

Quantum Hoare Type Theory

Kartik Singhal

University of Chicago

ks@cs.uchicago.edu

As quantum computers become real, it is high time we come up with effective techniques that help programmers write correct quantum programs. Inspired by Hoare Type Theory in classical computing, we propose Quantum Hoare Type Theory (QHTT), in which precise specifications about the modification to the quantum state can be provided within the type of computation. These specifications within a Hoare type are given in the form of Hoare-logic style pre- and postconditions following the propositions-as-types principle. The type-checking process verifies that the implementation conforms to the provided specification. QHTT has the potential to be a unified system for programming, specifying, and reasoning about quantum programs.

Contents

1	Introduction	2
2	Background	4
2.1	Floyd–Hoare logic	4
2.2	Hoare Type Theory	6
2.3	Quantum computation	7
2.4	Quantum Hoare logic with ghost variables	8
3	Quantum Hoare Type Theory	11
3.1	Primitive quantum commands	11
3.1.1	Initialization	11
3.1.2	Measurement	12
3.1.3	Unitary application	12
3.2	Syntax	12
3.3	Examples	15
3.3.1	Bell states	15
3.3.2	Quantum teleportation	16
3.3.3	Modular teleportation with classical control	17
3.3.4	Deutsch’s algorithm	22
4	Discussion	22
5	Conclusions and perspectives	25
	References	26
A	Grammar	31

B Statics	32
B.1 Valid types	32
B.2 Terms	32
B.3 Computation rules	33
B.4 Rules for primitive quantum commands	33

List of Tables

1 Common unitary gates in quantum computation	9
2 Propositional logic connectives	14

List of Figures

1 Circuit to prepare the first Bell state	15
2 Quantum teleportation circuit	16
3 Quantum circuit for Deutsch’s algorithm	22

Listings

1 Generating the first Bell state	15
2 Bell state program annotated with propositions	16
3 Quantum teleportation	17
4 Annotated quantum teleportation program	18
5 Alice’s portion of teleportation	19
6 Bob’s portion of teleportation	19
7 Modular teleportation	19
8 Annotated code for alice	20
9 Annotated code for bob	21
10 Annotated code for modular teleportation	21
11 Deutsch’s algorithm	23
12 Fair coin toss with an invalid postcondition	24

1 Introduction

Quantum computation is fast becoming a reality with the advent of real machines and Google’s recent breakthrough [Arute et al. 2019] in demonstrating so-called *quantum supremacy*.¹ With advances in hardware, there is an acceleration in the development of software for quantum machines, which in turn requires the development of special-purpose programming languages. Several efforts in this direction, both from academia and industry, have recently led to many programming languages aimed at quantum computing such as Q# [Svore et al. 2018], Quipper [Green et al. 2013], Scaffold [Javadi-Abhari et al. 2012], and Silq [Bichsel et al. 2020].

¹A term coined by Preskill [2012] to describe the moment when a quantum computer could perform a computational task that no classical computer could. Alternatives include *quantum advantage* [Palacios-Berraquero, Mueck, and Persaud 2019] and *quantum primacy* [Durham, Garisto, and Wiesner 2021].

There is also a sudden need to prepare and train the upcoming generation of quantum programmers who sometimes have little to no previous exposure to the fundamental concepts in quantum computation. If we were to put ourselves in the position of a beginner to quantum programming, we might start by learning basic concepts such as what is a qubit, what does it mean for a qubit to be in a superposition of 0 and 1, and the most perplexing of all: entanglement. We may then try to write some simple programs in a suitable quantum programming language to aid our understanding. We may even be able to run some of these programs on simulators and even on real machines. We will make mistakes when writing these programs, however, and while sometimes those mistakes may be apparent from the output, a lot of times (as a beginner), they may be not. What can we do?

A common method to test our assumptions while programming in the classical realm is to write assertions about what we believe to be the truth at a particular point in the program. These assertions are usually dynamic in nature, meaning they are only tested while running the program and involve a cycle of writing, running, debugging, and rerunning the program. While this may be a suitable approach in classical computing, it is quite wasteful of the limited quantum resources if we were to apply it to quantum programming. Furthermore, this approach is impossible to use if the number of qubits is just over 50 or so as we hit the limits of both current quantum machines and classical simulation. Despite these limitations, there have been many recent proposals to use run-time assertions for debugging quantum programs [Huang and Martonosi 2019; Li et al. 2020].

What can we from the programming languages community offer that can help the upcoming workforce of quantum programmers avoid mistakes in their programs?

In classical software development, the use of strong static type systems has been proven to be hugely beneficial in avoiding large classes of bugs before running the programs. Recent mainstream languages such as Rust [Matsakis and Klock II 2014] even manage to statically prevent memory safety bugs [Jung et al. 2017; 2021; Xu et al. 2021] that have been notorious for being the root causes of a large swath of security vulnerabilities found in production software written in unsafe languages such as C. There is a similar need for innovation in type systems for quantum programming languages.

Recent academic languages, such as QWire [Paykin 2018; Paykin, Rand, and Zdancewic 2017] and Silq [Bichsel et al. 2020], that employ a linear type system to statically enforce the no-cloning theorem of quantum mechanics show promise in the approach of statically preventing bugs in the quantum domain. However, semantic properties about the quantum state, such as whether a qubit is in superposition or whether it is entangled with some other qubit, have not been tackled with the help of types. Furthermore, assertions in the form of pre- and postconditions around small pieces of code have shown the most promise in the quantum domain [Huang and Martonosi 2018; 2019]. This seems fair: we, as humans, can only hold so much information about the quantum state in our brains, it makes sense to reason about pieces of quantum code in a modular manner.

The most common approach to reason with pre- and postconditions in the classical domain is to use variations of Floyd-Hoare logic that utilize specification triples of the form $\{P\} c \{Q\}$, where c is the program to be executed, P a proposition on the initial state associated with the program (precondition), and Q that on the final state (postcondition). The intuitive reading of such a triple is that if the precondition holds on the program state before executing the program, and the program terminates, then the postcondition must also hold. These triples, along with a set of axioms and inference rules for composing them together, have proven to be quite effective in proving the correctness of classical state-manipulating (imperative) programs over the last several decades. It is no surprise that similar logics have appeared in the quantum computing domain as well [Rand 2016; Unruh 2019; Ying 2012; Zhou, Yu, and Ying 2019].

Hoare Type Theory (HTT) [Nanevski, Morrisett, and Birkedal 2008] encodes Hoare-style specifications as types that can then be used to enforce desired properties about the program state statically.

This encoding is made possible with the use of dependent types that can depend on values instead of just other types and which add immense expressive power to a programming language. In effect, we obtain a functional programming language in which imperative code is encapsulated within a computation type which is indexed by the type of the result of the computation and its specification in the form of pre- and postconditions. We then obtain a system where the process of type checking a program becomes equivalent to verifying that the program meets the given specification. In this work, we consider the problem of reasoning about quantum programs in HTT and present our Quantum Hoare Type Theory (QHTT).

To obtain a quantum variant of HTT, we replace the classical primitive commands in its original formulation with quantum-specific primitive commands and, taking inspiration from a novel approach by Unruh [2019], employ ghost variables for reasoning locally about the quantum state in our quantum proposition logic. Further, the use of ghost variables opens the door to an expressive language of propositions in which we can specify predicates about the equality, superposition, and entanglement of qubits.

Our approach takes inspiration from previous foundational work in reasoning about quantum programs such as “Quantum Weakest Preconditions” [D’hondt and Panangaden 2006] and “Floyd–Hoare Logic for Quantum Programs” [Ying 2012] in the spirit of Hoare [1969], Dijkstra [1976], and Gries [1981], but attempts to bring those reasoning techniques into a type system. The hope is that programmers will be able to encode some semantic properties that they expect of their programs as specifications in their code, and that type checking will ensure the correctness of their implementation to those specifications. In the classical setting, Hoare Type Theory accomplishes exactly this goal. The work presented here is an attempt to merge these ideas for the emerging landscape of quantum computing.

We limit our presentation to quantum programs that do not involve iteration. This is not to suggest that iteration is not supported in the system, only that we are yet to study it.

Our investigation suggests that QHTT could be a unified system for programming, specifying, and reasoning about quantum programs.

2 Background

In this section, we provide background on the core ideas from the existing literature that form the foundation for our work—Hoare logic, Hoare Type Theory, quantum computation, and quantum Hoare logic with ghost variables.

2.1 Floyd–Hoare logic

In the late 1960s, Robert Floyd and Tony Hoare came up with a set of axioms and rules of inference that can be used to prove the properties of computer programs. This approach is commonly just known as Hoare logic since Hoare was the first to apply it to text-based programs [Hoare 1969], while Floyd had applied it to flowcharts [Floyd 1967]. Here we present the axioms and rules of Hoare logic in a natural deduction style.

As noted above, this technique involves the use of *Hoare triples* of the form:

$$\{P\} c \{Q\}$$

which can be interpreted as “if the precondition P is true before the execution of the program c , then the postcondition Q will be true on its completion.”

In the classical setting, the most important command is an assignment statement. Hoare logic states the following axiom for assignment:

$$\begin{array}{c} \text{AXIOM OF ASSIGNMENT} \\ \{P[x \rightarrow e]\} x := e \{P\} \end{array}$$

where x is an assignable (variable) and e is a pure expression. The meaning of this axiom is that the precondition is obtained from the postcondition P by substituting e for all free occurrences of x in P .

Furthermore, there are rules of inference that are required to derive new theorems from one or more axioms or theorems already proved. We first look at the two *rules of consequence*. In the following notation, propositions above the line are called *premises* and the proposition below the line is called *conclusion*.

$$\begin{array}{c} \text{WEAKENING THE POSTCONDITION} \\ \frac{\{P\} c \{Q\} \quad Q \Rightarrow R}{\{P\} c \{R\}} \end{array} \qquad \begin{array}{c} \text{STRENGTHENING THE PRECONDITION} \\ \frac{P \Rightarrow Q \quad \{Q\} c \{R\}}{\{P\} c \{R\}} \end{array}$$

The first rule says that if the execution of a program c ensures the truth of postcondition Q , then it also ensures the truth of any proposition R logically implied by Q . This effectively lets us weaken the postcondition. In the other direction, we can strengthen the precondition using the second rule, which says that if Q is a precondition of program c , then so is any other proposition P that logically implies Q .

Alternatively, we can provide a single rule combining the two above:

$$\begin{array}{c} \text{RULE OF CONSEQUENCE} \\ \frac{P \Rightarrow Q \quad \{Q\} c \{R\} \quad R \Rightarrow S}{\{P\} c \{S\}} \end{array}$$

The next rule lets us compose sequential programs together:

$$\begin{array}{c} \text{RULE OF COMPOSITION} \\ \frac{\{P\} c \{Q\} \quad \{Q\} d \{R\}}{\{P\} c ; d \{R\}} \end{array}$$

where the semicolon denotes the sequential composition of programs. The rule of composition states that if the postcondition of the first program is identical to the precondition of the second program under which it satisfies its postcondition R , then the composition of the programs will satisfy R , given that the first program satisfies its precondition P .

These are the most important rules that we need to understand the formal system presented in this work. We do not present other rules such as that for iteration (which we will not need) and others that can be derived. We also note that since we do not have an iteration construct in our system, we do not run into the issue of reasoning about the termination of programs. In other words, all programs presented in our system are terminating.

There is one other crucial concept we need to understand before moving to the next section. Let us motivate it with the help of an example that will also serve the purpose of demonstrating how to use Hoare logic. Suppose we are given the following simple triple:

$$\{x = 1\} x := 5 \{x > 0\}$$

This obviously seems correct. But how can we mechanically use the formal system just described to show this triple is valid? The missing piece is the idea of the *strongest postcondition*. Given the assignment

statement, we can compute the strongest possible postcondition to be $x = 5$, which implies $x > 0$, that is, the given postcondition. Now, we can use the first rule of consequence (weakening the postcondition) to derive the given triple.

In Hoare Type Theory that we describe next and its quantum variant later, inferring the strongest postcondition for each statement will be the most important step during the process of type checking.

We note that the style of reasoning using the strongest postconditions that we described above is also called the forward reasoning style. There is a dual theory of *weakest preconditions* that are computed backward starting from a postcondition. We prefer the forward style in this work as it is easier to compute the strongest postconditions automatically using symbolic execution.

2.2 Hoare Type Theory

A type theory is a formal system of term rewriting in which each term has a type, notated as $t : A$ for a term t and a type A . A type theory that admits dependent types can be called a dependent type theory. Dependent type theories can be viewed as logical systems as popularized by the Curry–Howard correspondence (also known as the principle of propositions-as-types or proofs-as-programs). Under this interpretation, logical quantifiers such as \forall (universal) and \exists (existential) correspond respectively to the dependent function (Π) and pair (Σ) types in type theory, for example. It turns out that with the expressive power that comes with dependent types, it becomes possible to encode the behavioral specifications of programs into their types. For a given function, the output type then corresponds to a theorem with the input types serving as its hypotheses. Proving a theorem corresponds to writing a program of that type, and checking the proof translates to type checking.

Hoare Type Theory or HTT [Nanevski, Morrisett, and Birkedal 2008] is a specific dependent type theory that introduces a type constructor called a *Hoare type* that both isolates effectful programs from the remaining functional core language and lets us specify their behavior. In its most basic form, a Hoare type is written as:

$$\Delta. \{P\} x : A \{Q\}$$

where $x : A$ is the return value and its type and P, Q are pre- and postconditions, respectively, thus incorporating the specification methodology of Hoare logic into types. An optional context, Δ , for declaring *ghost* (or *logical*) variables may be included in the Hoare type for specifying certain variables that can be used only in the pre- and postconditions but not in the implementation of the program. Hoare types can also be seen as *monads* [Moggi 1989] that are common in functional programming and are similarly used to isolate effects. Like other monads, such a *Hoare monad* also provides a type constructor, a return operation, and a bind operation that satisfy the usual monad laws. The main difference is that the Hoare monad is not only indexed by the type of the computation but also with the pre- and postconditions that must be satisfied by the given computation. A fairly accessible introduction to Hoare Type Theory is available in the lecture notes by Perconti [2012].

How do we verify the correctness of programs in HTT? A user may provide arbitrary pre- and postconditions as the type for their program. In HTT, the process of verifying that the given triple is valid is performed during type checking. There are axioms and typing rules in HTT that correspond to the axioms for primitive commands and inference rules such as those for weakening the postcondition² and composition in Hoare logic. Each primitive effectful command in HTT requires a rule for inferring its strongest postconditions. When a user-specified postcondition needs to be checked for a given command,

²Often just called the rule of consequence in the context of HTT.

we need to show that it can be obtained by weakening the inferred strongest postcondition. The rule of composition is essentially driven by the monadic bind operation. The type-checking process is carefully divided into two phases. The first phase performs basic type checking and generates verification conditions by computing the strongest postconditions at each step of the program. In the second phase, the verification conditions need to be proved. The first phase is decidable, but the second phase may not be, hence for the second phase, either the verification conditions are deferred to an automated theorem prover (such as Z3 [de Moura and Bjørner 2008]) or need to be validated manually using an interactive theorem prover (such as Coq [Coq Development Team 2020]).

Let us consider the example we saw in the previous section written in a typical syntax for HTT:

```
x := 5 : {x = 1} r:unit {x > 0}
```

where r is the variable that holds the result of the assignment statement. Since an assignment returns nothing, we see the type of output is *unit*, and the pre- and postconditions are the same as before.

We use an alternative syntax inspired by the F* programming language [Swamy et al. 2016] in our presentation. F* is an evolution of HTT in that it supports multiple monadic effects and takes a hybrid approach between automated and interactive theorem proving. In F* syntax, we can write the above triple as:

```
x := 5 : ST (r:unit) (requires {x = 1}) (ensures {x > 0})
```

where ST is the state monad that encapsulates computations that modify the state (here assignment), and it is indexed by the result of the computation and its two pre- and postconditions. The keywords *requires* and *ensures* provide extra clarity for people unfamiliar with Hoare logic, and explicitly naming the computational effect monad, ST , helps ensure multiple effects can be supported simultaneously. Inspired by this notation, we christen our quantum state monad as QST .

2.3 Quantum computation

It is well known that quantum computation can be described in terms of elementary linear algebra. We assume familiarity with basic notions such as vectors, matrices, vector spaces, inner products, bases, linear independence, etc.

Quantum computation is usually described over a specific vector space called the *Hilbert space*, \mathcal{H} , which for our purposes can be thought of as an m -dimensional complex vector space (\mathbb{C}^m) equipped with an inner product. In this work, we often use the term “state space” to refer to a finite-dimensional Hilbert space.

The basic unit of information in quantum computation is a qubit that can be represented by a unit vector in the two-dimensional complex vector space, \mathbb{C}^2 . Using Dirac’s bra-ket notation, we can describe the state of an arbitrary qubit, $|\psi\rangle$, as a linear combination (or *superposition*) of the computational basis vectors:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$$

where

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

and the two constants α, β are complex amplitudes that obey the *normalization constraint*, $|\alpha|^2 + |\beta|^2 = 1$. This means that when measured in the computational basis, the probability of the qubit being in the

classical state 0 is $|\alpha|^2$ and that for state 1 is $|\beta|^2$. Measurement is how we extract information from a quantum system, but its side effect is that it collapses the quantum state. For example, if the classical bit obtained was 1, the posterior state becomes $|1\rangle$, and we lose the superposition state. To work with multiple qubits, we need to take their tensor product, denoted by \otimes .³ The dimension, m , of the resulting state vector is exponential to the total number of qubits, n , i.e., $m = 2^n$. For example, if we take the tensor product of the basis states above, we get a four-dimensional state vector:

$$|0\rangle \otimes |1\rangle = |0\rangle |1\rangle = |01\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ 0 \begin{bmatrix} 0 \\ 1 \end{bmatrix} \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

Computation is performed by the application of unitary quantum gates, such as the NOT gate (**X**), the Hadamard (**H**) gate, or the controlled-NOT (**CX**) gate, that can be represented as square matrices. Mathematically, a gate application is equivalent to multiplying the matrix representing the gate with the state vector of the qubit(s). [Table 1](#) shows some commonly used gates along with their circuit notation and matrix representations. For example:

$$X|0\rangle = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

This basic unit of information is not the only reason why quantum computing is fundamentally different from classical computing—certain quantum mechanical phenomena are highly counterintuitive! One such phenomenon is *quantum entanglement*, which, at its most basic level, involves a correlation between two qubits such that when one is measured, the outcome obtained necessarily influences the measurement outcome of the other qubit (even if the two qubits are far apart). For example, a simple but well-known fully entangled quantum state over two qubits is the first Bell state, which is often written as $|\beta_{00}\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$, where the first 0 in $|00\rangle$ corresponds to the first qubit and the second 0 to the second qubit and similarly for $|11\rangle$. In this state, the probability of obtaining either 00 or 11 for the pair of qubits on measurement is 1/2 each. That is, if we were to measure only the first qubit and obtain 0, it is guaranteed that on measurement, we will find the other qubit to also be in the state 0. When a qubit is not entangled, it is said to be in a *separable* state.

We refer the reader to the standard textbook by [Nielsen and Chuang \[2010\]](#) for more background. We also recommend the excellent series of essays by [Matuschak and Nielsen \[2019\]](#) for a gentle introduction and to build intuition.

2.4 Quantum Hoare logic with ghost variables

We need a predicate logic analogous to Separation Logic [\[Reynolds 2002\]](#) for quantum computation to be able to reason about only the interesting portions of the quantum state while still ensuring the correctness of non-local effects such as entanglement. Various quantum Hoare logics that have been proposed [\[Rand 2016; Ying 2012; Zhou, Yu, and Ying 2019\]](#) do not support the frame rules that provide Separation Logic its power. The closest approach is that taken by [Unruh \[2019\]](#) that allows predicates on the quantum state similar to what we would like in our setting, such as whether certain qubits are in a classical state, or

³More precisely, we are taking a Kronecker product, which is a special case of the tensor product that is used to multiply vector spaces together. The \otimes symbol can often be omitted, as shown in the example that follows.

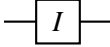
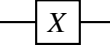
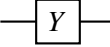
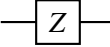
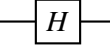
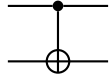
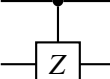
Gate	Name	Notation	Matrix
I	Identity		$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$
X	Pauli X		$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$
Y	Pauli Y		$\begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$
Z	Pauli Z		$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$
H	Hadamard		$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$
CX	controlled-NOT		$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
CZ	controlled-Z		$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$

Table 1: Common unitary gates in quantum computation

uniform superposition, entangled or separable.⁴ In this section, we provide an introduction to Unruh’s “Quantum Hoare Logic with Ghost Variables” that we use in our predicate language.

Let us make precise some mathematical definitions that we need in this section.

Definition 2.1 (Vector Subspace). A subspace of a vector space is a nonempty subset of its vectors that is closed under the operations of vector addition and scalar multiplication.

Remark. All vector spaces are equipped with at least two (*trivial*) subspaces: the singleton set, $\{0\}$, consisting of the zero vector and the vector space itself.

Definition 2.2 (Span of state vectors). For any subset P in a Hilbert space \mathcal{H} , written $(P \subseteq \mathcal{H})$, $\text{span}\{P\}$ is the smallest (closed) subspace of \mathcal{H} containing P . In other words, $\text{span}\{P\}$ is the set of linear combinations of vectors in P .

A core idea behind various quantum logics that comes from [Birkhoff and von Neumann \[1936\]](#) is that closed subspaces of a state space can be thought of as logical propositions with the inclusion relationship between them forming a partial order. Unruh’s quantum Hoare logic defines syntactic sugar for propositions and predicates on subspaces that we will find useful and that we describe next:

- \top is the proposition corresponding to the complete state space, \mathcal{H} , under consideration. It is logically equivalent to the always true proposition.
- \perp is the proposition corresponding to the zero-dimension subspace, $\{0\}$, and is logically equivalent to the absurd proposition.

⁴Thanks to Brad Lackey for pointing out that Unruh’s use of auxiliary systems to express properties of states is equivalent to the notion of *purification* [[Nielsen and Chuang 2010](#), p. 110] in quantum information theory.

- $P \wedge Q$ is the set intersection of the two subspaces: $P \cap Q$. This is equivalent to the logical *and* connective.
- $P \vee Q$ is the sum (linear combination) of the subspaces P and Q . In other words, it is the span of the union of the two subspaces: $\text{span}\{P \cup Q\}$. This is equivalent to the logical *or* connective.
- $U \cdot P$ corresponds to the application of the unitary operator U to a given subspace P . It is sugar for the subspace $\{U\psi : \psi \in P\}$.
- $X \in_q P$ lets us specify that the qubits referenced in the list of variables X lie in the subspace P (of a dimension same as or larger than 2^n , where n is the number of qubits in X). In other words, if V is the set of all quantum variables, then this predicate corresponds to the complete quantum state being $P \otimes \mathcal{H}[V \setminus X]$ where the brackets specify the domain of qubits that we are referring to and the backslash is the set difference operator.
- $X \equiv_q Y$, called the quantum equality between two disjoint lists of variables X and Y (of the same dimension), states that X and Y are quantumly equal if and only if swapping them in the complete quantum state will not change the state.
- $X \equiv_{cl} Y$, called the classical equality between two disjoint lists of variables X and Y (of the same dimension), states that X and Y are classically equal if and only if measuring them in the computational basis will always give the same outcome.
- $X =_q \psi$ specifies that variables X are in the specific state ψ . This is a sugar for $X \in_q \text{span}\{\psi\}$.

The last predicate, for example, lets us write $(a, b) =_q (|00\rangle + |11\rangle)/\sqrt{2}$ to specify that quantum variables (or qubits) a, b are in the first Bell state. It is easy to see that the span of the given state vector, $\text{span}\{(|00\rangle + |11\rangle)/\sqrt{2}\}$, only includes the set of all of its scalar multiples. Since state vectors need to obey the normalization constraint on their amplitudes, the scalar multiple can only be a complex number of the form $e^{i\theta}$ whose absolute value is always 1.⁵ Thus, the span of a single vector is just that vector, and hence, we can express that certain quantum variables correspond to a specific state vector using the notion of subspaces as propositions.

Unruh's logic further utilizes ghost variables (also known as logical variables) that can only appear in propositions to make it easy to state properties about the local state when a variable corresponding to that state may not be in scope. For example, it becomes possible to state that a single qubit input to a program is entangled with another qubit, even though we may never see that qubit locally. This is quite useful when proving the quantum teleportation protocol that we will see later.

Ghost variables in Unruh's set-up are identified by their mnemonic names. He uses g & G for a ghost variable and a list of ghost variables respectively, and similarly e & E for entangled ghosts and u & U for unentangled ghosts. For example, if we would like to use the qubit b from the above qubit pair individually, we can specify that it is in the first Bell state as $(e, b) =_q (|00\rangle + |11\rangle)/\sqrt{2}$, where e is a fresh entangled ghost variable. Intuitively, a ghost variable can assume whatever state is needed to make a given proposition true.

Unruh further defines some syntactic sugar for predicates that include ghosts:

- **uniform**(X), which states that X is uniformly distributed. It is sugar for $(X, e) =_q \sum_i \frac{1}{\sqrt{m}} (|i\rangle \otimes |i\rangle)$ where e is a fresh entangled ghost and i ranges from bitvectors 0 to $m - 1$. Note that it is not possible to express this without using ghosts. Consider, for example, a single qubit in \mathbb{C}^2 in a uniform distribution of its basis states. The only subspace that includes both $|0\rangle$ and $|1\rangle$ is \top .

⁵This scalar, called the *global phase factor*, can be ignored in most cases as it does not affect the computation.

- **separable**(X), which states that X is separable, that is, not entangled with any other system. This is sugar for $X \equiv_q u$ where u is a fresh unentangled ghost.
- **class**(X), which states that X holds a classical state. This is sugar for $X \equiv_{cl} u$ where u is a fresh unentangled ghost.

Unruh demonstrates the usefulness of these predicates over a simple imperative quantum language that includes primitives for initialization, unitary application, and measurement. We take inspiration from his logic and inference rules while designing the specification language of QHTT that we describe in the next section.

For reference, his rules for the unitary application and initialization commands are reproduced here:

$$\begin{array}{ll}
 \text{UNITARY} & \text{INITIALIZATION} \\
 \{P\} \text{ apply } U \text{ to } X \{(U \text{ on } X) \cdot P\} & \{P\} \text{ init } x \{P[x \rightarrow e], x =_q |0\rangle\}
 \end{array}$$

The unitary application rule simply produces the strongest postcondition from the precondition P by applying the given operator to the chosen qubit(s), X , in the subspace P . In $U \text{ on } X$, U is appropriately padded with the identity operation so that the operation $U \otimes I_{[V \setminus X]}$ is applied to the quantum state corresponding to all quantum variables, V . The initialization rule substitutes a fresh entangled ghost e for the existing quantum variable x and then initializes x to $|0\rangle$. As we will see, our unitary application rule is quite similar, but our initialization command has different semantics and hence a different rule. Another distinction from Unruh’s work is that he considers an infinite-dimensional Hilbert space, while in QHTT we consider a finite number of qubits that are determined via the program syntax.

3 Quantum Hoare Type Theory

To design a Hoare Type Theory for quantum computation, we need to consider some assumptions about the quantum state and its programming model.

Each program is assumed to have access to a pool of countably infinite number of logical qubits that can be initialized on demand and are returned to the pool (discarded) on measurement. With “logical”, we mean qubits with no errors; such qubits may be realized using quantum error correction techniques. We define primitive quantum commands for qubit initialization, unitary application, and measurement next.

3.1 Primitive quantum commands

Here we show our primitive quantum commands that are inspired by similar primitives in “The Quantum IO Monad” [Altenkirch and Green 2009] along with the Hoare types that encode the specifications for these commands. The predicates that we see are the same as we saw in the previous section. Each predicate implicitly takes the current local state of the program as an argument. These specifications are given in a small footprint style.

3.1.1 Initialization

```

1  init : (b: bit) → QST (q: qbit)
2           (requires {⊤})
3           (ensures {q =q |b})

```

We can initialize one qubit at a time in either of the two computational basis states, $|0\rangle$ or $|1\rangle$. For example, the initialization command $x \leftarrow \text{init } 0$ allocates a fresh qubit from the global pool, initializes it to $|0\rangle$ and returns a reference to the qubit and binds it to x .

The precondition is true for any local state, which suggests that this command may be called in any context. The postcondition states that the modified local state includes a new binding for the newly allocated qubit.

Note that our initialization command is different from Unruh’s work discussed in the previous section. The argument to our command is the initial classical value, while in Unruh’s *init* the argument is a qubit that gets initialized to $|0\rangle$.

3.1.2 Measurement

```

1 meas : (q: qbit) → QST (x: bit)
2           {ψ: vector, e_x: qbit}
3           (requires {(q ⊗ ℋ[V\q]) =_q ψ})
4           (ensures {class(q) ∧ (e_x ⊗ ℋ[V\q]) =_q ψ})

```

We assume all measurements to be in the computational basis. The measurement command, in effect, returns the qubit to the pool after returning the result of the measurement as a bit. We record the distribution of the outcome by introducing an entangled ghost variable that replaces the measured qubit in the global state. How this helps will be clear when we work through some examples later in the document. We continue to use the state space partition notation $(q \otimes \mathcal{H}[V\backslash q])$ first used in §2.4 to refer to only the local state under consideration. Note that this notation should not be confused to mean that q is separable from the rest of the system (it may or may not be).

3.1.3 Unitary application

```

1 apply _ to _ : (g: unitary) →
2               qs: (qbit ⊗ qbit) →
3               QST (_: unit)
4               {P: prop}
5               (requires {P})
6               (ensures {(g on qs) · P})

```

For convenience, we only show the rule for two-qubit gate application, but we support both one- and two-qubit gates. Since it modifies the state in place, the unitary application command returns the trivial unit output. The pre- and postconditions are the same as Unruh’s rule for unitary application command.

These primitive quantum commands and a specification language based on the predicates proposed by Unruh [2019] make the core of what is different in our Quantum Hoare Type Theory from the original formulation of HTT.

3.2 Syntax

Now, we are in a good position to talk about the raw syntax of QHTT. The following grammar shows the available base types and type constructors in the language:

Typ, A, B, C	$::=$	Types
	unit	
	bit	
	qbit	Logical qubit
	vector	State vectors
	complex	Complex scalars
	unitary	Unitaries
	prop	Propositions
	$x : A \rightarrow B$	Dependent function
	$A \otimes B$	Pair
	$\Gamma.\{P\}x : A\{Q\}$	Quantum Hoare Type

The base type *bit* inhabits the values 0 and 1, but the type *qbit* is just a reference to a logical qubit. The state is represented using the type *vector* that may further need the type *complex* to represent arbitrary quantum states. These last two types are only included for use in the specification language and are not to be used by the programming syntax. The type *unitary* includes all basic one- and two-qubit gates (such as the Pauli gates and CX) that we include in the language. The propositions formed using the specification language are given the *prop* type. \otimes is the pair type constructor. Note that we use the F*-inspired syntax QST ($x:A$) $\{ \Gamma \}$ (**requires** $\{P\}$) (**ensures** $\{Q\}$) in examples, which is just syntactic sugar for the Quantum Hoare Type shown in the formal syntax above. The optional context, Γ , is used to bind ghost variables and propositions that may appear in the pre- and postcondition to relate the beginning and the end state, as we saw in the rules for measurement and unitary application.

Next, we show the grammar for the terms of the language that include values for the base types and introduction and elimination forms for the type constructors:

Tm, K, L, M, N	$::=$	Terms
	x	Variable
	$()$	Unit
	0	0 value of type bit
	1	1 value of type bit
	ψ	Metavariable for quantum state vector
	$\lambda x.M$	Abstraction
	KN	Application
	(M,N)	Pair
	fst K	First projection from a pair
	snd K	Second projection from a pair
	do E	Suspended computation

We note that we do not provide a concrete syntax for specifying quantum state vectors that are only used in the propositions. The term *do E* represents a suspended computation and is the introduction form for the Hoare type.

We have already seen the primitive commands:

Cmd, c	$::=$	Commands
	init M	Initialization
	apply U to M	Unitary application
	meas M	Measurement
	if M then E_1 else E_2	Conditional

Connective	Formal name	Informal name
\perp	contradiction	false
\top	true	true
\wedge	conjunction	and
\vee	disjunction	or
\Rightarrow	implication	implies

Table 2: Propositional logic connectives

We did not mention the classical control construct if-then-else earlier. It branches to E_1 or E_2 based on the truth value of the bit M .

We next show the predicates of the specification language from [Unruh \[2019\]](#).

$Prop, P, Q, R$	$::=$	Propositions (prop)
	\top	Complete state space, trivial
	\perp	$\{0\}$, trivial
	$P \wedge Q$	Intersection of subspaces. P, Q
	$P \vee Q$	Sum of subspaces
	$U \cdot P$	Operator
	$x \in_q S$	x lies in subspace S
	$M =_q \psi$	M is equal to state ψ
	uniform (x)	x is uniformly distributed
	$x_1 \equiv_q x_2$	quantum equality
	separable (x)	x is separable
	$x_1 \equiv_{cl} x_2$	classical equality
	class (x)	x is classical

Apart from the quantum predicates that we see above, we also have access to the usual connectives from propositional logic, which we summarize in [Table 2](#). Some of these connectives have the same symbols as the quantum predicates, but their meaning will always be clear from the context. Additionally, we use $x =_c v$ to denote that a classical bit x is equal to a value v .

Throughout the type-checking process, we need to maintain some environment information in the form of contexts that track variable bindings and propositions available at a point in the program:

Ctx, Γ	$::=$	Contexts
	\cdot	empty context
	$\Gamma, x : A$	context extension
	Γ, P	context extension with prop

These contexts are especially relevant when calling a quantum function as the function may require the binding of the ghost variables in its specification from the environment that it is called in. Similarly, a *merge* of environments may be needed when a function call returns to determine the most precise propositions at the place of return.

Finally, we show the computations that represent the monadic unit and three kinds of bind operations:

$Comp, E, F$	$::=$	Computations
	return M	return
	$x \leftarrow K; E$	run computation K
	$x \leftarrow c; E$	run command c
	$x = M; E$	let binding




Figure 1: Circuit to prepare the first Bell state

```

1 bell100 : unit → QST (a, b): (qbit ⊗ qbit)
2           (requires {⊤})
3           (ensures {(a,b) =q |β00})
4 bell100 = λx.do
5     a ← init 0
6     apply H to a
7     b ← init 0
8     apply CX to (a, b)
9     return (a, b)

```

Listing 1: Generating the first Bell state

The complete grammar and typing rules are included in [Appendices A](#) and [B](#) of this document for easy reference. Most of the typing rules are similar to the classical HTT.

3.3 Examples

In this section, we show some simple examples along with their specifications as types. We also show the verification steps at each point in the program.

3.3.1 Bell states

Our simplest example involves the creation of a Bell (or EPR) state, which is one of the four maximally entangled quantum states of two qubits. Specifically, we create a circuit ([Fig. 1](#)) to produce the first Bell state, which we write in the mnemonic notation as $|\beta_{00}\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$.

[Listing 1](#) shows a circuit that corresponds to a suspended quantum computation that can be composed with other circuits. It takes no input and returns two qubits as output. The specification of the program is in its type (the first three lines). Here we see our first two propositions: top, \top , which is satisfied by any state space, and the $X =_q \psi$ predicate that is satisfied by quantum variables, X , that lie in the span of the given state ψ . Intuitively, this circuit does not require any inputs and can work in any given state space, and produces two qubits that are maximally entangled in the first Bell state $|\beta_{00}\rangle$.

[Listing 2](#) shows how we can prove whether this program conforms to the given specification. The initialization command (at lines 6 and 11) allocates a new qubit, initializes it to $|0\rangle$, and binds it to the variable on the left side of \leftarrow . Each occurrence of initialization adds its strongest postcondition, as shown in the form of comments. At line 8, the Hadamard gate application on $|0\rangle$ results in the equal superposition state, $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. The strongest postcondition for a unitary application simply involves applying the operator to the subspace specified by the precondition (on lines 9 and 14). Note that we use the \Leftrightarrow symbol in these annotations for a simplification step. Finally, the return operation lifts the return value

```

1 bell100 : unit → QST (a, b): (qbit ⊗ qbit)
2           (requires {⊤})
3           (ensures {(a,b) =q |β00⟩})
4 bell100 = λx.do
5           -- ⊤
6           a ← init 0
7           -- a =q |0⟩
8           apply H to a
9           -- (H on a) · (a =q |0⟩)
10          -- ⇔ a =q |+⟩
11          b ← init 0
12          -- a =q |+⟩ ∧ b =q |0⟩
13          apply CX to (a, b)
14          -- (CX on (a, b)) · (a =q |+⟩ ∧ b =q |0⟩)
15          -- ⇔ (a,b) =q |β00⟩
16          return (a, b)

```

Listing 2: Bell state program annotated with propositions




Figure 2: Quantum teleportation circuit

to the Quantum Hoare type along with its strongest postcondition (which is trivially the same as the precondition).

3.3.2 Quantum teleportation

Our second example shows how we can compose circuits defined in QHTT together and still maintain correctness. Figure 2 shows the circuit for quantum teleportation [Bennett et al. 1993] and Listing 3 shows the corresponding code, which uses the Bell pair generation circuit we defined above. This example also demonstrates how we can bind ghost variables in the Hoare type: $\{\psi : \text{vector}\}$ at line 2 states that ψ is a ghost variable of *vector* type. Recall that ghost variables can only be used in logical propositions and are not part of the program. We place the context declaring the ghost variables right after the output type to signal that the ghost variables are only bound in the *requires* and *ensures* constructors that follow next in the Hoare type. We see another predicate `class(q)` here that states that the given qubit *q* is in a classical state, which is what we expect during the quantum teleportation protocol. Or there would be no way to tell from the specification whether the implementation just returned the input qubit as the output.

Verification steps for this program are shown in Listing 4. In line 8, we see the call to the `bell100` function. Since the precondition of the function is trivial, there is nothing unusual, but after the function


```

1 teleport : q: qbit → QST b: qbit
2           {ψ: vector}
3           (requires {q =q ψ})
4           (ensures {class(q) ∧ b =q ψ})
5 teleport = λq.do
6   (a, b) ← bell100 ()
7   apply CX to (q, a)
8   apply H to q
9   x ← meas q
10  y ← meas a
11  apply CX to (y, b)
12  apply CZ to (x, b)
13  return b

```

Listing 3: Quantum teleportation

call, the calling environment is merged with the postcondition of the function. We see the expected postcondition on the newly bound qubit pair, (a, b) , just before the simplification step. In line 19, something more interesting happens! When we measure a quantum variable (in the default computational basis), its state is collapsed to either the classical bit 0 or 1 based on the distribution of either outcome. In our logic, however, both the outcome and the distribution are still maintained. This is achieved by replacing the measured qubit (which is implicitly discarded in the measurement operation) in the state space under consideration with an entangled ghost variable. Furthermore, we maintain the correspondence between the classical bit that stores the measurement result and the ghost variable that replaced its corresponding qubit. So at line 21, e_x corresponds to the classical bit x . Similarly, in lines 24–26, e_y is introduced corresponding to y . This helps us in the next two steps of the program where operations controlled on the previous measurement outcomes are applied. We recover the original state ψ at the end. Note that we use the comma symbol interchangeably with the conjunction symbol in the propositions.

This version of teleportation may not seem very interesting because it is essentially equivalent to teleportation performed without measurements (using the principle of deferred measurements). Hence, we turn to a more realistic variant in our next example.

3.3.3 Modular teleportation with classical control

As the teleportation protocol is usually stated, we have two agents Alice and Bob, who share an entangled Bell pair at the beginning of the protocol. Alice wants to communicate a message encoded in a qubit to Bob. Alice performs a Bell measurement on her two qubits, the one that encodes the message and the other which is part of the entangled Bell pair. Then she sends the result of this measurement—the two bits—to Bob. Bob receives those bits and performs correction steps on his qubit to obtain the message. We can naturally split the teleportation circuit into three parts: Bell state preparation, Alice’s portion of the circuit, and Bob’s portion of the circuit.

Listings 5–7 show the implementation of the latter two of these parts (we have already seen Bell state preparation) and the final `teleport` circuit (that composes the three parts together) in the form of separate modules (functions) along with their specifications as types. Alice’s circuit requires that her second input qubit be entangled in the Bell state. This precondition shows our first example of using Unruh’s entangled ghost variable, e , in our specification, which is declared in the context right before the `requires` construct along with other ghost variables. It further ensures that her two input qubits are both consumed in the

```

1 teleport : q: qbit → QST b: qbit
2           {ψ: vector}
3           (requires {q =q ψ})
4           (ensures {class(q) ∧ b =q ψ})
5 teleport = λq.do
6           -- q =q ψ
7           -- ⇔ q =q α|0⟩+β|1⟩
8           (a, b) ← bell100 ()
9           -- q =q ψ ∧ (a, b) =q |β00⟩
10          -- ⇔ (q, a, b) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|00⟩+|11⟩))
11          apply CX to (q, a)
12          -- (CX on (q, a)) · (q =q ψ ∧ (a, b) =q |β00⟩)
13          -- ⇔ (q, a, b) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|10⟩+|01⟩))
14          apply H to q
15          -- (H on q) · ((q, a, b) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|10⟩+|01⟩)))
16          -- ⇔ (q, a, b) =q 1/2(α(|0⟩+|1⟩)(|00⟩+|11⟩) + β(|0⟩-|1⟩)(|10⟩+|01⟩))
17          -- ⇔ (q, a, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
18          --           + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
19          x ← meas q
20          -- class(q) ∧
21          -- (ex, a, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
22          --           + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
23          y ← meas a
24          -- class(q) ∧ class(a) ∧
25          -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
26          --           + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
27          apply CX to (y, b)
28          -- class(q) ∧ class(a) ∧
29          -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|0⟩+β|1⟩)
30          --           + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|0⟩-β|1⟩))
31          apply CZ to (x, b)
32          -- class(q) ∧ class(a) ∧
33          -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|0⟩+β|1⟩)
34          --           + |10⟩(α|0⟩+β|1⟩) + |11⟩(α|0⟩+β|1⟩))
35          -- ⇔ (ex, ey) =q 1/2(|00⟩+|01⟩+|10⟩+|11⟩) ∧ b =q α|0⟩+β|1⟩
36          return b
37          -- class(q) ∧ b =q α|0⟩+β|1⟩
38          -- ⇔ class(q) ∧ b =q ψ

```

Listing 4: Annotated quantum teleportation program

```

1  alice : (q, a): (qbit ⊗ qbit) →
2      QST (x, y): (bit ⊗ bit)
3      {e: qbit, ex: qbit, ey: qbit, α: complex, β: complex}
4      (requires {q =q α|0⟩+β|1⟩ ∧ (a, e) = |β00⟩})
5      (ensures {class(q) ∧ class(a) ∧
6                  (ex, ey, e) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
7                  + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))})
8  alice = λqa.do
9      (q, a) = qa
10     apply CX to (q, a)
11     apply H to q
12     x ← meas q
13     y ← meas a
14     return (x, y)

```

Listing 5: Alice's portion of teleportation

```

1  bob : (x, y, b): (bit ⊗ bit ⊗ qbit) →
2      QST q: qbit
3      {c: complex, d: complex}
4      (requires {b =q c|0⟩+d|1⟩})
5      (ensures {y =c 0 ⇒ q =q (c|0⟩ + (-1)~x d|1⟩) ∧
6                  y =c 1 ⇒ q =q (c|1⟩ + (-1)~x d|0⟩})
7  bob = λxyb.do
8      (x, y, b) = xyb
9      if y then apply X to b else ()
10     if x then apply Z to b else ()
11     return b

```

Listing 6: Bob's portion of teleportation

```

1  teleport : q: qbit → QST b: qbit
2      {ψ: vector}
3      (requires {q =q ψ})
4      (ensures {class(q) ∧ b =q ψ})
5  teleport = λq.do
6      (a, b) ← bell100 ()
7      (x, y) ← alice (q, a)
8      b ← bob (x, y, b)
9      return b

```

Listing 7: Modular teleportation

```

1  alice : (q, a): (qbit ⊗ qbit) →
2      QST (x, y): (bit ⊗ bit)
3      {e: qbit, ex: qbit, ey: qbit, α: complex, β: complex}
4      (requires {q =q α|0⟩+β|1⟩ ∧ (a, e) = |β00⟩})
5      (ensures {class(q) ∧ class(a) ∧
6                (ex, ey, e) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
7                + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩)})})
8  alice = λqa.do
9      (q, a) = qa
10     -- q =q α|0⟩+β|1⟩ ∧ (a, e) =q |β00⟩}
11     -- ⇔ (q, a, e) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|00⟩+|11⟩))
12     apply CX to (q, a)
13     -- (CX on (q, a)) · (q =q ψ ∧ (a, e) =q |β00⟩)
14     -- ⇔ (q, a, e) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|10⟩+|01⟩))
15     apply H to q
16     -- (H on q) · ((q, a, e) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|10⟩+|01⟩)))
17     -- ⇔ (q, a, e) =q 1/2(α(|0⟩+|1⟩)(|00⟩+|11⟩) + β(|0⟩-|1⟩)(|10⟩+|01⟩))
18     -- ⇔ (q, a, e) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
19     --      + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
20     x ← meas q
21     -- class(q) ∧
22     -- (ex, a, e) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
23     --      + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
24     y ← meas a
25     -- class(q) ∧ class(a) ∧
26     -- (ex, ey, e) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
27     --      + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
28     return (x, y)

```

Listing 8: Annotated code for alice

process as they are asserted to be classical in the postcondition. Since Alice is dealing with entanglement, it is important to maintain the distribution over all qubits involved, including the ghost qubit e and the newly created ghosts e_x and e_y because of her measurement operations.

Bob's specification shows what correction is applied if his qubit is in an arbitrary state. Note that bob uses the classical control construct if-then-else in this variant of the circuit. Moreover, note that to distinguish between the input and output of the qubit, we are using two different names, but both b and q refer to the same qubit. The specification for the `teleport` function is the same as before, but the implementation is now much shorter.

We show the annotated verification steps in Listings 8–10. The verification conditions generated for `alice` are almost the same as those generated during the first part of our first teleportation example, except we do not have the qubit b in scope and instead are dealing with its ghost e .

In Bob's case, we start with his qubit b in some arbitrary quantum state and perform corrections as specified by the classical conditionals. We do not need entangled ghosts here as Bob is only dealing with two classical bits and his own qubit.

Finally, in the teleport code in Listing 10, while calling the function `alice` at line 11, the entangled ghost, e , in her type is bound to the corresponding qubit b in the calling environment as they are both part of the Bell pair along with a . After substituting b for e , we obtain the same propositions as before on return

```

1 bob : (x, y, b): (bit ⊗ bit ⊗ qbit) →
2   QST q: qbit
3   {c: complex, d: complex}
4   (requires {b =q c|0⟩+d|1⟩})
5   (ensures {y =c 0 ⇒ q =q (c|0⟩ + (-1)x d|1⟩} ∧
6           {y =c 1 ⇒ q =q (d|0⟩ + (-1)x c|1⟩}))
7 bob = λxyb.do
8   (x, y, b) = xyb
9   -- b =q c|0⟩+d|1⟩
10  if y then apply X to b else ()
11  -- y =c 0 ⇒ b =q (c|0⟩ + d|1⟩) ∧
12  -- y =c 1 ⇒ b =q (c|1⟩ + d|0⟩)
13  if x then apply Z to b else ()
14  -- y =c 0 ⇒ b =q (c|0⟩ + (-1)x d|1⟩) ∧
15  -- y =c 1 ⇒ b =q (c|1⟩ + (-1)x d|0⟩)
16  return b

```

Listing 9: Annotated code for bob

```

1 teleport : q: qbit → QST b: qbit
2           {ψ: vector}
3           (requires {q =q ψ})
4           (ensures {class(q) ∧ b =q ψ})
5 teleport = λq.do
6   -- q =q ψ
7   -- ⇔ q =q α|0⟩+β|1⟩
8   (a, b) ← bell100 ()
9   -- q =q ψ ∧ (a, b) =q |β00⟩
10  -- ⇔ (q, a, b) =q 1/√2(α|0⟩(|00⟩+|11⟩)+β|1⟩(|00⟩+|11⟩))
11  (x, y) ← alice (q, a)
12  -- class(q) ∧ class(a) ∧
13  -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
14  -- + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
15  b ← bob (x, y, b)
16  -- Merge:
17  -- class(q) ∧ class(a) ∧
18  -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|1⟩+β|0⟩)
19  -- + |10⟩(α|0⟩-β|1⟩) + |11⟩(α|1⟩-β|0⟩))
20  -- and
21  -- y =c 0 ⇒ b =q (c|0⟩ + (-1)x d|1⟩) ∧
22  -- y =c 1 ⇒ b =q (c|1⟩ + (-1)x d|0⟩)
23  -- ⇔
24  -- class(q) ∧ class(a) ∧
25  -- (ex, ey, b) =q 1/2(|00⟩(α|0⟩+β|1⟩) + |01⟩(α|0⟩+β|1⟩)
26  -- + |10⟩(α|0⟩+β|1⟩) + |11⟩(α|0⟩+β|1⟩))
27  -- ⇔ (ex, ey) =q 1/2(|00⟩+|01⟩+|10⟩+|11⟩) ∧ b =q α|0⟩+β|1⟩
28  return b
29  -- b =q α|0⟩+β|1⟩ ⇔ b =q ψ

```

Listing 10: Annotated code for modular teleportation




Figure 3: Quantum circuit for Deutsch’s algorithm

from the function. A non-trivial environment merge is needed after line 15 on return from the function bob. We have annotated this merge operation in more detail in lines 16–22 for clarity. After relating the existing ghosts e_x and e_y in the return environment with the classical bits x and y in the postcondition for Bob, we obtain the correct propositions. Simplification after the merge operation lets us recover the original message at line 27.

3.3.4 Deutsch’s algorithm

Our final example in Listing 11 demonstrates that our system can work in a higher-order setting where we can pass a circuit as input to another circuit and reason about the resulting computation. Here U_f is a quantum oracle whose implementation is unknown, but its type specifies everything we know about it—it is parameterized over a classical function, f , that takes a bit and returns a bit; it takes two qubits as input and returns two qubits as output; and finally, if the input qubits store classical bits x and y , then the output qubits store x and $y \oplus f(x)$. The circuit U_f is used as a black box in the implementation of Deutsch’s [1985] algorithm as shown in Fig. 3. The idea is to be able to reason about the algorithm’s correctness using the type of the black box alone.

Note that the second argument to the deutsch function carries its Hoare type (as specified above) along, even though we do not show it in the code at line 9 to avoid extra noise in the presentation. At line 25, the first argument, f , is passed to the oracle as a parameter without further calls to f . At lines 27–28, we show the standard case analysis for the application of the oracle to the given quantum states of a and b depending on the constant or balanced property of f . We note that this example requires some manual intervention in interpreting the conclusion of the algorithm as we are not allowed to inspect the behavior of f directly.

4 Discussion

In §3.3.2, when we presented the code for quantum teleportation, it was apparent from its type (specification) that the quantum variables q (input) and b (output) are different because of the proposition $\mathbf{class}(q)$, which suggests that q is discarded in the postcondition. It would be nicer if we could also specify information about the resource usage of the program; in teleportation, for example, an EPR pair gets consumed along with the input qubit.

A limitation of our system is that we currently do not support propositions about the distribution of measurement outcomes. For example, Listing 12 shows a program (along with its verification steps), which produces a uniformly random bit. At line 10, we can infer a weaker proposition on q , $\mathbf{uniform}(q)$, knowing its exact state. But the proposition $\mathbf{uniform}(b)$ in the postcondition of the program would be invalid as it is only defined for quantum variables and not classical bits.

```

1  U_f : (f: bit → bit) →
2      (a, b) : (qbit ⊗ qbit) →
3      QST (c, d) : (qbit ⊗ qbit)
4      {x: bit, y: bit}
5      (requires {(a, b) =q |x,y⟩})
6      (ensures {(c, d) =q |x,y ⊕ f(x)⟩})
7
8  deutsch : (f: bit → bit) →
9      ((bit → bit) → (qbit ⊗ qbit) → QST (qbit ⊗ qbit)) →
10     QST r: bit
11     (requires {⊤})
12     (ensures {(r =c 0 ⇒ f(0) = f(1)) ∧
13              (r =c 1 ⇒ f(0) ≠ f(1)) })
14 deutsch = λf.λU_f.do
15     -- ⊤
16     a ← init 0
17     -- a =q |0⟩
18     b ← init 1
19     -- a =q |0⟩ ∧ b =q |1⟩
20     apply H to a
21     -- (H on q) · (a =q |0⟩ ∧ b =q |1⟩)
22     -- ⇔ a =q |+⟩ ∧ b =q |1⟩
23     apply H to b
24     -- a =q |+⟩ ∧ b =q |-⟩
25     (a, b) ← U_f f (a, b)
26     -- (U_f on (a, b)) · (a =q |+⟩ ∧ b =q |-⟩)
27     -- ⇔ (a =q |+⟩ ∧ b =q |-⟩) -- f(0) = f(1)
28     -- ∨ (a =q |-⟩ ∧ b =q |-⟩) -- f(0) ≠ f(1)
29     apply H to a
30     -- (a =q |0⟩ ∧ b =q |-⟩) -- f(0) = f(1)
31     -- ∨ (a =q |1⟩ ∧ b =q |-⟩) -- f(0) ≠ f(1)
32     r ← meas a
33     -- class(a) ∧
34     -- (r =c 0 ∧ b =q |-⟩) -- f(0) = f(1)
35     -- ∨ (r =c 1 ∧ b =q |-⟩) -- f(0) ≠ f(1)
36     return r

```

Listing 11: Deutsch's algorithm

```

1 cointoss : unit → QST b: bit
2           (requires {⊤})
3           (ensures {uniform(b)})
4 cointoss = λx.do
5           -- ⊤
6           q ← init 0
7           -- q =q |0⟩
8           apply H to q
9           -- q =q 1/√2(|0⟩+|1⟩) ⇔ q =q |+⟩
10          -- ⇒ uniform(q)
11          x ← meas q
12          -- class(q) ∧ ex =q |+⟩
13          return x

```

Listing 12: Fair coin toss with an *invalid* postcondition

We would like to mention that the type system presented in this document is not bidirectional [Dunfield and Krishnaswami 2021] as we forego distinguishing between the introduction and elimination forms of terms while defining our syntax for simplicity. The original presentations of HTT [Nanevski, Ahmed, et al. 2007; Nanevski, Morrisett, and Birkedal 2008] (as do many other type theories) keep that distinction to be able to prove the decidability of type checking and provide a syntax-directed type system to aid implementation. We leave this for future work.

Finally, a note about scalability: there are three important aspects to scalability considerations: the expressive power of the programming language to support writing large quantum programs; compositional reasoning about programs for modularity and reuse of verified code; and lastly but importantly, the manual and computational effort required for verification. We have taken the first and second of these into account by intentionally designing the syntax of our language to be close to Quipper [Green et al. 2013], which is a proven, scalable quantum programming language. In particular, our Deutsch’s algorithm example shows that the language supports higher-order programming. We further benefit from building upon HTT that lets us seamlessly compose Hoare types together, hence permitting modular verification as we show in our quantum teleportation example. The third aspect of scalability dealing with the proof burden requires some comparison with existing work.

The full-fledged “Floyd–Hoare Logic for Quantum Programs” [Ying 2012] or QHL, which builds upon previous work on modeling quantum programs as superoperators [Selinger 2004] and quantum predicates as Hermitian operators [D’hondt and Panangaden 2006], suffers from two problems. First, it requires verbose specifications over the complete quantum state because it does not support Separation logic-style frame rules. We avoid this problem by adopting Unruh’s variant of QHL that effectively lets us get away with only specifying properties of the local state. Second, the verification process is manual, non-trivial, and does not scale to complicated algorithms such as Shor’s [1994], as a recent implementation effort using an interactive theorem prover shows [Liu et al. 2019]. A fundamental reason behind this second problem is the representation of quantum propositions as Hermitian operators. Indeed, because of this scalability limitation, there is a recent trend in research based on QHL to move away from general Hermitian operators and towards a special class of propositions based on projection operators [Li et al. 2020; Zhou, Yu, and Ying 2019], which are equivalent to our notion of closed subspaces as propositions. This prior research and our intention to use a hybrid system like F* (letting us rely more on automation) for implementing QHTT suggest that we are in the right direction. A full account of scalability considerations

of our approach requires further investigation, and along with an implementation, we leave it for future work.

5 Conclusions and perspectives

In this paper, we describe our work on Quantum Hoare Type Theory, which is a dependently typed functional programming language with support for quantum computation. We use the quantum Hoare logic by Unruh [2019] to compactly specify properties about the quantum state but do it in a type theoretical setting. We further demonstrate the usefulness of our approach using simple examples, including bell state preparation and quantum teleportation.

With future work on code extraction, our approach can become a unified system for programming, specification, and reasoning about quantum computation. There are several other avenues to explore:

Quantum Language Our programming language currently does not support any iterative or recursive constructs. We would like to add those in the future, along with the reasoning rules needed for loop invariants. Furthermore, currently, it is not possible to construct pure unitary operators directly in our language apart from those provided as constants. As Altenkirch and Green [2009] show, unitary operations form an algebraic monoidal structure, and it will be nice to be able to construct arbitrary unitaries in the language, which will greatly enhance its expressive power from current Clifford circuits to a universal language.

Mechanization There are multiple implementations of HTT in Coq such as Ynot [Nanevski, Morrisett, Shinnar, et al. 2008]. We would like to mechanize QHTT in Coq, F*, or some such dependent type theory for higher assurance of the usefulness and soundness of QHTT. This will also enable us to extract verified circuits in a lower level quantum language such as OpenQASM [Cross et al. 2017] for execution on real machines.

Linearity Peter Selinger and collaborators have recently proposed a linear dependent typed version of Proto-Quipper (dubbed Proto-Quipper-D) [Fu, Kishida, and Selinger 2020; Selinger 2020]. It is an interesting challenge to reconcile linearity in our theory based on their proposal.

Circuits as Arrows Further, Proto-Quipper treats quantum circuits as first-class citizens of the language. We would like to explore modifying our theory to treat Quantum Hoare types as arrows instead of as monads as was suggested by Vizzotto, Altenkirch, and Sabry [2006]. It makes sense from the perspective of sequential composition as arrows can have an arbitrary number of inputs/outputs as opposed to monads.

Behavioral Types Another avenue for exploration is to incorporate more precise types for qubits that can distinguish between qubits in the pure classical state vs. those in superposition vs. those in entanglement [Honda 2015; Jorrand and Perdrix 2009] such as those inspired by the various quantum resource theories or the Heisenberg representation of quantum mechanics [Rand, Sundaram, et al. 2019; 2020]. For example, in “Gottesman Types for Quantum Programs”, Rand, Sundaram, et al. [2020] assign types \mathbf{Z} and \mathbf{X} to single qubits in computational and Hadamard basis states respectively. Then, the Hadamard operator can be assigned the intersection of two arrow types, $(\mathbf{X} \rightarrow \mathbf{Z}) \cap (\mathbf{Z} \rightarrow \mathbf{X})$, as it converts between the computational and Hadamard basis states. Further, in that system, we can use the \times type

constructor to assign the type $\mathbf{X} \times \mathbf{Z}$ to the separable product state that we see for the pair of qubits (a, b) , right before applying the CX gate in Listing 1 to obtain the first Bell state. These and other such precise types can make our type system even more expressive.

Classical Effects Finally, it will be an interesting challenge to reconcile both classical and quantum effects together in a single system as many interesting quantum algorithms involve a mixture of quantum operations and classical processing. It will be important to reason about the complete behavior of an implementation (as opposed to only the quantum effects) to be able to treat it as a reusable module. We see promise in exploring this direction using the F* programming language, where there is already existing work that lets one define additional effects relatively easily [Ahman et al. 2017].

Acknowledgments

I would like to thank John Reppy and Robert Rand for their guidance. I would also like to thank the anonymous reviewers of QPL '20, and the larger QPL community for their feedback on an earlier presentation of this work [Singhal and Reppy 2020]. Casey Duckering, Anne Farrell, Sophia Lin, and Jocelyn Wilcox read the manuscript and provided helpful feedback.

This material is based upon work supported by EPiQC,⁶ an NSF Expedition in Computing, under Grant No. 1730449. Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

References

- Danel Ahman, Cătălin Hrițcu, Kenji Maillard, Guido Martínez, Gordon Plotkin, Jonathan Protzenko, Aseem Rastogi, and Nikhil Swamy (2017). “Dijkstra Monads for Free”. In: *Proc. POPL '17*. ACM, pp. 515–529. DOI: [10.1145/3009837.3009878](https://doi.org/10.1145/3009837.3009878). arXiv: [1608.06499](https://arxiv.org/abs/1608.06499) (cit. on p. 26).
- Thorsten Altenkirch and Alexander S. Green (2009). “The Quantum IO Monad”. In: *Semantic Techniques in Quantum Computation*. Ed. by Simon J. Gay and Ian Mackie. Cambridge: Cambridge University Press, pp. 173–205. DOI: [10.1017/CB09781139193313.006](https://doi.org/10.1017/CB09781139193313.006). URL: <https://www.cs.nott.ac.uk/~psztxa/g5xnsnc/chapter.pdf> (cit. on pp. 11, 25).
- 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

⁶EPiQC: Enabling Practical-scale Quantum Computation. URL: <https://epiqc.cs.uchicago.edu>

- Supremacy using a Programmable Superconducting Processor”. In: *Nature* 574.7779, pp. 505–510. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5) (cit. on p. 2).
- Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters (1993). “Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels”. In: *Phys. Rev. Lett.* 70 (13), pp. 1895–1899. DOI: [10.1103/PhysRevLett.70.1895](https://doi.org/10.1103/PhysRevLett.70.1895) (cit. on p. 16).
- Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin Vechev (2020). “Silq: A High-Level Quantum Language with Safe Uncomputation and Intuitive Semantics”. In: *Proc. PLDI '20*. ACM, pp. 286–300. DOI: [10.1145/3385412.3386007](https://doi.org/10.1145/3385412.3386007). URL: <https://files.sri.inf.ethz.ch/website/papers/pldi20-silq.pdf> (cit. on pp. 2, 3).
- Garrett Birkhoff and John von Neumann (1936). “The Logic of Quantum Mechanics”. In: *Annals of Mathematics* 37.4, pp. 823–843. DOI: [10.2307/1968621](https://doi.org/10.2307/1968621) (cit. on p. 9).
- The Coq Development Team (2020). *The Coq Proof Assistant*. Version 8.12.0. DOI: [10.5281/zenodo.4021912](https://doi.org/10.5281/zenodo.4021912) (cit. on p. 7).
- Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta (2017). *Open Quantum Assembly Language*. OpenQASM 2.0. arXiv: [1707.03429](https://arxiv.org/abs/1707.03429). URL: <https://github.com/Qiskit/openqasm/tree/OpenQASM2.x> (cit. on p. 25).
- Ellie D’hondt and Prakash Panangaden (2006). “Quantum Weakest Preconditions”. In: *Mathematical Structures in Computer Science* 16.3, pp. 429–451. DOI: [10.1017/S0960129506005251](https://doi.org/10.1017/S0960129506005251). URL: https://www.cs.mcgill.ca/~prakash/Pubs/weakest_mscs.pdf (cit. on pp. 4, 24).
- Leonardo de Moura and Nikolaj Bjørner (2008). “Z3: An Efficient SMT Solver”. In: *Tools and Algorithms for the Construction and Analysis of Systems*. Springer, pp. 337–340. DOI: [10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24) (cit. on p. 7).
- David Deutsch (1985). “Quantum theory, the Church–Turing principle and the universal quantum computer”. In: *Proc. R. Soc. Lond. A* 400.1818, pp. 97–117. DOI: [10.1098/rspa.1985.0070](https://doi.org/10.1098/rspa.1985.0070) (cit. on p. 22).
- Edsger W. Dijkstra (1976). *A Discipline of Programming*. Prentice-Hall (cit. on p. 4).
- Jana Dunfield and Neel Krishnaswami (2021). “Bidirectional Typing”. In: *ACM Comput. Surv.* 54.5. DOI: [10.1145/3450952](https://doi.org/10.1145/3450952) (cit. on p. 24).
- Ian Durham, Daniel Garisto, and Karoline Wiesner (2021). “Physicists Need to Be More Careful with How They Name Things”. In: *Scientific American*. URL: <https://www.scientificamerican.com/article/physicists-need-to-be-more-careful-with-how-they-name-things/> (cit. on p. 2).
- Robert W. Floyd (1967). “Assigning meanings to programs”. In: *Proc. Sympos. Appl. Math., Vol. XIX*. Amer. Math. Soc., Providence, R.I., pp. 19–32 (cit. on p. 4). Reprinted as “Assigning Meanings to Programs”. In: *Program Verification: Fundamental Issues in Computer Science*. Ed. by Timothy R. Colburn, James H. Fetzer, and Terry L. Rankin. Springer, 1993, pp. 65–81. DOI: [10.1007/978-94-011-1793-7_4](https://doi.org/10.1007/978-94-011-1793-7_4).
- Peng Fu, Kohei Kishida, and Peter Selinger (2020). “Linear Dependent Type Theory for Quantum Programming Languages: Extended Abstract”. In: *Proc. LICS '20*. ACM, pp. 440–453. DOI: [10.1145/3373718.3394765](https://doi.org/10.1145/3373718.3394765). arXiv: [2004.13472](https://arxiv.org/abs/2004.13472) (cit. on p. 25).
- Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron (2013). “Quipper: A Scalable Quantum Programming Language”. In: *Proc. PLDI '13*. ACM, pp. 333–342. DOI: [10.1145/2491956.2462177](https://doi.org/10.1145/2491956.2462177). arXiv: [1304.3390](https://arxiv.org/abs/1304.3390) (cit. on pp. 2, 24).
- David Gries (1981). *The Science of Programming*. Springer (cit. on p. 4).

- C. A. R. Hoare (1969). “An Axiomatic Basis for Computer Programming”. In: *Commun. ACM* 12.10, pp. 576–580. DOI: [10.1145/363235.363259](https://doi.org/10.1145/363235.363259) (cit. on p. 4).
- Kentaro Honda (2015). “Analysis of Quantum Entanglement in Quantum Programs using Stabilizer Formalism”. In: *Proc. QPL ’15*, pp. 262–272. DOI: [10.4204/eptcs.195.19](https://doi.org/10.4204/eptcs.195.19) (cit. on p. 25).
- Yipeng Huang and Margaret Martonosi (2018). “QDB: From Quantum Algorithms Towards Correct Quantum Programs”. In: *9th Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU ’18)*, 4:1–4:14. DOI: [10.4230/OASIcs.PLATEAU.2018.4](https://doi.org/10.4230/OASIcs.PLATEAU.2018.4) (cit. on p. 3).
- Yipeng Huang and Margaret Martonosi (2019). “Statistical Assertions for Validating Patterns and Finding Bugs in Quantum Programs”. In: *Proc. ISCA ’19*. ACM, pp. 541–553. DOI: [10.1145/3307650.3322213](https://doi.org/10.1145/3307650.3322213). arXiv: [1905.09721](https://arxiv.org/abs/1905.09721) (cit. on p. 3).
- Ali Javadi-Abhari, Arvin Faruque, Mohammad Javad Dousti, Lukas Svec, Oana Catu, Amlan Chakrabarti, Chen-Fu Chiang, Seth Vanderwilt, John Black, Frederic T. Chong, Margaret Martonosi, Martin Suchara, Ken Brown, Massoud Pedram, and Todd Brun (2012). *Scaffold: Quantum Programming Language*. Tech. rep. Princeton University. URL: <https://www.cs.princeton.edu/research/techreps/TR-934-12> (cit. on p. 2).
- Philippe Jorrand and Simon Perdrix (2009). “Abstract Interpretation Techniques for Quantum Computation”. In: *Semantic Techniques in Quantum Computation*. Ed. by Simon J. Gay and Ian Mackie. Cambridge: Cambridge University Press, pp. 206–234. DOI: [10.1017/CB09781139193313.007](https://doi.org/10.1017/CB09781139193313.007) (cit. on p. 25).
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer (2017). “RustBelt: Securing the Foundations of the Rust Programming Language”. In: *Proc. ACM Program. Lang.* 2.POPL. DOI: [10.1145/3158154](https://doi.org/10.1145/3158154) (cit. on p. 3).
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer (2021). “Safe Systems Programming in Rust”. In: *Commun. ACM* 64.4, pp. 144–152. DOI: [10.1145/3418295](https://doi.org/10.1145/3418295) (cit. on p. 3).
- Gushu Li, Li Zhou, Nengkun Yu, Yufei Ding, Mingsheng Ying, and Yuan Xie (2020). “Projection-Based Runtime Assertions for Testing and Debugging Quantum Programs”. In: *Proceedings of the ACM on Programming Languages* 4.OOPSLA, p. 150. DOI: [10.1145/3428218](https://doi.org/10.1145/3428218) (cit. on pp. 3, 24).
- Junyi Liu, Bohua Zhan, Shuling Wang, Shenggang Ying, Tao Liu, Yangjia Li, Mingsheng Ying, and Naijun Zhan (2019). “Formal Verification of Quantum Algorithms Using Quantum Hoare Logic”. In: *Proc. CAV ’19*. Springer, pp. 187–207. DOI: [10.1007/978-3-030-25543-5_12](https://doi.org/10.1007/978-3-030-25543-5_12) (cit. on p. 24). See also: “Quantum Hoare Logic”. In: *Archive of Formal Proofs* (2019). QHLProver-AFP, formal proof development in Isabelle. URL: <https://isa-afp.org/entries/QHLProver.html>.
- Nicholas D. Matsakis and Felix S. Klock II (2014). “The Rust Language”. In: *Proceedings of the ACM SIGAda Annual Conference on High Integrity Language Technology (HILT ’14)*. ACM, pp. 103–104. DOI: [10.1145/2663171.2663188](https://doi.org/10.1145/2663171.2663188) (cit. on p. 3).
- Andy Matuschak and Michael A. Nielsen (2019). *Quantum computing for the very curious*. URL: <https://quantum.country/qcvc> (cit. on p. 8).
- Eugenio Moggi (1989). “Computational Lambda-Calculus and Monads”. In: *Proc. LICS ’89*. IEEE Press, pp. 14–23. DOI: [10.1109/LICS.1989.39155](https://doi.org/10.1109/LICS.1989.39155). URL: <https://person.dibris.unige.it/moggi-eugenio/ftp/lics89.pdf> (cit. on p. 6).
- Aleksandar Nanevski, Amal Ahmed, Greg Morrisett, and Lars Birkedal (2007). “Abstract Predicates and Mutable ADTs in Hoare Type Theory”. In: *Programming Languages and Systems. ESOP 2007*. Springer, pp. 189–204. DOI: [10.1007/978-3-540-71316-6_14](https://doi.org/10.1007/978-3-540-71316-6_14). URL: <https://software.imdea.org/~aleks/http/esop07.pdf> (cit. on p. 24).

- Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal (2008). “Hoare Type Theory, Polymorphism and Separation”. In: *J. Funct. Program.* 18.5-6, pp. 865–911. DOI: [10.1017/S0956796808006953](https://doi.org/10.1017/S0956796808006953). URL: <https://software.imdea.org/~aleks/http/jfpsep07.pdf> (cit. on pp. 3, 6, 24).
- Aleksandar Nanevski, Greg Morrisett, Avraham Shinnar, Paul Govereau, and Lars Birkedal (2008). “Ynot: Dependent Types for Imperative Programs”. In: *Proc. ICFP '08*. ACM, pp. 229–240. DOI: [10.1145/1411204.1411237](https://doi.org/10.1145/1411204.1411237). URL: <https://software.imdea.org/~aleks/http/ynot08.pdf> (cit. on p. 25).
- Michael A. Nielsen and Isaac L. Chuang (2010). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press. DOI: [10.1017/CB09780511976667](https://doi.org/10.1017/CB09780511976667) (cit. on pp. 8, 9).
- Carmen Palacios-Berraquero, Leonie Mueck, and Divya M. Persaud (2019). “Instead of ‘supremacy’ use ‘quantum advantage’”. In: *Nature* 576.7786. DOI: [10.1038/d41586-019-03781-0](https://doi.org/10.1038/d41586-019-03781-0) (cit. on p. 2).
- Jennifer Paykin (2018). “Linear/non-Linear Types For Embedded Domain-Specific Languages”. Publicly Accessible Penn Dissertations. 2752. PhD thesis. University of Pennsylvania. URL: <https://repository.upenn.edu/edissertations/2752/> (cit. on p. 3).
- Jennifer Paykin, Robert Rand, and Steve Zdancewic (2017). “QWIRE: A Core Language for Quantum Circuits”. In: *Proc. POPL '17*. ACM, pp. 846–858. DOI: [10.1145/3009837.3009894](https://doi.org/10.1145/3009837.3009894). URL: https://jpaykin.github.io/papers/prz_qwire_2017.pdf (cit. on p. 3).
- James T. Perconti (2012). *Hoare Type Theory: Dependent Types for State*. URL: <https://www.ccs.neu.edu/home/amal/course/7480-s12/HTT-notes.pdf> (cit. on p. 6).
- John Preskill (2012). *Quantum computing and the entanglement frontier*. Rapporteur talk at the 25th Solvay Conference on Physics (“The Theory of the Quantum World”), 19-22 October 2011. arXiv: [1203.5813](https://arxiv.org/abs/1203.5813) (cit. on p. 2).
- Robert Rand (2016). *Verification Logics for Quantum Programs*. Tech. rep. Submitted as a qualifying examination (WPE-II) for the PhD program. University of Pennsylvania. arXiv: [1904.04304](https://arxiv.org/abs/1904.04304) (cit. on pp. 3, 8).
- Robert Rand, Aarthi Sundaram, Kartik Singhal, and Brad Lackey (2019). “A Type System for Quantum Resources”. Draft. URL: <https://ks.cs.uchicago.edu/publication/quantum-resource-types> (cit. on p. 25).
- Robert Rand, Aarthi Sundaram, Kartik Singhal, and Brad Lackey (2020). “Gottesman Types for Quantum Programs”. In: *Proc. QPL '20*, pp. 279–290. DOI: [10.4204/EPTCS.340.14](https://doi.org/10.4204/EPTCS.340.14). URL: <https://rand.cs.uchicago.edu/publication/rand-2020-gottesman/> (cit. on p. 25).
- John C. Reynolds (2002). “Separation Logic: A Logic for Shared Mutable Data Structures”. In: *Proc. LICS '02*. IEEE, pp. 55–74. DOI: [10.1109/LICS.2002.1029817](https://doi.org/10.1109/LICS.2002.1029817). URL: <https://www.cs.cmu.edu/~jcr/seplogic.pdf> (cit. on p. 8).
- Peter Selinger (2004). “Towards a Quantum Programming Language”. In: *Mathematical Structures in Computer Science* 14.4, pp. 527–586. DOI: [10.1017/S0960129504004256](https://doi.org/10.1017/S0960129504004256). URL: <https://www.mathstat.dal.ca/~selinger/papers/papers/qpl.pdf> (cit. on p. 24).
- Peter Selinger (2020). *Dependently Typed Quantum Programming in Proto-Quipper*. Invited talk at PPlanQC '20. URL: <https://popl20.sigplan.org/details/planqc-2020-papers/15/> (cit. on p. 25).
- Peter W. Shor (1994). “Algorithms for Quantum Computation: Discrete Logarithms and Factoring”. In: *Proc. 35th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, pp. 124–134. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700) (cit. on p. 24). Expanded version: “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In:

- SIAM J. Comput.* 26.5 (1997), pp. 1484–1509. DOI: [10.1137/s0097539795293172](https://doi.org/10.1137/s0097539795293172). arXiv: [quant-ph/9508027](https://arxiv.org/abs/quant-ph/9508027).
- Kartik Singhal and John Reppy (2020). “Quantum Hoare Type Theory: Extended Abstract”. In: *Proc. QPL '20*, pp. 291–302. DOI: [10.4204/EPTCS.340.15](https://doi.org/10.4204/EPTCS.340.15). URL: <https://ks.cs.uchicago.edu/publication/qhtt/> (cit. on p. 26).
- Krysta M. Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher E. 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: *Proceedings of the Real World Domain Specific Languages Workshop 2018*. RWDSL '18. ACM, 7:1–7:10. DOI: [10.1145/3183895.3183901](https://doi.org/10.1145/3183895.3183901). arXiv: [1803.00652](https://arxiv.org/abs/1803.00652) (cit. on p. 2).
- Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin (2016). “Dependent Types and Multi-Monadic Effects in F*”. In: *Proc. POPL '16*. ACM, pp. 256–270. DOI: [10.1145/2914770.2837655](https://doi.org/10.1145/2914770.2837655). URL: <https://www.fstar-lang.org/papers/mumon/> (cit. on p. 7).
- Dominique Unruh (2019). “Quantum Hoare Logic with Ghost Variables”. In: *Proc. LICS '19*, pp. 1–13. DOI: [10.1109/LICS.2019.8785779](https://doi.org/10.1109/LICS.2019.8785779). arXiv: [1902.00325](https://arxiv.org/abs/1902.00325) (cit. on pp. 3, 4, 8, 9, 12, 14, 25).
- Juliana Kaizer Vizzotto, Thorsten Altenkirch, and Amr Sabry (July 2006). “Structuring quantum effects: superoperators as arrows”. In: *Mathematical Structures in Computer Science* 16.3, pp. 453–468. DOI: [10.1017/S0960129506005287](https://doi.org/10.1017/S0960129506005287). arXiv: [quant-ph/0501151](https://arxiv.org/abs/quant-ph/0501151) (cit. on p. 25).
- Hui Xu, Zhuangbin Chen, Mingshen Sun, Yangfan Zhou, and Michael R. Lyu (2021). “Memory-Safety Challenge Considered Solved? An In-Depth Study with All Rust CVEs”. In: *ACM Trans. Softw. Eng. Methodol.* 31.1, 3. DOI: [10.1145/3466642](https://doi.org/10.1145/3466642). arXiv: [2003.03296](https://arxiv.org/abs/2003.03296) (cit. on p. 3).
- Mingsheng Ying (2012). “Floyd–Hoare Logic for Quantum Programs”. In: *ACM Transactions on Programming Languages and Systems* 33.6, 19:1–19:49. DOI: [10.1145/2049706.2049708](https://doi.org/10.1145/2049706.2049708). arXiv: [0906.4586](https://arxiv.org/abs/0906.4586) (cit. on pp. 3, 4, 8, 24).
- Li Zhou, Nengkun Yu, and Mingsheng Ying (2019). “An Applied Quantum Hoare Logic”. In: *Proc. PLDI '19*. ACM, pp. 1149–1162. DOI: [10.1145/3314221.3314584](https://doi.org/10.1145/3314221.3314584). URL: <https://opus.lib.uts.edu.au/bitstream/10453/140615/2/3314221.3314584.pdf> (cit. on pp. 3, 8, 24).

A Grammar

Typ, A, B, C	$::=$	Types
		unit
		bit
		qbit
		vector
		complex
		unitary
		prop
		$x : A \rightarrow B$
		$A \otimes B$
		$\Gamma. \{P\}x : A \{Q\}$
		Logical qubit
		State vectors
		Complex scalars
		Unitaries
		Propositions
		Dependent function
		Pair
		Quantum Hoare Type
$Prop, P, Q, R$	$::=$	Propositions (prop)
		\top
		\perp
		$P \wedge Q$
		$P \vee Q$
		$U \cdot P$
		$x \in_q S$
		$M =_q \psi$
		uniform (x)
		$x_1 \equiv_q x_2$
		separable (x)
		$x_1 \equiv_{cl} x_2$
		class (x)
		Complete state space, trivial
		$\{0\}$, trivial
		Intersection of subspaces. P, Q
		Sum of subspaces
		Operator
		x lies in subspace S
		M is equal to state ψ
		x is uniformly distributed
		quantum equality
		x is separable
		classical equality
		x is classical
$Term, K, L, M, N$	$::=$	Terms
		x
		$()$
		0
		1
		ψ
		$\lambda x.M$
		KN
		(M, N)
		fst K
		snd K
		do E
		Variable
		Unit
		0 value of type bit
		1 value of type bit
		Metavariable for quantum state vector
		Abstraction
		Application
		Pair
		First projection from a pair
		Second projection from a pair
		Suspended computation
Cmd, c	$::=$	Commands
		init M
		apply U to M
		meas M
		if M then E_1 else E_2
		Initialization
		Unitary application
		Measurement
		Conditional

$Unitary, U$::=	Unitaries
		I Identity
		X Pauli X
		Y Pauli Y
		Z Pauli Z
		H Hadamard
		CX CNOT
		CZ Controlled-Z
$Comp, E, F$::=	Computations
		return M return
		$x \leftarrow K; E$ run computation K
		$x \leftarrow c; E$ run command c
		$x = M; E$ let binding
Ctx, Γ	::=	Contexts
		\cdot empty context
		$\Gamma, x : A$ context extension
		Γ, P context extension with prop

B Statics

B.1 Valid types

$\Gamma \vdash A \text{ type}$

(type formation)

TY-UNIT $\frac{}{\Gamma \vdash \text{unit type}}$	TY-BIT $\frac{}{\Gamma \vdash \text{bit type}}$	TY-QBIT $\frac{}{\Gamma \vdash \text{qbit type}}$	TY-PROP $\frac{}{\Gamma \vdash \text{prop type}}$	TY-PI $\frac{\Gamma \vdash A \text{ type}}{\Gamma \vdash x : A \rightarrow B \text{ type}}$
TY-PAIR $\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash B \text{ type}}{\Gamma \vdash A \otimes B \text{ type}}$			TY-HOARE $\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash P : \text{prop}}{\Gamma, x : A \vdash Q : \text{prop}}}{\Gamma \vdash \{P\}x : A\{Q\} \text{ type}}$	

B.2 Terms

$\Gamma \vdash M : A$

(term typing)

TM-VARE $\frac{}{\Gamma, x : A \vdash x : A}$	TM-UNITI $\frac{}{\Gamma \vdash () : \text{unit}}$	TM-ZEROI $\frac{}{\Gamma \vdash 0 : \text{bit}}$	TM-ONEI $\frac{}{\Gamma \vdash 1 : \text{bit}}$	TM-PII $\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : x : A \rightarrow B}$
TM-PIE $\frac{\Gamma \vdash K : x : A \rightarrow B \quad \Gamma \vdash M : A}{\Gamma \vdash KM : B[M/x]}$	TM-PAIRI $\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash (M, N) : A \otimes B}$	TM-FSTE $\frac{}{\Gamma \vdash \text{fst } K : A}$	TM-SNDE $\frac{}{\Gamma \vdash \text{snd } K : B}$	

Most of these rules are standard, but the following rule for suspended computation, which introduces a Hoare type, requires a bit of explanation:

$$\frac{\text{TM-HOAREI} \quad \Gamma, P * \top \vdash E : (x : A. \{P \multimap Q\})}{\Gamma \vdash \mathbf{do} E : \Gamma. \{P\} x : A \{Q\}}$$

Here, the operator $*$ is analogous to a similar operator (*separating conjunction*) in Separation Logic that says that if $R * S$, then R and S hold of disjoint portions of the state. So, in the assumption of the premise, $P * \top$ is a large footprint proposition that separates the portion of the state that the computation E requires from the rest of the state using the trivial \top predicate (which signifies that the computation E knows nothing about the rest of the state). Further, the difference operator, $P \multimap Q$, says that the portion of the state satisfied by P is replaced by a portion of the state that satisfies Q . In other words, the computation E does not affect the rest of the state. This frame rule is essential to be able to provide small footprint style specifications in HTT, and hence, in our system. We leave a precise definition of the quantum analogs of separating conjunction and difference operator to future work.

B.3 Computation rules

$$\boxed{\Gamma, P \vdash E : (x : A. \{Q\})} \quad (\text{computations})$$

$$\frac{\text{COMP-CONSQ} \quad \begin{array}{l} \Gamma, P \vdash E : (x : A. \{Q\}) \\ \Gamma, x : A, Q \models R \end{array}}{\Gamma, P \vdash E : (x : A. \{R\})} \quad \frac{\text{COMP-RETURN} \quad \Gamma \vdash M : A}{\Gamma, P \vdash \mathbf{return} M : (x : A. \{P \wedge x = M\})}$$

Notice that the rule **COMP-CONSQ** requires that we show that the postcondition Q implies R (denoted by \models that denotes semantic consequence) in the appropriate context. Like the introduction rule **TM-HOAREI** for the Hoare type, its elimination rule also requires some explanation:

$$\frac{\text{COMP-HOAREE} \quad \begin{array}{l} \Gamma \vdash K : \{Q_1\} x : A \{Q_2\} \\ \Gamma, P \models Q_1 * \top \\ \Gamma, x : A, P \wedge (Q_1 \multimap Q_2) \vdash E : (y : B. \{R\}) \end{array}}{\Gamma, P \vdash x \leftarrow K; E : (y : B. \{R\})}$$

Here, K is a suspended computation (with a Hoare type) sequenced with another computation E . We need to show that the overall precondition, P , implies the large footprint proposition, $Q_1 * \top$, i.e., a portion of state satisfying Q_1 exists as required by K . The last premise says that in the expanded context, after executing K , computation E is well-typed.

B.4 Rules for primitive quantum commands

We showed the rules for primitive commands in the form of quantum Hoare types in §3.1, here we show them in a natural deduction style for completeness. The main thing to notice about these rules is that they specify the strongest postcondition that captures the most precise relationship between the initial state and the modified state after the command's execution.

$$\begin{array}{c}
\text{CMD-INIT} \\
\Gamma \vdash M : \mathbf{bit} \\
\frac{\Gamma, x : \mathbf{qbit}, P \wedge (x =_q \psi) \vdash E : (y : B.\{Q\})}{\Gamma, P \vdash x \leftarrow \mathbf{init} M; E : (y : B.\{Q\})}
\end{array}$$

where x is a fresh qubit variable and ψ is either $|0\rangle$ or $|1\rangle$ depending on the input bit.

$$\begin{array}{c}
\text{CMD-MEAS} \\
\Gamma \vdash M : \mathbf{qbit} \\
\Gamma, P \models M =_q \psi \\
\frac{\Gamma, x : \mathbf{bit}, e_x : \mathbf{qbit}, P \wedge (M =_q \psi \multimap (\mathbf{class}(M) \wedge e_x =_q \psi)) \vdash E : (y : B.\{Q\})}{\Gamma, P \vdash x \leftarrow \mathbf{meas} M; E : (y : B.\{Q\})}
\end{array}$$

where x and e_x are fresh variables and the strongest postcondition marks that the measured qubit is now in a classical state and a ghost qubit, e_x , now holds the state, ψ , previously held by the measured qubit.

Lastly, we show the rule for a single-qubit gate application:

$$\begin{array}{c}
\text{CMD-UNITARY} \\
\Gamma \vdash U : \mathbf{unitary} \\
\Gamma \vdash M : \mathbf{qbit} \\
\frac{\Gamma, x : \mathbf{unit}, (P \multimap U \cdot P) \vdash E : (y : B.\{Q\})}{\Gamma, P \vdash x \leftarrow \mathbf{apply} U \mathbf{to} M; E : (y : B.\{Q\})}
\end{array}$$

where $U \cdot P$ corresponds to the unitary application to the overall state by appropriate padding.