

Resource-Efficient Quantum Computing by Breaking Abstractions

Yunong Shi¹, Pranav Gokhale¹, Prakash Murali², Jonathan M. Baker¹, Casey Duckering¹, Yongshan Ding¹,
Natalie C. Brown³, Christopher Chamberland^{5,6}, Ali Javadi Abhari⁴, Andrew W. Cross⁴, David I. Schuster¹,
Kenneth R. Brown¹, Margaret Martonosi², and Frederic T. Chong¹

¹University of Chicago

²Princeton University

³Duke University

⁴IBM T.J. Watson Research Center

⁵AWS Center for Quantum Computing

⁶California Institute of Technology

Abstract—Building a quantum computer that surpasses the computational power of its classical counterpart is a great engineering challenge. Quantum software optimizations can provide an accelerated pathway to the first generation of quantum computing applications that might save years of engineering effort. Current quantum software stacks follow a layered approach similar to the stack of classical computers, which was designed to manage the complexity. In this review, we point out that greater efficiency of quantum computing systems can be achieved by breaking the abstractions between these layers. We review several works along this line, including two hardware-aware compilation optimizations that break the quantum Instruction Set Architecture (ISA) abstraction and two error-correction/information-processing schemes that break the qubit abstraction. Last, we discuss several possible future directions.

I. INTRODUCTION

Quantum computing has recently transitioned from a theoretical prediction to a nascent technology. With development of Noisy Intermediate-Scale Quantum (NISQ) devices, cloud-based Quantum Information Processing (QIP) platforms with up to 53 qubits are currently accessible to the public. It also has been recently demonstrated by the *Quantum Supremacy* experiment on the Sycamore quantum processor, a 54-qubit quantum computing device manufactured by Google, that quantum computers can outperform current classical supercomputers in certain computational tasks [7]. These developments suggest that the future of quantum computing is promising. Nevertheless, there is still a gap between the ability and reliability of current QIP technologies and the requirements of the first useful quantum computing applications. The gap is mostly due to the presence of systematic errors including qubit decoherence, gate errors, State Preparation And Measurement (SPAM) errors. As an example, the best reported qubit decoherence time on a superconducting QIP platform is around $500\mu\text{s}$ (meaning that in $500\mu\text{s}$, the probability of a logical 1 state staying unflipped drops to $1/e \approx 0.368$), the error rate of 2-qubit gates is around 1%-5% in a device, measurement error of a single qubit is between 2%- 5% [1]. In

addition to the errors in the elementary operations, emergent error modes such as crosstalk are reported to make significant contributions to the current noise level in quantum devices [62], [37], [18]. With these sources of errors combined, we are only able to run quantum algorithms of very limited size on current quantum computing devices. Thus, it will require tremendous efforts and investment to solve these engineering challenges and we cannot expect a definite timeline for its success. Because of the uncertainties and difficulties in relying on hardware breakthroughs, it will also be crucial in the near term to close the gap using higher-level quantum optimizations and software hardware co-design, which could maximally utilize noisy devices and potentially provide an accelerated pathway to real-world quantum computing applications.

Currently, major quantum programming environments, including Qiskit [6] by IBM, Cirq [3] by Google, PyQuil [60] by Rigetti and strawberry fields [68] by Xanadu, follow the quantum circuit model. These programming environments support users in configuring, compiling and running their quantum programs in an automated workflow and roughly follow a layered approach as illustrated in Fig. 1. In these environments, the compilation stack is divided into layers of subroutines that are built upon the abstraction provided by the next layer. This design philosophy is similar to that of its classical counterpart, which has slowly converged to this layered approach over many years to manage the increasing complexity that comes with the exponentially growing hardware resources. In each layer, burdensome hardware details are well encapsulated and hidden behind a clean interface, which offers a well-defined, manageable optimization task to solve. Thus, this layered approach provides great portability and modularity. For example, the Qiskit compiler supports both the superconducting QIP platform and the trapped ion QIP platform as the backend (Fig. 2). In the Qiskit programming environment, these two backends share a unified, hardware-agnostic programming frontend even though the hardware characteristics and the qubit control methods of the two platforms are rather different. Superconducting (SC) qubits are macroscopic LC circuits




Fig. 1: The workflow of the quantum computing stack roughly followed by current programming environments (*e.g.*, Qiskit, Cirq, ScaffCC) based on the quantum circuit model.

placed inside dilution fridges of temperature near absolute zero. These qubits can be regarded as artificial atoms and are protected by a metal transmission line from environmental noise. For SC QIP platforms, qubit control is achieved through sending microwave pulses into the transmission line that surrounds the LC circuits to change the qubit state and those operations are usually done within several hundreds of nanoseconds. On the other hand, trapped ion qubits are ions confined in the potential of electrodes in vacuum chambers. Trapped ion qubits have a much longer coherence (> 1 second) and quantum operations are performed by shining modulated laser beam. The quantum gates are also much slower than that of SC qubits but the qubit connectivity (for 2-qubit gates) are much better. In Qiskit, the hardware characteristics of the two QIP platforms are abstracted away in the quantum circuit model so that the higher level programming environment can work with both backends.

However, the abstractions introduced in the layered approach of current QC stacks restrict opportunities for cross-layer optimizations. For example, without accessing the lower level noise information, the compiler might not be able to properly optimize gate scheduling and qubit mapping with regard to the final fidelity. For near-term quantum computing, maximal utilization of the scarce quantum resources and reconciling quantum algorithms with noisy devices is of more importance

than to manage complexity of the classical control system. In this review, we propose a shift of the quantum computing stack towards a more vertical integrated architecture. We point out that breaking the abstraction layers in the stack by exposing enough lower level details could substantially improve the quantum efficiency. This claim is not that surprising — there are many supporting examples from the classical computing world such as the emergence of application specific architectures like the GPU and the TPU. However, this view is often overlooked in the software/hardware design in quantum computing.

We examine this methodology by looking at several previous works along this line. We first review two compilation optimizations that break the ISA abstraction by exposing pulse level information (Section II) and noise information (Section III), respectively. Then, we discuss an information processing scheme that improves general circuit latency by exposing the third energy level of the underlying physical space, *i.e.*, breaking the qubit abstraction using qutrits (Section IV). Then, we discuss the Gottesman-Kitaev-Preskill (GKP) qubit encoding in a Quantum Harmonic Oscillator (QHO) that exposes error information in the form of small shifts in the phase space to assist the upper level error mitigation/correction procedure (Section V).

At last, we envision several future directions that could further explore the idea of breaking abstractions and assist the realization of the first quantum computers for real-world applications.

II. BREAKING THE ISA ABSTRACTION USING PULSE-LEVEL COMPILATION

In this section, we describe a quantum compilation methodology proposed in [67], [28] that achieves an average of 5X speedup in terms of generated circuit latency, by employing the idea of breaking the ISA abstraction and compiling directly to control pulses.

A. Quantum compilation

Since the early days of quantum computing, quantum compilation has been recognized as one of the central tasks in realizing practical quantum computation. Quantum compilation was first defined as the problem of synthesizing quantum circuits for a given unitary matrix. The celebrated Solovay-Kitaev theorem [34] states that such synthesis is always possible if a universal set of quantum gates is given. Now the term of quantum compilation is used more broadly and almost all stages in Fig. 1 can be viewed as part of the quantum compilation process.

There are many indications that current quantum compilation stack (Fig. 1) is highly-inefficient. First, current circuit synthesis algorithms are far from saturating (or being closed to) the asymptotic lower bound in the general case [51], [34]. Also, the formulated circuit synthesis problem is based on the fundamental abstraction of quantum ISA (Section II-B) and largely discussed in a hardware-agnostic settings in previous work but the underlying physical operations cannot be directly described by the logical level ISA (as shown in Fig. 2). The translation from the logical ISA to the operations directly




Fig. 2: The same abstractions in the quantum computing stack on the logical level can be mapped to different physical implementations. Here we take the Superconducting (SC) QIP platform and the trapped ion QIP platform as examples of the physical implementations. (Left) In the quantum circuit model, both SC qubits and trapped-ion qubits are abstracted as 2-level quantum systems and their physical operations are abstracted as quantum gates, even though these 2 systems have different physical properties. (Middle) SC qubits are SC circuits placed inside a long, metal transmission line. The apparatus requires a dilution fridge of temperature near absolute zero. The orange standing waves are oscillations in the transmission line, which are driven by external microwave pulses and used to control the qubit states. (Right) Trapped ion qubits are confined in the potential of cylindrical electrodes. Modulated laser beam can provide elementary quantum operations for trapped ion qubits. The apparatus is usually contained inside a vacuum chamber of pressure around 10^{-8}Pa . The two systems require different high-level optimizations for better efficiency due to their distinct physical features.

supported by the hardware is typically done in an *ad-hoc* way. Thus, there is a mismatch between the expressive logical gates and the set of instructions that can be efficiently implemented on a real system. This mismatch significantly limits the efficiency of the current quantum computing stack thus the underlying quantum devices' computing ability and wastes precious quantum coherence. While improving the computing efficiency is always valuable, improving quantum computing efficiency is do-or-die: computation has to finish before qubit decoherence or the results will be worthless. Thus, improving the compilation process is one of the most, if not the most, crucial goals in near-term QC system design.

By identifying this mismatch and the fundamental limitation in the ISA abstraction, in [67], [28], we proposed a quantum compilation technique that optimizes across existing abstraction barriers to greatly reduce latency while still being practical for large numbers of qubits. Specifically, rather than limiting the compiler to use 1- and 2-qubit quantum instructions, our framework aggregates the instructions in the logical ISA into a customized set of instructions that correspond to optimized control pulses. We compare our methodology to the standard compilation workflow on several promising NISQ quantum

applications and conclude that our compilation methodology has an average speedup of $5\times$ with a maximum speedup of $10\times$. We use the rest of this section to introduce this compilation methodology, starting with defining some basic concepts.

B. Quantum ISA

In the quantum computing stack, a restricted set of 1- and 2-qubit quantum instructions are provided for describing the high-level quantum algorithms, analogous to the Instruction Set Architecture (ISA) abstraction in classical computing. In this paper, we call this instruction set the *logical ISA*. The 1-qubit gates in the logical ISA include the Pauli gates, $P = \{X, Y, Z\}$. It also includes the Hadamard H gate, whose symbol in the circuit model is given as an example in Fig. 2 on the left column. The typical 2-qubit instruction in the logical instruction set is the Controlled-NOT (CNOT) gate, which flips the state of the target qubit based on the state of the control qubit.

However, usually quantum computing devices does not directly support the logical ISA. Based on the system characteristics, we can define the *physical ISA* that can be directly mapped to the underlying control signals. For example, superconducting devices typically has Cross-Resonance (CR) gate

or *i*SWAP gate as their intrinsic 2-qubit instruction, whereas for trapped-ion devices the intrinsic 2-qubit instruction can be the Mølmer-Sørensen gate or the controlled phase gate.

C. Quantum Control

As shown in Fig. 2 and discussed in the last subsection, underlying physical operations in the hardware such as microwave control pulses and modulated laser beam are abstracted as quantum instructions. A quantum instruction are simply as pre-fined control pulse sequences.

The underlying evolution of the quantum system is continuous and so are the control signals. The continuous control signals offer much richer and flexible controllability than the quantum ISA. The control pulses can drive the quantum computing hardware to a desired quantum states by varying a system-dependent and time-dependent quantity called the Hamiltonian. The Hamiltonian of a system determines the evolution path of the quantum states. The ability to engineer real-time system Hamiltonian allows us to navigate the quantum system to the quantum state of interest through generating accurate control signals. Thus, quantum computation can be done by constructing a quantum system in which the system Hamiltonian evolves in a way that aligns with a quantum computing task, producing the computational result with high probability upon final measurement of the qubits. In general, the path to a final quantum state is not unique and finding the optimal evolution path is a very important but challenging problem [40], [64], [25].

D. The Mismatch Between ISA and Control

Being hardware-agnostic, the quantum operation sequences composed by logical ISA have limited freedom in terms of controllability and usually will not be mapped to the optimal evolution path of the underlying quantum system, thus there is a mismatch between the ISA and low-level quantum control. With two simple examples, we demonstrate this mismatch.

- We can consider the instruction sequence consists of a CNOT gate followed by a *X* gate on the control bit. In current compilation workflow, these two logical gates will be further decomposed into the physical ISA and be executed sequentially. However, on SC QIP platforms, the microwave pulses that implement these two instructions could in fact be applied simultaneously (because of their commutativity). Even the commutativity can be captured by the ISA abstraction, in the current compilation workflow, the compiled control signals are sub-optimal.
- SWAP gate is an important quantum instruction for circuit mapping. The SWAP operation is usually decomposed as 3 Controlled-NOT (CNOT) operations, as realized in the circuit below. This decomposition could be thought of the implementation of in-place memory SWAPs with three alternating XORs for classical computation. However, for systems like quantum dots [42], the SWAP operation is directly supported by applying particular constant control signals for a certain period of time. In this case, this decomposition of SWAP into three CNOTs introduces substantial overhead.



In experimental physics settings, equivalences from simple gate sequences to control pulses can be hand optimized [63]. However, when circuits become larger and more complicated, this kind of hand optimization become less efficient and the standard decomposition becomes less favorable, motivating a shift toward numerical optimization methods that are not limited by the ISA abstraction.

E. Quantum Optimal Control

Quantum Optimal Control (QOC) theory provides an alternative in terms of finding the optimal evolution path for the quantum compilation tasks. Quantum optimal control algorithms typically perform analytical or numerical methods for this optimization, among which, gradient ascent methods, such as the GRAdient Ascent Pulse Engineering (GRAPE) [33], [15] algorithm, are widely used. The basic idea of GRAPE is as follows: for optimizing the control signals of M parameters ($u_1, \dots, u_M$) for a target quantum state, in every iteration, GRAPE minimizes the deviation of the system evolution by calculating the gradient of the final fidelity with respect to the M control parameters in the M dimensional space. Then GRAPE will update the parameters in the direction of the gradient with adaptive step size [33], [15], [40]. With a large number of iterations, the optimized control signals are expected to converge and find optimized pulses.

In [67], we utilize GRAPE to optimize our aggregated instructions that are customized for each quantum circuit as opposed to selecting instructions from a pre-defined pulse sequences. However, one disadvantage of numerical methods like GRAPE is that the running time and memory use grow exponentially with the size of the quantum system for optimization. In our work, we are able to use GRAPE for optimizing quantum systems of up to 10 qubits with the GPU accelerated optimal control unit [40]. As shown in our result, the limit of 10 qubits does not put restrictions on the result of our compilation methodology.

F. Pulse-Level Optimization: A Motivating Example

Next, we will illustrate the workflow of our compilation methodology with a circuit instance of the Quantum Approximate Optimization Algorithm (QAOA) for solving the MAXCUT problem on the triangle graph (Fig. 3).¹ This QAOA circuit with logical ISA (or variants of it up to single qubit gates) can be reproduced by most existing quantum compilers. This instance of the QAOA circuit is generated by the ScaffCC compiler, as shown in Fig. 3 (a). We assume this circuit is executed on a superconducting architecture with 1D nearest neighbor qubit connectivity. A SWAP instruction is inserted in the circuit to satisfy the linear qubit connectivity constraints.

On the other hand, our compiler generates the aggregated instruction set $G_1 - G_5$ as illustrated in Fig. 3 (b) automatically,

¹The angle parameters γ and β can be determined by variational methods [46] and are set to 5.67 and 1.26.




Fig. 3: Example of a QAOA circuit. (a) The QAOA circuit with the logical ISA. (b) The QAOA circuit with aggregated instructions. (c) Generated control pulses for G_3 in the ISA-based compilation. (d) Control pulses for G_3 from aggregated instructions based compilation. Each curve is the amplitude of a relevant control signal. The pulse sequences in (d) provides a $3\times$ speedup comparing to the pulse sequences in (c). Pulse sequences reprinted with permission from [67]

and uses GRAPE to produce highly-optimized pulse sequences for each aggregated instruction. In this minimal circuit instance, our compilation method reduces the total execution time of the circuit by about $2.97\times$ comparing to compilation with restricted ISA. Fig. 3 (c) and (d) show the generated pulses for G_3 with ISA-based compilation and with our aggregated instruction based, pulse-level optimized compilation.

G. Optimized Pulse-Level Compilation Using Gate Aggregation: the workflow

Now we give a systematic view of the workflow of our compiler. First, at the program level, our compiler performs module flattening and loop unrolling to produce the quantum assembly (QASM), which represents a schedule of the logical operations. Next, the compiler enters the commutativity detection phase. Different from the ISA-based approach, in this phase, our compilation process converts the QASM code to a more flexible logical schedule that explores the commutativity between instructions. To further explore the commutativity in the schedule, the compiler aggregates instructions in the schedule to produce a new logical schedule with instructions that represents diagonal matrices (and are of high commutativity). Then the compiler enters the scheduling and mapping phase. Because of commutativity awareness, our compiler can generate a much more efficient logical schedule by re-arranging the aggregated instructions with high commutativity. The logical schedule is then converted to a physical schedule after the qubit mapping stage. Then the compiler generates the final aggregated instructions for pulse optimization and use GRAPE for producing the corresponding control pulses. The goal of the final aggregation is to find the optimal instruction set that produces the lowest-latency control pulses while preserving the parallelism in the circuit aggregations that are small as much as possible. Finally, our compiler outputs an optimized physical schedule along with the corresponding optimized control pulses. Fig. 4 shows the Gate Dependency Graph (GDG) of the QAOA circuit in Figure 5 in different compilation stages. Next, we walk through the compilation backend with this example, starting from the commutativity detection phase.

1) *Commutativity detection*:: In the commutativity detection phase, the false dependence between commutative instructions are removed and the GDG is re-structured. This is because if a pair of gates commutes, the gates can be scheduled in either order. Also, it can be further noticed that, in many NISQ quantum algorithms, it is ubiquitous that for instructions *within* an instruction block to not commute, but for the full instruction block to commute with each other [19], [38]. As an example, in Figure 5, the CNOT-Rz-CNOT instruction block commute with each other because these blocks correspond to diagonal unitary matrices. However, each individual instruction in these circuit blocks does not commute. Thus, after aggregating these instructions, the compiler is able to schedule new aggregated instructions in any order, which is impossible before. This commutativity detection procedure opens up opportunities for more efficient scheduling.

2) *Scheduling and mapping*:: **Commutativity-aware logical scheduling (CLS)** In our scheduling phase, our logical scheduling algorithm is able to fully utilize the detected commutativity in the last compilation phase. The CLS iteratively schedules the available instructions on each qubits. At each iteration, the CLS draws instruction candidates that can be executed in the earliest time step to schedule.

Qubit mapping In this phase of the compilation, the compiler transform the circuit to a form that respect the topological constraints of hardware connectivity [45]. To conform to the




Fig. 4: The example in Fig. 5 in the form of gate dependency graph.




Fig. 5: An example of commutativity-aware logical scheduling. With commutativity detected, the circuit depth can be shortened.

device topology, the logical instructions are processed in two steps. First, we place frequently interacting qubits near each other by bisecting the qubit interaction graph along a cut with few crossing edges, computed by the METIS graph partitioning library [32]. Once the initial mapping is generated, two-qubit operations between non-neighboring qubits are prepended with a sequence of SWAP rearrangements that move the control and target qubits to be adjacent.

3) *Instruction aggregation*:: In this phase, the compiler iterates with the optimal control unit to generate the final aggregated instructions for the circuit. Then, the optimal control unit optimizes each instruction individually with GRAPE.

4) *Physical Execution*:: Finally, the circuit will be scheduled again using the CLS from Section II-G2, the physical schedules will be sent to the control unit of the underlying quantum hardware and trigger the optimized control pulses at appropriate timing and the physical execution.

H. Discussion

In [67], we selected 9 important quantum /classical-quantum hybrid algorithms in the NISQ era as our benchmarks. Across all 9 benchmarks, our compilation scheme achieves a geometric mean of $5.07\times$ pulse time reduction comparing to the standard gate-based compilation. The result in [67] indicates that addressing the mismatch between quantum gates and the control pulses by breaking the ISA abstraction can greatly improve the compilation efficiency. Going beyond the ISA-based compilation, this work opens up a door to new QC system designs.

III. BREAKING THE ISA ABSTRACTION USING NOISE-ADAPTIVE COMPILATION

In recent years, QC compute stacks have been developed using abstractions inspired from classical computing. The instruction set architecture (ISA) is a fundamental abstraction which defines the interface between the hardware and software. The ISA abstraction allows software to execute correctly on any hardware which implements the interface. This enables application portability and decouples hardware and software development.

For QC systems, the hardware-software interface is typically defined as a set of legal instructions and the connectivity topology of the qubits [14], [60], [22], [21], [20], — it does not include information about qubit quality, gate fidelity or micro-operations used to implement the ISA instructions. While technology independent abstractions are desirable in the long




Fig. 6: Daily variations in qubit coherence time (larger is better) and gate error rates (lower is better) for selected qubits and gates in IBM’s 16-qubit system. The most or least reliable system elements change across days.

run, our work [48], [49] reveals that such abstractions are detrimental to program correctness in NISQ quantum computers. By exposing microarchitectural details to software and using intelligent compilation techniques, we show that program reliability can be improved significantly.

A. Noise Characteristics of QC Systems

QC systems have spatial and temporal variations in noise due to manufacturing imperfections, imprecise qubit control and external interference. To motivate the necessity for breaking the ISA abstraction barrier, we present real-system statistics of hardware noise in systems from three leading QC vendors – IBM, Rigetti and University of Maryland. IBM and Rigetti systems use superconducting qubits [12], [10] and University of Maryland (UMD) uses trapped ion qubits [16]. The gates in these systems are periodically calibrated and their error rates are measured.

Figure 6 shows the coherence times and two-qubit gate error rates in IBM’s 16-qubit system (*IBMQ16 Rueschlikon*). From daily calibration logs we find that, the average qubit coherence time is 40 microseconds, two-qubit gate error rate is 7%, readout error rate is 4% and single qubit error rate is 0.2%. The two-qubit and readout errors are the dominant noise sources and vary up to 9X across gates and calibration cycles. Rigetti’s systems also exhibit error rates and variations of comparable magnitude. These noise variations in superconducting systems emerge from material defects due to lithographic manufacturing, and are expected in future systems also [35], [36].

Trapped ion systems also have noise fluctuations even though the individual qubits are identical and defect-free. On a 5-qubit

trapped ion system from UMD, we observed up to 3x variation in the two-qubit gate error rates because of fundamental challenges in qubit control using lasers and their sensitivity to motional mode drifts from temperature fluctuations.

We found that these microarchitectural noise variations dramatically influence program correctness. When a program is executed on a noisy QC system, the results may be corrupted by gate errors, decoherence or readout errors on the hardware qubits used for execution. Therefore, it is crucial to select the most reliable hardware qubits to improve the *success rate* of the program (the likelihood of correct execution). The success rate is determined by executing a program multiple times and measuring the fraction of runs that produce the correct output. High success rate is important to ensure that the program execution is not dominated by noise.

B. Noise-Adaptive Compilation: Key Ideas

Our work breaks the ISA abstraction barrier by developing compiler optimizations which use hardware calibration data. These optimizations boost the success rate a program run by avoiding portions of the machine with poor coherence time and operational error rates.

We first review the key components in a QC compiler. The input to the compiler is a high-level language program (Scaffold in our framework) and the output is machine executable assembly code. First, the compiler converts the program to an intermediate representation (IR) composed of single and two-qubit gates by decomposing high-level QC operations, unrolling all loops and inlining all functions. Figure 7a shows an example IR. The qubits in the IR (program qubits) are **mapped** to distinct qubits in the hardware, typically in a way that reduces qubit communication. Next, gates are **scheduled** while respecting data dependencies. Finally, on hardware with limited connectivity, such as superconducting systems, the compiler **inserts SWAP operations** to enable 2-qubit operations between non-adjacent qubits.

Figure 7b and 7c show two compiler mappings for a 4-qubit program on IBM’s 16-qubit system. In the first mapping, the compiler must insert SWAPs to perform the two-qubit gate between p_1 and p_3 . Since SWAP operations are composed of three two-qubit gates, they are highly error prone. In contrast, the second mapping requires no SWAPs because the qubits required for the CNOTs are adjacent. While SWAP optimizations can be performed using the device ISA, the second mapping is also noise-optimized i.e., it uses qubits with high coherence time and low operational error rates. By considering microarchitectural noise characteristics, our compiler can determine such noise-optimized mappings to improve the program success rate.

We developed three strategies for noise optimization. First, our compiler maps program qubits onto hardware locations with high reliability, based on the noise data. We choose the initial mapping based on two-qubit and readout error rates because they are the dominant sources of error. Second, to mitigate decoherence errors, all gates are scheduled to finish before the coherence time of the hardware qubits. Third, our compiler optimizes the reliability of SWAP operations by minimizing the




Fig. 7: Figure (a) shows the intermediate representation of the Bernstein-Vazirani algorithm (BV4). Each horizontal line represents a program qubit. X and H are single qubit gates. The Controlled NOT (CNOT) gates from each qubit $p_{0,1,2}$ to p_3 are marked by vertical lines with XOR connectors. The readout operation is indicated by the meter. Figure (b) shows the qubit layout in IBMQ16, a naive mapping of BV4 onto this system. The black circles denote qubits and the edges indicate hardware CNOT gates. The edges are labelled with CNOT gate error ($\times 10^{-2}$). The hatched qubits and crossed gates are unreliable. In this mapping, a SWAP operation is required to perform the CNOT between p_1 and p_3 and error-prone operations are used. Figure (c) shows a mapping where qubit movement is not required and unreliable qubits and gates are avoided.




Fig. 8: Noise-adaptive compilation using SMT Optimization. Inputs are a QC program IR, details about the hardware qubit configuration, and a set of options, such as routing policy and solver options. From these, compiler generates a set of appropriate constraints and uses them to map program qubits to hardware qubits and schedule operations. The output of the optimization is used to generate an executable version of the program.

number of SWAPs whenever possible and performing SWAPs along reliable hardware paths.

C. Implementation using SMT Optimization

Our compiler implements the above strategies by finding the solution to a constrained optimization problem using a Satisfiability Modulo Theory (SMT) solver. The variables and constraints in the optimization encode program information, device topology constraints and noise information. The variables express the choices for program qubit mappings,

gate start times and routing paths. The constraints specify that qubit mappings should be distinct, the schedule should respect program dependencies and that routing paths should be non-overlapping. Fig. 8 summarizes the optimization-based compilation pipeline for IBMQ16.

The objective of our optimization is to maximize the success rate of a program execution. Since the success rate can only be determined from a real-system run, we model it at compile time as the *program reliability*. We define the reliability of a program as the product of the reliability of all gates in the program. Although this is not a perfect model for the success rate, it serves as a useful measure of overall correctness [7], [41]. For a given mapping, the solver determines the reliability of each two-qubit and readout operation and computes an overall reliability score. The solver maximizes the reliability score over all mappings by tracking and adapting to the error rates, coherence limits, and qubit movement based on program qubit locations.

In practice, we use the Z3 SMT solver to express and solve this optimization. Since the reliability objective is a non-linear product, we linearize the objective by optimizing for the additive logarithms of the reliability scores of each gate. We term this algorithm as R-SMT[★]. The output of the SMT solver is used to create machine executable code in the vendor-specified assembly language.

D. Real-System Evaluation

We present real-system evaluation on IBMQ16. Our evaluation uses 12 common QC benchmarks, compiled using R-SMT[★] and T-SMT[★] which are variants of our compiler and IBM's Qiskit compiler (version 0.5.7) [6] which is the default for this system. R-SMT[★] optimizes the reliability of the program using hardware noise data. T-SMT[★] optimizes the execution time of the program considering real-system gate durations and coherence times, but not operational error rates. IBM Qiskit is also noise-unaware and uses randomized algorithms for SWAP optimization. For each benchmark and compiler, we measured the success rate on IBMQ16 system using 8192 trials per program. Success rate of 1 indicates a perfect noise-free execution.




Fig. 9: For real data/experiment, on IBMQ16, qubit mappings for Qiskit and our compiler with three optimization objectives, varying the type of noise-awareness. The edge labels indicate the CNOT gate error rate ($\times 10^{-2}$), and the node labels indicate the qubit’s readout error rate ($\times 10^{-2}$). The thin red arrows indicate CNOT gates. The thick yellow arrows indicate SWAP operations. ω is a weight factor for readout error terms in the R-SMT[★] objective. (a) Qiskit finds a mapping which requires SWAP operations and uses hardware qubits with high readout errors (b), T-SMT[★] finds a mapping which requires no SWAP operations, but it uses an unreliable hardware CNOT between p_3 and p_0 . (c) Program qubits are placed on the best readout qubits, but p_0 and p_3 communicate using swaps. (d) R-SMT[★] finds a mapping which has the best reliability where the best CNOTs and readout qubits are used. It also requires no SWAP operations.




Fig. 10: Measured success rate of R-SMT[★] compared to Qiskit and T-SMT[★]. (Of 8192 trials per execution, success rate is the percentage that achieve the correct answer in real-system execution.) ω is a weight factor for readout error terms in the R-SMT[★] objective, 0.5 is equal weight for CNOT and readout errors. R-SMT[★] obtains higher success rate than Qiskit because it adapts the qubit mappings according to dynamic error rates and also avoids unnecessary qubit communication.

Figure 10 shows the success rate for the three compilers on all the benchmarks. R-SMT[★] has higher success rate than both baselines on all benchmarks, demonstrating the effectiveness of noise-adaptive compilation. Across benchmarks R-SMT[★] obtains geomean 2.9x improvement over Qiskit, with up to 18x gain. Figure 9 shows the mapping used by Qiskit, T-SMT[★] and R-SMT[★] for BV4. Qiskit places qubits in a lexicographic order without considering CNOT and readout errors and incurs extra swap operations. Similarly, T-SMT[★] is also unaware of noise variations across the device, resulting in mappings which use unreliable hardware. R-SMT[★] outperforms these baselines because it maximizes the likelihood of reliable execution

by leveraging microarchitectural noise characteristics during compilation.

Full results of our evaluation on 7 QC systems from IBM, Rigetti and UMD can be found in [48], [49].

E. Discussion

Our work represents one of the first efforts to exploit hardware noise characteristics during compilation. We developed optimal and heuristic techniques for noise adaptivity and performed comprehensive evaluations on several real QC systems [49]. We also developed techniques to mitigate crosstalk, another major source of errors in QC systems, using compiler techniques that schedule programs using crosstalk characterization data from the hardware [50]. In addition, our techniques are already being used in industry toolflows [56], [61]. Recognizing the importance of efficient compilation, other research groups have also recently developed mapping and routing heuristics [11], [74] and techniques to handle noise [70], [69].

Our noise-adaptivity optimizations offer large gains in success rate. These gains mean the difference between executions which yield correct and usable results and executions where the results are dominated by noise. These improvements are also multiplicative against benefits obtained elsewhere in the stack and will be instrumental in closing the gap between near-term QC algorithms and hardware. Our work also indicates that it is important to accurately characterize hardware and expose characterization data to software instead of hiding it behind a device-independent ISA layer. Additionally our work also proposes that QC programs should be compiled once-per-execution using the latest hardware characterization data to obtain the best executions.

Going beyond noise characteristics, we also studied the importance of exposing other microarchitectural information to software. We found that when the compiler has access to the native gates available in the hardware (micro operations used to implement ISA-level gates), it can further optimize programs and improve success rates. Overall, our work indicates that QC machines are not yet ready for technology independent abstractions that shield the software from hardware. Bridging the information gap between software and hardware by breaking abstraction barriers will be increasingly important on the path towards practically useful NISQ devices.

IV. BREAKING THE QUBIT ABSTRACTION VIA THE THIRD ENERGY LEVEL

While quantum computation is typically expressed with the two-level binary abstraction of qubits, the underlying physics of quantum systems are **not intrinsically binary**. Whereas classical computers operate in binary states at the physical level (e.g. clipping above and below a threshold voltage), quantum computers have natural access to an infinite spectrum of discrete energy levels. In fact, hardware must actively suppress higher level states in order to realize an *engineered* two-level qubit approximation. In this sense, using three-level *qutrits* (quantum trits) is simply a choice of including an additional discrete energy level within the computational space. Thus, it is appealing to explore what gains can be realized by breaking the binary qubit abstraction.

In prior work on qutrits (or more generally, d-level *qudits*), researchers identified only constant factor gains from extending beyond qubits. In general, this prior work [55] has emphasized the information compression advantages of qutrits. For example, N qubits can be expressed as $\frac{N}{\log_2(3)}$ qutrits, which leads to $\log_2(3) \approx 1.6$ -constant factor improvements in runtimes.

Recently however, our research group demonstrated a novel qutrit approach that leads to exponentially faster runtimes (i.e. shorter in circuit *depth*) than qubit-only approaches [27], [26]. The key idea underlying the approach is to use the third state of a qutrit as temporary storage. Although qutrits incur higher per-operation error rates than qubits, this is compensated by dramatic reductions in runtimes and quantum gate counts. Moreover, our approach only applies qutrit operations in an intermediary stage: the input and output are still qubits, which is important for initialization and measurement on practical quantum devices [59], [58].

The net result of our work is to extend the frontier of what quantum computers can compute. In particular, the frontier is defined by the zone in which every machine qubit is a data qubit, for example a 100-qubit algorithm running on a 100-qubit machine. In this frontier zone, we do not have space for non-data workspace qubits known as ancilla. The lack of ancilla in the frontier zone is a costly constraint that generally leads to inefficient circuits. For this reason, typical circuits instead operate below the frontier zone, with many machine qubits used as ancilla. Our work demonstrates that ancilla can be substituted with qutrits, enabling us to operate efficiently within the ancilla-free frontier zone.




Fig. 11: Reversible AND circuit using a single ancilla bit. The inputs are on the left, and time flows rightward to the outputs. This AND gate is implemented using a Toffoli (CCNOT) gate with inputs q_0 , q_1 and a single ancilla initialized to 0. At the end of the circuit, q_0 and q_1 are preserved, and the ancilla bit is set to 1 if and only if both other inputs are 1.

A. Qutrit-Assisted AND Gate

We develop the intuition for how qutrits can be useful by considering the example of constructing an AND gate. In the framework of quantum computing, which requires *reversibility*, AND is not permitted directly. For example, consider the output of 0 from an AND gate with two inputs. With only this information about the output, the value of the inputs cannot be uniquely determined (00, 01, and 10 all yield an AND output of 0). However, these operations can be made reversible by the addition of an extra, temporary workspace bit initialized to 0. Using a single additional such *ancilla*, the AND operation can be computed reversibly as in Figure 11. While this approach works, it is expensive—in order to decompose the *Toffoli* gate in Figure 11 into hardware-implementable one- and two- input gates, it is decomposed into at least six CNOT (controlled-NOT) gates.

However, if we break the qubit abstraction and allow occupation of a higher *qutrit* energy level, the cost of the Toffoli AND operation is greatly diminished. Before proceeding, we review the basics of qutrits, which have three computational basis states: $|0\rangle$, $|1\rangle$, and $|2\rangle$. A qutrit state $|\psi\rangle$ may be represented analogously to a qubit as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle$, where $\|\alpha\|^2 + \|\beta\|^2 + \|\gamma\|^2 = 1$. Qutrits are manipulated in a similar manner to qubits; however, there are additional gates which may be performed on qutrits. We focus on the X_{+1} and X_{-1} operations, which are addition and subtraction gates, modulo 3. For example X_{+1} elevates $|0\rangle$ to $|1\rangle$ and elevates $|1\rangle$ to $|2\rangle$, while wrapping $|2\rangle$ to $|0\rangle$.

Just as single-qubit gates have qutrit analogs, the same holds for two-qutrit gates. For example, consider the CNOT operation, where an X gate is performed conditioned on the control being in the $|1\rangle$ state. For qutrits, an X_{+1} or X_{-1} gate may be performed, conditioned on the control being in any of the three possible basis states. Just as qubit gates are extended to take multiple controls, qutrit gates are extended similarly.

In Figure 12, a Toffoli decomposition using qutrits is given. A similar construction for the Toffoli gate is known from past work [39], [57]. The goal is to perform an X operation on the last (target) input qubit q_2 if and only if the two control qubits, q_0 and q_1 , are both $|1\rangle$. First a $|1\rangle$ -controlled X_{+1} is performed on q_0 and q_1 . This elevates q_1 to $|2\rangle$ iff q_0 and q_1 were both $|1\rangle$. Then a $|2\rangle$ -controlled X gate is applied to q_2 . Therefore, X is performed only when both q_0 and q_1 were $|1\rangle$, as desired. The controls are restored to their original states by a $|1\rangle$ -controlled X_{-1} gate, which undoes the effect of the




Fig. 12: A Toffoli decomposition via qutrits. Each input and output is a qubit. The red controls activate on $|1\rangle$ and the blue controls activate on $|2\rangle$. The first gate temporarily elevates q_1 to $|2\rangle$ if both q_0 and q_1 were $|1\rangle$. We then perform the X operation only if q_1 is $|2\rangle$. The final gate restores q_0 and q_1 to their original state.

first gate. The key intuition in this decomposition is that the qutrit $|2\rangle$ state can be used instead of ancilla to store temporary information.

B. Generalized Toffoli Gate




Fig. 13: Our circuit decomposition for the Generalized Toffoli gate is shown for 15 controls and 1 target. The inputs and outputs are both qubits, but we allow occupation of the $|2\rangle$ qutrit state in between. The circuit has a tree structure and maintains the property that the root of each subtree can only be elevated to $|2\rangle$ if all of its control leaves were $|1\rangle$. Thus, the U gate is only executed if all controls are $|1\rangle$. The right half of the circuit performs uncomputation to restore the controls to their original state. This construction applies more generally to any multiply-controlled U gate. Note that the three-input gates are decomposed into 6 two-input and 7 single-input gates in our actual simulation, as based on the decomposition in [17].

The intuition of our technique extends to more complicated gates. In particular, we consider the Generalized Toffoli gate, a ubiquitous quantum operation which extends the Toffoli gate to have any number of control inputs. The target input is flipped if and only if all of the controls are activated. Our qutrit-based circuit decomposition for the Generalized Toffoli gate is presented in Figure 13. The decomposition is expressed

in terms of three-qutrit gates (two controls, one target) instead of single- and two-qutrit gates, because the circuit can be understood purely classically at this granularity. In actual implementation and in our simulation, we used a decomposition [17] that requires 6 two-qutrit and 7 single-qutrit physically implementable quantum gates.

Our circuit decomposition is most intuitively understood by treating the left half of the the circuit as a tree. The desired property is that the root of the tree, q_7 , is $|2\rangle$ if and only if each of the 15 controls was originally in the $|1\rangle$ state. To verify this property, we observe the root q_7 can only become $|2\rangle$ iff q_7 was originally $|1\rangle$ and q_3 and q_{11} were both previously $|2\rangle$. At the next level of the tree, we see q_3 could have only been $|2\rangle$ if q_3 was originally $|1\rangle$ and both q_1 and q_5 were previously $|2\rangle$, and similarly for the other triplets. At the bottom level of the tree, the triplets are controlled on the $|1\rangle$ state, which are only activated when the even-index controls are all $|1\rangle$. Thus, if any of the controls were not $|1\rangle$, the $|2\rangle$ states would fail to propagate to the root of the tree. The right half of the circuit performs *uncomputation* to restore the controls to their original state.

After each subsequent level of the tree structure, the number of qubits under consideration is reduced by a factor of ~ 2 . Thus, the circuit depth is logarithmic in N , which is **exponentially smaller than ancilla-free qubit-only circuits**. Moreover, each qutrit is operated on by a constant number of gates, so the total number of gates is linear in N .

We verified our circuits, both formally and via simulation. Our verification scripts are available on our GitHub [4].

C. Simulation Results

	QUBIT	QUBIT+ANCILLA	QUTRIT
Depth	$\sim 633N$	$\sim 76N$	$\sim 38 \log_2(N)$
Gate Count	$\sim 397N$	$\sim 48N$	$\sim 6N$

TABLE I: Scaling of circuit depths and two-qutrit gate counts for all three benchmarked circuit constructions for the N -controlled Generalized Toffoli.

Table I shows the scaling of circuit depths and two-qutrit gate counts for all three benchmarked circuits. The QUBIT-based circuit constructions from past work are linear in depth and have a high linearity constant. Augmenting with a single borrowed ancilla (QUBIT+ANCILLA) reduces the circuit depth by a factor of 8. However, both circuit constructions are significantly outperformed by our QUTRIT construction, which scales logarithmically in N and has a relatively small leading coefficient. While there is not an asymptotic scaling advantage for two-qutrit gate count, the linearity constant for our QUTRIT circuit is 70x smaller than for the equivalent ancilla-free QUBIT circuit.

We ran simulations under realistic superconducting and trapped ion device noise. The simulations were run in parallel on over 100 n1-standard-4 Google Cloud instances. These simulations represent over 20,000 CPU hours, which was sufficient to estimate mean fidelity to an error of $2\sigma < 0.1\%$ for each circuit-noise model pair.




Fig. 14: Circuit simulation results for all possible pairs of circuit constructions and noise models. Each bar represents 1000+ trials, so the error bars are all $2\sigma < 0.1\%$. Our QUTRIT construction significantly outperforms the QUBIT construction. The QUBIT+ANCILLA bars are drawn with dashed lines to emphasize that it has access to an extra ancilla bit, unlike our construction. Figure reprinted with permission from [27]

The full results of our circuit simulations are shown in Figure 14. All simulations are for the 14-input (13 controls, 1 target) Generalized Toffoli gate. We simulated each of the three circuit benchmarks against each of our noise models (when applicable), yielding the 16 bars in the figure. Notice that our qutrit circuit consistently outperforms qubit circuits, with advantages ranging from 2x to 10,000x.

D. Discussion

The results presented in our work in [27], [26] are applicable to quantum computing in the near term, on machines that are expected within the next five years. By breaking the qubit abstraction barrier, we extend the frontier of what is computable by quantum hardware right now, without needing to wait for better hardware. As verified by our open-source circuit simulator coupled with realistic noise models, our circuits are more reliable than qubit-only equivalents, suggesting that qutrits offer a promising path towards scaling quantum computers. We propose further investigation into what advantage qutrits or qudits may confer. More broadly, it is critical for quantum architects to bear in mind that standard abstractions in classical computing do not necessarily transfer to quantum computation. Often, this presents unrealized opportunities, as in the case of qutrits.

V. BREAKING THE QUBIT ABSTRACTION VIA THE GKP ENCODING

Currently, there are many competing physical qubit implementations. For example, the transmon qubits [2] are encoded in the lowest two energy levels of the charge states in superconducting LC circuits with Josephson junctions; Trapped ion qubits can be encoded in two ground state hyperfine levels [9] or a ground state level and an excited level of an ion [13]; quantum dot qubits use electron spin triplets [43]. These QIP platforms have rather distinct physical characteristics, but they are all exposed to the other layers in the stack as qubits and other implementation details are often hidden. This abstraction is natural for classical computing stack because

that the robustness of classical bits decouples the programming logic from physical properties of the transistors except the logical value. In contrast, qubits are fragile so there are more than (superpositions of) the logical values that we want to know about the implementation. For example, in the transmon qubits and trapped ion qubits, logical states can be transferred to higher levels of the physical space by unwanted operations and this can cause leakage errors [24], [73]. It will be useful for other layers in the stack to access this error information and develop methods to mitigate it. In the previous section we discussed the qutrit approach that directly uses the third level for information processing, however, it could be more interesting if we can encode the qubit (qudit) using the whole physical Hilbert space to avoid leakage errors systematically and use the redundant degrees of freedom to reveal information about the noise in the encoding. The encoding proposed by Gottesman, Kitaev and Preskill (GKP) [29] provides such an example. GKP encoding is free of leakage errors and other errors (in the form of small shifts in phase space) can be identified and corrected by Quantum Non-Demolition (QND) measurements and simple linear optical operations. In realistic implementations of approximate GKP states (Section V-C), there are leakage errors between logical states, but the transfer probability is estimated to be at the order of 10^{-10} with current technology, thus negligible.

A. The phase space diagram

We describe the GKP qubits in the phase space. For a comparison, we first discuss the phase space diagram for a classical harmonic oscillator and a superconducting qubit.

Classical Harmonic Oscillators Examples of Classical Harmonic Oscillators (CHO) include LC circuits, springs and pendulums with small displacement. The voltage/displacement (denoted as p) and the current/momentum (denoted as q) value completely characterize the dynamics of CHO systems. The phase space diagram plots p vs q , which for CHOs are circles (up to normalization) with the radius representing the system energy. The energy of CHOs can be any non-negative real value.




Fig. 15: Phase space diagrams for a Classical Harmonic Oscillator (CHO), the ground state and the first excited state of a Quantum Harmonic Oscillator (QHO) and the logic 0 and 1 state of the GKP qubit. For quantum phase space diagrams, the plotted distribution is the Wigner quasi-probability function, where red indicates positive values and blue indicates negative values.

Quantum Harmonic Oscillators The Quantum Harmonic Oscillator is the quantized version of the CHO and is the physical model for superconducting LC circuits and superconducting cavity modes. One of the values get quantized for QHOs is the system energy, which can only take equally-spaced non-zero discrete values. The lowest allowed energy is not 0 but $\frac{1}{2}$ (up to normalization). We call the quantum state with the lowest energy the ground state. For a motion with a certain energy, the phase space diagram is not a circle anymore but a quasi-distribution that can be described by the Wigner function. We say the distribution is a “quasi” distribution because the probability can be negative. The phase space diagram for the ground state and first excited state is plot in Fig. 15.




Fig. 16: Left: an LC circuit. In superconducting LC circuits, normal current becomes superconducting current. Right: the energy potential of a Harmonic Oscillator. In QHOs like the superconducting LC circuits, the system energy becomes equally spaced discrete values. The plotted two levels are the ground state and the first excited state

Superconducting Charge Qubits The QHO does not allow us selectively address the energy levels, thus leakage errors will occur if we use the lowest two levels as the qubit logic space. For example, a control signal that provides the energy difference ΔE enables the transition $|0\rangle \rightarrow |1\rangle$, but will also make the transition $|1\rangle \rightarrow |2\rangle$ which brings the state out of the logic space. To avoid this problem, the Cooper Pair Box (CPB) design of a superconducting charge qubit replaces the inductor (see Fig. 17) with a Josephson junction, making the circuit an anharmonic oscillator, in which the energy levels are not equally spaced anymore. The Wigner function for CPB eigenstates are visually similar to those of QHO and only differ from them to the first order of the anharmonicity, thus we do not plot them in Fig. 15 separately.




Fig. 17: Left: an LC circuit. In superconducting LC circuits, normal current becomes superconducting current. Right: the energy potential of a Harmonic Oscillator. In QHOs like the superconducting LC circuits, the system energy becomes equally spaced discrete values. The plotted two levels are the ground state and the first excited state.

B. The Heisenberg uncertainty principle

We hope that with utilizing the whole physical states (higher energy levels), we can use the redundant space to encode and extract error information. However, the Heisenberg uncertainty principle sets the fundamental limit on what error information we can extract from the physical states — the more we know about the q variable, the less we know about the p variable. For example, we can “squeeze” the ground state of the QHO (also known as the vacuum state) in the p direction, however, the distribution in the q direction spreads, as shown in Fig. 18. Usually, we have to know both the p and q value to characterize the error information unless we know the error is biased. Thus, it’s a great challenge to design encodings in the phase space to reveal error information.




Fig. 18: A squeezed vacuum state

C. The GKP encoding

The GKP states are also called the grid states because each of them is a rectangular lattice in the phase space (see Fig. 15). There are also other types of lattice in the GKP family, for example, the hexagonal GKP [29]. Intuitively, the GKP encoding “breaks” the Heisenberg uncertainty principle—we do not know what are the measured p and q values of the state (thus expected values of p and q remain uncertain), but we do know that they must be integer multiples of the spacing of the grid. Thus, we have access to the error information in both directions and if we measure values that are not multiples of the spacing of the grid, we know there must be errors. Formally, the ideal GKP logical states are given by,

$$\begin{aligned} |\bar{0}\rangle_{\text{gkp}} &= \sum_{k=-\infty}^{\infty} S_p^k |q=0\rangle \\ |\bar{1}\rangle_{\text{gkp}} &= \sum_{k=-\infty}^{\infty} S_p^k |q=\sqrt{\pi}\rangle, \end{aligned} \quad (1)$$

where $S_p = e^{-2i\sqrt{\pi}p}$ is the displacement operator in q space, which shift a wave function in the q direction by $2\sqrt{\pi}$. These definitions show that for GKP logical 0 and 1, the spacing of the grid in q direction is $2\sqrt{\pi}$ and the spacing in the p is $\sqrt{\pi}$. In q direction, the logical $|0\rangle$ state has peaks at even multiples of $\sqrt{\pi}$ and the logical $|1\rangle$ state has peaks at odd multiples of $\sqrt{\pi}$. For logical $|+\rangle$ and $|-\rangle$, the spacing in p and q grid s are switched.

Approximate GKP states The ideal GKP states require infinite energy thus are not realistic. In the lab, we can prepare approximate GKP states as illustrated in Fig. 19, where peaks and the envelope are Gaussian curve.




Fig. 19: Approximate GKP $|0\rangle$ state in q and p axis.

Error correction with GKP qubits GKP qubits are designed to correct shift errors in q and p axis. A simple decoding strategy will be shifting the GKP state back to the closest peak. For example, if we measure a q value to be $2\sqrt{\pi} + \Delta q$, where $\Delta q < \frac{\sqrt{\pi}}{2}$, then we can shift it back to $2\sqrt{\pi}$. With this simple decoding, GKP can correct all shift errors smaller than $\frac{\sqrt{\pi}}{2}$. While there are other proposals for encoding qubits in QHO [54], [52], [47] that are designed for realistic errors such as photon loss, it is shown that GKP qubits have the most error correcting ability in the regime of experimental relevance [5].

In addition, GKP qubits can also provide error correction information when concatenating with Quantum Error Correction Codes (QECC) and yield higher thresholds. For example, when combining the GKP qubits with a surface code, the measured continuous p and q value in the stabilizer measurement can

reveal more about the error distribution than traditional qubits [53], [71], [23].

Lastly, it has been shown that given a supply of GKP-encoded Pauli eigenstates, universal fault-tolerant quantum computation can be achieved using only Gaussian operations [8]. Comparing to qubit error correction codes, the GKP encoding enables much simpler fault-tolerant constructions.

D. Fault-Tolerant Preparation of Approximate GKP States

The GKP encoding has straightforward logical operation and promising error correcting performance. However, the difficulty of using GKP qubits in QIP platforms lies in its preparation since they live in highly non-classical states with relatively high mean photon number (*i.e.*, the average energy levels). Thus, reliable preparation of encoded GKP states is an important problem. In [66], we gave fault-tolerance definitions for GKP preparation in superconducting cavities and designed a protocol that fault-tolerantly prepares the GKP states. We briefly describe the main ideas.

1) Goodness of approximate GKP states: Naturally because of the finite width of the peaks of approximate GKP states, it will not be possible to correct a shift error in p or q of magnitude at most $\frac{\sqrt{\pi}}{2}$ with certainty. For example, suppose we have an approximate $|\bar{0}\rangle$ GKP state with a peak at $q=0$ subject to a shift error e^{-ivp} with $|v| \leq \frac{\sqrt{\pi}}{2}$. The finite width of the Gaussian peaks will have a non-zero overlap in the region $\frac{\sqrt{\pi}}{2} < q < \frac{3\sqrt{\pi}}{2}$ and $-\frac{3\sqrt{\pi}}{2} < q < -\frac{\sqrt{\pi}}{2}$. Thus with non-zero probability the state can be decoded to $|\bar{1}\rangle$ instead of $|\bar{0}\rangle$ (see Fig. 20 for an illustration).

In general, if an approximate GKP state is afflicted by a correctable shift error, we would like the probability of decoding to the incorrect logical state to be as small as possible. A smaller overlap of the approximate GKP state in regions in q and p space that lead to decoding the state to the wrong logical state will lead to a higher probability of correcting a correctable shift error by a perfect GKP state.




Fig. 20: Peaks centered at even integer multiples of $\sqrt{\pi}$ in q space. The peak on the left contains large tails that extend into the region where a shift error is decoded to the logical $|\bar{1}\rangle$ state. The peak on the right is much narrower. Consequently for some interval δ , the peak on the right will correct shift errors of size $\frac{\sqrt{\pi}}{2} - \delta$ with higher probability than the peak on the left.




Fig. 21: Phase estimation circuit with the flag qubit. The protocol is aborted if the flag qubit measurement is non-trivial.

2) *Preparation of approximate GKP states using Phase Estimation:* We observe that the GKP states are the eigenstates of the S_p operator, thus we can use phase estimation to gradually project a squeezed vacuum state to an approximate GKP state. The phase estimation circuit for preparing an approximate $|\tilde{0}\rangle$ GKP state is given in Fig. 21. The first horizontal line represents the cavity mode that we want to prepare the GKP states. The second line is a transmon ancilla initialized to $|+\rangle$. The third line is a transmon flag qubit initialized to $|0\rangle$. The H gate is the Hadamard gate. $\Lambda(e^{i\gamma}) = \text{diag}(1, e^{i\gamma})$ is the gate with a control parameter γ in each round of the phase estimation to. After applying several rounds of the circuit in Fig. 21, the input squeezed vacuum state is projected onto an approximate eigenstate of S_p with some random eigenvalue $e^{i\theta}$. Additionally, an estimated value for the phase θ is obtained. After computing the phase, the state can be shifted back to an approximate $+1$ eigenstate of S_p .

In our protocol, we use a flag qubit to detect any damping event during in the controlled-displacement gate, if a non-trivial measurement is obtained, we abort the protocol and start over. Using our simulation results, we also find a subset of output states that are robust to measurement errors in the transmon ancilla and only accept states in that subset. We proved that our protocol is fault-tolerant according to the definition we gave. In practice, our protocol produces “good” approximate GKP states with high probability and we expect to see experimental efforts to implement our protocol.

E. Discussion

The GKP qubit architecture is a promising candidate for both near-term and fault-tolerant quantum computing implementations. With intrinsic error correcting capabilities, the GKP qubit breaks the abstraction layer between error correction and the physical implementation of qubits. In [66], we discussed the fault-tolerant preparation of GKP qubits and realistic experimental difficulties. We believe that qubit encodings like the GKP encoding will be useful for reliable quantum computing.

VI. CONCLUSION AND FUTURE DIRECTIONS

In this review, we proposed that greater quantum efficiency can be achieved by breaking abstraction layers in the quantum

computing stack. We examined some previous work in this direction that are closing the gap between current quantum technology and real-world quantum computing applications. We would also like to briefly discuss some promising future directions along this line.

A. Noise-Tailoring Compilation

We can further explore the idea of breaking the ISA abstraction. Near-term quantum devices have errors from elementary operations like 1- and 2-qubit gates, but also emergent error modes like cross-talk. Emergent error modes are hard to characterize and to mitigate. Recently, it has been shown that randomized compiling could transform complicated noise channels including cross-talk, SPAM errors and readout errors into simple stochastic Pauli errors [72], which could potentially enable subsequent noise-adaptive compilation optimizations. We believe if compilation schemes that combine noise tailoring and noise adaptation could be designed, they will outperform existing compilation methods.

B. Algorithm-Level Error Correction

Near-term quantum algorithms such as Variational Quantum Eigensolver (VQE) and Quantum Approximate Optimization Algorithm (QAOA) are tailored for NISQ hardware, breaking the circuit/ISA abstraction. We could take a step further and look at high level algorithms equipped with customized error correction/mitigation schemes. Prominent examples of this idea are the Generalized Superfast Encoding (GSE) [65] and the Majorana Loop Stabilizer Code (MLSC) [30] for quantum chemistry. In GSE and MLSC, the overhead of mapping Fermionic operators onto qubit operators stays constant with the qubit number (as opposed to linear scaling in the usual Jordan-Wigner encoding or logarithmic in Bravyi-Kitaev encoding). On the other hand, qubit operators in these mappings are logical operators of a distance 3 stabilizer error correction code so that we can correct all weight 1 qubit errors in the algorithm with stabilizer measurements. These work are the first attempts to algorithm-level error correction and we are expecting to see more efforts of this kind to improve the robustness of near-term algorithms.

C. Dissipation-Assisted Error Mitigation

We generally think of dissipation as competing with quantum coherence. However, with careful design of the quantum system, dissipation can be engineered and used for improving the stability of the underlying qubit state. Previous work on autonomous qubit stabilization [44] and error correction [31] suggest that properly engineered dissipation could largely extend qubit coherence time. Exploring the design space of such systems and their associated error correction/mitigation schemes might provide alternative paths to an efficient and scalable quantum computing stack.

ACKNOWLEDGEMENTS

This work is funded in part by EPiQC, an NSF Expedition in Computing, under grants CCF-1730449/1832377/1730082; in part by STAQ, under grant NSF Phy-1818914; and in part by DOE grants DE-SC0020289 and DE-SC0020331.

REFERENCES

- [1] Cramming more power into a quantum device. <https://www.ibm.com/blogs/research/2019/03/power-quantum-device/>. Accessed: 2010-09-30.
- [2] Charge-insensitive qubit design derived from the Cooper pair box. *Physical Review A - Atomic, Molecular, and Optical Physics*, 76(4), oct 2007.
- [3] Cirq: A python framework for creating, editing, and invoking noisy intermediate scale quantum (NISQ) circuits. <https://github.com/quantumlib/Cirq>, 2018.
- [4] Code for asymptotic improvements to quantum circuits via qutrits. <https://github.com/epiqc/qutrits>, 2019.
- [5] Victor V Albert, Kyungjoo Noh, Kasper Duivenvoorden, Dylan J Young, R T Brierley, Philip Reinhold, Christophe Vuillot, Linshu Li, Chao Shen, S M Girvin, Barbara M Terhal, and Liang Jiang. Performance and structure of single-mode bosonic codes. *Phys. Rev. A*, 97(3):32346, mar 2018.
- [6] Gadi Aleksandrowicz, Thomas Alexander, Panagiotis Barkoutsos, Luciano Bello, Yael Ben-Haim, David Bucher, Francisco Jose Cabrera-Hernández, Jorge Carballo-Franquis, Adrian Chen, Chun-Fu Chen, Jerry M. Chow, Antonio D. Córcoles-Gonzales, Abigail J. Cross, Andrew Cross, Juan Cruz-Benito, Chris Culver, Salvador De La Puente González, Enrique De La Torre, Delton Ding, Eugene Dumitrescu, Ivan Duran, Pieter Eendebak, Mark Everitt, Ismael Faro Sertage, Albert Frisch, Andreas Fuhrer, Jay Gambetta, Borja Godoy Gago, Juan Gomez-Mosquera, Donny Greenberg, Ikko Hamamura, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Hiroshi Horii, Shaohan Hu, Takashi Imamichi, Toshinari Itoko, Ali Javadi-Abhari, Naoki Kanazawa, Anton Karazeev, Kevin Kruslich, Peng Liu, Yang Luh, Yunho Maeng, Manoel Marques, Francisco Jose Martín-Fernández, Douglas T. McClure, David McKay, Srujan Meesala, Antonio Mezzacapo, Nikolaj Moll, Diego Moreda Rodríguez, Giacomo Nannicini, Paul Nation, Pauline Ollitrault, Lee James O’Riordan, Hanhee Paik, Jesús Pérez, Anna Phan, Marco Pistoia, Viktor Prutyay, Max Reuter, Julia Rice, Abdón Rodríguez Davila, Raymond Harry Putra Rudy, Mingi Ryu, Ninad Sathaye, Chris Schnabel, Eddie Schoute, Kanav Setia, Yunong Shi, Adenilton Silva, Yukio Siraichi, Seyon Sivarajah, John A. Smolin, Mathias Soeken, Hitomi Takahashi, Ivano Tavernelli, Charles Taylor, Pete Taylor, Kenso Trabing, Matthew Treinish, Wes Turner, Desiree Vogt-Lee, Christophe Vuillot, Jonathan A. Wildstrom, Jessica Wilson, Erick Winston, Christopher Wood, Stephen Wood, Stefan Wörner, Ismail Yunus Akhalwaya, and Christa Zoufal. Qiskit: An open-source framework for quantum computing, 2019.
- [7] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhii, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [8] Ben Q. Baragiola, Giacomo Pantaleoni, Rafael N. Alexander, Angela Karanjai, and Nicolas C. Menicucci. All-Gaussian universality and fault tolerance with the Gottesman-Kitaev-Preskill code. *arXiv e-prints*, page arXiv:1903.00012, Feb 2019.
- [9] B. B. Blinov, D. Leibfried, C. Monroe, and D. J. Wineland. Quantum computing with trapped ion hyperfine qubits. In *Experimental Aspects of Quantum Computing*, pages 45–59. Springer US, 2005.
- [10] S. A. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Manenti, M. P. da Silva, R. Sinclair, E. Acala, N. Alidoust, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. El Bouayadi, D. Girshovich, S. Hong, K. Kuang, M. Lenihan, T. Manning, A. Marchenkov, J. Marshall, R. Maydra, Y. Mohan, W. O’Brien, C. Osborn, J. Otterbach, A. Papageorge, J.-P. Paquette, M. Pelstring, A. Polloreno, G. Prawiroatmodjo, V. Rawat, M. Reagor, R. Renzas, N. Rubin, D. Russell, M. Rust, D. Scarabelli, M. Scheer, M. Selvanayagam, R. Smith, A. Staley, M. Suska, N. Tezak, D. C. Thompson, T.-W. To, M. Vahidpour, N. Voderhals, T. Whyland, K. Yadav, W. Zeng, and C. Rigetti. Parametrically Activated Entangling Gates Using Transmon Qubits. *Phys. Rev. Applied*, 10:034050, Sep 2018.
- [11] Andrew M. Childs, Eddie Schoute, and Cem M. Unsal. Circuit transformations for quantum architectures. 2019.
- [12] Jerry M. Chow, A. D. Córcoles, Jay M. Gambetta, Chad Rigetti, B. R. Johnson, John A. Smolin, J. R. Rozen, George A. Keefe, Mary B. Rothwell, Mark B. Ketchen, and M. Steffen. Simple all-microwave entangling gate for fixed-frequency superconducting qubits. *Phys. Rev. Lett.*, 107:080502, Aug 2011.
- [13] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091–4094, 1995.
- [14] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta. Open quantum assembly language, 2017.
- [15] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov. Second order gradient ascent pulse engineering. *Journal of Magnetic Resonance*, 212:412–417, October 2011.
- [16] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536, Aug 2016.
- [17] Yao-Min Di and Hai-Rui Wei. Elementary gates for ternary quantum logic circuit. *arXiv preprint arXiv:1105.5485*, 2011.
- [18] Alexander Erhard, Joel James Wallman, Lukas Postler, Michael Meth, Roman Stricker, Esteban Adrian Martinez, Philipp Schindler, Thomas Monz, Joseph Emerson, and Rainer Blatt. Characterizing large-scale quantum computers via cycle benchmarking, 2019.
- [19] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm. *ArXiv e-prints*, November 2014.
- [20] X. Fu, L. Rieseboos, M. A. Rol, J. van Straten, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, V. Newsum, K. K. L. Loh, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels. eQASM: An Executable Quantum Instruction Set Architecture, 2018.
- [21] X. Fu, M. A. Rol, C. C. Bultink, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels. An experimental microarchitecture for a superconducting quantum processor. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-50 ’17, pages 813–825. ACM, 2017.
- [22] X. Fu, M. A. Rol, C. C. Bultink, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels. A microarchitecture for a superconducting quantum processor. *IEEE Micro*, 38(3):40–47, May 2018.
- [23] Kosuke Fukui, Akihisa Tomita, Atsushi Okamoto, and Keisuke Fujii. High-Threshold Fault-Tolerant Quantum Computation with Analog Quantum Error Correction. *Physical Review X*, 8(2), may 2018.
- [24] Joydip Ghosh, Austin G. Fowler, John M. Martinis, and Michael R. Geller. Understanding the effects of leakage in superconducting quantum-error-detection circuits. *Physical Review A - Atomic, Molecular, and Optical Physics*, 88(6), dec 2013.
- [25] Steffen J. Glaser, Ugo Boscarin, Tommaso Calarco, Christiane P. Koch, Walter Köckenberger, Ronnie Kosloff, Ilya Kuprov, Burkhard Luy, Sophie Schirmer, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. Training schrödinger’s cat: quantum optimal control. *The European Physical Journal D*, 69(12):279, Dec 2015.
- [26] Pranav Gokhale, Jonathan M Baker, Casey Duckering, Natalie C Brown, Kenneth R Brown, and Fred Chong. Extending the frontier of quantum computers with qutrits. *IEEE Micro*, 2020.
- [27] Pranav Gokhale, Jonathan M. Baker, Casey Duckering, Natalie C. Brown, Kenneth R. Brown, and Frederic T. Chong. Asymptotic improvements to quantum circuits via qutrits. In *Proceedings of the 46th International Symposium on Computer Architecture*, ISCA ’19, pages 554–566. New York, NY, USA, 2019. ACM. <http://doi.acm.org/10.1145/3307650.3322253>.
- [28] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. Partial compilation of variational algorithms for noisy intermediate-scale quantum machines. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO ’52, page 266–278, New York, NY, USA, 2019. Association for Computing Machinery.
- [29] Daniel Gottesman, Alexei Kitaev, and John Preskill. Encoding a qubit in an oscillator. *Phys. Rev. A*, 64(1):12310, jun 2001.

- [30] Zhang Jiang, Jarrod McClean, Ryan Babbush, and Hartmut Neven. Majorana loop stabilizer codes for error correction of fermionic quantum simulations, 2018.
- [31] Eliot Kapit. Hardware-efficient and fully autonomous quantum error correction in superconducting circuits. *Physical Review Letters*, 116(15), Apr 2016.
- [32] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. In *SIAM Journal on Scientific Computing*, volume 20, 02 1970.
- [33] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296 – 305, 2005.
- [34] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, USA, 2002.
- [35] 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. Fluctuations of energy-relaxation times in superconducting qubits. *Phys. Rev. Lett.*, 121:090502, Aug 2018.
- [36] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, , and W. D. Oliver. A quantum engineer’s guide to superconducting qubits. 2019.
- [37] Dylan Langharst and Kenneth Rudinger. Crosstalk and drift-detection on the IBM Quantum Experience. In *APS March Meeting Abstracts*, volume 2018 of *APS Meeting Abstracts*, page K39.009, Jan 2018.
- [38] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, A. Aspuru-Guzik, and A. G. White. Towards quantum chemistry on a quantum computer. *Nature Chemistry*, 2:106–111, February 2010.
- [39] Benjamin P. Lanyon, Marco Barbieri, Marcelo P. Almeida, Thomas Jennewein, Timothy C. Ralph, Kevin J. Resch, Geoff J. Pryde, Jeremy L. O’Brien, Alexei Gilchrist, and Andrew G. White. Simplifying quantum logic using higher-dimensional hilbert spaces. *Nature Physics*, 5:134 EP –, 12 2008.
- [40] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Physical Review A*, 95:042318, Apr 2017.
- [41] Norbert M. Linke, Dmitri Maslov, Martin Roetteler, Shantanu Debnath, Caroline Figgatt, Kevin A. Landsman, Kenneth Wright, and Christopher Monroe. Experimental comparison of two quantum computing architectures. *Proceedings of the National Academy of Sciences*, 114(13):3305–3310, 2017.
- [42] Daniel Loss and David P DiVincenzo. Quantum computation with quantum dots. *Physical Review A*, 57(1):120, 1998.
- [43] Daniel Loss and David P. DiVincenzo. Quantum computation with quantum dots. *Physical Review A - Atomic, Molecular, and Optical Physics*, 57(1):120–126, 1998.
- [44] Yao Lu, S. Chakram, Ngainam Leung, Nathan Earnest, Ravi Naik, Ziwen Huang, Peter Groszkowski, Eliot Kapit, Jens Koch, and David Schuster. Universal stabilization of a parametrically coupled qubit. *Physical Review Letters*, 119, 07 2017.
- [45] Dmitri Maslov, Sean M. Falconer, and Michele Mosca. Quantum circuit placement. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27:752–763, 2008.
- [46] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
- [47] Marios H. Michael, Matti Silveri, R. T. Brierley, Victor V. Albert, Juha Salmilehto, Liang Jiang, and S. M. Girvin. New class of quantum error-correcting codes for a bosonic mode. *Physical Review X*, 6(3), 2016.
- [48] Prakash Murali, Jonathan M. Baker, Ali Javadi-Abhari, Frederic T. Chong, and Margaret Martonosi. 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*, ASPLOS ’19, pages 1015–1029, New York, NY, USA, 2019. ACM.
- [49] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. Full-stack, Real-system Quantum Computer Studies: Architectural Comparisons and Design Insights. In *Proceedings of the 46th International Symposium on Computer Architecture*, ISCA ’19, pages 527–540, New York, NY, USA, 2019. ACM.
- [50] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. 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*, ASPLOS ’20, page 1001–1016, New York, NY, USA, 2020. Association for Computing Machinery.
- [51] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
- [52] K Noh, V V Albert, and L Jiang. Quantum Capacity Bounds of Gaussian Thermal Loss Channels and Achievable Rates With Gottesman-Kitaev-Preskill Codes. *IEEE Transactions on Information Theory*, 65(4):2563–2582, apr 2019.
- [53] Kyungjoo Noh and Christopher Chamberland. Fault-tolerant bosonic quantum error correction with the surface-gottesman-kitaev-preskill code. *Phys. Rev. A*, 101:012316, Jan 2020.
- [54] Nissim Ofek, Andrei Petrenko, Reinier Heeres, Philip Reinhold, Zaki Leghtas, Brian Vlastakis, Yehani Liu, Luigi Frunzio, S M Girvin, Liang Jiang, Mazhar Mirrahimi, M H Devoret, and R J Schoelkopf. Demonstrating Quantum Error Correction that Extends the Lifetime of Quantum Information. *arXiv:quant-ph/1602.04768*, 2016.
- [55] Archimedes Pavlidis and Emmanuel Floratos. Arithmetic circuits for multilevel qudits based on quantum fourier transform. *arXiv preprint arXiv:1707.08834*, 2017.
- [56] Qiskit. Qiskit NoiseAdaptiveLayout pass. <https://github.com/Qiskit/qiskit-terra/pull/2089>, 2019. Accessed: 2019-08-01.
- [57] T. C. Ralph, K. J. Resch, and A. Gilchrist. Efficient toffoli gates using qudits. *Phys. Rev. A*, 75:022313, Feb 2007.
- [58] J. Randall, A. M. Lawrence, S. C. Webster, S. Weidt, N. V. Vitanov, and W. K. Hensinger. Generation of high-fidelity quantum control methods for multilevel systems. *Phys. Rev. A*, 98:043414, 10 2018.
- [59] J. Randall, S. Weidt, E. D. Standing, K. Lake, S. C. Webster, D. F. Murgia, T. Navickas, K. Roth, and W. K. Hensinger. Efficient preparation and detection of microwave dressed-state qubits and qutrits with trapped ions. *Phys. Rev. A*, 91:012322, 01 2015.
- [60] Rigetti. PyQuil. <https://github.com/rigetticomputing/pyquil>, 2019. Accessed: 2019-08-01.
- [61] Rigetti Quilc. Use swap fidelity instead of gate time as a distance metric . <https://github.com/rigetti/quilc/pull/395>, 2019. Accessed: 2019-08-01.
- [62] Mohan Sarovar, Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik O. Nielsen, and Robin Blume-Kohout. Detecting crosstalk errors in quantum information processors. 2019.
- [63] N. Schuch and J. Siewert. Natural two-qubit gate for quantum computation using the XY interaction. *Physical Review A*, 67(3):032301, March 2003.
- [64] T. Schulte-Herbrüggen, A. Spoerl, and S. J. Glaser. Quantum CISC Compilation by Optimal Control and Scalable Assembly of Complex Instruction Sets beyond Two-Qubit Gates. *ArXiv e-prints*, December 2007.
- [65] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D. Whitfield. Superfast encodings for fermionic quantum simulation, 2018.
- [66] Yunong Shi, Christopher Chamberland, and Andrew Cross. Fault-tolerant preparation of approximate gkp states. *New Journal of Physics*, 21, 08 2019.
- [67] Yunong Shi, Nelson Leung, Pranav Gokhale, Zane Rossi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 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*, ASPLOS ’19, pages 1031–1044, New York, NY, USA, 2019. ACM. <http://doi.acm.org/10.1145/3297858.3304018>.
- [68] Strawberry fields, 2016.
- [69] Swamit S. Tannu and Moinuddin Qureshi. 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*, MICRO ’52, pages 253–265, New York, NY, USA, 2019. ACM.
- [70] Swamit S. Tannu and Moinuddin K. Qureshi. 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*, ASPLOS ’19, page 987–999, New York, NY, USA, 2019. Association for Computing Machinery.

- [71] Christophe Vuillot, Hamed Asasi, Yang Wang, Leonid P. Pryadko, and Barbara M. Terhal. Quantum error correction with the toric Gottesman-Kitaev-Preskill code. *Physical Review A*, 99(3), mar 2019.
- [72] Joel J. Wallman and Joseph Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Physical Review A*, 94(5), Nov 2016.
- [73] Christopher J. Wood and Jay M. Gambetta. Quantification and characterization of leakage errors. *Physical Review A*, 97(3), mar 2018.
- [74] Alwin Zulehner, Alexandru Paler, and Robert Wille. An efficient methodology for mapping quantum circuits to the ibm qx architectures, 2017.