



HAL
open science

Branch Sequentialization in Quantum Polytime

Emmanuel Hainry, Romain Péchoux, Mário Alberto Machado da Silva

► **To cite this version:**

Emmanuel Hainry, Romain Péchoux, Mário Alberto Machado da Silva. Branch Sequentialization in Quantum Polytime. 2024. hal-04823282v2

HAL Id: hal-04823282

<https://inria.hal.science/hal-04823282v2>

Preprint submitted on 11 Dec 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

Branch Sequentialization in Quantum Polytime

Emmanuel Hainry Romain Péchoux Mário Silva

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

December 11, 2024

Abstract

Quantum computation leverages the use of quantumly-controlled conditionals in order to achieve computational advantage. However, since the different branches in the conditional may operate on the same qubits, a typical approach to compilation involves performing the branches sequentially, which can easily lead to an exponential blowup of the program complexity. We introduce and study a compilation technique for avoiding branch sequentialization in a language that is sound and complete for quantum polynomial time, improving on previously existing polynomial-size bounds and showing the existence of techniques that preserve the intuitive complexity of the program.

1 Introduction

Quantum computing is an emerging paradigm of computation where quantum physical phenomena, such as entanglement and superposition, are used to obtain an advantage over classical computation. A testament to the richness of the field is the variety of computational models: quantum Turing machines [3], quantum circuits [18, 16], measurement-based quantum computation [4, 6], linear optical circuits [13], among others. Some of these models have been shown to be equivalent in terms of computational power and complexity. For instance, Yao’s equivalency result [18] shows that polynomial-time quantum Turing machine are computationally equivalent to uniform and poly-size quantum circuit families.

A lot of effort has been put on developing high-level quantum programming languages to allow programmers to abstract themselves from the technicalities of these low-level models. Towards that end, several verification techniques such as type systems [9] or categorical approaches for reasoning on programs semantics [2, 11] have been studied and developed to ensure the physical reality of compiled programs, for example, by ensuring that it preserves the main properties of quantum mechanics such as no-cloning theorem [1] or unitarity [8]. An important line of research in this area involves checking polytime termination of quantum programs [5, 17, 10].

By Yao’s Theorem, this property implies the feasibility of the corresponding quantum circuit by ensuring that its size is polynomially bounded in the program input size. However, there are still quite a few obstacles to the full use of these techniques. In particular, designing efficient compilation strategies is not trivial [10].

A prominent example is the time complexity of *quantum branching* in programs, i.e., when the flow in a loop or in a conditional is determined upon the state of a qubit. In the classical setting, the cost of branching is the maximum cost between the two branches (Figure 1). However, this is not necessarily the case in the quantum setting, as a consequence of no-cloning: in a quantum circuit, the two branches may contain operations on the same qubits, and thus require an implementation in series. This results in a circuit whose total depth is the sum of the depths of the two branches as illustrated by Figure 1. In [19], while trying to encode efficient operations over quantum data structures, the authors encounter this problem which they have coined *branch sequentialization*. While the authors provide a heuristic for avoiding branch sequentialization, it is only applicable in a few precise examples. Given the importance of preserving the time complexity of a program in its circuit implementation, there is an interest in discovering general techniques that avoid the problem of branch sequentialization altogether.

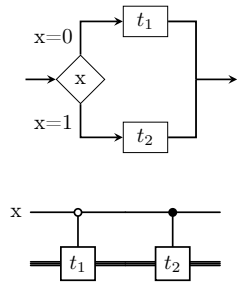


Figure 1: Classical vs. quantum branching.

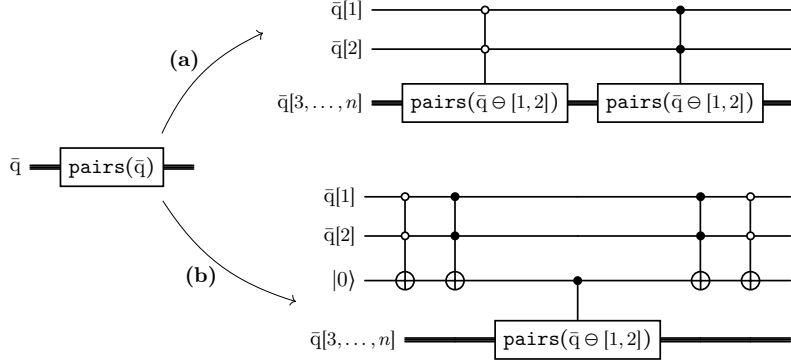


Figure 3: Compilation strategies: branch sequentialization (a) vs optimized approach (b).

1.1 Motivating example

Consider the program `PAIRS` defined in Figure 2. The procedure `pairs` takes as input a sorted set \bar{q} of qubits (i.e., a collection of pairwise distinct qubits) on which it will perform operations. By language design, `pairs` immediately terminates whenever \bar{q} is empty. First, `pairs` checks that the number of qubits in \bar{q} , given by its size $|\bar{q}|$, is larger than 1 to enter the recursive case, otherwise it applies a NOT gate to the remaining qubit (line 9). On line 3, the program will branch depending on the state $\bar{q}[1,2]$ of the first two qubits in \bar{q} . Out of all four cases (lines 4-7), `pairs` only performs an operation when the first two qubits are in state $|00\rangle$ or $|11\rangle$, in which case it performs a recursive call on $\bar{q} \ominus [1,2]$, the sorted set \bar{q} where the first and second qubits have been removed.

With $x \in \{0,1\}^*$ and $y \in \{0,1\}$, given the input state $|xy\rangle$, `pairs` will apply a NOT gate to y if and only if x is a string consisting only of sequences of 00 and 11. Put another way, `pairs` encodes a unitary transformation that inverts the state of the last qubit of an input when x belongs to the regular language defined by $(00 | 11)^*$.

Since `pairs` performs at most one call per branch, and consumes 2 qubits from its input while doing so, we conclude that its runtime complexity is bounded linearly.

Let us now turn to finding a circuit implementation for the recursive case of `pairs`. Consider the two compilation strategies (a) and (b) shown in Figure 3. While Strategy (a) could be considered the more direct approach to building the circuit, at each recursive call the size of the circuit for `pairs` is the sum of the sizes of each branch. On the other hand, while the strategy in (b) requires the creation of an ancilla and the use of extra Toffoli gates, it only requires the implementation of one call to `pairs`. As a consequence, the strategies (a) and (b) result in circuits of depth $\Theta(|\bar{q}|2^{|\bar{q}|})$ and $\Theta(|\bar{q}|)$, respectively, showing how implementing the branches sequentially can result in an exponential blowup in circuit size. It is simple enough to find a compilation strategy that prevents the duplication of `pairs` in the recursive case. However, this becomes much less trivial once we consider programs with more complex recursive calls. *Contribution.* In this paper, we study the problem of branch sequentialization and solve it in the case of quantum polynomial time, in the following way:

- We identify a programming language fragment BFOQ (for Basic FOQ) that is sound and complete for quantum polynomial time (Theorem 2). That is, any BFOQ program computes a function in FBQP, the class of functions computable in polynomial time by a quantum Turing machine with bounded error. Conversely, any function in FBQP can be computed by a BFOQ program. BFOQ is a strict but expressive subset of the PFOQ programming language of [10] whose expressive power is illustrated through many examples (see Table 1);
- We introduce a compilation strategy `compile`⁺ from PFOQ to quantum circuits based on two subroutines `compr`⁺ (Algorithm 1) and `optimize`⁺ (Algorithm 2): while `compr`⁺ just generates the compiled circuit by a simple structural induction on program statements, `optimize`⁺

```

1  decl pairs( $\bar{q}$ ) {
2    if  $|\bar{q}| > 1$  then
3      qcase  $\bar{q}[1,2]$  of {
4        00  $\rightarrow$  call pairs( $\bar{q} \ominus [1,2]$ );
5        01  $\rightarrow$  skip;
6        10  $\rightarrow$  skip;
7        11  $\rightarrow$  call pairs( $\bar{q} \ominus [1,2]$ );
8      }
9    else  $\bar{q}[1] \neq \text{NOT}$ ; }
10 :: call pairs( $\bar{q}$ );

```

Figure 2: Branching program PAIRS.

(Programs)	$P(\bar{q})$	\triangleq	$D :: S$
(Procedure declarations)	D	\triangleq	$\varepsilon \mid \mathbf{decl} \text{ proc}[x](\bar{q})\{S\}, D$
(Statements)	S	\triangleq	$\mathbf{skip}; \mid \bar{q}[i] \mathbf{*} = U^f(j); \mid S S \mid \mathbf{if} \ b \ \mathbf{then} \ S \ \mathbf{else} \ S$ $\mid \mathbf{qcase} \ \bar{q}[i] \ \mathbf{of} \ \{0 \rightarrow S, 1 \rightarrow S\} \mid \mathbf{call} \ \text{proc}[i](s);$

Figure 4: Syntax of FOQ programs.

perform some optimization by merging (recursive) procedure calls in different branches in the program.

- We show that the **compile**⁺ is sound, i.e., the generated circuit fairly simulates the input program: this correctness result lies on the orthogonality of the control structures used in the **optimize**⁺ subroutine (Lemma 2).
- On PFOQ programs, we exhibit a direct improvement on size complexity of the generated circuit with respect to the compilation algorithm studied in [10] (Theorem 3).
- We show that, on BFOQ programs, **compile**⁺ produces circuits whose size is asymptotically bounded by their level (Theorem 4), i.e., by the maximal number of consecutive procedure calls in all branches (including quantum ones) of a program execution, thus avoiding branch sequentialization on a sound and complete language for quantum polynomial time.

1.2 Related work

Resource optimization in quantum computing is a well-studied subject for low level computational models such as quantum circuits or ZX-diagrams: in this *constant-depth* scenario, (i.e., taking a specific and fixed circuit of constant size and, thus, constant depth), is it possible to reduce its number of gates [15, 14] (or at least its number of non-Clifford gates [12, 7]), with techniques such as gate substitution, graph-rewriting, among others.

Resource optimization for high-level quantum programs is still a relatively undeveloped research area as it involves the asymptotic consideration of families of circuits. Such an issue has strong connections with programming language-based characterizations of quantum polynomial time classes [17, 5, 10] as, by design, their set of programs is sound and complete for uniform families of quantum circuits of polynomial size, as per Yao’s equivalency theorem. While [17, 5] provide non-constructive proofs of the existence of quantum circuits of polynomial size, [10] introduces a programming language that avoids an exponential blowup in the complexity of recursive (quantum) branching with a direct compilation strategy for ensuring polysized circuit representations. However this strategy still performs branch sequentialization and generates polynomial bounds whose degree is not accurate.

2 First-Order Quantum Programming Language

We consider the FOQ (First-Order Quantum) programming language with quantum control, introduced in [10] to characterize quantum polynomial time. A complete account of its syntax and its semantics is given in Appendix A.

2.1 Syntax

A FOQ program $P \triangleq D :: S$ is defined in Figure 4 by a list of procedure declarations D and a program statement S . The language include 4 basic datatypes for expressions. *Sorted set* expressions s are either variables \bar{q} , the empty sorted set nil , or $s \ominus [i]$, the sorted set s where the i -th element has been removed. Intuitively, a sorted set is a list of unique (i.e., non-duplicable) qubit pointers. *Integer* expressions, noted i, j , are either an integer variable x , a constant n , an addition by a constant $i \pm n$ or the size of a sorted set $|s|$. *Boolean* expressions b are defined in a standard way using boolean operators and arithmetic operators, e.g., $i > j$. Finally, *qubit* expressions are of the shape $s[i]$ which denotes the i -th qubit pointed to in s . $s[i_1, \dots, i_n]$ is a shorthand for $s[i_1], \dots, s[i_n]$. Finally, we also allow for the syntactic sugar on sorted state of pointing to the n -th *last* qubit in the set, by defining for any $n \geq 1$, $\bar{q}[-n] \triangleq \bar{q}[|\bar{q}| - n + 1]$.

A procedure of name proc is defined by a procedure declaration $\mathbf{decl} \text{proc}[x](\bar{q})\{S^{\text{proc}}\}$ which takes a sorted set \bar{q} and an (optional) integer x as input parameters and executes the *procedure statement* S^{proc} . Let Procedures be an enumerable set of procedure names. We will write S instead of S^{proc} when the procedure is clear from context, and we denote by $\text{proc} \in P$ the fact that proc appears in D . Given two statements $S, S', S \in S'$ denotes the fact that S is a substatement of S' . Furthermore, we have that $\text{proc} \in S$ holds if there are i and s such that $\mathbf{call} \text{proc}[i](s); \in S$.

Statements include the no-op instruction, unitary operations, sequences, classical and quantum conditionals, and procedures calls. Of these, we highlight the quantum conditional $\mathbf{qcase} \bar{q}[i] \mathbf{of} \{0 \rightarrow S_0, 1 \rightarrow S_1\}$, which allows branching by executing statements S_0 and S_1 in superposition according to the state of qubit $\bar{q}[i]$, and also the procedure call $\mathbf{call} \text{proc}[i](s);$, which runs procedure proc with *integer expression* i and *sorted set expression* s , a list of unique qubit pointers. The quantum conditional can be extended to n qubits $\mathbf{qcase} \bar{q}[i_1, \dots, i_n] \mathbf{of} \{0^n \rightarrow S_{0^n}, \dots, 1^n \rightarrow S_{1^n}\}$ in a standard way as used in Figure 2.

In a statement $\bar{q}[i] *= U^f(j);$, if the integer expression j evaluates to n , then the unitary operator $\llbracket U^f \rrbracket(n)$ corresponding to the unary construct $U^f(j)$ is applied to qubit $\bar{q}[i]$. For expressivity purposes, these constructs are parameterized by some polynomial-time approximable total function $f \in \mathbb{Z} \rightarrow [0, 2\pi)$ and some integer expression j . For example, the gates of the quantum Fourier transform can be defined by $R_n \triangleq \llbracket \text{Ph}^{\lambda x \cdot \pi / 2^{x-1}} \rrbracket(n)$ with $\llbracket \text{Ph}^f \rrbracket(n) \triangleq \begin{pmatrix} 1 & 0 \\ 0 & e^{if(n)} \end{pmatrix}$. Other basic unary gates are the *NOT* and the *RY* gate (see [10]). We also make use of some syntactic sugar to describe statements encoding constant-time quantum operations. For instance, the *CNOT*, *SWAP*, and Toffoli gates can be defined by:

$$\begin{aligned} \text{CNOT}(\bar{q}[i], \bar{q}[j]) &\triangleq \mathbf{qcase} \bar{q}[i] \mathbf{of} \{0 \rightarrow \mathbf{skip};, 1 \rightarrow \bar{q}[j] *= \text{NOT}\} \\ \text{SWAP}(\bar{q}[i], \bar{q}[j]) &\triangleq \text{CNOT}(\bar{q}[i], \bar{q}[j]) \text{CNOT}(\bar{q}[j], \bar{q}[i]) \text{CNOT}(\bar{q}[i], \bar{q}[j]) \\ \text{TOF}(\bar{q}[i], \bar{q}[j], \bar{q}[k]) &\triangleq \mathbf{qcase} \bar{q}[i] \mathbf{of} \{0 \rightarrow \mathbf{skip};, 1 \rightarrow \text{CNOT}(\bar{q}[i], \bar{q}[j])\} \end{aligned}$$

We define notions of rank that provide quantitative information on the recursion level of a given program or procedure.

Definition 1 (Rank). *Given a FOQ program P , the rank of a procedure proc in P , denoted $rk_P(\text{proc})$, is defined as follows:*

$$rk_P(\text{proc}) \triangleq \begin{cases} 0, & \text{if } \neg(\exists \text{proc}', \text{proc} \geq_P \text{proc}'), \\ \max_{\text{proc} \geq_P \text{proc}'} \{rk_P(\text{proc}')\}, & \text{if } \exists \text{proc}', \text{proc} \geq_P \text{proc}' \wedge \neg(\text{proc} \sim_P \text{proc}'), \\ 1 + \max_{\text{proc} >_P \text{proc}'} \{rk_P(\text{proc}')\}, & \text{if } \text{proc} \sim_P \text{proc}, \end{cases}$$

where $\max(\emptyset) \triangleq 0$. The rank of a program is defined as the maximum rank among all procedures, i.e., for a program $P \triangleq D :: S$, we have that $rk(P) \triangleq \max_{\text{proc} \in D} rk_P(\text{proc})$.

Example 1. *The program PAIRS given in Figure 2 has rank 1, since*

$$rk(\text{PAIRS}) \triangleq \max_{\text{proc} \in \text{PAIRS}} rk_P(\text{proc}) = rk(\text{pairs}) = 1.$$

2.2 Semantics

Let \mathcal{H}_{2^n} denote the Hilbert space of n qubits \mathbb{C}^{2^n} , $\mathcal{L}(\mathbb{N})$ denote the set of lists of natural numbers, and $\mathcal{P}(\mathbb{N})$ denote the powerset of natural numbers.

Expressions

For $\mathbb{K} \in \{\mathbb{N}, \mathbb{Z}, \mathbb{C}^{2 \times 2}, \mathcal{L}(\mathbb{N})\}$, we write $(e, l) \Downarrow_{\mathbb{K}} v$ when the expression e evaluates to the value $v \in \mathbb{K}$ under the context $l \in \mathcal{L}(\mathbb{N})$. The context l is just the sorted set of qubit pointers into consideration when evaluating e . For example, we have that $(\bar{q}[2], [1, 4, 5]) \Downarrow_{\mathbb{N}} 4$ (the second qubit is of index 4), $(\bar{q}[4], [1, 4, 5]) \Downarrow_{\mathbb{N}} 0$ (index 0 is used for error), and $(\bar{q} \ominus [3], [1, 4, 5]) \Downarrow_{\mathcal{L}(\mathbb{N})} [1, 4]$ (the third qubit has been removed).

Statements

Let $\text{Conf}_n \triangleq (\text{Statements} \cup \{\top, \perp\}) \times \mathcal{H}_{2^n} \times \mathcal{P}(\mathbb{N}) \times \mathcal{L}(\mathbb{N})$ be the set of configurations of n qubits. In a configuration $c \triangleq (S, |\phi\rangle, \mathcal{S}, l)$, S is the statement to be executed, $|\phi\rangle$ is the quantum state, \mathcal{S} is a set of accessible (pointers to) qubits and l is the list of qubit pointers under consideration. In case of error the program exits and the two special symbols \top and \perp are markers for success (termination) and failure (error), respectively. The set \mathcal{S} of accessible qubits is used to ensure that unitary operations on qubits can be physically implemented. For example, statements S_0 and S_1 of a quantum branch **qcase** $\bar{q}[i]$ **of** $\{0 \rightarrow S_0, 1 \rightarrow S_1\}$ cannot access $\bar{q}[i]$ to ensure that the operation can be physically implemented by a controlled-circuit.

The big-step semantics $\cdot \xrightarrow{\cdot} \cdot$ is defined as a relation in $\bigcup_{n \in \mathbb{N}} \text{Conf}_n \times \mathbb{N} \times \text{Conf}_n$. When $c \xrightarrow{m} c'$ holds the *level* m is an integer corresponding to the maximum number of procedure calls performed over each (condition and quantum) branch during the evaluation of c . More formally, the level of a program $P = D :: S$ on n qubits, denoted $\text{level}_P(n)$, is defined by

$$\text{level}_P(n) \triangleq \max\{m \in \mathbb{N} \mid \exists |\phi\rangle, |\phi'\rangle \in \mathcal{H}_{2^n}, (S, |\phi\rangle, \mathcal{S}_n, l_n) \xrightarrow{m} (\top, |\phi'\rangle, \mathcal{S}_n, l_n)\},$$

with $\mathcal{S}_n \triangleq \{1, \dots, n\}$ and $l_n \triangleq [1, \dots, n]$.

Example 2. Consider the program *PAIRS* of Figure 2. We have that each procedure call removes two qubits until it reaches a case of input size $|\bar{q}|$ either 1 or 0 (depending on if n is odd or even) and for both sizes there are no more procedure calls. On an empty sorted set, the program exits after the first call. Then, $\text{level}_{\text{PAIRS}}(n) = \lfloor \frac{n}{2} \rfloor + 1 = O(n)$.

2.3 Polytime fragments of FOQ

In [10], the polynomial-time fragment of FOQ, denoted PFOQ, is defined by placing two restrictions on procedure calls: a well-foundedness criterion for termination and a restriction on the number of admissible recursive calls per (classical or quantum) branch to avoid exponentiation.

PFOQ and Basic FOQ. Given a program $P \triangleq D :: S$, the call relation $\rightarrow_P \subseteq \text{Procedures} \times \text{Procedures}$ is defined for any two procedures $\text{proc}_1, \text{proc}_2 \in S$ as $\text{proc}_1 \rightarrow_P \text{proc}_2$ whenever $\text{proc}_2 \in S^{\text{proc}_1}$. The partial order \geq_P is then the transitive closure of \rightarrow_P , and \sim_P denotes the equivalence relation where $\text{proc}_1 \sim_P \text{proc}_2$ if $\text{proc}_1 \geq_P \text{proc}_2$ and $\text{proc}_2 \geq_P \text{proc}_1$ both hold. **call** $\text{proc}' \in S^{\text{proc}}$ is a *recursive* procedure call whenever $\text{proc} \sim_P \text{proc}'$. The strict order $>_P$ is defined as $\text{proc}_1 >_P \text{proc}_2$ if $\text{proc}_1 \geq_P \text{proc}_2$ and $\text{proc}_1 \not\sim_P \text{proc}_2$ both hold.

A program P is in WF if it satisfies the constraint that each recursive procedure call removes at least one qubit in its parameter. Programs in WF are terminating (well-founded) but still allow the programmer to write programs with exponential runtime. To avoid such programs, we use the notion of *width of a procedure* proc in a program P , noted $\text{width}_P(\text{proc})$, defined inductively on procedure declarations by counting the number of recursive calls in the procedure body, taking the maximum of each (classical or quantum) conditional.

Definition 2 (Width of a procedure). Given $P \in \text{FOQ}$ and $\text{proc} \in P$, the width of proc in P , noted $\text{width}_P(\text{proc})$, is defined as $\text{width}_P(\text{proc}) \triangleq w_P^{\text{proc}}(S^{\text{proc}})$, where $w_P^{\text{proc}}(S)$ is the width of the procedure proc in P relative to statement S , defined inductively as:

$$\begin{aligned} w_P^{\text{proc}}(\text{skip};) &= w_P^{\text{proc}}(q \text{ *}= U^f(i);) \triangleq 0, \\ w_P^{\text{proc}}(S_1 S_2) &\triangleq w_P^{\text{proc}}(S_1) + w_P^{\text{proc}}(S_2), \\ w_P^{\text{proc}}(\text{if } b \text{ then } S_{\text{true}} \text{ else } S_{\text{false}}) &\triangleq \max(w_P^{\text{proc}}(S_{\text{true}}), w_P^{\text{proc}}(S_{\text{false}})), \\ w_P^{\text{proc}}(\text{qcase } q \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\}) &\triangleq \max(w_P^{\text{proc}}(S_0), w_P^{\text{proc}}(S_1)), \\ w_P^{\text{proc}}(\text{call } \text{proc}'[i](s);) &\triangleq \begin{cases} 1 & \text{if } \text{proc} \sim_P \text{proc}' \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Let $\text{WIDTH}_{\leq 1}$ be the set of programs with procedures of width at most 1. Finally, define $\text{PFOQ} \triangleq \text{FOQ} \cap \text{WF} \cap \text{WIDTH}_{\leq 1}$. We will show that PFOQ can be restricted to a fragment, denoted BFOQ (for Basic FOQ), that is still complete for polynomial time and where the compilation procedure conserves the intuitive complexity of the program. Let BASIC denote the set of programs where i) procedures do not use classical inputs ii) procedure call parameters are restricted to sorted set variables \bar{q} or to a fixed sorted set s (s is fixed for each BFOQ program). Then, we define $\text{BFOQ} \triangleq \text{BASIC} \cap \text{PFOQ}$. It trivially holds that $\text{BFOQ} \subsetneq \text{PFOQ} \subsetneq \text{FOQ}$.

Example 3 (Quantum Full Adder \in BFOQ). Let *ADD* encode the unitary transformation such that, for $a, b \in \{0, 1\}^n$ and $c_{\text{in}}, c_{\text{out}} \in \{0, 1\}$, *ADD* performs the following transformation, where η represents the n least significant bits of $(a + b + c_{\text{in}})$.

$$\text{ADD} |a_n b_n \dots a_1 b_1 0^n c_{\text{in}}\rangle \triangleq |a_n b_n \dots a_1 b_1 c_{\text{out}} \eta_1 \dots \eta_n\rangle,$$

where c_{in} and c_{out} encode the carry-in and carry-out values, respectively. The operator *ADD* corresponds to the circuit in Figure 6 and is described by the program *FA* below.

```

1  decl fullAdder( $\bar{q}$ ) {
2    if  $|\bar{q}| > 3$  then /*  $\bar{q}[1] = a, \bar{q}[2] = b, \bar{q}[-2] = |0\rangle$  and  $\bar{q}[-1] = c_{\text{in}}$  */
3      TOF( $\bar{q}[1], \bar{q}[2], \bar{q}[-2]$ )
4      CNOT( $\bar{q}[1], \bar{q}[2]$ )
5      TOF( $\bar{q}[2], \bar{q}[-1], \bar{q}[-2]$ ) /*  $c_{\text{out}} = (a \cdot b) \oplus (c_{\text{in}} \cdot (a \oplus b))$  */
6      CNOT( $\bar{q}[2], \bar{q}[-1]$ ) /*  $\eta = a \oplus b \oplus c_{\text{in}}$  */
7      CNOT( $\bar{q}[1], \bar{q}[2]$ )
8      call fullAdder( $\bar{q} \ominus [1, 2, -1]$ );
9    else skip; }
10 :: call fullAdder( $\bar{q}$ );

```

It holds that $\text{FA} \in \text{FOQ} \cap \text{WF}$ and that $\text{width}_{\text{FA}}(\text{fullAdder}) = 1$. Therefore, $\text{FA} \in \text{PFOQ}$. Also, *FA* is clearly in BFOQ as there is only one recursive call and no integer parameter.

Example 4 (Quantum Fourier Transform \in BFOQ). The quantum Fourier transform can be described by the program *QFT* below

```

1  decl qft( $\bar{q}$ ) {
2     $\bar{q}[1] *= \text{H}$ ;
3    call rot( $\bar{q}$ );
4    call shift( $\bar{q}$ );
5    call qft( $\bar{q} \ominus [-1]$ ); }
6  decl shift( $\bar{q}$ ) {
7    if  $|\bar{q}| > 1$  then
8      SWAP( $\bar{q}[1], \bar{q}[-1]$ )
9      call shift( $\bar{q} \ominus [-1]$ );
10   else skip; }
11 decl rot( $\bar{q}$ ) {
12   if  $|\bar{q}| > 1$  then
13     qcase  $\bar{q}[-1]$  of {
14       0  $\rightarrow$  skip;
15       1  $\rightarrow \bar{q}[1] *= \text{Ph}^{\lambda x \cdot \pi / 2^{x-1}}(|\bar{q}|)$ ; }
16     call rot( $\bar{q} \ominus [-1]$ );
17   else skip; }
18 :: call qft( $\bar{q}$ );

```

The program consists of three procedures, *qft*, *shift*, and *rot*, and is in PFOQ since it is in WF and $\text{width}_{\text{QFT}}(\text{qft}) = \text{width}_{\text{QFT}}(\text{shift}) = \text{width}_{\text{QFT}}(\text{rot}) = 1$. All procedure calls are performed on the set \bar{q} or $\bar{q} \ominus [-1]$, and therefore the program is also in BASIC. This program can be compiled to the circuit Figure 5 for input size 4, implementing the quantum Fourier transform. This circuit differs from (but is equivalent to) the standard implementation of the quantum Fourier transform. This is due to some of the restrictions put in BFOQ. The standard circuit can be obtained directly through compilation of a PFOQ program.

Properties of PFOQ and BFOQ programs.

PFOQ programs have a consecutive number of procedure calls in the distinct branches of their execution (i.e., level) that is bounded polynomially in the input size (number of qubits). The degree of the polynomial can be obtained from the rank, which can be inferred syntactically.

Lemma 1 (Polynomial level [10]). For any PFOQ program P , $\text{level}_P(n) = O(n^{rk(P)})$.

The set PFOQ of programs was shown to be sound and complete for the class FBQP, of function computable in quantum polynomial time [10].

Theorem 1 (PFOQ-Soundness and Completeness [10]). For every function f in FBQP, there is a PFOQ program P that computes f with probability $\frac{2}{3}$ using at most a polynomial number of extra ancilla. Conversely, given a program P in PFOQ, if P computes $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, with probability $p \in (\frac{1}{2}, 1]$ then $f \in \text{FBQP}$.

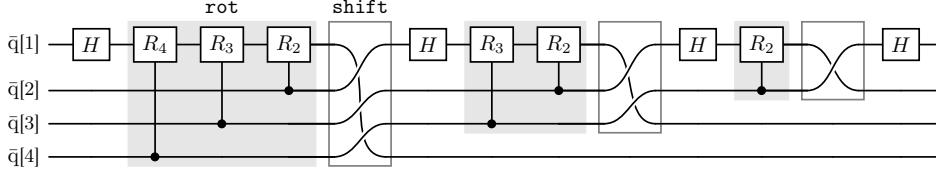


Figure 5: Circuit for the QFT as defined by the program in Example 4.

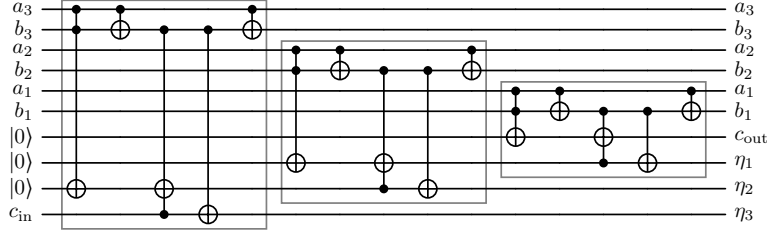


Figure 6: Quantum Full Adder circuit (Example 3) for input size 10.

Being a strict subset of PFOQ, BFOQ is trivially sound. Surprisingly, it is also complete.

Theorem 2 (BFOQ-Soundness and Completeness). *Theorem 1 holds for BFOQ.*

Proof. Soundness is trivial since BFOQ is contained in PFOQ, and completeness is given analogously to the proof of [10, Theorem 5], by noticing that all programs constructed in the proof are also contained in BFOQ. \square

3 Circuit compilation

In this section, we introduce a compilation strategy for PFOQ programs that strictly improves on the compilation algorithm of [10]. We also show that, in the BFOQ fragment, circuit complexity scales in such a way that the cost of branching is the maximum cost of each branch, thereby avoiding branch sequentialization.

3.1 A new compilation algorithm

The compilation algorithm **compile**⁺ takes as input a program P and a natural number n (the number of input qubits) and returns a circuit implementation of P for an input size of n qubits. **compile**⁺ is defined by its subroutine **compr**⁺ (Algorithm 1) in the following way:

$$\mathbf{compile}^+(P, n) \triangleq \mathbf{compr}^+(P, [1, \dots, n], \cdot, \{\}),$$

where P is the program to be compiled, $[1, \dots, n]$ is list of qubit pointers (initially all qubits), \cdot is an empty control structure, and $\{\}$ an empty dictionary. A *control structure* is a partial function in $\mathbb{N} \rightarrow \{0, 1\}$ mapping qubit pointers to their control values in a quantum case. For $n \in \mathbb{N}$ and $k \in \{0, 1\}$, $cs[n := k]$ is the control structure obtained from cs by setting $cs(n) \triangleq k$. We denote by $dom(cs)$ the domain of the control structure. For a given $x \in \{0, 1\}^*$, we say that state $|x\rangle$ *satisfies* cs if, $\forall n \in dom(cs)$, $cs(n) = k$ implies that $x_n = k$. The purpose of control structures in the algorithm is to preserve the information of each quantum branch and to allow for merging using ancillas.

The aim of subroutine **compr**⁺ is just to generate the quantum circuit corresponding to P on n qubits inductively on the statement of P. When the analyzed statement is a (possibly recursive) procedure call, **compr**⁺ calls the **optimize**⁺ subroutine (Algorithm 2) to perform an optimization of the generated quantum circuit. **optimize**⁺ has the same inputs as **compr**⁺ with the addition of a list of *controlled statements* l_{Cst} and the name proc of the procedure under analysis. A controlled statement is defined as a pair (cs, S) where cs is a control structure and S is a FOQ statement. This will allow us to generate the circuit implementation of S (which may contain multiple gates) while keeping track of the branch on which it is implemented.

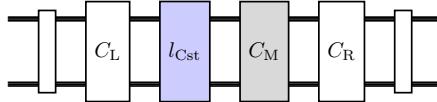
Algorithm 1 (`compr+`)**Input:** $(D :: S, l, cs, \text{Anc}) \in \text{Programs} \times \mathcal{L}(\mathbb{N}) \times (\mathbb{N} \rightarrow \{0, 1\}) \times \mathcal{D}$

```
1: if S = skip; then
2:   C ← 1 ▷ Identity circuit
3:
4: else if S = s[i] *= Uf(j); and (s[i], l) ↓ℕ n and (Uf(j), l) ↓ℂ2×2 M then
5:   C ← M(cs, [n]) ▷ Controlled gate
6:
7: else if S = S1 S2 then
8:   C ← compr+(D :: S1, l, cs, Anc) ◦ compr+(D :: S2, l, cs, Anc) ▷ Composition
9:
10: else if S = if b then Strue else Sfalse and (b, l) ↓ℝ b then
11:   C ← compr+(D :: Sb, l, cs, Anc) ▷ Conditional
12:
13: else if S = qcase s[i] of {0 → S0, 1 → S1} and (s[i], l) ↓ℕ n then ▷ Quantum case
14:   C ← compr+(D :: S0, l, cs[n := 0], Anc) ◦ compr+(D :: S1, l, cs[n := 1], Anc)
15:
16: else if S = call proc[i](s); and (s, l) ↓ℒ(ℕ) [] then
17:   C ← 1 ▷ Nil call
18:
19: else if S = call proc[i](s); and (s, l) ↓ℒ(ℕ) l' ≠ [] and (i, l) ↓ℤ n then ▷ Procedure call
20:   a ← new ancilla()
21:   Anc[proc', n, |l'|] ← (a, l');
22:   C ← optimize+(D, [(·[a = 1], Sproc{n/x}], proc, l', Anc)
23: end if
24: return C
```

The compilation algorithm `compile+` is based on the process of merging procedure calls by reasoning about the orthogonality relations within the circuit. It is similar to the compilation algorithm `compile` of [10] based on the subroutines `compr` and `optimize`, with the differences highlighted in the code of Algorithms 1 and 2: `optimize+` strictly improves on `optimize` (Theorem 3) by extending this analysis to procedures of different ranks.

3.2 Soundness and optimization

The correctness of `optimize+` is a consequence of an orthogonality property between elements of the circuit being compiled that remains invariant throughout the compilation. In Algorithm 2, a recursive procedure `proc` is compiled by generating three separate circuits C_L , C_M , and C_R . The compilation process makes use of a list l_{Cst} of controlled statements which have not yet been compiled. Given a program P and an input procedure `proc`, the list l_{Cst} contains controlled statements (cs, S) such that $w_P^{\text{proc}}(S) = 1$, whereas C_M contains circuits for statements such that $w_P^{\text{proc}}(S) = 0$ but that are nonetheless orthogonal to those in l_{Cst} , for which merging will be possible. Circuits C_L and C_R contain circuits for statements for which $w_P^{\text{proc}}(S) = 0$ and that are not orthogonal to the elements of l_{Cst} . The circuit below pictures how the output circuit is generated from the circuits C_L , C_M , and C_R and the list l_{Cst} .



The steps of `optimize+` are depicted in Figure 7, where in each case we treat a controlled statement $(cs, S) \in l_{\text{Cst}}$. A gate placed inside the violet box \blacksquare denotes the new controlled statement that replaces (cs, S) in l_{Cst} . A gate placed inside a grey box \blacksquare indicates a circuit that is compiled and added to C_M . The notation is agnostic to the precise placement of these objects within l_{Cst} and C_M , making use of the orthogonality relation described in Lemma 2 which renders the choice inconsequential. Figures 7a, 7b, and 7c contain two circuits consisting of the different possible cases of compilation:

Algorithm 2 (optimize^+) Build circuit for recursive procedure proc
Inputs: $(D, l_{\text{Cst}}, \text{proc}, l, \text{Anc}) \in \text{Decl} \times \mathcal{L}(\text{Cst}) \times \text{Procedures} \times \mathcal{L}(\mathbb{N}) \times \mathcal{D}$

```

1:  $C_L \leftarrow \mathbb{1}; C_R \leftarrow \mathbb{1}; C_M \leftarrow \mathbb{1}; P \leftarrow D :: \text{skip};$ 
2: while  $l_{\text{Cst}} \neq []$  do
3:    $(cs, S) \leftarrow \text{hd}(l_{\text{Cst}}); l_{\text{Cst}} \leftarrow \text{tl}(l_{\text{Cst}})$ 
4:
5:   if  $S = S_1 S_2$  then
6:      $\text{Anc}' \leftarrow \text{Anc.copy}()$  /* create copy of ancilla dictionary */
7:     if  $w_P^{\text{proc}}(S_1) = 1$  then
8:        $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs, S_1)]; C_M \leftarrow \text{compr}^+(D :: S_2, l, cs, \text{Anc}') \circ C_M$ 
9:     else
10:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs, S_2)]; C_M \leftarrow \text{compr}^+(D :: S_1, l, cs, \text{Anc}') \circ C_M$ 
11:    end if
12:  end if
13:
14:  if  $S = \text{if } b \text{ then } S_{\text{true}} \text{ else } S_{\text{false}}$  and  $(b, l) \Downarrow_{\mathbb{B}} b$  then
15:    if  $w_P^{\text{proc}}(S_b) = 1$  then
16:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs, S_b)]$ 
17:    else
18:       $C_M \leftarrow \text{compr}^+(D :: S_b, l, cs, \text{Anc}) \circ C_M$ 
19:    end if
20:  end if
21:
22:  if  $S = \text{qcase } s[i] \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\}$  and  $(s[i], l) \Downarrow_{\mathbb{N}} n$  then
23:    if  $w_P^{\text{proc}}(S_0) = 1$  and  $w_P^{\text{proc}}(S_1) = 1$  then
24:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs[n := 0], S_0), (cs[n := 1], S_1)]$ 
25:    else if  $w_P^{\text{proc}}(S_1) = 0$  then
26:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs[n := 0], S_0)];$ 
27:       $C_M \leftarrow \text{compr}^+(D :: S_1, l, cs[n := 1], \text{Anc}) \circ C_M$ 
28:    else if  $w_P^{\text{proc}}(S_0) = 0$  then
29:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [(cs[n := 1], S_1)];$ 
30:       $C_M \leftarrow \text{compr}^+(D :: S_0, l, cs[n := 0], \text{Anc}) \circ C_M$ 
31:    end if
32:  end if
33:
34:  if  $S = \text{call proc}'[i](s)$  and  $(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} l' \neq []$  and  $(i, l) \Downarrow_{\mathbb{Z}} n$  then
35:    if  $(\text{proc}', n, |l'|) \in \text{Anc}$  then
36:      Let  $(a, l'') = \text{Anc}[\text{proc}', n, |l'|]$  in
37:       $e \leftarrow \text{new ancilla}()$ ; /* compatible procedure already exists: merging case */
38:       $C_L \leftarrow C_L \circ \text{NOT}(cs, e) \circ \text{NOT}(\cdot[e = 1], a) \circ \text{SWAP}(\cdot[e = 1], l', l'')$ ;
39:       $C_R \leftarrow \text{SWAP}(\cdot[e = 1], l'', l') \circ \text{NOT}(\cdot[e = 1], a) \circ \text{NOT}(cs, e) \circ C_R$ 
40:    else
41:       $a \leftarrow \text{new ancilla}()$  /* no compatible procedure: create new ancilla */
42:       $\text{Anc}[\text{proc}', n, |l'|] \leftarrow (a, l')$ ;
43:       $C_L \leftarrow C_L \circ \text{NOT}(cs, a); C_R \leftarrow \text{NOT}(cs, a) \circ C_R;$ 
44:       $l_{\text{Cst}} \leftarrow l_{\text{Cst}} @ [\cdot[a = 1], S^{\text{proc}'}\{n/x\}]$ 
45:    end if
46:  end if
47: end while
48: return  $C_L \circ C_M \circ C_R$ 

```

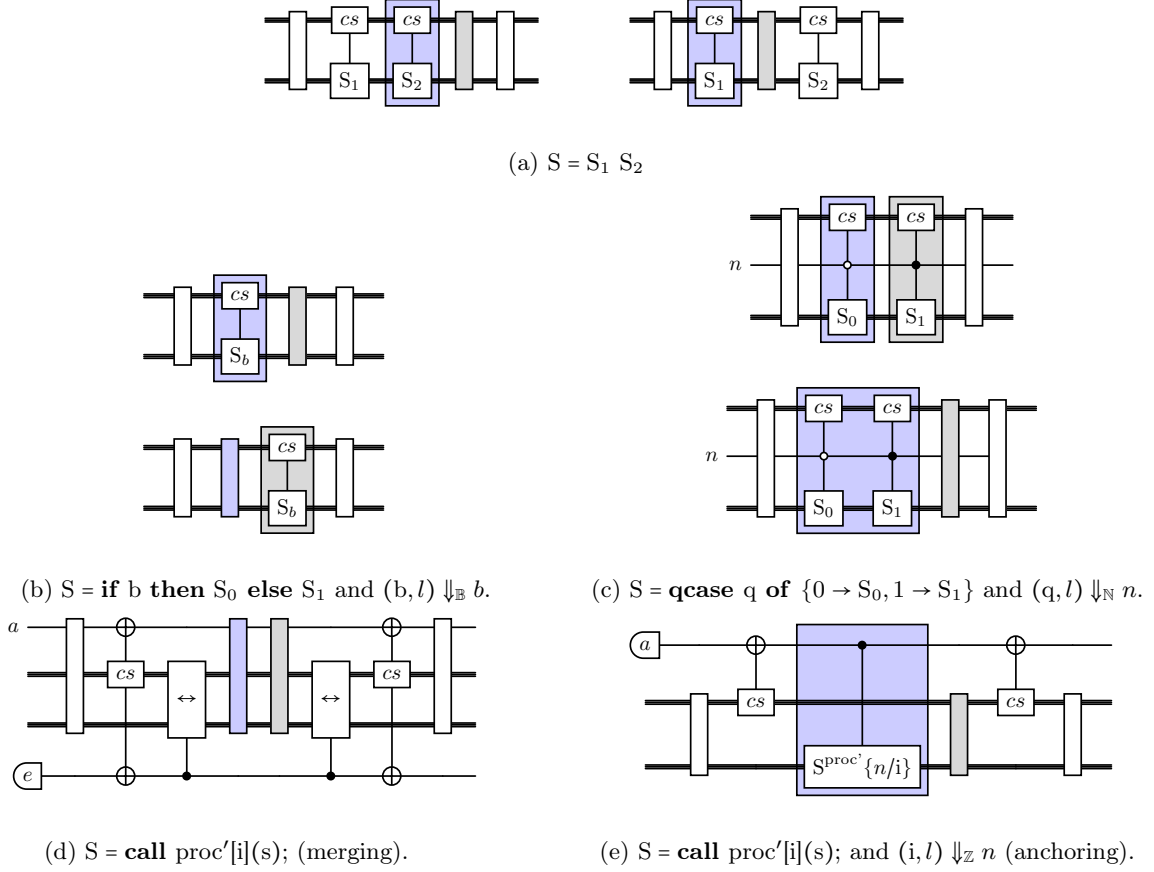


Figure 7: A step of the **optimize**⁺ subroutine.

- In Figure 7a, we consider a sequence $S_1 S_2$. This includes the case where S_2 is recursive (left) and the one where S_1 is recursive (right). Given the $\text{WIDTH}_{\geq 1}$ condition, there are no more cases.
- In Figure 7b, the case of classical control, we have the step where S_b contains a recursive call (above) and where it does not (below).
- For the case of quantum branching (Figure 7c), it is possible that only one of the two statements, say S_0 , contains a recursive call (left) and the case where both do (right).
- Finally, we consider the case of a procedure call. Either there is already a compatible procedure and merging is performed (Figure 7d) or a new ancilla is used to anchor the procedure (Figure 7e).

We formalize orthogonality between controlled statements as follows.

Definition 3 (Orthogonality between control structures). *We say that two control structures are orthogonal, also denoted $cs \perp cs'$, if $\exists i \in \mathbb{N}$ such that $i \in \text{dom}(cs) \cap \text{dom}(cs')$ and where $cs(i) + cs'(i) = 1$.*

Hence, two control structures are orthogonal if there is no base state that satisfies them both. We now show that the steps of **optimize**⁺ respect an orthogonality invariant.

Lemma 2 (Orthogonality invariant). *At each step of the subroutine **optimize**⁺, the list l_{Cst} and the circuit C_M satisfy the following properties:*

1. All controlled statements in l_{Cst} are mutually orthogonal:

$$\forall (cs, S), (cs', S') \in l_{\text{Cst}} \text{ such that } (cs, S) \neq (cs', S'), \text{ we have that } cs \perp cs'.$$

2. Any controlled statement in l_{Cst} commutes with any element of C_M :

$$\forall (cs, S) \in l_{\text{Cst}}, \forall M(cs', [n]) \in C_M \text{ we have that } cs \perp cs'.$$

Proof. We start **optimize**⁺ with a single procedure statement and an empty control structure, i.e. C_M is empty and $l_{\text{Cst}} = \{(\cdot, S^{\text{proc}})\}$, in which case the lemma is clearly true.

We now prove by induction that it is an invariant. Let $(cs, S) \in l_{\text{Cst}}$ be the controlled statement being treated. If $S = S_1 S_2$, let $w_P^{\text{proc}}(S_1) = 1$ and $w_P^{\text{proc}}(S_2) = 0$. Then, (cs, S) is replaced with (cs, S_1) in l_{Cst} and C_M is unchanged – therefore, the invariant property remains true. The case where $w_P^{\text{proc}}(S_1) = 0$ and $w_P^{\text{proc}}(S_2) = 1$ is analogous.

If $S = \mathbf{if\ b\ then\ } S_0 \mathbf{\ else\ } S_1$, then consider the case $(b, l) \Downarrow_{\mathbb{B}} 0$, where if $w_P^{\text{proc}}(S_0) = 1$ we have that (cs, S) is replaced with (cs, S_0) in l_{Cst} and C_M remains the same, and therefore the property remains true. If $w_P^{\text{proc}}(S_0) = 0$ we have that (cs, S) is removed from l_{Cst} and $C_M \leftarrow \llbracket (cs, S) \rrbracket \circ C_M$. By the induction hypothesis all other cs_i in l_{Cst} are orthogonal to cs and therefore the property is conserved. The same can be shown for the case of $(b, l) \Downarrow_{\mathbb{B}} 1$.

If $S = \mathbf{qcase\ } s[i] \mathbf{\ of\ } \{0 \rightarrow S_0, 1 \rightarrow S_1\}$ with $(s[i], l) \Downarrow_{\mathbb{N}} n$, notice first that $cs_i \perp cs_{i'}$ implies both $cs_i \perp cs_{i'}[n = 0]$ and $cs_i \perp cs_{i'}[n = 1]$, and that clearly $cs_{i'}[n = 0] \perp cs_{i'}[n = 1]$. Then, all cases preserve the condition.

If $S = \mathbf{call\ } \text{proc}'[i](s)$ with $(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} l' \neq []$ and $(i, l) \Downarrow_{\mathbb{Z}} n$, we consider two cases:

- (i) A corresponding ancilla a already exists (merging), which by the constraints of PFOQ implies that $\cdot[a = 1] \perp cs_i$ for all cs_i in l_{Cst} . Therefore, by the induction hypothesis, adding cs to a preserves the orthogonality conditions.
- (ii) No corresponding ancilla exists (anchoring), in which case the creation of the ancilla does not change the orthogonality between statements, as intended. \square

We now show that algorithm **compile**⁺ strictly improves on the asymptotic size of circuit, compared to the **compile** algorithm of [10].

Theorem 3. *For any PFOQ program P , $\#\mathbf{compile}^+(P, n) = O(\#\mathbf{compile}(P, n))$. Furthermore, there exist programs for which $\#\mathbf{compile}^+(P, n) = o(\#\mathbf{compile}(P, n))$.*

Proof. The circuit size in both cases is asymptotically bounded by the number of ancillas created. Since we do more merging than before the result follows. Table 1 provides some examples to show the second claim. \square

Theorem 4 (No branch sequentialization). *For $P \in \text{BFOQ}$ and $n \in \mathbb{N}$, $\#\mathbf{compile}^+(P, n) = O(\text{level}_P(n))$.*

Proof. The theorem can be shown by structural induction on the program body, by checking that it is the case in each scenario that the circuit size scales with the level of the program. All cases are straightforward except the one of the quantum control case, which is proven at the end. The BASIC restrictions give us the following two properties during the compilation: (a) merging can be done in constant time, since there is no need for controlled-swap gates, and (b) a call to a recursive function only result in at most $O(n)$ calls to procedures of the same rank with unique ancillas.

We proceed by structural induction on the program body, considering that the statement is part of a procedure call for procedure proc .

- ($S = \mathbf{skip}$; or $S = \bar{q}[i] \mathbf{*} = U$;) in this case we have that, $\text{level}_P(n) = 0$ and the the circuit is of constant size.
- ($S = S_1 S_2$) In this case, S_1 and S_2 are compiled in series (Figure 7a). The size of the circuit for S is then given by the sum of the sizes of the circuits of S_1 and S_2 , and by definition the level of S is the sum of the levels.
- ($S = \mathbf{if\ b\ then\ } S_0 \mathbf{\ else\ } S_1$) Depending on the value of b the circuit for S either it corresponds to the circuit for S_0 or S_1 . Therefore the size of the circuit is bounded by the maximum of between the two statements, as in the definition of level.
- ($S = \mathbf{call\ } \text{proc}[i](s)$;) This case also follows the definition of level since the circuit size is the one given inductively by the non-procedure-call operations (constant size) plus the circuit given by the procedure calls.

Notice that, for all statements besides the **qcase**, the size of the circuit follows the definition of level. We check that the number of ancillas created for S is bounded by the maximum number of ancillas for S_0 and S_1 separately. To show this, we proceed by induction on the rank r of the procedure.

The base case is given by (b), therefore we may consider $r > 1$. For the inductive case, we consider three possible scenarios:

Problem	Example	Circuit complexity	
		[10]	compile ⁺
Full Adder	Example 3	$\Theta(n)$	$\Theta(n)$
Quantum Fourier Transform	Example 4	$\Theta(n^2)$	$\Theta(n^2)$
Palindromes	Example 5	$\Theta(n)$	$\Theta(n)$
Chained Substring	Example 6	$\Theta(n^3)$	$\Theta(n)$
Sum(r), $r \geq 1$	Example 7	$\Theta(n^r)$	$\Theta(n)$

Table 1: Circuit size complexity bounds given by the compilation strategy in [10] and **compile**⁺ described in this work. For all of these problems, we give the corresponding programs in BFOQ.

- $w_{\text{proc}}^P(S_0) = w_{\text{proc}}^P(S_1) = 0$. Therefore, S_0 and S_1 contain only calls to procedures of rank strictly lower than r . This may only occur a constant number of times in the depth of a program, therefore we may simply consider the sum of the number of ancillas as a sufficient upper bound on the asymptotic number of ancillas for S .
- $w_{\text{proc}}^P(S_0) = w_{\text{proc}}^P(S_1) = 1$. In this case, S_0 and S_1 are of the same rank, r , and all their procedure calls may be merged. Therefore, the asymptotic number of such calls is bounded between the maximum between S_0 and S_1 (consider that, if there is no overlap between the ancillas needed, their number is still bounded linearly). Applying the IH on the procedures of rank $r - 1$ we obtain the desired result.
- $w_{\text{proc}}^P(S_0) = 0$ and $w_{\text{proc}}^P(S_1) = 1$. Therefore, S_0 contains calls to procedures of rank $r' < r$ whereas S_1 contains calls to procedures of rank r . The number of procedures of rank r' is bounded asymptotically by the maximum between those in S_0 and S_1 , therefore we obtain our result. \square

4 Examples

In this section, we provide several examples illustrating our results, including general examples on regular expressions. We show that any regular language can be decided by a BFOQ program whose compiled quantum circuit of linear size (Theorem 5). A benchmark, illustrating the difference between our compilation algorithm and the one in [10], is provided in Table 1.

Example 5 (Palindromes). *Consider the following BFOQ program PALINDROME.*

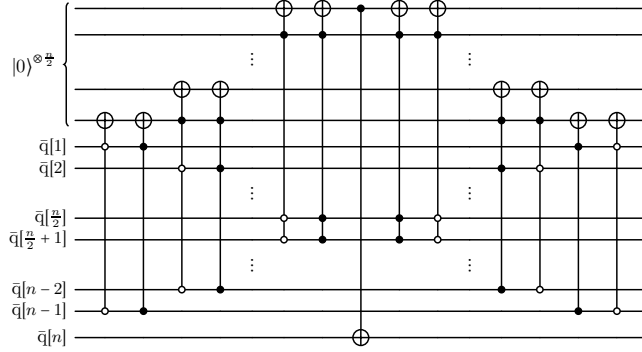
```

1  decl palindrome( $\bar{q}$ ) {
2    if  $|\bar{q}| > 2$  then
3      qcase  $\bar{q}[1, |\bar{q}| - 1]$  of {
4        00  $\rightarrow$  call palindrome( $\bar{q} \ominus [1, -2]$ );
5        01  $\rightarrow$  skip;
6        10  $\rightarrow$  skip;
7        11  $\rightarrow$  call palindrome( $\bar{q} \ominus [1, -2]$ );
8      }
9    else  $\bar{q}[-1] \ast = \text{NOT}$ ; }
10 :: call palindrome( $\bar{q}$ );

```

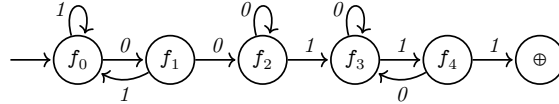
PALINDROME \in WF since all recursive procedure calls decrease the input sorted set. Furthermore, at most one recursive call is done per branch, and therefore PALINDROME \in WIDTH _{≤ 1} and so the program is also in PFOQ. Further checking that all procedure calls in the program are either of the form \bar{q} or $\bar{q} \ominus [1, |\bar{q}| - 1]$, we conclude that it is also in BFOQ. We are therefore in a position to apply Theorem 4.

Since $rk(\text{PALINDROME}) = rk_{\text{PALINDROME}}(\text{palindrome}) = 1$, by Lemma 1, we obtain the conclusion that $\#\text{optimize}^+(\text{P}, n) = O(n)$, i.e., the compilation procedure generates a circuit of size linear on the input. Indeed, for *PALINDROME*, compile^+ generates the following circuit in the case where n is even:



The circuit makes use of $n/2$ ancillas that are reset to zero, only applying a NOT gate to $\bar{q}[n]$ if $\bar{q}[1, \dots, n-1]$ forms a palindrome.

Example 6 (Chained substring). Let $s_0 = 001$ and $s_1 = 11$. Let \mathcal{L} be the regular language defined by identifying strings containing an instance of s_0 followed eventually by an instance of s_1 in a word, i.e., the language defined by the regular expression $*s_0*s_1*$. We can define a BFOQ program (Appendix C) that detects inputs in \mathcal{L} using the following call graph:



The program has as body a procedure call $\text{call } f_0(\bar{q})$; and consists of 5 procedures f_i and a terminating procedure \oplus . An arrow $s \xrightarrow{b} t$ with $b \in \{0, 1\}$ indicates a procedure call of the form $\text{call } t(\bar{q} \oplus [1])$; appears in the body of procedure s done in a branch with $\bar{q}[1]$ in state b .

The maximum rank of a procedure is 3 (for f_0 and f_1) and the circuit obtained by the technique in [10] gives a circuit of size $\Theta(n^3)$. On the other hand, the size of the circuit produced by compile^+ grows linearly on the input size.

In the previous example, the bound obtained by compile^+ was linear, which is the expected complexity in the case of detecting a regular language. It is straightforward to show that this is the case for any regular language, using the bound on the size of BFOQ circuits given in Theorem 4.

Definition 4. Let $A: \{0, 1\}^* \rightarrow \{0, 1\}$ be a decision problem. Given a FOQ program P , we say that P decides A if, for $\bar{x} \in \{0, 1\}^*$ and $y \in \{0, 1\}$, we have that $\llbracket P \rrbracket(|\bar{x}y\rangle) = |\bar{x}(y \oplus A(\bar{x}))\rangle$.

Theorem 5 (Regular languages). For any regular language \mathcal{L} , there exists a BFOQ program P that decides if $\bar{x} \in \mathcal{L}$, for any $\bar{x} \in \{0, 1\}^*$, such that $\#\text{compile}^+(P, n) \in O(n)$.

Proof sketch. Since \mathcal{L} is regular, there exists a deterministic finite automaton \mathcal{D} that decides it. It is relatively simple to construct from \mathcal{D} an BFOQ program that decides the language in the sense given in Definition 4. Since \mathcal{D} is deterministic, the level of the corresponding program is bounded linearly. Using Theorem 4 we obtain the desired result. \square

Example 7. Let SUM_r be the decision problem of checking if an input bitstring contains precisely r 1s. This corresponds to identifying bitstrings in the regular expression $(0^*1)^r 0^*$ and therefore, by Theorem 5, there exists a BFOQ program deciding SUM_r such that compile^+ outputs a family of circuits of linear size.

5 Conclusions and Future Work

In this paper, we have delineated an expressive fragment, named BFOQ, of the first-order quantum programming language with quantum control of [10]. We have shown that BFOQ is sound and complete for polynomial time computation (Theorem 2) and that the branch sequentialization problem introduced by [19] is solved for BFOQ programs: the compiled circuit has size upper-bounded by the maximal complexity of program branches (Theorem 4). As a consequence, the compilation procedure generates circuits with a better size complexity than the compilation algorithm of [10] (Theorem 3). This result and the expressivity of BFOQ are illustrated by the Examples of Table 1. A future and challenging research direction includes the extension of this work to higher-order.

References

- [1] Samson Abramsky. No-cloning in categorical quantum mechanics. *Semantic Techniques in Quantum Computation*, pages 1–28, 2009.
- [2] John Baez. Quantum quandaries: a category-theoretic perspective. *The structural foundations of quantum gravity*, pages 240–265, 2006.
- [3] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. doi:10.1137/S0097539796300921.
- [4] Hans J. Briegel, David E. Browne, Wolfgang Dür, Robert Raussendorf, and Maarten Van den Nest. Measurement-based quantum computation. *Nature Physics*, 5(1):19–26, 2009. doi:10.1038/nphys1157.
- [5] Ugo Dal Lago, Andrea Masini, and Margherita Zorzi. Quantum implicit computational complexity. *Theoretical Computer Science*, 411(2):377–409, 2010. doi:10.1016/j.tcs.2009.07.045.
- [6] Vincent Danos and Elham Kashefi. Determinism in the one-way model. *Physical Review A*, 74(5):052310, 2006. doi:10.1103/PhysRevA.74.052310.
- [7] Niel de Beaudrap, Xiaoning Bian, and Quanlong Wang. Fast and effective techniques for T-count reduction via spider nest identities. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICS.TQC.2020.11.
- [8] Alejandro Díaz-Caro, Mauricio Guillermo, Alexandre Miquel, and Benoît Valiron. Realizability in the unitary sphere. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13. IEEE, 2019.
- [9] 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.
- [10] Emmanuel Hainry, Romain Péchoux, and Mário Silva. A programming language characterizing quantum polynomial time. In *Foundations of Software Science and Computation Structures*, pages 156–175. Springer, 2023. doi:10.1007/978-3-031-30829-1_8.
- [11] Chris Heunen and Jamie Vicary. *Categories for Quantum Theory: an introduction*. Oxford University Press, 2019.
- [12] Aleks Kissinger and John van de Wetering. Reducing the number of non-Clifford gates in quantum circuits. *Phys. Rev. A*, 102:022406, Aug 2020. doi:10.1103/PhysRevA.102.022406.
- [13] Emanuel Knill, Raymond Laflamme, and Gerard J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, Jan 2001. doi:10.1038/35051009.
- [14] D. Maslov, G.W. Dueck, D.M. Miller, and C. Negrevergne. Quantum circuit simplification and level compaction. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(3):436–444, March 2008. doi:10.1109/tcad.2007.911334.

- [15] 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):23, May 2018. doi:10.1038/s41534-018-0072-4.
- [16] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2011.
- [17] Tomoyuki Yamakami. A Schematic Definition of quantum Polynomial Time Computability. *J. Symb. Log.*, 85(4):1546–1587, 2020. doi:10.1017/jsl.2020.45.
- [18] Andrew Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 352–361, 1993. doi:10.1109/SFCS.1993.366852.
- [19] Charles Yuan and Michael Carbin. Tower: Data Structures in Quantum Superposition. *Proceedings of the ACM on Programming Languages*, 6(OOPSLA2):259–288, Oct 2022. doi:10.1145/3563297.

A Semantics of FOQ programs

In Section 2, we have defined $\mathcal{L}(\mathbb{N})$ as the set of lists of natural numbers $[n_1, \dots, n_k]$ (the empty list being denoted by $[]$), which are used to represent list of (unique) qubit pointers in the semantics.

Basic data types τ are interpreted as follows:

$$\begin{aligned} \llbracket \text{Integers} \rrbracket &\triangleq \mathbb{Z} & \llbracket \text{Booleans} \rrbracket &\triangleq \mathbb{B} & \llbracket \text{SortedSets} \rrbracket &\triangleq \mathcal{L}(\mathbb{N}) \\ \llbracket \text{Qubits} \rrbracket &\triangleq \mathbb{N} & \llbracket \text{Operators} \rrbracket &\triangleq \tilde{\mathbb{C}}^{2 \times 2} \end{aligned}$$

Each basic operation $\text{op} \in \{+, -, >, \geq, =, \wedge, \vee, \neg\}$ of arity n , with $1 \leq n \leq 2$, has a type signature $\tau_1 \times \dots \times \tau_n \rightarrow \tau$ fixed by the program syntax. For example, the operation $+$ has signature $\text{Integers} \times \text{Integers} \rightarrow \text{Integers}$. A total function $\llbracket \text{op} \rrbracket \in \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$ is associated to each basic operation op .

A function $\llbracket U^f \rrbracket \in \mathbb{Z} \rightarrow \tilde{\mathbb{C}}^{2 \times 2}$ is associated to each U^f as follows:

$$\llbracket \text{NOT} \rrbracket(n) \triangleq \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \llbracket R_Y^f \rrbracket(n) \triangleq \begin{pmatrix} \cos(f(n)) & -\sin(f(n)) \\ \sin(f(n)) & \cos(f(n)) \end{pmatrix}, \quad \llbracket \text{Ph}^f \rrbracket(n) \triangleq \begin{pmatrix} 1 & 0 \\ 0 & e^{if(n)} \end{pmatrix},$$

where $\tilde{\mathbb{C}}$ is the set of polynomial time computable complex numbers, i.e., complex numbers whose both real and imaginary part are in $\tilde{\mathbb{R}}$. Each of the above matrices is unitary, i.e., the matrix M satisfies $M^* \circ M = M \circ M^* = I$, with M^* being the conjugate transpose of M and I being the identity matrix.

For each basic type τ , the reduction $\Downarrow_{\llbracket \tau \rrbracket}$ is a map in $\tau \times \mathcal{L}(\mathbb{N}) \rightarrow \llbracket \tau \rrbracket$. Intuitively, it maps an expression of type τ to its value in $\llbracket \tau \rrbracket$ for a given list l of pointers in memory. These reductions are defined in Figure 8, where e and d denote either an integer expression i or a Boolean expression b .

$$\begin{array}{c} \frac{(e, l) \Downarrow_{\llbracket \tau_1 \rrbracket} m \quad (d, l) \Downarrow_{\llbracket \tau_2 \rrbracket} n}{(e \text{ op } d, l) \Downarrow_{\llbracket \text{op} \rrbracket(\llbracket \tau_1 \rrbracket, \llbracket \tau_2 \rrbracket)} \llbracket \text{op} \rrbracket(m, n)} \text{ (Op)} \quad \frac{(i, l) \Downarrow_{\mathbb{Z}} n}{(U^f(i), l) \Downarrow_{\mathbb{C}^{2 \times 2}} \llbracket U^f \rrbracket(n)} \text{ (Unit)} \\ \\ \frac{}{(n, l) \Downarrow_{\mathbb{Z}} n} \text{ (Cst)} \quad \frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_m] \quad (i, l) \Downarrow_{\mathbb{Z}} k \in [1, m]}{(s \oplus [i], l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_m]} \text{ (Rm}_\epsilon) \\ \\ \frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_n]}{(|s|, l) \Downarrow_{\mathbb{Z}} n} \text{ (Size)} \quad \frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_m] \quad (i, l) \Downarrow_{\mathbb{Z}} k \notin [1, m]}{(s \oplus [i], l) \Downarrow_{\mathcal{L}(\mathbb{N})} []} \text{ (Rm}_\notin) \\ \\ \frac{}{(\text{nil}, l) \Downarrow_{\mathcal{L}(\mathbb{N})} []} \text{ (Nil)} \quad \frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_m] \quad (i, l) \Downarrow_{\mathbb{Z}} k \in [1, m]}{(s[i], l) \Downarrow_{\mathbb{N}} x_k} \text{ (Qu}_\epsilon) \\ \\ \frac{}{(\bar{q}, l) \Downarrow_{\mathcal{L}(\mathbb{N})} l} \text{ (Var)} \quad \frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} [x_1, \dots, x_m] \quad (i, l) \Downarrow_{\mathbb{Z}} k \notin [1, m]}{(s[i], l) \Downarrow_{\mathbb{N}} 0} \text{ (Qu}_\notin) \end{array}$$

Figure 8: Semantics of expressions

Recall from Section 2 that the set of *configurations* over n qubits, denoted Conf_n , is defined by

$$\text{Conf}_n \triangleq (\text{Statements} \cup \{\top, \perp\}) \times \mathcal{H}_{2^n} \times \mathcal{P}(\mathbb{N}) \times \mathcal{L}(\mathbb{N}),$$

where $\mathcal{P}(\mathbb{N})$ being the powerset over \mathbb{N} and where \top and \perp are two special symbols for termination and error, respectively. Let \diamond stand for a symbol in $\{\top, \perp\}$.

A configuration $c = (S, |\psi\rangle, \mathcal{S}, l) \in \text{Conf}_n$ contains a statement S to be executed (provided that $S \notin \{\top, \perp\}$), a quantum state $|\psi\rangle$ of length n , a set \mathcal{S} containing the qubit pointers that are allowed to be accessed by statement S , and a list l of qubit pointers.

The program big-step semantics $\cdot \longrightarrow \cdot$, described in Figure 9, is defined as a relation in $\bigcup_{n \in \mathbb{N}} \text{Conf}_n \times \mathbb{N} \times \text{Conf}_n$.

$$\begin{array}{c}
\frac{}{(\text{skip}, |\psi\rangle, \mathcal{S}, l) \xrightarrow{0} (\top, |\psi\rangle, \mathcal{S}, l)} \text{(Skip)} \quad \frac{(s[i], l) \Downarrow_{\mathbb{N}} n \notin \mathcal{S}}{(s[i] \text{ * } = U^f(j);, |\psi\rangle, \mathcal{S}, l) \xrightarrow{0} (\perp, |\psi\rangle, \mathcal{S}, l)} \text{(Asg}_{\perp}) \\
\frac{(s[i], l) \Downarrow_{\mathbb{N}} n \in \mathcal{S} \quad (U^f(j), l) \Downarrow_{\mathbb{C}^{2 \times 2}} M}{(s[i] \text{ * } = U^f(j);, |\psi\rangle, \mathcal{S}, l) \xrightarrow{0} (\top, I_{2^{n-1}} \otimes M \otimes I_{2^{l(\psi)-n}} |\psi\rangle, \mathcal{S}, l)} \text{(Asg}_{\top}) \\
\frac{(S_1, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m_1} (\top, |\psi'\rangle, \mathcal{S}, l) \quad (S_2, |\psi'\rangle, \mathcal{S}, l) \xrightarrow{m_2} (\diamond, |\psi''\rangle, \mathcal{S}, l)}{(S_1 S_2, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m_1+m_2} (\diamond, |\psi''\rangle, \mathcal{S}, l)} \text{(Seq}_{\diamond}) \\
\frac{(S_1, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m} (\perp, |\psi\rangle, \mathcal{S}, l)}{(S_1 S_2, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m} (\perp, |\psi\rangle, \mathcal{S}, l)} \text{(Seq}_{\perp}) \\
\frac{(b, l) \Downarrow_{\mathbb{B}} b \in \mathbb{B} \quad (S_b, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m_b} (\diamond, |\psi'\rangle, \mathcal{S}, l)}{(\text{if } b \text{ then } S_{\text{true}} \text{ else } S_{\text{false}}, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m_b} (\diamond, |\psi'\rangle, \mathcal{S}, l)} \text{(If)} \\
\frac{(s[i], l) \Downarrow_{\mathbb{N}} n \in \mathcal{S} \quad (S_k, |\psi\rangle, \mathcal{S} \setminus \{n\}, l) \xrightarrow{m_k} (\top, |\psi_k\rangle, \mathcal{S} \setminus \{n\}, l)}{(q\text{case } s[i] \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\}, |\psi\rangle, \mathcal{S}, l) \xrightarrow{\max_k m_k} (\top, \sum_k |k\rangle_n \langle k|_n |\psi_k\rangle, \mathcal{S}, l)} \text{(Case}_{\top}) \\
\frac{(s[i], l) \Downarrow_{\mathbb{N}} n \in \mathcal{S} \quad (S_k, |\psi\rangle, \mathcal{S} \setminus \{n\}, l) \xrightarrow{m_k} (\diamond_k, |\psi_k\rangle, \mathcal{S} \setminus \{n\}, l) \quad \perp \in \{\diamond_0, \diamond_1\}}{(q\text{case } s[i] \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\}, |\psi\rangle, \mathcal{S}, l) \xrightarrow{\max_k m_k} (\perp, |\psi\rangle, \mathcal{S}, l)} \text{(Case}_{\perp}) \\
\frac{(s[i], l) \Downarrow_{\mathbb{N}} n \notin \mathcal{S}}{(q\text{case } s[i] \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\}, |\psi\rangle, \mathcal{S}, l) \xrightarrow{0} (\perp, |\psi\rangle, \mathcal{S}, l)} \text{(Case}_{\notin}) \\
\frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} l' \neq [] \quad (i, l) \Downarrow_{\mathbb{Z}} n \quad (S^{\text{proc}}\{n/x\}, |\psi\rangle, \mathcal{S}, l') \xrightarrow{m} (\diamond, |\psi'\rangle, \mathcal{S}, l')}{(\text{call proc}[i](s);, |\psi\rangle, \mathcal{S}, l) \xrightarrow{m+1} (\diamond, |\psi'\rangle, \mathcal{S}, l)} \text{(Call}_{\diamond}) \\
\frac{(s, l) \Downarrow_{\mathcal{L}(\mathbb{N})} []}{(\text{call proc}[i](s);, |\psi\rangle, \mathcal{S}, l) \xrightarrow{1} (\top, |\psi\rangle, \mathcal{S}, l)} \text{(Call}_{\top})
\end{array}$$

Figure 9: Semantics of statements

B Proofs

In this section, we provide the full proof of Theorem 5. Towards that end, we first define a notion of call-graph.

Definition 5 (Call graph). *A call graph \mathcal{G} is a triple (proc, V, E) where*

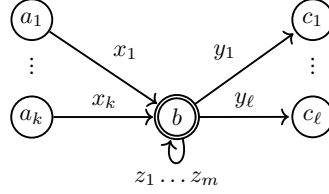
- $\text{proc} \in V$ is the entry node;
- $V \subseteq \text{Procedures}$ is a set of nodes containing a special procedure \oplus ;
- $E \subseteq V \times L \times V$ is a set of labeled directed edges, where labels correspond to combinations of quantum and classical conditionals.

Procedure \oplus only applies a NOT gate to the last qubit in the input and terminates. Labels L are defined as follows: values $\{0, 1\}$ denote a quantum if statement on the first qubit, and $|\bar{q}| = n$ or $|\bar{q}| > n$ denotes a boolean condition on the size of the input.

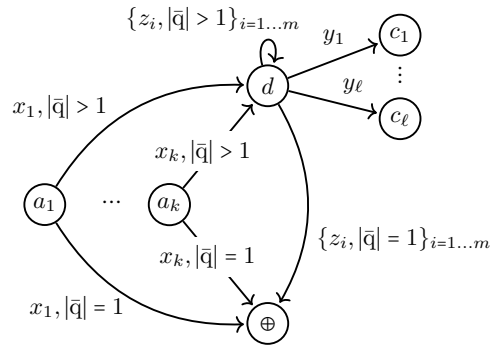
Theorem 5 (Regular languages). *For any regular language \mathcal{L} , there exists a BFOQ program P that decides if $\bar{x} \in \mathcal{L}$, for any $\bar{x} \in \{0, 1\}^*$, such that $\#\text{compile}^+(P, n) \in O(n)$.*

Proof. Let \mathcal{D} be a deterministic finite automaton deciding \mathcal{L} . We will construct the PFOQ program P by using \mathcal{D} to define the call graph for P . The subtlety in the transformation is in the difference between the acceptance condition in \mathcal{D} (i.e., termination in an accept state) and the *acceptance nodes* of the call graph, referring here to the \oplus nodes. For a base state $|\bar{x}y\rangle$, the program P outputs $|\bar{x}(-y)\rangle$ iff it *ever* reaches a \oplus node, at which point P terminates.

The call graph is then defined as follows. The call graph contains a node for each state of \mathcal{C} , with the same transitions, except for those that constitute incoming or outgoing edges of an accept state, i.e., edges x_i, y_i, z_i in the following diagram:



These edges are encoded in the call graph as follows:



with an extra procedure d for each accept state and using edges with classical conditions to handle the acceptance condition of the program. P is then defined as the program given by the call graph where the procedures consist only of the procedure calls defined in the graph. For a set of conditions c_i , with $i = 1 \dots m$, we denote by $\{c_i\}_{i=1 \dots m}$ a set of m edges each labelled with a conditions c_i .

Since each procedure performs at most one procedure call per branch (recursive or otherwise) its level will be linear on the input size. \square

C Regular languages (Example 6)

The following is the program defined by the call graph given in Example 6.

```

1  decl  $f_0(\bar{q})\{$ 
2    qcase  $\bar{q}[1]$  of  $\{0 \rightarrow \text{call } f_1(\bar{q} \ominus [1]), 1 \rightarrow \text{call } f_0(\bar{q} \ominus [1])\},$ 

3  decl  $f_1(\bar{q})\{$ 
4    qcase  $\bar{q}[1]$  of  $\{0 \rightarrow \text{call } f_2(\bar{q} \ominus [1]), 1 \rightarrow \text{call } f_0(\bar{q} \ominus [1])\},$ 

5  decl  $f_2(\bar{q})\{$ 
6    qcase  $\bar{q}[1]$  of  $\{0 \rightarrow \text{call } f_2(\bar{q} \ominus [1]), 1 \rightarrow \text{call } f_3(\bar{q} \ominus [1])\},$ 

7  decl  $f_3(\bar{q})\{$ 
8    qcase  $\bar{q}[1]$  of  $\{0 \rightarrow \text{call } f_3(\bar{q} \ominus [1]), 1 \rightarrow \text{call } f_4(\bar{q} \ominus [1])\},$ 

9  decl  $f_4(\bar{q})\{$ 
10   qcase  $\bar{q}[1]$  of  $\{0 \rightarrow \text{call } f_3(\bar{q} \ominus [1]), 1 \rightarrow \text{call } \oplus(\bar{q} \ominus [1])\},$ 

11 decl  $\oplus(\bar{q})\{$ 
12    $\bar{q}[-1] \text{ *= NOT; }$ 

13 :: call  $f_0(\bar{q});$ 

```

It is straightforward to verify that all conditions of BFOQ are met and that the level of the program is linearly bounded.