


# QDB: From Quantum Algorithms Towards Correct Quantum Programs

Yipeng Huang<sup>1</sup>

Department of Computer Science, Princeton University

yipeng@cs.princeton.edu

 <https://orcid.org/0000-0003-3171-6901>

Margaret Martonosi

Department of Computer Science, Princeton University

mrm@princeton.edu

---

## Abstract

With the advent of small-scale prototype quantum computers, researchers can now code and run quantum algorithms that were previously proposed but not fully implemented. In support of this growing interest in quantum computing experimentation, programmers need new tools and techniques to write and debug QC code. In this work, we implement a range of QC algorithms and programs in order to discover what types of bugs occur and what defenses against those bugs are possible in QC programs. We conduct our study by running small-sized QC programs in QC simulators in order to replicate published results in QC implementations. Where possible, we cross-validate results from programs written in different QC languages for the same problems and inputs. Drawing on this experience, we provide a taxonomy for QC bugs, and we propose QC language features that would aid in writing correct code.

**2012 ACM Subject Classification** Computer systems organization → Quantum computing

**Keywords and phrases** Correctness, debugging

**Digital Object Identifier** 10.4230/OASICS.PLATEAU.2018.

## 1 Introduction

Quantum computing is reaching an inflection point. After years of work on both QC algorithms and low-level QC devices, small but viable QC prototypes are now available to run programs. These QC prototypes are increasing in size, with much research attention being placed on improving their reliability and increasing the counts of qubits (quantum bits), the fundamental building block for QC [11, 19, 31].

With small-scale machines available to run real code, a natural challenge lies in creating correct and useful programs to run on them [3, 12]. Until recently, QC algorithms were rarely programmed for actual execution, and therefore relatively little QC debugging has ever occurred. Furthermore, QC debugging faces challenges beyond that of classical computing. In particular, typical debugging approaches based on printing out variable values during program execution do not easily apply to QC programs, because program states in QC “collapse” to classical values when observed. Second, QC’s “no cloning rule” precludes us from making a spare copy of variables to observe them elsewhere. Third, while we have more freedom to observe states in QC simulations on classical computers, the massive state spaces of QC executions limits this approach to small programs. Finally, even when limited simulations are tractable, it can be difficult to interpret the simulation results.

---

<sup>1</sup> This work is funded in part by EPiQC, an NSF Expedition in Computing, under grant 1730082



© Yipeng Huang and Margaret Martonosi;  
licensed under Creative Commons License CC-BY

The 9th Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU 2018).

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

41 This paper surveys a range of QC algorithms and programs and offers a set of empirical  
 42 and experiential insights on today’s state-of-the-art in QC debugging. For three benchmarks  
 43 representing different application areas, we perform detailed debugging based on small-scale  
 44 simulations. For each, we give case studies of the types of bugs we found. Most importantly,  
 45 we use these experiences to assemble a set of “design patterns for QC programming” and  
 46 related best practices in QC debugging.

47 In particular, the contributions of this paper are as follows:

- 48 ■ We specifically explore three major areas: quantum chemistry, integer factorization, and  
 49 database search. This is a broad spectrum of QC algorithms across not just application  
 50 domains, but also problem size and algorithm strategies. This allows us to point out  
 51 particular domain-specific challenges or opportunities.
- 52 ■ Where available, we study the same algorithm implemented in different languages or  
 53 infrastructures. From this, we draw comparative insights regarding how programming  
 54 language or environment support can be useful in QC programming and debugging.
- 55 ■ From these insights and experiences, we lay out a plan for debugging support in QC pro-  
 56 gramming environments to aid users in creating quantum code. These include assertions,  
 57 unit testing, code reuse, polymorphism, and QC-specific language types and syntax.

58 Overall, while QC programming has received significant prior attention and QC debugging  
 59 has received some as well, our work offers steps forward in its detailed and comparative  
 60 assessment across problem types and languages. We see our work offering useful insights for  
 61 QC programmers themselves, as well as language and system designers interested in building  
 62 next-generation compilers and debuggers.

## 63 **2 Background on QC programming**


64 First, we review the principles of quantum computing [14, 22, 23, 26], in order to understand  
 65 how writing correct quantum programs is different from classical programming.

### 66 **2.1 Qubits, superpositions, and entanglement**

67 The basic unit of information in QC is the qubit, which can take on values of  $|0\rangle$  and  $|1\rangle$  like  
 68 bits in classical computing, but can also be viewed as a probabilistic “superposition” between  
 69 the two values. Quantum computers can also “measure” the value of a qubit, forcing it to  
 70 collapse out of superposition into a classical value such as ‘0’ or ‘1’. Measurement disturbs  
 71 the values of variables in a quantum computer, so we cannot easily pause execution and  
 72 observe the values of qubits as a quantum program runs.

73 The state of individual qubits can be “entangled” together. For this reason, as more  
 74 qubits come into play in a quantum computer, the number of states that data can be in grows  
 75 exponentially. For example, a two-qubit system can take on the values  $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ ,  
 76 along with superpositions among these values; furthermore, the two qubits can even be in a  
 77 state of entanglement where the two cannot be treated as independent pieces of information.  
 78 A three qubit system has potential superpositions of eight states, and so on. This exponential  
 79 growth of possible values underlies the power of QC.

80 As a result of this large number of possible states, running a quantum program in  
 81 simulation on a classical computer is costly. Naive simulation of a 20-qubit quantum  
 82 computer, for example, needs  $2^{20}$  or roughly one million floating point numbers just to store  
 83 the program state at any instant. For this reason, testing and debugging quantum programs  
 84 in simulation is only possible for toy-sized programs.



■ **Figure 1** Decomposition of a simple QC program. Time flows left to right, showing sequences of operations applied to qubits  $q_0$  and  $q_1$ . The left program is a “controlled” arbitrary operation  $U$ , which means whether the operation  $U$  works on  $q_1$  is dependent on the value of  $q_0$ . The left sequence decomposes into the equivalent right sequence of more basic operations. The basic operations include single-qubit “rotations”  $A$  through  $D$  that alter the probability distribution of qubit values. The operations also include two two-qubit “CNOT” operations that flip a qubit (denoted  $\oplus$ ) contingent on the value of another qubit (denoted  $\bullet$ ) [26].

## 85 2.2 Quantum computer operations, programs, and a taxonomy for bugs

86 The process of quantum computing involves applying operations on qubits. We use diagrams  
87 such as Figure 1 to represent sequences of quantum operations. Looking at Figure 1 we see  
88 that quantum programs consist of three conceptual parts [8]:

- 89 1. **Inputs** to quantum algorithms include *classical* input parameters such as coefficients for  
90 rotations  $A$  through  $D$ , and *quantum* initial values for qubits such as  $q_0$  and  $q_1$ .
- 91 2. **Operations**, such as the specification of how a complex operation such as controlled  
92 arbitrary operation  $U$  (Figure 1, left) decomposes into basic operations  $A$  through  $D$  and  
93 CNOTs (Figure 1, right). Additionally, both basic and complex operations can be further  
94 composed according to patterns such as iteration, recursion, and mirroring.
- 95 3. **Outputs** of quantum algorithms are the final classical measurement values of qubits such  
96 as  $q_0$  and  $q_1$ . Furthermore, any temporary variables used in the course of a program have  
97 to be safely disentangled from the rest of the quantum state and discarded.

98 Bugs in quantum programs can crop up due to mistakes made in any of these three parts  
99 of a QC program. We will give examples of each kind of bug along with how to prevent them,  
100 using detailed case studies in the rest of this paper.

## 101 2.3 QC algorithm primitives, benchmarks, and open source frameworks

102 Given the rapid growth of QC infrastructure, we now have a chance to test a variety of  
103 quantum algorithms written in many languages [18]. Many different quantum algorithms rely  
104 on a handful of QC algorithm primitives to get speedups relative to classical algorithms [4, 24,  
105 25]. Table 1 classifies canonical quantum algorithms according to their algorithm primitives,  
106 and cites example implementations in different QC languages and tool chains.

107 This paper specifically focuses on program bugs and defenses in three areas: a quantum  
108 chemistry problem that uses quantum phase estimation, integer factorization using Shor’s  
109 order finding algorithm, and Grover’s database search algorithm.

110 Using programs written in the Scaffold language as a starting point [13], we compile  
111 Scaffold code to OpenQASM, a QC assembly language [5]. Then, we simulate the programs  
112 operation-by-operation in the QX simulator [15], in order to see their intermediate states and  
113 outputs. We cross reference the programs’ results against implementations in other languages,  
114 such as LIQWi> [32], ProjectQ [10, 36] and Q# [37]. From this debugging experience we  
115 identify possible bugs and defenses. Furthermore, we review the codes across languages to  
116 understand the relative merits of different QC language features.

■ **Table 1** Quantum algorithm primitives and open source benchmarks in open source tool chains.

Primitives	Quantum algorithms	Benchmark implementations
Entanglement protocols	superdense coding / quantum teleportation	Q# teleportation [37] pyQuil teleportation [35]
Quantum (random) walks	tree traversal	Scaffold / Quipper binary welded tree [6, 13, 39]
	graph traversal	Scaffold / Quipper triangle finding problem [6, 13, 39]
	satisfiability	Scaffold / Quipper Boolean formula [6, 13, 39]
Adiabatic	Ising spin model	Scaffold / Q# adiabatic Ising model [13, 37]
	quantum approximate optimization algorithm	QISKit Aqua QAOA pyQuil QAOA ansatz [35]
Variational Quantum Eigensolver	Hamiltonian simulation	QISKit Aqua quantum chemistry Q# H <sub>2</sub> simulation [37] Rigetti Grove VQE [35]
Quantum Fourier Transform (QFT)	phase estimation	Scaffold / Quipper ground state estimation [6, 13, 39]
	period finding	Scaffold class number [13]
	order finding	Scaffold / ProjectQ / Q# Shor's factoring [13, 36, 37]
	hidden subgroup problem	Quipper unique shortest vector [6, 39]
	linear algebra	Quipper quantum linear systems [6, 39]
Amplitude amplification	database search	Scaffold square root [13] ProjectQ / Q# Grover's database search [36, 37]

### 117 3 Case study: Quantum chemistry

118 First, we discuss our experience building up and debugging a simple quantum chemistry  
 119 program. Quantum chemistry problems entail finding properties of molecules from theoretical  
 120 first principles [20, 27]. Researchers anticipate these will be the first applications for QC due  
 121 to the relatively few number of qubits they need to surpass classical computer algorithms.  
 122 Debugging these problems is distinctively challenging, due to the importance of getting a  
 123 large number of classical input parameters all correct, and because of the dearth of physically  
 124 meaningful intermediate states we can check in the course of algorithm execution.

#### 125 3.1 Bug type 1: Incorrect classical input parameters

126 A key part of quantum chemistry programs is in correctly building up a “Hamiltonian”  
 127 subroutine that simulates inter-electron forces. The procedure for doing this was laid out in  
 128 detail by Whitfield [41]. We followed this procedure to create a subroutine for simulating  
 129 the hydrogen molecule, but we needed additional validation from several other sources to  
 130 get a bug-free subroutine [40]. These resources include raw chemistry data found in open  
 131 source repositories for the LIQUi> framework<sup>2</sup>. The final parameters for actual operations  
 132 on qubits were validated against a follow-up paper [33] and an implementation in the  
 133 QISKit framework<sup>3</sup>. Because the procedure for preparing these quantum chemistry models  
 134 involves many steps and needs domain expertise, software packages such as OpenFermion  
 135 now automate this process [21]. Nonetheless, there is room for improvement in standardizing  
 136 input data formats to eliminate bugs in this process.

137 Once the Hamiltonian subroutine is built, we can use the model in a variety of quantum  
 138 algorithms spanning different primitives in Table 1. These include phase estimation (an

<sup>2</sup> [https://github.com/StationQ/Liquid/blob/master/Samples/h2\\_sto3g\\_4.dat](https://github.com/StationQ/Liquid/blob/master/Samples/h2_sto3g_4.dat)

<sup>3</sup> <https://github.com/Qiskit/aqua/blob/master/test/H2-0.735.json>

■ **Table 2** QC calculated energy for  $H_2$  (bond length = 73.48 pm) for different electron assignments.

	Electron assignments				QC calculated energy (relative)
	Bonding		Antibonding		
	↑	↓	↑	↓	
3 <sup>rd</sup> excited state (E3)	0	0	1	1	-0.164
2 <sup>nd</sup> excited state (E2)	0	1	1	0	-0.217
	1	0	0	1	
1 <sup>st</sup> excited state (E1)	0	1	0	1	-0.244
	1	0	1	0	
Ground state (G)	1	1	0	0	-0.295

139 application of quantum Fourier transforms) [28], variational quantum eigensolvers [30], and  
 140 adiabatic algorithms [1]. In this paper, we use iterative phase estimation to find the ground  
 141 state energy of our  $H_2$  model, validating results published by Lanyon [17].

## 142 3.2 Bug type 2: Incorrect quantum initial values

143 The correct preparation of qubit initial values is important. Incorrect initial values would  
 144 cause the program to find solutions to different problems altogether. In this quantum  
 145 chemistry problem, the initial values control the locations of the two electrons in  $H_2$ . As  
 146 shown in Table 2, we need the qubit assignment for finding the ground energy of  $H_2$ , while  
 147 other assignments lead to results for other energy levels.

148 The symmetry of  $H_2$  allows us to perform a sanity check, to make sure the Hamiltonian  
 149 and the iterative phase estimation subroutines are working correctly. Though there are six  
 150 ways to assign two electrons to four locations, there are in fact only four distinct energy  
 151 levels, as shown in the experimental data. Checking that the two different ways to obtain E1  
 152 (and E2) give the same energy levels validates that the model correctly preserves symmetry.

## 153 3.3 Defense type 1: Assertions on algorithm preconditions

154 Given how important correct initial values are for all quantum algorithms, it is worthwhile  
 155 to explicitly check for these algorithm preconditions before continuing with execution or  
 156 simulation. What the preconditions should be depends on the type of algorithm. For example,  
 157 the phase estimation subroutine in this case study (along with other algorithms relying on  
 158 quantum Fourier transforms), expect inputs that are maximally in superposition among all  
 159 possible values. Likewise, “ancillary qubits” such as the inputs to the Hamiltonian subroutine  
 160 take on completely classical (integer) initial values. Lastly, quantum protocols often need to  
 161 start with entangled states. These required input states are among the few places in quantum  
 162 algorithms where we can check states for specific values. We can check these preconditions  
 163 by running or simulating programs up to the entry point of subroutines, and performing a  
 164 premature measurement to check for these anticipated states, finally restarting the program  
 165 knowing that execution is correct up to that point. Thus far, the Q# framework has the  
 166 most extensive support for precondition checking [37].

## 167 3.4 Defense type 2: Assertions on algorithm progress

168 Unlike the other two case studies later in this paper, the debugging process for the quantum  
 169 chemistry benchmark is coarse-grained. That is because the Hamiltonian subroutine is a  
 170 monolithic block of code whose components do not have obvious expected outputs—its

■ **Table 3** Shor’s factorization algorithm subroutines [23, p. 25].

Program subroutine code	Shared library code
Shor’s routine for factoring 15; calculating powers of a number ■ controlled modular multiplication ■ controlled modular addition ■ controlled addition	■ quantum Fourier transform ■ controlled controlled rotation ■ controlled rotation ■ controlled swap ■ swap

■ **Table 4** Correct and incorrect code for rotation decomposition. Using the Scaffold language [13] as an example, we code out the controlled operation U in Figure 1 where U is a rotation in just one axis. Because only one axis is needed, we can drop either operation A or C, paying attention to the sign on the angles. Reordering the lines of code or signs results in a rotation in the wrong direction.

Correct, operation A unneeded	Correct, operation C unneeded	Incorrect, angles flipped
<code>Rz(q1,+angle/2); // C</code> <code>CNOT(q0,q1);</code> <code>Rz(q1,-angle/2); // B</code> <code>CNOT(q0,q1);</code> <code>Rz(q0,+angle/2); // D</code>	<code>CNOT(q0,q1);</code> <code>Rz(q1,-angle/2); // B</code> <code>CNOT(q0,q1);</code> <code>Rz(q1,+angle/2); // A</code> <code>Rz(q0,+angle/2); // D</code>	<code>Rz(q1,-angle/2);</code> <code>CNOT(q0,q1);</code> <code>Rz(q1,+angle/2);</code> <code>CNOT(q0,q1);</code> <code>Rz(q0,+angle/2); // D</code>

171 components represent pair-wise electron interactions, and do not have inherent physical  
 172 meaning. So how do we debug this program? The preconditions in the last section make sure  
 173 the inputs to the algorithm are correct; the other observable state we have for debugging is  
 174 to check the behavior of the algorithm as a whole.

175 In this quantum chemistry program, we can check for two types of overall algorithm  
 176 behavior. One is the solution should converge to a steady value as finer Trotter time steps (a  
 177 kind of numerical approximation) are chosen; a lack of this type of convergence indicates  
 178 a bug in the Hamiltonian subroutine. The other algorithm behavior is when we vary the  
 179 precision of the phase estimation algorithm, the most significant bits of the measurement  
 180 output sequences should be the same—in other words, rounding the output of a high-precision  
 181 experiment should yield the same output as a lower-precision experiment. a lack of this  
 182 convergence indicates a bug in the iterative phase estimation subroutine. These checks for  
 183 expected algorithm progress also apply to other algorithms.

## 184 4 Case study: Shor’s algorithm for integer factorization

185 While our debugging strategy for quantum chemistry had to be coarse-grained, the debugging  
 186 process for Shor’s algorithm in this section allows us to look inside the program one subroutine  
 187 at a time, where we can compare the intermediate results against known expected values.

188 Shor’s factorization algorithm uses a quantum computer to factor a composite number  
 189 in polynomial time complexity, providing exponential speedup relative to the best known  
 190 classical algorithms [34]. We follow an example for an implementation that minimizes the  
 191 qubit cost [2], and replicate results for factoring 15, the simplest example [16] [26, p. 235].

### 192 4.1 Bug type 3: Incorrect operations and transformations

193 In order to correctly implement Shor’s algorithm we first have to build up the quantum  
 194 subroutines shown in Table 3. These basic subroutines can be tricky to get right. Take the

■ **Listing 1** Controlled adder subroutine using Fourier transform in the Scaffold language [13].

```

1 // outputs a + b, where a is a 'width' bit constant integer
2 // b is an integer encoded on 'width' qubits in Fourier space
3 module cADD (
4     const unsigned int c_width, // number of control qubits
5     qbit ctrl0, qbit ctrl1, // control qubits
6     const unsigned int width, const unsigned int a, qbit b[]
7 ) {
8     for (int b_indx=width-1; b_indx>=0; b_indx--) {
9         for (int a_indx=b_indx; a_indx>=0; a_indx--) {
10            if ((a >> a_indx) & 1) { // shift out bits in constant a
11                double angle = M_PI/pow(2,b_indx-a_indx); // rotation angle
12                switch (c_width) {
13                    case 0: Rz ( b[b_indx], angle ); break;
14                    case 1: cRz ( ctrl0, b[b_indx], angle ); break;
15                    case 2: ccRz ( ctrl0, ctrl1, b[b_indx], angle ); break;
16                }
17            }
18        }
19    }
20 }

```

195 controlled rotation in Figure 1 as an example: Table 4 shows multiple ways to code the  
 196 decomposition of the controlled rotation, and small mistakes can lead to incorrect behavior.

## 197 4.2 Defense type 3: Language support for subroutines / unit tests

198 An obvious defense against coding mistakes in basic subroutines is to use a library of shared  
 199 code. Doing so helps ensure program correctness by allowing programmers to exhaustively  
 200 validate small subroutines, in order to bootstrap larger subroutines. Unit testing is especially  
 201 important in QC as running or simulating large quantum programs is impossible for now.

202 An additional benefit is logically structured code allows compilers to select the best con-  
 203 crete implementation for the abstract functionality the programmer needs, based on hardware  
 204 constraints and input parameters [8]. For example, the most cost-efficient implementation for  
 205 modular exponentiation in Shor's factorization algorithm depends on how many qubits are  
 206 available: the compiler can choose from minimum-qubit [2, 9, 38] or minimum-operation [29]  
 207 implementations for the arithmetic subroutines.

## 208 4.3 Bug type 4: Incorrect composition of operations using iteration

209 Once we have built our basic subroutines, a common pattern in quantum programs is to use  
 210 iterations to compose subroutines. Listing 1 shows the iteration code for a constant-value  
 211 adder, showing tricky places in lines 8 through 11 for bugs to crop up, including indexing  
 212 errors, bit shifting errors, endian confusion, and mistakes in rotation angles. In general this  
 213 type of iteration code is commonplace in programs that rely on quantum Fourier transforms.

## 214 4.4 Defense type 4: Language support for numerical data types

215 One way to defend against bugs in iteration code is to introduce QC data types for numbers,  
 216 providing greater abstraction than working with raw qubits. For example, ProjectQ has  
 217 quantum integer data types [36], while Q# [37] and Quipper [6, 39] offer both big endian  
 218 and little endian versions of subroutines involving iterations. These QC data types permit  
 219 useful operators (e.g., checking for equality) that help with debugging and writing assertions.

■ **Table 5** Correct classical input  $a$  and  $a^{-1}$  to Shor’s algorithm for factoring 15, using 7 as a guess.

$k$ , the algorithm iteration	0	1	2	3	...
$a = 7^{2^k} \pmod{15}$	7	4	1	1	...
$a^{-1}; a \times a^{-1} \equiv 1 \pmod{15}$	13	4	1	1	...

■ **Table 6** Probability of measuring values of outputs and ancillary qubits of Shor’s algorithm, with incorrect inputs ( $a^{-1} = 12$  instead of 13 on first iteration). If the ancillary qubits collapse to zero on measurement, the algorithm still succeeds, returning correct outputs of 0, 2, 4, 6 [26, p. 235]. However, the possibility of measuring non-zero for the ancillary qubits indicates a bug.

Probability	Output measurement								
	0	1	2	3	4	5	6	7	
Ancillary qubit measurement	0	1/8	0	1/8	0	1/8	0	1/8	0
	4	1/64	1/64	1/64	1/64	1/64	1/64	1/64	1/64
	7	1/64	1/64	1/64	1/64	1/64	1/64	1/64	1/64
	8	1/64	1/64	1/64	1/64	1/64	1/64	1/64	1/64
	13	1/64	1/64	1/64	1/64	1/64	1/64	1/64	1/64

220 **4.5 Bug type 5: Incorrect deallocation of qubits**

221 Variable scoping is an important language feature in classical computing that ensures proper  
 222 data encapsulation. In QC, scoping is similarly important for temporary variables known  
 223 as “ancillary qubits.” Anything that happens to a subroutine’s ancillary qubits—such as  
 224 measurement, reinitialization, or lapsing into decoherence—may have unintended effects on  
 225 the subroutine’s outputs<sup>4</sup>. Because improper ancillary qubit deallocation can lead to wrong  
 226 results, it is important for subroutines to reverse their operations on their ancillary qubits,  
 227 so that they properly undo any entanglement between the ancillary and output qubits.

228 We can demonstrate a bug involving incorrect qubit deallocation, by deliberately making  
 229 a mistake while reversing operations in a subroutine. For example, Shor’s algorithm relies on  
 230 correct pairs modular inverse numbers as input parameters, such as those in Table 5. By  
 231 feeding an incorrect pair of inputs (e.g., replacing 13 with a 12), the algorithm proceeds to  
 232 possibly give us wrong output values, as shown in Table 6. At the same time, the mistake  
 233 prevents the modular multiplication operation from being properly reversed, which has the  
 234 effect of preventing the ancillary qubits from properly disentangling with other qubits, so  
 235 they fail to return to their initial values at the end of the algorithm.

236 **4.6 Defense type 5: Assertions on algorithm postconditions**

237 We can use postconditions at the end of algorithms to detect bugs that lead to incorrect  
 238 deallocation of ancillary qubits. Continuing with our example in Table 6, we see that the  
 239 cases where ancillary qubits collapse to anything other than zero correspond to cases where  
 240 the outputs are wrong. That is because the ancillary qubits remain improperly entangled  
 241 with the output qubits at the end of the algorithm. We can detect these buggy outputs by  
 242 asserting that ancillary qubits should always return to their initial values. The significance  
 243 of these observations is that when algorithms work correctly, we typically do not care to  
 244 measure the value of ancillary qubits as they do not contain information. But in buggy QC  
 245 algorithm implementations, they are useful side channels for debugging.

<sup>4</sup> As an analogy in classical computing, it is as if accessing an out-of-scope variable can still affect program state; while such behavior is unintuitive, it is a result of how entanglement works in QC.



■ **Table 7** Grover’s amplitude amplification subroutine in two languages, showcasing QC-specific language syntax for reversible computation (rows 2 & 6) and controlled operations (rows 3 & 5).

	Scaffold (C syntax) [13]	ProjectQ (Python syntax) [36]
1	<code>int j; qbit ancilla[n-1]; // scratch register for(j=0; j&lt;n-1; j++) PrepZ(ancilla[j],0);</code>	<code># reflection across # uniform superposition</code>
2	<code>// Hadamard on q for(j=0; j&lt;n; j++) H(q[j]); // Phase flip on q = 0...0 so invert q for(j=0; j&lt;n; j++) X(q[j]);</code>	<code>with Compute(eng):   All(H)   q   All(X)   q</code>
3	<code>// Compute x[n-2] = q[0] and ... and q[n-1] CCNOT(q[1], q[0], ancilla[0]); for(j=1; j&lt;n-1; j++)   CCNOT(ancilla[j-1], q[j+1], ancilla[j]);</code>	<code>with Control(eng, q[0:-1]):</code>
4	<code>// Phase flip Z if q=00...0 cZ(ancilla[n-2], q[n-1]);</code>	<code>Z   q[-1]</code>
5	<code>// Undo the local registers for(j=n-2; j&gt;0; j-)   CCNOT(ancilla[j-1], q[j+1], ancilla[j]); CCNOT(q[1], q[0], ancilla[0]);</code>	<code># ProjectQ automatically # uncomputes control</code>
6	<code>// Restore q for(j=0; j&lt;n; j++) X(q[j]); for(j=0; j&lt;n; j++) H(q[j]);</code>	<code>Uncompute(eng)</code>

## 246 5 Case study: Grover’s algorithm for database search

247 So far, we have presented defenses against bugs following two general strategies. One is  
248 to use assertions to detect when and where the program has a bug. The other is to use  
249 quantum-specific programming language features to prevent bugs altogether: these features  
250 include support for subroutines and numerical types for quantum data. Here in this section,  
251 we use the Grover’s benchmark to showcase two more language features for common QC  
252 program patterns: reversible computation and controlled operations.

253 Grover’s search algorithm finds an entry that matches search criteria, among an input  
254 data set of size  $N$ , with a time cost on the order of  $\sqrt{N}$ . That represents a polynomial  
255 speedup relative to the linear time cost in a classical computer [7].

256 The Grover’s algorithm comprises three parts. First, the input qubits representing the  
257 indices of the matching entries are put in a state of superposition, akin to querying all entries  
258 at once. Second, the queries are put through a subroutine that checks for the search criteria.  
259 In this case study, our criteria is to find the square root of a number in a Galois field of  
260 two elements, a simple abstract algebra setting. Finally in the critical step, the amplitude  
261 amplification algorithm primitive amplifies the index that matches the criteria while damping  
262 out those that do not. The operations in this final step are prime examples of two QC  
263 program patterns, reversible computation and controlled operations. We show in Table 7  
264 how these code patterns are written in two languages, Scaffold [13] and ProjectQ [36].

### 265 5.1 Bug type 6: Incorrect composition of operations using mirroring

266 Section 4.5 discussed how bugs in deallocating ancillary qubits can happen due to bad  
267 parameters. Here we see how bugs in deallocating ancillary qubits can happen due to  
268 incorrect composition of operations following a mirroring pattern. For example, in Table 7,

■ **Table 8** Applicability of defense strategies (down) against location of QC program bugs (across).

		input		operations				output
		classical params. §3.1	qubit alloc. §3.2	basic §4.1	iterate §4.3	mirror §5.1	recurse §5.3	qubit dealloc. §4.5
QC specific lang. features	unit testing §4.2		✓	✓	✓	✓	✓	✓
	data types §4.4				✓			
	reverse comp. §5.2				✓	✓		✓
	controlled ops. §5.4				✓	✓	✓	
Assertion checks	preconditions §3.3		✓					
	algo progress §3.4	✓	✓	✓	✓	✓	✓	✓
	postconds. §4.6	✓	✓	✓	✓	✓	✓	✓

269 the operations in rows 2 and 3 are respectively mirrored and undone in rows 6 and 5. These  
270 lines of code need careful reversal of every loop and every operation.

## 271 5.2 Defense type 6: Language support for reversible computation

272 Syntax support for reversible computation, such as that in ProjectQ [36], automatically  
273 mirrors and inverts sequences of operations, shortening code and reducing mistakes.

## 274 5.3 Bug type 7: Incorrect composition of operations using recursion

275 A common pattern in quantum programs involves performing operations (e.g., add), contingent  
276 on a set of qubits known as control qubits. Without language support, this pattern needs  
277 many lines of code and manual allocation of ancillary qubits. In the Scaffold code example  
278 in Table 7, rows 3 and 5 are just computing the intersection of qubits  $q$ , with the help of  
279 ancillary qubits initialized in row 1, in order to realize the controlled rotation operation in  
280 row 4. Furthermore, quantum algorithms often need varying numbers of control qubits in  
281 different parts of the algorithm, leading to replicated code from multiple versions of the same  
282 subroutine differing only by the number of control qubits<sup>5</sup>.

## 283 5.4 Defense type 7: Language support for controlled operations

284 Language support for controlled operations (e.g. ProjectQ) shortens code, preventing mistakes.

## 285 6 Conclusion

286 For the first time, we have access to comprehensive and representative program benchmarks  
287 for all major areas of quantum algorithms, implemented in multiple languages, along with  
288 input datasets and outputs that are detailed enough to permit cross-validation. Using  
289 our experience running and debugging these programs, we presented in this paper defense  
290 strategies that facilitate writing bug-free QC code, summarized in Table 8. Successful  
291 transplantation of these ideas from classical languages to QC languages can pave the way  
292 towards correct and useful quantum programs.

<sup>5</sup> An example appeared in the Shor's case study Listing 1. The addition operation was contingent on control qubits taken as parameters in lines 4 and 5. Depending on how many control qubits were needed, the switch statement in lines 12 through 15 applied the correct operation.

293 — **References** —

- 294 1 R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. Las Heras, R. Babbush,  
295 A. G. Fowler, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey,  
296 E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana,  
297 P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, E. Solano, H. Neven, and  
298 John M. Martinis. Digitized adiabatic quantum computing with a superconducting circuit.  
299 *Nature*, 534:222 EP –, 06 2016. URL: <http://dx.doi.org/10.1038/nature17658>.
- 300 2 Stephane Beauregard. Circuit for Shor's algorithm using  $2n+3$  qubits. *Quantum Info. Com-*  
301 *put.*, 3(2):175–185, March 2003. URL: [http://dl.acm.org/citation.cfm?id=2011517.](http://dl.acm.org/citation.cfm?id=2011517.2011525)  
302 2011525.
- 303 3 Frederic T. Chong, Diana Franklin, and Margaret Martonosi. Programming languages and  
304 compiler design for realistic quantum hardware. *Nature*, 549:180 EP –, 09 2017. URL:  
305 <http://dx.doi.org/10.1038/nature23459>.
- 306 4 Patrick J. Coles, Stephan Eidenbenz, Scott Pakin, Adetokunbo Adedoyin, John Am-  
307 brosiano, Petr M. Anisimov, William Casper, Gopinath Chennupati, Carleton Coffrin,  
308 Hristo Djidjev, David Gunter, Satish Karra, Nathan Lemons, Shizeng Lin, Andrey Y.  
309 Lokhov, Alexander Malyzhenkov, David Mascarenas, Susan M. Mniszewski, Balu Nadiga,  
310 Dan O'Malley, Diane Oyen, Lakshman Prasad, Randy Roberts, Philip Romero, Nan-  
311 dakishore Santhi, Nikolai Sinitsyn, Pieter Swart, Marc Vuffray, Jim Wendelberger, Bo-  
312 ram Yoon, Richard J. Zamora, and Wei Zhu. Quantum algorithm implementations for  
313 beginners. *CoRR*, abs/1804.03719, 2018. URL: <http://arxiv.org/abs/1804.03719>,  
314 arXiv:1804.03719.
- 315 5 A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta. Open Quantum Assembly  
316 Language. *ArXiv e-prints*, July 2017. arXiv:1707.03429.
- 317 6 Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît  
318 Valiron. Quipper: A scalable quantum programming language. In *Proceedings of the 34th*  
319 *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*  
320 *'13*, pages 333–342, New York, NY, USA, 2013. ACM. URL: [http://doi.acm.org/10.](http://doi.acm.org/10.1145/2491956.2462177)  
321 [1145/2491956.2462177](http://doi.acm.org/10.1145/2491956.2462177), doi:10.1145/2491956.2462177.
- 322 7 Lov K Grover. From Schrödinger's equation to the quantum search algorithm. *Pramana*,  
323 56(2-3):333–348, 2001.
- 324 8 T. Häner, T. Hoefler, and M. Troyer. Using Hoare logic for quantum circuit optimization.  
325 *ArXiv e-prints*, September 2018. arXiv:1810.00375.
- 326 9 Thomas Häner, Martin Roetteler, and Krysta M. Svore. Factoring using  $2n + 2$  qubits  
327 with toffoli based modular multiplication. *Quantum Info. Comput.*, 17(7-8):673–684, June  
328 2017. URL: <http://dl.acm.org/citation.cfm?id=3179553.3179560>.
- 329 10 Thomas Häner, Damian S. Steiger, Mikhail Smelyanskiy, and Matthias Troyer. High per-  
330 formance emulation of quantum circuits. In *Proceedings of the International Conference for*  
331 *High Performance Computing, Networking, Storage and Analysis, SC '16*, pages 74:1–74:9,  
332 Piscataway, NJ, USA, 2016. IEEE Press. URL: [http://dl.acm.org/citation.cfm?id=](http://dl.acm.org/citation.cfm?id=3014904.3015003)  
333 [3014904.3015003](http://dl.acm.org/citation.cfm?id=3014904.3015003).
- 334 11 Aram Harrow. Why now is the right time to study quantum computing. *XRDS*, 18(3):32–  
335 37, March 2012. URL: <http://doi.acm.org/10.1145/2090276.2090288>, doi:10.1145/  
336 [2090276.2090288](http://doi.acm.org/10.1145/2090276.2090288).
- 337 12 Thomas Häner, Damian S Steiger, Krysta Svore, and Matthias Troyer. A software method-  
338 ology for compiling quantum programs. *Quantum Science and Technology*, 3(2):020501,  
339 2018. URL: <http://stacks.iop.org/2058-9565/3/i=2/a=020501>.
- 340 13 Ali JavadiAbhari, Shruti Patil, Daniel Kudrow, Jeff Heckey, Alexey Lvov, Frederic T.  
341 Chong, and Margaret Martonosi. ScaffCC: A framework for compilation and analy-  
342 sis of quantum computing programs. In *Proceedings of the 11th ACM Conference on*

- 343 *Computing Frontiers*, CF '14, pages 1:1–1:10, New York, NY, USA, 2014. ACM. URL:  
344 <http://doi.acm.org/10.1145/2597917.2597939>, doi:10.1145/2597917.2597939.
- 345 14 Phillip Kaye, Raymond Laflamme, and Michele Mosca. *An Introduction to Quantum Com-*  
346 *puting*. Oxford University Press, Inc., New York, NY, USA, 2007.
- 347 15 N. Khammassi, I. Ashraf, X. Fu, C. G. Almudever, and K. Bertels. QX: A high-performance  
348 quantum computer simulation platform. In *Proceedings of the Conference on Design, Au-*  
349 *tomation & Test in Europe*, DATE '17, pages 464–469, 3001 Leuven, Belgium, Belgium,  
350 2017. European Design and Automation Association. URL: [http://dl.acm.org/citation.](http://dl.acm.org/citation.cfm?id=3130379.3130487)  
351 [cfm?id=3130379.3130487](http://dl.acm.org/citation.cfm?id=3130379.3130487).
- 352 16 B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James,  
353 A. Gilchrist, and A. G. White. Experimental demonstration of a compiled version  
354 of Shor's algorithm with quantum entanglement. *Phys. Rev. Lett.*, 99:250505, Dec  
355 2007. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.99.250505>, doi:10.  
356 1103/PhysRevLett.99.250505.
- 357 17 B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D.  
358 Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, A. Aspuru-Guzik, and A. G. White.  
359 Towards quantum chemistry on a quantum computer. *Nature Chemistry*, 2:106 EP –, 01  
360 2010. URL: <http://dx.doi.org/10.1038/nchem.483>.
- 361 18 R. LaRose. Overview and Comparison of Gate Level Quantum Software Platforms. *ArXiv*  
362 *e-prints*, July 2018. arXiv:1807.02500.
- 363 19 Norbert M. Linke, Dmitri Maslov, Martin Roetteler, Shantanu Debnath, Caroline Fig-  
364 gatt, Kevin A. Landsman, Kenneth Wright, and Christopher Monroe. Experimen-  
365 tal comparison of two quantum computing architectures. *Proceedings of the National*  
366 *Academy of Sciences*, 114(13):3305–3310, 2017. URL: [http://www.pnas.org/content/](http://www.pnas.org/content/114/13/3305)  
367 [114/13/3305](http://www.pnas.org/content/114/13/3305), arXiv:<http://www.pnas.org/content/114/13/3305.full.pdf>, doi:10.  
368 1073/pnas.1618020114.
- 369 20 S. McArdle, S. Endo, A. Aspuru-Guzik, S. Benjamin, and X. Yuan. Quantum computa-  
370 tional chemistry. *ArXiv e-prints*, August 2018. arXiv:1808.10402.
- 371 21 J. R. McClean, I. D. Kivlichan, K. J. Sung, D. S. Steiger, Y. Cao, C. Dai, E. Schuyler  
372 Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, C. Huang,  
373 J. Izaac, Z. Jiang, X. Liu, M. Neeley, T. O'Brien, I. Ozfidan, M. D. Radin, J. Romero,  
374 N. Rubin, N. P. D. Sawaya, K. Setia, S. Sim, M. Steudtner, Q. Sun, W. Sun, F. Zhang, and  
375 R. Babbush. OpenFermion: The Electronic Structure Package for Quantum Computers.  
376 *ArXiv e-prints*, October 2017. arXiv:1710.07629.
- 377 22 N.D. Mermin. *Quantum Computer Science: An Introduction*. Cambridge University Press,  
378 2007.
- 379 23 Tzvetan S. Metodi, Arvin I. Faruque, and Frederic T. Chong. Quantum computing  
380 for computer architects, second edition. *Synthesis Lectures on Computer Architecture*,  
381 6(1):1–203, 2011. URL: <https://doi.org/10.2200/S00331ED1V01Y201101CAC013>,  
382 arXiv:<https://doi.org/10.2200/S00331ED1V01Y201101CAC013>, doi:10.2200/  
383 S00331ED1V01Y201101CAC013.
- 384 24 Ashley Montanaro. Quantum algorithms: an overview. *npj Quantum Information*, 2:15023,  
385 2016.
- 386 25 Michele Mosca. Quantum algorithms. In *Encyclopedia of Complexity and Systems Science*,  
387 pages 7088–7118. Springer, 2009.
- 388 26 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information:*  
389 *10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition,  
390 2011.
- 391 27 Jonathan Olson, Yudong Cao, Jonathan Romero, Peter Johnson, Pierre-Luc Dallaire-  
392 Demers, Nicolas Sawaya, Prineha Narang, Ian Kivlichan, Michael Wasielewski, and Alán

- 393 Aspuru-Guzik. Quantum information and computation for chemistry. *arXiv preprint*  
394 *arXiv:1706.05413*, 2017.
- 395 **28** S. Patil, A. JavadiAbhari, C. Chiang, J. Heckey, M. Martonosi, and F. T. Chong. Char-  
396 acterizing the performance effect of trials and rotations in applications that use quantum  
397 phase estimation. In *2014 IEEE International Symposium on Workload Characterization*  
398 *(IISWC)*, pages 181–190, Oct 2014. doi:10.1109/IISWC.2014.6983057.
- 399 **29** Archimedes Pavlidis and Dimitris Gizopoulos. Fast quantum modular exponentiation archi-  
400 tecture for Shor’s factoring algorithm. *Quantum Info. Comput.*, 14:649–682, May 2014.  
401 URL: <http://dl.acm.org/citation.cfm?id=2638682.2638690>.
- 402 **30** Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J.  
403 Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on  
404 a photonic quantum processor. *Nature Communications*, 5:4213 EP –, 07 2014. URL:  
405 <http://dx.doi.org/10.1038/ncomms5213>.
- 406 **31** John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79,  
407 August 2018. URL: <https://doi.org/10.22331/q-2018-08-06-79>, doi:10.22331/  
408 q-2018-08-06-79.
- 409 **32** M. Roetteler, K. M. Svore, D. Wecker, and N. Wiebe. Design automation for quantum  
410 architectures. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*,  
411 pages 1312–1317, March 2017. doi:10.23919/DATE.2017.7927196.
- 412 **33** Jacob T. Seeley, Martin J. Richard, and Peter J. Love. The Bravyi-Kitaev transforma-  
413 tion for quantum computation of electronic structure. *The Journal of Chemical Physics*,  
414 137(22):224109, 2012. URL: <https://doi.org/10.1063/1.4768229>, arXiv:[https://doi.](https://doi.org/10.1063/1.4768229)  
415 [org/10.1063/1.4768229](https://doi.org/10.1063/1.4768229), doi:10.1063/1.4768229.
- 416 **34** Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms  
417 on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997. URL: <http://dx.doi.org/10.1137/S0097539795293172>, doi:10.1137/S0097539795293172.
- 418 **35** R. S. Smith, M. J. Curtis, and W. J. Zeng. A Practical Quantum Instruction Set Architec-  
419 ture. *ArXiv e-prints*, August 2016. arXiv:1608.03355.
- 420 **36** Damian S. Steiger, Thomas Häner, and Matthias Troyer. ProjectQ: an open source software  
421 framework for quantum computing. *Quantum*, 2:49, January 2018. URL: [https://doi.](https://doi.org/10.22331/q-2018-01-31-49)  
422 [org/10.22331/q-2018-01-31-49](https://doi.org/10.22331/q-2018-01-31-49), doi:10.22331/q-2018-01-31-49.
- 423 **37** Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bet-  
424 tina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler.  
425 Q#: Enabling scalable quantum computing and development with a high-level DSL. In  
426 *Proceedings of the Real World Domain Specific Languages Workshop 2018*, RWDSL2018,  
427 pages 7:1–7:10, New York, NY, USA, 2018. ACM. URL: [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/3183895.3183901)  
428 [3183895.3183901](http://doi.acm.org/10.1145/3183895.3183901), doi:10.1145/3183895.3183901.
- 429 **38** Yasuhiro Takahashi and Noboru Kunihiro. A quantum circuit for Shor’s factoring algorithm  
430 using  $2n + 2$  qubits. *Quantum Info. Comput.*, 6(2):184–192, March 2006. URL: <http://dl.acm.org/citation.cfm?id=2011665.2011669>.
- 431 **39** Benoît Valiron, Neil J. Ross, Peter Selinger, D. Scott Alexander, and Jonathan M. Smith.  
432 Programming the quantum future. *Commun. ACM*, 58(8):52–61, July 2015. URL: <http://doi.acm.org/10.1145/2699415>, doi:10.1145/2699415.
- 433 **40** Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer.  
434 Gate-count estimates for performing quantum chemistry on small quantum comput-  
435 ers. *Phys. Rev. A*, 90:022305, Aug 2014. URL: [https://link.aps.org/doi/10.1103/](https://link.aps.org/doi/10.1103/PhysRevA.90.022305)  
436 [PhysRevA.90.022305](https://link.aps.org/doi/10.1103/PhysRevA.90.022305), doi:10.1103/PhysRevA.90.022305.
- 437 **41** J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik. Simulation of electronic structure  
438 Hamiltonians using quantum computers. *Molecular Physics*, 109:735–750, March 2011.  
439 arXiv:1001.3855, doi:10.1080/00268976.2011.552441.
- 440  
441  
442