

Q-Pilot: Field Programmable Qubit Array Compilation with Flying Ancillas

Hanrui Wang^{1*}, Daniel Bochen Tan^{2*}, Pengyu Liu³, Yilian Liu⁴, Jiaqi Gu⁵, Jason Cong², Song Han¹

¹MIT, ²University of California, Los Angeles, ³Carnegie Mellon University, ⁴Cornell University, ⁵Arizona State University, *Equal Contributions

ABSTRACT

Neutral atom arrays have become a promising platform for quantum computing, especially the *field programmable qubit array* (FPQA) endowed with the unique capability of atom movement. This feature allows dynamic alterations in qubit connectivity during runtime, which can reduce the cost of executing long-range gates and improve parallelism. However, this added flexibility introduces new challenges in circuit compilation. Inspired by the placement and routing strategies for FPGAs, we propose to map all data qubits to fixed atoms while utilizing movable atoms to route for 2-qubit gates between data qubits. Coined *flying ancillas*, these mobile atoms function as ancilla qubits, dynamically generated and recycled during execution. We present Q-Pilot, a scalable compiler for FPQA employing flying ancillas to maximize circuit parallelism. For two important quantum applications, quantum simulation and the Quantum Approximate Optimization Algorithm (QAOA), we devise domain-specific routing strategies. In comparison to alternative technologies such as superconducting devices or fixed atom arrays, Q-Pilot effectively harnesses the flexibility of FPQA, achieving reductions of 1.4 \times , 27.7 \times , and 6.3 \times in circuit depth for 100-qubit random, quantum simulation, and QAOA circuits, respectively.

1 INTRODUCTION

Quantum computing (QC) hardware has seen rapid scaling, with superconducting systems offering up to 433 qubits [1–4], and neutral atom arrays reaching 1000+ qubits [5, 50]. Utilizing these machines requires mapping qubits in a quantum program/circuit to physical qubits on the QPU, typically constrained by limited connectivity given by a coupling graph. For example, Fig. 1a illustrates a simple QPU with four physical qubits connected in a ring. 2-Q entangling gates, crucial for quantum programs, are restricted to adjacent physical qubits (e.g., (p_0, p_1)). Consider a quantum program with gates $CZ(q_0, q_1)$, $CZ(q_1, q_2)$, and $CZ(q_2, q_0)$. In Fig. 1b, the initial qubit mapping is $q_i \mapsto p_i$ for $i = 0, 1, 2$. While this mapping supports the first two gates, $CZ(q_2, q_0)$ involves non-adjacent p_2 and p_0 . Here, a SWAP gate is inserted to *route* qubits, transforming the mapping. However, SWAP is costly: it can increase circuit depth, leading to more decoherence noise, and typically requires three 2-Q entangling gates, accumulating gate errors. Given the current QPUs’ relatively high noise levels, as quantum circuits grow, it becomes crucial that compilers minimize the overheads incurred by mapping and routing to optimize performance [7, 8, 20, 29, 42, 44, 46, 51, 56, 59, 60, 76, 77, 85, 86].

A recent breakthrough enables atom movement during quantum circuit execution [11], profoundly impacting compilation by introducing dynamic coupling graphs for QPUs, as opposed to static configurations (Fig.1). In this work, we focus on a field programmable qubit array (FPQA) architecture that incorporates this technology. FPQA features two atom types (Fig.2): SLM atoms (blue) are *fixed* atoms




Figure 1: (a) The coupling graph of a QPU. (b) Qubit mapping and routing. The initial mapping is annotated at the beginning of each wire/qubit. A SWAP gate changes the mapping. (c) Using an ancilla and two more CNOTs to implement $CZ(q_0, q_2)$.

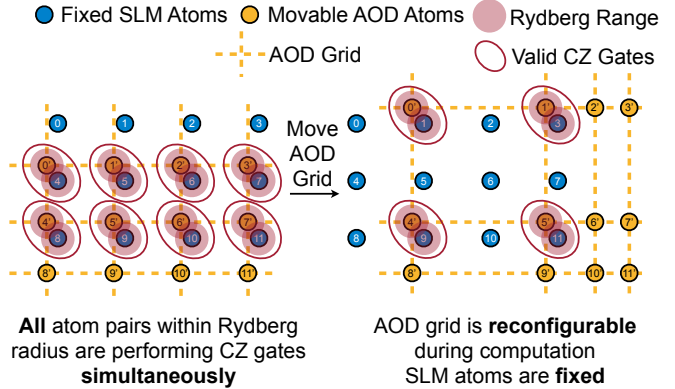


Figure 2: Field Programmable Qubit Array (FPQA).

in traps generated by a spatial light modulator (SLM); AOD atoms (yellow) are *movable* atoms in traps generated by a 2D acousto-optic deflector (AOD). The 2D AOD, a product of two 1D AODs, allows us to specify X coordinates for columns (yellow dashes) and Y coordinates for rows. Consequently, AOD qubits move *by entire rows and columns*. To avoid non-deterministic behavior from trap overlap, we *prohibit AOD rows/columns from moving over others*. Physically, atom movement is a high-fidelity operation primarily constrained by coherence time: with only 0.1% coherence time, an atom can traverse a region for $\sim 2,000$ qubits [11]. These movements are explicitly applied for 2-Q gates, which are induced by a *global* Rydberg laser activating all atoms. If two qubits are within the *Rydberg radius* r_b , they become ‘coupled,’ enabling the application of a CZ gate by the Rydberg laser. Moving atoms between circuit stages couples different qubit pairs, resulting in a dynamic coupling graph. To avoid unintended 2-Q gates, other atoms must be *sufficiently separated* ($> 2.5r_b$). The global Rydberg laser requires less control and calibration, enhancing the *scalability* of FPQA compared to prior works [22, 55] where the laser *individually* address qubits for 2-Q gates.

We introduce Q-Pilot, the first scalable FPQA router drawing inspiration from FPGA placement and routing. Our approach, termed “routing with flying ancillas,” involves qubit mapping akin to cell placement and the use of movable ancilla qubits to bridge fixed atoms, similar to FPGA routing. The advantages of flying ancillas include 1) high-parallelism circuit execution, 2) scalable compilation, and 3) no atom transfer required during computation. To boost parallelism and thus reduce circuit depth, we implement a high-parallelism generic router, dynamically arranging AODs and scheduling 2-Q gates. The router heuristically schedules as many parallel executable gates as possible in one laser stage up to AOD movement constraints. We

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DAC ’24, June 23–27, 2024, San Francisco, CA, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0601-1/24/06.

<https://doi.org/10.1145/3649329.3658470>

also devise *application-specific* strategies: for each Pauli string in quantum simulation, we create multiple ancillas for a “root” qubit and employ graph algorithm to find the longest chain in the SLM array to perform the gates; for QAOA, we create ancilla per qubit instead of per gate, and leverage the commutation of gates to maximize parallel execution and reduce depth. Extensive experiments demonstrate that our FPQA compilation framework outperforms the best baseline, achieving 1.4×, 27.7×, and 6.7× smaller circuit depth for 100-qubit random, quantum simulation, and QAOA circuits.

2 FLYING ANCILLAS

2.1 Motivating Example of Routing CZ

Revisiting the issue of the last gate in Fig. 1, (c) introduces an alternative using ancilla qubit q_A at p_3 instead of the SWAP: q_A is initialized to $|0\rangle$, the three-qubit initial state (in order $q_0q_Aq_2$) can be written as $a|000\rangle + b|001\rangle + c|100\rangle + d|101\rangle$. After the first CNOT, it becomes $a|000\rangle + b|001\rangle + c|110\rangle + d|111\rangle$. After the CZ, it becomes $a|000\rangle + b|001\rangle + c|110\rangle - d|111\rangle$. After the second CNOT, it is $a|000\rangle + b|001\rangle + c|100\rangle - d|101\rangle$, which is the same as the case where CZ(q_0, q_2) acts on the initial state. In this process, q_A acts as a ‘fan-out’ of q_0 . However, note that it is only on the Z basis. Hence, this method’s effectiveness hinges on the targeted 2-Q gate, specifically CZ in our case (and ZZ later on), but it is not universally applicable. Thus, we decompose other 2-Q gates using CZ or ZZ beforehand. Some previous works [21, 25] leveraged these fan-outs to reduce cost in circuit synthesis, but we apply them in routing because uniquely in FPQA, the fan-out qubits can move physically. If we rely on SWAPs for the routing, the depth increases by 3 because we need 3 CNOT for 1 SWAP, yet the new approach only increases depth by 2.

2.2 General Theory of Routing CZs with Ancillas

We prove a general result independent of the coupling graph. Given an arbitrary n -qubit state $\Psi = C_0|0\rangle + C_1|1\rangle + \dots + C_{2^n-1}|2^n-1\rangle$, and a set of qubit pairs C , applying CZ $_{j,j'}$ $\forall (j, j') \in C$ yields

$$\Psi' = \left(\prod_{(j,j') \in C} \text{CZ}_{j,j'} \right) \Psi = \sum_{x=0}^{2^n-1} C_x \prod_{(j,j') \in C} (-1)^{x_j x_{j'}} |x\rangle, \quad (1)$$

where x_j is the j -th bit of x . We consider an alternative procedure as illustrated in Fig. 3 where we 1) apply transversal CNOTs from the n qubits to n fresh ancillas yielding Φ_1 , 2) apply one of the four possibilities (2 choices of whether to $+n$ for 2 indices) CZ $_{j(+n),j'(+n)}$ $\forall (j, j') \in C$ yielding Φ_2 , and 3) apply transversal CNOTs again yielding Φ_3 . We prove that $\Phi_3 = \Psi' \otimes |0^n\rangle$, so our procedure is equivalent to applying the original CZs in Eq. 1.

For every basis state $|x\rangle$ appended with n fresh ancillas, applying transversal CNOTs flips the ancilla state to $|x\rangle$. Thus,

$$\Phi_1 = \left(\prod_{i=0}^{n-1} \text{CNOT}_{i,i+n} \right) (\Psi \otimes |0^n\rangle) = \sum_{x=0}^{2^n-1} C_x |\overline{xx}\rangle, \quad (2)$$

where an overhead line denotes the concatenation of bit-strings. Then, for every pair $(j, j') \in C$, we apply one of the 4 possible CZs,

$$\Phi_2 = \left(\prod_{(j,j') \in C} \text{CZ}_{j(+n),j'(+n)} \right) \Phi_1 = \sum_{x=0}^{2^n-1} C_x \prod_{(j,j') \in C} \left((-1)^{\overline{xx}_j(+n) \cdot \overline{xx}_{j'(+n)}} |\overline{xx}\rangle \right) = \sum_{x=0}^{2^n-1} C_x \prod_{(j,j') \in C} (-1)^{x_j x_{j'}} |\overline{xx}\rangle, \quad (3)$$

where we use the fact that both the j -th bit (from the left) and the $(j+n)$ -th bit of \overline{xx} equals the j -th bit of x , similarly for j' . Applying




Figure 3: The general case of routing CZs with ancillas. The 3 CZs on the right can be executed simultaneously.

transversal CNOTs, again, flips every state $|\overline{xx}\rangle$ back to $|x\rangle|0^n\rangle$, i.e.,

$$\Phi_3 = \left(\prod_{i=0}^{n-1} \text{CNOT}_{i,i+n} \right) \Phi_2 = \Psi' \otimes |0^n\rangle, \quad (4)$$

which finishes our proof.

Note that CZ gates are commutable, so the ones in Eq. 3 can be applied in any order, which may unlock some freedom in scheduling. Moreover, for each $(j, j') \in C$, we have 4 possible CZs to use in Eq. 3 and many of them can be parallelized. For example, in Fig. 3, $n = 3$, and the original CZs are $C = \{(0, 1), (1, 2), (2, 0)\}$ which takes at least 3 steps. Using the procedure just presented, the CZs on $(0+n, 1)$, $(1+n, 2)$, $(2+n, 0)$ can be scheduled to just one step.

2.3 Flying Ancillas in FPQA

The flying ancillas scheme proves particularly advantageous for FPQA over other QC platforms, owing to its high-fidelity movements. The most similar setting is in a multi-chain ion trap QPU [6], where chains of ions are laid out in 1D, and two chains can be moved to merge or split again. However, because there is no distinction between stationary and movable qubits like in FPQA, moving a regular qubit in the ion trap quantum computer has the same cost as moving an ancilla, so the flying ancilla scheme does not hold a big advantage. Additionally, the limited number of qubits available on ion trap QPUs discourages leveraging numerous ancillas. Flying qubits, typically optical, are also employed as communication resources between individual superconducting QPUs but face challenges, including a low interfacing fidelity of approximately 80% per flying qubit [41]. In contrast, in FPQA, the two extra CNOTs required by flying ancilla can achieve 99.5% fidelity, and the ancilla movement has negligible error [19]. Despite this high fidelity, the state-of-the-art FPQA compilation work [61] primarily utilizes only the movement of data qubits for routing, overlooking the potential advantages of routing via flying ancillas.

3 ROUTING FRAMEWORK

3.1 Overview

Given a target problem, the input values to the router are (1) the SLM array parameters (#rows, #columns, and locations), (2) the AOD array parameters (#rows and #columns), and (3) the qubit mapping. We focus on routing in this work, so we simply map qubits in reading order throughout. We refer to them as *configurations* of the FPQA. Based on a configuration and the target problem, we leverage a high-parallelism router to generate an optimized schedule. A fast performance evaluator can efficiently return the corresponding performance metric or cost, including the number of 1-Q gates, 2-Q




Figure 4: The flowchart of the FPQA compilation framework.

Algorithm 1: Generic router for arbitrary quantum circuits

Input : Gates G in a quantum circuit
 $s \leftarrow \emptyset$; Initial schedule;
 $g \leftarrow G$; Initial gate candidates;
 $i \leftarrow 0$;
while $g \neq \emptyset$ **do**
 $p_i = \text{FrontLayer}(g)$;
 $q_i = \emptyset$;
 for $p \in p_i$ **do**
 if $\text{IsLegal}(q_i \cup p)$ **then**
 $q_i \leftarrow q_i \cup p$;
 $g \leftarrow g \setminus p$;
 end
 end
 $s \leftarrow s \cup \text{GenerateSchedule}(q_i)$;
 $i \leftarrow i + 1$;
end
Output: Schedule s with maximum depth i

gates, the circuit depth, and movement distance, which are closely related to the circuit fidelity.

With this performance evaluator, our compiler also supports router-in-the-loop FPQA architecture design space exploration. We can use the evaluated cost as feedback to optimize a configuration that targets higher circuit fidelity iteratively. After certain epochs, the compiler will output the best configuration and optimized schedule.

3.2 Compilation of General Quantum Circuit

The compilation process is shown in Fig. 4. Given a quantum circuit, we first decompose the target circuit into 1-Q rotations and 2-Q CZ gates. Then, the gates are performed in alternating 1-Q and 2-Q stages. In the 1-Q stages, we turn on the individual addressable Raman laser to perform the desired gates on the target qubits. After all the available 1-Q gates are done, we move to 2-Q stages. In such stages, we select a set of CZ gates from the non-dependent front-layer of the circuit that can be performed in parallel. Then, we create flying ancillas from the control qubits and move the ancillas close to their corresponding target qubits. We turn on the Rydberg laser so that the ancillas will perform CZ gates with the target qubits. At last, we recycle these ancillas with CNOT gates and repeat this process. After all gates are done, we perform measurements and get results.

We first introduce the *generic router* for arbitrary quantum circuits. We first transpile a given circuit with the universal gate set: CZ+1Q gates. Then, we iteratively extract the front layer of that circuit. If the front layer contains 1Q gates, we perform the 1Q gates first. After that, the front layer contains only a set of CZ gates.

To maximize the parallelism of the generated schedule while maintaining good scalability, we propose a heuristic-based scheduler that selects as many CZ gates as possible while honoring the constraint of FPQA in a single 2-Q gate stage described in Alg. 1. We show an example in Fig. 5. Given a quantum circuit, we first detect the *source layer of the dependency graph* according to the gate dependency, defined as the maximum potentially parallelizable gates, e.g., (g_0, g_1, g_2, g_3) . From the i -th source layer, we use a greedy algorithm to decide (1) the maximum legal subset of the source layer gates and (2) the grid locations of the AOD ancilla qubits. The key intuition behind this greedy heuristic is that the order of rows and columns cannot be reversed. In the first search step, the router adds g_0 and g_1 to the subset and checks the legality according to the qubit ordering in rows and columns. Since the row/column order of g_0 and g_1 all satisfies $g_0 \leq g_1$, this is confirmed to be a legal subset. Then, in the next step, the router tries to add g_2 to the subset. However, in the column dimension, the 1st qubit is of order $g_0 \leq g_1 \leq g_2$, which conflicts with the 2nd qubit order $g_2 \leq g_0 \leq g_1$. In other words, it is impossible to move AOD ancilla qubits to enable parallel execution of those three gates. This violation of the order rule will kick gate g_2 out of the legal subset. Similarly, the router adds g_3 into the subset and ultimately finds the maximum legal subset (g_0, g_1, g_3) and the corresponding grid locations of three AOD ancilla qubits, i.e., $(0,0)$, $(1,1)$, and $(2,2)$. This is the end of the sub-schedule for the i -th stage. To physically execute this i -th sub-schedule, we perform a three-phase operation: (1) move the AOD ancilla qubits near the target SLM qubits and *copy* their states via CNOTs; (2) move the ancilla qubits to the target qubits to parallel execute the gates in the legal subset; and (3) finally recycle the states from ancilla qubits to the original SLM qubits via CNOTs. After performing a set of 2-Q gates, the router will repeat the process on the next *source layer* of the circuit g_4, g_5, g_2, g_6 . The router will repeat this legal subset selection process on the remaining gates until all gates are executed in the schedule and ultimately output the complete schedule.

Note that the AOD configuration between two iterations can be different. For our example, the ancillas live on different crosses. This can be achieved by transferring a set of spare SLM atoms to arbitrary rows and columns on the AOD grid. The atom loss is not fatal in

Algorithm 2: Customized router for quantum simulation

Data: List P : qubits in the Pauli string with non- I paulis
 $s \leftarrow \emptyset$; Schedule for the compiled program
 $P \leftarrow P \setminus P[0]$; $P[0]$ is the root qubit
 $g \leftarrow \emptyset$; Directed compatibility graph. Two qubits are compatible if and only if there is a path between them.
for $q_i \in P$ **do**
 $g.nodes \leftarrow g.nodes \cup q_i$;
for $q_i \in P$ **do**
 for $q_j \in P \setminus q_i$ **do**
 if $q_j.row \geq q_i.row \ \& \ q_j.col \geq q_i.col$ **then**
 $g.edges \leftarrow g.edges \cup (q_i, q_j)$;
while $g \neq \emptyset$ **do**
 Find the longest path l in g ;
 $s \leftarrow s \cup \text{GenerateSchedule}(l)$;
 $g \leftarrow g \setminus \{n | n \in l\}$;
Result: Schedule s

our scheme because they do not contain any quantum information during transfer.

3.3 Customized Router for Quantum Simulation

For specific applications, we propose domain-specific routing strategies for higher parallelism. The first application is quantum simulation. To simulate the evolution under a Pauli string, the core part of the simulation algorithm works as follows: First, select a starting qubit inside the given Pauli string and then perform CNOTs on all pairs between the starting one and other qubits in the string.

We propose a longest path-based algorithm to compile this problem on the FPQA, described in Alg. 2. We configure the AOD array so that all ancilla qubits are on the *diagonal* of the grid and can be moved with the best flexibility. Then, we select the qubit i with the smallest index and fan-out its state to all AOD ancilla qubits. To maximize the parallelism, we need to find the longest legal path in the directed dependency graph, where each qubit points to all other qubits in its lower-right corner, as shown in Fig. 6.

Given this longest path, we move the AOD ancilla qubits to their target SLM qubits and perform CNOTs in parallel. Further, those executed qubits will be removed from the candidate set, and the longest path-finding procedure will repeat until all gates are executed. Note that this longest path-finding can be implemented efficiently with dynamic programming. Compared to the generic router, which applies atom transfer to create and recycle ancilla qubits at each stage, this specialized quantum simulation router will maintain the states on the ancilla qubits across stages for one Pauli string, thus having a lower overhead. To generate N fan-out qubits for N non-identity Pauli operators in a string, we initiate a fan-out operation by relocating the root SLM qubit near the target and executing a CNOT gate, as indicated by the underline X below. Additional fan-out qubits can be generated by performing CNOT operations with adjacent qubits. This leads to a geometric progression in the number of new fan-out qubits—1, 2, 4, 6, 8, etc.—yielding a circuit depth of $O(\sqrt{N})$.

3.4 Customized Router for QAOA

Another task that can be highly parallel is Quantum Approximate Optimization Algorithm (QAOA). In QAOA, we are given a graph, and our target is to perform 2-Q gates on every edge in the graph shown in Fig. 7.

```

O O O O O O O O O O O O O O O O O O O O O O O O O O O O
O O O O O O O O O O O O O X O O O O O O O O O O O O O +1
O O O O O X O O O O O O X X O O O O O O O O O O O O +2
O O O O O X X O O O O X X X X O O O X O O O O O O +4
O O O O X X X X O O X X X X X X O O X X O X O O +6
O X O X X X X X X X X X X X X X X X X X X X X X O +8
X X X X X X X X X X X X X X X X X X X X X X X X X X X
```

Algorithm 3: QAOA compilation with flying ancilla and high-parallelism router

Input: edges: List $[(q_1, q_2)]$ is a list of edges
Program: an empty quantum program
while edges is not empty **do**
 cancel_pairs_first_row=[]
 Find e_0 in edges with smallest $e_0.q_1$.
 cancel_pairs_first_row.append(e_0)
 Program.cancel(e_0)
 while True **do**
 Find e in edges with smallest $e[0]$, which is compatible and $e.q_1.y = e_0.q_1.y$.
 if can't find then
 Break
 cancel_pairs_first_row.append(e)
 Program.cancel(e)
 for $i = 1 \dots n_row, i' = i \dots n_row$ **do**
 match=True
 for e in cancel_pairs_first_row **do**
 replace $e.q_1.y$ with i , replace $e.q_2.y$ with i'
 if e not in edges **then**
 match=False
 if match **then**
 for e in cancel_pairs_first_row **do**
 replace $e.q_1.y$ with i , replace $e.q_2.y$ with i'
 Program.cancel(e)
Output: Compiled Program

First, we create one ancilla for each qubit. These ancillas will be recycled once the whole graph is done. Then, our router completes this task in a multi-stage way. Each stage will perform one or multiple 2-Q gates corresponding to some edges in the graph. We illustrate the detailed procedure of the first stage in Fig. 7. Among all the qubit pairs, we select the one with the smallest index as the highest-priority pair to begin with, e.g., $(0', 1)$. Since each AOD row and column must move simultaneously, we check which pair can be performed inside the same row, e.g., $(1', 3)$ is matched in this case. The rest of the AOD columns have already moved outside the SLM array. Then, they will not interact with any SLM atoms. Once the locations of all ancilla qubits in the first row are determined, other ancilla qubits on the rest of the rows can only move vertically due to the grid constraint. Then, we need to determine the vertical location of each row one by one. For the second AOD row, we found the best vertical location that can allow the most pairs to interact, e.g., in this case, we match two pairs $(4', 9)$ and $(5', 11)$. Note that any undesired interaction is illegal and thus should be avoided. This process will repeat until no rows can legally interact with any SLM atoms. Then, we can determine the locations of all AOD qubits and turn on the Rydberg laser to perform the parallel 2-Q gates. This greedy algorithm, details

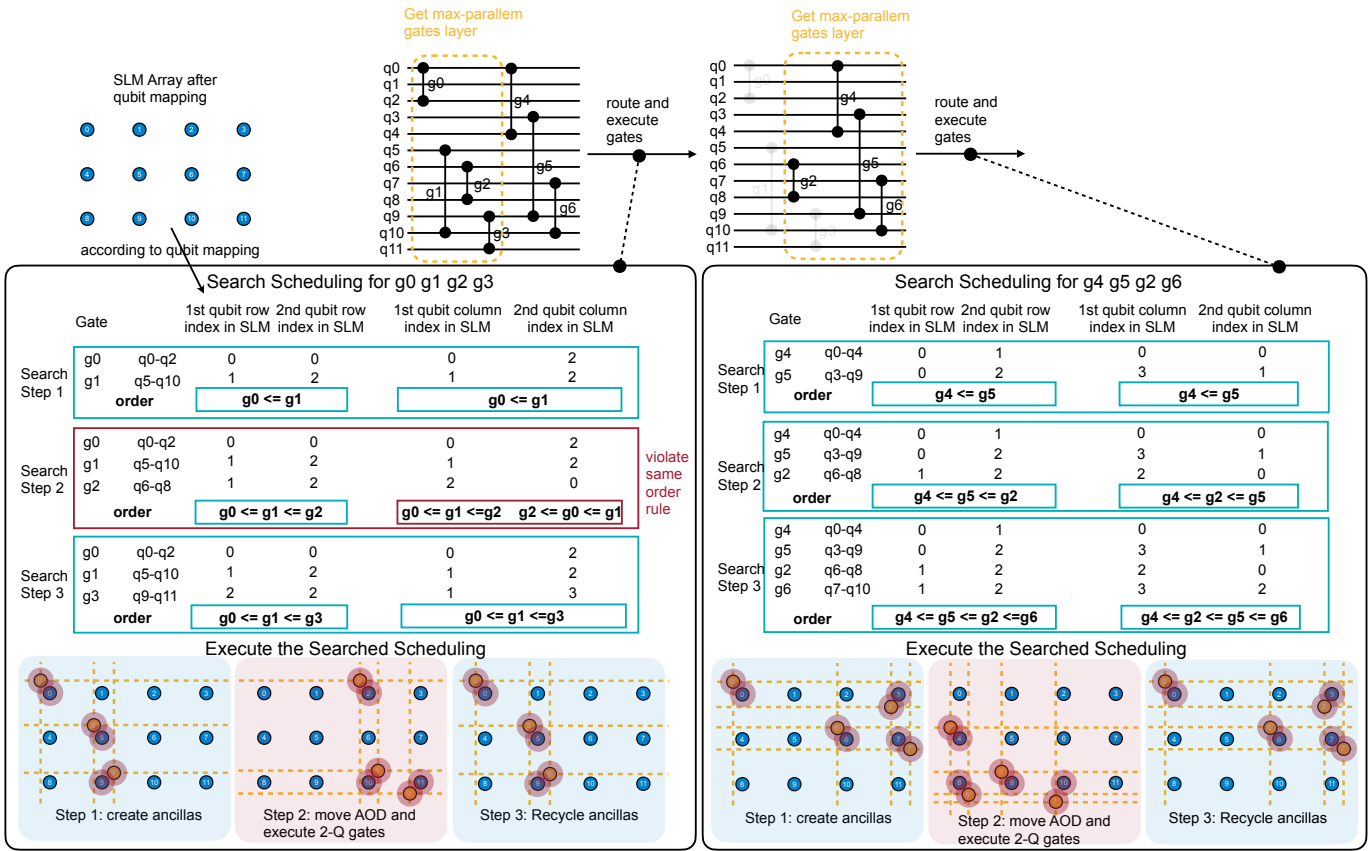


Figure 5: Routing process of an example circuit using the generic router. The router adds as many as possible to one stage to execute them simultaneously.

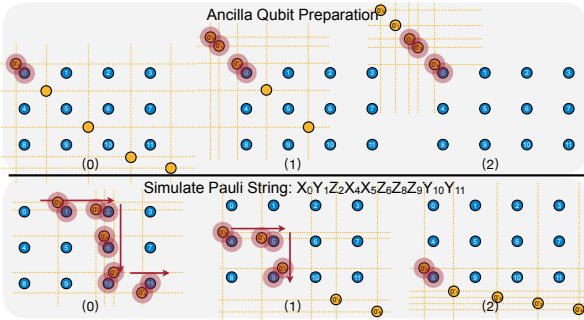


Figure 6: Q-Pilot routing quantum simulation circuits.

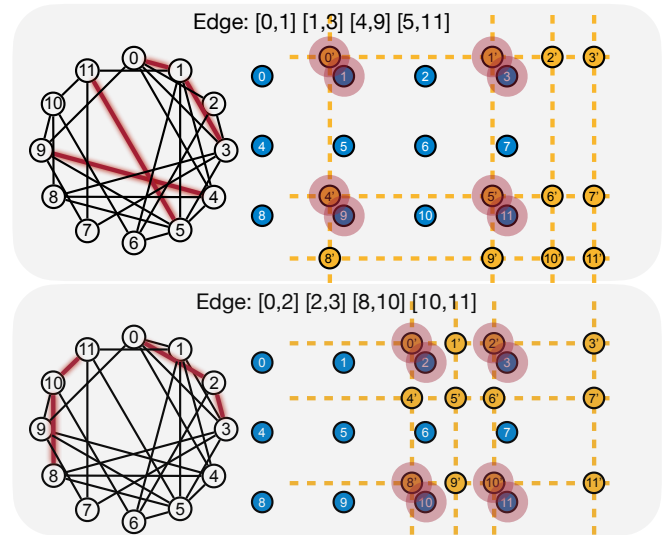


Figure 7: Scheduling of a QAOA circuit using Q-Pilot. The graph representation of the circuit is shown on the left. The edges correspond to interactions between two qubits.

in Alg. 3, always tries to achieve the maximum matching on the first row and ultimately reaches a schedule with max parallelism.

So far, we have finished the first stage of the schedule with four 2-Q gates being performed. In the second stage, the highest-priority pair now becomes (0',2), and the same procedure as stage one can be applied to find a legal schedule with maximal parallelism. After t stages, the compilation flow ends with a t -stage legal schedule where all 2-Q gates are performed. Lastly, we recycle the ancillas and complete the task.

Fig.9 depicts the spatiotemporal dynamics of a 100-qubit QAOA circuit, revealing a periodic pattern in atom positions due to iterative

optimization for parallel gate execution. The figure also includes histograms detailing the frequency of movements, distances traversed, and average speeds. Based on realistic parameters from [11], the typical speed is measured at 0.15 m/s. Fig. 10 shows the execution




Figure 8: Q-Pilot compiling one stage of an example QAOA circuit.

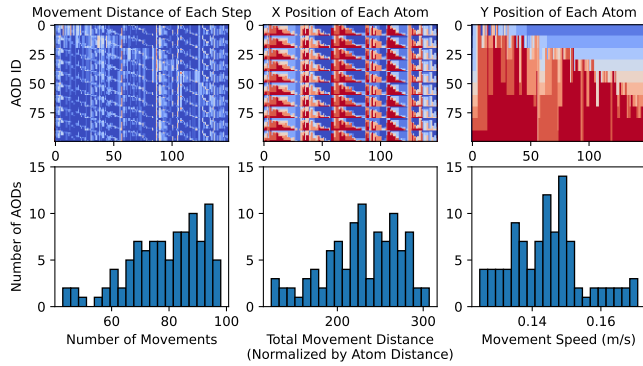


Figure 9: Movement spatiotemporal patterns.

timeline of a program compiled with the FPQA where movements are the largest part.

4 EVALUATION

4.1 Evaluation Methodology

Benchmarks. We utilize three benchmark types: random, quantum simulation, and QAOA circuits. Benchmarks were created for 5, 10, 20, 50, and 100 qubits. Random circuits were generated with Qiskit’s `random_circuit` function, which randomly places 1-Q and 2-Q gates on qubits. The number of CNOT gates is set at $2x$, $5x$, $10x$, $20x$, and $50x$ the qubit count. Quantum simulation circuits were formed from 100 random Pauli strings. The probability p of a qubit having a Pauli operator X , Y , or Z varies from 0.1 to 0.5. QAOA circuits were constructed using ZZ gates between random qubit pairs. These pairs had an edge probability p of 0.1 to 0.5. We also designed specific QAOA circuits based on 3-regular and 4-regular graphs. These circuits also used 5, 10, 20, 50, and 100 qubits.

Baselines. 3 devices as baselines were chosen: the 127-qubit IBM Washington machine, a 16×16 square lattice, and a 16×16 triangular lattice of fixed neutral atoms, following Ref. [61]. The IBM machine features a heavy hexagon coupling graph. The square lattice’s atoms connect to four nearest neighbors, while the triangular lattice’s atoms connect to six. Qiskit’s transpiler compiled benchmark circuits onto these devices at optimization level 3. Circuit depth, defined as the number of parallel 2-Q gate layers, was a key comparison metric, alongside the number of 2-Q gates in each *compiled* circuit for the baseline devices and Q-Pilot. Additionally, Q-Pilot was benchmarked




Figure 10: Detailed execution of compiled program.




Figure 11: Comparison of compiled 2-Q gate count and circuit depth between Q-Pilot and the three baselines on random circuits. The random circuits vary in size, from 5-Q to 100-Q, and have 2-Q gate count between $2 \times$ and $10 \times$ qubit count.

against the solver-based compiler from Ref. [61], used for QAOA problems on 3- and 4-regular graphs. Comparisons included circuit depths and compilation times, with a 4,000s timeout (\sim an hour) set for the solver-based compiler due to its exponential runtime scaling.

4.2 Main Results

Results on random circuits. Fig. 11 shows the results of compiling random circuits. Compared with three baseline devices, for 100 qubits Q-Pilot shows an average of $4.2 \times$ reduction in the compiled 2-Q gate count, as well as an average of $1.4 \times$ reduction in compiled circuit depth compared with the best-performing baseline approach.

Results on quantum simulation circuits. Fig. 12 shows the results of compiling quantum simulation circuits. For Pauli probabilities 50%, Q-Pilot shows an average of $6.9 \times$ reduction in the compiled 2-Q gate count and an average of $27.7 \times$ reduction in compiled circuit depth compared with the best-performing baseline on 100-qubit circuits compared with the best baseline. Besides the random Pauli strings, we also test with the Pauli strings used in some molecule simulation problems [30]. As shown in Table 1, Q-Pilot shows an average $1.36 \times$ reduction in the 2-Q gate count and average $2.60 \times$ circuit depth reduction over the best baseline.

Results on QAOA circuits. Fig. 13 shows the results of compiling Max-Cut QAOA circuits for 4-regular graphs and random graphs

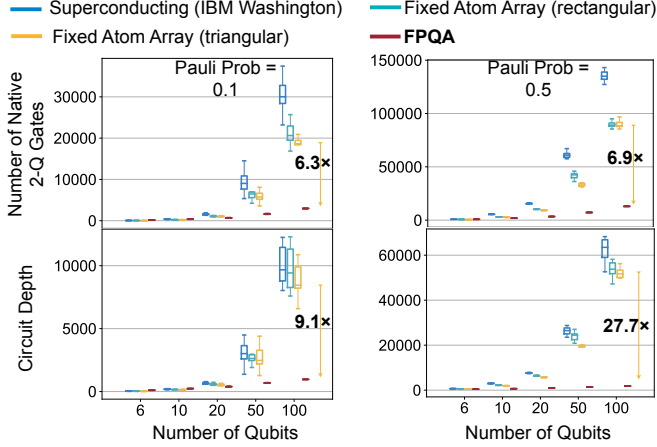


Figure 12: Comparison of compiled 2-Q gate count and circuit depth between Q-Pilot and the three baselines on quantum simulation circuits from 5-Q to 100-Q. The circuits are generated with Pauli probability $p = 0.1$ and 0.5 .

Table 1: Quantum Simulation for Molecule Pauli strings.

Benchmark	Device	H2	LiH_UCCSD	H2O	BeH2
Depth	FAA(rectangular)	76	2,772	31,087	43,919
	FAA(triangular)	61	2,052	26,189	37,314
	Superconducting	77	3,403	40,080	59,259
	Ours	61	849	7,585	10,617
#2Q Gate	FAA(rectangular)	82	3,577	41,306	58,720
	FAA(triangular)	73	2,616	35,353	51,699
	Superconducting	85	5,082	67,247	103,594
	Ours	94	2,130	20,966	29,518

with edge occupancy 30%. Q-Pilot again shows an average of 10.0 \times reduction in compiled 2-Q gate count and an average of 6.7 \times reduction in the compiled circuit depth.

Comparison with the Solver-Based Compiler. As illustrated in Table 2, we compare Q-Pilot against the solver-based methods [61, 62] in compiling QAOA circuits for regular graphs. Ref. [62] relaxes the formulation of Ref. [61] to tradeoff compilation time and quality. While these methods achieve better solutions, they struggle with larger problems, often failing to find a solution within an hour due to its exponential runtime scaling. In contrast, Q-Pilot efficiently compiles all these problems in under 1 second, with the compiled circuit depth not exceeding 4 \times the optimal depth.

4.3 Analysis

Impact of Array Size on Circuit Depth. Fig. 14 shows how array sizes affect the compiled circuit depth. We organized the qubits into rectangular arrays of varying widths (8, 16, 32, 64, 128), with the optimal array widths marked by stars in the figure. Optimal array widths vary across different problems, highlighting a tradeoff between greater parallelism within a row and across different rows. Specifically, in Fig. 14, we observe that QAOA circuits achieve optimal performance with large array width (128), while random circuits and quantum simulation problems are best served with moderate array widths (64 or 32 in our study). The insight here is while larger array widths offer more parallel execution paths, they might not always correspond to increased efficiency for all types of problems, possibly due to overheads or specific characteristics of the circuit structure.

How does the 2-Q gate error rate affect the overall error rate? Fig. 15 (a) shows the relation between the overall error rate and the




Figure 13: Comparison of compiled 2-Q gate count and circuit depth between Q-Pilot and the three baselines on QAOA circuits. The QAOA circuits vary in size, from 6-Q to 100-Q, and are generated with edge $p = 0.3$ and 4-regular graphs.

Table 2: Comparison of Q-Pilot with solver based method.

Benchmark		6Q	10Q	20Q	50Q	100Q	
3-reg.	runtime(s)	solver [61]	0.173	0.381	74.5	timeout	timeout
		iter-p [62]	0.509	2.16	14.6	966	timeout
		Ours	5.57E-3	9.89E-3	1.07E-2	7.52E-2	1.77E-1
	depth	solver [61]	3	3	3	-	-
		iter-p [62]	3	5	6	10	-
		Ours	5	7	11	24	45
4-reg.	runtime(s)	solver [61]	18.1	3.93E3	timeout	timeout	timeout
		iter-p [62]	0.852	2.64	23.4	2.34E3	timeout
		Ours	6.25E-3	9.31E-3	2.10E-2	7.23E-2	3.42E-1
	depth	solver [61]	5	5	-	-	-
		iter-p [62]	5	6	8	15	-
		Ours	6	9	15	32	60

2-Q gate error rate. We model the circuit error with the equation introduced [61], where ϵ is the overall error rate:

$$\epsilon = 1 - f_2^{NT} f_1^{G_1} \exp\left(-N \frac{\sum_i T_0 \sqrt{D_i}}{T_2}\right), \quad (5)$$

N is the maximum number of qubits used (including AOD and SLM), and T is the circuit depth. G_1 is the number of 1-Q gates. f_1 and f_2 are the fidelity of 1-Q and 2-Q gates, respectively. T_2 is the coherence time of the qubit, and T_0 is the characteristic time of atom movement. D_i is the maximum distance atoms moved in stage i . In our estimation, we choose $f_1 = 99.9\%$, $T_2 = 1.5s$, and $T_0 = 300\mu s$ [61]. The three benchmarks used here are 1) quantum simulation circuits with 5 qubits and 100 non-trivial Pauli strings with $p = 0.1$, 2) random 5Q circuits with an average of two 2-Q gates per qubit, and 3) QAOA circuits for random 3-regular graphs. The error rates are below 0.5 when the 2-Q gate has an error rate below 10^{-3} .

What is the distribution of the parallelism? Fig. 15 (b) shows the percentage of stages with the number of 2-Q gates simultaneously executed for QAOA problems. The average parallelism of 20Q, 50Q, and 100Q problems are 3.32, 4.13, and 4.90, respectively. As the problem scales up, the parallelism of the problem is also increased.

Whether the application-specific compilers bring better performance for quantum simulation and QAOA? Fig. 16 shows the advantage of the domain-specific compiler compared to the general compiler. For quantum simulation, the domain-specific compiler reduces the 2-Q gate count by 1.5 \times and the circuit depth by 8.8 \times . For




Figure 14: Circuit depth vs array width of SLM and AOD arrays in FPQA. The three benchmark circuits are shown here for 50 qubits and 100 qubits. The star in each graph marks the optimal array width for the smallest circuit depth.




Figure 15: (a) Overall error rate vs. 2-Q gate error rate for random 6Q circuits with two 2Q gates per qubit, QAOA circuits based on random 3-regular graphs, and 5Q quantum simulation circuits with 100 Pauli strings and $p = 0.1$. (b) Ratio of total stages vs number of parallel 2-Q gates in a stage using Q-Pilot on QAOA circuits with 20-Q, 50-Q, and 100-Q.

QAOA, the domain-specific compiler reduces the 2-Q gate count by 2.8 \times and the circuit depth by 10.1 \times . The advantages come from domain-specific heuristics that minimize the circuit depth.

How scalable is the Q-Pilot? We test Q-Pilot with a large number of qubits to show its scalability. For the QAOA problem, we choose random graphs with edge $p = 0.5$. It takes 1.51s, 10.75s, and 129.50s to compile 500, 1000, and 2000 qubits. For quantum simulation problems, we choose 100 random Pauli strings. It takes 6.91s, 14.28s, and 30.48s to compile 500, 1000, and 2000 qubits. We generate random circuits with a depth of 10 for general circuits, and it takes 2.64s, 8.70s, and 32.31s to compile 500, 1000, and 2000 qubits. The fast speed proves that Q-Pilot is scalable and can handle large-scale problems.

5 RELATED WORKS

Compilers for Neutral Atom Arrays. Previously, research focused on fixed atom arrays with static coupling. Ref. [9] introduced the first compiler framework for this architecture, extending existing techniques for superconducting devices and addressing unique constraints, including long-range interaction restriction zones and sporadic atom loss. Ref. [32] further considers gate durations in compilation. Geyser (Ref. [52]) leverages native 3-Q operations by

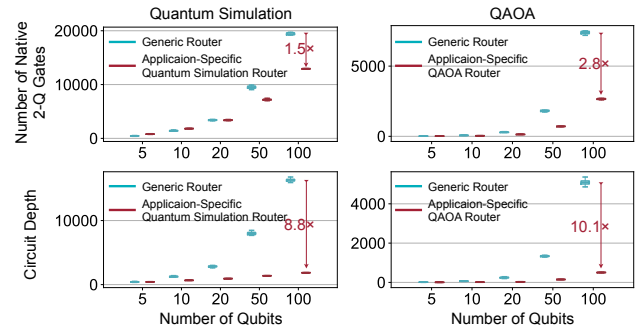


Figure 16: Advantage of our application specific Quantum Simulation and QAOA router comparing to the generic router.

blocking 3-Q sub-circuits and re-synthesizing them. A concurrent work FPQA-C [74] proposed a MAX k-Cut-based algorithm to perform qubit mapping and developed a router to schedule the gates. It does not leverage any ancilla qubits and thus is sub-optimal on circuit depth.

Ref. [12] first considered atom movement, exploring a hypothetical architecture with ‘1D displacement’ reconfigurability, more restrictive than FPQA. Ref. [61] presented the initial FPQA compiler, formulating constraints and using an SMT solver for qubit mapping and routing. However, the solver-based method’s scalability is limited by the exponential SMT solving. Their updated work [62], sacrificing some optimality for scalability, still struggles to handle a 100-qubit circuit within a day. Additionally, it is worth noting that they employ *atom transfer* operations, moving an AOD atom to an empty SLM trap when in proximity and vice versa. While atom transfer is already utilized in experiments [18], frequent transfers ‘heat up’ the atoms, potentially resulting in atom loss errors.

Compilers for Emerging Quantum Architectures. Superconducting quantum computers have been widely used in quantum computing research. However, there has been a growing interest in exploring alternative quantum computing hardware, such as neutral atom machines and trapped ion machines. Our work delves into the compiler design for neutral atom systems, building upon insights from previous works [9, 52, 61]. Furthermore, we have observed a surge in compiler designs tailored for trapped ion systems in recent years. Notable among these are TILT [78], designed for linear chain trapped ion systems; compilers for Quantum Charge Coupled Device-based trapped ion architectures as discussed in Ref. [48]; and compilers for shuttling-based trapped ion architectures as explored in Ref. [26]. Additionally, a variety of compilation techniques applicable to diverse quantum computing settings have been proposed [10, 13–17, 23, 24, 27, 28, 30, 37, 39, 40, 43, 45, 47, 49, 53, 54, 63–67, 78–81]. In this work, we target an emerging FPQA (Field Programmable Quantum Array) device implemented with dynamically reconfigurable atom arrays. Given the dynamic nature of the coupling map in such devices, traditional compiler techniques cannot be directly applied. To address this challenge, we introduce a novel flying ancilla-based compilation framework. This framework is designed to generate low-depth compiled circuits that maintain high scalability, catering specifically to the unique architecture of FPQA devices.

Qubit Mapping and Instruction Scheduling. Since noise forms the bottleneck of NISQ machines, many noise-adaptive quantum compilation techniques have been proposed [16, 53, 54, 63, 68–73, 75, 84]. Examples include various gate errors which can be suppressed by qubit mapping [28, 39, 45, 47, 64, 81], composite pulses [13, 40, 43,

79], dynamical decoupling [10, 15, 23, 37, 66], randomized compiling [67], hidden inverses [80], instruction scheduling [49, 78, 83], frequency tuning [17, 24, 65], parallel execution on multiple machines [58], algorithm-aware compilation [14, 27, 30], qubit specific basis gate [31, 38] and various pulse level optimizations [33–36]. Qubit mapping and routing, also named quantum layout synthesis or qubit allocation/placement, has been a popular research topic in the QC community [20, 29, 42, 44, 46, 56, 59, 60, 76, 82, 85, 86]. Instruction scheduling has also been widely explored [31, 38, 49, 57, 58, 78]. The most relevant previous work is Brandhofer *et al.* [12], which considers a hypothetical atom array architecture with ‘1D displacement’ reconfigurability, a much more restrictive architecture than FPQA. They construct a *potential* coupling graph of all potential connections and leverage an exact router to insert SWAPs, which limits the scalability.

6 CONCLUSION AND OUTLOOK

We architect the emerging neutral atom arrays – the field programmable quantum array (FPQA) with runtime-movable atoms. On this highly flexible FPQA, we propose mapping all the qubits to the fixed atoms and then leveraging movable atoms as ancillas to route between fixed qubits and perform 2-Q gates. To increase the parallelism of 2-Q gates, we design a high-parallelism generic router for arbitrary circuits and application-specific routing strategies for quantum simulation and QAOA. We hope this work can open up the avenue for future research on FPQA. Specifically, although the heuristics presented yield good results with short runtimes, a more general search framework where one can trade time for even higher solution quality would be valuable. It is also exciting to explore other domains like circuits involved in quantum error correction protocols, which is the building block for future fault-tolerant QPUs made with FPQA. Furthermore, given a scalable compiler, one can use the compilation results to guide hardware developments regarding potential features, e.g., multiple AODs or multiple ‘zones’ in FPQA, some for Rydberg interaction and others only for storage.

ACKNOWLEDGEMENT

This work is partially supported by NSF grant 2313083, MIT-IBM Watson AI Lab, and Qualcomm Innovation Fellowship. The authors would like to thank Dolev Bluvstein, Mikhail D. Lukin, and Hengyuan Zhou for valuable discussions on neutral atom arrays.

REFERENCES

- [1] [n. d.]. <https://newsroom.ibm.com/2022-11-09-IBM-Unveils-400-Qubit-Plus-Quantum-Processor-and-Next-Generation-IBM-Quantum-System-Two>.
- [2] [n. d.]. <https://www.rigetti.com/>.
- [3] [n. d.]. <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>.
- [4] [n. d.]. <https://spectrum.ieee.org/tech-talk/computing/hardware/intels-49qubit-chip-aims-for-quantum-supremacy>.
- [5] [n. d.]. <https://www.quera.com/aquila>.
- [6] [n. d.]. <https://ionq.com/posts/august-25-2021-deep-dive-reconfigurable-multicore-quantum-architecture>.
- [7] M. Alam *et al.* 2020. An efficient circuit compilation flow for quantum approximate optimization algorithm. DAC’20.
- [8] M. Alam *et al.* 2020. Noise resilient compilation policies for quantum approximate optimization algorithm. ICCAD’20.
- [9] J. M. Baker *et al.* [n. d.]. Exploiting long-distance interactions and tolerating atom loss in neutral atom quantum architectures. ISCA’21.
- [10] Michael J Biercuk, Hermann Uys, Aaron P VanDevender, Nobuyasu Shiga, Wayne M Itano, and John J Bollinger. 2009. Optimized dynamical decoupling in a model quantum memory. *Nature* 458, 7241 (2009), 996–1000.
- [11] D. Bluvstein *et al.* 2022. A quantum processor based on coherent transport of entangled atom arrays. *Nature* 604, 7906 (2022), 451–456.
- [12] S. Brandhofer *et al.* 2021. Optimal mapping for near-term quantum architectures based on Rydberg atoms. ICCAD’21.
- [13] Kenneth R Brown, Aram W Harrow, and Isaac L Chuang. 2004. Arbitrarily accurate composite pulse sequences. *Physical Review A* 70, 5 (2004), 052318.
- [14] Jinglei Cheng, Hanrui Wang, Zhiding Liang, Yiyu Shi, Song Han, and Xuehai Qian. 2022. TopGen: Topology-Aware Bottom-Up Generator for Variational Quantum Circuits. *arXiv preprint arXiv:2210.08190* (2022).
- [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*. 950–962.
- [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*. 937–949.
- [17] Yongshan Ding, Pranav Gokhale, Sophia Fuhui Lin, Richard Rines, Thomas Propson, and Frederic T Chong. 2020. Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 201–214.
- [18] S. Ebadi *et al.* 2022. Quantum optimization of maximum independent set using Rydberg atom arrays. *Science* 376, 6598 (2022), 1209–1215.
- [19] S. J. Evered *et al.* 2023. High-fidelity parallel entangling gates on a neutral atom quantum computer. *Nature* 622 (2023), 268–272.
- [20] H. Fan *et al.* 2022. Optimizing quantum circuit placement via machine learning. DAC’22.
- [21] P. Gokhale *et al.* 2021. Quantum fan-out: circuit optimizations and technology. QCE’21.
- [22] T. M. Graham *et al.* 2022. Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. *Nature* 604, 7906 (2022), 457–462.
- [23] Erwin L Hahn. 1950. Spin echoes. *Physical review* 80, 4 (1950), 580.
- [24] Ferdinand Helmer, Matteo Mariantoni, Austin G Fowler, Jan von Delft, Enrique Solano, and Florian Marquardt. 2009. Cavity grid for scalable quantum computation with superconducting circuits. *EPL (Europhysics Letters)* 85, 5 (2009), 50007.
- [25] P. Hoyer *et al.* 2005. Quantum fan-out is powerful. *Theory of Computing* 1, 5 (2005).
- [26] Fabian Kreppel, Christian Melzer, Janis Wagner, Janine Hilder, Ulrich Poschinger, Ferdinand Schmidt-Kaler, and André Brinkmann. 2022. Quantum Circuit Compiler for a Shuttling-Based Trapped-Ion Quantum Computer. (7 2022). [arXiv:2207.01964](https://arxiv.org/abs/2207.01964) [quant-ph]
- [27] 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*. 351–365.
- [28] G. Li *et al.* [n. d.]. Tackling the qubit mapping problem for NISQ-era quantum devices. ASPLOS’19.
- [29] G. Li *et al.* 2019. Tackling the qubit mapping problem for NISQ-era quantum. ASPLOS’19.
- [30] G. Li *et al.* 2022. Paulihedral: a generalized block-wise compiler optimization framework for quantum simulation kernels. ASPLOS’22.
- [31] Gushu Li, Yunong Shi, and Ali Javadi-Abhari. 2021. Software-hardware co-optimization for computational chemistry on superconducting quantum processors. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 832–845.
- [32] Y. Li *et al.* 2023. Timing-aware qubit mapping and gate scheduling adapted to neutral atom quantum computing. *TCAD* 42, 11 (2023).
- [33] Zhiding Liang, Jinglei Cheng, Hang Ren, Hanrui Wang, Fei Hua, Zhixin Song, Yongshan Ding, Fred Chong, Song Han, Yiyu Shi, and Xuehai Qian. 2023. PAN: Pulse Ansatz on NISQ Machines. [arXiv:2208.01215](https://arxiv.org/abs/2208.01215) [quant-ph]
- [34] Zhiding Liang, Jinglei Cheng, Zhixin Song, Hang Ren, Rui Yang, Hanrui Wang, Kecheng Liu, Peter Kogge, Tongyang Li, Yongshan Ding, *et al.* 2023. Towards Advantages of Parameterized Quantum Pulses. *arXiv preprint arXiv:2304.09253* (2023).
- [35] Zhiding Liang, Zhixin Song, Jinglei Cheng, Zichang He, Ji Liu, Hanrui Wang, Ruiyang Qin, Yiru Wang, Song Han, Xuehai Qian, and Yiyu Shi. 2023. Hybrid Gate-Pulse Model for Variational Quantum Algorithms. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1109/DAC56929.2023.10247923>
- [36] Zhiding Liang, Hanrui Wang, Jinglei Cheng, Yongshan Ding, Hang Ren, Xuehai Qian, Song Han, Weiwen Jiang, and Yiyu Shi. 2022. Variational quantum pulse learning. *arXiv preprint arXiv:2203.17267* (2022).
- [37] Daniel A Lidar. 2014. Review of decoherence free subspaces, noiseless subsystems, and dynamical decoupling. *Adv. Chem. Phys* 154 (2014), 295–354.
- [38] Sophia Fuhui Lin, Sara Sussman, Casey Duckering, Pranav S Mundada, Jonathan M Baker, Rohan S Kumar, Andrew A Houck, and Frederic T Chong. 2022. Let Each Quantum Bit Choose Its Basis Gates. *arXiv preprint arXiv:2208.13380* (2022).
- [39] Lei Liu and Jinglei 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>
- [40] Guang Hao Low, Theodore J Yoder, and Isaac L Chuang. 2014. Optimal arbitrarily accurate composite pulse sequences. *Physical Review A* 89, 2 (2014), 022341.
- [41] P. Magnard *et al.* 2020. Microwave quantum link between superconducting circuits housed in spatially separated cryogenic systems. *PRL* 125 (2020), 260502.
- [42] D. Maslov *et al.* 2008. Quantum Circuit Placement. *TCAD* 27, 4 (2008), 752–763.
- [43] J True Merrill and Kenneth R Brown. 2014. Progress in compensating pulse sequences for quantum computation. *Quantum Information and Computation for Chemistry* (2014), 241–294.
- [44] Abtin Molavi *et al.* [n. d.]. Qubit mapping and routing via MaxSAT. MICRO’22.

- [45] Abtin Molavi, Amanda Xu, Martin Diges, Lauren Pick, Swamit Tannu, and Aws Albarghouthi. 2022. Qubit Mapping and Routing via MaxSAT. *arXiv preprint arXiv:2208.13679* (2022).
- [46] P. Murali et al. 2019. Full-stack, real-system quantum computer studies: architectural comparisons and design insights. ISCA'19.
- [47] 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*. 1015–1029.
- [48] Prakash Murali, Dripto M. Debroy, Kenneth R. Brown, and Margaret Martonosi. 2020. Architecting Noisy Intermediate-Scale Trapped Ion Quantum Computers. In *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture (Virtual Event) (ISCA '20)*. IEEE Press, 529–542. <https://doi.org/10.1109/ISCA45697.2020.00051>
- [49] 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*. 1001–1016.
- [50] M. A. Norcia et al. [n. d.]. Iterative assembly of ^{171}Yb atom arrays in cavity-enhanced optical lattices. *arXiv:2401.16177*.
- [51] S. Park et al. 2022. A fast and scalable qubit-mapping method for noisy intermediate-scale quantum computers. DAC'22.
- [52] T. Patel et al. 2022. Geyser: a compilation framework for quantum computing with neutral atoms. ISCA '22.
- [53] 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>
- [54] Gokul Subramanian Ravi, Kaitlin N Smith, Pranav Gokhale, Andrea Mari, Nathan Earnest, Ali Javadi-Abhari, and Frederic T Chong. 2022. Vaqem: A variational approach to quantum error mitigation. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 288–303.
- [55] M Saffman. 2016. Quantum computing with atomic qubits and Rydberg interactions: progress and challenges. *Journal of Physics B* 49, 20 (oct 2016), 202001.
- [56] M. Y. Siraichi et al. 2018. Qubit allocation. CGO'18.
- [57] Kaitlin N. Smith, Gokul Subramanian Ravi, Prakash Murali, Jonathan M. Baker, Nathan Earnest, Ali Javadi-Abhari, and Frederic T. Chong. 2021. Error Mitigation in Quantum Computers through Instruction Scheduling. <https://doi.org/10.48550/ARXIV.2105.01760>
- [58] 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*. 59–71.
- [59] B. Tan et al. 2020. Optimal layout synthesis for quantum computing. ICCAD'20.
- [60] B. Tan et al. 2021. Optimal qubit mapping with simultaneous gate absorption. ICCAD'21.
- [61] B. Tan et al. 2022. Qubit mapping for reconfigurable atom arrays. (2022). ICCAD'22.
- [62] D. B. Tan et al. 2024. Compiling quantum circuits for dynamically field-programmable neutral atoms array processors. *Quantum* 8 (2024).
- [63] Swamit Tannu, Poulami Das, Ramin Ayanzadeh, and Moinuddin Qureshi. 2022. HAMMER: Boosting Fidelity of Noisy Quantum Circuits by Exploiting Hamming Behavior of Erroneous Outcomes. 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, 529–540. <https://doi.org/10.1145/3503222.3507703>
- [64] 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*. 987–999.
- [65] Richard Versluis, Stefano Poletto, Nader Khammassi, Brian Tarasinski, Nadia Haider, David J Michalak, Alessandro Bruno, Koen Bertels, and Leonardo DiCarlo. 2017. Scalable quantum circuit and control for a superconducting surface code. *Physical Review Applied* 8, 3 (2017), 034021.
- [66] Lorenza Viola, Emanuel Knill, and Seth Lloyd. 1999. Dynamical decoupling of open quantum systems. *Physical Review Letters* 82, 12 (1999), 2417.
- [67] Joel J Wallman and Joseph Emerson. 2016. Noise tailoring for scalable quantum computation via randomized compiling. *Physical Review A* 94, 5 (2016), 052325.
- [68] Hanrui Wang, Yongshan Ding, Jiaqi Gu, Yujun Lin, David Z Pan, Frederic T Chong, and Song Han. 2022. Quantumas: Noise-adaptive search for robust quantum circuits. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 692–708.
- [69] Hanrui Wang, Jiaqi Gu, Yongshan Ding, Zirui Li, Frederic T Chong, David Z Pan, and Song Han. 2022. QuantumNAT: quantum noise-aware training with noise injection, quantization and normalization. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*. 1–6.
- [70] Hanrui Wang, Zirui Li, Jiaqi Gu, Yongshan Ding, David Z Pan, and Song Han. 2022. QOC: quantum on-chip training with parameter shift and gradient pruning. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*. 655–660.
- [71] Hanrui Wang, Pengyu Liu, Jinglei Cheng, Zhiding Liang, Jiaqi Gu, Zirui Li, Yongshan Ding, Weiwen Jiang, Yiyu Shi, Xuehai Qian, et al. 2022. Quest: Graph transformer for quantum circuit reliability estimation. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*.
- [72] Hanrui Wang, Pengyu Liu, Yilian Liu, Jiaqi Gu, Jonathan Baker, Frederic T. Chong, and Song Han. 2023. DGR: Tackling Drifted and Correlated Noise in Quantum Error Correction via Decoding Graph Re-weighting. In *arXiv preprint*.
- [73] Hanrui Wang, Pengyu Liu, Kevin Shao, Dantong Li, Jiaqi Gu, David Z Pan, Yongshan Ding, and Song Han. 2023. TransformerQEC: Transferable Transformer for Quantum Error Correction Code Decoding. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD), FastML for Science Workshop*.
- [74] Hanrui Wang, Pengyu Liu, Bochen Tan, Yilian Liu, Jiaqi Gu, David Z. Pan, Jason Cong, Umut Acar, and Song Han. 2023. FPQA-C : A Compilation Framework for Field Programmable Qubit Array. In *arXiv preprint*.
- [75] Hanrui Wang, Yilian Liu, Pengyu Liu, Jiaqi Gu, Zirui Li, Zhiding Liang, Jinglei Cheng, Yongshan Ding, Xuehai Qian, Yiyu Shi, David Z. Pan, Frederic T. Chong, and Song Han. 2023. RobustState: Boosting Fidelity of Quantum State Preparation via Noise-Aware Variational Training. *arXiv preprint* (2023).
- [76] R. Wille et al. 2019. Mapping quantum circuits to IBM QX architectures using the minimal number of SWAP and H operations. DAC'19.
- [77] T.-A. Wu et al. 2022. A robust quantum layout synthesis algorithm with a qubit mapping checker. ICCAD'22.
- [78] Xin-Chuan Wu, Dripto M Debroy, Yongshan Ding, Jonathan M Baker, Yuri Alexeev, Kenneth R Brown, and Frederic T Chong. 2020. TILT: Achieving Higher Fidelity on a Trapped-Ion Linear-Tape Quantum Computing Architecture. *arXiv preprint arXiv:2010.15876* (2020).
- [79] Lei Xie, Jidong Zhai, ZhenXing Zhang, Jonathan Allcock, Shengyu Zhang, and Yi-Cong Zheng. 2022. Suppressing ZZ crosstalk of Quantum computers through pulse and scheduling co-optimization. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 499–513.
- [80] Bichen Zhang, Swarnadeep Majumder, Pak Hong Leung, Stephen Crain, Ye Wang, Chao Fang, Dripto M Debroy, Jungsang Kim, and Kenneth R Brown. 2021. Hidden Inverses: Coherent Error Cancellation at the Circuit Level. *arXiv preprint arXiv:2104.01119* (2021).
- [81] C. Zhang et al. 2021. Time-optimal qubit mapping. ASPLOS'21.
- [82] C. Zhang et al. 2021. Time-optimal qubit mapping. ASPLOS'21.
- [83] Junyao Zhang, Hanrui Wang, Gokul Subramanian Ravi, Frederic T Chong, Song Han, Frank Mueller, and Yiran Chen. 2023. DISQ: Dynamic Iteration Skipping for Variational Quantum Algorithms. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*.
- [84] Han Zheng, Gokul Subramanian Ravi, Hanrui Wang, Kanav Setia, Frederic T Chong, and Junyu Liu. 2023. SnCQA: A hardware-efficient equivariant quantum convolutional circuit architecture. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*.
- [85] X. Zhou et al. 2020. A Monte Carlo tree search framework for quantum circuit transformation. ICCAD'20.
- [86] A. Zulehner et al. 2018. Efficient mapping of quantum circuits to the IBM QX architectures. DATE'18.