Quantitative Robustness Analysis of Quantum Programs

SHIH-HAN HUNG, KESHA HIETALA, and SHAOPENG ZHU, University of Maryland, College Park, USA

MINGSHENG YING, University of Technology Sydney, Australia, State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, China, and Tsinghua University, China

MICHAEL HICKS and XIAODI WU, University of Maryland, College Park, USA

Quantum computation is a topic of significant recent interest, with practical advances coming from both research and industry. A major challenge in quantum programming is dealing with errors (quantum noise) during execution. Because quantum resources (e.g., qubits) are scarce, classical error correction techniques applied at the level of the architecture are currently cost-prohibitive. But while this reality means that quantum programs are almost certain to have errors, there as yet exists no principled means to reason about erroneous behavior. This paper attempts to fill this gap by developing a semantics for erroneous quantum while programs, as well as a logic for reasoning about them. This logic permits proving a property we have identified, called ϵ -robustness, which characterizes possible “distance” between an ideal program and an erroneous one. We have proved the logic sound, and showed its utility on several case studies, notably: (1) analyzing the robustness of noisy versions of the quantum Bernoulli factory (QBF) and quantum walk (QW); (2) demonstrating the (in)effectiveness of different error correction schemes on single-qubit errors; and (3) analyzing the robustness of a fault-tolerant version of QBF.

CCS Concepts: • **Theory of computation** → **Denotational semantics**; **Quantum information theory**;

Additional Key Words and Phrases: quantum programming, quantum noise, approximate computing

ACM Reference Format:

Shih-Han Hung, Kesha Hietala, Shaopeng Zhu, Mingsheng Ying, Michael Hicks, and Xiaodi Wu. 2019. Quantitative Robustness Analysis of Quantum Programs. *Proc. ACM Program. Lang.* 3, POPL, Article 31 (January 2019), 29 pages. <https://doi.org/10.1145/3290344>

1 INTRODUCTION

Quantum programming has been actively investigated for the past two decades. Early work on semantics and language design [Grattage 2005; Ömer 2003; Sabry 2003; Sanders and Zuliani 2000; Selinger 2004b] has been followed up, in the last few years, by the development of a number of mature languages, including Quipper [Green et al. 2013], Scaffold [Abhari et al. 2012], LIQW [Wecker and Svore 2014], Q# [Svore et al. 2018], and QWIRE [Paykin et al. 2017]. Various program logics have also been extended for verification of quantum programs [Baltag and Smets 2011; Brunet and Jorrand 2004; Chadha et al. 2006; Feng et al. 2007; Kakutani 2009; Ying 2011; Ying et al. 2017]. For detailed surveys, see Selinger [2004a], Gay [2006], and Ying [2016].

Authors’ addresses: Shih-Han Hung; Kesha Hietala; Shaopeng Zhu, University of Maryland, College Park, USA; Mingsheng Ying, University of Technology Sydney, Australia, State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, China, Tsinghua University, China; Michael Hicks; Xiaodi Wu, University of Maryland, College Park, USA.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART31

<https://doi.org/10.1145/3290344>

A major practical challenge in implementing quantum programs is dealing with errors (aka quantum noise) during execution. Most existing work on algorithms and programming languages assumes this problem will be solved by the hardware, as in classical computers, or with fault-tolerant protocols that are designed independently of any particular application [Chong et al. 2017]. As such, the semantics of programs is defined in a manner that ignores the possibility of errors [Green et al. 2013; Paykin et al. 2017; Wecker and Svore 2014].

Unfortunately, providing such a general-purpose, fault-tolerant quantum computing abstraction appears to be impractical for near-term quantum devices, for which precisely controllable qubits are expensive, error-prone, and scarce. Existing error correction techniques consume a substantial number of qubits, severely limiting the range of possible computations. For example, one logical qubit may require $10^3 - 10^4$ physical qubits [Fowler et al. 2012]. Furthermore, fault-tolerant operations on these logical qubits require many more physical operations than their non-fault-tolerant counterparts.

As such, research on *practical* quantum computation must focus on Noisy Intermediate-Scale Quantum (NISQ) computers (as phrased by Preskill [2018]), which will lack general-purpose fault tolerance. While some particular algorithms have been developed to reflect this reality [Moll et al. 2018; Peruzzo et al. 2014], there is as yet no principled method to reason about the error-affected performance of quantum applications. Such methods are needed to help guide the design of practical applications for near-term devices.

Contributions. This paper extends the quantum **while**-language [Ying 2011] with a semantics that accounts for the possibility of error, and defines an accompanying logic for reasoning about erroneous executions. Our work constitutes an alternative to the common, but impractical, one-size-fits-all approach to fault tolerance and instead elevates the question of errors to the level of the programming language. Our approach is inspired by the work of Carbin et al. [2013], which reasons about classical programs running on unreliable hardware. We make four main contributions.

First, we present the syntax and semantics (both operational and denotational) of the quantum **while**-language extended to include noisy operations. In particular, we modify unitary application to allow the noisy operation Φ (a superoperator) to occur with probability p . This approach permits modeling any local noise occurring during the execution of a quantum program, which is the standard noise model considered in the study of quantum error correction and fault-tolerant quantum computation [Gottesman 2010]. This error model is also used by experimental physicists for building and benchmarking quantum devices in both academia and industry.

Second, we define a notion of *quantum robustness*. In particular, we say that a noisy program $\tilde{P}$ is ϵ -robust under (Q, λ) if it computes a quantum state at most ϵ distance away from that of that of its ideal equivalent P when starting both from states satisfying quantum predicate Q [D’Hondt and Panangaden 2006] to degree λ . Our definition makes use of the so-called *diamond norm* [Gilchrist et al. 2005] to account for the potential enlarging effect of entanglement on the distance. We generalize the diamond norm to what we call the (Q, λ) -diamond norm, which allows us to consider only input states that satisfy a quantum predicate Q to degree λ . Doing so obtains more accurate bounds when considering specific quantum devices and/or knowledge of states owing to the use of classical control operators. We show that the (Q, λ) -diamond norm can be computed by a semidefinite program (SDP) by extending the algorithm of Watrous [2009].

Third, we define a logic for reasoning about quantum robustness, with the following judgment¹

$$(Q, \lambda) \vdash \tilde{P} \leq \epsilon. \quad (1.1)$$

¹Our notation draws an analogy with the typing judgment $\Gamma \vdash P : t$. In particular, Q and λ are “assumptions” about inputs just as Γ represents assumptions about input (their types); $\tilde{P}$ is the program we are reasoning about; and ϵ is the proved robustness of this program, just as t is the proved type of the program. Our rules are compositional like those of typing.

This judgment states that for any input state that satisfies a quantum predicate Q to degree λ , the distance between $\llbracket \widetilde{P} \rrbracket$ and $\llbracket P \rrbracket$ is then bounded by ϵ .

We prove our logic is sound. A particular challenge is the rule for loops, owing to the *termination problem* first studied by Li and Ying [2018]. In particular, if the loop body generates some error and the loop does not terminate, it is hard to prove any non-trivial bound on the final accumulated error. To avoid this difficulty in the setting of approximate computing, Carbin et al. [2013] simply assume that the loop will terminate within a bounded number of iterations or a trivial upper bound will be applied. To capture more complicated cases, we introduce a concept called the (a, n) -boundedness of the loop. Intuitively, a loop is (a, n) -bounded if, for every input state, after n iterations it is guaranteed that with probability at least $1 - a$ it has exited the loop. The probability is due to the quantum measurement in the loop guard. It is easy to see that (a, n) -boundedness implies the termination of the loop. Pleasantly, (a, n) bounding the ideal loop is sufficient to reason about a noisy loop with any error model Φ in the loop body.

Our fourth and final contribution is to develop several case studies that demonstrate the utility of reasoning about errors at the program level.

- We start with important examples in the quantum **while**-language, the *quantum Bernoulli factory* (QBF) and the *quantum walk* (QW), and directly analyze the robustness of noisy versions $\widetilde{QBF}$ and $\widetilde{QW}$ using our logic. In the process, we prove the (a, n) -boundedness of the loops in $\widetilde{QBF}$ and $\widetilde{QW}$ using both analytical and numerical methods.
- We also demonstrate the use of our semantics to show the efficiency of different error correction schemes. Consider the error correction of a single qubit in the environment where a single bit flip error happens with probability $0 < p < 1/2$. We consider three schemes: (1) P_1 without any error correction; (2) P_2 with error correction for bit flips; (3) P_3 with error correction for phase flips. We prove that their corresponding robustnesses $\epsilon_1, \epsilon_2, \epsilon_3$ (under trivial precondition $Q = I, \lambda = 0$) satisfy $\epsilon_2 < \epsilon_1 < \epsilon_3$. In other words, we conclude that an error correction scheme that is appropriate for the error model can reduce noise in a program, while an inappropriate error correction scheme may do the opposite.
- Combining these ideas, we further analyze the robustness of a fault-tolerant version of QBF and demonstrate that the use of appropriate fault-tolerant gadgets makes QBF more robust.

Organization. We introduce preliminaries about quantum information and the quantum **while**-language in Section 3 and Section 4 respectively. We present the noisy quantum **while**-language (syntax, semantics) in Section 5 and define quantum robustness and a logic for bounding it in Section 6. We conclude the paper with case studies in Section 7. We discuss related work, next.

2 RELATED WORK

2.1 Reasoning about Errors in Classical Programs

There has been significant work on reasoning about errors in classical software, which we describe here. We believe that reasoning about errors is even more important in a quantum setting, where we can expect that errors will be prevalent and that the nature of errors will change rapidly as hardware progresses.

2.1.1 Faulty Hardware. One example of an error that a classical computer may experience is a *transient hardware fault*. Previous work has demonstrated that programming language techniques can help give safety guarantees even in the presence of such errors. For example, Walker et al. [2006] present a type-theoretic framework for analyzing fault-tolerant lambda calculus in the presence of transient faults. Their type system guarantees that well-typed programs can tolerate

any single data fault. [Perry et al. \[2007\]](#) extend those ideas to typed assembly language. Similar to this work, we present a semantics that tracks errors during computation.

2.1.2 Program Continuity. Our work is also similar to existing work on verifying program continuity and robustness [[Chaudhuri et al. 2010, 2011](#)]. In order to prove that a program is continuous or robust, it is necessary to show that the output of that program will not change significantly given small changes to the program's inputs. We are interested in bounding the distance between the output of a noisy program and its corresponding ideal program given the same input.

2.1.3 Approximate and Probabilistic Computing. In recent years, there have been many advances in programming language support for approximate [[Baek and Chilimbi 2010; Boston et al. 2015; Carbin et al. 2012; Park et al. 2015; Sampson et al. 2011](#)] and probabilistic [[Bornholt et al. 2014; Sampson et al. 2014](#)] computing. Both of these styles of computing rely on uncertainty during computation, either due to hardware errors or randomness, but still require that programs satisfy some correctness properties. Programming language tools in this area have aimed to provide these correctness guarantees.

[Carbin et al. \[2013\]](#) present a tool for verifying *reliability conditions*, which indicate the probability that a computation produces the correct result. Given a hardware specification, which lists the probability with which an operation executes correctly, their analysis computes a conservative probability that a program value is computed correctly. We extend this notion to a quantum setting by computing the *robustness* of a quantum program, which relates to the probability that the state of the system after executing the program is correct. Also, our hardware specification records not only the probability of error, but also the nature of the errors that may occur (operator Φ). Doing so is more computationally expensive, but necessary to be able to reason about the effects of error correction schemes in quantum programs. [Carbin et al. \[2013\]](#) does not consider error correction; instead, errors are considered permanent and only probabilities of errors are used for characterizations.

2.2 Characterizing Error in Quantum Programs

To our knowledge, we are presenting the first semantics and logic for quantum computation that considers errors. Existing work on characterizing error in quantum programs has focused primarily on dynamic approaches such as simulation [[Gutiérrez et al. 2013](#)] or physical experimentation [[Chuang and Nielsen 1997; Emerson et al. 2005; Knill et al. 2008; Magesan et al. 2011](#)]. Resource estimation tools like QuRE can statically produce error estimates for algorithms given hardware specifications [[Suchara et al. 2013](#)]. However, these tools are targeted at quantum *circuits* (i.e. quantum programs without conditionals or loops) and typically assume that a single error model and single type of error correction will be used throughout the circuit.

3 QUANTUM INFORMATION: PRELIMINARIES AND NOTATIONS

This section presents background and notation on quantum information and quantum computation. For an extended background, we recommend notes by [Watrous \[2006\]](#) and the textbook by [Nielsen and Chuang \[2000\]](#). A summary of notation we use appears in Table 1.

3.1 Preliminaries

For any finite integer n , an n -dimensional Hilbert space $\mathcal{H}$ is essentially the space $\mathbb{C}^n$ of complex vectors. We use Dirac's notation, $|\psi\rangle$, to denote a complex vector in $\mathbb{C}^n$. The inner product of two vectors $|\psi\rangle$ and $|\phi\rangle$ is denoted by $\langle\psi|\phi\rangle$, which is the product of the Hermitian conjugate of $|\psi\rangle$, denoted by $\langle\psi|$, and vector $|\phi\rangle$. The norm of a vector $|\psi\rangle$ is denoted by $\| |\psi\rangle \| = \sqrt{\langle\psi|\psi\rangle}$.

Table 1. A brief summary of notation used in this paper

Hilbert Spaces:	$\mathcal{H}, \mathcal{A}$
States:	(pure states) $ \psi\rangle, \phi\rangle$ (metavariables); $ 0\rangle, 1\rangle, +\rangle, -\rangle$ (notable states) (density operators) ρ, σ (metavariables); $ \psi\rangle\langle\psi $ (as outer product)
Operations:	(unitaries) U, V (metavariables); H, X, Z (notable operations) (superoperators) $\mathcal{E}, \mathcal{F}$ (general); Φ (used to represent noise)
Measurements:	M $\{M_m\}_m$ (general); $\{M_0 = 0\rangle\langle 0 , M_1 = 1\rangle\langle 1 \}$ (example)

We define (linear) *operators* as linear mappings between Hilbert spaces. Operators between n -dimensional Hilbert spaces are represented by $n \times n$ matrices. For example, the identity operator $I_{\mathcal{H}}$ can be identified by the identity matrix on $\mathcal{H}$. The Hermitian conjugate of operator A is denoted by $A^\dagger$. Operator A is *Hermitian* if $A = A^\dagger$. The trace of an operator A is the sum of the entries on the main diagonal, i.e., $\text{tr}(A) = \sum_i A_{ii}$. We write $\langle\psi|A|\psi\rangle$ to mean the inner product between $|\psi\rangle$ and $A|\psi\rangle$. A Hermitian operator A is *positive semidefinite* (resp., *positive definite*) if for all vectors $|\psi\rangle \in \mathcal{H}$, $\langle\psi|A|\psi\rangle \geq 0$ (resp., > 0). This gives rise to the *Löwner order* $\sqsubseteq$ among operators:

$$A \sqsubseteq B \text{ if } B - A \text{ is positive semidefinite, } A \subset B \text{ if } B - A \text{ is positive definite.} \quad (3.1)$$

3.2 Quantum States

The state space of a quantum system is a Hilbert space. The state space of a *qubit*, or quantum bit, is a 2-dimensional Hilbert space. One important orthonormal basis of a qubit system is the *computational* basis with $|0\rangle = (1, 0)^\dagger$ and $|1\rangle = (0, 1)^\dagger$, which encode the classical bits 0 and 1 respectively. Another important basis, called the $\pm$ basis, consists of $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. The state space of multiple qubits is the *tensor product* of single qubit state spaces. For example, classical 00 can be encoded by $|0\rangle \otimes |0\rangle$ (written $|0\rangle|0\rangle$ or even $|00\rangle$ for short) in the Hilbert space $\mathbb{C}^2 \otimes \mathbb{C}^2$. The Hilbert space for an m -qubit system is $(\mathbb{C}^2)^{\otimes m} \cong \mathbb{C}^{2^m}$.

A *pure* quantum state is represented by a unit vector, i.e., a vector $|\psi\rangle$ with $\| |\psi\rangle \| = 1$. A *mixed* state can be represented by a classical distribution over an ensemble of pure states $\{(p_i, |\psi_i\rangle)\}_i$, i.e., the system is in state $|\psi_i\rangle$ with probability p_i . One can also use *density operators* to represent both pure and mixed quantum states. A density operator ρ for a mixed state representing the ensemble $\{(p_i, |\psi_i\rangle)\}_i$ is a positive semidefinite operator $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, where $|\psi_i\rangle\langle\psi_i|$ is the outer-product of $|\psi_i\rangle$; in particular, a pure state $|\psi\rangle$ can be identified with the density operator $\rho = |\psi\rangle\langle\psi|$. Note that $\text{tr}(\rho) = 1$ holds for all density operators. A positive semidefinite operator ρ on $\mathcal{H}$ is said to be a *partial* density operator if $\text{tr}(\rho) \leq 1$. The set of partial density operators is denoted by $\mathcal{D}(\mathcal{H})$.

3.3 Quantum Operations

Operations on quantum systems can be characterized by unitary operators. An operator U is *unitary* if its Hermitian conjugate is its own inverse, i.e., $U^\dagger U = U U^\dagger = I$. For a pure state $|\psi\rangle$, a unitary operator describes an *evolution* from $|\psi\rangle$ to $U|\psi\rangle$. For a density operator ρ , the corresponding evolution is $\rho \mapsto U\rho U^\dagger$. Common single-qubit unitary operators include

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (3.2)$$

The *Hadamard* operator H transforms between the computational and the $\pm$ basis. For example, $H|0\rangle = |+\rangle$ and $H|1\rangle = |-\rangle$. The *Pauli X* operator is a bit flip, i.e., $X|0\rangle = |1\rangle$ and $X|1\rangle = |0\rangle$. The *Pauli Z* operator is a phase flip, i.e., $Z|0\rangle = |0\rangle$ and $Z|1\rangle = -|1\rangle$.

More generally, the evolution of a quantum system can be characterized by an *admissible superoperator* $\mathcal{E}$, which is a *completely-positive* and *trace-non-increasing* linear map from $\mathcal{D}(\mathcal{H})$ to $\mathcal{D}(\mathcal{H}')$ for Hilbert spaces $\mathcal{H}, \mathcal{H}'$. A superoperator is positive if it maps from $\mathcal{D}(\mathcal{H})$ to $\mathcal{D}(\mathcal{H}')$ for Hilbert spaces $\mathcal{H}, \mathcal{H}'$. A superoperator $\mathcal{E}$ is k -positive if for any k -dimensional Hilbert space $\mathcal{A}$, the superoperator $\mathcal{E} \otimes I_{\mathcal{A}}$ is a positive map on $\mathcal{D}(\mathcal{H} \otimes \mathcal{A})$. A superoperator is said to be completely positive if it is k -positive for any positive integer k . A superoperator $\mathcal{E}$ is trace-non-increasing if for any initial state $\rho \in \mathcal{D}(\mathcal{H})$, the final state $\mathcal{E}(\rho) \in \mathcal{D}(\mathcal{H}')$ after applying $\mathcal{E}$ satisfies $\text{tr}(\mathcal{E}(\rho)) \leq \text{tr}(\rho)$.

For every superoperator $\mathcal{E} : \mathcal{D}(\mathcal{H}) \rightarrow \mathcal{D}(\mathcal{H}')$, there exists a set of Kraus operators $\{E_k\}_k$ such that $\mathcal{E}(\rho) = \sum_k E_k \rho E_k^\dagger$ for any input $\rho \in \mathcal{D}(\mathcal{H})$. Note that the set of Kraus operators is finite if the Hilbert space is finite-dimensional. The *Kraus form* of $\mathcal{E}$ is written as $\mathcal{E} = \sum_k E_k \circ E_k^\dagger$. A unitary evolution can be represented by the superoperator $\mathcal{E} = U \circ U^\dagger$. An identity operation refers to the superoperator $I_{\mathcal{H}} = I_{\mathcal{H}} \circ I_{\mathcal{H}}$. The Schrödinger-Heisenberg *dual* of a superoperator $\mathcal{E} = \sum_k E_k \circ E_k^\dagger$, denoted by $\mathcal{E}^*$, is defined as follows: for every state $\rho \in \mathcal{D}(\mathcal{H})$ and any operator A , $\text{tr}(A\mathcal{E}(\rho)) = \text{tr}(\mathcal{E}^*(A)\rho)$. The Kraus form of $\mathcal{E}^*$ is $\sum_k E_k^\dagger \circ E_k$.

3.4 Quantum Measurements

The way to extract information about a quantum system is called a quantum *measurement*. A quantum measurement on a system over Hilbert space $\mathcal{H}$ can be described by a set of linear operators $\{M_m\}_m$ with $\sum_m M_m^\dagger M_m = I_{\mathcal{H}}$. If we perform a measurement $\{M_m\}$ on a state ρ , the outcome m is observed with probability $p_m = \text{tr}(M_m \rho M_m^\dagger)$ for each m . A major difference between classical and quantum computation is that a quantum measurement changes the state. In particular, after a measurement yielding outcome m , the state collapses to $M_m \rho M_m^\dagger / p_m$. For example, a measurement in the computational basis is described by $M = \{M_0 = |0\rangle\langle 0|, M_1 = |1\rangle\langle 1|\}$. If we perform the computational basis measurement M on state $\rho = |+\rangle\langle +|$, then with probability $\frac{1}{2}$ the outcome is 0 and ρ becomes $|0\rangle\langle 0|$. With probability $\frac{1}{2}$ the outcome is 1 and ρ becomes $|1\rangle\langle 1|$.

4 QUANTUM PROGRAMS

Our work builds on top of the quantum **while**-language developed by Ying [2011, 2016]. Here we review the syntax and semantics of this language.

4.1 Syntax

Define Var as the set of quantum variables. We use the symbol q as a metavariable ranging over quantum variables and define a *quantum register* $\bar{q}$ to be a finite set of distinct variables. For each $q \in Var$, its state space is denoted by $\mathcal{H}_q$. The quantum register $\bar{q}$ is associated with the Hilbert space $\mathcal{H}_{\bar{q}} = \bigotimes_{q \in \bar{q}} \mathcal{H}_q$. If $\text{type}(q) = \mathbf{Bool}$ then $\mathcal{H}_q$ is the two-dimensional Hilbert space with basis $\{|0\rangle, |1\rangle\}$. If $\text{type}(q) = \mathbf{Int}$ then $\mathcal{H}_q$ is the Hilbert space with basis $\{|n\rangle : n \in \mathbb{Z}\}$. The syntax of a quantum **while** program P is defined as follows.

$$\begin{aligned} P ::= & \mathbf{skip} \mid q := |0\rangle \mid \bar{q} := U[\bar{q}] \mid P_1; P_2 \mid \\ & \mathbf{case } M[\bar{q}] = \overline{m \rightarrow P_m} \mathbf{end} \mid \mathbf{while } M[\bar{q}] = 1 \mathbf{do } P_1 \mathbf{done} \end{aligned} \quad (4.1)$$

The language constructs above are similar to their classical counterparts. (1) **skip** does nothing. (2) $q := |0\rangle$ sets quantum variable q to the basis state $|0\rangle$. (3) $\bar{q} := U[\bar{q}]$ applies the unitary U to the qubits in $\bar{q}$. (4) Sequencing has the same behavior as its classical counterpart. (5) **case** $M[\bar{q}] = \overline{m \rightarrow P_m}$ **end** performs the measurement $M = \{M_m\}$ on the qubits in $\bar{q}$, and executes program P_m if the outcome of the measurement is m . The bar over $\overline{m \rightarrow P_m}$ indicates that there may be one or more repetitions

of this expression. ² (6) **while** $M[\bar{q}] = 1$ **do** P_1 **done** performs the measurement $M = \{M_0, M_1\}$ on the qubits in $\bar{q}$, and executes P_1 if measurement produces the outcome corresponding to M_1 or terminates if measurement produces the outcome corresponding to M_0 .

We highlight two differences between quantum and classical while languages: (1) Qubits may only be initialized to the basis state $|0\rangle$. There is no quantum analogue for initialization to any expression (i.e. $x := e$) because of the no-cloning theorem of quantum states. Any state $|\psi\rangle \in \mathcal{H}_q$, however, can be constructed by applying some unitary U to $|0\rangle$. ³ (2) Evaluating the guard of a case statement or loop, which performs a measurement, potentially disturbs the state of the system.

We now present an example program written in the quantum **while**-language syntax. The *quantum walk* [Aharonov et al. 2001] is a widely considered example in quantum programming, quantum algorithms, and quantum simulation literature [Georgescu et al. 2014; Ying et al. 2017]. Here we consider a quantum walk on a circle with n points. We let the initial position of the walker be 0, and say that the program halts if and only if the walker arrives at position 1.

EXAMPLE 4.1 (QUANTUM WALK). Define the coin (or "direction") space $\mathcal{H}_c$ to be the 2-dimensional Hilbert space with orthonormal basis states $|L\rangle$ and $|R\rangle$, for Left and Right respectively. Define the position space $\mathcal{H}_p$ to be the n -dimensional Hilbert space with orthonormal basis states $|0\rangle, |1\rangle, \dots, |n-1\rangle$, where vector $|i\rangle$ represents position i for $0 \leq i < n$. Now the state space of the walk is $\mathcal{H} = \mathcal{H}_c \otimes \mathcal{H}_p$ and the initial state is $|L\rangle|0\rangle$. In each step of the walk:

- (1) Measure the position of the system to determine whether the walker has reached position 1. If the walker has reached position 1, the walk terminates. Otherwise, it continues. We use the measurement $M = \{|1\rangle\langle 1|, \sum_{i \neq 1} |i\rangle\langle i|\}$.
- (2) Apply the "coin-tossing" operator H to the coin space $\mathcal{H}_c$.
- (3) Perform the shift operator S defined by $S|L, i\rangle = |L, i-1(\text{mod } n)\rangle$, $S|R, i\rangle = |R, i+1(\text{mod } n)\rangle$ for $i = 0, 1, \dots, n-1$ to the space $\mathcal{H}$. The S operator can be written as

$$S = \sum_{i=0}^{n-1} |L\rangle\langle L| \otimes |i-1(\text{mod } n)\rangle\langle i| + \sum_{i=0}^{n-1} |R\rangle\langle R| \otimes |i+1(\text{mod } n)\rangle\langle i|.$$

In this algorithm, the walker takes one step left or one step right corresponding to the coin flip result $|L\rangle$ or $|R\rangle$. However, unlike the classical case, the result of the coin flip may be a superposition of $|L\rangle$ and $|R\rangle$, allowing the walker to take a step to the left and right simultaneously. This quantum walk can be described by the following program

$$QW_n \equiv p := |0\rangle; c := |L\rangle; \textbf{while } M[p] = 1 \textbf{ do } c := H[c]; c, p := S[c, p] \textbf{ done.} \quad (4.2)$$

4.2 Operational Semantics

The operational semantics of the quantum **while**-language are presented in Figure 1a. $\langle P, \rho \rangle \rightarrow \langle P', \rho' \rangle$, where $\langle P, \rho \rangle$ and $\langle P', \rho' \rangle$ are quantum configurations. In configurations, P (or P') could be a quantum program or the empty program E , and ρ and ρ' are partial density operators representing the current state. Intuitively, in one step, we can evaluate program P on input state ρ to program P' (or E) and output state ρ' . In order to present the rules in a non-probabilistic manner, the probabilities associated with each transition are encoded in the output partial density operator. ⁴

²Our syntax for conditional/case statements differs from that presented by Ying [2016] to make it more clear that there are multiple programs P_m .

³In our examples, we may write $q := |\psi\rangle$ for some fixed basis state $|\psi\rangle$. What we mean in this case is $q := |0\rangle; q := U[q]$ where U is the unitary operation that transforms $|0\rangle$ into $|\psi\rangle$.

⁴If we had instead considered a probabilistic transition system, then the transition rule for case statements could have been written as $\langle \text{case } M[\bar{q}] = \bar{m} \rightarrow P_m \textbf{ end}, \rho \rangle \xrightarrow{p_m} \langle P_m, \rho_m \rangle$ where $p_m = \text{tr}(M_m \rho M_m^\dagger)$ and $\rho_m = M_m \rho M_m^\dagger / p_m$.

$$\begin{array}{l}
\text{(Skip)} \quad \overline{\langle \mathbf{skip}, \rho \rangle \rightarrow \langle E, \rho \rangle} \\
\text{(Initialization)} \quad \overline{\langle q := |0\rangle, \rho \rangle \rightarrow \langle E, \rho_0^q \rangle} \\
\quad \text{where } \rho_0^q = \begin{cases} \mathcal{E}_{q \rightarrow 0}^{\text{bool}}(\rho) & \text{if } \text{type}(q) = \mathbf{Bool} \\ \mathcal{E}_{q \rightarrow 0}^{\text{int}}(\rho) & \text{if } \text{type}(q) = \mathbf{Int} \end{cases} \\
\text{(Unitary)} \quad \overline{\langle \bar{q} := U[\bar{q}], \rho \rangle \rightarrow \langle E, U\rho U^\dagger \rangle} \\
\text{(Sequence E)} \quad \overline{\langle E; P_2, \rho \rangle \rightarrow \langle P_2, \rho \rangle} \\
\text{(Sequence)} \quad \overline{\langle P_1; P_2, \rho \rangle \rightarrow \langle P_1', P_2, \rho' \rangle} \\
\text{(Case } m) \quad \overline{\langle \mathbf{case } M[\bar{q}] = m \rightarrow P_m \mathbf{end}, \rho \rangle \rightarrow \langle P_m, M_m \rho M_m^\dagger \rangle} \\
\quad \text{for each outcome } m \text{ of measurement } M = \{M_m\} \\
\text{(While 0)} \quad \overline{\langle \mathbf{while } M[\bar{q}] = 1 \mathbf{do } P_1 \mathbf{done}, \rho \rangle \rightarrow \langle E, M_0 \rho M_0^\dagger \rangle} \\
\text{(While 1)} \quad \overline{\langle \mathbf{while } M[\bar{q}] = 1 \mathbf{do } P_1 \mathbf{done}, \rho \rangle \rightarrow \langle P_1; \mathbf{while } M[\bar{q}] = 1 \mathbf{do } P_1 \mathbf{done}, M_1 \rho M_1^\dagger \rangle}
\end{array}
\tag{a}$$

$$\begin{array}{l}
\llbracket \mathbf{skip} \rrbracket \rho = \rho \\
\llbracket q := |0\rangle \rrbracket \rho = \begin{cases} \mathcal{E}_{q \rightarrow 0}^{\text{bool}}(\rho) & \text{if } \text{type}(q) = \mathbf{Bool} \\ \mathcal{E}_{q \rightarrow 0}^{\text{int}}(\rho) & \text{if } \text{type}(q) = \mathbf{Int} \end{cases} \\
\llbracket \bar{q} := U[\bar{q}] \rrbracket \rho = U\rho U^\dagger \\
\llbracket P_1; P_2 \rrbracket \rho = \llbracket P_2 \rrbracket (\llbracket P_1 \rrbracket \rho) \\
\llbracket \mathbf{case } M[\bar{q}] = m \rightarrow P_m \mathbf{end} \rrbracket \rho = \sum_m \llbracket P_m \rrbracket (M_m \rho M_m^\dagger) \\
\llbracket \mathbf{while } M[\bar{q}] = 1 \mathbf{do } P_1 \mathbf{done} \rrbracket \rho = \bigsqcup_{k=0}^{\infty} \llbracket \mathbf{while}^{(k)} \rrbracket \rho
\end{array}
\tag{b}$$

Fig. 1. quantum **while** programs: (a) operational semantics (b) denotational semantics.

In the *Initialization* rule, the superoperators $\mathcal{E}_{q \rightarrow 0}^{\text{bool}}(\rho)$ and $\mathcal{E}_{q \rightarrow 0}^{\text{int}}(\rho)$, which initialize the variable q in ρ to $|0\rangle\langle 0|$, are defined by $\mathcal{E}_{q \rightarrow 0}^{\text{bool}}(\rho) = |0\rangle_q\langle 0|\rho|0\rangle_q\langle 0| + |0\rangle_q\langle 1|\rho|1\rangle_q\langle 0|$ and $\mathcal{E}_{q \rightarrow 0}^{\text{int}}(\rho) = \sum_{n=-\infty}^{\infty} |0\rangle_q\langle n|\rho|n\rangle_q\langle 0|$. Here, $|\psi\rangle_q\langle\phi|$ denotes the outer product of states $|\psi\rangle$ and $|\phi\rangle$ associated with variable q ; that is, $|\psi\rangle$ and $|\phi\rangle$ are in $\mathcal{H}_q$ and $|\psi\rangle_q\langle\phi|$ is a matrix over $\mathcal{H}_q$. It is a convention in the quantum information literature that when operations or measurements only apply to part of the quantum system (e.g., a subset of quantum variables of the program), one should assume that an identity operation is applied to the rest of quantum variables. For example, applying $|\psi\rangle_q\langle\phi|$ to ρ means applying $|\psi\rangle_q\langle\phi| \otimes I_{\mathcal{H}_{\bar{q}}}$ to ρ , where $\bar{q}$ denotes the set of all variables except q . The identity operation is usually omitted for simplicity.

We do not explain the rules in detail, but hope their meaning is self-evident given the description of the language in [Section 4.1](#).

To illustrate the use of the operational semantics, we consider the classic *beam splitter experiment*, which demonstrates the difference between quantum and classical mechanics. In this experiment,

Fig. 2. The beam splitter experiment. The ‘0’ path corresponds to the photon having been transmitted at the first beam splitter, and the ‘1’ path corresponds to the photon having been reflected at the first beam splitter.

as shown in Figure 2, a photon source sends photons through two beam splitters. The final locations of the photons are determined using photon detectors. A classical analysis would assume that, for each photon, each beam splitter flips a fair coin to either reflect the photon or allow it to pass through. The full mirrors reflect incoming photons with probability one. As a result, we might expect half of the photons to reach one detector and half to reach the other. However, this is not what is observed in experiments. On the contrary, all photons will reach one detector.

EXAMPLE 4.2 (BEAM SPLITTER EXPERIMENT). *Model each beam splitter as a Hadamard gate and say that the photon source produces photons on the ‘0’ path, which corresponds to the $|0\rangle$ state. Then the beam splitter experiment (BSE) corresponds to the program*

$$BSE \equiv q_1 := |0\rangle; q_1 := H[q_1]; q_1 := H[q_1] \quad (4.3)$$

where $\text{type}(q_1) = \mathbf{Bool}$. Let $\rho = |1\rangle_{q_1}\langle 1|$. Then the evaluation of program BSE on input ρ proceeds as follows.

$$\begin{aligned} \langle BSE, \rho \rangle &= \langle q_1 := |0\rangle; q_1 := H[q_1]; q_1 := H[q_1], |1\rangle_{q_1}\langle 1| \rangle \\ &\rightarrow \langle q_1 := H[q_1]; q_1 := H[q_1], |0\rangle_{q_1}\langle 0| \rangle \\ &\rightarrow \langle q_1 := H[q_1], |+\rangle_{q_1}\langle +| \rangle \\ &\rightarrow \langle E, |0\rangle_{q_1}\langle 0| \rangle \end{aligned}$$

At the first beam splitter, the photons may either continue on their current path (corresponding to $|0\rangle$) or be reflected to the ‘1’ path (corresponding to $|1\rangle$). In quantum mechanics, both possibilities happen simultaneously, resulting in each photon continuing on a superposition of both paths. The superposition of both paths corresponds to the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. At the second beam splitter, the paths in superposition interfere with each other, resulting in the ‘1’ path being cancelled.

4.3 Denotational Semantics

The denotational semantics of a quantum **while** program is given in Figure 1b. It defines $\llbracket P \rrbracket$ as a superoperator that acts on $\rho \in \mathcal{H}_{Var}$ [Ying 2016]. The semantics of each term is given compositionally. We write $\mathbf{while}^{(k)}$ for the k th syntactic approximation (i.e., unrolling) of **while** and $\sqcup$ for the least upper bound operator in the complete partial order generated by Löwner comparison. For more detail on the semantics of loops, we refer the reader to Ying [2011, 2016].

We connect the denotational semantics to the operational semantics through the following proposition.

PROPOSITION 4.1 ([YING 2016]). *For any program P*

$$\llbracket P \rrbracket \rho \equiv \sum \{ |\rho' : \langle P, \rho \rangle \rightarrow^* \langle E, \rho' \rangle| \}, \quad (4.4)$$

where $\rightarrow^*$ is the reflexive, transitive closure of $\rightarrow$ and $\{ | \cdot | \}$ denotes a multi-set.

In short, the meaning of running program P on input state ρ is the sum of all possible output states, weighted by their probabilities.

The semantics presented so far assume that no noise will occur during computation. In Section 5, we extend the semantics to include possible errors that may occur during unitary application.

4.4 Quantum Predicates and Hoare Logic

A *quantum predicate* is a Hermitian operator M such that $0 \sqsubseteq M \sqsubseteq I$ [D'Hondt and Panangaden 2006]. For a predicate M and state ρ , $\text{tr}(M\rho)$ is the expectation of the truth value of predicate M in state ρ . Restricting M to be between 0 and I ensures that $0 \leq \text{tr}(M\rho) \leq 1$ for any $\rho \in \mathcal{D}(\mathcal{H})$.

The identity matrix corresponds to the **true** predicate because for any density operator ρ , $\text{tr}(I\rho) = 1$. The zero matrix corresponds to the **false** predicate because for any density operator ρ , $\text{tr}(0\rho) = 0$. $|0\rangle\langle 0|$ is the predicate that says that a state is in the subspace spanned by $|0\rangle$. As an example, the density operator ρ_0 corresponding to the state $|0\rangle$ is such that $\text{tr}(|0\rangle\langle 0|\rho_0) = 1$, and the density operator ρ_1 corresponding to the state $\sqrt{1/3}|0\rangle + \sqrt{2/3}|1\rangle$ is such that $\text{tr}(|0\rangle\langle 0|\rho_1) = \frac{1}{3}$.

Ying [2011, 2016] uses quantum predicates as the basis for defining quantum preconditions and postconditions in his quantum Hoare logic. Let M and N be quantum predicates and let P be a quantum **while** program. Then M is a precondition of N with respect to P , written $\{M\}P\{N\}$, if

$$\forall \rho. \text{tr}(M\rho) \leq \text{tr}(N\llbracket P \rrbracket \rho). \quad (4.5)$$

This inequality can be seen as the probabilistic version of the following statement: if state ρ satisfies predicate M , then after applying the program P the resulting state will satisfy predicate N . If we include an auxiliary space $\mathcal{A}$, then the equivalent statement is

$$\forall \rho. \text{tr}((M \otimes I_{\mathcal{A}})\rho) \leq \text{tr}((N \otimes I_{\mathcal{A}})(\llbracket P \rrbracket \otimes I_{\mathcal{A}})(\rho)). \quad (4.6)$$

5 NOISY QUANTUM PROGRAMS

In this section, we present the syntax and semantics for the quantum **while**-language with noise, as an extension of the quantum **while**-language. Our syntax allows one to explicitly encode any error model that describes local noise during the execution of a quantum program.

5.1 Noise in Quantum Computation

Here we briefly discuss how noise is modeled in the study of quantum error correction and fault-tolerant quantum computation [Gottesman 2010], which in turn comes from the noise model in quantum physical experiments. It is a convention to only consider “local” noise rather than correlated noise, because benign white noise is more likely than “adversarial” noise in actual quantum devices. A few types of natural local noise arise in realistic quantum systems, which generalize classical bit-flip errors, including:

- The *bit flip* noise flips the state with probability p , and can be represented by

$$\Phi_{p,\text{bit}} = (1-p)I \circ I + pX \circ X. \quad (5.1)$$

- The *phase flip* noise flips the phase with probability p , and can be represented by

$$\Phi_{p,\text{phase}} = (1-p)I \circ I + pZ \circ Z. \quad (5.2)$$

Other types of noise include depolarization, amplitude damping, and phase damping [Nielsen and Chuang 2000].

This model of noise is used by experimental physicists for building and benchmarking quantum devices in both academia and industry [Terhal 2015]. Noisy information of specific quantum devices can also be publicly available (e.g., the IBM Q-experience⁵).

5.2 Syntax

The syntax of a noisy quantum program $\tilde{P}$ is defined as follows

$$\begin{aligned} \tilde{P} ::= & \text{skip} \mid q := |0\rangle \mid \bar{q} :=_{p,\Phi} U[\bar{q}] \mid \tilde{P}_1; \tilde{P}_2 \mid \\ & \text{case } M[\bar{q}] = m \rightarrow \tilde{P}_m \text{ end} \mid \text{while } M[\bar{q}] = 1 \text{ do } \tilde{P}_1 \text{ done} \end{aligned} \quad (5.3)$$

This syntax is identical to that of the standard quantum while language described in Section 4, except that we have annotated the unitary application construct with an error probability p and an error model Φ , which is the superoperator of the noisy operation. The statement $\bar{q} :=_{p,\Phi} U[\bar{q}]$ will apply the correct operation U on $\bar{q}$ with probability $1 - p$ and will apply the noisy (or erroneous) operation Φ on $\bar{q}$ with probability p . The nature of Φ will depend on the underlying hardware, and is a parameter to our language.

For any noisy program $\tilde{P}$, its corresponding *ideal* program can be obtained by simply replacing any noisy unitary operations by their ideal versions (i.e., we ignore p and Φ). We write $\text{ideal}(\tilde{P})$ for this program, or simply P when there is no ambiguity.

We remark that noisy unitary operations are already expressive enough to capture many types of noise. First, any noise can depend on the quantum state of the system by the nature of modeling it as a quantum operation Φ . Second, noisy initialization can be modeled as initialization followed by application of a noisy identity operation, and noisy measurement can be modeled as application of a noisy identity operation followed by measurement. Third, errors that occur between applications of subsequent unitaries can also be modeled by noisy identity operations.

5.3 Semantics

The operational and denotational semantics of noisy quantum **while** programs are also identical to those of the standard quantum **while** programs, except that the rules related to unitary application now include an error term, as shown in Figure 3a and Figure 3b. Note that we do not require p and Φ to be the same in every instance of a unitary application. They may depend on the type of unitary being applied, or on other features of the program such as the number of operations performed so far. We use *explicit* characterizations of the noise given by Φ to enable us to argue about the effect of error-correcting gadgets explicitly written in the program. If we only considered error probability (like Carbin et al. [2013]) then we could only reason about error as accumulating throughout the program, and we would not be able to show that error-correcting gadgets reduce noise by cancelling previous errors.

EXAMPLE 5.1 (BEAM SPLITTER EXPERIMENT WITH ERRORS). *Consider the following noisy version of the beam splitter program.*

$$\widetilde{BSE} \equiv q_1 := |0\rangle; q_1 :=_{p,\Phi} H[q_1]; q_1 :=_{p,\Phi} H[q_1].$$

⁵<https://www.research.ibm.com/ibm-q/technology/devices/>.

$$\begin{aligned}
\text{(Unitary-Noisy)} \quad & \overline{\langle \bar{q} : \cong_{p,\Phi} U[\bar{q}], \rho \rangle \rightarrow \langle E, (1-p)U\rho U^\dagger + p\Phi(\rho) \rangle} \\
& \text{(a)} \\
& \llbracket \bar{q} : \cong_{p,\Phi} U[\bar{q}] \rrbracket \rho = (1-p)U\rho U^\dagger + p\Phi(\rho) \\
& \text{(b)}
\end{aligned}$$

Fig. 3. (a) Transition rule for noisy unitary application. (b) Denotation of noisy unitary application.

Note that the error probability p and error model Φ are the same in both applications of H . Let $\rho = |1\rangle_{q_1}\langle 1|$. Then the evaluation of program $\tilde{P}$ on input ρ proceeds as follows.

$$\begin{aligned}
\langle \widetilde{BSE}, \rho \rangle &= \langle q_1 := |0\rangle; q_1 : \cong_{p,\Phi} H[q_1]; q_1 : \cong_{p,\Phi} H[q_1], |1\rangle_{q_1}\langle 1| \rangle \\
&\rightarrow \langle q_1 : \cong_{p,\Phi} H[q_1]; q_1 : \cong_{p,\Phi} H[q_1], |0\rangle_{q_1}\langle 0| \rangle \\
&\rightarrow \langle q_1 : \cong_{p,\Phi} H[q_1], (1-p)|+\rangle_{q_1}\langle +| + p\Phi(|0\rangle_{q_1}\langle 0|) \rangle \\
&\rightarrow \langle E, (1-p)^2|0\rangle_{q_1}\langle 0| + p(1-p)H\Phi(|0\rangle_{q_1}\langle 0|)H + p\Phi(\rho_1) \rangle
\end{aligned}$$

where $\rho_1 = (1-p)|+\rangle_{q_1}\langle +| + p\Phi(|0\rangle_{q_1}\langle 0|)$. Here, the desired output state $|0\rangle_{q_1}\langle 0|$ is in superposition with error terms $H\Phi(|0\rangle_{q_1}\langle 0|)H$ and $\Phi(\rho_1)$. This means that not all of the photons will necessarily remain on the ‘0’ path. Some may end up on the ‘1’ path.

As an example, consider the case where the error probability p is 0.1 and the error model is defined by $\Phi(\rho) = X\rho X$. This means that with probability 0.1, an X gate is applied instead of an H gate. With this model, the final state of the system will be $0.91|0\rangle_{q_1}\langle 0| + 0.09|1\rangle_{q_1}\langle 1|$. So there is a 9% chance that a photon will end up on the ‘1’ path.

6 QUANTUM ROBUSTNESS

This section defines a notion of *quantum robustness*, which bounds the distance between the output of a noisy execution of a program and the ideal (noise-free) execution of the same program. We first introduce distance measures in quantum information, and then present the semantic definition of robustness and define a logic for reasoning about it, which we prove sound.

6.1 Distance Measures of Quantum States and Superoperators

In the context of computation with noise, we wish to measure the distance between the ideal state and the state influenced by noise. In classical probabilistic computation, the output can be described as a probability distribution over all possible outputs. A common measure is the *total variation* distance of two distributions p, q , defined by $|p - q| = 1/2 \sum_x |p(x) - q(x)|$.

In quantum computation, we define the *trace distance* as the quantum generalization of the total variation distance. For Hermitian operator A , let the trace norm $\|A\|_1$ be the summation of the absolute value of all its eigenvalues. When A is positive semidefinite, one has $\|A\|_1 = \text{tr}(A)$. It is worth noting that for any operators A and B the following triangle inequality holds:

$$\|A + B\|_1 \leq \|A\|_1 + \|B\|_1. \quad (6.1)$$

For two states $\rho_1, \rho_2 \in \mathcal{D}(\mathcal{H})$, the *trace distance* between ρ_1 and ρ_2 , denoted $T(\rho_1, \rho_2)$, is defined to be $\frac{1}{2}\|\rho_1 - \rho_2\|_1$. It can be shown that the trace distance satisfies

$$T(\rho_1, \rho_2) \equiv \frac{1}{2}\|\rho_1 - \rho_2\|_1 = \max_{0 \subseteq P \subseteq I} \text{tr}(P(\rho_1 - \rho_2)), \quad (6.2)$$

where the maximization is over all possible measurements P (i.e., $0 \sqsubseteq P \sqsubseteq I$). By definition, $0 \leq T(\rho, \sigma) \leq 1$ for any pair of states ρ, σ . One can also interpret the trace distance as the advantage with which one can distinguish two states ρ and σ . For example, the states $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ are perfectly distinguishable by measurement $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$, so $T(|0\rangle\langle 0|, |1\rangle\langle 1|) = 1$.

One can further define the distance between two different superoperators. For any two superoperators $\mathcal{E}$ and $\mathcal{E}'$ over $\mathcal{H}$, an intuitive way to define their distance is to find the input state to both superoperators that maximizes the trace distance of their output states. However, it turns out that this definition alone is insufficient to capture the distance between superoperators because of a unique quantum feature called *entanglement*. The distinguishability between $\mathcal{E}$ and $\mathcal{E}'$ can be significantly enlarged⁶ when one introduces an auxiliary Hilbert space $\mathcal{A}$ and tries to distinguish $\mathcal{E} \otimes I_{\mathcal{A}}$ and $\mathcal{E}' \otimes I_{\mathcal{A}}$ with some *entangled* input state over $\mathcal{H} \otimes \mathcal{A}$ [Gilchrist et al. 2005].

To account for this, we define the *diamond* norm between two superoperators $\mathcal{E}, \mathcal{E}'$ as

$$\|\mathcal{E} - \mathcal{E}'\|_{\diamond} \equiv \max_{\rho \in \mathcal{D}(\mathcal{H} \otimes \mathcal{A}) : \text{tr}(\rho)=1} T(\mathcal{E} \otimes I_{\mathcal{A}}(\rho), \mathcal{E}' \otimes I_{\mathcal{A}}(\rho)) \quad (6.3)$$

for any auxiliary space $\mathcal{A}$. Without loss of generality, one can assume $\mathcal{A}$ is a copy of $\mathcal{H}$, i.e., $\mathcal{A} = \mathcal{H}$ [Watrous 2018]. Given the representations of $\mathcal{E}, \mathcal{E}'$, one can efficiently calculate the diamond norm $\|\mathcal{E} - \mathcal{E}'\|_{\diamond}$ by a semidefinite program (SDP) [Watrous 2009].

Imagine superoperators that represent different components of a (noisy or ideal) quantum program, which may act on different parts of the quantum system. The diamond norm will allow us to address potential entanglement between different parts of the state and ensure that we can compose the distances computed from different components of the program.

6.2 Definition of Quantum Robustness

To capture how noise (error) impacts the execution of quantum program $\tilde{P}$, we want to compare $\llbracket \tilde{P} \rrbracket$ and $\llbracket P \rrbracket$, which are superoperators representing the execution of quantum program P with and without noise respectively. A natural candidate is to use the aforementioned diamond norm to measure the distance between $\llbracket \tilde{P} \rrbracket$ and $\llbracket P \rrbracket$. To account for prior knowledge of the input state, we extend the definition of the diamond norm to consider only input states that satisfy predicate Q to degree at least λ . More explicitly, we have

DEFINITION 6.1 ((Q, λ) -DIAMOND NORM). *Given superoperators $\mathcal{E}, \mathcal{E}'$, quantum predicate Q over $\mathcal{H}$, and $0 \leq \lambda \leq 1$, the (Q, λ) -diamond norm between $\mathcal{E}$ and $\mathcal{E}'$, denoted $\|\mathcal{E} - \mathcal{E}'\|_{Q, \lambda}$, is defined by*

$$\|\mathcal{E} - \mathcal{E}'\|_{Q, \lambda} \equiv \max_{\rho \in \mathcal{D}(\mathcal{H} \otimes \mathcal{A}) : \text{tr}(\rho)=1, \text{tr}(Q\rho) \geq \lambda} T(\mathcal{E} \otimes I_{\mathcal{A}}(\rho), \mathcal{E}' \otimes I_{\mathcal{A}}(\rho)), \quad (6.4)$$

where $\mathcal{A}$ is any auxiliary space.

We remark that $\mathcal{A}$ can be assumed to be $\mathcal{H}$ without loss of generality due to a similar reason for the original diamond norm (see, for example, Watrous [2018, Chap 3]).

We argue that (Q, λ) -diamond norm is a seminorm in the appendix of the extended version [Hung et al. 2018]. Intuitively, this is conceivable since we only restrict the input state from all density operators to a convex subset satisfying $\text{tr}(Q\rho) \geq \lambda$. Note that, by definition, $\|\cdot\|_{\diamond} \equiv \|\cdot\|_{I, \lambda}$ for any $0 \leq \lambda \leq 1$.

⁶For example, consider a 2-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ and two superoperators $\mathcal{E}, \mathcal{E}'$ over $\mathcal{H}$ with $\mathcal{E}(\rho) = \frac{1}{3}\text{tr}(\rho)I_{\mathcal{H}} + \frac{1}{3}\rho^T$ and $\mathcal{E}'(\rho) = \text{tr}(\rho)I_{\mathcal{H}} - \rho^T$, where ρ^T is the transpose of ρ . Without introducing an auxiliary space, direct calculation gives $T(\mathcal{E}(\rho), \mathcal{E}'(\rho)) = \frac{1}{3}\|I - 2\rho\|_1 \leq \frac{2}{3}$ for any single qubit state ρ . However, if we use an auxiliary qubit and apply $\mathcal{E}$ to the maximally entangled state $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, we have $T(\mathcal{E}(|\phi^+\rangle\langle\phi^+|), \mathcal{E}'(|\phi^+\rangle\langle\phi^+|)) = 1$.

Let $\mathcal{E}$ and $\mathcal{E}'$ be superoperators over $\mathcal{H}$. By extending Watrous [2009], we show that $\|\mathcal{E} - \mathcal{E}'\|_{Q,\lambda}$ can be efficiently computed by the following semidefinite program (SDP):

$$\max \quad \text{tr}(J(\Phi)W) \quad (6.5)$$

$$\text{s.t.} \quad W \leq I_{\mathcal{H}} \otimes \rho, \text{tr}(Q\rho) \geq \lambda, \quad (6.6)$$

$$\rho \in \mathcal{D}(\mathcal{H}), W \text{ is a positive semidefinite operator over } \mathcal{H} \otimes \mathcal{H}, \quad (6.7)$$

where $\Phi = \mathcal{E} - \mathcal{E}'$ and $J(\Phi)$ is the Choi-Jamiołkowski representation of Φ . Note that the above SDP is identical to the SDP used in Watrous [2009, Section 4] to compute $\|\mathcal{E} - \mathcal{E}'\|_{\diamond}$, except that we have added the additional constraint $\text{tr}(Q\rho) \geq \lambda$ to capture the requirement on input states. The correctness of the above SDP then basically follows from the analysis of Watrous [2009] and the definition of (Q, λ) -diamond norm.

The standard diamond norm and (Q, λ) -diamond norm can be significantly different for the same pair of superoperators $\mathcal{E}, \mathcal{E}'$. For example, consider $\mathcal{E} = H \circ H$ and $\mathcal{E}' = HZ \circ ZH$. We can show⁷ that $\|\mathcal{E} - \mathcal{E}'\|_{\diamond} = 1$, whereas $\|\mathcal{E} - \mathcal{E}'\|_{|0\rangle\langle 0|, \frac{3}{4}} = \sqrt{3}/2$. Thus, the (Q, λ) -diamond norm can help us leverage prior knowledge about input states to obtain more accurate bounds.

Using the (Q, λ) -diamond norm, we define a notion of quantum robustness as follows.

DEFINITION 6.2 (QUANTUM ROBUSTNESS). *The noisy program $\tilde{P}$ (over $\mathcal{H}$, having ideal program $P = \text{ideal}(\tilde{P})$) is ϵ -robust under (Q, λ) if and only if*

$$\|[\tilde{P}] - [P]\|_{Q,\lambda} \leq \epsilon. \quad (6.8)$$

Here, Q is a quantum predicate over $\mathcal{H}$ and $0 \leq \lambda, \epsilon \leq 1$.

By the definition of the (Q, λ) -diamond norm and the trace distance, (6.8) can be equivalently stated as the following.

$$\forall \rho \in \mathcal{D}(\mathcal{H} \otimes \mathcal{H}), \text{tr}(Q\rho) \geq \lambda \text{tr}(\rho) \Rightarrow \frac{1}{2} \|[\tilde{P}] \otimes I_{\mathcal{H}}(\rho) - [P] \otimes I_{\mathcal{H}}(\rho)\|_1 \leq \epsilon \text{tr}(\rho). \quad (6.9)$$

Since ϵ measures the distance between $[\tilde{P}]$ and $[P]$, the smaller ϵ is, the closer the noisy program $\tilde{P}$ is to the ideal program P . One could think of ϵ as measuring both the probability that noise can happen and intensity of that noise. When the noise is strong, a noisy program $\tilde{P}$ being ϵ -robust implies that the probability of noise is at most ϵ . When the noise is weak, it could occur with greater probability, but its effect will be much smaller.

Intuitively, the use of precondition Q can help us obtain more accurate bounds. For example, one can use Q to characterize prior information about the input state due to the nature of underlying physical systems. Even without any prior knowledge about the input state, preconditions can still be leveraged for different branches of the program in case statements and loops.

6.3 Logic for Quantum Robustness

A program's robustness can be proved by working out the (denotational) semantics of programs P and $\tilde{P}$ and applying Definition 6.2 directly. However, this computation may be difficult. As an

⁷For the normal diamond norm, consider the input state to be $\rho = |+\rangle\langle +| \otimes \sigma$ for some ancilla state σ . It is easy to see that $\mathcal{E}(\rho) = |0\rangle\langle 0|$ and $\mathcal{E}'(\rho) = |1\rangle\langle 1|$, which are perfectly distinguishable. For the $(|0\rangle\langle 0|, \frac{3}{4})$ -diamond norm, without loss of generality, consider any input state $|\psi\rangle = \cos \theta |0\rangle|\psi_0\rangle + \sin \theta |1\rangle|\psi_1\rangle$ where $\theta \in [0, \frac{\pi}{2}]$. Requiring $|\psi\rangle$ to satisfy $|0\rangle\langle 0|$ to degree at least $\frac{3}{4}$, we have $\cos^2 \theta \geq \frac{3}{4}$ and therefore $\theta \in [0, \frac{\pi}{6}]$. Simple calculation gives $H|\psi\rangle = \cos \theta |+\rangle|\psi_0\rangle + \sin \theta |-\rangle|\psi_1\rangle$ and $HZ|\psi\rangle = \cos \theta |+\rangle|\psi_0\rangle - \sin \theta |-\rangle|\psi_1\rangle$. The projector that maximally distinguishes these states is $|0\rangle\langle 0| \otimes I$, so we have

$$\|\mathcal{E} - \mathcal{E}'\|_{|0\rangle\langle 0|, 3/4} = \frac{1}{2} ((\cos \theta + \sin \theta)^2 - (\cos \theta - \sin \theta)^2) = \sin 2\theta \leq \frac{\sqrt{3}}{2},$$

the equality of which holds when $\theta = \pi/6$.

$$\begin{array}{c}
\frac{}{(Q, \lambda) \vdash \mathbf{skip} \leq 0} \text{ (Skip)} \quad \frac{}{(Q, \lambda) \vdash (q := |0\rangle) \leq 0} \text{ (Init)} \\
\frac{\|U \circ U^\dagger - \Phi\|_{Q, \lambda} \leq \epsilon}{(Q, \lambda) \vdash (\bar{q} :=_{p, \Phi} U[\bar{q}]) \leq p\epsilon} \text{ (Unitary)} \\
\frac{(Q', \lambda') \vdash \tilde{P} \leq \epsilon' \quad \epsilon' \leq \epsilon \quad Q \sqsubseteq Q' \quad \lambda' \leq \lambda}{(Q, \lambda) \vdash \tilde{P} \leq \epsilon} \text{ (Weaken)} \\
\frac{(Q/\delta, \lambda/\delta) \vdash \tilde{P} \leq \epsilon \quad 0 \sqsubseteq Q, Q/\delta \sqsubseteq I \quad 0 \leq \lambda, \lambda/\delta \leq 1}{(Q, \lambda) \vdash \tilde{P} \leq \epsilon} \text{ (Rescale)} \\
\frac{(Q_1, \lambda) \vdash \tilde{P}_1 \leq \epsilon_1 \quad (Q_2, \lambda) \vdash \tilde{P}_2 \leq \epsilon_2 \quad \{Q_1\}P_1\{Q_2\}}{(Q_1, \lambda) \vdash (\tilde{P}_1; \tilde{P}_2) \leq \epsilon_1 + \epsilon_2} \text{ (Sequence)} \\
\frac{\forall m, (Q_m, 1 - \delta) \vdash \tilde{P}_m \leq \epsilon \quad t, \delta \in [0, 1]}{(\sum_m M_m^\dagger Q_m M_m, 1 - t\delta) \vdash (\mathbf{case } M[\bar{q}] = m \rightarrow \tilde{P}_m \mathbf{end}) \leq (1 - t)\epsilon + t} \text{ (Case)} \\
\frac{\begin{array}{c} (Q, \lambda) \vdash \tilde{P}_1 \leq \epsilon \quad \{Q\}P_1\{\lambda M_0^\dagger M_0 + M_1^\dagger Q M_1\} \\ \tilde{P} \equiv \mathbf{while } M[\bar{q}] = 1 \mathbf{do } \tilde{P}_1 \mathbf{done} \quad P \text{ is } (a, n)\text{-bounded} \end{array}}{(\lambda M_0^\dagger M_0 + M_1^\dagger Q M_1, \lambda) \vdash \tilde{P} \leq n\epsilon/(1 - a)} \text{ (While-Bounded)} \\
\frac{}{(Q, \lambda) \vdash (\mathbf{while } M[\bar{q}] = 1 \mathbf{do } \tilde{P}_1 \mathbf{done}) \leq 1} \text{ (While-Unbounded)}
\end{array}$$

Fig. 4. Rules for logic of quantum robustness.

alternative, we present a logic for proving judgments of the form $(Q, \lambda) \vdash \tilde{P} \leq \epsilon$, meaning that program $\tilde{P}$ is ϵ -robust under (Q, λ) . In [Section 6.4](#) we prove the logic is sound. In [Section 7](#) we demonstrate the use of the logic, alongside direct proofs of robustness for some cases (e.g., error correction).

The rules for our logic are given in [Figure 4](#).

6.3.1 Simple Rules. The *Skip* and *Init* rules say that the skip and initialization operations are always error-free. These operations will not increase the distance between $\llbracket P \rrbracket$ and $\llbracket \tilde{P} \rrbracket$. The *Unitary* rule says that if we can bound the (Q, λ) -diamond norm between the intended operation U and the noise operation Φ by ϵ , then we can bound the total distance by $p\epsilon$. The *Weaken* rule says that we can always safely make the precondition more restrictive, increase the degree to which an input state must satisfy the predicate, or increase the upper bound on the distance between the noisy and ideal programs. The *Rescale* rule says that equivalent forms of our judgment can be obtained by rescaling Q and λ . Note that the *Rescale* rule does not weaken the judgment, but rather provides some flexibility in choosing Q, λ compatible with other rules; one can scale by δ so long as Q/δ and λ/δ are still well defined. The *Sequence* rule allows us to compose two judgments by summing their computed upper bounds. Note that in the *Sequence* and *While-Bounded* rules we define Hoare triples as in [Section 4.4](#).

6.3.2 The Case Rule. The *Case* rule says that, given appropriate bounds for every branch of a case statement, we can bound the error of the entire case statement. Note that $\sum_m M_m^\dagger Q_m M_m$ is the weakest precondition of the case construct in quantum Hoare logic [[Ying 2016](#)]. In a logic for classical programs, one might expect each branch of a case statement to satisfy the precondition

perfectly. However, in a quantum logic, as we will discuss in the soundness proof (see [Section 6.4](#)), this is not necessarily true. To see this, note that in our rule we start with the precondition $\sum_m M_m^\dagger Q_m M_m$ and $\lambda = 1 - t\delta$ on the input state to the case statement, but we can only guarantee that a weighted fraction of $1 - t$ of the branches satisfy a weaker precondition Q_m and $\lambda' = 1 - \delta$ for some choice of $t \in [0, 1]$. (Note that $1 - \delta \leq 1 - t\delta$ for $t, \delta \in [0, 1]$.)

When applying this rule, one can make $1 - \delta$ and ϵ the same for every $\widetilde{P}_m$ by applying the *Weaken* and/or *Rescale* rules. The choice of t will represent a tradeoff between a lower error bound and a more restrictive requirement on the satisfaction of the predicate (see [Example 6.2](#)).

EXAMPLE 6.1 (SIMPLE CASE STATEMENT). *Consider the following program.*

$$\begin{aligned} \widetilde{P} &\equiv \text{case } M[\bar{q}] = 0 \rightarrow \bar{q} : \cong_{p,\Phi} H[\bar{q}] \\ &\quad 1 \rightarrow \text{skip} \\ &\text{end} \end{aligned}$$

This program performs measurement in the $\{|0\rangle, |1\rangle\}$ basis and either applies a noisy Hadamard H gate or does nothing depending on the measurement outcome. In order to apply the Case rule, we must bound the error of both branches. By the Skip rule, the second branch will have zero error, i.e., $(Q, \lambda) \vdash \text{skip} \leq 0$ for any Q and λ . We will choose $Q = I$ and $\lambda = 1$. By the Unitary rule, the error of the first branch will depend on Φ .

Consider an error model given by $\Phi(\rho) = HZ\rho ZH$. This means that with probability $1 - p$, $\bar{q} : \cong_{p,\Phi} H[\bar{q}]$ applies an H gate, and with probability p it applies a Z gate followed by an H gate. Recall that $Z|0\rangle\langle 0|Z = |0\rangle\langle 0|$. This means that a Z error will not affect the state in the first branch (because the first branch corresponds to having measured a 0). So we have that $(|0\rangle\langle 0|, 1) \vdash (\bar{q} : \cong_{p,\Phi} H[\bar{q}]) \leq 0$, which says that if the input state is the $|0\rangle$ state, then with the error model defined, there will be no error during an application of H (i.e. $\epsilon = 0$). Note that without the precondition $|0\rangle\langle 0|$, we would have $\epsilon > 0$.

Given $(|0\rangle\langle 0|, 1) \vdash (\bar{q} : \cong_{p,\Phi} H[\bar{q}]) \leq 0$ and $(I, 1) \vdash \text{skip} \leq 0$, we can conclude that $(I, 1) \vdash \widetilde{P} \leq t$ by the Case rule with the following simplification,

$$M_0^\dagger Q_0 M_0 + M_1^\dagger Q_1 M_1 = |0\rangle\langle 0||0\rangle\langle 0| + |1\rangle\langle 1|I|1\rangle\langle 1| = |0\rangle\langle 0| + |1\rangle\langle 1| = I.$$

Because $\delta = 0$, we can choose $t = 0$ to get an overall error bound of zero. However, in general, the choice of t will represent a tradeoff between a lower error bound and a more restrictive requirement on the satisfaction of the predicate (i.e. a larger λ value).

Note that the precondition for each branch does not need to directly relate to the measurement basis. For example, consider instead the following program.

$$\begin{aligned} &\text{case } M[\bar{q}] = + \rightarrow \bar{q} : \cong_{p,\Phi} H[\bar{q}] \\ &\quad - \rightarrow \text{skip} \\ &\text{end} \end{aligned}$$

This program performs measurement in the $\{|+\rangle, |-\rangle\}$ basis. For this program, we can still use the judgments $(|0\rangle\langle 0|, 1) \vdash (\bar{q} : \cong_{p,\Phi} H[\bar{q}]) \leq 0$ and $(I, 1) \vdash \text{skip} \leq 0$ described in the previous example, but now our conclusion will be $(\frac{1}{2}|+\rangle\langle +| + |-\rangle\langle -|, 1) \vdash \widetilde{P} \leq t$, by the following simplification,

$$M_0^\dagger Q_0 M_0 + M_1^\dagger Q_1 M_1 = |+\rangle\langle +| + |0\rangle\langle 0||+\rangle\langle +| + |-\rangle\langle -||-\rangle\langle -| = \frac{1}{2}|+\rangle\langle +| + |-\rangle\langle -|.$$

Note that $\frac{1}{2}|+\rangle\langle +| + |-\rangle\langle -| \sqsubseteq I$, so we have restricted the set of states for which this error bound will hold. This may be useful in situations where it is difficult to compute an error bound given

certain preconditions. For example, it may be difficult to show that $(|+\rangle\langle+|, 1) \vdash (\bar{q} : \cong_{p,\Phi} H[\bar{q}]) \leq \epsilon$ for a sufficiently low ϵ , but it is easy to show that $(|0\rangle\langle 0|, 1) \vdash (\bar{q} : \cong_{p,\Phi} H[\bar{q}]) \leq 0$.

EXAMPLE 6.2 (T-VALUE TRADEOFFS). *Consider the following program.*

$$\begin{aligned} \tilde{P} &\equiv \text{case } M[\bar{q}] = 0 \rightarrow \tilde{P}_0 \\ &\quad 1 \rightarrow \tilde{P}_1 \\ &\text{end} \end{aligned}$$

Say that we have that $(Q_0, 1 - 0.05) \vdash \tilde{P}_0 \leq 0.01$ and $(Q_1, 1 - 0.25) \vdash \tilde{P}_1 \leq 0.1$ for programs $\tilde{P}_0$ and $\tilde{P}_1$. Now we can use the Case rule to conclude that $(M_0^\dagger Q_0 M_0 + M_1^\dagger Q_1 M_1, 1 - t\delta) \vdash \tilde{P} \leq (1 - t)\epsilon + t$ for $\delta = \min(0.05, 0.25) = 0.05$, $\epsilon = \max(0.01, 0.1) = 0.1$, and $t \in [0, 1]$.

If we choose $t = 0$ then we have the lowest possible error bound, but $1 - t\delta = 1$, so the produced error bound will only apply to states that completely satisfy the predicate $M_0^\dagger Q_0 M_0 + M_1^\dagger Q_1 M_1$.⁸ If we choose $t = 1$ then we have the least restrictive condition on input states, but $(1 - t)\epsilon + t = 1$, which is the trivial error bound. By choosing t to be between 0 and 1 we can trade between a low error bound and a less restrictive constraint on the input state. For example, for $t = 0.25$ we have that $1 - t\delta = 0.9875$ and $(1 - t)\epsilon + t = 0.325$. For $t = 0.5$ we have that $1 - t\delta = 0.975$ and $(1 - t)\epsilon + t = 0.55$.

6.3.3 *The Loop Rule.* Finally, the *While-Bounded* and *While-Unbounded* rules allow us to compute an upper bound for the distance between the noisy and ideal versions of a while loop. The *While-Unbounded* rule is a trivial bound on the distance. The *While-Bounded* rule, however, demonstrates a non-trivial upper bound with an assumption called (a, n) -boundedness. Such an assumption is necessary for us to get around the potential issue of termination and to be able to reason about interesting programs in our case study.

Intuitively, (a, n) -boundedness is a condition on how fast the *ideal* loop will converge, which is inherent to the control flow of the program and does not depend on any specific error model. We view this as an advantage as we do not need a new analysis for every possible noise model.

DEFINITION 6.3 ((a, n) -BOUNDEDNESS). *A while loop $P \equiv \text{while } M[q] = 1 \text{ do } P_1 \text{ done}$ is said to be (a, n) -bounded for $0 \leq a < 1$ and integer $n \geq 1$ if*

$$(\mathcal{E}^*)^n(M_1^\dagger M_1) \subseteq aM_1^\dagger M_1 \quad (6.10)$$

where the linear map $\mathcal{E}(\rho)$ is defined as $\llbracket P_1 \rrbracket(M_1 \rho M_1^\dagger)$ and $\mathcal{E}^*$ is the dual map of $\mathcal{E}$.⁹

Intuitively speaking, a loop is (a, n) -bounded if, for every state ρ , after n iterations it is guaranteed that at least a $(1 - a)$ -fraction of the state has exited the loop. A while loop with this nice property is guaranteed to terminate with probability 1 on all input states, which helps avoid the termination issue. As we will show in the examples that follow and in Section 7, specific (a, n) can be derived analytically or numerically for concrete programs.¹⁰ We also remark that by assuming (a, n) -boundedness, we avoid weakening λ like in the *Case* rule.

In Carbin et al. [2013], loops are assumed to have a bounded number of iterations or a trivial upper bound will be used (i.e., our rule *While-Unbounded*). Because of the use of (a, n) -boundedness, we can handle more complicated loops. One also has the freedom to choose appropriate values of Q for different purposes. A simple choice is $Q = \lambda I$. A less trivial choice of Q is shown in the following example.

⁸This might even be impossible when $M_0^\dagger Q_0 M_0 + M_1^\dagger Q_1 M_1 \sqsubset I$. In that case, the choice $t = 0$ becomes useless.

⁹If $\llbracket P_1 \rrbracket$ can be written as $\sum_k F_k \circ F_k^\dagger$ for some set of Kraus operators $\{F_k\}_k$, the Kraus form of $\mathcal{E}^*$ is $\sum_k M_1^\dagger F_k^\dagger \circ F_k M_1$.

¹⁰The rough idea is to guess (or enumerate) n and prove a either analytically or numerically (with a simple SDP).

EXAMPLE 6.3 (SLOW STATE PREPARATION). *In this example, we consider the following program which prepares the standard basis state $|1\rangle$:*

$$SSP \equiv q := |0\rangle; \text{while } M[q] = 0 \text{ do } q := H[q]; q := I[q] \text{ done}, \quad (6.11)$$

where M is the standard basis measurement $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. We consider the case where there is a bit flip error with probability 0.01 when applying the I gate, i.e., the ideal and the noisy loop bodies are $P_1 \equiv q := H[q]; q := I[q]$; and $\tilde{P}_1 \equiv q := H[q]; q := \approx_{0.01, X} I[q]$. Then $\llbracket P_1 \rrbracket = H \circ H$ and $\llbracket \tilde{P}_1 \rrbracket = 0.99H \circ H + 0.01XH \circ HX$. Consider $Q = |0\rangle\langle 0|$ and $\lambda = 1$. Since $\llbracket P_1 \rrbracket(|0\rangle\langle 0|) = \llbracket \tilde{P}_1 \rrbracket(|0\rangle\langle 0|) = |+\rangle\langle +|$, we have that $\|\llbracket P_1 \rrbracket - \llbracket \tilde{P}_1 \rrbracket\|_{|0\rangle\langle 0|, 1} = 0$, and by the Unitary rule, $(|0\rangle\langle 0|, 1) \vdash \tilde{P}_1 \leq 0$.

We can use this choice of Q and λ when applying the While-Bounded rule. First, note that the statement $\{Q\}P_1\{\lambda M_0^\dagger M_0 + M_1 Q M_1^\dagger\}$ holds because $\lambda M_0^\dagger M_0 + M_1 Q M_1^\dagger = I$ for our choice of M_0, M_1, Q, λ . Next, we need to show (a, n) -boundedness of the loop. This requires us to consider the behavior of $\mathcal{E}^*$ where $\mathcal{E}^*$ is the dual of $\mathcal{E} = HM_1 \circ M_1 H$. We claim that the while loop is $(1/2, 1)$ -bounded because

$$\mathcal{E}^*(M_1^\dagger M_1) = |0\rangle\langle 0|H|0\rangle\langle 0|H|0\rangle\langle 0| = \frac{1}{2}|0\rangle\langle 0| = \frac{1}{2}M_1^\dagger M_1.$$

Now by the While-Bounded rule, we have that $(I, 1) \vdash \widetilde{SSP} \leq 0$, i.e., the program is perfectly robust.

In Example 6.3, if we use the precondition I in our judgment of the robustness of the loop body, the best upper bound we can argue is $\epsilon = 0.01$, i.e., $(I, 1) \vdash \tilde{P}_1 \leq 0.01$. Then, applying the While-Bounded rule yields $(I, 1) \vdash \widetilde{SSP} \leq 0.02$ since $\frac{n\epsilon}{1-a} = 2\epsilon = 0.02$. Therefore, restricting the state space with the predicate $Q = |0\rangle\langle 0|$, as we did in Example 6.3, was a better choice, as it yielded a better bound. Moreover, this choice of Q is natural since the post-measurement state entering the loop satisfies Q perfectly. We note that the program SSP might look contrived for preparing $|1\rangle$ when compared to the more straightforward program $SP \equiv q := |0\rangle; q := X[q]$. We argue that SSP might be preferred over SP in the presence of noise. Consider the same noise model as above. Namely, let $\tilde{SP} \equiv q := |0\rangle; q := X[q]; q := \approx_{0.01, X} I[q]$. We have shown that $(I, 1) \vdash \widetilde{SSP} \leq 0$ given this noise model, which says that $\widetilde{SSP}$ is perfectly robust. By directly applying Definition 6.2, we can show that $\tilde{SP}$ is 0.01-robust given the same noise model,¹¹ which suggests that $\tilde{SP}$ is less desirable in this situation.

6.4 Soundness

In this section, we show that logic given in Figure 4 is sound.

THEOREM 6.4 (SOUNDNESS). *If $(Q, \lambda) \vdash \tilde{P} \leq \epsilon$ then $\tilde{P}$ is ϵ -robust under (Q, λ) .*

PROOF. The proof proceeds by induction on the derivation $(Q, \lambda) \vdash \tilde{P} \leq \epsilon$. We will work primarily from definition of robustness given in (6.9). We note that superoperators $\llbracket \tilde{P} \rrbracket$ and $\llbracket P \rrbracket$ apply on the same space. By the definition of the (Q, λ) -diamond norm, we need to consider $\llbracket \tilde{P} \rrbracket \otimes I$ and $\llbracket P \rrbracket \otimes I$. However, to simplify the presentation, we will omit “ $\otimes I$ ” in the following proof whenever there is no ambiguity.

Now we consider each possible rule used in the final step of the derivation.

- (1) *Skip*: This rule holds by observing $\llbracket \widetilde{\text{skip}} \rrbracket = \llbracket \text{skip} \rrbracket$, i.e., they refer to the same superoperator. Thus, any (Q, λ) -diamond norm between them is 0. We choose $Q = I$ and $\lambda = 0$.
- (2) *Init*: for the same reason as the proof for *Skip*.

¹¹ Note that $\llbracket \tilde{SP} \rrbracket = 0.99X \circ X + 0.01I \circ I$, and hence we have $\|\llbracket SP \rrbracket - \llbracket \tilde{SP} \rrbracket\|_{I, 1} = 0.01\|X \circ X - I \circ I\|_{I, 1} = 0.01$.

(3) *Unitary*: For every state ρ satisfying $\text{tr}(Q\rho) \geq \lambda$,

$$\begin{aligned} \frac{1}{2} \|\llbracket \tilde{q} \rrbracket \rho - \llbracket \tilde{q} := U[\tilde{q}] \rrbracket \rho\|_1 &= \frac{1}{2} \|(1-p)U\rho U^\dagger + p\Phi(\rho) - U\rho U^\dagger\|_1 \\ &= \frac{p}{2} \|U\rho U^\dagger - \Phi(\rho)\|_1 \leq p\|U \cdot U^\dagger - \Phi\|_{Q,\lambda} \leq p\epsilon. \end{aligned}$$

The second from last inequality holds by the definition of the (Q, λ) -diamond norm and the last inequality follows from the premise.

(4) *Weaken*: By induction, the premise $(Q', \lambda') \vdash \tilde{P} \leq \epsilon'$ implies

$$\forall \rho, \text{tr}(Q'\rho) \geq \lambda' \text{tr}(\rho) \Rightarrow \frac{1}{2} \|\llbracket \tilde{P} \rrbracket \rho - \llbracket P \rrbracket \rho\|_1 \leq \epsilon' \text{tr}(\rho).$$

For any density matrix ρ , constants $0 \leq \lambda' \leq \lambda \leq 1$, and predicates $Q \sqsubseteq Q'$, $\text{tr}(Q\rho) \geq \lambda \text{tr}(\rho)$ implies that $\text{tr}(Q'\rho) \geq \lambda' \text{tr}(\rho)$. And for $0 \leq \epsilon' \leq \epsilon \leq 1$, $\frac{1}{2} \|\llbracket \tilde{P} \rrbracket \rho - \llbracket P \rrbracket \rho\|_1 \leq \epsilon' \text{tr}(\rho)$ implies that $\frac{1}{2} \|\llbracket \tilde{P} \rrbracket \rho - \llbracket P \rrbracket \rho\|_1 \leq \epsilon \text{tr}(\rho)$. Therefore, we have that

$$\forall \rho, \text{tr}(Q\rho) \geq \lambda \text{tr}(\rho) \Rightarrow \frac{1}{2} \|\llbracket \tilde{P} \rrbracket \rho - \llbracket P \rrbracket \rho\|_1 \leq \epsilon \text{tr}(\rho).$$

So $\tilde{P}$ is ϵ -robust under (Q, λ) .

(5) *Rescale*: This rule follows by observing that the condition $\text{tr}(\rho Q) \geq \lambda$ is equivalent to the condition $\text{tr}(\rho Q/\delta) \geq \lambda/\delta$ for $\delta > 0$. We only require that $Q, Q/\delta$ and $\lambda, \lambda/\delta$ are well defined, namely, $0 \sqsubseteq Q, Q/\delta \sqsubseteq I$ and $0 \leq \lambda, \lambda/\delta \leq 1$.

(6) *Sequence*: For every state ρ ,

$$\begin{aligned} \|\llbracket \tilde{P}_1; \tilde{P}_2 \rrbracket \rho - \llbracket P_1; P_2 \rrbracket \rho\|_1 &= \|\llbracket \tilde{P}_2 \rrbracket \llbracket \tilde{P}_1 \rrbracket \rho - \llbracket P_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1 \\ &\leq \|\llbracket \tilde{P}_2 \rrbracket \llbracket \tilde{P}_1 \rrbracket \rho - \llbracket \tilde{P}_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1 + \|\llbracket \tilde{P}_2 \rrbracket \llbracket P_1 \rrbracket \rho - \llbracket P_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1 \\ &\leq \|\llbracket \tilde{P}_1 \rrbracket \rho - \llbracket P_1 \rrbracket \rho\|_1 + \|\llbracket \tilde{P}_2 \rrbracket \llbracket P_1 \rrbracket \rho - \llbracket P_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1. \end{aligned}$$

The inequality $\|\llbracket \tilde{P}_2 \rrbracket \llbracket \tilde{P}_1 \rrbracket \rho - \llbracket \tilde{P}_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1 \leq \|\llbracket \tilde{P}_1 \rrbracket \rho - \llbracket P_1 \rrbracket \rho\|_1$ follows because quantum superoperators are contractive.

Now assume that $\text{tr}(Q_1\rho) \geq \lambda \text{tr}(\rho)$. By induction, the premise $(Q_1, \lambda) \vdash \tilde{P}_1 \leq \epsilon_1$ implies that $\frac{1}{2} \|\llbracket \tilde{P}_1 \rrbracket \rho - \llbracket P_1 \rrbracket \rho\|_1 \leq \epsilon_1 \text{tr}(\rho)$. Also, by the premise $\{Q_1\}P_1\{Q_2\}$, we have that $\text{tr}(Q_2 \llbracket P_1 \rrbracket \rho) \geq \text{tr}(Q_1\rho) \geq \lambda \text{tr}(\rho) \geq \lambda \text{tr}(\llbracket P_1 \rrbracket \rho)$. Now we can use our induction hypothesis and the premise $(Q_2, \lambda) \vdash \tilde{P}_2 \leq \epsilon_2$ to conclude $\frac{1}{2} \|\llbracket \tilde{P}_2 \rrbracket \llbracket P_1 \rrbracket \rho - \llbracket P_2 \rrbracket \llbracket P_1 \rrbracket \rho\|_1 \leq \epsilon_2 \text{tr}(\llbracket P_1 \rrbracket \rho)$. So, finally, we have that

$$\frac{1}{2} \|\llbracket \tilde{P}_1; \tilde{P}_2 \rrbracket \rho - \llbracket P_1; P_2 \rrbracket \rho\|_1 \leq \epsilon_1 \text{tr}(\rho) + \epsilon_2 \text{tr}(\llbracket P_1 \rrbracket \rho) \leq (\epsilon_1 + \epsilon_2) \text{tr}(\rho).$$

(7) *Case*: Let $\tilde{P} = \text{case } M[\tilde{q}] = \overline{m \rightarrow \tilde{P}_m} \text{ end}$. Assume the input state ρ to the case statement satisfies $\sum_m M_m^\dagger Q_m M_m$ to degree λ' , i.e., $\sum_m \text{tr}(M_m^\dagger Q_m M_m \rho) \geq \lambda' \text{tr}(\rho)$. To leverage the premise $(Q_m, \lambda) \vdash \tilde{P}_m \leq \epsilon$ and the induction hypothesis to conclude that $\frac{1}{2} \|\llbracket \tilde{P}_m \rrbracket \rho - \llbracket P_m \rrbracket \rho\|_1 \leq \epsilon \text{tr}(\rho)$, one must show the precondition (Q_m, λ) holds for state ρ entering branch m . A naive approach is to show that $\text{tr}(M_m^\dagger Q_m M_m \rho) \geq \lambda \text{tr}(M_m^\dagger M_m \rho)$ holds for every branch m , which implies that $\text{tr}(Q_m M_m \rho M_m^\dagger) \geq \lambda \text{tr}(M_m \rho M_m^\dagger)$ where $M_m \rho M_m^\dagger$ is the (sub-normalized) post-measurement state entering branch m . We argue that this is in general impossible when $\lambda' = \lambda$. For instance, consider a collection of projective measurement operators $\{M_m\}_m$. If there exists a branch i and a state ρ supported on M_i such that $\text{tr}(M_i^\dagger Q_i M_i \rho) \geq \lambda \text{tr}(\rho)$ for some $\lambda > 0$, obviously $\sum_m M_m^\dagger Q_m M_m \rho \geq \lambda \text{tr}(\rho)$, but none of the preconditions is satisfied except for the one for branch i since $\text{tr}(M_j^\dagger Q_j M_j \rho) = 0$ for each $j \neq i$.

Instead, we show that for a majority of the clauses $\text{tr}(M_m^\dagger Q_m M_m \rho) \geq \lambda \text{tr}(M_m^\dagger M_m \rho)$ holds for some λ strictly less than λ' . To that end, let $p_m = \text{tr}(M_m^\dagger M_m \rho)$ and $q_m = \text{tr}(M_m^\dagger Q_m M_m \rho)$. Define δ_m to be such that $q_m = (1 - \delta_m)p_m$. Note that $0 \leq \delta_m \leq 1$ because $0 \leq q_m \leq p_m$ for every m . Without loss of generality we assume $\text{tr}(\rho) = 1$. Let $S(\rho)$ denote the collection of branches such that the precondition (Q_m, λ) holds, i.e., $S(\rho) = \{m : q_m \geq \lambda p_m, \text{ i.e., } \delta_m \leq 1 - \lambda\}$. We will determine a lower bound for $\sum_{m \in S(\rho)} p_m$ for each state ρ using a probabilistic argument.

First, note that $\{p_m\}$ is a probability distribution since $\sum_m p_m = 1$ and $p_m \geq 0$ for each m . Also, since (by our assumption) $\sum_m q_m \geq \lambda' \sum_m p_m = \lambda'$, we have that $\sum_m \delta_m p_m \leq (1 - \lambda')$. Now we can define a random variable Δ to be such that $\Pr[\Delta = \delta_m] = p_m$ for each m . The expected value of Δ is $\mathbb{E}[\Delta] \leq (1 - \lambda')$, and Markov's inequality yields

$$\sum_{m \notin S(\rho)} p_m = \Pr[\Delta \geq 1 - \lambda] \leq \frac{\mathbb{E}[\Delta]}{1 - \lambda} \leq \frac{1 - \lambda'}{1 - \lambda} =: t. \quad (6.12)$$

This says that a weighted fraction t of the branches will not satisfy the precondition (Q_m, λ) . Note that $t \in [0, 1]$ since $\lambda \leq \lambda'$. For $m \notin S(\rho)$, only the trivial upper bound $\epsilon_m = 1$ is guaranteed. Therefore, for each state ρ ,

$$\frac{1}{2} \|\llbracket \widetilde{P} \rrbracket \rho - \llbracket P \rrbracket \rho\|_1 = \frac{1}{2} \left\| \sum_m (\llbracket \widetilde{P}_m \rrbracket (M_m \rho M_m^\dagger) - \llbracket P_m \rrbracket (M_m \rho M_m^\dagger)) \right\|_1 \quad (6.13)$$

$$\leq \frac{1}{2} \sum_m \|\llbracket \widetilde{P}_m \rrbracket (M_m \rho M_m^\dagger) - \llbracket P_m \rrbracket (M_m \rho M_m^\dagger)\|_1 \quad (6.14)$$

$$\leq \sum_{m \in S(\rho)} \text{tr}(M_m \rho M_m^\dagger) \epsilon + \sum_{m \notin S(\rho)} \text{tr}(M_m \rho M_m^\dagger) \quad (6.15)$$

$$\leq (1 - t)\epsilon + t = ((1 - t)\epsilon + t)\text{tr}(\rho). \quad (6.16)$$

Rewriting $\lambda = 1 - \delta$ for ease of notation, we have $\lambda' = 1 - t\delta$. Finally, we note that the case $t = 0$ implies that the precondition of each branch is satisfied, and the error of the case statement can be bounded by ϵ .

- (8) *While-Bounded:* Let $P \equiv \mathbf{while} \ M[q] = 1 \ \mathbf{do} \ P_1 \ \mathbf{done}$. Let S_k be the bounded while loop of k iterations. Define the linear maps $\mathcal{E}(\rho) := \llbracket P_1 \rrbracket (M_1 \rho M_1^\dagger)$ and $\widetilde{\mathcal{E}}(\rho) := \llbracket \widetilde{P}_1 \rrbracket (M_1 \rho M_1^\dagger)$ and let $\llbracket S_k \rrbracket \rho = M_0 \rho M_0^\dagger + \llbracket S_{k-1} \rrbracket (\mathcal{E}(\rho))$ for $k \geq 1$ (with $\llbracket S_0 \rrbracket \rho = \rho$). In order to bound the distance between $\llbracket P \rrbracket$ and $\llbracket \widetilde{P} \rrbracket$, we first upper bound the distance between $\llbracket S_k \rrbracket$ and $\llbracket \widetilde{S}_k \rrbracket$ and then take the limit as $k \rightarrow \infty$. We then have

$$\frac{1}{2} \|\llbracket \widetilde{S}_k \rrbracket \rho - \llbracket S_k \rrbracket \rho\|_1 \leq \frac{1}{2} \|\llbracket \widetilde{S}_{k-1} \rrbracket (\widetilde{\mathcal{E}}(\rho)) - \llbracket S_{k-1} \rrbracket (\mathcal{E}(\rho))\|_1 \quad (6.17)$$

$$\leq \frac{1}{2} \|\widetilde{\mathcal{E}}(\rho) - \mathcal{E}(\rho)\|_1 + \frac{1}{2} \|\llbracket \widetilde{S}_{k-1} \rrbracket (\mathcal{E}(\rho)) - \llbracket S_{k-1} \rrbracket (\mathcal{E}(\rho))\|_1 \quad (6.18)$$

$$\leq \frac{1}{2} \sum_{i=0}^{k-1} \|\widetilde{\mathcal{E}}(\mathcal{E}^i(\rho)) - \mathcal{E}^{i+1}(\rho)\| \quad (6.19)$$

$$\leq \epsilon \sum_{i=0}^{k-1} \text{tr}(M_1^\dagger M_1 \mathcal{E}^i(\rho)) \quad (6.20)$$

$$\leq n \epsilon \text{tr}(M_1^\dagger M_1 \rho) \frac{1 - a^{\lceil k/n \rceil}}{1 - a}. \quad (6.21)$$

The second inequality (6.18) follows from a technique similar to the one used in the proof of the *Sequence* rule. We bound the first term in (6.18) by $\epsilon \text{tr}(M_1 \rho M_1^\dagger)$ by applying the premise $(Q, \lambda) \vdash P_1 \leq \epsilon$ and the induction hypothesis. We will prove the post-measurement state $M_1 \rho M_1^\dagger$ indeed satisfies the precondition (Q, λ) later. Similarly, each term in (6.19) is bounded above by $\epsilon \text{tr}(M_1^\dagger M_1 \mathcal{E}^i(\rho))$ and thus the inequality in (6.20) holds.

To establish the inequality in (6.21), let $b_i := \text{tr}(M_1^\dagger M_1 \mathcal{E}^i(\rho))$. Then the sequence $\{b_k\}_k$ is non-negative and non-increasing. We now prove an upper bound of the series. Since P is (a, n) -bounded, we know that

$$\text{tr}(M_1^\dagger M_1 \mathcal{E}^n(\sigma)) = \text{tr}((\mathcal{E}^*)^n(M_1^\dagger M_1) \sigma) \leq a \text{tr}(M_1^\dagger M_1 \sigma), \text{ where } a < 1 \quad (6.22)$$

for every state σ , and therefore $b_{i+n} \leq ab_i$ for every i . Since b_k is non-increasing, we know that

$$\sum_{i=0}^{k-1} b_i = \sum_{m=0}^{\lceil k/n \rceil - 1} (b_{nm} + \dots + b_{nm+n-1}) \quad (6.23)$$

$$\leq n \sum_{m=0}^{\lceil k/n \rceil - 1} b_{nm} \leq n \sum_{m=0}^{\lceil k/n \rceil - 1} a^m b_0 = \frac{n(1 - a^{\lceil k/n \rceil})b_0}{1 - a}. \quad (6.24)$$

Thus the inequality in (6.21) holds. Since $b_0 = \text{tr}(M_1^\dagger M_1 \rho) \leq \text{tr}(\rho)$, we have

$$\frac{1}{2} \|\llbracket \widetilde{S}_k \rrbracket \rho - \llbracket S_k \rrbracket \rho\|_1 \leq \frac{n\epsilon(1 - a^{\lceil k/n \rceil})}{1 - a} \text{tr}(\rho). \quad (6.25)$$

Taking the limit as $k \rightarrow \infty$, we have that $a^{\lceil k/n \rceil} \rightarrow 0$, which shows that $\frac{1}{2} \|\llbracket \widetilde{S}_k \rrbracket \rho - \llbracket S_k \rrbracket \rho\|_1 \leq \frac{n\epsilon}{1-a} \text{tr}(\rho)$, as desired.

In order to apply the premise $(Q, \lambda) \vdash P_1 \leq \epsilon$ to states of the form $M_1 \mathcal{E}^i(\rho) M_1^\dagger$, we need to show that $\text{tr}(QM_1 \mathcal{E}^i(\rho) M_1^\dagger) \geq \lambda \text{tr}(M_1 \mathcal{E}^i(\rho) M_1^\dagger)$ for each i . We can prove this by induction. For the base case $i = 0$, by the precondition $(\lambda M_0^\dagger M_0 + M_1^\dagger Q M_1, \lambda)$ on the input state to the loop, we have $\text{tr}(M_1^\dagger Q M_1 \rho) \geq \lambda \text{tr}(\rho) - \lambda \text{tr}(M_0^\dagger M_0 \rho) = \lambda \text{tr}(M_1^\dagger M_1 \rho)$. Therefore, $\text{tr}(QM_1 \mathcal{E}^0(\rho) M_1^\dagger) \geq \lambda \text{tr}(M_1 \mathcal{E}^0(\rho) M_1^\dagger)$.

For the inductive step, observe that the Hoare triple $\{Q\}P_1\{R\}$ yields $\text{tr}(R \llbracket P_1 \rrbracket \sigma) \geq \text{tr}(Q\sigma)$ for $R \equiv \lambda M_0^\dagger M_0 + M_1^\dagger Q M_1$ and all states σ . By the induction hypothesis, we have

$$\text{tr}(R \llbracket P_1 \rrbracket (M_1 \mathcal{E}^i(\rho) M_1^\dagger)) \geq \text{tr}(QM_1 \mathcal{E}^i(\rho) M_1^\dagger) \geq \lambda \text{tr}(M_1 \mathcal{E}^i(\rho) M_1^\dagger) \geq \lambda \text{tr}(\mathcal{E}^{i+1}(\rho)),$$

where the last inequality holds because quantum operations are trace-non-increasing (applied to $\llbracket P_1 \rrbracket$). Note that $\llbracket P_1 \rrbracket (M_1 \mathcal{E}^i M_1^\dagger) = \mathcal{E}^{i+1}$. Now by substituting R , we have

$$\text{tr}(M_1^\dagger Q M_1 \mathcal{E}^{i+1}(\rho)) \geq \lambda \text{tr}(\mathcal{E}^{i+1}(\rho)) - \lambda \text{tr}(M_0^\dagger M_0 \mathcal{E}^{i+1}(\rho)) = \lambda \text{tr}(M_1^\dagger M_1 \mathcal{E}^{i+1}(\rho)). \quad (6.26)$$

Or equivalently, $\text{tr}(QM_1 \mathcal{E}^{i+1}(\rho) M_1^\dagger) \geq \lambda \text{tr}(M_1 \mathcal{E}^{i+1}(\rho) M_1^\dagger)$, which concludes the proof.

(9) *While-Unbounded*: the proof is trivial as we use the trivial upper bound 1.

□

7 CASE STUDIES

In this section, we apply our robustness definition and logic to two example quantum programs: the quantum Bernoulli factory and the quantum walk. We also demonstrate the (in)effectiveness

of different error correction schemes on single-qubit errors and analyze the robustness of a fault-tolerant version of QBF. Some computational details are deferred to appendices in the supplemental material.

7.1 Quantum Bernoulli Factory

The quantum Bernoulli factory (QBF) [Dale et al. 2015] is the quantum equivalent of the classical Bernoulli factory problem [Keane and O'Brien 1994]. In the classical Bernoulli factory problem, given a function $f : [0, 1] \mapsto [0, 1]$ and a coin that returns heads with unknown probability p , the goal is to simulate a new coin that returns head with probability $f(p)$. In QBF, the goal is to generate the state $|f(p)\rangle$ given a description of f and a *quantum coin* described by the state

$$|p\rangle := \sqrt{p}|0\rangle + \sqrt{1-p}|1\rangle. \quad (7.1)$$

QBF is interesting because it can simulate a strictly larger class of functions f than can be simulated by its classical counterpart. One example of a function that can be simulated by QBF, but not by the classical Bernoulli factory, is the probability amplification function. The key to simulating the probability amplification function is producing the state $|f(p)\rangle = (2p-1)|0\rangle + 2\sqrt{p(1-p)}|1\rangle$. This state can be prepared by (the ideal version of) the following program [Li and Ying 2018]:

```

 $\widetilde{QBF} \equiv q_1 := |1\rangle; q_2 := |1\rangle;$ 
  while  $M[q_2] = 1$  do
     $q_1 := |0\rangle; q_2 := |0\rangle;$ 
     $q_1 := \cong_{p_V, \Phi_V} V[q_1]; q_2 := \cong_{p_V, \Phi_V} V[q_2];$ 
     $q_1, q_2 := \cong_{p_U, \Phi_U} U[q_1, q_2]$  done,

```

where M is the standard basis measurement $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. The unitary U is defined by

$$U = |01\rangle\langle\phi^+| + |00\rangle\langle\phi^-| + |10\rangle\langle\psi^+| + |11\rangle\langle\psi^-| \quad (7.2)$$

where $|\phi^\pm\rangle = \frac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle)$ and $|\psi^\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle)$. The unitary $V := \begin{bmatrix} \sqrt{p} & -\sqrt{1-p} \\ \sqrt{1-p} & \sqrt{p} \end{bmatrix}$ acting on $|0\rangle$ generates the state $|p\rangle$. We denote the error due to noisy unitary application of V and U by $\epsilon_V = p_V \|\Phi_V - V \circ V^\dagger\|_\diamond$ and $\epsilon_U = p_U \|\Phi_U - U \circ U^\dagger\|_\diamond$ respectively. Note that to simplify our discussion we compute ϵ_V and ϵ_U using the trivial precondition $(I, 0)$. Recall that $\|\cdot\|_{I,0} = \|\cdot\|_\diamond$.

Now we prove that $(I, 0) \vdash \widetilde{QBF} \leq 4\epsilon_V + 2\epsilon_U$. First, by the *Sequence* and *Unitary* rules, we bound the error in the loop body by $2\epsilon_V + \epsilon_U$. To show (a, n) -boundedness of the loop, we must consider the behavior of $\mathcal{E}^*$ where

$$\mathcal{E}^* = (M_1 \circ M_1^\dagger)^* \circ \llbracket q_1 := |p\rangle; q_2 := |p\rangle \rrbracket^* \circ \llbracket U[q_1, q_2] \rrbracket^*. \quad (7.3)$$

We evaluate $\mathcal{E}^*(M_1^\dagger M_1)$ as follows. Recall that $M_1 = I \otimes |1\rangle\langle 1|$.

$$\begin{aligned} M_1^\dagger M_1 &= I \otimes |1\rangle\langle 1| \xrightarrow{\llbracket U[q_1, q_2] \rrbracket^*} |\phi^+\rangle\langle\phi^+| + |\psi^-\rangle\langle\psi^-| \\ &\xrightarrow{\llbracket q_1 := |p\rangle; q_2 := |p\rangle \rrbracket^*} \frac{1}{2} I \otimes I \end{aligned} \quad (7.4)$$

$$\xrightarrow{(M_1 \circ M_1^\dagger)^*} \frac{1}{2} M_1^\dagger M_1. \quad (7.5)$$

This shows that $\mathcal{E}^*(M_1^\dagger M_1) = \frac{1}{2} M_1^\dagger M_1$. Therefore, the while loop is $(\frac{1}{2}, 1)$ -bounded. Now we can use the *While-Bounded* rule to conclude that the error bound for the loop is $\frac{2\epsilon_V + \epsilon_U}{1 - \frac{1}{2}} = 4\epsilon_V + 2\epsilon_U$,

and therefore $(I, 0) \vdash (\text{while } M[q_2] = 1 \text{ do } \widetilde{P}_1 \text{ done}) \leq 4\epsilon_V + 2\epsilon_U$ where $\widetilde{P}_1$ is the body of the loop. Two additional applications of the *Sequence* rule conclude the proof.

We can compute ϵ_V and ϵ_U by identifying the error probabilities p_V, p_U and error models Φ_V, Φ_U , and numerically calculating the diamond norm. For example, consider the case where the state preparation is ideal, i.e., $p_V = 0$, and noise in the application of U is characterized by $p_U = 10^{-5}$ and $\Phi_U = (\frac{X \circ X + Y \circ Y + Z \circ Z + I \circ I}{4})^{\otimes 2}$, the 4-dimensional depolarizing channel. Then $\epsilon_V = 0$ and $\epsilon_U = p_U \|\Phi_U - U \circ U^\dagger\|_\diamond$. Applying an SDP solver [da Silva 2015; Watrous 2009] for the calculation of the diamond norm, we find that $\epsilon_U = 1 \times 10^{-5} \times 0.9375 = 9.375 \times 10^{-6}$, and therefore the error bound of $\widetilde{QBF}$ is 1.875×10^{-5} . More details can be found in the appendix of the extended version [Hung et al. 2018].

7.2 Quantum Walk on a Circle

Here we analyze the error of the quantum walk algorithm introduced in Section 4. The noisy quantum walk on a circle with n points can be written as the following program:

$$\widetilde{QW}_n \equiv p := |0\rangle; c := |L\rangle; \text{while } M[p] = 1 \text{ do } c \stackrel{\cong_{p_H, \Phi_H}}{=} H[c]; c, p \stackrel{\cong_{p_S, \Phi_S}}{=} S[c, p] \text{ done.} \quad (7.6)$$

Now we show that $(I, 0) \vdash \widetilde{QW}_6 \leq 30(\epsilon_H + \epsilon_S)$ where $\epsilon_H = p_H \|\Phi_H - H \circ H^\dagger\|_\diamond$ and $\epsilon_S = p_S \|\Phi_S - S \circ S^\dagger\|_\diamond$ are the errors due to noisy application of H and S respectively. First, by the *Sequence* and *Unitary* rules, we bound the error in the loop body by $\epsilon_H + \epsilon_S$. Next, we numerically test increasing values of (a, n) until we find a pair that satisfies (6.10). For this program, we find that the pair $(\frac{5}{6}, 5)$ satisfies the inequality, i.e., $(\mathcal{E}^*)^5 (M_1^\dagger M_1) \sqsubseteq \frac{5}{6} M_1^\dagger M_1$ where $\mathcal{E}^* = (M_1^\dagger \circ M_1) \circ \llbracket c, p := S[c, p] \rrbracket^* \circ \llbracket c := H[c] \rrbracket^*$. This implies that the loop is $(\frac{5}{6}, 5)$ -bounded. Note that here, unlike in the previous example, we have computed the (a, n) values numerically. This may be useful in cases where direct deduction of a and n is difficult. Now we can use the *While-Bounded* rule to conclude that the error bound for the loop is $\frac{5(\epsilon_H + \epsilon_S)}{1 - \frac{5}{6}} = 30(\epsilon_H + \epsilon_S)$, and therefore $(I, 0) \vdash (\text{while } M[p] = 1 \text{ do } \widetilde{P}_1 \text{ done}) \leq 30(\epsilon_H + \epsilon_S)$ where $\widetilde{P}_1$ is the body of the loop. Two additional applications of the *Sequence* rule conclude the proof.

As an example, consider the program $\widetilde{QW}_6$ where only the Hadamard gate may be faulty, i.e., $p_S = 0$. Say that noisy Hadamard application is characterized by $p_H = 5 \times 10^{-5}$ and $\Phi_H(p) = \frac{X \circ X + Y \circ Y + Z \circ Z + I \circ I}{4}$, the 2-dimensional depolarizing channel. Applying an SDP solver [da Silva 2015; Watrous 2009, 2013] for the calculation of the diamond norm, we find that $\epsilon_H = 5 \times 10^{-5} \times 0.75 = 3.75 \times 10^{-5}$. Therefore the error of $\widetilde{QW}_6$ is $30\epsilon_H = 1.125 \times 10^{-3}$. More details can be found in the appendix of the extended version [Hung et al. 2018].

7.3 Error Correction

In this section we use our semantics to show that an error correction scheme that is appropriate for the error model can reduce noise in a program, while an inappropriate error correction scheme may do the opposite. We consider three programs: P_1, P_2 , and P_3 . P_1 performs the identity operation, i.e., $P_1 \equiv q := I[q]$. The noisy version of $P_1, \widetilde{P}_1$, allows error during the application of I . P_2 and P_3 are semantically equivalent to P_1 (in the sense that they correspond to the same superoperator¹²), but both employ error correction. P_2 uses a three-qubit repetition code that can correct a bit flip (X error) on a single qubit and P_3 uses a three-qubit repetition code that can correct a phase flip (Z

¹²To make this mathematically rigorous, one needs to distinguish between variables and ancillas in the program. The semantics refers to the superoperator that traces out the ancilla part.

error) on a single qubit. P_2 and P_3 both have the following form:

$$\begin{aligned}\bar{q} &:= \text{ENCODE}[q]; \\ \bar{\bar{q}} &:= \bar{I}[\bar{q}]; \\ \bar{q} &:= \text{CORRECT}[\bar{q}]; \\ q &:= \text{DECODE}[\bar{q}]\end{aligned}$$

where we use *ENCODE* as a stand-in for the operations associated with turning the qubit q into its encoded counterpart $\bar{q}$, *CORRECT* as a stand-in for the operations associated with syndrome detection and error correction, and *DECODE* as a stand-in for the operations associated with converting the encoded qubit $\bar{q}$ back into q . We use $\bar{I}$ to represent a fault-tolerant version of the identity operation. In the noisy programs $\tilde{P}_2$ and $\tilde{P}_3$, we only allow error during application of $\bar{I}$. To simplify the calculation, we will assume that encoding, decoding, and error correction are all ideal (i.e. they have an error probability of 0). For a description of *ENCODE*, *CORRECT*, and *DECODE* see the appendix of the extended version [Hung et al. 2018].

For all three programs, we define the error model acting on a single qubit by $\Phi(\rho) = X\rho X$. With this error model, application of an I gate to a qubit will succeed with probability $1 - p$, and will instead become an application of a X gate with probability p . Now we can use our definition of the denotational semantics from Section 5 to directly compute $\llbracket \tilde{P}_1 \rrbracket$, $\llbracket \tilde{P}_2 \rrbracket$, and $\llbracket \tilde{P}_3 \rrbracket$. We find that:

$$\llbracket \tilde{P}_1 \rrbracket \rho = (1 - p)I\rho I + pX\rho X \quad (7.7)$$

$$\llbracket \tilde{P}_2 \rrbracket \rho = ((1 - p)^3 + 3p(1 - p)^2)I\rho I + (3p^2(1 - p) + p^3)X\rho X \quad (7.8)$$

$$\llbracket \tilde{P}_3 \rrbracket \rho = ((1 - p)^3 + 3p^2(1 - p)^2)I\rho I + (3p(1 - p) + p^3)Z\rho Z \quad (7.9)$$

Now we can directly compute $\|\llbracket P_1 \rrbracket - \llbracket \tilde{P}_1 \rrbracket\|_\diamond$, $\|\llbracket P_2 \rrbracket - \llbracket \tilde{P}_2 \rrbracket\|_\diamond$, and $\|\llbracket P_3 \rrbracket - \llbracket \tilde{P}_3 \rrbracket\|_\diamond$ to determine error rates for these programs. We find that:

$$\|\llbracket P_1 \rrbracket - \llbracket \tilde{P}_1 \rrbracket\|_\diamond = p \quad (7.10)$$

$$\|\llbracket P_2 \rrbracket - \llbracket \tilde{P}_2 \rrbracket\|_\diamond = 3p^2 - 2p^3 \quad (7.11)$$

$$\|\llbracket P_3 \rrbracket - \llbracket \tilde{P}_3 \rrbracket\|_\diamond = 3p(1 - p)^2 + p^3 \quad (7.12)$$

This allows us to conclude that, under $Q = I$ and $\lambda = 0$,

$$\tilde{P}_1 \text{ is } p \text{ robust} \quad \tilde{P}_2 \text{ is } 3p^2 - 2p^3 \text{ robust} \quad \tilde{P}_3 \text{ is } 3p(1 - p)^2 + p^3 \text{ robust}$$

Note that for $0 < p < \frac{1}{2}$, we have that $3p^2 - 2p^3 < p$ and $p < 3p(1 - p)^2 + p^3$, which tells us that, in the presence of X errors, correcting for bit flips will improve the error rate while correcting for phase flips will make the error rate worse. For computational details see the appendix of the extended version [Hung et al. 2018].

Here we invoke the semantics of the program to prove quantum robustness by definition. This is an expensive calculation. However, it is necessary to account for the effect of error correction. Ideally, we imagine a combination of uses of the semantics and the rules to trade off between the cost of the calculation and the accuracy of the bounds.

7.4 Fault-tolerant Quantum Bernoulli Factory

In this section, we consider a fault-tolerant implementation of the quantum Bernoulli factory. The fault-tolerant *QBF*, denoted by $\widetilde{QBF}$, can be written as follows.

```

 $\widetilde{QBF} \equiv q_1 := |1\rangle; q_2 := |1\rangle;$ 
while  $M[q_2] = 1$  do
   $q_1 := |0\rangle; q_2 := |0\rangle;$ 
   $\bar{q}_1 := \text{ENCODE}[q_1]; \bar{q}_2 := \text{ENCODE}[q_2];$ 
   $\bar{q}_1 := \cong_{1, \Phi_{\bar{V}}} \bar{V}[\bar{q}_1];$ 
   $\bar{q}_1 := \text{CORRECT}[\bar{q}_1];$ 
   $\bar{q}_2 := \cong_{1, \Phi_{\bar{V}}} \bar{V}[\bar{q}_2];$ 
   $\bar{q}_2 := \text{CORRECT}[\bar{q}_2];$ 
   $\bar{q}_1, \bar{q}_2 := \cong_{1, \Phi_{\bar{U}}} \bar{U}[\bar{q}_1, \bar{q}_2];$ 
   $\bar{q}_1 := \text{CORRECT}[\bar{q}_1]; \bar{q}_2 := \text{CORRECT}[\bar{q}_2];$ 
   $q_1 := \text{DECODE}[\bar{q}_1]; q_2 := \text{DECODE}[\bar{q}_2]$ 
done

```

where *ENCODE*, *CORRECT*, and *DECODE* have the same meanings as in the previous section. As before, we assume that encoding, decoding, and error correction are not affected by noise. $\bar{V}$ and $\bar{U}$ are the fault-tolerant operators that correspond to V and U respectively. We define the error model associated with $\bar{V}$ by $\Phi_{\bar{V}}(\rho) = \Phi_V^{\otimes 3}(\bar{V}\rho\bar{V}^\dagger)$ where $\Phi_V(\rho) = (1-p_V)I\rho I + p_V X\rho X$. Similarly, we define the error model associated with $\bar{U}$ by $\Phi_{\bar{U}}(\rho) = \Phi_U^{\otimes 6}(\bar{U}\rho\bar{U}^\dagger)$ where $\Phi_U(\rho) = (1-p_U)I\rho I + p_U X\rho X$. Note that because we are using a three-qubit repetition code, $\bar{V}$ and $\bar{U}$ will be 3-qubit and 6-qubit unitaries respectively. These noise models are applied with probability one.

Using the definitions from [Section 5](#) we can show that for $i \in \{1, 2\}$,

$$\llbracket \bar{q}_i := \cong_{1, \Phi_{\bar{V}}} \bar{V}[\bar{q}_i]; \bar{q}_i := \text{CORRECT}[\bar{q}_i] \rrbracket \rho = (1 - q_V) \bar{V} \rho \bar{V}^\dagger + q_V X^{\otimes 3} \bar{V} \rho \bar{V}^\dagger X^{\otimes 3}, \quad (7.13)$$

where $q_V = 3p_V^2(1 - p_V) + p_V^3$. Note that (7.13) is obtained through a calculation similar to the one used to compute $\llbracket P_2 \rrbracket$ in [Section 7.3](#). This says applying $\bar{V}$ followed by error correction is equivalent to applying $\bar{V}$ with probability $1 - q_V$, and applying $\bar{V}$ followed by $X^{\otimes 3}$ with probability q_V . Thus we can use the *Unitary* rule to compute the error of the noisy application of $\bar{V}$ followed by error correction:

$$\epsilon_{\bar{V}} = q_V \|(X^{\otimes 3} \circ X^{\otimes 3}) \circ (\bar{V} \circ \bar{V}^\dagger) - (\bar{V} \circ \bar{V}^\dagger)\|_\diamond = q_V \|(X^{\otimes 3} \circ X^{\otimes 3}) - (I \circ I)\|_\diamond = q_V. \quad (7.14)$$

The final equality holds because the unitaries are perfectly distinguishable. Note that $\|\Phi(U \circ U^\dagger) - U \circ U^\dagger\|_\diamond = \|\Phi - I \circ I\|_\diamond$ holds for any superoperator Φ and unitary U . Similarly,

$$\llbracket \bar{q}_1, \bar{q}_2 := \cong_{1, \Phi_{\bar{U}}} \bar{U}[\bar{q}_1, \bar{q}_2]; \bar{q}_1 := \text{CORRECT}[\bar{q}_1]; \bar{q}_2 := \text{CORRECT}[\bar{q}_2] \rrbracket \rho \quad (7.15)$$

$$= ((1 - q_U)I \circ I + q_U X^{\otimes 3} \circ X^{\otimes 3})^{\otimes 2} (\bar{U} \rho \bar{U}^\dagger), \quad (7.16)$$

where $q_U = 3p_U^2(1 - p_U) + p_U^3$. Thus, by the *Unitary* rule, the error of the noisy application of $\overline{U}$ followed by error correction is

$$\epsilon_{\overline{U}} = \|((1 - q_U)I \circ I + q_U X^{\otimes 3} \circ X^{\otimes 3})^{\otimes 2} - I \circ I\|_{\diamond} \quad (7.17)$$

$$= \|q_U^2(X^{\otimes 6} \circ X^{\otimes 6}) + q_U(1 - q_U)(I^{\otimes 3} \otimes X^{\otimes 3}) \circ (I^{\otimes 3} \otimes X^{\otimes 3}) \quad (7.18)$$

$$+ q_U(1 - q_U)(X^{\otimes 3} \otimes I^{\otimes 3}) \circ (X^{\otimes 3} \otimes I^{\otimes 3}) - (2q_U - q_U^2)I \circ I\|_{\diamond} \quad (7.19)$$

$$= 2q_U - q_U^2. \quad (7.20)$$

The last equality above can be proved as follows. For ease of notation, we define the superoperator

$$\mathcal{E}(\rho) := \frac{1}{2q_U - q_U^2} (q_U^2(X^{\otimes 6}\rho X^{\otimes 6}) + q_U(1 - q_U)(I^{\otimes 3} \otimes X^{\otimes 3})\rho(I^{\otimes 3} \otimes X^{\otimes 3}) \quad (7.21)$$

$$+ q_U(1 - q_U)(X^{\otimes 3} \otimes I^{\otimes 3})\rho(X^{\otimes 3} \otimes I^{\otimes 3})). \quad (7.22)$$

The superoperator $\mathcal{E}$ is trace-preserving, so $\epsilon_{\overline{U}} = (2q_U - q_U^2)\|\mathcal{E} - I \circ I\|_{\diamond} \leq (2q_U - q_U^2)$. To show $\epsilon_{\overline{U}}$ is also lower bounded by $2q_U - q_U^2$, it suffices to consider the input state $\rho = (|0\rangle\langle 0|)^{\otimes 6}$ and the projector $\Pi = I - (|0\rangle\langle 0|)^{\otimes 6}$. By definition, $\epsilon_{\overline{U}} = (2q_U - q_U^2)\|\mathcal{E} - I \circ I\|_{\diamond} \geq (2q_U - q_U^2)\text{tr}(\Pi(\mathcal{E}(\rho) - \rho)) = 2q_U - q_U^2$.

Now using the argument given in [Section 7.1](#), we can show that $(I, 0) \vdash \widetilde{QBF} \leq 4\epsilon_V + 2\epsilon_U$. Note that *ENCODE* and *DECODE* do not impact the robustness. Without error correction, as shown in [Section 7.1](#), $(I, 0) \vdash \widetilde{QBF} \leq 4\epsilon_V + 2\epsilon_U$ where $\epsilon_V = \|\Phi_V - I \circ I\|_{\diamond} = p_V$ and $\epsilon_U = \|\Phi_U^{\otimes 2} - I \circ I\|_{\diamond} = 2p_U - p_U^2$. In the case where the error rate is constant, i.e., $p_V = p_U = p$ for some probability p , $(I, 0) \vdash \widetilde{QBF} \leq O(p)$ and $(I, 0) \vdash \widetilde{QBF} \leq O(p^2)$. This shows that the error of $\widetilde{QBF}$ is suppressed by a factor of p with a fault-tolerant implementation of the loop body.

REMARK 7.1. *Throughout our example, we make the assumption that operations like ENCODE, CORRECT and DECODE are noise-free. In the actual setting of fault-tolerant quantum computation, they can also contain noise. In order to suppress the error rate, one may need to use the construction of fault-tolerant gadgets in the proof of the threshold theorem [Aharonov and Ben-Or 1997]. We remark that the assumption we made is just to simplify the example and ease presentation. It would in fact be possible to prove a threshold theorem in our formalism without this assumption, although the calculation could be much more complicated. It is an interesting open question to see whether one can simplify this calculation by adding more rules to our logic.*

8 CONCLUSIONS AND FUTURE WORK

We have presented a semantics for describing quantum computation with errors and an analysis that bounds the distance between the result of a noisy program and its corresponding ideal program on the same input. We used our analysis to compute error bounds for noisy versions of the quantum Bernoulli factory and quantum walk programs. We also showed how our analysis can be used to compute the error bounds for small circuits with and without error correction, showing examples of when using error correction is beneficial and when there are tradeoffs between the efficiency of error corrections and related costs.

A natural next step for our work is to encode the rules from [Section 6.3](#) in a proof assistant so that they can be applied in an automated fashion. We have shown that the (Q, λ) -diamond norm can be computed by an SDP, and the (a, n) values for loop boundedness can be computed analytically or numerically, so implementing our rules is feasible. Given an implementation, we are interested how our analysis may be used to construct circuits with lower error rates. This application is inspired by work by [Misailovic et al. \[2014\]](#), which uses classical reliability analysis to determine which

operations can be replaced by their noisy counterparts. It could hence be used to inform decisions about which implementations of quantum algorithms are practical for near-term use.

ACKNOWLEDGMENTS

We would like to thank Andrew Childs for helpful discussions on quantum walks. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research Quantum Testbed Pathfinder Program under Award Number DE-SC0019040, and the Canadian Institute for Advanced Research.

REFERENCES

- Ali Javadi Abhari, Arvin Faruque, Mohammad Javad Dousti, Lukas Svec, Oana Catu, Amlan Chakrabati, Chen-Fu Chiang, Seth Vanderwilt, John Black, Fred Chong, Margaret Martonosi, Martin Suchara, Ken Brown, Massoud Pedram, and Todd Brun. 2012. *Scaffold: Quantum Programming Language*. Technical Report TR-934-12. Princeton University.
- Dorit Aharonov, Andris Ambainis, Julia Kempe, and Umesh Vazirani. 2001. Quantum Walks on Graphs. In *STOC*.
- Dorit Aharonov and Michael Ben-Or. 1997. Fault-tolerant Quantum Computation with Constant Error. In *STOC*.
- Woongki Baek and Trishul M. Chilimbi. 2010. Green: A Framework for Supporting Energy-conscious Programming Using Controlled Approximation. In *PLDI*.
- Alexandru Baltag and Sonja Smets. 2011. Quantum Logic as a Dynamic Logic. *Synthese* 179, 2 (2011).
- James Bornholt, Todd Mytkowicz, and Kathryn S. McKinley. 2014. Uncertain $\langle T \rangle$: A First-order Type for Uncertain Data. In *ASPLOS*.
- Brett Boston, Adrian Sampson, Dan Grossman, and Luis Ceze. 2015. Probability Type Inference for Flexible Approximate Programming. In *OOPSLA*.
- Olivier Brunet and Philippe Jorrand. 2004. Dynamic Quantum Logic for Quantum Programs. *International Journal of Quantum Information* 2, 1 (2004).
- Michael Carbin, Deokhwan Kim, Sasa Misailovic, and Martin C. Rinard. 2012. Proving Acceptability Properties of Relaxed Nondeterministic Approximate Programs. In *PLDI*.
- Michael Carbin, Sasa Misailovic, and Martin C. Rinard. 2013. Verifying Quantitative Reliability for Programs That Execute on Unreliable Hardware. In *OOPSLA*.
- Rohit Chadha, Paulo Mateus, and Amílcar Sernadas. 2006. Reasoning About Imperative Quantum Programs. *Electronic Notes in Theoretical Computer Science* 158 (2006).
- Swarat Chaudhuri, Sumit Gulwani, and Roberto Lubliner. 2010. Continuity Analysis of Programs. In *POPL*.
- Swarat Chaudhuri, Sumit Gulwani, Roberto Lubliner, and Sara Navidpour. 2011. Proving Programs Robust. In *ESEC/FSE*.
- Frederic T. Chong, Diana Franklin, and Margaret Martonosi. 2017. Programming Languages and Compiler Design for Realistic Quantum Hardware. *Nature* 549 (2017).
- Isaac L. Chuang and Michael A. Nielsen. 1997. Prescription for Experimental Determination of the Dynamics of a Quantum Black Box. *Journal of Modern Optics* 44, 11-12 (1997).
- Marcus P. da Silva. 2015. matlab-diamond-norm: A MATLAB function for computing the diamond norm (completely bounded induced 1-norm on linear superoperators). <https://github.com/BBN-Q/matlab-diamond-norm>
- Howard Dale, David Jennings, and Terry Rudolph. 2015. Provable Quantum Advantage in Randomness Processing. *Nature Communications* 6 (2015).
- Ellie D'Hondt and Prakash Panangaden. 2006. Quantum Weakest Preconditions. *Mathematical Structures in Computer Science* 16, 3 (2006).
- Joseph Emerson, Robert Alicki, and Karol Życzkowski. 2005. Scalable Noise Estimation with Random Unitary Operators. *Journal of Optics B: Quantum and Semiclassical Optics* 7, 10 (2005).
- Yuan Feng, Runyao Duan, Zhengfeng Ji, and Mingsheng Ying. 2007. Proof Rules for the Correctness of Quantum Programs. *Theoretical Computer Science* 386, 1-2 (2007).
- Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. 2012. Surface Codes: Towards Practical Large-scale Quantum Computation. *Physical Review A* 86, 3 (2012).
- Simon J. Gay. 2006. Quantum Programming Languages: Survey and Bibliography. *Mathematical Structures in Computer Science* 16, 4 (2006).
- Iulia Georgescu, Sahel Ashhab, and Franco Nori. 2014. Quantum Simulation. *Reviews of Modern Physics* 86, 1 (2014).
- Alexei Gilchrist, Nathan K. Langford, and Michael A. Nielsen. 2005. Distance Measures to Compare Real and Ideal Quantum Processes. *Physical Review A* 71, 6 (2005).
- Daniel Gottesman. 2010. An Introduction to Quantum Error Correction and Fault-tolerant Quantum Computation. *Proceedings of Symposia in Applied Mathematics: Quantum Information Science and Its Contributions to Mathematics* 68

(2010).

- Jonathan Grattage. 2005. A Functional Quantum Programming Language. In *LICS*.
- Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. 2013. Quipper: A Scalable Quantum Programming Language. In *PLDI*.
- Mauricio Gutiérrez, Lukas Svec, Alexander Vargo, and Kenneth R. Brown. 2013. Approximation of Realistic Errors by Clifford Channels and Pauli Measurements. *Physical Review A* 87, 3 (2013).
- Shih-Han Hung, Kesha Hietala, Shaopeng Zhu, Mingsheng Ying, Michael Hicks, and Xiaodi Wu. 2018. Quantitative Robustness Analysis of Quantum Programs (Extended Version). *CoRR* abs/1811.03585 (2018). arXiv:1811.03585
- Yoshihiko Kakutani. 2009. A Logic for Formal Verification of Quantum Programs. In *ASIAN*.
- M. S. Keane and George L. O'Brien. 1994. A Bernoulli Factory. *ACM Transactions on Modeling and Computer Simulation* 4, 2 (1994).
- Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R. Brad Blakestad, John D. Jost, Chris Langer, Roee Ozeri, S. Seidelin, and David J. Wineland. 2008. Randomized Benchmarking of Quantum Gates. *Physical Review A* 77, 1 (2008).
- Yangjia Li and Mingsheng Ying. 2018. Algorithmic Analysis of Termination Problems for Quantum Programs. In *POPL*.
- Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. 2011. Scalable and Robust Randomized Benchmarking of Quantum Processes. *Physical Review Letters* 106, 18 (2011).
- Sasa Misailovic, Michael Carbin, Sara Achour, Zichao Qi, and Martin C. Rinard. 2014. Chisel: Reliability- and Accuracy-aware Optimization of Approximate Computational Kernels. In *OOPSLA*.
- Nikolaj Moll, Panagiotis Barkoutsos, Lev S. Bishop, Jerry M. Chow, Andrew Cross, Daniel J. Egger, Stefan Filipp, Andreas Fuhrer, Jay M. Gambetta, Marc Ganzhorn, Abhinav Kandala, Antonio Mezzacapo, Peter Müller, Walter Riess, Gian Salis, John Smolin, Ivano Tavernelli, and Kristan Temme. 2018. Quantum Optimization Using Variational Algorithms on Near-term Quantum Devices. *Quantum Science and Technology* 3, 3 (2018).
- Michael A. Nielsen and Isaac Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge University Press.
- Bernhard Ömer. 2003. *Structured Quantum Programming*. Ph.D. Dissertation. Vienna University of Technology.
- Jongse Park, Hadi Esmailzadeh, Xin Zhang, Mayur Naik, and William Harris. 2015. FlexJava: Language Support for Safe and Modular Approximate Programming. In *ESEC/FSE*.
- Jennifer Paykin, Robert Rand, and Steve Zdancewic. 2017. QWIRE: A Core Language for Quantum Circuits. In *POPL*.
- Frances Perry, Lester Mackey, George A. Reis, Jay Ligatti, David I. August, and David Walker. 2007. Fault-tolerant Typed Assembly Language. In *PLDI*.
- Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien. 2014. A Variational Eigenvalue Solver on a Photonic Quantum Processor. *Nature Communications* 5, 4213 (2014).
- John Preskill. 2018. Quantum Computing in the NISQ Era and Beyond. *Quantum* 2, 79 (2018).
- Amr Sabry. 2003. Modeling Quantum Computing in Haskell. In *The Haskell Workshop*.
- Adrian Sampson, Werner Dietl, Emily Fortuna, Danushen Gnanapragasam, Luis Ceze, and Dan Grossman. 2011. EnerJ: Approximate Data Types for Safe and General Low-power Computation. In *PLDI*.
- Adrian Sampson, Pavel Panchekha, Todd Mytkowicz, Kathryn S. McKinley, Dan Grossman, and Luis Ceze. 2014. Expressing and Verifying Probabilistic Assertions. In *PLDI*.
- Jeff W. Sanders and Paolo Zuliani. 2000. Quantum Programming. In *MPC*.
- Peter Selinger. 2004a. A Brief Survey of Quantum Programming Languages. In *FLOPS*.
- Peter Selinger. 2004b. Towards a Quantum Programming Language. *Mathematical Structures in Computer Science* 14, 4 (2004).
- Martin Suchara, John Kubitowicz, Arvin I. Faruque, Frederic T. Chong, Ching-Yi Lai, and Gerardo Paz. 2013. QuRE: The Quantum Resource Estimator Toolbox. In *ICCD*.
- Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. 2018. Q#: Enabling Scalable Quantum Computing and Development with a High-level DSL. In *RWDSDL*.
- Barbara M. Terhal. 2015. Quantum Error Correction for Quantum Memories. *Reviews of Modern Physics* 87, 2 (2015).
- David Walker, Lester Mackey, Jay Ligatti, George A. Reis, and David I. August. 2006. Static Typing for a Faulty Lambda Calculus. In *ICFP*.
- John Watrous. 2006. Introduction to Quantum Computation. <https://cs.uwaterloo.ca/~watrous/LectureNotes/CPSC519.Winter2006/all.pdf>. Course notes.
- John Watrous. 2009. Semidefinite Programs for Completely Bounded Norms. *Theory of Computing* 5, 11 (2009).
- John Watrous. 2013. Simpler Semidefinite Programs for Completely Bounded Norms. *Chicago Journal of Theoretical Computer Science* 2013, 8 (2013).
- John Watrous. 2018. *The Theory of Quantum Information*. Cambridge University Press.

- Dave Wecker and Krysta Svore. 2014. LIQ| $\rangle$: A Software Design Architecture and Domain-Specific Language for Quantum Computing. *CoRR* abs/1402.4467 (2014). arXiv:[1402.4467](#)
- Mingsheng Ying. 2011. Floyd–Hoare Logic for Quantum Programs. *ACM Transactions on Programming Languages and Systems* 33, 6 (2011).
- Mingsheng Ying. 2016. *Foundations of Quantum Programming*. Morgan Kaufmann.
- Mingsheng Ying, Shenggang Ying, and Xiaodi Wu. 2017. Invariants of Quantum Programs: Characterisations and Generation. In *POPL*.