

ReQWIRE: Reasoning about Reversible Quantum Circuits

Robert Rand Jennifer Paykin Dong-Ho Lee Steve Zdancewic
 {rrand, jpaykin, dongle, stevez}@cis.upenn.edu

University of Pennsylvania

Common quantum algorithms make heavy use of ancillae, scratch qubits that initialized at some state and later returned to that state and discarded. Existing quantum circuit languages allow programmers to assert that a qubit has been returned to the $|0\rangle$ state before it is discarded, allowing for a range of optimizations. However, existing languages cannot verify these assertions, which introduces a potential source of errors. In this paper, we present methods for verifying that ancillae are discarded in the desired state, and use these methods to implement a verified compiler from classical functions to quantum oracles.

1 Introduction

Many quantum algorithms rely heavily on *quantum oracles*, classical programs executed inside quantum circuits. Toffoli proved that any classical, boolean-valued function $f(x)$ can be implemented as a unitary circuit f_u satisfying $f_u(x, z) = (x, z \oplus f(x))$ [15]. Toffoli’s construction for quantum oracles is used in many quantum algorithms, such as the modular arithmetic of Shor’s algorithm [12]. As a concrete example, Figure 1 shows quantum circuits that implement the boolean functions and (\wedge) and or (\vee).

Unfortunately, Toffoli’s construction introduces significant overhead. Consider a circuit meant to compute the boolean formula $(a \vee b) \wedge (c \vee d)$. The circuit needs two additional scratch wires, or *ancillae*, to carry the outputs of $(a \vee b)$ and $(c \vee d)$, as seen in Figure 2. The annotation 0 at the start of a wire means that qubit is initialized in the state $|0\rangle$. When constructed in this naive way, the resulting circuit no longer corresponds to a unitary transformation, and cannot be safely used in a larger quantum circuit.

The solution is to *uncompute* the intermediate values $a \vee b$ and $c \vee d$ and then discard them at the end of the quantum circuit (Figure 3). The annotation 0 at the end of a wire is an *assertion* that the qubit at that point is in the zero state, at which point we can safely discard it without affecting the remainder of the state.

How can we verify that such an assertion is actually true? We cannot dynamically check the assertion, since we can only access the value of a qubit by measuring it, collapsing the qubit in question to a 0 or 1 state. However, we can statically reason that the qubit must be in the state $|0\rangle$ by analyzing the circuit semantics.

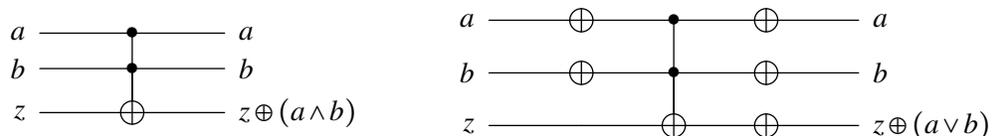


Figure 1: Quantum oracles implementing the boolean \wedge and \vee . The \oplus gates represent negation, and \bullet represents control.

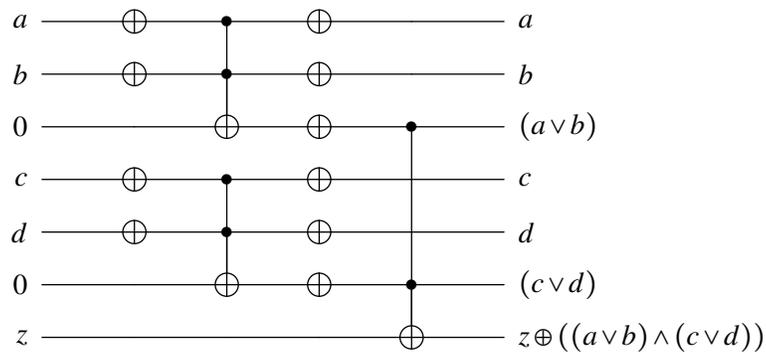


Figure 2: An non-unitary quantum oracle for $(a \vee b) \wedge (c \vee d)$

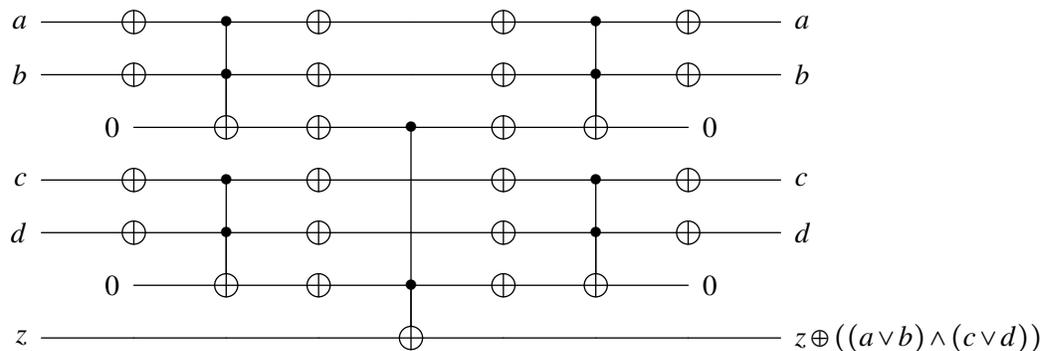


Figure 3: A unitary quantum oracle for $(a \vee b) \wedge (c \vee d)$ with ancillae

The claim that a qubit is in the 0 state is a semantic assertion about the behavior of the circuit. Unfortunately, this makes it hard to verify—computing the semantics of a quantum program is computationally intractable in general. Circuit programming languages often allow users to make such assertions, but not to verify that they are true. For example, Quipper [6] allows programmers to make assertions about the state of ancillae, but these assertions are never checked. The QCL quantum circuit language [7] provides a built-in method for creating reversible circuits from classical functions, but the programmer must trust this method to safely manage ancillae. The REVERC compiler [1] for the (non-quantum) reversible computing language REVS [8] provides a similar method, and verifies that it correctly uncomputes its ancilla. However, other assertions in REVS that a wire is correctly in the 0 state are ignored if they cannot be automatically verified.

In this paper, we develop verification techniques safely working with ancillae. Our approach allows the programmer to discard qubits that are in the state $|0\rangle$ or $|1\rangle$ provided she first formally proves that the qubits are in the state specified. Inspired by the REVERC compiler [1], we also provide syntactic conditions that the programmer may satisfy to guarantee that her assertions are true. However, our quantum circuits do not need to match this syntactic specification: a programmer may instead manually prove that her circuit safely discards qubits using the denotational semantics of the language. This gives the programmer the flexibility to use ancillae where the proof of such assertions are non-trivial.

We develop these techniques in the context of *QWIRE* (as in “require”), a domain-specific programming language for describing and reasoning about quantum circuits [9]. *QWIRE* is implemented as an embedded language inside the Coq proof assistant [3], which allows us to formally verify properties of our circuits. These properties can range from coarse-grained (“this circuit corresponds to a unitary transformation”) to precise (“this teleport circuit is equal to the identity”) [11]. *QWIRE* is an ongoing project and available for public use at <https://github.com/jpaykin/QWIRE>.

This paper reports on work-in-progress that makes the following contributions:¹

- We extend *QWIRE* with assertion-bearing ancillae.
- We give semantic conditions for the closely related properties of (a) when a circuit is reversible, and (b) when a circuit contains only valid assertions about its ancillae.
- We provide syntactic conditions that guarantee the correctness of these assertions for common use-cases.
- We implement a compiler to transform boolean expressions into reversible *QWIRE* circuits, and prove its correctness.
- We show how this compilation can be used to perform quantum arithmetic via a quantum adder.

2 The *QWIRE* Circuit Language

QWIRE [9] is a small quantum circuit language designed to be embedded in a larger, functional programming language. We have implemented *QWIRE* in the Coq proof assistant, which provides access to dependent types and the Coq interactive proof system. We use these features to type check *QWIRE* circuits and verify properties about their semantics [11]. In this section we give a brief introduction to the syntax and semantics of *QWIRE*, including the new ancilla assertions.

A *QWIRE* circuit consists of a sequence of gate applications terminated with some output wires.

$$\text{Circuit } W ::= \text{output } p \mid \text{gate } p' \leftarrow g \ p ; \text{Circuit } W$$

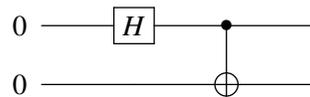
The parameter W refers to a *wire type*: `Bit`, `Qubit` or some tuple of `Bits` and `Qubits` (including the empty tuple `One`). A *pattern of wires*, denoted p , can be a bit-valued wire `bit v`, a qubit-valued wire `qubit v`, a pair of wires $(p1,p2)$ or an empty tuple $()$. Gates g are either unitary gates U , drawn from a universal gate set, or members of a small set of non-unitary gates:

$$\begin{aligned} W &:= \text{Bit} \mid \text{Qubit} \mid \text{One} \mid W \otimes W \\ g &:= U \mid \text{init_0} \mid \text{init_1} \mid \text{meas} \mid \text{discard} \mid \text{assert_0} \mid \text{assert_1} \end{aligned}$$

The `init` and `meas` gates initialize and measure qubits, respectively; `meas` results in a bit, which can be discarded by the `discard` bit or used as a control. The `assert_0` and `assert_1` gates take a qubit as input and discard it, provided that it is in the state $|0\rangle$ or $|1\rangle$ respectively. We will discuss the semantics of these gates, and how to verify assertions, in Sections 3 and 4.

As an example, the following *QWIRE* circuit prepares a Bell state:

```
gate p1 ← init_0 ();
gate p2 ← init_0 ();
gate p1 ← H p1;
gate (p1, p2) ← CNOT (p1, p2);
output (p1,p2)
```

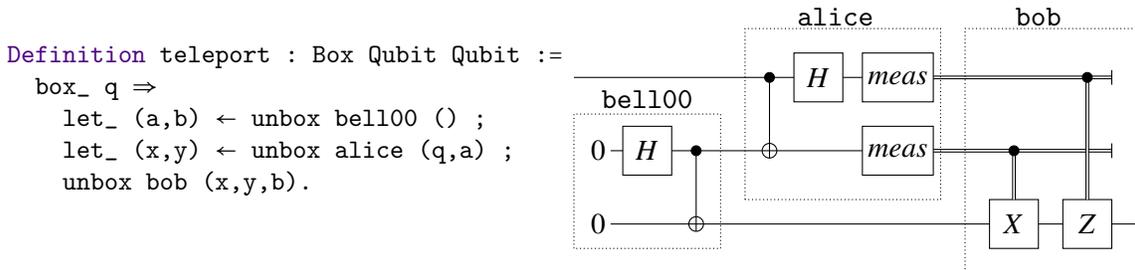


¹Section 7 elaborates on the state of the Coq development that underlies this work.

QWIRE also includes some more powerful functionality for constructing circuits. Circuits can be composed via a let binding $\text{let_} p \leftarrow C; C'$, where the output of the first circuit C is plugged into the wires p in C' . It's worth highlighting two useful instances of composition: The $\text{inSeq} (;)$ function takes a $\text{Box } W' W'$ and a $\text{Box } W' W''$ and composes them sequentially to return a $\text{Box } W' W''$. The inPar function likewise takes a circuit $c1$ of type $\text{Box } W_1 W_1'$ and $c2$ of type $\text{Box } W_2 W_2'$ and composes them in parallel, producing $c1 \parallel c2$ of type $\text{Box } (W_1 \otimes W_2) (W_1' \otimes W_2')$.

Circuits can also be *boxed* by collecting the input of a circuit in an input pattern $\text{box_} p \Rightarrow C$, creating a closed term of type $\text{Box } W' W'$ in the host language. Here, the input wire type W matches the type of the input wire p , and the output type W' is the same as that of the underlying circuit. Such a boxed circuit can be *unboxed* to be used again in other circuits.

Boxing, unboxing, and composing circuits is illustrated by the *teleport* circuit:



3 A Safe and Unsafe Semantics

As in prior work, QWIRE's semantics is given in terms of *density matrices*, denoted ρ , that represent distributions over pure quantum states known as *mixed states*. A QWIRE circuit of type $\text{Box } W' W'$ maps a $2^{\llbracket W \rrbracket} \times 2^{\llbracket W \rrbracket}$ density matrix to a $2^{\llbracket W' \rrbracket} \times 2^{\llbracket W' \rrbracket}$ density matrix, where $\llbracket W \rrbracket$ is the size of a wire type:²

$$\llbracket \text{One} \rrbracket = 0 \quad \llbracket \text{Qubit} \rrbracket = \llbracket \text{Bit} \rrbracket = 1 \quad \llbracket W_1 \otimes W_2 \rrbracket = \llbracket W_1 \rrbracket + \llbracket W_2 \rrbracket$$

In this work we use mixed states only to refer to total, rather than partial distributions. This means that all of our mixed states have a trace equal to 1.

In this work, we give two different semantics for quantum circuits that differ in how they treat assertions. The *safe* semantics corresponds to an operational model that does not trust assertions, so an assert_x gate first measures the input qubit before discarding the result. The *unsafe* semantics assumes that assertions are accurate, so an assert_x gate simply discards its input qubit without measuring it. The two semantics coincide exactly when all assertions in a circuit are accurate, in which case we call the circuit *valid*.

In the safe semantics, assertions are identical to the discard gate, which measures and then discards the qubit.

$$\begin{aligned} \text{denote_safe } U \rho &= \llbracket U \rrbracket \rho \llbracket U \rrbracket^\dagger \\ \text{denote_safe } \text{init}_0 \rho &= |0\rangle \rho \langle 0| \\ \text{denote_safe } \text{init}_1 \rho &= |1\rangle \rho \langle 1| \\ \text{denote_safe } \text{meas } \rho &= |0\rangle \langle 0| \rho |0\rangle \langle 0| + |1\rangle \langle 1| \rho |1\rangle \langle 1| \\ \text{denote_safe } \{\text{discard}, \text{assert}_0, \text{assert}_1\} \rho &= \langle 0| \rho |0\rangle + \langle 1| \rho |1\rangle \end{aligned}$$

²In practice, the semantics must be "padded" by an additional type so that it can be applied in a larger quantum system.

Here $\llbracket U \rrbracket$ is the unitary matrix corresponding to the gate U ; multiplying by $\llbracket U \rrbracket$ and $\llbracket U \rrbracket^\dagger$ is equivalent to applying $\llbracket U \rrbracket$ to all the pure states in the distribution. The initialization gates init_0 and init_1 both add a single qubit to the system, in the $|0\rangle$ and $|1\rangle$ states respectively. The meas gate produces a mixed state corresponding to a probability distribution over the measurement result ($|0\rangle$ or $|1\rangle$). The discard gate removes a classical-valued bit from the state.

Under the safe semantics, the assertions assert_0 and assert_1 are treated as a measurement followed by a discard. This is semantically the same as the denotation of discard , except that discard is guaranteed by the type system to only throw away a classically valued bit. This operation on qubits is safe even if the qubit is in a superposition of $|0\rangle$ and $|1\rangle$, due to the implicit measurement.

The *unsafe* semantics is the same as the safe semantics, except for assert_0 and assert_1 :

$$\begin{aligned} \text{denote_unsafe } \text{assert}_0 \rho &= \langle 0 | \rho | 0 \rangle \\ \text{denote_unsafe } \text{assert}_1 \rho &= \langle 1 | \rho | 1 \rangle \end{aligned}$$

It should be immediately clear why this is unsafe: if ρ isn't in the zero state (in the first case), then an assertion produces a density matrix with a trace less than 1. Operationally, this corresponds to the instruction “throw away this qubit in the zero state”, which is quantum-mechanically impossible in the general case. However, this semantics corresponds to the intended meaning of assert_x when we know the assertion is true. It also ensures that the composition of init_x with assert_x is equivalent to the identity, which allows us to optimize away qubit initialization and discarding.

We can now define what it means for the ancilla assertions in a circuit to be valid.

Definition $\text{valid_ancillae } W (c : \text{Circuit } W) : \mathbb{P} := (\text{denote } c = \text{denote_unsafe } c)$.

An equivalent definition states that the unsafe semantics preserves the trace of its input (which is always 1) and therefore maps it to a total probability distribution.

Definition $\text{valid_ancillae}' W (c : \text{Circuit } W) : \mathbb{P} :=$
 $\forall \rho, \text{Mixed_State } \rho \rightarrow \text{trace } (\text{denote_unsafe } c \rho) = 1$.

The second definition follows from the first because the safe semantics is trace preserving. The first follows from the second since the $\text{denote_unsafe } c \rho$ corresponds to a sub-distribution of $\text{denote_safe } c \rho$. If its trace is one, they must then represent the same distribution.

These two definitions precisely characterize what it means for circuits to have valid annotations. In the next section we define syntactic conditions that are sufficient but not necessary for validity. Programmers will often write syntactically valid circuits (such as the compiler in Section 5) but in the edge case, the semantic definitions of validity is still available.

An important property related to the validity of a circuit is its *reversibility*. We say that c and c' are *equivalent*, written $c \equiv c'$, if both their safe and unsafe denotations are equal. Reversibility says that a circuit has a left and right inverse:

Definition $\text{reversible } \{W_1 W_2\} (c : \text{Box } W_1 W_2) : \mathbb{P} :=$
 $(\text{exists } c', c' ;; c \equiv \text{id_circ}) \wedge (\text{exists } c', c ;; c' \equiv \text{id_circ})$

In Section 5, the compiler produces circuits that are their own inverse:

Definition $\text{self_inverse } \{W\} (c : \text{Box } W W) : \mathbb{P} := c ;; c \equiv \text{id_circ}$.

We can now show that in any reversible circuit, all the ancilla assertions hold.

Lemma 3.1. *If c is reversible, then it is valid.*

Proof. Let c' be c 's inverse. By the second definition of validity, it suffices to show that the trace of $\text{denote_unsafe } c \rho$ is equal to 1 for every initial mixed state ρ . We know that the trace of $\text{denote_unsafe id_circ } \rho$ is 1, hence

$$1 = \text{trace } (\text{denote_unsafe } (c;;c') \rho) = \text{trace } (\text{denote_unsafe } c' (\text{denote_unsafe } c \rho))$$

Because the unsafe semantics is trace-non-increasing, it must be the case that the trace of $\text{denote_unsafe } c \rho$ is 1 as well. \square

4 Syntactically Valid Ancillae

Let c be a circuit made up only of classical gates: the initialization gates, the not gate X, the controlled-not gate CNOT, and the Toffoli gate T. Let c' be the result of reversing the order of the gates in c and swapping every initialization with an assertion of the corresponding boolean value. Then every assertion in $c;;c'$ (where semicolons denote sequencing) is valid.

Unfortunately, every circuit of this form is also equivalent to the identity circuit, so as a syntactic condition of validity, this is much too restrictive. In practice, the quantum oracles discussed in the introduction are mostly symmetric, but they introduce key pieces of asymmetry to compute meaningful results. In REVERC, this construction is called the *restricted inverse*; QCL [7] and Quipper [6] take similar approaches.

Let c be a circuit with an equal number of input and output wires whose qubits can be broken up into two disjoint sets: the first n qubits are called the *source*, and the last t qubits are called the *target*. That is, $c : \text{Box } (n+t \otimes \text{Qubit}) (n+t \otimes \text{Qubit})$. The syntactic condition of *source symmetry* on circuits guarantees that c is the identity on all source qubits. In addition, it guarantees that assertions are only made on source qubits with a corresponding initialization.

A classical gate *acts on the qubit i* if it affects the value of that qubit in an m -qubit system: X acts on its (only) argument, CNOT acts on its second argument (the target) and Toffoli acts on its third argument.

The property of *source symmetry* on circuits is defined inductively as follows:

- The identity circuit is source symmetric.
- If g is a classical gate and c is source symmetric, then $g;;c$ is source symmetric.
- If g is a classical gate that acts on a qubit in the target, and c is source symmetric, then both $g;;c$ and $c;;g$ are source symmetric.
- If c is source symmetric and i is in the source of c , then $\text{init_at } b \ i;;c;;\text{assert_at } b \ i$ is source symmetric.

The key property of a source symmetric circuit is that it does not affect the value of its source qubits. We say that a circuit c is a *no-op* at qubit i if, when initialized with a boolean b , the qubit is still equal to b after executing the circuit. We could define this as $\llbracket c \rrbracket (\rho_1 \otimes |b\rangle \langle b| \otimes \rho_2) = \rho_1' \otimes |b\rangle \langle b| \otimes \rho_2'$ for some $\rho_1, \rho_2, \rho_1', \rho_2'$, but this would require ρ_1 and ρ_2 (and ρ_1' and ρ_2') to be separable, which is an unnecessary restriction. Instead, we use the `valid_ancillae` predicate and say if we initialize an ancilla in state x at i , apply b , and then assert that $i = x$, our assertion will be valid:

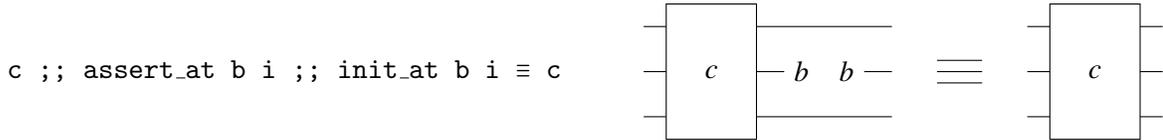
Definition `noop_on` $(m \ k : \mathbb{N}) (c : \text{Box } (\text{Qubits } (1 + m)) (\text{Qubits } (1+m))) : \mathbb{P} :=$
 $\forall b, \text{valid_ancillae } (\text{init_at } b \ i;;c;;\text{assert_at } b \ i).$

We similarly define a predicate `noop_on_source n`, that says that a given circuit is a no-op on each of its first n inputs.

These inductive definitions allow us to state a number of closely related lemmas about symmetric circuits:

Lemma 4.1. *If the classical gate g acts on the qubit k and $i \neq k$, then g is a no-op on i .*

Lemma 4.2. *Let c be a circuit such that `c ; ; assert_at b i` is a valid assertion. Then*



Lemma 4.3. *If c and c' are both no-ops on qubit i , then `c ; ; c'` is also a no-op on qubit i .*

Conjecture 4.4. *If c is source symmetric, then it is a no-op on its source.*

These lemmas have been admitted, rather than proven, in the Coq development (`Symmetric.v`). Conjecture 4.4 is labeled as a *conjecture* rather than a lemma, since we do not yet have a paper proof of the statement. It may be the case that we need to strengthen our definition of no-op for this conjecture to hold.

Since all ancillae in a source symmetric circuit occur on sources, we can prove from the statements above that source symmetric circuits are valid.

Theorem 4.5. *If c is source symmetric, then all its assertions are valid.*

Source symmetric circuits also satisfy a more general property—they are reversible.

The inverse of a source symmetric circuit is defined by induction on source symmetry:

- The inverse of the identity circuit is the identity;
- The inverse of `g ; ; c ; ; g` is `g ; ; c-1 ; ; g`;
- The inverse of `c ; ; g` is `g ; ; c-1` and vice versa; and
- The inverse of `init_at b i ; ; c ; ; assert_at b i` is `init_at b i ; ; c-1 ; ; assert_at b i`.

Clearly the inverse of any source symmetric circuit is also source symmetric, and the inverse is involutive, meaning $(c^{-1})^{-1} = c$.

Theorem 4.6. *If c is source symmetric, then `c-1 ; ; c` is equivalent to the identity circuit.*

Proof. By induction on the proof of source symmetry. The only interesting case is the case for ancilla, showing

$$\text{init_at } b \ i \ ; ; \ c^{-1} \ ; ; \ \text{assert_at } b \ i \ \circ \ \text{init_at } b \ i \ ; ; \ c \ ; ; \ \text{assert_at } b \ i \ \equiv \ \text{id_circ}.$$

From Theorem 4.5 we know that the circuit `init_at b i ; ; c-1 ; ; assert_at b i` is valid. Then, by Lemma 4.2, we know that `init_at b i ; ; c-1 ; ; assert_at b i ; ; init_at b i` is equivalent to `init_at b i ; ; c-1`. Thus the goal reduces to `init_at b i ; ; c-1 ; ; c ; ; assert_at b i`. This is equivalent to the identity by the induction hypothesis as well as the fact that `init_at b i ; ; assert_at b i` is the identity. \square

We can now say that any circuit followed by its inverse is valid. But this theorem is easily extensible. For instance, we can add the following to our inductive definition of symmetric and the theorem will still hold:

- If c is source symmetric and $c \equiv c'$ then c' is source symmetric.

This extension allows us to use existing (semantic) equivalences to satisfy our (syntactic) source symmetry predicate, which in turn proves the semantic property of validity. For example, because teleportation is semantically equivalent to the identity circuit, we know trivially that it is valid, even though it is not source symmetric. The Coq development provides many useful compiler optimizations in the file `Equations.v` that can now be used in establishing source symmetry.

5 Compiling Oracles

Now that we have syntactic guarantees for circuit validity, we consider a compiler from boolean expressions to source-symmetric circuits, producing the quantum oracles described in the introduction. The resulting circuits will all be source symmetric, so it follows from the previous section that their use of ancillae are valid.

We begin with a small boolean expression language, borrowed from Amy *et al.* [1], with variables x , constants, negation \neg , conjunction \wedge , and exclusive or \oplus .

$$b ::= x \mid t \mid f \mid \neg b \mid b_1 \wedge b_2 \mid b_1 \oplus b_2$$

The *interpretation function* $\llbracket b \rrbracket_f$ takes a boolean expression b and a valuation function $f : \text{Var} \rightarrow \text{bool}$ and returns the value of the boolean expression with the variables assigned as in f .

The compiler takes a boolean expression b and a *map* Γ from the variables of b to the wire indices³. The resulting circuit `compile b Γ` has $|\Gamma|+1$ qubit-valued input and output wires, where $|\Gamma|$ is the number of variables in the scope of b .

The compiler uses `init_at`, `assert_at`, `X_at`, `CNOT_at` and `Toffoli_at` circuits, each of which applies the corresponding gate to the given index in the list of n wires. Due to space constraints, we show only the cases for `true`, variables, and $b_1 \wedge b_2$; the other cases are analogous.

```
Fixpoint compile (b : bexp) ( $\Gamma$  : Ctx) : Box (Qubits (1 +  $|\Gamma|$ )) (Qubits (1 +  $|\Gamma|$ )) :=
  match b with
  | b_t           => X_at 0
  | b_var v       => CNOT_at (1 + get_index  $\Gamma$  v) 0
  | b_and b1 b2   => init_at false 1                ;;
                    id_circ || compile b1  $\Gamma$       ;;
                    init_at false 2                ;;
                    id_circ || id_circ || compile b2  $\Gamma$  ;;
                    Toffoli_at 1 2 0                ;;
                    id_circ || id_circ || compile b2  $\Gamma$  ;;
                    assert_at false 2              ;;
                    id_circ || compile b1  $\Gamma$      ;;
                    assert_at false 1
  | ...
  end.
```

We make heavy use of sequencing (`;`) and parallel (`||`) operators in defining this circuit. The `true` case outputs the exclusive-or of `true` with the target wire, which is equivalent to simply negating the target wire. The variable case `b_var` applies a CNOT gate from the variable's associated wire to the target, thereby sharing its value.

³In the Coq development, these maps are represented by linear typing contexts.

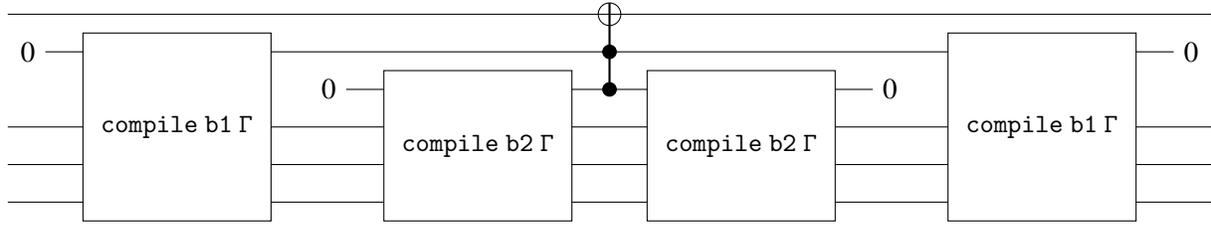


Figure 4: Compiling $b_1 \wedge b_2$ on 3 qubits. The top wire is the target.

The $b_1 \wedge b_2$ case (Figure 4) is the most interesting. We first initialize a qubit in the 0 state and recursively compile the value of b_1 to it. We then do the same for b_2 . We apply a Toffoli gate from b_1 and b_2 , now occupying the 1 and 2 positions in our list, to the target qubit at 0. We then reapply the symmetric functions $\text{compile } b_2 \Gamma$ and $\text{compile } b_1 \Gamma$ to their respective wires, returning the ancillae to their original states and discarding them. We are left with the target wire z holding the boolean value $b_z \oplus (b_1 \wedge b_2)$ and $|\Gamma|$ wires retaining their initial values. Note that this entire circuit is source symmetric and therefore our assertions are guaranteed to hold by Theorem 4.5.

We can now go about proving the correctness of this compilation.

```
Theorem compile_correct :  $\forall$  (b : bexp) ( $\Gamma$  : Ctx) (f : Var  $\rightarrow$  bool) (z : bool),
  vars b  $\subseteq$  domain  $\Gamma$   $\rightarrow$ 
   $\llbracket \text{compile } b \Gamma \rrbracket$  (bool_to_matrix t  $\otimes$  basis_state  $\Gamma$  f) =
  bool_to_matrix (z  $\oplus$   $\llbracket b \rrbracket_f$ )  $\otimes$  basis_state  $\Gamma$  f.
```

The function `basis_state` takes the wires referenced by Γ and the assignments of f and produces the corresponding basis state. This forms the input to the compiled boolean expression along with the target, a classical qubit in the $|0\rangle$ or $|1\rangle$ state. The statement of `compile`'s correctness says that when we apply $\llbracket \text{compile } b \Gamma \rrbracket$ to this basis state with an additional target qubit, we obtain the same matrix with the result of the boolean expression on the target. The proof follows by induction on the boolean expression.

6 Quantum Arithmetic in QWIRE

In this section we show how to use the compiler from the previous section to implement a quantum adder, which has applications in many quantum algorithms, including Shor's algorithm. A verified quantum adder is therefore an important step towards verifying a variety of quantum programs.

The input to an adder consists of two n -bit numbers represented as sequences of bits $x_{1:n}$ and $y_{1:n}$, as well as a carry-in bit c_{in} . The output consists of the sum $sum_{1:n}$ and the carry-out c_{out} .

To begin, consider a simple 1-bit adder that takes in three bits c_{in}, x and y , and computes their sum and carry-out values. The sum is equal to $x \oplus y \oplus c_{in}$ and the carry is $(c_{in} \wedge (x \oplus y)) \oplus (x \wedge y)$. The expressions can be compiled to 4- and 5-qubit circuits `adder_sum` and `adder_carry`, respectively, where the order of qubits is c_{out}, sum, y, x , and c_{in} .

```
Definition adder_sum : Box (4  $\otimes$  Qubit) (4  $\otimes$  Qubit) :=
  compile ((c_in  $\wedge$  (x  $\oplus$  y))  $\oplus$  (x  $\wedge$  y)) (list_of_Qubits 4).
```

```
Definition adder_carry : Box (5  $\otimes$  Qubit) (5  $\otimes$  Qubit) :=
  compile (x  $\oplus$  y  $\oplus$  c_in) (list_of_Qubits 5).
```

```
Definition adder_1 : Box (5  $\otimes$  Qubit) (5  $\otimes$  Qubit) :=
  adder_carry ;; (id_circ || adder_sum).
```

Here, `adder_sum` computes the sum of its three input bits and `adder_carry` computes the carry, ignoring the result of `adder_sum`. Semantically, the adder should produce the appropriate boolean values; the operation `bools_to_matrix` converts a list of booleans to a density matrix.

```

Lemma adder_1_spec : ∀ (cin x y sum cout : bool),
  [[adder_1]] (bools_to_matrix [cout; sum; y; x; cin])
= (bools_to_matrix [ cout ⊕ (c_in ∧ (x ⊕ y) ⊕ (x ∧ y));
                    ; sum ⊕ (x ⊕ y ⊕ c_in)
                    ; y; x; cin]).

```

Next, we extend the 1-qubit adder to n qubits. The n -qubit adder contains two parts—`adder_left` and `adder_right`—defined recursively using padded `adder_1` and `adder_carry` circuits. The left part computes the sum and carry sequentially from the least significant bit, initializing an ancilla for the carry in each step. When it reaches the most significant bit, it computes the most significant bit of the sum and carry-out using the 1-qubit adder. The right part of the adder uncomputes the carries and discards the ancillae. The definitions of the circuits are shown below and illustrated in Figure 5.

```

Fixpoint adder_left (n : ℕ) : Box ((1+3*n) ⊗ Qubit) ((1+4*n) ⊗ Qubit) :=
  match n with
  | S n' ⇒ (id_circ || (id_circ || (id_circ || (adder_left n')))) ;;
            (init_at false (4*n) 0) ;;
            (adder_1_pad (4*n')) end.
Fixpoint adder_right (n : ℕ) : Box ((1+4*n) ⊗ Qubit) ((1+3*n) ⊗ Qubit) :=
  match n with
  | 0 ⇒ id_circ
  | S n' ⇒ (adder_carry_pad (4*n')) ;;
            (assert_at false (4*n) 0) ;;
            (id_circ || (id_circ || (id_circ || (adder_right n')))) end.
Fixpoint adder_circ (n : ℕ) : Box ((2+3*n) ⊗ Qubit) ((2+3*n) ⊗ Qubit) :=
  match n with
  | 0 ⇒ id_circ
  | S n' ⇒ (id_circ || (id_circ || (id_circ || (id_circ || (adder_left n'))))) ;;
            (adder_1_pad (4*n')) ;;
            (id_circ || (id_circ || (id_circ || (id_circ || (adder_right n'))))) end.

```

We now state the correctness of the n -qubit adder:

```

Lemma adder_circ_n_spec : ∀ (n : ℕ) (f : Var → bool),
  let li := list_of_Qubits (2 + 3 * n) in
  [[adder_circ_n n]] (ctx_to_matrix li f)
= (ctx_to_matrix li (compute_adder_n n f)).

```

Like `bools_to_matrix` above, `ctx_to_matrix` takes in a context and an assignment f of variables to booleans, and constructs the corresponding density matrix. The function `compute_adder_n` likewise takes a function f that assigns values to each of the $3 * n + 2$ input variables and returns a boolean function f' representing the state of the same variables after addition (computed classically). The specification states that the n -bit adder circuit computes the state corresponding to the function `compute_adder_n` for any initial assignment.

Note that the lemma gives a correspondence between the denotation of the circuit and functional computation on the assignment. This can reduce the time required to verify more complex arithmetic circuits. A natural next step is to verify the correspondence between our functions on lists of booleans and Coq's binary representations of natural numbers, thereby grounding our results in the Coq standard library and allowing us to easily move between numerical representations.

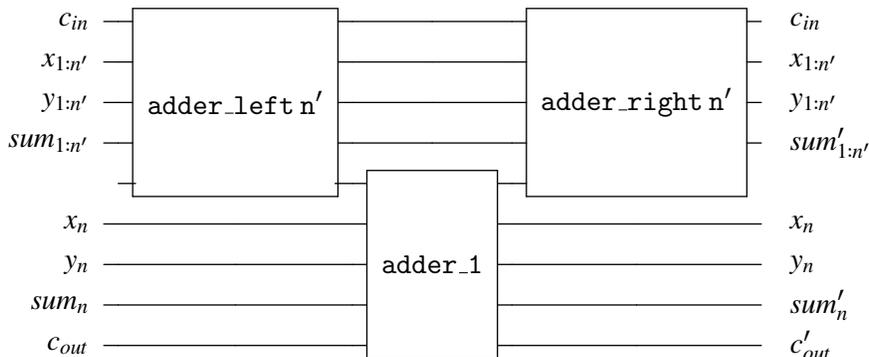


Figure 5: A quantum circuit for the n -adder where $n' = n - 1$. The n' ancillae created in `adder_left` are all terminated inside `adder_right`.

7 Related and Future Work

The area of reversible computation well predates quantum computing. [2], Bennett [2] introduced the reversible Turing machine, with the intent of designing a computer with low energy consumption, since destroying information necessarily dissipates energy. Toffoli designed the general approach for converting classical circuits to reversible ones presented in our introduction. While these ideas strongly influenced quantum computation, reversible computation is a subject of great interest in its own right, and we refer the interested reader to a standard text on the subject [4, 10].

This work builds heavily on the Quipper quantum programming language [6, 5], which includes ancillae terminations that are optimized away by joining them to corresponding initializations. Unfortunately, as is noted in the introduction, the language has no way of checking its “assertive terminations”:

The first thing to note is that it is the programmer, and not the compiler, who is asserting that the qubit is in state $|0\rangle$ before being terminated. In general, the correctness of such an assertion depends on intricacies of the particular algorithm, and is not something that the compiler can verify automatically. It is therefore the programmer’s responsibility to ensure that only correct assertions are made. The compiler is free to rely on these assertions, for example by applying optimizations that are only correct if the assertions are valid. [6]

This work was motivated precisely by the desire to fill in this gap, and by Quipper’s demonstration of the power of assertive terminations.

The other important work in this space is Amy *et al.*’s REVERC [1], which builds upon the REVS programming language [8], a small heavily-optimized language for reversible computing. REVERC verifies many of the optimizations from REVS and includes a compiler from boolean expressions to reversible circuits. The validity of this compilation is verified in the F^* programming language [14]. One key challenge in this paper was to port that compiler from a language that uses only classical operations on numbered registers (and whose semantics are therefore in terms of boolean expressions), to a language using higher-order abstract syntax whose denotation is in terms of density matrices (representing pure and mixed quantum states).

The State of QWIRE This paper, and the whole QWIRE project, is a work in progress. QWIRE has been used to verify some interesting programs, including quantum teleportation, Deutsch’s algorithm

and a variety of random number generators (see `HOASProofs.v` in the development). It can also be used to prove the validity of a number of circuit optimizations, such as those of Staton [13] (see `Equations.v`). However, much remains to be done. The authors’ goal is to formally verify all of the claims in this paper, though some work still remains.

In particular, the following lemmas remain to be proved in Coq, by section:

- In Section 3, the proof of the equivalence of the two definitions of `valid_ancillae` (though the this paper does not build on that equivalence); and the proof of Lemma 3.1.
- In Section 4, Lemmas 4.1 to 4.3 and Conjecture 4.4.
- In Section 5, that the `CNOT_at` and `Toffoli_at` circuits, as well as sequencing `;;` and parallel `||` combinators, match their intended semantics.

The next step for QWIRE is to implement and verify circuit optimizations. We already have a number of equivalences we can in principle use to rewrite our circuits, and this work introduces new possible optimization, like reusing ancillae. It also allows us to treat circuits that properly initialize and dispose of ancilla as unitary circuits, allowing for further optimizations. Given that much of the progress towards practical quantum computing comes from increasingly clever optimizations (in tandem with more powerful quantum computers), verified compilation should play an important and exciting role in the near future.

References

- [1] Matthew Amy, Martin Roetteler & Krysta M Svore (2017): *Verified compilation of space-efficient reversible circuits*. In: *International Conference on Computer Aided Verification*, Springer, pp. 3–21, doi:10.1007/978-3-319-63390-9_1.
- [2] Charles H Bennett (1973): *Logical reversibility of computation*. *IBM journal of Research and Development* 17(6), pp. 525–532.
- [3] Coq Development Team (2017): *The Coq Proof Assistant Reference Manual, Version 8.6*. Electronic resource, available from <http://coq.inria.fr>.
- [4] Alexis De Vos (2011): *Reversible computing: fundamentals, quantum computing, and applications*. John Wiley & Sons.
- [5] Alexander Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger & Benoît Valiron (2013): *An Introduction to Quantum Programming in Quipper*. In: *Proceedings of the 5th International Conference on Reversible Computation, Lecture Notes in Computer Science* 7948, pp. 110–124, doi:10.1007/978-3-642-38986-3_10.
- [6] Alexander Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger & Benoît Valiron (2013): *Quipper: A Scalable Quantum Programming Language*. In: *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2013*, pp. 333–342, doi:10.1145/2491956.2462177.
- [7] Bernhard Ömer (2005): *Classical concepts in quantum programming*. *International Journal of Theoretical Physics* 44(7), pp. 943–955.
- [8] Alex Parent, Martin Roetteler & Krysta M Svore (2017): *REVS: A Tool for Space-Optimized Reversible Circuit Synthesis*. In: *International Conference on Reversible Computation*, Springer, pp. 90–101.
- [9] Jennifer Paykin, Robert Rand & Steve Zdancewic (2017): *QWIRE: A Core Language for Quantum Circuits*. In: *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017*, ACM, New York, NY, USA, pp. 846–858, doi:10.1145/3009837.3009894.
- [10] Kalyan S Perumalla (2013): *Introduction to reversible computing*. CRC Press.

- [11] Robert Rand, Jennifer Paykin & Steve Zdancewic (2017): *QWIRE Practice: Formal Verification of Quantum Circuits in Coq*. In: *Quantum Physical and Logic*.
- [12] P. W. Shor (1999): *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. *SIAM Review* 41, pp. 303–332, doi:[10.1137/S0036144598347011](https://doi.org/10.1137/S0036144598347011).
- [13] Sam Staton (2015): *Algebraic Effects, Linearity, and Quantum Programming Languages*. In: *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '15*, ACM, New York, NY, USA, pp. 395–406, doi:[10.1145/2676726.2676999](https://doi.org/10.1145/2676726.2676999).
- [14] Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss et al. (2016): *Dependent types and multi-monadic effects in F**. In: *ACM SIGPLAN Notices*, 51, ACM, pp. 256–270.
- [15] Tommaso Toffoli (1980): *Reversible computing*. In: *International Colloquium on Automata, Languages, and Programming*, Springer, pp. 632–644.