

MIT Open Access Articles

Tower: Data Structures in Quantum Superposition

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Yuan, Charles and Carbin, Michael. 2022. "Tower: Data Structures in Quantum Superposition." PACM on Programming Languages.

As Published: <https://doi.org/10.1145/3563297>

Publisher: ACM

Persistent URL: <https://hdl.handle.net/1721.1/146520>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license





Tower: Data Structures in Quantum Superposition

CHARLES YUAN, MIT CSAIL, USA

MICHAEL CARBIN, MIT CSAIL, USA

Emerging quantum algorithms for problems such as element distinctness, subset sum, and closest pair demonstrate computational advantages by relying on *abstract data structures*. Practically realizing such an algorithm as a program for a quantum computer requires an efficient implementation of the data structure whose operations correspond to *unitary operators* that manipulate quantum *superpositions* of data.

To correctly operate in superposition, an implementation must satisfy three properties — *reversibility*, *history independence*, and *bounded-time execution*. Standard implementations, such as the representation of an abstract set as a hash table, fail these properties, calling for tools to develop specialized implementations.

In this work, we present Core Tower, the first language for quantum programming with random-access memory. Core Tower enables the developer to implement data structures as pointer-based, linked data. It features a reversible semantics enabling every valid program to be translated to a unitary quantum circuit.

We present Boson, the first memory allocator that supports reversible, history-independent, and constant-time dynamic memory allocation in quantum superposition. We also present Tower, a language for quantum programming with recursively defined data structures. Tower features a type system that bounds all recursion using classical parameters as is necessary for a program to execute on a quantum computer.

Using Tower, we implement Ground, the first quantum library of data structures, including lists, stacks, queues, strings, and sets. We provide the first executable implementation of sets that satisfies all three mandated properties of reversibility, history independence, and bounded-time execution.

CCS Concepts: • **Theory of computation** → *Data structures design and analysis*; • **Computer systems organization** → **Quantum computing**; • **Software and its engineering** → Language features; Formal language definitions; Memory management.

Additional Key Words and Phrases: quantum programming, data structures, quantum random-access memory, reversible programming, history independence

ACM Reference Format:

Charles Yuan and Michael Carbin. 2022. Tower: Data Structures in Quantum Superposition. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 134 (October 2022), 30 pages. <https://doi.org/10.1145/3563297>

1 INTRODUCTION

Algorithms for quantum computers promise computational advantages in areas ranging from integer factorization [Shor 1997] and search [Grover 1996] to cryptography [Bennett and Brassard 2014], linear algebra [Harrow et al. 2009], physical simulation [Childs et al. 2018; Kassal et al. 2011], and combinatorial optimization [Farhi et al. 2014]. Enabling practical realization of these algorithms on quantum hardware is a major aim of *quantum programming languages*, in which a developer may write programs that execute on a quantum computer.

The principles of quantum mechanics specify how a quantum computer manipulates *quantum states* consisting of *qubits*, the quantum analogs of classical bits. A quantum state exists in a

Authors' addresses: Charles Yuan, MIT CSAIL, 32 Vassar St, Cambridge, MA, 02139, USA, chenhuiy@csail.mit.edu; Michael Carbin, MIT CSAIL, 32 Vassar St, Cambridge, MA, 02139, USA, mcabin@csail.mit.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/10-ART134

<https://doi.org/10.1145/3563297>

superposition, a weighted sum over classical states. The computer may evolve the state using *unitary operators*, or it may *measure* the state, which causes a superposition to assume a classical state with probability derived mathematically from the weight ascribed to that state in the sum.

1.1 Data Structures and Quantum Algorithms

Existing quantum programming languages [Altenkirch and Grattage 2005; Bichsel et al. 2020; Green et al. 2013; Paykin et al. 2017; Selinger and Valiron 2005; Svore et al. 2018; Wecker et al. 2014] build abstractions for individual qubits and primitive data types such as integers. By contrast, emerging quantum algorithms incorporate abstract data structures for sake of implementation and efficiency. The concrete representation of such a data structure classically is a linked structure in random-access memory, defined recursively and constructed by dynamic allocation.

For example, the set data structure can maintain a duplicate-free collection of items, check an element for membership in the collection, and add and remove elements from the collection. Examples of quantum algorithms that depend critically on sets include Ambainis [2003], who presents an $O(n^{2/3})$ quantum algorithm for element distinctness that improves over the $\Omega(n)$ classical bound; Bernstein et al. [2013], who present a quantum algorithm for subset sum that beats the classical complexity of $O(2^{n/4})$; and Aaronson et al. [2019], who present an $\tilde{O}(n^{2/3})$ quantum algorithm for closest pair, improving over $O(n \log n)$ classically.

Each of these algorithms relies on an implementation of sets whose operations execute in $O(\log^c n)$ time. Classically, random-access memory is used to attain this efficient time complexity, and these algorithms similarly invoke its quantum analogue. However, to date, no quantum programming language provides abstractions for constructing and manipulating sets and other abstract data types in random-access memory, a hurdle to the realization of these algorithms.

1.2 Three Requirements for Data Structures

Importantly, not all classical data structure implementations are compatible with quantum programming. Ambainis [2003] identifies a data structure appropriate for use in a quantum algorithm as one that satisfies three properties: *reversibility*, *history independence* and *bounded-time execution*.

Reversible Operation. The operations of the data structure must each be *reversible*, meaning it can be executed forwards or backwards to produce or undo its effect on the program state.

In order to execute on a quantum computer, a program must first be translated into a sequence of primitive *quantum gates*. Each gate corresponds to either a measurement or a unitary operator over the program state. Conventional operations over a data structure in superposition cannot safely invoke measurement, because measurement does not preserve the superposition of program data. Instead, it irreversibly collapses the data into a classical state, which would cause algorithms such as Aaronson et al. [2019]; Ambainis [2003]; Bernstein et al. [2013] to return incorrect results.

Thus, when a program executes an operation on data in superposition, the operation must correspond to a unitary operator, which has an inverse by definition. As a consequence, the operations of a data structure in superposition must be reversible.

History-Independent Representation. The representation of the data structure must be *history-independent*, meaning that each instance of the data structure is assigned a single unique physical representation in memory that is insensitive to the history of operations performed on it.

History independence is essential for the correctness or performance of algorithms such as Aaronson et al. [2019]; Ambainis [2003]; Bernstein et al. [2013]. The reason is that they rely on quantum *interference*, the phenomenon in which weights of identical states in a superposition combine or cancel, altering the outcomes of measurements and the output of the algorithm.

For example, the algorithm of [Ambainis \[2003\]](#) generates a superposition over sets of integers. Suppose that among these are two instances of the set $\{1, 2\}$ whose weights in the superposition have equal complex magnitude but opposite sign. Then, the correctness of the algorithm relies on quantum interference to guarantee that the probability of observing the set $\{1, 2\}$ cancels to zero.

However, suppose that sets are concretely represented using linked lists, such that the first instance is represented as $[1, 2]$ and the second as $[2, 1]$. Because the memory representations of these two instances are physically distinct, their weights do not cancel, leading the algorithm to produce incorrect outputs. Thus, for the algorithm to execute correctly, each semantic instance of an abstract data structure must be assigned a unique underlying physical representation.

Bounded-Time Execution. A data structure operation must execute in *bounded time*, meaning its execution time is bounded by classical parameters not dependent on quantum data.

A recursive definition of a data structure typically corresponds to recursive implementations of its operations. Recursion enables an operation to use variable time complexity to traverse and manipulate a variable-sized data structure. However, a quantum computer does not natively support recursive control flow and must first classically convert the program to a loop-free form. Hence, all recursion in the program must be bounded by classically determined parameters.

1.3 Tower: Data Structures in Quantum Superposition

Core Tower. We present Core Tower, the first language that enables quantum programming with data in random-access memory. It enables the construction and manipulation of pointer-based, linked data, by means of operations for integers, pointers, and control flow.

Reversibility. In Core Tower, every program that can execute forwards can also execute backwards. A developer may compute an expression by assigning it to a variable, or *uncompute* an expression, which reverses the computation, by *un-assigning* it from a variable. Other features enable the developer to swap two variables, or to swap a variable with the content of memory at an address.

We present a fully reversible operational semantics of Core Tower. Each operator has a reverse within the syntax, meaning a developer may write a program and derive another program with reversed semantics using the syntactic transformation of *program inversion*.

We provide a translation from a valid Core Tower program to a unitary operator that can execute on a quantum computer. Specifically, the operator corresponding to a Core Tower program acts uniformly across each classical program state in a superposition. The operator may then be embedded within a larger quantum algorithm that creates and manipulates weights of superpositions.

Boson. We present Boson, the first memory allocator that enables reversible, history-independent, and constant-time dynamic memory allocation in superposition. A developer may invoke Boson to write a history-independent implementation of a data structure.

History Independence. To leverage Boson, the developer constructs dynamically allocated, linked data whose linking structure is in one-to-one correspondence with semantic data structure instances. At runtime, Boson stores data into superpositions of allocation sites that are chosen to guarantee that the instance's physical memory representation, including the values of pointers, is unique.

Classically, a memory allocator may establish history independence by randomizing allocation addresses. The implementation of Boson lifts randomization to quantum superposition via *symmetrization*, an approach from quantum simulation [[Abrams and Lloyd 1997](#); [Berry et al. 2017](#)].

Tower. We combine Core Tower and Boson to build Tower, a programming language for quantum programming with data structures. Tower features co-recursive types and recursive functions that enable the developer to recursively define data structures and operations.

Bounded-Time Execution. The type system of Tower ensures that all recursive functions are bounded in size by classical parameters. A Tower program lowers to Core Tower by unrolling recursive calls, meaning that it corresponds to a bounded-size quantum circuit.

1.4 Ground: Data Structure Library

We use Tower to implement Ground, the first quantum library of data structures, consisting of lists, stacks, queues, strings, and sets. Of note, we demonstrate the first executable implementation of sets that is reversible, history-independent, and bounded-time, using radix trees [Morrison 1968] as proposed by Bernstein et al. [2013]. Prior to this work, we are not aware of any existing quantum implementation of the set data structure that satisfies all three properties.¹

As a case study, we describe three challenges we encountered during the development process that do not arise in classical programming, as well as techniques to overcome them:

- *Recursive uncomputation* means that functions that uncompute the results of recursive calls may incur exponential slowdown in time complexity. We instead structure a program to avoid having recursive calls be temporary values.
- *Mutated uncomputation* means that directly uncomputing temporary variables after the program state has been mutated may result in incorrect outcomes. We instead correctly uncompute values using context from the updated program state.
- *Branch sequentialization* means that the time complexity of a conditionally branching quantum program is equal to the sum and not the maximum of its branches. To prevent inefficiency, we structure a program to avoid recursion under branches.

1.5 Contributions

In this work, we present the following contributions:

- *Core Tower* (Section 4). We present Core Tower, the first language for quantum programming with random-access memory. We formalize its type system, reversible operational semantics, program inversion transformation, and translation into a unitary quantum circuit.
- *Boson* (Section 5). We present Boson, the first allocator that enables reversible, history-independent, and constant-time dynamic memory allocation in superposition. We present the implementation of Boson using techniques from quantum simulation.
- *Tower* (Section 6). We present Tower, a language that enables bounded-time quantum programming with recursively defined data structures. We formalize its type system, which verifies that recursion is classically bounded, and its lowering into Core Tower.
- *Ground* (Section 7.2). We present Ground, the first quantum library of abstract data structures, including lists, stacks, queues, strings, and sets. We present the first executable set implementation satisfying reversibility, history independence, and bounded-time execution.
- *Case Study* (Sections 7.3 to 7.5). We characterize three challenges we faced in the library implementation that do not arise in classical programming – recursive uncomputation, mutated uncomputation, and branch sequentialization – and techniques to overcome them.

Summary. Our work establishes programming abstractions for data structures in quantum superposition in the form of language design and provides the first quantum library of fundamental data structures. Together, Tower, Boson, and Ground enable developers to implement algorithms and construct quantum software utilizing familiar and useful data structures.

¹Concurrent to this work, Gidney [2022] provides an implementation of a quantum dictionary data structure. Their implementation does not use random-access memory and achieves linear time complexity, whereas our implementation achieves poly-logarithmic time complexity. Their work was made public while this work was under conference review.

2 BACKGROUND ON QUANTUM COMPUTATION

The following is an overview of key concepts in quantum computation relevant to this work. A comprehensive reference may be found in [Nielsen and Chuang \[2010\]](#).

Qubit. The basic unit of quantum information is the *qubit*, a linear combination $\gamma_0 |0\rangle + \gamma_1 |1\rangle$ known as a *superposition*, where $|0\rangle$ and $|1\rangle$ are *basis states* and $\gamma_0, \gamma_1 \in \mathbb{C}$ are *amplitudes* satisfying $|\gamma_0|^2 + |\gamma_1|^2 = 1$ describing relative weights of basis states. Examples of qubits include the classical zero bit $|0\rangle$, classical one bit $|1\rangle$, and the superposition states $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $\frac{1}{\sqrt{2}}(i|0\rangle - |1\rangle)$.

Quantum State. A 2^n -dimensional *quantum state* $|\psi\rangle$ is a superposition over n -bit strings. For example, $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ is a quantum state over two qubits. Multiple qubits form a quantum state system by means of the tensor product \otimes , e.g. the state $|01\rangle$ is equal to the product $|0\rangle \otimes |1\rangle$. As is standard in quantum computing, we also use the notation $|x, y\rangle$ to represent $|x\rangle \otimes |y\rangle$.

Unitary Operators. A *unitary operator* U is a linear operator on quantum states that preserves inner products and whose inverse is its Hermitian adjoint. Examples of unitary operators include single-qubit *quantum gates* such as:

- X – bit-flip (NOT) gate, which maps $|0\rangle \mapsto |1\rangle$ and vice versa;
- Z – phase-flip gate, which leaves $|0\rangle$ unchanged and maps $|1\rangle \mapsto -|1\rangle$;
- H – Hadamard gate, which maps $|0\rangle \mapsto \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle \mapsto \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.

Other examples include two-qubit gates, such as controlled-NOT, controlled-Z, and SWAP. The controlled gates perform NOT or Z on their target qubit if their control qubit is in state $|1\rangle$, and the SWAP gate swaps two qubits in a quantum state.

Measurement. A quantum *measurement* is a probabilistic operation over quantum states. When a qubit $\gamma_0 |0\rangle + \gamma_1 |1\rangle$ is measured,² the outcome is $|0\rangle$ with probability $|\gamma_0|^2$ and $|1\rangle$ with probability $|\gamma_1|^2$. Measuring a qubit within a quantum state causes the entire state to probabilistically assume one of two outcome states. To define the outcomes of measuring the first qubit of $|\psi\rangle$, we first rewrite the state into the form $|\psi\rangle = \gamma_0 |0\rangle \otimes |\psi_0\rangle + \gamma_1 |1\rangle \otimes |\psi_1\rangle$. Then, with probability $|\gamma_0|^2$ the measurement outcome is $|0\rangle \otimes |\psi_0\rangle$, and with probability $|\gamma_1|^2$ the outcome is $|1\rangle \otimes |\psi_1\rangle$. Outcomes of measuring multiple qubits or qubits other than the first are defined analogously.

Random Access. Quantum random-access memory (qRAM)³ [[Giovannetti et al. 2008](#)] enables manipulating data at addresses in superposition. Let k be the system word size. The single-qubit *quantum random-access gate* is defined to be the unitary operation that maps [[Ambainis 2003](#)]:

$$|i, b, z_1, \dots, z_m\rangle \mapsto |i, z_i, z_1, \dots, z_{i-1}, b, z_{i+1}, \dots, z_m\rangle$$

where i is a k -bit integer, b is a bit, and $|z_i, \dots, z_m\rangle$ is the qRAM – an array of m qubits. The effect of this gate is to swap the data at position i in the array z with the data in the register b .

In this work, we utilize a multi-qubit random-access gate, in which b and each z_i is a k -bit string rather than a single bit. Such a gate may be implemented as a straightforward extension of leading proposals for hardware qRAM [[Arunachalam et al. 2015](#); [Matteo et al. 2020](#); [Paler et al. 2020](#)].

We account for the cost of data structure operations using a model in which quantum random access takes $O(1)$ time. This accounting facilitates a direct comparison to the classical RAM cost model, in which dereferencing a pointer is modeled as constant-time regardless of the underlying hardware design. It is possible to use an alternative cost model in which quantum random access

²In this work, we refer only to projective computational (i.e. 0/1) basis measurements.

³The abbreviation qRAM is not to be confused with the QRAM model of computation by [Knill \[1996\]](#), which does not incorporate quantum random access.

takes $O(\log m)$ time, which is used by Aaronson et al. [2019]; Ambainis [2003]; Bernstein et al. [2013]; Buhrman et al. [2021]. To account under this model, the reported complexity of each data structure operation should be multiplied by a poly-logarithmic time factor.

3 QUANTUM DATA STRUCTURES IN TOWER

In this section, we present data structures in quantum superposition and justify the properties of reversibility, history independence, and bounded-time execution. We illustrate how a developer uses Tower to implement fundamental data structures satisfying these properties.

3.1 Linked Data in Superposition

In classical computing, linked data forms the basis of numerous abstract data structures. For example, linked lists may be used to concretely implement abstract stacks, queues, and sets. We use the notation $l \mapsto [1, 2]$ to represent a classical program state⁴ in which the variable l denotes a linked list storing the integers 1 and 2.

Superposition. In a Tower program, the state of all program variables and memory exists in quantum superposition. For example, the quantum state of the program may be in a uniform superposition of three classical states in which l denotes different linked lists:


$$\frac{1}{\sqrt{3}} \left(|l \mapsto [] \rangle + |l \mapsto [1, 2] \rangle + |l \mapsto [1, \dots, 100] \rangle \right)$$

Data Structure Operations. We may lift a classical operation on a list to a unitary operation on a list in superposition. For example, the operation `push_front` prepends an element onto every list in the superposition. Executing `push_front(l, 6)` would produce a new quantum state:

$$\frac{1}{\sqrt{3}} \left(|l \mapsto [6] \rangle + |l \mapsto [6, 1, 2] \rangle + |l \mapsto [6, 1, \dots, 100] \rangle \right)$$

Memory Representation. The memory representation of a singly-linked list is a collection of nodes, each consisting of content data and a pointer to the next node, or null if there is none.

Within a classical program state, `null` denotes the empty list, and a pointer to a memory location at which a pair of l and `null` are stored denotes the singleton list containing the element 1. The program state $l \mapsto [1, 2]$ is depicted by the following diagram, in which \square denotes `null`:



3.2 Implementing Linked Data Structures in Tower

Tower makes it possible to implement linked data structures in quantum superposition. A key aspect of the design of Tower is a set of programming primitives that are each *reversible*, meaning a program can be executed forwards or backwards to produce or reverse its computational effect.

Reversibility. Reversible primitives correspond to unitary operators that can be realized as gates in a quantum computer (Section 2). Though a quantum computer is also capable of performing irreversible measurement, measurement collapses the program state from superposition. Algorithms such as Aaronson et al. [2019]; Ambainis [2003]; Bernstein et al. [2013] are therefore designed to avoid performing measurement on the data structure in superposition, which if performed would cause the algorithm to produce an incorrect output or lose its efficiency advantage.

Correspondingly, we designed Tower with reversible primitives that make it possible to manipulate data structures without collapsing their superposition.

⁴A classical program state is a bit string, meaning that the set of all program states is finite.

```

1 type list = (uint, ptr<list>);
2 fun push_front(l: ptr<list>, x: uint) {
3   let head <- alloc<list>;
4   l <-> head;          /* Swap */
5   let node <- (x, head);
6   let head -> node.2;  /* Uncompute */
7   *l <-> node;        /* Swap */
8   let node -> default<list>; /* Uncompute */
9   return ();
10 }

```

Fig. 1. Implementation of push_front.

```

1 fun pop_front(l: ptr<list>) -> uint {
2   let node <- default<list>;
3   *l <-> node;
4   let head <- node.2;
5   let x <- node.1; /* Get output */
6   let node -> (x, head);
7   l <-> head;
8   let head -> alloc<list>;
9   return x;
10 }

```

Fig. 2. Implementation of pop_front.




Fig. 3. Program states after executing each line of `push_front(1, 6)`, where `l` denotes `[1, 2]`. In each step, newly introduced features are shown in red, removed features are shown in gray, and swapped features are shown in green. Features used for computing or uncomputing other features are shown in blue.

Implementing push_front. In Figure 1, we present the Tower implementation of the operation `push_front` on linked lists. When executed forwards, `push_front(1, 6)` prepends 6 to the list `l`, and when executed backwards, it removes 6 from the front of `l`. At a high level, the function initializes a new list node that stores `x` and points to the input list, and then points `l` to this node.

In Figure 3, we depict the effect of each line of the procedure diagrammatically. In Figure 3a, we show the initial program state of `push_front(1, 6)`, assuming `l` denotes `[1, 2]`. In Figure 3b, we depict the effect of line 3, which uses the *assignment* operator `<-` to assign the variable `head` to a new list node in memory, allocated using the `alloc` operator and initialized to an empty default value. Line 4 (Figure 3c) uses the *swap* operator `<->` to reversibly swap the values of `head` and `l`. Line 5 (Figure 3d) assigns `node` to be a pair of the input `x` and `head`.

Uncomputation. At this point, the variable `head` is no longer useful. Line 6 (Figure 3e) uses the *un-assignment* operator `->` to *un-assign* `head` from the second field of `node`, meaning that it *uncomputes* the value of the variable `head` and then discards the variable from the program.

Uncomputation [Bennett 1973] is the action of reversing the computation that produces the value of a variable, restoring that variable's value to zero. The reason uncomputation is necessary here is that by the design of Tower, every program is reversible, meaning that 1) a program may only discard a variable after restoring its value to zero, and 2) no language primitive enables the program to generally erase any unspecified value to zero. If the forward execution of the program

were permitted to discard or erase an unspecified value, then the reverse execution would be forced to correctly re-materialize that discarded or erased value, which is impossible.

On line 6, the program uses the un-assignment operator to restore the value of head to zero.⁵ Specifically, the effect of line 6 is defined as the reverse of that of the statement:

```
let head <- node.2;
```

An assignment such as this one may be interpreted as a reversible sequence of steps. First, head is initialized to zero. Then, each bit of head is flipped if and only if the corresponding bit of node.2 is set. The result is that the values of head and node.2 are equal.

The un-assignment on line 6 may be interpreted as the above sequence of steps in reverse. It starts with head and node.2 having equal values, and performs the same bit flips on head conditioned on node.2. The result is that head has a value of zero, which the program safely discards.

Continuing on, line 7 (Figure 3f) swaps node with the contents of memory at address 1. Recalling that after line 4, 1 stores an empty default node, it is then the case that after line 7, 1 points to the new list head and node stores an empty node, which is then uncomputed by line 8 (Figure 3g).

Reversing Operation. The reversibility of Tower programs makes it simple to implement the closely related operation of pop_front, which removes the first element from a list and returns the element. Indeed, the implementation of pop_front is nearly the syntactic reverse of push_front.

In Figure 2, we present pop_front. It first reverses line 8 of push_front, replacing un-assignment with assignment, and then line 7, which remains identical – the reverse of a swap operation is itself. It reverses line 6 of push_front, assigning the second field of node to head. Line 5 of pop_front is new to this function – it obtains x from the first field of node so that x may be returned. The following lines reverse the first three lines of push_front, and the function returns x.

3.3 Dynamic Memory Allocation with Boson

Sets are foundational to programming and critical to algorithms. Surprisingly, many implementations of sets such as lists, hash tables, and binary trees are unsuitable for quantum programming.

Interference. Algorithms such as Aaronson et al. [2019]; Ambainis [2003]; Bernstein et al. [2013]; Buhrman et al. [2021] rely on *interference*, the constructive or destructive interaction of amplitudes in a quantum state. For example, Ambainis [2003] generates a superposition over program states $s \hookrightarrow s_i$, meaning that variable s stores a set of integers s_i . The correctness of the algorithm relies on the fact that if s_1 and s_2 denote sets containing the same integers, then their amplitudes interfere. In particular, suppose that s_1 and s_2 both denote the set $\{1, 2\}$. Then, the algorithm requires that:

$$\gamma |s \hookrightarrow s_1\rangle - \gamma |s \hookrightarrow s_2\rangle = 0$$


History Independence. However, suppose that sets are implemented using linked lists, where s_1 is represented as $[1, 2]$ and s_2 as $[2, 1]$. Because $s \hookrightarrow [1, 2]$ and $s \hookrightarrow [2, 1]$ are physically distinct states, these terms do not cancel, causing the algorithm to produce incorrect outputs.

For the algorithm to produce correct outputs, the implementation of the data structure must satisfy the property of *history independence*. We adapt its definition from Naor and Teague [2001].⁶


Definition 3.1 (History Independence). An implementation of an abstract data structure is *history-independent* if and only if, when any two sequences of abstract operations yield semantically equivalent data structure instances, then the physical representations of the instances are identical.

⁵Tower also supports other operators that abstract the management of temporary values, which we discuss in Section 6. We focus here on the un-assignment operator for explanatory purposes.

⁶This definition is termed the “strong” variant of history independence in the original work. Conceptually similar definitions are found in cryptographic models of computing, such as Wang et al. [2014].



(a) One possible memory representation of $s \leftrightarrow [1, 2]$. (b) A distinct representation of the same state.



(c) Boson's unique superposition representation.

Fig. 4. Distinct physical memory representations of the same program state in a qRAM of size 6.

```

1 fun find_pos[n](l: ptr<list>, x: uint,
2   acc: uint) -> uint {
3   with {
4     let node <- default<list>;
5     *l <-> node;
6     let (this_x, next) <- node;
7     let r <- acc + 1;
8     let eq <- this_x == x;
9   } do if eq {
10    let out <- acc;
11  } else {
12    let out <- find_pos[n-1](next, x, r);
13  }
14  return out;
15 }

```

Fig. 5. Implementation of find_pos.

The implementation of a set as an unsorted linked list is not history-independent, as the lists $[1, 2]$ and $[2, 1]$ both represent the same set $\{1, 2\}$. Similarly, implementations of sets such as balanced binary trees or hash tables⁷ are not history-independent and are thus unsuitable for use in quantum algorithms that leverage interference on data structures.

Unique Linking Structure. A history-independent implementation must assign each semantic instance of the data structure a unique linking structure in the underlying representation. For example, suppose a set is instead implemented as a sorted list without duplicates. The set $\{1, 2\}$ is then uniquely represented as the list $[1, 2]$, ostensibly satisfying history independence.

Memory Addresses. However, a unique linking structure does not by itself guarantee history independence. In the example, the issue is that $s \leftrightarrow [1, 2]$ may still refer to physically distinct program states in which the pointers to the linked list nodes take on different values.

For example, Figure 4a depicts a program state in which s points to memory location 1, which stores value 1 and a pointer to memory location 3, which stores value 2 and null. Figure 4b depicts a state in which s points to memory location 5, which stores 1 and a pointer to memory location 1, which stores 2 and null. These two program states are physically distinct but semantically equivalent, causing quantum interference to fail.

Boson Allocator. We resolve this issue by developing Boson, a memory allocator that ensures that addresses of all dynamic allocations in the program are history-independent. Under Boson, the single unique program state in which s denotes $[1, 2]$ is Figure 4c, in which the linked list nodes are stored in a uniform superposition of all possible allocation sites.

Radix Trees. The sorted list implementation of sets has $O(n)$ operation complexity, which is asymptotically inefficient. By contrast, Bernstein et al. [2013] propose using a radix tree [Morrison 1968], whose linking structure is in one-to-one correspondence with sets and which supports operations in $O(\log n)$ time. In this work, we construct an efficient and history-independent implementation of sets by combining the radix tree with the Boson dynamic memory allocator.

⁷Other reasons why hash tables are unsuitable are 1) removal from a hash table is not directly reversible, and 2) all operations assume their worst case $O(n)$ rather than average case $O(1)$ complexity when all recursion is classically unfolded.

3.4 Recursive Data Structure Operations in Tower

The last tool needed to implement operations on linked data structures is recursion. We demonstrate recursion in Tower by implementing the function `find_pos`, which returns the position of an integer x in a list l , assuming it is present in the list. This operation traverses to the end of the list, and a natural implementation performs this traversal recursively.

Recursion Bounds. Classically, a recursive call pushes a return address onto the program stack and sets the program counter to the entry point of the function. By contrast, a quantum processor does not possess a program counter or the ability to loop potentially infinitely. All recursive procedures must instead be unrolled to a fixed, classically known depth, and inlined.

Tower guarantees that fully unrolling a recursive function is always possible using a classical *recursion bound* specifying the number of recursion levels. The Tower type system verifies that in a program, all recursive function calls satisfy their respective recursion bound.

Implementing find_pos. In Figure 5, we implement `find_pos` in Tower by recursively traversing l and comparing each element with x . The implementation takes an accumulator argument `acc`, such that calling the function with zero for `acc` returns the position of x .

On line 1, the bound $[n]$ denotes that `find_pos` may recursively call itself at most n times. Lines 3–13 feature a `with-do` statement, syntactic sugar that facilitates uncomputation. Tower executes this statement by executing the body of `with`, then the `do`, then the reverse of the `with`. In this program, reversing the `with`-block uncomputes `node`, `this_x`, `next`, `r`, and `eq`.

Line 5 swaps the contents of memory at address l with an empty node `node`, and line 6 extracts from `node` the content `this_x` and pointer `next` to the tail. Line 8 computes `eq`, a flag for whether `this_x` is equal to x . If it is, line 10 returns the current position, and if it is not, line 12 recursively calls `find_pos` on the tail of the list, with a recursion bound that has decreased by one.

4 CORE TOWER: REVERSIBLE AND CIRCUIT-BASED SEMANTICS

This section presents Core Tower, an imperative programming language that enables quantum programming with random-access memory. The operational semantics of the language is fully reversible, enabling a program to execute forwards or backwards. A Core Tower program admits *inversion*, which syntactically transforms it into another program with reversed semantics. A program also admits a syntactic translation into a unitary quantum circuit. In Section 6, we extend Core Tower with co-recursive data types and recursion to produce Tower.

4.1 Syntax

The syntax of Core Tower consists of types, values, expressions, and statements.

Type $\tau ::= () \mid \text{uint} \mid \text{bool} \mid (\tau_1, \tau_2) \mid \text{ptr}(\tau)$

Value $v ::= x \mid () \mid (x_1, x_2) \mid \bar{n} \mid \text{true} \mid \text{false} \mid \text{null}_\tau \mid \text{ptr}_\tau[p] \quad (n \in \text{UInt}, p \in \text{Addr})$

Expression $e ::= v \mid \pi_1(x) \mid \pi_2(x) \mid \text{uop } x \mid x_1 \text{ bop } x_2$

Operator $\text{uop} ::= \text{not} \mid \text{test} \quad \text{bop} ::= \&\& \mid \mid \mid + \mid - \mid *$

Statement $s ::= \text{skip} \mid s_1; s_2 \mid x \leftarrow e \mid x \rightarrow e \mid x_1 \Leftrightarrow x_2 \mid *x_1 \Leftrightarrow x_2 \mid \text{if } x \text{ then } s \text{ end}$

Types. The types τ are unit, unsigned integers, Booleans, products, and pointers.

Values. The language distinguishes syntactic values v , which are variables, unit, pairs of variables, numeric, Boolean, and pointer literals. Numeric literals n are elements of a set UInt of unsigned integers in the range $[0, 2^k - 1]$ and pointer literals p are elements of a set Addr of addresses in the range $[1, 2^k - 1]$, where the parameter k is the word size of the system.

$\frac{\text{TV-VAR}}{\Gamma, x : \tau \vdash x : \tau}$	$\frac{\text{TV-PAIR} \quad \Gamma \vdash x_1 : \tau_1 \quad \Gamma \vdash x_2 : \tau_2}{\Gamma \vdash (x_1, x_2) : (\tau_1, \tau_2)}$	$\frac{\text{TV-NUM}}{\Gamma \vdash \bar{n} : \text{uint}}$	$\frac{\text{TV-BOOL} \quad b \in \{\text{true}, \text{false}\}}{\Gamma \vdash b : \text{bool}}$
$\frac{\text{TV-NUL}}{\Gamma \vdash \text{null}_\tau : \text{ptr}(\tau)}$	$\frac{\text{TV-Ptr}}{\Gamma \vdash \text{ptr}_\tau[p] : \text{ptr}(\tau)}$	$\frac{\text{TE-PROJ}}{\Gamma \vdash x : (\tau_1, \tau_2)} \quad \Gamma \vdash \pi_i(x) : \tau_i$	$\frac{\text{TE-NOT}}{\Gamma \vdash x : \text{bool}} \quad \Gamma \vdash \text{not } x : \text{bool}$
$\frac{\text{TE-TEST} \quad \Gamma \vdash x : \tau \quad \tau \in \{\text{uint}, \text{ptr}(\tau')\}}{\Gamma \vdash \text{test } x : \text{bool}}$	$\frac{\text{TE-LOP} \quad \Gamma \vdash x_1 : \text{bool} \quad \Gamma \vdash x_2 : \text{bool} \quad bop \in \{\&\&, \}}{\Gamma \vdash x_1 \text{ } bop \text{ } x_2 : \text{bool}}$	$\frac{\text{TE-AOP} \quad \Gamma \vdash x_1 : \text{uint} \quad \Gamma \vdash x_2 : \text{uint} \quad bop \in \{+, -, *\}}{\Gamma \vdash x_1 \text{ } bop \text{ } x_2 : \text{uint}}$	

Fig. 6. Selected typing rules in Core Tower. The full definition is presented in Appendix A.

$\frac{\text{S-SKIP}}{\Gamma \vdash \text{skip} \dashv \Gamma}$	$\frac{\text{S-SEQ} \quad \Gamma \vdash s_1 \dashv \Gamma' \quad \Gamma' \vdash s_2 \dashv \Gamma''}{\Gamma \vdash s_1; s_2 \dashv \Gamma''}$	$\frac{\text{S-ASSIGN} \quad \Gamma \vdash e : \tau \quad x \notin \Gamma}{\Gamma \vdash x \leftarrow e \dashv \Gamma, x : \tau}$	$\frac{\text{S-UNASSIGN} \quad \Gamma \vdash e : \tau \quad x \notin \Gamma}{\Gamma, x : \tau \dashv x \rightarrow e \dashv \Gamma}$
$\frac{\text{S-SWAP} \quad \Gamma \vdash x_1 : \tau \quad \Gamma \vdash x_2 : \tau}{\Gamma \vdash x_1 \Leftrightarrow x_2 \dashv \Gamma}$	$\frac{\text{S-MEMSWAP} \quad \Gamma \vdash x_1 : \text{ptr}(\tau) \quad \Gamma \vdash x_2 : \tau}{\Gamma \vdash *x_1 \Leftrightarrow x_2 \dashv \Gamma}$	$\frac{\text{S-IF} \quad \Gamma \vdash s \dashv \Gamma \quad \Gamma \vdash x : \text{bool} \quad x \notin \text{mod}(s)}{\Gamma \vdash \text{if } x \text{ then } s \text{ end} \dashv \Gamma}$	

Fig. 7. Well-formation rules for statements in Core Tower.

Expressions. Expressions e are values, projections, and unary and binary operators. These operators are standard logical and arithmetic operators on Booleans and integers. The operation `test` tests whether an integer is equal to zero or a pointer is equal to null.

Statements. Statements s include `skip`, which has no effect, and composition $s_1; s_2$ of statements. The assignment statement $x \leftarrow e$ initializes a variable x and evaluates e , storing the result in x . The un-assignment statement $x \rightarrow e$ is unique to reversible programming — it evaluates e in reverse to restore x back to its default state, and discards x . The swap statement $x_1 \Leftrightarrow x_2$ exchanges the contents of variables x_1 and x_2 . The swap-with-memory statement $*x_1 \Leftrightarrow x_2$ exchanges the contents of memory at address x_1 with the contents of x_2 . Finally, the statement `if x then s end` has the effect of s whenever the condition x is true.

4.2 Type System

The type system of Core Tower assigns a type to a value or expression and determines whether a statement is well-formed and corresponds to a valid quantum program.

Typing. In Figure 6, we define the typing judgments for values and expressions. A context Γ is a mapping from variables x to types τ . The judgment $\Gamma \vdash v : \tau$ states that under context Γ , the value v has type τ . The judgment $\Gamma \vdash e : \tau$ states that under Γ , the expression e has type τ .

In Figure 7, we define the judgment $\Gamma \vdash s \dashv \Gamma'$, stating that under Γ , statement s is well-formed and yields new context Γ' . The statement `skip` has no effect, and sequencing of s_1 and s_2 composes their effect on the context. Assignment $x \leftarrow e$ is well-defined when Γ does not contain x , and produces a context giving x the type of e . Un-assignment is the reverse, removing x from the context. A swap requires its arguments to have the same type, and a swap with memory requires the first argument to be a pointer to the type of the second.

$$\begin{array}{c}
\text{SE-VAR} \\
\frac{R[x'] = v \quad \cdot \vdash v : \tau}{\langle x', R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle x', R[x \leftarrow v] \rangle} \\
\\
\text{SE-PAIR} \\
\frac{R[x_1] = v_1 \quad R[x_2] = v_2 \quad \cdot \vdash v_1 : \tau_1 \quad \cdot \vdash v_2 : \tau_2}{\langle (x_1, x_2), R[x \leftarrow (z(\tau_1), z(\tau_2))] \rangle \Downarrow_x \langle (x_1, x_2), R[x \leftarrow (v_1, v_2)] \rangle} \\
\\
\text{SE-VAL} \\
\frac{v \neq x' \quad v \neq (x_1, x_2) \quad \cdot \vdash v : \tau}{\langle v, R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle v, R[x \leftarrow v] \rangle} \\
\\
\text{SE-PROJ} \\
\frac{R[x'] = (v_1, v_2) \quad \cdot \vdash v_i : \tau}{\langle \pi_i(x'), R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle \pi_i(x'), R[x \leftarrow v_i] \rangle} \\
\\
\text{SE-NOT} \\
\frac{R[x'] = b \quad b \in \{\text{true}, \text{false}\}}{\langle \text{not } x', R[x \leftarrow \text{false}] \rangle \Downarrow_x \langle \text{not } x', R[x \leftarrow \neg b] \rangle} \\
\\
\text{SE-LOP} \\
\frac{R[x_1] = b_1 \quad R[x_2] = b_2 \quad \text{bop} \in \{\&\&, ||\} \quad b_1, b_2 \in \{\text{true}, \text{false}\}}{\langle x_1 \text{ bop } x_2, R[x \leftarrow \text{false}] \rangle \Downarrow_x \langle x_1 \text{ bop } x_2, R[x \leftarrow b_1 \text{ bop } b_2] \rangle} \\
\\
\text{SE-AOP} \\
\frac{R[x_1] = \bar{n}_1 \quad R[x_2] = \bar{n}_2 \quad \text{bop} \in \{+, -, *\}}{\langle x_1 \text{ bop } x_2, R[x \leftarrow \bar{0}] \rangle \Downarrow_x \langle x_1 \text{ bop } x_2, R[x \leftarrow n_1 \text{ bop } n_2] \rangle}
\end{array}$$

Fig. 8. Selected forward step rules of expressions in Core Tower. The full definition is presented in Appendix A.

An if-statement requires its branch s to have no net effect on the context. For reversibility, it requires the condition to be a Boolean that is not modified by s , as defined by the judgement:

$$\begin{array}{ll}
\text{mod}(\text{skip}) = \emptyset & \text{mod}(x_1 \Leftrightarrow x_2) = \{x_1, x_2\} \\
\text{mod}(s_1; s_2) = \text{mod}(s_1) \cup \text{mod}(s_2) & \text{mod}(*x_1 \Leftrightarrow x_2) = \{x_2\} \\
\text{mod}(x \leftarrow e) = \text{mod}(x \rightarrow e) = \{x\} & \text{mod}(\text{if } x \text{ then } s \text{ end}) = \text{mod}(s)
\end{array}$$

4.3 Reversible Operational Semantics

The operational semantics of Core Tower describes reversible operations that may be lifted to act in superposition over a quantum state. The semantics specifies how the machine state evolves during program evaluation while preserving sufficient information to enable reversibility.

Machine State. The semantics transitions over *machine states* $\langle s, R, M \rangle$. Machine states consist of a program s , a *register file* R mapping variables to values, and a *memory* M mapping addresses to values. The register file R corresponds to the main quantum registers over which we may perform arbitrary gates, while the memory M corresponds to the qRAM. We use Γ_R to denote the typing context formed by mapping each variable in R to the type of the value stored.

Default Value. The default value of a type τ , denoted $z(\tau)$, is the initial value of a variable of type τ and corresponds to a convenient memory representation of constant zeroes of appropriate size:

$$\begin{array}{lll}
z(()) = () & z(\text{bool}) = \text{false} & z(\text{ptr}(\tau)) = \text{null}_\tau \\
z(\text{uint}) = \bar{0} & z((\tau_1, \tau_2)) = (z(\tau_1), z(\tau_2)) &
\end{array}$$

The number of bits required to represent a type is the same across a superposition of values of that type, and can be determined statically. When τ_1 and τ_2 are represented using the same number of bits, their default values $z(\tau_1)$ and $z(\tau_2)$ are physically identical. For example, $\bar{0}$ and null_τ are identical, assuming that integers and pointers are represented using the same number of bits.

Forward Evaluation. In Figure 8, we define the forward expression evaluation judgment $\langle e, R[x \leftarrow z] \rangle \Downarrow_x \langle e, R[x \leftarrow v] \rangle$. This judgment states that expression e evaluates to value v , which is stored in register x , whose initial state was z . This judgment maintains reversibility by ensuring that 1) e is not destroyed by evaluation, 2) v is stored into a particular location x , and 3) the initial state z of variable x is a constant that remains known after this operation.

$$\begin{array}{c}
\text{SS-SEQSTEP} \\
\frac{\langle s_1, R, M \rangle \mapsto \langle s'_1, R', M' \rangle}{\langle s_1; s_2, R, M \rangle \mapsto \langle s'_1; s_2, R', M' \rangle} \\
\\
\text{SS-UNASSIGN} \\
\frac{\Gamma_R \vdash e : \tau \quad \langle e, R[x \leftarrow v] \rangle \uparrow_x \langle e, R[x \leftarrow z(\tau)] \rangle}{\langle x \rightarrow e, R[x \leftarrow v], M \rangle \mapsto \langle \text{skip}, R, M \rangle} \\
\\
\text{SS-MEMSWAPNULL} \\
\frac{R[x_1] = \text{null}_\tau}{\langle *x_1 \Leftrightarrow x_2, R, M \rangle \mapsto \langle \text{skip}, R, M \rangle} \\
\\
\text{SS-ASSIGN} \\
\frac{\Gamma_R \vdash e : \tau \quad \langle e, R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle e, R[x \leftarrow v] \rangle}{\langle x \leftarrow e, R, M \rangle \mapsto \langle \text{skip}, R[x \leftarrow v], M \rangle} \\
\\
\text{SS-SWAP} \\
\frac{}{\langle x_1 \Leftrightarrow x_2, R[x_1, x_2 \leftarrow v_1, v_2], M \rangle \mapsto \langle \text{skip}, R[x_1, x_2 \leftarrow v_2, v_1], M \rangle} \\
\\
\text{SS-MEMSWAPPTR} \\
\frac{R[x_1] = \text{ptr}_\tau[p]}{\langle *x_1 \Leftrightarrow x_2, R[x_2 \leftarrow v], M[p \leftarrow v'] \rangle \mapsto \langle \text{skip}, R[x_2 \leftarrow v'], M[p \leftarrow v] \rangle}
\end{array}$$

Fig. 9. Selected forward step rules of statements in Core Tower. The full definition is presented in Appendix A.

$$\begin{array}{c}
\text{RS-SEQSTEP} \\
\frac{\langle s_2, R', M' \rangle \mapsto_R \langle s'_2, R, M \rangle}{\langle s_1; s_2, R', M' \rangle \mapsto_R \langle s_1; s'_2, R, M \rangle} \\
\\
\text{RS-UNASSIGN} \\
\frac{\Gamma_R \vdash e : \tau \quad \langle e, R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle e, R[x \leftarrow v] \rangle}{\langle x \rightarrow e, R, M \rangle \mapsto_R \langle \text{skip}, R[x \leftarrow v], M \rangle} \\
\\
\text{RS-ASSIGN} \\
\frac{\Gamma_R \vdash e : \tau \quad \langle e, R[x \leftarrow v] \rangle \uparrow_x \langle e, R[x \leftarrow z(\tau)] \rangle}{\langle x \leftarrow e, R[x \leftarrow v], M \rangle \mapsto_R \langle \text{skip}, R, M \rangle} \\
\\
\text{RS-SWAP} \\
\frac{}{\langle x_1 \Leftrightarrow x_2, R[x_1, x_2 \leftarrow v_2, v_1], M \rangle \mapsto_R \langle \text{skip}, R[x_1, x_2 \leftarrow v_1, v_2], M \rangle}
\end{array}$$

Fig. 10. Selected reverse step rules of statements in Core Tower. The full definition is provided in Appendix A.

The first evaluation rule specifies that evaluating a variable x' copies its value into a destination variable x that starts in default state. The following three rules similarly specify copying a known constant v , or a pair of values, or one field of a pair, into x . The rules for unary and binary operators update x based on a criterion computed from the arguments, and the rule for arithmetic operators updates x to an arithmetic result over the arguments.

Reverse Evaluation. In Appendix A, we define the reverse expression evaluation judgment $\langle e, R[x \leftarrow v] \rangle \uparrow_x \langle e, R[x \leftarrow z] \rangle$. This judgment states that given the result v of evaluating e , stored in variable x , we may reverse the evaluation of e and restore x back to its initial state z . All rules defining this judgment are reverses of the rules defining the forward evaluation.

Forward Step. In Figure 9, we define the forward statement step judgment $\langle s, R, M \rangle \mapsto \langle s', R', M' \rangle$. This judgment states that given machine state $\langle R, M \rangle$, executing statement s results in new machine state $\langle R', M' \rangle$ and new statement s' to be executed next.

The first two rules skip the skip statement and execute sequences left to right. The rule $x \leftarrow e$ first allocates a variable x in default state, and then forward evaluates e , storing the result v into x . The rule for un-assignment $x \rightarrow e$ reverses the evaluation of e to restore the state of x from v back to the default state, and then deallocates x . The rule for $x_1 \Leftrightarrow x_2$ exchanges the values of x_1 and x_2 in R . The rule for swapping with memory at a null pointer has no effect. The rule to swap x_2 with memory at a valid pointer p swaps the value of x_2 in R with the contents of M at address p . Finally, the two rules for `if` either transition to the body s if the condition passes or to `skip` if it fails.

Reverse Step. In Figure 10, we define the reverse statement step judgment $\langle s, R', M' \rangle \mapsto_R \langle s', R, M \rangle$. This judgment states that given machine state $\langle R', M' \rangle$, executing statement s in reverse results in new machine state $\langle R, M \rangle$ and new statement s' to be executed next.

Its rules are the reverse of the forward step. Sequential statements execute from right to left. The roles of $x \leftarrow e$ and $x \rightarrow e$ are flipped, with the latter now evaluating e forward and the former evaluating backward. Because the reverse of a swap operation is itself, the rules for swap statements are the same. Executing an `if` in reverse requires reversing its body.

4.4 Soundness

We now show that Core Tower satisfies key soundness properties in both execution directions:

Definition 4.1 (Valid Machine State). A machine state $\langle s, R, M \rangle$ is *valid* under initial context Γ and final context Γ' , denoted $\Gamma \vdash \langle s, R, M \rangle \dashv \Gamma'$, if and only if $\Gamma_R = \Gamma$ and $\Gamma \vdash s \dashv \Gamma'$.

THEOREM 4.2 (PRESERVATION). *If $\Gamma \vdash \langle s, R, M \rangle \dashv \Gamma'$ and $\langle s, R, M \rangle \mapsto \langle s', R', M' \rangle$, then $\Gamma' \vdash \langle s', R', M' \rangle \dashv \Gamma''$ for some context Γ'' .*

PROOF. By induction over the forward operational semantics of Core Tower. \square

A valid machine state $\langle s, R, M \rangle$ fails to make forward progress only if it must execute statement $x \rightarrow e$ where e does not compute the value of x , i.e. $R = R' [x \leftarrow v]$ and $\langle e, R' [x \leftarrow v] \rangle \uparrow_x \langle e, R' [x \leftarrow z(\tau)] \rangle$ fails to hold. We denote this condition $\langle s, R, M \rangle \text{err}$ and define it in Appendix A.

THEOREM 4.3 (PROGRESS). *Given a machine state $\langle s, R, M \rangle$ that is valid under some initial and final contexts, either s is `skip` or $\langle s, R, M \rangle \text{err}$ or $\langle s, R, M \rangle \mapsto \langle s', R', M' \rangle$ for some s', R', M' .*

PROOF. By induction over the definition of valid machine states. \square

THEOREM 4.4 (REVERSE PRESERVATION). *If $\Gamma' \vdash \langle s, R', M' \rangle \dashv \Gamma''$ and $\langle s, R', M' \rangle \mapsto_R \langle s', R, M \rangle$, then $\Gamma \vdash \langle s', R, M \rangle \dashv \Gamma'$ for some context Γ .*

PROOF. By induction over the reverse operational semantics of Core Tower. \square

A valid machine state $\langle s, R', M' \rangle$ fails to make reverse progress only if it must execute statement $x \leftarrow e$ where e does not compute the value of x , defined analogously to the forward case. We denote this condition $\langle s, R', M' \rangle \text{err}_R$ and define it in Appendix A.

THEOREM 4.5 (REVERSE PROGRESS). *Given a state $\langle s, R', M' \rangle$ that is valid under some initial and final contexts, either s is `skip` or $\langle s, R', M' \rangle \text{err}_R$ or $\langle s, R', M' \rangle \mapsto_R \langle s', R, M \rangle$ for some s', R, M .*

PROOF. By induction over the definition of valid machine states. \square

4.5 Program Reversibility and Inversion

Let the notation \mapsto^* denote the reflexive transitive closure of \mapsto , and similarly \mapsto_R^* for \mapsto_R . We say that a program s *terminates* in the forward direction when, given a register file R and memory M , we have $\langle s, R, M \rangle \mapsto^* \langle \text{skip}, R', M' \rangle$ for some final R' and M' . We say that s terminates in the reverse direction when, given R' and M' , we have $\langle s, R', M' \rangle \mapsto_R^* \langle \text{skip}, R, M \rangle$ for some R and M .

The following theorem states that every Core Tower program is *reversible*, meaning that it terminates in the forward direction if and only if it does so in the reverse direction.

THEOREM 4.6 (REVERSIBILITY). $\langle s, R, M \rangle \mapsto^* \langle \text{skip}, R', M' \rangle$ iff $\langle s, R', M' \rangle \mapsto_R^* \langle \text{skip}, R, M \rangle$.

PROOF. Follows from forward and reverse progress and preservation. \square

The type system does not statically enforce that a program terminates in the forward direction, and it is the responsibility of the program to correctly uncompute temporary values.

Inversion. Given a program s , its *inversion* $I[s]$ is a program whose semantics are reversed:

$$\begin{aligned} I[\text{skip}] &= \text{skip} & I[x \leftarrow e] &= x \rightarrow e & I[x_1 \Leftrightarrow x_2] &= x_1 \Leftrightarrow x_2 \\ I[s_1; s_2] &= I[s_2]; I[s_1] & I[x \rightarrow e] &= x \leftarrow e & I[*x_1 \Leftrightarrow x_2] &= *x_1 \Leftrightarrow x_2 \\ & & & & I[\text{if } x \text{ then } s \text{ end}] &= \text{if } x \text{ then } I[s] \text{ end} \end{aligned}$$

The following theorem states that the forward semantics of a program is equivalent to the reverse semantics of its syntactic inversion.

THEOREM 4.7 (INVERTIBILITY). $\langle s, R, M \rangle \mapsto^* \langle \text{skip}, R', M' \rangle$ iff $\langle I[s], R', M' \rangle \mapsto_R^* \langle \text{skip}, R, M \rangle$.

PROOF. By induction over the definition of inversion. \square

4.6 Quantum Circuit Semantics

We now develop a semantics of Core Tower that translates a program s to a unitary quantum circuit in order to execute s on a quantum computer.

Expression Semantics. Given an expression e whose forward operational semantics is defined as $\langle e, R[x \leftarrow z(\tau)] \rangle \Downarrow_x \langle e, R[x \leftarrow v] \rangle$, we lift it to a unitary gate U_e that operates on register file R and a new register x and stores v in this register, that is, $U_e |R, z(\tau)\rangle = |R, v\rangle$. Each possible expression – production of constant value, variable reference, projection from a pair, integer arithmetic, and Boolean logic – can be implemented as a unitary gate manipulating x :

- A constant value v can be reversibly produced by applying the NOT gate on each bit of x whose corresponding bit is set in the bit representation of v .
- A variable x' may be reversibly copied into x by applying the controlled-NOT gate on each bit in x conditioned on the corresponding bit of x' .
- Projection is analogous to copying a variable, except only one field of the pair is copied.
- Integer arithmetic may be reversibly performed through circuit designs such as [Cheng and Tseng \[2002\]](#); [Draper \[2000\]](#); [Islam et al. \[2009\]](#).
- Boolean logic may be embedded reversibly into NOT and Toffoli (controlled-controlled-NOT) gates, as established by [Fredkin and Toffoli \[1982\]](#).

Statement Semantics. In Figure 11, we translate a Core Tower statement s to a circuit fragment $C \llbracket s \rrbracket$ operating on the register file R and memory M encoded as quantum states. In the circuits, a wire depicted as ---^k denotes a k -bit register representing an individual program value. A wire depicted as --- denotes a collection of such values, such as R and M . A circuit fragment may be expanded to operate on the entire program state by tensor product with the identity gate.

The skip statement translates to the identity circuit that maps a quantum state containing R and M to itself. Statements s_1 and s_2 are sequenced by concatenating their corresponding circuits $C \llbracket s_1 \rrbracket$ and $C \llbracket s_2 \rrbracket$. An expression e is assigned to x by initializing a zero register, executing U_e on $|R\rangle$ and this register, and naming the resulting register x . An expression is un-assigned by executing U_e^\dagger , the inverse of U_e , producing a register of zeroes that is discarded from the circuit. Variables x_1 and x_2 are swapped using swap gates across their entire bit register representations. A swap with memory is performed by invoking the quantum random access gate (Section 2) on the address x_1 , value x_2 , and memory M . For an if-statement, its body $C \llbracket s \rrbracket$ is executed, conditioned on the bit x .

Soundness. The circuit semantics defines a unitary operator $C \llbracket s \rrbracket : |R, M\rangle \mapsto |R', M'\rangle$ for each statement s . It faithfully represents the operational semantics of the statement in Core Tower:

THEOREM 4.8 (SEMANTICS EQUIVALENCE). $C \llbracket s \rrbracket |R, M\rangle = |R', M'\rangle$ iff $\langle s, R, M \rangle \mapsto^* \langle \text{skip}, R', M' \rangle$

PROOF. By induction on the structure of s . \square




Fig. 11. Definition of quantum circuit semantics of Core Tower.

5 BOSON: HISTORY-INDEPENDENT QUANTUM MEMORY ALLOCATION

In this section, we present *Boson*, a memory allocator for quantum programs. The task of the allocator is to manage dynamic memory – when invoked by the program, it returns a pointer to an available allocation site and updates internal metadata to mark that site as being in use.

Boson achieves the goals of reversibility, history independence, and bounded-time execution. Its allocation is reversible, with the reverse being deallocation. It guarantees that allocated data structures are history-independent – a data structure with a unique logical linking structure has a unique physical memory representation. Finally, it performs allocation in constant time.

To develop *Boson*, we present in Section 5.1 a core algorithm for allocation based on a *free list* data structure that is reversible and constant-time. Though this algorithm without modification is not history-independent, as we demonstrate in Section 5.2, we then present *symmetrization*, *Boson*'s modification to this algorithm that enables history-independent allocation, in Section 5.3.

5.1 Core Allocation Algorithm

To illustrate the core allocation algorithm, we present an example that stores several values in dynamically allocated memory. We assume that the memory consists of a heap of m words of size k , that integers and pointers are word-sized, and that `null` and zero have identical bit representations.

Free List. *Boson*'s allocation algorithm relies on a *free list* data structure that tracks the free blocks of memory available to the program and is stored in the unallocated portion of the heap. A pointer to the head of the free list is maintained in a designated *allocation register* denoted `alloc_reg`.

In Figure 12, we depict the initial state of the allocation register and heap. Before the program executes, *Boson* constructs this initial state by initializing `alloc_reg` to `null` and the blocks of the heap to a sequence $[\text{ptr}[1], \text{ptr}[2], \text{ptr}[3], \dots, \text{ptr}[m-1], \text{null}]$.⁸ Then, it swaps `alloc_reg`




Fig. 12. Program state of allocation register and heap before program execution. Blocks in gray indicate memory not allocated to the program.

```

1 let p1 <- null;
2 p1 <-> alloc_reg;
3 *p1 <-> alloc_reg;
4 let n1 <- 1;
5 *p1 <-> n1;
6 let p2 <- null;
7 p2 <-> alloc_reg;
8 *p2 <-> alloc_reg;
9 let n2 <- 2;
10 *p2 <-> n2;
    
```

Fig. 13. Program storing two values in allocated memory. Highlighted operations of the Boson allocator obtain pointers to allocation sites from the free list.




Fig. 14. Program states while executing the program in Figure 13 starting from Figure 12. Each allocation yields a block for the program to use and updates the allocation register to the new head of the free list.




Fig. 15. Program states while executing the program in Figure 13 from an alternative initial free list. Compared to Figure 14, this execution begins with a free list that is the reverse of Figure 12. The final state in Figure 15c is semantically equivalent to but physically distinct from Figure 14c, violating history independence.




Fig. 16. Disruption of interference by distinct representations of the empty heap. The program begins from a superposition over $[1, 2]$ and $[2, 1]$ (normalizing amplitudes not shown). It invokes `remove` to remove 1 and 2 in sequence. In both terms of the superposition, the heap is then empty. However, the free lists have distinct structures. Consequently, interference fails to occur and the program outputs may be corrupted.

with the first block of the heap. The effect is that at the start of program execution, the whole heap constitutes the free list, and a pointer to the first free block is stored in `alloc_reg`.

⁸The value of the last block is defined to be `null` so that further allocations return `null` once memory has been exhausted.




Fig. 17. Result of symmetrization on the initial program state from Figure 12 (normalizing amplitudes not shown). The symmetrized free list exists in a superposition of all possible permutations.




Fig. 18. Unique physical representation of state $1 \leftrightarrow [1, 2]$ (normalizing amplitudes not shown), which stores data in a superposition of all possible allocation sites and is history-dependent.

Allocation. In Figure 13, we present a program that stores the values 1 and 2 in memory. Lines 1 to 3 invoke the Boson allocator to obtain a pointer $p1$ to space allocated for the first value:

- (1) `let p1 <- null` initializes $p1$ to `null`, and then
- (2) `p1 <-> alloc_reg` stores in $p1$ the address of the first available free block and stores zero in `alloc_reg`, and then
- (3) `*p1 <-> alloc_reg` stores into `alloc_reg` the value at memory address $p1$, which is the address of the next available block, and stores zero into memory at address $p1$.

In Figure 14a, we depict the program state after executing line 3. The allocation register now points to the new head of the free list, and $p1$ points to the newly allocated, zero-initialized word.

On lines 6 to 8, the program invokes the allocator again to obtain space to store the second value. In Figure 14b, we depict the state after executing line 8, in which the block to which $p1$ points has been updated to store 1, and $p2$ points to another newly allocated and zero-initialized word. Finally, in Figure 14c, we depict the state after line 10, in which $p2$ now points to the value 2.

By virtue of using the free list, allocation takes constant time. Boson deallocates memory by reversing the allocation operation, which has the effect of putting the newly freed block at the head of the free list. To support variable-sized allocation, Boson partitions the heap into sections of blocks whose sizes are different powers of two, and allocates from the appropriate free list.

5.2 Core Algorithm Is Not History-Independent

Without modification, the core algorithm does not ensure that data structure representations are history-independent, because the addresses of individual allocations are dependent on the contents of the free list, which vary as the program executes.

In Figure 15, we depict how a free list structured differently from Figure 12, which may possibly be found in the middle of executing a program, would cause the execution of Figure 13 to produce a final state distinct from Figure 14c. Quantum interference (Section 3.3) would fail between Figure 14c and Figure 15c even though these program states are semantically equivalent.

Interference after Deallocation. As a more extreme and unintuitive example, even after all dynamic memory is deallocated, the representation of the empty heap may cause quantum interference to fail over the remaining register file, leading program outputs to be corrupted.⁹

⁹An alternative and equivalent mathematical description of the failure is that quantum entanglement [Nielsen and Chuang 2010] exists between the memory and the register file, meaning that discarding and measuring the memory would cause the quantum state of the register file to catastrophically collapse from superposition.

In Figure 16, we present this example. The initial program state is a superposition:

$$\frac{1}{\sqrt{2}} \left(|1 \leftrightarrow [1, 2]\rangle - |1 \leftrightarrow [2, 1]\rangle \right)$$

The program invokes the operation `remove`, which removes and deallocates an element from `l`, on the values 1 and 2 in sequence. The result is that all memory is returned to the free list. Because `l` now always points to an empty list, we would expect quantum interference to occur:

$$\frac{1}{\sqrt{2}} \left(|1 \leftrightarrow []\rangle - |1 \leftrightarrow []\rangle \right) = 0$$

However, the physical representations of `l`, which necessarily incorporate `alloc_reg` and the free list stored in memory, differ across the two terms. Thus, quantum interference fails.

As demonstrated by these examples, the core allocation algorithm must be modified to prevent the program from producing an incorrect output after memory allocations and deallocations.

5.3 Symmetrization Achieves History Independence

We now develop *symmetrization*, the operation that Boson performs to modify the core allocation algorithm to guarantee history independence. A key idea is that classically, a *statistical* form of history independence is achieved by making probabilistic choices in the data structure:

Definition 5.1 (Statistical History Independence). An implementation of an abstract data structure is *statistically history-independent* if and only if, when any two sequences of abstract operations yield equivalent data structures, then the distributions of their physical representations are identical.

Randomization. Conceivably, the core allocation algorithm of Section 5.1 may be made statistically history-independent by choosing new allocations uniformly at random. Doing so would require initializing the free list to a uniformly random permutation of the sequence $[\text{ptr}[1], \dots, \text{ptr}[m-1]]$.¹⁰

Superposition. Statistical history independence guarantees a unique probability distribution for the physical representation of the heap. However, this property is insufficient for the support of quantum algorithms, since probability distributions do not exist on equal basis with superpositions in quantum mechanics – interference only occurs between equal superpositions, not probability distributions, of program states [Bell 1964]. Thus, instead of randomization, Boson uses superposition to achieve full, rather than statistical, history independence for the heap.

Symmetrization. A single time before the program executes, Boson initializes the free list to a uniform superposition of permutations by invoking a unitary operator that we call *symmetrization*:

$$|1, \dots, n\rangle \mapsto \frac{1}{\sqrt{n!}} \sum_{\pi \in \text{Sym}(n)} |\pi(1), \dots, \pi(n)\rangle$$

The notation $|1, \dots, n\rangle$ denotes a bit string of length kn constituting the concatenation of the k -bit binary representations of the integers 1 through n .

THEOREM 5.2. *If the heap is symmetrized before program execution, then after any sequence of allocation and deallocation operations, the representation of the heap is history-independent.*

PROOF. Follows from Theorem 9 in Naor and Teague [2001], replacing equality of probability distributions with equality of superpositions. \square

¹⁰Invoking randomness only at initialization means that an allocation immediately following a deallocation will yield the just-deallocated address. This correlation does not violate statistical history independence because the heap always remains in a uniformly random distribution over all possible choices of allocation sites.

In Figure 17, we illustrate the result of symmetrization on an empty heap state. Before symmetrization, the free list follows the particular sequence depicted in Figure 12. Afterward, the free list exists in the unique superposition of all possible permutations, preventing the issue from Section 5.2.

In Figure 18, we present the unique physical representation of the program state $1 \leftrightarrow [1, 2]$ under symmetrization, which stores data in a superposition of all possible allocation sites. This uniqueness satisfies history independence and means that quantum interference occurs correctly.

Implementation. Symmetrization is a central component of quantum simulation, in which it imposes bosonic¹¹ exchange symmetry [Liboff 1980] on the list. A multitude of implementations of this operator exist in the simulation literature, including Abrams and Lloyd [1997] and Barenco et al. [1997], who achieve implementations in $O(n^2 \log^2 n)$ gates; Chiew et al. [2019], in $O(n^2 \log n)$ gates; and Berry et al. [2017], in $O(n \log n \log \log n)$ gates and parallel time $O(\log n \log \log n)$.¹²

6 TOWER: RECURSIVE DATA STRUCTURE OPERATIONS

This section presents Tower, a language with co-recursive types and bounded recursive functions.

6.1 Syntax

The syntax of Tower augments Core Tower with new types, expressions and statements, as well as function declarations d , recursion bounds β , and programs p . Recursion bounds enable Tower to guarantee that all programs are bounded by classically known parameters.

$\tau ::= \dots \mid t \mid vt.\tau$	Declaration $d ::= \text{fun } f^{\beta?} (x_1 : \tau_1, \dots, x_k : \tau_k) \rightarrow \tau \text{ is } s$
$e ::= \dots \mid \text{alloc}(\tau) \mid f^{\beta?}(x_1, \dots, x_k)$	Bound $\beta ::= b \mid n \mid b - n' \quad (n \in \mathbb{N}, n' \in \mathbb{N} \setminus \{0\})$
$s ::= \dots \mid \text{return } x$	Program $p ::= d; p \mid \text{fun main } (x : \tau) \rightarrow \tau' \text{ is } s$

Types. The syntax of Tower introduces type variables denoted t as well as co-recursive types $vt.\tau$ in which the variable t may appear in τ .

Expressions. The expression $\text{alloc}(\tau)$ allocates zero-initialized space for a value of type τ and returns a pointer to that type. The expression $f^{\beta?}(x_1, \dots, x_k)$ invokes function f on arguments x_1, \dots, x_k with recursion bound β , or no bound if β is absent.

Statements. The statement $\text{return } x$ returns the value of x to the caller.

Declarations. The declaration $\text{fun } f^{\beta?} (x_1 : \tau_1, \dots, x_k : \tau_k) \rightarrow \tau \text{ is } s$ declares a function named f with recursion bound variable b (or no bound if absent), arguments x_1, \dots, x_k with types τ_1, \dots, τ_k respectively, return type τ , and body s . A function body must contain exactly one return statement, as the final statement. Tower permits recursive, but not mutually recursive, functions.¹³

Bounds. A recursion bound may be a variable b , a classical integer n , or their difference.

Programs. A program is a series of function declarations followed by the entry point main .

6.2 Type System

The type system defines well-formed types and enforces that all recursion is classically bounded.

¹¹Fermionic antisymmetry, which is closely related, additionally flips the sign according to the parity of the permutation π .

¹²Comparable to the complexity of the quantum Fourier transform [Coppersmith 1994].

¹³The presented type system of Tower could be modified straightforwardly to support mutual recursion.

$$\begin{array}{c}
\text{TE-CALLNONRECURSIVE} \\
\frac{f_1 \neq f_2 \quad \Phi[f_2] = (\text{None}, ((\tau_1, \dots, \tau_k), \tau)) \quad \Gamma \vdash x_1 : \tau_1 \quad \dots \quad \Gamma \vdash x_k : \tau_k}{\Phi; \Gamma \vdash_{f_1} f_2(x_1, \dots, x_k) : \tau} \\
\\
\text{TE-CALLBOUNDED} \\
\frac{\beta = n \text{ or } \beta = b \text{ or } b - n \text{ where } \text{fst}(\Phi[f_1]) = b \quad f_1 \neq f_2 \quad \Phi[f_2] = (\text{Some } b_2, ((\tau_1, \dots, \tau_k), \tau)) \quad \Gamma \vdash x_1 : \tau_1 \quad \dots \quad \Gamma \vdash x_k : \tau_k}{\Phi; \Gamma \vdash_{f_1} f_2^\beta(x_1, \dots, x_k) : \tau} \\
\\
\text{TE-CALLSELF} \\
\frac{\Phi[f] = (\text{Some } b, ((\tau_1, \dots, \tau_k), \tau)) \quad \Gamma \vdash x_1 : \tau_1 \quad \dots \quad \Gamma \vdash x_k : \tau_k}{\Phi; \Gamma \vdash_f f^{b-n}(x_1, \dots, x_k) : \tau}
\end{array}$$

Fig. 19. Selected typing rules for expressions in Tower. The full definition is presented in Appendix A.

Type Formation. In Appendix A, we define the well-formation of types. A context Δ is a set of type variables currently in scope. The judgment $\Delta \vdash \tau$ ok states that under context Δ , the type τ is well-formed. All rules are standard except the case for co-recursive types, which additionally requires that a type variable t may only appear under a pointer indirection.

For example, the type of lists of integers is $\text{list} \triangleq \text{vt}.\text{(uint, ptr}(t))$, in which we require that the self-reference t must occur under a pointer, similarly to structure definitions in C. We interpret co-recursive types equi-recursively. For example, we equate the list type $\text{list} \triangleq \text{vt}.\text{(uint, ptr}(t))$ with its unfolding $\text{(uint, ptr}(\text{list}))$. In Appendix A, we define the judgment $\tau_1 = \tau_2$, an equivalence relation that is compatible with unfolding co-recursive types. The default value of a co-recursive type, $z(\text{vt}.\tau)$, is defined to be the default value of its unfolding, $z(\tau[t \leftarrow \text{vt}.\tau])$.

Recursive Functions. The type system of Tower introduces a function context Φ that maps functions to their recursion bound variable (if present), argument list, and return type, and augments the typing judgments with the function context Φ and the function f being type-checked.

In Figure 19, we define the new typing judgment for expressions. The judgment $\Phi; \Gamma \vdash_f e : \tau$ states that under function context Φ and context Γ , inside function f , expression e has type τ . The type system assigns $\text{alloc}(\tau)$ the type $\text{ptr}(\tau)$. Most of the remaining rules simply add Φ and f , and we show only the new rules for function calls.

If the function being called, f_2 , is not recursive, it must be called with no bound and the correct arguments. If it is recursive, it must be called with a bound, which in turn may refer only to the bound variable of the caller function f_1 . In the last rule, the function f is making a recursive call to itself, which requires the bound to be smaller than the bound variable b .

In Appendix A, we define the new judgment for well-formation of statements. The judgment $\Phi; \Gamma \vdash_f s \dashv \Gamma'$ states that under Φ and Γ , inside function f , statement s is well-formed and yields new context Γ' . Most of the rules simply add Φ and f . The new rule for return requires that the context contains only the returned variable, and yields an empty context.

6.3 Translation into Core Tower

Tower's dynamic semantics is defined by translation into Core Tower. Each statement invoking alloc is replaced with the three Core Tower statements defined in Section 5.1. Each function call, starting from main , is translated by recursively inlining the function body as follows:

In the general case, a function call has a specified recursion bound. Given a function call $x \leftarrow f^n(x_1, \dots, x_k)$, the translator looks up the definition of f and renames its arguments and return value to be $\text{fun } f^b(x_1 : \tau_1, \dots, x_k : \tau_k) \rightarrow \tau$ is s ; return x . It also renames the local variables of f to avoid conflicts with the caller. It then inlines s into the caller, replacing instances of b with n . It recursively expands introduced function calls until the recursion bound reaches zero, and it

replaces expressions $f^0(x_1, \dots, x_k)$ with $z(\tau)$. A reverse function call $x \rightarrow f^n(x_1, \dots, x_k)$ is similar, except that the body of f is inverted (Section 4.5) before it is inlined.

Additional features of Tower, such as with-do blocks, nested expressions, and patterns, also lower into Core Tower and are described in Appendix B. In Appendix C, we provide a detailed example for the translation of a recursive Tower program into Core Tower by the procedure above.

THEOREM 6.1. *In a well-typed program, the process of inlining recursive calls always terminates.*

PROOF. Every recursive call has a bound strictly smaller than the bound of the function definition. A call with bound zero terminates immediately. By induction, all other recursive calls terminate. \square

As a result, a well-typed Tower program can be converted into a bounded-size quantum circuit.

6.4 Translation Soundness

The soundness and reversibility properties of Tower follow from those of Core Tower. A well-typed Tower program translates to a Core Tower program that is well-typed if *coercion* between compatible pointer types is additionally permitted.

Coercion is used by the implementation of $\text{alloc}(\tau)$ in Section 5.1 to assign a pointer to a free block into a variable of type $\text{ptr}(\tau)$. Since a free block is bitwise compatible with a value of type τ , such coercion is always safe with respect to the operational and circuit semantics of Core Tower.

THEOREM 6.2. *The inlining of recursive calls and alloc operators in a well-typed Tower program produces a Core Tower program that is well-typed if the coercion of compatible pointers is permitted.*

PROOF. The inlining of $\text{alloc}(\tau)$ is type-correct if coercions are permitted, as discussed above. The type-correctness of the inlining of recursive functions holds by induction on the bound. \square

7 GROUND: DATA STRUCTURE LIBRARY AND CASE STUDY

This section presents Ground, a library of abstract data structures we implemented in Tower, as well as a case study of development challenges during its implementation.

7.1 Language Implementation

To develop the library, we implemented a Tower interpreter in OCaml that classically executes the operational semantics. The type checker ensures that all recursion is bounded, and the interpreter verifies at runtime that 1) values are uncomputed correctly and 2) all heap memory is freed.

We also implemented a compiler for Tower into quantum circuits. Given a Tower program, the compiler generates for each function a reversible logic circuit that when executed on a bit-representation of its input stores its output into an auxiliary register. The compiler outputs assembly code that invokes primitive logical, arithmetic, and memory operations.

The compiler can further instantiate these primitive operations into multi-controlled-NOT gates. The current implementations of these operations prioritize simplicity; more optimized designs are possible but not yet implemented. The output is in a simple textual netlist format that is amenable to conversion for use in other quantum computation frameworks.

The source code of the Tower interpreter and compiler and the Ground library and tests is available in the artifact of this paper [Yuan and Carbin 2022] and on GitHub.¹⁴

7.2 Summary of Data Structure Library

Table 1 summarizes the Ground data structure library. We implemented the abstract data structures of lists, stacks, queues, strings, and sets.¹⁵ In the table, the name of an abstract data structure

¹⁴<https://github.com/psg-mit/tower-oopsla22>

```

1 fun length[n](l: ptr<list>) -> uint {
2   with {
3     let l_empty <- l == null;
4   } do if l_empty {
5     let out <- 0;
6   } else with {
7     let temp <- default<list>;
8     *l <-> temp;
9     let next <- temp.2;
10    let r <- length[n-1](next);
11  } do {
12    let out <- r + 1;
13  }
14  return out;
15 }

```

Fig. 20. $O(2^n)$ -time list length.

```

1 fun length[n](l: ptr<list>, acc: uint) -> uint {
2   with {
3     let l_empty <- l == null;
4   } do if l_empty {
5     let out <- acc;
6   } else with {
7     let temp <- default<list>;
8     *l <-> temp;
9     let next <- temp.2;
10    let r <- acc + 1;
11  } do {
12    let out <- length[n-1](next, r);
13  }
14  return out;
15 }

```

Fig. 21. $O(n)$ -time, avoiding recursive uncomputation.

is followed by its underlying implementation in parentheses. All implementations are history-independent except for hash table-based sets.

The time complexity for certain pointer-based data structures is a function of n , the number of elements in the data structure. For word-based data structures, the complexity of operations instead scales with the system word size k ,¹⁶ and becomes polynomial in k if the operation performs recursion bounded by k and performs operations over words of size k at each level.

We report the counts of qubits and gates used by the compiled circuit for each operation, including temporary qubits, after minimal optimization. To focus on the fundamental complexity of each operation, these counts exclude 1) the qubits constituting the qRAM, 2) qubits and gates internal to the implementation of primitive logical, arithmetic, and memory operations, and 3) gates to copy and swap qubits, all of which are highly dependent on parameter and optimization choices. For operations with variable complexity, these counts are functions of n or k . Word sizes $4 \leq k \leq 12$ are assumed for operations whose complexity depends on k , and $k = 8$ is assumed otherwise.

In Appendix D, we detail the characteristics of each data structure. During the development process, we encountered several challenges that do not arise in classical programming, and share them with future quantum developers. In the following sections, we describe three challenges: recursive uncomputation, mutated uncomputation, and branch sequentialization.

7.3 Challenge 1: Recursive Uncomputation

Recursive uncomputation arises when the result of a recursive call is uncomputed, and causes certain programs to incur exponential slowdown in asymptotic time complexity. We must be careful to appropriately craft the implementation to avoid naive implementations that cause this slowdown.

In Figure 20, we present a correct but inefficient way to compute the length of a linked list. Line 3 determines if the list l is empty. If so, line 5 sets the output to zero, and if not, lines 7–9 obtain a pointer $next$ to the tail of the list. Line 10 recursively computes the length of the tail. Line 12 adds one to the result of the recursion. After the `do`-block completes, Tower reverses the effects of both prior `with`-blocks to uncompute all temporary variables and restore the original state of the list.

Recursive Uncomputation. Though the program in Figure 20 is correct, its time complexity is $O(2^n)$. The reason is that the recursive call on line 10 produces a temporary value r that is uncomputed at

¹⁵Maps may be implemented as a straightforward augmentation of sets with values stored alongside each key.

¹⁶The word size k is in turn logarithmic in the maximum number of elements n of the data structure.

Table 1. Overview of Ground library. All operations are reversible. “Recursion” denotes whether the operation uses bounded recursion. “Mutation” indicates whether it leaves input in a mutated state. “Complexity” indicates its time complexity in terms of the number of elements n or the word size k . “LoC” counts the lines of code implementing the operation and its dedicated helpers, excluding calls to other operations listed here. “Qubits” counts qubits used, excluding those internal to the qRAM and primitive logic/arithmetic/memory operations. “Gates” counts primitive logic/arithmetic/memory gates invoked, excluding copy/swap gates.

Data Structure	Reversible	Recursion	Mutation	Complexity	LoC	Qubits	Gates
List							
– length	Yes	Yes	No	$O(n)$	20	$34n + 32$	$23n + 3$
– sum	Yes	Yes	No	$O(n)$	20	$34n + 40$	$21n + 3$
– find_pos	Yes	Yes	No	$O(n)$	20	$42n + 31$	$19n + 3$
– remove	Yes	Yes	Yes	$O(n)$	48	$26n + 56$	$42n + 3$
Stack (list)							
– push_front	Yes	No	Yes	$O(1)$	8	40	4
– pop_front	Yes	No	Yes	$O(1)$	8	48	4
Queue (list)							
– push_back	Yes	Yes	Yes	$O(n)$	21	$34n + 32$	$24n$
– pop_front	Yes	No	Yes	$O(1)$	8	48	4
String (word)							
– is_empty	Yes	No	No	$O(1)$	2	25	3
– length	Yes	No	No	$O(1)$	2	24	1
– get_prefix	Yes	No	No	$O(k)$	8	$11k$	52
– get_substring	Yes	No	No	$O(k)$	8	$12k$	54
– get	Yes	No	No	$O(k)$	7	$6k + 1$	19
– is_prefix	Yes	Yes	No	$O(\text{poly}(k))$	26	$k^2 + 11k$	$98k + 3$
– num_matching	Yes	Yes	No	$O(\text{poly}(k))$	42	$k^2 + 13k + 4$	$110k + 127$
– equal	Yes	No	No	$O(k)$	8	$6k + 3$	5
– concat	Yes	No	No	$O(k)$	9	$11k$	8
– compare	Yes	Yes	No	$O(\text{poly}(k))$	27	$5k^2 + 12k$	$108k + 3$
Set (radix tree)							
– insert	Yes	Yes	Yes	$O(\text{poly}(k))$	136	$13k^2 + 21k + 9$	$1440k^2 + 5056k$
– contains	Yes	Yes	No	$O(\text{poly}(k))$	334	$17k^2 + 18k + 2$	$784k^2 + 1612k + 1$
Set (hash table)*							
– insert	Yes	Yes	Yes	$O(n)$	63	$52n + 72$	$68n + 15$
– contains	Yes	Yes	No	$O(n)$	7	$52n + 81$	$136n + 39$

* Hash table-based sets are not history-independent.

the end of the do-block. As a result, length makes both a forward and a reverse recursive call at each level, resulting in exponential complexity.

Avoiding recursive uncomputation requires not uncomputing the result of recursive calls. One possible approach is inspired by tail recursion from classical programming, where the result of a recursive call is returned without any intervening operations.

In Figure 21, we avoid recursive uncomputation using an extra accumulator argument. Line 5 returns the accumulator in the base case. Line 10 adds one to the accumulator, and line 12 recursively calls the function with the new accumulator and directly assigns the result to the output.

Calling this function with an initial accumulator of zero returns the length. In the new implementation, no second recursive call is necessary, and the implementation achieves $O(n)$ complexity.

We encountered this challenge in the implementations of find_pos, length, remove, and sum for lists as well as is_prefix, num_matching, and compare for strings.

```

1 fun contains(t: tree, k: uint) -> bool {
2   if t == null { let out <- false; }
3   else if t.1 == k { let out <- true; }
4   else if t.1 < k {
5     let out <- contains(t.2, k);
6   } else {
7     let out <- contains(t.3, k);
8   }
9   return out;
10 }

```

Fig. 23. $O(n)$ -time binary tree traversal.

```

1 fun contains(t: tree, k: uint) -> bool {
2   if t == null { let out <- false; }
3   else if t.1 == k { let out <- true; }
4   else {
5     if t.1 < k { let child <- t.2; }
6     else { let child <- t.3; }
7     let out <- contains(child, k);
8   }
9   return out;
10 }

```

Fig. 24. $O(\log n)$ -time, avoiding branch sequentialization.

7.4 Challenge 2: Mutated Uncomputation

Mutated uncomputation refers to when uncomputing a temporary value is not straightforward because the program mutates data under conditional branches. The typical strategy of uncomputing a temporary using the expression that originally computed it fails when variables in that expression have been mutated by intervening statements.

In Figure 22, we implement the `push_back` operation, which adds `x` to the end of `l`. Line 2 computes the flag `l_empty`. If it is true, lines 4–7 allocate a new node storing `x` and set `l` to point to it. Otherwise, line 12 extracts the `tail` of the list, and line 15 recursively calls `push_back` on it.

Mutated Uncomputation. Line 18 uncomputes `l_empty` differently from its computation on line 2. The reason is that `l` may have been modified by line 5, which would break the relationship between `l` and `l_empty`. Instead, the program uncomputes `l_empty` based on whether the tail of the new list stored in node is `null` on line 18, which is true if and only if `l_empty` was true after line 2.

As shown, correctly uncomputing certain values requires reasoning about the invariants that hold on the program state after mutation. We encountered this challenge in `remove` for lists, `push_back` for queues, and `insert` for sets implemented using both hash tables and radix trees. By contrast, operations that do not mutate data may typically uncompute temporaries straightforwardly.

7.5 Challenge 3: Branch Sequentialization

Branch sequentialization denotes the fact that the time complexity of a conditionally branching quantum program may be larger than its classical equivalent and the developer may need to structure the program differently to avoid this slowdown.

In Figure 23, we present a pseudo-Tower program that traverses a binary search tree. Classically, if `t` is balanced, then the time complexity of the program is $O(\log n)$.

Branch Sequentialization. However, if Figure 23 is compiled to a quantum circuit (Section 4.6), each `if`-statement becomes a gate that conditions over the gate translation of the `if`-body. Each

```

1 fun push_back[n](l: ptr<list>, x: uint) {
2   let l_empty <- l == null;
3   if l_empty {
4     let head <- alloc<list>;
5     l <-> head;
6     let head -> null;
7     let node <- (x, null);
8   } else {
9     let node <- default<list>;
10    *l <-> node;
11    with {
12      let (h, tail) <- node;
13      let node -> (h, tail);
14    } do {
15      let () <- push_back[n-1](tail, x);
16    }
17  }
18  let l_empty -> node.2 == null;
19  *l <-> node;
20  let node -> default<list>;
21  return ();
22 }

```

Fig. 22. Implementation of `push_back`.

recursive level of the resulting circuit contains a gate that traverses the left child and a gate that traverses the right child, resulting in $O(n)$ total complexity. Put simply, the time complexity of a quantum program is the sum and not the maximum of the complexities of its possible branches.

Efficiently traversing a branching data structure requires making only one recursive call in the function. In Figure 24, we present a program that performs such a traversal in $O(\log n)$.

We encountered this challenge in radix tree `insert` and `contains`, which use a similar approach to attain poly-logarithmic complexity in the number of elements.

8 LIMITATIONS

Error Checking. Programming with memory causes numerous errors such as null pointer dereference, use after free, and out of memory. Detecting such errors is difficult in quantum programs, which cannot use exceptional control flow. Consequently, Tower defines a dereference of a null pointer to be a no-op and does not dynamically enforce memory safety. Developments in quantum runtime error checking such as Li et al. [2020] may enable detection of memory errors.

Memory Management. In Tower, the developer must manually reset and deallocate memory, which raises the possibility of automatic memory reclamation. Unfortunately, classical garbage collection traverses the heap, which has a high time cost. Also, erasing information cannot be performed reversibly without uncomputation, which is difficult after memory has become garbage.

Boson incurs internal and external heap fragmentation when allocating variable-sized blocks of memory. Practical dynamic memory management in quantum programs must reconcile space inefficiency of simple allocation schemes with the expense of traversing the heap.

9 RELATED WORK

Quantum Programming with Data Structures. The three requirements for data structures in superposition were introduced by Ambainis [2003] and elaborated by Bernstein et al. [2013] and Aaronson et al. [2019], who propose using radix trees [Morrison 1968] to implement sets. These authors also define quantum operations such as creating a superposition of all sets.

Other examples of data structures in superposition have been proposed by Booth et al. [2021]; Chen and Movassagh [2021]; Jeffery et al. [2013]; Shi [2021]. They may be distinguished from data structures *not* in superposition that are accessed by quantum algorithms for tasks such as simulation [Chakraborty et al. 2019] and machine learning [Kerenidis and Prakash 2016].

Péchox et al. [2020] develop a quantum programming language with inductive data types and recursion. However, in this language it is not possible to construct a superposition over differently linked structures. For example, it is not possible to construct a superposition over linked lists with different lengths, nor over trees with different shapes. Such a restriction is not present in Tower.

In fact, this restriction prevents the implementation of a set that satisfies history independence. To illustrate, suppose we implement sets concretely as sorted, duplicate-free linked lists, as in Section 3. Then, if x denotes a set whose state is a superposition of $[1, 2]$ and $[1, 3]$, invoking `insert(x, 2)` should yield a superposition of $[1, 2]$ and $[1, 2, 3]$. However, this superposition is over two lists with different lengths, which is not permitted by Péchox et al. [2020].

Reversible Programming. Janus [Lutz 1986; Yokoyama and Glück 2007] introduced reversible programming and conditions that are checked after an `if`-block completes. Consequent developments include Bowman et al. [2011]; Haulund [2017]; Thomsen and Axelsen [2015]. Other applications include bidirectional lenses [Bohannon et al. 2008] and reversible hardware [Vieri 1995].

Reversible Memory Management. Uncomputation [Bennett 1989] has become a central feature of quantum programming languages such as Abraham et al. [2019]; Bichsel et al. [2020]; Svore

et al. [2018]. The fact that recursive uncomputation may lead to asymptotic slowdown was noted by Aaronson [2002]. Proposed solutions include whole-program uncomputation [Paradis et al. 2021] and heuristics trading off between space and time [Ding et al. 2020].

Axelsen and Glück [2013] present a reversible free list memory allocator, which inspired parts of Boson but is not history-independent. Buhrman et al. [2022] propose a quantum allocator that traverses a prefix tree stored in qRAM for every allocation. By contrast, Boson symmetrizes the heap once and then performs allocation in constant time.

History Independence. Naor and Teague [2001] introduced history independence in the context of privacy and cryptography, in which it guarantees that an adversary with access to the physical memory of a system cannot recover sensitive information about the sequence of events that produced the current state. They present history-independent versions of hash tables and a general scheme for memory allocation that inspired parts of Boson. Later work [Bajaj et al. 2016; Hartline et al. 2002] elaborates the deterministic and statistical forms of history independence and provides additional schemes for history-independent data structure construction.

By contrast, oblivious data structures [Wang et al. 2014] guarantee that the memory access patterns, e.g. the sequence of dereferenced pointers, of semantically equivalent data structure operations are identical. Unlike history independence, obliviousness is not necessary for a data structure to operate in superposition. Nevertheless, oblivious techniques such as static and garbled circuits [Yao 1982; Zahur and Evans 2013] and oblivious RAM [Goldreich and Ostrovsky 1996] enable construction of fixed-control-flow data structures and may be useful in the quantum domain.

10 CONCLUSION

Data structures establish abstractions over mathematical objects, such as lists and sets, that are taken for granted in the design of algorithms and ubiquitous in classical programming. However, data structure implementations suitable under quantum superposition must satisfy additional properties, and existing quantum programming frameworks do not support their construction.

In this work, we introduce Tower, a language for quantum programming with data structures, and Boson, a memory allocator that operates in quantum superposition, along with Ground, a library of abstract data structures, including the first implementation of the set data structure satisfying the requirements laid out by quantum algorithms. We hope that the library will be useful for quantum developers and that Tower and Boson will lead to even more useful abstractions.

Unlocking the potential of quantum computation rests on our ability to leverage the accumulated wisdom of algorithm and software design while grappling with physical and practical constraints that are unfamiliar to classical developers. As we have shown, understanding classically-motivated principles of computing, such as history independence, allows us to build bridges to abstractions such as data structures that will enable scalable and accessible quantum software.

ACKNOWLEDGMENTS

We would like to thank Jesse Michel, Logan Weber, Tian Jin, and anonymous reviewers who provided feedback on drafts of this paper. We thank Scott Aaronson for directing us to the quantum algorithms that were critical to the motivation of this work. We also give special thanks to Chris McNally for invaluable discussions on the nature of quantum random-access memory, and introducing to the authors the physical concept of exchange symmetry.

This work was supported in part by the MIT-IBM Watson AI Lab and the Sloan Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- Scott Aaronson. 2002. Quantum Lower Bound for Recursive Fourier Sampling. <https://doi.org/10.48550/ARXIV.QUANT-PH/0209060> arXiv:0209060 [quant-ph]
- Scott Aaronson, Nai-Hui Chia, Han-Hsuan Lin, Chunhao Wang, and Ruizhe Zhang. 2019. On the Quantum Complexity of Closest Pair and Related Problems. arXiv:1911.01973 [quant-ph]
- Héctor Abraham et al. 2019. Qiskit: An Open-source Framework for Quantum Computing.
- Daniel S. Abrams and Seth Lloyd. 1997. Simulation of Many-Body Fermi Systems on a Universal Quantum Computer. *Phys. Rev. Lett.* 79 (Sep 1997). Issue 13. <https://doi.org/10.1103/PhysRevLett.79.2586>
- T. Altenkirch and J. Grattage. 2005. A Functional Quantum Programming Language. In *IEEE Symposium on Logic in Computer Science*. <https://doi.org/10.1109/LICS.2005.1>
- Andris Ambainis. 2003. Quantum walk algorithm for element distinctness. arXiv:0311001 [quant-ph]
- Srinivasan Arunachalam, Vlad Gheorghiu, Tomas Jochym-O'Connor, Michele Mosca, and Priyaa Varshinee Srinivasan. 2015. On the robustness of bucket brigade quantum RAM. *New Journal of Physics* 17, 12 (Dec 2015). <https://doi.org/10.1088/1367-2630/17/12/123010>
- Holger Bock Axelsen and Robert Glück. 2013. Reversible Representation and Manipulation of Constructor Terms in the Heap. In *Reversible Computation*, Gerhard W. Dueck and D. Michael Miller (Eds.).
- Sumeet Bajaj, Anrin Chakraborti, and Radu Sion. 2016. Practical Foundations of History Independence. *IEEE Transactions on Information Forensics and Security* 11, 2 (2016). <https://doi.org/10.1109/TIFS.2015.2491309>
- Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. 1997. Stabilization of Quantum Computations by Symmetrization. *SIAM J. Comput.* 26, 5 (1997). <https://doi.org/10.1137/S0097539796302452>
- J. S. Bell. 1964. On the Einstein Podolsky Rosen paradox. *Physics* 1 (Nov 1964). Issue 3. <https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195>
- C. H. Bennett. 1973. Logical Reversibility of Computation. *IBM Journal of Research and Development* 17, 6 (1973). <https://doi.org/10.1147/rd.176.0525>
- Charles H. Bennett. 1989. Time/Space Trade-Offs for Reversible Computation. *SIAM J. Comput.* 18, 4 (Aug 1989). <https://doi.org/10.1137/0218053>
- Charles H. Bennett and Gilles Brassard. 2014. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science* 560 (2014). <https://doi.org/10.1016/j.tcs.2014.05.025>
- Daniel J. Bernstein, Stacey Jeffery, Tanja Lange, and Alexander Meurer. 2013. Quantum Algorithms for the Subset-Sum Problem. In *Post-Quantum Cryptography*.
- Dominic W. Berry, Mária Kieferová, Artur Scherer, Yuval R. Sanders, Guang Hao Low, Nathan Wiebe, Craig Gidney, and Ryan Babbush. 2017. Improved techniques for preparing eigenstates of fermionic Hamiltonians. *npj Quantum Information* 4 (2017). <https://doi.org/10.1038/s41534-018-0071-5>
- Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin Vechev. 2020. Silq: A High-Level Quantum Language with Safe Uncomputation and Intuitive Semantics. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. <https://doi.org/10.1145/3385412.3386007>
- Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. 2008. Boomerang: Resourceful Lenses for String Data. In *ACM SIGPLAN Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/1328438.1328487>
- Kyle E. C. Booth, Bryan O’Gorman, Jeffrey Marshall, Stuart Hadfield, and Eleanor Rieffel. 2021. Quantum-accelerated constraint programming. *Quantum* (2021). <https://doi.org/10.22331/q-2021-09-28-550>
- William Bowman, Roshan James, and Amr Sabry. 2011. Dagger Traced Symmetric Monoidal Categories and Reversible Programming. In *Workshop on Reversible Computation*.
- Harry Buhman, Bruno Loff, Subhasree Patro, and Florian Speelman. 2021. Limits of quantum speed-ups for computational geometry and other problems: Fine-grained complexity via quantum walks. arXiv:2106.02005 [quant-ph]
- Harry Buhman, Bruno Loff, Subhasree Patro, and Florian Speelman. 2022. Memory Compression with Quantum Random-Access Gates. <https://doi.org/10.48550/ARXIV.2203.05599> arXiv:2203.05599 [quant-ph]
- Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. 2019. The Power of Block-Encoded Matrix Powers: Improved Regression Techniques via Faster Hamiltonian Simulation. <https://doi.org/10.4230/LIPICS.ICALP.2019.33> arXiv:1804.01973 [quant-ph]
- Lijie Chen and Ramis Movassagh. 2021. Quantum Merkle Trees. <https://doi.org/10.48550/ARXIV.2112.14317> arXiv:2112.14317 [quant-ph]
- K-W Cheng and C-C Tseng. 2002. Quantum full adder and subtractor. *Electron. Lett.* 38 (2002). <https://doi.org/10.1049/el:20020949>
- Mitchell Chiew, Kooper de Lacy, Chao-Hua Yu, Sam Marsh, and Jingbo Wang. 2019. Graph comparison via nonlinear quantum search. *Quantum Information Processing* 18 (08 2019). <https://doi.org/10.1007/s11128-019-2407-2>

- Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. 2018. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences* 115, 38 (2018). <https://doi.org/10.1073/pnas.1801723115>
- D. Coppersmith. 1994. An approximate Fourier transform useful in quantum factoring.
- Yongshan Ding, Xin-Chuan Wu, Adam Holmes, Ash Wiseth, Diana Franklin, Margaret Martonosi, and Frederic T. Chong. 2020. SQUARE: Strategic Quantum Ancilla Reuse for Modular Quantum Programs via Cost-Effective Uncomputation. In *International Symposium on Computer Architecture*. <https://doi.org/10.1109/ISCA45697.2020.00054>
- Thomas G. Draper. 2000. Addition on a Quantum Computer. arXiv:quant-ph/0008033 [quant-ph]
- Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. <https://doi.org/10.48550/ARXIV.1411.4028> arXiv:1411.4028 [quant-ph]
- Edward Fredkin and Tommaso Toffoli. 1982. Conservative logic. *International Journal of Theoretical Physics* 21 (1982). <https://doi.org/10.1007/BF01857727>
- Craig Gidney. 2022. Quantum Dictionaries without QRAM. <https://doi.org/10.48550/ARXIV.2204.13835> arXiv:2204.13835 [quant-ph]
- Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. 2008. Quantum Random Access Memory. *Physical Review Letters* 100, 16 (Apr 2008). <https://doi.org/10.1103/physrevlett.100.160501>
- Oded Goldreich and Rafail Ostrovsky. 1996. Software Protection and Simulation on Oblivious RAMs. *J. ACM* 43, 3 (May 1996). <https://doi.org/10.1145/233551.233553>
- Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. 2013. Quipper: A Scalable Quantum Programming Language. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. <https://doi.org/10.1145/2491956.2462177>
- Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *ACM Symposium on Theory of Computing*. <https://doi.org/10.1145/237814.237866>
- Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. 2009. Quantum Algorithm for Linear Systems of Equations. *Physical Review Letters* 103, 15 (Oct 2009). <https://doi.org/10.1103/physrevlett.103.150502>
- Jason D. Hartline, Edwin S. Hong, Alexander E. Mohr, William R. Pentney, and Emily C. Rocke. 2002. Characterizing History Independent Data Structures. In *Algorithms and Computation*, Prosenjit Bose and Pat Morin (Eds.).
- Tue Haulund. 2017. Design and Implementation of a Reversible Object-Oriented Programming Language. <https://doi.org/10.48550/ARXIV.1707.07845> arXiv:1707.07845 [cs]
- Md Saiful Islam, Muhammad Mahbubur Rahman, Zerina Begum, and Mohd Z Hafiz. 2009. Low cost quantum realization of reversible multiplier circuit. *Information technology journal* 8, 2 (2009). <https://doi.org/10.3923/itj.2009.208.213>
- Stacey Jeffery, Robin Kothari, and Frederic Magniez. 2013. Nested Quantum Walks with Quantum Data Structures. In *ACM-SIAM Symposium on Discrete Algorithms*. <https://doi.org/10.1137/1.9781611973105.106>
- Ivan Kassal, James D. Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Alán Aspuru-Guzik. 2011. Simulating Chemistry Using Quantum Computers. *Annual Review of Physical Chemistry* 62, 1 (2011). <https://doi.org/10.1146/annurev-physchem-032210-103512>
- Iordanis Kerenidis and Anupam Prakash. 2016. Quantum Recommendation Systems. <https://doi.org/10.48550/ARXIV.1603.08675> arXiv:1603.08675 [quant-ph]
- E Knill. 1996. Conventions for quantum pseudocode.
- Gushu Li, Li Zhou, Nengkun Yu, Yufei Ding, Mingsheng Ying, and Yuan Xie. 2020. Projection-Based Runtime Assertions for Testing and Debugging Quantum Programs. In *ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*. <https://doi.org/10.1145/3428218>
- Richard L. Liboff. 1980. *Introductory quantum mechanics*. Addison-Wesley.
- Christopher Lutz. 1986. Janus: a time-reversible language. (1986). *Letter to R. Landauer*.
- Olivia Di Matteo, Vlad Gheorghiu, and Michele Mosca. 2020. Fault-Tolerant Resource Estimation of Quantum Random-Access Memories. *IEEE Transactions on Quantum Engineering* 1 (2020). <https://doi.org/10.1109/tqe.2020.2965803>
- Donald R. Morrison. 1968. PATRICIA—Practical Algorithm To Retrieve Information Coded in Alphanumeric. *J. ACM* 15, 4 (Oct 1968). <https://doi.org/10.1145/321479.321481>
- Moni Naor and Vanessa Teague. 2001. Anti-persistence: History independent data structures. In *ACM Symposium on Theory of Computing*. <https://doi.org/10.1145/380752.380844>
- Michael A. Nielsen and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press.
- Alexandru Paler, Oumarou Oumarou, and Robert Basmadjian. 2020. Parallelizing the queries in a bucket-brigade quantum random access memory. *Physical Review A* 102, 3 (Sep 2020). <https://doi.org/10.1103/physreva.102.032608>
- Anouk Paradis, Benjamin Bichsel, Samuel Steffen, and Martin Vechev. 2021. Unqomp: Synthesizing Uncomputation in Quantum Circuits. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. <https://doi.org/10.1145/3453483.3454040>

- Jennifer Paykin, Robert Rand, and Steve Zdancewic. 2017. QWIRE: A Core Language for Quantum Circuits. In *ACM SIGPLAN Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/3009837.3009894>
- Romain Péchoux, Simon Perdrix, Mathys Rennela, and Vladimir Zamdzhiev. 2020. Quantum Programming with Inductive Datatypes: Causality and Affine Type Theory. In *Lecture Notes in Computer Science*. Springer International Publishing. https://doi.org/10.1007/978-3-030-45231-5_29
- Peter Selinger and Benoît Valiron. 2005. A Lambda Calculus for Quantum Computation with Classical Control. *Typed Lambda Calculi and Applications* (2005). <https://doi.org/10.1017/S0960129506005238>
- Run-Hua Shi. 2021. Quantum Bloom Filter and Its Applications. *IEEE Transactions on Quantum Engineering* (2021). <https://doi.org/10.1109/TQE.2021.3054623>
- Peter W. Shor. 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 5 (Oct 1997). <https://doi.org/10.1137/S0097539795293172>
- Krysta Svore, Martin Roetteler, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, and Andres Paz. 2018. Q#: Enabling Scalable Quantum Computing and Development with a High-level DSL. In *Real World Domain Specific Languages Workshop*. <https://doi.org/10.1145/3183895.3183901>
- Michael Kirkedal Thomsen and Holger Bock Axelsen. 2015. Interpretation and Programming of the Reversible Functional Language RFUN. In *Symposium on the Implementation and Application of Functional Programming Languages*. <https://doi.org/10.1145/2897336.2897345>
- Carlin Vieri. 1995. Pendulum—a reversible computer architecture.
- Xiao Shaun Wang, Kartik Nayak, Chang Liu, T-H. Hubert Chan, Elaine Shi, Emil Stefanov, and Yan Huang. 2014. Oblivious Data Structures. In *ACM SIGSAC Conference on Computer and Communications Security*. <https://doi.org/10.1145/2660267.2660314>
- Dave Wecker, Krysta M. Svore, and Krysta M. Svore. 2014. LIQUi>: A Software Design Architecture and Domain-Specific Language for Quantum Computing. (February 2014).
- Andrew C. Yao. 1982. Protocols for Secure Computations. In *Symposium on Foundations of Computer Science*. <https://doi.org/10.5555/1382436.1382751>
- Tetsuo Yokoyama and Robert Glück. 2007. A Reversible Programming Language and Its Invertible Self-Interpreter. In *ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*. <https://doi.org/10.1145/1244381.1244404>
- Charles Yuan and Michael Carbin. 2022. Tower: Data Structures in Quantum Superposition. <https://doi.org/10.5281/zenodo.6819031>
- Samee Zahur and David Evans. 2013. Circuit Structures for Improving Efficiency of Security and Privacy Tools. In *IEEE Symposium on Security and Privacy*. <https://doi.org/10.1109/SP.2013.40>