

# TILT: Achieving Higher Fidelity on a Trapped-Ion Linear-Tape Quantum Computing Architecture

Xin-Chuan Wu  
Department of Computer Science  
University of Chicago  
Chicago, USA  
xinchuan@uchicago.edu

Dripto M. Debroy  
Department of Physics  
Duke University  
Durham, USA  
dripto@phy.duke.edu

Yongshan Ding  
Department of Computer Science  
University of Chicago  
Chicago, USA  
yongshan@uchicago.edu

Jonathan M. Baker  
Department of Computer Science  
University of Chicago  
Chicago, USA  
jmbaker@uchicago.edu

Yuri Alexeev  
Computational Science Division  
Argonne National Laboratory  
Lemont, USA  
yuri@anl.gov

Kenneth R. Brown  
Department of Electrical  
and Computer Engineering  
Duke University  
Durham, USA  
kenneth.r.brown@duke.edu

Frederic T. Chong  
Department of Computer Science  
University of Chicago  
Chicago, USA  
chong@cs.uchicago.edu

**Abstract**—Trapped-ion qubits are a leading technology for practical quantum computing. In this work, we present an architectural analysis of a linear-tape architecture for trapped ions. In order to realize our study, we develop and evaluate mapping and scheduling algorithms for this architecture.

In particular, we introduce TILT, a linear “Turing-machine-like” architecture with a multilaser control “head,” where a linear chain of ions moves back and forth under the laser head. We find that TILT can substantially reduce communication as compared with comparable-sized Quantum Charge Coupled Device (QCCD) architectures. We also develop two important scheduling heuristics for TILT. The first heuristic reduces the number of swap operations by matching data traveling in opposite directions into an “opposing swap,” and also avoids the maximum swap distance across the width of the head, as maximum swap distances make scheduling multiple swaps in one head position difficult. The second heuristic minimizes ion chain motion by scheduling the tape to the position with the maximal executable operations for every movement. We provide application performance results from our simulation, which suggest that TILT can outperform QCCD in a range of NISQ applications in terms of success rate (up to 4.35x and 1.95x on average). We also discuss using TILT as a building block to extend existing scalable trapped-ion quantum computing proposals.

**Index Terms**—Quantum Computing, Trapped-ion Architecture, Circuit Optimization

## I. INTRODUCTION

Quantum computing (QC) aims to solve certain computational problems beyond the capabilities of even the largest classical high-performance computers. By leveraging the quantum mechanical principles of superposition and entanglement, QC algorithms have the potential to revolutionize areas

such as machine learning [11], quantum chemistry [43,69], and cryptography [76]. Recently, IBM, Rigetti, and Google demonstrated their QC devices up to 72 superconducting transmon qubits [29,38,72], while IonQ and Honeywell have recently made significant steps with trapped-ion QC devices [39,70]. Current QC machines in the Noisy Intermediate-Scale Quantum (NISQ) era [71] are too small for large benchmarks and unable to support quantum error correction (QEC) [8,30]. We can, however, run on the order of hundreds of quantum operations using on the order of hundreds of quantum bits (qubits).

Trapped-ion technologies are among the most promising systems for scalable QC for many reasons. While most current technologies have limited hardware qubit connectivity, two-qubit gates on trapped-ion computers can be executed on arbitrary pairs of qubits [32]. In a trapped-ion quantum computer, the internal states of the ions form the qubit states and quantum gates are implemented through laser-based operations [18,32,52,87], where acousto-optic modulators (AOMs) apply lasers with carefully modulated amplitudes and frequencies in order to generate different quantum gates. In many trapped-ion technologies, the ions themselves must be physically moved, for example, shuttled across the linear trap storing all ions. Previous studies have shown that this transportation in a linear array can be done with minimal energy gain and without the loss of qubit coherence [7,12,35,37,44,73,79].

To scale the QC device, we can add ions to the chain, without fear of frequency collision since all ions are identical. While adding ions is relatively simple, controlling all the

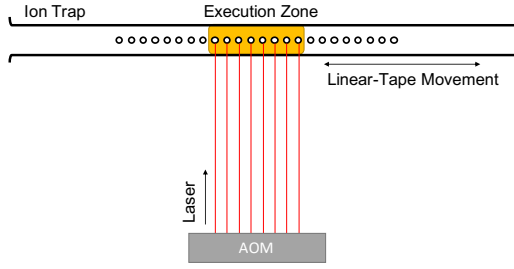


Fig. 1. A trapped-ion linear-tape quantum computing architecture. Acousto-optic modulators (AOMs) target laser beams for quantum operations, which can be applied to ions in the execution zone. In order to perform gate operations on the other qubits, the entire ion chain is translated until the target qubit is moved into the execution zone.

ions simultaneously is challenging. Recently, a linear-tape architecture was suggested by Monroe [62], where the device had a fixed set of lasers (often much smaller than the total number of ions in the trap) used to operate all qubits. While the lasers are fixed in place, the ions can be shuttled across the tape to allow them to be operated on. Figure 1 illustrates this concept, where the *tape* is a string of ions (each representing a qubit) arranged in a single linear trap and the *tape head* is the set of control lasers used to manipulate the qubits. The location of the tape head designates the *execution zone*, or the set of locations where qubits must be located in order for gates to be executed. If a qubit is not located in the execution zone, before applying a gate we must physically move the ions until they are aligned with the tape head. Since the number of ions in the execution zone is limited at one time, the entire ion chain will be moved back and forth during the execution of quantum circuits, similar to a Turing machine [34]. While the linear-tape model is inefficient in classical computing, our evaluation shows that a quantum version, with modest parallelism in the tape head, is competitive with contemporary 2D quantum computing designs. The key is that the qubits under the tape head form a completely connected graph and that this powerful communication structure can slide across the entire tape. Shuttling a small chain of ions has been demonstrated in [26], and larger tape-like demonstrations are planned [15].

By sharing laser controls, a trapped-ion linear-tape (TILT) architecture has easier calibration and simplified optics, as well as reduced hardware costs. In a TILT system, gates are implemented only on the qubits within the execution zone, and the entire chain shifts back and forth to operate all qubits. Such a machine does not require the difficult shuttling primitives of a quantum charge-coupled device (QCCD) architecture [45], since full chain shuttling is not nearly as difficult as split/merge operations or shuttling over junctions. Moreover, since the ions in the center of a trap are more evenly spaced (which can be accentuated by trap design), such an architecture has fewer issues with individual addressing and laser pointing errors. Consequently, gate fidelities would improve in a TILT architecture, and control optics would be required only for a

much smaller region of the chain, resulting in reduced control complexity and cost.

The construction of trapped-ion quantum computers is heavily influenced by commodity technology, just as classical computing architectures have been. In particular, trapped-ion systems exploit lasers and AOMs used for writing chip masks for commodity silicon chips. The frequency of these commodity lasers determines the ions used (Ytterbium); but more significantly, the size of the AOMs limits the size of the control head to 32 lasers.

In a TILT system, since the tape head is not covering the entire ion chain, full connectivity is no longer supported among all qubits. For a two-qubit gate, if the distance between the two involved qubits is less than or equal to the tape head size, this two-qubit gate is executable. It may require the whole chain to be moved under the head, but no other gates are required. However, if the distance between the pair of qubits is more than the tape head size, both swap gates and tape movements are necessary, as in Figure 2(a).

A linear architecture makes communication via qubit swaps more efficient by channeling movement in one path, forcing data to pass in opposite directions; and requiring fewer swap gates involved in the program execution, achieving higher circuit fidelity. Utilizing this architectural feature, we can frequently pair up regular swaps to create *opposing swaps*. As Figure 2(c) shows, an opposing swap combines two swaps in opposite directions, and hence each opposing swap is equivalent to two regular swaps. This effect, combined with modest tape head sizes, can substantially reduce swaps as compared with 2D architectures.

Regarding the ion chain movement, our scheduling objective is to maximize the overall quantum program success rate. Previous noise analysis studies showed that motional excitations during shuttling are strongest when the acceleration of an ion is the highest [82,90]. As a result, the distance traveled, which is generally traversed at a nearly constant velocity, is not particularly relevant to the heating. Instead, there is a chance of heating when the ion begins or ends its motion. When the ion chain starts and stops moving, the movement will trigger the possibility of an increase in thermal motion. This thermal motion can introduce errors during the application of multiqubit gates, as discussed in Section IV-E. Thus, minimizing the number of movements can improve the overall circuit success rate.

To support our analysis on TILT, we have developed *LinQ*, a toolflow to compile from high-level quantum programs to a TILT architecture that minimizes the number of swaps and tape moves and hence improves the expected success rate. We implement a simulator for TILT to evaluate application metrics by using a suite of NISQ applications. We demonstrate that the TILT architecture outperforms QCCD architectures in terms of program success rate in a range of NISQ benchmarks. Figure 4 shows the *LinQ* overview.

The main contributions of this work are summarized as follows.

- To evaluate the feasibility of TILT systems, we present,

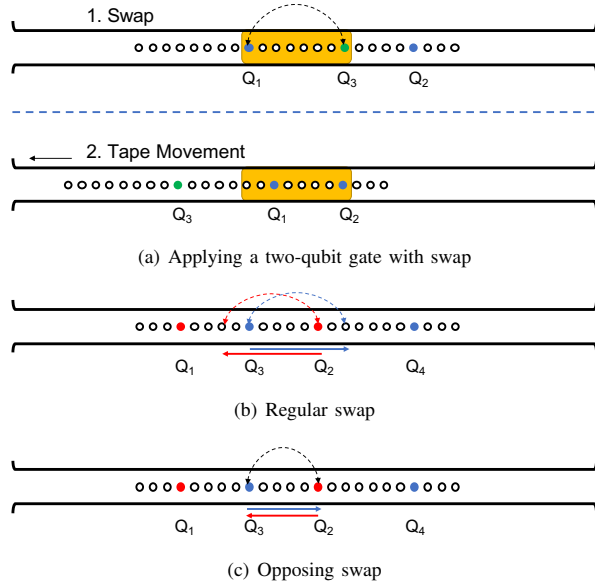


Fig. 2. (a) A two-qubit gate is applied to  $Q_1$  and  $Q_2$  (marked by blue). Since the distance between  $Q_1$  and  $Q_2$  is longer than the execution zone, a swap between  $Q_1$  and  $Q_3$  (marked by green) is needed. After the swap, a tape movement brings  $Q_1$  and  $Q_2$  in the execution zone for the two-qubit gate execution. (b) Suppose a two-qubit gate is applied to  $\{Q_1, Q_2\}$  (marked by red) and another two-qubit gate is applied to  $\{Q_3, Q_4\}$  (marked by blue). Two regular swaps are needed to make the two-qubit gates executable. (c) An opposing swap combines two regular swaps in opposite directions, and this swap makes both two-qubit gates executable. Thus, creating opposing swaps can reduce the total number of swaps.

for the first time, a comprehensive design and evaluation of the TILT architecture targeting systems with 64 qubits.

- To evaluate the performance of TILT architectures, we develop our toolflow, LinQ, that provides an optimizing compiler and simulator. We develop swap insertion and shuttling strategies to improve the overall quantum program success rates.
- To precisely understand the impact of thermal heating, we simulate using a noise model extracted from realistic experimental studies to estimate the application reliability.
- Our simulations show that TILT can achieve higher program success rates on a range of NISQ benchmarks compared to QCCD systems (up to 4.35x and 1.95x on average). Our results suggest that TILT provides a viable path toward scalable quantum computers.

## II. BACKGROUND

In this section, we give a brief overview of quantum computation. We then present the relevant background on trapped-ion systems.

### A. Principles of Quantum Computation

**Qubit.** A qubit is a two-level quantum system usually defined by two computational orthonormal basis states  $|0\rangle$  and  $|1\rangle$ . A quantum state can be expressed by any linear combination of

$|0\rangle$  and  $|1\rangle$ :  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , where  $\alpha$  and  $\beta$  are complex amplitudes satisfying  $|\alpha|^2 + |\beta|^2 = 1$ .

**Quantum Gates.** Quantum gates are unitary operations applied on qubits to map from one quantum state to another. Gates are applied to one or multiple qubits simultaneously. Arbitrary single-qubit gates and two-qubit controlled gates are known to be universal [25,66,88]. A quantum circuit (application) consists of a sequence of quantum gates to evolve the quantum state toward the solution state.

### B. Trapped-Ion QC Systems

In a trapped-ion quantum computer, information is stored in the internal states of atomic ions which are trapped in an oscillatory radio-frequency magnetic field. This field constrains the positions of the ions, and almost all current proposals have the ions arranged in a linear crystal, commonly referred to as the *ion chain*. The internal states of the ions can be manipulated by shining lasers on individual ions in order to implement single-qubit gates, and the motion of the chain is used as a bus in order to mediate entangling operations between multiple qubits. The qubit can be defined by either the hyperfine states or Zeeman states of the ion, each of which has benefits beyond the scope of this study [17].

**Ion Chain and Qubit Characteristics.** The ions in most modern ion traps are spaced approximately 5 microns apart. We do not consider  $T_1$  or  $T_2$  times in this study, but for a qubit built off of the hyperfine states of  $Yb^{171}$  ions, the former is on the order of  $10^{11}s$  and the latter is generally on the order of seconds, but measurements of up to ten minutes have been achieved [83]. The state of the qubits in this chain can be measured by shining a laser on the beam, which causes ions in the  $|1\rangle$  state to fluoresce. The ion crystal can then be imaged, with ions in the  $|1\rangle$  state appearing as bright spots, and ions in the  $|0\rangle$  state appearing as dark gaps [68].

**Laser-Based Operation.** In a trapped-ion quantum computer, gates are applied via laser pulses aimed at the target ions. Different operations can be implemented by varying the intensity, frequency, phase, or length of these pulses. For example, consider a single qubit gate. A laser is aimed at the ion which causes its state to slowly rotate, in the Bloch sphere, about the desired Pauli axis. The rotation speed can be increased by increasing the intensity of the beam. By carefully monitoring the intensity of the laser and tracking the timing, the desired rotation can be implemented. Two qubit gates are more complex; for trapped-ion quantum computers, the canonical two-qubit gate is the Mølmer-Sørensen gate, which implements an  $XX(\theta) = \exp(i\frac{\theta}{2}XX)$  operation by using lasers individually addressed on the desired pair of qubits, along with a global beam that hits all qubits in the chain. In order to create long-range two-qubit entanglement, this gate entangles each qubit with the motional state, applies state-dependent forces, and then unentangles the qubits from the motion. If done perfectly, this process leads to ions whose internal states are entangled without any residual entanglement between the internal states and the motion. In modern experiments, raw single-qubit gates have error rates around  $10^{-3}$ , although by using composite

pulses these rates can be improved significantly [16,86]. Error rates on two-qubit gates often depend on the length of the chain; but in small experiments, numbers as low as  $10^{-3}$  have been reported, although most experiments are still on the order of  $10^{-2}$  [5,50,87].

**Tape Shuttling.** In the proposed TILT architecture, the physical ion chain takes the place of the linear tape. By modulating the DC voltages of the electrodes in the trap, the ion chain can be moved along the axis of the trap. As explained above, lasers are used to process the information in a trapped-ion quantum computer. The idea behind an TILT architecture is that operations are applied only to ions near the middle of the trap, in the *execution zone*, and the chain is shifted back and forth in order to make long-range interactions occur. This strategy is counter to another architectural proposal, the quantum charge-coupled device [45]. A QCCD architecture comprises many smaller trapping zones within a single chip, and ions are shuttled around individually. Recently, Honeywell built the first QCCD system with four qubits [70]. While such an architecture allows for more parallelization and flexibility, it has high costs in terms of shuttling complexity since it requires expensive split/merge and junction crossing maneuvers, which lead to more thermal energy entering the system.

Thermal energy is stored in the motion of the chain, and it is also a significant factor in TILT architectures. The Mølmer-Sørensen gate mentioned in the preceding section is designed to be insensitive to the motional energy of the chain; however, this is true only for a perfectly applied gate. In reality, there is a small contribution to the infidelity of the gate, which is caused by residual entanglement between the internal state of the ions and the motional state, which was used to mediate the interaction. This error is due to the imperfect closure of a loop in phase space (the space of the positions and momenta of the particle). The error caused by this imperfect closure scales as  $\exp(2n + 1)$ , where  $n$  is the number of motional quanta in the chain. As a result, as a chain heats up, it becomes more sensitive to imperfections in the application of the gate.

### III. TILT ARCHITECTURE

TILT architectures have a linear chain of ions trapped within the oscillating potential which shuttles as one large block. We call this linear chain a *tape*. Most shuttling studies focus on split and merge operations or shuttling single qubits through junctions. A primary benefit of the linear tape architecture is that these difficult maneuvers are unnecessary.

#### A. Feasibility

Prior works [82,90] show that shuttling distance of ions is irrelevant to the heating since ions are shuttled at a nearly constant velocity. However, the start and stop of the ion shuttling will add thermal heating to the ion chain. Based on previous noise analysis studies [90], we derive a heating model in which every time the chain is shuttled its motional energy increases by some fixed amount  $k$ . This quantity is affected by how free the chain is to oscillate in its center of mass mode after shuttling. As a result, we scale this value by  $k \sim \sqrt{n}$

when moving to a larger chain, where  $n$  is the number of ions of the chain. Whenever a two-qubit gate is enacted, its fidelity is dependent on the total motional energy of the chain. High motional energy leads to larger gate errors, so shuttles have a negative impact on the overall success rate of the future gates. As a result, compiling techniques minimizing the number of tape moves are favored.

All components of this architecture have been developed to some extent, however ion trap quantum computing experiments have yet to be attempted at the point where our techniques are necessary. Chains as large as 79 ions have been demonstrated by ionQ [39], however they didn't individually address all ions. Using a TILT architecture would make this addressability concern much simpler. Shuttling techniques like the ones needed for our work have already been demonstrated for small segments of crystals, and would be scaled up to larger chains without an increase in complexity [33,70,82]. The benefit of the TILT architecture is that it does not require any pieces that have not been demonstrated already at the level of complexity necessary for the architecture, as opposed to QCCD architectures which will need more complex junction shuttles and trap designs as they scale up from their small demonstrations to larger quantum devices.

#### B. Gate Selection

There are multiple gate implementations for trapped-ion systems, including Frequency Modulated (FM), Phase Modulated (PM), and Amplitude Modulated (AM) gates [4,13,64, 87]. The gate time for FM gates is proportional to the total number of ions in a chain [49,50]. However, the operation time of PM and AM gates is only proportional to the distance between two involved ions [59,81,90]. Since TILT has a long chain with a smaller region for operations, using FM gates would not take advantage of our systems structure, leading to extra gate error. Therefore, PM or AM gates are proper gate implementations for TILT. In our study, we consider amplitude modulated (AM) gates for our two-qubit gate implementation.

#### C. Compared with QCCD

Another proposed model to scale trapped-ion QC systems is the quantum charge coupled device (QCCD), which requires complex junction shuttles and trap designs as they scale up [45]. Additionally, QCCD requires ion chain split/merge operations during the process of a computation [64] for communication with qubits in other traps.

Figure 3 shows the operations required for a short-distance (within the laser head size) communication pattern. QCCD requires a series of operations to perform all two-qubit interactions. However, TILT only needs a single tape movement to achieve the operations. As a result, TILT is expected to have higher success rates. Furthermore, if there are multiple short-distance communications, a QCCD system would require multiple split/merge/shuttle operations, but it is possible for TILT to achieve these operations with only one tape movement through a well-designed tape move scheduling. In this case, TILT will outperform QCCD architectures. For long-distance

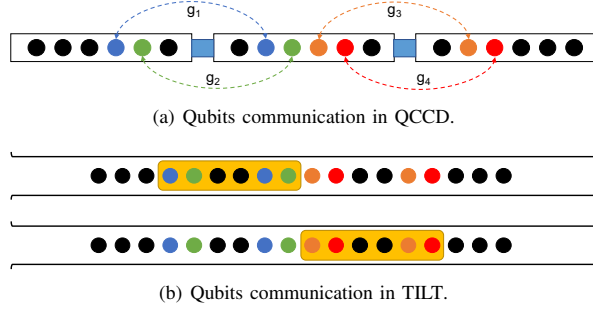


Fig. 3. Compare TILT with QCCD.  $g_1$ ,  $g_2$ ,  $g_3$ , and  $g_4$  are two-qubit gates on four pairs of qubits (marked by blue, green, orange, and red). (a) For QCCD, when a qubit communicates with another qubit in the different trap, QCCD requires the following operations: i) Swap the qubit to the end of the current trap. ii) Split from the current chain. iii) Shuttle to the target trap. iv) Merge to the target chain. v) Interact with the target qubit. To execute the gates  $g_1$ ,  $g_2$ ,  $g_3$ , and  $g_4$ , QCCD needs to repeat the above operations four times. (b) For TILT, performing a tape movement would allow the operations of the desired two-qubit gates.

communication, QCCD requires almost the same operations for short-distance communication if there is a shuttling pathway, but for TILT, the system might require multiple swaps and tape shuttles. Thus, QCCD can perform better than TILT in dealing with a long-distance communication.

In the current state of QC, most quantum applications heavily rely on short-distance communication patterns, such as the Variational Quantum Eigensolver (VQE) [43], Quantum Approximate Optimization Algorithm (QAOA) [27], Quantum Error Correction (QEC) like the surface code [28], Quantum Generative Adversarial Networks (QGAN) [23,36,54,78,91], and the Ising model solver [6]. Therefore, TILT is a promising design to achieve higher fidelity on these applications with a trapped-ion quantum computing architecture.

#### D. TILT Challenges

Under the device limitations of TILT, there are two major challenges to performing quantum circuits. Since the qubits are not fully connected, two-qubit gates on pairs of qubits whose distance is greater than the head size will need swap gates in order to make two qubits close enough to become executable. Consequently, we need qubit mapping and swap insertion techniques which minimize the number of swaps in order to achieve high quantum program success rates. Second, since only a part of the ion chain can be moved into the execution zone and since the tape movement will cause a heating error, minimizing the number of tape moves is essential to success.

### IV. LINQ FOR TILT ARCHITECTURE

To evaluate the TILT architecture, we develop our toolflow, LinQ, including the compiler and simulator. LinQ allows us to compile a quantum program written in a high-level programming language to TILT-architecture-level instructions. LinQ takes the device specification as input, including the number of qubits and the tape head size, and returns a compiled circuit, optimizing for program success rate. Since the TILT

machine is not yet realized, we use simulation to predict the performance of this novel architecture. We summarize the notations used in this paper in Table I.

TABLE I  
NOTATIONS USED IN THIS PAPER.

Notation	Definition
$n$	The number of ions of a chain.
$g$	A two-qubit gate.
$g.q_1, g.q_2$	A pair of qubits that $g$ is applied to.
$d_g$	The distance between $g.q_1$ and $g.q_2$
$M$	A mapping from logical to physical
$M_{q_i, q_j}$	A qubit mapping after swapping $(q_i, q_j)$
$G$	A set of two-qubit gates
$D(g, M_{q_i, q_j})$	$d_g$ under the mapping $M_{q_i, q_j}$
$\Gamma$	Background heating rate of the trap
$\tau$	Gate time
$k$	Amount of heating added during each shuttle
$\epsilon$	Gate error due to residual entanglement

#### A. LinQ Overview

Figure 4 shows an overview of our toolflow. LinQ contains three core compiler modules: trapped-ion native gate decomposition, qubit mapping and swap insertion, and tape movement scheduling. Since our target device size for TILT is more than 60 qubits, using one procedure for optimizing qubit mapping, swap insertion, and tape movement scheduling will have a long compilation time. Therefore, we separate the optimization problem into two steps and address them independently. However, when we perform swap insertion, we should take tape movement scheduling into consideration, because the swap patterns will affect the number of tape movements. In Section VI-B we will show that this two-step optimization can generate success rates similar to those of ideal trapped-ion systems.

The inputs for LinQ consist of a high-level quantum program and the device specification, including the number of qubits in the tape and the tape head size. First, LinQ decomposes the quantum program to trapped-ion native gates. Next, LinQ maps the logical qubits to physical qubits and resolves all long-distance two-qubit gates by inserting swaps gates. Then, LinQ performs our tape movement scheduling algorithm. The output is the executable circuit consisting of a sequence of scheduled gates and tape movements.

To evaluate the performance, such as success rates and the execution time, we pass the sequence of scheduled gates and tape movements to our simulator, which uses noise models extracted from realistic experimental studies.

#### B. Native Gate Decomposition

In order to generate executable code, gates from high-level quantum programming languages must be decomposed into the device-dependent native gates. LinQ decomposes the universal gates, such as CNOTs, CZs, and other operations into the TILT native gate set [57]. For example, we decompose a CNOT  $q_1, q_2$  gate into a sequence of rotations and XX



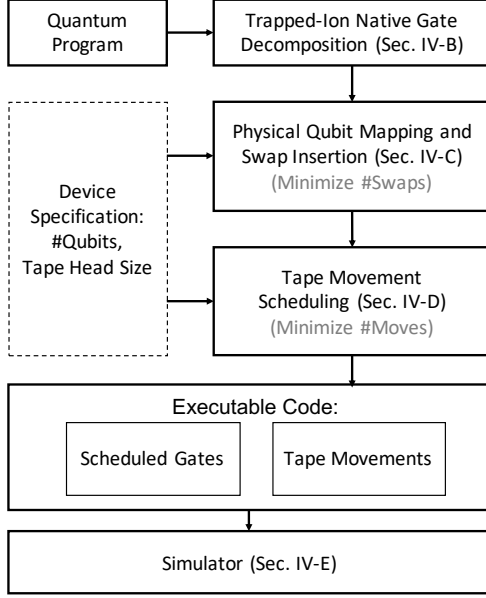


Fig. 4. LinQ overview. Taking a quantum program and device specification as input, LinQ generates the executable codes, including a sequence of scheduled gates and tape movements, and then runs the simulation to output the success rate and run-time.

operations:  $R_y(\pi/2) q_1$ ;  $XX(\pi/4) q_1, q_2$ ;  $R_x(-\pi/2) q_1$ ;  $R_x(-\pi/2) q_2$ ;  $R_y(-\pi/2) q_1$ .

### C. Qubit Mapping and Swap Insertion

Unlike fully connected trapped-ion systems, TILTsystems cannot operate long-distance two-qubit gates directly if the distance between the pair of qubits is longer than the tape head size. These long-distance two-qubit gates become unexecutable gates and to resolve them, we must insert swap gates. For a two-qubit gate  $g$  applied on  $q_1$  and  $q_2$ , we use  $d_g$  to denote the distance between  $q_1$  and  $q_2$ . The swap insertion problem is to find a sequence of intermediate qubits, for example,  $\{(q_1, q_i), (q_i, q_j), \dots, (q_k, q_2)\}$ , such that any pair  $(q_i, q_j)$  in the sequence has distance smaller than the tape head size.

Several approaches of swap insertion have been proposed. One common approach is to formulate the problems in a mathematical form, such as integer linear programming, and then utilize software solvers to find the optimal solutions [3,55,58,85,89]. This method is guaranteed to provide an optimal solution. However, since the time-to-solution grows exponentially with the number of qubits and the circuit depth (or the number gates), scaling this approach to large NISQ programs is infeasible. In our study, we focus on 60+ qubits applications. Hence, we need a practical solution for the swap insertion problem.

**Baseline Approach.** A straightforward idea is to allow swap with the tape head size distance to ensure a minimal number of swaps required for a two-qubit gate. We firstly establish our baseline implementation by using IBM Qiskit *Stochastic-*




Fig. 5. Considering two-qubit gates  $g_1$  and  $g_2$  on a system with the tape head size of  $L$ . Since  $d_{g_1}$  is  $L - 1$ , only one tape head position is allowed for  $g_1$  to be executed. As for  $g_2$ , there are three valid tape head positions because  $d_{g_2}$  is  $L - 3$ .

Swap [21], which is commonly used for swap insertion and it is also used in the highest optimization level of Qiskit circuit compilation pass.

However, two problems arise. First, using the tape head size as the swap distance will force the tape to move exactly once for every swap. If  $d_g$  is less than the tape head size, the tape movement scheduler will have multiple choices to schedule the tape position for gate  $g$ ; If  $d_g$  is equal to the tape head size, the scheduler must move the tape to the exact position for the two-qubit gate, and hence this will increase the number of tape movements in many cases. Second, this swap insertion method does not consider opposing swaps. An opposing swap (Figure 2(c)) combines two regular swaps in opposite directions. As a result, creating opposing swaps frequently can reduce the total number of swaps. This reduced swap count results in greater circuit success rates.

A good qubit mapping algorithm can also reduce the number of swap gates. Previous studies have proposed two types of solutions for qubit mapping and swap insertion for 1D and 2D architectures. One is to formulate the problem as an equivalent mathematical problem and then use a solver to find the optimal mapping to the problem [19,58]. This approach cannot scale to solve large problem for similar reasons to the swap insertion problem. The other solution is to use a heuristic search to obtain a mapping [40,51,67]. Since our target size is 64 qubits, we use a heuristic search to address this problem.

**LinQ Approach.** We adopt the existing heuristic algorithms for initial qubit mapping [40,51]. After we have an initial qubit mapping  $M$ , we improve the swap insertion. For resolving an unexecutable two-qubit gate  $g$  on  $q_1$  and  $q_2$ , we insert swaps by applying our heuristic algorithm, shown in Algorithm 1. The idea is to create opposing swaps to reduce the overall swap count. For each two-qubit gate applied on  $q_1$  and  $q_2$ , for all  $q_i$  between  $q_1$  and  $q_2$ , if the distance between  $(q_i, q_1)$  or  $(q_i, q_2)$  is less than the maximal swap length ( $MaxSwapLen$ ), then we put the swap candidate into a list  $S$ . Next, we compute the score for each swap candidate and select the swap with the minimal score to update the mapping and insert the swap to the circuit. If the tape head size is  $L$ ,  $MaxSwapLen$  can be  $L - 1$ . For a swap gate  $g$ , if  $d_g$  is equal to  $L - 1$ , the tape must be moved to the exact position for the gate execution, as shown in Figure 5. As a result, the circuit might require more tape movements. To mitigate this problem, we

**Algorithm 1** Swap Insertion

---

```

1:  $S \leftarrow \phi$ 
2: while  $g$  is unexecutable do
3:    $C \leftarrow \phi$ 
4:   for  $q_i$  from  $g.q_1$  to  $g.q_2$  do
5:     if  $\text{dist}(q_i, g.q_1) < \text{MaxSwapLen}$  then
6:        $C \leftarrow C \cup \{(q_i, g.q_1)\}$ 
7:     end if
8:     if  $\text{dist}(q_i, g.q_2) < \text{MaxSwapLen}$  then
9:        $C \leftarrow C \cup \{(q_i, g.q_2)\}$ 
10:    end if
11:  end for
12:  for  $(q_i, q_j)$  in  $C$  do
13:    Compute  $\text{Scores}(M_{q_i, q_j})$ 
14:  end for
15:  Find the swap with minimal score
16:   $S.append(\text{Swap}(q_i, q_j))$ 
17:   $M \leftarrow M_{q_i, q_j}$ 
18: end while
19: Insert the swap gate sequence  $S$  to the circuit

```

---

restrict  $\text{MaxSwapLen}$  to a certain number less than  $L - 1$  to create the flexibility for our tape movement scheduling. As long as the swap gate count is not increased significantly, by shorting the  $\text{MaxSwapLen}$ , we can increase the available head positions to potentially operate more gates in one tape movement. The best choice of the  $\text{MaxSwapLen}$  depends on the structure of the target application. An ideal  $\text{MaxSwapLen}$  will limit the overall  $d_g$  but will not increase the total number of swap gates. We can repeatedly run the procedure with different parameters to select the best  $\text{MaxSwapLen}$ .

The heuristic score function sums  $d_g$  for each two-qubit gate  $g$  under a mapping  $M$ . If a swap pair  $(q_i, q_j)$  leads to the minimal sum of all  $d_g$ , this swap pair will be selected. The swap score of the qubit pair  $(q_i, q_j)$  is defined as follows:

$$\text{Score}(M_{q_i, q_j}) = \sum_{g \in G} D(g, M_{q_i, q_j}) \times \alpha^{\Delta(g)}, \quad (1)$$

where  $G$  is the set of remaining two-qubit gates,  $M_{q_i, q_j}$  is the qubit mapping after swapping  $q_i$  and  $q_j$ ,  $D(g, M_{q_i, q_j})$  denotes  $d_g$  under the mapping  $M_{q_i, q_j}$ ,  $\alpha$  is a parameter ( $0 < \alpha < 1$ ) used to prioritize two qubit gates, and  $\Delta(g)$  is the circuit depth distance between the gate  $g$  and the current resolved gate. We use the sum of  $d_g$  to consider the impact for the future swaps. The optimal swap can impact the swaps in the future to reduce the overall qubit distance between the pair of qubits. Thus, we choose the swap candidates have the minimal scores.

The computational complexity of our heuristic algorithm can be estimated by the worst case, where the two qubits are separated in two ends and the swap insertion is performed repeatedly until two qubits are close enough to make the two-qubit gate become executable. The complexity is  $O(NG)$  for a two-qubit gate, where  $N$  is the total number of qubits in a tape and  $G$  is the number of remaining two-qubit gates.

**Algorithm 2** Tape Movement Scheduling

---

```

1:  $T \leftarrow \phi$ 
2:  $E \leftarrow \phi$ 
3:  $R \leftarrow G$ 
4: while  $R$  is not empty do
5:   for each tape head position  $p_i$  do
6:      $F_i \leftarrow \{\text{ExecutableGates}\}$ 
7:     Compute  $\text{Score}(p_i)$ 
8:   end for
9:   Find the tape head position  $p$  with the maximal score;
10:   $T.append(p)$ 
11:   $E.append(F_p)$ 
12:   $R \leftarrow R - F_p$ 
13: end while
14: Output the tape movement scheduling  $T$  and gate execution sequence  $E$ 

```

---

**D. Tape Movement Scheduling**

In TILT architectures, since the tape head is shared among all qubits, the tape will be moved back and forth frequently to execute a quantum application. Such tape movement/shuttling can introduce noise into the systems. Aware of issues like this, an effective tape movement scheduler can increase circuit fidelity by mitigating shuttling noise. In this work, we model the fidelity of the operations in order to improve the success rate. Our noise model considers the impact of thermal heating due to shuttling [90]. Since every tape move introduces thermal heating to the TILT system, this introduced heating leads to a lower gate fidelity. Hence, The design of our scheduling algorithm is to perform as many operations as possible before a tape movement is performed. Algorithm 2 shows the pseudocode of our scheduling algorithm. The *ExecutableGates* follows the gate dependency order. We compute the scores for each tape head position and select the position with the maximal score to schedule; we then repeat the procedure until all gates are scheduled and executed.

The objective of the heuristic cost function is to indicate the number of executable gates for a tape head position. As a result, the number of executable gates is the score for each tape head position. The general form of our heuristic cost function for a tape head position  $p$  is shown as follows:

$$\text{Score}(p) = n_p, \quad (2)$$

where  $n_p$  is the number of executable gates at the position  $p$ .

The time complexity of our heuristic scheduling algorithm is  $O((N - L)DL)$ , where  $N$  is the number of qubits in a trap,  $L$  is the tape head size, and  $D$  is the circuit depth.

**E. TILT Simulation**

To understand the impact of shuttling noise on TILT architectures, we perform noisy circuit simulation using realistic noise models. In our simulations we consider amplitude modulated (AM) gates for our two-qubit gate implementation [81]. These gates have a gate time of

$$\tau(d) = 38 \times d + 10 \quad (3)$$

where  $d$  is the distance between the two ions, in units of ion spacings, and the time is given in microseconds. This same model was studied in [64].

In an ion-trap quantum computer, single-qubit gate fidelities are independent of the thermal energy in the trap. Additionally, a perfectly applied Mølmer-Sørensen gate is independent of the thermal energy of the chain. As the vibrations increase, however, the gate becomes more sensitive to any imperfections in the laser pulses, which we quantify through the following model of gate fidelities after  $m$  moves have occurred, inspired by the work done in [90]:

$$F_m = 1 - \Gamma\tau + (1 - (1 + \epsilon)^{2mk+1}) \quad (4)$$

where  $\Gamma$  is the background heating rate of the trap,  $\tau$  is the gate time defined in microseconds as a function of number of intra-ion spacings between the involved qubits [64],  $k$  is the average amount of heating added to the chain during each shuttle, as discussed in Section III; and  $\epsilon$  is the error in each two-qubit gate due to residual entanglement with the motion. A gate can be thought of as a loop in phase space, where a perfect gate is a closed loop. An imperfect gate, due to a bad clock, rounding error in the pulse compilation, or laser intensity instability, can lead to this loop not perfectly closing. As the motional mode heats up, the gates become more sensitive to this kind of noise, as explained in [90]. An error model very similar to this is used in [64]. We do not use a linear approximation for the exponential in [90] to more accurately model higher motional mode excitations in our device.

In Honeywell's 4-qubit (8-ion) device [70], they report the average heating contribution of their shuttling operations as 2 quanta. This includes primitives such as split/merge and swap, as well as the simpler linear shuttles that the TILT architecture requires. Since split/merge operations and swap operations require complex and precise electrode voltages, these operations have significantly higher contributions to the heating rate than ion chain shuttling. As a result, we can fairly assume that the heating rate due to linear shuttles is lower than this number.

To scale this up to our experiments, we note that the primary consideration for heating is the softening of the common (center of mass) mode as the number of ions in the shuttled chain increases. The stopping force is still caused by the electrodes affecting the end of the chains, and is consequently constant [82]. However the mass of the chain increases, and as a result the noise in the system scales like a simple harmonic oscillator, with the frequency scaling by  $\sqrt{n/F} \sim \sqrt{n}$ .

## V. EXPERIMENTAL SETUP

### A. Benchmarks

To evaluate TILT architectures against real applications, we choose multiple important quantum applications as our benchmarks (Table II). The benchmarks are selected to have different program characteristics in order to show TILT properties for arbitrary quantum circuits. According to communication patterns, there are three categories of applications.

TABLE II  
LIST OF BENCHMARKS.

Application	Qubits	2Q Gates	Communication
ADDER	64	545	Short-distance gates
BV	64	64	Long-distance gates
QAOA	64	1260	Nearest-neighbor gates
RCS	64	560	Nearest-neighbor gates
QFT	64	4032	Long-distance gates
SQRT	78	1028	Long-distance gates

One consists of the applications with long-distance two-qubit gates, which require swap gates to become executable. The second category is for applications with short-distance two-qubit gates and hence do not necessarily need swap gates on TILT. The last category is nearest-neighbor two-qubit gates when the topology is in a 2D grid structure. This category will have short-distance two-qubit communication pattern when mapping on a linear topology.

The adder benchmark is based on the Cucarro adder [22]. Adders are important kernel functions in many quantum algorithms. Bernstein-Vazirani (BV) is a NISQ application commonly used to benchmark devices [9,24,87]. Quantum Approximate Optimization Algorithm (QAOA) [27] benchmark in our evaluation is a hardware-efficiency ansatz [60] for MaxCut problem. QAOA is a hybrid quantum-classical variational algorithm, and it is one of the most important quantum algorithms in the NISQ era. Random Circuit Sampling (RCS) has been proposed by Google to show quantum supremacy [2,14,56]. Quantum Fourier transform (QFT) [41] is an important function in many quantum algorithms (Shor's algorithm [76], phase estimation algorithm [20], and the algorithm for hidden subgroup problem [42]). This application uses Grover's search algorithm to find the square root number [41]. Grover's search algorithm is for database search, and it can achieve significant speedups compared with classical search algorithms [31,41].

### B. Simulation Parameters

We evaluate architectures with 60+ qubits and consider tape head sizes of 16 and 32. Using the noise model shown in Section IV-E, we estimate the fidelity of each gate after  $m$  tape moves.

All experiments are performed on a Ubuntu 16.04 system (Linux kernel 4.4-0-141-generic) with Intel Xeon Silver 4110 32-core CPU at 2.1 GHz and 128 GB of physical memory.

## VI. EVALUATION

In our evaluation, we first show the impact of swap insertion. We then compare TILT systems to QCCD systems, and present the compilation and estimated execution times of each application.

### A. LinQ Swap Insertion Performance

In this section, we show the importance of swap insertion choices. We use only those applications with long-distance



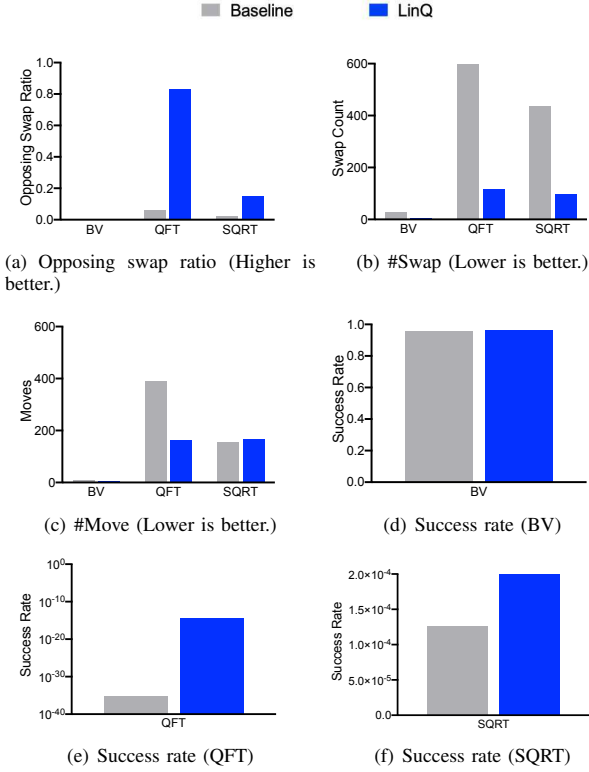


Fig. 6. Comparing LinQ swap insertion with the baseline swap insertion.

two-qubit gates (BV, QFT and SQRT) because no swap gates are needed for the other set of applications. We demonstrate the swap impact from two perspectives. First, we compare our LinQ swap insertion heuristic algorithm with a baseline method, which applies Qiskit StochasticSwap compilation tools to swap gate insertion for each unexecutable two-qubit gate. Second, we apply our swap insertion heuristic algorithm with a restricted swap distance to improve the tape movement scheduling. The results with our algorithm show increased circuit success rates.

We demonstrate the results with laser head size of 16. Figure 6(a) shows the opposing swap ratio, and Figure 6(b) shows swap counts with baseline and LinQ swap insertion. Our LinQ heuristic search reduces the total number of swaps and also increases the opposing swap ratio. For BV applications, however, LinQ does not create any opposing swaps because of the BV circuit structure [9,87].

Figure 6(c) shows the number of tape moves (lower is better). Compared with the baseline, LinQ can synthesize circuits with fewer swaps and consequently fewer gates in total. As a result, fewer tape moves are required for the circuit execution. The success rate is related to the total number of gates and tape moves. Fewer gates and tape moves results in a higher success rate. Thus, LinQ can also achieve higher success rates, as shown in Figure 6(d) - 6(f).

Figure 7 demonstrates the effect of limiting swap length.

Here we show the results of benchmarks on the devices with tape head size of 16. With LinQ swap insertion, while the swap is restricted to a shorter distance, the swap gate count could be slightly increased; but this restriction can potentially reduce the number of tape moves, leading to an improved success rate. Figure 7(a) shows the results of BV. The success rates are almost the same for the *MaxSwapLen* of 15 through 13, because there is no difference in the number of swaps and tape moves. While the max swap length is limited to a lower number, the number of swap gates and moves are increased, and the success rate drops. Figure 7(b) shows the number of swaps and tape moves for QFT under different swap length restrictions. When the maximal swap length is 14, the number of tape moves is the lowest. Although the swap count is increased compared with *MaxSwapLen* = 15, the overall success rate is increased. Therefore, we can get the highest success rate when the maximal swap length is shortened to 14. For SQRT, shown in Figure 7(c), when *MaxSwapLen* = 13, the swap count is slightly increased, but the tape moves can be significantly reduced. Hence, this application reaches the highest success rate at this configuration. For different applications, the best maximal swap length varies. We can iterate the LinQ procedure to find the best choice.

### B. Architecture Comparison

In this section, we compare the TILT architecture with head sizes of 16 and 32 with ideal trapped-ion (Ideal TI) devices and the QCCD system presented in [64].

**Ideal TI.** An ideal trapped-ion device means that the device has enough laser controls for each qubit. Consequently, two-qubit gates can be performed on an arbitrary pair of qubits. As a result ion shuttling and swap gate insertion are unnecessary under our error model. With comparing to this ideal case, we can learn how close we are to the optimal solution.

**QCCD.** A recent study on QCCD is presented in [64]. In our evaluation, we apply the same noise model and topology evaluated in the study. The QCCD configurations are 15-35 ions per trap in a linear topology, and we select the highest reported fidelity with AM gates for the comparison.

Figure 8 shows the results from our simulation. The success rate of TILT is related to the tape head size. For different tape head sizes, a larger tape head can cover a wider range of execution zone, reducing the number of swaps as well as tape moves and achieving a higher success rate. For ADDER and BV benchmarks, TILT has the same performance as QCCD. The reason is that the circuits are primarily operating on a few qubits. For QCCD, only a few shuttles have to be performed. Similarly for TILT, only a few tape moves are required to perform the circuit execution. However, for QAOA and RCS applications, the success rates of TILT are significantly higher than QCCD. As discussed in Section III-C, frequent short-distance two-qubit gates will cause QCCD to perform split/merge and shuttling operations periodically if the involved qubits are not in the same trap, while TILT can possibly schedule those operations in one tape move. Since QAOA and RCS repeat the operations on the pair of short-

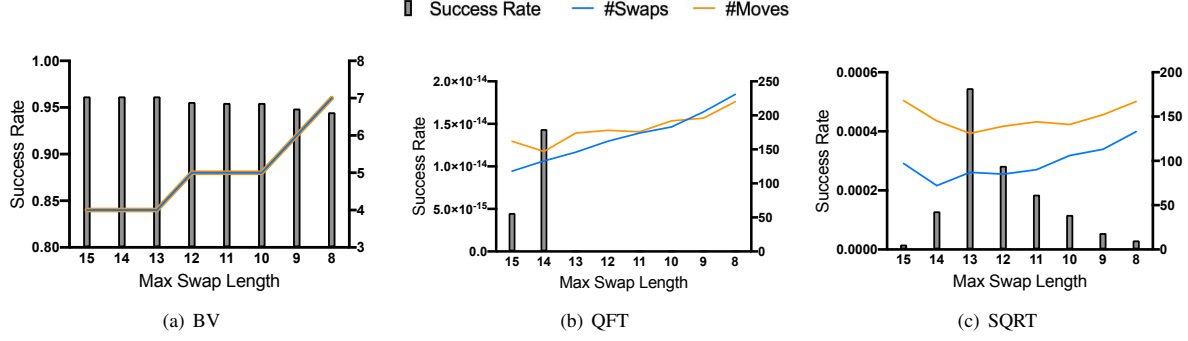


Fig. 7. Success rates under different *MaxSwapLen* restriction. For larger swap length, the circuit may need more tape moves. If the swap length is too short, however, the required number of tape moves increases as the total number of gates is increased. There is a swap length sweet spot depending on applications. For BV, the swap counts and move counts are the same so the lines are overlapped.




Fig. 8. Success rates on different device configurations (higher is better).

distance qubits, TILT outperforms QCCD in these results. For the circuits composed of short- and long-distance two-qubit gates (QFT and SQRT), TILT has lower success rate on QFT. This is as expected because the long-distance operations require multiple swaps and tape moves for TILT.

Our evaluation results suggest that TILT's success rate is better than that of QCCD for applications with a majority of the communication occurring within the width of a tape head.

### C. Compilation Time and Execution Time

Table III shows the compilation results. The column labels  $t_{\text{swap}}$  and  $t_{\text{move}}$  are the compilation times for swap insertion and tape move scheduling for the benchmarks using our LinQ toolflow. The  $t_{\text{swap}}$  is short for each application. As discussed in Section IV-C, the worst-case time complexity of searching for swap insertion is  $O(NG)$ . As a result, this heuristic search is scalable for larger applications. For tape move scheduling, as discussed in Section IV-D, the worst-case scheduling time is  $O((N-L)DL)$ . Since the circuit depth is increased when more swap gates are inserted, the scheduling time for QFT and SQRT is longer than for other applications. The number of swaps can be reduced by using a larger tape head, leading to lower circuit depth and a reduced  $t_{\text{move}}$ . The results and complexity analysis suggest that LinQ is applicable to larger circuits.

To estimate the program execution time, we analyze the compiled circuit depth and tape move distance. We model the gate time as shown in Equation 3, and the shuttling rate ( $t_m$ ) as  $1\mu\text{m}/\mu\text{s}$  [18]. We obtain the program execution time as following,

$$t_{\text{exe}} = t_m \times \text{dist} + \sum_d t_d, \quad (5)$$

where  $t_d$  is the maximum gate time for the  $d$ -th depth. From our simulation, the applications can be finished within seconds. The results show that TILT machines are viable for running large applications, like our benchmarks, since trapped-ion qubits have long coherence times [80].

## VII. TRAPPED-ION SCALING

Since ions are fundamentally identical and the same trapping parameters can handle a range of ion counts, scaling qubit count is not as significant a barrier as in other architectures. As we increase the number of ions, however, the heating cost of shuttling will also increase. This will lead to lower gate reliability, so there is an effective limitation on the number of ions in a single trap.

Various trapped-ion scaling techniques can be combined with TILT. In this section we discuss some additional techniques and technologies which could be used in conjunction with TILT to further scale trapped ion quantum computers.

TABLE III  
LINQ COMPILATION RESULTS.

Application	Head Size: 16					Head Size: 32				
	$t_{swap}(s)$	$t_{move}(s)$	#moves	dist( $\mu m$ )	$t_{exec}(s)$	$t_{swap}(s)$	$t_{move}(s)$	#moves	dist( $\mu m$ )	$t_{exec}(s)$
ADDER	0.002	6.214	10	104	2.967	0.002	3.248	5	68	3.252
BV	0.004	1.575	4	49	0.856	0.003	0.899	2	33	0.987
QAOA	0.011	2.750	18	232	1.564	0.011	1.101	4	72	1.357
RCS	0.001	3.148	65	992	1.704	0.001	0.680	11	214	0.856
QFT	9.206	55.617	162	2002	24.820	6.135	31.206	69	1276	33.876
SQRT	1.344	129.901	168	1816	46.554	1.075	55.234	76	1068	40.817

$t_{swap}$ : compile time for swap insertion.  $t_{move}$ : compile time for tape move scheduling. **#moves**: total number of tape moves. **dist**: total tape move distance.  $t_{exec}$ : program execution time.

**Sympathetic Cooling.** Gate reliability depends on the thermal energy in the chain. Sympathetic cooling is a technique to mitigate thermal heating, compensating for the increased heating rates expected in large-scale trapped-ion devices [18,63]. A dual-species ion chain is composed of two different types of ions, one for storing information and which can be cooled by another laser-cooled ion during circuit execution without damaging the stored information. TILT architectures are compatible with sympathetic cooling techniques, which would reduce the heating due to shuttling and allow for longer circuits. With cooling ions, a few modifications would be made for LinQ toolflow to consider the tape scheduling under the head of cooling laser beams. However, the fundamental design goal is still the same, minimizing the number of tape moves.

**QCCD Architectures.** The QCCD architecture is a well-known design for constructing a large-scale quantum computer [45,64]. QCCD systems have multiple small traps that are interconnected by segments and junctions. Before shuttling an ion from one trap to another, the ion must be split from the source chain. The split ion then is shuttled to the target trap and merged into the destination chain.

The operations necessary for individual ion shuttling have high costs in terms of shuttling complexity. Individual ion shuttling requires expensive split/merge and junction crossing maneuvers, which lead to more thermal energy entering the system. QCCD architectures could be combined with TILT as a combined architecture where the TILT systems discussed in this work would be a primitive for constructing the QCCD. Trap capacities could become larger, which might be useful for applications in which the circuit naturally breaks up into larger densely-communicating chunks.

Such an architecture could combine the strengths of each system, allowing for circuits which have significant medium to long-range communication to live within a single trap, while also allowing other sections of the circuit which might not require such distant communication to occur within smaller higher-fidelity trapping zones. Heterogeneous systems like this, where different traps serve different purposes, would be a logical evolution of trapped ion quantum computers into a more modular framework similar to modern classical computers.

**Photonic Interconnects.** TILT can also be utilized in modular quantum computer architectures such as the modular universal

scalable ion-trap quantum computer (MUSIQ) architecture proposed in [46,61]. The element logic units (ELUs) in this architecture consist of arrays of ions, and one or more ions in each ELU are coupled to photonic quantum channels through photonic interconnects. With these interconnects, we can use TILT architectures as the fundamental building blocks of ELUs to achieve modular TILT architectures and further scale QC devices, similar to the QCCD discussion above.

## VIII. RELATED WORK

Several studies of qubit mapping and swap insertion problems have been carried out for superconducting systems. One common approach is to formulate the problems in mathematical form, such as integer linear programming, and then utilize software solvers to find the optimal solutions [3,55,58,85,89]. This method is guaranteed to provide an optimal solution. However, since the time to solution grows exponentially with the number of qubits and gates, so scaling this approach to large NISQ programs is infeasible. In [89], a circuit partitioning approach is proposed. However, There is a large amount of wasted tape movement between each circuit block. Thus, our scheduling algorithm outperforms the ILP-based approach. Another approach is to apply dynamic programming to find the optimal solutions [40,77]. However, this method only works for circuits with 10 or fewer qubits since the time scaling is exponential. To avoid long execution times, several researchers have chosen to apply heuristic search algorithms [1,10,40,47,51,74,84,92]. Overall, most of the previous studies focus on compilation techniques for superconducting systems. In our work, we develop a scalable software toolflow to compile high-level quantum programs for TILT architecture, a novel trapped-ion system.

Previous studies have evaluated the performance of real devices with less than 20 qubits. In one case a fully connected 5-qubit trapped-ion system is compared with a 5-qubit superconducting transmon system [53]. In another study, several NISQ benchmarks are performed, comparing the trapped-ion system with superconducting systems [65]. These studies show that trapped-ion systems provide high program success rates because of the dense connectivity and higher gate fidelities compared with superconducting systems. In our work, we compare 64-qubit TILT systems to QCCD systems using our simulation tools and published results.

Several studies have proposed different strategies to scale trapped-ion quantum computers. Sargaran et al. proposed the SAQIP to apply reconfigurable optical interconnects for scalable trapped-ion systems [75]. Similarly, [46,61] proposed the MUSIQ architecture that uses photonic interconnects to link modular elementary logic units. This architecture can scale to thousands of qubits and support fault-tolerant error correction. Lekitsch et al. [48] proposed a blueprint for scalable trapped-ion systems that use microwave-based quantum gates with on-chip control electronics to control an arbitrary number of qubits. While these studies focus on very large systems that are unlikely to be built in the next decade, TILT architectures provide a practical near-term implementation to scale trapped-ion systems.

A recent study focuses on 50–100 qubit scale modular QCCD-based trapped-ion devices [64]. They propose the use of simulation techniques to study the impact of trap sizes, topology, and gate implementations. In our work, we provide simulation techniques to study TILT architectures. Architectural simulations allow us to evaluate the performance of different approaches of scaling trapped-ion systems before building them.

## IX. CONCLUSION

Trapped-ion technologies are a promising implementation for building practical quantum computers. Since all ions are identical, ions can be added to a long chain in order to scale the QC device. The TILT architecture is able to avoid issues with addressability as the number of ions addressed is held constant independent of chain length. Shuttling a small chain of ions has been demonstrated in [26], and larger TILT demonstrations are planned [15]. In this work, we show that TILT architectures offer a viable path toward QC devices approaching 60+ qubits. We present our optimizing compiler and simulator for TILT architectures. With a gate fidelity model derived from real experiments, our compiler performs qubit mapping, swap insertion, and tape move scheduling for NISQ applications within a few minutes. The results suggest that TILT can outperform QCCD in a range of NISQ applications. We also discuss using TILT as a building block to extend other scalable trapped-ion quantum computing proposals. Our evaluation of the TILT architecture motivates further experimental work in this direction and offers insights for future architecture design.

## ACKNOWLEDGMENTS

This work is funded in part by EPiQC, an NSF Expedition in Computing, under grants CCF-1730449/1832377; in part by STAQ, under grant NSF Phy-1818914; in part by NSF-OMA-2016136; and in part by DOE grants DE-SC0020289 and DE-SC0020331. This material is based upon work supported by the U.S. Department of Energy, Office of Science and Technology, under contract DE-AC02-06CH11357. DMD is also funded in part by an NSF QISE-NET fellowship (1747426). We thank Prakash Murali for fruitful discussions. We also thank the anonymous reviewers for their valuable comments and suggestions.

## REFERENCES

- [1] M. AlFailakawi, I. Ahmad, and S. Hamdan, "Lnn reversible circuit realization using fast harmony search based heuristic," in *Asia-Pacific Conference on Computer Science and Electrical Engineering*, 2014.
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [3] T. Bahreini and N. Mohammadzadeh, "An MINLP model for scheduling and placement of quantum circuits with a heuristic solution approach," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 12, no. 3, pp. 1–20, 2015.
- [4] H. Ball, M. J. Biercuk, A. Carvalho, R. Chakravorty, J. Chen, L. A. de Castro, S. Gore, D. Hover, M. Hush, P. J. Liebermann et al., "Software tools for quantum control: Improving quantum computer performance through noise and error suppression," *arXiv preprint arXiv:2001.04060*, 2020.
- [5] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, "High-fidelity quantum logic gates using trapped-ion hyperfine qubits," *Phys. Rev. Lett.*, vol. 117, p. 060504, Aug 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.117.060504>
- [6] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. Las Heras, R. Babbush, A. G. Fowler, B. Campbell, Y. Chen et al., "Digitized adiabatic quantum computing with a superconducting circuit," *Nature*, vol. 534, no. 7606, pp. 222–226, 2016.
- [7] M. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. Itano, J. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri et al., "Deterministic quantum teleportation of atomic qubits," *Nature*, vol. 429, no. 6993, p. 737, 2004.
- [8] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, "Mixed-state entanglement and quantum error correction," *Physical Review A*, vol. 54, no. 5, p. 3824, 1996.
- [9] E. Bernstein and U. Vazirani, "Quantum complexity theory," *SIAM Journal on computing*, vol. 26, no. 5, pp. 1411–1473, 1997.
- [10] A. Bhattacharjee, C. Bandyopadhyay, R. Wille, R. Drechsler, and H. Rahaman, "A novel approach for nearest neighbor realization of 2D quantum circuits," in *2018 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 2018, pp. 305–310.
- [11] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, p. 195, 2017.
- [12] R. Blakestad, C. Ospelkaus, A. VanDevender, J. Amini, J. Britton, D. Leibfried, and D. J. Wineland, "High-fidelity transport of trapped-ion qubits through an X-junction trap array," *Physical Review Letters*, vol. 102, no. 15, p. 153002, 2009.
- [13] R. Blumel, N. Grzesiak, and Y. Nam, "Power-optimal, stabilized entangling gate between trapped-ion qubits," *arXiv preprint arXiv:1905.09292*, 2019.
- [14] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," *Nature Physics*, vol. 14, no. 6, p. 595, 2018.
- [15] K. R. Brown. (2019) Software-tailored architectures for quantum codesign. [Online]. Available: <http://staq.pratt.duke.edu/>
- [16] K. R. Brown, A. W. Harrow, and I. L. Chuang, "Arbitrarily accurate composite pulse sequences," *Physical Review A*, vol. 70, no. 5, p. 052318, 2004.
- [17] N. C. Brown and K. R. Brown, "Comparing Zeeman qubits to hyperfine qubits in the context of the surface code:  $^{174}\text{Yb}^+$  and  $^{171}\text{Yb}^+$ ," *Phys. Rev. A*, vol. 97, p. 052301, May 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.97.052301>
- [18] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Applied Physics Reviews*, vol. 6, no. 2, p. 021314, 2019.
- [19] A. Chakrabarti, S. Sur-Kolay, and A. Chaudhury, "Linear nearest neighbor synthesis of reversible circuits by graph partitioning," *arXiv preprint arXiv:1112.0564*, 2011.
- [20] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, "Quantum algorithms revisited," *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1969, pp. 339–354, 1998.
- [21] A. Cross, "The IBM Q experience and QISKit open-source quantum computing software," *Bulletin of the American Physical Society*, vol. 63, 2018.

- [22] S. A. Cuccaro, T. G. Draper, S. A. Kutin, and D. P. Moulton, "A new quantum ripple-carry addition circuit," *arXiv preprint quant-ph/0410184*, 2004.
- [23] P.-L. Dallaire-Demers and N. Killoran, "Quantum generative adversarial networks," *Physical Review A*, vol. 98, no. 1, p. 012324, 2018.
- [24] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Demonstration of a small programmable quantum computer with atomic qubits," *Nature*, vol. 536, no. 7614, pp. 63–66, 2016.
- [25] D. P. DiVincenzo, "Two-bit gates are universal for quantum computation," *Physical Review A*, vol. 51, no. 2, p. 1015, 1995.
- [26] S. Fallek, C. Herold, B. McMahon, K. Maller, K. Brown, and J. Amini, "Transport implementation of the Bernstein–Vazirani algorithm with ion qubits," *New Journal of Physics*, vol. 18, no. 8, p. 083030, 2016.
- [27] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [28] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.
- [29] "A Preview of Bristlecone, Google's New Quantum Processor," <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>, accessed: 2019-03-29.
- [30] D. Gottesman, "An introduction to quantum error correction and fault-tolerant quantum computation," in *Quantum information science and its contributions to mathematics, Proceedings of Symposia in Applied Mathematics*, vol. 68, 2010, pp. 13–58.
- [31] L. K. Grover, "A fast quantum mechanical algorithm for database search," *arXiv preprint quant-ph/9605043*, 1996.
- [32] N. Grzesiak, R. Blümel, K. Beck, K. Wright, V. Chaplin, J. M. Amini, N. C. Pistenti, S. Debnath, J.-S. Chen, and Y. Nam, "Efficient arbitrary simultaneously entangling gates on a trapped-ion quantum computer," *arXiv preprint arXiv:1905.09294*, 2019.
- [33] M. Gutiérrez, M. Müller, and A. Bermúdez, "Transversality and lattice surgery: Exploring realistic routes toward coupled logical qubits with trapped-ion quantum processors," *Phys. Rev. A*, vol. 99, p. 022330, Feb 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.99.022330>
- [34] F. C. Hennie, "One-tape, off-line Turing machine computations," *Information and Control*, vol. 8, no. 6, pp. 553–578, 1965.
- [35] W. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk, "T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation," *Applied Physics Letters*, vol. 88, no. 3, p. 034101, 2006.
- [36] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou *et al.*, "Quantum generative adversarial learning in a superconducting quantum circuit," *Science advances*, vol. 5, no. 1, p. eaav2761, 2019.
- [37] G. Huber, T. Deuschle, W. Schnitzler, R. Reichle, K. Singer, and F. Schmidt-Kaler, "Transport of ions in a segmented linear paul trap in printed-circuit-board technology," *New Journal of Physics*, vol. 10, no. 1, p. 013004, 2008.
- [38] "IBM Announces Advances to IBM Quantum Systems and Ecosystem," <https://www-03.ibm.com/press/us/en/pressrelease/53374.wss>, accessed: 2019-03-29.
- [39] IONQ, "IonQ harnesses single-atom qubits to build the world's most powerful quantum computer," <https://ionq.com/news/december-11-2018>, 2018, online.
- [40] T. Itoko, R. Raymond, T. Imamichi, and A. Matsuo, "Optimization of quantum circuit mapping using gate transformation and commutation," *Integration*, vol. 70, pp. 43–50, 2020.
- [41] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi, "ScaffCC: Scalable compilation and analysis of quantum programs," *Parallel Computing*, vol. 45, pp. 2–17, 2015.
- [42] R. Jozsa, "Quantum factoring, discrete logarithms, and the hidden subgroup problem," *Computing in science & engineering*, vol. 3, no. 2, p. 34, 2001.
- [43] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, p. 242, 2017.
- [44] H. Kaufmann, T. Ruster, C. T. Schmiegelow, M. A. Luda, V. Kaushal, J. Schulz, D. Von Lindenfels, F. Schmidt-Kaler, and U. Poschinger, "Scalable creation of long-lived multipartite entanglement," *Physical Review Letters*, vol. 119, no. 15, p. 150503, 2017.
- [45] D. Kielpinski, C. Monroe, and D. J. Wineland, "Architecture for a large-scale ion-trap quantum computer," *Nature*, vol. 417, no. 6890, pp. 709–711, 2002. [Online]. Available: <https://doi.org/10.1038/nature00784>
- [46] J. Kim, P. Maunz, T. Kim, J. Hussman, R. Noek, A. Mehta, and C. Monroe, "Modular universal scalable ion-trap quantum computer (MUSIQ)," in *AIP Conference Proceedings*, vol. 1363, no. 1. American Institute of Physics, 2011, pp. 190–193.
- [47] A. Kole, K. Datta, and I. Sengupta, "A heuristic for linear nearest neighbor realization of quantum circuits by SWAP gate insertion using N-gate lookahead," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 6, no. 1, pp. 62–72, 2016.
- [48] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, "Blueprint for a microwave trapped ion quantum computer," *Science Advances*, vol. 3, no. 2, p. e1601540, 2017.
- [49] P. H. Leung and K. R. Brown, "Entangling an arbitrary pair of qubits in a long ion crystal," *Physical Review A*, vol. 98, no. 3, p. 032318, 2018.
- [50] P. H. Leung, K. A. Landsman, C. Figgatt, N. M. Linke, C. Monroe, and K. R. Brown, "Robust 2-qubit gates in a linear ion crystal using a frequency-modulated driving force," *Physical review letters*, vol. 120, no. 2, p. 020501, 2018.
- [51] 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*, 2019, pp. 1001–1014.
- [52] G.-D. Lin, S.-L. Zhu, R. Islam, K. Kim, M.-S. Chang, S. Korenblit, C. Monroe, and L.-M. Duan, "Large-scale quantum computation in an anharmonic linear ion trap," *EPL (Europhysics Letters)*, vol. 86, no. 6, p. 60004, 2009.
- [53] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Experimental comparison of two quantum computing architectures," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3305–3310, 2017.
- [54] S. Lloyd and C. Weedbrook, "Quantum generative adversarial learning," *Physical Review Letters*, vol. 121, no. 4, p. 040502, 2018.
- [55] A. Lye, R. Wille, and R. Drechsler, "Determining the minimal number of swap gates for multi-dimensional nearest neighbor quantum circuits," in *The 20th Asia and South Pacific Design Automation Conference*. IEEE, 2015, pp. 178–183.
- [56] I. L. Markov, A. Fatima, S. V. Isakov, and S. Boixo, "Quantum supremacy is both closer and farther than it appears," *arXiv preprint arXiv:1807.10749*, 2018.
- [57] D. Maslov, "Basic circuit compilation techniques for an ion-trap quantum machine," *New Journal of Physics*, vol. 19, no. 2, p. 023035, 2017.
- [58] D. Maslov, S. M. Falconer, and M. Mosca, "Quantum circuit placement," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, no. 4, pp. 752–763, 2008.
- [59] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavadia, and M. J. Biercuk, "Phase-modulated entangling gates robust to static and time-varying errors," *Physical Review Applied*, vol. 13, no. 2, p. 024022, 2020.
- [60] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.
- [61] C. Monroe, R. Raussendorf, A. Ruthven, K. Brown, P. Maunz, L.-M. Duan, and J. Kim, "Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects," *Physical Review A*, vol. 89, no. 2, p. 022317, 2014.
- [62] C. Monroe. (2019) 2019 CLEO Plenary: Christopher Monroe: Quantum Computing with Atoms. [Online]. Available: <https://www.youtube.com/watch?v=Nwo15la67C4>
- [63] M. Mudrich, S. Kraft, K. Singer, R. Grimm, A. Mosk, and M. Weidemüller, "Sympathetic cooling with two atomic species in an optical trap," *Physical review letters*, vol. 88, no. 25, p. 253001, 2002.
- [64] P. Murali, D. M. Debroy, K. R. Brown, and M. Martonosi, "Architecting noisy intermediate-scale trapped ion quantum computers," in *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture*, ser. ISCA '20. IEEE Press, 2020, p. 529–542. [Online]. Available: <https://doi.org/10.1109/ISCA45697.2020.00051>
- [65] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete, "Full-stack, real-system quantum computer studies: architectural comparisons and design insights," in *Proceedings of the*

- 46th International Symposium on Computer Architecture, 2019, pp. 527–540.
- [66] M. A. Nielsen and I. Chuang, “Quantum computation and quantum information,” 2002.
  - [67] S. Nishio, Y. Pan, T. Satoh, H. Amano, and R. Van Meter, “Extracting success from ibm’s 20-qubit machines using error-aware compilation,” *arXiv preprint arXiv:1903.10963*, 2019.
  - [68] R. Ozeri, “The trapped-ion qubit tool box,” *Contemporary Physics*, vol. 52, no. 6, pp. 531–550, 2011.
  - [69] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor,” *Nature Communications*, vol. 5, p. 4213, 2014.
  - [70] J. Pino, J. Dreiling, C. Figgatt, J. Gaebler, S. Moses, C. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson *et al.*, “Demonstration of the QCCD trapped-ion quantum computer architecture,” *arXiv preprint arXiv:2003.01293*, 2020.
  - [71] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
  - [72] “Rigetti Announces Its Hybrid Quantum Computing Platform and a 1m Prize,” <https://techcrunch.com/2018/09/07/rigetti-announces-its-hybrid-quantum-computing-platform-and-a-1m-prize/>, accessed: 2019-03-29.
  - [73] M. A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenkovic *et al.*, “Transport of quantum states and separation of ions in a dual RF ion trap,” *Quantum Information and Computation*, vol. 2, no. 4, p. 15, 2002.
  - [74] M. Saeedi, R. Wille, and R. Drechsler, “Synthesis of quantum circuits for linear nearest neighbor architectures,” *Quantum Information Processing*, vol. 10, no. 3, pp. 355–377, 2011.
  - [75] S. Sargaran and N. Mohammadzadeh, “SAQIP: A scalable architecture for quantum information processors,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 16, no. 2, pp. 1–21, 2019.
  - [76] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
  - [77] M. Y. Siraichi, V. F. d. Santos, S. Collange, and F. M. Q. Pereira, “Qubit allocation,” in *Proceedings of the 2018 International Symposium on Code Generation and Optimization*, 2018, pp. 113–125.
  - [78] H. Situ, Z. He, L. Li, and S. Zheng, “Quantum generative adversarial network for generating discrete data,” *arXiv preprint arXiv:1807.01235*, 2018.
  - [79] D. Stick, W. Hensinger, S. Olmschenk, M. Madsen, K. Schwab, and C. Monroe, “Ion trap in a semiconductor chip,” *Nature Physics*, vol. 2, no. 1, p. 36, 2006.
  - [80] N. Timoney, I. Baumgart, M. Johanning, A. Varón, M. B. Plenio, A. Retzker, and C. Wunderlich, “Quantum gates and memory using microwave-dressed states,” *Nature*, vol. 476, no. 7359, pp. 185–188, 2011.
  - [81] C. J. Trout, M. Li, M. Gutiérrez, Y. Wu, S.-T. Wang, L. Duan, and K. R. Brown, “Simulating the performance of a distance-3 surface code in a linear ion trap,” *New Journal of Physics*, vol. 20, no. 4, p. 043038, 2018.
  - [82] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger, “Controlling fast transport of cold trapped ions,” *Physical Review Letters*, vol. 109, no. 8, p. 080501, 2012.
  - [83] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum, and K. Kim, “Single-qubit quantum memory exceeding ten-minute coherence time,” *Nature Photonics*, vol. 11, no. 10, p. 646, 2017.
  - [84] R. Wille, O. Keszocze, M. Walter, P. Rohrs, A. Chattopadhyay, and R. Drechsler, “Look-ahead schemes for nearest neighbor optimization of 1D and 2D quantum circuits,” in *2016 21st Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2016, pp. 292–297.
  - [85] R. Wille, A. Lye, and R. Drechsler, “Optimal SWAP gate insertion for nearest neighbor quantum circuits,” in *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2014, pp. 489–494.
  - [86] S. Wimperis, “Broadband, narrowband, and passband composite pulses for use in advanced nmr experiments,” *Journal of Magnetic Resonance, Series A*, vol. 109, no. 2, pp. 221–231, 1994.
  - [87] K. Wright, K. Beck, S. Debnath, J. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. Pseni, M. Chmielewski, C. Collins *et al.*, “Benchmarking an 11-qubit quantum computer,” *Nature communications*, vol. 10, no. 1, pp. 1–6, 2019.
  - [88] X.-C. Wu, S. Di, E. M. Dasgupta, F. Cappello, H. Finkel, Y. Alexeev, and F. T. Chong, “Full-state quantum circuit simulation by using data compression,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–24.
  - [89] X.-C. Wu, Y. Ding, Y. Shi, Y. Alexeev, H. Finkel, K. Kim, and F. T. Chong, “ILP-based scheduling for linear-tape model trapped-ion quantum computers,” 2019.
  - [90] Y. Wu, S.-T. Wang, and L.-M. Duan, “Noise analysis for high-fidelity quantum entangling gates in an anharmonic linear Paul trap,” *Physical Review A*, vol. 97, no. 6, p. 062325, 2018.
  - [91] J. Zeng, Y. Wu, J.-G. Liu, L. Wang, and J. Hu, “Learning and inference on generative adversarial quantum circuits,” *Physical Review A*, vol. 99, no. 5, p. 052306, 2019.
  - [92] A. Zulehner, A. Paler, and R. Wille, “Efficient mapping of quantum circuits to the ibm qx architectures,” in *2018 Design, Automation & Test in Europe Conference & Exhibition*. IEEE, 2018, pp. 1135–1138.