

An Abstract Model and Efficient Routing for Logical Entangling Gates on Zoned Neutral Atom Architectures

Yannick Stade*, Ludwig Schmid*, Lukas Burgholzer*, and Robert Wille*†

*Chair for Design Automation, Technical University of Munich, Munich, Germany

†Software Competence Center Hagenberg GmbH, Hagenberg, Austria

{yannick.stade, ludwig.s.schmid, lukas.burgholzer, robert.wille}@tum.de

www.cda.cit.tum.de/research/quantum

Abstract—Recent experimental achievements have demonstrated the potential of neutral atom architectures for fault-tolerant quantum computing. These architectures feature the dynamic rearrangement of atoms during computation—enabling nearly arbitrary two-dimensional rearrangements. Additionally, they employ a zoned layout with dedicated regions for entangling, storage, and readout. This architecture requires design automation software that efficiently compiles quantum circuits to this hardware and takes care that atoms are in the right place at the right time. In this paper, we initiate this line of work by providing, (1) an abstract model of the novel architecture and, (2) an efficient solution to the routing problem of entangling gates. By this, we aim to maximize the parallelism of entangling gates and minimize the overhead caused by the routing of atoms between zones. In addition to that, we keep the realm of fault-tolerant quantum computing in mind and consider logical qubit arrays, each of which encodes one logical qubit. We implemented the proposed idea as a tool called NALAC and demonstrated its effectiveness and efficiency by showing that it can significantly reduce the routing overhead of logical entangling gates compared to the naive approach. As part of the *Munich Quantum Toolkit* (MQT), NALAC is publicly available as open-source at <https://github.com/cda-tum/mqt-qmap>.

Index Terms—quantum computing, compilation, quantum circuit routing, neutral atoms, quantum error correction, fault tolerance

I. INTRODUCTION

Over the past decade, significant efforts have been devoted to developing first, intermediate-scale quantum computing demonstrators and mitigating physical errors [1]—advancing both hardware and software development. Despite ongoing reductions in hardware errors, achieving the requisite error rates, on the order of 10^{-10} [2], necessary for executing large-scale meaningful algorithms such as integer factorization will likely require the implementation of sophisticated *Quantum Error Correction* (QEC) protocols [1]. This transition to *Fault-Tolerant Quantum Computing* (FTQC) necessitates advanced experimental setups that meet the novel requirements and constraints imposed by FTQC, alongside appropriate software support, to maximize the utilization of available hardware capabilities.

Experimental progress has demonstrated elementary FTQC operations on various hardware platforms, including superconducting chips [3]–[6], trapped ions [7]–[10], and neutral atoms [11]. However, these experiments often involve small, hand-constructed examples. In order to scale up to sizes relevant for practical use, sophisticated design automation methods and software are essential [12].

Neutral atoms have emerged as a promising platform for universal quantum computing [11], [13]–[16], offering long coherence times, arbitrary connectivity through dynamic atom rearrangement, and highly parallel gate execution between sets of atoms. Recent experimental breakthroughs were achieved using a novel zoned architecture, where different functionalities (storage, entangling, measurement) are performed in designated spatially separated zones, with shuttling facilitating the transfer of atoms between those zones [11]. Despite the rapid development of neutral atom-specific compilation software [17] and routing of atom movements [18]–[22], the absence of an appropriate abstract model and the lack of dedicated routing solutions for logical qubit arrays hinders the efficient utilization of this novel architecture and its capabilities.

In this work, we pioneer this line of work two-fold by

- 1) presenting an abstract model of the zoned architecture introduced in [11], including its capabilities, constraints, as well as the resulting problem of routing logical entangling gates, and
- 2) proposing an efficient approach to tackle this routing problem by reducing routing overhead and increasing gate parallelism.

To this end, we utilize graph theoretical techniques to derive a highly efficient solution for the routing problem. The core of the proposed approach is a specific coloring of the interaction graph representing the entangling gates in the quantum circuit. Thereby, compared to a naive solution, we schedule multiple entangling gates in parallel and ensure that multiple sets of entangling gates are applied without intermediate time-expensive shuttling of atoms between zones. We implemented the approach in the tool NALAC (*Neutral Atom Logical Array Compiler*) and benchmarked it on different quantum circuits, demonstrating its efficiency and effectiveness.

The model of the neutral atom architecture and the proposed compiler NALAC pave the road for further compiler development on similar architectures and introduce the first software tool providing an automated solution to the routing problem of transversal logical gates on zoned neutral atom architectures. The tool NALAC is integrated into the *Munich Quantum Toolkit* (MQT, [23]) that comprises simulators [24], validation tools [25], and tools to automatically select optimization passes [26] among many others. The full code of the proposed approach is publicly available at <https://github.com/cda-tum/mqt-qmap>.

The remainder of this paper is structured as follows: Section II provides background on FTQC and a brief introduction to neutral atoms. In Section III, we introduce

the model of the zoned architecture together with its capabilities and constraints. Section IV discusses the architecture-specific routing problem and illustrates the proposed solution compared to a naive approach. Technical implementation details are provided in Section V. Section VI evaluates the proposed approach on a set of benchmark circuits, demonstrating reduced routing overhead. Finally, Section VII concludes the paper.


II. PRELIMINARIES

In order to keep the remainder of this work self-contained, this section briefly describes the concept of logical quantum computing, which is essential when employing *Fault-Tolerant Quantum Computing* (FTQC). It then delves into the physical background of quantum computing using neutral atoms.

A. Fault-Tolerant Quantum Computing


The overarching principle of FTQC is the idea of distributing the information of a single *logical* qubit q^L across multiple *physical* qubits q_1, \dots, q_n . The resulting redundancy enables the detection of individual errors at the physical level and appropriate corrections [27], [28]. We henceforth refer to a (*logical*) *qubit array* as the set of physical qubits used to encode a single logical qubit. This is illustrated by means of the popular Steane code as the simplest example of the widely used color code [29].

Example 1 (Logical Qubit Arrays). *The Steane code requires $n = 7$ physical qubits q_1, \dots, q_7 to encode a single logical qubit q^L . Below, we illustrate the physical qubits arranged, e.g., in a 2×4 array. Throughout this paper, we will generally use squares to represent logical qubits and circles for physical qubits.*



The execution of operations on these encoded logical qubits (indicated by superscripts, e.g., CZ^L) depends on the specific QEC code and is generally highly complex—often requiring procedures such as braiding or magic state injection [30]. This study focuses on the special case of *transversal* gates, where the logical operation is executed by performing the same operation on each physical qubit individually. This prevents the possible spreading of errors, making transversal gates inherently fault-tolerant. In particular, we focus on codes that provide transversal entangling gates (such as CX, CZ, ...), encompassing surface codes [30], color codes [29], and, in general, the family of *Calderbank-Shor-Steane* (CSS) codes [27].

Example 2 (Transversal CZ gate). *To implement a logical CZ^L gate that is transversal for the respective code, a physical CZ gate must be applied between all pairs of qubits within the two logical arrays:*



Importantly, transversal two-qubit gates alone are insufficient for achieving universal fault-tolerant computation [31].




Figure 1. Zoned architecture considered within this work based on the experimental setup from [11].

Nevertheless, due to their inherent fault tolerance, they are a fundamental building block, essential for large-scale fault-tolerant computing. Therefore, correct handling of transversal gates constitutes the first important step toward a full compilation stack for FTQC.

B. Neutral Atom Quantum Computing

Quantum computing based on *neutral atoms* [3], [13]–[17] relies on qubits encoded into long-lived atomic states of alkali or alkaline-earth(-like) atoms such as Rb, Sr, or Yb. Atoms are confined within optical dipole traps generated by optical lattices or optical tweezers and, subsequently, laser-cooled to their motional ground state, facilitating the efficient trapping of thousands of atoms [32]–[34]. Recent experiments demonstrate the continuous reloading of atoms [35], [36] to compensate for potential losses.

Single-qubit gates are implemented by inducing state transitions using either global or focused laser beams, addressing the entire register or individual atoms [11], [16], [37], [38], respectively. Entangling gates exploit high-lying Rydberg states and their long-range dipole-dipole interactions among proximal atoms via Rydberg blockade [14], [39], [40]. Notably, parallel entanglement of numerous qubits is achieved by globally illuminating the qubit register, inducing two-qubit gates between pairs of qubits within the interaction range of the Rydberg blockade [41], [42]. Measurements are realized using fluorescence imaging or other non-destructive measurement techniques to read out qubit states in the computational basis [43]–[45].

Example 3 (Native Gates). *Figure 1 illustrates the implementation of single- and two-qubit gates utilizing Raman (blue, middle) and Rydberg (yellow, top) lasers, respectively. Regarding the Rydberg beam, it is worth noting that qubits nearby execute a CZ gate (top left), while isolated qubits undergo an identity operation (top right).*

Novel experimental techniques [11], [46] have demonstrated the ability to *shuttle* atoms during computation, i.e., to change the position of trapped atoms, allowing almost arbitrary two-dimensional rearrangements (this will be covered in more detail in Section III-C). This is achieved by introducing two possible types of optical traps between which atoms can switch: First, static *Spatial Light Modulator* (SLM) traps afford placement in nearly arbitrary yet fixed configurations; secondly, a dynamically




Figure 2. One execution cycle up to the execution of the CZ gate as described in Example 4.

adjustable 2D optical lattice, employing two *Acousto-Optic Deflectors* (AODs), arranged along the x- and y-axes, enables the parallel rearrangement of multiple qubits.

Overall, neutral atom architectures offer great parallelism for gate execution, which is favorable for transversal gates. The laser beams can illuminate entire qubit registers, inducing parallel operations on all illuminated qubits. Thus, a transversal operation on a logical qubit can be applied to all corresponding atoms in parallel.

III. ABSTRACT MODEL OF ZONED NEUTRAL ATOM ARCHITECTURES

So far, previous work has considered a monolithic architecture where all operations are performed in the same region [46]. In the recent experimental milestone [11], different operations are conducted in spatially separated zones. Quantum circuits are executed by alternately shuttling atoms between zones and illuminating them with the corresponding laser beams. In order to scale to large quantum circuits, the compilation process cannot be performed manually anymore, and design automation software is needed. However, this is still lacking as no abstract model for zoned architecture that captures its capabilities and constraints exists. To this end, we first review the zoned architecture based on the experimental setup from [11] and the execution steps. Based on that, we then formalize the resulting shuttling constraints and requirements for executing logical, transversal gates—eventually resulting in the required model.

A. Zoned Structure

In general, zoned architectures spatially separate certain operations, allowing favorable hardware setups and optimization. As depicted in Figure 1, the architecture considered in this work consists of three separate zones, each specifically designed for a particular purpose:

- a) **Entangling zone:** This zone is dedicated to the execution of entangling operations, i. e., CZ gates in the case of [11]. Using a global Rydberg beam applied to the whole zone, the atoms are excited to the Rydberg state. This applies parallel CZ entangling gates between all qubit pairs within the range of the Rydberg blockade of each other. Qubits not supposed to interact are placed sufficiently far from each other so that the Rydberg interaction can be neglected. A single qubit without another qubit nearby is still excited to the Rydberg state, but without an interaction partner,

it returns to its original state, effectively performing an identity operation. However, the Rydberg decay might possibly introduce an error similar to an actual CZ gate.

- b) **Storage zone:** This zone allows for densely packaging qubits and application of single-qubit gates while maintaining high coherence times. In this zone, the qubits are shielded from beams that could potentially disturb the quantum state (Rydberg or measurement). A large number of static SLM traps provide the necessary storing opportunities. Single-qubit rotations are realized using global and/or individually addressable laser beams.
- c) **Readout zone:** In this zone, the state of a set of qubits can be read out without disturbing the quantum state of other qubits. In particular, this enables mid-circuit measurements, which are essential for full-stack FTQC. Applying a readout beam measures all qubits within the readout zone simultaneously.

B. Quantum Circuit Execution Steps

While zoned architectures can yield improved hardware characteristics by shielding stored and operated qubits from each other, shuttling is required to move the respective qubits into the desired zones. Initially, the qubits are located in the storage zone, providing, e. g., the best coherence times. Given a sequence of transversal gates, single-qubit operations can be executed immediately without additional shuttling. However, as described above, CZ gates can only be performed in the entangling zone. Therefore, the task at hand is to shuttle qubits from the storage to the entangling zone, perform the CZ gate, and shuttle them back to the storage zone. In particular, for logical qubits, this requires the movement of all physical qubits corresponding to the logical qubit. One shuttle operation that moves (multiple) qubits from one SLM trap to another consists of the following steps:

- 1) **Loading:** Qubits are loaded from static SLM traps into adjustable AOD traps.
- 2) **Moving:** The loaded qubits can be moved in the inter-atomic space without disturbing other atoms.
- 3) **Storing:** Finally, the qubits are transferred from AOD traps back to SLM traps.

If not required, the last step can also be skipped, and qubits can remain in adjustable AOD traps. After the required qubits are shuttled to the entangling zone and arranged

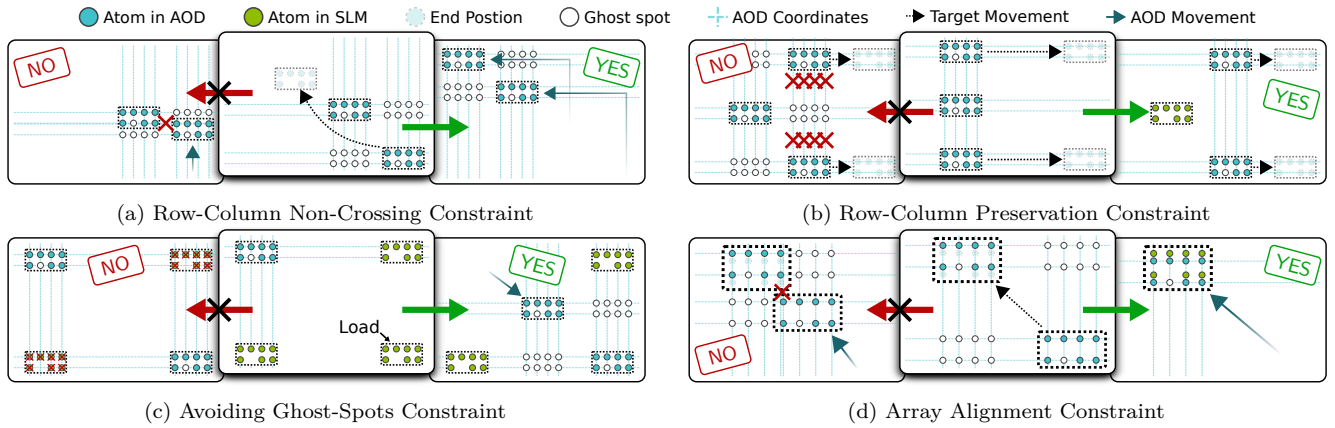


Figure 3. Illustration of the four constraints on shuttling operations.

such that pairs of atoms that are supposed to interact are located in proximity, the CZ gate is applied by switching on the Rydberg beam.

Example 4 (One Execution Cycle). *Figure 2 sketches the steps necessary to perform two parallel CZ gates between the qubit-pairs (1, 3) and (2, 4):*

- t = 1** *Qubit arrays 1 and 2 are transferred from the storage to the entangling zone in parallel. This includes loading, moving, and storing (the necessity of storing is discussed in Section III-C).*
- t = 2** *The next set of qubit arrays corresponding to qubits 3 and 4 are moved to the entangling zone in parallel.*
- t = 3** *The global Rydberg beam (yellow) performs entangling CZ gates between all qubit pairs of the arrays (1, 3) and (2, 4), respectively. Afterward, the previous steps can be applied in reverse order to move the atoms back to the storage zone.*

Similarly, that process can also transport qubits between the storage and readout zones. By consistently applying steps (1)–(3), it becomes possible to execute arbitrary sequences of transversal gates through array shuttling between zones and activation of laser beams.

C. Shuttling Constraints

However, an AOD cannot load and move the qubits arbitrarily; the shuttling of qubits must follow certain constraints [46]. As discussed in Section II-B, atoms can be confined in static SLM traps or dynamically adjustable AODs. Each AOD generates a series of laser beams (rows and columns) along the x- or y-axis at precise coordinates, capable of individual activation and deactivation. The intersections of these orthogonal beams form a grid that defines trap coordinates within the 2D plane, allowing for the manipulation of trap positions by modulating beam coordinates and consequently moving the trapped atoms. This method facilitates the trapping and simultaneous movement of a larger quantity of atoms, provided they adhere to the following four constraints, which later will be illustrated in Example 5:

- a) **Row-Column Non-Crossing Constraint:** AOD rows (columns) must maintain a minimum separation during movements and must not intersect with each

other. In particular, the order among rows (columns) must be preserved during movements.

- b) **Row-Column Preservation Constraint:** If atoms are in the same row (column), they remain in the same row (column) throughout a movement.
- c) **Avoiding Ghost-Spots Constraint:** Each intersection of AOD rows and columns generates a potential trap, resulting in a grid of traps. Some of those traps might be unintended, so-called *ghost spots* and can unintentionally trap or disturb other atoms.
- d) **Array Alignment Constraint:** To perform a transversal logical CZ gate between two arrays, they must overlap to bring the respective qubit pairs together. With both arrays loaded in the same AOD, this would violate Constraint (a). Therefore, one of the two arrays has to be placed back in SLM traps, which are not subject to the AOD constraints.

Example 5 (Shuttling Constraints). *Figure 3a depicts the case where the lower right array is supposed to move to the upper left corner. In order to make this possible, the middle atom array cannot remain at its position and has to be moved to the upper left as well to avoid AOD crossing. Alternatively, the middle array could be (temporarily) stored in SLM traps.*

Figure 3b illustrates the case where the middle array is supposed to stay at its position while the other two should move right. Before moving the other two arrays, the middle array must be stored in SLM traps.

Loading the lower right array shown in Figure 3c cannot be done directly without simultaneously loading the other two arrays. An additional offset movement is required to move the ghost spots into the inter-atomic space to load the array in a second step.

Figure 3d illustrates that at least one of the two arrays of a logical entangling gate has to be placed in SLM traps for the arrays to overlap. Otherwise, the AOD crossing constraint will be violated.

With the zoned structure and the constraints in place, the zoned neutral atom architecture model is complete. Now, the model can be used to formulate the routing problem on this architecture. For the special case of FTQC, we describe in the next section how the physical architecture can be abstracted into a logical model.




Figure 4. The physical architecture and the configuration (left), i.e., array size for the code, serve as the input. A logical architecture is derived from this and used to find a logical routing (middle). This logical routing is then translated into a logical array routing (right).

D. Towards Fault-Tolerant Quantum Computing

We want to pave the way to FTQC, supporting larger atom arrays that encode a logical qubit. The logical arrays must be shuttled between the zones to execute the different operations and establish the necessary connectivity for entangling gates. For the considered sequence of transversal gates, it suffices to compute the movements on the logical level, i.e., for the logical qubits representing an array of physical qubits. This is possible since all physical qubits of one logical array can be shuttled simultaneously.

Using a configuration that specifies the array size for the code, the physical architecture can be abstracted to a logical architecture. The physical architecture is parqueted with non-overlapping arrays of qubits, and only the left upper qubit of each array is taken as a reference for the logical qubit. After all movements necessary for the circuit execution have been computed, the logical qubits are replaced by their corresponding array of physical qubits. This process is illustrated in Figure 4.

IV. ROUTING OF LOGICAL ENTANGLING GATES FOR ZONED ARCHITECTURES

Shuttling qubits into proximity for entangling gates while simultaneously minimizing overhead through loading, storing, and shuttling constitutes a significant focus of ongoing research in neutral atom compilation. Different approaches have been proposed for different, non-zoned architectures [18]–[22], [47]. Other zone-based solutions [48]–[51] are targeted towards trapped ion hardware which has fundamentally different shuttling constraints [52], [53]. For neutral atoms, however, we first define the considered problem, present the straightforward, naive solution, and then introduce the proposed solution to the routing problem.

A. Considered Problem

The model proposed in the previous section can now be used to formulate the problem of *routing of logical, transversal entangling gates* that compilers have to solve for zoned neutral atom architectures to generate a sequence of execution steps. More precisely, given a sequence of entangling (2-qubit), transversal gates, the goal is to move the qubit arrays to the required zone for the subsequent operation with minimal *routing time overhead*—that is, the additional time required for the qubit movement—while adhering to the constraints outlined in Section III-C. According to Section III-D, it suffices to solve the routing problem on the logical level; hence, whenever we refer to qubits below, we mean logical qubits.

B. Naive Solution

A simple solution to generate an executable sequence of instructions avoiding AOD crossings is to execute one entangling operation at a time and shuttle the qubits to the entangling zone one by one. Idling qubits are kept in the storage zone to shield them from the interaction with the Rydberg beam. For each entangling operation on a pair of qubits, the naive solution moves them to the entangling zone, executes the entangling operation, and moves them back to the storage zone. In order to satisfy the AOD constraints, one of the qubits is moved first and stored in an SLM trap before the other one is loaded. The latter is moved close to the first to be in the interaction radius.

This naive solution serializes the execution of the entangling operations. Consequently, it does not exploit the potential parallelism of the neutral atom architecture.

C. Proposed Solution

Motivated by the shortcomings of the naive solution, this work introduces an alternative approach that tries to maximize the number of parallel entangling operations. This section outlines the general idea of the proposed solution while the details of its realization follow in Section V. We improve the naive solution by addressing two primary objectives:

- 1) **Parallel gate execution:** Execute as many independent entangling operations as possible per Rydberg beam to achieve high gate parallelism.
- 2) **Parallel AOD shuttling:** Load, store, or move as many qubits in parallel to minimize routing time overhead while fulfilling all AOD constraints from Section III-C.

We accomplish this by splitting the execution of the entangling gates into multiple *runs*. A run starts with shuttling a set of qubits into the entangling zone and ends with shuttling them back to the storage zone. Recall that having one qubit of each interacting pair in the SLM trap is necessary but also sufficient to comply with the array alignment constraint (Constraint (d)). This fact can be exploited to execute multiple entangling operations in one run without additional load or store operations. Therefore, a run is divided into *steps*, such that one step executes a set of independent entangling gates, where independent means that no two entangling gates share any qubits. The qubits held in the AOD move between the steps to their next interaction partner or resting position. This way, a two-fold increase in parallelism is achieved: First, by executing multiple entangling operations in parallel within one step, and second, by shuttling additional qubits in parallel into the entangling zone required for the steps in one run.




Figure 5. Comparison of sequential entangling operations (left) with proposed parallel execution and efficient logical qubit movement (right).

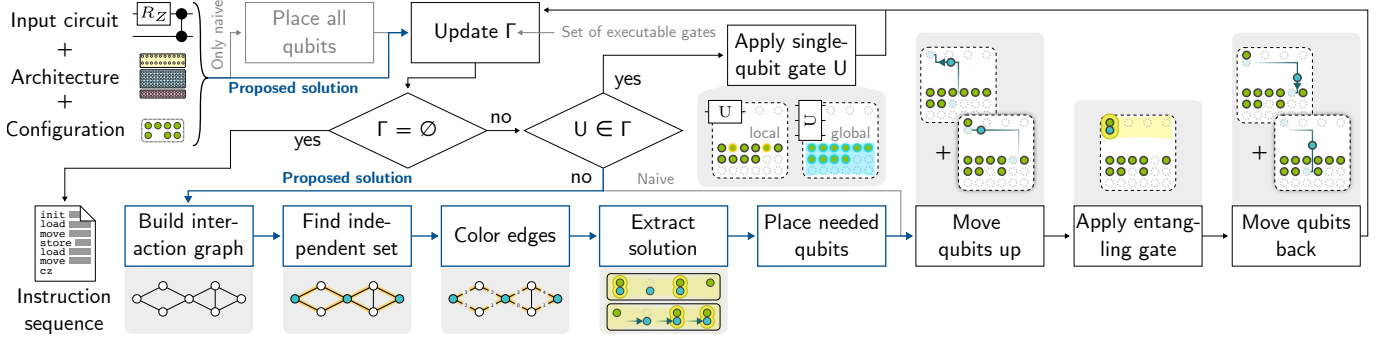


Figure 6. This flowchart illustrates the combination of the naive and the proposed solution. Where both solutions differ, the naive one is drawn in light grey and the proposed one in dark blue; common parts in black. All gates (CZ, single-qubit) are assumed to be transversal.

Example 6 (Naive vs. Proposed Solution). *Figure 5 (left) illustrates the naive solution, which executes only one entangling operation at a time. In contrast, Figure 5 (right) illustrates the proposed solution, which executes two entangling operations in parallel. The proposed solution has also placed further qubits into SLM traps in the entangling zone to execute more entangling operations without additional load or store operations. For the next step, the qubits move one position to the right to execute two more entangling operations. Afterward, all qubits will be moved back to the storage zone.*

Independent operations must be identified to realize this parallel execution of entangling gates. To this end, an interaction graph with qubits as nodes and edges between qubits that should interact is used and provides the core of the proposed approach. Based on this graph, we compute a maximal independent set that decides which qubits are placed in SLM and AOD traps. In order to obtain a feasible run regarding the AOD constraints, the edges of the interaction graph are colored such that the color determines the step in which the entangling operation represented by the edge is executed. Generally, the proposed approach proceeds as illustrated in Figure 6. The following section will explain each step in detail.

V. IMPLEMENTATION DETAILS

This section provides technical details of the proposed approach. We detail the construction of the interaction graph and how we obtain a schedule for the entangling operations from it by coloring its edges. Finally, we derive the exact positions of the qubits for each step.

A. Interaction Graph

The interaction graph is an undirected graph, where the nodes represent qubits that are involved in an entangling operation, and an edge connects two nodes if the corresponding qubits share one entangling operation that is executable at this moment. An entangling operation is executable if all operations in front are either already executed or commute with the entangling operation.

Example 7 (Interaction Graph). *On the left, Figure 7 shows part of a quantum circuit acting on eight (logical) qubits. All hidden gates to the left have already been executed, and all currently executable gates are highlighted in orange. Note that two of the three CZ gates to the right of the layer*

of single-qubit gates are executable because no single-qubit gate acts on their qubits. On the right, the figure shows the resulting circuit's interaction graph.

B. Independent Set

For every entangling gate, i.e., edge, that we execute, one of the adjacent qubits needs to be in an SLM trap and the other in an AOD trap. While trying to maximize the number of entangling gates, this leads to a max-cut problem. In general, a cut of a graph is a partition of the nodes into two sets. Accordingly, a max-cut is a cut that maximizes the number of edges between the two sets.

In our case, one set of nodes in the resulting cut are the qubits held in SLM traps, and the other set are qubits in AOD traps. The problem of finding a max-cut is NP-hard, and we use a maximal independent set instead to find a sufficiently good solution.

An independent set of nodes in a graph is a set of nodes that are not adjacent to each other. An independent set is maximal if it cannot be extended by adding another node without violating the independence property. The independent set then constitutes one set of the cut, and the remaining nodes the other set. Such a maximal independent set can be calculated by iterating over the nodes in the interaction graph and adding one node at a time to the set if it is not adjacent to any of the already selected nodes. To maximize the number of edges in the resulting cut, we order the nodes by their degree in descending order before iterating over them. All nodes ending up in the independent set correspond to the qubits kept in the AOD.




Figure 7. Part of a quantum circuit with its currently executable gates highlighted in orange and its corresponding interaction graph on the right.




Figure 8. The interaction graph from Example 7 with the nodes of the independent set colored in blue. All edges covered by this independent set are highlighted in orange.


The edges covered by the independent set, i. e., those that are adjacent to one node in the independent set, are the ones that can be executed in one run.

Example 8 (Independent Set). Consider again the interaction graph discussed in Example 7. Here, the node with the highest degree is the middle one and is selected first. The two nodes with degree three are not selected since they are adjacent to the previously selected one. Finally, the algorithm selects the two nodes with degree two that are not adjacent to the selected node. The resulting independent set is shown in Figure 8, where the selected nodes are colored in blue. Those are the ones that will be held in the AOD, whereas the green ones will be placed in SLM traps.

C. Color Edges

To determine the order in which the entangling operations are executed, we color the covered edges of the interaction graph. A color is a non-negative integer that is assigned to an edge and corresponds to the time step in which the entangling operation is executed. Consequently, the number of required colors corresponds to the number of time steps needed to execute all entangling operations. Two adjacent edges, i. e., edges that share a node, cannot have the same color. However, another constraint is necessary for a valid coloring to prevent AOD crossings.

Example 9 (Order Preservation of AOD Qubits). Consider four qubits where each pair of (q_1, q_2) , (q_2, q_3) , (q_3, q_4) , and (q_4, q_1) should undergo a CZ operation. The upper time sequence in the figure below is impossible because the two AOD qubits would need to cross each other, conflicting with the AOD Constraint (a) from Section III-C. As shown in the lower time sequence, the AOD qubits must maintain their order to prevent crossings, which makes three steps necessary to execute those four CZ operations.



In terms of the coloring of the interaction graph's edges, this can be achieved by adding the following condition:

- (C) On all paths of length two between two blue nodes, the inequality relation between the assigned colors must point towards the same (blue) end node.

Example 10 (Coloring Constraint). The two execution sequences above correspond to the two possible colorings below. The left coloring violates Condition (C) whereas the right one satisfies it.




Figure 9. The coloring requires eight iterations, and the result is shown in the bottom row. The top row shows selected intermediate steps. The “<”-signs denote the extra constraints on the coloring to be valid.



To satisfy this additional constraint while keeping the total number of colors (i. e., time steps) low, we employ a variation of the DSatur algorithm [54] that maintains its polynomial time complexity. The DSatur algorithm is a greedy algorithm that iteratively selects the edge with the most different adjacent colors and selects the least admissible color for it. The least admissible color is the smallest color that is different from all adjacent colors of the edge and larger than any color of an adjacent edge that does not share the same node in the independent set. This ensures that the Condition (C) is satisfied. We modify the DSatur algorithm accordingly as defined in Algorithm 1. The following Example 11 illustrates its application.

Algorithm 1: Modified DSatur Algorithm

Input: Interaction graph $G = (V, E)$ (undirected)
Independent Set $I \subseteq V$
Output: Valid coloring $c : E \rightarrow \mathbb{N}$ of the edges

- 1 **for** $e \in E$ **do**
- 2 $e.color \leftarrow \text{undef.}$
- 3 **for** $v \in I$ sorted by degree in descending order **do**
- 4 $S \leftarrow$ set of edges adjacent to v
- 5 **for** $e \in S$ lexicographically ordered by the number of different adjacent colors (desc.) and degree (desc.) **do**
- 6 $e.color \leftarrow$ least admissible color
 ▷ must be different from all adjacent colors and larger than any color of an adjacent edge not adjacent to v
- 7 $S \leftarrow S \setminus \{e\}$
- 8 **return** $c : e \mapsto e.color$

Example 11 (Coloring Edges). Algorithm 1 starts with the node of the independent set with the highest degree and colors all its adjacent edges. Since, at this stage, no other edges are colored, they are colored with 1, 2, 3, and 4, respectively (see steps $t = 1$ and 4 in Figure 9). The two remaining nodes in the independent set have the same degree; we assume that it is the left one's turn next. For the edge colored in Step 5, the algorithm selects Color 2, the least color larger than the adjacent 1. The algorithm continues coloring the remaining edges, which results in the coloring shown in the bottom row in Figure 9.




Figure 10. Visualization of all time steps corresponding to the run stemming from the interaction graph in Example 7.




Figure 11. The grey arrows depict a directed graph that represents a partial order on the set of SLM qubits. The order of the qubits in which they are placed in the entangling zone, from left to right, is given by a topological ordering of the graph.

With the coloring at hand, we can now calculate the exact positions of each qubit in every step and their movements in between.

D. Positioning the Qubits

Recall that the independent set determines which qubits are held in AOD traps, and the coloring determines the steps at which an entangling gate is executed. We still need to determine the exact positions of the qubits in the entangling zone for each step to facilitate the correct entangling operation at the right time. In the following, we will refer to qubits held in the AOD as AOD qubits and to those kept in SLM traps as SLM qubits.

The AOD qubits will move from left to right during one run. Hence, the SLM qubit adjacent to one AOD qubit must be placed in ascending order with respect to the color assigned to the edge from the AOD to the SLM qubit. This defines a total order on the SLM qubits in the neighborhood of one AOD qubit. Combining the total orders induced by all AOD qubits gives a partial order on the set of all SLM qubits.¹ We compute a total order on the SLM qubits by topologically sorting the directed graph induced by the partial order. Thereby, the computed total order satisfies the partial order in the sense that the relation is a superset of the partial order.


Example 12 (Topological Order). *The colors around q_7 in Figure 11 induce an order on the SLM qubits. The order is depicted with grey arrows. The orderings induced by q_1 and q_3 are already subsumed by the existing order (also see*

¹It is not self-evident that the directed graph induced by combining all total orders around an SLM qubit is acyclic, which is a requirement for the partial order to be well-defined. This limitation must already be taken care of during coloring: The inner for-loop actually orders the edges first by the partial order induced by the already colored edges and then by the number of different adjacent colors and their degree. Then, the function `leastAdmissibleColor` ensures that the new color does not introduce any cycles.

Footnote 1). Following the topological order of the SLM qubits, we place the qubits in the entangling zone from left to right. Consequently, we place from left to right the qubits q_5 , q_6 , q_2 , and q_4 . This way, q_7 can first interact with q_5 in time step $t = 0$ and afterward with the other SLM qubits by moving one position right between the steps.

The order of the adjustable qubits in the AOD must be the same as the one used for the coloring, whereby the AOD qubit with the highest degree is placed rightmost. By keeping an equal and sufficient distance between the SLM qubits to avoid any undesired interaction, we could, in principle, calculate the exact positions of all qubits in the entangling zone for each step. However, there are time steps in which some AOD qubits shall not interact with any SLM qubit. In this case, we need to ensure with so-called *resting positions* that these AOD qubits are outside any other qubit's interaction radius.

Example 13 (Resting Positions). *According to the coloring in Figure 9, the qubit q_1 does not interact with any other qubit during time step $t = 1$. However, the order of the AOD qubits used for the coloring implies that q_1 is located in between q_7 and q_3 . In time step $t = 1$ those interact with q_5 and q_6 , respectively. To avoid any interaction, we need to create a resting position for q_1 between q_5 and q_6 . Those resting positions can be reused later for other qubits.*



These resting positions shift all subsequent SLM qubits by one position to the right. After this, the positions of the SLM qubits will be known. Relative to those, the AOD qubits' positions in each time step are calculated according to the coloring.

Example 14 (Complete Run). *As seen in previous examples, in particular, Example 7, q_7 interacts with four qubits. Those are not part of the independent set and, hence, are placed in SLM traps. From one time step to the next, the AOD qubits are moved from left to right to match the position of their interacting partner. In time step $t = 1$, q_1 stays at the resting position between qubit q_5 and q_6 . This resting position will be reused for qubit q_3 in the next step. All resulting steps of this run are shown in Figure 10.*

E. Finalizing the Solution

We are now able to perform one run of multiple entangling operations without additional load and store




Figure 12. Comparison of the execution time of the circuit of the naive approach and NALAC.

operations. The only step missing is moving the affected qubits from the storage zone to the entangling zone and moving them back after the run.

To move the qubits to the entangling zone and store them in the previously computed order, we opt for a procedure that tries to maximize the number of qubits that can be loaded in parallel. Unlike the naive solution, the proposed solution does not place the qubits from the start. Instead, it follows a demand-driven approach: When a qubit is loaded into an AOD, its position can either be defined due to a previous run or be undefined. In the latter case, we find a free spot in the storage zone to maximize the number of qubits that can be loaded in parallel. Qubits can be loaded in parallel if they are in the same row and in the correct order, i.e., the order in which they are needed in the entangling zone. Hence, the proposed solution finds the next free spot in the storage zone for a qubit with an undefined position satisfying both conditions.² To keep the space consumed horizontally outside of the SLM traps in the storage zone small, we first load qubits into the AOD with a large misplacement, meaning they are located far left and must travel far right or vice versa. After a run has finished, the used qubits are moved back to the storage zone. For this, we look for a minimal number of rows in the storage that fit all qubits, and the free spots in those rows are filled with qubits from the entangling zone.

Finally, as described in Section III-D, the solution for the logical qubit is translated to a solution for arrays of atoms, each representing a logical qubit. Therefore, every qubit in the logical solution is replaced by its respective array of atoms, where the logical qubit determines the upper left atom in the array.

VI. EVALUATION

In order to evaluate the effectiveness of the proposed solution compared to the naive approach, both techniques have been implemented on top of the tool *MQT QMAP* publicly available at <https://github.com/cda-tum/mqt-qmap>. The resulting tool is called NALAC, short for Neutral Atom Logical Array Compiler.

As a basis for our experiments, we used the architecture as illustrated in Figure 13. To compute the different routing overheads produced by both approaches, we used the following parameters based on the specifications of the neutral atom architecture in [11]:

Load/Store	Duration [μs]	20
Shuttling	Speed [$\frac{\mu\text{m}}{\mu\text{s}}$]	0.55
CZ Gate	Duration [μs]	0.2

²We only execute circuits with as many qubits as the storage zone can fit. Hence, there will always be a free spot.

As input circuits, we use preprocessed versions of circuits taken from the MQT Bench library [55] grouped into circuit families, such as qaoa, vqe, qft, etc. The pre-processing consists of translating the original circuits to the gate set supported

by the neutral atom architecture, i.e., local as well as global rotations, and CZ gates, using Qiskit [56] and a custom Python script. All experiments were conducted on a machine with an Apple M3 chip and 16 GB of RAM.

A. Comparison of the Naive Solution to NALAC

To demonstrate the advantage of NALAC in terms of the two objectives discussed in Section IV-C, namely 1) improved parallel gate execution, and 2) improved parallel AOD shuttling, we evaluated the time taken for loading, shuttling, and storing when executing the resulting circuits—short the *routing time overhead*. Figure 12 shows the obtained loading (including storing) and shuttling times separately for four different circuit families in various sizes. As one can see, NALAC produces circuits with significantly shorter loading and shuttling times compared to the naive solution. Table I summarizes the results for those four and additional circuit families with a fixed number of 20 logical qubits. The last column shows the average number of entangling operations performed in parallel per Rydberg beam for circuits compiled with NALAC. Note that the corresponding number for the naive solution is always one. The low numbers for the average parallel CZ gates for some




Figure 13. Layout of the zoned architecture used for the experiments.

Table I
NAIVE SOLUTION VS. NALAC

Circuit (20 Qubits)	Num. CZ-Gates	Routing Overhead [ms]		∅ Parallel CZ-Gates (NALAC)
		Naive	NALAC	
ae	380	130	32	1.1
dj	19	7	1	1.0
ghz	19	6	7	1.0
graphstate	20	7	1	3.3
portfoliovqe	570	195	31	1.0
qft	408	139	22	1.0
qftentangled	429	147	31	2.8
qnn	778	266	57	3.6
qpeexact	406	139	43	3.7
qpeinexact	407	139	43	3.7
realamprandom	570	195	27	1.2
su2random	570	195	27	1.2
twolocalrandom	570	195	27	1.2
wstate	38	13	7	1.0

circuits are not surprising, as the CZ gates in these circuits are interleaved with one-qubit gates, which prevent the parallel execution of many CZ gates. In general, it can be said that the more parallelism the structure of a circuit allows, the more NALAC can exploit its advantage over the serial naive approach. In the ghz circuit, for example, NALAC even has a non-significant larger routing overhead because—due to the structure of the circuit—the execution of the CZ gates is interrupted by one-qubit gates, which forces the CZ gates to be executed serially. Overall, those figures demonstrate that NALAC is superior to the naive solution because it yields substantially less routing time overhead and higher gate parallelism.

B. Comparison of Compilation Times

The improved execution times of NALAC come at a price: The proposed approach is more complex than the naive one and requires a longer compilation time. To evaluate the cost of the improvement, we compared the compilation time of both approaches. Figure 14 shows the compilation time for the naive solution and NALAC for circuits of the `qftentangled` family with different numbers of qubits. As expected, the runtime of NALAC is significantly higher than that of the naive solution. However, the performance remains polynomial; even for larger circuits, the runtime of NALAC is still within the range of milliseconds, making it a suitable approach even for larger numbers of qubits.





Figure 14. Compilation time of the naive approach and NALAC.


C. Influence of Array Size

The experiments above were all performed with a logical array size of one, i. e., one atom represents one logical qubit. As argued in Section III-D, NALAC also supports larger atom arrays representing a logical qubit. To evaluate that, also the impact of the array size on the routing overhead is considered, shown in Figure 15b. As can be seen, the qubits need to move further with larger array sizes, increasing the shuttling times. However, the loading times remain roughly the same since all qubits in one array can be loaded in parallel. The only effect that can happen here is that fewer logical qubits fit in one row, which requires more loading steps and consequently slightly increases the loading time.



(a) The size of the arrays for different color codes. (b) Impact of increasing array sizes on the routing overhead.

Figure 15. Comparing different array sizes.



(a) Architecture with a narrow storage zone. (b) Impact of a wide and a narrow storage zone on the routing overhead.

Figure 16. Comparing different shapes of the storage zone.

D. Supporting Hardware Design

Finally, in a last series of experiments, we investigated how the proposed approach can even be used to evaluate different hardware designs. To this end, we exemplarily considered another architecture, namely the one shown in Figure 16a, which has a narrower storage zone than the one from Figure 13 considered before. Figure 16b shows the corresponding resulting routing time overhead—clearly showcasing the impact of this design. In fact, the results show that the narrow storage zone leads to both longer loading and shuttling times. This can be explained by the fact that (1) the qubits need to travel longer distances on average to reach the entangling zone, and (2) fewer qubits fit in one row in the storage zone. This demonstrates how the proposed approach can help hardware designers to analyze trade-offs between different hardware parameters and settings.

VII. CONCLUSIONS

In this work, we considered the development of target-specific quantum compilers specialized for zoned neutral atom architectures. We first provided an abstract model of zoned neutral atom architectures, which are a promising candidate for fault-tolerant gate execution and, therefore, general FTQC. Then, we proposed a novel solution to the routing problem of entangling operations to this architecture. This solution minimizes the time required for loading, shuttling, and storing the qubits while maximizing the gate parallelism of entangling gates. We implemented the proposed solution in the tool NALAC and compared it to a naive solution. Evaluations showed that NALAC efficiently routes entangling operations of even larger quantum circuits to the zoned neutral atom architecture, applies to various array sizes, and can even support designers evaluating different hardware designs. NALAC is publicly available in open-source as part of the *Munich Quantum Toolkit* (MQT, [23]) at <https://github.com/cda-tum/mqt-qmap>. For full FTQC, this tool needs to be extended with further features, like magic state injection [30] and error correction cycles [27], [28].

Acknowledgements

We want to thank Johannes Zeiher for providing us with helpful comments on a draft of this article and the anonymous reviewers for their helpful feedback.

This work received funding from the European Research Council (ERC) under the European Union’s Horizon 2020

research and innovation program (grant agreement No. 101001318), was part of the Munich Quantum Valley, which the Bavarian state government supports with funds from the Hightech Agenda Bayern Plus, and has been supported by the BMWK based on a decision by the German Bundestag through project QuaST, as well as by the BMK, BMDW, and the State of Upper Austria in the frame of the COMET program (managed by the FFG).

REFERENCES

- [1] John Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018. DOI: 10.22331/q-2018-08-06-79.
- [2] Craig Gidney and Martin Ekerå, “How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits,” *Quantum*, vol. 5, p. 433, 2021. DOI: 10.22331/q-2021-04-15-433.
- [3] Rajeev Acharya *et al.*, “Suppressing quantum errors by scaling a surface code logical qubit,” *Nature*, vol. 614, no. 7949, pp. 676–681, 2023. DOI: 10.1038/s41586-022-05434-1.
- [4] R. Barends *et al.*, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, pp. 500–503, 2014. DOI: 10.1038/nature13171.
- [5] Maika Takita, Andrew W. Cross, A. D. Córcoles, Jerry M. Chow, and Jay M. Gambetta, “Experimental Demonstration of Fault-Tolerant State Preparation with Superconducting Qubits,” *Phys. Rev. Lett.*, vol. 119, no. 18, p. 180501, 2017. DOI: 10.1103/PhysRevLett.119.180501.
- [6] Youwei Zhao *et al.*, “Realization of an Error-Correcting Surface Code with Superconducting Qubits,” *Phys. Rev. Lett.*, vol. 129, no. 3, p. 030501, 2022. DOI: 10.1103/PhysRevLett.129.030501.
- [7] C. Ryan-Anderson *et al.*, *Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code*, 2022. arXiv: 2208.01863.
- [8] Laird Egan *et al.*, “Fault-tolerant control of an error-corrected qubit,” *Nature*, vol. 598, no. 7880, pp. 281–286, 2021. DOI: 10.1038/s41586-021-03928-y.
- [9] Lukas Postler *et al.*, “Demonstration of fault-tolerant universal quantum gate operations,” *Nature*, vol. 605, no. 7911, pp. 675–680, 2022. DOI: 10.1038/s41586-022-04721-1.
- [10] Ivan Pogorelov *et al.*, *Experimental fault-tolerant code switching*, 2024. arXiv: 2403.13732.
- [11] Dolev Bluvstein *et al.*, “Logical quantum processor based on reconfigurable atom arrays,” *Nature*, pp. 1–3, 2023. DOI: 10.1038/s41586-023-06927-3.
- [12] Lukas Burgholzer, “Design Automation Tools and Software for Quantum Computing,” Ph.D. dissertation, Johannes Kepler University Linz, 2023.
- [13] Loïc Henriët *et al.*, “Quantum computing with neutral atoms,” *Quantum*, vol. 4, p. 327, 2020. DOI: 10.22331/q-2020-09-21-327.
- [14] M. Saffman, T. G. Walker, and K. Mølmer, “Quantum information with rydberg atoms,” *Rev. Mod. Phys.*, vol. 82, pp. 2313–2363, 3 2010. DOI: 10.1103/RevModPhys.82.2313.
- [15] Mark Saffman, “Quantum computing with neutral atoms,” *National Science Review*, vol. 6, no. 1, pp. 24–25, 2019. DOI: 10.1093/nsr/nwy088.
- [16] T. M. Graham *et al.*, “Multi-qubit entanglement and algorithms on a neutral-atom quantum computer,” *Nature*, vol. 604, no. 7906, pp. 457–462, 2022. DOI: 10.1038/s41586-022-04603-6.
- [17] Ludwig Schmid *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. 033001, 2024. DOI: 10.1088/2058-9565/ad33ac.
- [18] Daniel Bochen Tan, Dolev Bluvstein, Mikhail D. Lukin, and Jason Cong, “Compiling Quantum Circuits for Dynamically Field-Programmable Neutral Atoms Array Processors,” *Quantum*, vol. 8, p. 1281, 2024. DOI: 10.22331/q-2024-03-14-1281.
- [19] Hanrui Wang *et al.*, “Atomique: A quantum compiler for reconfigurable neutral atom arrays,” 2024. arXiv: 2311.15123.
- [20] Hanrui Wang *et al.*, “Q-Pilot: Field Programmable Quantum Array Compilation with Flying Ancillas,” in *Design Automation Conf.*, 2024. arXiv: 2311.16190.
- [21] Natalia Nottingham *et al.*, *Decomposing and Routing Quantum Circuits Under Constraints for Neutral Atom Architectures*, 2023. arXiv: 2307.14996.
- [22] Ludwig Schmid, Sunghye Park, Seokhyeong Kang, and Robert Wille, “Hybrid Circuit Mapping: Leveraging the Full Spectrum of Computational Capabilities of Neutral Atom Quantum Computers,” in *Design Automation Conf.*, 2024. arXiv: 2311.14164.
- [23] Robert Wille *et al.*, “The MQT Handbook: A Summary of Design Automation Tools and Software for Quantum Computing,” in *Int’l Conf. on Quantum Software*, 2024. arXiv: 2405.17543, A live version of this document is available at <https://mqt.readthedocs.io>.
- [24] Lukas Burgholzer, Alexander Ploier, and Robert Wille, “Simulation paths for quantum circuit simulation with decision diagrams what to learn from tensor networks, and what not,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 42, no. 4, pp. 1113–1122, 2023. DOI: 10.1109/TCAD.2022.3197969.
- [25] Tom Peham, Lukas Burgholzer, and Robert Wille, “Equivalence checking of parameterized quantum circuits: Verifying the compilation of variational quantum algorithms,” in *Asia and South Pacific Design Automation Conf.*, Association for Computing Machinery, 2023, pp. 702–708. DOI: 10.1145/3566097.3567932.
- [26] N. Quetschlich, L. Burgholzer, and R. Wille, “MQT Predictor: Automatic Device Selection with Device-Specific Circuit Compilation for Quantum Computing,” *ACM Transactions on Quantum Computing*, 2024. arXiv: 2310.06889.
- [27] P.W. Shor, “Fault-tolerant quantum computation,” in *Conference on Foundations of Computer Science*, 1996, pp. 56–65. DOI: 10.1109/SFCS.1996.548464.
- [28] Daniel Gottesman, *Stabilizer Codes and Quantum Error Correction*, 1997. arXiv: quant-ph/9705052.

- [29] Héctor Bombín, “Gauge color codes: Optimal transversal gates and gauge fixing in topological stabilizer codes,” *New Journal of Physics*, vol. 17, no. 8, p. 083002, 2015. DOI: 10.1088/1367-2630/17/8/083002.
- [30] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A*, vol. 86, no. 3, p. 032324, 2012. DOI: 10.1103/PhysRevA.86.032324.
- [31] Bryan Eastin and Emanuel Knill, “Restrictions on Transversal Encoded Quantum Gate Sets,” *Phys. Rev. Lett.*, vol. 102, no. 11, p. 110502, 2009. DOI: 10.1103/PhysRevLett.102.110502.
- [32] Hannah J. Manetsch *et al.*, *A tweezer array with 6100 highly coherent atomic qubits*, 2024. arXiv: 2403.12021.
- [33] Daniel Barredo, Sylvain de Léséleuc, Vincent Lienhard, Thierry Lahaye, and Antoine Browaeys, “An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays,” *Science*, vol. 354, no. 6315, pp. 1021–1023, 2016. DOI: 10.1126/science.aah3778.
- [34] Lars Pause *et al.*, “Supercharged two-dimensional tweezer array with more than 1000 atomic qubits,” *Optica*, vol. 11, no. 2, pp. 222–226, 2024. DOI: 10.1364/OPTICA.513551.
- [35] Flavien Gyger *et al.*, “Continuous operation of large-scale atom arrays in optical lattices,” *Phys. Rev. Res.*, vol. 6, p. 033104, 3 2024. DOI: 10.1103/PhysRevResearch.6.033104.
- [36] M. A. Norcia *et al.*, “Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices,” *PRX Quantum*, vol. 5, p. 030316, 3 2024. DOI: 10.1103/PRXQuantum.5.030316.
- [37] Guoqing Wang, Wenchao Xu, Changhao Li, Vladan Vuletić, and Paola Cappellaro, *Individual-atom control in array through phase modulation*, 2023. arXiv: 2310.19741.
- [38] Harry Levine *et al.*, “Dispersive optical systems for scalable Raman driving of hyperfine qubits,” *Phys. Rev. A*, vol. 105, no. 3, p. 032618, 2022. DOI: 10.1103/PhysRevA.105.032618.
- [39] M. Müller, I. Lesanovsky, H. Weimer, H. P. Büchler, and P. Zoller, “Mesoscopic Rydberg Gate Based on Electromagnetically Induced Transparency,” *Phys. Rev. Lett.*, vol. 102, no. 17, p. 170502, 2009. DOI: 10.1103/PhysRevLett.102.170502.
- [40] L. Isenhower *et al.*, “Demonstration of a Neutral Atom Controlled-NOT Quantum Gate,” *Phys. Rev. Lett.*, vol. 104, no. 1, p. 010503, 2010. DOI: 10.1103/PhysRevLett.104.010503.
- [41] Harry Levine *et al.*, “Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms,” *Phys. Rev. Lett.*, vol. 123, no. 17, p. 170503, 2019. DOI: 10.1103/PhysRevLett.123.170503.
- [42] Simon J. Evered *et al.*, “High-fidelity parallel entangling gates on a neutral-atom quantum computer,” *Nature*, vol. 622, no. 7982, pp. 268–272, 2023. DOI: 10.1038/s41586-023-06481-y.
- [43] T. M. Graham *et al.*, “Midcircuit measurements on a single-species neutral alkali atom quantum processor,” *Phys. Rev. X*, vol. 13, p. 041051, 4 2023. DOI: 10.1103/PhysRevX.13.041051.
- [44] M. A. Norcia *et al.*, “Midcircuit Qubit Measurement and Rearrangement in a ^{171}Yb Atomic Array,” *Phys. Rev. X*, vol. 13, no. 4, p. 041034, 2023. DOI: 10.1103/PhysRevX.13.041034.
- [45] Emma Deist *et al.*, “Mid-Circuit Cavity Measurement in a Neutral Atom Array,” *Phys. Rev. Lett.*, vol. 129, no. 20, p. 203602, 2022. DOI: 10.1103/PhysRevLett.129.203602.
- [46] Dolev Bluvstein *et al.*, “A quantum processor based on coherent transport of entangled atom arrays,” *Nature*, vol. 604, no. 7906, pp. 451–456, 2022. DOI: 10.1038/s41586-022-04592-6.
- [47] Sebastian Brandhofer, Ilia Polian, and Hans Peter Büchler, “Optimal Mapping for Near-Term Quantum Architectures based on Rydberg Atoms,” in *Int’l Conf. on CAD*, 2021, pp. 1–7. DOI: 10.1109/ICCAD51958.2021.9643490.
- [48] Daniel Schoenberger, Stefan Hillmich, Matthias Brandl, and Robert Wille, “Shuttling for Scalable Trapped-Ion Quantum Computers,” in *Design, Automation and Test in Europe*, 2024. arXiv: 2402.14065.
- [49] Fabian Kreppel *et al.*, “Quantum Circuit Compiler for a Shuttling-Based Trapped-Ion Quantum Computer,” *Quantum*, vol. 7, p. 1176, 2023. DOI: 10.22331/q-2023-11-08-1176.
- [50] Daniel Schoenberger, Stefan Hillmich, Matthias Brandl, and Robert Wille, “Using Boolean Satisfiability for Exact Shuttling in Trapped-Ion Quantum Computers,” in *Asia and South Pacific Design Automation Conf.*, 2024, pp. 127–133. DOI: 10.1109/ASP-DAC58780.2024.10473902.
- [51] Mark Webber, Steven Herbert, Sebastian Weidt, and Winfried K. Hensinger, “Efficient Qubit Routing for a Globally Connected Trapped Ion Quantum Computer,” *Advanced Quantum Technologies*, vol. 3, no. 8, p. 2000027, 2020. DOI: 10.1002/qute.202000027.
- [52] Winfried K. Hensinger, “Quantum computer based on shuttling trapped ions,” *Nature*, vol. 592, no. 7853, pp. 190–191, 2021. DOI: 10.1038/d41586-021-00844-z.
- [53] Prakash Murali, Dripto M. Debroy, Kenneth R. Brown, and Margaret Martonosi, “Architecting Noisy Intermediate-Scale Trapped Ion Quantum Computers,” in *Int’l Symposium on Computer Architecture*, 2020, pp. 529–542. DOI: 10.1109/ISCA45697.2020.00051.
- [54] R. M. R. Lewis, “Bounds and Constructive Heuristics,” in *Guide to Graph Colouring*. Springer International Publishing, 2021, pp. 39–76. DOI: 10.1007/978-3-030-81054-2_3.
- [55] Nils Quetschlich, Lukas Burgholzer, and Robert Wille, “MQT Bench: Benchmarking Software and Design Automation Tools for Quantum Computing,” *Quantum*, vol. 7, p. 1062, 2023. DOI: 10.22331/q-2023-07-20-1062.
- [56] Qiskit contributors, *Qiskit: An open-source framework for quantum computing*, 2023. DOI: 10.5281/zenodo.2573505.