

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

DasAtom: A Divide-and-Shuttle Atom Approach to Quantum Circuit Transformation

Yunqi Huang, Dingchao Gao, Shenggang Ying, and Sanjiang Li*

Abstract—Neutral atom (NA) quantum systems are emerging as a leading platform for quantum computation, offering superior or competitive qubit count and gate fidelity compared to superconducting circuits and ion traps. However, the unique features of NA devices, such as long-range interactions, long qubit coherence time, and the ability to physically move qubits, present distinct challenges for quantum circuit compilation. In this paper, we introduce DasAtom, a novel divide-and-shuttle atom approach designed to optimise quantum circuit transformation for NA devices by leveraging these capabilities. DasAtom partitions circuits into subcircuits, each associated with a qubit mapping that allows all gates within the subcircuit to be directly executed. The algorithm then shuttles atoms to transition seamlessly from one mapping to the next, enhancing both execution efficiency and overall fidelity. For a 30-qubit Quantum Fourier Transform (QFT), DasAtom achieves a 415.8x improvement in fidelity over the move-based algorithm Enola and a 10.6x improvement over the SWAP-based algorithm Tetris. Notably, this improvement is expected to increase exponentially with the number of qubits, positioning DasAtom as a highly promising solution for scaling quantum computation on NA platforms.

Keywords: neutral atom quantum computing, quantum circuit transformation, divide-and-conquer (DAC), subgraph isomorphism

I. INTRODUCTION

QUANTUM computing has the potential to revolutionise various fields, including cryptography [1], chemistry [2], and machine learning [3]. Among the diverse quantum hardware platforms under development, neutral atom (NA) quantum systems have garnered significant attention due to their inherent advantages in scalability, qubit connectivity, and gate fidelity [4]–[6]. Unlike other quantum platforms such as superconducting circuits, NA systems can leverage long-range interactions and native multi-qubit gates, which allow for the execution of complex quantum operations with fewer resources. Additionally, the ability to physically move qubits within NA systems introduces a new dimension in quantum circuit optimisation that is not present in more rigid architectures [5], [7].

Work partially supported by the Australian Research Council (DP220102059), the National Science Foundation of China (12071271, 12471437), the Beijing Nova Program (20220484128), the Innovation Program for Quantum Science and Technology (2021ZD0302901) and the Young Scientists Fund of the Natural Science Foundation of Jiangsu Province (BK20240536)

Yunqi Huang and Sanjiang Li are with Centre for Quantum Software and Information (QSI), Faculty of Engineering and Information Technology, University of Technology Sydney, NSW 2007, Australia. Dingchao Gao and Shenggang Ying are with Key Laboratory of System Software (Chinese Academy of Sciences) and State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences.

Corresponding author (E-mail: Sanjiang.Li@uts.edu.au)

Despite these advantages, the unique characteristics of NA devices present novel challenges in quantum circuit compilation. Specifically, the need to efficiently map qubits and execute operations while improving overall fidelity is a non-trivial problem. Existing quantum circuit compilation methods, such as SWAP-based [8]–[10] and move-based [11]–[15] algorithms, are not fully optimised for the capabilities of NA systems, often leading to suboptimal performance in terms of overall fidelity.

The SWAP-based method, exemplified by Tetris [10], employs a fixed-atom array and transforms quantum circuits in a manner similar to compilers for superconducting devices [16]–[20]. This approach leverages the long-range interactions of NA systems to achieve denser qubit connectivity, which contrasts with state-of-the-art IBM superconducting quantum computers that typically have sparse qubit connections, with an average degree between 2 and 3. Tetris operates as a heuristic greedy algorithm, addressing both qubit connection constraints and parallel execution constraints.

On the other hand, Enola [12], the most recent and advanced move-based compiler, assumes sufficient separation between atoms to eliminate parallel execution constraints. As [11]–[14], Enola exploits gate commutability in QAOA [21] circuits, aiming to scheduling gates from a commutation group in near-optimal number of Rydberg stages. For generic quantum circuits like Quantum Fourier Transform (QFT), CZ gates are often blocked by single-qubit gates (cf. Fig. 4). Enola executes each layer of CZ gates *collectively*. Empirical results on QFT circuits indicate that the slow atom movement (partially due to the large atom distance) is the primary factor contributing to overall fidelity loss.

However, even these advanced methods do not fully capitalise on the strengths of the NA platform, particularly the synergistic use of long-range interactions and atom shuttling, which can significantly enhance the efficiency and fidelity of quantum circuit execution. This gap highlights the need for a more comprehensive approach, leading to the development of DasAtom—a divide-and-shuttle atom algorithm specifically designed for NA systems. DasAtom combines the advantages of both Tetris and Enola while avoiding their shortcomings. Tetris effectively uses long-range interactions to enable dense qubit connectivity, but it does not utilise the ability to move atoms, limiting its flexibility. On the other hand, Enola leverages atom shuttling to adapt qubit mappings dynamically, yet it cannot take advantage of long-range interactions as atoms are far apart. DasAtom integrates both of these key

capabilities: it partitions circuits into subcircuits, assigns an optimal qubit mapping for each subcircuit, and then shuttles atoms to smoothly transition between mappings. By doing so, it ensures that every gate in a subcircuit is directly executable, thereby enhancing overall fidelity and efficiency.

We conducted empirical comparisons of DasAtom’s performance against Tetris, Enola, and Atomique using the benchmark circuits used in [10]. This set of 33 benchmark circuits includes RevLib circuits as well as the quantum circuits Bernstein–Vazirani (BV), Quantum Volume (QV), and QFT, which range from 5 to 16 qubits with up to 3,089 CZ gates. In addition, we evaluated Deutsch–Jozsa (DJ), 3-regular MaxCut QAOA, Greenberger–Horne–Zeilinger (GHZ), QFT, QV, two-local ansatz, and W-state circuits with qubit counts ranging from 5 to 50. Our experiments demonstrate that DasAtom consistently delivers significant performance gains in both overall fidelity and compiler runtime. For instance, DasAtom achieves a 415.8x, 10.6x, 16.1x improvement in fidelity over, respectively, Enola, Tetris, and Atomique on a 30-qubit QFT, while the runtimes of Enola, Tetris, and Atomique are, respectively, 9,851x, 384x, 2.6x that of DasAtom. Moreover, these improvements are expected to scale exponentially with the number of qubits, making DasAtom a highly promising solution for quantum computation on NA platforms.

The remainder of this paper is organised as follows: Section II provides a brief background and discusses related work. Sections III and IV offer a detailed description of the DasAtom algorithm, including its implementation and a comprehensive performance analysis across various quantum circuits. In Section V, we explain why DasAtom outperforms other algorithms and explore potential avenues for optimisation. The last section concludes the paper.

II. BACKGROUND AND RELATED WORK

A. Neutral atom quantum hardware

In neutral atom quantum hardware, neutral atoms are trapped in arrays of tweezers and the computational states $|0\rangle$ and $|1\rangle$ are encoded in the hyperfine ground states of an alkali or alkaline-earth-like atom [4], [22], [23]. These atoms can be arranged in one, two, or even three-dimensional configurations [24]–[27]. In this work, we focus on a $b \times b$ regular grid $G(b, b)$ with constant distance $d > 0$. Fig. 1 shows a 3×3 grid $G(3, 3)$. In the following, we write q, q', q_i for program qubits in a circuit, and write p, p', p_i for nodes or their coordinates in $G(b, b)$.

Single-qubit gates are implemented through individual or global optical addressing of the atoms, and two-qubit gates are realised by exciting the atoms into a Rydberg state using laser beams. The excitation to Rydberg states induces a strong dipole-dipole interaction between the atoms [28]. This interaction is governed by an interaction radius R_{int} , within which a CZ gate on two atoms p_i, p_j can be performed if the distance $D(p_i, p_j) \leq R_{\text{int}}$, where D represents the Euclidean distance.

One significant advantage of NA platform is its ability for long-range interaction. Two qubits can interact even if they




Fig. 1: The architecture graph of a neutral atom quantum hardware, where (x, y) ($0 \leq x, y \leq 2$) denotes the location of an atom, and two atoms are connected if their distance is smaller than $R_{\text{int}} = r_{\text{int}} \times d$, where $r_{\text{int}} \in \{1, \sqrt{2}, 2\}$.

are not neighbours in the grid. For instance, the interaction radius R_{int} can be $r_{\text{int}} \times d$ for some $r_{\text{int}} \geq 1$. Qubit connectivity of a quantum architecture is often captured by its architecture graph $AG = (V, E)$, where the nodes in V correspond to the physical qubits (i.e., trapped atoms here), while the edges indicate the qubits capable of interacting with each other. For NA platform, each physical qubit p in V is assigned to the trap coordinates (x, y) , and the edges E can be defined as:

$$E = \{(p_i, p_j) | p_i, p_j \in V, D(p_i, p_j) \leq R_{\text{int}}\}. \quad (1)$$

In the 3×3 grid shown in Fig. 1, the black, red dotted, and blue edges are, respectively, those with distances $d, \sqrt{2}d, 2d$.

To minimise crosstalk between gates, *parallel gate execution* is feasible only if qubits corresponding to different CZ gates maintain a distance of at least the restriction radius $R_{\text{restr}} \geq R_{\text{int}}$ from all qubits involved in other simultaneously executed two-qubit gates. Specifically, for two CZ gates g on p_i, p_j and g' on p_a, p_b to be executed in parallel, the conditions

$$D(p_u, p_v) > R_{\text{restr}} \quad (2)$$

must be satisfied for any $u \in \{i, j\}$ and any $v \in \{a, b\}$. For example, suppose $R_{\text{int}} = R_{\text{restr}} = d$ in Fig. 1. Let $p_{i,j}$ be a physical qubit located at (i, j) . Then $\text{CZ}(p_{0,2}, p_{1,2})$ can be executed in parallel with $\text{CZ}(p_{0,0}, p_{1,0})$, but not in parallel with $\text{CZ}(p_{1,1}, p_{1,0})$.

The ability to move qubits is the most distinctive advantage of NA platforms [5], [7], [23]. In NA systems, qubits are captured in two types of traps (cf. Fig. 2). A spatial light modulator (SLM) generates an array of static traps, while a 2D acousto-optic deflector (AOD) creates mobile traps that can move within the plane. The AOD traps are formed at the intersection of a set of rows and columns. Atom movement is a high-fidelity operation and an atom can traverse a region for 2,000 qubits with only 0.1% coherence time [5]. Fig. 2 illustrates the procedure of atom movement in neutral atom quantum computing (NAQC): At each movement stage, a $k \times k$ AOD array is activated for some integer $k > 0$, atoms are loaded from the SLM array to the AOD array, and the AOD array is moved, subject to the constraint that different columns (rows) must not cross each other. After the movement, the

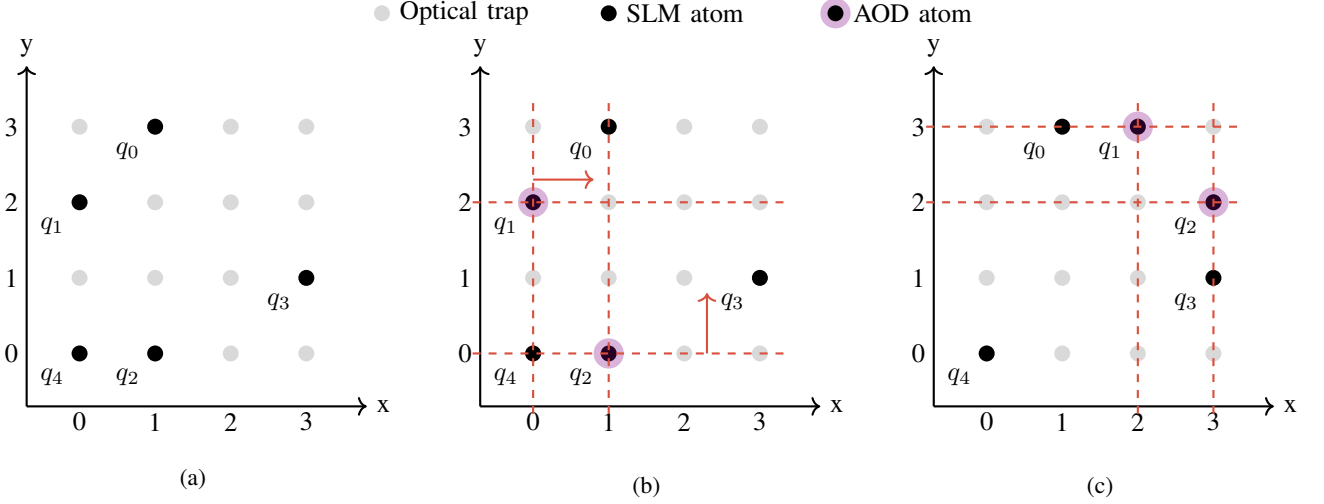


Fig. 2: Illustration of atom movement in NAQC: (a) program qubits are mapped to a 4×4 SLM array, (b) a 2×2 AOD array is activated, and two SLM atoms, carrying q_1, q_2 , are loaded (i.e., transferred) to the AOD array, (c) AOD rows (columns) are moved to the right (upward), carrying q_1 and q_2 to positions (2,3) and (3,2). To apply a CZ gate between q_0 and q_1 , we offload q_1 to the SLM array and then apply an individual Rydberg laser to interact q_0 and q_1 . If $R_{\text{restr}} = R_{\text{int}} = d$, then $CZ(q_0, q_1)$ and $CZ(q_2, q_3)$ can be applied in parallel. However, this is not possible when $R_{\text{restr}} = \sqrt{2}d$ and $R_{\text{int}} = d$.




Fig. 3: The QFT-5 circuit, where H is the Hadamard gate, $P(\theta)$ is a phase gate with phase $\exp(i\theta)$.




Fig. 4: The QFT-5 circuit decomposed in gate set $\{CZ, R_x, R_y, R_z\}$, where each R denotes an R_x, R_y , or R_z gate.

AOD atoms are offloaded to the SLM array and the AOD array is deactivated. During each movement, the relative positions of AOD atoms are preserved, but the relative positions of SLM atoms are changed after the movement.

B. Quantum circuit compilation

In quantum computing, a quantum circuit serves as a model that represents the flow of information and operations within a quantum algorithm. A quantum circuit is composed of qubits and quantum gates that manipulate the states of these qubits. In this work, we denote a quantum circuit by C , which consists of a sequence of quantum gates $\{g_1, \dots, g_m\}$ acting on qubits in $Q = \{q_1, \dots, q_n\}$. Fig. 3 shows the well-known QFT circuit on five qubits. Note that, for clarity, we have omitted the SWAP gates at the end of the circuit.

Circuits like QFT-5 in Fig. 3 often contain gates that are not native in a target quantum device. This means we need to decompose non-native gates in C into native gates. For

example, let $\mathcal{G} = \{R_x, R_y, R_z, CZ\}$ be the set of native gates of the target quantum device. For QFT-5 in Fig. 3, we need to decompose H and control phase gate $CP(\theta)$ in gates in \mathcal{G} . The result is shown in Fig. 4.

Due to limited qubit connectivity, it is often not possible to execute the synthesised circuit directly, as some CZ gates may act on two far away device qubits. For example, Fig. 2a illustrates a 4×4 SLM array and an initial mapping f that maps program qubits q_i ($0 \leq i \leq 4$) within $G(4, 4)$. Assuming $R_{\text{int}} = 2d$, a CZ gate $CZ(q_i, q_j)$ can be executed if the distance $D(f(q_i), f(q_j)) \leq 2d$. Clearly, $CZ(q_1, q_2)$ cannot be executed directly in this case, as $D(f(q_1), f(q_2)) > 2d$.

We next outline the general procedures for quantum circuit transformation. To execute a decomposed circuit C on an NA device, each program qubit q in C is first mapped to a physical atom located at $p_{i,j}$ in the device's SLM array. With this initial mapping f , not all CZ gates are directly executable. If a CZ




Fig. 5: The decomposed QFT-5 circuit with all single-qubit gates removed. The CZ circuit is partitioned in 12 CZ layers and divided into two subcircuits with the second consisting of the last four CZ gates.

gate g cannot be executed directly, the two qubits involved must be moved close together. In NA platforms, this can be achieved in two ways: by inserting SWAP gates or by atom shuttling (cf. [23]). For instance, in the mapping shown in Fig. 2a, executing $CZ(q_1, q_2)$ requires either swapping q_1 with q_4 or q_2 with q_4 using SWAP gates. Alternatively, we can move the atom carrying q_1 from $p_{0,2} = f(q_1)$ to one of the three neighbouring positions around $p_{1,0} = f(q_2)$. If a neighbouring position is occupied, the occupying atom must be moved first.

C. Related work

Quantum circuit transformation (QCT) is the process of converting a program circuit into a form that is executable on a target quantum device, whether it be a superconducting or NA quantum device. QCT is a crucial component of quantum circuit compilation. Since IBM launched its cloud-based quantum computing services, numerous algorithms have been proposed in the literature, including [16]–[20], [29]–[31].

The unique features of NA devices introduce distinct challenges for quantum circuit compilation. Early efforts sought to leverage these features within fixed atom arrays. Baker et al. proposed the first compiler for NA devices, which accounted for long-range interaction [8]. Multi-qubit gate support was subsequently incorporated in [9]. Inspired by the popular block puzzle game, Li et al. [10] proposed the heuristic algorithm *Tetris*, which effectively reduces qubit idle time while exploiting the rich qubit connections in NA device and adhering to parallel execution constraints.

Unlike the above SWAP-gate based shuttling approach, Tan et al. [11], [12] developed OLSQ-DPQA, an SMT solver-based compiler for the dynamically field-programmable qubit array (DPQA) architecture [5], which employs atom shuttling. To address the scalability challenge of OLSQ-DPQA, several follow-up compilers have been introduced, including Atomique [13], Q-Pilot [14], and Enola [15], all designed for DPQA. Enola’s compilation process is divided into scheduling, placement, and routing. It schedules a commutation group of CZ gates in a near-optimal number of Rydberg stages and, for generic quantum circuits, schedules each layer of parallel CZ gates as a Rydberg stage. Enola offers two placement strategies: dynamic placement, which generates a new qubit mapping for each Rydberg stage, and static placement, which retains the same mapping throughout. Compared to OLSQ-DPQA, Atomique, and Q-Pilot, Enola demonstrates superior or comparable fidelity improvements [15].

Combining the two methods described above is natural. Brandhofer et al. [32] proposed a compiler that integrates SWAP gates with a special atom movement technique called one-

dimensional displacements, aimed at reducing circuit depth and improve circuit fidelity. Another hybrid compiler was proposed in [33], where the NA hardware is grid-based, and atom can be moved from one grid point to another only if the target grid point is unoccupied. Experiments in [33] show that, for shuttling-favoured NA devices, the hybrid approach is essentially the same as the atom shuttling method in performance.

Interested readers may consult the recent review paper [23] for more information.

Divide-and-conquer (DAC) is a widely used approach in quantum circuit transformation that several algorithms employ to enhance efficiency. Siraichi et al. [34] introduced the BMT algorithm, which partitions a circuit’s gate list into sublists, aligns them through subgraph isomorphism and token swapping, and determines an optimal transformation path with dynamic programming. Similarly, Wu et al. [35] developed an SMT-based DAC method that partitions circuits and links sublist embeddings with SWAP gates to minimise transformation costs. Both methods, however, rely on exhaustive search, limiting scalability to circuits with 20 or more qubits. Recently, Huang et al. [36] proposed the adaptive DAC algorithm, which dynamically partitions circuits and optimises routing heuristics, achieving substantial performance gains, including nearly 50% improvement on the IBM Tokyo architecture.

III. DASATOM TRANSFORMATION FRAMEWORK

Before describing our algorithm, we recall that qubit interaction in a quantum circuit can be represented by a graph.

Definition 1 (Interaction Graph). For a quantum circuit C with qubit set Q , its interaction graph, denoted as $IG(C)$, is an undirected graph (Q, E_{int}) , where:

- each node represents a qubit of C .
- Two nodes $q, q' \in Q$ are connected if there is a two-qubit gate in C acting on q, q' .

If the interaction graph of C matches well with the architecture graph $AG = (V, E)$, there is no need to route the qubits: what we need is an embedding f from $IG(C)$ to AG . Here a 1-1 mapping $f : Q \rightarrow V$ is an embedding if $(f(q), f(q'))$ is an edge in AG for any edge (q, q') in $IG(C)$. In this case, f is also called a subgraph isomorphism.

A. Overview

Let \mathcal{G} be an NA-native gate set that includes the two-qubit CZ gate and local addressed rotation gates R_z , as well as (locally or globally addressed) R_x and R_y . Given an input circuit C on n qubits, we first synthesise C using gates from \mathcal{G} (cf. [37]).




Fig. 6: The flowchart of DasAtom

We assume that single-qubit gates and CZ gates are executed in separate stages, meaning that in each stage, either single-qubit gates or CZ gates are executed, but not both. Whenever there is a single-qubit gate in the front layer, we execute (and remove) it to form a new front layer. This ensures that, at some point, we obtain a front layer consisting only of CZ gates, which can be executed collectively, provided the parallel gate execution constraints are satisfied.

In the following, we assume single-qubit gates are removed before calling DasAtom, a convention also adopted in OLSQ-DPQA, Atomique, Q-Pilot, and Enola. For simplicity, we continue to refer to the resulting CZ-only circuit as C .

Let d denote the atom distance, $R_{\text{int}} \geq d$ the interaction radius, and $R_{\text{restr}} \geq R_{\text{int}}$ the restriction radius. We assume that CZ gates are executed individually using Rydberg lasers. Our algorithm, DasAtom, proceeds as follows (cf. Fig. 6 for the flowchart):

- 1) Partition the CZ circuit C into layers L_1, \dots, L_m , where each L_i contains only CZ gates.
- 2) Divide the CZ circuit C into subcircuits such that (i) each layer L_i is fully contained within a subcircuit; (ii) the interaction graph of each subcircuit is embeddable in $\text{Grid}(b, b)$, where $b = \lceil \sqrt{n} \rceil$.
- 3) For each subcircuit C_i , construct a mapping f_i that embeds the interaction graph $IG(C_i)$ to $\text{Grid}(b, b)$. For each CZ gate $g = CZ(q, q')$ in C_i , we execute g by interacting $f_i(q)$ and $f_i(q')$. Since $D(f_i(q), f_i(q')) \leq R_{\text{int}}$, the gate can be executed directly. However, due to parallel execution constraints, not all gates in the same layer can be executed simultaneously.
- 4) Perform routing based on atom shuttling (instead of inserting SWAP gates). This can be achieved by modifying the routing algorithm provided in [12], [15] or [33]. This process transforms f_i into f_{i+1} .

In the following, we discuss the implementation of the above procedures in detail.

B. Circuit division

After decomposing gates in circuit C into NA-native gates and removing single-qubit gates, we partition C into CZ layers L_1, \dots, L_m . We further divide C into subcircuits C_1, \dots, C_k , where each C_i consists of consecutive CZ layers from C , and the interaction graph of C_i can be embedded in AG with a mapping f_i . We denote $C_i = C[\ell_i : m_i]$ if C_i is composed of layers $L_{\ell_i}, \dots, L_{m_i}$. Thus, we have $1 = \ell_1 \leq m_1 < \ell_2 \leq m_2 < \dots < \ell_k \leq m_k = m$.

Unlike the gate-by-gate partitioning methods used in [34] and [35], our division scheme is coarse-grained. This design has two implications: first, it significantly reduces the number of costly subgraph isomorphism checks required by our algorithm; second, it presents opportunities for further optimisation. For circuit division, we employ the Rustworkx implementation of VF2 [38]. When partitioning the circuit into layers, the number of VF2 calls is at most $O(\text{depth})$, where depth is the number of CZ layers in C . In contrast, when using a gate-by-gate partition, the number of VF2 calls grows exponentially in depth .

Algorithm 1 Circuit division

Require: A CZ circuit C and an architecture graph AG
Ensure: Subcircuits SC and embeddings EM

- 1: Initialize two empty sequences SC and EM
- 2: $i_{\text{last}} \leftarrow 0$
- 3: **for** $0 < i \leq \text{depth of } C$ **do**
- 4: **if** $IG(C[i_{\text{last}} : i])$ is embeddable in AG **then**
- 5: continue
- 6: **else**
- 7: Append $C[i_{\text{last}} : i - 1]$ to SC
- 8: Find an embedding f_i for $IG(C[i_{\text{last}} : i - 1])$ and append it to EM
- 9: $i_{\text{last}} \leftarrow i$
- 10: **end if**
- 11: **end for**
- 12: return SC and EM

For our running example, if we set $R_{\text{int}} = \sqrt{2}d$, then the QFT-5 circuit can be partitioned into two subcircuits, see Fig. 5. The interaction graphs and their embeddings to $AG = G(3, 3)$ of these two subcircuits are shown in Fig. 7. To minimise qubit idle time, it is desirable to execute as many CZ gates in parallel as possible, ensuring they satisfy the previously mentioned restriction constraint. In our example, $R_{\text{restr}} = 2 \times R_{\text{int}} = 2\sqrt{2}d$. The first sub-circuit includes, for instance, a CZ layer consisting of gates $CZ(q_1, q_3)$ and $CZ(q_0, q_4)$, with the qubit mapping f shown in Fig. 7(a). Since $D(f(q_3), f(q_4)) = d < R_{\text{restr}}$, $CZ(q_1, q_3)$ and $CZ(q_0, q_4)$ cannot be executed in parallel due to the restriction constraint.

C. Atom shuttling

Suppose we have divided the CZ circuit C into subcircuits C_1, \dots, C_m with corresponding embeddings f_1, \dots, f_m . How




Fig. 7: The embeddings in $G(3,3)$ and interaction graphs of the two subcircuits of QFT-5 (shown in Fig. 5).

do we transition from one mapping to the next? In superconducting quantum devices, this is typically done by inserting SWAP gates. In this section, we show how this can be completed in NAQC by atom shuttling. The routing method described below is adapted from [15] for our purpose.

Let $Q = \{q_0, \dots, q_{n-1}\}$ be the set of program qubits in the input circuit. Suppose f, f' are two 1-1 mappings from Q to grid points in $G(b, b)$ (where $b = \lceil \sqrt{n} \rceil$). Our goal is to transition from f to f' . To do this, we need to move each $q_i \in Q$ from its current position $f(q_i) = (x_i, y_i)$ in $G(b, b)$ to a new position $f'(q_i) = (x'_i, y'_i)$, also within $G(b, b)$.

Formally, we represent each move from $f(q_i)$ to $f'(q_i)$ as a 4-tuple $m_i \triangleq (x_i, y_i, x'_i, y'_i)$. We consider a move m_i to be trivial if $x'_i = x_i$ and $y'_i = y_i$, meaning the qubit remains stationary. Define $M = \{m_i \mid m_i \text{ is non-trivial}, 0 \leq i \leq n-1\}$. Moves in M can be implemented *serially* (cf. Fig. 2): for each move m_i , we activate a 1×1 AOD at the starting point (x_i, y_i) , load q_i , move the AOD to (x'_i, y'_i) , and then offload q_i at its new position. If the target position is occupied by another qubit q_j (for some $j > i$), we temporarily move q_j aside, place q_i at the target, and delay movement m_j until later.

While this approach accomplishes the transition, it is not ideal since atom movement is slow and can reduce overall fidelity. In many cases, moves can be executed in parallel. Fig. 2 actually illustrates how two moves, $m_1 = (0, 2, 2, 3)$ and $m_2 = (1, 0, 3, 2)$, can be implemented simultaneously. However, not all moves are eligible for parallel execution; legal parallel moves must satisfy the constraint that different AOD columns (or rows) do not cross or overlap.

To formalise the concept of legal parallel moves, we introduce the notion of *conflicting moves*. Two moves $m_i \triangleq (x_i, y_i, x'_i, y'_i)$ and $m_j \triangleq (x_j, y_j, x'_j, y'_j)$ in M are considered in *conflict* if they fail to meet the following condition for some $*$ in $\{<, =, >\}$:

- $(x_i * x_j) \iff (x'_i * x'_j)$ and $(y_i * y_j) \iff (y'_i * y'_j)$.

If m_i and m_j do not conflict, they are deemed *compatible*. A set of moves $M' = \{m_{i_1}, \dots, m_{i_\ell}\}$ is compatible if every pair within the set is compatible. In this case, the moves in M' can be implemented in parallel by activating a $k \times k$ AOD array for some $k > 0$, positioning all qubits corresponding to the moves in M' at the grid points, loading them from the SLM

to the AOD, moving them in parallel to their destinations, and offloading them. Our objective is to find a sequence of parallel, compatible moves that transports each (x_i, y_i) to (x'_i, y'_i) efficiently, minimising the total movement time.

Example 1. Fig. 7 shows two mappings f and f' in $G(3,3)$ from $Q = \{q_1, \dots, q_5\}$. In this instance, we have $M = \{m_0 \triangleq (0, 2, 0, 1), m_1 \triangleq (1, 0, 0, 2), m_4 \triangleq (0, 1, 1, 0)\}$. The three moves m_0, m_1, m_4 in M are pairwise conflicting. For example, $m_0 = (0, 2, 0, 1)$ and $m_1 = (1, 0, 0, 2)$ are conflicting as $x_0 = 0 < 1 = x_1$ but $x'_0 = 0 = 0 = x'_1$, violating the constraint $(x_i * x_j) \iff (x'_i * x'_j)$ for $*$ being $<$ or $=$. This means that we should implement them sequentially. Fig. 8a shows such a legal implementation, where m_1 is implemented first, followed by m_0 , and lastly, m_4 .

To address the general problem, we construct a conflict graph for M following the approach in [15]. In this graph, each node represents a move in M , and an edge connects two nodes if their moves are incompatible. Fig. 8b illustrates the conflict graph for the movement set in Example 1. Any set of compatible moves forms an independent set of vertices, reducing the problem to finding a maximum independent set (MIS) in the conflict graph. Given that the MIS problem is NP-hard, an approximate solution can be obtained by searching for a maximal independent set [15]. This approach allows us to identify a sequence of compatible moves that transition from f to f' in a reasonably short total movement time, as the longest move in each compatible set determines the AOD movement time for the set [15].

Remark 1. Several existing routing algorithms, such as those in [15], [33], are suitable for our purposes. To facilitate comparison with Enola, we adapt the routing algorithm from [15] to meet our requirements, with only a minor adjustment. In Enola, the routing task is slightly different: at each Rydberg stage, a set S of CZ gates is selected, where no two gates share a common qubit. The objective is to shuttle atoms so that the two qubits p_i, p_j involved in each CZ gate in S are aligned at the same grid point.¹ Given a current mapping f , this alignment only requires moving $f(p_i) \triangleq (x_i, y_i)$ to $f(p_j) \triangleq (x_j, y_j)$, or vice versa. Thus, any CZ gate $CZ(p_i, p_j)$ in S involves two dual moves: $m = (x_i, y_i, x_j, y_j)$ and $m' = (x_j, y_j, x_i, y_i)$. By contrast, DasAtom does not include

¹In DPQA, each grid point in the SLM array can hold two traps.

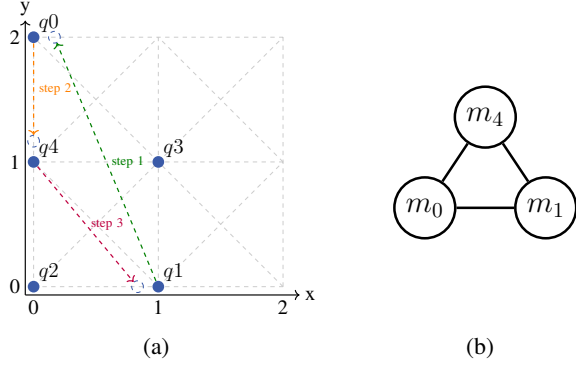


Fig. 8: (a) Atom movement between the two mappings in Fig. 7, (b) the conflict graph for $M = \{m_0, m_1, m_4\}$, where $m_0 = (0, 2, 0, 1)$, $m_1 = (1, 0, 0, 2)$, and $m_4 = (0, 1, 1, 0)$.

dual movements within its conflict graph.

IV. EVALUATION AND ANALYSIS

We implemented our proposed algorithm in Python. All experiments for DasAtom and Enola were conducted on a MacBook Pro featuring a 2.3 GHz Intel Core i5 processor and 16 GB memory, while all experiments for Tetris and Atomique were conducted on an Ubuntu 20.04 server with 40 cores of Intel Xeon Gold 5215 @ 2.50GHz, 512 GB of RAM. DasAtom is open source at <https://github.com/Huangyunqi/DasAtom>.

In our evaluation, we assume that the input circuit is decomposed in the same manner, with all single-qubit gates removed before applying the algorithms.

For readers' convenience, we summarise all key parameters, symbols, and their default values (if applicable) in Table I.

TABLE I: Key parameters, symbols, and default values

Symb.	Meaning	default value	Symb.	Meaning
T_2	coherence time	1.5s	n	qubit count
f_{cz}	CZ gate fidelity	0.995	m	CZ gate count
f_{trans}	atom transfer fidelity	1	s	atom transfer count
t_{cz}	CZ gate duration	0.2 μ s	h	compiled circuit depth
t_{trans}	atom transfer duration	20 μ s	D	total maximal movement distance
v	atom moving speed	0.55 μ m/ μ s	T	compiled circuit execution time
d	atom distance	3 μ m	T_{idle}	qubit idle time
R_{int}	interaction radius	6 μ m	r_{int}	interaction factor
R_{restr}	restriction radius	12 μ m	r_{restr}	restriction factor

A. Fidelity

For easy of comparison across algorithms on different NAQC platforms, and following [23], we use the *approximate success probability* as a proxy for circuit fidelity. We assume single-qubit gates to be error-free in this study, disregarding their contribution to the overall error analysis. This is justified by their typically much higher fidelity compared to CZ gates.

Let T_2 be the atom dephasing time, n, m the numbers of qubits and CZ gates in the (compiled) circuit, T the total circuit execution time, s the number of atom transfers, v the speed of atom movement, and t_{cz} and t_{trans} the durations of a CZ gate execution and atom transfer, f_{cz} and f_{trans} the fidelities of gate CZ and atom transfer. The approximate success probability of the compiled circuit is

$$P(C) = \exp\left(-\frac{T_{idle}}{T_2}\right) \times f_{cz}^m \times f_{trans}^s, \quad (3)$$

where

$$T_{idle} = n \times T - m \times t_{cz}. \quad (4)$$

We note that T is calculated as

$$T = h \times t_{cz} + s \times t_{trans} + D/v, \quad (5)$$

where h is the depth of the compiled circuit and D the sum of the maximal movement distance of each (parallel) movement over all atom movements. As Tetris uses no atom movement, the execution time of Tetris' compiled circuit is simply $h \times t_{cz}$.

B. Architecture assumptions

Various architectures have been proposed for NAQC [5], [6], [39], [40], each with distinct features. Some architectures support long-range interaction using individually addressed Rydberg lasers for CZ gate execution [6], [40], while others enable dynamic atom shuttling [5], [7], or employ a global Rydberg laser to illuminate the entire plane, allowing parallel execution of CZ gates [5]. These architectural differences have led to distinct compiler designs that leverage each platform's unique capabilities.

Tetris [10] operates with individually addressed CZ gates. To better exploit long-range interaction, atoms are compactly arranged in a *fixed* 2-dimensional array with distance d . For any given interaction radius $R_{int} \triangleq r_{int} \times d \geq d$, a smaller d implies a larger interaction factor (i.e., r_{int}) and greater qubit connectivity (cf. Fig. 1). Notably, Tetris adopts a variable restriction area that depends on the inter-qubit distance of the gate, and this area is strictly smaller than the restriction area we define when R_{restr} is set equal to R_{int} . This difference allows Tetris to achieve more parallel CZ gates than DasAtom, even when we set R_{restr} to its minimal.

Enola [15] and Atomique [13] are designed for DPQA architectures, which support atom shuttling during computation and use a global Rydberg laser to apply CZ gates in parallel. To facilitate parallel execution and avoid unwanted interactions, both compilers assume an atom distance of 2.5 times the interaction radius. Enola uses one SLM (static) array and one AOD (movable) array, applying CZ gates between qubits on the SLM array and allowing atoms to transfer between arrays. This enables qubit routing to be fully accomplished via atom shuttling. Atomique assumes one SLM and multiple AOD arrays, applying CZ gates between atoms across different arrays and allowing SWAP gates during compilation, without requiring atoms to be transferred between arrays.

TABLE II: Comparison of DasAtom, Tetris, Enola, and Atomique on benchmark circuits from [10]. ‘Depth’ represents the number of CZ layers in the input circuit, ‘F’ and ‘RT’ denote the fidelity and runtime of the corresponding algorithm, ‘D’ indicates the CZ-depth of the compiled circuit, ‘M’ is the number of move stages, ‘P’ is the number of subcircuits, and ‘SW’ represents the number of SWAPs. Circuits are ordered by depth.

Circuit Info.				DasAtom				Tetris				Enola (static)				Atomique				
Name	#Q	#CZ	Depth	F	M	D	P	RT (s)	F	SW	D	RT (s)	F	M	RT (s)	F	SW	M	D	RT(s)
4mod5-v1_22	5	11	10	0.9463	0	11	1	0.01	0.9322	1	14	1.9	0.9311	20	46.38	0.9378	0	2	11	0.433
alu-v0_27	5	17	15	0.9183	0	17	1	0.01	0.8911	2	20	4.37	0.896	32	61.92	0.9081	0	5	16	0.434
bv_n16	16	15	15	0.9195	9	15	2	0.03	0.8224	8	21	0.61	0.8573	30	56.54	0.9244	0	15	15	0.434
qv_n16_d5	16	120	15	0.5436	8	120	2	0.69	0.3762	25	82	13.47	0.3435	132	126.23	0.5221	2	43	46	0.907
mod5mils_65	5	16	16	0.9229	0	16	1	0.01	0.9091	1	19	5.2	0.9005	32	56.91	0.9127	0	5	16	0.63
ising_model_13	13	120	20	0.5479	0	120	1	0.06	0.5159	4	77	30.62	0.3974	80	154.6	0.5305	0	41	51	0.715
ising_model_10	10	90	20	0.6368	0	90	1	0.05	0.5907	5	57	24.18	0.5342	80	131.1	0.6187	0	28	41	0.905
ising_model_16	16	150	20	0.4713	0	150	1	0.1	0.4117	9	88	118.1	0.2941	120	140.9	0.4473	0	38	52	0.792
decod24-v2_43	4	22	22	0.8956	0	22	1	0.02	0.7705	10	52	7.62	0.8717	44	126.1	0.8813	0	2	22	0.67
4gt13_92	5	30	26	0.8604	0	30	1	0.02	0.8604	0	30	5.22	0.8221	58	108.29	0.8424	0	5	26	0.641
qv_n12_d10	12	180	30	0.4032	8	180	2	0.27	0.2622	29	119	15.43	0.2383	204	209	0.3334	5	48	89	1.35
qft_16	16	240	58	0.2883	41	240	6	0.32	0.1549	44	204	118.1	0.0957	380	238	0.1476	17	125	217	1.486
rd84_142	15	154	81	0.4591	7	154	2	0.31	0.3369	21	146	60.23	0.2448	238	161.2	0.338	5	55	118	1.168
sf_274	6	336	300	0.1855	0	336	1	0.49	0.1721	5	343	218	0.1012	636	879.3	0.1186	3	84	305	1.83
sf_276	6	336	301	0.1855	0	336	1	0.42	0.1645	8	341	147.4	0.102	620	693.3	0.1060	5	93	309	1.89
con1_216	9	415	346	0.1237	17	415	3	0.48	0.0897	22	415	147.6	0.0399	776	965.1	0.0462	12	146	378	2.295
wim_266	11	427	352	0.1166	12	427	3	0.5	0.0819	24	436	163.9	0.029	758	549.7	0.0359	10	124	386	2.294
rd53_130	7	448	383	0.1054	9	448	3	0.56	0.0807	18	459	206.8	0.0404	852	943.2	0.0185	20	129	434	2.615
f2_232	8	525	449	0.0715	13	525	3	0.54	0.0583	14	523	251.6	0.0206	978	900.2	0.016	13	151	468	2.727
cm152a_212	12	532	461	0.0691	7	532	2	1.22	0.0571	13	498	295.8	0.0101	1010	598	0.0218	6	184	474	2.643
rd53_251	8	564	492	0.0587	16	564	4	1.07	0.0438	20	577	275.2	0.0149	1089	911.3	0.0110	16	178	527	3.039
hwb5_53	6	598	535	0.0495	23	598	6	0.58	0.0375	19	622	337.2	0.0168	1136	1354	0.004	27	216	609	3.356
pm1_249	14	771	634	0.0205	23	771	4	1.8	0.0122	36	763	385.1	8.18E-04	1378	755	0.0013	15	231	686	3.944
cm42a_207	14	771	634	0.0205	23	771	4	1.32	0.012	37	754	0.61	8.93E-04	1428	802.5	0.0013	15	231	686	4.215
dc1_220	11	833	705	0.0152	19	833	4	2.02	0.0083	41	838	380.6	9.51E-04	1548	1194	0.0010	15	189	750	4.25
squar5_261	13	869	720	0.0125	31	869	6	1.91	5.35E-03	58	866	409.9	4.11E-04	1630	900.2	3.95E-04	24	290	790	4.664
z4_268	11	1343	1112	0.0012	46	1343	8	2.22	3.41E-04	83	1369	687.2	1.31E-05	2500	2017	7.90E-07	44	381	1227	8.308
radd_250	13	1405	1210	8.36E-04	54	1405	9	2.19	2.06E-04	96	1430	700.6	3.08E-06	2658	1826	3.05E-07	44	427	1306	8.75
adr4_197	13	1498	1249	5.25E-04	52	1498	8	2.01	1.06E-04	109	1519	754.3	1.39E-06	2786	1682	1.59E-07	44	446	1372	9.669
sym6_145	7	1701	1499	0.0002	38	1701	9	1.84	9.47E-05	49	1729	947.5	5.02E-06	3286	3210	4.66E-07	37	462	1582	11.47
misex1_241	15	2100	1797	2.54E-05	56	2100	7	3.78	6.40E-06	95	2040	1017	1.50E-09	3936	2144	9.75E-10	29	625	1852	15.52
rd73_252	10	2319	1963	8.47E-06	87	2319	14	4.024	8.79E-07	154	2317	1130	6.50E-09	4930	3011	1.04E-12	63	731	2161	18.57
square_root_7	15	3089	2520	1.64E-07	147	3089	17	8.855	5.99E-09	229	2940	1583	1.56E-13	5716	2940	9.73E-15	41	835	2677	25.28
Geo.Mean	-	-	-	0.03458	-	-	-	0.3824	0.01951	-	-	84.21	0.0033	-	472.6	0.0018	-	-	-	2.345

To leverage both long-range interaction and atom shuttling, DasAtom assumes an NAQC architecture that supports in-computation atom shuttling, similar to DPQA, but with individually addressable Rydberg lasers for executing CZ gates. This setup requires one SLM array and one AOD array, with atoms able to transfer between layers. To facilitate long-range interaction, we assume atoms are initially arranged in a 2-dimensional SLM array with distance $d \leq R_{\text{int}}$. While no current NAQC devices combine both capabilities, devices demonstrating each individually exist [5], [6], and there are no fundamental physical limitations preventing the future integration of these two features.

C. Overall fidelity comparison

In our experiments, the common NA hardware parameters we assumed are summarised in Table I. In addition, we assume

- Atom distance d : $3\mu\text{m}$ for Tetris and DasAtom, $15\mu\text{m}$ for Enola and Atomique.
- $R_{\text{int}} = 6\mu\text{m}$ and $R_{\text{restr}} = 12\mu\text{m}$.

For DasAtom and Tetris, the default coupling graph is set to grid of $\lceil\sqrt{n}\rceil \times \lceil\sqrt{n}\rceil$. For Atomique, we use the default configuration of two AOD arrays.

1) *Results on benchmark circuits used in [10]*: In [10], the authors evaluated, among others, 33 benchmark circuits from RevLib and IBM Qiskit. After proper decomposition, these circuits include qubit counts ranging from 5 to 16 and up to 3,089 CZ gates. We empirically evaluated the performance of the algorithms on these benchmarks, and the results are summarised in Table II.

During compilation, we assume that each SWAP gate inserted by Tetris and Atomique is decomposed into three CZ gates and several single-qubit gates. In our calculation of their overall fidelity, we completely disregard the presence of additional single-qubit gates in the compiled circuit. As a result, the calculated fidelity is slightly higher than the actual value.

Recall that Enola operates in both static and dynamic modes. For these benchmark circuits, the two modes produce similar overall fidelity, with the dynamic mode performing slightly better. This improvement comes at the cost of a significantly increased runtime. Indeed, for all circuits with CZ depth ≥ 30 , the dynamic method does not complete within 1 hour. So we only show results for the static method in Table II.

Atomique allocates qubits to distinct arrays and does not permit atom transfer. Instead, it may introduce SWAP gates during compilation process. Compilation results also indicate the use of ancilla qubits.

From Table II, we observe the following:

- 1) **DasAtom consistently outperforms Tetris and Enola in fidelity and surpasses Atomique in all but one case.** On average, DasAtom’s fidelity is 1.77, 10.48, and 19.21 times that of Tetris, Enola, and Atomique, respectively. The largest improvement occurs on the circuit ‘square_root_7’, with fidelity ratios of 27.38, 1.05E06, and 1.69E07, respectively.
- 2) For each algorithm and circuit, compilation can be completed within one hour, with DasAtom and Atomique being significantly faster than Tetris and Enola.

TABLE III: Comparison among DasAtom, Tetris, Enola, and Atomique on 20-qubit quantum benchmark circuits with varying topologies, where the columns have the same interpretation as Table II and ‘a.c.’ denotes ‘almost complete’.

Circuit Info.				DasAtom				Tetris			Enola (dynamic)		Atomique			
Name	#CZ	Depth	topology	F	M	D	P	F	SW	D	F	M	F	SW	M	D
QV	600	60	a.c.	0.0456	61	579	6	0.0055	146	356	0.0039	508	0.027	22	185	230
2-local random	570	77	a.c.	0.0467	166	540	18	0.0013	250	549	0.0037	603	0.0208	37	285	341
QFT	410	77	complete	0.1176	68	398	8	0.0237	112	324	0.0178	424	0.0606	18	157	252
W-state	38	21	linear	0.8265	0	38	1	0.6696	14	42	0.6845	55	0.7873	0	28	38
3-regular graph	30	8	3-regular	0.8603	0	30	1	0.5997	24	35	0.7522	32	0.8502	0	9	11
DJ	19	19	star	0.9025	6	19	2	0.7255	15	31	0.8044	37	0.9024	0	19	19
GHZ	19	19	linear	0.9091	0	19	1	0.7365	14	28	0.8119	37	0.8868	0	14	19
Geo.Mean	-	-	-	0.2832	-	-	-	0.0865	-	-	0.0979	-	0.2104	-	-	-




Fig. 9: Interaction graphs of (a) QFT-5 and two-local random, (b) QV, (c) 3-regular graph, (d) DJ, and (e) GHZ and W-state

2) *Results on practical quantum circuits with varying qubit number and topologies:* To assess the scalability of our algorithm, we also evaluated the performance of the four algorithms on larger circuits, extracted from MQTBench² and [15]. These circuits include DJ, GHZ, QFT, W-state, QV, two-local random circuits, and 3-regular MaxCut QAOA circuits, which have very different interaction graphs (cf. Fig. 9). Since the dynamic mode of Enola performs exponentially better than the static mode as the number of qubits increases, we focus solely on the dynamic mode in these experiments.

Table III summarises the circuit information and performance details of the algorithms on seven 20-qubit circuits. As shown in the table, DasAtom consistently outperforms the other three in overall fidelity. In addition, the fidelity of DasAtom is strongly correlated with the number of CZ gates in the circuit. Table IV also presents the geometric mean of fidelity for each type of circuit with varying qubit counts. The table indicates that (i) DasAtom significantly outperforms both Tetris and Enola, and (ii) Atomique performs comparably to, or even slightly better than, DasAtom on DJ, W-state, and 3-regular circuits. However, on QFT circuits, DasAtom’s fidelity is, on average, 7.9x of that of Atomique. Notably, Atomique also outperforms Enola on these generic circuits; as shown in [15], Enola excels when leveraging gate commutativity.

TABLE IV: Geometric mean fidelity of various circuits

Circuit	#Q	DasAtom	Tetris	Enola	Atomique
QFT	5-50	0.0066	3.65E-04	2.63E-06	8.31E-04
two local	5-30	0.0613	0.0044	0.0044	0.0265
QV	5-25	0.1411	0.0393	0.0398	0.1002
DJ	5-50	0.8412	0.6791	0.6413	0.8615
W-state	5-50	0.7665	0.5424	0.4763	0.6945
3-regular	10-60	0.8548	0.2547	0.2785	0.7890

We now provide a detailed comparison of these algorithms’ performance across the different circuit types.

a) *QFT circuits:* These circuits also have complete interaction graphs, see Fig. 9a. In Fig. 10a, the y-axis denotes the ratio of DasAtom’s fidelity compared to that of Tetris, Enola, and Atomique, plotted on a logarithmic scale. The results clearly demonstrate that the ratio increases exponentially with the number of qubits, ranging from 5 to 50. In terms of runtime, Fig. 10b shows that Enola and Tetris are three and two orders of magnitude slower than DasAtom, respectively, while Atomique’s runtime is closer to that of DasAtom.

b) *Two-local random circuits:* These circuits have complete interaction graphs. Fig. 11a shows that DasAtom’s fidelity outperforms of Tetris, Enola, and Atomique, with performance ratios increasing exponentially as the qubit count ranges from 5 to 30. For Atomique, the increase in the performance ratio is more gradual. Note that only results for circuits with up to 30 qubits are shown here, as Enola’s runtime is very slow for larger circuits.

c) *Quantum Volume circuits:* These circuits have (nearly) complete interaction graphs. Fig. 11b shows that DasAtom’s fidelity outperforms that of Tetris and Enola, with performance ratios increasing exponentially with qubit count. Meanwhile, DasAtom achieves an average fidelity improvement of 1.4 times over Atomique.

d) *3-Regular MaxCut QAOA circuits:* These circuits feature 3-regular interaction graphs. Fig. 11c shows that the fidelity of all four algorithms decrease as the qubit count increases. DasAtom significantly outperforms both Tetris and Enola, while its fidelity remains close to that of Atomique.

e) *Deutsch-Jozsa circuits:* These circuits have star-like interaction graphs. Fig. 11d shows how fidelity decreases as the qubit count increases for the four algorithms. DasAtom outperforms both Tetris and Enola by a large margin, but is slightly outperformed by Atomique. This may be because the

²<https://www.cda.cit.tum.de/mqtbench/>

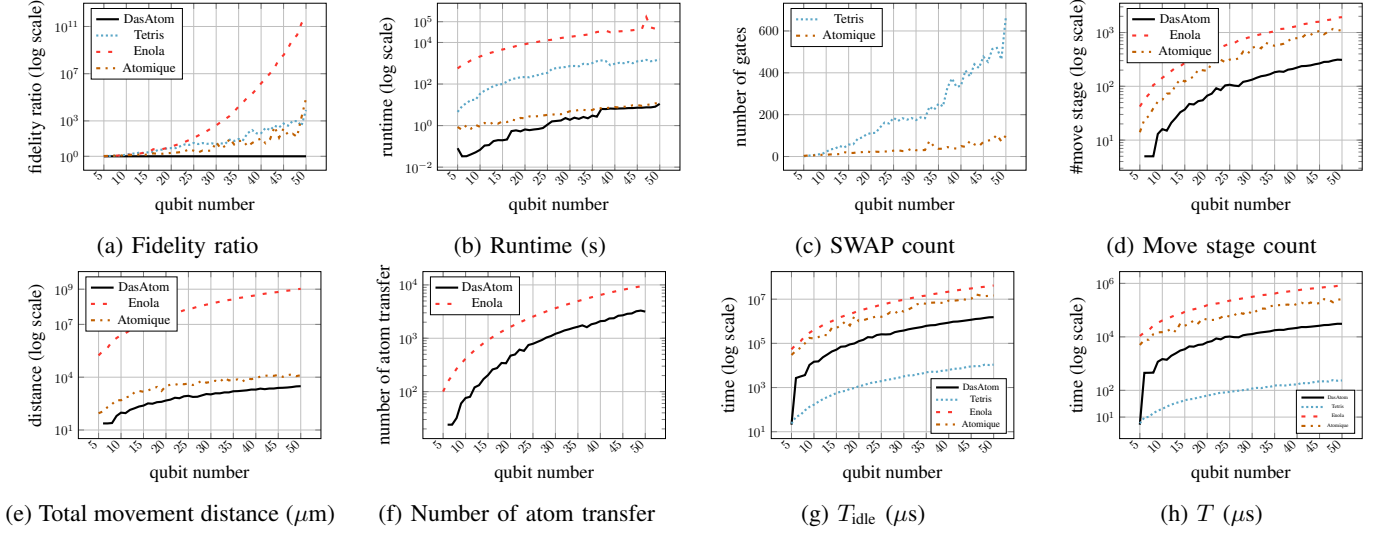


Fig. 10: Change of key performance indicators of QFT circuits under different algorithms

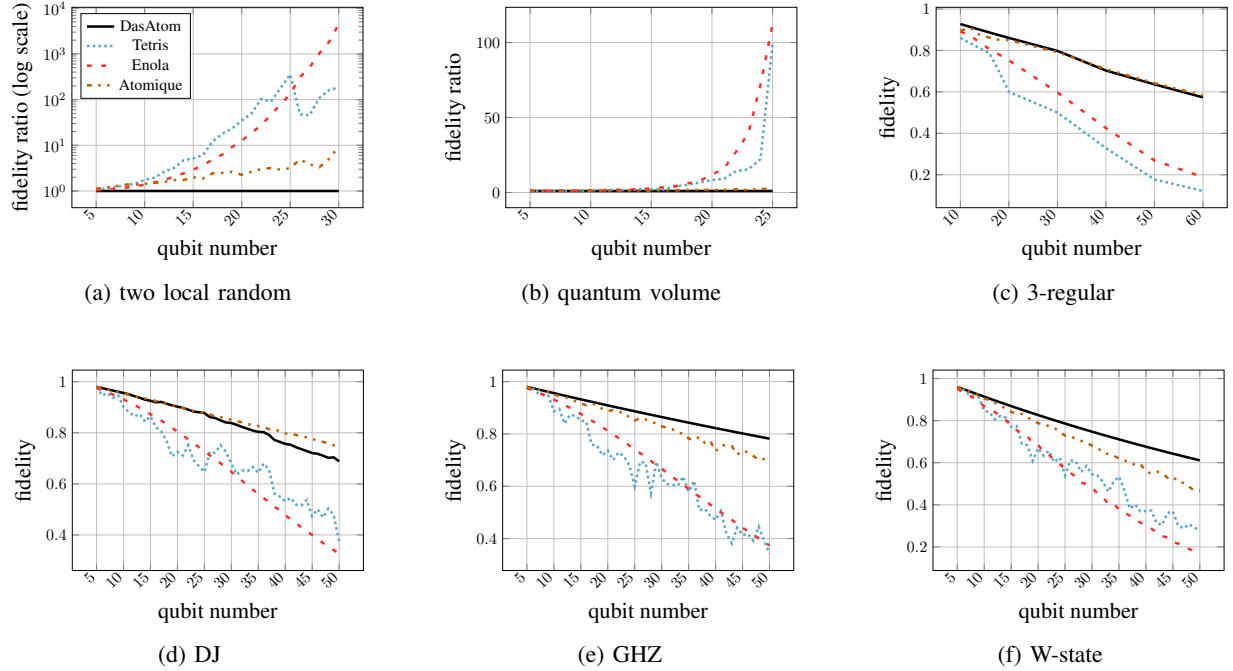


Fig. 11: Fidelity comparison on two-local random, QV, 3-regular, DJ, GHZ, and W-state circuits

star-like graph is a bipartite graph, allowing the central qubit to be placed on one AOD array, with all other qubits on the SLM array, thus avoiding intra-array two-qubit interactions.

f) *GHZ circuits and W-state circuits*: These circuits have linear interaction graphs. Figs. 11e and 11f show how the fidelity decreases as the qubit count increases for all algorithms. DasAtom clearly outperforms the other three.

D. Analysis

The experiments demonstrate that DasAtom outperforms the other three algorithms, achieving exponentially better results in QFT and QV circuits. In this subsection, we address two key questions: Which hardware parameters most significantly im-

pact DasAtom's performance? What factors enable DasAtom to outperform the other algorithms?

1) Ablation studies on hardware parameters:

a) *Interaction and restriction radii*: We start by evaluating DasAtom's performance with different interaction radii $r_{\text{int}} \in \{1, \sqrt{2}, 2\}$ on DJ circuits (5 to 45 qubits) and QFT circuits (5 to 20 qubits). The results indicate that setting $r_{\text{int}} = 2$ yields slightly better performance than other values, justifying our choice of $R_{\text{int}} = 2d$ in the evaluations presented in Sec. IV-C.

To analyse the effects of different interaction and restriction radii on DasAtom, we examined the QFT-30 circuit, varying r_{int} and r_{restr} from 1 to 8, with the condition that




Fig. 12: Fidelity comparison on QFT-30 with various interaction radii, where values at $r_{\text{int}} = 1, 5, 8$ are explicitly shown.

$r_{\text{restr}} \geq r_{\text{int}}$. Notably, when $r_{\text{int}} \geq 9$, the architecture graph becomes complete on $G(6, 6)$, rendering further increases in r_{int} unnecessary. As shown in Fig. 12, increasing r_{int} from 1 to 8 has a more pronounced effect on Tetris than on DasAtom: the fidelity of Tetris rises from $2.3\text{E-}06$ to 0.01, while DasAtom’s fidelity increases from 0.007 to 0.01. This difference is partly related to the fact that Tetris uses a smaller restriction area than DasAtom, even when $r_{\text{restr}} = r_{\text{int}}$.

For DasAtom, we also experimented with different values of r_{restr} for each r_{int} . The findings show that as r_{restr} increases, fidelity gradually decreases. Once r_{restr} reaches a certain threshold, the fidelity stabilises because all gates must be executed serially. However, the reduction in fidelity due to increasing r_{restr} is nearly negligible.

b) The most important factors: From Eq. 3, the approximate success probability of the compiled circuit, $P(C) = \exp(-\frac{T_{\text{idle}}}{T_2}) \times f_{\text{cz}}^m \times f_{\text{trans}}^s$ exponentially depends on T_{idle} and polynomially on f_{cz} and f_{trans} , raised to the powers of the CZ count m and atom transfer count s , respectively. To identify the parameters that most affect performance, we conducted ablation studies on the QFT-20 circuit, focusing on hardware parameters f_{trans} , f_{cz} , T_2 , and atom distance d . The results, shown in Fig. 13, reveal the following:

- f_{cz} has the most significant impact on all algorithms.
- f_{trans} also affects DasAtom and Enola significantly, but has no impact on the other two algorithms.
- Coherence time T_2 has a strong influence on all algorithms except Tetris, when T_2 varies from 0.15s to 15s.
- Increasing the atom distance from 1 to 20 reduces DasAtom’s fidelity by only 3.4%. This parameter clearly affects Atomique and Enola, but has no impact on Tetris’ fidelity.

Like atom distance, other parameters, including atom movement speed v , and CZ and atom transfer durations (t_{cz} and t_{trans}), have relatively limited impact on DasAtom’s fidelity. Specifically, reducing v from 0.55 to 0.25, increasing t_{cz} from 0.2 to 0.5, or increasing t_{trans} from 20 μs to 40 μs decreases fidelity by approximately 1.4%, 0.15%, and 11.8%, respectively.

Similar findings are observed for QV-20. However, due to the simpler topologies of the 20-qubit GHZ and 3-regular graph circuits, the impact of f_{cz} and f_{trans} is somewhat less pronounced, though the overall trend remains similar to that

for QFT-20 and QV-20.

2) *Why DasAtom performs better?:* We analyse here why DasAtom outperforms the other algorithms.

a) SWAP gates have low fidelity: Based on our hardware parameter settings (see Table. I), the fidelity cost of inserting a SWAP gate can be approximated as $f_{\text{cz}}^3 = 0.995^3 = 0.985$. In contrast, the fidelity cost of moving an atom at a speed of $0.55 \mu\text{m}/\mu\text{s}$ over a distance $x \mu\text{m}$ can be estimated as $\exp(-\frac{2 \times T_{\text{trans}} + x/0.55}{T_2})$, where $2 \cdot T_{\text{trans}}$ represents the time for the atom to transition between SLM to AOD traps at both the initial and target positions. A simple calculation shows that the fidelity drops to 0.985 when $x \approx 12,384 \mu\text{m}$. Additionally, for swapping the positions of two atoms, the distance can be up to 6,202 μm for the fidelity to remain at 0.985. This illustrates that SWAP gates incur significant fidelity costs. As shown in Fig. 10c, Tetris has a much rapid increase in SWAP count compared to Atomique.

b) Atom transfer and atom shuttling are slow: In contrast, atom transfer and shuttling can be up to 100x slower than SWAP gates. This implies that the overall circuit execution time T , and consequently the qubit idle time T_{idle} , may be significantly prolonged for algorithms dependent on atom shuttling. As shown in Fig. 10e-h, Enola exhibits the fastest increases in total movement distance, atom transfer count, T_{idle} , and T among the three algorithms that use atom shuttling.

Remark 2. Fig. 10a demonstrates an exponential increase in DasAtom’s fidelity ratio compared to Tetris, Enola, and Atomique for QFT circuits. Since this evaluation assumes $f_{\text{trans}} = 1$, the approximate success probability (cf. Eq. 3) is essentially $P(C) = \exp(-\frac{T_{\text{idle}}}{T_2}) \times f_{\text{cz}}^m$. Here, m (the CZ gate count) generally increases with the qubit count n . DasAtom’s exponential fidelity improvement over Tetris (Atomique) is primarily (partly) due to the additional SWAP gates inserted, which increase at least linearly with n (cf. Fig. 10c). For Enola (Atomique), the increase is primarily (partly) driven by higher T_{idle} values as n grows (cf. Eq. 4 and Fig. 10g), caused by increased atom transfers and total movement distance (cf. Eq. 5 and Fig. 10e,f). Note that if $f_{\text{trans}} < 1$, DasAtom’s fidelity ratio compared to Enola will be even higher, as Enola’s increasing atom transfer count s (appearing as an exponent of f_{trans} in Eq. 3) grows faster than that of DasAtom (cf. 10f).

V. FURTHER DISCUSSIONS

The evaluations in the previous section demonstrate that, despite its simplicity, the DAC scheme is highly effective. This effectiveness can be attributed to two primary factors: first, atom shuttling incurs significantly lower fidelity costs than SWAP gate insertion (cf. Sec. IV-D2a); second, long-range interactions enable dense qubit connectivity, minimising the need for mapping transitions and ultimately lowering atom movement costs.

To implement DasAtom, we assume the NAQC architecture has the following capabilities: high-fidelity atom shuttling during computation, long-range interaction, and individually applied CZ gates. Notably, a significant portion of DasAtom’s

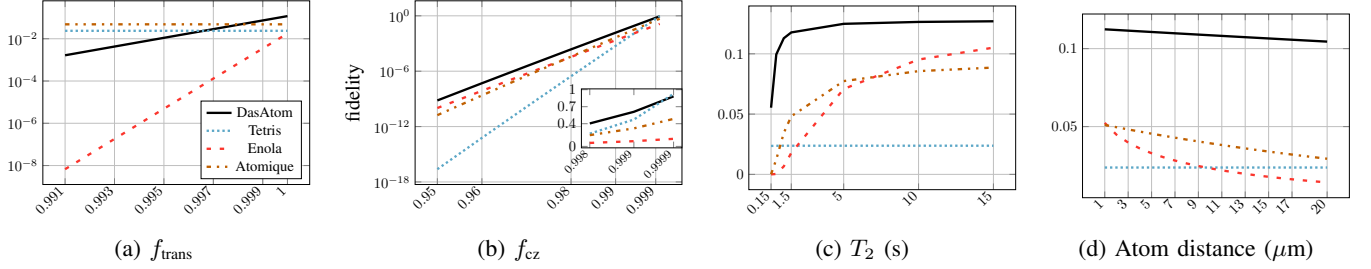


Fig. 13: Ablation studies on QFT-20 for hardware parameters f_{cz} , f_{trans} , T_2 , and atom distance d .

improvement over Tetris stems from the hardware capability of in-computation atom shuttling. A rough estimate indicates that, when simulating SWAP gates with atom shuttling (cf. [23]), DasAtom outperforms Tetris by 3.14x on QFT-50. By contrast, Fig. 10a shows an improvement of 8,108.5x under the default settings.

DasAtom assumes a compact atom arrangement with atom distance smaller than the interaction radius, while Atomique and Enola assume an atom distance that is 2.5x the interaction radius. This configuration results in shorter atom shuttling distances for DasAtom. Recall that the total movement distance D can contribute significantly to the circuit execution time T and qubit idle time T_{idle} (cf. Eq. 4 and Eq. 5). One might wonder if this compact atom spacing is the primary reason for DasAtom’s exponential performance advantage over Enola as the qubit count increases on QFT circuits. However, as shown in Fig. 13d, increasing the atom distance from 1 to 20 μm only reduces DasAtom’s fidelity on QFT-20 by about 7%. The primary reason for DasAtom’s success over Enola lies in its exponentially shorter total movement distance and lower atom transfer count (cf. Fig. 10e,f), which together result in reduced circuit execution time T and shorter qubit idle time T_{idle} .

A. Implications

Although a denser atom configuration is generally advantageous, atom distance does not significantly impact DasAtom’s improvement over other algorithms. In fact, a smaller atom distance translates to a larger interaction factor r_{int} . As shown in Fig. 12 and Fig. 13d, changes in fidelity remain modest as r_{int} or atom distance varies for DasAtom.

Our chosen architecture parameters are optimistic, and there are scenarios where Tetris may outperform DasAtom. For instance, Fig. 14 shows that for the QFT-50 circuit, Tetris achieves higher fidelity than DasAtom under certain conditions, such as when $f_{trans} \leq 0.99$ and $f_{cz} \geq 0.98$. In general, Tetris performs comparatively better when f_{cz} is close to 1, f_{trans} is low, T_2 is small, and a larger r_{int} is available.

Enola and Atomique are suitable when using a DPQA architecture, where individually applied CZ gates are unavailable. This is particularly advantageous for certain circuits in QAOA and quantum simulation, where gate commutativity can be exploited. For generic circuits, Enola’s fidelity improves faster with higher values of f_{trans} and T_2 (cf. Fig. 13a,c). However, due to the exponential increase in total movement distance

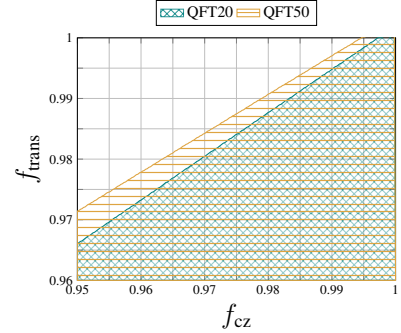


Fig. 14: Shaded areas show where Tetris achieves higher fidelity than DasAtom for QFT-20 (teal cross-hatch) and QFT-50 (orange horizontal lines).

and atom transfer count with the qubit count (cf. Fig. 10e,f), Enola is unlikely to outperform DasAtom when CZ gates are individually addressable.

When atom transfer fidelity is relatively low but CZ gate fidelity is high, Atomique is a more suitable choice, especially for circuits with moderate depth (e.g., ≤ 100 , cf. Table II). Additionally, Atomique is scalable, but it is also worth noting that it requires multiple AOD arrays and ancilla qubits.

B. Limitation and potential improvement

It is important to note that our implementation is not yet fully optimised. Specifically, the mappings for each subcircuit could be designed to enable smoother transitions, making the process more efficient. For example, the mapping of the second part of QFT-5 shown in Fig. 7c is not optimal. A better mapping would be to place q_4 at (0, 2) and q_1 at (1, 0), requiring only an exchange between the positions of q_0 and q_4 .

Although our algorithm relies on subgraph isomorphism checks, which are theoretically NP-hard, this has not proven to be a significant obstacle in most cases. By leveraging the Rustworkx implementation of VF2, we successfully compiled the 500-qubit QFT in about one hour. However, this implementation does not perform as well on certain circuits, like QV, especially when the number of qubits exceeds 25. To handle circuits with thousands of qubits, more efficient (though potentially approximate) subgraph isomorphism algorithms (like PathLAD+ [41]) will be necessary. Given that our target architecture is grid-based, this should not pose a major challenge. However, as shown in the fidelity ratio curves

in Fig. 10a, the primary concern for compilation on NAQC hardware should be low fidelity rather than scalability.

Our divide-and-shuttle method could be integrated with the multi-AOD array approach introduced in Atomique [13]. In Atomique, program qubits are distributed across different arrays. Interactions between qubits in different arrays are performed via atom shuttling, while SWAP gates are used for interactions between qubits within the same array. However, as shown in Table II, Atomique’s performance degrades significantly even compared to Enola when the circuit depth exceeds 300. A hybrid strategy might involve partitioning the circuit into subcircuits, employing Atomique within each subcircuit, and utilising atom shuttling to connect consecutive subcircuits. Although this would introduce atom transfers, it has the potential to enhance overall fidelity by reducing the number of inserted SWAP gates.

VI. CONCLUSION

We have introduced DasAtom, a novel algorithm designed to optimise the execution of quantum circuits on neutral atom platforms by exploiting long-range interactions and atom shuttling. Unlike previous methods such as SWAP-based Tetris and atom move-based Enola, DasAtom achieves significant fidelity improvements through partitioning circuits into subcircuits and dynamically adjusting qubit mappings. Our results demonstrate that for a 30-qubit QFT, DasAtom outperforms Enola by 415.8x and Tetris by 10.6x, illustrating the substantial gains in fidelity that can be achieved with this method.

DasAtom’s ability to scale its performance with the increasing number of qubits highlights its potential as a crucial tool in the future of NAQC. As quantum circuits grow in complexity and size, the exponential improvement in fidelity offered by DasAtom will become increasingly valuable. This work not only provides a pathway to more efficient quantum circuit execution on NAQC platforms but could also inspire hardware experts to explore and experiment with new architectures. Future work will focus on expanding DasAtom’s capabilities to handle even larger and more complex quantum circuits.

ACKNOWLEDGEMENT

We gratefully acknowledge the anonymous reviewers for their constructive feedback, which has significantly enhanced the presentation of this work.

REFERENCES

- [1] P. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 1994, pp. 124–134.
- [2] Y. Cao, J. Romero, J. P. Olson, *et al.*, “Quantum chemistry in the age of quantum computing,” *Chemical Reviews*, vol. 119, no. 19, pp. 10856–10915, 2019.
- [3] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [4] L. Henriët, L. Beguin, A. Signoles, *et al.*, “Quantum computing with neutral atoms,” *Quantum*, vol. 4, p. 327, 2020.
- [5] D. Bluvstein, H. Levine, G. Semeghini, *et al.*, “A quantum processor based on coherent transport of entangled atom arrays,” *Nature*, vol. 604, no. 7906, pp. 451–456, 2022.
- [6] A. G. Radnaev, W. C. Chung, D. C. Cole, *et al.*, “A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout,” 2024. arXiv: 2408.08288 [quant-ph].
- [7] A. Lengwenus, J. Kruse, M. Schlosser, S. Tichelmann, and G. Birkel, “Coherent transport of atomic quantum states in a scalable shift register,” *Phys. Rev. Lett.*, vol. 105, p. 170502, 17 Oct. 2010.
- [8] J. M. Baker, A. Litteken, C. Duckering, H. Hoffmann, H. Bernien, and F. T. Chong, “Exploiting long-distance interactions and tolerating atom loss in neutral atom quantum architectures,” in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, 2021, pp. 818–831.
- [9] T. Patel, D. Silver, and D. Tiwari, “Geyser: A compilation framework for quantum computing with neutral atoms,” in *Proceedings of the 49th Annual International Symposium on Computer Architecture*, ser. ISCA’22, New York, New York: Association for Computing Machinery, 2022, pp. 383–395.
- [10] Y. Li, Y. Zhang, M. Chen, X. Li, and P. Xu, “Timing-aware qubit mapping and gate scheduling adapted to neutral atom quantum computing,” *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 42, no. 11, pp. 3768–3780, 2023.
- [11] B. Tan, D. Bluvstein, M. D. Lukin, and J. Cong, “Qubit mapping for reconfigurable atom arrays,” in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, ser. ICCAD ’22, San Diego, California: Association for Computing Machinery, 2022.
- [12] D. B. Tan, D. Bluvstein, M. D. Lukin, and J. Cong, “Compiling Quantum Circuits for Dynamically Field-Programmable Neutral Atoms Array Processors,” *Quantum*, vol. 8, p. 1281, Mar. 2024.
- [13] H. Wang, P. Liu, D. Tan, *et al.*, “Atomique: A quantum compiler for reconfigurable neutral atom arrays,” in *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*, Los Alamitos, CA, USA: IEEE Computer Society, Jul. 2024, pp. 293–309.
- [14] H. Wang, D. Tan, P. Liu, *et al.*, “Q-pilot: Field programmable qubit array compilation with flying ancillas,” in *Proceedings of the 61st ACM/IEEE Design Automation Conference (DAC)*, ACM/IEEE, 2024.
- [15] D. B. Tan, W.-H. Lin, and J. Cong, “Compilation for dynamically field-programmable qubit arrays with efficient and provably near-optimal scheduling,” *arXiv:2405.15095*, 2024.
- [16] A. Zulehner, A. Paler, and R. Wille, “An efficient methodology for mapping quantum circuits to the IBM QX architectures,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, no. 7, pp. 1226–1236, 2018.
- [17] G. Li, Y. Ding, and Y. Xie, “Tackling the qubit mapping problem for NISQ-era quantum devices,” in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2019, Providence, RI, USA, April 13-17, 2019*, I. Bahar, M. Herlihy, E. Witchel, and A. R. Lebeck, Eds., ACM, 2019, pp. 1001–1014.
- [18] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, “T|ket>: A retargetable compiler for NISQ devices,” *Quantum Science and Technology*, vol. 6, no. 1, p. 014003, 2020.
- [19] S. Li, X. Zhou, and Y. Feng, “Qubit mapping based on subgraph isomorphism and filtered depth-limited search,” *IEEE Trans. Computers*, vol. 70, no. 11, pp. 1777–1788, 2021.
- [20] X. Zhou, Y. Feng, and S. Li, “A monte carlo tree search framework for quantum circuit transformation,” in *2020 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, IEEE, 2020, pp. 1–7.
- [21] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” *arXiv:1411.4028*, 2014.

- [22] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, “Optical dipole traps for neutral atoms,” in *Advances in atomic, molecular, and optical physics*, vol. 42, Elsevier, 2000, pp. 95–170.
- [23] L. Schmid, D. F. Locher, M. Rispler, *et al.*, “Computational capabilities and compiler development for neutral atom quantum processors—connecting tool developers and hardware experts,” *Quantum Science and Technology*, vol. 9, no. 3, p. 033 001, Apr. 2024.
- [24] M. Endres, H. Bernien, A. Keesling, *et al.*, “Atom-by-atom assembly of defect-free one-dimensional cold atom arrays,” *Science*, vol. 354, no. 6315, pp. 1024–1027, 2016.
- [25] D. Barredo, S. De Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, “An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays,” *Science*, vol. 354, no. 6315, pp. 1021–1023, 2016.
- [26] D. Barredo, V. Lienhard, S. De Leseleuc, T. Lahaye, and A. Browaeys, “Synthetic three-dimensional atomic structures assembled atom by atom,” *Nature*, vol. 561, no. 7721, pp. 79–82, 2018.
- [27] M. Schlosser, S. Tichelmann, D. Schäffner, *et al.*, “Scalable multilayer architecture of assembled single-atom qubit arrays in a three-dimensional talbot tweezer lattice,” *Phys. Rev. Lett.*, vol. 130, p. 180 601, 18 May 2023.
- [28] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, “Fast quantum gates for neutral atoms,” *Physical Review Letters*, vol. 85, no. 10, p. 2208, 2000.
- [29] X. Zhou, S. Li, and Y. Feng, “Quantum circuit transformation based on simulated annealing and heuristic search,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 12, pp. 4683–4694, 2020.
- [30] L. Lao, H. van Someren, I. Ashraf, and C. G. Almudever, “Timing and resource-aware mapping of quantum circuits to superconducting processors,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 2, pp. 359–371, 2022.
- [31] P. Zhu, S. Feng, and Z. Guan, “An iterated local search methodology for the qubit mapping problem,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 8, pp. 2587–2597, 2022.
- [32] S. Brandhofer, I. Polian, and H. P. Büchler, “Optimal mapping for near-term quantum architectures based on rydberg atoms,” in *IEEE/ACM International Conference On Computer Aided Design, ICCAD 2021, Munich, Germany, November 1-4, 2021*, IEEE, 2021, pp. 1–7.
- [33] L. Schmid, S. Park, S. Kang, and R. Wille, “Hybrid circuit mapping: Leveraging the full spectrum of computational capabilities of neutral atom quantum computers,” in *Proceedings of the 61st ACM/IEEE Design Automation Conference (DAC)*, ACM/IEEE, 2024.
- [34] M. Y. Siraichi, V. F. d. Santos, C. Collange, and F. M. Q. Pereira, “Qubit allocation as a combination of subgraph isomorphism and token swapping,” *Proceedings of the ACM on Programming Languages*, vol. 3, no. OOPSLA, pp. 1–29, 2019.
- [35] T.-A. Wu, Y.-J. Jiang, and S.-Y. Fang, “A robust quantum layout synthesis algorithm with a qubit mapping checker,” ser. ICCAD ’22, San Diego, California: Association for Computing Machinery, 2022.
- [36] Y. Huang, X. Zhou, F. Meng, and S. Li, “Qubit mapping: The adaptive divide-and-conquer approach,” *arXiv*, vol. 2409.04752, 2024.
- [37] N. Nottingham, M. A. Perlin, D. Shah, *et al.*, “Circuit decompositions and scheduling for neutral atom devices with limited local addressability,” in *IEEE International Conference on Quantum Computing and Engineering (QCE) 2024*, 2024.
- [38] L. P. Cordella, P. Foggia, C. Sansone, and M. Vento, “A (sub)graph isomorphism algorithm for matching large graphs,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 10, pp. 1367–1372, 2004.
- [39] S. J. Evered, D. Bluvstein, M. Kalinowski, *et al.*, “High-fidelity parallel entangling gates on a neutral-atom quantum computer,” *Nature*, vol. 622, no. 7982, pp. 268–272, Oct. 2023.
- [40] T. M. Graham, Y. Song, J. Scott, *et al.*, “Multi-qubit entanglement and algorithms on a neutral-atom quantum computer,” *Nature*, vol. 604, no. 7906, Apr. 2022.
- [41] Y. Wang, C. Jin, and S. Cai, “Pathlad+: Towards effective exact methods for subgraph isomorphism problem,” *Artificial Intelligence*, vol. 337, p. 104 219, 2024.