



HAL
open science

Formal Methods for Quantum Algorithms

Christophe Chareton, Sébastien Bardin, Dong Ho Lee, Benoît Valiron,
Renaud Vilmart, Zhaowei Xu

► **To cite this version:**

Christophe Chareton, Sébastien Bardin, Dong Ho Lee, Benoît Valiron, Renaud Vilmart, et al.. Formal Methods for Quantum Algorithms. Handbook of Formal Analysis and Verification in Cryptography, 1, CRC Press; Taylor and Francis, pp.319-422, 2023, Prospect in Pure and Applied Mathematics, 9780367546656. 10.1201/9781003090052-7 . cea-04479879

HAL Id: cea-04479879

<https://cea.hal.science/cea-04479879v1>

Submitted on 27 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Formal methods for quantum algorithms: A Survey

Christophe Chareton⁴, Sébastien Bardin⁴, Dongho Lee^{2,4}, Benoit Valiron²,
Renaud Vilmart³, and Zhaowei Xu^{1,5}

¹Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes
Formelles, 91190, Gif-sur-Yvette, France

²Université Paris-Saclay, Inria, CentraleSupélec, CNRS, ENS Paris-Saclay,
Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

³Université Paris-Saclay, Inria, CNRS, ENS Paris-Saclay, Laboratoire Méthodes
Formelles, 91190, Gif-sur-Yvette, France

⁴Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

⁵University of Tartu, Estonia

Abstract

While recent progress in quantum hardware open the door for significant speedup in cryptography as well as additional key areas (biology, chemistry, optimization, machine learning, etc), quantum algorithms are still hard to implement right, and the validation of such quantum programs is a challenge. Moreover, importing the testing and debugging practices at use in classical programming is extremely difficult in the quantum case, due to the destructive aspect of quantum measurement. As an alternative strategy, formal methods are prone to play a decisive role in the emerging field of quantum software. Recent works initiate solutions for problems occurring at every stage of the development process: high-level program design, implementation, compilation, etc. We review the induced challenges for an efficient use of formal methods in quantum computing and the current most promising research directions.

Contents

1	Introduction	5
2	General Background in Quantum Computing	6
2.1	Hybrid Computational Model	7
2.1.1	Hybrid Circuit Model	7
2.1.2	Quantum Data Registers	7
2.1.3	Separable and Entangled States	9
2.1.4	Quantum Circuits	9
2.1.5	Quantum Matrix Semantics and density operators	10
2.1.6	Other Models for Quantum Computations	13
2.2	Algorithms	13
2.2.1	Quantum Algorithms and Complexity	13
2.2.2	Quantum Algorithm Design	14
2.3	Challenges for Quantum Computation	16
2.3.1	Destructive Measurement and Non-Determinism	16
2.3.2	Quantum Noise	16
2.3.3	Efficient Compilation on Constrained Hardware	17
3	General Background on Formal Methods	17
3.1	Introduction	17
3.2	Principles	18
3.3	The Formal Method Zoo	18
4	Overview of Formal Methods in Quantum Computing	19
4.1	The Need for Formal Methods in Quantum Computing	20
4.2	Typology of Properties to Verify	20
4.2.1	Functional Specifications	20
4.2.2	Complexity specifications	21
4.2.3	Structural Constraints	21
4.2.4	Circuit Equivalence	22
5	Low-Level Verification: Compilation and Equivalence	23
5.1	ZX-Calculus and Quantomatic/PyZX	23
5.1.1	Semantical Model	23
5.1.2	Verified Properties	25
5.1.3	Algorithms and Tools	28
5.2	Path-Sum circuit Equivalence Verification	28
5.2.1	Semantical Model	28
5.2.2	Path-Sums Reduction	29
5.2.3	Verified Properties	30
5.3	Quantum abstract interpretation	30
5.4	Toward Integrated Verified Optimization: VOQC	31
5.4.1	Architecture	31
5.4.2	Optimization Procedure	32
5.5	Formally verified quantum compilation in an imperative setting: CertiQ	33

6	Formal Quantum Programming Languages	33
6.1	Quantum Programming Languages Design	33
6.1.1	Structure of Quantum Algorithms	33
6.1.2	Requirements for Quantum Programming Languages	35
6.1.3	Review of the Existing Approaches	35
6.2	Formalizing the Operational Semantics	37
6.2.1	Quantum Lambda-Calculi	37
6.2.2	Monadic Semantics	38
6.3	Type Systems	40
6.3.1	Quantum Data and Type Linearity	40
6.3.2	Example: Quantum Teleportation	41
6.3.3	Extending the Type System to Support Circuits	42
6.3.4	Dependent types and Proofs of Programs	43
6.3.5	Discussion	44
7	High and Mid-Level Verification: Algorithms and Programs	44
7.1	Quantum Hoare Logic	44
7.1.1	Quantum Programming Language: Quantum WHILE-Programs	45
7.1.2	Quantum States, Operations and Predicates	46
7.1.3	Quantum Program Verification	46
7.1.4	Implementations and Extensions	48
7.1.5	Other Quantum Hoare Logics	48
	Quantum Hoare Logic with ghost variables.	49
	Quantum Hoare Logic with classical variables.	49
	Robustness analysis	49
	Quantum separation logic	49
	Quantum relational Hoare logic	50
	Quantum Hoare type theory	50
	Quantum dynamic logic.	50
7.2	QBRICKS	51
7.2.1	Writing Quantum Circuits Functions in QBRICKS: QBRICKS-DSL	51
7.2.2	Parametrized Path-Sums	52
7.2.3	From Quantum Circuits to Path-Sums	53
7.2.4	Probabilistic Reasoning	53
7.2.5	Verified Properties	53
7.2.6	Deduction and Proof Support	54
7.2.7	Implementation and Case Studies	54
7.3	SQIR	55
7.3.1	Programming Language	55
7.3.2	Matrix Semantics and Specifications	56
7.3.3	Implementation and Case Studies	56
7.3.4	Comparison Between QBRICKS and SQIR	56
7.4	Conclusion about Formal Verification of Quantum Programs	57

8	Discussion and Bibliographical Notes	58
8.1	Deductive Verification	58
8.2	Model Checking	58
8.3	Type Checking	58
8.4	Runtime assertion Checking	59
8.5	Verification of Quantum Communication Protocols	60
9	Conclusion	60
9.1	Summary	60
9.2	Main Current Challenges	61

1 Introduction

Cryptography and quantum information. Quantum computing dates back to 1982, when Richard Feynman raised the idea [69] of simulating quantum mechanics phenomena by storing information in particles and controlling them according to the laws of quantum mechanics. In the brief history of quantum computing, the description in 1994 by Peter Shor of an algorithm [174], performing the decomposition of prime integers in polynomial time on the size of the input, plays a major role. Indeed, it was the first-ever described quantum algorithm with a practical utility—breaking the RSA public key cryptosystems in a tractable manner.

In an asymmetric cryptosystem such as RSA, information is encrypted via a key that is a solution for a given mathematical function—the decomposition of a given integer into prime factors for the case of RSA. The security of such a protocol is based on the fundamental assumption that no potential eavesdropper has the means to compute this solution efficiently. Shor’s algorithm is based on (1) a reduction of the prime factor decomposition problem into the *order-finding* problem and (2) an adequate use of quantum parallelism to perform modular exponentiation of integers over many different inputs in a single row, enabling a polynomial resolution of the *order-finding*. Thus, the computation time for performing the prime decomposition is reduced from exponential to polynomial, and therefore breaks RSA’s fundamental assumption. Shor’s original article [174] also presents a variation of the order-finding resolution algorithm, solving the discrete logarithm problem with similar performances. Doing so, it brings a procedure for breaking elliptic curve cryptosystems.

Symmetric-key cryptosystems are also challenged by quantum computing [163]. As an example, Simon’s quantum algorithm [176] brings an exponential speedup for computing the period of a function (given the promise that this period indeed exists). Several applications in public-key cryptosystems were described, providing exponential gain in, e.g., distinguishing a three-round Feistel construction [118, 163], key recovering in the Evan Mansour encryption scheme [119] and attacking the CBC-MAC message authentication scheme [163].

Finally, Grover search quantum algorithm [84] brings a quadratic speedup in the search for a distinguished element in unstructured databases. Hence, while in this case the complexity gain is less decisive than for the procedures introduced above, its potential for cryptography is significant as it weakens *any* symmetric-key encryption system.

Thus, quantum computing challenges current cryptographic uses and practices. Shor’s algorithm opened a research program for cryptographic solutions resisting the power of quantum computation, called *post-quantum cryptography* [39].

Interestingly the induced challenge also received answers from quantum information theory itself. Indeed, one of the major distinctive features of quantum information is that it cannot be read without being affected. This entails that an eavesdropper trying to access a quantum information exchange cannot help betraying her attempt. Based on this feature, one can encode a cryptographic key in a quantum message and, in case of eavesdropping, detect it *a posteriori*, renounce this particular key and try another sending. The study of *Quantum key distribution protocols* is an active research area [164, 22, 128, 134].

Quantum computing and quantum software. These cryptographic aspects are one of many applications studied in the young research field of quantum computing. Others are, e.g., machine learning [26, 166, 133], optimization [64], solving linear systems [87], etc. In all these domains there are quantum algorithms beating the best known classical algorithms by quadratic or even exponential factors, complexity-wise.

These algorithms are based on laws and phenomena specific to quantum mechanics (such as quantum superposition, entanglement, unitary operations). Therefore, implementing them requires

a framework consisting of both a *dedicated hardware* (quantum computers) and a *dedicated software* (quantum programming languages and compilation toolchains).

In the last 20 years, several such languages have been proposed, such as QISKIT [153], Liqui \rangle [193], Q# [181], Quipper [83], PROJECTQ [180], *etc.* Still, the field is in its infancy, and many questions still need to be answered before we can reach the level of maturity observed for classic programming languages. Standing questions include, for example, introducing a foundational computing model and semantics for quantum programming languages, adequate programming abstractions and type systems, or the ability to interact with severely constrained hardware in an efficient way (optimizing compilers).

Verification of quantum programs. While testing and debugging are the common verification practice in classical programming, they become extremely complicated in the quantum case. Indeed, debugging and assertion checking are *a priori* very complicated due to the destructive aspects of quantum measurement (see Section 2.3.1 below). Moreover, the probabilistic nature of quantum algorithms seriously impedes system-level quantum testing. Finally, classical emulation of quantum algorithms is (strongly believed to be) intractable.

On the other hand, nothing prevents *a priori* the *formal verification* [44] of quantum programs, i.e. *proving* by (more or less) automated reasoning methods that a given quantum program behaves as expected for any input, or at least that it is free from certain classes of bugs.

Interestingly, while formal methods were first developed for the classical case where they are still used with parsimony—mainly for safety-critical domains—as testing remains the main validation methods, their application to quantum computing could become more mainstream, due to the inherent difficulties of testing quantum programs.

Goal of this survey This survey introduces both the requirements and challenges for formal methods in quantum programs specification and verification, and the existing propositions to overcome these challenges.

The first sections give the general background. In Section 2 we introduce the main concepts at stake with quantum computing and quantum algorithms. We provide a state of the art introduction for formal methods, given in Section 3. The specific requirements for formal reasoning in the quantum case are then developed in Section 4. Then we come to concrete quantum programming and formal verification material. In Section 5 we introduce several existing solutions for the formal verification of quantum compilation and the equivalence of quantum program runs. Generating such runs requires specific programming languages. The formal interpretation of quantum languages is introduced in Section 6. Then in Section 7 we present the main existing solutions for formally verified quantum programming languages. In Section 8 we introduce references for further usage of formal methods linked with quantum information, and we conclude this survey with a discussion in Section 9.

2 General Background in Quantum Computing

By many aspects, quantum computing constitutes a new paradigm. Making great use of quantum superposition and quantum entanglement, it requires to define proper versions for such fundamental concepts as data structures at stake in computation, or the elementary logical operations at use. We introduce the well-known hybrid quantum computation model in Section 2.1.

Quantum computers are not intended to, and will not, replace classical ones. One should better see the opening of a new field, with possibilities to solve new problems. Section 2.2 presents these new problems and introduces quantum algorithms design.

As a new software technology, quantum computing comes with specific challenges and difficulties. These specificities are closely related to the particular needs for formal reasoning in quantum computing. They are introduced in Section 2.3.

2.1 Hybrid Computational Model

Let us first introduce the main concepts at stake in quantum programming. They concern the architecture of quantum computers, the structure of quantum information and quantum programs, and their formal interpretation.

2.1.1 Hybrid Circuit Model

The vast majority of quantum algorithms are described within the context of the *quantum co-processor model* [113], i.e. an hybrid model where a *classical* computer controls a *quantum* co-processor holding a quantum memory, as shown in Figure 1. In particular, the classical computer performs control operations (**if ... else** statements, loops, etc). The co-processor can apply a fixed set of elementary operations (buffered as *quantum circuits*) to update and query (*measure*) the quantum memory. Importantly, while measurement allows retrieving classical (probabilistic) information from the quantum memory, it also modifies it (*destructive effect*).




Figure 1: Scheme of the hybrid model

Major *quantum programming languages* such as Quipper [83], Liqui|>[193], Q# [181], PROJECTQ [180], SILQ [27], and the rich ecosystem of existing quantum programming frameworks [154] follow this hybrid model.

2.1.2 Quantum Data Registers

The following paragraphs introduce several definitions and notations for quantum data registers. In particular, we follow the standard *Dirac notation*. For more details about this content, we refer the reader to the standard literature [140].

Kets and basis-kets. While in classical computing the state of a bit is one between two possible states (0 or 1), in quantum computing the state of a *quantum bit* (or *qubit*) is described by *amplitudes* over the two elementary values 0 and 1 (denoted $|0\rangle$ and $|1\rangle$), i.e. linear combinations of vectors $\alpha_0|0\rangle + \alpha_1|1\rangle$ where α_0 and α_1 are any *complex values* satisfying $|\alpha_0|^2 + |\alpha_1|^2 = 1$. In a sense, amplitudes are generalization of probabilities.

More generally, quantum states are defined in complex finite-dimensional Hilbert spaces¹: the state of a *qubit register* of n qubits (called a *ket* of length n –dimension 2^n) is a column vector with 2^n rows, formed as a *superposition* of the 2^n elementary basis vectors of length n (the “basis kets”), *i.e.* a ket is any linear combination of the form

$$|u\rangle_n = \sum_{k=0}^{2^n-1} \alpha_k |k\rangle_n \quad (1)$$

such that $\sum_{k=0}^{2^n-1} |\alpha_k|^2 = 1$.

Bit-vectors and basis kets. Depending on the context, it may be more convenient to index the terms in the sum above with bit vectors instead of integers. We call *bit vector of length n* any sequence $x_0x_1\dots x_{n-1}$ of elements in $\{0, 1\}$. Along this survey, we assume the implicit casting of these values to/from booleans (with the least significant bit on the right). For any positive n , we denote the set of bit vectors of size n by BV_n . We also surcharge notation $|j\rangle_n$ shown above with bit vector inputs. Formally, for any bitvector \vec{x} of length n , $|\vec{x}\rangle_n = |\sum_{i=0}^{n-1} x_i * 2^{n-i-1}\rangle_n$. Hence, one can write state $|u\rangle_n$ from (1) as

$$|u\rangle_n = \sum_{\vec{x} \in BV_n} \alpha_{\sum_{i=0}^{n-1} x_i * 2^{n-i-1}} |\vec{x}\rangle_n$$

It may also be convenient to represent basis kets through their index’s binary writing. For example, the two qubits kets basis is equivalently given as $\{|0\rangle_2, |1\rangle_2, |2\rangle_2, |3\rangle_2\}$ or as $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$.

We omit the length index n from notation $|u\rangle_n$ when it is either obvious from the context or irrelevant. We also adopt the implicit convention of writing basis kets with either integer indexes k, i, j or bit-vector \vec{x} and general kets with indexes u, v, w . Hence, in the following $|u\rangle, |v\rangle, |i\rangle$ and $|\vec{x}\rangle$ all designate kets, the last two having the additional characteristics of being basis kets.

When considering two registers of respective size m and n , the state of the compound system lives in the *Kronecker product*² –or *tensor product*–of the original state spaces: a general state is then of the form

$$\sum_{\vec{x} \in BV_m, \vec{y} \in BV_n} \alpha_{\vec{x}, \vec{y}} |\vec{x}\rangle_m \otimes |\vec{y}\rangle_n.$$

In particular, the state of a qubit register of n qubits lives in the tensor product of n state-spaces of one single qubit.

Adjointness. In the following we also use the adjoint transformation for matrices. The adjoint of matrix M with r rows and c columns is the matrix M^\dagger , with c rows and r columns and such that for any indexes $j, k \in \llbracket 0, c \rrbracket \times \llbracket 0, r \rrbracket$ ³, cell $M^\dagger(j, k)$ holds the conjugate value $M(k, j)^*$ of $M(k, j)$ (for any complex number c , its conjugate c^* is the complex number with the same real part and the opposite imaginary part as c). The adjoint of a ket $|u\rangle_n$ is called a *bra*. It is a row vector with 2^n columns denoted $\langle u|_n$ –or simply $\langle u|$ –and with indices the conjugates of those of $|u\rangle_n$. This bra-ket notation is particularly convenient for representing operations over vectors. Given a ket $|u\rangle$ and a bra $\langle v|$,

¹In the finite-dimensional case, Hilbert spaces are vector spaces equipped with an *inner* (scalar) product.

²Given two matrices A (with r rows and c columns) and B , their Kronecker product is the matrix $A \otimes B = \begin{pmatrix} a_{11}B & \dots & a_{1c}B \\ \vdots & \ddots & \vdots \\ a_{r1}B & \dots & a_{rc}B \end{pmatrix}$. This operation is central in quantum information representation. It enjoys a number of useful algebraic properties such as associativity, bilinearity or the equality $(A \otimes B) \cdot (C \otimes D) = (A \cdot C) \otimes (B \cdot D)$, where \cdot denotes matrix multiplication.

³where, for any two integers i, j with $i < j$, $\llbracket i, j \rrbracket$ denote the induced interval, that is the set of integers k such that $i \leq k < j$.

$|u\rangle\langle v|$ denotes their Kronecker product –or *outer product*. Furthermore, if $|u\rangle$ and $\langle v|$ have the same length, then $\langle v|u\rangle$ denotes their scalar product –also called *inner product*. In particular, in the case of basis states $|i\rangle$ and $\langle j|$, $\langle i|j\rangle = 1$ if $i = j$ and 0 otherwise and $|i\rangle\langle j|$ is the square matrix of width 2^n with null coefficient everywhere except for cell (i, j) with coefficient 1. If $i = j$, then $|i\rangle\langle j|$ operates as the projector upon $|i\rangle$.

Quantum measurement and Born rule. The probabilistic law for measurement of kets is given by the so-called *Born rule*: for any $k \in \llbracket 0, 2^n \rrbracket$, measuring state $|u\rangle_n$ from Equation 1 results in k with probability $|\alpha_k|^2$. The measurement is destructive: if the result were k , the state of the register is now $|k\rangle_n$ (with amplitude 1).

2.1.3 Separable and Entangled States

From Section 2.1.2, a quantum state vector of length n is a superposition of basis elements with coefficients whose squared moduli sum to one. Then, tensoring n quantum states $|u_j\rangle_1$ of length 1 results in a state $|u\rangle_n = \bigotimes_{j \in \llbracket 0, n \rrbracket} |u_j\rangle_1$ of length n . One can decompose back $|u\rangle_n$ into the family $\{|u_j\rangle\}_{j \in \llbracket 0, n \rrbracket}$: we say that $|u\rangle_n$ is a *separable* state. Note that the structure of quantum information introduced above contains states missing the property of being separable. As an example, the state

$$|\beta_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

cannot be written as a tensor product of two single-qubit states. This phenomenon is called *entanglement*, and $|\beta_{00}\rangle$ is an *entangled state*. It induces that one can store more quantum information in n qubits altogether than separately.

Example 2.1 (Bell states). *State $|\beta_{00}\rangle$ is a construction of particular interest in quantum mechanics and quantum computing. In their famous 1935 article [62], Einstein, Podolsky and Rosen argued for the incompleteness of quantum mechanics, based on considerations upon $|\beta_{00}\rangle$. In 1964 [21], J.S. Bell proposed an experiment to test the argument. It was based on statistics over experiments on the four following states:*

$$\begin{aligned} |\beta_{00}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle_2 + |11\rangle_2) & |\beta_{01}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle_2 + |10\rangle_2) \\ |\beta_{10}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle_2 - |11\rangle_2) & |\beta_{11}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle_2 - |10\rangle_2) \end{aligned}$$

These four states are now known as the Bell states (notation β stands for the initial B) and are used in many quantum protocols, such as teleportation or superdense coding (see Section 5.1). We use them and their generation as a running example in the rest of this survey.

2.1.4 Quantum Circuits

Three kinds of operations may be applied to quantum memory, exemplified in Figure 2 with the circuit generating and measuring Bell states:

- the *initialization* phase allocates and initializes quantum registers (arrays of qubits) from classical data. In Figure 2 it is represented on the two first qubit wires as \vdash , indexed by value i_w . It creates a quantum register in one of the basis states (that is, in the case of a two-qubit register, in one of the four basis states $|00\rangle, |01\rangle, |10\rangle, |11\rangle$),

- the actual quantum computing part consists in transforming an initialized state. This is performed by applying a sequence of proper quantum operations, structured in a so-called *quantum circuit*. In Figure 2 this part is identified with a dashed box (itself sequenced with dotted boxes a. and b.),
- the extraction of useful information from a quantum computation is performed through the *measurement* operation, by which one probabilistically gets classical data from the quantum memory register. Measurement is represented, on each qubit it is applied to, as \square .
- In generalized circuits, not all wires in a register need to be initialized and measured. Hence in Figure 2 the transformations are performed over the two first wires of a wider register, and the additional wires are left untouched.




Figure 2: Generalized circuit to create and measure Bell states

Note that we reserve the term circuit for the pure quantum part (the dashed box in Figure 2). We call generalized circuit a process made of a circuit together with, possibly, initialization and measurements.

Quantum circuits are built by *combining*, either in *sequence* or in *parallel*, a given set of elementary operations called *quantum gates*. In addition to sequence and parallelism, derived circuit combinators (controls, reversion, ancillas, etc) are often used in quantum circuit design (See Figure 9 in Section 6.1.1 for details). The circuit part of Figure 2 uses two different quantum gates, drawn in dotted boxes:

- the Hadamard gate H (a.), which induces a state superposition on a given qubit,
- the *control not* gate, often written $CNOT$ (b.) and represented as \oplus . It is a binary gate, flipping the *target qubit* (in wire 2 in our case) when the *control qubit* (wire 1) has value 1.

2.1.5 Quantum Matrix Semantics and density operators

The transformation operated by a quantum circuit on a quantum register is commonly interpreted as a matrix. In this setting, the parallel combination of circuits is interpreted by the *Kronecker product* and the sequential combination by the *matrix multiplication*.

Quantum circuits. Quantum circuits happen to operate as *unitary* operators (preserving the inner product between vectors). A set of elementary gates is (pseudo-) *universal* if, by a combination of parallel and sequential composition, one can synthesize (or approximate) all *unitary* operations. Examples for elementary gates are given in Table 1, with their matrix semantics interpretation. Apart from the already encountered gates H and $CNOT$, it features two additional families of gates, $Ph(\theta)$ and $R_Z(\theta)$, where θ is an angle inducing a so-called *phase factor* $e^{i\theta}$. $Ph(\theta)$ operates a simple scalar multiplication by a phase factor, while $R_Z(\theta)$ operates as a rotation. Table 1 is given with indexes ranging over any angle θ , making the set of gates universal. Usually, we restrict it to angles of measure $\frac{\pi}{2^n}$, with n ranging over integers. This restriction makes the resulting set of gates pseudo-universal.

Table 1: Elementary gates and their matrix semantics

H	$CNOT$	$Ph(\theta)$	$R_Z(\theta)$
$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{i\theta} \end{pmatrix}$	$\begin{pmatrix} e^{-i\theta} & 0 \\ 0 & e^{i\theta} \end{pmatrix}$

Example 2.2 (Semantics for the Bell generating circuit). *Let us look at Figure 2 again. First, an Hadamard gate is applied to the first wire and nothing happens to the second wire (it stays untouched, which is represented by the identity matrix). The matrix for the first column—the dotted box indexed with a.—of Figure 2 is*

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix}$$

Then gate CNOT is applied, with matrix $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ and the sequential combination of the two sub-circuits translates, in the matrix semantics, as their usual product (mind the reverse ordering, wrt the figure):

$$\mathbf{Mat}(\text{Bell-circuit}) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \cdot \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & -1 \\ 1 & 0 & -1 & 0 \end{pmatrix}$$

Application to input initialized kets. In the matrix formalism, we interpret $|0\rangle$ as the column vector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $|1\rangle$ as $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and the concatenation $|ij\rangle$, where i and j both are sequences of 0 or 1, as the Kronecker product $|i\rangle \otimes |j\rangle$. For example, the two qubits basis kets $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ are represented, respectively, as

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

The transformation performed by a quantum circuit C upon a quantum state $|\psi\rangle$ is interpreted as the matrix product $\mathbf{Mat}(C) \cdot |\psi\rangle$ of the matrix for this circuit by the column vector for this quantum state. By notation abuse, we also simply write it $C|\psi\rangle$.

Example 2.3. *One can now directly verify that the Bell generating circuit from Figure 2 generates the Bell states from Example 2.1: for any $a, b \in \{0, 1\}$,*

$$\text{Bell-circuit} \cdot |ab\rangle = |\beta_{ab}\rangle$$

Measurement. Last, measurement is performed over an orthonormal basis of the Hilbert space. For sake of simplicity, we only consider the case of measurements in the computational basis. Hence, measuring a quantum register results in a basis state, with probabilities following the Born rule introduced in Section 2.1.2: measuring any state $|u\rangle_n$ results in basis state $|k\rangle_n$ with probability (written $\text{proba_measure}(|u\rangle_n, |k\rangle_n)$) $|\alpha_k|^2$, where α_k is the amplitude of $|k\rangle_n$ in $|u\rangle_n$. Applying this rule to the Bell state, one easily state that for any $a, b, i, j \in \{0, 1\}$,

$$\text{proba_measure}(|\beta_{ab}\rangle, |ij\rangle) = \frac{1}{2}(\text{if } b = 0 \text{ then } i \oplus j \text{ else } 1 - (i \oplus j))$$

where \oplus denotes addition modulo 2. Note, from Example 2.1, that index a in notation $|\beta_{ab}\rangle$ only accounts for a -1 factor in the second term of the state superposition. Hence, since measurement is ruled by the Born rule and since this rule ignores negation (see Section 2.1.2), then index a does not appear in the expression of $\text{proba_measure}(|\beta_{ab}\rangle, |ij\rangle)$.

Discussion over the matrix semantics. Matrix semantics is the usual standard formalism for quantum computing (see [140] for example). Still, the size of matrices grows exponentially with the width (number of qubits) of circuits, so it is often cumbersome when addressing circuits from non-trivial algorithm instances. Furthermore, algorithms usually manipulate parametrized families of circuits. The resulting parametrized families of matrices may not be conveniently writable.

Hence, a more compact interpretation for quantum circuits may be helpful. In particular, path-sum semantics [5, 4] directly interprets quantum circuits by the input/output function they induce over kets—corresponding, in matrix terms, for any circuit C of width n , to the function $|u\rangle_n \mapsto \mathbf{Mat}(C) \cdot |u\rangle_n$. To do so, it exhibits a generic form for quantum registers description, which is generated by a restricted number of parameters and composes nicely with sequence and parallel compositions. Path-sum semantics plays a growing role in formal specification and verification. It is introduced with further details in Section 5.2.

Density operators. In the preceding paragraph, we introduced measurement as a non-deterministic operation over quantum states. Another strategy consists in dealing with a notion of states featuring this non-determinism. A *mixed state* (as opposed to a *pure state*) is a probability distribution over several states. Alternatively, it can be seen as an incomplete description of a state, featuring the incomplete knowledge one may have about it. Then, measurement can be characterized as a simple transition between mixed states.

In quantum processes, this view is formalized by density operators, that extends matrices formalism with the characterization of probabilistic states. For sake of brevity, in this paragraph we give only a short introduction to the density operator formalism. Our aim here is only to provide the required definitions and notations for this review. For further detail about density operators and for the related soundness proofs, we refer the interested reader either to [168] or [139] (Section 2.4).

Basically, the density operator for a pure state $|x\rangle$ is the reflexive outer product $|x\rangle\langle x|$. Given a set S of indices and a distribution of states $\{|x_k\rangle\}_{k \in S}$, each occurring with probability p_k , we represent the overall mixed state as the density operator

$$\rho := \sum_{k \in S} p_k |x_k\rangle\langle x_k|$$

By linearity, the result of applying a unitary U to $|x\rangle\langle x|$ is given by the product $U|x\rangle\langle x|U^\dagger$. A measurement of a quantum register q of size n may be described by the collection of possible projectors it realizes, that is the set $M = \{M_k := |k\rangle_n \langle k|_n\}_{k \in \llbracket 0, 2^n - 1 \rrbracket}$. In the density operators formalism,

the action of M over a state ρ may result in any state $M_k \rho M_k^\dagger$, with probability $\text{tr}(M_k M_k^\dagger \rho)$, where the trace $\text{tr}(M)$ of a square matrix M with n rows and columns is defined as the sum $\sum_{j \in \llbracket 0, n \rrbracket} M(j, j)$ of its diagonal cell values. Then, the overall action of a measurement M over a density operator ρ results in $\rho' = \sum_{k \in \llbracket 0, 2^n - 1 \rrbracket} M_k \rho M_k^\dagger$.

Now, measurement description generalizes to the case of *partial measurements*, where only a sub-register is measured. Let us consider the case of a quantum register q of size $n = n_1 + n_2$. We write q_1 and q_2 for the concatenated sub-registers and H, H_1, H_2 for the respectively induced Hilbert spaces. To simplify the notations we consider the case of measuring the first n qubits. For any density operator ρ , if it is separable as $\rho = |x_1\rangle_{n_1} \langle y_1|_{n_1} \otimes |x_2\rangle_{n_2} \langle y_2|_{n_2}$, then the partial trace of ρ over H_2 is defined as $\text{tr}_2(\rho) = \langle y_1 | x_1 \rangle_{x_2} \langle y_2 |$ and the definition generalizes by linearity to any density operator ρ . Then, $\text{tr}_2(\rho)$ equivalently represents the result of:

- (1) measuring register q_2 from the mixed state ρ and (2) forgetting the measured qubits while conserving memory of the unmeasured subregister q_1 state;
- or just forgetting about (*discarding*) register q_2 in the description of ρ . Then, $\text{tr}_2(\rho)$ is the description of the sub-system held by q_1 . We call $\text{tr}_2(\rho)$ a *partial density operator* on H_1 .

2.1.6 Other Models for Quantum Computations

Many alternatives are currently explored for physical implementations of quantum computing machines and worth mentioning. Some of them (such as Measurement Based Quantum Computing [157, 32], topological quantum computations [76], linear optical networks [1], adiabatic quantum computing [65], *etc.*) differing on rather fundamental aspects (like, e.g., the elementary operations constituting computations). Nevertheless, currently, formal methods developments mainly address the standard circuit model introduced above.

The ZX-Calculus [47] also provides an alternative graphical formalism to reason about quantum processes. Basically, in this setting, quantum operations are represented by diagrams and their composition through sequence or parallelism corresponds to graphical compositions in the calculus. This language comes with a series of enabled transformations over graphs, preserving computational equivalence. ZX-Calculus is presented in Section 5.1.

2.2 Algorithms

As previously introduced, quantum computers are meant to perform calculations that classical computers are *a priori* not able to perform in a *reasonable* time. We give the formal complexity theory characterization for this point in Section 2.2.1, then Section 2.2.2 discusses the usual conventions for quantum algorithms descriptions.

2.2.1 Quantum Algorithms and Complexity

It is commonly assumed that formal problems are *tractable* by a computer if there exists an algorithm to solve this problem in time (measured by the number of elementary operations it requires) that is bounded by a polynomial over the size of the input parameters. Formal problems satisfying this criterion for classical computers form a *complexity class* usually referred to as \mathbf{P} . It is schematically represented in Figure 3.

As introduced in Section 2.1.2, extracting useful information from a quantum register requires a measurement, ruled by the Born law. Therefore, a quantum computation is an alternation of non-deterministic (measurement) and deterministic (circuit unitary application, classical post-treatment,

etc) operations. Since such computations are probabilistic, the tractability criterion from above needs to be slightly adapted. Instead of considering problems for which a polynomial algorithm brings a solution with certainty, we consider those for which a polynomial algorithm brings a solution with an error probability of at most $\frac{1}{3}$. The corresponding class of problems for quantum computers is called *bounded error quantum polynomial time* (**BQP**).

In addition to **P** and **BQP**, Figure 3 represents the *non-deterministic polynomial time* class **NP**. It gathers formal problems \mathcal{P} for which there is an algorithm that, given a candidate solution, checks whether this candidate is an actual solution for \mathcal{P} in polynomial time. It is trivial that **P** is included in **NP** and it is also proved that $\mathbf{P} \subseteq \mathbf{BQP}$. There are good reasons to believe that these inclusions are strict. Nevertheless, strictness is not formally proved and there are a variety of problems that belong to **NP** without a known tractable resolution algorithm.




Figure 3: Complexity classes P, NP, BQP

Hence, quantum algorithm performance is not to be evaluated against the best *possible* performance of any classical computation (which depends on whether $\mathbf{BQP} = \mathbf{P}$) but, more pragmatically, against the best classical *known* equivalent.

Then, quantum computing is relevant for problems that are polynomially solvable by a quantum computer (with a given probability of success) but *intractable* by a classical one. They appear in the gray zone in Figure 3. The question whether the dark gray part ($\mathbf{BQP} \setminus \mathbf{NP}$) is empty or not depends on whether $\mathbf{BQP} \subseteq \mathbf{NP}$, which is unknown, but several **BQP**-complete problems have been described through the literature [200]. These are neither easily computable nor verifiable with known classical means, but may be computed with quantum means: since quantum computers output a right solution for them with probability $> \frac{2}{3}$, after several runs one can select the best represented output as the sought solution.

Let us stress that the correspondence between polynomial solvability and tractability is not strict. For a problem, not belonging to a polynomial class bounds the size of input parameters concretely tractable by a computer, but does not absolutely forbid computation for any instance of it. Hence, for some problems, the quantum advantage does not consist in providing a polynomial resolution, but in reducing computation time to extend the set of tractable inputs. A typical example is the Grover search algorithm [84], searching for a distinguished element in an unstructured data set, providing a quadratic acceleration against classical procedures.

2.2.2 Quantum Algorithm Design

Before introducing the challenges at stake with implementations of quantum algorithms and their formal solutions, we make a few observations on the usual format used to describe quantum algorithms, based on an example: Figure 4 reproduces the core quantum part for Shor's algorithm [140, p. 232]—certainly the most emblematic of all quantum algorithms.

The following observations generally hold for other quantum algorithms in the literature. We give them together with illustrations (within parenthesis) from the example of Figure 4.

1. The algorithm is structured in two main parts: a specification preamble and a **Procedure** description. The preamble indicates the minimal specification an implementation should satisfy. Note that it figures as an actual part of the algorithm itself. It contains three types of entries:
 - a description of the **Inputs**, giving a signature for the parameters (a black-box circuit U , integers x, N, L and two quantum registers of sizes t and L) and some preconditions for these elements (e.g. x co-prime to L , L being N bits long, *etc*);
 - a description of the **Outputs** of the algorithm. It contains, again, a signature (an integer) and a success condition for these **Outputs** (to be equal to the sought modular order);
 - a **Runtime** specification, containing: (1) a probability of success for each run of the **Procedure** ($O(1)$), (2) resources specifications. In the example, the latter consists in bounding the number of required elementary operations. Further metrics are also often used (the maximal *width* of required quantum circuits—the number of qubits a circuits requires, the maximal depth of a circuit—the maximal number of operations performed on a given qubit, *etc*).
2. The **Procedure** itself consists of a sequence of declared operations, interspersed with formal descriptions of the state of the system along with the performance of these operations. (in Figure 4 these elements are given in parallel, declarations of operations constitute the right-hand side and intermediate formal assertions are on the left-hand side). These assertions serve as specifications for the declared operations. For instance, operation “create superposition” has precondition the formal expression of Line 1, left (framed in blue) and postcondition the one of Line 2, left (framed in red). They serve as arguments to convince the reader that the algorithm **Outputs** conditions are met at the end of the **Procedure** (notice that the ultimate such postcondition—the measured state being r — corresponds to the success condition for the overall algorithm). But we can also interpret them as *contracts* for the programmer, committing her to implement each function in any way provided that whenever its inputs satisfy the preconditions, then its outputs satisfy the postconditions.
3. The algorithm description is parametric, and so should be any program implementing it. Hence, the quantum programming paradigm is higher-order: a quantum program is a function from (classical data) input parameters to quantum circuits. Then, each instance of a quantum circuit behaves as a function from its (quantum data) inputs to its (quantum data) outputs.

Most of the current quantum programs [153, 83, 193, 181, 180, 27, 41] proceed as implementations for functions such as those declared in Figure 4 (right part of the **Procedure** part), providing no guarantee over the algorithm specifications, be it about their functional behavior or their resource specifications.

Based on the preceding comments, the purpose of formal verification can be summarized as completing algorithm implementations with their proved specifications. In other words:

Formal verification in quantum programming aims at providing solutions to furnish, in addition to quantum programs, evidence that these programs meet their specifications, in terms of both success probability and resource usage.

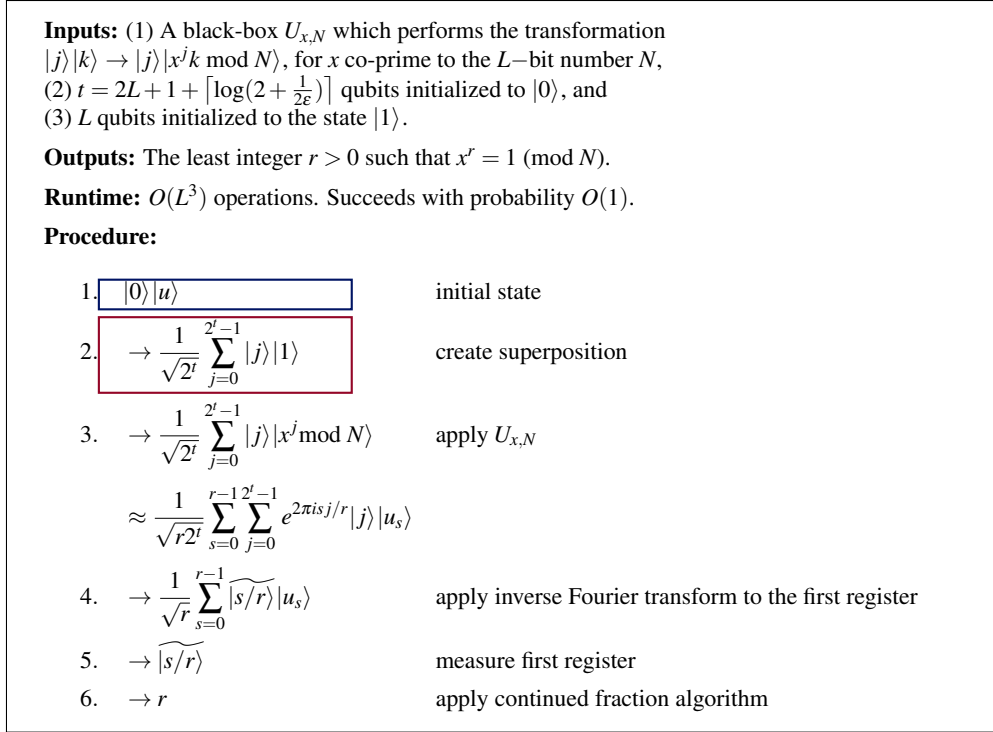


Figure 4: Bird-eye view of the circuit for Shor’s factoring algorithm [174] (as presented in [140, p. 232])

2.3 Challenges for Quantum Computation

Let us now introduce some particularities of quantum programming with regards to classical computing. They raise design challenges that are particular to this programming paradigm.

2.3.1 Destructive Measurement and Non-Determinism

One of the main particularities of quantum programming is that the output produced by the quantum memory device follows the probabilistic Born rule (see Section 2.1.2). So, in the general case, the result of quantum computation is non-deterministic.

Furthermore, a computation in the model from Figure 1 contains both probabilistic quantum computations and classical control structures, performed by the classical controller. Hence, control itself may depend on the probabilistic data received from the quantum device and the execution flow itself is probabilistic.

2.3.2 Quantum Noise

Another particularity comes from the difficulty to maintain big quantum systems in a given state and to control the evolution of this state through time. Along with a quantum computation, uncontrolled modifications (bit or phase flip, amplitude damping, *etc.*) of the quantum state may occur.

To overcome this phenomenon, one solution consists in integrating error correction mechanisms into the compilation. Error correction design is an active research fields [40, 129, 81, 75]. Many propositions have been developed. They mainly consist in designing redundant quantum circuits (one logical qubit is implemented by many different physical qubits). Then the main challenge is to design a solution for testing the reliability along with a computation without losing the state of the register due to destructive measurement.

Since error correction requires bigger quantum registers (due to redundancy), its possible implementation is conditioned by the design, elaboration and availability of large quantum processors.

An alternative strategy is to not correct quantum errors, but design computations to limit their effect. John Preskill introduced the notion of *Noisy-Intermediate Scale Quantum (NISQ)* technologies [152]. The formal analysis of error propagation requires the identification of possible errors together with rules specifying how their probability of occurrences propagates along with quantum circuits [97, 155].

2.3.3 Efficient Compilation on Constrained Hardware

Languages such as *Liqui|*), *Q#*, *Quipper*, *etc.* enable the description and building of quantum circuits for so-called *logical qubits*. In practice, realizing a quantum circuit on an actual quantum machine (*physical qubits*) requires several compilation passes, in addition to the error correction mentioned in the preceding paragraph. Among others:

- the physical realization should respect the physical constraints of its target architecture, which concerns, e.g. connectivity of qubits or register size limits. Considering this point requires qubit reordering intermediary operations and an adequate mapping between theoretical and physical qubits;
- the set of possible quantum operations over physical qubits may not correspond to the set of elementary gates from the logical circuit description, which would require an adequate gate synthesis and circuit rewriting;
- last but not least, physical realization of an algorithm should be as resource frugal as possible, requiring the development of circuit optimization techniques.

Each of these steps consists of low-level operations over circuits. They must all preserve functional equivalence while reaching their proper purpose. For these low-level developing layers, one requires tools and languages formalizing functionally equivalent circuit transformation operations. *ZX-Calculus* [47] (Section 5.1) is particularly well-fitted for such design, other propositions include the proof of path-sum semantics equivalence [5, 4] (Section 5.2) and formally verified circuit optimization [91] (Section 5.4).

3 General Background on Formal Methods

We now present a brief overview of formal methods. While the domain is old and has led to rich literature, we try to highlight the underlying main principles and to quickly describe the most popular classes of techniques so far.

3.1 Introduction

Formal methods and formal verification [44] denote a wide range of techniques aiming at *proving* the correctness of a a system with a *mathematical guarantee*—reasoning over *all* possible inputs and

paths of the system, with methods drawn from logic, automated reasoning and program analysis. The last two decades have seen an extraordinary blooming of the field, with significant case studies ranging from pure mathematics [80] to complete software architectures [112, 124] and industrial systems [20, 107]. In addition to offering an alternative to testing, formal verification has in principle the decisive additional advantages to both enable parametric proof certificates and offer once-for-all absolute guarantees for the correctness of programs.

3.2 Principles

Formal methods' main principles were laid mostly in the 1970s. Pioneers include Floyd [74], Hoare [94], Dijkstra [59], Cousot [49] and Clark [42]. While there is a wide diversity of approaches in the field, any formal method builds upon the following three key ingredients:

- a formal semantics M representing the possible behaviors of a system or program— M is typically equipped with an operational semantics and behaviors are often represented as a set of traces $L(M)$;
- a formal specification φ of the acceptable or correct behaviors— φ is typically a logical formula or an automaton representing a set of traces L_φ ;
- a (semi-)decision procedure verifying that possible behaviors are indeed correct, denoted $M \models \varphi$ —typically a semi-algorithm to check whether $L(M) \subseteq L_\varphi$ holds or not.

Regarding the complexity of realistic systems and programs, the verification problem is usually undecidable, hence the impossibility to have a fully automated and perfectly precise decision procedure for it. The different formal method communities bring different responses to get around this fundamental limitation, yielding different trade-offs in the design space, favoring either restriction of the classes of systems under analysis, restrictions of the classes of properties, human guidance or one-sided answer (over-approximations or under-approximations).

Overall, after two decades of maturation, formal methods have made enough progress to be successfully applied to (mostly safety-critical) software [89, 191, 12, 31, 50, 116, 107].

3.3 The Formal Method Zoo

We present now in more detail the main classes of formal methods. While recent techniques tend to blur the lines and combine aspects from several main approaches, this classification is still useful to understand the trade-off at stake in the field.

- Type checking and unification: at the crossroad of programming language design and formal methods, type systems [150] allow forbidding by design certain classes of errors or bad code patterns (such as trying to add together a number and a boolean in Java or, in a quantum setting, trying to apply a unitary operation over a classical data register). Traditionally, type systems focus on simple “well-formedness” properties (good typing), but they scale very well (modular reasoning) and require only a little manual annotation effort (type inference). While first type systems were based on basic unification [93, 136], advanced type systems with dependent types or flow-sensitivity come closer and closer to full-fledged verification techniques;
- Model checking and its many variants: while initially focused on finite-state systems [42] (typically, idealized protocols or hardware models) and complex temporal properties—with

essentially graph-based and automata-based decision procedures, model checking [43] has notably evolved along the year to cope with infinite-state systems, either through specific decidable classes (e.g., Petri Nets or Timed Automata) or through abstractions. The current approaches to software model checking include notably symbolic bounded verification [116, 35] for bug finding and counter-example guided abstraction refinement [89] for bug finding and proof of invariants (but it may loop forever). Usually, model checking relies on specifications expressed in a variant of modal logic such as temporal logic [42, 151], dynamic logic [86, 122, 72] or mu-calculus [114, 167];

- **Abstract Interpretation:** Generally speaking, Abstract Interpretation [49] is a general theory of abstraction for fixpoint computations. Abstract Interpretation-based static analysis [161] builds over Abstract Interpretation to effectively compute sound (i.e. overapproximated) abstractions of all reachable states of a program. Hence, these techniques are well suited for proving invariants. More precisely, Abstract Interpretation provides a systematic recipe to design sound abstract computation over sets of program states, by connecting the concrete domain (e.g., sets of states) to a given abstract domain (e.g., interval constraints) through a Galois Connexion between an abstraction and a concretization functions. In practice, Abstract Interpretation comes down to computing the fixpoint over the abstract domain, ensuring termination but losing precision. Different abstract domains provide different trade-offs between cost and precision. Historically speaking, the approach targets full automation and implicit properties (e.g., runtime error);
- **Deductive verification and first-order reasoning:** Deductive program verification [17, 71, 94, 182] is probably the oldest formal method technique, dating back to 1969 [94] and the development of Hoare logic. In this approach, programs are annotated with logical assertions, such as pre- and postconditions for operations or loop invariants, then so-called proof obligations are automatically generated (e.g., by the weakest precondition algorithm) in such a way that proving (a.k.a. discharging) them ensures that the logical assertions hold along any execution of the program. These proof obligations are commonly expressed in first-order or separation logic [158] and proven by the help of proof assistants [145, 141] or automatic solvers lying on Satisfiability Modulo Theory [18] or Automated Theorem Proving [73];
- **Interactive proof and second-order reasoning:** some techniques completely drop the hope for automation in favor of expressivity, relying on 2nd order or even higher-order specification and proofs languages—typically in Coq [145] or ISABELLE/HOL [141]. Once programmed and proved in the language, a certified functional program can then often be extracted. This family of approaches is very versatile and almost any problem or specification can be encoded, yet it requires lots of manual effort, both for the specification and the proofs—higher-order proofs can be automatically checked but not found. Still, the technique has been for example used for certified compilers or operating systems [124, 112].

4 Overview of Formal Methods in Quantum Computing

As presented in Section 3, formal methods for proving properties of classical algorithms, programs and systems are well-developed and versatile. In this section, we present the needs for formal methods in the realm of quantum computation, while the later sections are devoted to answering them.

4.1 The Need for Formal Methods in Quantum Computing

As introduced in Section 2, the data structures at stake in quantum computing make the computations hard to represent for developers. Intermediary formal languages are of great help for understanding what quantum programs do and describing their functional behavior.

Furthermore, as introduced in Section 1, directly importing the testing and debugging practices at use in classical programming is extremely difficult in the quantum case⁴, due to the destructive aspect of quantum measurement. Moreover, the probabilistic nature of quantum algorithms seriously impedes system-level quantum testing. As a consequence, test-based programming strategies do not seem adequate in the quantum case and quantum computing needs alternative debugging strategies and methodologies.

So far, existing quantum processors were small enough that their behavior could be entirely simulated on a classical device. Hence, a short-term solution for overcoming the debugging challenge relied on classical simulations of quantum programs. Since quantum computer prototypes are now reaching the size limit over which this simulation will not be possible anymore (among others, [9] is often referred to as the milestone for this context change, referred to as *quantum supremacy*), more robust solutions must be developed.

On the other hand, nothing prevents *a priori* the formal verification of quantum programs. In addition to constituting alternative debugging strategies, formal methods have several additional decisive advantages. In particular:

- They enable parametrized reasoning and certification in the higher-order quantum programming context introduced in Section 2.2.2: formal certification of a parametrized program holds for any circuit generated by this program, whatever the value of its parameters. In contrast, testing certification holds for the particular values of these parameters that are used in a test. Formal certification is not limited by the size of the parameters;
- It provides once for all an absolute, mathematically proven, certification of a program's specifications, whereas testing furnishes at best only statistical arguments based on bounded-size input samples.

4.2 Typology of Properties to Verify

In this section we detail the different properties one has to mind for developing correct quantum programs. The goal of formal certification is to provide solutions for their verification.

4.2.1 Functional Specifications

A major challenge is to give assurance on the input/output relationship computed by a given program, that is verifying whether a given program implements an intended function f . Functional specifications are two-layered :

High-level specifications. We give a circuit an overall specification independent from the concrete implementation. It is made of a success condition and a minimum probability, for any run of this circuit, to result in an output satisfying the success condition.

In the hybrid model, circuits are run on a quantum co-processor but controlled by a classical computer, performing control operations (such as **if** and **while** instructions, simple sequence, *etc.*). As a simple example, an algorithm such as the one from Figure 4 outputs a success with probability

⁴It requires major adaptations and redefinitions, see Section 8.4 for details.

p . It can be included in a control structure including k iterations of it. This higher-level procedure has probability $[1 - (1 - p)^k]$ to output a success at least once, which can be made arbitrarily close to 1.

More generally, a high-level quantum verification framework [131, 194] considers an algorithm as a controlled sequence of quantum functions. There, one considers quantum operations as primitives and composes them together via controlled sequence operations. These operations are interpreted as functions in the semantical formalism. For example, [168, 131, 194] formalize quantum programs in the *density operators formalism*. This view is introduced in detail in Section 7.1 together with the Quantum Hoare Logic (QHL) [131, 194].

Intermediate specifications. In the **Procedure** section from Figure 4, each mid-level step of the algorithm is given a formal specification, that is a description of the state of the system. These intermediate specifications are deterministic and concern quantum data.

In a lower-level verification approach view, instead of inputting quantum operations as primitive functions, one builds quantum circuit implementations of these operations, by adequately combining quantum gates. Such a framework [91, 156, 38] relies on a circuit description language such as Quipper or QWIRE. Then, an adequate semantics characterization for the built circuits enables to reason about the quantum data received as inputs and delivered as outputs. A certification solution for quantum circuits enables us to reason compositionally about their semantics. This programming view is explored in Sections 7.2 and 7.3.

4.2.2 Complexity specifications

The major reason for developing quantum computers and quantum algorithms is to lower computing complexity specifications, *w.r.t.* classical computing solutions (see Section 2.2.1 for precisions). Therefore, the relevance of a quantum implementation relies on the fact it satisfies low complexity specifications. As introduced in Section 2.2.2, they may be formulated through different metrics, such as the width and/or depth of quantum circuits, their number of elementary gates or more complex metrics such as *quantum volume* [125].

The complexity specification is also crucial for another reason: remind from Section 2.3.2 that quantum computation is subject to noise: the bigger a quantum circuit is, the most prone to error it is. Functional specifications introduced so far reason about the theoretical output of quantum computations, in the absence of errors. The risk of error in a circuit is closely related to the structural characteristics of this circuit, among which are the different measures of complexity. Therefore, the information provided by these measures is also crucial to appreciate the functional trustfulness of an implementation.

4.2.3 Structural Constraints

Quantum circuit design must also consider various *structural constraints* that discriminate through several criteria:

1. They can be either relative to a target architecture or absolute (induced by quantum physics laws). The first category comprises, for example, the number of available qubits in a processor, the connectivity between physical qubits, the set of available elementary gates, *etc.* The second category mainly deals with aspects induced by quantum Calculus unitarity (no cloning theorem, ancilla management, quantum control, *etc.*);

2. Now, depending on the programming language at stake, absolute structural constraints may either be taken into charge by the language design or left to the user's responsibility. For example, the *no-cloning rule* is derived from the unitarity of quantum processes. It forbids using the same quantum data register twice:
 - in languages where quantum data registers are full right objects (eg: Quipper, QWIRE, *etc*), caring for the respect of no-cloning is left to the user. In this case, formal verification may help her to do so. Solutions like ProtoQuipper [162] or QWIRE [148, 156] tackle this problem through linear type systems (see Section 6.3);
 - another possibility is to reduce the expressivity of the language (eg: QFC/QPL [168], SQIR [91], QBRICKS [38]), to prevent any possible violation of no-cloning. In SQIR or QBRICKS, quantum data registers are addressed via integer indexes, but the quantum data they hold are not directly accessible from the programming language itself. These data concern the semantics of the language and they are formalized only in the specification language. Hence, the respecting conditions for the no-cloning theorem are reduced to simple indexing rules for quantum circuits.
3. Last, structural program constraints can be either syntactic or semantic. The first category contains, for example, all constraints that are linked with qubit identification (eg. : *do not control an operation by the value of a qubit it is acting on*). The most representative example for semantic constraints concerns a particular aspect of quantum computing that we do not detail in this survey: the management of ancilla qubits. Ancilla qubits provide additional memory for some sections of quantum circuits, the content of which is then discharged at some stage of the computation. Discharging a part of a register is possible (without affecting the rest of the memory) only if there is no interaction between the memory to discharge and the rest of the memory (See [140] for further details). Hence, ancilla management is possible modulo some non-entanglement specifications, regarding the semantics of quantum circuits.

4.2.4 Circuit Equivalence

Compilation of quantum programs contains many circuit rewriting operations (see Section 2.3.3). They concern the implementation of logical qubits in a physical framework and require certification for functional behavior preservation. Concretely, given a logical circuit C , compiling C traces as a chain of circuits, starting from C and each obtained from the precedent by applying a circuit rewriting operation. Each such rewriting must preserve the input/output relation, to ensure that, provided C fits its functional specifications, then so does the final physical qubits circuit. In Section 5 we present two tools enabling the verification of circuit equivalence: the ZX-Calculus (Section 5.1) and the path-sum equivalence verification (Section 5.2).

Further formal comparisons between quantum processes. Different notions of equivalence between quantum processes are also at stake with further uses of quantum information, such as communication protocols. Recent developments [185, 19] generalize the equivalence specification to further comparison predicates between quantum processes. Since they are not designed for the formalization of algorithms, which is the scope of the present survey, we do not detail these propositions in the present survey.

5 Low-Level Verification: Compilation and Equivalence

Realizing logical circuits into physical devices (*circuit compilation*) requires to deal with severe constraints: the number of available qubits, their connectivity, the set of elementary operations, the instability of quantum information—requiring the insertion of error correction mechanisms, *etc.* As mentioned in Section 4.2.4, the underlying circuit transformations must preserve functional equivalence with the initial circuit representation, all along the compilation process. In the present section we introduce formal tools for checking such equivalences and certifying compilation correctness.

5.1 ZX-Calculus and Quantomatic/PyZX

ZX-Calculus [46] is a powerful graphical language for representing and manipulating quantum information. This language historically stems from category theory applied to quantum mechanics, through the program Categorical Quantum Mechanics initiated by Samson Abramsky and Bob Coecke [2].

For our purposes, it is interesting to see ZX-diagrams as a lax version of quantum circuits. This laxness on the one hand implies that not all ZX-diagrams are implementable with physical qubits, but on the other hand, it allows formalism to get powerful results on the underlying equational theory (rewriting rules, pseudo-normal forms).

The level of abstraction provided by the language allows the user to reason about quantum programs or protocols while significantly alleviating the "bureaucracy checks" typically coming with circuit-level reasoning, in particular, checking sub-circuit equivalence in the presence of ancillas. It also allows unifying different models of quantum computation (circuits, measurement-based quantum computing, lattice surgery, *etc.*), as well as to provide optimization strategies for these models. Last but not least, it can be used to formally (yet, graphically) verify properties on protocols or programs—all that based on simple graph-based manipulations.

5.1.1 Semantical Model

The ZX-diagrams are generated from a set of primitives:

$$\left\{ \begin{array}{l} |, \text{X}, \cup, \cap, \text{green circle } \alpha, \text{red circle } \alpha, \text{yellow square } \alpha \end{array} \right\}_{\substack{n, m \in \mathbb{N} \\ \alpha \in \mathbb{R}}}$$

which can be composed either:

- sequentially: $\begin{array}{c} | \dots | \\ \boxed{D_1} \\ | \dots | \\ \boxed{D_2} \\ | \dots | \end{array}$
- or in parallel: $\begin{array}{cc} | \dots | & | \dots | \\ \boxed{D_1} & \boxed{D_2} \\ | \dots | & | \dots | \end{array}$

where D_1 and D_2 are both ZX-diagrams (themselves composed of the above primitives). We denote by \mathbf{ZX} the set of ZX-diagrams. In these, information flows from top to bottom, which is in contrast with quantum circuits where it flows from left to right. This is only a matter of convention, as string diagrams, on which the ZX-Calculus formalism relies upon, are oriented vertically.

These diagrams are used to represent linear maps, thanks to the so-called *standard interpretation* of ZX-diagrams as complex number matrices $[\![\cdot]\!] : \mathbf{ZX} \rightarrow \mathcal{M}(\mathbb{C})$ ⁵. It is inductively defined as:

$$\begin{aligned}
 \begin{array}{|c|} \hline \dots \\ \hline D_1 \\ \hline \dots \\ \hline D_2 \\ \hline \dots \\ \hline \end{array} &= \begin{array}{|c|} \hline \dots \\ \hline D_2 \\ \hline \dots \\ \hline \end{array} \circ \begin{array}{|c|} \hline \dots \\ \hline D_1 \\ \hline \dots \\ \hline \end{array} & \quad \begin{array}{|c|c|} \hline \dots & \dots \\ \hline D_1 & D_2 \\ \hline \dots & \dots \\ \hline \end{array} &= \begin{array}{|c|} \hline \dots \\ \hline D_1 \\ \hline \dots \\ \hline \end{array} \otimes \begin{array}{|c|} \hline \dots \\ \hline D_2 \\ \hline \dots \\ \hline \end{array} \\
 \begin{array}{|c|} \hline \dots \\ \hline 1 \\ \hline \dots \\ \hline \end{array} &= id_{\mathbb{C}^2} = |0\rangle\langle 0| + |1\rangle\langle 1| & \quad \begin{array}{|c|} \hline \dots \\ \hline \text{X} \\ \hline \dots \\ \hline \end{array} &= \sum_{i,j \in \{0,1\}} |ji\rangle\langle ij| \\
 \begin{array}{|c|} \hline \dots \\ \hline \text{C} \\ \hline \dots \\ \hline \end{array} &= \begin{array}{|c|} \hline \dots \\ \hline \text{C} \\ \hline \dots \\ \hline \end{array}^\dagger = |00\rangle + |11\rangle & \quad \begin{array}{|c|} \hline \dots \\ \hline \text{H} \\ \hline \dots \\ \hline \end{array} &= |+\rangle\langle 0| + |-\rangle\langle 1| \\
 \begin{array}{|c|} \hline \dots \\ \hline \alpha \\ \hline \dots \\ \hline \end{array} &= |0^m\rangle\langle 0^n| + e^{i\alpha} |1^m\rangle\langle 1^n| & \quad \begin{array}{|c|} \hline \dots \\ \hline \alpha \\ \hline \dots \\ \hline \end{array} &= |+\rangle\langle +^n| + e^{i\alpha} |-\rangle\langle -^n|
 \end{aligned}$$

where $|+\rangle := \frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and $|-\rangle := \frac{|0\rangle-|1\rangle}{\sqrt{2}}$ and $|u\rangle\langle v|$ is the ket bra outer product from Section 2.1.2.

For example, $id_{\mathbb{C}^2} = |0\rangle\langle 0| + |1\rangle\langle 1| = (1\ 0) \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + (0\ 1) \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$.

Notice that the green (light) and red (dark) nodes only differ from the basis in which they are defined (as $(|+\rangle, |-\rangle)$ defines an orthonormal basis of \mathbb{C}^2) and that they can have an arbitrary number of inputs and outputs. It often happens that a green or red node has a parameter of value 0. In this case, by convention, this angle 0 is omitted. Finally, notice that H represents exactly the Hadamard gate of quantum circuits. This is not a coincidence, as ZX-diagrams can be seen as a generalization of quantum circuits. In particular, we can map any quantum circuit to a ZX-diagram that represents exactly the same quantum operator:

$$\begin{aligned}
 Ph(\theta) &\mapsto \begin{array}{|c|} \hline \text{red} & \text{red} & \pi \\ \hline \text{green} & \text{green} & 2\theta \\ \hline \end{array} & \quad R_Z(\theta) &\mapsto \begin{array}{|c|} \hline \text{green} & \text{red} & \text{red} & \pi \\ \hline \text{green} & \text{green} & \text{green} & -2\theta \\ \hline \end{array} \\
 H &\mapsto \begin{array}{|c|} \hline \text{H} \\ \hline \end{array} & \quad CNOT &\mapsto \begin{array}{|c|} \hline \text{green} & \text{red} & \text{red} \\ \hline \text{green} & \text{green} & \text{green} \\ \hline \end{array}
 \end{aligned}$$

and that preserves sequential and parallel compositions. The elementary gates given above are the ones detailed in Table 1.

We can actually map any *generalized* quantum circuit (i.e. circuit including measure) into a ZX-diagram. Indeed, initializations of qubits are easy to represent: $|0\rangle \mapsto \begin{array}{|c|} \hline \text{red} & \text{red} \\ \hline \text{green} & \text{green} \\ \hline \end{array}$, and there exists an extension of the ZX-Calculus [48, 36] that allows the language to represent measurements. In this extension, we represent the environment as \perp , which becomes an additional generator of the diagrams (we denote by \mathbf{ZX}^\perp this updated set of ZX-diagrams). This generator can also be understood as discarding a qubit. However, contrary to classical data, this action affects the rest of the system. Introducing \perp forces us to change the codomain of the standard interpretation, but we will not give the details here. Simply keep in mind that the measurement in the computational basis ($|0\rangle, |1\rangle$) is

$$\text{represented by } \begin{array}{|c|} \hline \text{green} \\ \hline \perp \\ \hline \end{array} .$$

In this way, we can (fairly) easily represent any generalized quantum circuit as a ZX-diagram. But we can actually represent more, and this is an active field of research to try and characterize diagrams that can be put in circuit form (we talk about “extracting a circuit”). First was introduced the notion of causal flow [52] which was then extended to that of “gflow” (for generalized flow) [33]. Some other variations exist [11].

⁵To be more precise, the standard interpretation associates to any ZX-diagram in $\mathbf{ZX}[n,m]$ (i.e. with n inputs and m outputs) a complex matrix of dimension $2^m \times 2^n$ i.e. in $\mathcal{M}_{2^m \times 2^n}(\mathbb{C})$.

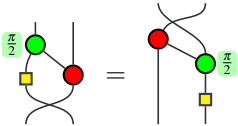
Quantum circuits, however, are not the only computational model one might want ZX-diagrams to compile to. Indeed, it so happens that the primitives of the ZX-Calculus quite naturally match those of *lattice surgery* [55], a scheme for error correction [75, 95]. In particular, ZX-diagrams implementing a (physical) lattice surgery procedure features a special notion of flow, the *PF flow* (for Pauli fusion flow) [54].

5.1.2 Verified Properties

The strength of ZX-Calculus comes from its powerful equational theory. This equational theory allows to define equivalence classes of ZX-diagrams and to conveniently decide whether two different diagrams represent the same quantum operator.

This question can be asked for quantum circuits as well, as two different circuits may represent the same operator (e.g. $H^2 = Id$). Some such equational theories exist for quantum circuits [169, 6], but none, as of today, for a universal fragment of quantum mechanics (notice that in [179] a completeness theorem is given for the interaction between measurements and pure parts of the circuit, but crucially one for the pure part itself is not provided but only assumed).

The main difference between the two formalisms is that the equational theory of the ZX-Calculus allows for a powerful result in this language, aggregated under the paradigm “*only connectivity matters*”. This result states that we can treat any ZX-diagram as an undirected open graph, where the Hadamard box and the green and red nodes are considered as vertices. In particular, any (open) graph isomorphism is an allowed transformation.

Example 5.1.  because the two diagrams can be obtained from one another by simply “moving their nodes around” (while keeping inputs and outputs fixed).

This result also allows us to unambiguously represent a horizontal wire. For instance:



This “meta”-rule, that all isomorphisms of open graphs are allowed, constitutes the backbone of the ZX-Calculus. In what follows, different sets of axioms, that satisfy different needs, will be presented, but this meta-rule will always be there implicitly.

When two diagrams D_1 and D_2 are proven to be equal using the equational theory \mathcal{T} , we write $\mathcal{T} \vdash D_1 = D_2$. The axiomatization \mathbf{zx} for the ZX-Calculus can be found in Figure 5, and it was recently proven to be complete for the standard interpretation $\llbracket \cdot \rrbracket$ [189]:

Theorem 1. \mathbf{ZX}/\mathbf{zx} is complete with respect to $\llbracket \cdot \rrbracket$:

$$\forall D_1, D_2 \in \mathbf{ZX}, \llbracket D_1 \rrbracket = \llbracket D_2 \rrbracket \iff \mathbf{zx} \vdash D_1 = D_2$$

Here \mathbf{ZX}/\mathbf{zx} represents the quotient of \mathbf{ZX} by the equational theory \mathbf{zx} . The completeness property is fundamental. It allows us to reason on quantum processes through diagrammatic transformations rather than by matrix computations. In particular, it tells us that whenever two diagrams represent the same quantum operator, they can be turned into one another using only the rules of \mathbf{zx} .

It is customary in quantum computing to work with particular (restricted) sets of gates. For a lot of such restrictions, there exist complete axiomatizations [85, 101, 102, 100, 103, 104]. The ZX-




Figure 5: The equational theory zx ⁶. All rules – provably – hold in their upside-down and color-swapped (between green and red) versions.

Calculus with measurements has a similar completeness result for ZX^\pm / zx^\pm [36] with an updated set of rules zx^\pm which we will not give here for conciseness purposes.

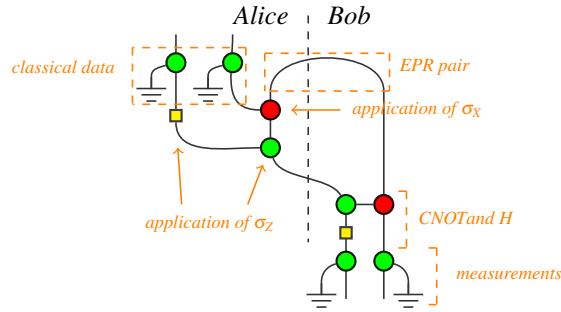
Some properties of quantum protocols or algorithms can then be verified by diagram transformations. To give the reader the flavor of such verifications, we detail the example of superdense coding.

Example 5.2 (Superdense Coding). *The idea of the superdense coding protocol is to transmit two classical bits using a single qubit. This is not possible in general, but it is when the two parties initially share an entangled pair of qubits. The protocol goes as follows:*

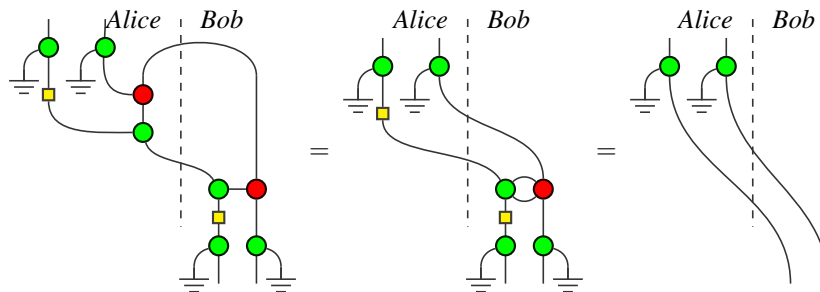
- Alice and Bob initially share the (previously defined) EPR pair $\beta_{00} = \frac{|00\rangle + |11\rangle}{\sqrt{2}}$, and Alice moreover has two bits she wants to send to Bob
- Alice applies $\sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ to her qubit if her first bit is 1, then $\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ if her second bit is 1
- Alice sends her qubit to Bob
- Bob applies a CNOT between his qubit and the one he received from Alice, then a H gate on his qubit, and finally measures his two qubits

⁶The way we denote the equations relates to their names in broader literature if such exists, and more informal names if not. (S) is the spider rule, (I_g) and (I_r) are the green and red identity rules, (E) is the empty diagram rule, (CP) the copy rule, (B) the bialgebra rule, (EU) the Euler angles rule, (HD) the Hadamard decomposition and (H) the Hadamard colour change rule.

This protocol can be represented with a ZX-diagram as follows.




It is then possible to verify that Bob eventually does get (copies of) Alice's bits, using the equational theory (although the whole derivation is not given here):



We can hence see that data is transmitted from Alice to Bob, without any loss. Interestingly, this protocol can be extended for secure communication between the two parties [192]. The larger protocol uses instances of the smaller one to also check whether an eavesdropper has tried intercepting or copying data.

If Bob is aware that the data he received was compromised, he can abort everything by simply discarding his qubits, so that the eavesdropper (Eve) gets absolutely no information:



where U denotes an unknown operator applied by Eve to the qubit she intercepted. Notice here how no information can pass from Alice to Eve. No information is retrieved by the latter.

A plethora of quantum protocols have been verified with ZX-Calculus in a similar manner [92]. Note however that the theory is ever evolving, and in particular, \perp was not introduced in the language at that time, so the author had to use a trick to make up for the absence of measurement (namely case-based reasoning).

5.1.3 Algorithms and Tools

It is possible to manipulate ZX-diagrams in a computer-verified way. For instance, Quantomatic [110] allows the users to define at the same time diagrams in graphical form and equational theories. It is also possible to work with user-defined nodes in the diagram so that even though $\frac{\perp}{\perp}$ is not part of the “vanilla” ZX-Calculus, it can be defined as a new node. It is then possible in the tool to manipulate diagrams in a way that satisfies the equational theory, and even to define rewriting strategies that can be then applied in an automated way.

The verification of protocols and programs using the ZX-Calculus relies on diagrammatic equivalence. This problem, in general, is at least QMA-complete, [29, 98] (the quantum counterpart of NP-completeness). This problem is linked to the one of simplification/optimization, which asks how a quantum operator can be simplified, given a particular metric (e.g. the number of non-Clifford gates). Indeed, for instance if D_1 and D_2 are two diagrams representing the same *unitary* (i.e. $\llbracket D_1 \rrbracket = \llbracket D_2 \rrbracket$), then simplifying $D_2^\dagger \circ D_1$ should ideally get us to the identity.

In the case of the Clifford fragment (obtained when the angles in $\begin{pmatrix} n \\ \vdots \\ \bullet \\ \vdots \\ m \end{pmatrix} \alpha$ and $\begin{pmatrix} n \\ \vdots \\ \bullet \\ \vdots \\ m \end{pmatrix} \alpha$ are restricted to multiples of $\frac{\pi}{2}$), there exists a strategy that reduces the diagram in (pseudo-)normal form [10]. When this algorithm terminates, the resulting diagram is of size $O(n^2)$ where n is the number of inputs and outputs in the diagram. The algorithm is polynomial in the overall size of the diagram it is applied on.

Turning an arbitrary diagram into a normal form can be done in principle [103], however, the complexity of this algorithm is EXPSPACE for universal fragments. So this approach is obviously not preferred in general. However, one can use the ideas of the algorithm for the Clifford fragment as a starting point to get a rewriting strategy for the general case. Applications to quantum circuits and improvements on this strategy can be found in the literature [109, 53, 60, 11], and implementations in the PyZX tool [108]. This tool can in particular be used to tackle circuit equivalence verification, using a different but related approach to that of Section 5.2 below. The formalism used later is that of *path-sums*, where morphisms were showed in [123, 190] to be essentially equivalent to ZX-diagrams, allowing us to apply strategies for path-sums to the ZX-Calculus and vice-versa. Next Section is devoted to we introducing the path-sum formalism and its use for circuits verification.

5.2 Path-Sum circuit Equivalence Verification

Path-sums are a recent alternative direction for verifying the equivalence between quantum circuits [4, 5]. In this section we briefly present it, together with the main verification related achievements.

Note that a generalization of path-sum semantics is introduced in Section 7.2, for parametrized families of circuits. For sake of readability, conciseness and coherence with this further content, in the coming paragraphs we slightly simplify path-sums related notations. We refer the reader to the original definitions [4, 5] for the full formalism and underlying mathematical structures.

5.2.1 Semantical Model

The standard semantics for quantum circuits is the matrix formalism, introduced in Section 2.1.5. It associates to each quantum circuit C a matrix $\mathbf{Mat}(C)$ and it interprets the behavior of this circuit as a function $|x\rangle \rightarrow \mathbf{Mat}(C) \cdot |x\rangle$ from kets to kets, where \cdot stands for the usual matrix product.

Notice that this standard semantics builds on an intermediary object—the matrix—to derive and interprets the functional behavior of circuits. And it does so by use of a higher-order function—the matrix product. Contrarily, path-sums are a straight construction of the *input/output* function

performed by circuits, enabling compositional reasoning.

Concretely, a path-sum $PS(x)$ is a quantum register state (a ket), parametrized by an input basis ket $|x\rangle$ and defined as the sum of kets

$$PS(x) ::= \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} e^{\frac{2\pi \cdot i \cdot P_k(x)}{2^m}} |\phi_k(x)\rangle \quad (2)$$

where the $P_k(x)$ are called *phase polynomials* while the $|\phi_k(x)\rangle$ are *basis-kets*. This representation is *closed* under functional composition and Kronecker product. For instance, if

$$\begin{aligned} C : |x\rangle &\mapsto PS(x) = \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} e^{\frac{2\pi \cdot i \cdot P_k(x)}{2^m}} |\phi_k(x)\rangle, \\ C' : |y\rangle &\mapsto PS'(y) = \frac{1}{\sqrt{2^{n'}}} \sum_{k=0}^{2^{n'}-1} e^{\frac{2\pi \cdot i \cdot P'_k(y)}{2^{m'}}} |\phi'_k(y)\rangle, \end{aligned}$$

then their parallel combination $\text{parallel}(C, C')$ sends $|x\rangle \otimes |y\rangle$ to:

$$\frac{1}{\sqrt{2^{n+n'}}} \sum_{j=0}^{2^{n+n'}-1} e^{\frac{2\pi \cdot i \cdot (2^{m'} \cdot P_{j/2^n}(x) + 2^m \cdot P'_{j\%2^n}(y))}{2^{m+m'}}} |\phi_{j/2^n}(x)\rangle \otimes |\phi'_{j\%2^n}(y)\rangle$$

The sequential combination of quantum circuits C and C' receives a similar compositional definition, parametrized by path-sums components for circuits C and C' .

5.2.2 Path-Sums Reduction

While path-sums compose nicely, a given linear map (eg. the input/output function for a quantum circuit) does not have a unique representative path-sum. Hence, given two different path-sum, how to decide whether they both encode the behavior of a given circuit? To tackle this problem, an equivalence relation is defined with a few, simple rules that can be oriented. As an example, the *HH* rule enables to simplify a path-sum expression over the reduction from a sequence of two consecutive Hadamard gates to the identity—see Figure 6. All these rules transform a path-sum into an equivalent one, with a lower number of *path variables* (parameter n in the notation of Equation 2).

$$\frac{PS(x) = \frac{1}{\sqrt{2^{n+1}}} \sum_{j=0}^2 \sum_{k=0}^{2^n} e^{\frac{1}{2} \cdot \pi \cdot i \cdot (\frac{1}{2} j(k_i + Q_k(x)) + P_k(x))} |\phi_k(x)\rangle}{PS(x) = \frac{1}{\sqrt{2^{n+1}}} \sum_{k=0}^{2^n} e^{\frac{1}{2} \cdot \pi \cdot i \cdot (P[i:=Q_k(x)]_k(x))} |\phi[i := Q_k(x)]_k(x)\rangle} \quad HH$$

Figure 6: The *HH* path-sum transformation rule

We refer the desirous reader to [5, 4] for an exhaustive exposition of the corresponding proof system. It was proved strongly normalizing, meaning that every sequence of reduction rules application terminates with an irreducible path-sum. Furthermore, finding and applying such a normalizing sequence is feasible in time polynomial in the width of the circuit at stake, which makes the overall reduction procedure tractable.

5.2.3 Verified Properties

Hence, path-sums provide a human-readable formalism for the interpretation of quantum circuits as ket data functions. Furthermore, it is given a polynomial normalization procedure, based on a restricted set of rewriting rules. The method was probed against both circuit equivalence and functional specifications verification. More precisely:

Translation validation consists, for a given quantum algorithm, in (1) computing the path-sums for both a non-optimized and an optimized circuit realization and (2) using the normalization procedure for the automatic checking of their equivalence. It was performed on various quantum routine instances (Grover, modular adder, Galois field multiplication, *etc*) of various size (up to several dozens of qubits). Interestingly, the methods proved as efficient for identifying non-equivalence (over erroneous instances) as for checking equivalence.

Quantum algorithms verification consists in verifying whether a given circuit instance respects its functional description. It was performed for instances of similar case studies as for the translation validation (QFT, Hidden Shift [187]), with up to a hundred qubits.

As conclusion, the path-sum formalism provides a fully automatized procedure for verifying the equivalence between two circuits. Hence, given two quantum circuits, the latter being a supposed optimized version of the former, path-sums treatment enables us to verify that they implement the same quantum function. Since path-sums perform internal complexity reduction, an open direction is a search for efficient heuristics extracting an optimized quantum circuit from reduced path-sums. As mentioned in Section 5.1, this problem is closely linked to the reduction of ZX-Calculus diagrams. So, in its general form, it faces the same complexity limitations. The search for efficient reduction procedure applying to identified useful fragments could then both benefit from and feed advances in the ZX setting.

In its present state of development, using path-sums for quantum algorithm verification is restricted to compilation time, when program parameters are instantiated. Furthermore, a new run of the path-sum reduction is required at each new call of a given quantum function. In Section 7.2 we introduce the QBRICKS language, whose semantics is based on a parametrized extension of path-sums. It enables the verification, once for all, of parametrized programs, holding for any possible future parameter instances—yet at the price of full automation as the manipulation of parametrized path-sums requires first-order logic reasoning.

5.3 Quantum abstract interpretation

The techniques presented so far target an exhaustive functional description of a quantum circuit. Because of the intrinsic complexity of quantum computing, their use is limited, in particular by the size of circuits. As illustrated in the benchmarks summed up in Section 5.2.3, path-sums reduction enabled, for instance, the verification of quantum circuits up to a hundred of qubits.

To push the boundary, a possible strategy comes from Abstract Interpretation (see Section 3.3). There, one does not target an exhaustive description of the functional behavior of a circuit, but an over-approximation of it. Such a framework relies on:

- the identification of a conveniently structured set of properties of interest (the *abstract domain*);
- sound abstraction and concretization to/from this abstract domain;
- reasoning tools for the abstract domain.

In [149], the author introduces such an abstract interpretation of quantum states based on their entanglement structure. The technique enables to identify mutually separable sub-registers. It is useful, for example, for uncomputation, qubit discarding or identifying a convenient decomposition for further analysis of a state.

A more recent development on abstract interpretation applied to quantum process is [199], in which the authors define abstract domains made of tuples of partial projections over quantum sub-registers. Intuitively, the idea is to overcome the exponential complexity of quantum states by decomposing them into sub-spaces. Interestingly, the method was implemented and evaluation results are provided. In particular, this abstraction enabled to characterize the invariant in the main loop structuring the Grover search algorithm and to prove it for instances of width up to 300 qubits.

5.4 Toward Integrated Verified Optimization: VOQC

A noticeable effort for an integrated verified quantum optimization was recently led through the development of a *Verified Optimizer for Quantum Circuits* (VOQC—pronounced ‘vox’) [91]. As main aspects, in comparison to ZX- and path-sum calculus, VOQC:

- is integrated into a core programming environment and applies on circuits issued from parametrized programs;
- not only validates the equivalence between an input quantum circuit and a candidate optimized version of it, but also provides the formally verified optimization procedure, directly generating this optimized version.




Figure 7: A simplified view of VOQC architecture

5.4.1 Architecture

In Figure 7 we give a simplified view of VOQC architecture: it relies on an Intermediate Representation Language named SQIR. Since it may also be used as a verified programming environment, SQIR is introduced and developed *per se* in a dedicated section below (Section 7.3). In the present context, we can think of it as a core language generating parametrized quantum circuits. Then, for any instance of the parameters, VOQC extracts the corresponding sequence of operations, applies an optimization procedure upon this sequence of operations and extracts back an optimized proved equivalent SQIR quantum circuit. Furthermore,

- in addition to optimization, VOQC also contains some circuit mapping functionalities. They perform further circuit transformation so that the output circuits fit to specific quantum architecture qubit connectivity constraints. This functionality is a preliminary address to the problem raised as first bullet in Section 2.3.3,
- VOQC environment also provides both ways compilations between SQIR and the standard assembly language OPENQASM [51]. Hence, it opens the way for a modular easy integration in any standard programming environment, in particular QISKIT [153], which uses OPENQASM as assembly language.

5.4.2 Optimization Procedure

VOQC optimization process provides two functionalities, one is deterministic (optimization by propagation and cancellation) and the other one requires a replacing circuit input.

Optimization by propagation and cancellation is based on local circuit rewriting schemes and self composition properties of elementary gates, borrowed from [138] (eg., sequences of an even number of either H , $CNOT$ or X gate annihilate as the identity, successive occurrences of R_z gates melt by summing their angle parameters, *etc*). Hence, the procedure consists of two successive steps:

- propagate: for any elementary gate, it
 - considers several identified patterns enabling gate commutations,
 - finds all occurrences of these patterns,
 - and performs the related commutation, pushing any occurrence of this elementary gate to the end of the computation,
- cancellation then consists in deleting the resulting repetitive occurrences of the elementary gate at stake.

Optimization by circuit replacement consists in substituting a part of a quantum circuit (a sub-circuit) by another one that is proven to be functionally equivalent. In this case, the equivalence proof is led by help of the path-sums semantics (see Section 5.2). Hence, in its most general form, the process requires an external equivalence proof oracle. Nevertheless, VOQC provides some instances of such proved equivalence patterns (eg, rotation merging), whose application is automatic.

Performance and Achievements. VOQC performance has been evaluated against several standard quantum computation routines, and compared with several existing optimizers [7, 153, 138, 178, 108]. Note that in VOQC, since optimization is performed as a succession of rewriting operations, the formal verification consists in assessing the equivalence between the input and the output of the optimization, it does not address the optimization performance. Still, on reported experiments, VOQC performance competes with other existing non-verified optimizers, both regarding computation time and circuit complexity reduction—precise performance comparison tables appear in [91]. Thus, in the current state of the art, the benefits of formally verified circuit optimization comes for free.

5.5 Formally verified quantum compilation in an imperative setting: CertiQ

CertiQ [173] is another noticeable effort for verified compilation. Interestingly, it applies to QISKIT language, which is certainly *the most widely used quantum compiler*. CertiQ proceeds similarly as VOQC, by applying successive optimization *passes*, each consisting in applying local circuit equivalence transformations over quantum circuits. CertiQ has been evaluated against the very compilation environment provided by IBM QISKIT framework. And verified optimization went through 26 out of the 30 compilation passes the framework offered at writing time. Hence, the current limitations of formally verified compilation is mainly inherited from the state of the art in (non verified) quantum compilation.

The application to QISKIT also inevitably brings an additional drawback: while CertiQ can bring insurance that compilation respects the functional equivalence between an input and an output quantum circuits, the initial circuit building brick still lacks formal verification. Indeed, the environment is not provided formal means to assess that the input circuit meets a given functional specification. An open work direction then is to develop formal circuit building verification methods similar as SQIR or QBRICKS but applying to (mostly imperative) widely used development environments such as QISKIT, so as to complete the verified development chain.

6 Formal Quantum Programming Languages

For reasoning on a concrete *runs* of quantum algorithms solving a specific problem *instances*, the notion of quantum circuit is natural. Yet, a quantum circuit is only a by-product of a specific run of a quantum algorithm, holding only for a specific instantiation of the algorithm parameters.

Hence, a quantum algorithm is not reducible to a quantum circuit. To run quantum algorithms and reason about them, one needs *quantum programming languages* (QPL): this is the topic of the current section.

6.1 Quantum Programming Languages Design

In Section 5, circuits were merely seen as sequences of elementary gates. However, in most quantum algorithms circuits follow a more complex structure: they are built compositionally from smaller sub-circuits and circuits combinators. Circuits are usually *static* objects, buffered until completion before being flushed to the quantum co-processor. Still, in some algorithms, they are *dynamically* generated: the tail of the circuit depends on the result of former measurements.

In this section, we discuss the high-level structure of quantum algorithms, the requirements for a quantum programming language, and review some of the existing proposals.

6.1.1 Structure of Quantum Algorithms

The usual model for quantum computation was depicted in Figure 1: a classical computer controls a quantum co-processor, whose role is to hold a quantum memory. A programmatic interface for interacting with the co-processor is provided to the programmer sitting in front of the classical computer. The interface gives methods to send instructions to the quantum memory to allocate and initialize new quantum registers, apply unitary gates on qubits, and eventually perform measurements. Even though the set of instructions is commonly represented as a circuit, it is merely the result of a *trace of classical execution* of a classical program on the classical computer.

Figure 8 presents two standard workflows with a quantum co-processor. In Figure 8a, the classical execution inputs some (classical) parameters, performs some pre-processing, generates a circuit,

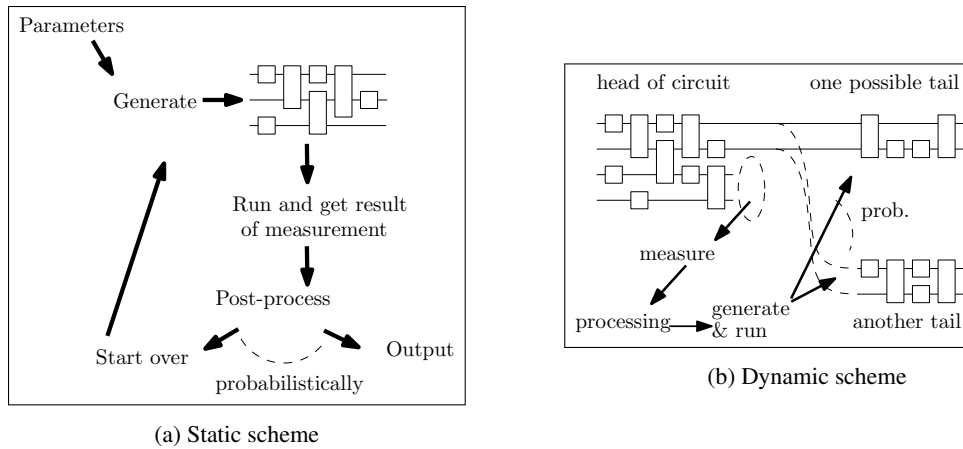


Figure 8: Workflows for quantum algorithm

sends the circuit to the coprocessor, collects the result of the measurement, and finally performs some post-processing to decide whether an output can be produced or if one needs to start over. Shor’s factoring algorithm [175] or Grover’s algorithm [84] fall into this scheme: the circuit is used as a fancy probabilistic oracle. Most of the recent variational algorithms [135] also fall into this scheme, with the subtlety that the circuit might be updated at each step. The other, less standard workflow is presented in Figure 8b. In this scheme, the circuit is built “on the fly”, and measurements might be performed on a sub-part of the memory along the course of execution of the circuit. The latter part of the circuit might then depends on the result of classical processing in the middle of the computation.

Understanding a quantum circuit as a by-product of the execution of a classical program shines a fresh view on quantum algorithms: it cannot be identified with a quantum circuit. Instead, in general, at the very least a quantum algorithm describes a *family* of quantum circuits. Indeed, consider the setting of Figure 8a. The algorithm is fed with some parameters and then builds a circuit: the circuit will depend on the shape of the parameters. If for instance, we were using Shor’s factoring algorithm, we would not build the same circuit for factoring 15 or 114,908,028,227. The bottom line is that a quantum programming language should be able to describe *parametrized* families of circuits.

The circuits described by quantum algorithms are potentially very large—a concrete instance of the HHL algorithm [87] for solving linear systems of equations has been shown [165] to count as much as $\sim 10^{40}$ elementary gates, if not optimized. Unlike the circuit-construction schemes hinted at in Section 5, this circuit is not uniquely given as a list of elementary gates: it is built from sub-circuits—possibly described as a list of elementary gates but not only—and from high-level circuit combinators. These combinators build a circuit by (classically) processing a possibly large sub-circuit. Some standard such combinators are shown in Figure 9 (where we represent inverse with reflected letters). There is a distinction to be made between the combinator, applied on a sub-circuit, and its semantics, which is an action on each elementary gate. Combinators are abstractions that can be composed to build larger combinators, such as the one presented in Figure 10, built from inversion, controlling and sequential composition.




Figure 9: Standard Circuit Combinators

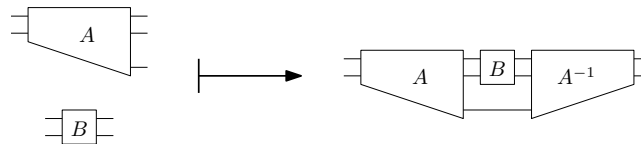


Figure 10: Example of a derived circuit combinator

6.1.2 Requirements for Quantum Programming Languages

Any scalable quantum programming language should therefore allow the following operations within a common framework.

- Manipulation of quantum registers and quantum circuits as first-class objects. The programmer should both be able to refer to “wires” in a natural manner and handle circuits as independent objects.
- Description of *parametric families* of quantum circuits, both in a procedural manner as a sequence of operations–gates or subcircuits–and in an applicative manner, using circuit combinators;
- Classical processing. In our experience [83], quantum algorithms mostly consists of classical processing–processing the parameters, building the circuits, processing the result of the measurement.

This broad description of course calls for refinements. For instance some of the classical processing might be performed on the quantum co-processor–typically the simple classical controls involved in quantum error correction. The level of classical processing performed on the classical computer and performed on the quantum co-processor is dependent on the physical implementation. Even though some recent proposals such as Quingo [183] discuss the design of quantum programming languages aware of the two levels of classical processing–in and out of the co-processor, this is still a work in progress.

6.1.3 Review of the Existing Approaches

Most of the current existing quantum programming languages follow the requirements discussed in Section 6.1.2. In this section, we review some typical approaches followed both in academic and

industrial settings. This review is by no means meant to be exhaustive: its only purpose is to discuss the possible strategies for the design of QPLs.

When designing a realistic programming language from scratch, the main problem is access to existing libraries and tools. In the context of quantum computation, one would need for instance to access the file system, make use of specific libraries such as Lapack or BLAS, *etc.* One can also rely on the well-maintained and optimized compiler or interpreter of the host language. To quickly come up with a scalable language, the easiest strategy consists in *embedding* the target language in a host language. Indeed, a quantum programming language can be seen as a domain-specific language (DSL), and it can be built over a regular language.

Even though the advantages of working inside a host language are clear, there are two main drawbacks. The first one is the potential rigidity of the host language: there might be constructs natural to the DSL that are hardly realizable inside the host language. The second drawback has to do with the compilation toolchain: the shallow embedding of the DSL makes it impossible to access its abstract syntax tree, rendering specific manipulation thereof impossible.

Embedded QPLs. The first scalable embedded proposal is Quipper [83, 82]. Embedded in Haskell, it capitalizes on *monads* to model the interaction with the quantum co-processor. Quipper’s monadic semantics is meant to be easily abstracted and reasoned over: it is the subject of Section 6.2.2. Since Quipper, there has been a steady stream of embedded quantum programming languages, often dedicated to a specific quantum co-processor or attached to a specific vendor, and mostly in Python: QISKIT [153] and PROJECTQ [180] for IBMQ, CirQ [41] for Google, Strawberry Fields [106] for Xanadu, AQASM for Atos, *etc.* From a language-design point of view, most of these approaches heavily rely on Python objects to represent circuits and operations: the focus is on usability and versatility rather than safety and well-foundedness.

Standalone QPLs. On the other side of the spectrum, some quantum programming languages have been designed as standalone languages, with their own parser and abstract syntax tree. Maybe the first proposed scalable language was Ömer’s QCL [142]. Ömer experimented several features such as circuit-as-function, automatic inversion and oracle generation. However, due to its non-modular approach, the language did not have successors.

Liqui|) [193] and its sequel Q# [181] developed by Microsoft are good examples of an attempt at building a standalone language while keeping a tight link with an existing programming environment, as Liqui|) and Q# are tightly linked with the F# framework (itself embedded in the whole .NET framework), making it possible to easily “reuse” library functions from within a Q# piece of code. On the other hand, Q# has its own syntax and type system, to capture run-time errors specific to quantum computation.

ScaffCC [99] is another example of a standalone QPL. Even though the language is rather low-level its compiler has been heavily optimized and experimented over, and it serves as a support for a long stream of research on quantum compiler optimizations.

The last noteworthy language to cite in the series is SILQ [27], as it serves as a good interface with the next paragraph: aimed at capturing most of the best practices in terms of soundness and safety, it is nonetheless targeted toward usability.

Formal QPLs. The last line of works on QPLs we would like to mention here is formal languages aimed at exploring and understanding the design principles and the semantics of quantum algorithms. We shall be brief as this is the topic of the remainder of this survey. The initial line of work

was initiated by Selinger [168] with the study of a small flow-chart language with primitive constructs to interact with the quantum co-processor: qubit initialization, elementary gate application and measurement. This language was later extended to a simple Lambda-Calculus with similar primitive quantum features [170, 172]. Even though the language is not aimed at full-scale quantum algorithms, it is nonetheless enough to serve as a testbed for experimenting type systems and many operational [170, 57, 120] and denotational [143, 171, 58, 88] semantics.

The study of formal QPLs took a turn toward circuit-description languages *à la* Quipper with the development of scalable quantum languages. One of the first proposals of formalization is QWIRE [148], embedded in the proof assistant Coq. QWIRE uses Coq expressive system to encode the sophisticated typing rules of QWIRE. In a sense, Coq type system is expressive enough to use Coq as a host language *and* still be able to manipulate the abstract syntax tree of a program. The main design choice for QWIRE is to separate pure quantum computation with its constraints such as no-cloning, from classical computation.

Albeit disconnected from Coq, the formalization of Quipper has followed a similar root. This development is based on the formal language ProtoQuipper [162], which extracts the critical features out of Quipper: the creation and manipulation of circuits using a minimal Lambda-Calculus. The language is equipped with a linear type system and a simple operational semantics based on circuit construction. The simple core proposed by ProtoQuipper has stirred a line of research on the topic, including the formalization of inductive datatypes, recursion and dependent types in this context [160, 159, 130, 77].

The last class of formal programming language we want to mention focuses on the specification and verification of high-level properties of programs, and are solely based on circuit manipulation: unlike QWIRE or Quipper, qubits are not first-class objects and circuits are simple “bricks” to be horizontally or vertically stacked. In this class of languages, one can mention qPCF [144], mainly a theoretical exploration of dependent type systems in this setting, and QBRICKS, presented in Section 7.2.

6.2 Formalizing the Operational Semantics

In order to reason on quantum programming languages, one needs to have a formal understanding of their operational semantics.

6.2.1 Quantum Lambda-Calculi

The Lambda-Calculus [16] is a formal language encapsulating the main property of higher-order functional languages: functions are first-class citizens that can be passed as arguments to other functions. Lambda-Calculus features many extensions to model and reason about side-effects such as probabilistic or non-deterministic behaviors, shared memory, read/write, *etc.*

One of the first formal proposals of a quantum, functional language has precisely been a *quantum* extension of Lambda-Calculus [170]. On top of the regular Lambda-Calculus constructs, the *quantum Lambda-Calculus* features constants to name the operations of qubit initialization, unitary maps and measurements. A minimal system consists of the following terms:

$$M, N ::= x \mid \lambda x.M \mid MN \mid \\ \text{tt} \mid \text{ff} \mid \text{if } M \text{ then } N_1 \text{ else } N_2 \mid \\ \text{qinit} \mid U \mid \text{meas.}$$

Terms are represented with M and N , while variables x range over an enumerable set of identifiers. The term $\lambda x.M$ is an *abstraction*: it stands for a function of argument x and of body M .

The application of a function M to an argument N is represented with MN . To this core Lambda-Calculus, we can add constructs to deal with booleans: tt and ff are the boolean constant values, while $\text{if } M \text{ then } N_1 \text{ else } N_2$ is the usual test. Finally, qinit stands for qubit initialization, meas for measurement and U ranges over a set of unitary operations. These three constants are functions: for instance, $\text{qinit } \text{tt}$ corresponds to $|1\rangle$ and $\text{qinit } \text{ff}$ to $|0\rangle$, while meas applied to a qubit stands for the measurement of this qubit. A fair coin can then be represented by the term

$$\text{meas}(\text{H}(\text{qinit } \text{tt})), \quad (3)$$

where H stands for the Hadamard gate.

The question is now: how do we formalize the evaluation of a piece of code? In the regular Lambda-Calculus, evaluation is performed with *substitution* as follows:

$$(\lambda x.M)N \rightarrow M[x := N]$$

where $M[x := N]$ stands for M where all free occurrences of x —i.e., those corresponding to the argument of the function—have been replaced by N . Even though we can still require such a rule in the context of the quantum Lambda-Calculus, this does not say how to deal with the term $\text{qinit } \text{tt}$.

In order to give an operational semantics to the Lambda-Calculus, a naive idea could be to add yet another construction: a set of constants $c_{|\phi\rangle}$, once for every possible qubit state $|\phi\rangle$. If—as shown by van Tonder [188]—this can somehow be made to work, a more natural presentation consists in mimicking the behavior of a quantum co-processor, in the style of Knill’s QRAM model [113]: we define an *abstract machine* $(|\phi\rangle_n, L, M)$ consisting of a finite memory state $|\phi\rangle_n$ of n qubits, a term M with n free variables x_1, \dots, x_n , and a *linking function* L , bijection between $\{x_1, \dots, x_n\}$ and the qubit indices $\{1, \dots, n\}$. Variables of M captured by L are essentially pointers to qubits standing in the quantum memory. The fair coin of Eq. (3) then evaluates as follows.

$$\begin{aligned} & (| \rangle_0, \{ \}, \text{meas}(\text{H}(\text{qinit } \text{tt}))) \\ & \rightarrow (|1\rangle_1, \{x \mapsto 1\}, \text{meas}(\text{H}x)) \\ & \rightarrow \left(\frac{1}{\sqrt{2}}(|0\rangle_1 - |1\rangle_1), \{x \mapsto 1\}, \text{meas}x\right) \\ & \rightarrow \begin{cases} (|0\rangle_1, \{ \}, \text{ff}) & \text{with prob. } 0.5 \\ (|1\rangle_1, \{ \}, \text{tt}) & \text{with prob. } 0.5. \end{cases} \end{aligned}$$

In this evaluation, most of quantum computation has been exemplified: initialization of qubits, unitary operations and measurements. Handling the latter in particular requires a *probabilistic evaluation*, and this requires some care—we invite the interested reader to consult for example Selinger & Valiron [170] for details.

6.2.2 Monadic Semantics

The operational semantics of the quantum Lambda-Calculus is very limited. Indeed, as discussed in Section 6.1, quantum algorithms do not in general send operations one by one to the quantum co-processor: instead, a quantum program must build circuits (or pieces thereof) before sending them to the co-processor as batch jobs. The quantum Lambda-Calculus does not allow to build circuits: operations can only be sent one at a time. In particular, there is no possibility to create, manipulate and process circuits: circuit generation in the quantum Lambda-Calculus is a *side-effect* that is external to the language. One cannot interfere with it, and embedding the quantum Lambda-Calculus as it stands inside a host language, as suggested in Section 6.1.3, would not help.

The solution devised by Quipper consists in relying on a special language feature from Haskell called *monad*. A monad is a type operator encapsulating a side effect. Consider for instance a

probabilistic side effect. There are therefore two classes of terms: terms without side-effect, with types e.g. `Bool`, or `Int`, and terms with side-effect, with types e.g. `P(Bool)` or `P(Int)` standing for “term evaluating to a boolean/integer, possibly with a probabilistic effect”. The operator `P(-)` captures the probabilistic side effect.

A monad comes with two standard maps. In the case of `P` we would have:

```
return :: A -> P(A)
bind   :: P(A) -> (A -> P(B)) -> P(B)
```

The `return` operation says that an effect-free term can be considered as having an effect—in the case of the probabilistic effect, it just means “with probability 1”. The `bind` operation⁷ says how to compose effectful operations: given a function inputting `A` and returning an object of type `B` with a probabilistic side-effect, how to apply this function to a term of type `A` also having a probabilistic effect? We surely get something of type `P(B)`, but the way to construct it is described by `bind`. A few equations have to be satisfied by `return` and `bind` for them to describe a monad. For instance, `bind return` is the identity on `P(A)`. There can of course be more operations: for instance, we can add to the signature of `P` an operator `coin` of type `() -> P(Bool)`⁸.

A nice property of monads is that effectful operations can be written with syntactic sugar in an imperative style:

```
do
  x <- coin ()
  if x then return 0 else return 1
```

is a term of type `P(Int)` equals to

```
bind (coin ()) (\x.if x then return 0 else return 1)
```

once the syntactic sugar has been removed.

Following this approach, quantum computation can be understood as side-effect: it combines both (1) Read/Write effect, since gates are sent to the coprocessor, and results of measurements are received; (2) Probabilistic effects, since measurement is a probabilistic operation. The first attempt at formalizing this monad is Green’s quantum IO monad [3]: it has then been further developed in Quipper [83].

Internally, circuits in Quipper are represented using a simple inductive datatype akin to a list of gates⁹. The interaction with the quantum-coprocessor is modeled using a specific I/O monad `Circ`. This monad encapsulates the construction of circuits featuring wires holding qubits but also wires holding bits —results of measurements. A bit is however uniquely useable “inside” the monad: to use it in Haskell—in an if-then-else construct for instance—we need to “lift” it into a regular Boolean. The signature of the monad in particular includes

```
qinit :: Bool -> Circ(Qubit)
measure :: Qubit -> Circ(Bit)
hadamard :: Qubit -> Circ(Qubit)
```

A fair coin can be implemented as a circuit returning a bit, of type `() -> Circ Bit`.

⁷In Haskell, this map is denoted with `>>=`. For the sake of legibility, here we denote it with `bind`.

⁸In Haskell, the unit type is denoted with `()`.

⁹Technically a tree structure, as measurements entails branching.


```

bitcoin () = do
  q <- qinit True
  q' <- hadamard q
  r <- measure q'
  return r

```

The function `bitcoin` will merely generate a computation —producing a circuit— waiting to be executed. To implement the toss-coin of Section 6.2.2, we then need to “run” the circuit for lifting the bit into a Boolean. Provided that we have a function

```
run :: Circ Bit -> P Bool
```

one can then implement `coin ()` as `run (bitcoin ())`.

Thanks to the monadic encapsulation, circuits can be manipulated within Haskell. For instance, inversion and control can be coded in Haskell as circuit combinators with the following types.

```

inverse :: (a -> Circ b) -> (b -> Circ a)
control :: (a -> Circ b) -> ((a,Qubit) -> Circ (b,Qubit))

```

The compositionality of the monadic semantics also makes it possible to automatically construct oracles out of classical specification [83, 186].

Compared with the quantum λ -Calculus discussed in Section 6.2.1, where the program can only send gates one by one to the co-processor, Quipper gives to the programmer the ability to manipulate circuits. Although both the quantum λ -Calculus and Quipper represent quantum computations mathematically, Quipper provides a richer model, better-suited for program specification and verification than the plain λ -Calculus.

6.3 Type Systems

In Section 6.2, we discussed how to model the operational behavior of a quantum program. We have however not mentioned yet the run-time errors inherent to quantum computation. In the classical world, type systems are a standard strategy to catch run-time errors at compile-time. Several run-time errors specific to quantum computation can also be caught with a type system, with a few specificities that we discuss in this section.

6.3.1 Quantum Data and Type Linearity

The main problem with quantum information is that it is *non-duplicable* (a.k.a. *non clonable*, see Section 4.2.3). In terms of quantum programming language, this means that a program cannot duplicate a quantum bit: if U is a unitary map acting on two qubits, the function $\lambda x.U(x \otimes x)$ trying to feed U with two copies of its argument makes no sense. Similarly, it is not possible to control a gate acting on a qubit with the same qubit. The Quipper code of Figure 11 is, therefore, buggy.

Type systems, in a broad sense, provide a predicate that says that a well-typed program does not have a certain class of bugs: in the case of quantum programming languages, a large class of bugs comes from duplicating non-duplicable objects. This calls for a *linear* type system, enforcing at least the non-duplication of qubits. This has been taken into account in recent scalable implementations such as Silq [27].

On the theoretical side, type systems for quantum Lambda-Calculi and ProtoQuipper [162]—the formalization of Quipper—are typically based on linear logic. Originally designed by Girard [79] (as a continuation of [121]), linear logic assumes that formulas are linear—i.e. non-duplicable and non-erasable—by default, and the logic comes equipped with a logic constructor “!” to annotate duplicable

```

exp :: Circ Qubit
exp = do
  q1 <- qinit True
  q2 <- qinit True
  r <- qnot q1 'controlled' q1
  return r

```

Figure 11: An example of a buggy Quipper program

and erasable formulas. Linear logic also proposes a special pairing constructor \otimes replacing the usual product and compatible with both the linearity constructs and the (linear) implication \multimap .

A core type system for a quantum Lambda-Calculus with pairing therefore consists of the following grammar:

$$A, B ::= \text{qubit} \mid \text{bool} \mid A \multimap B \mid A \otimes B \mid !A.$$

Type $A \multimap B$ represents the type of (linear) functions, using their argument only once. Type $A \otimes B$ represents the pair of a term of type A and a term of type B . Type $!A$ stands for a duplicable term of type A . We give a few examples as follows.

- The identity function $\lambda x.x$ is of type $A \multimap A$, but also of type $!(A \multimap A)$ as it is duplicable (since it does not contain any non-duplicable object);
- If the pairing construct is represented with $\langle -, - \rangle$, the function $\lambda x.\langle x, x \rangle$ is of type $!A \multimap (A \otimes A)$: it asks for a duplicable argument;
- The operator `qinit` is of type $!(\text{bool} \multimap \text{qubit})$: it is duplicable but it does *not* generate a duplicable qubit;
- The operator `meas` can however be typed with $!(\text{qubit} \multimap \text{bool})$ as a boolean should be duplicable;
- Provided that U is a unitary acting on two qubits, one can type it in a functional manner with $\text{qubit} \otimes \text{qubit} \multimap \text{qubit} \otimes \text{qubit}$: it inputs two (non-duplicable) qubits and outputs the (still non-duplicable) modified qubits;
- In particular, provided that we assume implicit dereliction, casting duplicable elements of type $!A$ to A , the term $\lambda x.U \langle x, x \rangle$ can only be typed with $!\text{qubit} \multimap \text{qubit} \otimes \text{qubit}$: its argument has to be duplicable. The fact that this program can never be actually used on a concrete qubit is a property of the type system (intuitively, `qinit` only generates non-duplicable qubits).

6.3.2 Example: Quantum Teleportation

Mixing quantum computation and higher-order objects can yield non-trivial objects. The scheme of teleportation is given in Figure 12. It consists of three steps: (A) creation of a Bell state, (B) measure in the Bell basis to retrieve two booleans and (C) application of a gate $U_{b_1 b_2}$ dependent on the result of the measure. The state of the top wire is then “teleported” to the bottom wire. It is possible to understand the pieces of the quantum teleportation protocol as three (duplicable) functions:

- (A) $!(() \multimap \text{qubit} \otimes \text{qubit})$
- (B) $!(\text{qubit} \multimap (\text{qubit} \multimap \text{bool} \otimes \text{bool}))$
- (C) $!(\text{qubit} \multimap (\text{bool} \otimes \text{bool} \multimap \text{qubit}))$

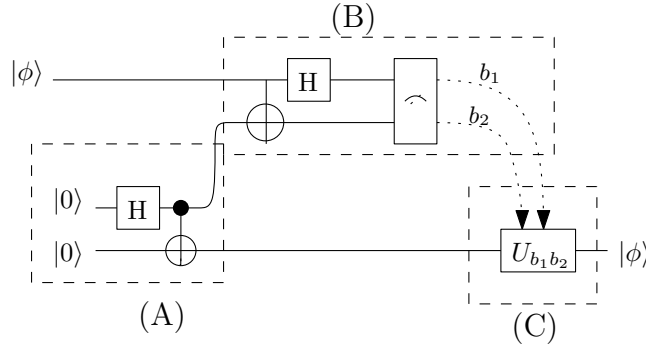


Figure 12: Teleportation algorithm

The parts (B) and (C) are duplicable functions producing two non-duplicable functions of type $\text{qubit} \multimap \text{bool} \otimes \text{bool}$ and $\text{bool} \otimes \text{bool} \multimap \text{qubit}$. The teleportation algorithm then feeds the two qubits emitted by (A) to (B) and (C); this gives a general type

$$!(\text{() } \multimap (\text{qubit} \multimap \text{bool} \otimes \text{bool}) \otimes (\text{bool} \otimes \text{bool} \multimap \text{qubit}))$$

for the protocol. It can be used several times (as it is duplicable). Each time it is run, it generates a pair of *non-duplicable* functions $\langle f, g \rangle$, and the specification of the protocol states that these two functions are inverse one of the other.

These two functions f and g of the pair are *non-duplicable*. Indeed, each of them holds a (non-duplicable) qubit coming from (A). Moreover, in a sense, these functions are entangled, since the Bell state from (A) is entangled.

6.3.3 Extending the Type System to Support Circuits

Let us assume that the type system of the quantum Lambda-Calculus is extended with lists: $[A]$ stands for lists of elements of type A (see e.g. [170] to see how to do this). A function $[\text{qubit}] \multimap [\text{qubit}]$ inputs a list of qubits. It can apply unitary gates to these qubit arguments: It can in fact describe different circuits, depending on the size of the list. Such a function, therefore, describes a *family* of circuits. The quantum Lambda-Calculus is not expressive enough to extract one circuit out of this family of circuits and operate on it (e.g. by inverting or controlling it).

The ProtoQuipper language [162] and its successors [160, 159, 77, 130] are formalized fragments of the programming language Quipper. They enforce structural properties of quantum programs using the linear type system of the quantum Lambda-Calculus, yet extending it to support circuit manipulation. ProtoQuipper comes with a new type construct $\text{Circ}(A, B)$: the type of circuits from A to B , and two functions:

- `box` sends $A \multimap B$ to $\text{Circ}(A, B)$. It takes a function $A \multimap B$, partially evaluates it, and stores the emitted circuit in an object of type $\text{Circ}(A, B)$.
- `unbox` sends $\text{Circ}(A, B)$ to $A \multimap B$. It takes a circuit from A to B and reads it as a function of input A and output B .

The behavior of `box` and `unbox` is specified by the operational semantic of the language [162].

One of the subtleties is the fact that `box` turns a function—possibly representing a family of many circuits—to *one* circuit. In the case of a function of type $[\text{qubit}] \multimap [\text{qubit}]$, this corresponds

to choosing one size of list and building the circuit for this input size. Whenever the type system supports inductive types such as lists, the operator box then also takes a *shape* as a second argument, for deciding on the shape of the circuit to build. In recent works [77], ProtoQuipper’s type system has been extended to very expressive *dependent types*, in order to characterize with a very fine-grain the shape structure of a family of circuits.

For instance, suppose that the program P sends `[qubit]` to `[qubit]`. It corresponds to a family of circuits, but if we pick a choice of input size n , the type gives no information on the output shape of the circuit—that is, the number of output wires. Maybe P duplicates each input wire? With dependent types, we can for instance index list-types with size and type P with

$$\forall n. [\text{qubit}]_n \multimap [\text{qubit}]_{2n}.$$

This type tells us that P corresponds to a family of circuits of even output wires. This makes it possible to catch errors when using circuit combinators: for instance, the inverse operator can be typed with

$$\forall n m. \text{Circ}([\text{qubit}]_m, [\text{qubit}]_n) \multimap \text{Circ}([\text{qubit}]_n, [\text{qubit}]_m)$$

The inverse of P then becomes a function of type

$$\forall n. [\text{qubit}]_{2n} \multimap [\text{qubit}]_n.$$

In particular, this function can only be applied to lists of even sizes. This run-time error cannot be checked without shape information.

Although such a type system becomes very expressive, in general, it fails to feature a type inference algorithm, as this would require to be able to solve arbitrary arithmetic equations.

6.3.4 Dependent types and Proofs of Programs

To build a dependently typed language, an alternative approach is to embed it inside an existing host language with this feature: QWIRE [148, 147] follows this route and relies on the language and proof-assistant Coq [25]. While Quipper uses Haskell’s monads to encode circuits, QWIRE capitalizes on Coq’s inductive types within the formalism of the Calculus of Inductive Constructions (CiC) [146], the logical framework of Coq.

In Haskell, inductive types are limited in expressivity: without dependent types, it is not possible to impose constraints on the content of a datastructure. As circuits in Quipper are internally made of elements of an inductive type, it makes it impossible to forbid ill-defined constructions such as re-using an erased wire, or using twice a wire on a controlled-not, as shown in Fig. 11.

QWIRE can instead rely on dependent, inductive types to enforce such constraints on circuit constructions: instead of simply considering circuits as lists of gates, a circuit in QWIRE consists in a list of gates together with proofs that the gates are added in a sensible way—in other words, the constructors of the inductive type of circuits in QWIRE corresponds to a set of *typing rules* for writing valid circuits.

Interestingly enough, QWIRE does not have to rely on “!” type constructor to distinguish between duplicable and non-duplicable data. The idea is that instead of working in a linear-logic based type-system, QWIRE considers an equivalent linear-non-linear model [23]. In this paradigm, there are two intertwined languages:

- A linear language aimed at qubit manipulation and gate application: a program in this language is a circuit. The linearity of the type system enforces the necessary constraints so that e.g. Fig. 11 is indeed invalid;

- A high-level, regular language—typically a lambda-calculus, with a regular type system. This language represents the “usual” programming paradigm where classical, conventional computation happens;
- Then there are two operations, akin to `box` and `unbox`, to move from the linear language to the classical language.

In the context of QWIRE, the classical, regular language is Coq and the linear language is encoded using the inductive type of circuits. As Coq features dependent types, QWIRE can then be regarded as a dependently-typed quantum programming language. However, due to the sophisticated type system, QWIRE does not feature a type inference algorithm.

6.3.5 Discussion

Type systems for quantum programming languages provide very efficient ways to encode and—whenever featuring a type inference algorithm—*automatically* verify some important properties of programs, and in particular to rule out at compile-time large classes of run-time errors specific to quantum computation. In particular, type systems have been used to characterize and enforce

- structure of parametric families of circuits;
- linearity of non-duplicable elements;
- control and inversion of only purely quantum circuits.

However, to be able to go further and characterize functional correctness with respect to specification, or validate the number of gates of a circuit, or catch subtle bugs involving concatenation of inverted circuits, one needs to move away from the simple linear type systems of quantum lambda-calculi and shift towards sophisticated dependent types, such as the extension ProtoQuipper-D [77] of ProtoQuipper, or the approach of QWIRE. The gain in expressiveness is then at the expense of automation [177, 148].

The quest for a finer trade-off, permitting automation while capturing some of what is currently only available with dependent type system is an active research area in the community.

7 High and Mid-Level Verification: Algorithms and Programs

Most *quantum programming languages* (QISKIT [153] Quipper [83], Liqui|) [193], Q# [181], PROJECTQ [180], SILQ [27], *etc*) embed features for quantum circuit manipulations within a standard classical programming language. Such *circuit-building quantum languages* is the current consensus for high-level executable quantum programming languages. A current major challenge is to link this language design paradigm with formally verified programming. In the present section, we introduce the main existing propositions in that direction.

7.1 Quantum Hoare Logic

Quantum Hoare logic (QHL) [37, 66, 105, 194, 185, 184, 19] is a general framework for reasoning about the classical control instructions over unitary operations in quantum algorithms. Referring again to Figure 1, the focus is on the interaction between the classical computer and the quantum co-processor, instead of the circuit analysis as in Section 5, or the gate to gate circuit building

functions as in Sections 7.2 and 7.3. Therefore, we consider it as a high-level description of quantum algorithms.

It is based on the assertion method of Floyd and Hoare [74, 94]—attach each program point with an assertion and whenever the data flow reaches a program point the attached assertion should be satisfied—which was originated with Alan Turing [8]. Hoare’s approach enables (interactive) theorem proving for high-level algorithmic description verification that proceeds at the same abstraction level as the language itself. This makes verification more human-friendly than lower-level (machine-friendly) verification.

7.1.1 Quantum Programming Language: Quantum WHILE-Programs

The guideline for the hybrid model introduced in Section 2.1.1 is summed up by the slogan “*quantum data and classical control*” [168]: quantum data can be superposed and entangled, they are manipulated by basic quantum operations—unitary evolution and measurement, but the high-level control is still classical (e.g. branch, loops, *etc.*).

In light of this slogan, QHL introduces a minimal programming language for describing quantum algorithms [195]. We follow [194] for the introduction of QHL technical environment. Let q (resp. \bar{q}) be a quantum variable (resp. a list of quantum variables); let U be a unitary operator acting on the qubits \bar{q} and let $M \triangleq \{M_m\}_m$ with $\sum_m M_m^\dagger M_m = I$ be a measurement on the qubits \bar{q} , each M_m corresponding to a measurement result m (see Section 2.1.5). As a special case, let $M' \triangleq \{M_0, M_1\}$ with $M_0^\dagger M_0 + M_1^\dagger M_1 = I$. Then quantum WHILE-programs are generated by the following syntax.

$S \triangleq$	skip	No operation
	$q := 0\rangle$	Initialization
	$\bar{q} *= U$	Unitary operation
	$S_1; S_2$	Sequential composition
	if $\square m \cdot M[\bar{q}] = m \rightarrow S_m$ fi	Probabilistic branching
	while $M'[\bar{q}] = 1$ do S_0 od	Probabilistic while loop

The intended semantics of language constructs above is similar to that of their classical counterparts. To illustrate the quantum features contained in these constructs, we make the following comments:

- (i) in the initialization, the choice of a fixed state $|0\rangle$ is without loss of generality, since any known quantum state can be prepared by applying a unitary operator to $|0\rangle$;
- (ii) according to Born rule, measurement results follow a probabilistic law. Since they lie on measurement result observations, branching (resp. while loops) is therefore probabilistic. It creates different branches $\{S_m\}_m$ (resp. $\{\mathbf{skip}, S_0\}$), chosen according to outcomes of the measurement M (resp. M') on the qubits \bar{q} .

We refer the reader to [194] for a detailed exposition of the syntax above.

Example 7.1 (Preparation of the Bell state $|\beta_{00}\rangle$). *Let p and q be quantum variables, each denoting one qubit. Then the following program initiates them to $|0\rangle$ and implements the circuit from Figure 2, preparing state $|\beta_{00}\rangle$ from Example 2.1.:*

$$\beta_{00} \triangleq p := |0\rangle; q := |0\rangle; p *= H; (p, q) *= CNOT$$

Note that the quantum programming language defined above is in the spirit of the hybrid circuit model presented in Section 2. Indeed, the basic sequence of quantum operations (initialization, unitary operation, and measurement) are meant to be interpreted as a generalized quantum circuit to be executed on a quantum co-processor; and post-measurement branchings (in, e.g., probabilistic branching and while loop) are meant to be controlled by a classical computer.

7.1.2 Quantum States, Operations and Predicates

Measuring a quantum state transforms it, following the Born rule (see Section 2.1.2). The resulting probability distribution over quantum states is formalized as a *mixed state* (as opposed to *pure states*, see Section 2.1.5). For example, a measurement on any pure quantum state $|+\rangle \triangleq \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ or $|-\rangle \triangleq \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ will result in the mixed state $\mathbb{E} = \{(\frac{1}{2}, |0\rangle), (\frac{1}{2}, |1\rangle)\}$, with states $|0\rangle$ and $|1\rangle$ occurring with an equal probability of $\frac{1}{2}$ (notice that this observation makes both states $|+\rangle$ $|-\rangle$ impossible to distinguish by simple measurement).

In this way, the representation of the final state of applying a series of measurements to a quantum state could expand exponentially. To address this issue, a square-matrix representation of quantum states, i.e. partial density operator, is adopted instead. For example, pure quantum state $|+\rangle$ is represented as $|+\rangle \langle +|$, and mixed quantum state \mathbb{E} as $\frac{1}{2}|0\rangle \langle 0| + \frac{1}{2}|1\rangle \langle 1|$. See Section 2.1.5 for a brief introduction to partial density operators, and [168, 195, 139] for further details.

If we see the matrix representation of a quantum state (partial density operator) as a linear operator, then a quantum operation—initialization, unitary evolution and measurement—can be thought of as a super operator, i.e. a function from linear operators to linear operators. What's interesting is that every quantum WHILE-program defined above can be interpreted as a quantum operation, and partial density operators are closed under quantum operations. This justifies the success of representing quantum states as partial density operators and defining the denotational semantics of quantum programs as quantum operations [168, 195].

Following [56], a quantum predicate on vector space H is defined as a Hermitian operator M between the zero operator 0_H (representing the contradiction) and the identity operator I_H (representing the tautology). Instead of the usual binary satisfaction judgment, QHL evaluates the satisfaction of a predicate by a state as a real value between 0 (false) and 1 (true). It is defined as the trace $tr(M\rho)$ of the product $M\rho$. Intuitively, it represents the expectation for the truth value of M in the mixed state ρ (which is, again, a probability distribution over pure states).

Then, the intuition of implication between predicates is also probabilistic. It is filled by the *Löwner order* $M \sqsubseteq N$, relating operators M and N , if and only if, for any state ρ , the expectation truth value of N in ρ is more or equal to that of M in ρ . This condition is formalized as $tr(M\rho) \leq tr(N\rho)$ for all states ρ (See, e.g., [194, Lemma 2.1]).

Adopting such quantum predicates as assertions, among many others (e.g., interpreted as physical observables), provides simple expression means for many properties of quantum effects. For example, quantum predicate $|+\rangle \langle +|$ expresses that a state ρ is in the equal superposition $|+\rangle$ with probability $tr(|+\rangle \langle +|\rho)$; quantum predicate $|\beta_{00}\rangle \langle \beta_{00}|$ expresses that a state ρ is in the maximal entanglement $|\beta_{00}\rangle$ with probability $tr(|\beta_{00}\rangle \langle \beta_{00}|\rho)$, etc.

7.1.3 Quantum Program Verification

For now, a quantum (partial) correctness formula can be the Hoare's triple $\{P\} S \{Q\}$, where S is a quantum WHILE-program, and P, Q are quantum predicates. To define the partial-correctness semantics of quantum Hoare's triples, in the sequel, let $\llbracket S \rrbracket$ denote the semantic function of S (Note that $\llbracket S \rrbracket$ is a quantum operation defined by induction on S , cf. [195]), and $\llbracket S \rrbracket(\rho)$ the output of S on the input ρ .

Definition 7.1 (Semantics of partial correctness, cf. [194]). *Let P, Q be quantum predicates and S a quantum WHILE-program. We say that S is (partially) correct w.r.t. precondition P and postcondition Q , written $\models \{P\} S \{Q\}$, if*

$$\forall \rho, \quad tr(P\rho) \leq tr(Q\llbracket S \rrbracket(\rho)) + [tr(\rho) - tr(\llbracket S \rrbracket(\rho))]. \quad (4)$$

Note that Inequality (4) can be seen as a probabilistic version of the following statement: if state ρ satisfies predicate P , then, executing program S on input ρ , either S fails to terminate or the resulting state $\llbracket S \rrbracket(\rho)$ satisfies predicate Q .

Table 2: Proof system for partial correctness.

(Skip Axiom)	$\{P\} \text{ skip } \{P\}$
(Init Axiom)	$\{\sum_i i\rangle_q \langle 0 P 0\rangle_q \langle i \} q := 0\rangle \{P\}$
(Unit Axiom)	$\{U^\dagger P U\} \bar{q} * = U \{P\}$
(Comp Rule)	$\frac{\{P\} S_1 \{Q\} \quad \{Q\} S_2 \{R\}}{\{P\} S_1; S_2 \{R\}}$
(If Rule)	$\frac{\{P_m\} S_m \{Q\} \text{ for all } m}{\{\sum_m M_m^\dagger P_m M_m\} \text{ if } \square m \cdot M[\bar{q}] = m \rightarrow S_m \text{ fi } \{Q\}}$
(Par Loop Rule)	$\frac{\{P\} S_0 \{M_0^\dagger Q M_0 + M_1^\dagger P M_1\}}{\{M_0^\dagger Q M_0 + M_1^\dagger P M_1\} \text{ while } M'[\bar{q}] = 1 \text{ do } S_0 \text{ od } \{Q\}}$
(Order Rule)	$\frac{P \sqsubseteq P' \quad \{P'\} S \{Q'\} \quad Q' \sqsubseteq Q}{\{P\} S \{Q\}}$

The axiom system for proving partial correctness of quantum WHILE-programs is composed of axioms and inference rules manipulating quantum Hoare's triples [194]. It is shown in Table 2 (where $\{|i\rangle_q\}$ is the computational basis for quantum variable q). Remark that each of these rules and axioms follows the assertion method. Here we only show how to derive the most complex rule (Par Loop Rule). The derivation of other proof rules can be done similarly.

Intuition of (Par Loop Rule). To derive (Par Loop Rule), by assertion method, we attach each program point, say l_1, l_2, l_3 , of a WHILE-statement with an assertion, say R, P, Q , respectively:

$$\{l_1 : R\} \text{ while } M'[\bar{q}] = 1 \text{ do } \{l_2 : P\} S \text{ od } \{l_3 : Q\}$$

Fix the input ρ at the program point l_1 satisfying the assertion R . By semantics of a WHILE loop, after the measurement M' , one part $M_1 \rho M_1^\dagger$ of the input will go to the loop body through the program point l_2 where the assertion P will be satisfied; the other part $M_0 \rho M_0^\dagger$ will leave the while loop through the program point l_3 in which the assertion Q will be satisfied. Hence:

$$\text{tr}(R\rho) \leq \text{tr}(Q(M_0 \rho M_0^\dagger)) + \text{tr}(P(M_1 \rho M_1^\dagger))$$

Due to the arbitrariness of ρ , by properties of the trace function and Löwner order, we have that $R \sqsubseteq M_0^\dagger Q M_0 + M_1^\dagger P M_1$. Then, by weakening R to $M_0^\dagger Q M_0 + M_1^\dagger P M_1$ and lifting the above reasoning process into an inference rule, (Par Loop Rule) follows.

The following example illustrates how to derive a partially correct quantum Hoare triple using the axioms and inference rules presented above.

Example 7.2 (Specification and correctness proof for the $|\beta_{00}\rangle$ state construction program). *Recall from Example 7.1 the definition of quantum program β_{00} :*

$$\beta_{00} \triangleq p := |0\rangle; q := |0\rangle; p * = H; (p, q) * = CNOT$$

To show (partial) correctness of this program, it suffices to prove

$$\{I_p \otimes I_q\} \beta_{00} \{|\beta_{00}\rangle_{p,q} \langle\beta_{00}|\} \quad (5)$$

This can be done as follows. By sake of space we need to decompose the derivation tree. We first derive specifications for the initialization instructions:

$$\frac{\frac{\text{Init Axiom}}{\{I_p \otimes I_q\} p := |0\rangle \{ |0\rangle_p \langle 0| \otimes I_q \}} \quad \frac{\text{Init Axiom}}{\{ |0\rangle_p \langle 0| \otimes I_q \} q := |0\rangle \{ |0\rangle_p \langle 0| \otimes |0\rangle_q \langle 0| \}}}{\{I_p \otimes I_q\} p := |0\rangle; q := |0\rangle \{ |0\rangle_p \langle 0| \otimes |0\rangle_q \langle 0| \}} \text{Comp Rule} \quad (i)$$

Then come the unitary application instructions. We set the following abbreviations for Hoare triples:

$$\begin{aligned} (ii) & := \{ |0\rangle_p \langle 0| \otimes |0\rangle_q \langle 0| \} p * = H \{ |+\rangle_p \langle +| \otimes |0\rangle_q \langle 0| \} \\ (iii) & := \{ |+\rangle_p \langle +| \otimes |0\rangle_q \langle 0| \} (p, q) * = CNOT \{ |\beta_{00}\rangle_{p,q} \langle \beta_{00}| \} \end{aligned}$$

They are instances of axioms and we can combine them:

$$\frac{\frac{\text{Init Axiom}}{(ii)} \quad \frac{\text{Init Axiom}}{(iii)}}{\{ |0\rangle_p \langle 0| \otimes |0\rangle_q \langle 0| \} p * = H; (p, q) * = CNOT \{ |\beta_{00}\rangle_{p,q} \langle \beta_{00}| \}} \text{Comp Rule} \quad (iv)$$

Finally, the two preceding proof tree branch together via the sequential composition rule, which achieves the derivation:

$$\frac{(i) \quad (iv)}{\{I_p \otimes I_q\} \beta_{00} \{ |\beta_{00}\rangle_{p,q} \langle \beta_{00}| \}} \text{Comp Rule}$$

7.1.4 Implementations and Extensions

Several works have taken advantage of extended Quantum Hoare Logic, e.g. algorithmic analysis of termination problem [127] or characterization and generation of loop invariants (i.e. $M_0^\dagger Q M_0 + M_1^\dagger P M_1$ in (Par Loop Rule)) [198].

The practical illustration of QHL can be found in Liu et al. paper [132], containing an implementation in ISABELLE/HOL together with a formalization of Grover [84] and Quantum Phase Estimation (QPE [111]) algorithms. Nevertheless in these examples, the central verification part is assumed through Python libraries uses.

More recent work [131] includes full proof for a parametrized version of Grover's search algorithm. It constitutes an illustration of QHL use on a non-trivial example.

7.1.5 Other Quantum Hoare Logics

In addition to the framework introduced above, applying Hoare Logic methods to quantum processes brought several additional developments, focusing on different aspects of quantum computations. We introduce a few of these QHL related framework in the following paragraphs.

Quantum Hoare Logic with ghost variables. One of the principal shortcuts of [194] comes from the limitations of the specification language. In [184], the author in particular targets the possibility to characterize probabilistic distributions of values. His proposition introduces *ghost* variables in the specifications language: a *ghost* variable does not occur in the program but only in its specifications. In [184], a ghost variable is interpreted under an implicit existential quantification. Ghost variables enable, in particular, to explicitly refer to discarded, measured or overwritten qubits. In addition to probability distribution definitions, it brings several features to the expressive power of the specifications language, such as separability–unentanglement–of variables or the fact, for a given variable, to hold for a classical data value.

Quantum Hoare Logic with classical variables. With similar concern, [67] extends QHL specifications with classical variables. In the specification and verification of algorithms of practical use, holding classical information is indeed crucial. It enables, for example, describing and specifying an algorithm holding classical parameters or a hybrid program intertwining classical and quantum instructions (see Section 2.1.1).

Hence, the semantics relies on so-called *cq-states* (classical/quantum states) made of both:

- a classical variable assignment for the interpretation of classical variables;
- a density operator interpreting quantum variables.

This extension of the semantics induces an extension of the proof system, which is proved sound and complete—with respect to partial correctness. Interestingly, the paper presents detailed specified case studies, including Grover’s and Shor’s algorithms.

Robustness analysis is another important line of work, initiated in [97] as a continuation of QHL. The different solutions presented so far rely on the implicit assumption that quantum gates are applied deterministically as indicated by their matrix semantics. Still, as stated in Section 2.3.2, this noise-free modeling is not perfectly accurate in the NISQ era.

A more realistic description of the behavior of a gate would consider several different possible behaviors, weighed with their respective probabilities of occurrence: the *intended* one, expected by the semantics, and one or several additional erroneous behaviors¹⁰.

Interestingly, in this setting, the application of a quantum gate is formally represented as a probability distribution over unitaries. Hence, a non-deterministic gate application naturally formalizes in the formalism of super operators acting over density operators.

Then, a metrics is defined for measuring the difference between the behavior of a quantum system under a given error scheme (for each gate application, the mention of a possible erroneous behavior together with its probability of occurrence) and its intended behavior. It is called the *trace distance* and serves as an evaluation for the robustness of the implementation.

Among other case studies, the method is illustrated by evaluating a minimal quantum error correcting scheme.

Quantum separation logic [201] aims at simplifying quantum programs specifications. The leading observation is that, while quantum programs often manipulate big quantum circuits, with matrix semantics growing exponentially over their width, many quantum algorithms proceed via

¹⁰Strictly speaking, the language in 2.3.2 is limited to considering up to one possible error per gate application. But this is without loss of generality since more sophisticated scenarios can be encoded through, eg., the replacement of gate application intervals by potentially erroneous identity transformations.

sequences of local manipulations over quantum sub-registers. The authors exhibit the examples of Quantum Machine Learning [26] and Variational Quantum Algorithms (VQA [135]), which are among the most promising classes of algorithms in the NISQ era (see Section 2.3.2).

To efficiently reason about state evolution in such implementations, a Quantum Separation Logic is proposed, together with a dedicated proof system. This logic allows the expression of local manipulation on separated quantum registers while maintaining the state of the rest of the register. The approach is probed with case studies from both quantum programming analysis (VQA) and communication protocols security checks (one-time pad and secret sharing).

Quantum relational Hoare logic [185, 19] allows to reason about how the outputs of two quantum programs relate to each other given a relation between their inputs, which can be used to analyze security of post-quantum cryptography and quantum protocols.

Quantum Hoare type theory [177] is inspired by classical Hoare type theory and extends the Quantum IO Monad [3] by indexing it with pre- and postconditions that serve as program specifications, which has the potential to be a unified system for programming, specifying, and reasoning about quantum programs.

Quantum dynamic logic. We end up this section by mentioning a line of work that is not a strict extension or application of Hoare Logic [94], but shares similar concerns and related solutions. Just as QHL and its extensions, dynamic logic formalizes the evolution of a state along the execution of a process acting over it.

Dynamic logic inherits from the modal logic apparatus where, in addition to propositional logic connectives, modalities \Box and \Diamond are introduced. Intuitively, given a formula ϕ , $\Box\phi$ means that ϕ is *necessarily* true and $\Diamond\phi$ means that ϕ is *possibly* true. The standard semantics, based on Kripke models [115], helps catching this intuition. A Kripke model \mathcal{K} is made of a set of states S holding, each, a valuation $P : S \rightarrow \{\mathbf{true}, \mathbf{false}\}$ for a set of propositional variables P . And the modality is interpreted through an accessibility relation $R \subseteq S \times S$. For example, given a propositional variable p , $\Box p$ (resp. $\Diamond p$) is true in state s , written $\mathcal{K}, s \models \Box p$ (resp. $\mathcal{K}, s \models \Diamond p$) iff p (resp. $\mathcal{K}, s \models \Diamond p$) is true in every (resp. at least one) state s' such that $(s, s') \in R$.

In dynamic logic, several modalities coexist, formalizing a set of *actions*. Hence, given an action a , $\text{formula}[a]\phi$ is true iff ϕ is always true after performing action a . Actions can combine together, for example $[a;a']\phi$ means that ϕ is always true after the successive performance of actions a and a' .

In their *Logic of Quantum Programs* (LQP [14, 15]), Baltag and Smets designed a quantum version of dynamic logic. Here, states correspond to one-dimensional subspaces of a Hilbert space for a quantum register. Actions are of two kinds: a *test* encodes a measurement through the set of projection transitions corresponding to the different possible outcomes (for example, given formulas ϕ and ψ , $\phi?\psi$ is satisfied in the state of evaluation if and only if any state satisfying ϕ after measurement—any successful test for ϕ —also satisfies ψ), and proper *actions* deterministically encode unitary state transformations). In [14], the framework is illustrated by a correction proof for the teleportation and quantum secret sharing protocols. Baltag and Smets link their work to similar quantum logics as in [34].

Note that to treat measurement, LQP formalizes informations about the set of states that are possibly reachable, but the language does not hold any notion of probability. This shortcut is overcome in [13], where the authors introduce a probabilistic test modality characterizing, for any formula ϕ and rational $x \in [0, 1]$, that a test for ϕ has probability at least x to succeed from the evaluation state.

In [13], the authors also provide a decidability proof for their setting and illustrate its expressive power by formally specifying the Grover’s search algorithm and a leader election protocol. Further work in the community has used this probabilistic logic setting for the verification of the BB84 Quantum Key Distribution algorithm [24].

7.2 QBRICKS

QBRICKS [38] is a recently proposed circuit description language together with a deductive verification framework. It enables automated proof support for program specifications, reducing the required human effort for the development of verified programs.

QBRICKS object language (QBRICKS-DSL) consists in a minimal functional language with features for the design of circuit families. Similarly to the formal contract style of algorithm descriptions (see Section 2.2.2), QBRICKS functions are written with explicit pre- and postconditions, specifying their complexity and the parametrized input-output quantum data registers function they implement. These specifications are written in a dedicated formal language, called QBRICKS-SPEC.

To support proofs, QBRICKS is given a Hoare style derivation rules system including rules for each parametrized circuit constructor. These rules are enriched with equational theories enabling, in particular, reasoning about measurement and probabilities.

QBRICKS is a domain-specific language, embedded in the Why3 [70] deductive verification framework : programs are written in ML language and annotated with specifications in QBRICKS-DSL (pre- and postconditions, loop invariants, calls for lemmas, *etc.*). Compiling a QBRICKS program interprets these specifications as proof obligations. Then, a dedicated interface enables to directly access these proof obligations and either send them to a set of automatic SMT-solvers (CVC4, Alt-Ergo, Z3, *etc.*), or enter some interactive proof transformation commands (additional calls for lemmas or hypotheses, term substitutions, *etc.*) or even to proof assistants (Coq, ISABELLE/HOL).

7.2.1 Writing Quantum Circuits Functions in QBRICKS: QBRICKS-DSL

QBRICKS-DSL makes use of a regular inductive datatype for circuits, where the data constructors are elementary gates, sequential and parallel composition, and ancilla creation. In particular, unlike in e.g. Quipper or QWIRE, a quantum circuit in QBRICKS is not a function acting on qubits: it is a simple, static object. Nonetheless, for the sake of implementing quantum circuits from the literature, this does not restrict expressiveness as they are usually precisely represented as sequences of blocks.

The core of QBRICKS-DSL is presented in Figure 13. It is a small first-order functional, call-by-value language. To the elementary gates presented in Section 2.1.5, QBRICKS adds the qubit swapping gate SWAP and the identity ID. The constructors for high-level circuit operations are sequential composition SEQ, parallel composition PAR and ancilla creation/termination ANC.

Expression	e	$::=$	$x \mid c \mid f(e_1, \dots, e_n) \mid \text{let } \langle x_1, \dots, x_n \rangle = e \text{ in } e' \mid$ $\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \text{iter } f \ e_1 \ e_2$
Data Constructor	c	$::=$	$\underline{n} \mid \text{tt} \mid \text{ff} \mid \langle e_1, \dots, e_n \rangle \mid \text{CNOT} \mid \text{SWAP} \mid \text{ID} \mid \text{H} \mid \text{Ph}(e) \mid \text{R}_z(e) \mid$ $\text{ANC}(e) \mid \text{SEQ}(e_1, e_2) \mid \text{PAR}(e_1, e_2)$
Function	f	$::=$	$f_d \mid f_c$
Declaration	d	$::=$	$\text{let } f_d(x_1, \dots, x_n) = e$

Figure 13: The syntax for QBRICKS-DSL

The term constructs are limited to function calls, `let`-style composition, `test` with the ternary construct `if-then-else` and simple iteration: `iter f n a` stands for $f(f(\dots f(a)\dots))$, a succession of n calls to f .

Even though the language does not feature measurement, it is nonetheless possible to *reason* on probabilistic outputs of circuits, if we were to measure its output. This is expressed in a regular theory of real and complex numbers in the specification language (see Section 7.2.4 below for details).

7.2.2 Parametrized Path-Sums

To interpret circuit description functions, QBRICKS uses parametrized path-sums (pps), that is an extension of path-sums [5] (Section 5.2).

In QBRICKS setting, a path-sum P is an object from an opaque type with four parameters, presented in Table 3 with their types and identifier shortcuts: two integer constants `pps_width(P)` (the width of the target circuit) and `pps_range(P)` (the *range*, meaning that the output sum of kets has a term for each bit vector \vec{y} of length `pps_range(P)`—written $\vec{y} \in \text{BV}_{p,r(P)}$) and two functions `pps_angle(P)` and `pps_ket(P)`: for any input bit vector \vec{x} of length `pps_width(p)` (standing for a basis ket input to the target circuit) and for any index bit vector \vec{y} of length `pps_range(p)`, functions `pps_angle(P)` and `pps_ket(P)` respectively define a real scalar and a bit vector of length `pps_width(p)` (standing for a basis ket output to the target circuit)¹¹.

Table 3: Pps accessors and types

Identifier	Type	Id-abbreviation
<code>pps_width</code>	<code>int</code>	<code>p-w</code>
<code>pps_range</code>	<code>int</code>	<code>p-r</code>
<code>pps_angle</code>	<code>bit_vector → bit_vector → real</code>	<code>p-a</code>
<code>pps_ket</code>	<code>bit_vector → bit_vector → bit_vector</code>	<code>p-k</code>

Then for any bit vector \vec{x} of size `p-w(P)`, the expression

$$Ps(h, |\vec{x}\rangle) = \frac{1}{\sqrt{2^{p_r(P)}}} \sum_{\vec{y} \in \text{BV}_{p,r(P)}} e^{2\pi i \cdot \text{p-a}(P)(\vec{x}, \vec{y})} |\text{p-k}(P)(\vec{x}, \vec{y})\rangle_{p-w(P)} \quad (6)$$

combines these different element together to define a linear application for quantum state vectors. Function Ps is extended by linearity to any ket $|u\rangle$ of length `p-w(p)`¹².

A path-sum P is said to *correctly interpret a given circuit* C (written $(C \triangleright P)$) if and only if C has width $n = \text{pps_width}(P)$ and for any bit vector \vec{x} of length n , $\mathbf{Mat}(C) \cdot |\vec{x}\rangle = Ps(P, |\vec{x}\rangle)$. The relation $(_ \triangleright _)$ enjoys nice composition properties along QBRICKS-DSL circuit constructors (see [38]).

QBRICKS-SPEC generalizes path-sums by introducing Parametrized path-sums (pps). A pps is a function that inputs a set of parameters and outputs a path-sum. Then, it can be seen as a family of path-sums (one for each possible value of its parameters) describing the effects of the different members in a family of quantum circuits. Hence, it is well-fitted for the specification of parametrized algorithms such as Shor order finding (Shor-OF see Figure 4).

¹¹In the rest of this section, type `bit_vector` corresponds to bit vectors \vec{x} . They are encoded in an abstract type that provides a positive integer `length(x)` (the length of the vector) and a value function `get_bv(x)`: `int → int`. For any integer i , we commonly abbreviate `get_bv(x)i` as \vec{x}_i ; if $i \in \llbracket 0, \text{length}(\vec{x}) \rrbracket$ (that is, if i is actually in the range of the bit vector) then it is such that $0 \leq \vec{x}_i < 2$.

¹²That is, to any linear combination of basis kets $|\vec{x}\rangle$ of length `p-w(p)`.

The main strength of pps semantics, with regards to formal verification, is that each of the path-sum parametrized accessors (see Table 3) combines compositionally along with circuit constructors.

Hence, it enables reasoning about parametrized quantum circuits and their semantics without manipulating sum terms or other higher-order objects. Thanks to this tool, the automatic generation of proof obligations for QBRICKS specifications results in only first-order formulas, enabling a high level of automation when sent to SMT-solvers.

7.2.3 From Quantum Circuits to Path-Sums

The specification language for QBRICKS is a first-order predicate language, equipped with various equational theories. For any quantum circuit family C , QBRICKS-DSL enables to identify a pps $\text{circ_to_pps}(C)$. Each of its parametrized accessors is defined inductively upon the structure of C and it is proved that $(C \triangleright \text{circ_to_pps}(C))$, for any instance of circuit. These accessors are listed and given abbreviations in Table 4.

Table 4: Function `circ_to_pps` and accessors

Accessor	Abbreviation
<code>pps_width(circ_to_pps(C))</code>	<code>C_width(C)</code>
<code>pps_range(circ_to_pps(C))</code>	<code>C_range(C)</code>
<code>pps_angle(circ_to_pps(C))</code>	<code>C_angle(C, -, -)</code>
<code>pps_ket(circ_to_pps(C))</code>	<code>C_ket(C, -, -)</code>

7.2.4 Probabilistic Reasoning

QBRICKS-DSL does not contain any constructor for the measurement of quantum registers. Nevertheless, QBRICKS-SPEC provides reasoning tools about it. In particular, function

$$\text{proba_measure} : \text{circ} \times \text{ket} \times \text{int} \rightarrow \text{real}$$

inputs a circuit C , a quantum data register $|v\rangle_n$ and an index $j \in \llbracket 0, 2^n \rrbracket$. It outputs the probability, for one measuring the quantum register resulting from applying circuit C to $|v\rangle$, to get the bit vector representing integer j as a result. Function `proba_measure` is defined, for any input ket $|u\rangle$ of length `pps_width(C)`, by application of the Born rule (see Section 2.1.2), as

$$\text{proba_measure}(C, |i\rangle_n, j) = |(Ps(\text{circ_to_pps}, |i\rangle_n)(j))|^2.$$

QBRICKS-SPEC also provides similar reasoning ghost functions for discussing the effect of partial measurement over quantum sub-registers.

7.2.5 Verified Properties

The QBRICKS framework aims at providing tools for writing and verifying the standard format of quantum algorithm specifications as they appear in algorithms (see, e.g., Figure 4): to perform a given computation task with a given amount of resources. Hence, as illustrated above, QBRICKS-SPEC is designed for the formalization of both:

- parametrized input/output relations for families of circuits. For a family of circuits, they typically consist in characterizing their parametrized output ket vector. Thanks to function

proba_measure, QBRICKS enables to identify the probability to get a given result after measurement and derivated probability reasoning (such as bounding the parametrized probability of success of a computation, wrt a pre-defined success condition). Proof support for these specifications is processed through the pps formalism,

- complexity specifications: QBRICKS enables to specify the parametrized width, number of required ancilla qubits and number of elementary gates of a circuit family.

Example 7.3 (Pps specifications for the Bell generating circuit). *For example, a specification for the Bell generating circuit can be written using accessors of pps circ_to_pps (Bell-circuit), as below*

$$\Gamma, \vec{x}, \vec{y} : \text{bit_vector}, j : \text{int} \vdash$$

$$\left\{ \begin{array}{l} \text{bv_length}(\vec{x}) = 2 \quad \wedge \quad \text{bv_length}(\vec{y}) = 1 \quad \wedge \quad j \in \llbracket 0, 2 \rrbracket \\ \text{Bell-circuit} \\ \left\{ \begin{array}{l} \text{C_width}(\text{result}) = 2 \quad \wedge \quad \text{C_ket}(\text{result}, \vec{x}, \vec{y}) = \overline{x(0) \cdot (1 - x(1))} \quad \wedge \\ \text{C_range}(\text{result}) = 1 \quad \wedge \quad \text{C_angle}(\text{result}, \vec{x}, \vec{y}) = x(0) * y(0) \quad \wedge \\ \text{width}(\text{result}) = 2 \quad \wedge \quad \text{size}(\text{result}) = 2 \quad \wedge \quad \text{ancillas}(\text{result}) = 0 \end{array} \right\} \end{array} \right\}$$

This Hoare style notation, $\{Pre\}p\{Post\}$ states that whenever Pre is satisfied, then running p ensures that $Post$ is satisfied. Formula $Post$ uses result as a variable standing for program p . The specification uses free-ghost-bit vectors variables \vec{x} and \vec{y} and integer variable j . It requires \vec{x} and \vec{y} to have respective lengths 2 and 1 and j to be in $\llbracket 0, 2 \rrbracket$. Given these preconditions, the specifications ensures that the angle C_angle outputs $\vec{x}_0 * \vec{y}_0$ for inputs \vec{x} and \vec{y} and the ket function C_ket outputs $\vec{x}_0 \cdot (1 - \vec{x}_1)$ for inputs \vec{x}, \vec{y} and j . Then, one easily verifies that, applying equation 6 on each bit-vector $a \cdot b$ of length 2 results in the corresponding output $|\beta_{ab}\rangle$, formally:

$$Ps(\text{circ_to_pps}(\text{Bell-circuit}), |ab\rangle) = \mathbf{Mat}(\text{Bell-circuit}) \cdot |ab\rangle$$

So the postcondition, by use of C_width , C_range , C_angle and C_ket , enables a complete characterization of the input/output function performed by the Bell circuit. In addition, the specification brings some complexity related postcondition: the circuit has width 2 and size 2 (*i.e.*, length of the required quantum register and number of performed elementary operations), and does not use any additional ancilla qubit.

7.2.6 Deduction and Proof Support

Proof support in QBRICKS strongly relies on the compositional structure of quantum circuits, enabling compositional reasoning on both pps and complexity features. As an example, in Figure 14 we give some of the rules used for the characterization of circuits size. Gates ID and SWAP are considered as free, so they count for null. The other elementary gates have size 1, both sequence and parallel compositions sum the size of their components and ancilla creation/termination does not affect circuit size.

Similar deduction rules are defined for circuit width, number of ancilla qubits and pps accessors. We do not introduce them here for sake of concision but we refer the desirous reader to [38].

7.2.7 Implementation and Case Studies

QBRICKS enabled to implement, specify and formalize parametrized versions for Quantum Fourier Transform, Grover search algorithm, QPE, Shor order finding, *etc.* The main specificity of these QBRICKS implementations is to hold a high level of proof support automation. This is largely due to the reduction of proof obligations to first-order logic predicates by use of pps characterizations.

$$\begin{array}{c}
\frac{C \in \{\text{ID}, \text{SWAP}\}}{\text{width}(C) = 0} \text{ (Id-SWAP-size)} \qquad \frac{C \in \{\text{H}, \text{Ph}(n), \text{R}_z(n)\}}{\text{size}(C) = 1} \text{ (H-Ph-R}_z\text{-size)} \\
\\
\frac{\Gamma \vdash \text{size}(C_1) = n_1 \quad \Gamma \vdash \text{size}(C_2) = n_2 \quad \Gamma \vdash \text{width}(C_1) = \text{width}(C_2)}{\Gamma \vdash \text{size}(\text{SEQ}(C_1, C_2)) = n_1 + n_2} \text{ (seq-size)} \\
\\
\frac{\Gamma \vdash \text{size}(C_1) = n_1 \quad \Gamma \vdash \text{size}(C_2) = n_2}{\Gamma \vdash \text{size}(\text{PAR}(C_1, C_2)) = n_1 + n_2} \text{ (par-size)} \\
\\
\frac{\Gamma \vdash \text{size}(C) = n}{\Gamma \vdash \text{size}(\text{ANC}(C)) = n} \text{ (anc-size)}
\end{array}$$

Figure 14: Deduction rules for QBRICKS: size (number of gates)

7.3 SQIR

The SQIR language [91, 90] is the representation language used by the VOQC optimizer (see Section 5.4). On its own, it also constitutes a solution for formally proved correct quantum programs. It is developed concurrently with QBRICKS and holds similar concerns: basically, to reduce the expressivity of programming languages such as Quipper and QWIRE so as to (1) still enable the whole implementation of emblematic algorithms (2) enable formal proof of specifications.

The development of SQIR followed that of QWIRE (see Section 6.1.3) when the authors observed the difficulty to hold formal verification, mainly linked to the management of memory wires. Hence SQIR and QWIRE have overlapping author and developer ship, they are both deeply embedded in the Coq proof assistant and they share the same mathematical libraries for, e.g., matrices and complex numbers. Schematically, compared to QWIRE, SQIR has a reduced expressivity (disabling, eg., the identification of qubit held values), making tractable the formal verification of functional program properties.

7.3.1 Programming Language

Just as QBRICKS, the programming part of SQIR is reduced to the minimum enabling implementation of main quantum programming features. It is two-layered: the *unitary part* corresponds to the design of unitary circuits and a generalized circuit layer adds branching measurement and ket initialization.

In SQIR, quantum circuits have type $\mathbb{N} \rightarrow \text{Set}$, with a positive integer parameter corresponding to their width. Quantum memory wires are identified by integer indexes, bounded by a global register size parameter d . Then, the type for unitary operators in SQIR (ucom) features the application of a circuit to a quantum register with d wires. It is defined inductively as follows.

```

Inductive ucom (U:  $\mathbb{N} \rightarrow \text{Set}$ ) (d:  $\mathbb{N}$ ) : Set :=
| useq : ucom U d  $\rightarrow$  ucom U d  $\rightarrow$  ucom U d
| uapp1 : U 1  $\rightarrow$   $\mathbb{N} \rightarrow$  ucom U d
| uapp2 : U 2  $\rightarrow$   $\mathbb{N} \rightarrow$   $\mathbb{N} \rightarrow$  ucom U d
| uapp3 : U 3  $\rightarrow$   $\mathbb{N} \rightarrow$   $\mathbb{N} \rightarrow$   $\mathbb{N} \rightarrow$  ucom U d

```


This definition holds two kinds of operations:

- the sequential composition, `useq`, which inputs two circuits and outputs their sequential composition,
- the application of an elementary gates to (a) given wire(s), depending on the width of this gate. There are three different versions of this operations, for gates of width 1, 2 or 3, named respectively `uapp1`, `uapp2` and `uapp3`. As an example, `uapp1` inputs a gate `U` of width 1 and a parameter `i`. It outputs the result of applying `U` on wire `i` in a register of size `d`¹³.

Then, SQIR provides a generalized circuit building layer, enabling the sequential composition of unitary commands, their initialization and branching measurement. Again, a circuit is an object `com` of type $\mathbb{N} \rightarrow \text{Set}$ applied on a register of specified size `d`. It is defined, inductively, as follows.

```

Inductive com (U:  $\mathbb{N} \rightarrow \text{Set}$ ) (d:  $\mathbb{N}$ ) : Set :=
| uc : ucom U d  $\rightarrow$  com U d  $\rightarrow$  ucom U d
| skip : com U d
| meas :  $\mathbb{N} \rightarrow$  com U d  $\rightarrow$  com U d  $\rightarrow$  com U d
| seq : com U d  $\rightarrow$  com U d  $\rightarrow$  com U d

```

A generalized circuit is built as either the lifting `uc` of a circuit into a generalized circuit, the empty `skip` operation, the branching measurement `meas` (with inputs a wire identifier for a qubit to measure and two generalized circuits to execute, depending on the measurement result) and the sequence `seq` of two different generalized circuits.

7.3.2 Matrix Semantics and Specifications

The semantics for SQIR programs is based on the standard matrix apparatus. For the unitary fragment, it uses the matrix semantics presented in Section 2.1.5. It is extended for generalized circuits by density operators semantics, similarly as in Section 7.1.

In its present state of development, SQIR enables to specify functional properties, describing the input-output relation, similarly to the one for QBRICKS introduced in Section 7.2.5

7.3.3 Implementation and Case Studies

SQIR is implemented as an embedded DSL into the Coq proof assistant. It was illustrated with specified and proved parametrized implementations of Simon's algorithm, QPE and Grover's Search Algorithm.

7.3.4 Comparison Between QBRICKS and SQIR

QBRICKS and SQIR are being developed concurrently, with very similar objectives. In particular they both trade-off between offering user-friendly programming features and reducing the language expressivity to the minimal, to enable functional formal verification. The solutions they provide share many common points. We discuss their main design differences.

- SQIR elementary operations consist in applying quantum gates on given wires of a quantum register, whereas QBRICKS proceeds by assembling quantum gates together into a quantum

¹³Note that, for this construction to make sense, the parameter `i` should not be greater than `d`. This condition is encoded by the semantics of SQIR.

circuit, just as bricks of a wall. Both views are inter-simulable: QBRICKS provides a macro place with integer parameters specifying the wire identifier a given sub-circuit should be applied to and the size of the overall circuit (corresponding to the size of the available quantum register). This macro is built by the parallel combination of its sub-circuit arguments with the appropriate number of occurrences of ID gate. It is of similar use as SQIR function `uapp`. On the other hand, QBRICKS gates assemblage is trivially simulable through SQIR `uapp`;

- On the other hand, SQIR provides a generalized circuit building layer, including measurement and classical control. Nevertheless, this upper layer is formalized through density operators, which are cumbersome objects for formal reasoning. So far, this part of the language only received illustrations with toy examples, such as superdense coding or quantum teleportation. In more involved implementations (such as QPE or Grover), SQIR authors followed a specification and proof strategy similar to that of QBRICKS: by reasoning on the quantum data outputs of circuits, specifying over the probability distribution of result if a measurement were performed. Hence, designing a generalized circuit building language semantics probed against actual implementations of real usage algorithms is still an open challenge;
- As introduced in Section 4.2.2, complexity properties of circuits constitute a fundamental aspect of quantum certification, decisive with regards to both the physical reliability of a computation and the quantum advantage it may provide. In the present state of development, SQIR does not offer a solution for this type of specifications. Still, this could be merely implemented in SQIR as additional functionality.

7.4 Conclusion about Formal Verification of Quantum Programs

In Table 5 we sum up the main concrete case study realizations of formally verified quantum algorithms: instances of Grover algorithm from QBRICKS, SQIR and QHL, Deutsch-Jozsa and QPE instances from QBRICKS and SQIR, and Shor-0F implementation from QBRICKS. For each of these implementations, we give the length of the code (column LoC) and a measure of the human proof effort required for the specification proofs. It was obtained by adding the length of the program specifications (Spec stands for the number of lines of specifications and intermediary lemmas) and the number of proof commands that were required to prove these specifications (column Cmd).

Table 5: Compared implementations of formally verified quantum algorithms

	QBRICKS[38]		SQIR [90, 91]		QHL[131]	
	LoC	Spec+Cmd	LoC	Spec+Cmd	LoC	Spec+Cmd
DJ	11	85	10	261		
Grover	39	279	15	926	90	2975
QPE	23	246	40	812		
Shor-0F	132	1212				

To the best of our knowledge Table 5 is comprehensive regarding parametrized formally proved quantum algorithm ¹⁴. Let us stress out how young the field is (in complement to the reduced

¹⁴An additional formalization of Deutsch Jozsa algorithm is presented by Bordg et al. [30]. We do not include it in Table 5 since it is not generated by a programming language but directly led as an algebraic proof. The total length of the proof is over 1 700 lines. Additionally, the online SQIR repository contains a Shor algorithm folder, but to the best of our knowledge, it was never explicitly presented nor described. It seems, by the way, hardly comparable with the QBRICKS implementation since it focuses on different aspects by bringing additional classical post-processing functions but lacking the oracle implementation.

number of concrete realizations, note that none of them is dated earlier than 2019). Nevertheless, it has already brought promising results.

One of the main challenges for formal verification is to reduce the human proof effort that is required for the certification of programs. As Table 5 shows, comparing this effort to the length of effective programs, QBRICKS offers a quite stable ratio $\simeq 10$. In a quite regular way, SQIR adds a $\simeq 3.5$ factor to this ratio and QHL, for the case of Grover algorithm, requires $\simeq 10$ times more human effort ¹⁵

8 Discussion and Bibliographical Notes

We end up this survey by providing some additional references for usage of formal methods in quantum information and quantum computing.

8.1 Deductive Verification

Deductive verification appears to be the most promising direction for the development of formal methods in quantum computing. In particular, it is particularly adequate for the formal verification of functional specifications, which is crucial for quantum programming and prone to play, there, a role similar to that of testing and debugging in classical computing (see Section 4.1). It is worth noting that, currently, all existing formally verified quantum algorithms descriptions [131] or implementations [38, 91] are based either on deductive verification or interactive proofs.

8.2 Model Checking

Attempts for functional verification of quantum algorithm with model checking techniques were also led before these developments [78, 197, 196]. They enabled to verify toy examples of quantum processes in a completely automated way. Nevertheless, this direction is limited by its high scale sensitivity, which is specifically problematic for quantum programs, since they are designed to tackle large problems instances.

8.3 Type Checking

Apart from functional specification, specialized type systems for quantum programming languages also facilitate programming and debugging. In Section 4.1 we introduced the verification of structural constraints and the non-duplicability of quantum information. Type checking may also have further use in quantum computing.

Recently, the SILQ language was proposed. It is based on a linear type system which, upon other features, enables to verify, for any quantum circuit, whether it can be *uncomputed*, a computation feature required at many stages of quantum implementations. Based on this type system, SILQ enables to automatically generate the uncomputation steps of circuits. This partial automation of the development lowers the expertise requirements for developers and the length of programs with regards to languages such as Q# or Quipper.

¹⁵Note that the QHL implementation of Grover algorithm concerns a restricted case, with regards to the two others figuring in Table 5. Furthermore, it does not contain the gate-to-gate circuit building but uses large circuit portions as primitives instead. Therefore, factor $\simeq 10$ is actually an underestimation.

8.4 Runtime assertion Checking

First, recall from Section 1 that proving quantum programs is mainly meant to replace the standard classical debugging method of testing and assertion checking. Apart from the development of formal proofs as an alternative, efforts are led to adapt this classical strategy to the quantum case. There, we still decorate programs with formal specifications (called *assertions* in this context) describing the evolution of the system state through an execution. But instead of mathematical proofs, these assertions are probed by statistical testing over program fragments. The challenges faced by such methods are mainly twofold:

Destructive measurement : memory reading destructs the superposition of a quantum state, therefore one cannot continue the execution after checking. Hence, assertion checking can be applied only to fragments of an execution.

Non-determinism : what we aim to check is a superposition of states, which induces a probability distribution of outcomes when a measurement is performed. Then, checking an assertion requires a number of testing runs large enough to build up a representative statistical distribution.

To overcome these difficulties, a first strategy is to reduce the specifications so as to express only properties that may be handled by assertion checking. Huang and Martonosi propose a “runtime-monitoring like” verification method for quantum circuits [96]. The annotation language is reduced so as to specify, for a quantum register, to be either in a classical state, in a superposition or entangled, without any concern about further description of the state.

More recently, Li et al. [126] developed an assertion-checking based method for the verification of fine quantum registers states properties, including functional descriptions of circuit behaviors. This method is based on:

- an assertion language, based on QHL projections (see Section 7.1), enabling functional specifications about computations at stake;
- the use and formalization of *gentle measurements*: quantum registers are not measured in the usual computation basis, but in a basis containing an output that is very close to the expected state;
- an appropriate formalization of the notion of distance between quantum states, enabling verification in terms of confidence interval between the current state of the system and its expected value.

Hence, gentle measurement enables to test executions over full functional state specifications. In addition to bringing expressivity to the specifications, it lowers the undesired effects of destructive measurements: since a gentle measurement operator contains an eigenstate that is an approximation of the expected quantum state, in most cases the measurement effect on the system state can be considered negligible and the execution can be pursued. Furthermore, the test result probability distribution is centered on a specific value. Therefore, a much-reduced set of runs bring valuable statistical conclusions.

However, verification following this strategy only holds for a particular instance of a circuit, instead of a family of quantum circuits with unassigned parameters as in propositions from Section 7. Furthermore, in the general case, gentle measurements are implemented by applying (1) a unitary U uncomputing the system state into a ket of the computational basis (2) a measurement in the computational basis (3) unitary U^\dagger , to recover the initial state. Then, assertion checking inputs unitaries

U and U^\dagger that are themselves prone to error. More precisely, to test against the exact expected value of the state, operator U^\dagger should be equivalent to the computation under test. Practically, gentle measurements approximate the state under test. This enables a simplification of the measurement operator, at the cost of the robustness of the procedure.

8.5 Verification of Quantum Communication Protocols

Another challenge for formal methods is the verification of quantum information processes concerns quantum protocols.

Quantum key distribution protocols [164, 22, 128, 134] enable secured information exchanges between two parties. These protocols exploit the fact that, due to the destructive measurement, physics laws prevent almost any ¹⁶ possibility for a potential eavesdropping in the exchange of quantum information. In particular, several formally verified implementations of the BB84 quantum key distribution protocol [22] have been proposed in the literature, based either on process calculus [137, 117], formalization in Coq [28] or model checking [63, 68].

Bordg et al. [30] propose a formalization of quantum information in the ISABELLE/HOL proof assistant. They illustrate their methods through the cases of quantum teleportation and the quantum prisoner dilemma. This work also contains a formally verified implementation of the Deutsch-Jozsa algorithm in ISABELLE/HOL.

In a more fundamental prospect, Echenim et al. [61] provide an ISABELLE/HOL proof for the CHSH inequalities [45]. These are probability distributions about crossed measurement results of quantum observables. They provide a proof for the Bell theorem [21], inducing that no classical theory could account for the entanglement phenomenon (hence that quantum physics cannot be reduced to local classical theories).

9 Conclusion

9.1 Summary

Throughout this survey, we introduced the context, the main challenges and the most promising results in formally verified quantum programming. The current state of affairs in this emerging domain can be summed up as follows:

- Quantum computing is an emerging domain, with huge potential application fields and promises. Progresses in the development of concrete machines are reaching the practical relevance landmark: prototypes are getting powerful enough to overpass classical computers. Consequently, quantum software development is becoming a crucial industrial short-term need;
- Quantum software deals with an entirely new programming paradigm. Upon its main particularities are the dual nature of information (either classical or quantum), the destructive measurement and irreducibly probabilistic computations;
- These specificities make programming particularly non-intuitive and prone to error. Furthermore, they make it very hard to directly import usual debugging methods from the classical

¹⁶These protocols are based on the fact that, in the general case, a potential eavesdropper destructs a quantum message he attempts to intercept, so that the parties can detect the attempt. Nevertheless, it is based on probabilities and there is always a chance for an eavesdropper to perform only conservative measurements. The corresponding probability is bounded by r^n , with $r \in]0, 1[$ and n the length of the sent message.

practice (based on test and assertion checking). The technique presented in Section 8.4 might bring some hope, but this is still so far at a very preliminary stage;

- Formal methods appear as the privileged alternative for debugging strategies. Apart from providing solutions to the destructive measurement challenge, they have additional decisive advantages: mainly, they provide absolute guarantee of the correction of programs, and they hold for any instance of programs they verify;
- During the last ten years (the *genesis* of formal quantum programming), this new field has shown promising results in the different stages of software development: high-level program designs, circuit building languages, verification, compilation, optimization, *etc.*

9.2 Main Current Challenges

Although encouraging, these early successes draw the road map for rising the field from academic proof of concepts to practically usable programming solutions. We present the main coming challenges for this development in three categories: providing relevant integrated development solutions for the NISQ era, offering practical wide spreadable user experience and developing a full-fledged formally verified quantum compilation toolchain.

Provide relevant integrated development solutions for the NISQ era. So far, quantum formal verification mainly proved its relevance by offering solutions for the unitary core of quantum computations. As a matter of fact, illustrations and concrete implementations using these techniques primarily treated historical emblematic algorithms such as Shor [174], QPE [111] or Grover search [84]. The classical treatment in these algorithms can be completely decoupled from their quantum core.

However, the first generation of quantum computing machine (the *NISQ era*, see Section 2.3.2) hints towards a radically distinct mode of operation. NISQ machines will have limited, noisy resources. A major consequence is that these quantum processors are too small to support the error correction mechanisms required for Shor's and Grover's algorithms. Such NISQ processors aim instead at different kinds of algorithms: hybrid algorithms such as variational algorithms [135]. Hybrid algorithms tightly mix quantum and classical data treatments: one cannot decouple the quantum part of the algorithm from its classical part.

Adapting the quantum formal methods to the NISQ setting to support hybrid quantum/classical computation is a challenge and an active current research avenue.

Offer practical wide spreadable user experience: formalism, language design, automation.

In the present state of development, programming languages enabling formal verification usually sacrifice their expressivity to formal reasoning. In particular, all the solutions that have been probed against actual parametric quantum algorithms fail to satisfy essential elements of the requirements listed in Section 6.1.2 for scalable quantum programming languages: in addition to classical processing mentioned above, they cannot manipulate quantum registers and wire references. Verified programming should address these limitations.

Another issue, while concerning any method of quantum programming, is especially critical in the case of formally verified programming. It concerns the level of qualifications required from developers. The interpretation of quantum computations indeed requires unusual and non-intuitive mathematical formalism (including Kronecker products, complex phase amplitudes, probabilistic reasoning, *etc.*). While the need for qualified programmers is prone to grow rapidly in the coming

years, integrating formal verification should come with the development of user-friendly specification languages and highly automated mathematical reasoning engines.

Formally verified quantum compilation toolchain. In addition to the preceding considerations, in its early academic ages, quantum formal verification focused on idealized representations of quantum circuits, directly extracted from algorithm descriptions. As introduced in Section 2.3.2, these logical qubits are merely abstract models for actual computations to be run. A major addition that is left for future works in quantum formal verification is error correction, with formal verification that the state of the system preserves the functional correctness of computations (assuming a given error model, and possibly with probabilistic specifications). Presented in Section 5.4, VOQC is a first step towards this goal.

Another future direction concerns integration in a widespread classical development environment. Recall from Section 6.1.3 that many widespread quantum programming languages benefit from embeddings in usual programming languages—such as Python. Today most formal verification solutions are embedded in a more academic functional development environment (such as Haskell, Ocaml, proof assistants or deductive verification environments). Interfaces should be developed to integrate formally verified quantum computations into comprehensive projects.

At the other extremity of the development stack, formal verification should accompany the implementation of compiled programs on concrete machines. This induces verified solutions for the qubit mapping problem and gate simulation (see Section 2.3.3), depending on the particular target material and its proper constraints.

Acknowledgments

This work was supported in part by the French National Research Agency (ANR) under the research project SoftQPRO ANR-17-CE25-0009-02, by the DGE of the French Ministry of Industry under the research project PIA-GDN/QuantEx P163746-484124, by the STIC-AmSud project Qapla' 21-SITC-10 and by the Carnot project Qstack. We thank the anonymous reviewers for helpful comments on earlier drafts of the manuscript.

References

- [1] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011.
- [2] Samson Abramsky and Bob Coecke. Categorical quantum mechanics. In *Handbook of Quantum Logic and Quantum Structures*, pages 261–323. Elsevier, 2009.
- [3] Thorsten Altenkirch and Alexander S Green. The quantum IO monad. *Semantic Techniques in Quantum Computation*, pages 173–205, 2010.
- [4] Matthew Amy. *Formal Methods in Quantum Circuit Design*. PhD thesis, University of Waterloo, Ontario, Canada, 2019.
- [5] Matthew Amy. Towards large-scale functional verification of universal quantum circuits. In Peter Selinger and Giulio Chiribella, editors, *Proceedings 15th International Conference on*

- Quantum Physics and Logic, QPL 2018*, volume 287 of *Electronic Proceedings in Theoretical Computer Science*, pages 1–21, Halifax, Canada, 2019. EPTCS.
- [6] Matthew Amy, Jianxin Chen, and Neil J. Ross. A finite presentation of CNOT-dihedral operators. In Bob Coecke and Aleks Kissinger, editors, *Proceedings 14th International Conference on Quantum Physics and Logic, Nijmegen, The Netherlands, 3-7 July 2017*, volume 266 of *Electronic Proceedings in Theoretical Computer Science*, pages 84–97, 2018.
- [7] Matthew Amy, Dmitri Maslov, and Michele Mosca. Polynomial-time T-depth optimization of Clifford+T circuits via matroid partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 33(10):1476–1489, 2014.
- [8] Krzysztof R. Apt and Ernst-Rüdiger Olderog. Fifty years of Hoare’s logic. *Formal Aspects Comput.*, 31(6):751–807, 2019.
- [9] 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, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [10] Miriam Backens. The ZX-calculus is complete for stabilizer quantum mechanics. In *New Journal of Physics*, volume 16, page 093021. IOP Publishing, Sep 2014.
- [11] Miriam Backens, Hector Miller-Bakewell, Giovanni de Felice, Leo Lobski, and John van de Wetering. There and back again: A circuit extraction tale. *Quantum*, 5:421, 2021.
- [12] Thomas Ball, Byron Cook, Vladimir Levin, and Sriram K. Rajamani. Slam and static driver verifier : Technology transfer of formal methods inside microsoft. In *Integrated Formal Methods, 4th International Conference, IFM 2004*. Springer, 2004.
- [13] Alexandru Baltag, Jort Bergfeld, Kohei Kishida, Joshua Sack, Sonja Smets, and Shengyang Zhong. Plqp & company: decidable logics for quantum algorithms. *International Journal of Theoretical Physics*, 53(10):3628–3647, 2014.
- [14] Alexandru Baltag and Sonja Smets. LQP: the dynamic logic of quantum information. *Mathematical Structures in Computer Science*, 16(3):491–525, 2006.
- [15] Alexandru Baltag and Sonja Smets. The logic of quantum programs. *CoRR*, abs/2109.06792, 2021.
- [16] Henk P. Barendregt. *The Lambda-Calculus, its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundation of Mathematics*. North Holland, second edition, 1984.
- [17] Mike Barnett, Manuel Fähndrich, K. Rustan M. Leino, Peter Müller, Wolfram Schulte, and Herman Venter. Specification and verification: the SPEC# experience. *Commun. ACM*, 54(6):81–91, 2011.
- [18] Clark Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of model checking*, pages 305–343. Springer, 2018.
- [19] Gilles Barthe, Justin Hsu, Mingsheng Ying, Nengkun Yu, and Li Zhou. Relational proofs for quantum programs. *Proc. ACM Program. Lang.*, 4(POPL):21:1–21:29, 2020.

- [20] Patrick Behm, Paul Benoit, Alain Faivre, and Jean-Marc Meynadier. Météor: A successful application of B in a large project. In *In Proceedings of the World Congress on Formal Methods in the Development of Computing Systems (FM'99)*. Springer, 1999.
- [21] John S Bell. On the Einstein Podolsky Rosen paradox. *Physica Physique Fizika*, 1(3):195–200, 1964.
- [22] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, Bangalore, India, 1984.
- [23] Nick Benton. A mixed linear and non-linear logic: Proofs, terms and models (extended abstract). In Leszek Pacholski and Jerzy Tiuryn, editors, *Computer Science Logic, Eighth International Workshop, CSL'94, Selected Papers*, volume 933 of *Lecture Notes in Computer Science*, pages 121–135, 1994.
- [24] Jort Martinus Bergfeld and Joshua Sack. Deriving the correctness of quantum protocols in the probabilistic logic for quantum programs. *Soft Comput.*, 21(6):1421–1441, 2017.
- [25] Yves Bertot and Pierre Castéran. *Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions*. Springer Science & Business Media, 2013.
- [26] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017.
- [27] Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin T. Vechev. Silq: a high-level quantum language with safe uncomputation and intuitive semantics. In Alastair F. Donaldson and Emina Torlak, editors, *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, pages 286–300. ACM, 2020.
- [28] Jaap Boender, Florian Kammüller, and Rajagopal Nagarajan. Formalization of quantum protocols using Coq. In Chris Heunen, Peter Selinger, and Jamie Vicary, editors, *Proceedings of the 12th International Workshop on Quantum Physics and Logic (QPL 2015)*, volume 195 of *Electronic Proceedings in Theoretical Computer Science*, pages 71–83, Oxford, UK, 2015. EPTCS.
- [29] Adam D Bookatz. QMA-complete problems. *Quantum Information & Computation*, 14(5&6):361–383, 2014.
- [30] Anthony Bordg, Hanna Lachnitt, and Yijun He. Certified quantum computation in Isabelle/HOL. *Journal of Automated Reasoning*, 65(5):691–709, 2021.
- [31] E. Bounimova, P. Godefroid, , and D. Molnar. Billions and billions of constraints : Whitebox fuzz testing in production. In *35th International Conference on Software Engineering (ICSE)*, pages 122–131. IEEE/ACM, 2013.
- [32] Hans J. Briegel, Dan E. Browne, Wolfgang Dür, Robert Raußendorf, and Maarten Van den Nest. Measurement-based quantum computation. *Nature Physics*, 5(1):19–26, January 2009.
- [33] Daniel E Browne, Elham Kashefi, Mehdi Mhalla, and Simon Perdrix. Generalized flow and determinism in measurement-based quantum computation. *New Journal of Physics*, 9(8):250–250, aug 2007.

- [34] Olivier Brunet and Philippe Jorrand. Dynamic quantum logic for quantum programs. *International Journal of Quantum Information*, 2(01):45–54, 2004.
- [35] Cristian Cadar and Koushik Sen. Symbolic execution for software testing: three decades later. *Commun. ACM*, 56(2):82–90, 2013.
- [36] Titouan Carette, Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. Completeness of Graphical Languages for Mixed States Quantum Mechanics. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 108:1–108:15, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [37] Rohit Chadha, Paulo Mateus, and Amílcar Sernadas. Reasoning about imperative quantum programs. *Electronic Notes in Theoretical Computer Science*, 158:19–39, 2006.
- [38] Christophe Chareton, Sébastien Bardin, François Bobot, Valentin Perrelle, and Benoît Valiron. An automated deductive verification framework for circuit-building quantum programs. In Nobuko Yoshida, editor, *Programming Languages and Systems - 30th European Symposium on Programming, ESOP 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings*, volume 12648 of *Lecture Notes in Computer Science*, pages 148–177. Springer, 2021.
- [39] Lily Chen, Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray Perlner, and Daniel Smith-Tone. *Report on post-quantum cryptography*, volume 12. US Department of Commerce, National Institute of Standards and Technology, 2016.
- [40] John Chiaverini, Dietrich Leibfried, Tobias Schaetz, Murray D Barrett, RB Blakestad, J Britton, Wayne M Itano, Juergen D Jost, Emanuel Knill, Christopher Langer, et al. Realization of quantum error correction. *Nature*, 432(7017):602–605, 2004.
- [41] Cirq Developers. Cirq, July 2018. See full list of authors on Github: <https://github.com/quantumlib/Cirq/graphs/contributors>.
- [42] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In *Logics of Programs, Workshop, LNCS 131*, pages 52–71. Springer, 1981.
- [43] Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem. *Handbook of Model Checking*. Springer, 2018.
- [44] Edmund M. Clarke and Jeannette M. Wing. Formal methods: State of the art and future directions. *ACM Computing Surveys (CSUR)*, 28(4):626–643, 1996.
- [45] John F Clauser, Michael A Horne, Abner Shimony, and Richard A Holt. Proposed experiment to test local hidden-variable theories. *Physical review letters*, 23(15):880, 1969.
- [46] Bob Coecke and Ross Duncan. Interacting quantum observables: Categorical algebra and diagrammatics. *New Journal of Physics*, 13(4):043016, Apr 2011.
- [47] Bob Coecke and Aleks Kissinger. *Picturing quantum processes*. Cambridge University Press, Cambridge, United Kingdom, 2017.

- [48] Bob Coecke and Simon Perdrix. Environment and classical channels in categorical quantum mechanics. *Log. Methods Comput. Sci.*, 8(4), 2010.
- [49] Patrick Cousot and Radhia Cousot. Abstract interpretation : A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Fourth ACM Symposium on Principles of Programming Languages (POPL)*, pages 238–252. ACM, 1977.
- [50] Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. The ASTRÉE analyzer. In *European Symposium on Programming Languages and Systems, ESOP 2005*. Springer, 2005.
- [51] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta. Open quantum assembly language, 2017.
- [52] Vincent Danos and Elham Kashefi. Determinism in the one-way model. *Phys. Rev. A*, 74:052310, Nov 2006.
- [53] Niel de Beaudrap, Xiaoning Bian, and Quanlong Wang. Fast and effective techniques for T-count reduction via spider nest identities. In Steven T. Flammia, editor, *15th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2020, June 9-12, 2020, Riga, Latvia*, volume 158 of *LIPICs*, pages 11:1–11:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [54] Niel de Beaudrap, Ross Duncan, Dominic Horsman, and Simon Perdrix. Pauli fusion: a computational model to realise quantum transformations from ZX terms. In *QPL'19 : International Conference on Quantum Physics and Logic*, Los Angeles, United States, June 2019.
- [55] Niel de Beaudrap and Dominic Horsman. The ZX calculus is a language for surface code lattice surgery. *Quantum*, 4:218, January 2020.
- [56] Ellie D’Hondt and Prakash Panangaden. Quantum weakest preconditions. *Mathematical Structures in Computer Science*, 16(3):429–451, 2006.
- [57] Alejandro Díaz-Caro. A lambda calculus for density matrices with classical and probabilistic controls. In Bor-Yuh Evan Chang, editor, *Proceedings of the 15th Asian Symposium on Programming Languages and Systems (APLAS'17)*, volume 10695 of *Lecture Notes in Computer Science*, pages 448–467, Suzhou, China, 2017. Springer.
- [58] Alejandro Díaz-Caro and Octavio Malherbe. A concrete categorical semantics of lambda-S. In Beniamino Accattoli and Carlos Olarte, editors, *Proceedings of the 13th Workshop on Logical and Semantic Frameworks with Applications, LSFA 2018, Fortaleza, Brazil, September 26-28, 2018*, volume 344 of *Electronic Notes in Theoretical Computer Science*, pages 83–100. Elsevier, 2019.
- [59] Edsger W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
- [60] Ross Duncan, Aleks Kissinger, Simon Perdrix, and John van de Wetering. Graph-theoretic simplification of quantum circuits with the ZX-calculus. *Quantum*, 4:279, 2020.
- [61] Mnacho Echenim. Quantum projective measurements and the CHSH inequality. *Arch. Formal Proofs*, 2021, 2021.

- [62] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 47(10):777, 1935.
- [63] Mohamed Elboukhari, Mostafa Azizi, and Abdelmalek Azizi. Verification of quantum cryptography protocols by model checking. *Int. J. Network Security & Appl*, 2(4):43–53, 2010.
- [64] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. Technical Report MIT-CTP/4610, MIT, 2014.
- [65] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science*, 292(5516):472–475, 2001.
- [66] Yuan Feng, Runyao Duan, Zhengfeng Ji, and Mingsheng Ying. Proof rules for the correctness of quantum programs. *Theoretical Computer Science*, 386(1-2):151–166, 2007.
- [67] Yuan Feng and Mingsheng Ying. Quantum hoare logic with classical variables. *CoRR*, abs/2008.06812, 2020.
- [68] Verónica Fernández, María-José García-Martínez, Luis Hernández-Encinas, and Agustín Martín. Formal verification of the security of a free-space quantum key distribution system. In *Proc. World Congr. Comput. Sci. Comput. Eng. Appl. Comput.(WORLDCOMP) Int. Conf. Security Manag.(SAM)*, 2011.
- [69] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6–7):467–488, 1982.
- [70] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 - where programs meet provers. In Matthias Felleisen and Philippa Gardner, editors, *Proceedings of the 22nd European Symposium on Programming Languages and Systems (ESOP 2013), Held as Part of the European Joint Conferences on Theory and Practice of Software (ETAPS 2013)*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128, Rome, Italy, 2013. Springer.
- [71] Jean-Christophe Filliâtre. Deductive software verification. *STTT*, 13(5):397–403, 2011.
- [72] Michael J Fischer and Richard E Ladner. Propositional dynamic logic of regular programs. *Journal of computer and system sciences*, 18(2):194–211, 1979.
- [73] Melvin Fitting. *First-order logic and automated theorem proving*. Springer Science & Business Media, 2012.
- [74] R. W. Floyd. Assigning meanings to programs. In *Mathematical Aspects of Computer Science, Proceedings of Symposia in Applied Mathematics*, pages 19–32. American Mathematical Society, 1967.
- [75] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.
- [76] Michael Freedman, Alexei Kitaev, Michael Larsen, and Zhenghan Wang. Topological quantum computation. *Bulletin of the American Mathematical Society*, 40(1):31–38, 2003.

- [77] Peng Fu, Kohei Kishida, and Peter Selinger. Linear dependent type theory for quantum programming languages. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 440–453, 2020.
- [78] Simon J. Gay, Rajagopal Nagarajan, and Nikolaos Papanikolaou. QMC: a model checker for quantum systems. In Aarti Gupta and Sharad Malik, editors, *Proceeding of the 20th International Conference on Computer Aided Verification (CAV 2008)*, volume 5123 of *Lecture Notes in Computer Science*, pages 543–547, Princeton, NJ, USA, 2008. Springer.
- [79] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–101, 1987.
- [80] Georges Gonthier. Formal proof – the four-color theorem. *Notices of the AMS*, 55(11):1382–1393, 2008.
- [81] Daniel Gottesman. *Stabilizer codes and quantum error correction*. PhD thesis, Caltech, 1997.
- [82] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. An introduction to quantum programming in Quipper. In Gerhard W. Dueck and D. Michael Miller, editors, *Proceedings of the 5th International Conference on Reversible Computation (RC’13)*, volume 7948 of *Lecture Notes in Computer Science*, pages 110–124, Victoria, BC, Canada, 2013. Springer.
- [83] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. Quipper: A scalable quantum programming language. In Hans-Juergen Boehm and Cormac Flanagan, editors, *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation, (PLDI’13)*, pages 333–342, Seattle, WA, USA, 2013. ACM.
- [84] Lov K. Grover. A fast quantum mechanical algorithm for database search. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 212–219, Philadelphia, Pennsylvania, USA, 1996. ACM.
- [85] Amar Hadzihasanovic, Kang Feng Ng, and Quanlong Wang. Two complete axiomatisations of pure-state qubit quantum computing. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS ’18*, pages 502–511, New York, NY, USA, 2018. ACM.
- [86] David Harel, Dexter Kozen, and Jerzy Tiuryn. Dynamic logic. In *Handbook of philosophical logic*, pages 99–217. Springer, 2001.
- [87] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103:150502, Oct 2009.
- [88] Ichiro Hasuo and Naohiko Hoshino. Semantics of higher-order quantum computation via geometry of interaction. *Annals of Pure and Applied Logic*, 168(2):404 – 469, 2017.
- [89] Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Grégoire Sutre. Software verification with Blast. In *Proceedings of the 10th International Conference on Model Checking Software, SPIN’03*. Springer, 2003.
- [90] Kesha Hietala, Robert Rand, Shih-Han Hung, Liyi Li, and Michael Hicks. Proving quantum programs correct. In Liron Cohen and Cezary Kaliszyk, editors, *12th International Conference on Interactive Theorem Proving, ITP 2021, June 29 to July 1, 2021, Rome, Italy (Virtual Conference)*, volume 193 of *LIPICs*, pages 21:1–21:19, 2021.

- [91] Kesha Hietala, Robert Rand, Shih-Han Hung, Xiaodi Wu, and Michael Hicks. A verified optimizer for quantum circuits. *Proc. ACM Program. Lang.*, 5(POPL):1–29, 2021.
- [92] Anne Hillebrand. Quantum protocols involving multiparticle entanglement and their representations. Master’s thesis, University of Oxford, 2011.
- [93] R. Hindley. The principal type-scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, 1969.
- [94] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969.
- [95] Clare Horsman, Austin G. Fowler, Simon Devitt, and Rodney Van Meter. Surface code quantum computing by lattice surgery. *New Journal of Physics*, 14(12):123011, December 2012.
- [96] Yipeng Huang and Margaret Martonosi. Statistical assertions for validating patterns and finding bugs in quantum programs. In Srilatha Bobbie Manne, Hillery C. Hunter, and Erik R. Altman, editors, *Proceedings of the 46th International Symposium on Computer Architecture (ISCA 2019)*, pages 541–553, Phoenix, AZ, USA, 2019. ACM.
- [97] Shih-Han Hung, Kesha Hietala, Shaopeng Zhu, Mingsheng Ying, Michael Hicks, and Xiaodi Wu. Quantitative robustness analysis of quantum programs. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–29, 2019.
- [98] Dominik Janzing, Pawel Wocjan, and Thomas Beth. “non-identity-check” is QMA-complete. *International Journal of Quantum Information*, 03(03):463–473, 2005.
- [99] Ali JavadiAbhari, Shruti Patil, Daniel Kudrow, Jeff Heckey, Alexey Lvov, Frederic T. Chong, and Margaret Martonosi. ScaffCC: Scalable compilation and analysis of quantum programs. *Parallel Computing*, 45:2–17, 2015.
- [100] Emmanuel Jeandel. The rational fragment of the ZX-calculus, 2018.
- [101] Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. A complete axiomatisation of the ZX-calculus for Clifford+T quantum mechanics. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS ’18*, pages 559–568, New York, NY, USA, 2018. ACM.
- [102] Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. Diagrammatic reasoning beyond Clifford+T quantum mechanics. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS ’18*, pages 569–578, New York, NY, USA, 2018. ACM.
- [103] Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. A generic normal form for ZX-diagrams and application to the rational angle completeness. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–10, 2019.
- [104] Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. Completeness of the ZX-calculus. *Logical Methods in Computer Science*, 16(2), 2020.
- [105] Yoshihiko Kakutani. A logic for formal verification of quantum programs. In *Annual Asian Computing Science Conference*, pages 79–93. Springer, 2009.

- [106] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry Fields: A software platform for photonic quantum computing. *Quantum*, 3:129, 2019.
- [107] Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles, and Boris Yakobowski. Frama-C : A software analysis perspective. *Formal Asp. Comput.*, 27(3):573–609, 2015.
- [108] Aleks Kissinger and John van de Wetering. PyZX: Large scale automated diagrammatic reasoning. *Electronic Proceedings in Theoretical Computer Science*, 318:229–241, May 2020.
- [109] Aleks Kissinger and John van de Wetering. Reducing the number of non-clifford gates in quantum circuits. *Phys. Rev. A*, 102:022406, Aug 2020.
- [110] Aleks Kissinger and Vladimir Zamdzhiev. Quantomatic: A proof assistant for diagrammatic reasoning. In Amy P. Felty and Aart Middeldorp, editors, *Proceedings for the 25th International Conference on Automated Deduction (CADE-25)*, volume 9195 of *Lecture Notes in Computer Science*, pages 326–336, Berlin, Germany, 2015. Springer.
- [111] A Yu Kitaev. Quantum measurements and the abelian stabilizer problem. Available online as [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026), 1995.
- [112] Erwin Klein, June Andronick, Kevin Elphinstone, Gernot Heiser, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. seL4: formal verification of an operating-system kernel. *Commun. ACM*, 53(6):107–115, 2010.
- [113] Emmanuel Knill. Conventions for quantum pseudocode. Technical Report LA-UR-96-2724, Los Alamos National Laboratory, June 1996.
- [114] Dexter Kozen. Results on the propositional μ -calculus. *Theoretical computer science*, 27(3):333–354, 1983.
- [115] Saul A. Kripke. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94, 1963.
- [116] Daniel Kroening and Michael Tautschnig. Cbmc - c bounded model checker. In *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014*, pages 389–391. Springer, 2014.
- [117] Takahiro Kubota, Yoshihiko Kakutani, Go Kato, Yasuhito Kawano, and Hideki Sakurada. Semi-automated verification of security proofs of quantum cryptographic protocols. *Journal of Symbolic Computation*, 73:192–220, 2016.
- [118] Hidenori Kuwakado and Masakatu Morii. Quantum distinguisher between the 3-round Feistel cipher and the random permutation. In *2010 IEEE International Symposium on Information Theory*, pages 2682–2685. IEEE, 2010.
- [119] Hidenori Kuwakado and Masakatu Morii. Security on the quantum-type Even-Mansour cipher. In *2012 International Symposium on Information Theory and its Applications*, pages 312–316. IEEE, 2012.

- [120] Ugo Dal Lago, Claudia Faggian, Benoît Valiron, and Akira Yoshimizu. The geometry of parallelism: classical, probabilistic, and quantum effects. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'17)*, pages 833–845, Paris, France, 2017. ACM.
- [121] Joachim Lambek. The mathematics of sentence structure. *The American Mathematical Monthly*, 65(3):154–170, 1958.
- [122] Martin Lange. Model checking propositional dynamic logic with all extras. *Journal of Applied Logic*, 4(1):39–49, 2006.
- [123] Louis Lemonnier, John van de Wetering, and Aleks Kissinger. Hypergraph simplification: Linking the path-sum approach to the ZH-calculus. *arXiv preprint arXiv:2003.13564*, 2020.
- [124] Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009.
- [125] Jerzy Lewandowski. Volume and quantizations. *Classical and Quantum Gravity*, 14(1):71, 1997.
- [126] Gushu Li, Li Zhou, Nengkun Yu, Yufei Ding, Mingsheng Ying, and Yuan Xie. Projection-based runtime assertions for testing and debugging quantum programs. *Proceedings of the ACM on Programming Languages*, 4(OOPSLA):1–29, 2020.
- [127] Yangjia Li and Mingsheng Ying. Algorithmic analysis of termination problems for quantum programs. *Proceedings of the ACM on Programming Languages*, 2(POPL):1–29, 2017.
- [128] Sheng-Kai Liao, Wen-Qi Cai, Wei-Yue Liu, Liang Zhang, Yang Li, Ji-Gang Ren, Juan Yin, Qi Shen, Yuan Cao, Zheng-Ping Li, et al. Satellite-to-ground quantum key distribution. *Nature*, 549(7670):43–47, 2017.
- [129] Daniel A Lidar and Todd A Brun. *Quantum Error Correction*. Cambridge University Press, 2013.
- [130] Bert Lindenhovius, Michael Mislove, and Vladimir Zamdzhiev. Enriching a linear/non-linear lambda calculus: A programming language for string diagrams. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 659–668, 2018.
- [131] Junyi Liu, Bohua Zhan, Shuling Wang, Shenggang Ying, Tao Liu, Yangjia Li, Mingsheng Ying, and Naijun Zhan. Formal verification of quantum algorithms using quantum Hoare logic. In Isil Dillig and Serdar Tasiran, editors, *Computer Aided Verification*, pages 187–207, Cham, 2019. Springer International Publishing.
- [132] Tao Liu, Yangjia Li, Shuling Wang, Mingsheng Ying, and Naijun Zhan. A theorem prover for quantum Hoare logic and its applications. Available as [arXiv:1601.03835](https://arxiv.org/abs/1601.03835), 2016.
- [133] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.
- [134] Hoi-Kwong Lo, Xiongfeng Ma, and Kai Chen. Decoy state quantum key distribution. *Physical review letters*, 94(23):230504, 2005.

- [135] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
- [136] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, 1978.
- [137] Rajagopal Nagarajan and Simon Gay. Formal verification of quantum protocols. Available online as [arXiv:quant-ph/0203086](https://arxiv.org/abs/quant-ph/0203086), 2002.
- [138] Yunseong Nam, Neil J Ross, Yuan Su, Andrew M Childs, and Dmitri Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1):1–12, 2018.
- [139] Michael A. Nielsen and Isaac Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge, United Kingdom, 2002.
- [140] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016.
- [141] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: a proof assistant for higher-order logic*. Springer, 2002.
- [142] Bernhard Ömer. *Structured Quantum Programming*. PhD thesis, TU Wien, 2003.
- [143] Michele Pagani, Peter Selinger, and Benoît Valiron. Applying quantitative semantics to higher-order quantum computing. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14*, page 647–658, New York, NY, USA, 2014. Association for Computing Machinery.
- [144] Luca Paolini and Margherita Zorzi. qPCF: A language for quantum circuit computations. In T. V. Gopal, Gerhard Jäger, and Silvia Steila, editors, *Proceedings of the 14th Annual Conference on Theory and Applications of Models of Computation (TAMC'17)*, volume 10185 of *Lecture Notes in Computer Science*, pages 455–469, Bern, Switzerland, 2017.
- [145] Christine Paulin-Mohring. Introduction to the calculus of inductive constructions. In *volume 55 of Studies in Logic (Mathematical logic and foundations)*. College Publications, 2015.
- [146] Christine Paulin-Mohring. Introduction to the calculus of inductive constructions. In Bruno Woltzenlogel Paleo and David Delahaye, editors, *All about Proofs, Proofs for All*, volume 55 of *Studies in Logic (Mathematical logic and foundations)*. College Publications, 2015.
- [147] Jennifer Paykin. *Linear/non-Linear Types For Embedded Domain-Specific Languages*. PhD thesis, University of Pennsylvania, 2018.
- [148] Jennifer Paykin, Robert Rand, and Steve Zdancewic. QWIRE: a core language for quantum circuits. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'17)*, pages 846–858, Paris, France, 2017. ACM.
- [149] Simon Perdrix. Quantum entanglement analysis based on abstract interpretation. In *International Static Analysis Symposium*, pages 270–282. Springer, 2008.

- [150] Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [151] Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 46–57. iee, 1977.
- [152] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [153] Qiskit Community. Qiskit: An open-source framework for quantum computing, March 2017. The official way to cite Qiskit is using the bibentry at <https://github.com/Qiskit/qiskit/blob/master/Qiskit.bib>.
- [154] Quantum Computing Report. List of tools. Available online¹⁷, 2019.
- [155] Robert Rand, Kesha Hietala, and Michael Hicks. Formal verification vs. quantum uncertainty. In *3rd Summit on Advances in Programming Languages (SNAPL 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [156] Robert Rand, Jennifer Paykin, and Steve Zdancewic. QWIRE practice: Formal verification of quantum circuits in Coq. In Bob Coecke and Aleks Kissinger, editors, *Proceedings 14th International Conference on Quantum Physics and Logic (QPL 2017)*, volume 266 of *Electronic Proceedings in Theoretical Computer Science*, pages 119–132, Nijmegen, The Netherlands, 2017. EPTCS.
- [157] Robert Raussendorf, Daniel E Browne, and Hans J Briegel. Measurement-based quantum computation on cluster states. *Physical review A*, 68(2):022312, 2003.
- [158] John C Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74. IEEE, 2002.
- [159] Francisco Rios. *On a Categorically Sound Quantum Programming Language for Circuit Description*. PhD thesis, Dalhousie University, Halifax, Nova Scotia, Canada, August 2021.
- [160] Francisco Rios and Peter Selinger. A categorical model for a quantum circuit description language. In Bob Coecke and Aleks Kissinger, editors, *Proceedings 14th International Conference on Quantum Physics and Logic (QPL 2017)*, volume 266 of *Electronic Proceedings in Theoretical Computer Science*, pages 164–178, Nijmegen, The Netherlands, 2018.
- [161] Xavier Rival and Kwangkeun Yi. *Introduction to Static Analysis: An Abstract Interpretation Perspective*. The MIT Press, 2020.
- [162] Neil J. Ross. *Algebraic and Logical Methods in Quantum Computation*. PhD thesis, Dalhousie University, 2015. Available as [arxiv:1510.02198](https://arxiv.org/abs/1510.02198).
- [163] Thomas Santoli and Christian Schaffner. Using Simon’s algorithm to attack symmetric-key cryptographic primitives. *Quantum Inf. Comput.*, 17(1&2):65–78, 2017.
- [164] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Reviews of modern physics*, 81(3):1301, 2009.

¹⁷<https://quantumcomputingreport.com/resources/tools/>

- [165] Artur Scherer, Benoît Valiron, Siun-Chuon Mau, Scott Alexander, Eric Van den Berg, and Thomas E Chapuran. Concrete resource analysis of the quantum linear-system algorithm used to compute the electromagnetic scattering cross section of a 2d target. *Quantum Information Processing*, 16(3):60, 2017.
- [166] Maria Schuld. *Supervised Learning with Quantum Computers*. Springer, 2018.
- [167] Dana Scott and Jacobus Willem de Bakker. A theory of programs. *Unpublished manuscript, IBM, Vienna*, 1969.
- [168] Peter Selinger. Towards a quantum programming language. *Mathematical Structures in Computer Science*, 14(4):527–586, 2004.
- [169] Peter Selinger. Generators and relations for n -qubit Clifford operators. *Logical Methods in Computer Science*, 11(2), Jun 2015.
- [170] Peter Selinger and Benoît Valiron. A lambda calculus for quantum computation with classical control. In Paweł Urzyczyn, editor, *Typed Lambda Calculi and Applications*, pages 354–368, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [171] Peter Selinger and Benoît Valiron. A linear-non-linear model for a computational call-by-value lambda calculus (extended abstract). In Roberto M. Amadio, editor, *Proceedings of the 11th International Conference on Foundations of Software Science and Computational Structures (FOSSACS'08)*, volume 4962 of *Lecture Notes in Computer Science*, pages 81–96, Budapest, Hungary, 2008. Springer.
- [172] Peter Selinger and Benoît Valiron. Quantum Lambda Calculus. In Simon J. Gay and Ian Mackie, editors, *Semantic Techniques in Quantum Computation*, pages 135–172. Cambridge University Press, Cambridge, November 2009.
- [173] Yunong Shi, Runzhou Tao, Xupeng Li, Ali Javadi-Abhari, Andrew W Cross, Frederic T Chong, and Ronghui Gu. Certiq: A mostly-automated verification of a realistic quantum compiler. *arXiv preprint arXiv:1908.08963*, 2019.
- [174] Peter W. Shor. Algorithms for quantum computation: Discrete log and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS'94)*, pages 124–134, Santa Fe, New Mexico, US., 1994. IEEE, IEEE Computer Society Press.
- [175] Peter W Shor. Scheme for reducing decoherence in quantum computer memory. *Physical review A*, 52(4):R2493, 1995.
- [176] Daniel R Simon. On the power of quantum computation. *SIAM journal on computing*, 26(5):1474–1483, 1997.
- [177] Kartik Singhal and John Reppy. Quantum hoare type theory: Extended abstract. In Benoît Valiron, Shane Mansfield, Pablo Arrighi, and Prakash Panangaden, editors, *Proceedings 17th International Conference on Quantum Physics and Logic, QPL 2020, Paris, France, June 2 - 6, 2020*, volume 340 of *EPTCS*, pages 291–302, 2020.
- [178] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. t|ket>: a retargetable compiler for NISQ devices. *Quantum Science and Technology*, 6(1):014003, 2020.

- [179] Sam Staton. Algebraic effects, linearity, and quantum programming languages. *SIGPLAN Not.*, 50(1):395–406, January 2015.
- [180] Damian S Steiger, Thomas Häner, and Matthias Troyer. ProjectQ: an open source software framework for quantum computing. *Quantum*, 2(49):10–22331, 2018.
- [181] Krysta M Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. Q#: Enabling scalable quantum computing and development with a high-level domain-specific language. Available online as [arXiv:1803.00652](https://arxiv.org/abs/1803.00652), 2018.
- [182] Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, et al. Dependent types and multi-monadic effects in f. In *Proceedings of the 43rd annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 256–270, 2016.
- [183] Quingo Development Team. Quingo: A programming framework for heterogeneous quantum-classical computing with NISQ features. Draft available as [arXiv:2009.01686](https://arxiv.org/abs/2009.01686), 2020.
- [184] Dominique Unruh. Quantum Hoare logic with ghost variables. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13. IEEE, 2019.
- [185] Dominique Unruh. Quantum relational Hoare logic. *Proc. ACM Program. Lang.*, 3(POPL):33:1–33:31, 2019.
- [186] Benoît Valiron. Generating reversible circuits from higher-order functional programs. In Simon J. Devitt and Ivan Lanese, editors, *Reversible Computation - 8th International Conference, RC 2016, Bologna, Italy, July 7-8, 2016, Proceedings*, volume 9720 of *Lecture Notes in Computer Science*, pages 289–306. Springer, 2016.
- [187] Wim Van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *SIAM Journal on Computing*, 36(3):763–778, 2006.
- [188] André van Tonder. A lambda calculus for quantum computation. *SIAM Journal on Computing*, 33(5):1109–1135, 2004.
- [189] Renaud Vilmart. A near-minimal axiomatisation of ZX-calculus for pure qubit quantum mechanics. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–10, June 2019.
- [190] Renaud Vilmart. The structure of sum-over-paths, its consequences, and completeness for Clifford, 2020. Available online as [arXiv:2003.05678](https://arxiv.org/abs/2003.05678).
- [191] Willem Visser, Corina S. Pasareanu, and Sarfraz Khurshid. Test input generation with Java PathFinder. In *2004 ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA '04*. ACM, 2004.
- [192] Chuan Wang, Fu-Guo Deng, Yan-Song Li, Xiao-Shu Liu, and Gui Lu Long. Quantum secure direct communication with high-dimension quantum superdense coding. *Phys. Rev. A*, 71:044305, Apr 2005.

- [193] Dave Wecker and Krysta M Svore. LIQ*U*i}: A software design architecture and domain-specific language for quantum computing. Available online as [arXiv:1402.4467](https://arxiv.org/abs/1402.4467), 2014.
- [194] Mingsheng Ying. Floyd-Hoare logic for quantum programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 33(6):19:1–19:49, 2011.
- [195] Mingsheng Ying. *Foundations of Quantum Programming*. Morgan Kaufmann, 2016.
- [196] Mingsheng Ying and Yuan Feng. *Model Checking Quantum Systems: Principles and Algorithms*. Cambridge University Press, 2021.
- [197] Mingsheng Ying, Yangjia Li, Nengkun Yu, and Yuan Feng. Model-checking linear-time properties of quantum systems. *ACM Transactions on Computational Logic*, 15(3):22:1–22:31, 2014.
- [198] Mingsheng Ying, Shenggang Ying, and Xiaodi Wu. Invariants of quantum programs: characterisations and generation. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 818–832. ACM, 2017.
- [199] Nengkun Yu and Jens Palsberg. Quantum abstract interpretation. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, pages 542–558, 2021.
- [200] Shengyu Zhang. BQP-complete problems. In Grzegorz Rozenberg, Thomas Bäck, and Joost N. Kok, editors, *Handbook of Natural Computing*, pages 1545–1571. Springer, 2012.
- [201] Li Zhou, Gilles Barthe, Justin Hsu, Mingsheng Ying, and Nengkun Yu. A quantum interpretation of bunched logic for quantum separation logic. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–14. IEEE, 2021.