OpenQASM 3: A Broader and Deeper Quantum Assembly Language

ANDREW CROSS, ALI JAVADI-ABHARI, and THOMAS ALEXANDER, IBM Quantum,

IBM T. J. Watson Research Center

NIEL DE BEAUDRAP, Dept. Informatics, University of Sussex

LEV S. BISHOP, IBM Quantum, IBM T. J. Watson Research Center

STEVEN HEIDEL, COLM A. RYAN, and PRASAHT SIVARAJAH, AWS Center for

Quantum Computing

JOHN SMOLIN, JAY M. GAMBETTA, and BLAKE R. JOHNSON, IBM Quantum, IBM T. J.

Watson Research Center

12

Quantum assembly languages are machine-independent languages that traditionally describe quantum computation in the circuit model. Open quantum assembly language (OpenQASM 2) was proposed as an imperative programming language for quantum circuits based on earlier QASM dialects. In principle, any quantum computation could be described using OpenQASM 2, but there is a need to describe a broader set of circuits beyond the language of qubits and gates. By examining interactive use cases, we recognize two different timescales of quantum-classical interactions: real-time classical computations that must be performed within the coherence times of the qubits, and near-time computations with less stringent timing. Since the near-time domain is adequately described by existing programming frameworks, we choose in OpenQASM 3 to focus on the real-time domain, which must be more tightly coupled to the execution of quantum operations. We add support for arbitrary control flow as well as calling external classical functions. In addition, we recognize the need to describe circuits at multiple levels of specificity, and therefore we extend the language to include timing, pulse control, and gate modifiers. These new language features create a multi-level intermediate representation for circuit development and optimization, as well as control sequence implementation for calibration, characterization, and error mitigation.

CCS Concepts: • **Software and its engineering** → *Software notations and tools; Context specific languages; Domain specific languages;*

Additional Key Words and Phrases: Quantum computing, quantum programming languages, quantum circuits, dynamic circuits, circuit scheduling

NdB's contributions build on work with Jacob Wilkins, as part of the EPSRC-funded NQIT project and as the Oxford-Tencent Fellow in the Department of Computer Science at Oxford. Quantum circuit drawings were typeset using yquant [32].

Authors' addresses: A. Cross, A. Javadi-Abhari, T. Alexander, L. S. Bishop, J. Smolin, J. M. Gambetta, and B. R. Johnson, IBM Quantum, IBM T.J. Watson Research Center, 1101 Kitchawan Rd, Yorktown Heights, NY 10598; N. de Beaudrap, Dept. Informatics, University of Sussex, UK; S. Heidel, C. A. Ryan, and P. Sivarajah, AWS Center for Quantum Computing, 391 S. Holliston Ave, Pasadena, CA 91106.

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

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

2643-6817/2022/09-ART12

<https://doi.org/10.1145/3505636>

ACM Reference format:

Andrew Cross, Ali Javadi-Abhari, Thomas Alexander, Niel de Beaudrap, Lev S. Bishop, Steven Heidel, Colm A. Ryan, Prasahnt Sivarajah, John Smolin, Jay M. Gambetta, and Blake R. Johnson. 2022. OpenQASM 3: A Broader and Deeper Quantum Assembly Language. *ACM Trans. Quantum Comput.* 3, 3, Article 12 (September 2022), 50 pages.

<https://doi.org/10.1145/3505636>

1 INTRODUCTION

Open quantum assembly language (OpenQASM 2) [27] was proposed as an imperative programming language for quantum circuits based on earlier QASM dialects [15, 23, 26, 35, 76], and is one of the programming interfaces of IBM Quantum Services [1]. In the period since OpenQASM 2 was introduced, it has become something of a *de facto* standard, allowing a number of independent tools to inter-operate using OpenQASM 2 as the common interchange format. It also exists within the context of significant other activity in exploring and defining quantum assembly languages for practical use [14, 30, 37, 48, 55, 58, 73]. The features of OpenQASM 2 were strongly influenced by the needs of the time particularly with respect to the capabilities of near-term quantum hardware. We deem it time to update OpenQASM to better support the needs of the next phase of quantum system development as well as to incorporate some of the best ideas that have arisen in the other circuit description languages. One common feature of most of these machine-independent languages is that they describe quantum computation in the quantum circuit model [16, 33, 63], which includes non-unitary primitives such as teleportation [17] and the measurement-model of quantum computing [44, 67]. For practical purposes, we see advantage in a language that goes beyond the circuit model, extending the definition of a circuit to involve a significant amount of classical control flow.

We define an extended quantum circuit as a computational routine, consisting of an ordered sequence of quantum operations on quantum data—such as gates, measurements, and resets—and concurrent real-time classical computation. Dataflows back and forth between the quantum operations and real-time classical computation, in that the classical compute can depend on measurement results, and the quantum operations may involve or be conditioned on data from the real-time classical computation. A practical quantum assembly language should be able to describe such versatile circuits.

Recognizing that OpenQASM is also a resource for describing circuits that characterize, validate, and debug quantum systems, we also introduce instruction semantics that allow control of gate scheduling in the time domain, and that provide the ability to describe the microcoded implementations of gates in a way that is extensible to future developments in quantum control and is platform agnostic. The extension considered here is intended to be backwards compatible with OpenQASM 2, except in some uncommonly used cases where breaking changes were necessary.

As quantum applications are developing where substantial classical processing may be used along with the extended quantum circuits, we allow for a richer family of computation that incorporates interactive use of a **quantum processor (QPU)**. We refer to this abstraction as a quantum program. For example, the variational quantum eigensolver depends on classical computation to calculate an energy or cost function and uses classical optimization to tune circuit parameters [66]. By examining the broader family of interactive use cases, we recognize two different timescales of quantum-classical interactions: *real-time* classical computations that must be performed within the coherence times of the qubits, and *near-time* computations with less stringent timing. The real-time computations are critical for error correction and circuits that take advantage of feedback or

feedforward. However, the demands of the real-time, low-latency domain might impose considerable restrictions on the available compute resources in that environment. For example, real-time computations might be constrained by limited memory or reduced clock speeds. Whereas in the near-time domain, we assume more generic compute resources, including access to a broad set of libraries and runtimes. Since the near-time domain is already adequately described by existing programming tools and frameworks, we choose in OpenQASM 3 to focus on the real-time domain, which must be more tightly coupled to the execution of gates and measurement on qubits.

We use a dataflow model [43] where each instruction may be executed by an independent processing unit when its input data become available. Concurrency and parallelism are essential features of quantum control hardware. For example, parallelism is necessary for a fault-tolerance accuracy threshold to exist [12, 74], and we expect concurrent quantum and classical processing to be necessary for quantum error-correction [28, 31, 36].

Our extended OpenQASM language expresses a quantum circuit, or any equivalent representation of it, as a collection of (quantum) basic blocks and flow control instructions. This differs from the role of a high-level language. High-level languages may include mechanisms for quantum memory management and garbage collection, quantum data types with non-trivial semantics, or features to specify classical reversible programs to synthesize into quantum oracle implementations. Optimization passes at this high level work with families of quantum operations whose specific parameters might not be known yet. These passes could be applied once at this level and benefit every program instance we execute later. High-level intermediate representations may differ from OpenQASM until the point in the workflow where a specific circuit is generated as an instance from the family.

On the other hand, while quantum assembly languages are considered low-level programming languages, they are not often directly consumed by quantum control hardware. Rather, they sit at multiple levels in the middle of the stack where they might be hand-written, generated by scripts (meta-programming), or targeted by higher level software tools and then are further compiled to lower-level instructions that operate at the level of analog signals for controlling quantum systems.

Our choice of features to add to OpenQASM is guided by use cases. Although OpenQASM is not a high-level language, many users would like to write simple quantum circuits by hand using an expressive domain-specific language. Researchers who study circuit compiling need high-level information recorded in the intermediate representations to inform the optimization and synthesis algorithms. Experimentalists prefer the convenience of writing circuits at a relatively high level but often need to manually modify timing or pulse-level gate descriptions at various points in the circuit. Hardware engineers who design the classical controllers and waveform generators prefer languages that are practical to compile given the hardware constraints and make explicit circuit structure that the controllers can take advantage of. Our choice of language features is intended to acknowledge all of these potential audiences.

What follows is a review of the major new features introduced in OpenQASM 3 with a focus on explaining the ideas that motivated those changes and the design intent of how the features work. This manuscript is not a formal specification of the language. Rather, the formal specification exists as a live document [2] backed by a GitHub repository where people may comment or suggest changes through issues and pull requests. While this document hopefully serves as a useful guide to understanding OpenQASM 3, readers interested in writing circuits with compliant syntax are strongly encouraged to refer to the specification [2]. To facilitate continued evolution and refinement of the language, a governance structure has been established [18] including a technical steering committee to formalize how changes will be accepted and to spur the development of working groups to study specific issues of language design.

2 DESIGN PHILOSOPHY AND EXECUTION MODEL

In OpenQASM 3, we aim at describing a broader set of quantum circuits with concepts beyond simple qubits and gates. Chief among them are arbitrary classical control flow, gate modifiers (e.g., control and inverse), timing, and microcoded pulse implementations. While these extensions are not strictly necessary from a theoretical point of view—any quantum computation could in principle be described using OpenQASM 2—in practice they greatly expand the expressivity of the language.

2.1 Versatility in Describing Logical Circuits

One key motivation for the new language features is the ability to describe new kinds of circuits and experiments. For example, repeat-until-success algorithms [65] or magic state distillation protocols [20] have non-deterministic components that can be programmed using the new classical control flow instructions. There are other examples, which show benefits to circuit width or depth from incorporating classical operations. For instance, with measurement and feedforward any Clifford operation can be executed in constant depth [44]. The same is true of the quantum Fourier transform [79]. This extension also allows description of teleportation and feedforward in the measurement model [44, 67]. This collection of examples suggests the potential of the extended quantum circuit family to reduce the requirements needed to achieve quantum advantage in the near term.

It is not our intent; however, to transform OpenQASM into a general-purpose programming language suitable for classical programming. A general quantum application includes classical computations that need to interact with quantum hardware. For instance, some applications require pre-processing of problem data, such as choosing a co-prime in Shor’s algorithm [71], post-processing to compute expectation values of Pauli operators, or further generation of circuits like in the outer loop of many variational algorithms [66]. However, these classical computations do not have tight deadlines requiring execution within the coherence time of the quantum hardware. Consequently, we think of their execution in a *near-time* context, which is more conveniently expressed in existing classical programming frameworks. Our classical computation extensions in OpenQASM 3 instead focus on the *real-time* domain, which must be tightly coupled to quantum hardware. Even in the real-time context, we enable re-use of existing tooling by allowing references to externally specified real-time functions defined outside of OpenQASM itself.

2.2 Versatility in the Level of Description for Circuits

We wish to use the same tools for circuit development *and* for lower-level control sequences needed for calibration, characterization, and error mitigation. Hence, we need the ability to control timing and to connect quantum instructions with their pulse-level implementations for various qubit modalities. For instance, dynamical decoupling [78] or characterizations of decoherence and crosstalk [39] are all sensitive to timing, and can be programmed using the new timing features. Pulse-level calibration of gates can be fully described in OpenQASM 3.

Another design goal for the language is to create a suitable **intermediate representation (IR)** for quantum circuit compilation. Recognizing that quantum circuits have multiple levels of specificity from higher-level theoretical computation to lower-level physical implementation, we envision OpenQASM as a multi-level IR where different levels of abstraction can be used to describe a circuit. We can speak broadly of two levels of abstraction: logical and physical levels. A logical-level circuit (see Section 4) is described by abstract, discrete-time gates, and real-time classical computations. This is lowered by a compiler to a physical-level circuit (see Section 5) where quantum operations are implemented by continuous, time-varying signals with specific timing constraints. We introduce semantics that allow the compiler to optimize and re-write the circuit at different levels. For example, the compiler can reason directly on the level of controlled gates or gates raised

to some power without decomposing these operations, which may obscure some opportunities for optimization. Furthermore, the new semantics allow capturing intent in a portable manner without tying it to a particular implementation. For example, the timing semantics in OpenQASM are designed to allow higher level descriptions of timing constraints without being specific about individual gate durations.

2.3 Continuity with OpenQASM 2

In the design of interfaces for quantum computing, particularly at the early stage that we are currently in, backwards compatibility is not to be prized above all else. However, in recognition of the current importance of the preceding version of OpenQASM as a *de facto* standard, we made design choices to introduce as few incompatibilities between OpenQASM 2 and OpenQASM 3 as practically possible.

Our intent is that an OpenQASM 2 circuit can be interpreted as an OpenQASM 3 circuit with no change in meaning. In doing so, we hope to introduce as little disruption to software projects, which rely on OpenQASM, and to make OpenQASM 3 easier to adopt for those software developers who are already familiar with OpenQASM 2. We describe the supported features of OpenQASM 2 in more detail in Section 3.

2.4 Execution Model

We illustrate a potential application compilation and execution flow for quantum programs in Figure 1. An application might be composed of several quantum programs including near-time calculations and quantum circuits. The quantum program interacts with quantum hardware by emitting OpenQASM3 circuits and external real-time functions. Initial generation of these circuits and **extern** functions require a deep compilation stack to transform the circuits and **extern** into a form that is executable on the QPU. We represent these multiple stages by logical OpenQASM and physical OpenQASM. The logical OpenQASM represents the intent of the extended quantum circuit, whereas the physical OpenQASM is the lowest level where the circuit is mapped and scheduled on specific qubits. A quantum program might work with even higher level circuit representations than those expressible with OpenQASM 3, but in this flow, manipulations of these higher order objects would be contained within the quantum program. A quantum program is not constrained to produce only logical OpenQASM. Sometimes it might emit physical OpenQASM, or later stages of a program might take advantage of lower-cost ways to produce binaries executable by the QPU. For instance, a program describing the optimization loop of a variational algorithm might directly manipulate a data section in the binaries to describe circuits with updated circuit parameters [46]. We generically expect quantum programs to be able to enter the compilation toolchain at varying levels of abstraction as appropriate for different phases of program execution.

Our language model further assumes that the QPU has a global controller that orchestrates execution of the circuit. This global controller centralizes control flow decisions and has the ability to execute **extern** computations emitted by the quantum program. Some QPUs may also have a collection of local controllers that interact with a subset of qubits. Such segmentation of the QPU could enable concurrent execution of independent code segments.

3 COMPARISON TO OPENQASM 2

OpenQASM 3 is designed to allow quantum circuits to be expressed in greater detail, and also in a more versatile way. However, we considered it important in extending this functionality to remain as close to the language structure of OpenQASM 2 as practically possible.

Throughout this article, we occasionally note similarities or differences between OpenQASM versions 2 and 3. In this section, for the convenience of developers who are familiar with

Fig. 1. The compilation and execution model of a quantum program, and OpenQASM’s place in the flow. An application is written in terms of a quantum program. Certain regions of a quantum program are purely classical and can be written in classical programming languages for execution in a near-time context. The quantum program emits payloads for execution on a QPU. This payload consists of extended quantum circuits and external real-time classical functions. OpenQASM is the language to describe the quantum circuits, which includes interface calls to the external classical functions. There may be higher-order aspects of the circuit payload, which are optimized before OpenQASM is generated. An OpenQASM compiler can transform and optimize all aspects of the circuits described with the IR, including basis gates used, qubit mapping, timing, pulses, and control flow. The final physical circuit plus extern functions are passed to a target code generator, which produces binaries for the QPU.

OpenQASM 2, we explicitly describe the OpenQASM 2 features, which developers may rely on to work identically in OpenQASM 3.

Circuit Header. Top-level circuits in OpenQASM 2 must begin with the statement “OpenQASM 2.x;”,¹ possibly preceded by one or more blank lines or lines of comments, where x is a minor version number (such as 0). Other source-files could be included using the syntax `include "filename"`, in the header or elsewhere in the circuit. Top-level circuits in OpenQASM 3 begin similarly, e.g., with a statement “OpenQASM 3.0;” and OpenQASM 3 supports `include` statements in the same way.

Circuit Execution. The top-level of an OpenQASM 2 circuit is given by a sequence of instructions at the top-level scope. No means of specifying an explicit entry point is provided: The circuit is

¹Throughout the rest of the article, to reduce clutter, we often omit terminal semicolons for in-line examples.

interpreted as a stream of instructions, with execution starting from the first instruction after the header. A top-level OpenQASM 3 circuit may also be provided in this way: as a sequence of instructions at the top-level scope, without requiring an explicit entry point.

Quantum and Classical Registers. In OpenQASM 2, the only storage types available are qubits, which are allocated as part of a **qreg** declaration, and classical bits, which are allocated as part of a **creg** declaration, for example:

```
1 qreg q[5]; // a register 'q' with five qubits
2 creg c[2]; // a register 'c' with two classical bits
```

This syntax is supported in OpenQASM 3, and is equivalent to the (preferred²) syntax,

```
1 qubit[5] q; // a register 'q' of five qubits
2 bit[2] c; // a register 'c' of two classical bits
```

which, respectively, declares *q* and *c* as registers of more primitive **qubit** or **bit** storage types. (OpenQASM 3 also introduces other storage types for classical data: see Section 4.4.)

Names of Gates, Variables, and Constants. In OpenQASM 2, the names of registers and gates must begin with a lower-case alphabetic ASCII character. This constraint is relaxed in OpenQASM 3: Identifiers may now begin with other characters, such as capital letters, underscores, and a range of unicode characters. For example, angular values used as gate arguments may now be represented by greek letters. (The identifier π is reserved in OpenQASM 3 to represent the same constant as *pi*.) Note that OpenQASM 3 has a slightly different set of keywords from OpenQASM 2. This may cause errors in OpenQASM circuits, which happen to use a keyword from version 3 as an identifier name: such identifiers would have to be renamed in order for the source to be valid OpenQASM 3.

Basic Operations. In OpenQASM 2, there were four basic instructions, which affected stored quantum data:

- Single-qubit unitaries specified with the syntax **U**(*a*, *b*, *c*) for some angular parameters *a*, *b*, and *c*, acting on a single-qubit argument. This syntax is also supported in OpenQASM 3, and specifies an equivalent unitary transformation.³
- Two-qubit controlled-NOT operations, using the keyword **CX**, acting on two single-qubit arguments. In OpenQASM 3, the **CX** instruction is *no longer* a basic instruction or a keyword, but may be defined from more primitive commands. To adapt an OpenQASM 2 circuits to be valid OpenQASM 3, one may add a single instruction to define **CX** using a “**ctrl** @” gate modifier:

```
1 gate CX c, t { ctrl @ U( $\pi$ , 0,  $\pi$ ) c, t }
```

A definition of this sort for **CX** is also provided in the standard gate library for OpenQASM 3, which we describe in Section 4.1.1. (The meaning of the “**ctrl** @” gate modifier is described in Section 4.2.)

- Non-unitary **reset** and **measure** operations, acting on a single-qubit argument (and producing a single bit outcome in the case of **measure**). The statement **reset** *q*[*j*] operation has the effect of discarding the data stored in *q*[*j*], replacing it with the state $|0\rangle$. The statement **measure** *q*[*j*] -> *c*[*k*] has the effect of measuring *q*[*j*] in the standard basis, i.e., projecting the state into one of the eigenstates of the Pauli *Z* operator, and storing the corresponding

²The **qreg** and **creg** keywords may not be supported in future versions of OpenQASM.

³The OpenQASM 3 specification for single-qubit unitaries, described in Equation (1), differs by a global phase from the specification in OpenQASM 2. This change has no effect on the results of an OpenQASM 2 circuits.

bit-value in a classical bit $c[k]$. OpenQASM 3 supports these operations without changes, and also supports the alternative syntax $c[k] = \text{measure } q[j]$ for measurements.

Gate Declarations. OpenQASM 2 supports **gate** declarations to specify user-defined gates. The definitions set out a fixed number of single-qubit arguments, and optionally some arguments, which are taken to be angular parameters. These declarations use syntax such as

```
1 gate h a { U( $\pi/2$ , 0,  $\pi$ ) a; }
2 gate p( $\theta$ ) a { U(0, 0,  $\theta$ ) a; }
3 gate cp( $\theta$ ) a, b { p( $\theta/2$ ) a; p( $\theta/2$ ) b; CX a,b; p( $-\theta/2$ ) b; CX a,b; }
```

The specification of these gates, in the code-blocks enclosed by braces $\{ \dots \}$, are by sequences of *unitary* operations, whether basic unitary operations or ones defined by earlier **gate** declarations. OpenQASM 3 supports **gate** declarations with this syntax, and also admits only unitary operations as part of the gate definitions—but has greater versatility in how those unitary operations may be described (see Section 4.2), and also provides other ways to define subroutines (see Section 4.4).

Implicit Iteration. Operations on one or more qubits could be repeated across entire registers as well, using implicit iteration. For instance, for any operation “qop $q[j]$ ”, which acts on one qubit $q[j]$, we can instead perform it independently on every qubit in q by omitting the index (as in “qop q ”). For operations on two qubits of the form “qop $q[j]$, $r[k]$ ”, one or both arguments could omit the index, in which case the instruction would be repeated either for $j = 1, 2, \dots$ or $k = 1, 2, \dots$ (if one index is omitted), or for $j = k = 1, 2, \dots$ (if both indices are omitted). This functionality also extends to operations on three qubits or more. OpenQASM 3 supports this functionality unchanged, for **gate** subroutines. (OpenQASM 3 also supports other kinds of subroutine beyond **gate** definitions, and these other subroutines do not support implicit iteration in this way.)

Control Flow. The only control flow supported by OpenQASM 2 are **if** statements. These can be used to compare the value of a classical bit-register (interpreted as a little-endian representation of an integer) to an integer, and conditionally execute a single gate. An example of such a statement is

```
1 if (c == 5) mygate q, r, s;
```

which would test whether a classical register c (of length three or more) stores a bit-string $c_{n-1} \dots c_3 c_2 c_1 c_0$ equal to $0 \dots 0101$, to determine whether to execute `mygate q, r, s`. If statements (and classical register comparisons) of this kind are supported in OpenQASM 3, as is more general syntax for **if** statements and other forms of control-flow (see Section 4.4).

Barrier Instructions. OpenQASM 2 provides a **barrier** instruction, which may be invoked with or without arguments. When invoked with arguments, either of individual qubits or whole quantum registers, it instructs the compiler not to perform any optimizations that involve moving or simplifying operations acting on those arguments, across the source line of the **barrier** statement; when invoked without arguments, it has the same effect as applying it to all of the quantum registers that have been defined. This operation is also supported in OpenQASM 3, with the same meaning.

Opaque Definitions. OpenQASM 2 supports opaque declarations, to declare gates whose physical implementation may be possible, but is not readily expressed in terms of unitary gates. OpenQASM 3 provides means of defining operations on a lower level than unitary gates (see Section 5.2); therefore, opaque definitions are not needed. OpenQASM 3 compilers will simply ignore any opaque declarations.

Circuit Output. The outputs of an OpenQASM 2 circuit are the values stored in any declared classical registers. OpenQASM 3 introduces a means of explicitly declaring, which variables are to be produced as output or taken as input (see Section 4.5). However, any OpenQASM 3 circuit, which does not specify either an output or input will, by default, also produce all of its classical stored variables (whether of type **creg**, or a different type) as output.

4 CONCEPTS OF THE LANGUAGE: THE LOGICAL LEVEL

OpenQASM 3 is a multi-level IR for quantum computations, which expresses concepts at both a logical and a more fine-grained physical level. In this section, we give an overview of the important features of the logical level of OpenQASM 3, presenting the more low-level features in Section 5. For a finer-grained specification, we direct readers to the live specification [2].

4.1 Continuous Gates and Hierarchical Library

We define a mechanism for parameterizing unitary matrices to define quantum gates. The parameterization uses a set of built-in single-qubit gates and a gate modifier to construct a controlled version to generate a universal gate set [16]. Early QASM languages assumed a discrete set of quantum gates, but OpenQASM is flexible enough to describe universal computation with a continuous gate set. This gate set was chosen for the convenience of defining new quantum gates and is not an enforced compilation target. Instead of allowing the user to write a unitary as an $n \times n$ matrix, we make gate definitions through hierarchical composition allowing for code reuse to define more complex operations [35, 42]. For many gates of practical interest, there is a circuit representation with a polynomial number of one- and two-qubit gates, giving a more compact representation than requiring the programmer to express the full $n \times n$ matrix. In the worst case, a general n -qubit gate can be defined using an exponential number of these gates.

We now describe this built-in gate set. Single-qubit unitary gates are parameterized as

$$U(\theta, \phi, \lambda) := \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\phi+\lambda)} \cos(\theta/2) \end{pmatrix}. \quad (1)$$

This expression specifies any element of $U(2)$ up to a global phase. The global phase is here chosen so that the upper left matrix element is real. For example, $U(\pi/2, 0, \pi)$ $q[0]$; applies the Hadamard gate $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ to qubit $q[0]$. For the convenience of the reader, we note two important decompositions of U , through an Euler angle decomposition using rotations $R_z(\alpha) = \exp(-i\alpha Z/2)$ and $R_y(\beta) = \exp(-i\beta Y/2)$, and a $\{H, P\}$ basis representation where $P(\alpha) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\alpha} \end{pmatrix}$:

$$U(\theta, \phi, \lambda) = e^{i(\phi+\lambda)/2} R_z(\phi) R_y(\theta) R_z(\lambda) = e^{i\theta/2} P(\phi+\pi/2) H P(\theta) H P(\lambda-\pi/2). \quad (2)$$

Note. Users should be aware that the definition in Equation (1) is scaled by a factor of $e^{i(\phi+\lambda)/2}$ when compared to the original definition in OpenQASM 2 [27]. This implies that the transformation described by $U(a, b, c)$ has determinant $e^{i(a+c)}$, and is not in general an element of $SU(2)$. This change in global phase has no observable consequence for OpenQASM 2 circuits. We also note that **U** is the only uppercase keyword in OpenQASM 3.

New gates are associated with a unitary transformation by defining them using a sequence of built-in or previously defined gates and gate modifiers. For example, the **gate** declaration

```
1 gate h q { U(pi/2, 0, pi) q; }
```

defines a new gate called “h” and associates it to the unitary matrix of the Hadamard gate. Once, we have defined “h”, we can use it in later **gate** blocks. The definition does not necessarily imply

that h is implemented by an instruction $U(\pi/2, \emptyset, \pi)$ on the quantum computer. The implementation is left to the user and/or compiler (see Section 5.2), given information about the instructions supported by a particular target.

Controlled gates can be constructed by attaching a control modifier to an existing gate. For example, the NOT gate (i.e., the Pauli X operator) is given by $U(\pi, \emptyset, \pi)$ and the block

```
1 gate CX c, t { ctrl @ U( $\pi$ ,  $\emptyset$ ,  $\pi$ ) c, t; }
2 CX q[1], q[0];
```

defines the gate

$$CX := I \oplus x = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad (3)$$

and applies it to $q[1]$ and $q[0]$. This gate applies a bit-flip to $q[0]$ if $q[1]$ is one, leaves $q[1]$ unchanged if $q[0]$ is 0, and acts coherently over superpositions. The control modifier is described in more detail in Section 4.2.

Remark. Throughout the document, we use a tensor order with higher index qubits on the left. This ordering labels states in a way that is consistent with how numbers are usually written with the most significant digit on the left. In this tensor order, $CX \ q[0], q[1]$; is represented by the matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \quad (4)$$

From a physical perspective, any unitary U is indistinguishable from another unitary $e^{i\gamma}U$, which only differs by a global phase. When we attach a control to these gates, however, the global phase becomes a relative phase that is applied when the control qubit is one. To capture the programmer's intent, a built-in global phase gate allows the inclusion of arbitrary global phases on circuits. The instruction `gphase( $\gamma$ )` adds a global phase of $e^{i\gamma}$ to the scope containing the instruction. For example,

```
1 gate rz( $\tau$ ) q { gphase( $-\tau/2$ ); U( $\emptyset$ ,  $\emptyset$ ,  $\tau$ ) q; }
2 ctrl @ rz( $\pi/2$ ) q[1], q[0];
```

constructs the gate

$$R_z(\tau) = \exp(-i\tau Z/2) = \begin{pmatrix} e^{-i\tau/2} & 0 \\ 0 & e^{i\tau/2} \end{pmatrix} = e^{-i\tau/2} \begin{pmatrix} 1 & 0 \\ 0 & e^{i\tau} \end{pmatrix}, \quad (5)$$

and applies the controlled form of that phase rotation gate

$$I \oplus R_z(\pi/2) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & e^{-i\pi/4} & 0 \\ 0 & 0 & 0 & e^{i\pi/4} \end{pmatrix}. \quad (6)$$

In OpenQASM 2, `gate` subroutines are defined in terms of built-in gates `U` and `CX`. In OpenQASM 3, we can define controlled gates using the control modifier, so it is no longer necessary to include a built-in `CX` gate. For backwards compatibility, we include the gate `CX` in the standard gate library (described below), but it is no longer a keyword of the language.

4.1.1 Standard Gate Library. We define a standard library of OpenQASM3 gates in a file, we call `stdgates.inc`. Any OpenQASM3 circuit that includes the standard library can make use of these gates.

```

1  // OpenQASM 3 standard gate library
2
3  // phase gate
4  gate p( $\lambda$ ) a { ctrl @ gphase( $\lambda$ ) a; }
5
6  // Pauli gate: bit-flip or NOT gate
7  gate x a { U( $\pi$ , 0,  $\pi$ ) a; }
8  // Pauli gate: bit and phase flip
9  gate y a { U( $\pi$ ,  $\pi/2$ ,  $\pi/2$ ) a; }
10 // Pauli gate: phase flip
11 gate z a { p( $\pi$ ) a; }
12
13 // Clifford gate: Hadamard
14 gate h a { U( $\pi/2$ , 0,  $\pi$ ) a; }
15 // Clifford gate: sqrt(2) or S gate
16 gate s a { pow(1/2) @ z a; }
17 // Clifford gate: inverse of sqrt(2)
18 gate sdg a { inv @ pow(1/2) @ z a; }
19
20 // sqrt(S) or T gate
21 gate t a { pow(1/2) @ s a; }
22 // inverse of sqrt(S)
23 gate tdg a { inv @ pow(1/2) @ s a; }
24
25 // sqrt(NOT) gate
26 gate sx a { pow(1/2) @ x a; }
27
28 // Rotation around X-axis
29 gate rx( $\theta$ ) a { U( $\theta$ ,  $-\pi/2$ ,  $\pi/2$ ) a; }
30 // rotation around Y-axis
31 gate ry( $\theta$ ) a { U( $\theta$ , 0, 0) a; }
32 // rotation around Z axis
33 gate rz( $\lambda$ ) a { gphase( $-\lambda/2$ ); U(0, 0,  $\lambda$ ) a; }
34
35 // controlled-NOT
36 gate cx c, t { ctrl @ x c, t; }
37 // controlled-Y
38 gate cy a, b { ctrl @ y a, b; }
39 // controlled-Z
40 gate cz a, b { ctrl @ z a, b; }
41 // controlled-phase
42 gate cp( $\lambda$ ) a, b { ctrl @ p( $\lambda$ ) a, b; }
43 // controlled-rx
44 gate crx( $\theta$ ) a, b { ctrl @ rx( $\theta$ ) a, b; }
45 // controlled-ry
46 gate cry( $\theta$ ) a, b { ctrl @ ry( $\theta$ ) a, b; }
47 // controlled-rz
48 gate crz( $\theta$ ) a, b { ctrl @ rz( $\theta$ ) a, b; }
49 // controlled-H
50 gate ch a, b { ctrl @ h a, b; }
51
52 // swap
53 gate swap a, b { cx a, b; cx b, a; cx a, b; }
54
55 // Toffoli
56 gate ccx a, b, c { ctrl @ ctrl @ x a, b, c; }
57 // controlled-swap
58 gate cswap a, b, c { ctrl @ swap a, b, c; }
59
60 // four parameter controlled-U gate with relative phase  $\gamma$ 
61 gate cu( $\theta$ ,  $\varphi$ ,  $\lambda$ ,  $\gamma$ ) c, t { p( $\gamma$ ) c; ctrl @ U( $\theta$ ,  $\varphi$ ,  $\lambda$ ) c, t; }
62
63 // Gates for OpenQASM 2 backwards compatibility
64 // CNOT
65 gate CX c, t { ctrl @ U( $\pi$ , 0,  $\pi$ ) c, t; }
66 // phase gate
67 gate phase( $\lambda$ ) q { U(0, 0,  $\lambda$ ) q; }
68 // controlled-phase
69 gate cphase( $\lambda$ ) a, b { ctrl @ phase( $\lambda$ ) a, b; }
70 // identity or idle gate
71 gate id a { U(0, 0, 0) a; }
72 // IBM Quantum experience gates
73 gate u1( $\lambda$ ) q { U(0, 0,  $\lambda$ ) q; }
74 gate u2( $\varphi$ ,  $\lambda$ ) q { gphase( $-(\varphi+\lambda)/2$ ); U( $\pi/2$ ,  $\varphi$ ,  $\lambda$ ) q; }
75 gate u3( $\theta$ ,  $\varphi$ ,  $\lambda$ ) q { gphase( $-(\varphi+\lambda)/2$ ); U( $\theta$ ,  $\varphi$ ,  $\lambda$ ) q; }

```

4.2 Gate Modifiers

We introduce a mechanism for modifying existing unitary gates g to define new ones. A modifier defines a new unitary gate from g , acting on a space of equal or greater dimension, that can be applied as an instruction. We add modifiers for inverting, exponentiating, and controlling gates. This can be useful for programming convenience and readability, but more importantly it allows the language to capture gate semantics at a higher level which aids in compilation. For example, optimization opportunities that exist by analysing controlled unitaries may be exceedingly hard to discover once the controls are decomposed.

The Control Modifier. The modifier `ctrl @ g` represents a controlled- g gate with one control qubit. This gate is defined on the tensor product of the control space and target space by the matrix $I \oplus g$, where I has the same dimensions as g . (We regard `ctrl @ gphase(a)` as a special case, which is defined to be the gate $1 \oplus e^{ia} = U(0, 0, a)$.) A compilation target may or may not be able to execute the controlled gate without further compilation. The user or compiler may rewrite the controlled gate to use additional ancillary qubits to reduce the size and depth of the resulting circuit.

For example, we can define the Fredkin (or controlled-swap) gate in two different ways:

```

1 // Some useful reversible gates
2 gate x q { U( $\pi$ , 0,  $\pi$ ) q; }
3 gate cx c, t { ctrl @ x c, t; }
4 gate toffoli c0, c1, t { ctrl @ cx c0, c1, t; }
5 gate swap a, b { cx a, b; cx b, a; cx a, b; }
6
7 // Fredkin definition #1
8 gate fredkin1 c, a, b {
9   cx b, a;
10  toffoli c, a, b;
11  cx b, a;
12 }
13
14 // Fredkin definition #2
15 gate fredkin2 c, a, b { ctrl @ swap c, a, b; }
```

The definitions are equivalent because they result in the same unitary matrix for both Fredkin gates. If we directly apply the definitions, `fredkin1` involves fewer controlled operations than `fredkin2`. A compiler pass may infer this fact or even rewrite `fredkin2` into `fredkin1`. Note that although the `toffoli` gate is well defined here, the compiler or user needs to do further work to synthesize it in terms of more primitive operations, such as 1- and 2-qubit gates, or pulses (see Section 5.2).

For a second example, if we go one step further, we can define a controlled Fredkin gate. If we use the $ABA^\dagger$ pattern in the definition of the `fredkin1` gate, we require only one controlled `toffoli` gate after we remove unneeded controls from the `cx` gates. Although the controlled `toffoli` gate is defined on four qubits, it may be beneficial for a compiler or user to synthesize the gate using extra qubits as scratch space to reduce the total gate count or circuit depth, as in the following example:

```

1 // Require: the scratch qubit is 'clean', i.e. initialized to zero
2 // Ensure: the scratch qubit is returned 'clean'
3 gate cfredkin2 c0, c1, a, b, scratch {
4   cx b, a; // (from fredkin1)
5   toffoli c0, c1, scratch; // implement ctrl @ toffoli c0, c1, a, b
6   toffoli scratch, a, b; // by computing (and uncomputing) an AND of the two controls
7   toffoli c0, c1, scratch; // in the scratch qubit
8   cx b, a; // (from fredkin1)
9 }
```

We further provide the `negctrl` modifier to control operations with *negative* polarity, i.e., conditioned on the control bit being 0 rather than 1. Both `ctrl` and `negctrl` optionally accept an

Fig. 2. Toffoli network to compute a Boolean function $f = \bar{a}_0\bar{a}_1a_2a_3a_4 \oplus \bar{a}_0a_1a_2\bar{a}_3 \oplus \bar{a}_0a_1a_3a_4 \oplus a_0\bar{a}_1a_3a_4 \oplus a_0a_1a_2 \oplus \bar{a}_0\bar{a}_1\bar{a}_2a_3$ in a target qubit. The control pattern may be expressed in OpenQASM via control gate modifiers.

argument n , which specifies the number of controls of the appropriate type, which are added to the front of the list of arguments (omission means $n = 1$). In each case, n must be a positive integer expression which is a “compile-time constant”.⁴ For example, this can be useful in writing classical Boolean functions [69], as demonstrated by the following example (illustrated in Figure 2).

```

1 // Reversible Boolean function expressed using the control modifier.
2 qubit a[5];
3 qubit f;
4 reset f;
5 negctrl(2) @ ctrl(3) @ x a, f[0];
6 negctrl(2) @ ctrl(2) @ x a[0], a[3], a[1], a[2], f[0];
7 negctrl @ ctrl(3) @ x a[0], a[1], a[3], a[4], f[0];
8 negctrl @ ctrl(3) @ x a[1], a[0], a[3], a[4], f[0];
9 ctrl(3) @ x a[0], a[1], a[2], f[0];
10 negctrl(3) @ ctrl @ x a[0], a[1], a[2], a[3], f[0];

```

The Inversion Modifier. The modifier `inv @ g` represents the inverse $g^\dagger$ of the gate g . This can be easily computed from the definition of the gate g as follows:

- The inverse of any unitary operation $U = U_m U_{m-1} \dots U_1$ can be defined recursively by reversing the order of the gates in its definition and replacing each of those with their inverse $U^\dagger = U_1^\dagger U_2^\dagger \dots U_m^\dagger$.
- The inverse of a controlled operation $I \oplus U$ can be defined by reversing the operation, which is controlled, as $(I \oplus U)^\dagger = I \oplus U^\dagger$ for any unitary U . That is, `inv @ ctrl @ g` is defined to have the same meaning as `ctrl @ inv @ g` (commuting the “`inv @`” modifier past the “`ctrl @`” modifier).
- The base case is given by replacing `inv @ U(a, b, c)` by `U(-a, -c, -b)`, and replacing `inv @ gphase(a)` by `gphase(-a)`.

For example,

```

1 gate rz(τ) q { gphase(-τ/2); U(0, 0, τ) q; }
2 inv @ rz(π/2) q[0];

```

applies the gate `inv @ rz(π/2)`, which is defined by `{ gphase(π/4); U(0, -π/2, 0); }`. For an additional example: The gate `inv @ ctrl @ rz(π/2) q[1], q[0]` would be rewritten as `ctrl @ inv @ rz(π/2) q[1], q[0]`, which is the same as `ctrl @ rz(-π/2) q[1], q[0]`.

⁴By “compile-time constant”, we mean an expression, which is actually constant, or which depends only on (a) variables, which take fixed values, or (b) iterator variables in `for` loops, which can be unrolled by the compiler (i.e., which have initial and final iterator values, which are themselves compile-time constants).

The Powering Modifier. The modifier `pow(r) @ g` represents the r th power g^r of the gate g , where r is an integer or floating point number. Every unitary matrix U has a unique principal logarithm $\log U = iH$, where H has only real eigenvalues E , which satisfy $-\pi < E \leq \pi$. For any real r , we define the r th power as $U^r = \exp(irH)$.

For example, consider the Pauli Z operator $z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. The principal logarithm is $\log z = \begin{pmatrix} 0 & 0 \\ 0 & i\pi \end{pmatrix}$, from which it follows that $z^{1/2} = \exp(\frac{1}{2} \log z) = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$. In this way, the definitions

```
1 gate z q { U(0, 0, π) q; }
2 gate s q { pow(1/2) @ z q; }
```

allow us to define the gate s as the square root of z . (A compiler pass is responsible for resolving S into gates.) For another example, the principal logarithm of the Pauli X operator $x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ is $\log x = \frac{i\pi}{2} \begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}$. One can confirm that this operator has eigenvalues 0 and $i\pi$, and $x^{1/2} := \exp(\frac{1}{2} \log x) = \frac{1}{2} \begin{pmatrix} 1+i & 1-i \\ 1-i & 1+i \end{pmatrix}$. Therefore, the definitions of the $\sqrt{X}$ gate

```
1 gate sx a { gphase(π/4); U(π/2, -π/2, π/2) a; }
2 gate sx a { pow(1/2) @ x a; }
```

are equivalent.

In the case where r is an integer, the r th power `pow(r) @ g` can be implemented simply (albeit less efficiently) as r repetitions of g when $r > 0$, or r repetitions of `inv @ u` when $r < 0$. Continuing our previous example,

```
1 pow(-2) @ s q[0];
```

defines and applies s^{-2} . This can be compiled to

```
1 inv @ pow(2) @ s q[0];
```

and further simplified into

```
1 inv @ z q[0];
```

without solving the problem of synthesizing s . In general, however, compiling `pow(r) @ g` involves solving a circuit synthesis or optimization problem, even in the integral case.

4.3 Non-unitary Operations

OpenQASM 3 includes two basic non-unitary operations, which are the same as those found in OpenQASM 2.

The statement `bit = measure q` measures a qubit q in the Z -basis, and assigns the measurement outcome, “0” or “1”, to the target bit variable (see the subsection on classical types in Section 4.4). Measurement is “non-destructive”, in that it corresponds to a *projection* onto one of the eigenstates of Z ; the qubit remains available afterwards for further quantum computation. The `measure` statement also works with an array of qubits, performing a measurement on each of them and storing the outcomes into an array of bits of the same length. For example, the example below initializes, flips, and measures a register of 10 qubits.

```
1 qubit[10] qubits;
2 bit[10] bits;
3 reset qubits;
4 x qubits;
5 bits = measure qubits;
```


For compatibility with OpenQASM 2, we also support the syntax `measure q -> r` for a qubit `q` (or quantum register of some length `n`), and a bit variable `r` (or a classical bit register of the same length).

The statement `reset q` resets a qubit `q` to the state $|0\rangle$. With idealised quantum hardware, this is equivalent to measuring `q` in the standard basis, performing a Pauli X operation on the qubit if the outcome is “1”, and then discarding the measurement outcome. Mathematically, it corresponds to a partial trace over `q` (i.e., discarding it) before replacing it with a new qubit in the state $|0\rangle$. The `reset` statement also works with an array of qubits, resetting them to the state $|0\rangle \otimes \cdots \otimes |0\rangle$.

4.4 Real-time Classical Computing

OpenQASM 2 primarily described *static* circuits in which the only mechanism for control flow were `if` statements that controlled execution of a single `gate`. This constraint was largely imposed by corresponding limitations in control hardware. *Dynamic* circuits—with classical control flow and concurrent classical computation—represent a richer model of computation, which includes features that are necessary for eventual fault-tolerant quantum computers that must interact with real-time decoding logic. In the near term, these circuits also allow for experimentation with qubit re-use, teleportation, and iterative algorithms.

Dynamic circuits have motivated significant advances in control hardware capable of moving and acting upon real-time values [22, 38, 68]. To take advantage of these advances in control, we extend OpenQASM with classical data types, arithmetic and logical instructions to manipulate the data, and control flow keywords. While our approach to selecting classical instructions was conservative, these extensions make OpenQASM Turing-complete for classical computations in principle, augmenting the prior capability to describe static circuits composed of unitary gates and measurement.

Classical Types: In considering what classical instructions to add, we felt it important to be able to describe arithmetic associated with looping constructs and the bit manipulations necessary to shuttle data back and forth between qubit measurements and other classical registers. As a result, simple classical computations can be directly embedded within an OpenQASM 3 circuit. To support these, we introduce classical data types like signed/unsigned integers and floating point values, to provide well-defined semantics for arithmetic operations.

The OpenQASM type system was designed with two distinct requirements in mind. When used within high-level logical OpenQASM circuits the type system must capture the programmer’s intent. It must also remain portable to different kinds of quantum computer controllers. When used within low-level physical OpenQASM circuits the type system must reflect the realities of particular controller, such as limited memory and well-defined register lengths. OpenQASM takes inspiration from type systems within classical programming languages but adds some unique features for dynamic quantum circuits.

For logical-level OpenQASM circuits, we introduce some common types, which will be familiar to most programmers: `int` for signed integers, `uint` for unsigned integers, `float` for floating point numbers, `bool` for boolean true or false values, and `bit` for individual bits. The precision of the integer and floating-point types is not specified within the OpenQASM language itself: rather the precision is presumed to be a feature of a particular controller. Simple mathematical operations are available on these types, for example:

```
1 int x = 4;
2 int y = 2 ** x; // 16
```

More complex classical operations are delegated to `extern` functions, as we describe below.

For low-level platform-specific OpenQASM circuits, we introduce syntax for defining arbitrary bit widths of classical types. In addition to the hardware-agnostic **int** type, OpenQASM 3 also allows programmers to specify an integer with exact precision **int**[*n*] for any positive integer *n*. For example,

```
1 uint[8] x = 0;
2 // some time later...
3 bit a = x[2]; // read the 3rd least-significant bit of 'x' and assign it to 'a'
4 x[7] = 1;     // update the most-significant bit of 'x' to 1
```

shows direct access to the underlying bits in classical values. This kind of flexibility is common in hardware description languages like VHDL and Verilog where manipulations of individual bits within registers are common. We expect similar kinds of manipulations to be common in OpenQASM 3, and so we allow for direct access to bit-level updates within classical values. In practice, we expect target platforms for OpenQASM 3 to only support particular bit widths of each classical type: e.g., some target hardware might only allow for **uint**[8] and **uint**[16], while other targets might only offer **uint**[20]. This is intended to allow OpenQASM 3 to describe platform-specific code, which is tailored to the registers and other resources of controllers of a quantum circuit unit. If the programmer does not require a particular bit width, they may leave the bit-width unspecified: Bit-level operations on the variable would no longer be permitted, but an appropriate choice of bit-width would then be deferred to a platform-specific compiler.

OpenQASM 3 also introduces a fixed-point **angle** type to represent angles or phases. These may also be defined to have a specific bit-width through the types **angle**[*n*]. Values stored in variables of the type **angle**[*n*] may be interpreted as representing values in the range $[0, 2\pi)$,⁵ with a fractional binary expansion of the form $2\pi \times 0.c_{n-1}c_{n-2} \dots c_0$. (For example, with an **angle**[4], the binary representation of π would be $c_3c_2c_1c_0 = 1000$, using a big-endian representation; the representation of $\pi/2$ would be 0100 ; and so forth.) This representation of phases is common to both quantum phase estimation [51] as well as numerically-controlled oscillators. The latter case is particularly relevant to real-time control systems that need to track qubit frames at run-time, where they can take advantage of numerical overflow to automatically constrain the phase representation to the domain $[0, 2\pi)$, obviating the need for expensive floating-point modular arithmetic. The confluence of this angle representation in controllers and quantum algorithms influenced us to add first-class support for this type in OpenQASM 3.

Classical Control Flow: OpenQASM 2 circuits were effectively limited to straight line code: any quantum algorithm with more involved constructs, such as loops, could only be specified through meta-programming in a more general tool such as Python. The resulting OpenQASM 2 circuits could end up being quite lengthy in some cases, and lost all the structure from their original synthesis. In OpenQASM 3, we introduce **while** loops and **for** loops. We also extend the **if** statement syntax to allow for multiple instructions within the body of the **if**, and to allow for an **else** block to accompany it.

As an example, consider an inverse quantum Fourier transform circuit as written in OpenQASM 2 (Figure 3):

```
1 // Fourier basis sampling procedure, with the loop unrolled
2 qreg q[4];
3 creg c[4];
4
5 // An example preparation of register q, to be sampled:
```

⁵The **angle** types do not actually commit to the value of the angle being positive, and can also be interpreted as representing values in the range $[-\pi, \pi)$; the description above is provided for the sake of concreteness.

Fig. 3. Fourier sampling circuit with the loop unrolled.

```

6  reset q;
7  h q;
8  barrier q;
9
10 // Fourier sampling circuit:
11 h q[0];
12 measure q[0] -> c[0];
13
14 if (c[0] == 1) p(-pi/2) q[1];
15 h q[1];
16 measure q[1] -> c[1];
17
18 if (c[0] == 1) p(-pi/4) q[2];
19 if (c[1] == 1) p(-pi/2) q[2];
20 h q[2];
21 measure q[2] -> c[2];
22
23 if (c[0] == 1) p(-pi/8) q[3];
24 if (c[1] == 1) p(-pi/4) q[3];
25 if (c[2] == 1) p(-pi/2) q[3];
26 h q[3];
27 measure q[3] -> c[3];

```

(The **barrier** statement is discussed in Section 5.1.) One can recognize a natural iterative structure to this circuit, but the lack of a looping construct has forced us to unroll this structure when writing it out. Furthermore, each iteration has a series of **if** statements that in principle could be combined, as the body of those statements is a Z rotation in all cases. In OpenQASM 3, we might instead write the above circuit as follows:

```

1  // Fourier sampling circuit, expressed with a for loop
2  qubit[4] q;
3  bit b;
4  angle[4]  $\theta$  = 0; // an angular parameter, for classically controlled Z rotations
5
6  // An example preparation of register q, to be sampled:
7  reset q;
8  h q;
9  barrier q;
10
11 // Fourier sampling circuit:
12 for i in [0:3] {
13    $\theta$  >>= 1; // Divide by 2, all the rotation angles controlled by past measurement outcomes
14   p(- $\theta$ ) q[i]; // Do the classically controlled Z rotations all at once
15   h q[i];
16   measure q[i] -> b;
17    $\theta$ [3] = b; // Set new highest bit of  $\theta$  to measurement outcome
18 }

```

which preserves the iterative structure of the inverse QFT and uses the **angle**[4] type (and operations on its bit-level representation) to directly convert measurement bits into the appropriate Z rotations.

Subroutines: Sufficiently simple functions (accompanied by sufficiently capable controllers) allow for classical subroutines to be fully implemented within the language itself. Consider a simple vote function, which takes the majority “vote” of three **bit** parameters. This could be implemented as follows:

```

1  def vote(bit a, bit b, bit c) -> bit {
2      int count = 0;
3      if (a == 1) count++;
4      if (b == 1) count++;
5      if (c == 1) count++;
6
7      if (count >= 2) {
8          return 1;
9      } else {
10         return 0;
11     }
12 }

```

We extend this functionality to quantum procedures as well. For instance, a repeat-until-success quantum circuit [65] is naturally described by a **while** loop involving measurements, where the computation terminates only after a particular measurement outcome is observed. That is, the classical controllers make dynamic decisions about which operations are to be performed. The following is one example:

```

1  // Repeat-until-success circuit for rz( $\pi + \theta$ ) where  $\cos(\theta)=3/5$ 
2
3  // This subroutine applies the identity to psi if the output is 01, 10, or 11.
4  // If the output is 00, it applies a Z-rotation by the angle  $\theta + \pi$  where  $\cos(\theta)=3/5$ .
5  def segment(qubit[2] anc, qubit psi) -> bit[2] {
6      bit[2] b;
7      reset anc;
8      h anc;
9      ccx anc[0], anc[1], psi;
10     s psi;
11     ccx anc[0], anc[1], psi;
12     z psi;
13     h anc;
14     b = measure anc;
15     return b;
16 }
17
18 // Main circuit for the repeat-until-success circuit
19 qubit in;
20 qubit[2] ancilla;
21 bit[2] flags = b"11";
22 bit out;
23
24 reset in;
25 h in;
26
27 // The segment subroutine returns the desired outcome 00 with probability 5/8,
28 // so we iterate a non-deterministic number of times until that outcome is seen.
29 while (flags != b"00") { // braces are optional in this case
30     flags = segment(ancilla, in);
31 }
32 rz( $\pi - \arccos(3 / 5)$ ) in; // total rotation of  $2\pi$ 
33 h in;
34 out = measure in; // output should equal zero

```

External Functions: Rather than bulk up the feature set of OpenQASM for classical computation further—which would require sophisticated classical compiler infrastructure to manage—we selected only those elements we thought likely to be used frequently. To provide access to more

sophisticated classical computations, we instead introduce a **extern** mechanism to connect OpenQASM 3 circuits to arbitrary, opaque classical computations.

An **extern** is declared similarly to a function declaration in a C header file. That is, rather than define commonly used simple subroutines in OpenQASM, the programmer may provide a signature of the form:

```
1 extern vote(bit a, bit b, bit c) -> bit;
```

and then use `vote` like a subroutine acting on classical bits. In order to connect to existing classical compiler infrastructure, the definition of an **extern** is not embedded within an OpenQASM circuit, but is expected to be provided to the compilation tool chain in some other language format such as Python, C, x86 assembly, or LLVM IR. This allows externs to be compiled with tools like GCC or LLVM with a minimal amount of additional structure to manage execution within the run-time requirements of a global controller. Importantly, this strategy of connecting to existing classical programming infrastructure implies that the programmer can use existing libraries without porting the underlying code into a new programming language.

Unlike other approaches to interfacing with classical computing, OpenQASM does not have an explicit synchronization primitive such as the `WAIT` instruction in QUIL [3, 73]. In particular, invoking an **extern** function does not imply a synchronization boundary at the call site.⁶ A compiler is free to use dataflow analysis to insert any required target-specific synchronization primitives at the point where outputs from extern call are used. That is, invoking an **extern** function *schedules* a classical computation, but does not wait for that calculation to terminate. The **extern** semantics abstract away the particulars of a vendor-specific **application binary interface (ABI)** for how data are moved to and from the classical function.

4.5 Input and Output Parameters

In addition to the real-time classical compute constructs of OpenQASM 3, we introduce features targeted to support near-time computation, in the form of OpenQASM programs with input and output parameters (described elsewhere as “parameterized circuits” [46]). This functionality is provided through keywords **input** and **output**, which act as modifiers on variable declarations, and allow OpenQASM 3 circuits to represent these variables as accepting input parameters and returning select output parameters.

The **input** modifier can be used to indicate that one or more variable declarations represent values, which will be provided at run-time, upon invocation. This allows the programmer to use the same compiled circuits, which only differ in the values of certain parameters. As OpenQASM 2 did not allow for input parameters or **input** declaration, for compatibility OpenQASM 3 does not require an **input** declaration to be provided: in this case it assumes that there are no input parameters. When an **input** declaration is provided, the compiler produces an executable that leaves these free parameters unspecified: A circuit run would take as input both the executable and some choice of the parameters.

The **output** modifier can be used to indicate that one or more variables are to be provided as an explicit output of the quantum procedure. Note that OpenQASM 2 did not allow the programmer to specify that only a subset of its variables should be returned as output, and so it would return all classical variables (which were all **creg** variables) as output. For compatibility, OpenQASM 3 does not require an **output** declaration to be provided: In this case it assumes that all of the declared variables are to be returned as output. If the programmer provides one or more **output** declarations,

⁶It is possible to achieve non-blocking semantics in QUIL using memory movement instructions to invoke function execution, and then reserving `WAIT` for synchronization upon completion.

then only those variables described as outputs will be returned as an output of the quantum process. A variable may not be marked as both **input** and **output**.

The **input** and **output** modifiers allow the programmer to more easily write variational quantum algorithms: A quantum algorithm with some free parameters, which may be run many times with different parameter values which are determined by a classical optimiser at near-time. Rather than write a circuit which generates a new sequence of operations for each run, OpenQASM 3 allows such circuits to be expressed as a single program with input parameters. This allows the programmer to communicate many different circuits with a single file, which only has to be compiled once, amortizing the cost of compilation across many runs. For an example, we may consider a circuit, which performs a measurement in a basis given by an input parameter:

```
1 input int basis; // 0 = X basis, 1 = Y basis, 2 = Z basis
2 output bit result;
3 qubit q;
4
5 // Some complicated circuit...
6
7 if (basis == 0) h q;
8 else if (basis == 1) rx( $\pi/2$ ) q;
9 result = measure q;
```

For a second example, consider the **Variable Quantum Eigensolver (VQE)** algorithm [66]. In this algorithm, the same circuit is repeated many times using different sets of free parameters to minimize an expectation value. The following is an example, in which there is also more than one input variable:

```
1 input angle param1;
2 input angle param2;
3 qubit q;
4
5 // Build an ansatz using the above free parameters, eg.
6 rx(param1) q;
7 ry(param2) q;
8
9 // Estimate the expectation value and store in an output variable
```

The following Python pseudocode illustrates the differences between using and not using input and output parameters in a quantum program for the case of the VQE. This pseudocode is fictional but demonstrates the near-term execution model: where there is some classical code coordinating the compilation and execution of OpenQASM programs.

When input and output parameters are not used, the program must be recompiled to account for every change of the input variable θ . Compilation is assumed to be expensive and slow so it will be the main bottleneck toward completion of the overall VQE algorithm.

```
for theta in thetas:
    # Create an OpenQASM circuit with  $\theta$  defined
    circuit = substitute_theta(read("circuit.qasm"))

    # The slow compilation step is run on each iteration of the inner loop
    binary = compile_qasm(circuit)
    result = run_program(binary)
```

When input and output parameters circuits are used, we can hoist compilation out of the loop. The only part of the program that is changing is the value of θ . The entire program does not need to be re-compiled each time.


```
# io_circuit.qasm begins with the line "input angle  $\theta$ ;"
circuit = read("io_circuit.qasm")
# The slow compilation step only happens once
binary = compile_qasm(circuit)

for theta in thetas:
    # Each iteration of the inner loop is reduced to only running the circuit
    result = run_program(binary,  $\theta$ =theta)
```

5 CONCEPTS OF THE LANGUAGE: THE PHYSICAL LEVEL

In addition to support for logical-level quantum computations, OpenQASM 3 has added support for quantum computing experiments and platform-dependent tuning of quantum instructions on the physical level. At the physical level, quantum operations are implemented by causing time-varying stimuli to be transmitted to the qubits, and capturing the time-dependent responses. For example, for superconducting transmon platforms this might take the form of shaped microwave pulses transmitted via a number of coaxial cables, and for trapped ions might take the form of modulated laser pulses. There are various tradeoffs involved in optimal selection of the mapping from logical quantum operations to specific pulse implementations. Therefore, there can be benefits in allowing a programmer who has knowledge of both the hardware limitations and the algorithm requirements to influence the physical implementation for key parts of the circuit. For example, choosing a gate implementation with shorter duration but less-accurate calibration might be a good choice for a variational ansatz where the specific unitary implemented is less important than minimizing decoherence. Another example could be choosing a specific ratio of delays for the idle periods of a qubit to implement spin-echo (and more generally, dynamical decoupling) sequences to cancel out errors caused by hardware parameters that are known to drift slowly compared to the timescale of circuit execution.

In this section, we give an overview of the important features of the physical, lower level of OpenQASM 3. For a detailed description we recommend reading the live specification [2].

5.1 Timing and Optimization

A key aspect of expressing code for quantum experiments is the ability to control the timing of gates and pulses. Examples include characterization of decoherence and crosstalk [39], dynamical decoupling [41, 61, 77], dynamically corrected gates [29, 49], and gate parallelism or scheduling [62]. We introduce such timing semantics in OpenQASM 3 to enable such circuits.

Delay Statements and Duration Types: We introduce a **delay** statement, which allows the programmer to specify relative timing of operations. These statements can take the form **delay**[t], where t is a value of type **duration**. We introduce the **duration** type to support such timing instructions: values of type **duration** represent amounts of time measured in seconds (typically also modified by an SI prefix). For example, a qubit's relaxation time (T_1) can be measured by an OpenQASM snippet of the following form:

```
1 duration stride = 1us; // a time duration, specified in SI units
2 int p = 3; // record the 3rd datapoint in the relaxation curve
3 reset q[0];
4 x q[0];
5 delay[p * stride] q[0];
6 c0 = measure q[0]; // survival probability of the excited state
```

Furthermore, any instruction may take a bracketed duration, such as `x[30ns] q[1]` for a gate `x` defined with a **gate** declaration. Instructions with bracketed durations are required to have exactly

that duration (although, see the **stretch** type below). If an instruction is not defined by a **defcal** consistent with the requested duration, then it may not be executable on hardware.

Boxes and Barrier Statements: OpenQASM3 includes features to constrain the reordering of gates, where the timing of those gates might otherwise be changed by the compiler (or the gates removed entirely as a valid optimization on the logical level).

We introduce a **box** statement, to help with scoping the timing of a particular part of the circuit. A boxed sub-circuit is different from a **gate** (see Section 4.1) or **def** subroutine (see Section 4.4), in that it is merely a pointer to a piece of code within the larger scope which contains it. This can be used to signal permissible logical-level optimizations to the compiler: optimizing operations within a **box** definition is permitted, and optimizations that move operations from one side to the other side of a box are permitted, but moving operations either into or out of the box as part of an optimization is forbidden. The compiler can also be given a description of the operation which a **box** definition is meant to realise, allowing it to re-order gates *around* the box. For example, consider a dynamical decoupling sequence inserted in a part of the circuit:

```

1 rx(2*π/12) q;
2 box {
3   delay[ddt] q;
4   x q;
5   delay[ddt] q;
6   x q;
7   delay[ddt] q;
8 }
9 rx(3*π/12) q;
```

By boxing the sequence, we create a box that implements the identity. The compiler is now free to commute a gate past the box by knowing the unitary implemented by the box:

```

1 rx(5*π/12) q;
2 box {
3   delay[ddt] q;
4   x q;
5   delay[ddt] q;
6   x q;
7   delay[ddt] q;
8 }
```

The compiler can thus perform optimizations without interfering with the implementation of the dynamical decoupling sequence.

As with other operations, we may use square brackets to assign a duration to a box: This can be used to put hard constraints on the execution of a particular sub-circuit by requiring it to have the assigned duration. This can be useful in scenarios where the exact duration of a piece of code is unknown (e.g., if it is runtime dependent), but where it would be helpful to impose a duration on it for the purposes of scheduling the larger circuit. For example, if the duration of the parameterized gates `mygate1(a, b)`, `mygate2(a, b)` depend on values of the variables `a` and `b` in a complex way, but an offline calculation has shown that the total will never require more than 150 ns for all valid combinations:

```

1 // some complicated circuit that gives runtime values to a, b
2 box [150ns] {
3   delay[str1] q1; // schedule as late as possible within the box
4   mygate1(a, a+b) q[0], q[1];
5   mygate2(a, a-b) q[1], q[2];
```

```

6   mygate1(a-b, b) q[0], q[1];
7 }

```

In the above example, `str1` is of **stretch** type, to be discussed in the next section, and in this the compiler will resolve it to the proper value to fill out the box to 150ns.

OpenQASM 3 also retains the **barrier** instruction of OpenQASM 2, which prevents gates from being reordered across its source line, either on a specified qubit (or register) `q`, or a sequence of qubits (or registers) `q, r, ...`. For example,

```

1  cx r[0], r[1];
2  h q[0];
3  h p[0];
4  barrier r, q[0];
5  h p[0];
6  cx r[1], r[0];
7  cx r[0], r[1];

```

This will prevent the compiler from performing optimizations, which involve combining or canceling the CNOT gates on the register `r`. However, it does not prevent similar optimizations of the `h p[0]` operations on either side of it. Outside of any **box**, a **barrier** statement may also be invoked without any arguments (as in the examples of classical control flow in Section 4.4), in which case it prevents gates from being reordered across its source line, on *all* qubits.

More generally, **delay** statements also prevent optimizations or re-orderings of gates, which would move a gate acting on some qubit past a **delay** statement on the same qubit. The **barrier** instruction in OpenQASM 3 is similar to a special case of **delay**, with a duration of zero. The difference between **barrier** and **delay[0]** is that the use of **delay** indicates a fully scheduled series of instructions, that any should not be altered by further scheduling, whereas **barrier** merely indicates an ordering constraint, and if necessary a scheduling pass may add finite-duration delays adjacent to **barriers**. In this way, **barrier** statements may be regarded as helping to provide time synchronization within the given scope and across the qubits on which they are applied. However, the constraints implied by **delay** and **barrier** statements do not leak out of any containing **box** declarations, and thus do not prevent commutation with uses of the **box**: They are treated as implementation details of the box, and are invisible to the outside scope.

Stretch Types: By specifying relative timing of operations rather than absolute timing, it is possible to use OpenQASM 3 to provides more flexible timing of operations, which can be helpful in a setting with a variety of calibrated gates with different durations. We introduce a new type **stretch**, representing a duration of time which is resolvable to a concrete duration at compile time once the exact duration of calibrated gates are known. A **stretch** quantity is a value which the compiler attempts to minimize, subject to timing constraints on one or more qubits. (We say that each a **stretch** value has a “natural duration” of 0, though constraints may impose a larger value on any given **stretch** value.) This makes it possible for a programmer/compiler to express gate designs such as evenly spacing gates (e.g., to implement a higher-order echo decoupling sequence), left-aligning a sequence of gates, or applying a gate for the duration of some subcircuit. We are motivated to include this feature by the importance of specifying gate timing and parallelism, in a way that is independent of the precise duration and implementation of gates at the pulse-level description. This increases circuit portability by decoupling circuit timing intent from the underlying pulses, which may change from machine to machine or even from day to day.

Fig. 4. Arbitrary alignment of gates in time using stretchy delays. (a) left-justified alignment intent, and timed instructions after stretches are resolved. (b) alignment intent of a short gate at the one-third point of a long gate, and after stretch resolution.

We refer to the problem of assigning definite values to **stretch** values, as “the stretch problem”. The **stretch** type itself, and our approach to the stretch problem, is inspired by how a similar problems are solved in $\text{\TeX}$ [52] using “glues”, where the compiler spaces out letters, words, and so on, appropriately given a target font. In quantum circuits, we have the additional challenge that timings on different qubits are not independent, e.g., if a two-qubit gate connects the qubits. This means that the choice of timing on one qubit line can have a ripple effect on other qubits. We describe how the stretch problem is solved through multi-objective linear programming in Section 5.3.

A simple use case for stretchable delays are gate alignments. We can use this feature to gain control on how gates are aligned in a circuit, regardless of the latencies of those gates on the machine. For example, consider the following code sample for left-aligning operations on a collection of qubits:

```

1 // left alignment
2 stretch g1, g2, g3;
3 barrier q;
4 cx q[0], q[1];
5 delay[g1] q[0], q[1];
6  $u(\pi/4, 0, \pi/2)$  q[2];
7 delay[g2] q[2];
8 cx q[3], q[4];
9 delay[g3] q[3], q[4];
10 barrier q;
```

The corresponding circuit is illustrated in Figure 4(a) (where the spring symbols on and across qubit wires denote delays with **stretch** parameters). After the circuit is lowered to a particular machine with known gate calibrations and durations, the compiler resolves every stretchy **delay** instruction into one with concrete timing, so that the operations performed on each qubit take the same cumulative amount of time.

We can further control the exact alignment by giving relative weights to the stretchy delays. This method is flexible, e.g., to align a gate at the $1/3$ point of another gate, we can use a **delay**[g]

Fig. 5. Dynamically corrected CNOT gate where the spectator has a rotary pulse. The rotary gates are stretchy, and the design intent is to interleave a “winding” and “unwinding” that is equal to the total duration of the CNOT. We do this without knowledge of the CNOT duration, and the compiler resolves them to the correct duration during lowering to the target backend.

instruction before, and a **delay**[2*g] after the gate to be aligned, as in the following example (illustrated as a circuit in Figure 4(b)):

```

1 // one-third alignment
2 stretch g;
3 barrier q;
4 cx q[0], q[1];
5 delay[g] q[2];
6 u(π/4, 0, π/2) q[2];
7 delay[2*g] q[2];
8 barrier q;

```

The scope of a **stretch** variable can be bounded by a **box** statement. If the **box** statement has been assigned a duration (e.g., to help with scheduling operations in the larger circuit), this imposes a limitation on the total duration of the operations to be performed within the **box**, which in turn, constrains the **stretch** values inside. For example:

```

1 // define a 1ms box whose content is just a centered CNOT
2 box [1ms] {
3   stretch a;
4   delay[a] q;
5   cx q[0], q[1];
6   delay[a] q;
7 }

```

The “natural” duration of instructions within a box (e.g., taking a hypothetical value of 0 for all stretch values) must be smaller than the declared box duration, otherwise a compile-time error will be raised. The **stretch** inside the box will always be set to fill the difference between the declared duration and the natural duration. If the instructions exceed the deadline at runtime, a runtime error is raised.

Instructions other than **delay** can also be “stretchy”, if they are explicitly defined as such. Consider for example a rotation called rotary that is applied for the entire duration of some other gate [75]. Once resolved to a concrete time, it can be passed to the gate definition (in terms of pulses) to realize the gate. This scenario is illustrated schematically in Figure 5 and can be expressed in OpenQASM 3 as:

```

1 // stretchy rotation gate
2 stretch g;
3 box {
4   cx control, target;
5   rotary(amp)[g] spectator; // rotate for some duration, with some pulse amplitude
6   rotary(-amp)[2*g] spectator; // rotate back, for twice that duration
7   rotary(amp)[g] spectator; // rotate forward again, making overall identity
8 }

```

In particular, the duration of a **box** statement can also be given as a **stretch**, in which case its “natural” duration is constrained by the natural durations of its operations, and the duration it takes at run-time depends on the constraints imposed in the larger circuit in which the **box** block occurs. For example,

```

1 // using box to label a sub-circuit, and later delaying for that entire duration
2 stretch mybox;
3 box [mybox] {
4   cx q[0], q[1];
5   delay[200ns] q[0];
6 }
7 delay[mybox] q[2], q[3];
8 cx q[2], q[3];

```

realizes a CNOT operation on qubits $q[0]$ and $q[1]$, followed by a delay of at least 200 ns, followed by a CNOT on qubits $q[2]$ and $q[3]$.

Because stretches must be resolved through solving a linear program, there are limitations on their use: They can appear only in linear combinations of stretch variables and durations, and their numerical value is only available at a late stage in the compilation process. Because **duration** types are known explicitly without solving a linear program, they can be used more freely in nonlinear expressions and control flow. The following example is valid with `pi_time` as **duration** but would be invalid if it were a **stretch**.

```

1 // sequence of 30ns  $\pi$  pulses to fill total duration
2 duration total_time = 1000ns;
3 duration pi_time = 30ns;
4 for _ in [1:floor(total_time/pi_time)] {
5   x[pi_time] q[0];
6 }

```

5.2 Calibrating Quantum Operations

OpenQASM describes circuits applying quantum operations such as unitary gates and projective measurements to qubits. Quantum operations are typically implemented with classical time-dependent signals coupled to quantum systems. Control hardware such as arbitrary waveform generators, flux sources, and lasers emit signals to orchestrate the synchronous emission of calibrated control fields. These guide the quantum system to implement the desired quantum operation [40]. Consequently, control system implementations at the hardware and software level are highly platform-dependent. For example, superconducting transmon qubits encode a qubit in a non-linear oscillator formed by a parallel circuit consisting of a Josephson junction and a capacitor. To manipulate the state of the qubit a series of shaped microwave control pulses are applied to it [53]. The design of these control signals is an active area of research, and developments in this area enable software updates to improve the performance of existing hardware systems [45]. The support for calibration grammars in OpenQASM is motivated by the rapid pace of development within the quantum hardware space and the need for hardware/software co-design to extract the maximum performance from today’s quantum computers. Similar considerations have motivated several software projects to directly support signal-level control of quantum hardware, including QGL [4], Artiq [19], OpenPulse [13, 60], JaqalPaw [56], QUIL-T [7], QCCS [8], Qua [6], and Pulser [72].

OpenQASM has added support for specifying instruction calibrations in the form of **defcal** (short for “define calibration”) declarations, which allow the programmer to specify a microcoded [80] implementation of a **gate**, **measure**, or **reset** instruction within a lower-level control grammar as implemented by the target hardware vendor. This particular language

keyword is borrowed from the recently proposed QUIL-T extension of the Quil language [7]. Support for **defcal** declarations is optional for OpenQASM compilers and target devices. The example below shows how the programmer can select the *openpulse* calibration grammar (introduced later in this section), and declare the calibration for an *rx* instruction.

```

1 // select the defcalgrammar to use
2 defcalgrammar "openpulse";
3
4 // define the rx gate
5 gate rx( $\theta$ ) q {
6     U( $\theta$ ,  $-\pi/2$ ,  $-\pi/2$ ) q;
7 }
8
9 // define the implementation of the rx-gate for physical qubit 0
10 defcal rx(angle[20]  $\theta$ ) $0 { ... }
11
12 rx( $\pi/2$ ) $0; // the defcal is automatically linked at compile time

```

The available calibration instructions are determined by the user's selected **defcalgrammar**. A **defcal** declaration has a similar form to a **gate** declaration, but with some subtle differences. Whereas a **gate** declaration defines a previously undefined gate in terms of other OpenQASM instructions, a **defcal** declares the implementation of an OpenQASM instruction for a target device. In a **gate**, parameters are infinite precision angles in the **gate**, whereas in the **defcal** a precision may be specified, enabling the definition of gate calibrations at the native precision of the target hardware. The **defcal** is defined for physical qubits rather than being the definition of a unitary gate applicable to any qubit. In this case, the **defcal** is for a fixed qubit ($\$0$), i.e., the zeroth physical qubit on the device. (OpenQASM reserves identifiers of the form $\$n$, where n is a non-negative integer for physical qubits of the corresponding integer index in the target device.) We assume that compiler mapping passes assign virtual qubits to particular physical qubits in the device.

Calibration grammars may require the ability to insert top-level configuration information and system setup outside of the **defcal** in order to share that information across multiple **defcals**. To facilitate this, OpenQASM 3 introduces **cal** blocks. Within **cal** blocks the semantics of the selected calibration grammar are valid. Values declared within the **cal** block may only be referenced from other **cal** blocks or **defcal** declarations that capture values from the containing scope. The example below uses a **cal** block to define frequencies, ports, and frames to be used in later **defcal** declarations.

```

1 cal {
2     // defined within `cal`, so it may not leak
3     // back out to the enclosing block's scope
4     float qubit_freq = 5.2e9;
5     // declare an external port (see below for more information)
6     extern port d0;
7     // declare a new frame for port d0
8     frame frd0 = newframe(d0, qubit_freq, 0.0);
9 }

```

Through **defcal** and **cal** declarations, the programmer may define the implementation of canonical instructions for the target hardware. These operation definitions are not typically invoked from a logical-level specification of an operation. Rather, a target-system compiler links the circuit-level OpenQASM circuit to these definitions, as shown in the *rx* definition example above. This embeds an extendable translation layer between the circuit model and control-hardware implementations of quantum instructions within OpenQASM. If a **defcal** does not have a corresponding **gate** declaration the compiler will treat the call to this **defcal** as a call to an opaque circuit operation of the same name and signature. In such a case, the compiler is

not able to reason about the state of the circuit across calls to these **defcal**s. Consequently, this functionality replaces the opaque keyword of OpenQASM 2.

Calibration Resolution. The signals required to enact gates on physical hardware may vary greatly based on the kind of gate, the physical qubit type, and the input gate parameters involved. The role of **defcal** declarations is to specify how a **gate** or primitive operation is performed, with **defcal** dispatch support for specialized definitions for specific physical qubits and gate parameters.

For instance, **defcal** **rx**(**angle**[20] θ) \$q defines the implementation of a rx gate of any angle on any physical qubit \$q, whereas **defcal** **rx**(π) \$0 specifically targets an X rotation of π radians on physical qubit \$0. The reference to physical rather than virtual qubits is critical, as quantum registers are normally not interchangeable from the perspective of the details of the control stimulus that enacts them.

At compile-time, specialized **defcal** declarations are greedily matched with qubit specializations selected before instruction arguments. If not all parameters are specialized, the declaration with the largest number of specialized parameters is matched first. In the event of a tie, the compiler will choose the **defcal** block that appeared first in the circuit.⁷ For instance, given:

```
1 defcal rx(angle[20]  $\theta$ ) $q { ... }
2 defcal rx(angle[20]  $\theta$ ) $0 { ... }
3 defcal rx( $\pi/2$ ) $0 { ... }
```

the operation **rx**($\pi/2$) \$0 would match the **defcal** declaration on line 3; **rx**(π) \$0 would match the declaration on line 2; and **rx**($\pi/2$) \$1 would match the declaration on line 1. Calibrations for measurements are specified similarly, apart from having a classical bit as part of its signature, e.g., **defcal** **measure** \$0 -> **bit**.

To integrate with OpenQASM's timing system when performing scheduling, the compiler requires the set of physical qubit resources used by the **defcal** implementation as well as the duration of its execution. The scheduler uses this information to resolve the scheduled circuit specified by the programmer. For this reason, all resolved **defcal** usages must have compile-time computable durations. If the duration of a **defcal** depends on its input parameters, the parameters and consequently the definition must be resolvable at compile-time to enable deterministic scheduling of the operations at the circuit level. Calibrations must also be independent of when the operation is invoked, to allow their substitution anywhere the corresponding circuit instruction is applied. Note that some instructions, such as a **reset** instruction implemented with measurement feedback, might require control flow. A calibration grammar should support such use cases while still guaranteeing the deterministic duration of the instruction.

Calibration Grammars. Describing the application of control fields (or other operation specifics) for various quantum platforms may be vastly different. For example, a SWAP gate within a superconducting qubit system may be actuated with a time-dependent microwave stimulus applied to the target qubits. In contrast, an ion-trap system may implement a SWAP gate by physically exchanging two ions through the application of a series of DC potentials and laser pulses, which may be global in nature [47, 50]. The control fields, and correspondingly, the semantics required to describe these two operations may vary significantly between hardware vendors.

In OpenQASM we do not attempt to capture all such possibilities, but rather embrace the role of a pluggable multi-level intermediate representation. Vendors are free to define their

⁷This strategy is loosely inspired by Haskell's pattern matching for resolving function calls [57] and a similar strategy appears in Quil-T's **defcal** [7].

calibration grammar and an accompanying implementation. OpenQASM, in turn, provides the `defcalgrammar <grammar>` declaration which allows the programmer to inform the compiler implementation of which calibration grammar should be selected for the following calibration declarations. Calibration support is optional, and OpenQASM compiler implementations that support defining calibrations should provide an interface for new grammars to plugin and extend the compilation infrastructure. It is then the responsibility of the hardware provider to specify a grammar and implement the required hooks into the compiler. This might be as simple as a `defcal` mapping to custom hardware calibrations within the target system, or as complex as a fully-fledged pulse-programming language as explored later in this section.

By allowing vendors to extend the calibration grammar, OpenQASM provides a standardized pipeline for vendors to lower from a hardware-agnostic representation into one that is hardware-aware. It is the responsibility of the hardware vendor to consume the calibrations along with the rest of the OpenQASM and to compile the circuit into a final executable for the target hardware.

OpenPulse Grammar. OpenPulse was initially specified as a RESTful API system that was later extended with an implementation for Qiskit [13, 60]. Within the OpenQASM specification, the OpenPulse grammar provides a hardware-independent representation for programming qubits at the level of microwave pulses [2]. In a typical quantum computing system, control electronics apply stimulus in the form of time-varying pulses to hardware ports. These enact control fields upon the target quantum system. For example, in a superconducting qubit system, smooth pulse envelopes with a carrier frequency on resonance with the superconducting qubit are applied through an arbitrary waveform generator to excite the qubit from its ground state to its excited state. OpenPulse was originally designed to target superconducting qubit systems. In general, it captures the control requirements of quantum systems defined by a time-dependent Hamiltonian of the form $H(t) = H_{\text{sys}} + \sum_k \alpha_k(t)H_k$ where H_{sys} is a time-independent drift Hamiltonian and H_k are control-fields modulated by the time-dependent control knobs $\alpha_k(t)$ described within the OpenPulse program in the form of pulses emitted to hardware ports. This system model describes many quantum technologies such as ion traps, quantum dots, and neutral atoms [21, 72, 81]. For this reason, it is the canonical `defcalgrammar` provided within the OpenQASM3 specification.

In OpenPulse all operations fundamentally operate on a `port` which is a software abstraction representing any input or output component meant to manipulate and observe qubits. As they are uniquely defined by the target system they are declared as `extern port d0;`. Ports are ultimately mapped to some combination of hardware resources. For instance, a port may directly correspond to a digital-to-analog converter. In the quantum system a port typically corresponds to a Hamiltonian term, H_k with pulses applied to the port defining the time-dependent $\alpha_k(t)$ control-term. For instance, a port could apply a time-dependent voltage to a qubit, or manipulate the coupling between two neighboring qubits. Ports are both vendor and device specific.

To emit pulses that track the frame of the qubit OpenPulse introduces the `frame` as the canonical value to which pulses are applied. It is akin to the rotating frame of a Hamiltonian and is responsible for tracking two properties of pulses emitted on a `port`. Firstly, the `frame` represents a carrier signal $e^{i(2\pi f t + \phi)}$, parameterized by a frequency f and phase ϕ . Frames are tracked throughout a circuit execution by the system's runtime. Any modifications to a frame within a calibration entry will persist to calls to the same frame elsewhere in the circuit. In this way, it is analogous to a numerically controlled oscillator (NCO). One motivation for keeping track of accrued phase is to allow the natural implementation of the “virtual Z-gate”. This does not require a physical pulse but rather shifts the phase of all future pulses on that frame. When declared a `frame` is tied to a `port`. Frames may be defined both within the program scope — `frame frd0 = newframe(d0, 5e9, 0.0);` — or as an external linkage to plugin to existing calibration frameworks — `extern frd0;`.

A frame contains two attributes, frequency and phase of types **float** and **angle** respectively. These may be set and modified through dot notation. For example, `frd0.frequency = 4.9e9; frd0.phase +=  $\pi$` ; sets the frequency for a frame and increments the phase of the frame. These updates will occur at the current time of the frame. Multiple frames may be defined for each port effectively allowing multiplexing over a port's IO for multiple carrier signals. Consequently, each frame maintains its pointer to an abstract internal clock which is incremented through operations such as **play** and **barrier**.

The following code sample demonstrates how to define an implementation of an rz gate for the physical qubit \$0.

```
1 defcal rz(angle  $\theta$ ) $0 {
2   driveframe0.phase -=  $\theta$ ;
3 }
```

In this example, the gate is realized by updating the frame tracking the qubits. All subsequent pulses that we apply to the frame will have an updated phase, allowing the implementation of the virtual Z-gate.

The manipulation of a **port** is channeled through the **frame** which provides a *view* on top of the **port**. The primary stimulus applied to a **frame** is the **waveform** type. These may be defined either as an array of complex samples which define the points for the waveform envelope — **waveform** `arb_waveform = [1+0im, 0+1im, 0.5+0.5im]`. Or as an **extern** declared waveform template representing a waveform, which will later be materialized into a list of complex samples, either by the compiler or the hardware — **extern** `gaussian(complex[float] amp, duration d, duration sigma) -> waveform;`. Waveforms are played on a frame through the **play** instruction

```
1 defcal sx $0 {
2   // emit a pulse with a carrier signal, time and
3   // port defined by frd0.
4   play(gaussian(0.1, 160dt, 40dt), frd0);
5 }
```

The **waveform** is therefore emitted with a carrier defined by the frequency and phase of `frd0` at the current time of `frd0` on port `d0`. The frame is blocked during execution of the **play** operation, and its time is incremented by `160dt` after execution completes.

It is also possible to calibrate the **measure** and **reset** operations. Due to the varied nature of readout protocols in quantum hardware we rely on vendors to define the appropriate external subroutines to access the readout chain. For example, an implementation of **measure** is shown below.

```
1 cal {
2   // Capture definition available in the hardware
3   extern capture(duration dur, frame capture_frame) -> bit;
4
5   extern port m0;
6   extern port cap0;
7
8   frame frm0 = newframe(m0, 5e9, 0.);
9   frame frcap0 = newframe(cap0, 5e9, 0.);
10 }
11
12 defcal measure $0 -> bit {
13   // measurement stimulus envelope
14   waveform meas_wf = gaussian_square(0.1, 16000dt, 262dt, 13952dt);
15
16   // play the stimulus
17   play(meas_wf, frm0);
18   // delay the capture frame to account for signal latencies
```

```

19  delay[320dt] frcap0;
20  // capture transmitted data after interaction with measurement resonator
21  return capture(16000dt, frcap0);
22  }

```

OpenPulse's timing model is *relative*. Within a program, there is no explicit reference to a global clock, but instead, only relative references to the starting time of a **defcal/cal** or the current relative time of other frames through the **barrier** instruction. This enables OpenPulse operations to be scheduled in a *position-independent* manner. Each **defcal** or **cal** block defines a relative zero-time, which is defined by when the OpenQASM scheduling layer schedules the execution of the block. When a **frame** is declared it is assigned an absolute time that is equal to the relative zero time for the enclosing block. Every **frame** used within a **cal** or **defcal** starts and ends with an implicit barrier across all used frames, establishing a unified start and end time for all frames within the block. The example below demonstrates the timing behaviour of OpenPulse with the implementation of an echoed cross-resonance gate [70].

```

1  cal {
2    extern port d0;
3    extern frame frd0;
4    extern frame frd1;
5    extern frame frd2;
6  }
7
8
9  defcal ecr $0 $1 {
10   // define and update frame at t=0
11   frame cr1_0 = newframe(d0, frd1.frequency, frd1.phase);
12   // synchronize time of frames at t1=max_time(frd0, frd1, frd2)
13   barrier frd0, frd1, frd2, cr1_0;
14   // play a waveform on frame cr1_0 incrementing time of cr1_0 by dur
15   play(gaussian_square(dur, -amp, sigma, square_width), cr1_0);
16   // synchronize frames before next call at t2=t1+dur
17   barrier frd0, frd1, frd2, cr1_0;
18   // call another defcal
19   x $0;
20   // synchronize frames before next call at t3=t2+duration(x)
21   barrier frd0, frd1, frd2, cr1_0;
22   play(gaussian_square(dur, -amp, sigma, square_width), cr1_0);
23   // synchronize frames before leaving defcal at t4=t3+dur
24   barrier frd0, frd1, frd2, cr1_0;
25  }

```

The example above demonstrates several additional features of the OpenPulse language. The notion and semantics of the **barrier** frd0, frd1, frd2, cr1_0; statement is borrowed from the circuit layer to allow the OpenPulse grammar extension to synchronize the time of all frames to a unique instant. A **defcal** may also call another **defcal** to construct a more sophisticated calibration as shown by the call of `x $0;`. Furthermore, the use of the external frame frd2 enables the calibration grammar to inform the circuit scheduler that qubit \$2 is indirectly used by the gate. The scheduler is therefore made aware that the qubits are indirectly coupled. This allows it to avoid scheduling them simultaneously, such as the case where many qubits are coupled through a resonator bus.

5.3 Solving the Stretch Problem

We now describe how the compiler determines the values of **stretch** values, in contexts involving operations, with other **duration** values, and **box** statements. We treat **duration** values to be variables in a linear system of inequalities that the compiler must solve, subject to certain constraints.

We call this problem the “stretch problem”. If we attempt to execute an OpenQASM circuit, a stretch solving algorithm must be applied to compute explicit durations for all delays.

Our goal is that any OpenQASM circuit with mapped qubits and **defcals** gives rise to a set of stretch problems that, if bounded and feasible, provides explicit timing for all operations within every basic block. The stretch problem is to determine definite values for all **stretch** values in a given scope — subject to constraints on gate ordering, such as those imposed by **delay** and **barrier** statements, and the range of optimization techniques — so that the operations performed on each qubit can be scheduled appropriately, and in particular take the same amount of time. The stretches get resolved to explicit delays in a single late-stage pass that solves the stretch problem, thereby specifying all delays and a complete schedule for the operations. We define the stretch problem in such a way that it is bounded, and has a unique optimum if it is feasible. If the problem is infeasible for a given target, a compiler error is issued, providing some insight about the unsatisfied constraints. In the case of a circuit without any stretchable delays, the stretch problem may be unsolvable in all but the most trivial cases.

The stretch problem is a lexicographic multi-objective linear programming problem [24], as follows: given a matrix A and vectors c_j and b , lexicographically minimize $c_1^\top x, c_2^\top x, \dots, c_m^\top x$ subject to $Ax \leq b$ and $x \geq 0$. This linear programming problem can be solved in polynomial time and efficiently in practice. For example, one may first solve “minimize $c_1^\top x$ subject to $Ax \leq b$ and $x \geq 0$ ” to obtain the cost β_1 , then solve “minimize $c_2^\top x$ subject to $Ax \leq b, c_1^\top x = \beta_1$, and $x \geq 0$ ” to obtain the cost β_2 , and for forth until all objectives are met. This is described in [24] and the references therein; multi-objective optimization is implemented in software such as CPLEX [25].

The stretch problem is formulated independently for each basic block or box B . The total duration of the block is mapped to some variable $T = x_0$, which we wish to minimize. The block B consists of instructions applied to qubits Q . An instruction is any task with a non-negative duration that involves a subset of qubits; the duration of each instruction is either known or unknown, and may be associated with variables x_1, x_2 , and so on. To clarify the presentation of the stretch problem, imagine that each qubit $q \in Q$ is associated with a duration value D_q . The value of D_q is determined by the durations of operations and delays on q , as well as constraints on when multi-qubit blocks can be performed (depending on the scheduling of operations on the that must come before, on each of the qubits involved). This defines D_q a function of the non-negative stretch variables x_i . We then impose the constraint that $x_0 = D_q$ for all $q \in Q$. A feasible solution to this assigns durations to each **stretch** variable, ensuring that for each qubit $q \in Q$, the operations on q occur at some time $t \in [0, T]$. If we lexicographically minimize the variables x_i , we single out a particular solution where, among other things, $T = x_0$ is minimized.

For example, the code below inserts a dynamical decoupling sequence where the *centers* of pulses are equidistant from each other. The circuit is depicted in Figure 6. We specify correct durations for the delays by using backtracking operations to properly take into account the finite duration of each gate.

```

1  stretch s, t, r, xlen, ylen;
2  stretch init = s - .5 * xlen;
3  stretch mid = s - .5 * xlen - .5 * ylen;
4  stretch fin = s - .5 * ylen;
5  box {
6    delay[init] $0;
7    x[xlen] $0;
8    delay[mid] $0;
9    y[ylen] $0;
10   delay[mid] $0;
11   x[xlen] $0;
12   delay[mid] $0;
13   y[ylen] $0;

```

Fig. 6. Dynamical decoupling of a spectator qubit using finite-duration pulses. This design intent can be expressed by defining a single stretch variable s that corresponds to the distance between equidistant gate centers. The other durations which correspond to actual circuit delays are derived by simple arithmetic on durations. Given a target system with calibrated X and Y gates, the solution to the stretch problem can be found.

```

14  delay[fin] $0;
15  delay[t] $1;
16  cx $2, $3;
17  cx $1, $2;
18  u $3;
19  delay[r] $3;
20  }

```

The **stretch** declarations impose the constraints: $s, t, r, x_{len}, y_{len} \geq 0$, $s - \frac{1}{2}x_{len} \geq 0$, $s - \frac{1}{2}x_{len} - \frac{1}{2}y_{len} \geq 0$, and $s - \frac{1}{2}y_{len} \geq 0$. Some additional constraints $x_{len} \geq A$ and $y_{len} \geq B$ are imposed by the target platform, where A and B are the minimum amount of time possible to realize the operations x $\$0$ and y $\$0$ respectively. Let x_{16} , x_{17} , and x_{18} denote the (fixed and not explicitly declared) durations of the operations on lines 16, 17, and 18. Then, we obtain the following durations of operations on qubits:

$$\begin{aligned}
 D_{q[0]} &= 5s, & D_{q[2]} &= x_{16} + x_{17}, \\
 D_{q[1]} &= t + x_{17}, & D_{q[3]} &= x_{16} + x_{18} + r.
 \end{aligned} \tag{7}$$

By setting $D_{q[i]} = T$ for all i , we obtain additional relationships and constraints on the durations above, specifically:

$$T = x_{16} + x_{17}, \quad s = \frac{1}{5}x_{16} + \frac{1}{5}x_{17}, \quad t = x_{16}, \quad r = x_{17} - x_{18}. \tag{8}$$

From the constraints on s and r , this is feasible if $x_{16} + x_{17} \geq \frac{5}{2}A + \frac{5}{2}B$ and $x_{17} \geq x_{18}$. The solution, which is obtained by the stretch problem in this case actually determines fixed values for T , s , t , and r , and sets x_{len} to its minimum value A , and y_{len} to its minimum value B .

Note that the order in which **stretch** variables are declared determine the order in which they are minimized. This can have a significant effect on scheduling. For a second example, consider the following code sample:

```

1  stretch a, b, c, d;
2  delay[a] q[0];
3  delay[b] q[1];
4  h q[0];

```

Fig. 7. Implicit vs. explicit delay. (a) An implicit delay exists on $q[0]$, but it is not part of the circuit description. Thus, this circuit does not care about timing and the RZ gate is free to commute on the top wire. (b) An explicit delay is part of the circuit description. The timing is consistent and can be resolved if and only if this delay is exactly the same duration as RY on $q[1]$. The delay is like a barrier in that it prevents commutation on that wire. However, RZ can still commute before the $CNOT$ if it has duration 0.

```

5 h q[1];
6 delay[c] q[0];
7 delay[d] q[1];

```

Let x_4 and x_5 stand for the (fixed and not explicitly declared) durations of the operations on lines 4 and 5. We have

$$\begin{aligned}
 &\text{lexmin } a, b, c, d \text{ s.t.}, \\
 &a, b, c, d \geq 0, \\
 &a + x_4 + c = b + x_5 + d.
 \end{aligned} \tag{9}$$

From this, optimising first for a , then b , and so forth, we would obtain $a = b = 0$, and either $c = x_5 - x_4$ and $d = 0$ if $x_4 \leq x_5$, or $c = 0$ and $d = x_4 - x_5$ otherwise. This yields a left-aligned schedule. If we had instead declared **stretch** d, c, b, a ; on line 1, we would have optimized the values in the opposite order (minimising first d , then c , etc.) to obtain a right-aligned schedule.

5.4 Multi-level Representation

In previous sections, we introduced the language features of OpenQASM, which can range from relatively high-level constructs such as multi-controlled gates to low-level timing and microcoded pulses. As such, OpenQASM is designed to be a multi-level intermediate representation, where the focus shifts from target-agnostic computation to concrete implementation as more hardware specificity is introduced.

An OpenQASM circuit can also mix different abstraction levels by introducing constraints where needed, but allowing the compiler to make decisions where there are no constraints. Examples of circuit constraints are tying a virtual qubit to a particular physical qubit, or left-aligning some parallel gates that have different durations.

We illustrate this point by considering timing constraints in an OpenQASM circuit. When a **delay** instruction is used, even though it implements the identity channel in the ideal case, it is understood to provide explicit timing. Therefore, an explicit **delay** instruction will prevent commutation of gates that would otherwise commute. For example, in Figure 7(a), there will be an implicit delay between the “ cx ” gates on qubit 0. However, the “ rz ” gate is still free to commute on that qubit, because the delay is implicit. Once, the delay becomes explicit (perhaps at lower stages of compilation), gate commutation is prohibited (Figure 7(b)).

Furthermore, even though **stretch** is used to specify constraints for the solver, its use in the circuit is like any other instruction. This has the benefit that design intents can be specified on a high-level circuit and the intent will be carried with the circuit, until resolved. The following example for simultaneous randomized benchmarking is an instance of this, where alignment is specified on general unitaries of the Clifford group, which may be decomposed into a variety of

Fig. 8. Simultaneous 2-qubit randomized benchmarking using stretchy delays to align gates. (a) Independent sequences of Clifford operations are applied to the top 2 and bottom 2 qubits, but varying decompositions and gate speeds cause misalignment. (b) A barrier can force coarse synchronization points. (c) Stretchy delays inserted before and after every Clifford can ensure better simultaneity. These delays will be carried as part of the circuit, no matter how the Cliffords are decomposed. (d) The decomposition of Cliffords themselves can have stretchy delays, ensuring even finer alignment of gates. The program can be written in a way that these delays override the previous ones in-between Cliffords.

ways. In addition, the stretchiness can be included as part of the definition of a gate, in which case whenever that gate is expanded those stretch constraints will be automatically inserted.

6 COMPILATION PHASES BY EXAMPLE

In this section, we present a full example using an OpenQASM circuit for an iterative phase estimation algorithm [34]. We show how the circuit may get transformed during multiple phases of compilation, and how OpenQASM can be used as the intermediate representation at each phase. This simple example also serves to highlight many of the features of the language: classical control flow, gate modifiers, virtual and physical qubits, timing and stretches, and pulse-defined calibrations.

OpenQASM itself is agnostic to any particular toolchain, so the compilation phases and transformations in this section are merely intended as a representative example of typical steps in circuit compilation.

Iterative phase estimation is an algorithm for calculating the eigenvalue (phase) of a unitary up to some number of bits of precision. In contrast to the textbook phase estimation, it uses fewer qubits (only one control qubit), but at the cost of multiple measurements. Each measurement adds one bit to the estimated phase. Importantly, all iterations must happen in real time during the coherence interval of the qubits, and the result of each measurement feeds forward to the following iteration to influence the angles of rotation.

Initial circuit. The initial circuit may have been produced by high level software tools and languages and may have already passed through stages of high level transformations.

```

1  /*
2  * Iterative phase estimation
3  */
4  OPENQASM 3;
5  include "stdgates.inc";

```

```

6
7  const n = 3;           // number of iterations
8  const  $\theta = 3 * \pi / 8$ ; // phase angle on target qubit
9
10 qubit q; // phase estimation qubit
11 qubit r; // target qubit for the controlled-unitary gate
12 angle[n] c = 0; // phase estimation bits
13
14 // initialize
15 reset q;
16 reset r;
17
18 // prepare uniform superposition of eigenvectors of phase
19 h r;
20
21 // iterative phase estimation loop
22 for k in [1:n] { // implicitly cast val to int
23   reset q;
24   h q;
25   ctrl @ pow(2**k) @ phase( $\theta$ ) q, r;
26   inv @ phase(c) q;
27   h q;
28   measure q -> c[0];
29   // newest measurement outcome is associated to a  $\pi/2$  phase shift
30   // in the next iteration, so shift all bits of c left
31   c <<= 1;
32 }
33
34 // Now c contains the n-bit estimate of  $\varphi$  in the
35 // eigenvalue  $e^{i\varphi}$  and qreg r is projected to an
36 // approximate eigenstate of the phase gate.

```

6.1 Target-independent Compilation Phase

The target-independent phase applies transformations that do not use information about any particular target system. The goal of these transformations is quantum circuit synthesis and optimization.

Constant propagation and folding. This transformation substitutes the values of constants as they occur in expressions throughout the circuit. We repeatedly apply this transformation as constants are exposed.

Gate modifier evaluation and synthesis. The control, power, and inverse gate modifiers are replaced by gate sequences, and those are simplified when possible based on the structure of the circuit. In this case, the inverse and power modifiers can be simplified by modifying the angle arguments of the phase and controlled-phase gates. For this example, we replace the control modifier and phase gate by a cphase gate from the standard library to emphasize the transformation. For further emphasis, we also change the circuit notation, replacing a controlled unitary gate (raised to a power) by a two-qubit gate with two connected control dots and an angle label.

```

1 // simplify gate modifiers and substitute constants
2 OPENQASM 3;
3 include "stdgates.inc";
4 qubit q;

```

```

5  qubit r;
6  angle[3] c = 0;
7  reset q;
8  reset r;
9  h r;
10 for k in [1:3] {
11   reset q;
12   h q;
13   cphase(2**k * 3*π/8) q, r;
14   phase(-c) q;
15   h q;
16   measure q -> c[0];
17   c <=<= 1;
18 }

```

Loop unrolling. The trip count of this loop is constant and statically known. Therefore, the loop can be unrolled. This is not necessary as the machine executing an OpenQASM circuit is assumed to have control flow capabilities. However, doing so in the compiler could expose certain optimization opportunities. For example, the following pass could remove the reset operation in the first iteration, because it is redundant with the initial circuit reset.

```

1  // loop unrolling
2  OPENQASM 3;
3  include "stdgates.inc";
4  qubit q;
5  qubit r;
6  angle[3] c = 0;
7  reset q;
8  reset r;
9  h r;
10 int k = 1;
11 h q;
12 cphase(2**k * 3*π/8) q, r;
13 phase(-c) q;
14 h q;
15 measure q -> c[0];
16 c <=<= 1;
17 k = 2;
18 reset q;
19 h q;
20 cphase(2**k * 3*π/8) q, r;
21 phase(-c) q;
22 h q;
23 measure q -> c[0];
24 c <=<= 1;
25 k = 3;
26 reset q;
27 h q;
28 cphase(2**k * 3*π/8) q, r;
29 phase(-c) q;
30 h q;
31 measure q -> c[0];
32 c <=<= 1;

```

Constant propagation and folding (repeat). We repeat this step to evaluate and substitute the value of the power variable.

Gate simplification. Gate simplification rules can be applied to reduce the gate count. In this example, rotations with angle 0 are identity gates and can be removed.

```

1 // constant propagation and simplification
2 OPENQASM 3;
3 include "stdgates.inc";
4 qubit q;
5 qubit r;
6 angle[3] c = 0;
7 reset q;
8 reset r;
9 h q;
10 h r;
11 cphase(3*π/8) q, r;
12 h q;
13 measure q -> c[0];
14 c <=< 1;
15 reset q;
16 h q;
17 cphase(6*π/8) q, r;
18 phase(-c) q;
19 h q;
20 measure q -> c[0];
21 c <=< 1;
22 reset q;
23 h q;
24 cphase(12*π/8) q, r;
25 phase(-c) q;
26 h q;
27 measure q -> c[0];
28 c <=< 1;

```

6.2 Target-dependent Compilation Phase

The target-dependent phase applies transformations that are generically necessary to lower a target-independent OpenQASM 3 circuit to an executable. Each transformation uses information that is specific to a class of target machines, rather than one specific target, and the class of potential targets is refined with each transformation.

Basis translation. As we enter the next phase of compilation, we need to use information about the target system. The basis translation step rewrites quantum gates in terms of a set of gates

available on the target. In this example, we assume a quantum computer with a calibrated set of basis gates including phase, h, and cx.

```

1 // basis translation
2 OPENQASM 3;
3 gate phase( $\lambda$ ) q { U( $\theta$ ,  $\theta$ ,  $\lambda$ ) q; }
4 gate cx c, t { ctrl @ U( $\pi$ ,  $\theta$ ,  $\pi$ ) c, t; }
5 gate h a { U( $\pi/2$ ,  $\theta$ ,  $\pi$ ) a; }
6 qubit q;
7 qubit r;
8 angle[3] c = 0;
9 reset q;
10 reset r;
11 h r;
12 h q;
13 phase( $3*\pi/16$ ) q;
14 cx q, r;
15 phase( $-3*\pi/16$ ) r;
16 cx q, r;
17 phase( $3*\pi/16$ ) r;
18 h q;
19 measure q -> c[0];
20 c <<= 1;
21 reset q;
22 h q;
23 phase( $3*\pi/8$ ) q;
24 cx q, r;
25 phase( $-3*\pi/8$ ) r;
26 cx q, r;
27 phase( $3*\pi/8$ ) r;
28 phase( $-c$ ) q;
29 h q;
30 measure q -> c[0];
31 c <<= 1;
32 reset q;
33 h q;
34 phase( $6*\pi/8$ ) q;
35 cx q, r;
36 phase( $-6*\pi/8$ ) r;
37 cx q, r;
38 phase( $6*\pi/8$ ) r;
39 phase( $-c$ ) q;
40 h q;
41 measure q -> c[0];
42 c <<= 1;

```

Gate simplification (repeat). We repeat the gate simplification transformations to reduce the gate count. In this example, phase gates commute through the controls of CNOT gates and can be merged with other phase gates by adding the angle arguments.

```

1 // gate commutation and merging
2 OPENQASM 3;
3 gate phase( $\lambda$ ) q { U( $\theta$ ,  $\theta$ ,  $\lambda$ ) q; }
4 gate cx c, t { ctrl @ U( $\pi$ ,  $\theta$ ,  $\pi$ ) c, t; }
5 gate h a { U( $\pi/2$ ,  $\theta$ ,  $\pi$ ) a; }
6 qubit q;
7 qubit r;

```



```

8  angle[3] c = 0;
9  reset q;
10 reset r;
11 h r;
12 h q;
13 cx q, r;
14 phase(-3*π/16) r;
15 cx q, r;
16 phase(3*π/16) r;
17 phase(3*π/16) q;
18 h q;
19 measure q -> c[0];
20 c <= 1;
21 reset q;
22 h q;
23 cx q, r;
24 phase(-3*π/8) r;
25 cx q, r;
26 phase(3*π/8) r;
27 phase(3*π/8-c) q;
28 h q;
29 measure q -> c[0];
30 c <= 1;
31 reset q;
32 h q;
33 cx q, r;
34 phase(-6*π/8) r;
35 cx q, r;
36 phase(6*π/8) r;
37 phase(6*π/8-c) q;
38 h q;
39 measure q -> c[0];
40 c <= 1;

```

Physical qubit mapping and routing. In this step, we need to use the physical connectivity of the target device to map virtual qubits of the circuit to physical qubits of the device. In this example, we assume a CNOT can be applied between physical qubits 0 and 1 in either direction, so no routing is necessary.

```

1 // layout on physical qubits
2 OPENQASM 3;
3 gate phase(λ) q { U(0, 0, λ) q; }
4 gate cx c, t { ctrl @ U(π, 0, π) c, t; }
5 gate h a { U(π/2, 0, π) a; }
6 angle[3] c = 0;
7 reset $0;
8 reset $1;
9 h $1;
10 h $0;
11 cx $0, $1;
12 phase(1.8125*π) $1; // mod 2*π
13 cx $0, $1;
14 phase(0.1875*π) $1;
15 phase(0.1875*π) $0;
16 h $0;
17 measure $0 -> c[0];

```

```

18 c <=<= 1;
19 reset $0;
20 h $0;
21 cx $0, $1;
22 phase(1.625* $\pi$ ) $1; // mod  $2\pi$ 
23 cx $0, $1;
24 phase(0.375* $\pi$ ) $1;
25 angle[32] temp_1 = 0.375* $\pi$ ;
26 temp_1 -= c; // cast and do arithmetic mod  $2\pi$ 
27 phase(temp_1) $0;
28 h $0;
29 measure $0 -> c[0];
30 c <=<= 1;
31 reset $0;
32 h $0;
33 cx $0, $1;
34 phase(1.25* $\pi$ ) $1; // mod  $2\pi$ 
35 cx $0, $1;
36 phase(0.75* $\pi$ ) $1;
37 angle[32] temp_2 = 0.75* $\pi$ ;
38 temp_2 -= c; // cast and do arithmetic mod  $2\pi$ 
39 phase(temp_2) $0;
40 h $0;
41 measure $0 -> c[0];
42 c <=<= 1;

```

Scheduling. We can enforce a scheduling policy without knowing or using concrete gate durations. We use stretchy delays for this purpose, which enact a “as late as possible” schedule. The actual gate durations will be available in the next stage where the circuit is tied to certain gate calibration data, and the stretch problem can be solved.

```

1 // scheduling intent without hardcoded timing
2 OPENQASM 3;
3 gate phase( $\lambda$ ) q { U( $\theta$ , 0,  $\lambda$ ) q; }
4 gate cx c, t { ctrl @ U( $\pi$ , 0,  $\pi$ ) c, t; }
5 gate h a { U( $\pi/2$ , 0,  $\pi$ ) a; }
6 angle[3] c = 0;
7 reset $0;
8 reset $1;
9 h $1;
10 h $0;
11 stretch s1, s2;
12 delay[s1] $0;
13 delay[s2] $1;
14 cx $0, $1;
15 phase(1.8125* $\pi$ ) $1;
16 cx $0, $1;
17 phase(0.1875* $\pi$ ) $1;
18 phase(0.1875* $\pi$ ) $0;
19 h $0;
20 measure $0 -> c[0];
21 c <=<= 1;
22 reset $0;
23 h $0;
24 stretch s3;
25 delay[s3] $1;
26 cx $0, $1;
27 phase(1.625* $\pi$ ) $1; // mod  $2\pi$ 
28 cx $0, $1;

```

```

29 phase(0.375* $\pi$ ) $1;
30 angle[32] temp_1 = 0.375* $\pi$ ;
31 temp_1 -= c; // cast and do arithmetic mod 2  $\pi$ 
32 phase(temp_1) $0;
33 h $0;
34 measure $0 -> c[0];
35 c <=< 1;
36 reset $0;
37 h $0;
38 stretch s4;
39 delay[s4] $1;
40 cx $0, $1;
41 phase(1.25* $\pi$ ) $1; // mod 2* $\pi$ 
42 cx $0, $1;
43 phase(0.75* $\pi$ ) $1;
44 angle[32] temp_2 = 0.75* $\pi$ ;
45 temp_2 -= c; // cast and do arithmetic mod 2  $\pi$ 
46 phase(temp_2) $0;
47 h $0;
48 measure $0 -> c[0];
49 c <=< 1;
50 stretch s5;
51 delay[s5] $1;

```

Calibration linking and stretch resolution. A pulse sequence is defined for each gate, reset, and measurement. Using their durations, the stretch variables are resolved into delays with concrete timing. Angles are rounded to the precision of the defcal arguments (i.e., to defcal precision) at this step, if they have not already been rounded to the appropriate precision by an earlier transformation.

```

1 // read calibration and enforce full schedule
2 OPENQASM 3;
3 defcalgrammar "openpulse";
4
5 cal {
6   extern gaussian_square(complex[float] amp, duration d, duration square_width, duration sigma) -> waveform;
7   extern drag(complex[float] amp, duration d, duration sigma, float beta) -> waveform;
8   extern capture(duration dur, frame capture_frame) -> bit;
9
10  extern port d0;
11  extern port d1;
12  extern port m0;
13  extern port cap0;
14
15  const float q0_freq = 5.0e9;
16  const float q1_freq = 5.1e9;
17  const float q0_mfreq = 4.1e9;
18
19  frame frd0 = newframe(d0, q0_freq, 0.0);
20  frame frd1 = newframe(d1, q1_freq, 0.0);
21  frame frm0 = newframe(m0, q0_mfreq, 0.0);
22  frame frcap0 = newframe(cap0, q0_mfreq, 0.0);
23 }
24
25 defcal y90p $0 {
26   play(drag(0.1, 160dt, 40dt, 0.05), frd0);
27 }
28

```

```

29 defcal y90p $1 {
30     play(drag(0.2, 160dt, 40dt, 0.1), frd1);
31 }
32
33 defcal cr90p $0, $1 {
34     frame cr0_1 = newframe(d0, frd0.frequency, frd0.phase);
35     play(gaussian_square(0.3, 480dt, 160dt, 40dt), cr0_1);
36 }
37
38 defcal phase(angle[20]  $\theta$ ) $0 {
39     frd0.phase -=  $\theta$ ;
40 }
41
42 defcal phase(angle[20]  $\theta$ ) $1 {
43     frd1.phase -=  $\theta$ ;
44 }
45
46 defcal cr90m $0, $1 {
47     phase( $-\pi$ ) $1;
48     cr90p $0, $1;
49     phase( $\pi$ ) $1;
50 }
51
52 defcal x90p $q {
53     phase( $\pi$ ) $q;
54     y90p $q;
55     phase( $-\pi$ ) $q;
56 }
57
58 defcal xp $q {
59     x90p $q;
60     x90p $q;
61 }
62
63 defcal h $q {
64     phase( $\pi$ ) $q;
65     y90p $q;
66 }
67
68 defcal cx $control $target {
69     phase( $-\pi/2$ ) $control;
70     xp $control;
71     x90p $target;
72     cr90p $control, $target;
73     xp $control;
74     cr90m $control, $target;
75 }
76
77 defcal measure $0 -> bit {
78     play(gaussian_square(0.1, 1600dt, 160dt, 40dt), frm0);
79     return capture(1600dt, frcap0);
80 }
81
82 angle[3] c = 0;
83 reset $0;
84 reset $1;
85 h $1; // 30 ns
86 h $0; // 20 ns
87 delay[10ns] $0;
88 cx $0, $1; // 300 ns
89 phase(1.8125* $\pi$ ) $1; // 0 ns
90 cx $0, $1;
91 phase(0.1875* $\pi$ ) $1;
92 phase(0.1875* $\pi$ ) $0; // 0 ns
93 h $0;
94 measure $0 -> c[0]; // 400 ns
95 delay[420ns] $1;
96 c <<= 1;
97 reset $0; // 200 ns
98 h $0;

```

```

99  delay[220ns] $1;
100 cx $0, $1;
101 phase(1.625* $\pi$ ) $1; // mod 2* $\pi$ 
102 cx $0, $1;
103 phase(0.375* $\pi$ ) $1;
104 angle[32] temp_1 = 0.375* $\pi$ ;
105 temp_1 -= c; // cast and do arithmetic mod 2  $\pi$ 
106 phase(temp_1) $0;
107 h $0;
108 measure $0 -> c[0];
109 delay[420ns] $1;
110 c <= 1;
111 reset $0;
112 h $0;
113 delay[220ns] $1;
114 cx $0, $1;
115 phase(1.25* $\pi$ ) $1; // mod 2* $\pi$ 
116 cx $0, $1;
117 phase(0.75* $\pi$ ) $1;
118 angle[32] temp_2 = 0.75* $\pi$ ;
119 temp_2 -= c; // cast and do arithmetic mod 2  $\pi$ 
120 phase(temp_2) $0;
121 h $0;
122 measure $0 -> c[0];
123 delay[420ns] $1;
124 c <= 1;

```

The compiled circuit is now submitted to the target machine code generator to produce binaries for the target control system of the quantum computer. These are then submitted to the execution engine to orchestrate the quantum computation. Additional phases of target-dependent compilation occur that are beyond the scope of OpenQASM 3.

7 RELATED WORK AND FUTURE EXTENSIONS

The design of OpenQASM 3 was naturally influenced by features present in other quantum programming languages, as well as constructions, which are standard concepts in the quantum computing literature and which were considered feasible to include. A number of other features or design choices were also considered. In some cases, these were set aside in the interests of some other design choice; in other cases, they were considered to be possibly worth pursuing but not a core functionality. In this section, we describe how some of the “new” features in OpenQASM 3 (i.e., which extend what existed in OpenQASM 2) relate to the features of other quantum programming languages, which choices were considered but not included in OpenQASM 3, and which choices are still under consideration for inclusion.

7.1 Comparison to Other Quantum Programming Languages

There has been significant prior work on quantum programming languages and intermediate representations, some of which have directly influenced OpenQASM. Many high-level quantum programming languages are embedded in classical languages. For example, Scaffold [42], QCL [64], and qcor [59] are embedded in C or C++, Quipper [11] in Haskell, ProjectQ [9] in Python, and QIR [5] in LLVM. On the other hand, languages such as OpenQASM and Quil [73] are designed to be standalone and portable intermediate representations.

Classical control flow exists in many of the above languages. These are sometimes only used as a means of code compression (i.e., fully unrollable at compile time) [9, 42], and sometimes are intended to mix classical and quantum compute [10, 73]. In embedded languages, the host language provides direct access to classical types and control flow. OpenQASM 3 expands upon these by introducing new types such as **angle**, and allowing the semantic definition of purely classical **extern** functions in the context of real-time computation.

Gate modifiers (at least for “control”), exist in most of the above programming languages. In OpenQASM 3, we utilize these modifiers for two reasons. First, they decouple semantics from implementation — many implementations could exist for a particular modified gate. Second, modifiers in conjunction with two simple language built-ins (**U** gate and **gphase**), generate the entire standard gate library.

The use of timing and pulse-level control within quantum circuits has long been present in software and languages used for controlling experiments. Rigorous specifications for this were introduced for example in QGL [4], OpenPulse [60], eQASM [37], Artiq [19], Pulser [72], Jaqal [54], Qua [6], and Quil-t [7]. OpenQASM 3 takes this further by streamlining the incorporation of timing in the circuit model, by allowing timing constraints to be expressed at a level that is decoupled from actual pulse implementations (using **stretch**). This further motivated the inclusion of a scoping mechanism in OpenQASM 3 in the form of **box**, which allows references to specific blocks of code and to guide compiler optimizations. This has been addressed before in a more limited form using barriers [27] and pragmas [73].

7.2 Features Considered but Not Adopted

In this article, we have presented the design decisions in OpenQASM 3 and the rationale behind them. In a similar vein, multiple other language features were contemplated but abandoned or left for future revisions of the language. In the following, we briefly touch upon these.

Some languages such as Q# and ProjectQ support resource and memory management, in the form of mid-circuit allocation and de-allocation of qubits. Some languages such as Scaffold allow for the expression of classical code that will then be compiled to reversible quantum circuits (e.g.,

oracles). In each of these cases, OpenQASM intentionally remains lower-level and more grounded in executable circuits, aim at keeping the compiler's complexity more manageable.

Conversely, some languages such as Quil have unstructured control flow in the form of explicit program labels and JUMP statements, which are in fact more in the style of a classical assembly language. In order to retain backwards compatibility with OpenQASM 2, we chose not to require any *explicit* entry points, preferring to allow the program to proceed from the first instruction in a global scope. OpenQASM 3 also adopted the use of `for` and `while` loops to express control flow, both to retain easier programmability by humans, and to allow for easier timing analysis by the compiler. These all motivated a design philosophy that it should not be necessary in OpenQASM 3 to manage the (classical) control flow, on a similarly low level to management of quantum operations and resources.

7.3 Considering Further Language Features for OpenQASM

The OpenQASM language will continue to evolve; however, within its governance model, and certain features may be added. An example of such a feature are unitary circuit families and generic subroutines, corresponding roughly to templated function definitions, which may be defined entirely in OpenQASM 3. (While such subroutines could be written in a separate host language, writing them in OpenQASM 3 allows for greater portability.) Another language feature, which is at an advanced stage of consideration at the time of writing it the inclusion of classical arrays as a container type; A third feature, which has been set aside, but has been re-proposed for consideration is a fixed-point real number type. The usefulness of such features, the way that they interoperate or conflict with other features, and how they may be incorporated into the language and supported as features is the subject of discussion in working groups established for this purpose. In this way, OpenQASM 3 will continue to evolve, incorporating features, which are sufficiently important and consistent with the other functionality which it provides, through input from the community.

8 CONCLUSION

OpenQASM 3 has introduced language features that both expand and deepen the scope quantum circuits that can be described with particular focus on their physical implementation and the interactions between classical and quantum computing. We have endeavored throughout to give a consistent “look and feel” to these features so that regardless of abstraction level, OpenQASM 3 still feels like one language. This manuscript has illustrated the primary new features introduced in OpenQASM 3 and put those features in context with examples. In particular, we have shown features and representations appropriate to a variety of abstraction levels and use cases, from high-level gate modifiers to low-level microcoded gate implementations.

Language design is an open-ended problem and we expect that as we attempt to use OpenQASM 3 to program circuits on our primitive quantum computers we will discover many awkward constructions, missing features, or incomplete specifications in the language. We fully expect the language to continue to evolve over time driven by real-world usage and hardware development. In particular, there are already proposals under consideration to modify implicit type conversions or enable code re-use through generic functions. Consequently, the formal language specification is posted as a live document [2] within the Qiskit project. We will be forming a formal process for reviewing proposed changes so that OpenQASM might continue to adapt to the needs of its users.

ACKNOWLEDGMENTS

We thank the community who gave early feedback on the OpenQASM 3 spec [2], especially Hossein Ajallooiean, Luciano Bello, Yudong Cao, Lauren Capelluto, Pranav Gokhale, Michael Healey, Hiroshi Horii, Sonika Johri, Peter Karalekas, Moritz Kirste, Kevin Krsulich, Andrew Landahl,

Prakash Murali, Salva de la Puente, Kenneth Rudinger, Zachary Schoenfeld, Yunong Shi, Robert Smith, Stefan Teleman, Ntwali Bashinge Toussaint, and Jack Woehr.

REFERENCES

- [1] [n.d.]. IBM Quantum Experience. Retrieved from <https://quantum-computing.ibm.com/>. accessed November 2020.
- [2] [n.d.]. OpenQASM 3.x Live Specification. Retrieved from <https://qiskit.github.io/openqasm/>. accessed February 2021.
- [3] [n.d.]. PyQuil. Retrieved from <https://github.com/rigetti/openqasm/pyquil>. accessed November 2020.
- [4] [n.d.]. Quantum Gate Language (QGL). Retrieved from <https://github.com/BBN-Q/QGL>. accessed November 2021.
- [5] [n.d.]. Quantum Intermediate Representation (QIR). Retrieved from <https://github.com/microsoft/qsharp-language/tree/main/Specifications/QIR>. accessed December 2021.
- [6] [n.d.]. Quantum Machines Qua. Retrieved from <https://www.quantum-machines.co/blog/keep-your-finger-on-the-pulse-with-qua-a-pulse-level-quantum-programming-language/>. accessed November 2021.
- [7] [n.d.]. Quil analog control RFC. Retrieved from <https://github.com/quil-lang/quil/blob/master/rfcs/analog/proposal.md>. accessed November 2021.
- [8] [n.d.]. Zurich Instruments QCCS. Retrieved from <https://github.com/zhinst/zhinst-toolkit>. accessed November 2021.
- [9] accessed November 2020. ProjectQ. <https://projectq.ch>.
- [10] accessed November 2020. Q#. Retrieved from <https://github.com/microsoft/qsharp-language>.
- [11] accessed November 2020. The Quipper Language. Retrieved from <http://www.mathstat.dal.ca/~selinger/quipper/>.
- [12] D. Aharonov and M. Ben-Or. 2008. Fault-tolerant quantum computation with constant error rate. *SIAM Journal on Computing* 38, 4 (2008), 1207–1282.
- [13] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. McKay. 2020. Qiskit Pulse: Programming quantum computers through the cloud with pulses. *Quantum Science and Technology* 5, 044006 (2020).
- [14] M. Amy. 2019. Sized types for low-level quantum metaprogramming. In *Proceedings of the Lecture Notes in Computer Science, Vol 11497. Reversible Computation*. RC 2019, Springer.
- [15] S. Balensiefer, L. Kreger-Stickles, and M. Oskin. 2005. QUALE: Quantum architecture layout evaluator. In *Proceedings of the Quantum Information and Computation III* 103 (2005).
- [16] A. Barenco, C. Bennett, R. Cleve, D. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. 1995. Elementary gates for quantum computation. *Physical Review A* 52, 5 (1995), 3457.
- [17] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. 1993. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters* 70, 13 (Mar 1993), 1895–1899. DOI: <https://doi.org/10.1103/PhysRevLett.70.1895>
- [18] Lev Bishop, Steven Heide, Bettina Heim, Ali Javadi, Blake Johnson, and Philipp Schindler. [n.d.]. Retrieved March 16, 2022 from <https://medium.com/qiskit/introducing-a-technical-steering-committee-for-openqasm3-f9db808108e1>.
- [19] Sébastien Bourdauducq, whitequark, Robert Jordens, Yann Sionneau, enjoy-digital, cjbe, David Nadlinger, hartytp, JBoulder, Daniel Slichter, Drew, mntng, r-srinivas, apatura-iris, Stewart Mackenzie, Z. Smith, Pawel K, Marius Weber, kemstevens, Felix Held, and David Leibbrandt. 2018. *m-labs/artiq: 4.0*. DOI: <https://doi.org/10.5281/zenodo.1492176>
- [20] S. Bravyi and A. Kitaev. 2005. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A* 71, 2 (2005), 022316.
- [21] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. 2019. Trapped-Ion quantum computing: Progress and challenges. *Applied Physics Reviews* 6, 2 (June 2019), 021314. DOI: <https://doi.org/10.1063/1.5088164>
- [22] A. Butko et al. 2020. Understanding Quantum Control Processor Capabilities and Limitations through Circuit Characterization. In *2020 International Conference on Rebooting Computing (ICRC)*. 66–75. DOI: [10.1109/ICRC2020.2020.00011](https://doi.org/10.1109/ICRC2020.2020.00011)
- [23] I. Chuang. 2005, accessed November 2020. qasm2circ. Retrieved from <http://www.media.mit.edu/quanta/qasm2circ/>.
- [24] M. Cococcione, M. Pappalardo, and Y. D. Sergeyev. 2018. Lexicographic multi-objective linear programming using grossone methodology: Theory and algorithm. *Applied Mathematics and Computation* 318, C (2018), 298–311.
- [25] IBM ILOG Cplex. 2009. V12. 1: User’s manual for CPLEX. *International Business Machines Corporation* 46, 53 (2009), 157.
- [26] A. W. Cross. 2005. qasm-tools. Retrieved from <http://www.media.mit.edu/quanta/quanta-web/projects/qasm-tools/>. accessed November 2020.
- [27] A. W. Cross, L. Bishop, J. Smolin, and J. Gambetta. 2017. Open quantum assembly language. arXiv:1707.03429. Retrieved from <https://arxiv.org/abs/1707.03429>.
- [28] P. Das, C. A. Pattison, S. Manne, D. Carmean, K. Svore, M. Qureshi, and N. Delfosse. 2020. A scalable decoder micro-architecture for fault-tolerant quantum computing. arXiv:2001.06598. Retrieved from <https://arxiv.org/abs/2001.06598>.
- [29] A. De and L. P. Pryadko. 2013. Universal set of scalable dynamically corrected gates for quantum error correction with always-on qubit couplings. *Physical Review Letters* 110, 7 (2013), 070503.
- [30] N. de Beaudrap. 2019. ReQASM: A recursive extension to OpenQASM. *Private Communication* (2019).

- [31] N. Delfosse. 2020. Hierarchical decoding to reduce hardware requirements for quantum computing. arXiv:2001.11427. Retrieved from <https://arxiv.org/abs/2001.11427>.
- [32] B. Deseif. 2020. yquant: Typesetting quantum circuits in a human-readable language. arXiv:2007.12931. Retrieved from <https://arxiv.org/abs/2007.12931>.
- [33] D. E. Deutsch. 1989. Quantum computational networks. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 425, 1868 (1989), 73–90.
- [34] M. Dobšiček, G. Johansson, V. Shumeiko, and G. Wendin. 2007. Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark. *Physical Review A* 76, 3 (2007), 030306.
- [35] M. Dousti, A. Shafaei, and M. Pedram. 2016. Squash 2: A hierarchical scalable quantum mapper considering ancilla sharing. *Quantum Information & and Computation* 16, 4 (2016), 332–356.
- [36] A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg. 2012. Towards practical classical processing for the surface code. *Physical Review Letters* 108, 180501 (2012), 180501.
- [37] X. Fu, L. Rieseboos, M. A. Rol, J. van Straten, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, V. Newsum, K. K. L. Loh, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels. 2019. eQASM: An executable quantum instruction set architecture. In *Proceedings of the 25th International Symposium on High-Performance Computer Architecture (HPCA19)* (2019).
- [38] X. Fu, M. A. Rol, C. C. Bultink, J. Van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels. 2017. An experimental microarchitecture for a superconducting quantum processor. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*. 813–825.
- [39] J. M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, J. A. Smolin, J. M. Chow, C. A. Ryan, C. Rigetti, S. Poletto, T. A. Ohki, Mark B. Ketchen, and M. Steffen. 2012. Characterization of addressability by simultaneous randomized benchmarking. *Physical Review Letters* 109, 24 (2012), 240504.
- [40] S. J. Glaser, U. Boscaín, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, et al. 2015. Training Schrödinger’s cat: Quantum optimal control. *The European Physical Journal D* 69, 12 (2015), 1–24.
- [41] E. L. Hahn. 1950. Spin echoes. *Physical Review* 80, 4 (1950), 580.
- [42] A. Javadi-Abhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. Chong, and M. Martonosi. 2014. ScaffCC: A framework for compilation and analysis of quantum computing programs. In *Proceedings of the ACM International Conference on Computing Frontiers (CF 2014)* (2014).
- [43] W. M. Johnston, J. R. Paul Hanna, and R. J. Millar. 2004. Advances in dataflow programming languages. *Computing Surveys* 36, 1 (2004), 1–34.
- [44] R. Josza. 2005. An introduction to measurement based quantum computation. arXiv:0508124 [quant-ph].
- [45] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. Borgorin, M. Brink, L. Capelluto, O. Gunluk, T. Itoko, N. Kanazawa, and A. Kandala. 2021. Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Science and Technology* 6, 2 (2021), 025020.
- [46] P. J. Karalekas, N. A. Tezak, E. C. Peterson, C. A. Ryan, M. P. da Silva, and R. S. Smith. 2020. A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Science and Technology* 5, 2 (Apr 2020), 024003. DOI : <https://doi.org/10.1088/2058-9565/ab7559>
- [47] Henning Kaufmann, Thomas Ruster, Christian T. Schmiegelow, Marcelo A. Luda, Vidyut Kaushal, Jonas Schulz, David von Lindenfels, Ferdinand Schmidt-Kaler, and Ulrich G. Poschinger. 2017. Fast ion swapping for quantum-information processing. *Physical Review A* 95, 5 (2017), 052319.
- [48] N. Khammassi, G. G. Guerreschi, I. Ashraf, J. W. Hogaboam, C. G. Almudever, and K. Bertels. 2018. cQASM v1.0: Towards a Common Quantum Assembly Language. arXiv:1805.09607. Retrieved from <https://arxiv.org/abs/1805.09607>.
- [49] K. Khodjasteh and L. Viola. 2009. Dynamically error-corrected gates for universal quantum computation. *Physical Review Letters* 102, 8 (2009), 080501.
- [50] D. Kielpinski, C. Monroe, and D. J. Wineland. 2002. Architecture for a large-scale ion-trap quantum computer. *Nature* 417, 6890 (2002), 709–711.
- [51] A. Kitaev. 1995. Quantum measurements and the abelian stabilizer problem. arXiv:quant-ph/9511026 (1995).
- [52] D. E. Knuth and D. Bibby. 1984. *The TeXbook*, Vol. 3. Addison-Wesley Reading.
- [53] J. Koch, M. Y. Terri, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. 2007. Charge-insensitive qubit design derived from the Cooper pair box. *Physical Review A* 76, 4 (2007), 042319.
- [54] Andrew J. Landahl, Daniel S. Lobser, Benjamin C. A. Morrison, Kenneth M. Rudinger, Antonio E. Russo, Jay W. Van Der Wall, and Peter Maunz. 2020. Jaqal, the quantum assembly language for QSCOUT. (2020), 1–14. arXiv:2003.09382. Retrieved from <http://arxiv.org/abs/2003.09382>.

- [55] A. J. Landahl, D. S. Lobser, B. C. A. Morrison, K. M. Rudinger, A. E. Russo, J. W. Van Der Wall, and P. Maunz. 2020. Jaqal, the quantum assembly language for QSCOUT. arXiv:2003.09382. Retrieved from <https://arxiv.org/abs/2003.09382>.
- [56] Daniel Lobser, Joshua Goldberg, Andrew J. Landahl, Peter Maunz, et al. 2021. JaqalPaw: A guide to defining pulses and waveforms for jaqal. Available online at the Sandia QSCOUT project website <https://www.sandia.gov/quantum/Projects/QSCOUT.html>. (2021). https://www.sandia.gov/quantum/Projects/Uploads/JaqalPaw__A_Guide_to_Defining_Pulses_and_Waveforms_for_Jaqal.pdf.
- [57] S. Marlow. 2010. Haskell 2010 language report. <https://www.haskell.org/onlinereport/haskell2010/>.
- [58] A. McCaskey, D. Lyakh, E. Dumitrescu, S. Powers, and T. Humble. 2020. XACC: A system-level software infrastructure for heterogeneous quantum-classical computing. *Quantum Science and Technology* 5, 2 (2020), 024002.
- [59] Alexander Mccaskey, Thien Nguyen, Anthony Santana, Daniel Claudino, Tyler Kharazi, and Hal Finkel. 2021. Extending C++ for heterogeneous quantum-classical computing. *ACM Transactions on Quantum Computing* 2, 2 (2021), 1–36. DOI: <https://doi.org/10.1145/3462670> arXiv:2010.03935.
- [60] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. Egger, S. Filipp, J. Gomez, M. Hush, A. Javadi-Abhari, D. Moreda, P. Nation, B. Paulovicks, E. Winston, C. J. Wood, J. Wootton, and J. M. Gambetta. 2018. Qiskit backend specifications for OpenQASM and OpenPulse experiments. arXiv:1809.03452. Retrieved from <https://arxiv.org/abs/1809.03452>.
- [61] S. Meiboom and D. Gill. 1958. Modified spin-echo method for measuring nuclear relaxation times. *Review of Scientific Instruments* 29, 8 (1958), 688–691.
- [62] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari. 2020. Software mitigation of crosstalk on noisy intermediate-scale quantum computers. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems*. 1001–1016.
- [63] M. A. Nielsen and I. L. Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge Univ. Press.
- [64] B. Omer. 2003. Structured quantum programming. Vienna University of Technology, Ph. D. Thesis (2003). <http://tph.tuwien.ac.at/~oemer/doc/structqprog.pdf>.
- [65] A. Paetznick and K. M. Svore. 2014. Repeat-Until-Success: Non-deterministic decomposition of single-qubit unitaries. *Quantum Information & Computation* 14, 15–16 (2014), 1277–1301.
- [66] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* 5, 1 (2014), 4213.
- [67] R. Raussendorf and H. J. Briegel. 2001. A one-way quantum computer. *Physical Review Letters* 86, 22 (May 2001), 5188–5191. DOI: <https://doi.org/10.1103/PhysRevLett.86.5188>
- [68] C. A. Ryan, B. R. Johnson, D. Ristè, B. Donovan, and T. A. Ohki. 2017. Hardware for dynamic quantum computing. *Review of Scientific Instruments* 88, 10 (2017), 104703.
- [69] B. Schmitt, M. Soeken, G. De Micheli, and A. Mishchenko. 2019. Scaling-up ESOP synthesis for quantum compilation. In *Proceedings of the 2019 IEEE 49th International Symposium on Multiple-Valued Logic*. IEEE, 13–18.
- [70] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta. 2016. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Physical Review A* 93, 6 (2016), 060302.
- [71] P. W. Shor. 1994. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE, 124–134.
- [72] Henrique Silvério, Sebastián Grijalva, Constantin Dalyac, Lucas Leclerc, Peter J. Karalekas, Nathan Shammah, Mourad Beji, Louis-Paul Henry, and Loic Henriët. 2021. Pulser: An open-source package for the design of pulse sequences in programmable neutral-atom arrays. arXiv:2104.15044. Retrieved from <https://arxiv.org/abs/2104.15044>.
- [73] R. Smith, M. Curtis, and W. Zeng. 2016. A practical quantum instruction set architecture. arXiv:1608.03355. Retrieved from <https://arxiv.org/abs/1608.03355>.
- [74] A. M. Steane. 1999. Space, time, parallelism and noise requirements for reliable quantum computing. *Fortschritte der Physik* 46, (4–5) (1999), 443–457.
- [75] N. Sundaresan, I. Lauer, E. Pritchett, E. Magesan, P. Jurcevic, and J. M. Gambetta. 2020. Reducing unitary and spectator errors in cross resonance with optimized rotary echoes. *PRX Quantum* 1, 2 (2020), 020318.
- [76] K. M. Svore, A. W. Cross, I. L. Chuang, A. V. Aho, and I. L. Markov. 2006. A layered software architecture for quantum computing design tools. *Computer* 39, 1 (2006), 74–83.
- [77] G. Uhrig. 2007. Keeping a quantum bit alive by optimized π -pulse sequences. *Physical Review Letters* 98, 10 (2007), 100504.
- [78] L. Viola, E. Knill, and S. Lloyd. 1999. Dynamical decoupling of open quantum systems. *Physical Review Letters* 82, 12 (1999), 2417.
- [79] N. Whitmeyer and J. Chen. 2020. Measurement based quantum computing. Retrieved from https://mwhitmeyer.github.io/pdfs/Measurement_Based_QC_Project.pdf.

- [80] M. V. Wilkes. 1989. The best way to design an automatic calculating machine. In *Proceedings of the Early British Computer Conferences*. MIT Press, Cambridge, MA, 182–184.
- [81] Floris A. Zwanenburg, Andrew S. Dzurak, Andrea Morello, Michelle Y. Simmons, Lloyd C. L. Hollenberg, Gerhard Klimeck, Sven Rogge, Susan N. Coppersmith, and Mark A. Eriksson. 2013. Silicon quantum electronics. *Reviews of Modern Physics* 85, 3 (July 2013), 961–1019. DOI: <https://doi.org/10.1103/RevModPhys.85.961>

Received April 2021; revised December 2021; accepted December 2021