

# QOS: A Quantum Operating System

Emmanouil Giortamis  
emmanouil.giortamis@tum.de  
TU Munich  
Germany

Nathaniel Tornow  
nathaniel.tornow@tum.de  
TU Munich and Leibniz Supercomputing Centre  
Germany

Francisco Romão  
francisco.romao@tum.de  
TU Munich  
Germany

Pramod Bhatotia  
pramod.bhatotia@tum.de  
TU Munich  
Germany

## Abstract

We introduce the Quantum Operating System (QOS), a unified system stack for managing quantum resources while mitigating their inherent limitations, namely their limited and noisy qubits, (temporal and spatial) heterogeneities, and load imbalance. QOS features the *QOS compiler*—a modular and composable compiler for analyzing and optimizing quantum applications to run on small and noisy quantum devices with high performance and configurable overheads. For scalable execution of the optimized applications, we propose the *QOS runtime*—an efficient quantum resource management system that multi-programs and schedules the applications across space and time while achieving high system utilization, low waiting times, and high-quality results.

We evaluate QOS on real quantum devices hosted by IBM, using 7000 real quantum runs of more than 70,000 benchmark instances. We show that the QOS compiler achieves 2.6–456.5× higher quality results, while the QOS runtime further improves the quality by 1.15–9.6× and reduces the waiting times by up to 5× while sacrificing only 1–3% of results quality (or fidelity).

## 1 Introduction

**Quantum Cloud Computing.** Quantum computing promises to solve computationally intractable problems with classical computers [2, 21]. Thanks to remarkable technological advances in materials science and engineering [31, 80], quantum hardware has become a reality in the form of quantum processing units (QPUs) that consist of quantum bits (qubits) [37]. Interestingly, QPUs are now readily available in a quantum-as-a-service fashion offered by all major cloud providers [3, 5, 28, 37].

While quantum hardware is now a reality, the associated quantum software systems are rudimentary. These QPUs face classical OS challenges that our systems community has tackled in the past, including scalability, performance, efficiency, faults (a.k.a. errors), scheduling, and utilization [10]. Unfortunately, no operating system exists to tackle these challenges holistically for modern quantum hardware.

**Fundamental Challenges of QPUs.** A natural tendency would be to treat these QPUs as yet another accelerator class


(e.g., GPU, TPU, FPGA) and manage them as *accelerator as-a-service* to offload compute-intensive tasks. We argue that this approach might be sub-optimal or even flawed!

The reason is that QPUs present fundamentally unique hardware-level challenges that the systems community has not considered and cannot be directly mapped to classical accelerator-oriented computing (we empirically detail these hardware challenges in § 3). In particular, QPUs operate in the NISQ-fashion (Noisy Intermediate-Scale Quantum [64]), leading to a *non-deterministic* computing platform, where even two QPUs with identical qubits exhibit completely different behaviors across space and time [57, 71].

More specifically, QPUs are inherently noisy and small in computational capacity [64], which limits the size of the problems they can solve. Second, the degree of noise differs across QPUs, even of identical architecture and model, making it difficult to decide which QPUs should execute a quantum program without compromising performance [72]. In addition, we can not trivially multi-program multiple quantum programs on the same QPU to increase utilization since QPU qubits can interfere with each other in undesirable and unpredictable ways [52], severely degrading performance [47]. Finally, it is generally impossible to save or copy a quantum program during execution [56], which further limits scheduling opportunities for preemption or resource sharing in general.

**State-of-the-Art of Quantum Software Systems.** The current state of software can be roughly compared to IBM mainframe batch OSes from the 60s, where the QPUs are managed through rudimentary interfaces. Researchers have proposed specialized approaches to address some of the aforementioned OS and QPU challenges individually, for instance, performance [84], multi-programming [17], or scheduling [73]. Unfortunately, these proposed approaches are designed to solve an individual issue, which prevents them from being composed together or with other OS mechanisms to create a holistic software stack. To leverage quantum computing practically, we must address the key challenge of combining such mechanisms in a unified software stack for quantum computing.

**Novelty.** However, designing a unified system stack that supports general OS abstractions while addressing the QPU challenges is not trivial. The system should support cross-stack



**Figure 1.** Example of a typical quantum algorithm (§ 2.1) (a) *Input graph for max-cut.* (b) *The quantum circuit encoding the formulation of max-cut for the graph.* (c) *The result of circuit execution is a probability distribution of bitstrings.* (d) *The result of (c) is interpreted as a max-cut between vertices  $\{0,1,4\}$  and  $\{2,3,5\}$ .*


software mechanisms, from the compiler level for quantum program optimization to the runtime level for QPU resource management. More specifically, we require a modular and extensible compiler infrastructure for increasing execution quality, multi-programming for increased utilization, and scheduling for load balancing, all in the presence of QPU noise and heterogeneities. This way, the system can achieve the cloud users’ goals, i.e., high-quality quantum computation and low waiting times, and the quantum cloud operator’s goals, i.e., QPU resource efficiency and scalability.

#### QOS: A Unified System Stack for Quantum Computing.

We propose QOS, an end-to-end system for holistically tackling quantum computing challenges. QOS provides a unified architecture for supporting compiler and OS mechanisms with pluggable and configurable policies. In QOS, we implement such policies to achieve the aforementioned users’ and operator’s goals. To achieve this, QOS builds on a unified abstraction and comprises two main components:

- **The Qernel Abstraction:** We introduce the Qernel abstraction that acts as a common denominator for the QOS mechanisms to apply their policies (§ 5.1). A Qernel contains the Qernel intermediate representation (QIR) and static and dynamic properties, leveraged by the QOS components to apply their policies.
- **QOS Compiler:** We introduce the QOS compiler (§ 5, 6, 7), an extensible and modular compiler workflow that leverages the QIR and static properties to optimize quantum programs for increased execution quality.
- **QOS Runtime:** We present the QOS runtime (§ 8), a scalable system for QPU resource efficiency. The system offers *automated QPU selection* to abstract heterogeneity away, *multi-programming* to increase QPU utilization, and *load-aware scheduling* to achieve low waiting times while maintaining high execution quality.

We implement QOS in Python by building on the Qiskit framework [65]. We evaluate QOS on IBM’s 27-qubit QPUs [37], using a dataset of more than 7000 quantum runs and 70,000 state-of-the-art quantum benchmark instances used in popular quantum algorithms [42, 67, 89]. Our evaluation shows that the QOS compiler improves the quantum program



**Figure 2.** Technical Foundations (§ 2.2) (a) *The quantum circuit of Figure 1.* (b) *The physical layout of an IBM Falcon QPU.* (c) *The transpiled circuit with the QPU’s noise sources.*

properties by 51% on average, which leads to 2.6–456.5× improvement in the quality of the results, depending on the problem size (§ 9.2). The QOS runtime increases the quality of the results by 1.15–9.6× for the same target utilization (§ 9.4) and reduces the waiting times by 5× while sacrificing at most 3% of the quality of the results (§ 9.5).


## 2 Background

### 2.1 Quantum Computing 101: An Example

Let us understand the basics of quantum computing using the classic max-cut problem. This simple combinatorial optimization problem is expressed in the quantum world as the Quantum Approximate Optimization Algorithm (QAOA) [21]. Figure 1 shows a high-level example of how QAOA solves a max-cut problem of the input graph of (a). To solve it, the problem must be first encoded as a quantum *circuit* (Figure 1 (b)), which consists of quantum bits (qubits) and quantum gates that exhibit quantum mechanical properties. Here, we use as many qubits as the number of nodes of the input graph, where each qubit  $q_i$  corresponds to a graph vertex  $i$ . To change the state of the qubits, we apply quantum gates over time, from left to right. There are two types of gates: 1-qubit gates (e.g., NOT gate) and 2-qubit gates (e.g., XOR gate). Finally, at the end of the circuit, we measure each qubit to read its value (0 or 1), which gives bitstrings as output.

Unlike classical circuits, which operate deterministically, quantum circuits are inherently probabilistic. The reason is that qubits exhibit quantum mechanical properties, such as *superposition*. In the superposition state, the qubit is not 0 or 1, but it is *both* simultaneously (recall Schrodinger’s cat experiment [77]). Therefore, quantum gates also have probabilistic effects; we can’t know the result until the final measurements (i.e., open the box and check the cat’s state). To obtain a meaningful result, we execute the circuit in many trials (“shots”), with each trial providing a specific bitstring from the qubit measurement. The solution of the quantum calculation is, therefore, a probability distribution over all possible bitstrings of the measured qubits (Figure 1 (c)).

In our example, the result of the final execution of the quantum circuit gives a probability distribution that represents the solutions of the max-cut problem. High probability maps to the solution, while low ( $\sim 0$ ) does not represent a solution. Figure 1 (d) shows a solution for our example. It corresponds



**Figure 3. (a) Challenge #1, Fidelity (§ 3.1). Impact of the number of qubits (circuit size) on fidelity. There is an average 98.9% reduction in fidelity from 4 to 24 qubits. (b) Challenge #2, Spatial heterogeneity (§ 3.2) Fidelity of a 12-qubit GHZ circuit on different IBM QPUs. There is a 38% fidelity difference from best to worst QPU.**

to the bitstring with the highest probability, 110010, which means that we have measured 1 for the qubits  $q_0$ ,  $q_1$ , and  $q_4$ ; therefore, a partition contains vertices  $\{0,1,4\}$ .

## 2.2 Technical Foundations

**Execution Model.** The technology and engineering required to build QPUs renders them an expensive resource, thus, QPUs are mainly offered in the cloud as a quantum-as-a-service model [3, 28, 37]. To run quantum programs, users typically write circuit-level code (Figure 2 (a)), which then *transpile* on the QPU to make it executable, send it to the cloud for execution, and finally get the results back. Specifically, the transpilation process performs three key steps: (1) converting the gates of the circuit to the native gate set of the QPU, (2) mapping the logical qubits of the circuit to the physical qubits of the QPU, (3) routing the qubits to the physical qubits with restrictive connectivity by inserting *SWAP* gates. Figure 2 (b) shows the physical layout of an IBM Falcon QPU. Vertices are the physical qubits, and the edges capture their connectivity, i.e., between which qubits we can apply 2-qubit gates. Figure 2 (c) shows the physical circuit after transpilation with the QPU’s noise characteristics, which we detail next.

**QPU Characteristics.** Today’s QPUs are described as noisy intermediate-scale quantum (NISQ) devices [64] since they exhibit low qubit numbers (e.g., up to a few 100s [37]) and are susceptible to hardware and environmental noise. Specifically, when measuring a qubit, there is a chance to read the opposite value, and when applying gates, there is a chance the gate performs a wrong operation [27]. On top of that, when qubits are left idle (no gates applied) for more than a few hundred microseconds, the superposition *decoheres* to the  $|0\rangle$  state [39], similar to resetting a register to 0. Lastly, qubits destructively interfere with each other via *crosstalk* effects [12]. Figure 2 (c) shows qubits  $Q_2$  and  $Q_3$  that influence each other via crosstalk, noisy gates, qubit  $Q_5$  that is left idle for long enough to decohere, and noisy measurements.

**QPU Heterogeneity.** Additionally, QPUs are vastly *heterogeneous* across space and time, unlike classical accelerators. Across space, QPUs vary in terms of technology, e.g., superconducting qubits [28, 37] or trapped ions [35], architectures of the same technology, e.g., Falcon or Osprey superconducting QPUs [37], and noise properties even for the same architecture [27], e.g., two identical QPUs exhibit different noise errors, etc. Across time, the QPUs are *calibrated* regularly to maintain their performance [36, 90, 94], a process that generates *calibration data*. These data quantify the noise errors, and change after each calibration cycle unpredictably.

**Execution Quality.** Lastly, to measure the quality of a circuit execution on NISQ QPUs, we use the **fidelity** metric [22], which measures the similarity between the noisy probability distribution and the ideal probability distribution that noiseless, ideal QPUs can obtain. Fidelity is a number in the  $[0,1]$  range, where a higher fidelity means a better quality result.


## 3 Motivation and Key Ideas

To motivate QOS, we present a set of *unique* challenges that distinguish QPUs from classical accelerators. We categorize our findings into four challenges that must be addressed to improve the practicality of quantum computing: fidelity, utilization, spatial and temporal heterogeneities, and load imbalance. The experimental methodology used is the same for the final system evaluation and is explained in detail in § 9.1.

### 3.1 Fidelity

Executing quantum programs with high fidelity is challenging since QPUs are characterized by relatively small numbers of qubits and noise, which leads to computation errors (§ 2.2). As the number of qubits and gates in a quantum circuit increases, the noise errors accumulate and the overall fidelity decreases.

**Results.** Our results are highlighted in Figure 3 (a). The  $x$  axis depicts the circuit size as the number of qubits while the  $y$  axis shows the fidelity, where higher is better. The experiment is run on the IBM Kolkata 27-qubit QPU. For each increase in



**Figure 4.** (a) Challenge #2, Temporal variance (§ 3.2) Fidelity of a 6-qubit GHZ circuit on IBM Perth, across 120 calibration days. There are 20 pairs of days with more than 5% difference in fidelity. (b) Challenge #3, Utilization (§ 3.3) Maximum utilization achieved on a 27-qubit QPU for nine benchmarks while maintaining at least 0.75 fidelity. The average utilization is 26.3%, and the max is 29.6%. (c) Challenge #4, QPU Load (§ 3.4) Number of pending jobs on different IBM QPUs. The groups separated by vertical red lines indicate QPUs of the same size. There is up to 57× difference in number of jobs between QPUs of the same size.

qubits, the average fidelity decreases, up to 98.9% from 4 to 24 qubits. Moreover, it is physically impossible to run circuits with a size larger than 27 qubits, since we cannot map them.

**Implication.** NISQ devices are limited due to size and noise and, therefore, cannot be practically used for large quantum circuits; either logically, because the circuit doesn’t fit in the device, or the execution results would be convoluted from noise errors, which translates to low fidelity.

**Key Idea #1: Circuit Optimizations:** To increase fidelity, we need a generic optimization infrastructure that transforms circuits into a physically and practically executable size.

### 3.2 Spatial and Temporal Heterogeneity

In the classical domain, two identical CPUs perform similarly for all applications, and at each point in time. In contrast, QPUs exhibit differences in the layout and connectivity of qubits [30] and variations in noise errors even for QPUs of the same model, which leads to spatial performance variance. Moreover, QPUs are calibrated regularly (§ 2.2), and after each calibration, the noise properties change [94]. As a result, the execution fidelity can vary across different calibration cycles, leading to temporal performance variance.

**Results.** Figure 3 (b) shows a 12-qubit GHZ circuit’s fidelity on different IBM QPUs. Fidelity varies across the QPUs, with a maximum difference of 38% from best to worst. Note that all six QPUs are of the same model (Falcon r5.11).

Figure 4 (a) shows a 6-qubit GHZ circuit’s fidelity over 120 calibration days executed on the IBM Perth 7-qubit QPU, where each data point represents a single day’s fidelity. The largest single-day difference in fidelity is 96.5%, and there are 20 instances of a single-day fidelity drop of more than 5%. Note that there is no way of predicting a QPU’s future calibration data to expect such performance differences.

**Implications.** Due to structural differences across QPUs, quantum circuits perform differently across them. Additionally, there is a high degree of temporal performance variance across calibration cycles, as the fidelity might change significantly from day to day with no discernible pattern.

**Key Idea #2: Performance Estimation:** We estimate a circuit’s potential performance on the available QPUs to automatically select the best-performing candidate(s).

### 3.3 Utilization

The fidelity of circuits decreases as their size increases (§ 3.1), and as a result, it becomes more challenging to utilize a QPU effectively. In contrast to the classical domain, where a CPU can be fully utilized, to get high-fidelity results in the quantum domain, we necessarily under-utilize QPUs.

**Results.** Figure 4 (b) shows the maximum utilization of the IBM Kolkata 27-qubit QPU for nine benchmarks while maintaining at least 0.75 fidelity. No benchmark exceeds 30% utilization, while the average is 26.3%. Higher fidelity values would yield even lower utilization and vice-versa.

**Implications.** There is a tradeoff between QPU utilization and performance (fidelity). In general, the lower utilization, the higher fidelity, and vice-versa. In contrast to the classical domain, the tension between these objectives is vastly larger.

**Key Idea #3: Multi-programming:** We spatially multiplex quantum circuits to increase system utilization (also known as *multi-programming* [17]), and when combined with circuit optimizations, it also increases fidelity.

### 3.4 QPU Load Imbalance

The quantum cloud faces QPU load imbalance. The root cause is spatiotemporal heterogeneity (§ 3.2), combined with the manual QPU selection offered by the current quantum cloud model [37]. This leads to users selecting the “best performant” QPU based on empirical or arbitrary metrics [71].



**Results.** Figure 4 (c) shows the average number of pending jobs for different IBM QPUs across October 2023. The groups of QPUs (separated by the red dashed line) have a size of 7, 27, and 127 qubits, respectively. There is a 49 $\times$ , 57 $\times$ , and 5.7 $\times$  maximum load difference across the groups, respectively.

**Implications.** Load imbalance leads to long waiting times for the users and thus, low quality of service. Additionally, there is no 1-1 mapping between the load and performance differences between QPUs. For instance, the 12-qubit GHZ circuit in Figure 3 (b) performs 1.1 $\times$  better on IBM Hanoi than IBM Cairo, yet the former exhibits 57 $\times$  higher load.

**Key Idea #4: Load-aware Scheduling:** We schedule (temporally multiplex) quantum circuits in a load-aware manner to balance the tradeoff between fidelity and waiting times.

## 4 Overview

Quantum computing is characterized by four main challenges that limit its practicality: (1) Execution fidelity is hindered by the small and noisy QPUs. (2) In contrast to classical accelerators, QPUs exhibit vast spatiotemporal heterogeneities, which renders their performance non-deterministic in both dimensions. (3) QPUs are heavily underutilized to give high-fidelity results. (4) QPUs face vast load imbalance, which leads to prolonged waiting times for the users.

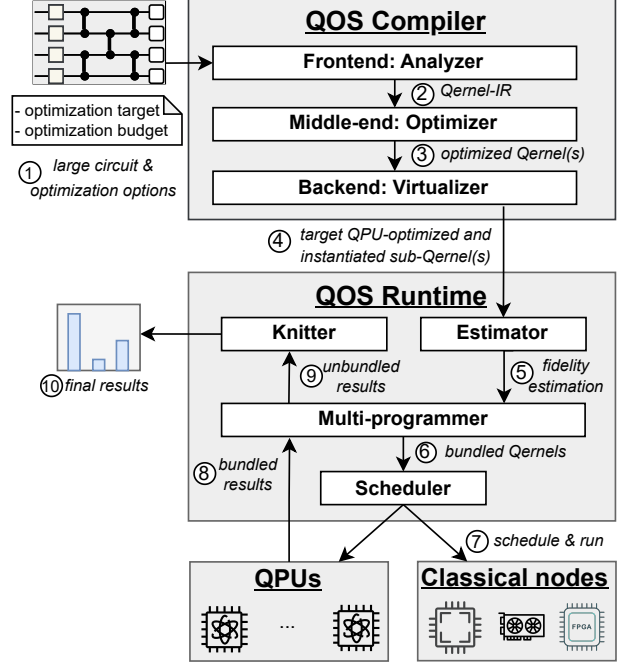
Existing work is narrow-scoped and focuses on tackling one challenge at a time, but unfortunately, there are two main issues with this point solution approach. Firstly, composing the individual mechanisms to address all challenges at once is impossible without a common and unified infrastructure. Secondly, without synergies between the individual mechanisms, it is hard to maximize the objectives of the users, i.e., high fidelity and low waiting times, and the objectives of the quantum cloud operator, i.e., resource efficiency.

To this end, we propose QOS, an end-to-end system that tackles the challenges of quantum computing holistically. QOS strives for three design goals: (1) A unified architecture that supports compiler and OS mechanisms with pluggable policies and tunable configuration for managing the tradeoffs of QC. (2) QOS should enable the execution of large quantum circuits with high fidelity and scale with increasing incoming workloads and additional QPUs. (3) QOS should be resource efficient by achieving high QPU utilization and balancing QPU load to minimize waiting times.

### 4.1 The QOS Architecture

Figure 5 shows the overview of our system’s design. QOS comprises a layered architecture that consists of two main components: the *QOS compiler* (top) and the *QOS runtime* (bottom), which we detail next.

**Qernel Abstraction.** QOS implements a wide range of mechanisms with different abstraction requirements, from the compilation to the execution runtime level. To enable the composability of these mechanisms in a unified architecture, we




**Figure 5.** QOS overview (§ 4): QOS consists of two main components: the QOS compiler (top) and the QOS runtime (bottom). Below QOS lie the QPU devices and classical nodes.

propose the *Qernel* abstraction that acts as a common denominator for the QOS mechanisms to apply their policies.

**QOS Compiler.** We propose the QOS compiler (Figure 5, top), a modular, extensible, and composable compiler infrastructure. It comprises three stages: (1) The frontend of the compiler, the *analyzer* (§ 5.3), accepts quantum circuits and lifts them to the *Qernel* abstraction, generates the intermediate representation (IR), and performs IR analysis passes to generate the IR static properties required by the next stages. (2) The middle-end, the *optimizer* (§ 6), is an extensible and composable set of optimization passes that leverages the IR and static properties to improve the execution fidelity of the quantum circuits with manageable overheads. (3) The backend, the *virtualizer* (§ 7), compiles the optimized *Qernels* for the target QPUs, similar to classical target code generation.

**QOS Runtime.** We propose the QOS runtime (Figure 5, bottom), a system that abstracts away the underlying heterogeneity and balances the tradeoff between the conflicting objectives of the cloud operator (resource efficiency) and the users (high fidelity and low waiting times). The runtime comprises four components: (1) The *estimator* predicts the fidelity of executing the optimized *Qernels* to guide scheduling decisions. (2) The *multi-programmer*, given the estimations, bundles low utilization *Qernels* to increase QPU utilization. (3) The *scheduler* multiplexes and runs the *Qernels* across space and time with the objective to maximize fidelity and minimize waiting times. Finally, (4) the *knitter* post-processes the *Qernel* execution results to return the final result to the user.



**Figure 6.** Compiler Frontend: Analyzer (§ 5). (a) An example quantum circuit of 4 qubits and 5 gates. (b) The Qernel intermediate representation (QIR) (§ 5.1). (c) The refined QIR (§5.1). (d) The Qernel’s static and dynamic properties (§5.2). During compilation, the dynamic properties are void, but they are initialized and used during the runtime.

## 4.2 Execution Workflow

First, users submit a circuit along with their optimization target and budget ①. The former represents the desired post-compilation circuit size, and the latter quantifies the additional overheads the user is willing to pay. The compiler’s frontend lifts the circuit to the Qernel abstraction and generates the IR and its static properties ②. The middle-end optimizes the Qernels through a modular set of passes ③, then the back-end generates the target QPU-optimized Qernels and submits them to the QOS runtime ④. The estimator predicts the fidelity of running the Qernel(s) on the QPUs to guide scheduling ⑤. The multi-programmer bundles Qernels with low utilization and sends them to the scheduler ⑥. The scheduler assigns and runs the bundled Qernels, optimizing for maximal fidelity and minimal waiting times ⑦. After the execution, the bundled results are retrieved by the multi-programmer ⑧ to be unbundled into separate results and are sent to the knitter ⑨. Finally, the knitter post-processes the separated results pass and returns them to the user ⑩.

## 5 Compiler Frontend: Qernel & Analyzer

### 5.1 The Qernel Abstraction

The Qernel is the unified abstraction acting as a common denominator for the QOS components. Specifically, a Qernel contains (1) the graph-based IR used by the QOS compiler and (2) the Qernel properties, which comprises static IR properties and dynamic properties used by the QOS runtime.

**Qernel Intermediate Representation (QIR).** Existing optimization techniques operate at the gate level, analogous to the instruction level in the classical domain. Therefore, we propose a graph-based Qernel Intermediate Representation (QIR) that captures the control flow of a quantum program, similar to the control flow graph of classical programs. By traversing the QIR, the compiler can identify important optimization opportunities, such as pairs of gates that cancel each other (like dead code elimination), gate dependencies (useful for gate scheduling, similar to instruction scheduling), or opportunities to remove hotspot gates (gates that contribute to noise errors in the computation). An example QIR is shown in Figure 6 (b), where the quantum circuit consisting of four qubits and five gates is lifted to the QIR.

Formally, a QIR is a directed acyclic graph (DAG)  $G = (V, E)$ , where  $V$  is the set of gates and every edge  $e_i \in E$  is the qubit the gate acts on. The edges’ directions reflect dependencies  $D = \{(V_i, V_j) \in V \times V\}$  between gates, i.e.,  $V_i$  must be scheduled before  $V_j$ . To identify hotspot nodes, we compute the degree of a node  $deg(V_i)$ , which reflects the number of control flow paths the gate  $V_i$  is part of. In the example of Figure 6, the QIR reveals four layers of gates:  $l_1 : (g_0, g_1)$ ,  $l_2 : g_2$ ,  $l_3 : (g_3, g_4)$ , and  $l_4 : M$ , which means that they have to be scheduled in this order, and the pairs in the same layers are susceptible to crosstalk noise (§ 2). Finally, the gate  $g_2$  is a hotspot node since its degree is 4, the highest in  $V$  (the measurements  $M$  are terminal nodes and do not count).

**Refined QIR.** Various optimizations require a simplified representation that captures only the connectivity between qubits. For instance, the connectivity structure might reveal hotspot qubits [4] that can be removed for fidelity improvement or opportunities for circuit cutting (§ 6). Figure 6 (c) shows the refined QIR, where we can see that  $q_1$  and  $q_2$  are only connected by a single gate, and removing it would split the circuit into two smaller circuits. Formally, a refined QIR is an acyclic, undirected, and weighted graph  $G = (V, E)$ , where  $V$  is the set of qubits and every edge  $e_i \in E$  between two qubits  $V_j, V_k$  has a weight  $w_i \in \mathbb{N}$  that represents the number of gates that act on  $V_j$  and  $V_k$ .

### 5.2 Qernel Properties

For applying the diverse set of its mechanisms, QOS requires data structures that keep up-to-date information about the quantum programs. Such information includes (1) the static properties, which are useful for the compiler, and (2) dynamic properties, which are useful for the runtime.

**Static Properties.** Apart from the IR, optimization passes leverage circuit properties to be more efficient and effective. The properties include the circuit’s size (number of qubits), depth, non-local gates and their types, the number of measurements, and others (Figure 6 (d)). Additionally, we include the features vectors defined in [89] since they are potentially useful for heuristic-based optimizations or regression-based prediction models [68].

**Dynamic Properties.** The Qernel also contains dynamic properties required by the QOS runtime. These include the

Qernel’s execution status (done, failed, running, scheduled), the estimator’s output, i.e., fidelity estimations (§ 8.1), and the final post-processed results (Figure 6 (d)).

### 5.3 Frontend: Analyzer

To discover and leverage optimization opportunities, we first need to perform circuit analysis, similar to classical program analysis. The analyzer transforms a quantum circuit into a Qernel and comprises an extensible set of passes that generate the QIR and static properties of the Qernel, which are then used by the optimizer and subsequently by the runtime.

**QIR Transformation Pass.** The first step in program analysis and optimization is to generate the QIR, implemented by the QIR transformation pass. Figure 6 (a)-(b) shows the generation of the QIR for an example quantum circuit. To generate the QIR, the pass iterates over each logical qubit of the circuit and each gate acting on that qubit. For each such gate, it creates a QIR vertex  $g_i$  and sets the qubits  $q_j, q_k$  the gate acts on as the edges  $e_j, e_k$  of  $g_i$ . By convention, the direction of the edge follows the direction from the control to the target qubits. When reaching measurement operations, it simply adds the terminal nodes  $\textcircled{M}$ . This process is repeated until all circuit qubits and gates are covered.

**QIR Refinement Pass.** To generate the refined QIR (Figure 6 (c)), we implement a transformation pass that traverses the QIR in a depth-first manner (breadth-first is equivalent). For each QIR node visited, i.e., a vertex  $g_i$  with a pair of edges  $q_j, q_k$ , it checks the current refined QIR for existing nodes with the same name. If true, it increments the weight of the edges between the nodes by one. Otherwise, it adds  $q_j$  or  $q_k$  or both as new nodes in the refined QIR and connects them with a weight of one.

**Analysis Passes.** We implement passes that analyze the QIR to identify optimization opportunities, such as gate dependencies (DependencyGraphPass), hotspot nodes (HotspotNodePass), and graph isomorphism (IsIsomorphicPass). We also implement properties passes that traverse the QIR to generate the Qernel static properties (§ 5.2) and comprise the *BasicAnalysisPass* and *SupermarqFeaturesPass*. Specifically, the former generates the key circuit properties while the latter computes the six feature vectors defined in [89], as explained in § 5.2. We show how the optimizer uses the information obtained from these passes in § 6.

## 6 Compiler Middle-end: Optimizer

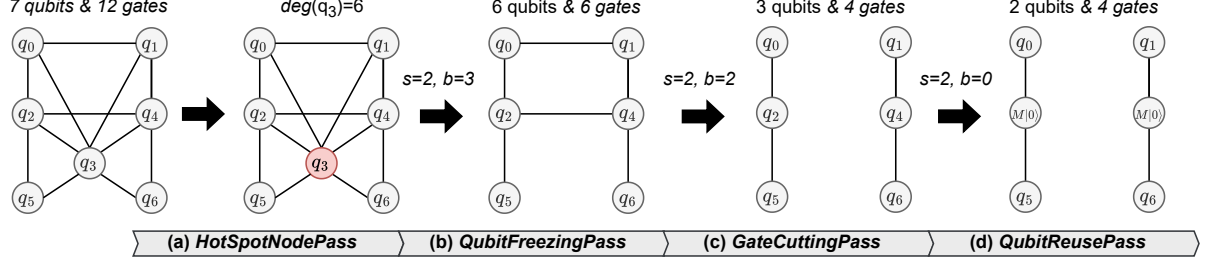
QPU’s comprise up to only a few 100s qubits, which sets the physical limit for circuit size and are noisy, which sets a practical limit to high fidelity execution (§ 3.1). To increase the scalability of circuits that run with high fidelity, we need a modular, extensible, and composable optimizer. Modular to support adding/removing optimization passes or changing their relative order, extensible to add new passes, and composable to chain the optimization improvements of the individual passes.

Such optimization passes are the circuit *compaction* techniques that, as the name suggests, reduce the circuit size, i.e., the number of qubits, rendering it executable on small QPU’s and at the same time, also simplify the circuit structure, i.e., remove noisy gates. Notably, we can compose more than one of these techniques to achieve even better results. The compiler’s middle-end, the optimizer, is a composable pipeline of transformation passes that *compact* the QIR to increase the scalability of quantum circuits running with high execution fidelity. **Challenges.** However, it is not trivial to implement such a pipeline. Currently, there is a plethora of individual compaction techniques that require their own sub-systems to operate, with no common infrastructure to compose them. Additionally, as we will show later, some techniques spawn an exponential number of sub-circuits (Table 1) and, after execution, require post-processing using classical hardware. We pose two questions: (1) How are the (exponentially) spawned circuits from different techniques handled? (2) How can we manage the tradeoff between fidelity improvement and exponential overheads from different techniques?

**Our Approach.** To this end, we design our optimizer with the goals of providing (1) a unified infrastructure for pluggable compaction mechanisms and (2) tunable knobs for configuring the tradeoff between overheads and performance improvement. For (1), we build and compose vastly different compaction techniques on the Qernel abstraction, specifically on the QIR and its refined form. For (2), we provide users with two knobs: the optimization budget (equivalent to optimization level)  $b \in \mathbb{N}$  and the size to reach  $s \in \mathbb{N}$ , which denotes the desired post-optimization QIR size (number of qubits). Since all overheads are exponential and the exponent’s base does not make a practical difference, the single budget knob,  $b$ , suffices.

**QIR Compaction Techniques.** There are two main QIR compaction techniques: circuit divide-and-conquer and qubit reuse. In the former category, the (large) QIR is cut into smaller fragments that are executed on small QPU’s, and the execution results are merged back to a single value. Circuit cutting and knitting [49, 63] belongs to this category. In our optimizer, we implement a pass that automatically cuts the QIR based on qubits (WireCuttingPass) and a pass that cuts the QIR based on gates (GateCuttingPass). To restrict the exponential overheads that scale with the number of cuts, we use the budget  $b$  to cut up to  $b$  times. At each cut location, the pass places a *virtual* gate, which must be later replaced with other gates that simulate the effects of the original (pre-cut) gate (§ 7.1).

We implement another circuit divide-and-conquer technique, also with exponential overheads, namely the *Qubit-FreezingPass*, which is limited to QAOA applications only [4]. At a high level, it removes qubits with significantly more noisy gates than other qubits, i.e., hotspot qubits, along with the gates. This reduces the number of physical qubits required by an underlying (small) QPU and greatly reduces the number of noisy gates. We use the same budget  $b = 3$  to remove the nodes with the highest degrees to restrict the exponential overheads.



**Figure 7.** Optimization workflow (§ 6). The initial refined QIR has 7 qubits and 12 gates. (a) The *HotSpotNodePass* identifies  $q_3$  as a hotspot node with a degree of 6. (b) We optimize the refined QIR with  $s=2, b=3$  by applying the *QubitFreezingPass* which reduces the qubit and gate counts to 6. (c) We apply the *GateCuttingPass* to further remove 2 gates. This gives two fragments of 3 qubits each. (d) We have depleted the budget, so we use the *QubitReusePass* to achieve the goal of  $s=2$ .

Lastly, in the qubit reuse category, we implement the *QubitReusePass*, which “compacts” multiple logical qubits into one to reduce the QIR’s size [19, 34, 75]. This process, however, increases the QIR’s depth; therefore, the tradeoff, in this case, is between QIR size and depth. To restrict the depth increase, we use qubit reuse as a last resort to achieve the user’s size requirement ( $s$ ) or to render the Qernel executable by at least one QPU in the system.

**Optimization Workflow.** Figure 7 shows the default optimization workflow for the refined QIR of a QAOA circuit (§ 2.1) with 7 qubits and 12 gates. The optimizer aims to achieve a maximum QIR size  $s = 2$  with an allowed budget  $b = 3$ . To achieve this, it takes the following steps:

**Step 1:** The optimizer calls the *HotspotNodePass* pass on the refined QIR to find a hotspot node. The pass identifies  $q_3$  as a hotspot with a degree of 6 gates (Figure 7, (a)).

**Step 2:** The optimizer applies the *QubitFreezingPass* to remove  $q_3$  and its gates. The new refined QIR size is 6 qubits with 6 gates. Then, it updates the budget to  $b = b - m$ , where  $m = 1$  is the number of qubits frozen (Figure 7, (b)).

**Step 3:** The optimizer applies either gate or wire cutting. To do so, it first computes the expected cost  $c_{gate}, c_{wire}$ , respectively, to achieve a circuit of size  $s = 2$ , and selects the one with the lowest cost  $c_{min} = \min(c_{gate}, c_{wire})$ . In this case, the cost for gate cutting is lower, so it applies the *GateCuttingPass* on 2 gates and updates the budget to  $b = 0$ . The new refined QIR size is two fragments of 3 qubits and 4 gates each (Figure 7, (c)).

**Step 4:** Since  $b = 0$  but  $s_{QIR} > s$ , the optimizer applies the *QubitReusePass* to achieve  $s = 2$ . The pass identifies qubits  $q_0, q_1$  as reusable and applies measurement and reset to them. The final refined QIR now has two fragments of  $s = 2$  and 4 gates (Figure 7, (d)).

The final optimizer’s output is a Qernel with a 42.8% smaller size and 66% less noisy gates. Note that each of the above passes alone wouldn’t achieve this result.

## 7 Compiler Backend: Virtualizer

The backend stage of the QOS compiler, the virtualizer, generates the final executable Qernels for the underlying runtime, similar to classical compilers that generate the target code.

Optimization Pass	# ISQs Generated	Post-processing
Wire Cutting	$O(4^k - 8^k)$	$O(4^k - 8^k)$
Gate Cutting	$O(6^k)$	$O(6^k)$
Qubit Freezing	$O(2^k)$	$O(1)$

**Table 1.** QOS Virtualizer (§ 7.1). Number of instantiated sub-Qernels (ISQs) generated and post-processing complexity for each optimization pass, as a function of the number of cuts  $k$ .

The virtualizer consists of two stages: (1) the instantiation, which replaces the virtual gates from the cutting optimization passes with the gates that simulate the original ones, and (2) target QPUs transpilation, which translates the high-level gates to the physical QPU gates and performs mapping and routing, as explained in § 2.2

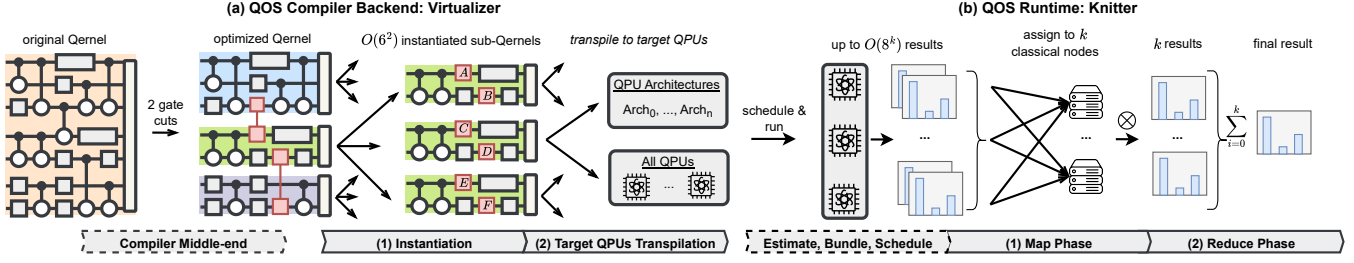
### 7.1 Instantiation

The circuit cutting and knitting passes we describe in § 6 cut a large Qernel into sub-Qernels by analyzing the QIR, identifying optimal cut locations, and then placing virtual gates there. However, to run the sub-Qernels, we must replace the virtual gates with a combination of 1-qubit gates that achieves the same computation as the original Qernel. The mapping between virtual and 1-qubit gates depends on the chosen cutting strategy, i.e., gate or wire cutting (§ 6). We refer to this process as *instantiation*.

The instantiation stage takes as input the optimized sub-Qernels with virtual gates and outputs instantiated sub-Qernels (ISQs) with the 1-qubit gates required to execute them. Similar to the general cutting approach, we implement a generic instantiation mechanism for supporting pluggable mappings from virtual to 1-qubit gates. The mappings depend on the cutting technique (i.e., gate and wire cutting require different mappings) but can also differ for the same technique. For instance, virtual gates from gate cutting can be mapped to different sets of 1-qubit gates that might be more optimal for specific QPU technologies.

By replacing a *single* virtual gate with *multiple* 1-qubit gates, the mapping function generates multiple ISQs that differ only by the 1-qubit gate. Then, replacing the next virtual gate in each ISQ generates even more copies, which leads to





**Figure 8. (a) Compiler Backend: Virtualizer (§ 7) and (b) QOS Runtime: Knitter (§ 8.4). (a) The Virtualizer consists of two stages: (1) Instantiation, which transforms the optimized Qernel to instantiated sub-Qernels (ISQs). This process generates  $O(6^k)$  ISQs for  $k$  cuts (here  $k=2$ ). (2) Target QPUs transpilation, where the ISQs are transpiled to every QPU architecture or every QPU in general. (b) The knitter consists of two stages as well: (1) The map phase and (2) the reduce phase.**

an exponential number of ISQs. Figure 8 (a) shows a Qernel optimized using two gate cuts. The red boxes are the two virtual gates that must be replaced with 1-qubit gates. In this example, the mapping function will replace the first virtual gate with six 1-qubit gates, creating six ISQs. For the next and final virtual gate, each of the six ISQs will produce six more ISQs, totaling 36 ISQs. Generally, in QOS, the exact overheads are  $O(2^k - 8^k)$  for  $k$  cuts for our optimization passes (Table 1).

## 7.2 Target QPUs Transpilation

Following instantiation, the ISQs must be transpiled (§ 2.2) to the target QPUs to be sent to the runtime for scheduling and execution. Since the number of ISQs might be large, depending on the optimization budget  $b$  used at the compiler middle-end (§ 6), we offer two transpilation modes that differ in granularity and overheads: (1) the coarse-grain per QPU architecture and (2) the fine-grain per QPU. We show evaluation results for their transpilation overheads in § 9.2.

**Per QPU-Architecture.** In the first mode, we transpile each ISQ for every type of QPU architecture available in the system. This coarse-grain approach bounds the transpilation overheads because typical quantum cloud providers have a limited number of architectures, e.g., up to five [37]. Since this mode does not scale with the number of QPUs, our experimentation shows that it is suitable for values of budget  $b \geq 5$ , which generate  $10^4 - 10^5$  ISQs.

**Per-QPU.** In the second mode, we transpile each ISQ to each available QPU in the system. This will enable the runtime components to make fine-grained decisions about the fidelity of running the ISQ on any QPU since they will have the exact noise information of this ISQ-QPU pair. The overheads are still bound since QPUs are constant in quantity in commercial clouds, e.g., up to 30 [37], in contrast to classical clouds that scale to thousands of classical nodes. Our experimentation showed that this transpilation mode is viable for  $b < 5$ .

## 8 QOS Runtime

The QOS runtime (Figure 5, bottom) schedules and executes Qernels across space and time in a scalable manner to achieve

the user’s goals, i.e., higher fidelity and lower waiting times, and the cloud operator’s goals, i.e., resource efficiency. It comprises four components, which we detail next. For simplicity, we use the general term *Qernel* throughout this Section.

### 8.1 Estimator

The estimator is responsible for predicting the fidelity of a given Qernel on the underlying QPUs without executing the Qernel. This prediction will be the leading decision factor for the scheduler when assigning the Qernel to a QPU. To achieve this, it computes a *score* for each Qernel-QPU assignment that captures the potential fidelity of that assignment and then uses the scores to rank the assignments. The estimator supports configurable scoring policies that consider (1) the Qernels’ properties generated from the compiler and (2) the QPUs’ calibration data, which are available to quantum cloud providers since they perform the calibration cycles.

For (1), important properties include the number and types of gates, depth, and the number of measurements (§ 5.2). For (2), recall that QPUs are characterized by calibration data that describe the exact error rates of the QPU for that calibration cycle (§ 2.2), specifically, the individual qubit readout errors, the individual gate errors, and the  $T_2$  coherence times. In this work, we implement two scoring policies: a numerical approach for fine-grained control over the estimations and a regression model approach for abstracting away the complexity of estimation.

**Numerical Cost Policy.** This policy estimates execution fidelity by leveraging the target-QPU transpilation output of the compiler backend (§ 7). Target transpilation enables fine-grained fidelity estimation by producing the mapping between logical and physical qubits and the gate (instruction) schedule. The mapping captures the expected readout and gate errors, while the gate schedule captures the order and exact timing that the gates will be applied on the qubits, which reveals the hardware decoherence and crosstalk errors, as explained in § 2.2.

Formally, for each qubit  $q_i$  the readout error is  $e_{r(i)}$ , for each gate  $g_j$  the error is  $e_{g(j)}$ , and the decoherence error is  $e_{d(i)} = 1 - e^{-t/T_2i}$ , where  $t$  is the idle time of the qubit  $q_i$  (no

gates act on it [15]) and  $T_2$  is the decoherence time of  $q_i$ . The crosstalk error between gates  $g_k$  and  $g_l$  is  $e_{ct}(k,l)$ . Putting it all together, the final fidelity score is computed as follows:  $f_{id} = 1 - \prod_{i=0}^N e_{r(i)} e_{d(i)} \prod_{j=0}^M e_{g(j)} \prod_{j=0, k=0}^{M \times M} e_{ct}(j,k)$ , where  $N$  is the circuit’s number of qubits and  $M$  is the number of gates. Since all hardware error information is known at-priori, and quantum errors accumulate multiplicatively, this policy produces high-accuracy estimations, as we show in § 9.3.

**Regression Model Policy.** As discussed in § 2.2, QPU noise errors are accurately measured at each calibration cycle, and their impact on fidelity during quantum computation can be described mathematically. Therefore, we can train a regression model to predict the fidelity of a transpiled Qernel on a possible QPU using the QPU’s calibration data and the Qernel’s static properties as features [68, 73]. Specifically, we use the aforementioned errors we defined in the numerical cost policy as QPU features and the static properties (§ 5.2) as Qernel features. Even simple regression models such as linear regression achieve high prediction accuracy, up to 99%. This policy is simple to use without detailed knowledge of the relationship between errors. However, in QOS, we use the numerical cost policy by default for estimation to have a clear understanding and full control of the process.

## 8.2 Multi-programmer

The size of quantum programs that run with high fidelity is small, leading to QPU underutilization (§ 3.3). To increase QPU utilization, QOS *multi-programs* two or more Qernels, potentially from different users, to run on the same QPU. We refer to this multi-programming as *bundling* the Qernels together. However, trivially bundling Qernels together will deteriorate fidelity because qubits interfere with each other via crosstalk errors (§ 2.2). On top of that, bundled Qernels that run for unequal durations do not necessarily increase utilization since QPU *effective* utilization is measured in space (number of QPU qubits allocated) and time (time qubits are performing actual computation).

For example, a 10-qubit Qernel  $Q_0$  running on a 20-qubit QPU gives 50% spatial utilization. However, assume that  $Q_0$  runs  $3\times$  longer than a 10-qubit Qernel  $Q_1$ . During  $\frac{2}{3}$  of  $Q_0$ ’s runtime, the qubits allocated to  $Q_1$  will be idle, decreasing the effective utilization to only 66%. Recall that it is impossible to schedule more Qernels during  $Q_0$ ’s runtime, unlike in a typical CPU (§ 2.2).

To minimize the fidelity impact and maximize the effective utilization of multi-programming, we utilize configurable Qernel *compatibility* functions that quantify how well-suited are two Qernels to run together.

**Qernel Compatibility Functions.** Compatibility functions measure the crosstalk errors and the effective utilization of a Qernel pair by considering the Qernels’ static properties (§ 5.2). To measure crosstalk effects, we identify pairs of 2-qubit gates that run in parallel during Qernel execution. To

quantify this without running the bundled Qernels, we use the *entanglement ratio* and *parallelism* Qernel static properties, where higher values indicate a higher chance for crosstalk errors [89]. Intuitively, the entanglement ratio captures the proportion of 2-qubit gates over all gates, and parallelism captures how many gates run in parallel per time unit, on average.

To measure the spatial dimension of effective utilization, it suffices to compute the ratio of allocated QPU qubits over the number of QPU qubits. To measure the temporal dimension, we compare the relative duration between two Qernels. The *depth* static property reflects the longest chain of gates that will be executed; therefore, it measures the Qernel’s duration. More technically, we define effective utilization as  $u_{eff} = \frac{N_{Cmax}}{N_{QPU}} * 100 + \sum_{n=1}^k \frac{D_n}{D_{max}} * \frac{N_{Cn}}{N_{QPU}} * 100$ , where  $N_{Cmax}$ ,  $N_{QPU}$  are the number of qubits of the longest Qernel and the QPU, respectively,  $k$  is the number bundled Qernels excluding the longest Qernel, and  $D$  is the depth of the Qernel.


To put everything together formally, we score a possible Qernel pair as follows:  $qc = \alpha u_{eff} + \beta ER_b + \gamma PA_b \mapsto [0,1]$ , where higher is better,  $\alpha + \beta + \gamma = 1$ , and  $c$  denotes *bundled*, i.e.,  $E_b$  is the entanglement ratio of the bundled Qernels. The four variables are tunable to give priorities on different objectives, e.g., prioritize effective utilization or minimize crosstalk. After experimenting and fine-tuning, we found that  $\alpha = 0.25, \beta = 0.25, \gamma = 0.5$ , and  $qc \geq 0.75$  gives balanced results, as we show in § 9.4.

Figure 9 shows an example workflow. The multi-programmer receives three Qernels with three estimations each and identifies *Qernel 0* and *Qernel 2* as a possible pair since their best QPU is the same (*QPU5*) (a). It computes their independent utilization, which is 31% and 37%, respectively, and the combined utilization is under 100% (b). It computes the compatibility score that surpasses the threshold ( $0.9 > 0.75$ ) (c). Next, we detail our multi-programming policies.

**Multi-programming Policies.** QOS supports pluggable multi-programming policies for maximizing effective utilization or minimizing fidelity penalties. In this work, we implement two multi-programming policies; the first is the *fast path* multi-programming, where we can immediately bundle two Qernels if there is no conflict between them, while the second requires re-compilation and re-estimation.

**Restrict Policy.** The restrict policy uses the target QPU transpilation output to bundle Qernels if there is no overlap in their layouts. Practically, this means that for Qernels  $Q_0$  and  $Q_1$ , their logical qubits are mapped to disjoint sets of physical qubits on the QPUs. In that case, the policy bundles the Qernels together, and fidelity loss is minimized through the aforementioned compatibility score.

**Re-evaluation Policy.** This policy is the fallback of the restrict policy. If the Qernel layouts overlap, the two Qernels are transpiled again for the target QPU, and their new fidelity is estimated. If the new fidelity is lower up to a fixed  $\epsilon > 0$



**Figure 9.** QOS multi-programmer example workflow (§ 8.2). (a) Use the estimator’s output to find Qernels with the same best QPU, (b) compute their independent utilization, and (c) compute their compatibility score. If compatible, (d) check for layout overlap, and (e) apply the appropriate multi-programming policy.

value compared to the original fidelities, the bundling is maintained. Otherwise, the multi-programmer selects the next most compatible Qernel pair.

Figure 9 (d) shows the check for layout overlap. In this example, yellow qubits belong to *Qernel 0* and green to *Qernel 1*. On the left, there is no overlap, while on the right, the red qubit is shared between the Qernels. (e) We apply the respective policy (in this example, re-evaluation).

### 8.3 Scheduler

Scheduling quantum programs involves fundamental trade-offs between conflicting objectives; specifically, users want maximal fidelity and minimal waiting times. However, to maximize fidelity, most programs must run on the same subset of QPUs that perform best in a given calibration cycle (§ 3.2). This will lead to large and growing queues on these QPUs, hence long waiting times for the users.

Our scheduler assigns and runs Qernels across space (which QPUs) and time (when) and supports pluggable policies for managing the aforementioned tradeoffs, prioritizing maximal fidelity, minimal waiting times, or a balanced approach. The scheduler assigns Qernels to QPUs based on the fidelity estimations provided by the estimator and the execution time estimations, which we detail next.

**Execution Time Estimation.** To optimize for minimal waiting times, the scheduler must first estimate each Qernel’s execution time and then aggregate the execution time estimations in each QPU’s queue to compute the total waiting times. To estimate the execution time, we iterate the longest path of the QIR of a Qernel (§ 5.1) that corresponds to the longest-duration gate chain and thus defines the Qernel’s execution time. By summing the gate durations of each node in the longest path, we get the Qernel’s total execution time.

**Formula-Based Policy.** Optimizing for conflicting objectives involves comparing two possible solutions (e.g., maximal fidelity vs. minimal waiting times). In the formula-based policy, we use a simple scoring formula (Equation 1) to compare and select between two possible assignments. This formula factors fidelity, waiting time, and utilization to determine which assignment is better, given priorities. The parameters are as follows:  $f_i$ : fidelity of the estimation result  $i$ ,  $t_i$ : waiting time for the QPU from estimation result  $i$ ,  $u_i$ : utilization of the

QPU for estimation result  $i$ ,  $c \in (0,1)$ : a system-defined constant that weighs the fidelity difference between estimations and finally,  $\beta$ : a system-defined constant acting as a weighting factor for utilization difference, balancing system throughput and fidelity. By selecting higher  $c$ , the system prioritizes fidelity over waiting times, and vice versa, and by selecting higher  $\beta$  the system prioritizes utilization over fidelity, and vice versa. By default,  $c = \beta = 0.5$ , which aims for balanced fidelity, waiting times, and utilization.


$$Score = c \frac{f_2 - f_1}{f_1} - (1 - c) \frac{t_2 - t_1}{t_1} + \beta \frac{u_2 - u_1}{u_1} \quad (1)$$

**Genetic Algorithm Policy.** Genetic algorithms excel at optimizing for conflicting objectives by efficiently searching over vast search spaces, and for that, they can be used in the context of QOS. We formulate a multi-objective optimization problem with the conflicting objectives of fidelity vs. waiting times and use the NSGA-II genetic algorithm [18] to solve it. The algorithm creates a Pareto front of possible solutions (schedules), each achieving a different combination of average fidelity and average waiting times. Then, to select one of those schedules, we use the formula described by Equation 1 to score each schedule and select the schedule with the highest score.

### 8.4 Knitter

Following scheduling and execution, the QOS runtime collects the results that are part of the initial circuit submitted by the user. Recall that the circuit is lifted to the QIR (§ 5.3), then optimized through divide-and-conquer techniques that place virtual gates inside the QIR (§ 6), and finally instantiated to replace the virtual gates with 1-qubit gates (§ 7.1). The instantiation process generates up to  $O(8^k)$  instantiated sub-Qernels (ISQs) for a single initial optimized Qernel (Figure 8 (a), Table 1). Finally, the ISQs are bundled with other ISQs, possibly from other users, for increased utilization (§ 8.2).

Therefore, to compute and return the final result to the user is not trivial; we must first *unbundle* the results from multi-programming and then merge the results from ISQs to the original Qernel, a process called *knitting*. The structure of the ISQs and their respective results resembles a tree structure, where the leaf nodes are up to  $O(8^k)$  results, and the root node is the final result. Therefore, we adopt the map-reduce pattern to perform knitting.



**Figure 10.** QOS Compiler (§ 9.2). *Impact of the QOS compiler on the circuit depth and the number of CNOTs. The circuits are optimized using budget  $b=3$ , and we compare against Qiskit (red horizontal line), FrozenQubits [4] and CutQC [85]. There is an average 46%, 38.6%, and 29.4% reduction in circuit depth, and 70.5%, 66% and 56.6% reduction in the number of CNOTs, respectively.*

**Unbundling for Multi-programming.** The results first pass through the multi-programmer to be unbundled. To do this, the multi-programmer keeps a record that maps the initial (solo) Qernel IDs to the new, bundled Qernel ID, as well as the Qernels’ sizes. Therefore, when receiving a new result from a Qernel with an ID  $i$ , it scans the record to find an entry  $i$ , and if found, it splits the probability distribution bitstrings (§ 2) into two parts: the left-most and the right-most bits based on the Qernel sizes. Then, it forwards the unbundled results to the knitter for the map-reduce phases.

**The Map Phase.** To efficiently process a large number of results (up to  $O(8^k)$ ), we follow a divide-and-conquer approach. Specifically, we split the results into  $k$  equal sizes and distribute them to  $k$  classical nodes to be processed in parallel (Figure 8 (b), step (1)). We parallelize across  $k$  to increase data locality and reduce communication overheads since all results for each of the  $k$  cuts will be in the memory of the same node. Locally, each node performs tensor product ( $\otimes$ ) operations on the probability distributions, which are parallelizable across the node’s threads. If available in the node, QOS leverages GPUs or TPUs to accelerate the tensor products. Following this process, the  $k$  nodes output  $k$  intermediate results, ready to be reduced into a single result.

**The Reduce Phase.** QOS selects any of the  $k$  nodes to perform the reduce step. The rest of the nodes send the intermediate results to this node, which performs a thread-parallel sum of  $k$  results. Equivalent to the map phase, the parallel sum can also be executed on GPUs. This produces the final output to be returned to the user (Figure 8 (b), step (2)).

## 9 Evaluation

### 9.1 Experimental Methodology

**Experimental Setup.** We conduct two types of experiments: (1) classical tasks, such as circuit transpilation and trace-based simulations, and quantum tasks (2), which run on *real* QPUs for measuring the circuits’ fidelities.

For (1), we use a server with a 64-core AMD EPYC 7713P processor and 512 GB ECC memory. For (2), we conduct our experiments on IBM Falcon r5.11 QPUs. Unless otherwise noted, we use the IBM Kolkata 27-qubit QPU.

**Framework and Configuration.** We use the Qiskit [65] Python SDK for compiling quantum circuits and running simulations. We compile quantum circuits with the highest

optimization level (3) and run with 8192 shots. Each data point presented in the figures is the median of five runs.

**Benchmarks.** We study QOS on a set of circuits used in state-of-the-art NISQ algorithms, adopted from the 3 benchmark suits of Supermarq [89], MQT-Bench [67] and QASM-Bench [43]. The algorithms’ circuits can be scaled by the number of qubits and depth. Specifically, we study 9 benchmarks: GHZ, W-State, Bernstein Vazirani (BV), Hamiltonian Simulation (HS- $t$ ), Quantum-enhanced Support Vector Machine (QSVM), Two Local Ansatz (TL- $n$ ), Variational Quantum Eigensolver (VQE- $n$ ), and Approximate Optimization Algorithm (QAOA-R/P), these benchmarks cover a wide range of relevant criteria for evaluating QOS.


For the TL and VQE circuits, we use circular and linear entanglement, respectively. The HS, VQE, and TL benchmarks are scalable by their circuit depth with the number of time-steps  $t$  and layers in the *ansatz*  $n$ . The QAOA-R/P circuits are initialized using regular/power-law graphs, respectively, with degree  $d \in \{1, 3\}$ .

**Metrics.** We evaluate the following metrics:

- **Fidelity:** We use the *Hellinger fidelity* as a measure of how close a noisy result is to the desired ground truth of a quantum circuit [22, 33]. The Hellinger fidelity is calculated as 
$$Fidelity(P_{ideal}, P_{noisy}) = \left(1 - H(P_{ideal}, P_{noisy})\right)^2 \mapsto [0, 1],$$
 where  $H$  is the Hellinger distance between two probability distributions, and  $P_{ideal}, P_{noisy}$  are the ideal and noisy probability distributions, respectively.
- **Circuit Properties:** Number of *CNOT* gates and *depth*. When a Qernel contains more than one sub-Qernel, we use the sub-Qernel with the maximum depth, amount of CNOTs, or an average of these two properties.
- **Waiting Time:** The time a circuit spends in a QPU’s queue, waiting for execution, in seconds.
- **Classical Overhead:** The optimization and post-processing overheads (§ 7) of the QOS compiler vs. Qiskit’s transpiler [66].
- **Quantum Overhead:** The number of additional quantum circuits we need to execute per original quantum circuit.

**Baselines.** We evaluate the QOS compiler against Qiskit v0.41, CutQC [85] and FrozenQubits [4]. QOS’s multi-programmer is evaluated against [17]. Regarding QOS scheduler, to the





**Figure 11.** QOS Compiler (§ 9.2). *Impact of the QOS compiler on the circuit fidelity against Qiskit [66], CutQC [85], and FrozenQubits [4]. The circuits are optimized using budget  $b=3$ . There is a mean 2.6 $\times$ , 1.6 $\times$ , and 1.11 $\times$  improvement for 12-qubit circuits, respectively. There is a 456.5 $\times$ , 7.6 $\times$ , and 1.67 $\times$  improvement for circuits of 24 qubits, respectively.*

best of your knowledge, [73] is the only peer-reviewed quantum scheduler, but it doesn’t provide source code or enough technical details to faithfully implement it.

## 9.2 QOS Compiler

**RQ1:** *How well does the QOS compiler improve the fidelity of circuits that run on NISQ QPUs?* We evaluate the performance of the QOS compiler w.r.t the post-optimization properties and fidelity of the circuits while also analyzing the classical and quantum costs of our approach.

**Effect on the Circuit Depth and Number of CNOTs.** In Figure 10, we show the performance of the QOS compiler on the circuits’ depth and number of CNOTs, where we plot the relative difference in post-optimization circuit depth and the number of CNOTs between Qiskit (the red horizontal line) and FrozenQubits [4], CutQC [85], and the QOS compiler. Figures 10 (a) and (c) show that the circuit depth decreases by 46%, 38.6%, and 29.4%, respectively. Figures 10 (b) and (d) show that the number of CNOTs decreases by 70.5%, 66%, and 56.6%, respectively. The improvement in both metrics against the baselines is attributed to the composability of our compiler; the combined effect of circuit compactions (§ 6) achieves better results than standalone techniques.

**Impact on Fidelity.** Figure 11 shows the QOS-optimized circuits’ fidelity against Qiskit [66], CutQC [85], and FrozenQubits [4]. The results show a mean 2.6 $\times$ , 1.6 $\times$ , and 1.11 $\times$  improvement for 12-qubit circuits, respectively, and a 456.5 $\times$ , 7.6 $\times$ , and 1.67 $\times$  improvement for circuits of 24 qubits, respectively. The fidelity improvement is a consequence of lower circuit depths and fewer CNOTs, as shown in Figure 10.

**Classical and Quantum Overheads.** Figure 12 (a) shows the average classical and quantum overheads of the QOS compiler. The classical overhead is 16.6 $\times$  and 2.5 $\times$  for 12 and 24 qubits, respectively, and the quantum overhead is 31.3 $\times$  and 12 $\times$  for 12 and 24 qubits, respectively. However, fidelity improves by 2.6 $\times$  and 456.5 $\times$  for 12 and 24 qubits, respectively; therefore, for larger circuits, the fidelity improvement is worth the cost. **Scalability.** To demonstrate that the QOS Compiler increases the scalability, we run the VQE-1 benchmark on a hypothetical 1000-qubit QPU with one-qubit gate errors of  $10^{-4}$ , two-qubit

gate errors of  $10^{-3}$ , and measurement errors of  $10^{-2}$ . We optimize with budget  $b \in \{0, 1, 4, 8\}$  and report the estimated fidelity. Figure 12 (b) shows that all budget  $b$  values improve the estimated fidelity, with a tradeoff of improvement vs. overheads.

**RQ1 takeaway:** The QOS compiler improves the properties of quantum circuits by 51% on average, leading to an improvement in fidelity of 2.6–456.5 $\times$ , while incurring acceptable classical and quantum overheads.

## 9.3 Estimator

**RQ2:** *How well does QOS’s estimator address spatial and temporal heterogeneities?* We evaluate the estimator’s precision in selecting the top-performing QPU for each benchmark. We establish a baseline using the on-average best-performing machine every calibration day. On the day of the experiment, IBM Auckland was the best-performing machine (also with the highest number of pending jobs).


**Estimator’s Accuracy.** Figure 12 (c) shows the fidelity of the eight benchmarks when run on QPUs selected by the estimator versus when run on the IBM Auckland QPU. The QPU selected for the BV benchmark is Auckland; therefore, we omit this result. For the rest of the benchmarks, the IBM Sherbrooke and Brisbane QPUs were automatically selected. Interestingly, the fidelity is on par or even higher than IBM Auckland, except for only one benchmark, the QAOA-P1.

**RQ2 takeaway:** QOS’s estimator automatically identifies QPUs with higher fidelity than the current standard practice.

## 9.4 Multi-programmer

**RQ3:** *How well does QOS’s multi-programmer increase QPU utilization with minimum fidelity penalties?* We evaluate the impact of the multi-programmer on the fidelity of co-scheduled circuits for certain utilization thresholds.

**Utilization vs. Fidelity.** Figure 13 (a) shows the average fidelity of nine benchmarks with utilization of 30%, 60%, and 88%. The three bars represent: no multi-programming (No



**Figure 12.** QOS Compiler (§ 9.2) and QOS Estimator (§ 8.1). (a) *Compiler: classical and quantum overheads and fidelity improvement as a relative factor to Qiskit. For 24 qubits, the improvement outweighs the overheads.* (b) *Compiler: scalability to a large, hypothetical 1000-qubit QPU. Any budget  $b > 0$  achieves higher quality results than using no optimizations.* (c) *Estimator’s performance: fidelity of IBM Auckland vs. the QPU automatically selected by the estimator.*

*M/P*) refers to large circuits that run solo, baseline multi-programming (*Baseline M/P*) refers to [17], and QOS’s multi-programming approach (*QOS M/P*). There is an average 9.6× improvement in fidelity compared to solo execution and an average 15% (1.15×) improvement compared to the baseline. **Effective Utilization.** The results in Figure 13 (b) show that QOS achieves, on average, a 7.2% higher effective utilization (§ 8.2), with a maximum improvement of 10.1%.

**Fidelity Penalty vs. Solo Execution.** In Figure 13 (c), we evaluate the fidelity penalty of multiprogramming vs. solo circuit execution for utilization of 30%, 60%, and 88%. The fidelity loss is 2%, 9%, and 18%, respectively. The average fidelity loss is 9.6% compared to solo execution, which is in line with previous studies [17, 47]. In the worst case (18%), the fidelity loss is caused by the restrictions in high-quality qubit allocations and the crosstalk errors.

**RQ3 takeaway:** The QOS multi-programmer improves fidelity by 1.15–9.6× and effective utilization by 7.2% compared to the baselines while incurring an acceptable fidelity penalty (< 10%) compared to solo execution.

## 9.5 Scheduler

**RQ4:** *How well does QOS’s scheduler balance fidelity vs. waiting times and balance the load across QPUs?* We evaluate our scheduler by generating a representative workload consisting of a dataset we collected during the development of QOS.

**Dataset Collection.** During our exploration of the motivational challenges (§ 3) and experimentation and evaluation of the QOS components and their policies, we collected a dataset of 70,000 benchmark circuits and more than 7000 job runs in the quantum cloud. We use this dataset to simulate representative workloads, as we detail next.

**Workload Generation.** To generate realistic workloads, we monitored all available QPUs on the IBM Quantum Cloud [37] for ten days in November 2023 to estimate the hourly job arrival rate. The average hourly rate is 1500 jobs per hour and is the baseline system workload for our evaluation.

**Fidelity vs. Waiting Time.** Figure 14 (a) shows the performance of the formula-based scheduling policy. We show the average fidelity and waiting time as the fidelity weight,  $c$ , changes (§ 8.3). A weight of 0.7 achieves  $\sim 5\times$  lower waiting times than full priority of fidelity while sacrificing only  $\sim 2\%$  fidelity. Figure 14 (b) shows the Pareto front of scheduling solutions generated by the genetic algorithm policy. A weight  $c = 0.5$  achieves  $2\times$  lower waiting times with 4% lower fidelity. **QPU Load Balancing.** Figure 14 (c) shows the QPU load as the total runtime each QPU was active, in seconds, for the formula-based policy. All QPUs handle similar loads, with a maximum difference of 15.2%.


**RQ4 takeaway:** QOS scheduler balances the trade-off between waiting times and fidelity by reducing them  $5\times$  and only 2%, respectively, while balancing the load across QPUs.

## 10 Related work


Quantum optimization techniques can be categorized as (1) qubit mapping and routing [6, 44, 50, 51, 58, 61, 81, 86, 88, 93, 96, 98], (2) instruction/pulse scheduling [15, 26, 52, 78, 82, 91, 97], (3) gate synthesis/decomposition [14, 46, 58, 62, 78, 79, 95], (4) execution post-processing and readout improvement [11, 13, 16, 48, 59, 60, 87], and (5) circuit *compaction* [4, 7, 9, 20, 34, 49, 54, 55, 63, 85]. These techniques are implemented standalone without a compiler infrastructure and, thus, are not composable. Instead, the QOS compiler offers a powerful IR that enables incorporating such techniques in a composable manner.

Moreover, application-specific optimizations focus only on specific algorithms to enhance fidelity but lack generality [1, 24, 25, 32, 40, 45, 69, 70, 74, 83, 92]. In contrast, the QOS compiler is a generic approach applicable to all applications.

The state-of-the-art in quantum multi-programming is fairly limited [17, 47] and focuses solely on high-quality mapping, overlooking the systematic selection of compatible programs for utilization or fidelity. Notably, key optimizations from [17] are integrated into the Qiskit transpiler workflow [66], therefore, are already used by the QOS Virtualizer (§ 7.2).



**Figure 13.** QOS Multi-programmer (§ 9.4). **(a)** Impact of multi-programming on fidelity. There is a 9.6× and 1.15× improvement compared to no multi-programming and the baseline, respectively. **(b)** Effective utilization. There is 7.2% higher effective utilization on average. **(c)** Relative fidelity w.r.t. solo circuit execution. There is an average 9.6% drop in fidelity due to QOS’s multi-programming.



**Figure 14.** QOS Scheduler (§ 9.5). **(a)** Formula-based scheduling policy: Average fidelity vs. average waiting times. A fidelity weight  $c = 0.7$  achieves  $\sim 5\times$  lower waiting time for only  $\sim 2\%$  lower fidelity. **(b)** Genetic algorithm policy: it creates a Pareto front of schedules, where a fidelity weight  $c = 0.5$  achieves  $2\times$  lower waiting times for  $\sim 4\%$  fidelity decrease. **(c)** QPU load as the total runtime of each QPU for the formula-based policy. The maximum load difference between any two QPUs is 15.2%.

Lastly, current quantum scheduling methods [8, 73, 76] are limited because they (1) schedule circuits one at a time, (2) neglect QPU utilization, (3) lack fine control over waiting times versus fidelity, and (4) require manual input for final scheduling decisions. Work in the quantum cloud computing area [38, 41, 72] and in quantum serverless [23, 29, 53]; describes quantum cloud characteristics or potential architectures, but QOS is the first end-to-end QPU management system.

## 11 Conclusion

We presented QOS, a system that composes cross-stack OS abstractions to address the challenges of quantum computing holistically. The synergy between compaction techniques, performance estimation, multi-programming, and scheduling systematically explores the tradeoff space associated with quantum. Specifically, QOS achieves up to 456.5× higher fidelity at a 12× overhead cost, up to 9.6× higher fidelity for a target utilization for 9.6% lower fidelity than solo execution, and up to 5× lower waiting times for 2% lower fidelity.

**Contributions.** Our main contributions include:

1. To our knowledge, QOS is the first attempt to combine circuit compaction with quantum resource management to tackle the challenges of QPUs holistically.
2. We leverage the QOS compiler infrastructure to compose optimizations that improve fidelity in a scalable manner, significantly outperforming their individual application (i.e., the current practice).

3. To our knowledge, we are the first to account for and improve both temporal and spatial QPU utilization when multi-programming quantum programs, while mitigating its associated fidelity penalties.
4. Our scheduler balances the inherent tradeoff between fidelity and waiting times, leading to better overall QoS.

## Acknowledgements

We thank Karl Jansen and Stefan Kühn from the Center for Quantum Technology and Applications (CQTA)- Zeuthen for supporting this work by providing access to IBM quantum resources. We also thank Ahmed Darwish and Dmitry Lugovoy for their contributions to this work. Funded by the Bavarian State Ministry of Science and the Arts as part of the Munich Quantum Valley (MQV).

## References

- [1] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. 2020. Circuit Compilation Methodologies for Quantum Approximate Optimization Algorithm. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 215–228. <https://doi.org/10.1109/MICRO50266.2020.00029>
- [2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (2019), 505–510.
- [3] awsQuantum [n. d.]. AWS Bracket. <https://aws.amazon.com/braket/>. Accessed: 2022-04-11.

- [4] Ramin Ayanzadeh, Narges Alavisamani, Poulami Das, and Moinuddin Qureshi. 2023. FrozenQubits: Boosting Fidelity of QAOA by Skipping Hotspot Nodes. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) (ASPLOS 2023). Association for Computing Machinery, New York, NY, USA, 311–324. <https://doi.org/10.1145/3575693.3575741>
- [5] azurequantum [n. d.]. Azure Quantum. <https://azure.microsoft.com/en-us/products/quantum>. Accessed: 2022-04-11.
- [6] Jonathan M. Baker, Andrew Litteken, Casey Duckering, Henry Hoffmann, Hannes Bernien, and Frederic T. Chong. 2021. Exploiting Long-Distance Interactions and Tolerating Atom Loss in Neutral Atom Quantum Architectures. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 818–831. <https://doi.org/10.1109/ISCA52012.2021.00069>
- [7] Luciano Bello, Agata M. Brańczyk, Sergey Bravyi, Almudena Carrera Vazquez, Andrew Eddins, Daniel J. Egger, Bryce Fuller, Julien Gacón, James R. Garrison, Jennifer R. Glick, Tanvi P. Gujarati, Ikko Hamamura, Areeq I. Hasan, Takashi Imamichi, Caleb Johnson, Ieva Liepuoniute, Owen Lockwood, Mario Motta, C. D. Pemmaraju, Pedro Rivero, Max Rossmannek, Travis L. Scholten, Seetharami Seelam, Iskandar Sitdikov, Dharmashankar Subramanian, Wei Tang, and Stefan Woerner. 2023. Circuit Knitting Toolbox. <https://github.com/Qiskit-Extensions/circuit-knitting-toolbox>. <https://doi.org/10.5281/zenodo.7987997>
- [8] Debasmita Bhounik, Ritajit Majumdar, Amit Saha, and Susmita Sur-Kolay. 2023. Distributed Scheduling of Quantum Circuits with Noise and Time Optimization. arXiv:2309.06005 [quant-ph]
- [9] Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin Vechev. 2020. Silq: A High-Level Quantum Language with Safe Uncomputation and Intuitive Semantics. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (PLDI 2020). Association for Computing Machinery, New York, NY, USA, 286–300. <https://doi.org/10.1145/3385412.3386007>
- [10] Sergey Bravyi, Oliver Dial, Jay M Gambetta, Dario Gil, and Zaira Nazario. 2022. The future of quantum computing with superconducting qubits. *Journal of Applied Physics* 132, 16 (2022), 160902.
- [11] Sergey Bravyi, Sarah Sheldon, Abhinav Kandala, David C. McKay, and Jay M. Gambetta. 2021. Mitigating measurement errors in multiqubit experiments. *Phys. Rev. A* 103 (Apr 2021), 042605. Issue 4. <https://doi.org/10.1103/PhysRevA.103.042605>
- [12] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. 2019. Validating quantum computers using randomized model circuits. *Phys. Rev. A* 100 (Sep 2019), 032328. Issue 3. <https://doi.org/10.1103/PhysRevA.100.032328>
- [13] Siddharth Dangwal, Gokul Subramanian Ravi, Poulami Das, Kaitlin N. Smith, Jonathan M. Baker, and Frederic T. Chong. 2023. VarSaw: Application-tailored Measurement Error Mitigation for Variational Quantum Algorithms. arXiv:2306.06027 [quant-ph]
- [14] Poulami Das, Eric Kessler, and Yunong Shi. 2023. The Imitation Game: Leveraging CopyCats for Robust Native Gate Selection in NISQ Programs. In *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 787–801. <https://doi.org/10.1109/HPCA56546.2023.10071025>
- [15] Poulami Das, Swamit Tannu, Siddharth Dangwal, and Moinuddin Qureshi. 2021. ADAPT: Mitigating Idling Errors in Qubits via Adaptive Dynamical Decoupling. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (Virtual Event, Greece) (MICRO '21). Association for Computing Machinery, New York, NY, USA, 950–962. <https://doi.org/10.1145/3466752.3480059>
- [16] Poulami Das, Swamit Tannu, and Moinuddin Qureshi. 2021. JigSaw: Boosting Fidelity of NISQ Programs via Measurement Subsetting. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (Virtual Event, Greece) (MICRO '21). Association for Computing Machinery, New York, NY, USA, 937–949. <https://doi.org/10.1145/3466752.3480044>
- [17] Poulami Das, Swamit S. Tannu, Prashant J. Nair, and Moinuddin Qureshi. 2019. A Case for Multi-Programming Quantum Computers. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Columbus, OH, USA) (MICRO '52). Association for Computing Machinery, New York, NY, USA, 291–303. <https://doi.org/10.1145/3352460.3358287>
- [18] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197. <https://doi.org/10.1109/4235.996017>
- [19] Matthew DeCross, Eli Chertkov, Megan Kohagen, and Michael Foss-Feig. 2022. Qubit-reuse compilation with mid-circuit measurement and reset. arXiv:2210.08039 [quant-ph]
- [20] Yongshan Ding, Xin-Chuan Wu, Adam Holmes, Ash Wiseth, Diana Franklin, Margaret Martonosi, and Frederic T. Chong. 2020. SQUARE: Strategic Quantum Ancilla Reuse for Modular Quantum Programs via Cost-Effective Uncomputation. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. 570–583. <https://doi.org/10.1109/ISCA45697.2020.00054>
- [21] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. arXiv:1411.4028 [quant-ph]
- [22] fidelity-qiskit [n. d.]. Qiskit Hellinger fidelity. [https://qiskit.org/documentation/stubs/qiskit.quantum\\_info.hellinger\\_fidelity.html](https://qiskit.org/documentation/stubs/qiskit.quantum_info.hellinger_fidelity.html). Accessed: 2022-04-11.
- [23] Jose Garcia-Alonso, Javier Rojo, David Valencia, Enrique Moguel, Javier Berrocal, and Juan Manuel Murillo. 2022. Quantum Software as a Service Through a Quantum API Gateway. *IEEE Internet Computing* 26, 1 (Jan 2022), 34–41. <https://doi.org/10.1109/MIC.2021.3132688>
- [24] Pranav Gokhale, Olivia Angiuli, Yongshan Ding, Kaiwen Gui, Teague Tomesh, Martin Suchara, Margaret Martonosi, and Frederic T. Chong. 2019. Minimizing State Preparations in Variational Quantum Eigensolver by Partitioning into Commuting Families. arXiv:1907.13623 [quant-ph]
- [25] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 2019. Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Columbus, OH, USA) (MICRO '52). Association for Computing Machinery, New York, NY, USA, 266–278. <https://doi.org/10.1145/3352460.3358313>
- [26] Pranav Gokhale, Ali Javadi-Abhari, Nathan Earnest, Yunong Shi, and Frederic T. Chong. 2020. Optimized Quantum Compilation for Near-Term Algorithms with OpenPulse. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 186–200. <https://doi.org/10.1109/MICRO50266.2020.00027>
- [27] google-nisq-properties [n. d.]. Quantum Computer Datasheet. <https://quantumai.google/hardware/datasheet/weber.pdf>. Accessed: 2023-07-17.
- [28] googleQuantum [n. d.]. Google Cirq. <https://quantumai.google/cirq>. Accessed: 2022-04-11.
- [29] M Grossi, I. Crippa, A Aita, G Bartoli, V Sammarco, E Picca, N Said, F Tramonto, and F Mattei. 2021. A Serverless Cloud Integration For Quantum Computing. arXiv preprint arXiv:2107.02007 (2021).
- [30] Laszlo Gyongyosi and Sandor Imre. 2019. A Survey on quantum computing technology. *Computer Science Review* 31 (2019), 51–71. <https://doi.org/10.1016/j.cosrev.2018.11.002>
- [31] David Hanneke, JP Home, John D Jost, Jason M Amini, Dietrich Leibfried, and David J Wineland. 2010. Realization of a programmable two-qubit quantum processor. *Nature Physics* 6, 1 (2010), 13–16.
- [32] Tianyi Hao, Kun Liu, and Swamit Tannu. 2023. Enabling High Performance Debugging for Variational Quantum Algorithms Using Compressed Sensing. In *Proceedings of the 50th Annual International*



- Symposium on Computer Architecture* (Orlando, FL, USA) (ISCA '23). Association for Computing Machinery, New York, NY, USA, Article 9, 13 pages. <https://doi.org/10.1145/3579371.3589044>
- [33] Ernst Hellinger. 1909. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik* 1909, 136 (1909), 210–271.
- [34] Fei Hua, Yuwei Jin, Yanhao Chen, Suhas Vittal, Kevin Krsulich, Lev S Bishop, John Lapeyre, Ali Javadi-Abhari, and Eddy Z Zhang. 2023. CaQR: A Compiler-Assisted Approach for Qubit Reuse through Dynamic Circuit. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*. 59–71.
- [35] H. Häffner, C.F. Roos, and R. Blatt. 2008. Quantum computing with trapped ions. *Physics Reports* 469, 4 (2008), 155–203. <https://doi.org/10.1016/j.physrep.2008.09.003>
- [36] ibmq-calibration [n. d.]. IBM Quantum calibration jobs. <https://quantum-computing.ibm.com/admin/docs/admin/calibration-jobs>. Accessed: 2022-04-11.
- [37] ibmQuantum [n. d.]. IBM Quantum. <https://www.ibm.com/quantum-computing/>. Accessed: 2022-04-11.
- [38] Peter J Karalekas, Nikolas A Tezak, Eric C Peterson, Colm A Ryan, Marcus P da Silva, and Robert S Smith. 2020. A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Science and Technology* 5, 2 (mar 2020), 024003. <https://doi.org/10.1088/2058-9565/ab7559>
- [39] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Yu Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, Erik Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, Daniel Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and John M. Martinis. 2018. Fluctuations of Energy-Relaxation Times in Superconducting Qubits. *Phys. Rev. Lett.* 121 (Aug 2018), 090502. Issue 9. <https://doi.org/10.1103/PhysRevLett.121.090502>
- [40] Lingling Lao and Dan E. Browne. 2022. 2QAN: A Quantum Compiler for 2-Local Qubit Hamiltonian Simulation Algorithms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture* (New York, New York) (ISCA '22). Association for Computing Machinery, New York, NY, USA, 351–365. <https://doi.org/10.1145/3470496.3527394>
- [41] Frank Leymann, Johanna Barzen, Michael Falkenthal, Daniel Vietz, Benjamin Weder, and Karoline Wild. 2020. Quantum in the Cloud: Application Potentials and Research Opportunities. *arXiv:2003.06256 [quant-ph]*
- [42] Ang Li, Samuel Stein, Sriram Krishnamoorthy, and James Ang. 2021. QASMBench: A Low-level QASM Benchmark Suite for NISQ Evaluation and Simulation. *arXiv preprint arXiv:2005.13018* (2021).
- [43] Ang Li, Samuel Stein, Sriram Krishnamoorthy, and James Ang. 2023. QASMBench: A Low-Level Quantum Benchmark Suite for NISQ Evaluation and Simulation. *ACM Transactions on Quantum Computing* 4, 2, Article 10 (feb 2023), 26 pages. <https://doi.org/10.1145/3550488>
- [44] Gushu Li, Yufei Ding, and Yuan Xie. 2019. Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (ASPLOS '19). Association for Computing Machinery, New York, NY, USA, 1001–1014. <https://doi.org/10.1145/3297858.3304023>
- [45] Gushu Li, Anbang Wu, Yunong Shi, Ali Javadi-Abhari, Yufei Ding, and Yuan Xie. 2022. Paulihedral: A Generalized Block-Wise Compiler Optimization Framework for Quantum Simulation Kernels. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (ASPLOS '22). Association for Computing Machinery, New York, NY, USA, 554–569. <https://doi.org/10.1145/3503222.3507715>
- [46] Andrew Litteken, Lennart Maximilian Seifert, Jason D. Chadwick, Natalia Nottingham, Tanay Roy, Ziqian Li, David Schuster, Frederic T. Chong, and Jonathan M. Baker. 2023. Dancing the Quantum Waltz: Compiling Three-Qubit Gates on Four Level Architectures. In *Proceedings of the 50th Annual International Symposium on Computer Architecture* (Orlando, FL, USA) (ISCA '23). Association for Computing Machinery, New York, NY, USA, Article 71, 14 pages. <https://doi.org/10.1145/3579371.3589106>
- [47] Lei Liu and Xinglei Dou. 2021. QuCloud: A New Qubit Mapping Mechanism for Multi-programming Quantum Computing in Cloud Environment. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 167–178. <https://doi.org/10.1109/HPCA51647.2021.00024>
- [48] Filip B Maciejewski, Zoltán Zimborás, and Michał Oszmaniec. 2020. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum* 4 (2020), 257.
- [49] Kosuke Mitarai and Keisuke Fujii. 2021. Constructing a virtual two-qubit gate by sampling single-qubit operations. *New Journal of Physics* 23, 2 (2021), 023021.
- [50] Abtin Molavi, Amanda Xu, Martin Diges, Lauren Pick, Swamit Tannu, and Aws Albarghouthi. 2022. Qubit Mapping and Routing via MaxSAT. In *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1078–1091. <https://doi.org/10.1109/MICRO56248.2022.00077>
- [51] Prakash Murali, Jonathan M. Baker, Ali Javadi-Abhari, Frederic T. Chong, and Margaret Martonosi. 2019. Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (ASPLOS '19). Association for Computing Machinery, New York, NY, USA, 1015–1029. <https://doi.org/10.1145/3297858.3304075>
- [52] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. 2020. Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (ASPLOS '20). Association for Computing Machinery, New York, NY, USA, 1001–1016. <https://doi.org/10.1145/3373376.3378477>
- [53] Hoa T. Nguyen, Muhammad Usman, and Rajkumar Buyya. 2022. QaaS: A Serverless Function-as-a-Service Framework for Quantum Computing. *arXiv:2205.14845 [quant-ph]*
- [54] Alexandru Paler, Robert Wille, and Simon J. Devitt. 2016. Wire recycling for quantum circuit optimization. *Phys. Rev. A* 94 (Oct 2016), 042337. Issue 4. <https://doi.org/10.1103/PhysRevA.94.042337>
- [55] Anouk Paradis, Benjamin Bichsel, Samuel Steffen, and Martin Vechev. 2021. Unqomp: Synthesizing Uncomputation in Quantum Circuits. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (Virtual, Canada) (PLDI 2021). Association for Computing Machinery, New York, NY, USA, 222–236. <https://doi.org/10.1145/3453483.3454040>
- [56] James L Park. 1970. The concept of transition in quantum mechanics. *Foundations of physics* 1, 1 (1970), 23–33.
- [57] Tirthak Patel, Abhay Potharaju, Baolin Li, Rohan Basu Roy, and Devesh Tiwari. 2020. Experimental Evaluation of NISQ Quantum Computers: Error Measurement, Characterization, and Implications. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–15. <https://doi.org/10.1109/SC41405.2020.00050>
- [58] Tirthak Patel, Daniel Silver, and Devesh Tiwari. 2022. Geyser: A Compilation Framework for Quantum Computing with Neutral Atoms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture* (New York, New York) (ISCA '22). Association for Computing Machinery, New York, NY, USA, 383–395. <https://doi.org/10.1145/3470496.3527428>
- [59] Tirthak Patel and Devesh Tiwari. 2020. DisQ: A Novel Quantum Output State Classification Method on IBM Quantum Computers Using Openpulse. In *Proceedings of the 39th International Conference on Computer-Aided Design* (Virtual Event, USA) (ICCAD '20). Association

- for Computing Machinery, New York, NY, USA, Article 139, 9 pages. <https://doi.org/10.1145/3400302.3415619>
- [60] Tirthak Patel and Devesh Tiwari. 2020. VERITAS: Accurately Estimating the Correct Output on Noisy Intermediate-Scale Quantum Computers. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–16. <https://doi.org/10.1109/SC41405.2020.00019>
- [61] Tirthak Patel and Devesh Tiwari. 2021. Qraft: Reverse Your Quantum Circuit and Know the Correct Program Output. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Virtual, USA) (ASPLOS '21). Association for Computing Machinery, New York, NY, USA, 443–455. <https://doi.org/10.1145/3445814.3446743>
- [62] Tirthak Patel, Ed Younis, Costin Iancu, Wibe de Jong, and Devesh Tiwari. 2022. Quest: systematically approximating quantum circuits for higher output fidelity. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 514–528.
- [63] Tianyi Peng, Aram W Harrow, Maris Ozols, and Xiaodi Wu. 2020. Simulating large quantum circuits on a small quantum computer. *Physical Review Letters* 125, 15 (2020), 150504.
- [64] John Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum* 2 (aug 2018), 79. <https://doi.org/10.22331/q-2018-08-06-79>
- [65] Qiskit contributors. 2023. Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2573505>
- [66] qiskit-transpiler [n. d.]. Qiskit Transpiler. <https://qiskit.org/documentation/apidoc/transpiler.html>. Accessed: 2022-06-09.
- [67] Nils Quetschlich, Lukas Burgholzer, and Robert Wille. 2022. MQT Bench: Benchmarking software and design automation tools for quantum computing. *arXiv preprint arXiv:2204.13719* (2022).
- [68] N. Quetschlich, L. Burgholzer, and R. Wille. 2023. MQT Predictor: Automatic Device Selection with Device-Specific Circuit Compilation for Quantum Computing. *arXiv:2305.02337*
- [69] Gokul Subramanian Ravi, Pranav Gokhale, Yi Ding, William Kirby, Kaitlin Smith, Jonathan M. Baker, Peter J. Love, Henry Hoffmann, Kenneth R. Brown, and Frederic T. Chong. 2022. CAFQA: A Classical Simulation Bootstrap for Variational Quantum Algorithms. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Vancouver, BC, Canada) (ASPLOS 2023). Association for Computing Machinery, New York, NY, USA, 15–29. <https://doi.org/10.1145/3567955.3567958>
- [70] Gokul Subramanian Ravi, Kaitlin Smith, Jonathan M. Baker, Tejas Kannan, Nathan Earnest, Ali Javadi-Abhari, Henry Hoffmann, and Frederic T. Chong. 2023. Navigating the Dynamic Noise Landscape of Variational Quantum Algorithms with QISMET. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) (ASPLOS 2023). Association for Computing Machinery, New York, NY, USA, 515–529. <https://doi.org/10.1145/3575693.3575739>
- [71] Gokul Subramanian Ravi, Kaitlin N. Smith, Pranav Gokhale, and Frederic T. Chong. 2021. Quantum Computing in the Cloud: Analyzing job and machine characteristics. In *2021 IEEE International Symposium on Workload Characterization (IISWC)*. 39–50. <https://doi.org/10.1109/IISWC53511.2021.00015>
- [72] Gokul Subramanian Ravi, Kaitlin N. Smith, Pranav Gokhale, and Frederic T. Chong. 2022. Quantum Computing in the Cloud: Analyzing job and machine characteristics. *Archive* (2022). *arXiv:2203.13121* [quant-ph]
- [73] Gokul Subramanian Ravi, Kaitlin N. Smith, Prakash Murali, and Frederic T. Chong. 2021. Adaptive job and resource management for the growing quantum cloud. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 301–312. <https://doi.org/10.1109/QCE52317.2021.00047>
- [74] Saloni Resch, Anthony Gutierrez, Joon Suk Huh, Srikant Bhargava, Yasuko Eckert, Gabriel Loh, Mark Oskin, and Swamit Tannu. 2021. Accelerating Variational Quantum Algorithms Using Circuit Concurrency. *arXiv:2109.01714* [cs.ET]
- [75] Movahhed Sadeghi, Soheil Khadirsharbiyani, and Mahmut Taylan Kandemir. 2022. Quantum Circuit Resizing. *arXiv:2301.00720* [cs.ET]
- [76] Marie Salm, Johanna Barzen, Frank Leymann, and Benjamin Weder. 2022. Prioritization of Compiled Quantum Circuits for Different Quantum Computers. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 1258–1265. <https://doi.org/10.1109/SANER53432.2022.00150>
- [77] schrodinger [n. d.]. Schrödinger's cat. [https://en.wikipedia.org/wiki/Schr%C3%B6dinger%27s\\_cat](https://en.wikipedia.org/wiki/Schr%C3%B6dinger%27s_cat). Accessed: 2022-04-11.
- [78] Yunong Shi, Nelson Leung, Pranav Gokhale, Zane Rossi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 2019. Optimized Compilation of Aggregated Instructions for Realistic Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (ASPLOS '19). Association for Computing Machinery, New York, NY, USA, 1031–1044. <https://doi.org/10.1145/3297858.3304018>
- [79] Yunong Shi, Nelson Leung, Pranav Gokhale, Zane Rossi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 2019. Optimized Compilation of Aggregated Instructions for Realistic Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (ASPLOS '19). Association for Computing Machinery, New York, NY, USA, 1031–1044. <https://doi.org/10.1145/3297858.3304018>
- [80] Irfan Siddiqi. 2021. Engineering high-coherence superconducting qubits. *Nature Reviews Materials* 6, 10 (2021), 875–891.
- [81] Marcos Yukio Siraichi, Vinicius Fernandes dos Santos, Caroline Collange, and Fernando Magno Quintao Pereira. 2018. Qubit Allocation. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization* (Vienna, Austria) (CGO 2018). Association for Computing Machinery, New York, NY, USA, 113–125. <https://doi.org/10.1145/3168822>
- [82] Kaitlin N. Smith, Gokul Subramanian Ravi, Prakash Murali, Jonathan M. Baker, Nathan Earnest, Ali Javadi-Cabbari, and Frederic T. Chong. 2022. TimeStitch: Exploiting Slack to Mitigate Decoherence in Quantum Circuits. *ACM Transactions on Quantum Computing* 4, 1, Article 8 (oct 2022), 27 pages. <https://doi.org/10.1145/3548778>
- [83] Samuel Stein, Nathan Wiebe, Yufei Ding, Peng Bo, Karol Kowalski, Nathan Baker, James Ang, and Ang Li. 2022. EQC: Ensembled Quantum Computing for Variational Quantum Algorithms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture* (New York, New York) (ISCA '22). Association for Computing Machinery, New York, NY, USA, 59–71. <https://doi.org/10.1145/3470496.3527434>
- [84] Wei Tang, Teague Tomesh, Martin Suchara, Jeffrey Larson, and Margaret Martonosi. 2021. CutQC: Using Small Quantum Computers for Large Quantum Circuit Evaluations. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Virtual, USA) (ASPLOS '21). Association for Computing Machinery, New York, NY, USA, 473–486. <https://doi.org/10.1145/3445814.3446758>
- [85] Wei Tang, Teague Tomesh, Martin Suchara, Jeffrey Larson, and Margaret Martonosi. 2021. Cutqc: using small quantum computers for large quantum circuit evaluations. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 473–486.
- [86] Swamit S. Tannu and Moinuddin Qureshi. 2019. Ensemble of Diverse Mappings: Improving Reliability of Quantum Computers by Orchestrating Dissimilar Mistakes. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Columbus,

- OH, USA) (*MICRO '52*). Association for Computing Machinery, New York, NY, USA, 253–265. <https://doi.org/10.1145/3352460.3358257>
- [87] Swamit S Tannu and Moinuddin K Qureshi. 2019. Mitigating measurement errors in quantum computers by exploiting state-dependent bias. In *Proceedings of the 52nd annual IEEE/ACM international symposium on microarchitecture*. 279–290.
- [88] Swamit S. Tannu and Moinuddin K. Qureshi. 2019. Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (*ASPLOS '19*). Association for Computing Machinery, New York, NY, USA, 987–999. <https://doi.org/10.1145/3297858.3304007>
- [89] Teague Tomesh, Pranav Gokhale, Victory Omole, Gokul Subramanian Ravi, Kaitlin N Smith, Joshua Vizslai, Xin-Chuan Wu, Nikos Hardavellas, Margaret R Martonosi, and Frederic T Chong. 2022. Supermarq: A scalable quantum benchmark suite. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 587–603.
- [90] Caroline Tornow, Naoki Kanazawa, William E. Shanks, and Daniel J. Egger. 2022. Minimum Quantum Run-Time Characterization and Calibration via Restless Measurements with Dynamic Repetition Rates. *Phys. Rev. Appl.* 17 (Jun 2022), 064061. Issue 6. <https://doi.org/10.1103/PhysRevApplied.17.064061>
- [91] Vinay Tripathi, Huo Chen, Mostafa Khezri, Ka-Wa Yip, E.M. Levenson-Falk, and Daniel A. Lidar. 2022. Suppression of Crosstalk in Superconducting Qubits Using Dynamical Decoupling. *Phys. Rev. Appl.* 18 (Aug 2022), 024068. Issue 2. <https://doi.org/10.1103/PhysRevApplied.18.024068>
- [92] Cenik Tüysüz, Giuseppe Clemente, Arianna Crippa, Tobias Hartung, Stefan Kühn, and Karl Jansen. 2023. Classical splitting of parametrized quantum circuits. *Quantum Machine Intelligence* 5, 2 (2023), 34.
- [93] Robert Wille, Lukas Burgholzer, and Alwin Zulehner. 2019. Mapping Quantum Circuits to IBM QX Architectures Using the Minimal Number of SWAP and H Operations. In *Proceedings of the 56th Annual Design Automation Conference 2019* (Las Vegas, NV, USA) (*DAC '19*). Association for Computing Machinery, New York, NY, USA, Article 142, 6 pages. <https://doi.org/10.1145/3316781.3317859>
- [94] Nicolas Wittler, Federico Roy, Kevin Pack, Max Werninghaus, Anurag Saha Roy, Daniel J. Egger, Stefan Filipp, Frank K. Wilhelm, and Shai Machnes. 2021. Integrated Tool Set for Control, Calibration, and Characterization of Quantum Devices Applied to Superconducting Qubits. *Phys. Rev. Appl.* 15 (Mar 2021), 034080. Issue 3. <https://doi.org/10.1103/PhysRevApplied.15.034080>
- [95] Amanda Xu, Abtin Molavi, Lauren Pick, Swamit Tannu, and Aws Albarghouthi. 2023. Synthesizing Quantum-Circuit Optimizers. *Proceedings of the ACM on Programming Languages* 7, PLDI (jun 2023), 835–859. <https://doi.org/10.1145/3591254>
- [96] Chi Zhang, Ari B. Hayes, Longfei Qiu, Yuwei Jin, Yanhao Chen, and Eddy Z. Zhang. 2021. Time-Optimal Qubit Mapping. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Virtual, USA) (*ASPLOS '21*). Association for Computing Machinery, New York, NY, USA, 360–374. <https://doi.org/10.1145/3445814.3446706>
- [97] Alexander Zlokapa and Alexandru Gheorghiu. 2020. A deep learning model for noise prediction on near-term quantum devices. arXiv:2005.10811 [quant-ph]
- [98] Alwin Zulehner, Alexandru Paler, and Robert Wille. 2019. An Efficient Methodology for Mapping Quantum Circuits to the IBM QX Architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 38, 7 (July 2019), 1226–1236. <https://doi.org/10.1109/TCAD.2018.2846658>