).

[6] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit, Phys. Rev. Lett. **113**, 220501 (2014).

[7] J. M. Chow, J. M. Gambetta, L. Tornberg, Jens Koch, Lev S. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Randomized benchmarking and process tomography for gate errors in a solid-state qubit, Phys. Rev. Lett. **102**, 090502 (2009).

[8] E. Magesan *et al.*, Efficient measurement of quantum gate error by interleaved randomized benchmarking, Phys. Rev. Lett. **109**, 080505 (2012).

[9] R. Barends *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, Nature (London) **508**, 500 (2014).

[10] C. A. Ryan, M. Laforest, and R. Laflamme, Randomized benchmarking of single- and multi-qubit control in liquid-state NMR quantum information processing, New J. Phys. **11**, 013034 (2009).

[11] F. Dolde *et al.*, High-fidelity spin entanglement using optimal control, Nat. Commun. **5**, 3371 (2014).

[12] M. Veldhorst *et al.*, An addressable quantum dot qubit with fault-tolerant control-fidelity, Nat. Nanotechnol. **9**, 981 (2014).

[13] J. T. Muhonen *et al.*, Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking, J. Phys. Condens. Matter **27**, 154205 (2015).

[14] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, J. Opt. B **7**, S347 (2005).

[15] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, Investigating the limits of randomized benchmarking protocols, Phys. Rev. A **89**, 062321 (2014).

[16] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, Phys. Rev. A **80**, 012304 (2009).

[17] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[18] J. J. Wallman and S. T. Flammia, Randomized benchmarking with confidence, New J. Phys. **16**, 103032 (2014).

[19] C. Granade, C. Ferrie, and D. G. Cory, Accelerated randomized benchmarking, New J. Phys. **17**, 013042 (2015).

[20] J. J. Wallman, M. Barnhill, and J. Emerson, Robust characterization of loss rates, arXiv:1412.4126.

[21] J. Wallman, C. Grenade, R. Harper, and S. T. Flammia, Estimating the Coherence of Noise, arXiv:1503.07865.

[22] H. Akaike, A new look at the statistical model identification, IEEE Trans. Autom. Control **19**, 716 (1974).

[23] M. D. Shulman, S. P. Harvey, J. M. Nichol, S. D. Bartlett, A. C. Doherty, V. Umansky, and A. Yacoby, Suppressing qubit dephasing using real-time hamiltonian estimation, Nat. Commun. **5**, 5156 (2014).

Chapter 9

Estimating the coherence of noise

Noise mechanisms in quantum systems can be broadly characterized as either coherent (i.e., unitary) or incoherent. For a given fixed average error rate, coherent noise mechanisms will generally lead to a larger worst-case error than incoherent noise. We show that the coherence of a noise source can be quantified by the *unitarity*, which we relate to the average change in purity averaged over input pure states. We then show that the unitarity can be efficiently estimated using a protocol based on randomized benchmarking that is efficient and robust to state-preparation and measurement errors. We also show that the unitarity provides a lower bound on the optimal achievable gate infidelity under a given noisy process.

IOP Publishing New J. Phys. 17 (2015) 113020 doi:10.1088/1367-2630/17/11/113020

New Journal of Physics
The open access journal at the forefront of physics

Deutsche Physikalische Gesellschaft DPG
IOP Institute of Physics

Published in partnership with: Deutsche Physikalische Gesellschaft and the Institute of Physics

OPEN ACCESS

RECEIVED
20 April 2015

REVISED
23 August 2015

ACCEPTED FOR PUBLICATION
5 October 2015

PUBLISHED
5 November 2015

PAPER

Estimating the coherence of noise

Joel Wallman[1,3], **Chris Granade**[2], **Robin Harper**[2] **and Steven T Flammia**[2]

[1] Institute for Quantum Computing and Department of Applied Mathematics, University of Waterloo, Waterloo, Canada
[2] Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, Australia
[3] Author to whom any correspondence should be addressed.

E-mail: joel.j.wallman@gmail.com

Keywords: characterization, randomized benchmarking, coherent errors

Supplementary material for this article is available online

Abstract

Noise mechanisms in quantum systems can be broadly characterized as either coherent (i.e., unitary) or incoherent. For a given fixed average error rate, coherent noise mechanisms will generally lead to a larger worst-case error than incoherent noise. We show that the coherence of a noise source can be quantified by the *unitarity*, which we relate to the average change in purity averaged over input pure states. We then show that the unitarity can be efficiently estimated using a protocol based on randomized benchmarking that is efficient and robust to state-preparation and measurement errors. We also show that the unitarity provides a lower bound on the optimal achievable gate infidelity under a given noisy process.

To harness the advantages of quantum information processing, quantum systems have to be controlled to within some maximum threshold error. Certifying whether the error is below the threshold is possible by performing full quantum process tomography [1, 2], however, quantum process tomography is both inefficient in the number of qubits and is sensitive to state-preparation and measurement errors (SPAM) [3].

Randomized benchmarking [4–9] and direct fidelity estimation [10, 11] have been developed as efficient methods for estimating the average infidelity of noise to the identity. However, the worst-case error, as quantified by the diamond distance from the identity, can be more relevant to determining whether an experimental implementation is at the threshold for fault-tolerant quantum computation [12]. The best possible bound on the worst-case error (without further information about the noise) scales as the square root of the infidelity and can be orders of magnitude greater than the reported average infidelity [13, 14].

However, this scaling of the worst-case bound is only known to be saturated by unitary noise. If the noise is known to be stochastic Pauli noise, the worst-case error is directly proportional to the average infidelity [9], vastly improving on the general bound. Consequently, quantifying the intermediate regime between unitary and fully incoherent noise may allow the bound on the worst-case error to be substantially improved.

Randomized benchmarking is also emerging as a useful tool for diagnosing the noise in an experiment [15, 16], which can then be used to optimize the implementation of gates by varying the experimental design. In this spirit, an experimental protocol for characterizing the coherence of a noise channel will be an important tool as the quest to build a fault-tolerant quantum computer progresses.

In this paper, we present a protocol for estimating a particular quantification of the coherence of noise, which we term the *unitarity*, in the experimental implementation of a unitary 2-design. Our protocol is efficient and robust against SPAM, and is a minor modification of randomized benchmarking. The unitarity is defined as the average change in the purity of a pure state after applying the noise channel, with the contributions due to the identity component subtracted off, (see equation (4)) and is closely related to the purity of the Jamiołkowski isomorphic state (see proposition 9). We show that the unitarity is invariant under unitary gates and attains its maximal value if and only if the noise is unitary. Furthermore, we show that the unitarity can be combined with the average gate fidelity to quantify how far a noise channel is from depolarizing noise. Finally, we show that the unitarity of a noise channel provides a lower bound on the best achievable gate infidelity assuming perfect unitary control.

Our approach to quantifying coherence complements other recent work on quantifying coherence since we focus on the coherence of quantum *operations* rather than the coherence of quantum *states* relative to a preferred basis [17].

1. Defining unitarity

We begin by defining the unitarity of a noise channel $\mathcal{E}$: $\mathcal{B}(\mathbb{C}^d) \to \mathcal{B}(\mathbb{C}^d)$, that is, a completely positive (CP) linear map that takes quantum states to quantum states. The purity of a quantum state ρ is $\mathrm{Tr}\ \rho^\dagger\rho \in [0,\ 1]$ with $\mathrm{Tr}\ \rho^\dagger\rho = 1$ if and only if ρ is a pure state. An initial candidate for a definition of the unitarity of $\mathcal{E}$ is

$$\int \mathrm{d}\psi\ \mathrm{Tr}\Big[\mathcal{E}(\psi)^\dagger \mathcal{E}(\psi)\Big], \tag{1}$$

that is, as the purity of the output states averaged over all pure state inputs. However, this definition is problematic, since it would lead to the nonunital state-preparation channel

$$\mathcal{E}_0(\rho) = \mathrm{Tr}(\rho)|0\rangle\langle 0| \tag{2}$$

having the same value of unitarity as a unitary channel, even though it does not preserve coherent superpositions. Similarly, the (trace-decreasing) filtering channel

$$\mathcal{E}_1(\rho) = |0\rangle\langle 0|\ \rho\ |0\rangle\langle 0| \tag{3}$$

does not preserve coherent superpositions and so should have the same unitarity value as a complete depolarizing channel. Both of these problematic channels arise when either the identity is mapped to coherent terms or *vice versa*.

To avoid these issues, we define the unitarity of a noise channel to be the average purity of output states, with the identity components subtracted, averaged over all pure states. That is, we define

$$u(\mathcal{E}) = \frac{d}{d-1}\int \mathrm{d}\psi\ \mathrm{Tr}\Big[\mathcal{E}'(\psi)^\dagger \mathcal{E}'(\psi)\Big], \tag{4}$$

where the normalization factor is chosen so that $u(\mathcal{I}) = 1$ and $\mathcal{E}'$ is defined so that $\mathcal{E}'(A) = \mathcal{E}(A) - [\mathrm{Tr}\ \mathcal{E}(A)/\sqrt{d}]\mathbb{1}$ for all traceless A (to account for trace-decreasing channels, such as in equation (2)) and $\mathcal{E}'(\mathbb{1}_d) = 0$ (to account for non-unital channels, such as in equation (3)). Equivalently, if $\{A_1, \ldots, A_{d^2}\}$ is any set of traceless and trace-orthonormal operators (e.g., the normalized Paulis), then we can define the generalized Bloch vector $\mathbf{n}(\rho)$ of a density operator ρ with unit trace to be the vector of $d^2 - 1$ expansion coefficients

$$\rho = \mathbb{1}_d/d + \sum_{k>1} n_k A_k. \tag{5}$$

Our definition of the unitarity is then equivalent to

$$u(\mathcal{E}) = \frac{d}{d-1}\int \mathrm{d}\psi\ \Big\|\mathbf{n}[\mathcal{E}(\psi)] - \mathbf{n}\Big[\mathcal{E}\Big(\mathbb{1}_d/d\Big)\Big]\Big\|^2, \tag{6}$$

that is, the average squared length (i.e., Euclidean norm) of the generalized Bloch vector after applying the map $\mathcal{E}$ with the component due to the identity subtracted off.

2. The estimation protocol

We now present a protocol for characterizing the unitarity of the noise in an experimental implementation of a unitary 2-design $\mathcal{G}$ under the assumption that the experimental implementation of any $U \in \mathcal{G}$ can be written as $\mathcal{U} \circ \mathcal{E}$ where $\mathcal{U}$ denotes the channel corresponding to conjugation by U and $\mathcal{E}$ is a completely positive, trace-preserving (CPTP) channel independent of U. (Note that, as in all randomized benchmarking papers, the assumption that $\mathcal{E}$ is independent of $\mathcal{U}$ can be relaxed without dramatically effecting the results [8, 13, 16].)

The protocol is to repeat many independent trials of the following.

- Choose a sequence $\mathbf{j} = (j_1, \ldots, j_m)$ of m integers in $\mathbb{N}_{|\mathcal{G}|} = \{1, \ldots, |\mathcal{G}|\}$ uniformly at random.
- Estimate the expectation value $Q_\mathbf{j}$ of an operator Q after preparing the state ρ and applying the sequence $U_\mathbf{j} = U_{j_m} U_{j_{m-1}} \ldots U_{j_1}$ of operators. In the ideal case that $\mathcal{E} = \mathcal{I}$, the expectation value is given by

$$Q_\mathbf{j} = \mathrm{Tr}\Big(QU_\mathbf{j}\ \rho U_\mathbf{j}^\dagger\Big). \tag{7}$$

We will show in section 4, theorem 5 that, under the above assumptions on the noise, the expected value of $Q_{\mathbf{j}}^2$ over all random sequences $\mathbf{j}$ obeys

$$\mathbb{E}_{\mathbf{j}}\left[Q_{\mathbf{j}}^2\right] = A + Bu(\mathcal{E})^{m-1} \tag{8}$$

for trace-preserving noise, where A and B are constants incorporating SPAM and the nonunitality of the noise and $u(\mathcal{E}) \in [0,\ 1]$ is the unitarity of the noise defined in equation (4), with $u(\mathcal{E}) = 1$ if and only if $\mathcal{E}$ is unitary.

Therefore estimating $\mathbb{E}_{\mathbf{j}}[Q_{\mathbf{j}}^2]$ for multiple values of m using the above protocol and fitting to equation (8) gives an efficient and robust estimator of the unitarity.

2.1. Estimators

Note that, as opposed to standard presentations of randomized benchmarking, we are considering the expectation of an operator Q rather than the probability of a single outcome. These descriptions are mathematically equivalent, though expressing the measurement as an observable allows a more concise description of measurement procedures that involve averaging over multiple observables. For example, we can average over the non-identity Pauli operators, while keeping the sequence the same. As will be discussed in section 4, this allows us to simulate a two-state measurement involving S, the SWAP operator. We term this the *purity measurement*, as it estimates the relevant state-dependent term in

$$\mathrm{Tr}\left(\rho_{\mathbf{j}}^2\right) = \frac{1}{d} + \|\mathbf{n}(\rho_{\mathbf{j}})\|^2, \tag{9}$$

the purity of the state $\rho_{\mathbf{j}}$ produced by the sequence $\mathbf{j}$. What we actually use is a shifted and rescaled version of this defined by

$$P_{\mathbf{j}} = \frac{d}{d-1}\|\mathbf{n}(\rho_{\mathbf{j}})\|^2, \tag{10}$$

which for physical states is always in the interval $[0,\ 1]$. For a single qubit, and measuring in the Pauli basis, this quantity is just $P_{\mathbf{j}} = \langle X\rangle^2 + \langle Y\rangle^2 + \langle Z\rangle^2$, where each expectation value is taken with respect to the state $\rho_{\mathbf{j}}$.

The purity measurement can be performed in one of two ways. The direct way involves using two copies of the experiment (with the same sequence) that are run in parallel and a SWAP gate applied immediately prior to measurement. A method using only one copy makes use of the expansion

$$S = \sum_k A_k \otimes A_k^\dagger \tag{11}$$

for any orthonormal operator basis $\{A_k\}$ (e.g., the normalized Paulis) by adding up the expectation values over measurements in the operator basis for the same sequences, that is, by estimating $\sum_k \mathbb{E}_{\mathbf{j}}[(A_k)_{\mathbf{j}}^2]$.

Implementing the purity measurement using this averaging reduces the between-sequence contribution to the uncertainty in our estimates of $\mathbb{E}_{\mathbf{j}}[Q_{\mathbf{j}}^2]$, since if the noise is approximately unitary, then the final state will be relatively pure but will generically overlap with all non-identity Paulis. We note, however, that the above summation over a trace-orthonormal basis is not scalable with the number of qubits, since there are exponentially many n-qubit Paulis. We leave possible optimizations and an analysis of the scalable two-copy protocol as an open problem.

Also note that unlike in standard randomized benchmarking, we do not require the unitary 2-design to be a group since we do not require an inverse operation, or even that the set $\mathcal{G}$ is closed under composition.

2.2. Trace-decreasing noise

More generally, some experimental noise $\mathcal{E}$ may be trace-decreasing with an average survival rate

$$S(\mathcal{E}) = \int \mathrm{d}\psi\ \mathrm{Tr}[\mathcal{E}(\psi)], \tag{12}$$

which is the amount of the trace of the quantum state ψ that survives the error channel $\mathcal{E}$, averaged over the Haar measure $\mathrm{d}\psi$. When $\mathcal{E}$ is itself the average noise over $\mathcal{G}$, the average loss rate can be estimated by

$$\mathbb{E}_{\mathbf{j}}\left[Q_{\mathbf{j}}\right] = CS(\mathcal{E})^{m-1} \tag{13}$$

where C is a constant determined by SPAM [16].

For trace-decreasing noise, the standard decay curve in equation (8) can be generalized to

$$\mathbb{E}_{\mathbf{j}}\left[Q_{\mathbf{j}}^2\right] = A\lambda_+^{m-1} + B\lambda_-^{m-1}, \tag{14}$$

for some constants A and B where

$$\lambda_+ + \lambda_- = S(\mathcal{E})^2 + u(\mathcal{E}). \tag{15}$$

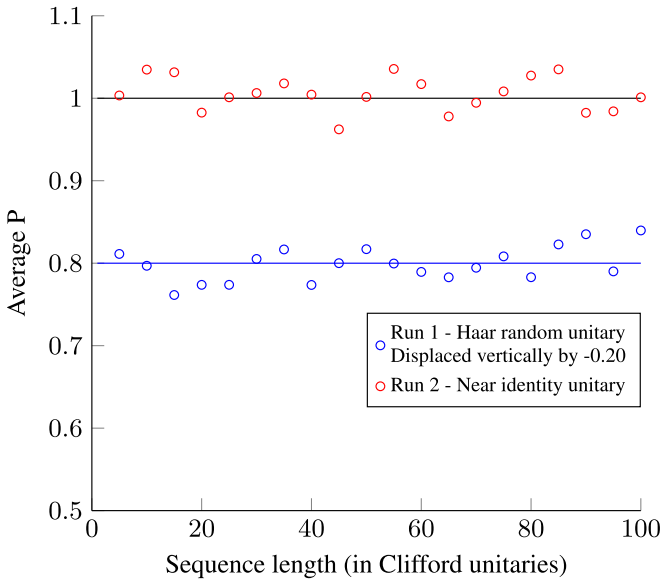

Figure 1. Comparison of our model (equation (8)) with simulated data, where the error channel is taken to be a fixed unitary drawn from the Haar measure (blue), or given by a rotation of 0.1 about σ_x (red). In the simulated data, the purity is measured following the formula of equation (10). The solid lines show the model equation (8) for the calculated values of A, B and $u(\mathcal{E})$ for each scenario. This shows that even for extreme values of unitary noise, our model correctly predicts the unitarity.

The above protocol is a variation of standard randomized benchmarking experiments, and is very similar to the protocol for estimating loss presented in [16]. In particular, one estimates an exponential decay rate in an exactly analogous manner (see equation (8)) and the result is obtained in a manner that is robust to SPAM.

However, there are three small but crucial differences to the experimental protocol presented in [16], leading to significant differences in the analysis and interpretation of the decay curves. Most importantly, the post-processing is different, since in the present paper the survival probabilities for the individual sequences are squared before they are averaged. Secondly, the preparation and measurement procedures in the loss protocol of [16] are ideally the maximally mixed state and the trivial (identity) measurement respectively, which is outperformed in this protocol by the use of the purity measurement. Finally, the current protocol requires a unitary 2-design, whereas the loss protocol only requires a unitary 1-design (although it also works for a unitary 2-design).

3. Numerical simulations

We now illustrate our experimental protocol by numerically simulating it for a variety of single-qubit noise models and showing that the results fit well to (8). In figure 1, we give an example of the correctness of our model equation (8) by showing that it agrees with simulated experimental data in the extreme case that the error channel is a fixed unitary. In figure 2, we demonstrate the utility of equation (8) as an estimation model by estimating $u(\mathcal{E})$ from simulated data drawn according to our protocol. We simulate measurement error on each measurement with small independent random orthogonal matrices, scaling the unital components with a random factor between 0.95 and 1.0. In both these simulations we simulate SPAM on the prepared state by applying a random near-identity unitary. We choose $\mathcal{G}$ to be the single-qubit Clifford group.

Concretely, in figure 1 we show two runs. In the first we set $\mathcal{E}$ to be some fixed (systematic) unitary chosen randomly according to the Haar measure (*Haar-random unitary*) and some near-identity unitary represented by a rotation of 0.1 radians around the X-axis of the Bloch sphere (*near-identity unitary*). As figure 1 demonstrates, unitary noise results in no decay in the number of sequences.

In figure 2, we show different types of unital noise composed with the nonunital amplitude-damping channel

$$\mathcal{E}_{\mathrm{d}}(\rho) = p\,|0\rangle\langle 0| + (1-p)\rho \tag{16}$$

to simulate relaxation to a ground state. The particular unital channels we consider are a Haar-random unitary and a gate-dependent noise channel corresponding to choosing a fixed perturbation of the eigenvalues of a unitary g by $\mathrm{e}^{\mathrm{i}\epsilon}$ to simulate over/under-rotation errors, where the perturbations ϵ are chosen independently and uniformly from $[-0.1,\ 0.1]$ radians for each gate (*rotation channel*).

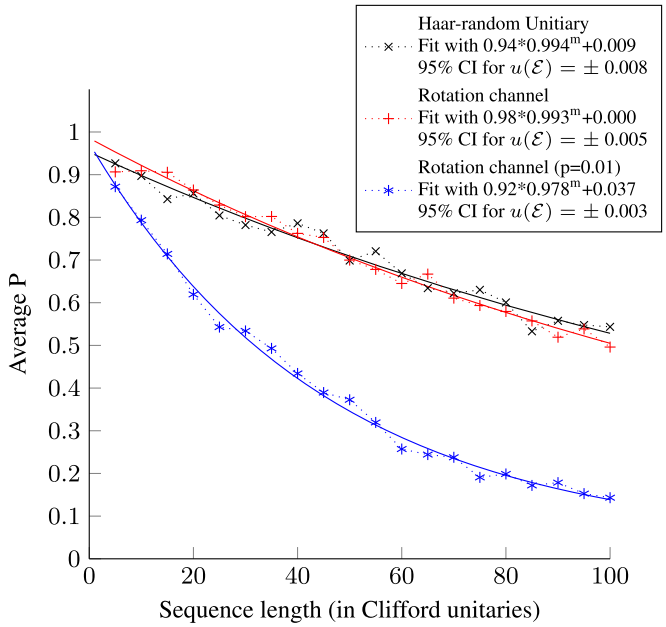

Figure 2. Numerical purity decay curve for our protocol when the noise consists of the non-unital channel in equation (16) with $p = 0.003$ composed with: a Haar-random unitary (black crosses), and fixed gate-dependent over/under-rotations where the eigenvalues for each gate are perturbed by $e^{\pm i\delta}$ for some $\delta \in [-0.01,\ 0.01]$ chosen independently and uniformly (red pluses). The purity is measured in the same manner as in figure 1. The third plot (blue stars) shows (gate-dependent) over/under-rotations composed with the nonunital channel with $p = 0.01$. The solid lines give the fit to equation (8), where the slope gives an estimate for $u(\mathcal{E})$ given by the values 0.994, 0.993 and 0.978, consistent with the theoretical values of 0.994, 0.994 and 0.981, respectively. Confidence intervals for $u(\mathcal{E})$ were calculated assuming Gaussian noise, such that the least-squares fit is a maximum likelihood estimator.

Note that the statistical fluctuations in figures 1 and 2 arise from between-sequence variations and within-sequence variations. The between-sequence variations arise from sampling a small number of random sequences (30 sequences in this case) relative to the total number. The between-sequence variations are minimised by measuring an observable for the purity. A perturbation expansion of the form $\boldsymbol{\mathcal{E}} = \boldsymbol{\mathcal{I}} - r\boldsymbol{\delta}$ (where r is the average gate infidelity of $\boldsymbol{\mathcal{E}}$ to the identity) together with appropriate bounds on the diamond norm can be used to bound these fluctuations and show that they must decrease with gate infidelity, as in [13]. However, a more detailed analysis is complicated by the complexity of the relevant representation theory (that is, four-fold tensor products). The within-sequence variations arise from the need to estimate the expectation values of the observables. For the purpose of figures 1 and 2, we used an unbiased estimator of the squared expectation values, simulating N measurements (with N set to 150).

Finally, we consider the unitarity of random channels drawn from the random ensemble of Bruzda *et al* [18], using the QuTiP software package [19] to draw channels and compute their unitarity (see supplemental material). As shown in figure 3, the distribution of unitarities depends strongly on the Kraus rank of the random channel. Moreover, as demonstrated in figure 4, this information is correlated with, but distinct from, the average gate fidelity.

4. Derivation of the fit models

We now derive the decay curve in equation (14) for trace-non-increasing noise and show how the decay curve in equation (8) emerges as a special case for trace-preserving noise.

Since we are dealing with sequences of channels, it will be convenient to work in the Liouville representation. Since a quantum channel is a linear map between finite-dimensional vector spaces, it is always possible to represent it as a matrix acting on basis coefficients in some given bases for the vector spaces.

In order to construct the Liouville representation of channels, let $\mathcal{A} = \{A_1,\ \ldots,\ A_{d^2}\}$ be an orthonormal basis of $\mathbb{C}^{d\times d}$ according to the Hilbert–Schmidt inner product $\langle A,\ B\rangle = \mathrm{Tr}\ A^\dagger B$. Any density matrix ρ can be expanded as $\rho = \sum_{k\in\mathbb{N}_{d^2}} \langle A_k,\ \rho\rangle A_k$ and so we can identify ρ with a column vector $|\rho) \in \mathbb{C}^{d^2}$ whose kth entry is $\langle A_k,\ \rho\rangle$. The Liouville representation of a channel $\mathcal{E}$ is then the unique matrix $\boldsymbol{\mathcal{E}} \in \mathbb{C}^{d^2\times d^2}$ such that $\boldsymbol{\mathcal{E}}\ |\rho) = |\mathcal{E}[\rho])$, which has entries $\boldsymbol{\mathcal{E}}_{kl} = \langle A_k,\ \mathcal{E}(A_l)\rangle = (A_k\ |\boldsymbol{\mathcal{E}}|\ A_l)$. An immediate consequence of the

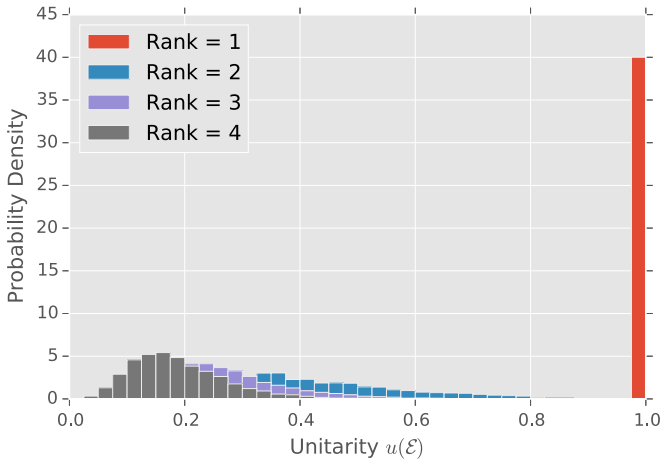

Figure 3. Unitarity of single-qubit CPTP channels chosen according to the random distributions of Bruzda *et al* [18] with varying ranks of the Kraus operators, demonstrating that the unitarity carries information about the structure of the channel. In particular, channels which require more Kraus operators to specify tend towards much smaller unitarity.

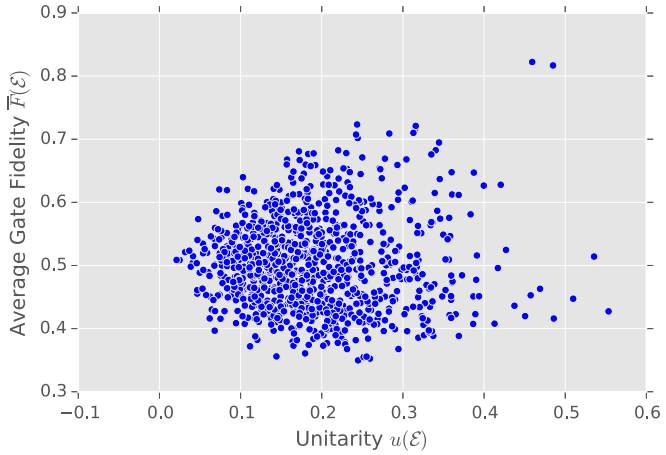

Figure 4. Unitarity of single-qubit CPTP channels chosen according to the random distributions of Bruzda *et al* [18], plotted versus their fidelities to the identity channel. This example shows that even though the two quantities are correlated, they are not redundant and give different insight into the structure of the noise.

uniqueness of $\boldsymbol{\mathcal{E}}$ is that the composition of abstract maps is represented in the Liouville representation by matrix multiplication.

The Liouville representation of unitary channels forms a unitary projective representation of the unitary group $U(d)$. When we wish to emphasize the Liouville representation as a formal representation (rep) of the unitary group $U(d)$ (or subgroups thereof), we will use the notation $\phi_L(U)$ instead of $\boldsymbol{\mathcal{U}}$. With this notation, it is easy to verify that ϕ_L is indeed a unitary representation of $U(d)$, since the Liouville representation of composition is matrix multiplication and it can easily be verified that $\phi_L(U^\dagger) = \phi_L(U)^\dagger$.

Any representation ϕ of a semisimple group $\mathcal{G}$ [such as $SU(d)$] over a vector space V can be unitarily decomposed into a direct sum of irreducible representations (irreps) $\oplus_l \ \phi_l \otimes \mathbb{1}_{n_l}$, where the l label the irreps and the n_l are the corresponding multiplicities and a rep ϕ over a vector space V is irreducible if there are no nontrivial subspaces of V that are invariant under the action of ϕ. A particularly important irrep for this paper is the trivial irrep ϕ_T such that $\phi_T(g) = 1$ for all $g \in \mathcal{G}$.

In the Liouville representation, vectors $b \in \mathbb{C}^{d^2}$ are in one-to-one correspondence with operators $B \in \mathbb{C}^{d\times d}$, so invariant (vector) subspaces under the Liouville representation can be identified with operator subspaces that are invariant under conjugation in the canonical (i.e., $d \times d$ matrix) representation. In particular, the identity operator $\mathbb{1}$ is invariant under conjugation by any unitary, so $|\mathbb{1}\rangle)$ is an invariant subspace of the Liouville representation corresponding to a trivial irrep. We now fix $A_1 = \mathbb{1}/\sqrt{d}$ (so that $\langle A_1, A_1\rangle = 1$), so that the Liouville representation of any unitary $\mathcal{U}$ is

$$\phi_L(U) = 1 \oplus \phi_u(U), \tag{17}$$

where $\oplus$ denotes the matrix direct sum and we refer to $\phi_u(U)$ as the *unital irrep*, which has dimension $d^2 - 1$. Furthermore, any CP channel $\mathcal{E}$ can be written in a corresponding block form as

$$\boldsymbol{\mathcal{E}} = \begin{pmatrix} S(\mathcal{E}) & \boldsymbol{\mathcal{E}}_{\text{sdl}} \\ \boldsymbol{\mathcal{E}}_{\text{n}} & \boldsymbol{\mathcal{E}}_{\text{u}} \end{pmatrix}, \tag{18}$$

where we refer to $\boldsymbol{\mathcal{E}}_{\text{sdl}}, \boldsymbol{\mathcal{E}}_{\text{n}}$ and $\boldsymbol{\mathcal{E}}_{\text{u}}$ as the *state-dependent leakage*, *nonunital* and *unital* blocks respectively. We now show how $\boldsymbol{\mathcal{E}}_{\text{u}}$ is related to the definition of the unitarity in equation (4).

Proposition 1. *The unitarity of a channel $\mathcal{E}$ is*

$$u(\mathcal{E}) = \frac{1}{d^2 - 1} \operatorname{Tr} \boldsymbol{\mathcal{E}}_{\text{u}}^{\dagger} \boldsymbol{\mathcal{E}}_{\text{u}} \tag{19}$$

Proof. For any operator A, $\operatorname{Tr} A^{\dagger} A = (A|A)$ and $\boldsymbol{\mathcal{E}}' = P_{\text{u}} \boldsymbol{\mathcal{E}} P_{\text{u}}$ where P_{u} is the projector onto the unital irrep, so equation (4) can be rewritten as

$$\begin{aligned} u(\mathcal{E}) &= \frac{d}{d-1} \int \mathrm{d}\psi \left(\psi \middle| P_{\text{u}} \boldsymbol{\mathcal{E}}^{\dagger} P_{\text{u}} \boldsymbol{\mathcal{E}} P_{\text{u}} \middle| \psi \right) \\ &= \frac{d}{d-1} \operatorname{Tr} \boldsymbol{\mathcal{E}}_{\text{u}}^{\dagger} \boldsymbol{\mathcal{E}}_{\text{u}} \mathcal{O}, \end{aligned} \tag{20}$$

where $\mathcal{O} = \int \mathrm{d}\psi \, |\psi)(\psi|$, with the slight abuse of notation $\boldsymbol{\mathcal{E}}_{\text{u}} = P_{\text{u}} \boldsymbol{\mathcal{E}} P_{\text{u}}$.

Since $\mathcal{O}$ commutes with the action of the unitary group, Schur's lemma implies that it is a weighted sum of projectors onto the irreps of ϕ_L

$$\mathcal{O} = \lambda_T P_T + \lambda_{\text{u}} P_{\text{u}}. \tag{21}$$

The projector onto the trivial irrep is $P_T = |A_1)(A_1|$ and so

$$\frac{\operatorname{Tr} P_T \mathcal{O}}{\operatorname{Tr} P_T} = \lambda_T = \int \mathrm{d}\psi \left| \left(A_1 \middle| \psi \right) \right|^2 = \frac{1}{d}. \tag{22}$$

Because $\operatorname{Tr}\mathcal{O} = 1$ from the normalization of the Haar measure, we can solve for λ_{u} in the expression

$$\operatorname{Tr} \mathcal{O} = 1 = \frac{1}{d}(1) + \lambda_{\text{u}} \left(d^2 - 1 \right) \tag{23}$$

and we find $\lambda_{\text{u}} = 1/d(d+1)$. Plugging this in and using $P_{\text{u}} P_T = 0$ gives the final result. □

Before we derive the decay curve in equation (14) using the expression for the unitarity from proposition 1, let us first simplify the quantity of interest. The expectation value of Q given that the sequence $\mathbf{j}$ was applied is

$$Q_{\mathbf{j}} = \left(Q \middle| \boldsymbol{\mathcal{U}}_{j_m} \boldsymbol{\mathcal{E}} \ldots \boldsymbol{\mathcal{U}}_{j_2} \boldsymbol{\mathcal{E}} \boldsymbol{\mathcal{U}}_{j_1} \middle| \rho \right), \tag{24}$$

where a residual noise term has been absorbed into the experimental state preparation ρ. Noting that

$$Q_{\mathbf{j}}^2 = \left(Q^{\otimes 2} \middle| \boldsymbol{\mathcal{U}}_{j_m}^{\otimes 2} \boldsymbol{\mathcal{E}}^{\otimes 2} \ldots \boldsymbol{\mathcal{U}}_{j_2}^{\otimes 2} \boldsymbol{\mathcal{E}}^{\otimes 2} \boldsymbol{\mathcal{U}}_{j_1}^{\otimes 2} \middle| \rho^{\otimes 2} \right), \tag{25}$$

the expected average of the squares is

$$\begin{aligned} \mathbb{E}_{\mathbf{j}} \left[Q_{\mathbf{j}}^2 \right] &= |\mathcal{G}|^{-m} \sum_{\mathbf{j}} Q_{\mathbf{j}}^2 \\ &= \left(Q^{\otimes 2} \middle| \left(\boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} \boldsymbol{\mathcal{E}}^{\otimes 2} \boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} \right)^{m-1} \middle| \rho^{\otimes 2} \right) \\ &= \left(Q^{\otimes 2} \middle| \boldsymbol{\mathcal{M}}^{m-1} \middle| \rho^{\otimes 2} \right), \end{aligned} \tag{26}$$

where $\boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} = |\mathcal{G}|^{-1} \sum_{g \in \mathcal{G}} \boldsymbol{g}^{\otimes 2}$, we define the averaged operator $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} \boldsymbol{\mathcal{E}}^{\otimes 2} \boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2}$, and we have used the fact that $|\mathcal{G}|^{-1} \sum_{g \in \mathcal{G}} \phi(g)$ is the projector onto the trivial subreps for any rep ϕ of a group $\mathcal{G}$ so that $\boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} = (\boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2})^2$ [20]. Thus, to derive the fit model we must first identify the trivial irreps of $\mathcal{G}$ in $\phi_L(U)^{\otimes 2}$, since this is where $\boldsymbol{\mathcal{M}}$ is supported.

Proposition 2. *The averaged operator $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2} \boldsymbol{\mathcal{E}}^{\otimes 2} \boldsymbol{\mathcal{U}}_{\text{avg}}^{\otimes 2}$ is supported on a two-dimensional subspace spanned by $|\mathbb{1}_{d^2})$ and $|S)$, where S is the SWAP operator.*

Proof. Define $\chi_R(g) = \mathrm{Tr}\, R(g)$ as the character of the rep R. Then we can use Schur's orthogonality relations to count the number of trivial irreps. Let $\langle\!\langle \chi_R, \chi_{R'} \rangle\!\rangle = |\mathcal{G}|^{-1} \sum_{g\in\mathcal{G}} \chi_R^*(g)\chi_{R'}(g)$ denote the character inner product for $\mathcal{G}$. From the direct sum structure in equation (17), the number of trivial irreps is

$$\left\langle\!\!\left\langle \chi_1, \chi_{L\otimes L} \right\rangle\!\!\right\rangle = \left\langle\!\!\left\langle \chi_1, \chi_1 + 2\chi_{\mathrm{u}} + \chi_{\mathrm{u}}^2 \right\rangle\!\!\right\rangle = 1 + 0 + \left\langle\!\!\left\langle \chi_1, \chi_{\mathrm{u}}^2 \right\rangle\!\!\right\rangle.$$

Since χ_{u} is real-valued, we have $\langle\!\langle \chi_1, \chi_{\mathrm{u}}^2 \rangle\!\rangle = \langle\!\langle \chi_{\mathrm{u}}, \chi_{\mathrm{u}} \rangle\!\rangle$. If $\mathcal{G}$ acts irreducibly on the unital block [21], then $\langle\!\langle \chi_{\mathrm{u}}, \chi_{\mathrm{u}} \rangle\!\rangle = 1$ and the number of trivial irreps is 2.

The two trivial irreps in $\phi_L^{\otimes 2}$ are spanned by the orthonormal vectors $|B_1)$ and $|B_2)$ where

$$B_1 = \mathbb{1}_{d^2}/d, \quad B_2 = \left(S - \mathbb{1}_{d^2}/d\right)/\sqrt{d^2-1}, \tag{27}$$

and S is the SWAP operator. To check this, note that

$$\left(B_j \middle| \phi_L(U)^{\otimes 2} \middle| B_k\right) = \mathrm{Tr}\left(B_j^\dagger U \otimes U B_k U^\dagger \otimes U^\dagger\right) = \delta_{jk},$$

since both identity and SWAP are invariant under conjugation by $U \otimes U$ and $S^2 = \mathbb{1}_{d^2}$. Since $\phi_L(U)^{\otimes 2}$ is a unitary rep, B_1 and B_2 are the first two elements of a two-qudit orthonormal Schur basis $\{B_j\}$ for $\phi_L^{\otimes 2}$ and so correspond to trivial irreps. Therefore $\boldsymbol{\mathcal{M}}$ is zero except for the 2×2 submatrix supported on $|B_1)$ and $|B_2)$. These vectors have the same span as $|\mathbb{1}_{d^2})$ and $|S)$. □

The next proposition characterizes the averaged operator on the supported subspace.

Proposition 3. *In the invariant* $\mathrm{B_i}$ *basis from equation (27), the averaged operator* $\boldsymbol{\mathcal{M}}$ *has the following matrix elements*

- $\boldsymbol{\mathcal{M}}_{11} = S(\mathcal{E})^2$,
- $\boldsymbol{\mathcal{M}}_{12} = (d^2-1)^{-1/2}\|\boldsymbol{\mathcal{E}}_{\mathrm{sdl}}\|^2$,
- $\boldsymbol{\mathcal{M}}_{21} = (d^2-1)^{-1/2}\|\boldsymbol{\mathcal{E}}_{\mathrm{n}}\|^2$, *and*
- $\boldsymbol{\mathcal{M}}_{22} = (d^2-1)^{-1}\|\boldsymbol{\mathcal{E}}_{\mathrm{u}}\|_F^2 = u(\mathcal{E})$.

Proof. We will establish the matrix elements with respect to $|\mathbb{1}_{d^2})$ and $|S)$, and the claims about the B_i basis will follow by taking appropriate linear combinations.

Because the B_i basis is invariant, we can ignore the average unitary terms in $\boldsymbol{\mathcal{M}}$. We first find that

$$\begin{aligned}\left(\mathbb{1}_{d^2} \middle| \boldsymbol{\mathcal{E}}^{\otimes 2} \middle| \mathbb{1}_{d^2}\right) &= \mathrm{Tr}\, \mathcal{E}^{\otimes 2}\left(\mathbb{1}_{d^2}\right) = \mathrm{Tr}\left(\mathcal{E}\left(\mathbb{1}_d\right)^{\otimes 2}\right) \\ &= \left(\mathrm{Tr}\, \mathcal{E}\left(\mathbb{1}_d\right)\right)^2 = d^2 S(\mathcal{E})^2. \end{aligned} \tag{28}$$

Next we can use the identity $\langle S, A \otimes B\rangle = \mathrm{Tr}\,[S(A\otimes B)] = \mathrm{Tr}(AB)$ and the fact that $\mathcal{E}(A^\dagger) = \mathcal{E}(A)^\dagger$ to find

$$\begin{aligned}\left(S \middle| \boldsymbol{\mathcal{E}}^{\otimes 2} \middle| \mathbb{1}_{d^2}\right) &= \mathrm{Tr}\left[S\mathcal{E}\left(\mathbb{1}_d\right)^{\otimes 2}\right] = \mathrm{Tr}\left[\mathcal{E}\left(\mathbb{1}_d\right)\mathcal{E}\left(\mathbb{1}_d\right)\right] \\ &= \left(\mathcal{E}\left(\tfrac{1}{d}\right)|\mathcal{E}\left(\tfrac{1}{d}\right)\right) = \|\boldsymbol{\mathcal{E}}_n\|^2. \end{aligned} \tag{29}$$

The expression for $(\mathbb{1}_{d^2}|\boldsymbol{\mathcal{E}}^{\otimes 2}|S)$ follows similarly using the adjoint channel. Finally, we can use the expansion $S = \sum_k A_k \otimes A_k^\dagger$ for any orthonormal operator basis A_k to obtain

$$\begin{aligned}\left(S \middle| \boldsymbol{\mathcal{E}}^{\otimes 2} \middle| S\right) &= \mathrm{Tr}\left[S\boldsymbol{\mathcal{E}}^{\otimes 2}(S)\right] = \sum_k \mathrm{Tr}\left[S\mathcal{E}\left(A_k\right) \otimes \mathcal{E}\left(A_k^\dagger\right)\right] \\ &= \sum_k \mathrm{Tr}\left[\mathcal{E}\left(A_k\right)^\dagger \mathcal{E}\left(A_k\right)\right] = \|\boldsymbol{\mathcal{E}}\|_F^2, \end{aligned} \tag{30}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The values of the matrix elements are then established by using the form of equation (18) and the definition of the B_i basis from equation (27) and taking various linear combinations. □

The final step in deriving the fit model is to analyze the eigenvalues and eigenvectors of the averaged operator.

Proposition 4. *The averaged operator $\mathcal{M}$ has two distinct nontrivial eigenvectors.*

Proof. Since the averaged operator vanishes almost everywhere, we only need to consider the 2×2 submatrix derived above. The nontrivial eigenvalues are

$$\lambda_{\pm} = \frac{1}{2}\mathcal{M}_{11} + \frac{1}{2}\mathcal{M}_{22} \pm \frac{1}{2}\sqrt{\left[\mathcal{M}_{11} - \mathcal{M}_{22}\right]^2 + 4\mathcal{M}_{12}\mathcal{M}_{21}}. \tag{31}$$

This spectrum is degenerate precisely when the terms under the square root both vanish (since both terms are nonnegative). Whenever the spectrum is nondegenerate, there are trivially two distinct eigenvectors, so we only need to deal with the degenerate case.

We will break the analysis for the degenerate spectrum into two nontrivial cases, $\mathcal{M}_{11} = \mathcal{M}_{22}$ and either $\mathcal{M}_{12} = 0$ or $\mathcal{M}_{21} = 0$, exclusive. There are also two trivial cases: when $\mathcal{M}_{12} = \mathcal{M}_{21} = 0$, the matrix $\mathcal{M}$ is already diagonal and we are done. We ignore the pathological case when $\mathcal{M}_{11} = 0$, since this corresponds physically to a state that is never observable. In both nontrivial cases, we will make use of the two-qudit state $\Pi_a = \dfrac{\mathbb{1} - S}{d(d-1)}$, the maximally mixed state on the antisymmetric subspace. Expanding this state in the B_i basis gives

$$\left|\Pi_a\right) = \pi_1 \left|B_1\right) + \pi_2 \left|B_2\right) = \frac{1}{d}\left|B_1\right) - \frac{\sqrt{d^2-1}}{d(d-1)}\left|B_2\right).$$

The key feature of this state is that $\pi_2 < 0$.

Case 1: $\mathcal{M}_{12} = 0$. In this case

$$\mathcal{M} = \begin{pmatrix} \lambda & 0 \\ y & \lambda \end{pmatrix}, \tag{32}$$

where $\lambda > 0$ and $y \geqslant 0$. Taking the mth power gives

$$\mathcal{M}^m = \lambda^{m-1}\begin{pmatrix} \lambda & 0 \\ my & \lambda \end{pmatrix}. \tag{33}$$

If we perform the measurement $\{\Pi_a, \mathbb{1} - \Pi_a\}$ on a system prepared in the state $\mathbb{1}_{d^2}/d^2$ which evolves under $\mathcal{M}^m$, then the probability of observing the outcome Π_a is

$$\left(\Pi_a\middle|\mathcal{M}^m\middle|\mathbb{1}_{d^2}/d^2\right) = \frac{\lambda^{m-1}}{d}\left(\lambda\pi_1 + my\pi_2\right). \tag{34}$$

Since $\lambda,\ \pi_1 > 0, y \geqslant 0$, and $\pi_2 < 0$, in order for this to be a probability for all m, we require $y = 0$ and so $\mathcal{M}$ is actually diagonal.

Case 2: $\mathcal{M}_{21} = 0$. In this case,

$$\mathcal{M} = \begin{pmatrix} \lambda & y \\ 0 & \lambda \end{pmatrix}, \tag{35}$$

where $\lambda > 0$ and $y \geqslant 0$. Taking the mth power gives

$$\mathcal{M}^m = \lambda^{m-1}\begin{pmatrix} \lambda & my \\ 0 & \lambda \end{pmatrix}. \tag{36}$$

Therefore the probability of detecting the system (i.e., measuring $\mathbb{1}_d^2$) when a system is prepared in the state Π_a and evolves under $\mathcal{M}^m$ is

$$\left(\mathbb{1}_{d^2}\middle|\mathcal{M}^m\middle|\Pi_a\right) = \lambda^{m-1}\left(\lambda\pi_1 + my\pi_2\right). \tag{37}$$

Again since $\lambda,\ \pi_1 > 0, y \geqslant 0$, and $\pi_2 < 0$, for this to be a valid probability for all m, we require $y = 0$ and so $\mathcal{M}$ is actually diagonal. □

We now have all the ingredients to derive the fit models of equations (8) and (14).

Theorem 5. *For time- and gate-independent noise, the expected value $\mathbb{E}_{\mathrm{j}}[Q_{\mathrm{j}}^2]$ obeys the decay equation*

$$\mathbb{E}_{\mathrm{j}}\left[Q_{\mathrm{j}}^2\right] = A + Bu(\mathcal{E})^{m-1}$$

for trace-preserving noise, and for trace-decreasing noise it obeys

$$\mathbb{E}_{\mathrm{j}}\left[Q_{\mathrm{j}}^2\right] = A\lambda_+^{m-1} + B\lambda_-^{m-1},$$

where $\lambda_\pm$ are given by equation (31), $\lambda_+ + \lambda_- = S(\mathcal{E})^2 + u(\mathcal{E})$, and the constants A and B depend only on state preparation and measurement errors and the unitary that diagonalizes $\mathcal{M}$.

Proof. Proposition 4 establishes that the matrix $\mathcal{M}$ is diagonalizable by a similarity transform with eigenvalues given by equation (31). From equation (26), we can diagonalize $\mathcal{M}$ and absorb the similarity transform into $|\rho^{\otimes 2})$ and $(Q^{\otimes 2}|$ as SPAM, yielding

$$\begin{aligned}\mathbb{E}_{\mathrm{j}}\left[Q_{\mathrm{j}}^2\right] &= \left(Q^{\otimes 2}\middle|\mathcal{M}^{m-1}\middle|\rho^{\otimes 2}\right)\\ &= A\lambda_+^{m-1} + B\lambda_-^{m-1}.\end{aligned} \tag{38}$$

Trace-preserving noise is a special case of this, since if $\mathcal{E}$ is TP, then by proposition 3 we have $\lambda_+ = S(\mathcal{E})^2 = 1$ and so $\lambda_- = u(\mathcal{E})$. □

We note that the unitary that diagonalizes $\mathcal{M}$ will in general depend on the noise channel, and hence will depend on u. We conflate this dependence with the SPAM errors in our fit model, as the diagonalization of $\mathcal{M}$ does not depend on the sequence length m. Neglecting the dependence on u thus results in a model that is correct, but is slightly less sensitive to u than is optimal.

We are now equipped to formalize the observation made in section 2 that the optimal observable for trace-preserving noise is an operator proportional to B_2. This follows from noting that such operators overlap fully with the component of $\mathcal{M}$ that give rise to the exponential term, as given by equation (30).

5. Properties of the unitarity

We now prove some properties of the unitarity for CPTP channels that make it a practical quantification of the coherence of a channel. We begin by proving that the unitarity and the average incoherent survival probability can be used to bound the nonunital and state-dependent leakage terms which are subtracted off in the definition of unitarity in equation (4).

Proposition 6. *For any channel $\mathcal{E}$,*

$$\max\left\{\|\mathcal{E}_{\mathrm{n}}\|^2, \|\mathcal{E}_{\mathrm{sdl}}\|^2\right\} \leqslant \tfrac{1}{2}\left(d^2-1\right)\left[S(\mathcal{E})^2 - u(\mathcal{E})\right]. \tag{39}$$

If $\mathcal{E}$ is trace-preserving, then $\|\mathcal{E}_{\mathrm{n}}\|^2 \leqslant (d-1)[1-u(\mathcal{E})]$.

Proof. Consider the maximally mixed states on the symmetric and antisymmetric subspaces, $\Pi_s = \frac{\mathbb{1}+S}{d(d+1)}$ and $\Pi_a = \frac{\mathbb{1}-S}{d(d-1)}$ respectively, and let E_s and E_a be the respective projectors onto these spaces. Expanding these states in the B_i basis gives

$$\begin{aligned}\left|\Pi_s\right) &= \frac{1}{d}\left|B_1\right) + \frac{\sqrt{d^2-1}}{d(d+1)}\left|B_2\right)\\ \left|\Pi_a\right) &= \frac{1}{d}\left|B_1\right) - \frac{\sqrt{d^2-1}}{d(d-1)}\left|B_2\right).\end{aligned} \tag{40}$$

Preparing the state Π_s (Π_a), evolving under $\mathcal{M}$ and then measuring the POVM $\{E_a, \mathbb{1}-E_a\}$ ($\{E_s, \mathbb{1}-E_s\}$) produces the outcomes E_a (E_s) with probabilities

$$\begin{aligned}p_{as} &= \left(E_a\middle|\mathcal{M}\middle|\Pi_s\right)\\ &= \frac{d-1}{2d}\left(S(\mathcal{E})^2 - u(\mathcal{E}) - \frac{\|\mathcal{E}_{\mathrm{n}}\|^2}{d-1} + \frac{\|\mathcal{E}_{\mathrm{sdl}}\|^2}{d+1}\right)\end{aligned} \tag{41}$$

$$\begin{aligned}p_{sa} &= \left(E_s\middle|\mathcal{M}\middle|\Pi_a\right)\\ &= \frac{d+1}{2d}\left(S(\mathcal{E})^2 - u(\mathcal{E}) + \frac{\|\mathcal{E}_{\mathrm{n}}\|^2}{d+1} - \frac{\|\mathcal{E}_{\mathrm{sdl}}\|^2}{d-1}\right)\end{aligned} \tag{42}$$

respectively, where we have used proposition 3. Since both these expressions are probabilities we have $0 \leqslant p_{as} \leqslant 1$ and $0 \leqslant p_{sa} \leqslant 1$. Taking appropriate linear combinations of these two inequalities will cancel the

dependence on either $\|\mathcal{E}_{\mathrm{n}}\|^2$ or $\|\mathcal{E}_{\mathrm{sdl}}\|^2$, isolating the other variable. Simplifying the resulting expressions gives the bound equation (39) for both quantities individually, hence the maximum holds as well.

Furthermore, if the noise is trace-preserving, then $\|\mathcal{E}_{\mathrm{sdl}}\|^2 = 0$ and $S(\mathcal{E}) = 1$, so $p_{as} \geqslant 0$ gives $\|\mathcal{E}_{\mathrm{n}}\|^2 \leqslant (d-1)[1-u(\mathcal{E})]$ for trace-preserving noise. □

We now prove that $u(\mathcal{E}) = 1$ if and only if $\mathcal{E}$ is unitary and that $u(\mathcal{E})$ is invariant under composition with unitaries.

Proposition 7. *For any channel $\mathcal{E}$, $u(\mathcal{E}) \leqslant 1$ with equality if and only if $\mathcal{E}$ is unitary. Furthermore, the unitarity satisfies $u(\mathcal{V} \circ \mathcal{E} \circ \mathcal{U}) = u(\mathcal{E})$ for any unitaries $U, V \in U(d)$.*

Proof. The unitary invariance $u(\mathcal{V} \circ \mathcal{E}\circ\mathcal{U}) = u(\mathcal{E})$ follows immediately from the invariance of the trace under cyclic permutations.

Since the norms of vectors are always nonnegative, $u(\mathcal{E}) = 1$ only if $\mathcal{E}$ is trace-preserving and unital by equation (39), in which case the adjoint channel $\mathcal{E}^\dagger$ is also a channel [22] and so the eigenvalues of $\mathcal{E}^\dagger\mathcal{E}$ (i.e., the singular values of $\mathcal{E}$) are all bounded by one [23]. Therefore $u(\mathcal{E}) = 1$ only if $\mathcal{E}$ is unital and all the eigenvalues of $\mathcal{E}$ have unit modulus and consequently if $|\det \mathcal{E}| = 1$. However, the only channels with $|\det \mathcal{E}| = 1$ are unitary channels [24]. Since $u(\mathcal{E})$ is unitarily invariant and $u(\mathcal{I}) = 1$, $u(\mathcal{E}) = 1$ if and only if $\mathcal{E}$ is unitary, as claimed. □

We now show that the unitarity can be used with the average gate infidelity to quantify the intermediate regime between incoherent and unitary errors. It is useful to define a notion of average gate infidelity that has been optimized to remove unitary noise. First recall the definition of average gate infidelity,

$$r(\mathcal{E}) = 1 - \int \mathrm{d}\psi \ \mathrm{Tr}[\psi \mathcal{E}(\psi)], \tag{43}$$

Then for any CPTP channel $\mathcal{E}$, define

$$R(\mathcal{E}) = \min_{U \in U(d)} r(\mathcal{E}\circ\mathcal{U}). \tag{44}$$

where the equivalance follows from the Haar invariance in equation (43). This quantity can be thought of as the best average gate infidelity that is achievable with perfect unitary control. For example, if $\mathcal{E}$ is a unitary channel, then $R(\mathcal{E}) = 0$.

Proposition 8. *For any CPTP channel $\mathcal{E}$ with average gate infidelity $r = r(\mathcal{E})$ to the identity and $R = R(\mathcal{E})$ as above, then the following inequalities hold*

$$u(\mathcal{E}) \geqslant [1 - dR/(d-1)]^2 \geqslant [1 - dr/(d-1)]^2. \tag{45}$$

The chain of inequalities is saturated if and only if $\mathcal{E}$ has a unital block $\mathcal{E}_{\mathrm{u}}$ that is a diagonal scalar matrix.

Proof. Any channel with infidelity r to the identity can be written as $\mathcal{E} = \mathcal{I} - r\Delta$ where the diagonal entries of Δ are nonnegative and $\mathrm{Tr}\ \Delta = d(d+1)$. We then have

$$\|\mathcal{E}_{\mathrm{u}}\|_2^2 \geqslant \sum_{k=2}^{d^2}\left(1 - \Delta_{kk} r\right)^2 = d^2 - 1 - 2r\ \mathrm{Tr}\ \Delta + \sum_{k=2}^{d^2}\Delta_{kk}^2 r^2,$$

with equality if and only if $\mathcal{E}_{\mathrm{u}}$ is diagonal. The term $\sum_k \Delta_{kk}^2$ is uniquely minimized for nonnegative Δ_{kk} subject to the constraint $\mathrm{Tr}\ \Delta = d(d+1)$ by setting $\Delta_{kk} = d/(d-1)$ (that is, by setting all the diagonal entries to be equal). This proves the weaker inequality bounding $u(\mathcal{E})$ in terms of r. To get the stronger inequality in terms of $R(\mathcal{E})$, we use the unitary invariance proven in proposition 7 and optimize the inequality over all unitary channels. □

We note that the first inequality in equation (45) is saturated at 1 when the noise channel is unitary, and the chain of inequalities is saturated for depolarizing noise, or depolarizing noise composed with amplitude damping.

An immediate corollary of proposition 8 is that the unitarity can be used to put a lower bound on the best possible average infidelity in the presence of perfect unitary control. Rearranging equation (45), we find a lower bound

$$\frac{d-1}{d}\left(1 - \sqrt{u(\mathcal{E})}\right) \leqslant R(\mathcal{E}) \leqslant r(\mathcal{E}). \tag{46}$$

The unitarity is also closely related to the purity of the Jamiłłkowski state associated to the noise channel.

Proposition 9. *The unitarity is related to the purity of the Jamiołkowski state by*

$$d^2 \operatorname{Tr}\left[J(\mathcal{E})^\dagger J(\mathcal{E})\right] = S(\mathcal{E})^2 + \|\boldsymbol{\mathcal{E}}_{\mathrm{sdl}}\|^2 + \|\boldsymbol{\mathcal{E}}_{\mathrm{n}}\|^2 + \left(d^2-1\right)u(\mathcal{E}), \tag{47}$$

where $J(\mathcal{E}) = (\mathcal{E}\otimes\mathcal{I})[\Phi]$ *and* $\Phi = \frac{1}{d}\sum_{j,k}|jj\rangle\langle kk|$.

Proof. We begin with an alternate representation of the maximally entangled state $\Phi = d^{-1}\sum_k A_k\otimes A_k$. By cycling the adjoint channel in the trace, the purity of $J(\mathcal{E})$ becomes

$$\begin{aligned}\operatorname{Tr}\left[J(\mathcal{E})^\dagger J(\mathcal{E})\right] &= \operatorname{Tr}\left[\Phi\left(\mathcal{E}^\dagger\mathcal{E}\otimes\mathcal{I}\right)[\Phi]\right]\\ &= \left(\Phi\middle|\boldsymbol{\mathcal{E}}^\dagger\boldsymbol{\mathcal{E}}\otimes\boldsymbol{\mathcal{I}}\middle|\Phi\right)\\ &= \frac{1}{d^2}\sum_{j,k}\left(A_j\otimes A_j\middle|\boldsymbol{\mathcal{E}}^\dagger\boldsymbol{\mathcal{E}}\otimes\boldsymbol{\mathcal{I}}\middle|A_k\otimes A_k\right)\\ &= \frac{1}{d^2}\sum_{j,k}\left(A_j\middle|\boldsymbol{\mathcal{E}}^\dagger\boldsymbol{\mathcal{E}}\middle|A_k\right)\left(A_j\middle|A_k\right).\end{aligned} \tag{48}$$

Since the A_k are a trace orthonormal basis, the last line simplifies to

$$\operatorname{Tr}\left[J(\mathcal{E})^\dagger J(\mathcal{E})\right] = \tfrac{1}{d^2}\operatorname{Tr}\boldsymbol{\mathcal{E}}^\dagger\boldsymbol{\mathcal{E}}. \tag{49}$$

Comparing this expression to the decomposition in equation (18) and using proposition 1 completes the proof. □

Finally, we give a simple example that shows that the unitarity is *not* a monotone, in the sense that it can oscillate under composition of channels. Consider the two (nearly) dual qubit channels,

$$\mathcal{E}_0(\rho) = \operatorname{Tr}(\rho)|0\rangle\langle 0| \quad\text{and}\quad \frac{1}{2}\mathcal{E}_0^\dagger(\rho) = \langle 0|\,\rho\,|0\rangle\frac{\mathbb{1}}{2}.$$

Then the unitarity of both $\mathcal{E}_0$ and $\frac{1}{2}\mathcal{E}_0^\dagger$ is zero, while the unitarity of the composed channel $\frac{1}{2}\boldsymbol{\mathcal{E}}_0\boldsymbol{\mathcal{E}}_0^\dagger$ is 1/12.

We note that for some restricted classes of channels the unitarity is indeed a monotone. For example, a trivial application of von Neumann's trace inequality shows that if the singular values of the unital block are all less than or equal to 1 (which holds for all qubit channels and all unital channels), then it is a monotone for trace-preserving channels.

6. Conclusion

In this paper, we have shown that the coherence of a noisy process can be quantified by the unitarity, which corresponds to the change in the purity (with the identity components subtracted off) averaged over pure states. We have presented a protocol for efficiently estimating the unitarity of the average noise in the implementation of a unitary 2-design.

We have also proven that the unitarity is 1 if and only if the noise source is unitary and provided a tight lower bound for the unitarity in terms of the infidelity (which can be estimated using randomized benchmarking [8]). This allows the intermediate regime between fully incoherent and unitary errors to be quantified, potentially allowing for improved bounds on the worst-case error. We have also shown that the unitarity provides a lower bound on the best achievable gate infidelity assuming perfect unitary control.

Our present results also have direct implications for the loss protocol when applied to a unitary 2-design, since the variance over random sequences of fixed length for the protocol in [16] is

$$\mathbb{V}_{\mathrm{j}}\left(Q_{\mathrm{j}}\right) = \mathbb{E}_{\mathrm{j}}\left(Q_{\mathrm{j}}^2\right) - \left[\mathbb{E}_{\mathrm{j}}\left(Q_{\mathrm{j}}\right)\right]^2, \tag{50}$$

which decays faster with m for fixed $S(\mathcal{E})$ if the unitarity is smaller (and hence the two decay rates in the fit curve for determining the unitarity, $\lambda_\pm$, are smaller). A lower variance over sequences allows a more precise estimation of the average incoherent survival probability for a fixed number of experiments. Similar implications may also hold for standard randomized benchmarking since $u(\mathcal{E})$ can easily be seen to be one of the eigenvalues of the averaged operator in [13] that determines the variance and is precisely the eigenvalue that determines the asymptotic variance. However, in order to establish a concrete bound, it would have to be shown that $u(\mathcal{E})$ is in fact the largest eigenvalue.

There are four important open problems raised by this work. First, while the unitarity is a monotone for unital noise, it is not a monotone for trace-decreasing noise. We leave open the problem of finding necessary and sufficient conditions for when $u(\mathcal{E})$ is a monotone, or finding other quantities that are monotonic in general.

Second, our protocol characterizes the unitarity of the average noise, but does not characterize the unitarity of the errors in the individual gate. While a variant of interleaved randomized benchmarking [25] should hold for the current protocol, obtaining reasonable bounds on the unitarity of the individual error is an open problem.

Third, the signal for our protocol is substantially improved by the purity measurement, but the method of performing the purity measurement via measuring Pauli operators is not scalable beyond a handful of qubits because of the exponential size of the Pauli group on n qubits. Moreover, measuring any single Pauli operator will in general give a small signal as the number of qubits grows, since we do not perform an inversion step. Directly using the SWAP operation on two copies of the system running in parallel is a mathematical solution, but the extra resources required to implement this might be prohibitive and an analysis of the role of crosstalk and correlations would be required to justify this idea. Thus, identifying efficient measurements that give a good signal on multi-qubit systems remains an open problem.

Finally, a pressing open problem identified in this paper is to obtain an improved bound on the worst-case error in terms of both the infidelity and the unitarity. Such a bound would substantially reduce the effort required to certify that an experimental implementation is near (or below) the threshold for fault-tolerant quantum computation.

Acknowledgments

We thank Jay Gambetta for pointing out that estimating purity would reduce variance. This work was supported by the ARC via EQuS project number CE11001013, by IARPA via the MQCO program, and by the US Army Research Office grant numbers W911NF-14-1-0098 and W911NF-14-1-0103. STF also acknowledges support from an ARC Future Fellowship FT130101744.

References

[1] Chuang I L and Nielsen M A 1997 Prescription for experimental determination of the dynamics of a quantum black box *J. Mod. Opt.* **44** 2455
[2] Poyatos J, Cirac J and Zoller P 1997 Complete characterization of a quantum process: the two-bit quantum gate *Phys. Rev. Lett.* **78** 390
[3] Merkel S T, Gambetta J M, Smolin J A, Poletto S, Córcoles A D, Johnson B R, Ryan C A and Steffen M 2013 Self-consistent quantum process tomography *Phys. Rev.* A **87** 062119
[4] Emerson J, Alicki R and Życzkowski K 2005 Scalable noise estimation with random unitary operators *J. Opt. B: Quantum Semiclass. Opt.* **7** S347
[5] Knill E, Leibfried D, Reichle R, Britton J, Blakestad R, Jost J D, Langer C, Ozeri R, Seidelin S and Wineland D J 2008 Randomized benchmarking of quantum gates *Phys. Rev.* A **77** 012307
[6] Dankert C, Cleve R, Emerson J and Livine E 2009 Exact and approximate unitary 2-designs and their application to fidelity estimation *Phys. Rev.* A **80** 012304
[7] López C C, Bendersky A, Paz J P and Cory D G 2010 Progress toward scalable tomography of quantum maps using twirling-based methods and information hierarchies *Phys. Rev.* A **81** 062113
[8] Magesan E, Gambetta J M and Emerson J 2011 Scalable and robust randomized benchmarking of quantum processes *Phys. Rev. Lett.* **106** 180504
[9] Magesan E, Gambetta J M and Emerson J 2012 Characterizing quantum gates via randomized benchmarking *Phys. Rev.* A **85** 042311
[10] Flammia S T and Yi-Kai L 2011 Direct fidelity estimation from few pauli measurements *Phys. Rev. Lett.* **106** 230501
[11] da Silva M P, Landon-Cardinal O and Poulin D 2011 Practical characterization of quantum devices without tomography *Phys. Rev. Lett.* **107** 210404
[12] Kitaev A 1997 Quantum computations: algorithms and error correction *Russ. Math. Surv.* **52** 1191–249
[13] Wallman J J and Flammia S T 2014 Randomized benchmarking with confidence *New J. Phys.* **16** 103032
[14] Sanders Y R, Wallman J J and Sanders B C 2015 Bounding quantum gate error rate based on reported gate fidelity (arXiv:1501.04932v2)
[15] O'Malley P J J *et al* 2015 Qubit metrology of ultralow phase noise using randomized benchmarking *Phys. Rev. Appl.* **3** 044009
[16] Wallman J J, Barnhill M and Emerson J 2015 Robust characterization of loss rates *Phys. Rev. Lett.* **115** 060501
[17] Baumgratz T, Cramer M and Plenio M B 2014 Quantifying coherence *Phys. Rev. Lett.* **113** 140401
[18] Bruzda W, Cappellini V, Hans-Jürgen S and Życzkowski K 2009 Random quantum operations *Phys. Lett.* A **373** 320–4
[19] Johansson J R, Nation P D and Nori F 2013 QuTiP 2: a Python framework for the dynamics of open quantum systems *Comput. Phys. Commun.* **184** 1234–40
[20] Goodman R and Wallach N R 2009 *Symmetry, Representations, and Invariants* (*Graduate Texts in Mathematics*) (Berlin: Springer)
[21] Gross D, Audenaert K and Eisert J 2007 Evenly distributed unitaries: On the structure of unitary designs *J. Math. Phys.* **48** 052104
[22] Watrous J 2011 *CS 766/QIC 820 Theory of Quantum Information* (*Fall 2011*)
[23] Evans D E and Hoegh-Krohn R 1978 Spectral properties of positive maps on C^*-algebras *J. London Math. Soc.* **17** 345–55
[24] Wolf M M and Cirac J I 2008 Dividing quantum channels *Commun. Math. Phys.* **279** 147–68
[25] Magesan E *et al* 2012 Efficient measurement of quantum gate error by interleaved randomized benchmarking *Phys. Rev. Lett.* **109** 080505

Bibliography

[1] R. Harper and S. T. Flammia, Quantum Science and Technology **2**, 015008 (2017).

[2] R. Harper *et al.*, Phys. Rev. A **95**, 042120 (2017).

[3] M. A. Fogarty *et al.*, Phys. Rev. A **92**, 022326 (2015).

[4] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, New Journal of Physics **17**, 113020 (2015).

[5] J. J. Wallman, ArXiv e-prints (2017), 1703.09835.

[6] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).

[7] P. W. Shor, SIAM Journal on Computing **26**, 1484 (1997).

[8] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pp. 212–219, New York, NY, USA, 1996, ACM.

[9] D. Gottesman, *Stabilizer codes and quantum error correction*, PhD thesis, California Institute of Technology, 1997.

[10] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, New York, NY, USA, 2011).

[11] A. Kitaev, Russ. Math. Surv. **52**, 1191 (1997).

[12] D. Leibfried *et al.*, Phys. Rev. Lett. **77**, 4281 (1996).

[13] Z. Hradil, Phys. Rev. A **55**, R1561 (1997).

[14] R. Blume-Kohout, New J. Phys. **12** (2010).

[15] C. Granade, J. Combes, and D. G. Cory, New Journal of Physics **18**, 033024 (2016).

[16] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Phys. Rev. Lett. **105**, 150401 (2010).

[17] I. L. Chuang and M. A. Nielsen, Journal of Modern Optics **44**, 2455 (1997).

[18] E. Mount *et al.*, Physical Review A **92**, 060301 (2015).

[19] M. Veldhorst *et al.*, Nat Nano **9**, 981 (2014), Letter.

[20] D. Greenbaum, ArXiv e-prints (2015), 1509.02921.

[21] R. Blume-Kohout *et al.*, Nature Communications **8** (2017).

[22] J. Emerson, R. Alicki, and K. Życzkowski, J. Opt. B Quantum Semiclassical Opt. **7**, S347 (2005).

[23] E. Knill *et al.*, Phys. Rev. A **77**, 012307 (2008).

[24] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Phys. Rev. A **80**, 012304 (2009).

[25] E. Magesan, J. M. Gambetta, and J. Emerson, Phys. Rev. Lett. **106**, 180504 (2011).

[26] N. S. Yanofsky, ArXiv e-prints (2007), 0708.0261.

[27] A. Einstein, B. Podolsky, and N. Rosen, Physical Review **47**, 777 (1935).

[28] D. Gottesman and I. L. Chuang, Nature **402**, 390 (1999).

[29] R. Koenig and J. A. Smolin, Journal of Mathematical Physics **55**, 122202 (2014).

[30] S. Aaronson and D. Gottesman, Phys. Rev. A **70**, 052328 (2004).

[31] R. Cleve, D. Leung, L. Liu, and C. Wang, Quant. Inf. Comput. **16**, 0721 (2016), 1501.04592.

[32] D. Gottesman, eprint arXiv (1998), quant-ph/9807006.

[33] M. A. Nielsen, Phys. Lett. A **303**, 249 (2002).

[34] E. Magesan, R. Blume-Kohout, and J. Emerson, Phys. Rev. A **84**, 012309 (2011).

[35] C. J. Wood, J. D. Biamonte, and D. G. Cory, Quantum Info. Comput. **15**, 759 (2015).

[36] J. C. Bridgeman and C. T. Chubb, Journal of Physics A: Mathematical and Theoretical **50**, 223001 (2017).

[37] W. F. Stinespring, Proceedings of the American Mathematical Society **6**, 211 (1955).

[38] J. J. Wallman and S. T. Flammia, New J. Phys. **16**, 103032 (2014).

[39] K. Kraus, *States, Effects and Operations: Fundamental Notions of Quantum Theory,* Vol 190*, Lecture Notes in Physics* (Springer, Berlin, 2011).

[40] A. Jamiokowski, Reports on Mathematical Physics **3**, 275 (1972).

[41] I. L. Chuang and M. A. Nielsen, J. Mod. Opt. **44**, 2455 (1997).

[42] C. King and M. B. Ruskai, IEEE Transactions on Information Theory **47**, 192 (2001).

[43] Ł. Rudnicki, Z. Puchała, and K. Zyczkowski, ArXiv e-prints (2017), 1707.06926.

[44] H. Zhu, R. Kueng, M. Grassl, and D. Gross, ArXiv e-prints (2016), 1609.08172.

[45] D. Gross, K. Audenaert, and J. Eisert, J. Math. Phys. **48**, 052104 (2007).

[46] B. Schumacher, Phys. Rev. A **51**, 2738 (1995).

[47] R. Jozsa, Journal of Modern Optics **41**, 2315 (1994).

[48] M. Horodecki, P. Horodecki, and R. Horodecki, Phys. Rev. A **60**, 1888 (1999).

[49] A. C. Dugas, J. J. Wallman, and J. Emerson, ArXiv e-prints (2016), 1610.05296.

[50] N. Johnston and D. W. Kribs, Journal of Physics A: Mathematical and Theoretical **44**, 495303 (2011).

[51] E. Magesan, J. M. Gambetta, and J. Emerson, Phys. Rev. A **85**, 042311 (2012).

[52] R. N. Alexander, P. S. Turner, and S. D. Bartlett, Phys. Rev. A **94**, 032303 (2016).

[53] J. P. Gaebler *et al.*, Phys. Rev. Lett. **108**, 260503 (2012).

[54] K. R. Brown *et al.*, Phys. Rev. A **84**, 030303 (2011).

[55] J. M. Chow *et al.*, Phys. Rev. A **82**, 040305 (2010).

[56] J. M. Chow *et al.*, Phys. Rev. Lett. **102**, 090502 (2009).

[57] C. A. Ryan, M. Laforest, and R. Laflamme, New Journal of Physics **11**, 013034 (2009).

[58] S. Olmschenk, R. Chicireanu, K. D. Nelson, and J. V. Porto, New Journal of Physics **12**, 113007 (2010).

[59] A. D. Córcoles *et al.*, Phys. Rev. A **87**, 030301 (2013).

[60] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, Phys. Rev. A **89**, 062321 (2014).

[61] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout, Physical Review Letters **119** (2017).

[62] R. Blume-Kohout *et al.*, ArXiv e-prints (2013), 1310.4492.

[63] H. Ball, T. M. Stace, S. T. Flammia, and M. J. Biercuk, Phys. Rev. A **93**, 022303 (2016).

[64] B. H. Fong and S. T. Merkel, ArXiv e-prints (2017), 1703.09747.

[65] E. Magesan *et al.*, Phys. Rev. Lett. **109**, 080505 (2012).

[66] R. Barends *et al.*, Nature **508**, 500 (2014).

[67] K. Takeda *et al.*, Science Advances **2** (2016).

[68] E. Kawakami *et al.*, Proceedings of the National Academy of Sciences **113**, 11738 (2016).

[69] J. T. Muhonen *et al.*, Journal of Physics: Condensed Matter **27**, 154205 (2015).

[70] S. Kimmel, M. P. da Silva, C. a. Ryan, B. R. Johnson, and T. Ohki, Phys. Rev. X **4**, 011050 (2014).

[71] A. Carignan-Dugas, J. J. Wallman, and J. Emerson, Phys. Rev. A **92**, 060302 (2015).

[72] A. W. Cross, E. Magesan, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, Npj Quantum Information **2**, 16012 EP (2016), Article.

[73] O. Moussa, M. P. da Silva, C. A. Ryan, and R. Laflamme, Phys. Rev. Lett. **109**, 070504 (2012).

[74] B. R. Johnson *et al.*, New Journal of Physics **17**, 113019 (2015).

[75] J. Helsen, J. J. Wallman, and S. Wehner, ArXiv e-prints (2016), 1609.08188.

[76] J. Helsen, J. J. Wallman, S. T. Flammia, and S. Wehner, ArXiv e-prints (2017), 1701.04299.

[77] C. Granade, C. Ferrie, and D. G. Cory, New Journal of Physics **17**, 013042 (2015).

[78] J. Combes, C. Granade, C. Ferrie, and S. T. Flammia, ArXiv e-prints (2017), 1702.03688.

[79] G. Feng *et al.*, Phys. Rev. Lett. **117**, 260501 (2016).

[80] D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error, in *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing*, pp. 176–188, 1997.

[81] E. Knill, R. Laflamme, and W. H. Zurek, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **454**, 365 (1998).

[82] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, New J. Phys. **18**, 012002 (2015), 1501.04932.

[83] J. J. Wallman, ArXiv e-prints (2015), 1511.00727.

[84] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, Physical Review Letters **117** (2016).

[85] J. J. Wallman, M. Barnhill, and J. Emerson, Phys. Rev. Lett. **115**, 060501 (2015), 1412.4126.

[86] Z. Chen *et al.*, Phys. Rev. Lett. **116**, 020501 (2016).

[87] J. J. Wallman, M. Barnhill, and J. Emerson, New Journal of Physics **18**, 043021 (2016).

[88] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, ArXiv e-prints (2016), 1601.02036.

[89] N. Wiebe, A. Kapoor, C. Granade, and K. M. Svore, ArXiv e-prints (2015), 1507.02642.

[90] D. Ristè *et al.*, npj Quantum Information **3**, 16 (2017).

[91] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Scientific Reports **7** (2017).

[92] G. Torlai *et al.*, ArXiv e-prints (2017), 1703.05334.

[93] N. L. Cramer, A representation for the adaptive generation of simple sequential programs, in *Proceedings of the 1st International Conference on Genetic Algorithms*, pp. 183–187, Hillsdale, NJ, USA, 1985, L. Erlbaum Associates Inc.

[94] J. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, 1st ed. (The MIT Press, Cambridge MA 02142-1209, 1992).

[95] C. Ryan, J. Collins, and M. O. Neill, *Grammatical evolution: Evolving programs for an arbitrary language* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1998), pp. 83–96.

[96] R. Harper and A. Blair, Dynamically defined functions in grammatical evolution, in *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on*, pp. 2638–2645, IEEE, 2006.

[97] R. Harper, Ge, explosive grammars and the lasting legacy of bad initialisation, in *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pp. 1–8, IEEE, 2010.

[98] R. Harper, Spatial co-evolution: quicker, fitter and less bloated, in *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pp. 759–766, ACM, 2012.

[99] M. Schmidt and H. Lipson, Science **324**, 81 (2009).

[100] D. Poulin, Surprising facts about quantum error correction, 2017.