

On the Principles of Differential Quantum Programming Languages

Shaopeng Zhu

Shih-Han Hung

Shouvanik Chakrabarti

Xiaodi Wu

University of Maryland, College Park, USA

Abstract

Variational Quantum Circuits (VQCs), or the so-called quantum neural-networks, are predicted to be one of the most important near-term quantum applications, not only because of their similar promises as classical neural-networks, but also because of their feasibility on near-term noisy intermediate-size quantum (NISQ) machines. The need for gradient information in the training procedure of VQC applications has stimulated the interest and development of auto-differentiation techniques on quantum circuits. We propose the first formalization of this technique, not only in the context of quantum circuits but also for imperative quantum programs (e.g., with controls), inspired by the success of differential programming languages in classical machine learning. In particular, we overcome a few unique difficulties caused by exotic quantum features (such as quantum no-cloning) and provide a rigorous formulation of differentiation applied to bounded-loop imperative quantum programs, its code-transformation rules, as well as a sound logic to reason about their correctness.

Moreover, we have implemented our code transformation in OCaml and demonstrated the resource-efficiency of our scheme both analytically and empirically. We also conduct a case study of training a VQC instance with controls, which shows the advantage of our scheme over existing auto-differentiation on quantum circuits without controls.

Keywords quantum programming languages, differential programming languages, auto differentiation, quantum machine learning, collection of semantics

1 Introduction

Background. Recent years have witnessed the rapid development of quantum computing, with practical advances coming from both research and industry. Quantum programming is one topic that has been actively investigated. Early work on language design [22, 32, 39, 40, 44] has been followed up recently by several implementations of these languages, including Quipper [24], Scaffold [1], LIQUi| \rangle [49], Q# [46], and QWIRE [33]. Extensions of program logics have also been proposed for verification of quantum programs [3, 9, 10, 17, 27, 52, 55]. See also surveys [18, 43, 53].

With the availability of prototypes of quantum machines, especially the recent establishment of quantum supremacy [2], the research of quantum computing has entered a new stage

where near-term Noisy Intermediate-Scale Quantum (NISQ) computers [37], e.g., the 53-qubit quantum machine from Google [2] and IBM [20], becomes the important platform for demonstrating quantum applications. Variational quantum circuits (VQCs) [15, 16, 35], or the so-called *quantum neural networks*, are predicted to be one of the most important applications on NISQ machines. This is not only because VQCs bear a lot of similar promises like classical neural networks as well as potential quantum speed-ups from the perspective of machine learning (e.g., see the survey [8]), but also because VQC is, if not the only, one of the few candidates that can be implemented on NISQ machines. Because of this, a lot of study has already been devoted to the design, analysis, and small-scale implementation of VQC (e.g., see the survey [6]).

Typical VQC applications replace classical neural networks, which are just parameterized classical circuits, by quantum circuits with *classically parameterized unitary gates*. Namely, one will have a "quantum" mapping from input to output replacing classical mapping in machine learning applications. An important component of these applications is a training procedure which optimizes a *loss function* the now depends on the *read-outs* and the *parameters* of VQCs.

It is well-known that gradient-based approaches are the best for the training procedure. However, computing the gradients of loss functions that depend on the read-outs of quantum circuits has a similar complexity of simulating quantum circuits, which is infeasible for classical computation. Thus, the ability of evaluating these "quantum" gradients efficiently by quantum computation is critical for the scalability of VQC applications, e.g., on the recent 53-qubit machines.

Fortunately, analytical formulas of gradients used in VQCs have been studied by [16, 19, 29, 41, 42]. In particular, Schuld et al. [41] proposed the so-called *phase-shift rule* that uses two quantum circuits to compute the partial derivative respective to one parameter for quantum circuits. One of the very successful tools for quantum machine learning, called PennyLane [7], implemented the phase-shift rule to achieve *auto differentiation* (AD) for the read-outs of quantum circuits. It also integrated automatic differentiation from established machine learning libraries such as TensorFlow or PyTorch for any additional classical calculation in the training procedure. However, none of these research was conducted from the perspective of programming languages and no rigorous foundation or principles have been formalized.

Motivations. An important motivation of this paper is to provide a *rigorous formalization* of the auto-differentiation

technique applied to quantum circuits. In particular, we will provide a formal formulation of quantum programs, their semantics, and the meaning of differentiation on them. We will also study the code-transformation rules for auto-differentiation and prove their correctness.

As we will highlight below, research on the formalization will encounter many new challenges that have not been considered and addressed by existing results [16, 19, 29, 41, 42]. Consider one of the basic requirements, e.g., compositionality. As we will show, differentiating the composition of quantum programs will necessarily involve running multiple quantum programs on copies of initial quantum states. How to represent the collection of quantum programs succinctly and also bound the number of required copies is a totally new question. Among our techniques to address this question, we also need to change the previously proposed construct, e.g., the phase-shift rule [41], to something else.

Moreover, we want to go beyond the restriction to quantum circuits in [16, 19, 29, 41, 42]. We are inspired by classical machine learning examples that demonstrate the advantage of neural-networks with program features (e.g. controls) over the plain ones (e.g., classical circuits) [23, 25], which is also the major motivation of promoting the paradigm shift from deep learning towards *differential programming* by prominent machine learning researchers.

Augmenting VQCs with controls is not only feasible on NISQ machines (at least for simple ones), but also now a logical step for the study of their applications in machine learning. (We indeed conduct one case study in Section 8.)

Thus, we are inspired to investigate the *principles of differential quantum imperative languages* beyond circuits.

Research challenges & Solutions. We will rely on a little notation that should be self-explanatory. Please refer to a detailed preliminary on quantum information in Section 2.

Let us start with a simple classical program

$$\text{MUL} \equiv v_3 = v_1 \times v_2, \quad (1.1)$$

where v_3 is the product of v_1 and v_2 . Consider the differentiation with respect to θ , we have

$$\frac{\partial}{\partial \theta}(\text{MUL}) \equiv v_3 = v_1 \times v_2; \quad (1.2)$$

$$\dot{v}_3 = \dot{v}_1 \times v_2 + v_1 \times \dot{v}_2, \quad (1.3)$$

where MUL keeps track of variables v_1, v_2, v_3 and their derivatives $\dot{v}_1, \dot{v}_2, \dot{v}_3$ at the same time. One simple yet important observation is that classical variables v_1, v_2, v_3 are real-valued and can be naturally differentiated, while quantum variables are quantum states represented by matrices.

What are hence reasonable quantities to compute derivatives about quantumly? One natural choice from the principles of quantum mechanics is the (classical) read-outs of quantum systems through measurements, which we formulate as the *observable semantics* of quantum programs. This

natural choice also serves the purpose of gradients computation of loss functions in quantum machine learning applications, which are typically phrased in terms of these read-outs. We also directly model the parameterization of quantum programs after VQCs, i.e., each unitary gate becomes classically parameterized. The above modeling of quantum programs is very different from classical ones. It is also unclear whether any reasonable analogues of classical chain-rule and forward/backward mode can exist within quantum programs.

To model the meaning of one quantum program computing the derivative of another, we define the *differential semantics* of programs. There is a subtle quantum-unique design choice. The observable semantics of any quantum program will depend on the observable and its input state. Thus, any program computing its derivative could potentially depend on these two extra factors. We find out this potential dependence is undesirable and propose the strongest possible definition: i.e., one derivative computing program should work for *any pair* of observable and input states. We demonstrate that this strong requirement is not only achievable but also critical for the composition of auto differentiation.

We are ready to describe the technical challenges for the compositionality. Consider the following quantum program:

$$\text{QMUL} \equiv U_1(\theta); U_2(\theta), \quad (1.4)$$

which performs $U_1(\theta)$ and $U_2(\theta)$ gates sequentially. Note that gate application is matrix multiplication in quantum. Roughly speaking, if the product rule of differentiation (as exhibited in (1.3)) remains in the quantum setting, at least symbolically, then one should expect $\frac{\partial}{\partial \theta}(\text{QMUL})$ contains

$$\frac{\partial}{\partial \theta}(U_1(\theta)); U_2(\theta) \text{ and } U_1(\theta); \frac{\partial}{\partial \theta}(U_2(\theta)) \quad (1.5)$$

two different parts as sub-programs similarly in (1.3).

However, we cannot run $\frac{\partial}{\partial \theta}(U_1(\theta)); U_2$ and $U_1(\theta); \frac{\partial}{\partial \theta}(U_2(\theta))$ together due to the quantum *no-cloning* theorem [51]. This is because they share the same initial state and we cannot clone two copies of it and maintain the correct correlation among different parts. Note that this is not an issue classically as we can store all v_i, \dot{v}_i at the same time as in (1.3). As a result, quantum differentiation needs to run *multiple* (sub-)programs on *multiple* copies of the initial state.

This poses a unique challenge for differentiation of quantum composition: (1) we hope to have a simple scheme of code transformation, ideally close-to-classical, for intuition and easy implementation of the compiler, whereas it needs to express correctly the collection of quantum programs during code transformation; (2) for the purpose of efficiency, we also want to reasonably bound the number of required copies of the initial states, which roughly refers to the number of different quantum programs in the collection.

We develop a few techniques to achieve both goals at the same time. First, we propose the so-called *additive* quantum programs as a *succinct* intermediate representation for the

collection of programs during the code transformation. Now the entire differentiation procedure will be divided into two steps: (1) all code transformations happen on additive programs and are very similar to classical ones (see Figure 4); (2) the collection of programs can be recovered by a compilation procedure from any additive program. Additive quantum programs are equipped with a new *sum* operation that models the multiple choices as exhibited in (1.5), which resembles a similar idea in differential lambda-calculus [12].

Second, we also design a new rule for $\frac{\partial}{\partial \theta}(U(\theta))$ which is slightly different from [16, 19, 29, 41, 42]. The existing phase-shift rule makes use of two quantum circuits for one differentiation, which causes a lot of inconvenience in the formulation and potential trouble for efficiency. We replace it by one circuit with one additional ancilla register for each parameter. We conduct a careful resource analysis of our differentiation procedure and show the number of required copies of initial states is reasonable comparing to the classical setting. The correctness of the code transformation of composition critically relies on our design choice as well as the strong definition related to the differential semantics.

With the previous setup, we can naturally build the differentiation for controls (i.e., the quantum condition statement). Note that a general solution for classical controls is still unknown [5] due to the non-smoothness of the guard. Similar to the classical setting [36], we do not have a solution to deal with general unbounded loop. We instead provide a solution to bounded loops which can be deemed as a macro consisting of simple quantum condition statements.

Contributions. We formulate the parameterized quantum bounded while-programs with classically parameterized unitary gates modelled after VQCs [15, 30, 35] and their realistic examples on ion-trap machines, e.g. [56], in Section 3.

In Section 4, we illustrate our design of *additive* quantum programs. Specially, we add the syntax $P_1 \dot{+} P_2$ to represent the either-or choice between P_1 and P_2 in (1.5). We formulate its semantics and the compilation rules that map any additive program into a collection of normal programs in the way to serve our differentiation purpose.

In Section 5, we formulate the observable and the differential semantics of quantum programs and formally define the meaning of program $S'(\theta)$ computing the differential semantics of $S(\theta)$ in the strongest possible sense. In particular, we require that $S'(\theta)$ should work for any pair of observable and input state, which is critical for the compositionality.

In Section 6, we show that such a strong requirement is indeed achievable by demonstrating the code-transformation rules for the differentiation procedure. Thanks to the use of additive quantum programs, the code transformation is much simplified and as intuitive as classical ones. We also develop a logic with the judgement $S'(\theta)|S(\theta)$ that states $S'(\theta)$ computes the differential semantics of $S(\theta)$. We prove

the soundness of our logic and use it to show the correctness of our code-transformation rules.

In Section 7, we conduct a resource analysis to further justify our design. We show that the *occurrence count* of parameters capture the extra resource required in both the classical auto-differentiation and our scheme. In this sense, our resource cost is reasonable comparing to classical.

Finally, in Section 8, we demonstrate the implementation of our code transformation in OCaml and apply it to the training of one VQC instance with controls via classical simulation. Specifically, this instance shows an advantage of controls in machine learning tasks, which implies the advantage of our scheme over previous ones that cannot handle controls. We have also empirically verified the resource-efficiency of our scheme on representative VQC instances.

Related classical work. There is an extensive study of automatic differentiation (AD) or differential programming in the classical setting (e.g., see books [11, 26]). The most relevant to us are those studies from the programming language perspective. AD has traditionally been applied to imperative programs in both the forward mode, e.g. [28, 50], and the reverse mode, e.g., [45]. The famous *backpropagation* algorithm [38] is also a special case of reserve-mode AD used to compute the gradient of a multi-layer perceptron. AD has also been recently applied to functional programs [13, 14, 34]. Motivated by the success of deep learning, there is significant recent interest to develop both the theory and the implementation of AD techniques. Please refer to the survey [4] and the keynote talk at POPL'18 [36] for more details.

2 Quantum Preliminaries

We present basic quantum preliminaries here (a summary of notation in Table 1). Details are deferred to Appendix A.

2.1 Math Preliminaries

Let n be a natural number. We refer to the complex vector space \mathbb{C}^n as an n -dimensional Hilbert space \mathcal{H} . We use $|\psi\rangle$ to denote a complex vector in \mathbb{C}^n . The Hermitian conjugate of $|\psi\rangle \in \mathbb{C}^n$ is denoted by $\langle\psi|$. The *inner product* of $|\psi\rangle$ and $|\phi\rangle$, defined as the product of $\langle\psi|$ and $|\phi\rangle$, is denoted by $\langle\psi|\phi\rangle$. The norm of a vector $|\psi\rangle$ is denoted by $\| |\psi\rangle \| = \sqrt{\langle\psi|\psi\rangle}$.

We define (linear) *operators* as linear maps between Hilbert spaces, which can be represented by matrices for finite dimension. Let A be an operator; its Hermitian conjugate is denoted by A^\dagger . A is *Hermitian* if $A = A^\dagger$. The *trace* of 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 of $|\psi\rangle$ and $A|\psi\rangle$. A Hermitian operator A is *positive semidefinite* if for all vectors $|\psi\rangle \in \mathcal{H}$, $\langle\psi|A|\psi\rangle \geq 0$.

2.2 Quantum States and Operations

The state space of a *qubit* is a 2-dimensional Hilbert space. Two important orthonormal bases of a qubit system are: the

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

Spaces	\mathcal{H}, \mathcal{A}	$L(\mathcal{H})$ (Linear operators)
States	(pure states)	$ \psi\rangle, \phi\rangle$ (metavariables); $ 0\rangle, 1\rangle, +\rangle, -\rangle$
	(density)	ρ, σ (meta); $ \psi\rangle\langle\psi $
Operations	(unitaries)	U, V, σ (meta); $H, X, Z, X \otimes X$
	(superoperators)	\mathcal{E}, \mathcal{F} (general); Φ (quantum channels)
Measurements	M	$\{M_m\}_m$ (general); $\{ 0\rangle\langle 0 , 1\rangle\langle 1 \}$
Observables	O	$\sum_m \lambda_m \phi_m\rangle\langle\phi_m $ (general); $ 0\rangle\langle 0 - 1\rangle\langle 1 $ (example)
Programs	(no parameters)	P, Q (meta)
	(parameterized)	$P(\theta), S(\theta)$ (meta); $R_\sigma(\theta)$ (notable examples)
	(collection)	$\underline{P(\theta)}, \underline{S(\theta)}$ (meta); $\frac{\partial}{\partial \theta}(QCT_{(7,6)}(\theta))$ (example)
Semantics	(operational)	$\langle P, \rho \rangle \rightarrow \langle Q, \rho' \rangle$ (steps)
	(denotational)	$\llbracket P \rrbracket \rho = \rho'$ (example)
	(observable)	$\llbracket (O, \rho) \rightarrow P(\theta) \rrbracket$ (example)
Code Process	(Transform)	$\frac{\partial}{\partial \theta}(S(\theta))$ (example)
	(Compile)	$\text{Compile}(S'(\theta))$ (example)
Logic	(Judgement)	$S'(\theta) S(\theta) \rangle$ (example)
Count	(Non-Abort)	$ \# \frac{\partial}{\partial \theta_j}(P(\theta)) $ (example)
	(Occurrence)	$\text{OC}_j(P(\theta))$ (example)

computational basis with $|0\rangle = (1, 0)^\dagger$ and $|1\rangle = (0, 1)^\dagger$; the \pm basis, consisting of $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.

A *pure* quantum state is 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 use *density operators* to represent states: a density operator representing the ensemble $\{(p_i, |\psi_i\rangle)\}_i$ is a positive semidefinite operator $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$. 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 on \mathcal{H} is denoted by $\mathcal{D}(\mathcal{H})$.

Operations on quantum systems can be characterized by unitary operators. Denoting the set of linear operators on \mathcal{H} as $L(\mathcal{H})$, an operator $U \in L(\mathcal{H})$ is *unitary* if $U^\dagger U = U U^\dagger = I_{\mathcal{H}}$. A unitary *evolves* a pure state $|\psi\rangle$ to $U|\psi\rangle$, or a density operator ρ to $U\rho U^\dagger$. Common unitary operators include: the *Hadamard* operator H , which transforms between the computational and the \pm basis via $H|0\rangle = |+\rangle$ and $H|1\rangle = |-\rangle$; the *Pauli X* operator which performs a bit flip, i.e., $X|0\rangle = |1\rangle$ and $X|1\rangle = |0\rangle$; *Pauli Z* which performs a phase flip, i.e., $Z|0\rangle = |0\rangle$ and $Z|1\rangle = -|1\rangle$; *Pauli Y* mapping $|0\rangle$ to $i|1\rangle$ and $|1\rangle$ to $-i|0\rangle$; *CNOT* gate mapping $|00\rangle \mapsto |00\rangle, |01\rangle \mapsto |01\rangle, |10\rangle \mapsto |11\rangle, |11\rangle \mapsto |10\rangle$.

More generally, evolution of a quantum system can be characterized by an *admissible superoperator* \mathcal{E} , namely a *completely-positive* and *trace-non-increasing* (See Appendix A) linear map from $\mathcal{D}(\mathcal{H})$ to $\mathcal{D}(\mathcal{H}')$ for Hilbert spaces $\mathcal{H}, \mathcal{H}'$.

For every superoperator \mathcal{E} , 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 $\rho \in \mathcal{D}(\mathcal{H})$. The *Kraus form* of \mathcal{E} is therefore $\mathcal{E} = \sum_k E_k \circ E_k^\dagger$. 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 $\sum_k E_k^\dagger \circ E_k$.

2.3 Quantum Measurements

Quantum *measurements* extracts classical information out of quantum systems. A quantum measurement on a system over Hilbert space \mathcal{H} can be described by a set of linear operators $\{M_m\}_m$ with $\sum_m M_m^\dagger M_m = I_{\mathcal{H}}$ (identity matrix on \mathcal{H}). If we perform a measurement $\{M_m\}_m$ on a state ρ , the outcome m is observed with probability $p_m = \text{tr}(M_m \rho M_m^\dagger)$ for each m , and the post-measurement state collapses to $M_m \rho M_m^\dagger / p_m$.

3 Parameterized Quantum Bounded While-Programs

We adopt the *bounded-loop* variant of the quantum **while**-language developed by Ying [53], and augment it by parameterizing the unitaries, as this provides sufficient expressibility for parameterized quantum operations: indeed, **abort**, **skip** and initialization behave independently of parameters, while “parameterized measurements” can be implemented with a regular measurement followed by a parameterized unitary.

From here onward, \bar{v} is a finite set of variables, and θ a length- k vector of real-valued parameters.

3.1 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 $q \in \text{Var}$, its state space is denoted by \mathcal{H}_q . The quantum register \bar{q} is associated with the Hilbert space $\mathcal{H}_{\bar{q}} = \bigotimes_{q \in \bar{q}} \mathcal{H}_q$.¹ A *T-bounded, k-parameterized quantum while-program* is generated by the following syntax:

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

¹If $\text{type}(q) = \mathbf{Bool}$ then $\mathcal{H}_q = \text{span}\{|0\rangle, |1\rangle\}$. If $\text{type}(q) = \mathbf{Bounded Int}$ then \mathcal{H}_q is with basis $\{|n\rangle : n \in [-N, N]\}$ ($N \in \mathbb{Z}^+$) for some finite N . We require the Hilbert space to be finite dimensional for implementation.

where

$$\begin{aligned} & \text{while}^{(1)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done} \\ \equiv & \text{case } M[\bar{q}] = \{0 \rightarrow \text{skip}, 1 \rightarrow P_1(\theta); \text{abort}\}, (3.1) \\ & \text{while}^{(T \geq 2)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done} \\ \equiv & \text{case } M[\bar{q}] = \{0 \rightarrow \text{skip}, 1 \rightarrow P_1(\theta); \text{while}^{(T-1)}\} \end{aligned}$$

Unparameterized programs can be obtained by fixing $\theta^* \in \mathbb{R}^k$ in some $P(\theta)$. We denote the set of variables accessible to $P(\theta)$ as $\text{qVar}(P(\theta))$; the collection of all “ T -bounded while-programs $P(\theta)$ s.t. $\text{qVar}(P(\theta)) = \bar{v}$ ” as $\mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$, and similarly, the unparameterized one as $\mathbf{q}\text{-while}_{\bar{v}}^{(T)}$.

Now let us formally define *parameterization* of unitaries: let $\theta := (\theta_1, \dots, \theta_k)$ ($k \geq 1$). A k -parameterized unitary $U(\theta)$ is a function $\mathbb{R}^k \rightarrow L(\mathcal{H})$ s.t. (1) given any $\theta^* \in \mathbb{R}^k$, $U(\theta^*)$ is an unitary on \mathcal{H} , and (2) the parameterized-matrix representation of $U(\theta)$ is entry-wise smooth.

A important family of examples is the *single-qubit rotations* about the Pauli axis X, Y, Z with angle θ :

$$R_{\sigma}(\theta) := \exp\left(\frac{-i\theta}{2}\sigma\right), \sigma \in \{X, Y, Z\} \quad (3.2)$$

One can also extend Pauli rotations to multiple qubits. For example, consider *two-qubit coupling* gates $\{R_{\sigma \otimes \sigma} := \exp(\frac{-i\theta}{2}\sigma \otimes \sigma)\}_{\sigma \in \{X, Y, Z\}}$. Note that these two-qubit gates can generate entanglement between two qubits. Combined with single-qubit rotations, they form a *universal gate set* for quantum computation. Another important feature of these gates is that they can already be reliably implemented in experiments, such as ion-trap quantum computers [56].

As a result, we will work mostly with these gates in the rest of this paper. However, note that one can easily add and study other parameterized gates in our framework as well.

The language constructed above is similar to their classical counterparts. (0) **abort** terminates the program, outputting $0 \in \mathcal{D}(\mathcal{H}_{\bar{q}})$. (1) **skip** does nothing to states in $\mathcal{D}(\mathcal{H}_{\bar{q}})$. (2) $q := |0\rangle$ sets quantum variable q to the basis state $|0\rangle$. (3) for any $\theta^* \in \mathbb{R}^k$, $\bar{q} := U(\theta^*)[\bar{q}]$ applies the unitary $U(\theta^*)$ to the qubits in \bar{q} . (4) Sequencing has the same behavior as its classical counterpart. (5) for $\theta^* \in \mathbb{R}^k$, **case** $M[\bar{q}] = \bar{m} \rightarrow P_m(\theta^*)$ **end** performs the measurement $M = \{M_m\}$ on the qubits in \bar{q} , and executes program $\overline{P_m(\theta^*)}$ if the outcome of the measurement is m . The bar over $\bar{m} \rightarrow P_m$ indicates that there may be one or more repetitions of this expression. (6) **while**^(T) $M[\bar{q}] = 1$ **do** $P_1(\theta^*)$ **done** performs the measurement $M = \{M_0, M_1\}$ on \bar{q} , and terminates if the outcome corresponds to M_0 , or executes $P_1(\theta^*)$ then reiterates ($T \geq 2$) / aborts ($T = 1$) otherwise. The program iterates at most T times.

We highlight two differences between quantum and classical while languages: (1) Qubits may only be initialized to the state $|0\rangle$. There is no quantum analogue for initialization to any expression (i.e. $x := e$) due to 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 $|0\rangle$. (2) Evaluating the guard of a case statement or loop, which performs a measurement, potentially disturbs the state of the system.

3.2 Operational and Denotational Semantics

We present the *operational semantics* of parameterized programs in Figure 1a. Transition rules are represented as $\langle P, \rho \rangle \rightarrow \langle P', \rho' \rangle$, where $\langle P, \rho \rangle$ and $\langle P', \rho' \rangle$ are quantum *configurations*.² In configurations, P (or P') could be a quantum program or the empty program \downarrow , 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 \downarrow) 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.³ For m , the superoperator \mathcal{E}_m is defined by $\mathcal{E}_m(\rho) = M_m \rho M_m^\dagger$, yielding the post-measurement state.

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 $\mathcal{E}_{q \rightarrow 0}^{\text{B-int}}(\rho) = \sum_{n=-N}^N |0\rangle_q \langle n| \rho |n\rangle_q \langle 0|$ ($N \in \mathbb{Z}^+$). 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 conventional 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 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 .

We present the *denotational semantics* of parameterized programs in 1b, defining $\llbracket P \rrbracket$ as a superoperator on $\rho \in \mathcal{H}_{\bar{v}}$ [53]. For more details we refer the reader to Ying [52, 53].

We have the following connection between the denotational semantics and operational for parameterized programs: in short, the meaning of running program $P(\theta^*)$ on input state ρ and any $\theta^* \in \mathbb{R}^k$ is the sum of all possible output states with multiplicity, weighted by their probabilities.

Proposition 3.1 ([53]). $\forall P(\theta) \in \mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$, and any specific $\theta^* \in \mathbb{R}^k$, $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$,

$$\llbracket P(\theta^*) \rrbracket(\rho) = \sum \{ |\rho' : (P(\theta^*), \rho) \rightarrow^* (\downarrow, \rho') \} \}. \quad (3.3)$$

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

We close the section with a notion arising from the following observation: some programs, while syntactically not

²Recall that, fixing arbitrary $\theta^* \in \mathbb{R}^k$, both semantics reduce to those of unparameterized programs, so for compactness we write P for $P(\theta^*)$, etc.

³If we had instead considered a probabilistic transition system, then one could write $\langle \text{case } M[\bar{q}] = \bar{m} \rightarrow P_m \text{ end}, \rho \rangle \xrightarrow{P_m} \langle P_m, \rho_m \rangle$ where $p_m = \text{tr}(M_m \rho M_m^\dagger)$ and $\rho_m = M_m \rho M_m^\dagger / p_m$.

$$\begin{aligned}
& \text{(Abort)} \quad \overline{\langle \text{abort}[\bar{q}], \rho \rangle} \rightarrow \langle \downarrow, \mathbf{0} \rangle \\
& \text{(Skip)} \quad \overline{\langle \text{skip}[\bar{q}], \rho \rangle} \rightarrow \langle \downarrow, \rho \rangle \\
& \text{(Initialization)} \quad \overline{\langle q := |0\rangle, \rho \rangle} \rightarrow \langle \downarrow, \rho_0^q \rangle \\
& \quad \text{where } \rho_0^q = \begin{cases} \mathcal{E}_{q \rightarrow 0}^{\text{bool}}(\rho) & \text{if } \text{type}(q) = \mathbf{Bool} \\ \mathcal{E}_{q \rightarrow 0}^{\text{B-int}}(\rho) & \text{if } \text{type}(q) = \mathbf{Bdd Int} \end{cases} \\
& \text{(Unitary)} \quad \overline{\langle \bar{q} := U(\boldsymbol{\theta}^*)[\bar{q}], \rho \rangle} \rightarrow \langle \downarrow, U(\boldsymbol{\theta}^*)\rho U^\dagger(\boldsymbol{\theta}^*) \rangle \\
& \quad \overline{\langle P_1(\boldsymbol{\theta}^*), \rho \rangle} \rightarrow \langle P_1'(\boldsymbol{\theta}^*), \rho' \rangle \\
& \text{(Sequence)} \quad \overline{\langle P_1(\boldsymbol{\theta}^*); P_2(\boldsymbol{\theta}^*), \rho \rangle} \rightarrow \langle P_1'(\boldsymbol{\theta}^*); P_2(\boldsymbol{\theta}^*), \rho' \rangle \\
& \text{(Case } m) \quad \overline{\langle \text{case } M[\bar{q}] = m \rightarrow P_m(\boldsymbol{\theta}^*) \text{ end}, \rho \rangle} \rightarrow \\
& \quad \langle P_m(\boldsymbol{\theta}^*), \mathcal{E}_m(\rho) \rangle, \forall \text{ outcome } m \text{ of } M = \{M_m\} \\
& \text{(While}^{(T)} 0) \quad \overline{\langle \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\boldsymbol{\theta}^*) \text{ done}, \rho \rangle} \rightarrow \\
& \quad \langle \downarrow, \mathcal{E}_0(\rho) \rangle \\
& \text{(While}^{(T)} 1) \quad \overline{\langle \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\boldsymbol{\theta}^*) \text{ done}, \rho \rangle} \rightarrow \\
& \quad \langle P_1(\boldsymbol{\theta}^*); \text{while}^{(T-1)}, \mathcal{E}_1(\rho) \rangle
\end{aligned}$$

(a)

$$\begin{aligned}
& \llbracket \text{abort}[\bar{q}] \rrbracket \rho = \mathbf{0} \\
& \llbracket \text{skip}[\bar{q}] \rrbracket \rho = \rho \\
& \llbracket q := |0\rangle \rrbracket \rho = \mathcal{E}_{q \rightarrow 0}^{\text{bool}} \rho \text{ or } \mathcal{E}_{q \rightarrow 0}^{\text{B-int}}(\rho) \\
& \llbracket \bar{q} := U(\boldsymbol{\theta}^*)[\bar{q}] \rrbracket \rho = U(\boldsymbol{\theta}^*)\rho U^\dagger(\boldsymbol{\theta}^*) \\
& \llbracket P_1(\boldsymbol{\theta}^*); P_2(\boldsymbol{\theta}^*) \rrbracket \rho = \llbracket P_2(\boldsymbol{\theta}^*) \rrbracket (\llbracket P_1(\boldsymbol{\theta}^*) \rrbracket \rho) \\
& \llbracket \text{case } M[\bar{q}] = m \rightarrow P_m(\boldsymbol{\theta}^*) \text{ end} \rrbracket \rho = \sum_m \llbracket P_m(\boldsymbol{\theta}^*) \rrbracket \mathcal{E}_m \rho \\
& \llbracket \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\boldsymbol{\theta}^*) \text{ done} \rrbracket \rho = \\
& \quad \sum_{n=0}^{T-1} \mathcal{E}_0 \circ (\llbracket P_1(\boldsymbol{\theta}^*) \rrbracket \circ \mathcal{E}_1)^n \rho
\end{aligned}$$

(b)

Figure 1. Parameterized T -bounded quantum **while** programs: (a) operational semantics (b) denotational semantics. We instantiate $\boldsymbol{\theta}^* \in \mathbb{R}^k$ since choice of parameters affects choices made at control gates. “**Bdd**” stands for **Bounded**.

“**abort**[\(\bar{q}\)]”, semantically aborts. Simple examples include $U(\boldsymbol{\theta})$; **abort** or a case sentence that has **abort** on each branch.

We formalize this concept (essential-abortion for unparameterized programs may be analogously defined):

Definition 3.2 (“Essentially Abort”). Let $P(\boldsymbol{\theta}) \in \mathbf{q\text{-}while}^{(T)}(\boldsymbol{\theta})$. $P(\boldsymbol{\theta})$ “essentially aborts” if one of the following holds:

1. $P(\boldsymbol{\theta}) \equiv \text{abort}[\bar{q}]$;
2. $P(\boldsymbol{\theta}) \equiv P_1(\boldsymbol{\theta}); P_2(\boldsymbol{\theta})$, and either $P_1(\boldsymbol{\theta})$ or $P_2(\boldsymbol{\theta})$ essentially aborts;
3. $P \equiv \text{case } M[\bar{q}] = m \rightarrow P_m(\boldsymbol{\theta}) \text{ end}$, and each $P_m(\boldsymbol{\theta})$ essentially aborts.

$$\begin{aligned}
& \text{(Sum Components)} \quad \overline{\langle P_1(\boldsymbol{\theta}^*) + P_2(\boldsymbol{\theta}^*), \rho \rangle} \rightarrow \langle P_1(\boldsymbol{\theta}^*), \rho \rangle, \\
& \quad \overline{\langle P_1(\boldsymbol{\theta}^*) + P_2(\boldsymbol{\theta}^*), \rho \rangle} \rightarrow \langle P_2(\boldsymbol{\theta}^*), \rho \rangle
\end{aligned}$$

Figure 2. additive parameterized quantum bounded while programs: operational semantics. We fix $\boldsymbol{\theta}^* \in \mathbb{R}^k$ and inherit all the other rules from parameterized programs in Fig. 1a.

4 Additive Parameterized Quantum Bounded While-Programs

As we discussed in the introduction, we will introduce a variant of *additive* quantum programs as a succinct way to describe the collection of programs that are necessary to compute the derivatives. To that end, we introduce our design of the syntax and the semantics of additive quantum programs as well as a compilation method that turns any additive quantum program into a collection of normal programs for the actual computation of derivatives.

4.1 Syntax

We adopt the convention to use underlines to indicate additive programs, such as $\underline{P(\boldsymbol{\theta})}$, to distinguish from normal program $P(\boldsymbol{\theta})$. The syntax of $\underline{P(\boldsymbol{\theta})}$ is given by

$$\begin{aligned}
\underline{P(\boldsymbol{\theta})} ::= & \underline{\text{abort}[\bar{q}]} \mid \underline{\text{skip}[\bar{q}]} \mid \underline{q := |0\rangle} \mid \underline{\bar{q} := U(\boldsymbol{\theta})[\bar{q}]} \mid \\
& \underline{P_1(\boldsymbol{\theta}); P_2(\boldsymbol{\theta})} \mid \underline{\text{case } M[\bar{q}] = m \rightarrow P_m(\boldsymbol{\theta}) \text{ end}} \mid \\
& \underline{\text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\boldsymbol{\theta}) \text{ done}} \mid \underline{P_1(\boldsymbol{\theta}) + P_2(\boldsymbol{\theta})},
\end{aligned}$$

where the only new syntax $+$ is the *additive* choice. Intuitively, $\underline{P_1(\boldsymbol{\theta}) + P_2(\boldsymbol{\theta})}$ refers to the program that either executes $\underline{P_1(\boldsymbol{\theta})}$ or $\underline{P_2(\boldsymbol{\theta})}$. This intuition will be implemented by the definition of its operational and denotational semantics. We assume $+$ has lower precedence order, and is left associative.⁴ If $\underline{P(\boldsymbol{\theta})} = \underline{P_1(\boldsymbol{\theta}) + P_2(\boldsymbol{\theta})}$, then $\text{qVar}(\underline{P(\boldsymbol{\theta})}) \equiv \text{qVar}(P_1(\boldsymbol{\theta})) \cup \text{qVar}(P_2(\boldsymbol{\theta}))$. Denote the collection of all non-deterministic $\underline{P(\boldsymbol{\theta})}$ s.t. $\text{qVar}(\underline{P(\boldsymbol{\theta})}) = \bar{v}$ as $\mathbf{add-q-while}_{\bar{v}}^{(T)}(\boldsymbol{\theta})$.

4.2 Operational and Denotational Semantics

We exhibit operational semantics in Figure 2 and define a similar denotational semantics for any $\underline{P(\boldsymbol{\theta}^*)} \in \mathbf{add-q-while}_{\bar{v}}^{(T)}(\boldsymbol{\theta})$.

Definition 4.1 (Denotational Semantics). Fix $\boldsymbol{\theta}^* \in \mathbb{R}^k$, $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$,

$$\llbracket \underline{P(\boldsymbol{\theta}^*)} \rrbracket(\rho) \equiv \{|\rho' : \langle \underline{P(\boldsymbol{\theta}^*)}, \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle|\} \quad (4.1)$$

We can now explain the name of $+$. Intuitively, the (S-C) rule in Figure 2 provides different choices of paths which will be eventually summed up in (4.1). This resembles the idea of the sum operator in differential lambda-calculus [12]. Interestingly, some slight change of our (S-C) rule, e.g., in the study of non-deterministic quantum programs [54], can

⁴E.g., $\underline{X + Y}; \underline{Z} = \underline{X + (Y; Z)}$, $X + Y + Z := (X + Y) + Z$.

also be used to describe the interlacing behavior between different components in parallel quantum programs.

4.3 Compilation Rules

We exhibit the compilation rules in Figure 3 as a way to transform a additive program $\underline{P}(\theta)$ into a multiset of normal programs. The compiled set of programs will be later used in the actual implementation of the differentiation procedure. We will show, in Section 7, that our compilation procedure is also efficient in the sense that the total number of non-essentially-aborting program in $\text{Compile}(\underline{P}_m(\theta))$ is reasonably bounded, in particular for the “fill and break” algorithm (FB(\bullet)) as the compilation of the case statement.

Our compilation rule is well-defined as it is compatible with the denotational semantics and operational semantics of $\underline{P}(\theta)$ in the following sense:

Proposition 4.2. *Denoting with \sqcup the union of multisets, then for any $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$,*

$$\{|\rho' : \rho' \neq \mathbf{0}, \rho' \in \llbracket \underline{P}(\theta^*) \rrbracket \rho\} = \bigsqcup_{Q(\theta) \in \text{Compile}(\underline{P}(\theta))} \{|\rho' \neq \mathbf{0} : \langle Q(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle\}. \quad (4.2)$$

Proof. Structural Induction. See Appendix C.1 for details. \square

Note that (4.2) removes $\mathbf{0}$ from the multi-set as we are only interested in non-trivial final states. Moreover, in $\text{Compile}(\underline{P}(\theta))$, some programs may essentially abort (Definition 3.2). For implementation, we are interested in the number of $Q(\theta) \in \text{Compile}(\underline{P}(\theta))$ that do not essentially abort:

Definition 4.3 (Number of Non-Aborting Programs). *The number of non-aborting programs of $\underline{P}(\theta)$, denoted as $\#P(\theta)$, is defined as*

$$\#P(\theta) = |\text{Compile}(\underline{P}(\theta)) \setminus \{Q(\theta) \in \text{Compile}(\underline{P}(\theta)) : Q(\theta) \text{ essentially aborts.}\}|$$

where $|C|$ is the cardinality of a multiset C and $C_0 \setminus C_1$ denotes the multiset difference of C_0 and C_1 .

Example 4.1 (Generic-Case). *Consider the following simple program with the case statement*

$$\underline{P}(\theta) \equiv \text{case } M[\bar{q}] = 0 \rightarrow \underline{P}_1(\theta) + \underline{P}_2(\theta), \\ 1 \rightarrow \underline{P}_3(\theta)$$

where $P_1(\theta), P_2(\theta), P_3(\theta) \in \mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$, none of them essentially aborts, and each of $P_1(\theta), P_2(\theta), P_3(\theta)$ contains no control

- (Atomic) $\text{Compile}(\underline{P}(\theta)) \equiv \{|\underline{P}(\theta)|\},$
if $\underline{P}(\theta) \equiv \text{abort}[\bar{v}] \mid \text{skip}[\bar{v}] \mid \underline{q} := |0\rangle$
 $|\bar{v} := U(\theta)[\bar{v}]$.
- (Sequence) $\text{Compile}(\underline{P}_1(\theta); \underline{P}_2(\theta)) \equiv$
 $\begin{cases} \{|\text{abort}|\}, & \text{if } \text{Compile}(\underline{P}_1(\theta)) = \{|\text{abort}|\}; \\ \{|\text{abort}|\}, & \text{if } \text{Compile}(\underline{P}_2(\theta)) = \{|\text{abort}|\}; \\ \{|\underline{Q}_1(\theta); \underline{Q}_2(\theta) : \underline{Q}_b(\theta) \in \text{Compile}(\underline{P}_b(\theta))|\}, & \text{otherwise.} \end{cases}$
- (Case m) $\text{Compile}(\text{case}) \equiv \text{FB}(\text{case})$, described in Fig.3b.
- (While^(T)) $\text{Compile}(\text{while}^{(T)})$: use (Case m) and (Sequence).
- (Sum) $\text{Compile}(\underline{P}_1(\theta) + \underline{P}_2(\theta)) \equiv$
 $\begin{cases} \text{Compile}(\underline{P}_1(\theta)) \sqcup \text{Compile}(\underline{P}_2(\theta)), & \text{if } \forall b \in \{1, 2\}, \text{Compile}(\underline{P}_b(\theta)) \neq \{|\text{abort}|\}; \\ \text{Compile}(\underline{P}_1(\theta)), & \text{if } \text{Compile}(\underline{P}_2(\theta)) = \{|\text{abort}|\}, \\ & \text{Compile}(\underline{P}_1(\theta)) \neq \{|\text{abort}|\}; \\ \text{Compile}(\underline{P}_2(\theta)), & \text{if } \text{Compile}(\underline{P}_1(\theta)) = \{|\text{abort}|\}, \\ & \text{Compile}(\underline{P}_2(\theta)) \neq \{|\text{abort}|\}; \\ \{|\text{abort}|\}, & \text{otherwise} \end{cases}$
- (a)
1. $\forall m \in [0, w]$, let C_m denote the sub-multiset of $\text{Compile}(\underline{P}_m(\theta))$ composed of programs that do *not* essentially abort; without loss of generality, assume $|C_0| \geq |C_1| \geq \dots \geq |C_w|$.
 2. If all C_m 's are empty, return $\text{FB}(\text{case}) \equiv \{|\text{abort}[\bar{v}]|\}$; else, pad each C_m to size $|C_0|$ by adding “ $\text{abort}[\bar{v}]$ ”.
 3. $\forall m \in [0, w]$, index programs in C_m as $\{|\underline{Q}_{m,0}(\theta), \dots, \underline{Q}_{m,|C_0|-1}(\theta)|\}$.
 4. Return $\text{FB}(\text{case}) \equiv \{|\text{case } M[\bar{q}] = \overline{m} \rightarrow \underline{Q}_{m,j^*} \text{ end } |\}_{j^*}$ with $0 \leq j^* \leq |C_0| - 1$.
- (b)

Figure 3. nondeterministic programs: (a) compilation rules. (b) “Fill and Break” (“FB(\bullet)”) procedure for computing $\text{Compile}(\text{case})$. case stands for case $M[\bar{q}] = \overline{m} \rightarrow \underline{P}_m(\theta)$ end; while^(T) stands for while^(T) $M[\bar{q}] = 1$ do $\underline{P}_1(\theta)$ done. Here \sqcup denotes union of multisets. One may observe from a routine structural induction and the definition of “essentially abort” that: for all $\underline{P}(\theta)$, either $\text{Compile}(\underline{P}(\theta)) = \{|\text{abort}|\}$, or $\text{Compile}(\underline{P}(\theta))$ does not contain essentially abort programs.

gates. Then for any $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$, fixing θ^* we have

$$\begin{aligned} \langle P(\theta^*), \rho \rangle &\xrightarrow{(\text{Case } m)} \langle \underline{P}_1(\theta^*) + \underline{P}_2(\theta^*), M_0 \rho M_0^\dagger \rangle \\ &\xrightarrow{(\text{Sum})} \langle \underline{P}_1(\theta^*), M_0 \rho M_0^\dagger \rangle \\ &\rightarrow^* \langle \downarrow, \llbracket \underline{P}_1(\theta^*) \rrbracket (M_0 \rho M_0^\dagger) \rangle; \\ \langle P(\theta^*), \rho \rangle &\xrightarrow{(\text{Case } m)} \langle \underline{P}_1(\theta^*) + \underline{P}_2(\theta^*), M_0 \rho M_0^\dagger \rangle \\ &\xrightarrow{(\text{Sum})} \langle \underline{P}_2(\theta^*), M_0 \rho M_0^\dagger \rangle \rightarrow^* \langle \downarrow, \llbracket \underline{P}_2(\theta^*) \rrbracket (M_0 \rho M_0^\dagger) \rangle; \\ \langle P(\theta^*), \rho \rangle &\xrightarrow{(\text{Case } m)} \langle \underline{P}_3(\theta^*), M_1 \rho M_1^\dagger \rangle \\ &\rightarrow^* \langle \downarrow, \llbracket \underline{P}_3(\theta^*) \rrbracket (M_1 \rho M_1^\dagger) \rangle \end{aligned}$$

Hence by Definition 4.1.

$$\begin{aligned} \llbracket P(\theta^*) \rrbracket \rho &= \{ \llbracket P_1(\theta^*) \rrbracket (M_0 \rho M_0^\dagger), \llbracket P_2(\theta^*) \rrbracket (M_0 \rho M_0^\dagger), \\ &\quad \llbracket P_3(\theta^*) \rrbracket (M_1 \rho M_1^\dagger) \} \end{aligned}$$

We verify computation results from the compilation rules are consist with this. Writing “compilation rule” as “CP” for short, one observes $\text{Compile}(P_1(\theta) + P_2(\theta)) \stackrel{\text{CP, Sum}}{=} \{P_1(\theta), P_2(\theta)\}$, while $\text{Compile}(P_3(\theta)) = \{P_3(\theta)\}$ since we assumed non-essentially-abortness. Apply our “fill and break” procedure to obtain $C_0 = \{P_1(\theta), P_2(\theta)\}$, $C_1 = \{P_3(\theta), \text{abort}[\bar{v}]\}$.

$$\text{Compile}(P(\theta)) = \left\{ \begin{array}{ll} \text{case } M[\bar{q}] = 0 \rightarrow & P_1(\theta), \\ & 1 \rightarrow P_2(\theta), \\ \text{case } M[\bar{q}] = 0 \rightarrow & P_2(\theta), \\ & 1 \rightarrow \text{abort}[\bar{v}] \end{array} \right\}$$

Evolving pursuant to the normal programs operational semantics (Fig 1) leads to the follow that agrees with $\llbracket P(\theta^*) \rrbracket \rho$.

$$\begin{aligned} &\coprod_{Q(\theta) \in \text{Compile}(P(\theta))} \{ \rho' : \langle Q(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle \} \\ &= \{ \llbracket P_1(\theta^*) \rrbracket (M_0 \rho M_0^\dagger), \llbracket P_2(\theta^*) \rrbracket (M_0 \rho M_0^\dagger), \\ &\quad \llbracket P_3(\theta^*) \rrbracket (M_1 \rho M_1^\dagger), 0 \}. \end{aligned}$$

5 Observable and Differential Semantics

To capture physically observable quantities from quantum systems, physicists propose the notation of *observable* which is a Hermitian matrix over the same space. Any observable O is a combination of information about quantum measurements and classical values for each measurement outcome. To see why, let us take its spectral decomposition of $O = \sum_m \lambda_m |\psi_m\rangle\langle\psi_m|$. Then $\{|\psi_m\rangle\langle\psi_m|\}_m$ form a projective measurement. We can design an experiment to perform this projective measurement and output λ_m when the outcome is m . The *expectation* of the output is exactly given by

$$\text{tr}(O\rho) = \sum_m a_m \text{tr}(M_m \rho M_m^\dagger). \quad (5.1)$$

Comparing to the random outcome of any quantum measurement, the expectation $\text{tr}(O\rho)$ given by any observable O represents meaningful classical information of quantum systems. In applications of quantum machine learning, it is also the quantity used in the loss function. Thus, given any observable O , we will define the *observable semantics* of quantum programs both the mathematical object to take derivatives from the original programs and the read-out of the programs that compute these derivatives.

Moreover, one should note the extra cost to approximate $\text{tr}(O\rho)$ to high precision. For example, one can repeat the aforementioned experiment with $\{|\psi_m\rangle\langle\psi_m|\}_m$ measurement and use the statistical information to recover $\text{tr}(O\rho)$. The number of iterations will depend on the additive precision δ and the norm of O . To simplify our presentation, also to make a precise resource count as detailed in Section 7, throughout

the paper, we assume that⁵

$$-I_{\mathcal{H}} \sqsubseteq O \sqsubseteq I_{\mathcal{H}}. \quad (5.2)$$

By statistically concentration bounds (e.g. the Chernoff bound), to approximate $\text{tr}(O\rho)$ with additive error δ , one needs to repeat $O(1/\delta^2)$ times with $O(1/\delta^2)$ copies of initial states.

5.1 Observable Semantics

We define the *observable semantics* of both normal (denoted by $P(\theta)$, $P'(\theta)$) and additive (denoted by $\underline{S}(\theta)$, $\underline{S}'(\theta)$) parameterized programs as follows.

Definition 5.1 (Observable Semantics). $\forall P(\theta) \in \mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$, any observable $O \in \mathcal{O}_{\bar{v}}$ and input state $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$, the observable semantics of P , denoted $\llbracket (O, \rho) \rightarrow P(\theta) \rrbracket$, is

$$\llbracket (O, \rho) \rightarrow P(\theta) \rrbracket (\theta^*) \equiv \text{tr}(O \llbracket P(\theta^*) \rrbracket \rho), \forall \theta^* \in \mathbb{R}^k. \quad (5.3)$$

Namely, $\llbracket (O, \rho) \rightarrow P(\theta) \rrbracket$ is a function from \mathbb{R}^k to \mathbb{R} whose value per point is given by (5.3).

Similarly, for any $\underline{S}(\theta) \in \mathbf{add}\text{-}\mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$ with $\text{Compile}(\underline{S}(\theta)) = \{P_i(\theta)\}_{i=1}^t$ where $P_i(\theta) \in \mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$, its observable semantics is given by

$$\llbracket (O, \rho) \rightarrow \underline{S}(\theta) \rrbracket (\theta^*) \equiv \sum_{i \in [1, t]} \llbracket (O, \rho) \rightarrow P_i(\theta) \rrbracket (\theta^*), \forall \theta^* \in \mathbb{R}^k. \quad (5.4)$$

To compute gradients of quantum observables for each parameter, one needs an ancilla variable as hinted by results in quantum information theory about gradient calculations for simple unitaries (e.g., Bergholm et al. [7], Schuld et al. [41]). To that end, we can easily extend quantum programs with ancilla variables. For each $j \in [1, k]$, the j -th *ancilla* of $\mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$ is a quantum variable denoted by $A_{j, \bar{v}}$ disjoint from \bar{v} . We write A instead of $A_{j, \bar{v}}$ when j, \bar{v} are clear from context. Ancilla A could consists of any number of qubits while we will mostly use *one-qubit* A in this paper.

We will only consider programs augmented with *one ancilla variable* A_j at any time. (So let us fix j for the following discussion). We will then consider programs that operate on the larger space $\mathcal{D}(\mathcal{H}_{\bar{v} \cup \{A\}})$ and an additional observable A to define the *observable semantics with ancilla*.

Definition 5.2 (Observable Semantics with Ancilla). Given any $P'(\theta) \in \mathbf{q}\text{-while}_{\bar{v} \cup \{A_{j, \bar{v}}\}}^{(T)}(\theta)$, any observable $O \in \mathcal{O}_{\bar{v}}$, input state $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$, and moreover the observable O_A on ancilla A , the observable semantics with ancilla of P , overloading the notation $\llbracket (O, \rho) \rightarrow P'(\theta) \rrbracket$, is

$$\begin{aligned} &\llbracket ((O, O_A), \rho) \rightarrow P'(\theta) \rrbracket (\theta^*) \equiv \\ &\text{tr} \left((O_A \otimes O) \llbracket P'(\theta^*) \rrbracket (|\bar{0}\rangle_{\{A\}} \langle \bar{0}|) \otimes \rho \right), \forall \theta^* \in \mathbb{R}^k. \end{aligned}$$

Again, $\llbracket ((O, O_A), \rho) \rightarrow P(\theta) \rrbracket$ is a function from \mathbb{R}^k to \mathbb{R} whose value per point is given by (5.5).

⁵ \sqsubseteq is defined by $A \sqsubseteq B \iff B - A$ positive semidefinite. See Appendix C.1.

Similarly, for $S'(\theta) \in \text{add-q-while}_{\bar{v} \cup \{A_j, \bar{v}\}}^{(T)}(\theta)$ such that $\text{Compile}(S'(\theta)) = \{|P'_i(\theta)|\}_{i=1}^t$ where $P'_i(\theta) \in \text{q-while}_{\bar{v} \cup \{A_j, \bar{v}\}}^{(T)}(\theta)$, its observable semantics is: $\forall \theta^* \in \mathbb{R}^k$,

$$\llbracket ((O, O_A), \rho) \rightarrow S'(\theta) \rrbracket(\theta^*) \equiv \sum_{i \in [1, t]} \llbracket ((O, O_A), \rho) \rightarrow P'_i(\theta) \rrbracket(\theta^*) \quad (5.5)$$

The only difference from the normal observable semantics lies in (5.5), where we initialize the ancilla with $|0\rangle$, which is a natural choice and evaluate the observable $O_A \otimes O$. As we will see in the technique, the independence between O_A and O in the form of $O_A \otimes O$ will help use obtain the strongest guarantee of our differentiation procedure.

5.2 Differential Semantics

Given the definition of observable semantics, its differential semantics can be naturally defined by

Definition 5.3 (Differential Semantics). *Given additive program $S(\theta) \in \text{add-q-while}_{\bar{v}}^{(T)}(\theta)$, its j -th differential semantics is defined by*

$$\frac{\partial}{\partial \theta_j} (\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket), \quad (5.6)$$

which is again a function from \mathbb{R}^k to \mathbb{R} . Moreover, for any $S'(\theta) \in \text{add-q-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$ with ancilla A , we say that “ $S'(\theta)$ computes the j -th differential semantics of $S(\theta)$ ” if and only if there exists an observable O_A on ancilla A for $S'(\theta)$ such that $\forall O \in \mathcal{O}_{\bar{v}}, \rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$,

$$\llbracket ((O, O_A), \rho) \rightarrow S'(\theta) \rrbracket = \frac{\partial}{\partial \theta_j} (\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket). \quad (5.7)$$

We remark that (5.6) is well defined because $\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket$ is a function from \mathbb{R}^k to \mathbb{R} . It is also a smooth function because we assume that parameterized unitaries are entry-wise smooth, and the observable semantics is obtained by multiplication and addition of such entries.

We also remark that the order of quantifiers in (5.7) makes the predicate the strongest that one can hope for. This is because the observable semantics of $S(\theta)$ will depend on O and ρ in general. Thus, the program to compute its differential semantics could also depend on O and ρ in general. However, in our definition, $S'(\theta)$ is a single fixed program that works for any O and ρ regardless of the seemingly complicated relationship. We remark that this definition is consistent with the classical case where a single program is used to compute the derivatives for any input. Indeed, we can achieve this definition in the quantum setting and it is also critical in the proof of Theorem 6.2 (item (5)).

6 Code Transformations and the Differentiation Logic

We describe the code transformation rules of the differentiation operator $\frac{\partial}{\partial \theta}(\cdot)$ in Section 6.1. We also define a logic and prove its soundness for reasoning about the correctness of these code transformations, with the following judgement

$$\underline{S'(\theta)} | \underline{S(\theta)}, \quad (6.1)$$

which states that $S'(\theta)$ computes the differential semantics of $S(\theta)$ in sense of Definition 5.3. We fix $\theta = \theta_j$ and hence A stands for A_j, \bar{v} and $\frac{\partial}{\partial \theta}$ for $\frac{\partial}{\partial \theta_j}$ through this section.⁶

6.1 Code Transformations

We first define some gates associated with the single-qubit rotation and the two-qubit coupling gates, which will appear in the code transformation rules. Let A be a single qubit.

Definition 6.1. 1. Consider $|q_1\rangle := R_\sigma(\theta)|q_1\rangle$, where $\sigma \in \{X, Y, Z\}$. Define

$$C - R_\sigma(\theta)|A, q_1\rangle \equiv \left((|0\rangle_A \langle 0|) \otimes R_\sigma(\theta) \right) + \quad (6.2)$$

$$\left((|1\rangle_A \langle 1|) \otimes R_\sigma(\theta + \pi) \right) \quad (6.3)$$

$$R'_\sigma(\theta)|A, q_1\rangle \equiv A := H|A\rangle; C - R_\sigma(\theta)|A, q_1\rangle; A := H|A\rangle. \quad (6.4)$$

2. Substituting $\sigma \otimes \sigma$ for σ and q_1, q_2 for q_1 in Eqns (6.3, 6.4), one defines $C - R_{\sigma \otimes \sigma}(\theta), R'_{\sigma \otimes \sigma}(\theta)$.

For 1-qubit rotation $R_\sigma(\theta)$, the “controlled-rotation” gate $C - R_\sigma(\theta)$ maps $|0, q_1\rangle \mapsto |0\rangle \otimes R_\sigma(\theta)|q_1\rangle$, and $|1, q_1\rangle \mapsto |1\rangle \otimes R_\sigma(\theta + \pi)|q_1\rangle$; $R'_\sigma(\theta)$ conjugates $C - R_\sigma(\theta)$ with Hadamard. Similarly for corresponding two-qubit coupling gates.

We exhibit our code transformation rules in Figure 4. For Unitary rules we only include 1-qubit rotations and two-qubit coupling gates, since they form a universal gate set and are easy to implement on quantum machines. It is also possible to include more unitary rules (e.g., by following the calculations in [41]), which we will leave as future directions.

6.2 The differentiation logic and its soundness

We develop the differentiation logic given in Figure 5 to reason about the correctness of code transformations. It suffices to show that our logic is sound. For ease of notation, in future analysis we write $\frac{\partial}{\partial \theta}(P(\theta))$ in place of $\frac{\partial}{\partial \theta_j}(P(\theta))$ when $P(\theta) \in \text{q-while}_{\bar{v}}^{(T)}(\theta)$.

Theorem 6.2 (Soundness). *Let $S(\theta) \in \text{add-q-while}_{\bar{v}}^{(T)}(\theta)$, $S'(\theta) \in \text{add-q-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$. Then,*

$$\underline{S'(\theta)} | \underline{S(\theta)} \implies \underline{S'(\theta)} \text{ computes the differential semantics of } \underline{S(\theta)}. \quad (6.5)$$

⁶If A already exists, i.e., $S(\theta) \in \text{add-q-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$, we treat $\bar{v}_{\text{new}} \leftarrow \bar{v}_{\text{old}} \cup A_{\text{old}}$ and add A_{new} . Any observable O on \bar{v}_{old} becomes $O_{A_{\text{old}}} \otimes O$ on \bar{v}_{new} . Both A_{old} and A_{new} are initialized to $|0\rangle$ in observable semantics.

$$\begin{aligned}
(\text{Trivial}) \quad & \frac{\partial}{\partial \theta}(\underline{\text{abort}}[\bar{v}]), \frac{\partial}{\partial \theta}(\underline{\text{skip}}[\bar{v}]), \frac{\partial}{\partial \theta}(q := |0\rangle) \equiv \underline{\text{abort}}[\bar{v} \cup \{A\}]. \\
(\text{Trivial-Unitary}) \quad & \frac{\partial}{\partial \theta}(\bar{v} := U(\theta)|\bar{v}\rangle) \equiv \underline{\text{abort}}[\bar{v} \cup \{A\}], \text{ if } \theta_j \notin \theta. \\
(1\text{-qb Rotation}) \quad & \frac{\partial}{\partial \theta}(|q_1\rangle := R_\sigma(\theta)|q_1\rangle) \equiv |A, q_1\rangle := R'_\sigma(\theta)|A, q_1\rangle. \\
(2\text{-qb Coupling}) \quad & \frac{\partial}{\partial \theta}(|q_1, q_2\rangle := R_{\sigma \otimes \sigma}(\theta)|q_1\rangle) \equiv |A, q_1, q_2\rangle := R'_{\sigma \otimes \sigma}(\theta)|A, q_1, q_2\rangle. \\
(\text{Sequence}) \quad & \frac{\partial}{\partial \theta}(\underline{S_1(\theta); S_2(\theta)}) \equiv \underline{S_1(\theta); \frac{\partial}{\partial \theta}(S_2(\theta))} + \underline{\frac{\partial}{\partial \theta}(S_1(\theta); S_2(\theta))}. \\
(\text{Case}) \quad & \frac{\partial}{\partial \theta}(\underline{\text{case } M[\bar{q}] = m \rightarrow S_m(\theta) \text{ end}}) \equiv \underline{\text{case } M[\bar{q}] = m \rightarrow \frac{\partial}{\partial \theta}(S_m(\theta)) \text{ end}}. \\
(\text{while}^{(T)}) \quad & \text{Use (Case) and (Sequence).} \\
(\text{Sum Components}) \quad & \frac{\partial}{\partial \theta}(\underline{S_1(\theta)} + \underline{S_2(\theta)}) \equiv \frac{\partial}{\partial \theta}(\underline{S_1(\theta)}) + \frac{\partial}{\partial \theta}(\underline{S_2(\theta)})
\end{aligned}$$

Figure 4. Code Transformation Rules. For (1-qb Rotation) and (2-qb Coupling), $(\sigma \in \{X, Y, Z\})$; $R'_\sigma(\theta), R'_{\sigma \otimes \sigma}(\theta)$ are as in Definition 6.1. $\theta_j \notin \theta$ means “the unitary $U(\theta)$ trivially uses θ_j ”: for example in $P(\theta)|q_1\rangle \equiv R_X(\theta_1); R_Z(\theta_2)$, $\theta = (\theta_1, \theta_2)$ and $R_X(\theta_1)$ trivially uses θ_2 .

Let us highlight the ideas behind the proof of the soundness and all detailed proofs are deferred to Appendix D. First remember that $\theta = \theta_j$ and for all the proofs we can choose $Z_A = |0\rangle\langle 0| - |1\rangle\langle 1|$ as the observable on the one-qubit ancilla A . Thus, we will omit Z_A and overload the notation, $\forall P'(\theta) \in \mathbf{q}\text{-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$:

$$\llbracket (O, \rho) \rightarrow P'(\theta) \rrbracket \text{ means } \llbracket ((O, Z_A), \rho) \rightarrow P'(\theta) \rrbracket, \quad (6.6)$$

to simplify the presentation. We make similar overloading convention for $S'(\theta) \in \mathbf{add-q-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$. Let us go through these logic rules one by one.

1. *Abort, Skip, Initialization, Trivial-Unitary* rules work because these statements do not depend on θ .
2. Since $\text{While}^{(T)}$ can be deemed as a macro of other statements, the correctness of $\text{While}^{(T)}$ rule follows by unfolding $\text{while}^{(T)}$ and applying other rules.
3. The *Sum Component* rule is due to the property of observable semantics $\llbracket \cdot \rrbracket$ and additive operator $(+)$:

$$\frac{\partial}{\partial \theta}(\llbracket P_1 + P_2 \rrbracket) = \llbracket \frac{\partial}{\partial \theta}(P_1) \rrbracket + \llbracket \frac{\partial}{\partial \theta}(P_2) \rrbracket, \quad (6.7)$$

which follows from our definition design.

4. Our *Rot-Couple* rule is different from the phase-shift rule in [41] by using only one circuit in derivative computing. However, the proof of the *Rot-Couple* rule is largely inspired by the one of the phase-shift rule.
5. The proof of the *Sequence* rule relies very non-trivially on our design of the observable semantics with ancilla

$$\begin{aligned}
(\text{Abort}) \quad & \frac{\partial}{\partial \theta}(\underline{\text{abort}}[\bar{v}])|\underline{\text{abort}}[\bar{v}] \quad (\text{Skip}) \quad \frac{\partial}{\partial \theta}(\underline{\text{skip}}[\bar{v}])|\underline{\text{skip}}[\bar{v}] \\
(\text{Initialization}) \quad & \frac{\partial}{\partial \theta}(q := |0\rangle)|q := |0\rangle \\
& \theta_j \notin \theta \\
(\text{Trivial-Unitary}) \quad & \frac{\partial}{\partial \theta}(\bar{q} = U(\theta)|\bar{v}\rangle)|\bar{q} = U(\theta)|\bar{v}\rangle \\
(\text{Rot-Couple}) \quad & \frac{\partial}{\partial \theta}(\bar{q} = R_\sigma(\theta)|\bar{v}\rangle)|\bar{q} = R_\sigma(\theta)|\bar{v}\rangle \\
& \frac{\partial}{\partial \theta}(S_0(\theta))|S_0(\theta) \quad \frac{\partial}{\partial \theta}(S_1(\theta))|S_1(\theta) \\
(\text{Sequence}) \quad & \frac{\partial}{\partial \theta}(S_0(\theta); S_1(\theta))|(S_0(\theta); S_1(\theta)) \\
& \forall m, \frac{\partial}{\partial \theta}(S_m(\theta))|S_m(\theta) \\
(\text{Case}) \quad & \frac{\partial}{\partial \theta}(\underline{\text{case } M[\bar{q}] = m \rightarrow S_m(\theta) \text{ end}})|\underline{\text{case } M[\bar{q}] = m \rightarrow S_m(\theta) \text{ end}} \\
& \frac{\partial}{\partial \theta}(S_1(\theta))|S_1(\theta) \\
(\text{While}^{(T)}) \quad & \frac{\partial}{\partial \theta}(\underline{\text{while}^{(T)} } M[\bar{q}] = 1 \text{ do } S_1(\theta) \text{ done})| \\
& \underline{\text{while}^{(T)} } M[\bar{q}] = 1 \text{ do } S_1(\theta) \text{ done} \\
& \frac{\partial}{\partial \theta}(S_0(\theta))|S_0(\theta) \quad \frac{\partial}{\partial \theta}(S_1(\theta))|S_1(\theta) \\
(\text{Sum Component}) \quad & \frac{\partial}{\partial \theta}(\underline{S_0(\theta)} + \underline{S_1(\theta)})|(\underline{S_0(\theta)} + \underline{S_1(\theta)})
\end{aligned}$$

Figure 5. The differentiation logic. Wherever applicable, $q \in \bar{v}, \bar{q} \subseteq \bar{v}, S_i(\theta) \in \mathbf{add-q-while}_{\bar{v}}^{(T)}(\theta)$. In (Rot-Couple), $\sigma \in \{X, Y, Z, X \otimes X, Y \otimes Y, Z \otimes Z\}$.

(Definition 5.2) and the strong requirement of computing differential semantics in Definition 5.3. Firstly,

$$\begin{aligned}
\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta); S_1(\theta)) \rrbracket &= \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)); \underline{S_1(\theta)} \rrbracket + \\
&\llbracket (O, \rho) \rightarrow \underline{S_0(\theta)}; \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket
\end{aligned}$$

We use the induction hypothesis to reason about each term above. Consider the case $S_0(\theta) = S_0(\theta)$ and $S_1(\theta) = S_1(\theta)$. We show the following for the second step:

$$\begin{aligned}
\llbracket (O, \rho) \rightarrow S_0(\theta); \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket &= \llbracket (O, \llbracket S_0(\theta) \rrbracket(\rho)) \rightarrow \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket \\
&\quad (6.8) \\
\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)); S_1(\theta) \rrbracket &= \llbracket (\llbracket S_1(\theta) \rrbracket^*(O), \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)) \rrbracket, \\
&\quad (6.9)
\end{aligned}$$

where we make critical use of the fact $S_0(\theta), S_1(\theta) \in \mathbf{q}\text{-while}_{\bar{v}}^{(T)}(\theta)$ and $\frac{\partial}{\partial \theta}(S_0(\theta)), \frac{\partial}{\partial \theta}(S_1(\theta)) \in \mathbf{add-q-while}_{\bar{v} \cup \{A\}}^{(T)}(\theta)$.

For (6.8), since $\frac{\partial}{\partial \theta}(S_1(\theta))$ computes the derivative for any input state and observable, we choose the input state $\llbracket S_0(\theta) \rrbracket(\rho)$ and observable O .

For (6.9), we don't change the state ρ but change the observable O by applying the *dual* super-operator $\llbracket S_1(\theta) \rrbracket^*$. Since $\frac{\partial}{\partial \theta}(S_0(\theta))$ computes the derivative for

any input state and any observable, we choose the input state ρ and observable $\llbracket S_1(\theta) \rrbracket^*(O)$.

The dual super-operator $\llbracket S_1(\theta) \rrbracket^*$ has the property that $\text{tr}(O \llbracket S_1(\theta) \rrbracket(\rho)) = \text{tr}(\llbracket S_1(\theta) \rrbracket^*(O) \rho)$, which corresponds to the Schrodinger picture (evolving states) and Heisenberg picture (evolving observables) respectively in quantum mechanics.

6. The proof of the *Case* rule basically follows from the linearity of the observable semantics and the smooth semantics of *Case*. It is interesting to compare with the classical case [5] where the non-smoothness of the guard causes an issue for auto differentiation.

Example 6.1 (Simple-Case). *Consider the following simple instantiating of Ex. 4.1*

$$P(\theta)|_{q_1} \equiv \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} R_X(\theta)|_{q_1}; R_Y(\theta)|_{q_1}, \\ 1 \rightarrow R_Z(\theta)|_{q_1} \end{array}$$

Let us apply code transformation and compilation. Let *CT*, *CP* to denote “code transformation” and “compilation”, and “*Seq*” and “*Rot*” denote Sequence and Rotation rules resp.

$$\begin{aligned} \frac{\partial}{\partial \theta}(P(\theta)) &\stackrel{CT, \text{case}}{=} \begin{array}{l} \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} \frac{\partial}{\partial \theta}(R_X(\theta)|_{q_1}); \\ R_Y(\theta)|_{q_1}), \\ 1 \rightarrow \frac{\partial}{\partial \theta}(R_Z(\theta)|_{q_1}) \end{array} \\ \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} (R'_X(\theta)|_{A, q_1}; \\ R_Y(\theta)|_{q_1}) + \\ (R_X(\theta)|_{q_1}; \\ R'_Y(\theta)|_{A, q_1}), \\ 1 \rightarrow R'_Z(\theta)|_{A, q_1} \end{array} \end{array} \\ &\stackrel{CT, \text{Seq} + \text{Rot}}{=} \begin{array}{l} \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} (R'_X(\theta)|_{A, q_1}; \\ R_Y(\theta)|_{q_1}), \\ 1 \rightarrow R'_Z(\theta)|_{A, q_1}, \end{array} \\ \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} R_X(\theta)|_{q_1}; \\ R'_Y(\theta)|_{A, q_1}), \end{array} \end{array} \\ &\stackrel{\text{Compile}(\bullet)}{\mapsto} \left\{ \begin{array}{l} \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} R'_X(\theta)|_{A, q_1}; \\ R_Y(\theta)|_{q_1}, \\ 1 \rightarrow R'_Z(\theta)|_{A, q_1}, \\ \text{case } M[q_1] = 0 \rightarrow \begin{array}{l} R_X(\theta)|_{q_1}; \\ R'_Y(\theta)|_{A, q_1}), \end{array} \\ 1 \rightarrow \text{abort}. \end{array} \right\} \end{aligned}$$

7 Execution and Resource Analysis

In this section we illustrate the execution of the entire differentiation procedure and analyze its resource cost. Consider any program $P(\theta) \in \mathbf{q}\text{-while}_{\frac{(T)}{\delta}}(\theta)$ and the parameter θ .

Execution. The first step in differentiation is to apply the code transformation rules (in Section 6) to $P(\theta)$ and obtain an additive program $\frac{\partial}{\partial \theta}(P(\theta))$. Then one needs to compile $\frac{\partial}{\partial \theta}(P(\theta))$ into a multiset $\{P'_i(\theta)\}_{i=1}^m$ of normal non-aborting programs $P'_i(\theta)$. The total count of these programs is given by $m = \|\# \frac{\partial}{\partial \theta}(P(\theta))\|$. Note that the above procedure could be done at the compilation time.

Given any pair of O and ρ , the real execution to compute the derivative of $\llbracket (O, \rho) \rightarrow P(\theta) \rrbracket$ is to approximate the observable semantics $\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(P(\theta)) \rrbracket$. By Definition 5.2,

we need to approximate

$$\sum_{i=1}^m \text{tr} \left((Z_A \otimes O) \llbracket P'_i(\theta) \rrbracket ((|\bar{0}\rangle_{\{A\}} \langle \bar{0}|) \otimes \rho) \right), \quad (7.1)$$

where each term is the observable $Z_A \otimes O$ on the output state of $P'_i(\theta)$ given input state ρ and the ancilla qubit $|0\rangle$.

To approximate the sum in (7.1) to precision δ , one could first treat the sum divided by m as the observable applied on the program that starts with a uniformly random choice of i from $1, \dots, m$ and then execute $P'_i(\theta)$. By Chernoff bound, one only needs to repeat this procedure $O(m^2/\delta^2)$ times.

Resource count. We are only interested in non-trivial (extra) resource that is something that you wouldn't need if you only run the original program. Ancilla qubits count as the non-trivial resource. However, for our scheme, the number of required ancilla is 1 qubit when applying $\frac{\partial}{\partial \theta}(\cdot)$ once.

The more non-trivial resource is the number of the copies of input state, which is directly related to the number of repetitions in the procedure, which again connects to $m = \|\# \frac{\partial}{\partial \theta}(P(\theta))\|$. We argue that our code transformation is efficient so that m is reasonably bounded. To that end, we show the relation between m and a natural quantity defined on the original program $P(\theta)$ (i.e., before applying any $\frac{\partial}{\partial \theta}(\cdot)$ operator) called the *occurrence count* of the parameter θ .

Definition 7.1. The “Occurrence Count for θ_j ” in $P(\theta)$, denoted $\text{OC}_j(P(\theta))$, is defined as follows:

1. If $P(\theta) \equiv \text{abort}[\bar{v}] \text{ skip}[\bar{v}] | q := |0\rangle$ ($q \in \bar{v}$), then $\text{OC}_j(P(\theta)) = 0$;
2. $P(\theta) \equiv U(\theta)$: if $U(\theta)$ trivially uses θ_j , then $\text{OC}_j(P(\theta)) = 0$; otherwise $\text{OC}_j(P(\theta)) = 1$.
3. If $P(\theta) \equiv U(\theta) = P_1(\theta); P_2(\theta)$ then $\text{OC}_j(P(\theta)) = \text{OC}_j(P_1(\theta)) + \text{OC}_j(P_2(\theta))$.
4. If $P(\theta) \equiv \text{case } M[\bar{q}] = m \rightarrow P_m(\theta) \text{ end}$ then $\text{OC}_j(P(\theta)) = \max_m \text{OC}_j(P_m(\theta))$.
5. If $P(\theta) \equiv \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done}$ then $\text{OC}_j(P(\theta)) = T \cdot \text{OC}_j(P_1(\theta))$.

Intuition of the “Occurrence Count” definition is clear: it basically counts the non-trivial number of occurrence of θ_j in the program, treating *case* as if it is deterministic. To see why this is a reasonable quantity, consider the auto-differentiation in the classical case. For any non-trivial variable v (i.e., v has some dependence on the parameter θ), we will compute both v and $\frac{\partial}{\partial \theta}(v)$ and store them both as variables in the new program. Thus, the classical auto-differentiation essentially needs the number of non-trivial occurrences more space and related resources. As we argued in the introduction, we cannot directly mimic the classical case due to the non-cloning theorem. The extra space requirement in the classical setting turns into the requirement on the extra copies of the input state in the quantum setting. Indeed, we can bound m by the occurrence count.

Proposition 7.2. $\|\# \frac{\partial}{\partial \theta_j}(P(\theta))\| \leq \text{OC}_j(P(\theta))$.

Proof. Structural induction. For details, see Appendix E.1. \square

8 Implementation and Case Study

We have built a compiler (written in OCaml) that implements our code transformation and compilation rules. We use it to train one VQC instance with controls and empirically verify its resource-efficiency on representative VQC instances. All codes and test programs are available in the supplementary materials. Complete details are in Appendix F. Experiments are performed on a MacBook Pro with a Dual-Core Intel Core i5 Processor clocked at 2.7 GHz, and 8GB of RAM.

8.1 Training VQC instances with controls

Consider a simple classification problem over 4-bit inputs $z = z_1 z_2 z_3 z_4 \in \{0, 1\}^4$ with true label given by $f(z) = \neg(z_1 \oplus z_4)$. We construct two 4-qubit VQCs P_1 (**no control**) and P_2 (**with control**) that consists of a single-qubit Pauli X, Y and Z rotation gate on each qubit and compare their performance.

For parameters $\Gamma = \{\gamma_1, \dots, \gamma_{12}\}$ define the program

$$Q(\Gamma) \equiv \begin{aligned} &R_X(\gamma_1)[q_1]; R_X(\gamma_2)[q_2]; R_X(\gamma_3)[q_3]; R_X(\gamma_4)[q_4]; \\ &R_Y(\gamma_5)[q_1]; R_Y(\gamma_6)[q_2]; R_Y(\gamma_7)[q_3]; R_Y(\gamma_8)[q_4]; \\ &R_Z(\gamma_9)[q_1]; R_Z(\gamma_{10})[q_2]; R_Z(\gamma_{11})[q_3]; R_Z(\gamma_{12})[q_4], \end{aligned}$$

where q_1, q_2, q_3, q_4 refer to 4 qubit registers. Given parameters $\Theta = \{\theta_1, \dots, \theta_{12}\}$, $\Phi = \{\phi_1, \dots, \phi_{12}\}$, define

$$P_1(\Theta, \Phi) \equiv Q(\Theta); Q(\Phi). \quad (8.1)$$

Similarly, for parameters $\Theta = \{\theta_1, \dots, \theta_{12}\}$, $\Phi = \{\phi_1, \dots, \phi_{12}\}$, $\Psi = \{\psi_1, \dots, \psi_{12}\}$, define

$$P_2(\Theta, \Phi, \Psi) \equiv Q(\Theta); \text{case } M[q_1] = 0 \rightarrow Q(\Phi) \\ 1 \rightarrow Q(\Psi). \quad (8.2)$$

Note that P_1 and P_2 execute the same number of gates for each run. To use P_i to perform the classification or in the training, we first initialize q_1, q_2, q_3, q_4 to the classical feature vector $z = z_1 z_2 z_3 z_4$ and then execute P_i . The predicted label y is given by measuring the 4th qubit q_4 in 0/1 basis.

We will conduct a supervised learning by minimizing the following loss function given by the average negative log likelihood of correctly classifying the inputs. Thus

$$\text{loss} = -\frac{1}{16} \sum_{z \in \{0,1\}^4} \log(\Pr[y = f(z) = \neg(z_1 \oplus z_4)]). \quad (8.3)$$

Note that loss is a function on Θ, Φ (or Θ, Φ, Ψ). More importantly, for each z , $\Pr[y = \neg(z_1 \oplus z_4)]$ can be represented by the observable semantics of P_1 (or P_2) with observable $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$. Thus, the gradient of loss can be obtained by using the collection of $\frac{\partial}{\partial \alpha}(P_i)$ for $\alpha \in \Theta, \Phi$ (or $\frac{\partial}{\partial \alpha}(P_2)$ for $\alpha \in \Theta, \Phi, \Psi$). We classically simulate the above training procedure with gradient descent⁷ (see Figure 6). After 1000

⁷We attempted to directly use PennyLane to train the system without control, but were not able to make it work since PennyLane does not implement the gradient of the log function.

epochs with manually optimized hyperparameters, the loss for P_1 (**no control**) attains a minimum of 0.6931 in less than 100 epochs and subsequently plateaus. The loss for P_2 (**with control**) continues to decrease and attains a minimum of 0.0936. It demonstrates the advantage of both controls in quantum machine learning and our scheme to handle controls, whereas previous schemes (such as PennyLane due to its quantum-node design [7]) fail to do so.




Figure 6. Training P_1 and P_2 to classify inputs according to the labelling function $f(z) = \neg(z_1 \oplus z_4)$.

8.2 Benchmark testing on representative VQCs

We also test our compiler on important VQC candidates such as quantum neural-networks (QNN) [16], quantum approximate optimization algorithms (QAOA) [15], and variational quantum eigensolver (VQE) [35]. We enrich these examples, by adding simple controls or 2-bounded loops and increasing the number of qubits to 18~40, to make them sufficiently sophisticated but yet realistic for near-term quantum applications. A selective output performance of our compiler is in Table 2, with full details in Appendix F. Our scheme is resource-efficient as $|\# \frac{\partial}{\partial \theta}(\cdot)|$ is always reasonably bounded.

$P(\theta)$	$OC(\cdot)$	$ \# \frac{\partial}{\partial \theta}(\cdot) $	#gates	#lines	#layers	#qb's
$QNN_{M,i}$	24	24	165	189	3	18
$QNN_{M,w}$	56	24	231	121	5	18
$QNN_{L,i}$	48	48	363	414	6	36
$QNN_{L,w}$	504	48	2079	244	33	36
$VQE_{M,i}$	15	15	224	241	3	12
$VQE_{M,w}$	35	15	224	112	5	12
$VQE_{L,i}$	40	40	576	628	5	40
$VQE_{L,w}$	248	40	1984	368	17	40
$QAOA_{M,i}$	18	18	120	142	3	18
$QAOA_{M,w}$	42	18	168	94	5	18
$QAOA_{L,i}$	36	36	264	315	6	36
$QAOA_{L,w}$	378	36	1512	190	33	36

Table 2. Output on selective examples. $\{M, L\}$ stands for “medium, large”; $\{i, w\}$ stands for including “if, while”.

Acknowledgments

SZ, SH, and XW were supported in part 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. SC and XW were supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams program. XW also received support from the National Science Foundation (grants CCF-1755800 and CCF-1816695).

References

- [1] 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.
- [2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (2019), 505–510.
- [3] Alexandru Baltag and Sonja Smets. 2011. Quantum Logic as a Dynamic Logic. *Synthese* 179, 2 (2011).
- [4] Atılım Günes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. 2017. Automatic Differentiation in Machine Learning: A Survey. *J. Mach. Learn. Res.* 18, 1 (Jan. 2017), 5595–5637. <http://dl.acm.org/citation.cfm?id=3122009.3242010>
- [5] Thomas Beck and Herbert Fischer. 1994. The if-problem in automatic differentiation. *J. Comput. Appl. Math.* 50, 1-3 (1994), 119–131.
- [6] Marcello Benedetti, Erika Lloyd, and Stefan Sack. 2019. Parameterized quantum circuits as machine learning models. *arXiv e-prints*, Article arXiv:1906.07682 (Jun 2019), arXiv:1906.07682 pages. arXiv:quant-ph/1906.07682
- [7] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, and Nathan Killoran. 2018. PennyLane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968* (2018).
- [8] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum machine learning. *Nature* 549, 7671 (2017), 195. arXiv:arXiv:1611.09347
- [9] Olivier Brunet and Philippe Jorrand. 2004. Dynamic Quantum Logic for Quantum Programs. *International Journal of Quantum Information* 2, 1 (2004).
- [10] Rohit Chadha, Paulo Mateus, and Amílcar Sernadas. 2006. Reasoning About Imperative Quantum Programs. *Electronic Notes in Theoretical Computer Science* 158 (2006).
- [11] George Corliss, Christèle Faure, Andreas Griewank, Lauren Hascoët, and Uwe Naumann (Eds.). 2002. *Automatic Differentiation of Algorithms: From Simulation to Optimization*. Springer-Verlag New York, Inc., New York, NY, USA.
- [12] Thomas Ehrhard and Laurent Regnier. 2003. The differential lambda-calculus. *Theoretical Computer Science* 309, 1-3 (2003), 1–41.
- [13] Conal M. Elliott. 2009. Beautiful Differentiation. In *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP '09)*. ACM, New York, NY, USA, 191–202. <https://doi.org/10.1145/1596550.1596579>
- [14] Conal M. Elliott. 2018. The Simple Essence of Automatic Differentiation. *Proc. ACM Program. Lang.* 2, ICFP, Article 70 (July 2018), 29 pages. <https://doi.org/10.1145/3236765>
- [15] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. (2014). arXiv:arXiv:1411.4028
- [16] Edward Farhi and Hartmut Neven. 2018. Classification with Quantum Neural Networks on Near Term Processors. *arXiv e-prints*, Article arXiv:1802.06002 (Feb 2018), arXiv:1802.06002 pages. arXiv:quant-ph/1802.06002
- [17] 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).
- [18] Simon J. Gay. 2006. Quantum Programming Languages: Survey and Bibliography. *Mathematical Structures in Computer Science* 16, 4 (2006).
- [19] Gian Giacomo Guerreschi and Mikhail Smelyanskiy. 2017. Practical optimization for hybrid quantum-classical algorithms. *arXiv e-prints*, Article arXiv:1701.01450 (Jan 2017), arXiv:1701.01450 pages. arXiv:quant-ph/1701.01450
- [20] Martin Giles. 2019. IBM's new 53-qubit quantum computer is the most powerful machine you can use. <https://www.technologyreview.com/f/614346/ibms-new-53-qubit-quantum-computer-is-the-most-powerful-machine-you-can-use/>
- [21] Pranav Gokhale. 2018. Variational Quantum Eigensolver Demo. *ISCA 2018* (2018).
- [22] Jonathan Grattage. 2005. A Functional Quantum Programming Language. In *LICS*.
- [23] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, Adrià Puigdomènech Badia, Karl Moritz Hermann, Yori Zwols, Georg Ostrovski, Adam Cain, Helen King, Christopher Summerfield, Phil Blunsom, Koray Kavukcuoglu, and Demis Hassabis. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538 (10 2016), 471.
- [24] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. 2013. Quipper: A Scalable Quantum Programming Language. In *PLDI*.
- [25] Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. Learning to Transduce with Unbounded Memory. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'15)*. MIT Press, Cambridge, MA, USA, 1828–1836. <http://dl.acm.org/citation.cfm?id=2969442.2969444>
- [26] Andreas Griewank. 2000. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- [27] Yoshihiko Kakutani. 2009. A Logic for Formal Verification of Quantum Programs. In *ASIAN*.
- [28] Gershon Kedem. 1980. Automatic Differentiation of Computer Programs. *ACM Trans. Math. Softw.* 6, 2 (June 1980), 150–165. <https://doi.org/10.1145/355887.355890>

- [29] Jin-Guo Liu and Lei Wang. 2018. Differentiable learning of quantum circuit Born machines. *Phys. Rev. A* 98 (Dec 2018), 062324. Issue 6. <https://doi.org/10.1103/PhysRevA.98.062324>
- [30] 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 Muller, 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), 030503. arXiv:arXiv:1710.01022
- [31] Michael A. Nielsen and Isaac Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge University Press.
- [32] Bernhard Ömer. 2003. *Structured Quantum Programming*. Ph.D. Dissertation. Vienna University of Technology.
- [33] Jennifer Paykin, Robert Rand, and Steve Zdancewic. 2017. QWIRE: A Core Language for Quantum Circuits. In *POPL*.
- [34] Barak A. Pearlmutter and Jeffrey Mark Siskind. 2008. Reverse-mode AD in a Functional Framework: Lambda the Ultimate Backpropagator. *ACM Trans. Program. Lang. Syst.* 30, 2, Article 7 (March 2008), 36 pages. <https://doi.org/10.1145/1330017.1330018>
- [35] 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 (2014), 4213. arXiv:arXiv:1304.3061
- [36] Gordon Plotkin. 2018. Some Principles of Differential Programming Languages. *POPL 2018* (2018).
- [37] John Preskill. 2018. Quantum computing in the NISQ era and beyond. *Quantum* 2 (2018), 79. arXiv:arXiv:1801.00862
- [38] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning representations by back-propagating errors. *Nature* 323, 6088 (1986), 533–536.
- [39] Amr Sabry. 2003. Modeling Quantum Computing in Haskell. In *The Haskell Workshop*.
- [40] Jeff W. Sanders and Paolo Zuliani. 2000. Quantum Programming. In *MPC*.
- [41] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. 2018. Evaluating analytic gradients on quantum hardware. *arXiv preprint arXiv:1811.11184* (2018).
- [42] Maria Schuld, Alex Bocharov, Krysta Svore, and Nathan Wiebe. 2018. Circuit-centric quantum classifiers. *arXiv e-prints*, Article arXiv:1804.00633 (Apr 2018), arXiv:1804.00633 pages. arXiv:quant-ph/1804.00633
- [43] Peter Selinger. 2004. A Brief Survey of Quantum Programming Languages. In *FLOPS*.
- [44] Peter Selinger. 2004. Towards a Quantum Programming Language. *Mathematical Structures in Computer Science* 14, 4 (2004).
- [45] Bert Speelpenning. 1980. *Compiling Fast Partial Derivatives of Functions Given by Algorithms*. Ph.D. Dissertation. Champaign, IL, USA. AAI8017989.
- [46] 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*.
- [47] Qingfeng Wang and Tauqir Abdullah. 2018. An Introduction to Quantum Optimization Approximation Algorithm.
- [48] John Watrous. 2006. Introduction to Quantum Computation. <https://cs.uwaterloo.ca/~watrous/LectureNotes/CPSC519.Winter2006/all.pdf>. Course notes.
- [49] Dave Wecker and Krysta Svore. 2014. LIQUi|>: A Software Design Architecture and Domain-Specific Language for Quantum Computing. *CoRR abs/1402.4467* (2014). arXiv:1402.4467
- [50] Robert Edwin Wengert. 1964. A Simple Automatic Derivative Evaluation Program. *Commun. ACM* 7, 8 (Aug. 1964), 463–464. <https://doi.org/10.1145/355586.364791>
- [51] William K. Wootters and Wojciech H. Zurek. 1982. A single quantum cannot be cloned. *Nature* 299, 5886 (1982), 802–803.
- [52] Mingsheng Ying. 2011. Floyd–Hoare Logic for Quantum Programs. *ACM Transactions on Programming Languages and Systems* 33, 6 (2011).
- [53] Mingsheng Ying. 2016. *Foundations of Quantum Programming*. Morgan Kaufmann.
- [54] Mingsheng Ying and Yangjia Li. 2018. Reasoning about Parallel Quantum Programs. *arXiv preprint arXiv:1810.11334* (2018).
- [55] Mingsheng Ying, Shenggang Ying, and Xiaodi Wu. 2017. Invariants of Quantum Programs: Characterisations and Generation. In *POPL*.
- [56] Daiwei Zhu, NM Linke, Marcello Benedetti, KA Landsman, NH Nguyen, CH Alderete, Alejandro Perdomo-Ortiz, Nathan Korda, A Garfoot, Charles Brecque, et al. 2018. Training of quantum circuits on a hybrid quantum computer. *arXiv preprint arXiv:1812.08862* (2018).

A Detailed Quantum Preliminary

This is a more detailed treatment of Section 2. For a further extended background, we recommend the notes by Watrous [48] and the textbook by Nielsen and Chuang [31].

A.1 Preliminaries

For any non-negative 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 $\| |\psi\rangle \| = \sqrt{\langle\psi|\psi\rangle}$.

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:

$$\begin{aligned} A \sqsubseteq B & \quad \text{if } B - A \text{ is positive semidefinite,} \\ A \sqsubset B & \quad \text{if } B - A \text{ is positive definite.} \end{aligned}$$

A.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 $|0\rangle = (1, 0)^\top$ and $|1\rangle = (0, 1)^\top$, 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 $\mathbb{C}^2 \otimes \mathbb{C}^2$. An important 2-qubit state is the EPR state $|\beta_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. 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 $\mathcal{D}(\mathcal{H})$.

A.3 Quantum Operations

Operations on quantum systems can be characterized by unitary operators. Denoting the set of linear operators on \mathcal{H} as $L(\mathcal{H})$, an operator $U \in L(\mathcal{H})$ is *unitary* if its Hermitian conjugate is its own inverse, i.e., $U^\dagger U = U U^\dagger = I_{\mathcal{H}}$. 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 unitary operators include

$$H = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \end{bmatrix}, \quad X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad Y = -iXZ$$

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 $Z|1\rangle = -|1\rangle$. *Pauli Y* maps $|0\rangle$ to $i|1\rangle$ and $|1\rangle$ to $-i|0\rangle$. The *CNOT* gate C maps $|00\rangle \mapsto |00\rangle$, $|01\rangle \mapsto |01\rangle$, $|10\rangle \mapsto |11\rangle$, $|11\rangle \mapsto |10\rangle$. One may obtain the EPR state $|\beta_{00}\rangle$ via $|00\rangle \xrightarrow{H_1} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|0\rangle \xrightarrow{C_{1,2}} \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.

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 $\sum_k E_k^\dagger \circ E_k$.

A.4 Quantum Measurements and Observables

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 $\{M_m\}_m$ 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 $|1\rangle\langle 1|$.

B More on the definition of parameterized quantum programs

B.1 Definition of qVar

Given $P(\theta)$ a T -bounded k -parameterized quantum **while**-program, let $qVar(P(\theta))$, read *the set of quantum variables accessible to P , be recursively defined as follows* [53]:

1. If $P(\theta) \equiv \text{skip}[\bar{q}], \text{abort}[\bar{q}]$ or $U[\bar{q}]$, then $qVar(P(\theta)) = \bar{q}$.
2. If $P(\theta) \equiv q := |0\rangle$, then $qVar(P(\theta)) = q$.
3. If $P(\theta) \equiv P_1(\theta); P_2(\theta)$, then $qVar(P(\theta)) = qVar(P_1(\theta)) \cup qVar(P_2(\theta))$. When analyzing $P_1(\theta); P_2(\theta)$, we identify $P_1(\theta)$ with $\mathcal{I} \otimes P_1(\theta)$, where $\mathcal{I} \equiv I_{qVar(P_2(\theta)) \setminus qVar(P_1(\theta))} \circ I_{qVar(P_2(\theta)) \setminus qVar(P_1(\theta))}$, the identity operation on the variables where P_1 originally has no access to.
4. If $P(\theta) \equiv \text{case } M[\bar{q}] = m \rightarrow P_m(\theta) \text{ end}$, then $qVar(P(\theta)) = \bar{q} \cup \bigcup_m qVar(P_i(\theta))$. We make the same identification for $P_i(\theta)$'s as the above.
5. If $P(\theta) \equiv \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done}$, then $qVar(P(\theta)) = \bar{q} \cup qVar(P_1(\theta))$. We make the same identification for $P_1(\theta)$ as the above.

One defines $qVar(P)$ for unparameterized P analogously.

C Detailed Proof from Section 4

C.1 Proof of Prop 4.2: Non-deterministic Compilation Rules Well-Defined

Proof. Structural induction.

1. (Atomic) and (Sequence \downarrow): as operational semantics of atomic operations are inherited from parameterized quantum while programs, they do not induce non-determinism. For these operations, $\llbracket P(\theta^*) \rrbracket \rho = \{\llbracket P(\theta^*) \rrbracket \rho\}$ as is the right hand side. (Sequence \downarrow) follows from the fact that $\downarrow; P(\theta) \equiv P(\theta)$.
2. (Sequence) If one of $\text{Compile}(P_b(\theta^*))$ is $\{\text{abort}\}$ the statement is immediately true, so we assume otherwise for below.
 - a. \subseteq : Let $\rho' \in \{|\rho' \neq 0 : \langle P_1(\theta^*); P_2(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle|\}$. By definition there exists $\rho_2 \neq 0$ s.t. $\langle P_1(\theta^*); P_2(\theta^*), \rho \rangle \rightarrow^* \langle P_2(\theta^*), \rho_2 \rangle$ and $\langle P_2(\theta^*), \rho_2 \rangle \rightarrow^* \langle \downarrow, \rho' \rangle$.

ρ' . By inductive hypothesis, $\rho' \in \bigsqcup_{Q_j(\theta) \in \text{Compile}(P_2(\theta))} \{|\rho' \neq 0 : \langle Q_j(\theta^*), \rho_2 \rangle \rightarrow^* \langle \downarrow, \rho' \rangle|\}$. Since $\langle P_1(\theta^*); P_2(\theta^*), \rho \rangle \rightarrow^* \langle P_2(\theta^*), \rho_2 \rangle$, there must be some $\langle P_3(\theta^*), \rho_3 \rangle \rightarrow \langle \downarrow, \rho_2 \rangle$ triggering the last transition according to our operational semantics. Inductively apply such an argument, knowing that all computation paths are finitely long,⁸ we conclude that $\rho_2 \in \{|\rho'_2 : \langle P_1(\theta), \rho \rangle \rightarrow^* \langle \downarrow, \rho'_2 \rangle|\}$. By the inductive hypothesis we have $\rho_2 \in \bigsqcup_{R_i(\theta) \in \text{Compile}(P_1(\theta))} \{|\rho'_2 \neq 0 : \langle R_i(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho'_2 \rangle|\}$. Thus

$$\rho' \in \bigsqcup_{R_i(\theta) \in \text{Compile}(P_1(\theta)), Q_j(\theta) \in \text{Compile}(P_2(\theta))} \{|\rho' \neq 0 : \langle R_i(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho_2 \rangle\} \quad (\text{C.1})$$

$$\{|\rho' \neq 0 : \langle R_i(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho_2 \rangle\} \quad (\text{C.2})$$

$$\{|\rho' \neq 0 : \langle R_i(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho_2 \rangle\} \quad (\text{C.3})$$

$$\bigwedge \langle Q_j(\theta^*), \rho_2 \rangle \rightarrow^* \langle \downarrow, \rho' \rangle \} \\ = \bigsqcup_{R_i(\theta) \in \text{Compile}(P_1(\theta)), Q_j(\theta) \in \text{Compile}(P_2(\theta))} \{|\rho' \neq 0 : \langle R_i(\theta^*), Q_j(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle|\} \quad (\text{C.4})$$

$$\{|\rho' \neq 0 : \langle R_i(\theta^*), Q_j(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle|\} \quad (\text{C.5})$$

as desired.

- b. \supseteq : Let ρ' be a member of the multiset in right hand side ("RHS") of C.5 for some $R_i(\theta) \in \text{Compile}(P_1(\theta))$, $Q_j(\theta) \in \text{Compile}(P_2(\theta))$. Then by inductive hypothesis, $\langle P_2(\theta), \llbracket R_i(\theta^*) \rrbracket \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle$ ($\rho' \neq 0$), and $\langle P_1(\theta), \rho \rangle \rightarrow^* \langle \downarrow, \llbracket R_i(\theta^*) \rrbracket \rho \rangle$ ($\llbracket R_i(\theta^*) \rrbracket \rho \neq 0$). One may start with $\langle P_2(\theta^*), \llbracket R_i(\theta^*) \rrbracket \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle$ then repeatedly apply the Sequential rule in the operational semantics, tracing back the path $\langle P_1(\theta), \rho \rangle \rightarrow^* \langle \downarrow, \llbracket R_i(\theta^*) \rrbracket \rho \rangle$, and conclude that $\langle P_1(\theta^*); P_2(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle$ ($\rho' \neq 0$).

The above said the two multisets have the same set of distinct elements, and furthermore each copy of $\rho' \in \llbracket P(\theta^*) \rrbracket \rho$ induces at least (by \subseteq , or by the deterministic nature of $R_i(\theta^*); Q_j(\theta^*)$) as well as at most (by \supseteq) one copy of the same element in RHS(C.5). This says the two multisets are equal.

3. (Case) Applying the operational transition rule for 1 step one may observe that (writing "IH" the inductive

⁸We don't count "transitions" like $\langle P(\theta^*), \rho \rangle \rightarrow \langle P(\theta^*), \rho \rangle$ when analyzing computation paths, as nothing evolves in these "trivial transitions".

hypothesis)

$$\begin{aligned}
& \{ |0 \neq \rho' \in \llbracket \text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end} \rrbracket \rho \} \\
&= \coprod_{m^*} \{ | \rho' \neq 0 : \langle P_m(\theta^*), M_m \rho M_m^\dagger \rangle \rightarrow \langle \downarrow, \rho' \rangle | \} \\
&\stackrel{IH}{=} \coprod_{m^*} \coprod_{Q_{m^*, i_{m^*}}(\theta) \in \text{Compile}(\underline{P_m(\theta)})} \quad (C.6) \\
&\{ | \rho' \neq 0 : \langle \llbracket Q_{m^*, i_{m^*}}(\theta^*) \rrbracket, M_{m^*} \rho M_{m^*}^\dagger \rangle \rightarrow^* \langle \downarrow, \rho' \rangle | \} \quad (C.7)
\end{aligned}$$

On the other hand we have

$$\begin{aligned}
& \coprod_{Q(\theta) \in \text{Compile}(\text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end})} \\
& \{ | \rho' \neq 0 : \langle Q(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle | \} \\
&\stackrel{\text{CP, Case}}{=} \coprod_{Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end})} \\
& \{ | \rho' \neq 0 : \langle Q(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle | \} \\
&\stackrel{(**)}{=} \coprod_{m^*} \coprod_{Q_{m^*, i_{m^*}}(\theta) \in \text{Compile}(\underline{P_m(\theta)})} \{ | \rho' \neq 0 : \langle \llbracket Q_{m^*, i_{m^*}}(\theta^*) \rrbracket, \\
& M_{m^*} \rho M_{m^*}^\dagger \rangle \rightarrow^* \langle \downarrow, \rho' \rangle | \}
\end{aligned}$$

as desired. Note that the last step (**) is due to the behavior of the superoperators in $\text{FB}(\text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end})$: when evolved by one transition, they either go to some *non-(essentially-aborting)* $\langle \llbracket Q_{m^*, i_{m^*}}(\theta^*) \rrbracket, M_{m^*} \rho M_{m^*}^\dagger \rangle$ for some $Q_{m^*, i_{m^*}}(\theta) \in \text{Compile}(\underline{P_m(\theta)})$ for some m^* , or go to $\langle \text{abort}, M_{m^*} \rho M_{m^*}^\dagger \rangle$. Besides, each copy of final state of non-(essentially-aborting) $\langle \llbracket Q_{m^*, i_{m^*}}(\theta^*) \rrbracket, M_{m^*} \rho M_{m^*}^\dagger \rangle$ appears exactly once in the multiset. Due to our construction of $\text{FB}(\text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end})$. On the other hand, computational paths starting with *essentially aborting* $\langle Q_{m^*, i_{m^*}}(\theta^*), M_{m^*} \rho M_{m^*}^\dagger \rangle$ always terminate in $\langle \downarrow, 0 \rangle$, therefore not counted on either side of (**).

4. (While^(T)) By definition, While^(T) reduces to (Case) and (Sequence).
5. (Sum Components) If one of $\text{Compile}(P_b(\theta^*))$ is $\{\text{abort}\}$ the statement is immediately true, so we assume otherwise for below. Like in the **case** case, applying the operational transition rule for 1 step one may observe

that (writing “IH” the inductive hypothesis)

$$\begin{aligned}
& 0 \neq \rho' \in \llbracket \underline{P_1(\theta^*)} + \underline{P_2(\theta^*)} \rrbracket \rho \quad (C.9) \\
&= \coprod_{m^*=1,2} \{ | \rho' \neq 0 : \langle P_m(\theta^*), \rho \rangle \rightarrow \langle \downarrow, \rho' \rangle | \} \\
&\stackrel{IH}{=} \coprod_{m^*=1,2} \{ | \rho' \neq 0 : \langle \llbracket Q_{m^*, i_{m^*}}(\theta^*) \rrbracket, \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle, Q_{m^*, i_{m^*}}(\theta) \in \text{Compile}(\underline{P_m(\theta)}) \} \\
&\stackrel{\text{CP, Sum Component}}{=} \coprod_{Q(\theta) \in \text{Compile}(\underline{P_1(\theta)} + \underline{P_2(\theta)})} \{ | \rho' \neq 0 : \langle Q(\theta^*), \rho \rangle \rightarrow^* \langle \downarrow, \rho' \rangle | \}, (C.11)
\end{aligned}$$

as desired!

□

D Detailed Proofs from Section 6

Throughout, we use LHS, RHS to denote “left hand side” and “right hand side” resp., and “IH” to denote “Inductive Hypothesis”. Wherever applicable, we adopt the overloading convention explained in the main text (See Eqn 6.6).

Before giving proofs, we record the full details for code transformation rule for while^(T) for your interest. Let us (C.8) denote $\frac{\partial}{\partial \theta}(\text{while}^{(T)} M[\bar{q}] = 1 \text{ do } S_1(\theta) \text{ done}) \equiv \text{Seq}^T$, with:

$$\begin{aligned}
\text{Seq}^{(1)} &\equiv \text{case } M[\bar{q}] = 0 \rightarrow \underline{\text{abort}}, \\
&1 \rightarrow \left(\frac{\partial}{\partial \theta_j}(\underline{P_1(\theta)}); \underline{\text{abort}} \right) + \left(\underline{P_1(\theta)}; \underline{\text{abort}} \right),
\end{aligned}$$

and $\text{Seq}^{(T \geq 2)}$ recursively defined as:

$$\begin{aligned}
\text{Seq}^{(T \geq 2)} &\equiv \text{case } M[\bar{q}] = 0 \rightarrow \underline{\text{abort}}, \\
&1 \rightarrow \left(\frac{\partial}{\partial \theta_j}(\underline{P_1(\theta)}); \text{while}^{(T-1)} \right) + \left(\underline{P_1(\theta)}; \text{Seq}^{(T-1)} \right)
\end{aligned}$$

Note that $\text{Seq}^{(T-1)}$ essentially aborts. Let us now introduce some helper lemmas for the soundness proof:

D.1 Technical Lemmas

Lemma D.1. $U(\theta) \in \{R_\sigma(\theta), R_{\sigma \otimes \sigma}(\theta)\}_{\sigma \in \{X, Y, Z\}}$. Let $\frac{d}{d\theta}U(\theta)$ denote the entry-wise derivative of $U(\theta)$. Then,

$$\frac{d}{d\theta}U(\theta) = \frac{1}{2}U(\theta + \pi); \quad (D.1)$$

$$\begin{aligned} & \frac{\partial}{\partial \theta}(\llbracket(O, \rho) \rightarrow U(\theta)\rrbracket) \\ &= \text{tr}(O \cdot U(\theta) \cdot \rho \cdot (\frac{d}{d\theta}U(\theta))^\dagger) + \\ & \text{tr}(O \cdot \frac{d}{d\theta}U(\theta) \cdot \rho \cdot U^\dagger(\theta)) \end{aligned} \quad (D.2)$$

$$\stackrel{D.1}{=} \frac{1}{2} \text{tr} \left(O \left(U(\theta) \rho U^\dagger(\theta + \pi) + U(\theta + \pi) \rho U^\dagger(\theta) \right) \right). \quad (D.3)$$

Proof. Let $U(\theta) \in \{R_\sigma(\theta), R_{\sigma \otimes \sigma}(\theta)\}_{\sigma \in \{X, Y, Z\}}$. Then $U(\theta) = \cos \frac{\theta}{2} I_\bullet - i \sin \frac{\theta}{2} \sigma$, where $\sigma \in \{X, Y, Z, X \otimes X, Y \otimes Y, Z \otimes Z\}$. In the cases where σ is 1-qubit gate, I_\bullet denotes identity on that one qubit; likewise for 2-qubit cases. Correctness of Eqns D.1, D.2 basically follows from straightforward computation.

1. Equation D.1:

$$\begin{aligned} & \frac{d}{d\theta}(U(\theta)) \\ &= \frac{1}{2} \left((-\sin \frac{\theta}{2}) I_\bullet - i \cos \frac{\theta}{2} \sigma \right) \end{aligned} \quad (D.4)$$

$$= \frac{1}{2} \left(\cos(\frac{\theta}{2} + \frac{\pi}{2}) I_\bullet - i \sin(\frac{\theta}{2} + \frac{\pi}{2}) \sigma \right) \quad (D.5)$$

$$= \frac{1}{2} \left(\cos(\frac{\theta + \pi}{2}) I_\bullet - i \sin(\frac{\theta + \pi}{2}) \sigma \right) \quad (D.6)$$

$$= \frac{1}{2} U(\theta + \pi) \quad (D.7)$$

2. Equation D.2:

$$\begin{aligned} & \frac{\partial}{\partial \theta}(\llbracket(O, \rho) \rightarrow U(\theta)\rrbracket) \\ &= \frac{\partial}{\partial \theta}(\text{tr}(OU(\theta)\rho U^\dagger(\theta))) \end{aligned} \quad (D.8)$$

$$\begin{aligned} &= \frac{\partial}{\partial \theta} \left(\text{tr} \left(O \cdot \left(\cos \frac{\theta}{2} I_\bullet - i \sin \frac{\theta}{2} \sigma \right) \cdot \rho \cdot \right. \right. \\ & \quad \left. \left. \left(\cos \frac{\theta}{2} I_\bullet + i \sin \frac{\theta}{2} \sigma^\dagger \right) \right) \right) \end{aligned} \quad (D.9)$$

$$= \frac{\partial}{\partial \theta} \left(\text{tr} \left(\cos^2 \frac{\theta}{2} O \rho + i \cos \frac{\theta}{2} \sin \frac{\theta}{2} O \rho \sigma^\dagger \right. \right. \quad (D.10)$$

$$\left. \left. - i \sin \frac{\theta}{2} \cos \frac{\theta}{2} O \sigma \rho + \sin^2 \frac{\theta}{2} \sigma \rho \sigma^\dagger \right) \right) \quad (D.11)$$

$$= \text{tr} \left(2 \cos \frac{\theta}{2} \cdot \left(\cos \frac{\theta}{2} \right)' O \rho \right) \quad (D.12)$$

$$+ \text{tr} \left(i \left[\cos \frac{\theta}{2} \left(\sin \frac{\theta}{2} \right)' + \left(\cos \frac{\theta}{2} \right)' \sin \frac{\theta}{2} \right] O (\rho \sigma^\dagger - \sigma \rho) \right)$$

$$+ \text{tr} \left(2 \sin \frac{\theta}{2} \cdot \left(\sin \frac{\theta}{2} \right)' O \sigma \rho \sigma^\dagger \right), \quad (D.13)$$

On the other hand,

$$\begin{aligned} & \text{tr}(O \cdot U(\theta) \cdot \rho \cdot (\frac{d}{d\theta}U(\theta))^\dagger) + \\ & \text{tr}(O \cdot \frac{d}{d\theta}U(\theta) \cdot \rho \cdot U^\dagger(\theta)) \end{aligned} \quad (D.14)$$

$$\begin{aligned} &= \text{tr} \left(O \cdot \left(\cos \frac{\theta}{2} I_\bullet - i \sin \frac{\theta}{2} \sigma \right) \cdot \rho \cdot \right. \\ & \quad \left. \left(\left(\cos \frac{\theta}{2} \right)' I_\bullet + i \left(\sin \frac{\theta}{2} \right)' \sigma^\dagger \right) \right) + \\ & \text{tr} \left(O \cdot \left(\left(\cos \frac{\theta}{2} \right)' I_\bullet - i \left(\sin \frac{\theta}{2} \right)' \sigma \right) \cdot \rho \cdot \right. \\ & \quad \left. \left(\cos \frac{\theta}{2} I_\bullet + i \sin \frac{\theta}{2} \sigma^\dagger \right) \right) \end{aligned} \quad (D.15)$$

$$= \text{tr} \left(2 \cos \frac{\theta}{2} \cdot \left(\cos \frac{\theta}{2} \right)' O \rho \right) \quad (D.16)$$

$$\begin{aligned} & + \text{tr} \left(i \left[\cos \frac{\theta}{2} \left(\sin \frac{\theta}{2} \right)' + \left(\cos \frac{\theta}{2} \right)' \sin \frac{\theta}{2} \right] O (\rho \sigma^\dagger - \sigma \rho) \right) \\ & + \text{tr} \left(2 \sin \frac{\theta}{2} \cdot \left(\sin \frac{\theta}{2} \right)' O \sigma \rho \sigma^\dagger \right), \end{aligned} \quad (D.17)$$

as desired. \square

Lemma D.2. Let $P(\theta) \in \mathbf{q}\text{-while}_{\overline{v}}^{(T)}(\theta)$, $P'(\theta) \in \mathbf{q}\text{-while}_{\overline{v} \cup \{A\}}^{(T)}(\theta)$. Then for arbitrary $\theta^* \in \mathbb{R}^k$, $\rho \in \mathcal{D}(\mathcal{H}_{\overline{v}})$, $O \in \overline{v}$,

1. $\llbracket(O, \rho) \rightarrow P(\theta^*); P'(\theta^*)\rrbracket = \llbracket(O, \llbracket P(\theta^*) \rrbracket \rho) \rightarrow P'(\theta^*)\rrbracket$.
2. $\llbracket(O, \rho) \rightarrow P'(\theta^*); P(\theta^*)\rrbracket = \llbracket(\llbracket P(\theta^*) \rrbracket^*(O), \rho) \rightarrow P'(\theta^*)\rrbracket$, with $\llbracket P(\theta^*) \rrbracket^*$ the dual of $\llbracket P(\theta^*) \rrbracket$.

Proof. Observe that both $P(\theta^*)$; $P'(\theta^*)$ and $P'(\theta^*)$; $P(\theta^*)$ lives in $\mathbf{q}\text{-while}_{\overline{v} \cup \{A\}}^{(T)}(\theta)$, so we identify the “smaller” program $P(\theta^*)$ with $\mathcal{I}_A \otimes P(\theta^*)$, where $\mathcal{I} \equiv I_A \circ I_A$ denotes the identity operation on Ancilla.

1. Unfolding definition 5.2,

$$\begin{aligned} & \llbracket(O, \rho) \rightarrow P(\theta^*); P'(\theta^*)\rrbracket \\ &= \text{tr}((Z_A \otimes O) \llbracket P(\theta^*) \rrbracket; P'(\theta^*) \llbracket (|0\rangle_A \langle 0|) \otimes \rho)) \\ \text{Fig1b, Sequence} \quad &= \text{tr}((Z_A \otimes O) \llbracket P'(\theta^*) \rrbracket \llbracket \mathcal{I}_A \otimes P(\theta^*) \rrbracket \llbracket (|0\rangle_A \langle 0|) \otimes \rho)) \\ &= \text{tr}((Z_A \otimes O) \llbracket P'(\theta^*) \rrbracket \llbracket (|0\rangle_A \langle 0|) \otimes (\llbracket P(\theta^*) \rrbracket \rho)) \rrbracket) \\ &= \llbracket(O, \llbracket P(\theta^*) \rrbracket \rho) \rightarrow P'(\theta^*)\rrbracket, \text{ as desired.} \end{aligned}$$

2. Observe that:

$$\begin{aligned} & \llbracket(O, \rho) \rightarrow P'(\theta^*); P(\theta^*)\rrbracket \\ &= \text{tr}((Z_A \otimes O) \llbracket P'(\theta^*) \rrbracket; P(\theta^*) \llbracket (|0\rangle_A \langle 0|) \otimes \rho)) \\ \text{Fig1b, Sequence} \quad &= \text{tr}((Z_A \otimes O) \llbracket \mathcal{I}_A \otimes P(\theta^*) \rrbracket \llbracket P'(\theta^*) \rrbracket \llbracket (|0\rangle_A \langle 0|) \otimes \rho)) \\ &= \text{tr}((Z_A \otimes O) \llbracket \mathcal{I}_A \otimes P(\theta^*) \rrbracket \zeta), \quad (D.18) \end{aligned}$$

while

$$\llbracket (P(\theta^*))^*(O), \rho \rrbracket \rightarrow P'(\theta^*) \quad (D.19)$$

$$\begin{aligned} &= \text{tr}((Z_A \otimes \llbracket P(\theta^*)^*(O) \rrbracket) \llbracket P'(\theta^*) \rrbracket (|0\rangle_A \langle 0| \otimes \rho)) \\ &= \text{tr}((Z_A \otimes \llbracket P(\theta^*)^*(O) \rrbracket) \zeta). \end{aligned} \quad (D.20)$$

Now RHS (Eqn D.18) equals RHS (Eqn D.20) via duality, as $I_A \otimes P(\theta^*)$ does nothing on the ancilla, while $\llbracket P(\theta^*)^*(O) \rrbracket \in O_{\bar{v}}$ and Z_A are compatible observables. \square

Lemma D.3. Let $S_m(\theta) \in \text{add-q-while}_{\bar{v}}^{(T)}(\theta)$ ($\forall m \in [0, w]$, where $w \geq 1$). Denote:⁹

$$\begin{aligned} &\text{Compile}(S_m(\theta)) \\ &\equiv \{|Q_{m,0}(\theta), \dots, Q_{m,s_m}(\theta)\}, \end{aligned} \quad (D.21)$$

$$\begin{aligned} &\text{Compile}\left(\frac{\partial}{\partial \theta}(S_m(\theta))\right) \\ &\equiv \{|P_{m,0}(\theta), \dots, P_{m,t_m}(\theta)\}. \end{aligned} \quad (D.22)$$

Then for arbitrary $\rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$, $O \in O_{\bar{v}}$, we have the following computational properties regarding the observable semantics:

$$\begin{aligned} &\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta); S_1(\theta)) \rrbracket \\ &= \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)); \underline{S_1(\theta)} \rrbracket \\ &\quad + \llbracket (O, \rho) \rightarrow \underline{S_0(\theta)}; \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket, \end{aligned} \quad (D.23)$$

$$\begin{aligned} &\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(\text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end}) \rrbracket \\ &= \sum_m \sum_{i_m} \llbracket (O, M_m \rho M_m^\dagger) \rightarrow P_{m,i_m}(\theta) \rrbracket, \\ &\quad \text{where } i_m \text{ runs through } [0, t_m]; \end{aligned} \quad (D.24)$$

$$\begin{aligned} &\llbracket (O, \rho) \rightarrow \left(\frac{\partial}{\partial \theta}(S_0(\theta) + S_1(\theta)) \right) \rrbracket \\ &= \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)) \rrbracket + \end{aligned} \quad (D.25)$$

$$\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket \quad (D.26)$$

$$\begin{aligned} &\llbracket (O, \rho) \rightarrow \underline{S_0(\theta) + S_1(\theta)} \rrbracket \\ &= \llbracket (O, \rho) \rightarrow \underline{S_0(\theta)} \rrbracket + \llbracket (O, \rho) \rightarrow \underline{S_1(\theta)} \rrbracket, \end{aligned} \quad (D.27)$$

and

$$\begin{aligned} &\llbracket (O, \rho) \rightarrow \text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end} \rrbracket \\ &= \sum_m \sum_{j_m \in [0, s_m]} \llbracket (O, \mathcal{E}_m \rho) \rightarrow Q_{m,j_m}(\theta) \rrbracket \end{aligned} \quad (D.28)$$

⁹As a reminder: we shall frequently refer to Notations in Eqns D.21 and D.22 for the rest of the section.

Proof. We first make some direct application of the code transformation (hereinafter “CT”, Fig 4) and compilation (hereinafter “CP”, Fig 3) rules; some results will immediately follow from these application.

$$\begin{aligned} &\text{Compile}\left(\frac{\partial}{\partial \theta}(S_0(\theta); S_1(\theta))\right) \\ &\stackrel{\text{CT,Sequence}}{=} \text{Compile}\left(\underline{S_0(\theta)}; \frac{\partial}{\partial \theta}(S_1(\theta))\right) + \\ &\quad \left(\frac{\partial}{\partial \theta}(S_0(\theta)); \underline{S_1(\theta)}\right) \\ &\stackrel{\text{CP,Sum Components}}{=} \text{Compile}\left(\underline{S_0(\theta)}; \frac{\partial}{\partial \theta}(S_1(\theta))\right) \quad (D.29) \\ &\quad \coprod \text{Compile}\left(\frac{\partial}{\partial \theta}(S_0(\theta)); \underline{S_1(\theta)}\right) \\ &\stackrel{\text{CP,Sequence}}{=} \{ |Q_{0,g}(\theta); P_{1,h}(\theta)| \}_{g \in [s_0], h \in [t_1]} \quad (D.30) \\ &\quad \coprod \{ |P_{0,i}(\theta); Q_{1,j}(\theta)| \}_{i \in [t_0], j \in [s_1]} \quad (D.31) \end{aligned}$$

, where $[s_0], [s_1], [t_0], [t_1]$ stand for e.g. $[0, s_0], [0, s_1], [0, t_0], [0, t_1]$ respectively; and

$$\begin{aligned} &\text{Compile}\left(\frac{\partial}{\partial \theta}(\text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end})\right) \\ &\stackrel{\text{CT,Case}}{=} \text{Compile}\left(\text{case } M[\bar{q}] = \overline{m \rightarrow \frac{\partial}{\partial \theta}(S_m(\theta))} \text{ end}\right) \\ &\stackrel{\text{CP,Case}}{=} \text{FB}\left(\text{case } M[\bar{q}] = \overline{m \rightarrow \frac{\partial}{\partial \theta}(S_m(\theta))} \text{ end}\right); \end{aligned}$$

also,

$$\begin{aligned} &\text{Compile}\left(\frac{\partial}{\partial \theta}(S_0(\theta) + S_1(\theta))\right) \quad (D.32) \\ &\stackrel{\text{CT,Sum Component}}{=} \text{Compile}\left(\frac{\partial}{\partial \theta}(S_0(\theta)) + \frac{\partial}{\partial \theta}(S_1(\theta))\right) \quad (D.33) \end{aligned}$$

$$\stackrel{\text{CP,Sum Component}}{=} \{ |P_{0,0}(\theta), \dots, P_{0,t_0}| \} \quad (D.34)$$

$$\coprod \{ |P_{1,0}(\theta), \dots, P_{1,t_1}| \}. \quad (D.35)$$

As promised, Eqns D.23 and D.26 immediately follows from the results above together with the corresponding definitions. So is Eqn D.27, following the same lines of unfolding by CP-(ND component) laid out in Eqns D.33~D.35 then a direct pattern matching by definition. To obtain D.24, observe

that

$$\begin{aligned} & \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta} (\text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end}) \rrbracket \\ &= \sum_{Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = \overline{m \rightarrow \frac{\partial}{\partial \theta} (S_m(\theta))} \text{ end})} \llbracket (O, \rho) \rightarrow Q(\theta) \rrbracket \end{aligned} \quad (\text{D.36})$$

$$\llbracket (O, \rho) \rightarrow Q(\theta) \rrbracket \quad (\text{D.37})$$

$$= \sum_{Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = \overline{m \rightarrow \frac{\partial}{\partial \theta} (S_m(\theta))} \text{ end})} \text{tr}((Z_A \otimes O) \llbracket Q(\theta) \rrbracket (|0\rangle_A \langle 0| \otimes \rho)) \quad (\text{D.38})$$

$$\text{tr}((Z_A \otimes O) \llbracket Q(\theta) \rrbracket (|0\rangle_A \langle 0| \otimes \rho)) \quad (\text{D.39})$$

$$\begin{aligned} & \stackrel{A \not\subseteq \bar{q} \subseteq \bar{v}}{=} \sum_{m^*} \sum_{i_{m^*} \in [0, t_{m^*}]} \text{tr}((Z_A \otimes O) \cdot \\ & \quad (\llbracket P_{m^*, i_{m^*}}(\theta) \rrbracket ((|0\rangle_A \langle 0|) \otimes M_{m^*} \rho M_{m^*}^\dagger))) \quad (\text{D.40}) \\ &= \sum_{m^*} \sum_{i_{m^*}} \llbracket (O, M_{m^*} \rho M_{m^*}^\dagger) \rightarrow P_{m^*, i_{m^*}}(\theta) \rrbracket \text{ as desired.} \end{aligned}$$

A few more words on Eqn D.40. Since A is disjoint from $\bar{v} \supseteq \bar{q}$, for each $Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = \overline{m \rightarrow \frac{\partial}{\partial \theta} (S_m(\theta))} \text{ end})$, evolving one step per the operational semantics will lead to the configuration $\langle \llbracket P_{m^*, i_{m^*}}(\theta) \rrbracket, ((|0\rangle_A \langle 0|) \otimes M_{m^*} \rho M_{m^*}^\dagger) \rangle$ for some non-(essentially)-aborting $P_{m^*, i_{m^*}}(\theta)$. Besides, each such configuration appears exactly once affording one summand, due to construction of $\text{FB}(\bullet)$. On the other hand, if $P_{m^*, i_{m^*}}(\theta)$ essentially aborts, the trace computed as a summand vanishes, making no contribution to the sum, thereby maintaining the equation.

Eqn D.28 is proved using the same lines of logic:

$$\begin{aligned} & \llbracket (O, \rho) \rightarrow \text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end} \rrbracket \quad (\text{D.41}) \\ &= \sum_{Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end})} \llbracket (O, \rho) \rightarrow Q(\theta) \rrbracket \\ &= \sum_{Q(\theta) \in \text{FB}(\text{case } M[\bar{q}] = \overline{m \rightarrow S_m(\theta)} \text{ end})} \text{tr}(O \llbracket Q(\theta) \rrbracket \rho) \end{aligned} \quad (\text{D.42})$$

$$\begin{aligned} & \stackrel{A \not\subseteq \bar{q} \subseteq \bar{v}}{=} \sum_{m^*} \sum_{j_{m^*} \in [0, s_{m^*}]} \text{tr}(O \llbracket Q_{m^*, j_{m^*}}(\theta) \rrbracket M_{m^*} \rho M_{m^*}^\dagger) \\ &= \sum_{m^*} \sum_{j_{m^*} \in [0, s_{m^*}]} \llbracket (O, \mathcal{E}_{m^*} \rho) \rightarrow Q_{m^*, j_{m^*}}(\theta) \rrbracket \end{aligned}$$

as desired. \square

Lemma D.4. Keeping the notations in Lemma D.3, we have $\forall \theta^* \in \mathbb{R}^k, O \in \mathcal{O}_{\bar{v}}, \rho \in \mathcal{D}(\mathcal{H}_{\bar{v}})$,

$$\left(\frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow S_0(\theta); S_1(\theta) \rrbracket) \right) \Big|_{\theta^*} = \sum_{g \in [0, s_0]} \left(\frac{\partial}{\partial \theta} (\llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow S_1(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.43})$$

$$+ \sum_{j \in [0, s_1]} \left(\frac{\partial}{\partial \theta} (\llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow S_0(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.44})$$

Proof. Observe:

$$\begin{aligned} & \stackrel{\text{LHS(D.43)}}{=} \stackrel{\text{CP, Sequence}}{=} \sum_{g \in [0, s_0], j \in [0, s_1]} \left(\frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow Q_{0,g}(\theta); Q_{1,j}(\theta) \rrbracket) \right) \Big|_{\theta^*} \\ & \stackrel{(***) \text{, See Below}}{=} \sum_{g \in [0, s_0], j \in [0, s_1]} \left(\frac{\partial}{\partial \theta} (\llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow Q_{1,j}(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.45}) \end{aligned}$$

$$+ \sum_{g \in [0, s_0], j \in [0, s_1]} \left(\frac{\partial}{\partial \theta} (\llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow Q_{0,g}(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.46})$$

$$\begin{aligned} &= \sum_{g \in [0, s_0]} \left(\frac{\partial}{\partial \theta} (\llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow S_1(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.47}) \\ &+ \sum_{j \in [0, s_1]} \left(\frac{\partial}{\partial \theta} (\llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow S_0(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.48}) \end{aligned}$$

Where the last step is direct application of definition of input-space observable semantics. For the step (***) claiming $\text{RHS(D.45)} = \text{RHS(D.45)} + \text{RHS(D.46)}$, observe that it boils down to the following: for arbitrary $(g^*, j^*) \in [0, s_0] \times [0, s_1]$,

$$\left(\frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow Q_{0,g^*}(\theta); Q_{1,j^*}(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.49})$$

$$= \left(\frac{\partial}{\partial \theta} (\llbracket (O, \llbracket Q_{0,g^*}(\theta^*) \rrbracket \rho) \rightarrow Q_{1,j^*}(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.50})$$

$$+ \left(\frac{\partial}{\partial \theta} (\llbracket (\llbracket Q_{1,j^*}(\theta^*) \rrbracket^*(O), \rho) \rightarrow Q_{0,g^*}(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.51})$$

We argue the correctness of D.49~D.51 below. Simplifying notations, write $Q_{0,g^*}(\theta)$, $Q_{1,j^*}(\theta)$ as $Q_0(\theta)$, $Q_1(\theta)$ respectively, then let

$$Q_0(\theta) \equiv \sum_x K_x(\theta) \circ K_x^\dagger(\theta), \quad (\text{D.52})$$

$$Q_1(\theta) \equiv \sum_y J_y(\theta) \circ J_y^\dagger(\theta) \quad (\text{D.53})$$

denote the kraus operator decomposition of $Q_0(\theta)$, $Q_1(\theta)$, each with finitely many summands.¹⁰

Then,

$$\left(\frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow Q_g(\theta); Q_j(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.56})$$

$$= \frac{\partial}{\partial \theta_j} \text{tr}(O \llbracket Q_0(\theta); Q_1(\theta) \rrbracket \rho) \Big|_{\theta^*} \quad (\text{D.57})$$

$$= \frac{\partial}{\partial \theta_j} \text{tr}(O \sum_y J_y(\theta) (\sum_x K_x(\theta) \rho K_x^\dagger(\theta)) J_y^\dagger(\theta)) \Big|_{\theta^*} \quad (\text{D.58})$$

$$= \sum_{x,y} \frac{\partial}{\partial \theta_j} \text{tr}(O J_y(\theta) K_x(\theta) \rho K_x^\dagger(\theta) J_y^\dagger(\theta)) \Big|_{\theta^*} \quad (\text{D.59})$$

$$\stackrel{(a)}{=} \sum_y \left[\left(\text{tr}(O \frac{d}{d\theta_j} J_y(\theta) (\sum_x K_x(\theta^*) \rho K_x^\dagger(\theta^*)) J_y^\dagger(\theta^*)) \right. \right. \\ \left. \left. + \text{tr}(O J_y(\theta^*) (\sum_x K_x(\theta^*) \rho K_x^\dagger(\theta^*)) \frac{d}{d\theta_j} J_y^\dagger(\theta^*)) \right) \right] \Big|_{\theta^*} \\ + \sum_y \left[\left(\text{tr}(O J_y(\theta^*) (\sum_x \frac{d}{d\theta_j} K_x(\theta) \rho K_x^\dagger(\theta^*)) J_y^\dagger(\theta^*)) \right) \right. \\ \left. + \text{tr}(O J_y(\theta^*) (\sum_x K_x(\theta^*) \rho K_x^\dagger(\theta^*)) \frac{d}{d\theta_j} J_y^\dagger(\theta^*)) \right) \Big|_{\theta^*}$$

¹⁰This is doable because for each $\theta^* \in \mathbb{R}^k$ we have

$$Q_0(\theta^*) = \sum_{k=1} K_x(\theta^*) \circ K_x^\dagger(\theta^*), \quad (\text{D.54})$$

$$Q_1(\theta^*) = \sum_{q=1} J_y(\theta^*) \circ J_y^\dagger(\theta^*). \quad (\text{D.55})$$

Since parameterized unitaries have a parameterized-matrix representation, the superoperator $\mathcal{U}(\theta) = U(\theta) \circ U^\dagger(\theta)$ for any parameterized unitaries also have a parameterized operator representation. In our parameterized quantum program syntax only unitaries are parameterized, thus the kraus operators $K_x(\theta)$, $J_y(\theta)$'s can be chosen uniformly. If one insists one may perform a standard structural induction to show the existence of such a parameterized Kraus operator decomposition for any $Q(\theta) \in \mathbf{q}\text{-while}_{\overline{v}}^{(T)}(\theta)$. For example, assuming each $S_m(\theta)$ decomposes to $\sum_{t_m,j}^{t_m,jm} K_{t_m,j}(\theta) \circ K_{t_m,j}^\dagger(\theta)$, then case $M[\overline{q}] = m \rightarrow S_m(\theta)$ decomposes into

$$\sum_m (\sum_{t_m,j} K_{t_m,j}(\theta) \circ K_{t_m,j}^\dagger(\theta)) \circ \mathcal{E}_m.$$

As in the unitary case, the linear Kraus operators are entry-wisely smooth, ensured by the entry-wise smoothness of matrix representations of parameterized unitaries.

$$+ \text{tr}(O J_y(\theta^*) (\sum_x K_x(\theta^*) \rho \frac{d}{d\theta_j} K_x^\dagger(\theta^*)) J_y^\dagger(\theta^*)) \Big|_{\theta^*} \quad (\text{D.60})$$

$$\stackrel{(b)}{=} \frac{\partial}{\partial \theta_j} \text{tr}(O \llbracket Q_1(\theta) \rrbracket \llbracket Q_0(\theta^*) \rrbracket \rho) \Big|_{\theta^*} \quad (\text{D.61})$$

$$+ \frac{\partial}{\partial \theta_j} \text{tr}(O \llbracket Q_1(\theta^*) \rrbracket \llbracket Q_0(\theta) \rrbracket \rho) \Big|_{\theta^*} \quad (\text{D.62})$$

$$\stackrel{\text{S-H Dual}}{=} \left(\frac{\partial}{\partial \theta_j} (\llbracket (O, \llbracket Q_0(\theta^*) \rrbracket \rho) \rightarrow Q_1(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.63})$$

$$+ \left(\frac{\partial}{\partial \theta_j} (\llbracket (\llbracket Q_1(\theta^*) \rrbracket^*(O), \rho) \rightarrow Q_0(\theta) \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.64})$$

where $\frac{d}{d\theta_j} \bullet$ again denotes the entry-wise derivative, and (a), (b) is obtained by unfolding the definition of trace, noticing that all entries of parameterized unitaries are smooth. S-H dual here is short for Schrödinger-Heisenberg dual, as is everywhere else in this section. \square

D.2 Proof Details of Theorem 6.2: Soundness of Differentiation Logic

Proof. As stated in the main text, throughout we fix $O_A \equiv Z_A$. Unfolding by definition, it suffices to show

$$\forall O \in \mathcal{O}_{\overline{v}}, \rho \in \mathcal{D}(\mathcal{H}_{\overline{v}}), \llbracket ((O, Z_A), \rho) \rightarrow \frac{\partial}{\partial \theta} (S(\theta)) \rrbracket \\ = \frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket). \quad (\text{D.65})$$

Let $O \in \mathcal{O}_{\overline{v}}$, $\rho \in \mathcal{D}(\mathcal{H}_{\overline{v}})$ be arbitrary. We consider each rule in Figure 5 possibly used as the final step of the derivation.

- (Abort)~(Trivial Unitary): for $S(\theta) \equiv \text{abort}[\overline{v}]$, $\text{skip}[\overline{v}]$ or $q := |0\rangle$, $\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket$ is a constant, so $\frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow S(\theta) \rrbracket) = 0$. On the other hand, $\frac{\partial}{\partial \theta} (S(\theta)) \stackrel{\text{CT, trivialize}}{=} \text{abort}[\overline{v} \cup \{A\}]$, meaning $\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta} (S(\theta)) \rrbracket = \text{tr}((Z_A \otimes O) \cdot 0) = 0$. In Trivial-Unitary, $\llbracket (O, \rho) \rightarrow U(\theta) \rrbracket$ doesn't depend on θ_j since $\theta_j \notin \theta$; hence $\frac{\partial}{\partial \theta_j} (\llbracket (O, \rho) \rightarrow U(\theta) \rrbracket) = 0$, as faithfully represented by the code transformation rule.
- (Unitary): Assume $S(\theta) \equiv U(\theta)$ with $U(\theta) \in \{R_\sigma(\theta), R_{\sigma \otimes \sigma}(\theta)\}_{\sigma \in \{X, Y, Z\}}$. Let σ range over $\{\sigma, \sigma \otimes \sigma\}_{\sigma \in X, Y, Z}$ (proof below works for all 6 cases where $\sigma^2 = I_\bullet$). Let $\frac{d}{d\theta} U(\theta)$ denote the entry-wise derivative of $U(\theta)$, then

recall from Lemma D.1 that

$$\frac{d}{d\theta}U(\theta) = \frac{1}{2}U(\theta + \pi); \quad (\text{D.66})$$

$$\begin{aligned} & \frac{\partial}{\partial\theta}(\llbracket(O, \rho) \rightarrow \underline{U(\theta)}\rrbracket) \\ &= \text{tr}(O \cdot U(\theta) \cdot \rho \cdot (\frac{d}{d\theta}U(\theta))^\dagger) \\ &+ \text{tr}(O \cdot \frac{d}{d\theta}U(\theta) \cdot \rho \cdot U^\dagger(\theta)). \quad (\text{D.67}) \end{aligned}$$

$$\stackrel{\text{D.66}}{=} \frac{1}{2}\text{tr}\left(O\left(U(\theta)\rho U^\dagger(\theta + \pi) + U(\theta + \pi)\rho U^\dagger(\theta)\right)\right)$$

Meanwhile,

$$\begin{aligned} & \llbracket(O, \rho) \rightarrow \frac{\partial}{\partial\theta}(\underline{U(\theta)})\rrbracket \\ & \stackrel{\text{CT,1-qb,2-qb}}{=} \text{tr}((Z_A \otimes O)\llbracket R'_\sigma(\theta) \rrbracket \cdot \\ & ((|0\rangle_A\langle 0|) \otimes \rho)) \quad (\text{D.68}) \end{aligned}$$

Let us unfold the definition of $R'_\sigma(\theta)$:

$$\begin{aligned} & \text{RHS(D.68)} \\ &= \frac{1}{2} \sum_{x,y \in \{0,1\}} \text{tr}((Z_A \otimes O)\llbracket C - R_\sigma(\theta); H_A \rrbracket \cdot \\ & ((|x\rangle_A\langle y|) \otimes \rho)) \\ &= \frac{1}{2}\text{tr}\left((Z_A \otimes O)\llbracket H_A \rrbracket(|0\rangle_A\langle 0| \otimes U(\theta)\rho U^\dagger(\theta) \right. \\ & + |0\rangle_A\langle 1| \otimes U(\theta)\rho U^\dagger(\theta + \pi) \quad (\text{D.69}) \\ & + |1\rangle_A\langle 0| \otimes U(\theta + \pi)\rho U^\dagger(\theta) \quad (\text{D.70}) \\ & \left. + |1\rangle_A\langle 1| \otimes U(\theta + \pi)\rho U^\dagger(\theta + \pi)\right) \quad (\text{D.71}) \end{aligned}$$

$$\begin{aligned} &= \frac{1}{2}\text{tr}\left((Z_A \otimes O)\left(\frac{\sum_{x,y} |x\rangle_A\langle y|}{2} \otimes U(\theta)\rho U^\dagger(\theta) \right. \quad (\text{D.72}) \\ & + \frac{\sum_x |x\rangle_A\langle 0| - \sum_x |x\rangle_A\langle 1|}{2} U(\theta)\rho U^\dagger(\theta + \pi) \quad (\text{D.73}) \\ & + \frac{\sum_x |0\rangle_A\langle x| - \sum_x |1\rangle_A\langle x|}{2} U(\theta + \pi)\rho U^\dagger(\theta) \quad (\text{D.74}) \\ & \left. + \frac{\sum_x |x\rangle_A\langle x| - (|1\rangle_A\langle 0| + |0\rangle_A\langle 1|)}{2} \cdot \right. \\ & \left. U(\theta + \pi)\rho U^\dagger(\theta + \pi)\right) \quad (\text{D.75}) \end{aligned}$$

Let's regroup the terms of D.72 ~ D.75, bearing in mind that when observing $Z_A \otimes O$ on the final state,

the “off-diagonal” part is killed by Z_A .¹¹ Hence,

$$\text{RHS(D.72} \sim \text{D.75)} \quad (\text{D.76})$$

$$= \frac{1}{4}\text{tr}\left((Z_A \otimes O)\left(|1\rangle_A\langle 1| + |0\rangle_A\langle 0|\right) \quad (\text{D.77})\right.$$

$$\otimes \left(U(\theta)\rho U^\dagger(\theta) + U(\theta + \pi)\rho U^\dagger(\theta + \pi)\right) \quad (\text{D.78})$$

$$+ \left(|0\rangle_A\langle 0| - |1\rangle_A\langle 1|\right) \quad (\text{D.79})$$

$$\left.\otimes \left(U(\theta)\rho U^\dagger(\theta + \pi) + U(\theta + \pi)\rho U^\dagger(\theta)\right)\right) \quad (\text{D.80})$$

Lastly, the behavior of $Z_A \otimes O$ ensures

$$\text{RHS(D.77} \sim \text{D.80)} = \frac{1}{2}\text{tr}\left(O\left(U(\theta)\rho U^\dagger(\theta + \pi) + U(\theta + \pi)\rho U^\dagger(\theta)\right)\right),$$

as desired. From this point, we again adopt the notations from Lemma D.3 Eqns D.21, D.22.

3. (Sum Component)

$$\llbracket(O, \rho) \rightarrow \left(\frac{\partial}{\partial\theta}(\underline{S_0(\theta) + S_1(\theta)})\right)\rrbracket$$

$$\stackrel{\text{Lem D.3}}{=} \llbracket(O, \rho) \rightarrow \frac{\partial}{\partial\theta}(\underline{S_0(\theta)})\rrbracket \quad (\text{D.81})$$

$$+ \llbracket(O, \rho) \rightarrow \frac{\partial}{\partial\theta}(\underline{S_1(\theta)})\rrbracket \quad (\text{D.82})$$

$$\stackrel{\text{IH}}{=} \frac{\partial}{\partial\theta}(\llbracket(O, \rho) \rightarrow \underline{S_0(\theta)}\rrbracket \quad (\text{D.83})$$

$$+ \llbracket(O, \rho) \rightarrow \underline{S_1(\theta)}\rrbracket) \quad (\text{D.84})$$

$$\stackrel{\text{Lem D.3}}{=} \frac{\partial}{\partial\theta}(\llbracket(O, \rho) \rightarrow \underline{S_0(\theta) + S_1(\theta)}\rrbracket), \quad (\text{D.85})$$

as desired.

4. (Sequence) Assume $\underline{S(\theta)} \equiv \underline{S_0(\theta)}; \underline{S_1(\theta)}$; It suffices to show that $\forall \theta^*$,

$$\llbracket(O, \rho) \rightarrow \frac{\partial}{\partial\theta}(\underline{S_0(\theta^*)}; \underline{S_1(\theta^*)})\rrbracket = \left(\frac{\partial}{\partial\theta}(\llbracket(O, \rho) \rightarrow \underline{S_0(\theta)}; \underline{S_1(\theta)}\rrbracket)\right)\Big|_{\theta^*} \quad (\text{D.86})$$

Let us first manipulate the equation using some computational properties developed from the helper lemmas:

$$\text{LHS(D.86)} \quad (\text{D.87})$$

$$\stackrel{\text{Lemma D.3, Eqn D.23}}{=} \llbracket(O, \rho) \rightarrow \frac{\partial}{\partial\theta}(\underline{S_0(\theta^*)}); \underline{S_1(\theta^*)})\rrbracket$$

$$+ \llbracket(O, \rho) \rightarrow \underline{S_0(\theta^*)}; \frac{\partial}{\partial\theta}(\underline{S_1(\theta^*)})\rrbracket$$

¹¹Namely, the trace computation can be decomposed into a finite linear combination of the form $\text{tr}((Z_A \otimes O)(|b\rangle_A\langle b'| \otimes \bullet))$ ($b, b' \in \{0, 1\}$), and this term vanishes whenever $b \neq b'$.

This equals, by Lem D.3 Eqns D.30,D.31,

$$\sum_{g \in [0, s_0], h \in [0, t_1]} \llbracket (O, \rho) \rightarrow Q_{0,g}(\theta^*); P_{1,h}(\theta^*) \rrbracket \quad (\text{D.88})$$

$$+ \sum_{i \in [0, t_0], j \in [0, s_1]} \llbracket (O, \rho) \rightarrow P_{0,i}(\theta^*); Q_{1,j}(\theta^*) \rrbracket \quad (\text{D.89})$$

Via Lem D.2, this translates to

$$\sum_{g \in [0, s_0], h \in [0, t_1]} \llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow P_{1,h}(\theta^*) \rrbracket \quad (\text{D.90})$$

$$+ \sum_{i \in [0, t_0], j \in [0, s_1]} \llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow P_{0,i}(\theta^*) \rrbracket$$

$$= \sum_{g \in s_0} \llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow \frac{\partial}{\partial \theta} (\underline{S_1(\theta^*)}) \rrbracket \quad (\text{D.91})$$

$$+ \sum_{j \in s_1} \llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow \frac{\partial}{\partial \theta} (\underline{S_0(\theta^*)}) \rrbracket, \quad (\text{D.92})$$

where the last step comes from definition of observable semantics with ancilla. Now apply the inductive hypothesis on $S_1(\theta)$, a universal statement for all $(\rho', O') \in \mathcal{D}(\mathcal{H}_{\bar{v}}) \times \mathcal{O}_{\bar{v}}$, to the instances $(\llbracket Q_{0,g}(\theta^*) \rrbracket \rho, O)$ ($\forall g \in [0, s_0]$), we have the first summand

$$\begin{aligned} & \text{RHS(D.91)} \\ &= \sum_g \left(\frac{\partial}{\partial \theta} (\llbracket (O, \llbracket Q_{0,g}(\theta^*) \rrbracket \rho) \rightarrow \underline{S_1(\theta)} \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.93}) \end{aligned}$$

Likewise, apply the inductive hypothesis on $S_0(\theta)$ to the instances $(\rho, \llbracket Q_{1,j}(\theta^*) \rrbracket^*(O))$ ($\forall j \in [0, s_1]$), we have the second summand

$$\begin{aligned} & \text{RHS(D.92)} \\ &= \sum_j \left(\frac{\partial}{\partial \theta} (\llbracket (\llbracket Q_{1,j}(\theta^*) \rrbracket^*(O), \rho) \rightarrow \underline{S_0(\theta)} \rrbracket) \right) \Big|_{\theta^*} \quad (\text{D.94}) \end{aligned}$$

Collecting everything, using again a technical lemma from above,

$$\begin{aligned} & \text{RHS(D.91)} + \text{RHS(D.92)} \\ &= \text{RHS(D.93)} + \text{RHS(D.94)} \quad (\text{D.95}) \end{aligned}$$

$$\stackrel{\text{Lem D.4}}{=} \text{RHS(D.86), as desired.} \quad (\text{D.96})$$

5. (Case) This follows again from the computational lemmas above, and applications of all inductive hypothesis. For conciseness, let us denote by $\text{IH}(m) \mapsto (M_m \rho M_m^\dagger, O)$ the application of the m -th ($m \in [0, w]$) inductive hypothesis to the instance $(M_m \rho M_m^\dagger, O) \in \mathcal{D}(\mathcal{H}_{\bar{v}}) \times \mathcal{O}_{\bar{v}}$. Then,

$$\begin{aligned} & \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta} (\underline{\text{case } M[\bar{q}] = m \rightarrow \underline{S_m(\theta)} \text{ end}}) \rrbracket \\ \stackrel{\text{Lem D.3}}{=} & \sum_m \sum_{i_m \in [0, t_m]} \llbracket (O, M_m \rho M_m^\dagger) \rightarrow P_{m,i_m}(\theta) \rrbracket \quad (\text{D.97}) \end{aligned}$$

$$= \sum_m \llbracket (O, M_m \rho M_m^\dagger) \rightarrow \underline{\frac{\partial}{\partial \theta} (S_m(\theta))} \rrbracket \quad (\text{D.98})$$

Applying the inductive hypothesis $\text{IH}(m) \mapsto (M_m \rho M_m^\dagger, O)$, $\forall m$, we get that the above equals

$$\begin{aligned} & \sum_m \frac{\partial}{\partial \theta} (\llbracket (O, (M_m \rho M_m^\dagger)) \rrbracket \rightarrow \underline{S_m(\theta)} \rrbracket) \\ &= \frac{\partial}{\partial \theta} (\sum_m \sum_{j_m \in [0, s_m]} \llbracket (O, \mathcal{E}_{m,j_m}) \rrbracket \rightarrow \underline{Q_{m,j_m}(\theta)} \rrbracket) \\ &\stackrel{\text{Lem D.3}}{=} \frac{\partial}{\partial \theta} (\llbracket (O, \rho) \rightarrow \text{case } M[\bar{q}] = m \rightarrow \underline{S_m(\theta)} \text{ end} \rrbracket) \end{aligned}$$

, as desired.

6. (While^T) This is verified by successive application of (Case) rule and (Sequence) rule T times. \square

E Detailed Proofs from Section 7

E.1 Proof for Proposition 7.2: Upper Bound of

$$|\# \frac{\partial}{\partial \theta_j} (P(\theta))|$$

- Proof.* 1. If $P(\theta) \equiv \text{abort}[\bar{v}] \text{ skip}[\bar{v}] | q := |0\rangle$ ($q \in \bar{v}$), then $\text{OC}_j(P(\theta)) = 0 = |\# \frac{\partial}{\partial \theta_j} (P(\theta))|$ by definition. Similarly, if $P(\theta) \equiv U(\theta)$ trivially used θ_j , no non-aborting programs exists in $\text{Compile}(\frac{\partial}{\partial \theta} (P(\theta)))$, yielding $|\# \frac{\partial}{\partial \theta_j} (P(\theta))| = 0$; otherwise, $|\# \frac{\partial}{\partial \theta_j} (P(\theta))|$ consists of a single $R'_\sigma(\theta)$, so both numbers are 1.
2. Assume $|\# \frac{\partial}{\partial \theta_j} (P_b(\theta))| \leq \text{OC}_j(P_b(\theta))$ for each $b \in \{1, 2\}$ and $P(\theta) \equiv P_1(\theta); P_2(\theta)$. Then

$$\begin{aligned} & |\# \frac{\partial}{\partial \theta_j} (P(\theta))| \\ &\stackrel{(*)}{=} |\#(P_1(\theta); \frac{\partial}{\partial \theta_j} (P_2(\theta)))| \\ &+ |\#(\frac{\partial}{\partial \theta_j} (P_1(\theta)); P_2(\theta))| \quad (\text{E.1}) \end{aligned}$$

$$\begin{aligned} & \leq |\#(\frac{\partial}{\partial \theta_j} (P_2(\theta)))| \\ &+ |\#(\frac{\partial}{\partial \theta_j} (P_1(\theta)))| \quad (\text{E.2}) \end{aligned}$$

$$\stackrel{\text{Inductive Hypothesis}}{\leq} \text{OC}_j(P_1(\theta)) + \text{OC}_j(P_2(\theta)) \quad (\text{E.3})$$

$$= \text{OC}_j(P(\theta)). \quad (\text{E.4})$$

where step $(*)$ is via rules CT,CP-(ND-Component).

3. Assume $|\# \frac{\partial}{\partial \theta_j} (P_m(\theta))| \leq \text{OC}_j(P_m(\theta))$ for each $m \in [0, w]$ and $P(\theta) \equiv \text{case } M[\bar{q}] = m \rightarrow \underline{P_m(\theta)} \text{ end}$. Then by the CT and CP rules of case, $\frac{\partial}{\partial \theta_j} (P(\theta))$ compiles to $\max_m |\# \frac{\partial}{\partial \theta_j} (P_m(\theta))|$ programs. Assume $m^* \in [0, w]$ is

s.t. $|\# \frac{\partial}{\partial \theta_j}(P_{m^*}(\theta))|$ attains that maximum. Then,

$$|\# \frac{\partial}{\partial \theta_j}(P(\theta))| = |\# \frac{\partial}{\partial \theta_j}(P_{m^*}(\theta))| \quad (\text{E.5})$$

$$\stackrel{\text{Inductive Hypothesis}}{\leq} \text{OC}_{m^*}(P_{m^*}(\theta)) \quad (\text{E.6})$$

$$\leq \max_m \text{OC}_m(P_m(\theta)) \quad (\text{E.7})$$

$$= \text{OC}_j(P(\theta)) \quad (\text{E.8})$$

, as desired.

4. If $P(\theta) \equiv \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done}$, it suffices to show $|\# \frac{\partial}{\partial \theta}(P(\theta))| \leq T \cdot |\# \frac{\partial}{\partial \theta}(P_1(\theta))|$. This is true for $T = 1$ as $\frac{\partial}{\partial \theta}(\text{while}^{(1)})$ essentially aborts. By induction if this is true for T , then for $T + 1$ we have

$$\begin{aligned} & |\# \frac{\partial}{\partial \theta}(\text{while}^{(T+1)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done})| \quad (\text{E.9}) \\ & \leq |\# \frac{\partial}{\partial \theta}(P_1(\theta); \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done})| \\ & \stackrel{\text{Eqn E.2}}{\leq} |\# \frac{\partial}{\partial \theta}(P_1(\theta))| \\ & \quad + |\# \frac{\partial}{\partial \theta}(\text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done})| \\ & \leq (T + 1) |\# \frac{\partial}{\partial \theta}(P_1(\theta))|, \text{ as desired.} \quad (\text{E.10}) \end{aligned}$$

□

F Evaluation Details

F.1 Training VQC instances with controls

For any parameter α , we exhibit the final results of $\text{Compile}(\frac{\partial P_1(\Theta, \Phi)}{\partial \alpha})$ and $\text{Compile}(\frac{\partial P_2(\Theta, \Phi, \Psi)}{\partial \alpha})$, with $P_1(\Theta, \Phi)$ and $P_2(\Theta, \Phi, \Psi)$ defined in the main text:

1. If $\alpha \in \Theta$: without loss of generality assume $\alpha = \theta_1$ (other situations are completely analogous). We denote:

$$Q'(\Theta) \equiv R'_X(\theta_1)|A, q_1\rangle; \dots; R_Z(\theta_{12})|q_{12}\rangle$$

then

$$\begin{aligned} \text{Compile}(\frac{\partial P_1(\Theta, \Phi)}{\partial \alpha}) &\equiv \{ | Q'(\Theta); Q(\Phi) | \}, \\ \text{Compile}(\frac{\partial P_2(\Theta, \Phi, \Psi)}{\partial \alpha}) &\equiv \{ | Q'(\Theta); \\ &\quad \text{case } M[q_1] = 0 \rightarrow Q(\Phi) \\ &\quad \quad 1 \rightarrow Q(\Psi) | \}. \end{aligned}$$

2. If $\alpha \in \Phi$: without loss of generality assume $\alpha = \phi_1$ (other situations are completely analogous). We denote:

$$Q'(\Phi) \equiv R'_X(\phi_1)|A, q_1\rangle; \dots; R_Z(\phi_{12})|q_{12}\rangle$$

$$\text{Compile}(\frac{\partial P_1(\Theta, \Phi)}{\partial \alpha}) \equiv \{ | Q(\Theta); Q'(\Phi) | \},$$

$$\begin{aligned} \text{Compile}(\frac{\partial P_2(\Theta, \Phi, \Psi)}{\partial \alpha}) &\equiv \{ | Q(\Theta); \\ &\quad \text{case } M[q_1] = 0 \rightarrow Q'(\Phi) \\ &\quad \quad 1 \rightarrow Q(\Psi) | \}. \end{aligned}$$

3. If $\alpha \in \Psi$: without loss of generality assume $\alpha = \psi_1$ (other situations are completely analogous). We denote:

$$Q'(\Psi) \equiv R'_X(\psi_1)|A, q_1\rangle; \dots; R_Z(\psi_{12})|q_{12}\rangle$$

$$\begin{aligned} \text{Compile}(\frac{\partial P_1(\Theta, \Phi)}{\partial \alpha}) &\equiv \{ | \text{abort}[A, q_1, \dots, q_{12}] | \}, \\ \text{Compile}(\frac{\partial P_2(\Theta, \Phi, \Psi)}{\partial \alpha}) &\equiv \{ | Q(\Theta); \\ &\quad \text{case } M[q_1] = 0 \rightarrow Q(\Phi) \\ &\quad \quad 1 \rightarrow Q'(\Psi) | \}. \end{aligned}$$

F.2 Benchmark testing on representative VQCs

In this section we introduce three examples from real-world quantum machine learning and quantum approximation algorithms, all of which are very promising candidates for implementation on near-term quantum devices. Without loss of generality, throughout this section θ denotes θ_1 .

Our first case, QNN_* , starts from a slightly simplified case from an actual quantum neural-network that has been implemented on ion-trap quantum machines [56]. QAOA_* is a quantum algorithm that produces approximate solutions for combinatorial optimization problems, which is regarded as one of the most promising candidates for demonstrating quantum supremacy [15, 47]. Quantum eigensolvers, crucial to quantum phase estimation algorithms, usually requires fully coherent evolution. In [35] Peruzzo et al. introduced an alternative approach that relaxes the fully coherent evolution requirement combined with a new approach to state preparation based on ansatz and classical optimization [21, 30, 35], namely our third example VQE_* . The control flow creating multi-layer structures for the three families of examples are very similar, and these algorithms mainly differ from each other in their basic “rotation-entanglement” blocks. Therefore I will introduce the basic blocks for all three families, then use QNN_* as an example to illustrate how the control flow (if, bounded while) works.

The basic “rotate-entangle” building block for QNN consists of a rotation stage and an entanglement stage – we consider these two stages together a single layer: in the rotation stage, one performs parameterized Z rotations followed by parameterized X rotations and then again parameterized Z rotations on the first 4 (small scale) or 6 qubits (medium or large scale); in the entanglement stage, one performs the parameterized $X \otimes X$ rotation on all pairs from the first 4 or 6 qubits. See Figure 7.

Basic block for VQE consists of three stages: the first stage is parameterized X followed by parameterized Z ; the second




Figure 7. Circuit representation of the basic building block of QNN. Figure credit: [56]

stage uses H and CNOT to entangle; the third stage performs parameterized Z, X, Z in that order. Basic block for QAOA, on the other hand, entangles using H and CNOT in the first stage, and then performs parameterized X rotations on the second stage.

The more interesting part lies in building multiple layers using control flow, which we will explain using $\text{QNN}_{*,w}$. The small scale if-controlled QNN (denoted as $\text{QNN}_{S,i}$) is two-layered. Let B denote the basic rotate-entangle block for QNN explained as above; B', B'' be two similarly-structured rotate-entangle blocks using different parameter-qubit combinations. $\text{QNN}_{S,i}$ performs B as the first layer, then measures on the first qubit, and performs B' or B'' as the second layer dependent on the measurement output. For medium and large scale if-controlled QNN (i.e. $\text{QNN}_{M,i}$ and $\text{QNN}_{L,i}$) we have larger rotate-entangle blocks, various parameter-qubit combinations (hence more of the B', B'' 's involving different parameter-qubit combinations), and more layers of measurement-based control.

(Bounded) while-controlled QNN works similarly: take $\text{QNN}_{S,w}$ as an example, it performs the rotate-entangle block B as the first layer, then measure the first qubit; if it outputs 0 we halt; otherwise we perform some B' , measure qubit $q1$ again, halts if outputs 0, performs B' the third time and aborts otherwise. Note that this is just a verbose way of saying “we wrapped B' with a 2-bounded while-loop”! And similarly, we build more layers on larger systems ($\text{QNN}_{M,w}, \text{QNN}_{L,w}$) using larger rotate-entangle blocks and more layers of bounded-while loops. One should note how bounded while is a succinct way to represent the circuits: as shown in Table 3, in $\text{QNN}_{L,w}$ we managed to represent 2079 unitary gates in just 244 lines of code.

We auto-differentiated the three families of quantum programs using our code transformation (hereinafter “CT”) and code compilation (hereinafter “CP”) rules. As shown in Table 3, the computation outputs the desired multiset of derivative programs, and the number of non-aborting programs ($|\# \frac{\partial}{\partial \theta}(P(\theta))|$) agrees with the upper bound described in Proposition 7.2. It should be noted that Table 3 indicates our auto-differentiation scheme works well for variously sized input programs, be the size measured by code length, gate count, layer count or qubit count, to name a few. In particular, a 40-qubit system is very close to a real world

quantum machine learning system developed by Google. For details of the code and the differentiation result, please see supplementary materials.

$P(\theta)$	$\text{OC}(\cdot)$	$ \# \frac{\partial}{\partial \theta}(\cdot) $	#gates	#lines	#layers	#qb's
$\text{QNN}_{S,b}$	1	1	20	24	1	4
$\text{QNN}_{S,s}$	5	5	20	24	1	4
$\text{QNN}_{S,i}$	10	10	60	67	2	4
$\text{QNN}_{S,w}$	15	10	60	66	3	4
$\text{QNN}_{M,i}$	24	24	165	189	3	18
$\text{QNN}_{M,w}$	56	24	231	121	5	18
$\text{QNN}_{L,i}$	48	48	363	414	6	36
$\text{QNN}_{L,w}$	504	48	2079	244	33	36
$\text{VQE}_{S,b}$	1	1	14	16	1	2
$\text{VQE}_{S,s}$	2	2	14	16	1	2
$\text{VQE}_{S,i}$	4	4	28	38	2	2
$\text{VQE}_{S,w}$	6	4	42	32	3	2
$\text{VQE}_{M,i}$	15	15	224	241	3	12
$\text{VQE}_{M,w}$	35	15	224	112	5	12
$\text{VQE}_{L,i}$	40	40	576	628	5	40
$\text{VQE}_{L,w}$	248	40	1984	368	17	40
$\text{QAOA}_{S,b}$	1	1	12	15	1	3
$\text{QAOA}_{S,s}$	3	3	12	15	1	3
$\text{QAOA}_{S,i}$	6	6	36	41	2	3
$\text{QAOA}_{S,w}$	9	6	36	29	3	3
$\text{QAOA}_{M,i}$	18	18	120	142	3	18
$\text{QAOA}_{M,w}$	42	18	168	94	5	18
$\text{QAOA}_{L,i}$	36	36	264	315	6	36
$\text{QAOA}_{L,w}$	378	36	1512	190	33	36

Table 3. Output on Test Examples. Note that: (1) $\{S, M, L\}$ stands for “small, medium, large”; $\{b, s, i, w\}$ stands for “basic, shared, if, while”; #lines is the number of lines for input programs. (2) gate count and layer count for T -bounded while is done by multiplying the corresponding count for loop body with T . (3) for $*,w$ we have $|\# \frac{\partial}{\partial \theta}(P(\theta))| < \text{OC}(P(\theta))$ because differentiating the unrolled bounded while generates essentially aborting programs which are optimized out of the final multiset.