

QisDAX: An Open Source Bridge from Qiskit to Trapped-Ion Quantum Devices

Kaustubh Badriker*[§], Aniket S. Dalvi[†], Filip Mazurek[†], Marissa D’Onofrio[†], Jacob Whitlow[†], Tianyi Chen[‡], Samuel Phiri[†], Leon Riesebois[†], Kenneth R. Brown[†] and Frank Mueller*[§]

*Department of Computer Science, North Carolina State University, Raleigh, NC

[†]Department of Electrical and Computer Engineering, Duke University, Durham, NC

[‡]Department of Physics, Duke University, Durham, NC

[§]Email: {kjbadrik, fmueller}@ncsu.edu

Abstract—Quantum computing has become widely available to researchers via cloud-hosted devices with different technologies using a multitude of software development frameworks. The vertical stack behind such solutions typically features quantum language abstraction and high-level translation frameworks that tend to be open source, down to pulse-level programming. However, the lower-level mapping to the control electronics, such as controls for laser and microwave pulse generators, remains closed source for contemporary commercial cloud-hosted quantum devices. One exception is the ARTIQ (Advanced Real-Time Infrastructure for Quantum physics) open-source library for trapped-ion control electronics. This stack has been complemented by the Duke ARTIQ Extensions (DAX) to provide modularity and better abstraction. It, however, remains disconnected from the wealth of features provided by popular quantum computing languages.

This paper contributes QisDAX, a bridge between Qiskit and DAX that fills this gap. QisDAX provides interfaces for Python programs written using IBM’s Qiskit and transpiles them to the DAX abstraction. This allows users to generically interface to the ARTIQ control systems accessing trapped-ion quantum devices. Consequently, the algorithms expressed in Qiskit become available to an open-source quantum software stack. This provides the first open-source, end-to-end, full-stack pipeline for remote submission of quantum programs for trapped-ion quantum systems in a non-commercial setting.

I. INTRODUCTION

Development of a practical quantum computer demands a well-designed, modular software architecture that considers all layers of a vertical stack, from the programming language to the qubit-specific hardware [1]–[3]. However, the field is currently dominated by a variety of architectures, where at least some components of their software stacks are customized, at lower levels in industry generally proprietary, and often without a path to cross-platform compatibility.

One source of this incompatibility comes from proprietary hardware components, such as microwave pulse generators and lasers, that are decoupled from the higher-level, hardware-agnostic layers of the software stack. Furthermore, their controls tend to be closed-source. An alternate, *open-source* design can instead be promoted by this type of decoupled stack, which enables better abstraction as well as cross-platform compatibility.

An increasingly popular, open-source quantum control system is ARTIQ (Advanced Real-Time Infrastructure for Quan-

tum physics), a software framework with dedicated, open-source control hardware [4], [5]. The ARTIQ stack is complemented by DAX (Duke ARTIQ extensions) [6], a device abstraction developed to provide high-level, modular utilities for controlling trapped-ion systems. However, in ARTIQ-based quantum computers, the rich capabilities offered by popular quantum computing languages are not accessible to lower layers of the computing stack.

This work contributes QisDAX to bridge this gap. QisDAX facilitates an interface between DAX and Qiskit [7], [8], allowing the entire quantum computation workflow for a trapped ion system to be incorporated into a single open-source, full-stack pipeline. A wide variety of Qiskit algorithms and frameworks therefore become accessible to DAX users and can be applied to a new set of backend devices, such as trapped-ion systems and simulators. By exporting results in Qiskit-compatible objects, QisDAX also facilitates classical result analysis or processing in a hybrid environment with repeated quantum kernel invocations.

Python programs written using Qiskit are transpiled via QisDAX for the DAX abstraction, which includes parallelization of gates wherever possible to reduce circuit depth. QisDAX then remotely submits these programs to the respective quantum device or runs them in simulation. One open-source backend accessible through QisDAX is CRYO-STAQ, a DAX-based trapped-ion quantum computing system hosted at Duke University.

The contributions of this work can be summarized as follows:

- We create a software bridge between Qiskit and DAX.
- We develop an algorithm to reduce circuit depth by parallelizing gates within the capabilities of a given backend device.
- We facilitate interactions between quantum and classical processors in order to evaluate results and realize hybrid quantum-classical computing.
- We allow verification of transpiled code by adding simulator backends.
- We evaluate our software stack both with a real quantum device and a simulator backend.
- We demonstrate the capability of the pipeline to remotely execute programs written in Qiskit on an academically

hosted quantum computer.

Overall, we provide an open-source software stack, spanning from quantum languages to low-level devices, that allows remote execution of programs on an academically-hosted quantum computer. The open-source software artifact is available to the public [9].

II. BACKGROUND

A. Qiskit

Qiskit is an open-source software development kit that simplifies the ability to compose, run, and analyze quantum circuits and programs [8], [10], [11]. Qiskit is currently the most widely used software stack for quantum cloud computing and is applied to a wide body of commercial and academic research [12], [13]. Recent extensions have added abstractions for entire algorithms plus domain-specific APIs (e.g., optimization, finance, machine learning and chemistry) [14], [15], as well as pulse-level programming [16].

The Qiskit software can be used for simulating circuits on classical devices via Qiskit Aer, as well as interfacing to a suite of IBM Quantum devices through the Qiskit Terra library. The Terra library provides transpilers for circuit optimization and translation to suitable data structures and interfaces. QisDAX, the contribution of this work, decouples the Qiskit Terra abstraction from IBM Quantum backends and instead transpiles down to DAX.

B. Duke ARTIQ Extensions (DAX) Architecture

DAX builds upon the ARTIQ infrastructure developed by M-Labs and NIST [5]. The program flow of ARTIQ and its dedicated FPGA hardware allows for real-time control with nanosecond precision over quantum physics experiments. However, due to ARTIQ’s genericity, quantum computing stacks based on this infrastructure often develop quite monolithic and system-specific control software.

DAX is a software framework that can reduce kernel overhead and increase modularity and portability between ARTIQ experiments [6]. DAX also provides high-level utilities and is currently used as the control system framework at Duke University, the University of Waterloo, and the University of Sydney, among others. As a result of the framework’s growing popularity in academic institutions, libraries such as QisDAX will make the systems more accessible to users.

In the DAX framework, users build modular control software by grouping system functionality into modules and services. Modules are self-contained and control zero or more related devices to perform basic procedures. E.g., a `trap` module may control the voltages applied to ion trap electrodes, in order to change the field shape about the ions. Modules are limited in control to the devices which they contain, i.e., they cannot share devices. Modules are added to a central registry of a system, so they can be found by services.

Services are components that control multiple modules. Any single module may be controlled by multiple services. E.g., a service that loads ions into a trap may control the `trap` module to set electrodes so that they form a suitable trapping

gradient, an `ablation` module that pulses a laser at an ablation target, and a `cw` module that controls the continuous-wave lasers used to ionize and cool the ablated atoms.

DAX clients further increase code reusability. Clients are generic experiments that, at runtime, combine with system-specific code for execution, allowing for high-level code transfer between systems. One such client is `DAX.program`, which implements an `Operation Interface` containing functions for common gate-level quantum operations with explicit timing control. The DAX architecture is depicted in Fig. 1.

`DAX.program-sim` is an addition to `DAX.program` allowing for classical simulation of quantum systems, with its pipeline designed to be identical to the one that runs on quantum hardware [17]. This simulator framework is considered a canonical backend for any program written using `DAX.program` and provides a reliable test bench for programs converted using `QisDAX`.

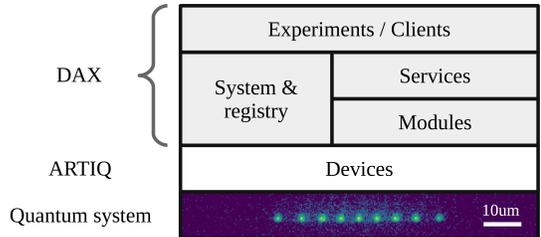


Fig. 1. Schematic overview of how DAX architecture combines with ARTIQ to control a quantum system, in this case a trapped ion device. This figure has been taken from [6].

III. DESIGN

The overall design objective of QisDAX is to provide users the experience of the Qiskit platform, which combines circuit abstractions with result evaluations. This means that the same data structures used by Qiskit should be made available by QisDAX, regardless of the underlying lower levels of the quantum software stack.

A. Software Design Challenges

Both Qiskit and DAX are designed as Python libraries, but they differ in their structure. These differences pose the following challenges:

- 1) A Qiskit program supports heterogeneous backends through providers. DAX, our immediate target, does not package a provider and only targets ARTIQ.
- 2) Qiskit represents circuits as Directed Acyclic Graphs (DAGs). As these DAGs do not translate directly to the DAX representation, we propose a novel, time-sliced approach. DAX provides various scoping constructs to more explicitly define the ordering of instructions within a circuit, indicating whether they may execute sequentially or in parallel. DAGs may be non-planar, but the scoping constructs expressed within a program are required to be planar.
- 3) The results of a DAX program are expressed as a vector of measurement values, many of which may be extending

across multiple channels simultaneously. For the results to be usable by any workflow that analyzes results from a Qiskit program, this representation must be converted to a `Qiskit Result` object representation.

- 4) Resource constraints for hardware controls are not explicit for Qiskit, as circuits operate on virtual qubits. In contrast, DAX programs operate on physical qubits with explicit specification of parallel execution of gate sets using shared resources, e.g., see resources in Sect. IV.

While these constraints are specific to DAX and ARTIQ, the aim of QisDAX is to provide a software layer that can be re-targeted to lower levels of other control stacks for ion traps, or even to control stacks using different a quantum device type such as neutral atoms.

B. Design Solutions

QisDAX provides the following solutions to the above challenges, while considering the design objectives:

- 1) We provide a transpilation component from a Qiskit program to a DAX program while considering the available resource types.
- 2) We provision the required interfaces and objects compatible with Qiskit for heterogeneous quantum/classical processing, namely provider and job abstractions.
- 3) We instantiate the DAX layer with hardware-specific options compatible to the Qiskit program, where a provider can be chosen from (i) the ion trap device or (ii) a simulator instance.
- 4) We facilitate the conversion of job results by transforming the DAX execution results through a component to Qiskit compatible objects.

We subsequently verify the correctness of this translation by validating the generated circuits under simulation via the `DAX.sim` [18] and the `DAX.program-sim` components. We further demonstrate the capability of our approach via circuit execution on trapped-ion quantum computer.

C. Circuit representations

Circuit representations must preserved in their semantics through the vertical layers of the software stack as part of the transpilation process, yet they should be compliant with existing Qiskit inspection and visualization capabilities.

1) *Visualizing a circuit*: Qiskit provides utility functions to display `QuantumCircuit` objects rendered as text, `matplotlib`, or even `LATEX`. The rendered circuit consists of a (time) line for every qubit, with gates as blocks spanning the lines for the qubits they operate on and additional representations for operations such as measuring (see Fig. 2) and barriers. This representation has to be preserved by lower layers, where virtual qubits can be mapped to physical ones.

2) *Qiskit DAG*: A circuit can be represented as a DAG consisting of inputs, outputs and operations as nodes and directed edges that correspond to gates and qubits (see Fig. 3). The depicted information may be enhanced by device-specific details such as operational parallelism, timing constraints and mappings to physical qubits when generating optimal DAX

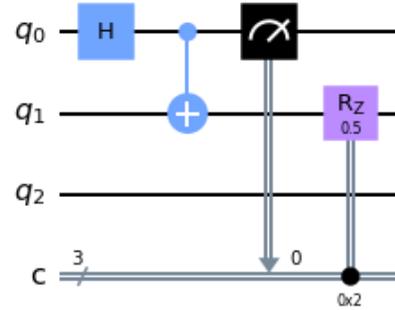


Fig. 2. A Qiskit circuit rendered using `matplotlib` renderer showing a Hadamard gate, a Controlled-X gate, a single qubit measurement to a classical register of size 3, and a Z rotation gate in order from left to right.

representations. On top of optimizations specified for the Qiskit transpiler, QisDAX provides only the trivial layout of mapping the i -th virtual qubit to the i th physical qubit. A cost (in terms of gates) and noise-aware mapping could be integrated as a post-processor at a later time to implement code optimization strategies.

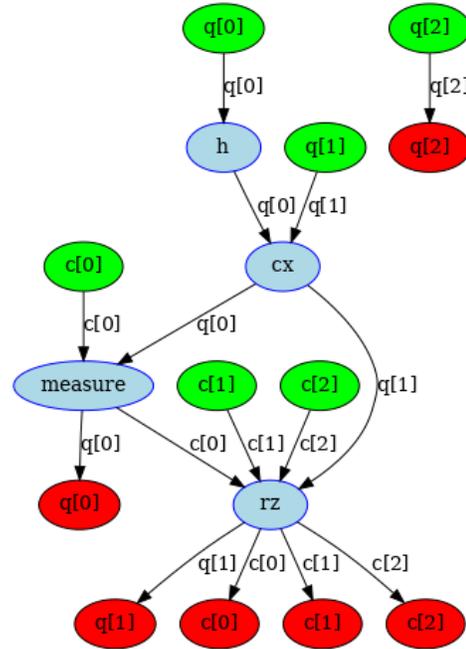


Fig. 3. Qiskit DAG

3) *DAX*: DAX extends the ARTIQ circuit representation as a program, i.e., a sequence of gates using the explicit program scoping constructs with `sequential` and with `parallel` to support the specification of gate parallelism at the level of physical qubits. Each scope may contain multiple instructions or other nested scopes. Instructions and scopes at the root level of a sequential scope are guaranteed to execute in the order they appear. For a parallel scope, the instructions and scopes at the root level may execute in parallel, subject to resource availability. In other words, any subset of gates with logical concurrency in a program can be executed utilizing physical parallelism. In a noise-free environment, this adjust-

ment would always result in the same quantum state regardless of the amount of actual parallelism (from none to all logically concurrent gates in parallel). However, in noisy environments, the result is a trade-off in which a system may be adversely affected by an increase in parallelism while simultaneously benefiting from lower decoherence due to decreased circuit depth.

IV. IMPLEMENTATION

Details specific to the software packages Qiskit, DAX, and, to a lesser extent, ARTIQ, influence implementation choices under the objectives of the QisDAX project.

A. Software Stack

QisDAX serves as a bridge between two projects, Qiskit and DAX. Qiskit serves as the input and surrounding driver program. DAX provides a number of utilities that interface with the lower-level ion-trap quantum hardware and simulators.

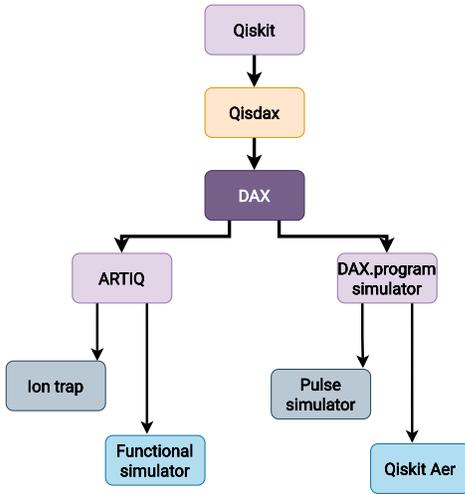


Fig. 4. Software stack

B. QisDAX Components

We ensure maximum interoperability with pre-existing Qiskit programs and minimum refactoring by providing the following utilities:

- **DAXProvider:** A counterpart to the Qiskit IBM [Q] provider, which ordinarily serves as a reference to access IBM Quantum backends, whereas ours refers to an ion trap device.
- **DAXSimulator:** An alternate backend that simulates the results of the DAX program execution transpiled down from Qiskit, by utilizing the DAX program simulator.
- **DAXPrinter:** A backend used for generating the DAX program without executing it. All results are reported in the ground state.
- **DAXArtiq:** A backend for executing circuits on supported quantum hardware. Backend can be configured through a resource configuration file for the network address and destination filesystem of the device controller.

- **DAXJob:** The base class for QisDAX jobs specifying the execution pipeline for circuits. It is also responsible for converting results back to Qiskit-compatible objects.
- **DAXSimJob:** Dispatches circuits to the DAX program simulator. It is derived from DAXJob.
- **DAXPrintJob:** Displays DAX code to stdout. It is also derived from DAXJob.
- **DAXArtiqJob:** Dispatches circuits to the configured ARTIQ-compatible backend. Requires network address information for the backend.
- **qobj_to_dax:** Converts a Qiskit QasmQobj to the equivalent DAX program.

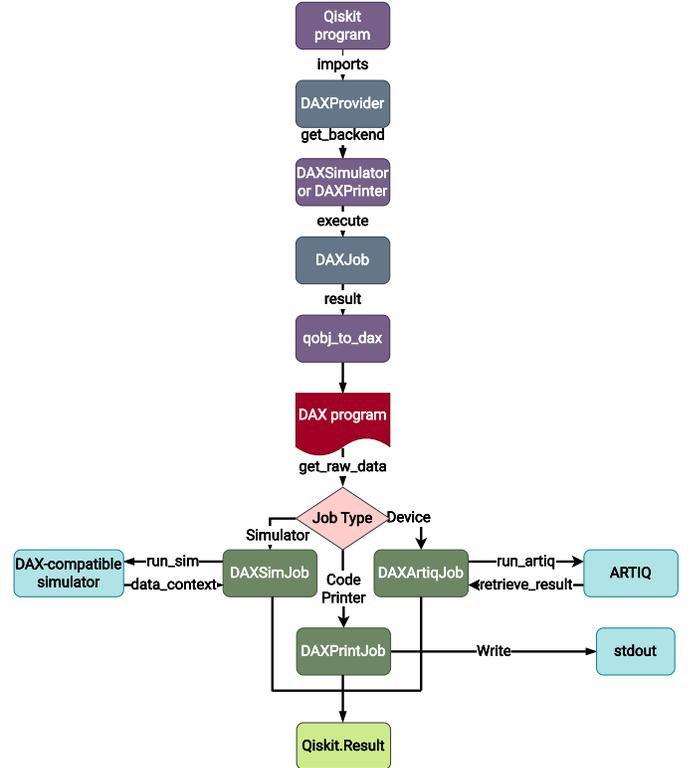


Fig. 5. QisDAX architecture

These utility classes are used to dispatch the circuits to the specified backends, including the circuit definition and all subsequent classical computations, which are handled as if processed by the IBM Quantum backend. As the backend is abstracted, both the input to the backend and the subsequent Qiskit result object do not need to be transformed but are fully compatible (see Fig. 5 for execution stages).

C. Resource Configuration

Resources availability can be configured through a resource specification file, resources.toml, in TOML format. The TOML specification provides an association through key-value pairs. QisDAX supports the following configuration options:

- **total_lasers:** Indicates the total number of lasers available for the trapped ion device. This number is assumed to account only for the lasers realizing gate operations, but

not others used for cooling, measurement, etc. Note that we assume no constraints on the other laser types. This allows flexibility for operations such as measurement, assuming no upper bound for simultaneous measurement operations and future enhancements for handling of mid-circuit measurements or qubit re-initializing.

- `total_mirrors`: Provides the total number of mirrors available to the trapped ion computer. (Note that mirrors are specific to ion traps utilizing Micro-electromechanical systems (MEMS) technology [19]).
- `relative_time`: Comma-separated list of relative times for executing the n -qubit gate, where n is the 1-based index of the timing value in the list.
- `lasers`: The number of lasers required to perform the gate on the circuit. Used in a TOML table for a particular gate.
- `mirrors`: The number of mirrors required to perform the gate on the circuit. Used in a TOML table for a particular gate.

D. Converting to DAX

1) *Restructuring Gates*: We restructure the linear timeline from Qiskit to a list of layers. Each layer in turn is a list of lists for each qubit. We utilize an approach similar to a breadth-first search over the DAG, adding parallel blocks to a sequential root block. Instructions are added to the parallel block spanning the entire circuit as individual sequential blocks for every qubit.

The pseudocode for the algorithm is shown in Algorithm 1: The algorithm uses the following functions:

- `get_qb_indices`: Returns a priority ordering of the qubit indices, ordered as the qubit with the longest remaining depth first. This also takes in account the relative time comparing single qubit and 2 qubit gates.
- `should_add`: Returns a tuple of a boolean and integer. If the first value is True, the gate is to be inserted in the current parallel layer. Note that in our implementation, the first value is always True for the first gate for every qubit in the layer. If the gate is not the first, the gate is added if the difference in depth for the particular qubit in the layer and the maximum depth for any qubit in the layer does not increase, ensuring uniform depth distribution across qubits. Adding the first gate for any qubit is the only exception, which ensures that the layer is not empty,
- `resource_cnt`: Returns a dictionary of the gate resources required to execute a particular layer. Since the gate execution may be concurrent, it is the type-wise (mirrors, lasers, etc.) total of resources across all the gates in a layer.
- `resource_check`: Returns True if the resources required for the layer may be fulfilled by the available resources.

The algorithm works in a step-by-step manner:

- Initialize `total_gates` as the count of all the gates to be scheduled and `next_indices` as the indices of the gates to be scheduled next.

Algorithm 1 QisDAX restructuring algorithm

```

1: procedure GET_PARALLELIZED_LAYERS(instrs, resources)
2:   parallelized_layers  $\leftarrow$  []
3:   while instrs are unvisited do
4:     while resources available and instr queues for all
5:       qbs have not been marked unavailable do
6:       layer  $\leftarrow$  get_next_layer(instrs, resources)
7:       parallelized_layers.append(layer)
8:   return parallelized_layers
9: procedure GET_NEXT_LAYER(instrs, resources)
10:  for qb_index =
11:    get_qb_indices(instrs, next_indices, resources) do
12:    layer  $\leftarrow$  []
13:    for qb  $\in$  1 to qbs do
14:      layer.append([])
15:    sequence  $\leftarrow$  instrs[qb_index]
16:    if instr is last for qb then
17:      mark instr queue as unavailable
18:      continue
19:    instr  $\leftarrow$  sequence[next_indices[qb_index]]
20:    flag  $\leftarrow$  True  $\triangleright$  True if the succeeding loop does not
    break
21:    for participant  $\in$  instr.qbs do
22:      if instr  $\neq$  participant.next then
23:        mark instr queue as unavailable
24:        flag  $\leftarrow$  False
25:        break
26:    if flag then
27:      is_first_gate  $\leftarrow$  instr = layer.first for all participants
28:      should_add, new_width  $\leftarrow$ 
29:        should_add(instr, layer, is_first_gate)
30:      if should_add then
31:        add instr to layer for all participants
32:        resource_cnt  $\leftarrow$  resource_cnt(layer, resources)
33:        resource_check  $\leftarrow$ 
34:          resource_check(resource_cnt, resources)
35:        if resource_check then
36:          for participant  $\in$  instr.qbs do
37:            next_indices[participant] += 1
38:            if new_width > max_width then
39:              max_width  $\leftarrow$  new_width
40:              mark participant layers as unavailable
41:            else if new_width = max_width then
42:              mark participant layers as unavailable
43:          else
44:            Remove instr from all participants
45:            resources unavailable
46:            break
47:          else
48:            mark participant layers as unavailable
49:
50:  return layer

```

- For generating each layer, keep track of the participation of each qubit (`first_gate`), whether the next gate has for each qubit has already been considered for the current layer (`width_checked`), the max depth for any qubit in the current layer (`max_width`), and whether the current layer has exhausted all available resources (`resource_exhausted`).
- Prioritize qubits with longer remaining depths to add to the current layer. Add the next gate for the highest priority qubit to the layer if there are enough resources and if

it is either the first gate for the qubit in the layer or if the difference between the layer depth and depth of the deepest layer does not change.

- iv. Continue adding gates to the layer, keeping track of the depth for each qubit. If the maximum depth for a layer changes, reset the depth tracker for all the layers.
- v. Continue adding layers to the root list until all the gates from the circuit have been included.

2) *Handling Multi-qubit Gates:* With the restructuring approach discussed above, we assume that each qubit has independent gate sequences. However, there may exist intersections in the form of multi-qubit gates. We mitigate this by splitting a layer into sub-layers, with a multi-qubit gate as the latter boundary of the preceding split.

3) *Serializing to DAX Code:* The QisDAX representation is that of a nested list of lists. At each level, we have:

- i. A collection of sub-layer organizer lists. The total time for each qubit in the layer is approximately equal to the other qubits, except for the last layer. They are realized as `with_parallel` scopes. The root context is assumed to be sequential.
- ii. A collection of sub-layer lists. They are realized as `with_sequential` scopes to ensure the relative order before and after a multi-qubit gate.
- iii. A collection of a list of gates for each qubit. They are realized as `with_parallel` scopes, as qubits in a sub layer are independent except for the final gate, which may be multi-qubit.
- iv. A collection of gates for each qubit. They are wrapped in a `with_sequential` scope, executed in the order they appear in the Qiskit circuit.

The DAX program is rendered through a Jinja template. However, the loops do not exist in the template itself. Instead, they are preprocessed and injected in the template as a string,

E. Measurement

DAX supports simultaneous measurement of multiple qubits while maintaining all intermediate measurement results in memory. A measurement extraction from a DAX data context simply consists of a nested list of integer values, one for each qubit channel being measured. When converting to DAX, the register information for the corresponding measurement gates is stored. Each value from the DAX data context can then be mapped to its corresponding Qiskit register in the Qiskit Result object.

V. EXAMPLE

A. Original circuit

As an example, we choose the Simon's algorithm [20] applied to a bitstring of 110 (see Fig. 6 for the quantum circuit).

B. Circuit after Transpilation

We transpile the above circuit via QisDAX with the resource configuration as specified by Listing 1.

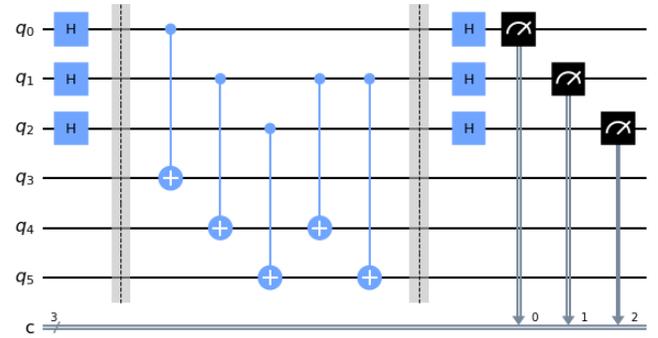


Fig. 6. Circuit for Simon's algorithm

Listing 1 Sample resource configuration for Simon's algorithm

```

1 total_lasers = 5
2 total_mirrors = 5
3 relative_time = '1,2,4'
4
5 [x]
6 lasers = 1
7 mirrors = 1
8
9 [h]
10 lasers = 2
11 mirrors = 2
12
13 [cx]
14 lasers = 2
15 mirrors = 2

```

We then obtain a representation for the circuit as seen in Fig. 7. Each colored box of gates is a layer executed within a `with_parallel` scope. This representation will then be serialized to DAX.

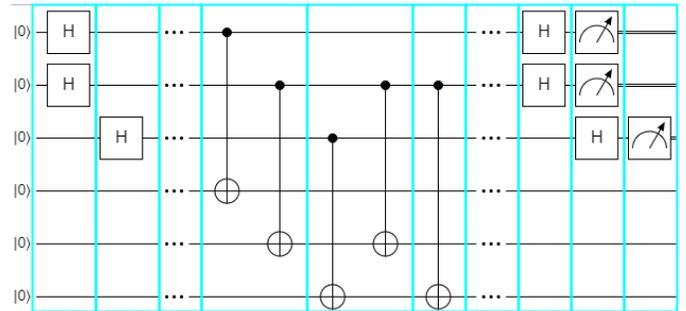


Fig. 7. DAX representation for Simon's algorithm

For the first layer, we schedule the Hadamard gates to be run on the first and second qubits from the top. When executed in parallel, these utilize 4 of the 5 available lasers. We then add the Hadamard on the third qubit in the next layer. Similarly, we obtain layers for the subsequent CNOTs. Barriers always terminate the current layer.

Note that the relative ordering of the gates may change, while preserving logical order. This is a result of the priority ordering of the qubits based on the length of the remaining circuit during processing the subsequent gates. Some of the qubits are non-engaged in some layers, even though gates are assigned to them in the immediately succeeding layer. This is

intentional, and inspecting the resource files reveals that the layers where qubits are inactive may have already exhausted the available resources.

C. Reshaping Raw Data to a Results Object

The data context stores measurement results as a nested list of values. For our example above, it stores results for each `store_measurement` invocation and for the qubit channel(s) specified by the parameter. For x shots and y measurements, we have $x * y$ measurements returned by the data context. Each measurement result returned by the data context is independent of the specified Qiskit register associated to store the result. Hence, we additionally maintain the order of registers in which measurements are to be reported in the Qiskit context, overwriting registers as necessary.

D. Understanding scoping constructs

While the previous example considers a small number of available resources on a simpler circuit, we get optimal parallelization by just scheduling gates using a first-come-first-served policy on each qubit, until we exhaust resources. However, for complicated circuits running on devices with a larger number of available resources, we also consider sub-circuit parallelization. Consider the resource configuration as outlined by Listing 2.

Listing 2 Resource configuration for demonstrating scoping structure

```

1 total_lasers = 20
2 total_mirrors = 20
3 relative_time = '1,4,10'
4
5 [x]
6 lasers = 1
7 mirrors = 1
8
9 [h]
10 lasers = 2
11 mirrors = 2
12
13 [cz]
14 lasers = 4
15 mirrors = 4
16
17 [ccx]
18 lasers = 6
19 mirrors = 6

```

This represents a circuit with a possible scoping arrangement overlaid as colored boxes, as depicted in Fig. 8.

- 1) The red boxes denote the parallel context of a layer, in a root sequential context.
- 2) The blue boxes are sequential contexts to organize the contents of a layer in the parallel context of the layer itself. These contexts are defined groups of qubits that execute completely independently of other sibling contexts in the layer.
- 3) The green boxes are parallel contexts within the parent sequential context (blue). These are used when the execution timelines of multiple qubits concur to execute a

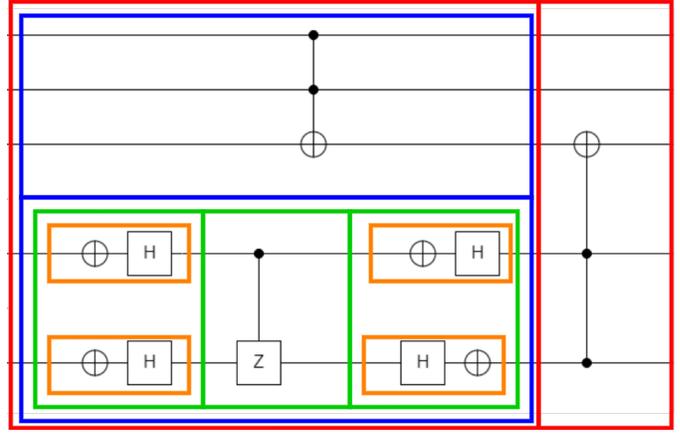


Fig. 8. Circuit with nested scopes

multi-qubit gate, after which they may resume independent execution.

- 4) The orange boxes are sequential contexts within the parent parallel context (green). These contain the gates in order for a single qubit.

We have a 3 qubit Toffoli on the first 3 qubits executing for 10 time units. We may schedule the Hadamard gate, the X gate, the CZ followed by another Hadamard and X gate on the 4th and 5th qubits, in parallel to the Toffoli. Further, for the 4th and 5th qubits, the X and Hadamard gates on either side of the CZ must be sequential. The maximum resource utilization in such a timeline would be when the Toffoli and both Hadamards on either side of the CZ execute concurrently.

While some boxes in Fig. 8 might not demonstrate all child nesting levels (see the top blue box, the horizontally middle green box, or the rightmost red box), it is deliberately drawn for ease of visualization. The implementation will only generate the innermost scope to reduce overhead. This results in the DAX program circuit as seen in Listing 3.

In Fig 8, the first red box from the left corresponds to the context at line 23, while the second red box corresponds to line 41. The top and bottom blue boxes are represented by the contexts at lines 24, and 25, respectively. The green boxes from left to right are represented by the contexts at lines 26, 33 and 34. Finally, each of the orange boxes appear top to bottom in the green boxes in the image exactly as their sequential contexts appear ordered in the generated DAX program.

VI. RESULTS

The QisDAX pipeline was demonstrated on a physical quantum computer and in simulation. The following subsections describe the results from these experiments.

A. Hardware Results

QisDAX was demonstrated on the CRYO-STAQ device, an experimental trapped-ion quantum computing system at Duke University [21]. CRYO-STAQ is designed to be a fully connected 32 qubit system with a cryogenic vacuum chamber. All-to-all connectivity of the qubits is enabled by a multi-channel acousto-optical modulator (AOM). CRYO-STAQ uses

Listing 3 DAX program to demonstrate scopes

```

1  from DAX.program import *
2
3  class QisDaxProgram(DaxProgram, Experiment):
4
5      def build(self):
6          # initialize program information
7
8      def run(self):
9          # Run the kernel
10         self._run()
11
12     @kernel
13     def _run(self):
14         self._qiskit_kernel()
15
16     @kernel
17     def _qiskit_kernel(self):
18         with self.data_context:
19             for _ in range(self._num_ iterations):
20                 self.core.reset()
21                 self.q.prep_0_all()
22
23             with parallel:
24                 self.q.ccx(0, 1, 2)
25                 with sequential:
26                     with parallel:
27                         with sequential:
28                             self.q.x(3)
29                             self.q.h(3)
30                         with sequential:
31                             self.q.x(4)
32                             self.q.h(4)
33                 self.q.cz(3, 4)
34                 with parallel:
35                     with sequential:
36                         self.q.x(3)
37                         self.q.h(3)
38                     with sequential:
39                         self.q.h(4)
40                         self.q.x(4)
41                 self.q.ccx(3, 4, 2)
42         # continues accordingly

```

a Kasli 2.0, from the ARTIQ hardware ecosystem, as the real-time control hardware solution. DAX is used as the control software solution.

To demonstrate the pipeline on CRYO-STAQ, we used QisDAX to remotely execute a series of single-qubit circuits written in Qiskit. Here, QisDAX was configured to use the DAX ARTIQ device backend, which appropriately generates the DAX.program circuit and executes it on a physical system using an ARTIQ based control system. The configuration file associated with this backend allows the user to enter the appropriate credentials required to access the remote system.

We ran a series of single-qubit circuits with increasing number of gates to benchmark the pipeline overhead with increasing circuit depth. This pipeline overhead captures the time it takes QisDAX to convert a Qiskit circuit to a DAX.program circuit, send the circuit to be remotely run on the physical device, and finally retrieve the results back to be returned to the user. It does not include the configurable wait time for the circuit to be executed on the physical device.

Fig. 9 shows the results from this experiment. As we scale the number of gates, and consequently the circuit depth (as

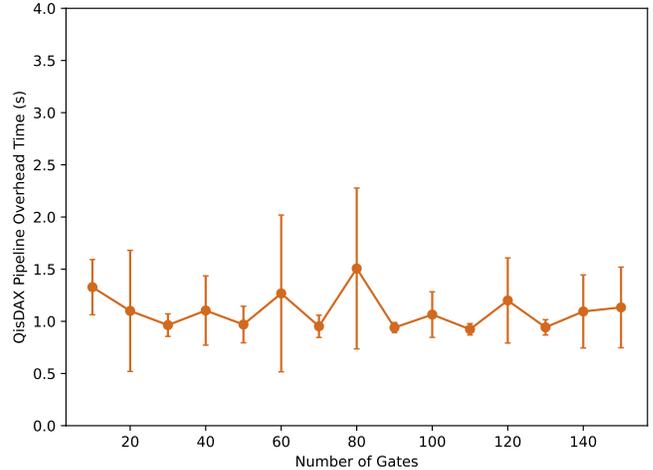


Fig. 9. QisDAX pipeline overhead. Each of these points indicates the average time taken by QisDAX to convert a Qiskit program into a DAX.program circuit, remotely submit it to CRYO-STAQ, and return the final result. Each data point on the plot was averaged over 10 samples, and the error bars indicate the standard deviation. The data shows a near-constant overhead of ~ 1.1 seconds across increasing circuit depths.

it is a single-qubit system), from 10 gates to 150 gates, the overhead time of the pipeline is constant at 1.1 seconds on average, with a standard deviation of 0.16 seconds. This demonstrates the favorable scaling of the pipeline overhead as the number of operations in a system increases.

The QisDAX pipeline also returns metadata about the executed job that the user may be interested in. The default metadata returned is described in Tab. I.

TABLE I
DEFAULT METADATA RETURNED BY QISDAX

Metadata	Description
RID	Unique ID of the job for lookup post execution
Arguments	Describes runtime arguments like name of generated DAX file
Queue time	Time stamp when job was queued
Run start time	Time stamp of when the circuit began execution
Run end time	Time stamp of when the circuit completed execution

B. Simulation Results

A number of benchmark programs were tested to analyze performance. Of these, we feature a subset that includes the Deutsch-Jozsa algorithm [22], Bernstein-Vazirani algorithm [23], Simon’s algorithm [20], Grover’s algorithm [24] and GHZ state generation [10]. All benchmarks are run using 3 qubits for consistency

1) *Simulation using DAX.program simulator:* At the time of writing, CRYO-STAQ was only capable of executing single-qubit circuits. Hence, we complemented these results by running additional benchmarks using the DAX.program-sim simulation [18] architecture. This was accomplished programmatically by selecting the DAX simulator backend provided by the QisDAX pipeline.

The benchmark setup measures the transpilation and execution time, not including the time to convert the results back

TABLE II

MEAN TRANSPILATION TIME [MS] WITH AND WITHOUT RESTRUCTURING

Benchmark	No restructuring	QisDAX	Slowdown
Bernstein-Vazirani algorithm	8104	13106	61.72%
Deutsch-Jozsa algorithm	9649	13917	44.24%
GHZ state	3063	6961	127.25%
Grover's algorithm	17584	21579	22.72%
Simon's algorithm	8400	11989	42.74%
Geometric Mean			50.77%

Listing 4 System configuration for transpilation time analysis

```

1 CPU Model name: Intel(R) Core(TM) i7-4820K CPU @
  ↪ 3.70GHz
2 Operating System: Ubuntu 22.04.2 LTS
3 GPU: NVIDIA Corporation GK104 [GeForce GTX 660
  ↪ OEM]
4 Memory: 7.7GiB
5 Swap: 2.0GiB
6 Disk: 2.0TB

```

to the `Qiskit Result` object. We compare this against a simplistic transpilation with no restructuring, where all instructions are assigned to a single `with_sequential` scope, which results in a longer circuit depth.

These results are shown in Tab. II, measured on a system with the configuration specified by Listing 4. The results depicted in Tab. II show average compilation of times (in milliseconds) of the benchmark circuits, each with 512 shots and repeated 16 times. The data indicate that QisDAX with its circuit transpilation adds an overhead in runtime cost of $\approx 50\%$ (geometric mean) over the simpler approach of scheduling the instructions sequentially with a standard deviation of 0.4. This cost is independent of the available quantum hardware, and may be minimized with advanced, performant hardware.

TABLE III

RUNTIME [μ S] WITH AND WITHOUT RESTRUCTURING, CRYO-STAQ CONFIGURATION

Benchmark	No restructuring	QisDAX	Speedup
Bernstein-Vazirani algorithm	400	345	13.75%
Deutsch-Jozsa algorithm	585	505	13.68%
GHZ state	360	340	5.56%
Grover's algorithm	1465	1315	10.24%
Simon's algorithm	895	520	41.90%
Geometric Mean			13.5%

2) *Statistical simulation:* We analyze the predicted execution time by considering the resource configuration outlined in Listing 5. The configuration option `relative_time` defines estimated execution times for single qubit gate operations as 5 units and two qubit gate operations to be 150 units. The estimated gate times here are based on measurements from the CRYO-STAQ machine, with units as μ s. We consider the runtime for a benchmark as the runtime on the simple path that takes the longest time. Here, runtime includes only circuit execution time, but neither device initialization nor measurement costs. These runtime estimates in Tab. III are

Listing 5 Resource configuration for runtime savings analysis

```

1 total_lasers = 5
2 total_mirrors = 5
3 relative_time = '5,150'
4
5 [id]
6 lasers = 1
7 mirrors = 1
8
9 [x]
10 lasers = 1
11 mirrors = 1
12
13 [z]
14 lasers = 1
15 mirrors = 1
16
17 [h]
18 lasers = 2
19 mirrors = 2
20
21 [cx]
22 lasers = 2
23 mirrors = 2
24
25 [cz]
26 lasers = 2
27 mirrors = 2

```

for a single shot. QisDAX results in a speedup of 13.5% (geometric mean) over a purely sequential approach.

We also compute benchmark times for a commercially available quantum system, the IonQ Aria [25]. For comparability, we keep the available resources the same, only changing the `relative_time` configuration option. Listing 6 highlights the changes made to the CRYO-STAQ configuration in Listing 5 to obtain the configuration for IonQ Aria.

Listing 6 Changes to Listing 5 to obtain IonQ Aria configuration

```

- relative_time = '5,150'
+ relative_time = '135,600'

```

TABLE IV

RUNTIME [μ S] WITH AND WITHOUT RESTRUCTURING, IONQ ARIA CONFIGURATION

Benchmark	No restructuring	QisDAX	Speedup
Bernstein-Vazirani algorithm	3900	2415	38.08%
Deutsch-Jozsa algorithm	5715	3150	44.88%
GHZ state	2820	2280	19.15%
Grover's algorithm	11955	7905	33.88%
Simon's algorithm	6915	3690	46.64%
Geometric Mean			34.89%

The results in Tab. IV demonstrate a higher relative savings. With single-qubit gates making a higher contribution to the circuit time, costs due to unparallelized two-qubit gates may be offset by parallelizing a larger set of single-qubit gates.

Savings are a function of a number of factors, including gate time, order of gate operations in the initial circuit and subcircuit optimizations. The current approach is greedy and does not consider alternative configurations for commutative operations, which may lead to different circuits.

TABLE V
PIPELINE RUNTIME [MS] WITH AND WITHOUT RESTRUCTURING, 10^6
SHOTS, CRYO_STAQ CONFIGURATION

Benchmark	No restructuring	QisDAX	Speedup
Bernstein-Vazirani algorithm	408104	358106	12.25%
Deutsch-Jozsa algorithm	594649	518917	12.74%
GHZ state	363063	346961	4.44%
Grover’s algorithm	1482584	1336579	9.85%
Simon’s algorithm	903399	531989	41.11%
Geometric Mean			12.29%

TABLE VI
PIPELINE RUNTIME [MS] WITH AND WITHOUT RESTRUCTURING, 10^6
SHOTS, IONQ ARIA CONFIGURATION

Benchmark	No restructuring	QisDAX	Speedup
Bernstein-Vazirani algorithm	3908104	2428106	37.87%
Deutsch-Jozsa algorithm	5724649	3163917	44.73%
GHZ state	2823063	2286961	18.99%
Grover’s algorithm	11972584	7926579	33.79%
Simon’s algorithm	6923399	3701989	46.53%
Geometric Mean			34.74%

Using Tab. II and Tab. III, the overall runtime for a pipeline leveraging QisDAX is determined, including time for transpilation followed by circuit execution for a million shots. The results in Tab. V and Tab. VI indicate an overall speedup of $\approx 12\%$ and $\approx 35\%$ (geometric mean) with the CRYO_STAQ and IonQ Aria configurations, respectively. With a reduction in circuit depth over a purely sequential approach by $\approx 36\%$ on average with a standard deviation of 0.05 (table omitted due to space), QisDAX provides considerable speedup with efficient utilization. QisDAX facilitates a trade-off between an increase in circuit execution time and decoherence in noisy quantum devices and a one-time transpilation cost. As the transpilation time is constant, the speedup improves proportional to number of shots. For repeated experiments, we may pre-transpile the circuits ahead-of-time.

These benefits also scale with the available resources in a device capable of parallelism. While our results assume a small number of lasers and mirrors for our benchmarks, increasing the available lasers and mirrors would further decrease overall circuit depth by allowing more qubits to be active per layer.

VII. RELATED WORK

As quantum computing workloads expand towards practical applications, we observe the emergence of multiple competing standards for implementation, both at a high abstraction level [10], [26], [27] and for low level controls [5]. The contexts for higher level quantum programming tools are similar, and efforts to enable interoperability have been forthcoming.

Quantastica [28], the closest related work, generates higher-level programs adhering to different quantum APIs from simpler circuit descriptions by providing proper API contexts in a template-like manner, similar to QisDAX’s translation from QisKit to DAX. However, QisDAX embeds critical circuit analysis within transpilation process to delimit serial

and parallel scopes critical for the vertical stack, and enables platform-aware optimizations with ability to execute on available quantum hardware.

Academic efforts towards open quantum computing have leveraged the availability of an open hardware ecosystem [5]. Attempts at creating an open, platform agnostic reference standard for quantum information have also been made [29]. Similar attempts for control hardware exist [30], with cross platform demonstrations [31]. Orchestrating heterogeneous systems as independent components in an application pipeline has introduced a need for platform-agnostic quantum-classical coupling [32] and verification systems [33].

Yet, these approaches still lack an open-source transpilation tool. QisDAX provides such capability by transpiling Qiskit programs into DAX code. Together with the underlying ARTIQ low-level controls, QisDAX provides the missing link that allows the wealth of quantum programs available in Qiskit to be automatically translated for non-IBM devices, as is demonstrated for the ion-trap CRYO_STAQ device at Duke University.

VIII. CONCLUSION

Interoperability between quantum computing stacks can be facilitated by adopting and integrating modular, open-source components. In this work, we have presented QisDAX, a bridge between two open-source quantum computing frameworks, Qiskit and DAX. QisDAX represents the first open-source, end-to-end, full-stack pipeline for remote submission of quantum programs for trapped ions in an academic setting. Its modular architecture also allows QisDAX to be re-targeted to any other control system, so that in the future it can support a variety of backend implementations, not limited to trapped-ion systems. QisDAX transpilation parallelizes gates wherever possible, maintaining circuit fidelity and result artifacts without developer overhead.

We demonstrate this transpilation procedure using backend implementations for simulators and trapped-ion devices, both local and remote. In doing so, we establish operational capabilities with algorithms from the Qiskit library, which includes parametrized quantum procedures as well as classical result analysis. The modular architecture of QisDAX also allows us to leverage advantages available only to the target system; i.e., we achieved parallel semantics that are trapped-ion specific and not readily available via Qiskit.

Single-qubit timing data from a trapped-ion device shows that the pipeline runtime overhead scales well with increasing circuit depth. A number of benchmark algorithms, run on a functional simulator, allow us to analyze the impact of the parallelization process by measuring the mean transpilation time and decrease in overall circuit depth. We use these results to benchmark runtimes for two trapped-ion systems, CRYO_STAQ and IonQ Aria. Though we incur a one-time transpilation cost, we calculate speedups of 12% and 34% for CRYO_STAQ and IonQ Aria, respectively, in the overall pipeline runtime due to shorter circuit depth, reducing the impact of decoherence and improving efficiency and throughput.

ACKNOWLEDGMENT

The work was funded in part by the National Science Foundation (NSF) projects STAQ Project (PHY-1818914), EPIQC — an NSF Expeditions in computing (CCF-1832377), NSF Quantum Leap Challenge Institute for Robust Quantum Simulation (OMA-2120757), NSF CROSS — Cross-layer Coordination and Optimization for Scalable and Sparse Tensor Networks (CCF-2217020), the Office of the Director of National Intelligence, Intelligence Advanced Research Projects Activity through ARO Contract W911NF-16-1-0082, and the U.S. Department of Energy, Office of Advanced Scientific Computing Research QSCOUT program. Support is also acknowledged from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. The following people contributed to earlier versions of QisDAX: Alexander Allen, Keith Mellendorf, and Quentin Sieredzki.

REFERENCES

- [1] D. F. Chong, Frederic T. and M. Martonosi, “Programming languages and compiler design for realistic quantum hardware,” *Nature*, vol. 549, no. 7671, pp. 180–187, 2017.
- [2] L. Riesebois, X. Fu, A. A. Moueddenne, L. Lao, S. Varsamopoulos, I. Ashraf, J. Van Someren, N. Khammassi, C. G. Almudever, and K. Bertels, “Quantum accelerated computer architectures,” in *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2019, pp. 1–4.
- [3] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete, “Full-stack, real-system quantum computer studies: Architectural comparisons and design insights,” in *Proceedings of the 46th International Symposium on Computer Architecture*, 2019, pp. 527–540.
- [4] S. Bourdeauducq, R. Jördens, P. Zotov, J. Britton, D. Slichter, D. Leibbrandt, D. Allcock, A. Hankin, F. Kermarec, Y. Sionneau, R. Srinivas, T. R. Tan, and J. Bohnet, “Artiq 1.0,” May 2016. [Online]. Available: <https://doi.org/10.5281/zenodo.51303>
- [5] G. Kaspruwicz, P. Kulik, M. Gaska, T. Przywozki, K. Pozniak, J. Jarosinski, J. W. Britton, T. Hartly, C. Balance, W. Zhang *et al.*, “Artiq and sinara: Open software and hardware stacks for quantum physics,” in *Quantum 2.0*. Optica Publishing Group, 2020, pp. QTu8B–14.
- [6] L. Riesebois, B. Bondurant, J. Whitlow, J. Kim, M. Kuzyk, T. Chen, S. Phiri, Y. Wang, C. Fang, A. V. Horn, J. Kim, and K. R. Brown, “Modular software for real-time quantum control systems,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2022, pp. 545–555.
- [7] M. S. ANIS, H. Abraham, AduOffei *et al.*, “Qiskit: An open-source framework for quantum computing,” 2021.
- [8] R. Wille, R. Van Meter, and Y. Naveh, “Ibm’s qiskit tool chain: Working with and developing for real quantum computers,” in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1234–1240.
- [9] “Qisdax, a qiskit to dax compiler (repository),” 2018. [Online]. Available: <https://gitlab.com/fmueller/qisdax/>
- [10]
- [11] 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, “Qiskit backend specifications for openqasm and openpulse experiments,” 2018.
- [12] P. Griffin and R. Sampat, “Quantum computing for supply chain finance,” in *2021 IEEE International Conference on Services Computing (SCC)*, 2021, pp. 456–459.
- [13] R. Semola, L. Moro, D. Bacciu, and E. Prati, “Deep reinforcement learning quantum control on ibmq platforms and qiskit pulse,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2022, pp. 759–762.
- [14] D. J. Egger, I. Hincks, H. Landa, M. Malekakhlagh, A. Parr, D. Puzzuoli, B. Rosand, R. K. Rupesh, M. Treinish, K. Ueda, and C. J. Wood, “Qiskit dynamics,” 2021. [Online]. Available: <https://github.com/Qiskit/qiskit-dynamics>
- [15] T. Q. N. developers and contributors, “Qiskit nature 0.6.0,” Apr. 2023, Qiskit Nature has some code that is included under other licensing. These files have been removed from the zip repository provided here and are only available via Github. See <https://github.com/Qiskit/qiskit-nature#license> for more details. [Online]. Available: <https://doi.org/10.5281/zenodo.7828768>
- [16] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. C. McKay, “Qiskit pulse: programming quantum computers through the cloud with pulses,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044006, aug 2020. [Online]. Available: <https://dx.doi.org/10.1088/2058-9565/aba404>
- [17] A. S. Dalvi, F. Mazurek, L. Riesebois, J. Whitlow, S. Majumder, and K. R. Brown, “Modular architecture for classical simulation of quantum circuits,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2022, pp. 810–812.
- [18] L. Riesebois and K. R. Brown, “Functional simulation of real-time quantum control software,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2022, pp. 535–544.
- [19] Y. Wang, S. Crain, C. Fang, B. Zhang, S. Huang, Q. Liang, P. H. Leung, K. R. Brown, and J. Kim, “High-fidelity two-qubit gates using a microelectromechanical-system-based beam steering system for individual qubit addressing,” *Physical Review Letters*, vol. 125, no. 15, p. 150505, 2020.
- [20] D. R. Simon, “On the power of quantum computation,” *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1474–1483, 1997. [Online]. Available: <https://doi.org/10.1137/S0097539796298637>
- [21] J. Kim, T. Chen, J. Whitlow, S. Phiri, B. Bondurant, M. Kuzyk, S. Crain, K. Brown, and J. Kim, “Hardware design of a trapped-ion quantum computer for software-tailored architecture for quantum co-design (staq) project,” in *Quantum 2.0*. Optical Society of America, 2020, pp. QM6A–2.
- [22] D. Deutsch and R. Jozsa, “Rapid solution of problems by quantum computation,” *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, vol. 439, no. 1907, pp. 553–558, 1992. [Online]. Available: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1992.0167>
- [23] E. Bernstein and U. Vazirani, “Quantum complexity theory,” *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1411–1473, 1997. [Online]. Available: <https://doi.org/10.1137/S0097539796300921>
- [24] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. New York, New York, USA: ACM, 1996, pp. 212–219.
- [25] IonQ, “Ionq aria: Practical performance,” <https://ionq.com/resources/ionq-aria-practical-performance>, Jul. 24, 2022.
- [26] C. Developers, “Cirq,” Dec. 2022, See full list of authors on Github: <https://github.com/quantumlib/Cirq/graphs/contributors>. [Online]. Available: <https://doi.org/10.5281/zenodo.7465577>
- [27] *Q# Language Specification*, Microsoft. [Online]. Available: <https://github.com/microsoft/qsharp-language/tree/main/Specifications/Language#q-language>
- [28] Quantastica, “Quantastica/qconvert-js: Quantum programming language converter.” [Online]. Available: <https://github.com/quantastica/qconvert-js>
- [29] A. Cross, A. Javadi-Abhari, T. Alexander, N. D. Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, P. Sivarajah, J. Smolin, J. M. Gambetta, and B. R. Johnson, “OpenQASM 3: A broader and deeper quantum assembly language,” *ACM Transactions on Quantum Computing*, vol. 3, no. 3, pp. 1–50, sep 2022. [Online]. Available: <https://doi.org/10.1145/3505636>
- [30] 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, “Qiskit backend specifications for OpenQASM and OpenPulse experiments,” *preprint arXiv:1809.03452*, 2018.
- [31] “Amazon braket python sdk,” GitHub, 10 2022. [Online]. Available: <https://github.com/aws/amazon-braket-sdk-python>

- [32] T. M. Mintz, A. J. Mccaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, "Qcor: A language extension specification for the heterogeneous quantum-classical model of computation," 2019.
- [33] A. Adams, E. Pinto, J. Young, C. Herold, A. McCaskey, E. Dumitrescu, and T. M. Conte, "Enabling a programming environment for an experimental ion trap quantum testbed," 2021.