

# XACC: A System-Level Software Infrastructure for Heterogeneous Quantum-Classical Computing\*

Alexander J. McCaskey,<sup>1, 2, †</sup> Dmitry I. Lyakh,<sup>1, 3</sup> Eugene F. Dumitrescu,<sup>1, 4</sup> Sarah S. Powers,<sup>1, 2</sup> and Travis S. Humble<sup>1, 4</sup>

<sup>1</sup>*Quantum Computing Institute, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*

<sup>2</sup>*Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*

<sup>3</sup>*National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*

<sup>4</sup>*Computational Sciences and Engineering Division,  
Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*

Quantum programming techniques and software have advanced significantly over the past five years, with a majority focusing on high-level language frameworks targeting remote REST library APIs. As quantum computing architectures advance and become more widely available, lower-level, system software infrastructures will be needed to enable tighter, co-processor programming and access models. Here we present XACC, a system-level software infrastructure for quantum-classical computing that promotes a service-oriented architecture to expose interfaces for core quantum programming, compilation, and execution tasks. We detail XACC’s interfaces, their interactions, and its implementation as a hardware-agnostic framework for both near-term and future quantum-classical architectures. We provide concrete examples demonstrating the utility of this framework with paradigmatic tasks. Our approach lays the foundation for the development of compilers, associated runtimes, and low-level system tools tightly integrating quantum and classical workflows.

## I. INTRODUCTION

The near-term availability of noisy quantum processing units (QPUs) has enabled a number of proof-of-principle demonstrations of quantum co-processing for existing domain computational science [1–5]. These demonstrations take advantage of sophisticated software frameworks enabling programming, compilation, and execution of quantum algorithms on remotely hosted QPUs. The majority of these efforts have thus far focused on the implementation of high-level, interpreted languages that enable quantum circuit composition, transpilation, and vendor-specific back-end execution [6–9]. These programming frameworks are well suited for near-term experimentation on remotely hosted quantum resources, as they enable quick prototyping and rely solely on existing remote communication libraries.

As QPUs scale, system designs are expected to more tightly integrate CPU and QPU interactions [10, 11]. A quantum-accelerated computing model should adopt best practices and mirror the design of modern heterogeneous high-performance computing, which offers the advantage of improved performance through more re-

fined control over execution scheduling and memory management. Along with tighter integration, the use of interpreted languages for controlling remote access execution must be replaced by lower-level, system software infrastructures, tools, and compilers for composing complex quantum-classical workflows. System-level software is necessary for robust mechanisms allowing quantum device driver execution as part of a tightly integrated co-processor programming model. Single-source quantum-classical compilers and tools, as well as methods for benchmarking, profiling, and debugging, will also be important for enabling domain computational science on these next-generation heterogeneous systems. Initial models for quantum system infrastructure can be realized using system-level languages such as C and C++ that have served as the foundation for high-performance application development on conventional hardware. Native languages may also expose bindings for higher-level languages, like Python, to enable experimentation and productivity.

To achieve these goals, we present the XACC system-level software framework. XACC, which stands for eXtreme-scale Accelerator programming framework, provides an open-source [12], low-level programming framework for the development of quantum-accelerated programs. The framework implements a hardware-agnostic approach to device integration that allows programmers to target multiple concrete hardware back-ends. XACC is designed around a familiar co-processor programming model that leverages an efficient, extensible, and modular service-oriented architecture. This is accomplished using a plug-and-play capability for a holistic quantum programming, compilation, and execution workflow. XACC source code can be found at <https://github.com/eclipse/xacc>.

In earlier work, we presented initial prototype designs

---

\* This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. (<http://energy.gov/downloads/doe-public-access-plan>).

† [mccaskeyaj@ornl.gov](mailto:mccaskeyaj@ornl.gov)

for XACC and its associated programming model [13]. Here we seek to provide a comprehensive and robust description of all aspects of the framework. We describe key design goals, high-level architectural considerations, intrinsic interface designs and specifications, and concrete implementations of these core interfaces. Our presentation of the XACC programming framework is organized as follows: in Sec. II we summarize the salient features of XACC. In Sec. III we detail the underlying service-oriented architecture and implementations. We describe the language interfaces for XACC in Sec. IV. We provide example demonstrations in Sec. V and offer conclusions in Sec. VI.

## II. FEATURES OF XACC

The purpose of the XACC framework is to provide a holistic, low-level software infrastructure that enables cross-platform programming, compilation, and execution of hybrid quantum-classical scientific applications. The framework adheres to the following design goals:

- 1) *Provide a quantum co-processor programming model*  
The XACC infrastructure manages QPUs as accelerators or co-processors to conventional (classical) computing systems. This design assumes that traditional scientific use cases will leverage QPUs by off-loading select classically intractable kernels to the available quantum resource and later process the results to inform remaining classical computations. We design the XACC framework to provide an extensible quantum accelerator back-end that defines common operations and communication protocols with a general, abstract quantum computer.
- 2) *Provide low-level system software interfaces* XACC exposes interfaces for the low-level system control of the CPU and QPU. These interfaces support a diversity of infrastructure tools, approaches, compilers, and languages for quantum-accelerated computing. The XACC framework is implemented as a C++ infrastructure in order to integrate existing quantum programming methods. Starting with a low-level, ubiquitous language like C++ enables extension to higher-level languages, such as Python, Julia, etc. through well-supported language binding libraries. C++ is also well-known for its portability, performance, and efficiency for conventional computing systems, and it is highly suitable for future hybrid heterogeneous systems with tighter quantum-classical integration. We therefore use C++ to lay the ground work for tighter integration models that move past current remote-process-invocation protocols to in-process, in-memory device driver access.
- 3) *Support cross-platform operation* The XACC framework is designed to be agnostic with respect to the

QPU hardware. Users are able to compose and compile applications for multiple quantum computing platforms using a single, common language. This feature supports portability as well as comparative analyses of commonly defined benchmarks. Since source cross-platform operation will prove critical in determining which hardware is best suited for distinct applications.

- 4) *Support modular and extensible functionality* The XACC framework supports a service-oriented, or plugin, architecture that enables modular and extensible functionality. Such a service-oriented software architecture is useful in integrating hardware back-ends, high-level programming approaches, intermediate-level quantum compilation, and error mitigation strategies. A plugin architecture directly enables researchers and programmers to efficiently swap out key aspects of the overall programming, compilation, and execution workflow with problem-specific implementations.

These features distinguish XACC from other programming frameworks. XACC is the only quantum-classical programming framework that provides a system-level software infrastructure permitting closer CPU-QPU integration. Here, system-level implies that the infrastructure is amenable for integration with low-level operating system calls and existing parallel runtime environments (e.g. classical heterogeneous programming models and architectures). Furthermore, a system-level infrastructure must provide appropriate abstractions for memory allocation (quantum and classical) as well as execution integration with provided device driver APIs. This separates XACC from existing quantum language efforts which sit at the application-level and enable quantum code expression without runtime memory allocation and execution. As such, we expect XACC to directly influence future QPU integration efforts for existing high-performance scientific computing codes written in Fortran, C, and C++. Moreover, as it is not a vendor-supplied framework, XACC further differentiates itself by promoting device interoperability as a core feature and design goal. By keeping the core XACC design abstract and extensible, it is also able to target different quantum computing models: gate-based and annealing.

## III. FRAMEWORK ARCHITECTURE

The architecture of the XACC programming framework exposes a series of interfaces for hardware-agnostic quantum programming, compilation, and execution. As shown in Fig. 1, the XACC framework uses a layered software architecture based upon the common compiler decomposition of a front-end, middle-end, and back-end layers. Each layer exposes different extension points, or interfaces, for implementing a variety of specific use cases. XACC currently includes interfaces for quantum language compilers, instructions, hardware devices (also

referred to as accelerators), compiler optimizations (alternatively, passes), observable measurements, and algorithms.

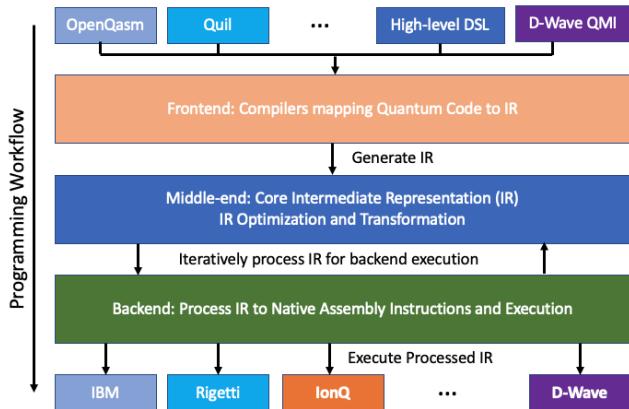


FIG. 1: XACC is composed of three high-level layers - the front-end, middle-end, and back-end. The overall workflow starts with quantum kernel programming at the frontend, followed by IR generation and processing, and ends with back-end execution. Each of these layers exposes a variety of critical extension points.

The front-end maps quantum code expressions, i.e. source code, into an XACC intermediate representation (IR) suitable for transformations and analysis by the middle layer. More specifically, the front-end layer maps source strings expressing a *quantum kernel* to an IR instance. For example, the XACC front-end may be extended to parse an input source code written in the IBM OpenQASM [14] dialect and enable translation into the XACC IR. Similar functionality is implemented for the Rigetti Quil [8] dialect and higher-level languages. As a framework, XACC provides an extensible quantum code transpilation or compilation interface that may be tailored to parse many different languages into a common IR for subsequent manipulation.

The middle layer exposes an IR object model while maintaining an application programming interface (API) for quantum programs that is agnostic to the underlying hardware. The structure of the IR provides an integration mechanism generalizing disparate programming approaches with multiple quantum devices. A critical middle layer concept is the IR transformation, which exposes a framework extension point for performing standard quantum compilation routines, including circuit synthesis, as well as the addition of error mitigation and correction techniques into the logical design.

The back-end layer exposes an abstract quantum computer interface that accepts instances of the IR and executes them on a targeted hardware device (see Sec. III E). The back-end layer provides further extension points mapping the IR to hardware-specific instruction sets and the low-level controls used to execute kernels. The former is performed by taking advantage of the middle-end's IR

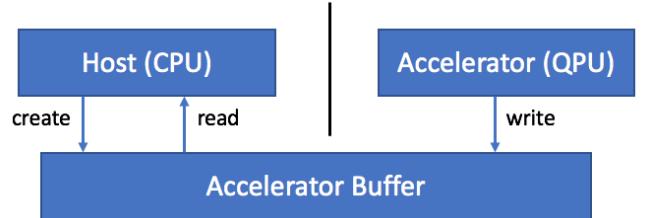


FIG. 2: The Accelerator Buffer serves as a memory object for addressing a quantum register and provides a mechanism for posting and retrieving execution results and associated metadata. XACC manages the interactions between the host system and the QPU to track the state of the Accelerator Buffer for both remote and local execution models.

transformation infrastructure. By using a common interface, quantum program execution via the back-end layer is easily extensible to new hardware.

With the layered architecture, XACC efficiently maps high-level programming languages to low-level hardware instruction sets. Where a direct mapping of  $N$  languages into  $M$  quantum devices would require  $NM$  separate language-to-hardware mappings, XACC reduces the amount of work necessary to  $N + M$  separate mapping implementations. The IR is the central construct connecting the front-end and back-end layers through a set of transformations.

### A. Accelerator Buffer

XACC models general quantum programs as operations on an allocated register of qubits. Specifically, XACC puts forward an *accelerator buffer* abstraction that models the underlying quantum register, or buffer, on which quantum operations are performed. As a memory buffer, it is allocated at run-time by the programmer to define the size of the register and to store results corresponding to measurement of the individual register elements. Subsequent execution of the resulting program will generate measurement results and metadata that are assigned to that buffer instance. XACC manages the reference to accelerator buffers throughout program execution and gathers the measurement results for assignment to the buffer. Since programmers allocate buffers, and buffers are operated on by some execution back-end, the buffer concept spans the front-end, middle-end, and back-end layers. Associated metadata tracks which qubits are operated upon as well as the aggregation of quantum execution measurement results.

To model the buffer abstraction, XACC defines an **AcceleratorBuffer** class that programmers must instantiate, containing at least the number of register elements required by the program. This buffer is passed by reference to the XACC accelerator execution back-end, which provides the execution infrastruc-

AcceleratorBuffer
<ul style="list-style-type: none"> <li>◊ <code>counts: map&lt;string,int&gt;</code></li> <li>◊ <code>nBits: int</code></li> <li>◊ <code>children : vector&lt;AcceleratorBuffer*&gt;</code></li> <li>◊ <code>metadata: map&lt;string, Variant&lt;...&gt;&gt;</code></li> </ul> <ul style="list-style-type: none"> <li>● <code>size() : int</code></li> <li>● <code>appendChild(name : string, child:AcceleratorBuffer*)</code></li> <li>● <code>getChildren() : vector&lt;AcceleratorBuffer*&gt;</code></li> <li>● <code>getInformation() : map&lt;string, Variant&lt;...&gt;&gt;</code></li> <li>● <code>getInformation(name: string): Variant&lt;...&gt;</code></li> <li>● <code>appendMeasurement(m:string, count:int)</code></li> <li>● <code>getExpectationValueZ() : double</code></li> <li>● <code>getMeasurementCounts():map&lt;string,int&gt;</code></li> <li>● <code>print(stream: ostream)</code></li> <li>● <code>load(stream: istream)</code></li> </ul>

FIG. 3: `AcceleratorBuffer` provides a data structure for addressing a quantum register and storing measurement results. The class provides a mechanism for posting and retrieving execution results and associated metadata.

ture (see Sec. III E). Methods and members of the `AcceleratorBuffer` base class are shown in Fig. 3.

An `AcceleratorBuffer` tracks the mapping of measured bit strings to the number of times they were observed (the `counts` map in Fig. 3), as well as an associated heterogeneous map (described in Sec. III I) of metadata. The buffer provides measurement results to the programmer by reference as shown in Fig. 2. The metadata maps string keys to a variant value type, which includes typical C++ types like integer, floating point, string, and vector data structures. The metadata member is intended to store execution-pertinent information that may be different across available quantum computing back-ends. For example, it may be populated by the back-end infrastructure with information about device properties such as noise, as well as variational parameters corresponding to the execution and expectation values.

The `AcceleratorBuffer` class tracks a list of children buffers and thereby enables a tree of execution data. This feature is well suited for variational methods, in which an initial buffer is appended as the subsequent executions are updated based on circuit refinement. New children buffers are appended to the global root buffer with each iteration, such that each child owns metadata and measurement results describing the iteration.

The code snippet in Fig. 4 shows a buffer being allocated, passed to the back-end execution infrastructure, and its results accessed.

## B. Kernels and Compilers

The XACC programming model stipulates that quantum code be described as standard C-like functions, called *quantum kernels*. Kernels must adhere to the the following requirements:

- *Annotation* - XACC quantum kernels must be annotated with the `__qpu__` function attribute to enable static, ahead-of-time compilation.
- *Unique Name* - Like standard C/C++ functions, kernels must have a unique function name.
- *AcceleratorBuffer* - Kernels must take an `AcceleratorBuffer` as their first argument.
- *Parameters* - Kernels can take any number of arguments after the `AcceleratorBuffer`.
- *Language* - The kernel function body must be written in a valid language, i.e. there must exists an XACC Compiler implementation for that language.

```

1 // User allocates buffer
2 // keeps reference to it throughout execution
3 auto buffer = xacc::qalloc(3);
4 // Execute some circuit / composite instruction
5 // this writes result data to the buffer
6 accelerator->execute(buffer, circuit);
7 // User still has that buffer, get results
8 auto results = buffer->getMeasurementCounts();
9
10 // Can add/get metadata, measurements,
11 // and expectation values
12 std::vector<double> fidelities =
13     buffer->getInformation("1q-gate-fidelities")
14     .as<std::vector<double>>();
15 auto counts = buffer->getMeasurementCounts();
16 auto expVal = buffer->getExpectationValueZ();

```

FIG. 4: Example of allocating a three qubit `AcceleratorBuffer` and executing a compiled circuit which persists results and metadata to the buffer. Note the `qalloc()` utility function for allocating quantum memory (`AcceleratorBuffer`).

```

1 // Annotated
2 // AcceleratorBuffer first
3 // parameters after
4 __qpu__ kernel(AcceleratorBuffer q, double x) {
5     // Function body written in
6     // some 'known' language, here xasm
7     X(q[0]);
8     Ry(q[1], x);
9     CX(q[1], q[0]);
10    Measure(q[0]);
11 }

```

FIG. 5: Prototypical XACC Quantum Kernel with function body written in the default XACC assembly language (XASM, see Sec. IV)

The final requirement introduces another abstraction that XACC exposes - the `Compiler` interface. `Compilers` take kernel source strings and map them to the XACC intermediate representation. This compile step takes the source string itself and the back-end accelerator that the programmer is compiling to, thereby enabling connectivity information, noise models, etc. at compile time. `Compilers` can be implemented for any number of languages with varying levels of language abstraction (low-level assembly, high-level domain specific languages, etc.), but must be able to parse the argument structure and function body to a valid instance of the IR. Since `Compilers` create IR from language source strings, they are also in position to map IR back to the their specific source language. XACC `Compilers` expose IR translation capabilities that enable this, thereby providing quantum source-to-source translation (for instance, mapping Quil to OpenQasm).

Using the base XACC API, programmers construct quantum kernel source code as standard strings or string literals, then get reference to the correct `Compiler` implementation for the source string, and compile it to an instance of the XACC IR. This is demonstrated in the code snippet in Fig. 6.

### C. Intermediate Representation

The XACC intermediate representation provides the architectural glue that connects the front-end programmer to back-end native assembly execution. The IR is a polymorphic data model residing at a slightly higher-level of abstraction than typical quantum assembly (QASM) and provides a manipulable, in-memory representation of quantum programs that is amenable to transformation and optimization. Moreover, a standard IR efficiently integrates multiple languages with various hardware instruction sets. The XACC IR architecture is shown in Fig. 8 and consists of three primary interfaces: `Instruction`, `CompositeInstruction`, and `IR`.

The IR begins with the definition of a general QASM-level instruction. The XACC `Instruction` interface describes general assembly instructions that have a unique name, operate on a set of qubits, are parameterized by concrete or variable parameters, and can be enabled/disabled. This interface is implemented for a set of instructions that are common in the digital gate or annealing models of quantum computation. Parameterized instructions are enabled via a variant type called `InstructionParameter`, which can be any of `int`, `double`, or `string` types. This allows instructions to depend on runtime parameters, providing a compile-once-and-update approach for iterative algorithms.

To compose instructions, XACC defines a `CompositeInstruction` interface which inherits from `Instruction`, but also contains a list of children instructions, modeling the familiar composite, or tree design pattern [15]. `CompositeInstructions`

```

1  auto qpu = xacc::getAccelerator("ibm:back-end");
2  auto quil = xacc::getCompiler("quil");
3  auto ir = quil->compile(R"
4    __qpu__ ansatz(AcceleratorBuffer q, double x) {
5      X 0
6      Ry(x) 1
7      CX 1 0
8    }
9    __qpu__ X0X1(AcceleratorBuffer q, double x) {
10      ansatz(q, x);
11      H 0
12      H 1
13      MEASURE 0 [0]
14      MEASURE 1 [1]
15    }
16  )", qpu);
17  auto x0x1 = ir->getComposite("X0X1");
18  // Translate quil kernel to openqasm...
19  auto openqasm = xacc::getCompiler("openqasm");
20  auto oqasmSrc = openqasm->translate(x0x1);

```

FIG. 6: Compiling Quil kernels for the IBM back-end.

```

1  auto xasm = xacc::getCompiler("xasm");
2  auto uccsd = xasm->compile(R"
3    __qpu__ void ansatz(qbit q,
4      std::vector<double> x) {
5      uccsd(q, x);
6    }")->getComposite("ansatz");
7  // Persist to string here -> nothing
8  auto emptyString = uccsd->toString();
9  // Dynamically generate the circuit
10 uccsd->expand({std::make_pair("nq", 4),
11   std::make_pair("ne", 2)});
12 // Persist to string here and you should see
13 // a UCCSD circuit on 4 qubits for 2 electrons
14 auto circuitString = uccsd->toString();

```

FIG. 7: Demonstration of using the `CompositeInstruction::expand()` method for complex circuit generation.

are therefore *n*-ary trees where nodes are `CompositeInstructions` and leaves are concrete `Instructions`. `CompositeInstructions` expose methods for adding, removing, and replacing children instructions, as well as keeping track of any variables that children instructions may depend on. These variables are represented as string `InstructionParameter` instances. For example `Ry(t0)` on qubit 0 depends on the `InstructionParameter` `t0`, thereby requiring any `CompositeInstruction` containing it to know about `t0`.

A `CompositeInstruction` can be dynamic in that its children may not be known until runtime. To accommodate this, `CompositeInstruction` exposes an `expand()` method which takes a heterogeneous

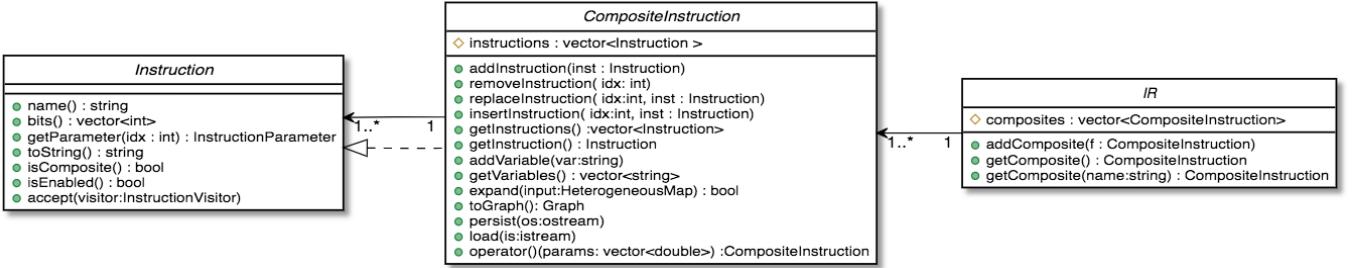


FIG. 8: The XACC IR architecture, composed of `Instruction`, `CompositeInstruction`, and `IR` interfaces, enables a polymorphic object-model for the description of quantum programs that spans available quantum back-ends.

mapping of key-value pairs that seed the generation of children instructions. Implementations of this method should use the input map to direct the generation of concrete `Instructions` and call `CompositeInstruction::addInstruction` to grow the `CompositeInstruction` dynamically. A prototypical example of this would be in programming recursive circuits such as the quantum Fourier transform, or a chemistry unitary coupled cluster circuit based on the number of qubits and electrons in the system (example usage shown in Fig. 7, and for example source code see [16]). Also, `CompositeInstructions` expose a `toGraph()` method call mapping the list of children instructions to a corresponding directed acyclic graph representation, useful in compilation and optimization routines. Finally, the `CompositeInstruction` can be evaluated at a set of concrete floating point parameters, each corresponding to one of the `CompositeInstruction`'s exposed variables.

To store and collect `CompositeInstructions`, XACC defines an `IR` interface which essentially aggregates constructed `CompositeInstructions`. In this way, the `IR` may be thought of as modeling a forest of *n-ary* trees, where each tree can reference others in the same `IR` context. Both `IR` and `CompositeInstructions` can be persisted to a human-readable JSON string, as well as exposing capabilities to load itself from those persisted JSON strings. This enables caching of compilation results for reuse in later compilation contexts. As seen in Fig. 6, `IR` also exposes an API for querying or getting reference to compiled `CompositeInstructions`.

Complementing the `IR` interface, XACC defines an `IRTransformation` interface, which exposes a general `apply()` method taking a `CompositeInstruction` as input for modification. This interface, and implementations of it, are critical for quantum compilation routines that seek to optimize quantum programs with regards to available back-end resources, enable error mitigation routines, or ensuring that quantum program connectivity is amenable for execution on the specific back-end connectivity.

To illustrate the `Instruction` and `CompositeInstruction` concepts, see the quantum kernel construction in Fig. 6. The Quil compiler implementation in this example will parse the input source

string, instantiate concrete low-level instructions for each line, and add them to a `CompositeInstruction`. Here each concrete instruction is provided to XACC as a subclass of the `Instruction` interface. The process begins by creating an `X` instruction, and setting its `bits` vector to `{0}`. Next the compiler creates a parameterized `Ry` gate, with a string `InstructionParameter x`. Finally a `CNOT` instruction is created with its `bits` as `{1,0}`. These will be added to a `CompositeInstruction` (with `variables = {x}`) and stored in the to-be-returned `IR` instance. Once a `CompositeInstruction` is created and stored, it can be referenced by other `CompositeInstructions`. This is how the second kernel is able to make reference to `ansatz` and add a couple of the `X` basis measurements.

A common pattern accompanying a tree of similar types is the visitor pattern, which enables type-specific operations to be added to a tree of common types at runtime via double dispatch. Essentially, the visitor pattern provides a mechanism for separating operations on a set of common objects from the structure of the objects themselves. XACC provides an implementation of the visitor pattern for `Instruction`, and in doing so provides an extensible mechanism for operations on the `IR`. Specifically, all XACC Instructions expose an `accept()` method that takes as input an `InstructionVisitor`. These visitors define a set of `visit` operations for each concrete `Instruction` subtype it would like to visit, and a null implementation

```

1 auto provider = xacc::getIRProvider("quantum");
2 auto kernel =
3     provider->createComposite("foo", {"theta"});
4 auto x = provider->createInstruction("X", {0});
5 auto ry = provider->createInstruction("Ry",
6                                         {1}, {"theta"});
7 auto cx =
8     provider->createInstruction("CX", {1,0});
9 kernel->addInstructions({x, ry, cx});

```

FIG. 9: Leveraging the `IRProvider` to create components of the `IR` without knowledge of underlying types.

for all those it does not implement. Each `Instruction` subclass implements the `accept` call which invokes `visit` on the input `InstructionVisitor` giving itself as the argument to the `visit` call, thereby giving the visitor type information about the `Instruction`. This pattern enables a number of operations on the IR tree that is pertinent to quantum compilation and execution. First, specific back-ends can implement the `InstructionVisitor` and its desired `visit` methods to map IR to the corresponding native assembly. Moreover, simulation back-ends can implement the visitor to execute specific simulation methods, e.g. performing unitary operations on a state vector simulator. The extensibility of this pattern gives XACC developers maximal flexibility to design operations on the IR tree.

Finally, the XACC IR package exposes a factory for the general creation of `Instructions`, `Composites`, and `IR`. The `IRProvider` interface exposes `createInstruction`, `createComposite`, and `createIR` methods that create subtypes of these interfaces. The code snippet in Fig. 9 demonstrates how the `IRProvider` decouples programmers from underlying concrete type information for instructions, composite instructions, and IR.

#### D. Observable and Observable Transform

The observable concept and transformations on observables have been proposed in the QCOR language specification [17]. In part, this specification describes a quantum observable as a rule over a range of qubits defining the measurement bases to be appended after an input circuit. It stipulates that implementations of the observable concept expose an `observe()` method which takes an unmeasured circuit or program and returns a list of (partially) measured circuits based on the observable's algebraic structure. The specification leaves the exact description of the circuit or kernel data structures to specification implementors. Moreover, the specification mandates that the observable concept expose methods for constructing the observable from both a string-like representation and a heterogeneous map of input parameters. Transformations on observables simply map one observable to another, and implementations are free

```

1 // Observable from string
2 auto x0x1 = xacc::getObservable("pauli");
3 x0x1->fromString("X0 X1");
4
5 // Measure, adds hadamards on 0,1
6 // and measure instructions
7 auto measured_circuit =
8     x0x1->observe(circuit)[0];
9
10 // Observable from options
11 auto h2 = xacc::getObservable("chemistry");
12 h2->fromOptions({
13     {"basis", "sto-3g"}, 
14     {"geometry", R"(
15         2
16             H 0.0 0.0 0.0
17             H 0.0 0.0 0.75
18     )"
19 });

```

FIG. 11: Demonstration of how programmers may create XACC `Observables`. Observables can be defined through appropriate subclasses of the `Observable`. Here we show general Pauli `Observables`, as well as a more complex chemistry `observable`.

to propose any number of transformations (like Jordan-Wigner, Bravyi-Kitaev, etc.).

We have implemented this concept in XACC via the definition of an `Observable` interface. XACC implements the `observe` method to take as input an unmeasured `CompositeInstruction` and return a list of measured `CompositeInstructions`. We also define the `ObservableTransform` to take an Observable and map it to a new Observable. This transformation interface enables an extensible mechanism for contributing fermion-spin mappings (Jordan-Wigner, Bravyi-Kitaev, etc.) and operator symmetry reductions (e.g. qubit tapering [18]). The class structure for these interfaces is shown in Fig. 10. Fig. 11 demonstrates how Observables may be created in code.

#### E. Accelerator

The primary interface put forward by the XACC back-end layer is the `Accelerator`. This concept is intended to be implemented for physical and virtual quantum computing back-ends that can be local or remote. Accelerators expose an `execute()` method that takes as input an `AcceleratorBuffer` instance and the `CompositeInstruction` containing the compiled representation of the quantum program. Accelerators are free to implement this as necessary for executing the program on their representative hardware or simulator, but must persist all measurement results to the provided buffer instance. Implementations can also take advantage of

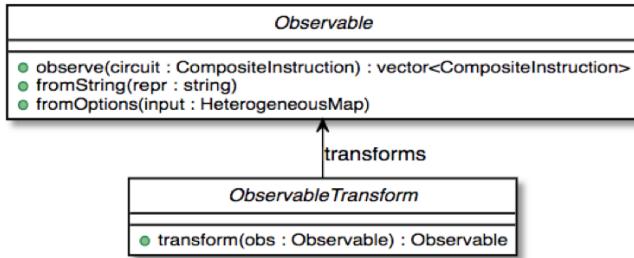


FIG. 10: The Observable and ObservableTransform interfaces.

the buffer heterogeneous metadata member to persist any other pertinent information (such as readout error probabilities,  $T_1$  and  $T_2$  times, or expectation values). The `execute` method is overloaded to enable the execution of multiple `CompositeInstructions` at once to improve on the number of remote process invocations required to execute a given problem.

`Accelerators` can be initialized with a heterogeneous map of options, a useful feature for tweaking the number of shots, the desired remote back-end, and any other execution parameters. `Accelerators` can return a list of hardware-dependent `IRTransformations` that will be run on the incoming `CompositeInstruction` to ensure that the program is amenable for execution on the accelerator. `Accelerators` also expose the underlying qubit connectivity via the `getConnectivity()` method. This can be used in compilation routines or other `IRTransformation` implementations.

XACC defines a decorator pattern on the `Accelerator` interface that promotes extensible pre- and post-processing of `Accelerator` execution input and results. The `AcceleratorDecorator` interface is an `Accelerator` sub-type but also delegates to a concrete `Accelerator` (see Fig. 12). It implements `Accelerator::execute()` to insert pre-processing of execution input, execution on the delegated `Accelerator`, and post-processing of measurement results and metadata. This design is particularly useful in near-term quantum computation as it enables an extension point for general error mitigation strategies. For example, one could implement automated qubit measurement readout-error mitigation by implementing the `AcceleratorDecorator` to post-process computed expectation values according to shift rules based on known readout error probabilities [1]. The snippet in Fig. 13 demonstrates how one might leverage this capability to automate certain error mitigation strategies. It is useful to note that this pattern enables a chain of decorators, with the final decorator delegating to

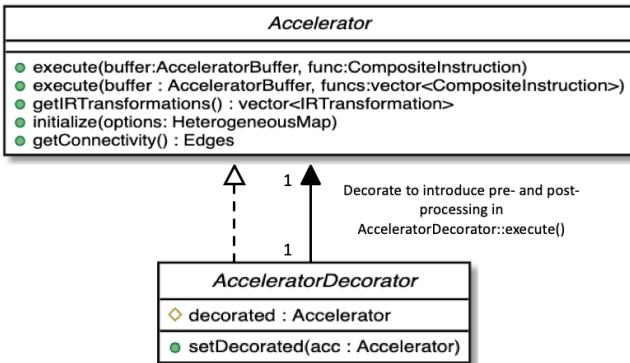


FIG. 12: The class architecture for the `Accelerator` and `AcceleratorDecorator`. The `Accelerator` provides an extension point for general quantum back-ends, while the decorator enables general pre- and post-processing around quantum back-end execution.

```

1 // Get a concrete accelerator
2 auto ibm = xacc::getAccelerator("ibm:tokyo",
3                                 {"shots", 8192});
4
5 // get an accelerator decorator
6 auto ro_decorator =
7     xacc::getAcceleratorDecorator("ro-error");
8
9 // decorate
10 ro_decorator->decorate(ibm);
11
12 // execute just like you
13 // would any accelerator
14 ro_decorator->execute(buffer, circuit);
15
16 // buffer has readout-error
17 // corrected results...
18 auto ro_corrected =
19     buffer->getInformation("exp-val");

```

FIG. 13: A code snippet demonstrating how one might leverage `Accelerators` and `AcceleratorDecorators` to automate certain error mitigation strategies.

a concrete `Accelerator`, thereby enabling a composition of different pre- and post-processing strategies (multiple error mitigation strategies).

## F. Algorithms

XACC defines a general `Algorithm` interface to enable the injection of pre-defined quantum-classical algorithms and therefore relieve programmers of the burden of having to program them from scratch. XACC defines an `Algorithm` as a concept or abstraction that takes as input a heterogeneous map of input data, and exposes a mechanism for execution with a provided `AcceleratorBuffer`. Examples of algorithm implementations provided by XACC are the variational quantum eigensolver, reduced density matrix element generation for quantum chemistry calculations, and a data-driven circuit learning algorithm for machine learning tasks that leverage quantum back-ends. Sec. V demonstrates the utility of this `Algorithm` interface.

## G. Optimizer

The QCOR specification puts forward a general optimizer concept [17]. This abstraction is meant to provide a general extension point for methods that optimize general multi-variate functions, and specifically for quantum computing, functions that internally make calls to a quantum back-end. XACC implements this concept via an `Optimizer` interface, which exposes an `optimize()` method that takes as input

a general `OptFunction`. The `OptFunction` is a thin class that wraps a C++ `std::function<double(const std::vector<double>, std::vector<double> &)>` instance representing the function to be optimized. This function instance must take as its first argument a `const std::vector<double>` representing the current optimization iteration parameters, and as its second input a `std::vector<double>&` representing the mutable function gradient with respect to the current iterate's parameters. It returns a double representing the evaluation of the function at the provided parameters. `Optimizer` implementations are free to ignore the second argument in derivative-free optimization routines. `Optimizers` also expose a `setOptions()` method that takes as input a heterogeneous map of input data, enabling programmers to configure the underlying optimization routines being leveraged.

XACC has concrete implementations of the `Optimizer` that delegate to the well-known `NLOpt` [19] and `mlpack` [20] libraries. The code snippet in Fig. 14 demonstrates the utility of this extension point.

```

1  auto optimizer =
2      xacc::getOptimizer("nlopt");
3
4  OptFunction f(
5      [](const std::vector<double> &x,
6          std::vector<double> &grad) {
7
8      // any preprocessing needed before
9      // back-end execution
10
11     // Execute on back-end Accelerator
12     qpu->execute(buffer, circuits);
13
14     // post-processing on results
15     // compute f_val scalar, compute
16     // gradient if necessary
17
18     return f_val;
19 },
20 n_params);
21
22 optimizer->setOptions(
23     HeterogeneousMap{
24         std::make_pair("nlopt-maxeval", 200),
25         std::make_pair("nlopt-optimizer",
26                         "l-bfgs")
27     });
28
29 auto result = optimizer->optimize(f);

```

FIG. 14: A code snippet demonstrating how one might leverage `Optimizers`.

## H. The XACC Service Registry

At its core, XACC architecture puts forward interfaces for accelerators, accelerator decorators, instructions, composite instructions, IR transformations, compilers, algorithms, observables, and observable transforms. XACC sees these interfaces as extension points for the framework, with implementations serving as available services for the various aspects of the quantum programming and execution workflow. In an effort to ensure modularity and extensibility, XACC only handles service implementations as standard (shared) pointers to the interface being implemented. In this way, XACC never knows about underlying types, implying that a single XACC install can be built but the underlying functionality can change without a total re-build of the framework.

To enable this, XACC requires a robust factory pattern that promotes this sort of interface-based programming. This is common in the Java, where major application frameworks provide modularity and extensibility through interface registration, dynamic loading, and reflection. To promote this in a standard way, the open-source community put forward the Open Services Gateway Initiative (OSGi) [21], which defines several abstractions common to modular frameworks with runtime-extensibility. The OSGi specification defines the concept of a *bundle* containing service (interface) implementations that are declared to the framework at runtime via standard metadata files. These services are loaded and registered with a core framework instance, thus enabling clients to request an instance of an interface implementation given by a specified string name. This technology underlies most graphical integrated development environments (IDE), and enables integration of programming tools across a number of languages.

XACC picks up on these concepts, and leverages a C++ implementation of the OSGi specification called `CppMicroServices` [22]. This framework enables runtime registration of interface implementations against their interface type and unique string identifiers. These implementations are provided to the framework as plugins, or code that sits outside the core XACC context, but provides an implementation of a core XACC interface, and can be added or removed without affecting the core framework. These plugins are built and installed as standard shared libraries, and dropped into a specific folder that XACC defines. At runtime, XACC loads the plugins (shared libraries) into the core CppMicroServices framework and all available interface implementations are registered. Clients can then request specific services by providing the interface type and the name of the service. The code in Fig. 15 demonstrates this process. First, a CppMicroServices framework is created and initialized, then all plugin libraries are loaded and all bundles are started (thus registering all service implementations). Later, clients can request a service (here `Compiler`) of a specified unique name (`openqasm`).

```

1 // Create CppMicroServices Framework
2 auto framework =
3     FrameworkFactory().NewFramework();
4
5 // Initialize the Framework
6 framework.Init();
7
8 // Get the Bundle Context
9 auto context = framework.GetBundleContext();
10
11 // Load and install plugins
12 auto plugins =
13     loadPluginLibraries("~/xacc/plugins");
14 for (auto& plugin : plugins) {
15     context.InstallBundles(plugin);
16 }
17
18 // Start the framework.
19 framework.Start();
20
21 // Get all installed bundles and install them
22 auto bundles = context.GetBundles();
23 for (auto b : bundles) {
24     b.Start();
25 }
26
27 // ... Later requesting a certain Service
28
29 std::string name = "openqasm";
30 xacc::Compiler compiler;
31 auto compilers =
32     context.GetServiceReferences<Compiler>();
33 for (auto& s : compilers) {
34     auto service = context.GetService(s);
35     if (service &&
36         service->name() == name) {
37         compiler = service;
38         break;
39     }
40 }

```

FIG. 15: Example code snippet demonstrating the CppMicroServices OSGi implementation and its utility to XACC. The CppMicroServices framework is started, bundles loaded as shared libraries, and services are requested, all at runtime without direct dependence on third-party contributions or framework re-builds.

To handle all of this, XACC leverages a class called the `ServiceRegistry`, which exposes templated methods to retrieve implementations corresponding to a certain interface type and the unique string name. This class exposes a `getService<TYPE>(name:string):TYPE` method allowing programmers to request a certain XACC interface implementation at the given name. The previously demonstrated code snippets in this work reference API calls like `xacc::getAccelerator()`, `xacc::getCompiler()`, etc. These are just wrappers around calls to the

```

1 // Create a map and use insert()
2 // method leveraging template
3 // type deduction
4 HeterogeneousMap m;
5 m.insert("int-key", 1);
6 m.insert("double-key", 2.0);
7 m.insert("vector-key",
8         std::vector<double>{1.0,2.0});
9 m.insert("string-key", std::string("hello"));
10
11 // Create from initializer list
12 HeterogeneousMap m2{
13     std::make_pair("int-key",1),
14     std::make_pair("double-key",2.0)
15 };

```

FIG. 16: Demonstration of the XACC `HeterogeneousMap`. This utility class is used throughout XACC as a mechanism for handling heterogeneous, problem- and back-end-specific information.

`ServiceRegistry::getService<T>(name:string)`, which returns the appropriate instance of type T.

## I. Heterogeneous Map

A crucial utility data structure that has been mentioned throughout the previous discussion and code snippets, but not yet elaborated on, is the XACC `HeterogeneousMap`. This abstraction is described in the QCOR specification [17] as a mechanism or data structure for mapping string keys to any value type. This type of construct is commonplace in Python (dict is a heterogeneous mapping), but is more difficult to implement in a system-level, statically typed language like C++. The standard library implements a map data structure, but it must be instantiated with the concrete key and value types. In C++17, there are two ways how this can be achieved. One way to achieve this heterogeneous value-type behavior would be to use type erasure via C++17 `std::any<T>`. Another way would be to create a standard map with string keys and variant-like value (`std::variant` in C++17). The XACC `HeterogeneousMap` leverages the C++14 standard, specifically static template class members, in order to achieve heterogeneous value types, thus enabling the data structure and corresponding API demonstrated in Fig. 16.

## J. XACC Framework API

The public XACC API has been demonstrated (in part) throughout the code snippets in the previous sections. Here we elaborate on the calls exposed in the API

and detail the functionality each provides. This API is the primary interface clients/programmers leverage when interacting with XACC and its underlying service framework. The API is demonstrated graphically in Fig. 17.

All applications leveraging the XACC framework must initialize the framework before invocation of any XACC API calls. XACC provides an `Initialize()` function for this, which takes as input the standard `argc`, `argv` command line argument variables. The primary function of this method is to parse and analyze any application command line options, and initialize the underlying XACC CppMicroServices plugin framework. This involves the searching, loading, and registering of shared libraries exposing core XACC service implementations. A corresponding `Finalize` API call is provided for invocation at the end of XACC usage that tears down the service registry and cleans up allocated memory.

The next API call that programmers will interact with is an allocator of quantum memory. In a similar vein as the C `malloc` call, XACC exposes a `qalloc` API call that allocates an instance of the `AcceleratorBuffer`. Note that at this level, XACC defines a `typedef` for `AcceleratorBuffer` called `qbit`. The primary purpose for this `typedef` is for readability, and it is leveraged in the QCOR single-source quantum-classical C++ compiler [23]. `qbit` and `AcceleratorBuffer` can be used interchangeably.

XACC exposes public methods for getting references to implementations of core interfaces, such as accelerator, compiler, IR provider, algorithm, IR transformation, and observable implementations. Each implementation exposes a unique name (`ibm` for the IBM Accelerator, for instance), so each of these API calls (`getAccelerator`, `getCompiler`, etc.) takes as input the name of the service implementation, and returns a shared pointer to that instantiated service. Programmers can also directly leverage a templated `getService<T>(name:string) : T` call to instantiate a service of type `T` with the given string name.

To improve programming efficiency, readability, and

XACC
<ul style="list-style-type: none"> <li>● <code>Initialize(argc:int, argv:char**)</code></li> <li>● <code>qalloc(size:int) : qbit</code></li> <li>● <code>getAccelerator(name:string, HeterogeneousMap options) : Accelerator</code></li> <li>● <code>AcceleratorDecorator(name:string) : AcceleratorDecorator</code></li> <li>● <code>getCompiler(name:string) : Compiler</code></li> <li>● <code>getIRProvider(name:string) : IRProvider</code></li> <li>● <code>getAlgorithm(name:string) : Algorithm</code></li> <li>● <code>getObservable(name:string) : Observable</code></li> <li>● <code>getObservableTransform(name:string) : ObservableTransform</code></li> <li>● <code>qasm(src:string)</code></li> <li>● <code>getCompiled(name:string) : CompositeInstruction</code></li> <li>● <code>getService&lt;T&gt;(name:string) : T</code></li> <li>● <code>Finalize()</code></li> </ul>

FIG. 17: The public XACC API provides programmers with an easy mechanism for interacting with the core XACC service registry, as well as allocating `AcceleratorBuffers` and compiling quantum kernels.

```

1 // JIT map XASM Ansatz to IR
2 xacc::qasm(R"( 
3   .compiler xasm
4   .circuit deuteron_ansatz
5   .parameters theta
6   .qbit q
7   X(q[0]);
8   Ry(q[1], theta);
9   CNOT(q[1],q[0]);
10  )");
11 auto ansatz =
12   xacc::getCompiled("deuteron_ansatz");
13
14 // Quil example, multiple kernels
15 xacc::qasm(R"(.compiler quil
16   .circuit ansatz
17   .parameters theta, phi
18   X 0
19   H 2
20   CNOT 2 1
21   CNOT 0 1
22   Rz(theta) 0
23   Ry(phi) 1
24   H 0
25   .circuit x0x1
26   ansatz(theta, phi)
27   H 0
28   H 1
29   MEASURE 0 [0]
30   MEASURE 1 [1]
31  )");
32 auto x0x1 = xacc::getCompiled("x0x1");

```

FIG. 18: Demonstration of the XACC `qasm()` call. Note multiple kernel strings can be compiled and reference each other.

utility of the quantum kernel string compilation (like in Fig. 6), XACC exposes a `qasm()` function. This function takes as input an enhanced quantum kernel source string syntax and compiles it to XACC IR. This source string is *enhanced* in that it requires that extra metadata be present in order to adequately compile the quantum code. Specifically, the source string must contain the following key words:

- a single `.compiler NAME`, to indicate which XACC compiler implementation to use.
- one or many `.circuit NAME` calls to give the created `CompositeInstruction` (circuit) a name.
- one `.parameters PARAM_NAME_1, PARAM_NAME_2, ..., PARAM_NAME_N` for each parameterized circuit to tell the Compiler the names of the parameters.
- A `.qbit NAME` optional keyword can be provided when the source code itself makes reference to the `qbit` or `AcceleratorBuffer`

Running this command with the appropriately provided keywords will compile the source string to XACC IR and store it in an internal compilation database (standard map of CompositeInstruction names to CompositeInstructions), and users can get reference to the individual CompositeInstructions via an exposed `getCompiled()` XACC API call. The code in Fig. 18 demonstrates how one would use `qasm()` and its overall utility.

#### IV. INTERFACE IMPLEMENTATIONS FOR QUANTUM COMPUTING

Having introduced the core interfaces exposed by core XACC, we now turn to detailing concrete implementations of these interfaces for quantum computation. Specifically, we will discuss gate and annealing model implementations of the core XACC IR interfaces, Pauli and Fermion implementations of the `Observable` interface (and Jordan-Wigner, Bravyi-Kitaev implementations of `ObservableTransform`), `Compiler` implementations for standard low-level assembly languages (Quil, OpenQasm, XACC XASM) using Antlr [24], `Accelerators` for available QPUs and simulators, and `Algorithm` implementations. We note that for most use cases programmers do not need to know the concrete type of these implementations as instances are created and provided as pointers to the interface itself via the XACC public API. The implementations to be discussed here are available in the `xacc::quantum` package.

XACC defines `Gate` and `Circuit` classes implementing `Instruction` and `CompositeInstruction`, respectively. To hide this type information from programmers, XACC defines a `QuantumIRProvider` which implements the `IRProvider` interface for the creation of `Gates` and `Circuits` (returned as pointers to `Instruction` and `CompositeInstruction`). The `Gate` implements the aforementioned `Instruction` API and is further subclassed by types providing specific digital gate implementations, such as Hadamard, CNOT, Rz, etc. The `Circuit` class implements the aforementioned `CompositeInstruction` API to provide a container for `Instructions` while also remaining a valid `Instruction` itself. The `Circuit` may contain further `Circuits`, enabling the familiar *n*-ary tree pattern. `Circuits` keep track of string variables that it and sub-`Instructions` depend on, and implements `CompositeInstruction::operator()()` to take concrete double values and evaluate all sub-`Instructions` dependent on a specific `Circuit` variables. A key aspect of `Circuits` is that they can be *dynamic*, i.e. sub-types of `Circuit` can override the `CompositeInstruction::expand` method taking an `HeterogeneousMap` of input data and expanding itself with concrete `Gate` instructions. This is useful for dynamic, pre-defined composite instructions that can be programmed at compile-time but expanded at runtime. A few examples of this, which are implemented in XACC,

```

1 // JIT map XASM Ansatz to IR
2 xacc::qasm(R"(
3   .compiler xasm
4   .circuit deuteron_ansatz
5   .parameters x
6   .qbit q
7   for (int i = 0; i < 2; i++) {
8     H(q[0]);
9   }
10  exp_i_theta(q, x, {"pauli", "X0 Y1 - Y0 X1"});
11  );
12  auto ansatz =
13    xacc::getCompiled("deuteron_ansatz");

```

FIG. 19: Demonstration of the XACC XASM Compiler through the `qasm()` call.

are the following `Circuit` sub-types:

- `range` Takes as input a `Gate` name and the range of qubits (given as a standard range `[start_idx, end_idx]`, or the total number of qubits). Based on this input it adds that `Gate` to all qubits in the range. Typical invocation using the XASM compiler looks like `range(buffer, {"gate", "H"}, {"nq", 4});`.
- `qft` Takes as input the range of qubits (given as a standard range `[start_idx, end_idx]`, or the total number of qubits) upon which to append a quantum Fourier transform. Typical invocation using the XASM compiler looks like `qft(buffer, {"nq", 4});`.
- `exp_i_theta` Takes as input a representation of a spin or fermionic operator  $H$  and provides a first order trotterization of  $\exp(i\theta H)$ . Typical invocation using the XASM compiler looks like `exp_i_theta(buffer, theta, {"pauli", "X0 X1 - Y0 Y1"});`.

An example of using these dynamic circuit generators with the `xacc::qasm()` call is shown in Fig. 19.

XACC defines a number of `Compiler` implementations, primary of which are the Quil, OpenQasm, and XASM compilers. All of these implementations leverage

```

1 // Note Accelerators can be initialized
2 // with a HeterogeneousMap.
3 auto ibm = xacc::getAccelerator("ibm",
4   {"shots", 8192},
5   {"backend", "ibmq_poughkeepsie"});
6 ibm->execute(buffer, circuits);

```

FIG. 20: Demonstration of configuring, getting, and executing a remote Accelerator.

```

1  __qpu__ foo(qbit q, double h, double j) {
2    qmi(q[0],q[0], h);
3    qmi(q[1],q[1], h);
4    qmi(q[0],q[1], j);
5  }

```

FIG. 21: An example quantum kernel for the D-Wave architecture.

the `Antlr` parser generation library for creating parse-tree listeners that map individual `Antlr` nodes to XACC IR instructions and composite instructions. The default `Compiler` in XACC is the XASM compiler (XACC assembly language compiler). This compiler puts forward a grammar specification for a C or C++ style low-level assembly language for quantum computing. This language is QASM-like in nature, but inherits the features of the XACC IR static and dynamic instruction set. In this way, XASM provides a dynamic, living language (living in the sense that low-level instructions that can be expressed are in fact XACC Instruction plugins, and can be added/removed at runtime). The XASM grammar specifies instructions with a unique name, a set of bits to operate on, an optional set of gate parameters, and an optional heterogeneous map of input parameters. Moreover, the XASM grammar supports simple `for` loops that enable programming multiple instructions in a more efficient and familiar manner. These XASM language source strings are parsed via auto-generated `Antlr` data structures and mapped to XACC IR via an `Antlr` tree-node visitation event listener implementation. The OpenQasm and Quil compilers are implemented in a similar way, but are based off of vendor-supplied grammar specifications.

XACC defines `Accelerator` implementations for IBM, Rigetti, D-Wave, IonQ, and a number of back-end simulators (TNQVM [25], C++ local IBM noise-aware simulator, etc). Each of these physical QPU implementations actually subclass a `RemoteAccelerator` class, which further subclass `Accelerator`. This remote subtype provides key functionality (that is common across currently available physical QPUs) for remote REST API calls (GET, POST, etc.). The base implementation relies on the Curl for People (CPR) implementation hosted at [26]. These `Accelerators` are further parameterized by pertinent execution metadata such as number of shots, specific back-end, etc. We note that these `Accelerators` do not delegate to provided Pythonic frameworks (qiskit, pyquil, etc.) but instead make direct calls to the exposed REST API provided by the vendors. An example of getting reference to the IBM `Accelerator` targeting the Poughkeepsie back-end with 8192 shots is provided in Fig. 20.

XACC provides two implementations of the `Observable` interface: the `PauliOperator` and the `FermionOperator`. The `PauliOperator` keeps track internally of a mapping of terms, with each term encapsulating a coefficient, a map of qubit sites to

string X, Y, Z, and an optional variable string (for parameterized operators). The `FermionOperator` internally keeps a map of terms, each term containing a coefficient, map of fermionic sites to bool indicating creation or annihilation, and an optional variable string. Both operators expose an API for common algebraic operations on operators (addition, subtraction, multiplication). Both leverage a custom `Antlr` parser for enabling the `Observable::fromString` functionality (`PauliOperators` can be created from strings like X0 X1, `FermionOperators` can be created from strings like 0^ 1). `PauliOperator` implements `observe()` to take an unmeasured `CompositeInstruction`, and for each term the operator contains append appropriate measurement gates, and return a list of measured `CompositeInstructions`. The `FermionOperator` implements `observe()` to first transform itself via an appropriate `ObservableTransform` (Jordan-Wigner, for instance) and then call and return the result of the `observe()` call invoked on the transformed `Observable`.

XACC provides support for the D-Wave quantum computing architecture by implementing the XACC core interfaces for expressing low-level machine instructions that describe biases and couplers for the D-Wave Ising Hamiltonian. Specifically we implement the `Instruction` interface for D-Wave quantum machine instructions that describe the bias or coupling strength on a given qubit or pair of qubits. XACC defines the `AnnealingInstruction` sub-type of `Instruction` which keeps track of two qubit indices and an `InstructionParameter` parameter for describing the bias or coupler (which can be a concrete double or variable). If the two qubit indices are equal, the instruction describes a bias value, and if they are different, it describes a coupler. Fig. 21 provides an example of programming an XACC quantum kernel with these D-Wave instructions. Of course, one could also subclass `CompositeInstruction` and implement `expand()` to auto-generate these instructions at runtime.

XACC also defines an `EmbeddingAlgorithm` interface for the injection of minor graph embedding strategies used in the mapping of these quantum machine instruction representations onto the physical hardware.

## V. DEMONSTRATION

Here we seek to demonstrate the utility of the XACC framework through a few concrete, common use cases. Specifically, we attempt to demonstrate use cases through a variety of XACC abstraction levels, beginning with low-level use cases depicting the base XACC API, and ending with complex hybrid quantum-classical algorithms that leverage custom `Observables` and `Algorithms`. The example code shown here is meant to present complete, working programs.

## A. Base API

Fig. 22 demonstrates how one might use the lowest level of the XACC API to construct a parameterized quantum program and execute it on a desired back-end. All XACC programs simply include the public XACC header file which provides access to the public XACC API discussed in Section IIIJ. In order to start requesting services of a given type, pro-

```

1  #include "xacc.hpp"
2  int main(int argc, char **argv) {
3      xacc::Initialize(argc, argv);
4      // Get reference to the Accelerator
5      auto accelerator =
6          xacc::getAccelerator("tnqvm");
7
8      // Get the IRProvider and create an
9      // empty CompositeInstruction
10     auto provider =
11         xacc::getIRProvider("quantum");
12     auto program =
13         provider->createComposite("foo", {"t"});
14
15     // Create X, Ry, CX, and Measure gates
16     auto x =
17         provider->createInstruction("X", {0});
18     auto ry =
19         provider->createInstruction("Ry", {1},
20                                     {"t"});
21     auto cx =
22         provider->createInstruction("CNOT", {1,0});
23     auto m0 =
24         provider->createInstruction("Measure",
25                                     {0});
26
27     // Add them to the CompositeInstruction
28     program->addInstructions({x, ry, cx, m0});
29
30     // Loop over [-pi, pi] and compute <Z0>
31     auto angles = xacc::linspace(-pi, pi, 20);
32     for (auto &a : angles) {
33         auto buffer = xacc::qalloc(2);
34         auto evalued = program->operator()({a});
35         accelerator->execute(buffer, evalued);
36         std::cout << "<Z0>(" << a << ") = "
37             << buffer->getExpectationValueZ()
38             << "\n";
39     }
40
41     xacc::Finalize();
42     return 0;
43 }
```

FIG. 22: Demonstration of programmatically constructing quantum programs as XACC IR. Here we construct a parameterized gate model circuit, and evaluate it at a number of concrete angles for execution on the TNQVM [25] simulation back-end.

grammers must first initialize the framework with the `xacc::Initialize()` call. This loads all available plugins (implementations of the core XACC interfaces) and parses command line options. Programmers next get reference to the desired `Accelerator` back-end through the `getAccelerator()` call. Next, to programmatically construct XACC IR, we get reference to an instance of the `IRProvider` and use it to create `Instructions` and `CompositeInstructions`. All concrete `Instructions` are added to the `CompositeInstruction` after creation. We now have a parameterized instance of the XACC IR (parameterized on the variable `t`, the rotation angle for the `Ry` instruction). Now, we are free to execute this `CompositeInstruction` at concrete instances of the variable `t`. We note that there is no reconstruction of the program or circuit at each `t`. We simply evaluate this `CompositeInstruction` at the given variable value and execute that concrete program. Execution is affected by a call to `Accelerator::execute()` which takes as input the user-provided buffer and the concrete, evaluated program. The `Accelerator::execute()` implementation is responsible for persisting all execution results to the pro-

```

1  #include "xacc.hpp"
2  int main(int argc, char **argv) {
3      xacc::Initialize(argc, argv);
4
5      // Get reference to the Accelerator
6      auto accelerator =
7          xacc::getAccelerator("ibm:ibmq_valencia");
8
9      // Allocate some qubits
10     auto buffer = xacc::qalloc(2);
11
12     // Compile
13     auto quil = xacc::getCompiler("quil");
14     auto ir = quil->compile(
15         R"(_qpu_ void bell(qbit q) {
16             H 0
17             CX 0 1
18             MEASURE 0 [0]
19             MEASURE 1 [1]
20         })", accelerator);
21
22     // Execute
23     accelerator->execute(buffer,
24                           ir->getComposites()[0]);
25
26     // View results
27     buffer->print();
28
29     xacc::Finalize();
30     return 0;
31 }
```

FIG. 23: Demonstration of creating raw quantum kernel source strings and creating IR via compilation with an appropriate XACC `Compiler` instance.

vided buffer. Afterwards, we are then free to process the results, here we get the expected value in the computational basis. All XACC programs end with a call to `xacc::Finalize()` which tears down the core service registry and cleans up allocated memory.

## B. Kernel Source Compilation and Execution

We note that the example in Fig. 22 is verbose and requires a lot of boilerplate code to express relatively simple concepts. So now we move up to a higher level of XACC abstraction, specifically, compiling quantum kernel source strings via `Compilers` and the `xacc::qasm()` call.

As a dual-source approach to quantum-accelerated computing, programmers express quantum code intended for compilation and execution as separate kernel source strings. The code in Fig. 23 demonstrates this functionality, specifically in programming a Bell state computation on the remotely hosted IBM Valencia backend. Programmers begin with the usual `Initialize` and `getAccelerator` calls, and instead of programmatically constructing XACC IR, get reference to a specific `Compiler` implementation to be used in kernel compilation. Programmers compose the quantum kernel code as a string (or string literal) and compile it with the correct `Compiler` instance, which creates and returns an instance of the IR, ready to be used in back-end execution. This compilation step can optionally take the `Accelerator` as input, thereby exposing connectivity, noise, or other back-end-specific properties at compile time.

Of course, we can replace this get-compiler, compile, get-IR, workflow with a single call to the `xacc::qasm()` function, depicted in Fig. 19.

## C. Hybrid Quantum-Classical Algorithms

Here we demonstrate leveraging the highest-level utilities in XACC to execute a typical quantum-classical algorithm. Specifically, we consider two common applications: (1) computing the binding energy of deuteron using the  $N = 3$  Hamiltonian from [1] using the variational quantum eigensolver (VQE), and (2) training a quantum circuit to learn a provided target probability distribution.

First we consider the application from nuclear physics, with XACC code demonstrated in Fig. 24. This example demonstrates the utility of the various high-level abstractions discussed in the preceding discussion, specifically the `Observable`, `Optimizer`, and `Algorithm`. Specifically we seek to compute the ground state energy of

$$H_3 = 15.531709I + 0.218291Z_0 - 6.125Z_1 - 2.143304(X_0X_1 + Y_0Y_1) - 9.625Z_2 - 3.913119(X_1X_2 + Y_1Y_2). \quad (1)$$

```

1 #include "xacc.hpp"
2 int main(int argc, char **argv) {
3     xacc::Initialize(argc, argv);
4
5     // Get the desired Accelerator and Optimizer
6     auto qpu =
7         xacc::getAccelerator("ibm:ibmq_valencia");
8     auto optimizer = xacc::getOptimizer("nlopt");
9
10    // Create the N=3 deuteron Hamiltonian
11    auto H_N_3 = xacc::quantum::getObservable(
12        "pauli",
13        std::string("15.531709 - 2.1433 X0X1 - "
14        "2.1433 Y0Y1 + .21829 Z0 - 6.125 Z1 - "
15        "9.625 Z2 - 3.91 X1 X2 - 3.91 Y1 Y2"));
16
17    // JIT map Quil QASM Ansatz to IR
18    xacc::qasm(R"(
19        .compiler xasm
20        .circuit ansatz
21        .parameters t0, t1
22        .qbit q
23        X(q[0]);
24        exp_i_theta(q, t0,
25            {"pauli", "X0 Y1 - Y0 X1"});
26        exp_i_theta(q, t1,
27            {"pauli", "X0 Z1 Y2 - X2 Z1 Y0"});
28    )");
29    auto ansatz = xacc::getCompiled("ansatz");
30
31    // Get the VQE Algorithm and initialize it
32    auto vqe = xacc::getAlgorithm("vqe");
33    vqe->initialize({
34        std::make_pair("ansatz", ansatz),
35        std::make_pair("observable", H_N_3),
36        std::make_pair("accelerator", qpu),
37        std::make_pair("optimizer", optimizer)
38    });
39
40    // Allocate some qubits and execute
41    auto buffer = xacc::qalloc(3);
42    vqe->execute(buffer);
43
44    // Get the result
45    auto energy =
46        (*buffer)[("opt-val"] .as<double>();
47
48    xacc::Finalize();
49    return 0;
50 }
```

FIG. 24: Demonstration of a complex quantum-classical algorithm using `Observable`, `Optimizer`, and `Algorithm`. Specifically, this code snippet executes the variational quantum eigensolver to compute the ground state energy of the three qubit deuteron Hamiltonian in [1].

```

1 #include "xacc.hpp"
2 int main(int argc, char **argv) {
3     xacc::Initialize(argc, argv);
4
5     // Get reference to the Accelerator
6     auto qpu =
7         xacc::getAccelerator("qcs:Aspen-4-2Q-A");
8
9     auto opt =
10        xacc::getOptimizer("mlpack");
11    xacc::qasm(R"
12 .compiler xasm
13 .circuit ansatz
14 .parameters x
15 .qbit q
16 U(q[0], x[0], -pi/2, pi/2 );
17 U(q[0], 0, 0, x[1]);
18 U(q[1], x[2], -pi/2, pi/2);
19 U(q[1], 0, 0, x[3]);
20 CNOT(q[0], q[1]);
21 U(q[0], 0, 0, x[4]);
22 U(q[0], x[5], -pi/2, pi/2);
23 U(q[1], 0, 0, x[6]);
24 U(q[1], x[7], -pi/2, pi/2);
25)");
26     auto ansatz = xacc::getCompiled("ansatz");
27
28     std::vector<double> target_distribution
29                     {.5, .5, .5, .5};
30
31     auto ddcl = xacc::getAlgorithm("ddcl");
32     ddcl->initialize({
33         std::make_pair("ansatz", ansatz),
34         std::make_pair("target_dist",
35             target_distribution),
36         std::make_pair("accelerator", qpu),
37         std::make_pair("loss", "js"),
38         std::make_pair("gradient",
39             "js-parameter-shift"),
40         std::make_pair("optimizer", opt)});
41
42     // Allocate some qubits and execute
43     auto buffer = xacc::qalloc(2);
44     ddcl->execute(buffer);
45
46     // Print the result
47     std::cout << "Loss: " <<
48         (*buffer)["opt-val"].as<double>()
49             << "\n";
50 }

```

FIG. 25: Demonstration of using XACC and the `Algorithm` interface for quantum-enhanced machine learning tasks.

with the variational ansatz given by

$$U(\eta, \theta) \equiv e^{-i\frac{\eta}{2}(X_0Y_1 - X_1Y_0)} e^{-i\frac{\theta}{2}(X_0Z_1Y_2 - X_2Z_1Y_0)}. \quad (2)$$

The code snippet begins with the usual initialization and getting reference to the desired `Accelerator`. Next, we create an `Observable` that models the three qubit Hamiltonian. Note specifically how we are able to express this Hamiltonian as it would be written in text. Next we get reference to the default `NLOpt Optimizer` for our VQE execution. Next, we program and compile the VQE state preparation circuit via the `xacc::qasm()` call, and we leverage the dynamic runtime `exp_i_theta CompositeInstruction`. This too can be expressed very simply, with a straightforward translation from the equation specified in Eq. 2. The dynamic exponentiation instruction will be just-in-time expanded to the appropriate circuit decomposition. Next, we get reference to the VQE `Algorithm`, which must be initialized with a valid `ansatz`, `observable`, `optimizer`, and `accelerator`. We then allocate three qubits via the `xacc::qalloc()` call, and execute the `Algorithm` with the allocated buffer. This call kicks off the VQE algorithm, and the `Optimizer` will iteratively converge to the optimal energy over the two circuit parameters `t0, t1`. The results of the computation can be retrieved from the `AcceleratorBuffer`, with execution results from every iteration and for every measured quantum program persisted to the provided `AcceleratorBuffer` as children `AcceleratorBuffers`.

Next we turn to an application of quantum computation to classical machine learning. Our goal is to train a parameterized quantum circuit to reproduce a target probability distribution [3]. For this data-driven circuit learning (DDCL) process, XACC provides an `Algorithm` implementation called `DDCL`. This implementation takes as input the target probability distribution, a `CompositeInstruction` representing the parameterized circuit, the back-end `Accelerator` to run on, the `Optimizer` used for the learning process, and loss function and gradient strategies. With this input, `DDCL` orchestrates the workflow necessary for to find the optimal circuit parameters that lead to a probability distribution that mimics the given target distribution. Fig. 25 demonstrates how one would use this `Algorithm`. The programmer starts by getting reference to the desired `Accelerator` and `Optimizer`, here a 2-qubit lattice on the Rigetti QCS platform and the ML Pack optimizer [20], respectively. Next we compile an XASM source code with 8 parameters using general, three-parameter IBM U rotation gates and a CNOT. Note the use of IBM (U, CNOT) instructions for Rigetti (RZ, CZ), XACC handles this cross-compilation via subsequent transformations of the IR. We define a target probability distribution and get reference to the `DDCL` algorithm. We initialize the algorithm with the required input data, specifically requesting that the algorithm compute the Jansen-Shannon divergence of the circuit and target distributions. Moreover, we specify a gradient computation strategy that evaluates the parameterized circuit at  $x + \pi/2$  for all  $x$ . Finally, we allocate a two qubit buffer and execute. Upon completion, the buffer provides all pertinent execution information, as well as any algorithm-specific results.

## VI. DISCUSSION

We have presented a system-level framework for heterogeneous quantum-classical computing. Our goal with this work is to provide a future-proof approach for enabling software technologies targeting quantum-accelerated heterogeneous computing architectures. XACC provides a modular and extensible approach to hardware-agnostic quantum computation that can serve as a foundation for a variety of software tools, compilers, and benchmarks. We have demonstrated a layered architecture that promotes a dual-source, co-processor programming model that enables quantum code to be expressed as high-level kernels. Our approach provides extension points for mapping these kernel expressions to a core, polymorphic intermediate representation, which is amenable to low-level quantum program optimizations, reductions, and analysis. Our extensible back-end infrastructure supports both gate model and annealing quantum computing technologies, and currently supports processors and simulators from IBM, Rigetti, D-Wave, and IonQ. As quantum processor technologies improve, our approach can serve as the foundation of a comprehensive integration framework for all aspects of the quantum programming software stack. This framework can serve as a means for tighter integration models that promote performant workflows leveraging quantum information processing.

## ACKNOWLEDGEMENTS

This work has been supported by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, the US Department of Energy (DOE) Office of Science Early Career Research Award, and the DOE Office of Science Advanced Scientific Computing Research (ASCR) Quantum Computing Application Teams, Quantum Testbed Pathfinder, and Accelerated Research in Quantum Computing programs, and DOE Office of Science High Energy Physics QuantISED. This work was also supported by the ORNL Undergraduate Research Participation Program, which is sponsored by ORNL and administered jointly by ORNL and the Oak Ridge Institute for Science and Education (ORISE). ORNL is managed by UT-Battelle, LLC, for the US Department of Energy under contract no. DE-AC05-00OR22725. ORISE is managed by Oak Ridge Associated Universities for the US Department of Energy under contract no. DE-AC05-00OR22750. The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan.

---

[1] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, *Phys. Rev. Lett.* **120**, 210501 (2018).

[2] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. Morris, T. S. Humble, and R. C. Pooser, arXiv e-prints , arXiv:1905.01534 (2019), [arXiv:1905.01534 \[quant-ph\]](https://arxiv.org/abs/1905.01534).

[3] K. E. Hamilton, E. F. Dumitrescu, and R. C. Pooser, *Phys. Rev. A* **99**, 062323 (2019).

[4] N. Klco, E. Dumitrescu, A. McCaskey, T. Morris, R. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. Savage, arXiv preprint arXiv:1803.03326 (2018).

[5] K. Yeter-Aydeniz, E. F. Dumitrescu, A. J. McCaskey, R. S. Bennink, R. C. Pooser, and G. Siopsis, *Phys. Rev. A* **99**, 032306 (2019).

[6] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, S. D. L. P. González, E. D. L. Torre, D. Ding, E. Dumitrescu, I. Duran, P. Eendebak, M. Everitt, I. F. Sertage, A. Frisch, A. Fuhrer, J. Gambetta, B. G. Gago, J. Gomez-Mosquera, D. Greenberg, I. Hamamura, V. Havlicek, J. Hellmers, L. Herok, H. Horii, S. Hu, T. Imamichi, T. Itoko, A. Javadi-Abhari, N. Kanazawa, A. Karazeev, K. Krsulich, P. Liu, Y. Luh, Y. Maeng, M. Marques, F. J. Martín-Fernández, D. T. McClure, D. McKay, S. Meesala, A. Mezzacapo, N. Moll, D. M. Rodríguez, G. Nannicini, P. Nation, P. Ollitrault, L. J. O'Riordan, H. Paik, J. Pérez, A. Phan, M. Pistoia, V. Prutyanov, M. Reuter, J. Rice, A. R. Davila, R. H. P. Rudy, M. Ryu, N. Sathaye, C. Schnabel, E. Schoute, K. Setia, Y. Shi, A. Silva, Y. Siraichi, S. Sivarajah, J. A. Smolin, M. Soeken, H. Takahashi, I. Tavernelli, C. Taylor, P. Taylor, K. Trabing, M. Treinish, W. Turner, D. Vogt-Lee, C. Vuillot, J. A. Wildstrom, J. Wilson, E. Winston, C. Wood, S. Wood, S. Wörner, I. Y. Akhallowaya, and C. Zoufal, “[Qiskit: An open-source framework for quantum computing](https://qiskit.org),” (2019).

[7] C. Gidney, “[Cirq and openfermion hands on tutorial](https://cirq.readthedocs.io/en/stable/tutorials/openfermion.html),” (2018).

[8] R. S. Smith, M. J. Curtis, and W. J. Zeng, “A practical quantum instruction set architecture,” (2016).

[9] J. R. McClean, K. J. Sung, I. D. Kivlichan, Y. Cao, C. Dai, E. Schuyler Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, O. Higgott, C. Huang, J. Izaac, Z. Jiang, X. Liu, S. McArdle, M. Neeley, T. O’Brien, B. O’Gorman, I. Ozfidan, M. D. Radin, J. Romero, N. Rubin, N. P. D. Sawaya, K. Setia, S. Sim, D. S. Steiger, M. Steudtner, Q. Sun, W. Sun, D. Wang, F. Zhang, and R. Babbush, arXiv e-prints , arXiv:1710.07629 (2017), [arXiv:1710.07629 \[quant-ph\]](https://arxiv.org/abs/1710.07629).

[10] K. A. Britt and T. S. Humble, ACM Journal on Emerging Technologies in Computing Systems (JETC) **13**, 39 (2017).

- [11] A. McCaskey, E. Dumitrescu, D. Liakh, and T. Humble, in *2018 IEEE International Conference on Rebooting Computing (ICRC)* (IEEE, 2018) pp. 1–12.
- [12] “Eclipse xacc,” <http://github.com/eclipse/xacc> (), accessed: 2020-01-06.
- [13] A. McCaskey, E. Dumitrescu, D. Liakh, M. Chen, W. Feng, and T. Humble, *SoftwareX* **7**, 245 (2018).
- [14] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, arXiv e-prints , arXiv:1707.03429 (2017), [arXiv:1707.03429 \[quant-ph\]](https://arxiv.org/abs/1707.03429).
- [15] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-oriented Software* (Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995).
- [16] “Eclipse xacc,” <https://github.com/eclipse/xacc/blob/master/quantum/plugins/circuits/exp/exp.cpp#L33> (), accessed: 2020-01-08.
- [17] T. M. Mintz, A. J. Mccaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, arXiv e-prints , arXiv:1909.02457 (2019), [arXiv:1909.02457 \[cs.PL\]](https://arxiv.org/abs/1909.02457).
- [18] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, arXiv e-prints , arXiv:1701.08213 (2017), [arXiv:1701.08213 \[quant-ph\]](https://arxiv.org/abs/1701.08213).
- [19] “Nlopt,” <https://github.com/stevengj/nlopt>, accessed: 2019-11-01.
- [20] R. R. Curtin, M. Edel, M. Lozhnikov, Y. Mentekidis, S. Ghaisas, and S. Zhang, *Journal of Open Source Software* **3**, 726 (2018).
- [21] “Osgi service platform,” <http://www.osgi.org/>, accessed: 2019-11-01.
- [22] “Cppmicroservices,” <http://github.com/CppMicroServices/CppMicroServices>, accessed: 2019-11-01.
- [23] A. McCaskey, (in-preparatoin).
- [24] T. J. Parr and R. W. Quong, *Softw. Pract. Exper.* **25**, 789 (1995).
- [25] A. McCaskey, E. Dumitrescu, M. Chen, D. Lyakh, and T. Humble, *PLOS ONE* **13**, 1 (2018).
- [26] “Curl for people,” <http://github.com/whoshuu/cpr>, accessed: 2019-10-17.