Featured in Physics

How to Wire a 1000-Qubit Trapped-Ion Quantum Computer

M. Malinowski,[1,*] D.T.C. Allcock,[1,2] and C.J. Ballance[1,3]

[1] *Oxford Ionics, Oxford OX5 1PF, United Kingdom*

[2] *Department of Physics, University of Oregon, Eugene, Oregon 97403, USA*

[3] *Department of Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom*

(Received 28 May 2023; accepted 28 August 2023; published 19 October 2023)

One of the most formidable challenges of scaling up quantum computers is that of control-signal delivery. Today's small-scale quantum computers typically connect each qubit to one or more separate external signal sources. This approach is not scalable due to the input/output (I/O) limitations of the qubit chip, necessitating the integration of control electronics. However, it is no small feat to shrink control electronics into a small package that is compatible with qubit-chip fabrication and operational constraints without sacrificing performance. This so-called "wiring challenge" is likely to impact the development of more powerful quantum computers even in the near term. In this paper, we address the wiring challenge of trapped-ion quantum computers. We describe a control architecture called WISE (Wiring using Integrated Switching Electronics), which significantly reduces the I/O requirements of ion-trap quantum computing chips without compromising performance. Our method relies on judiciously integrating simple switching electronics into the ion-trap chip—in a way that is compatible with its fabrication and operation constraints—while the complex electronics remain external. To demonstrate its power, we describe how the WISE architecture can be used to operate a fully connected 1000-qubit trapped-ion quantum computer using approximately 200 signal sources at a speed of approximately 40–2600 quantum gate layers per second.

DOI: 10.1103/PRXQuantum.4.040313

I. INTRODUCTION

Trapped-ion qubits are one of the most promising approaches to quantum computing, especially in the moisy intermediate-scale quantum (NISQ) era [1,2]. One of their main superpowers is ion transport, i.e., the ability to physically move ions in space [3]. This enables two powerful features:

(1) *Qubit reconfiguration*. Ion transport allows for flexible qubit routing—i.e., changing which qubit is coupled to which other qubits—without relying on error-prone multiqubit gates [4]. This allows for effective all-to-all connectivity even in a system composed of many uncoupled qubit registers.

(2) *Transport-assisted gates*. Ion transport allows for local control of quantum operation Rabi frequency even when the qubit drive operates at a fixed amplitude. This method can be employed with both laser and microwave qubit drives—as long as they are spatially inhomogeneous—and has been used in, e.g., Refs. [5–8].

Qubit reconfiguration is the basis of the quantum charge-coupled device (QCCD) architecture [9,10], which is one of the most promising approaches to trapped-ion quantum computing. The effective all-to-all connectivity—when combined with excellent coherence times [11–13]—is one of the reasons why trapped-ion systems achieve such high quantum volumes compared to other platforms [14]. On the other hand, quantum gates in today's QCCD systems are typically implemented by delivering localized externally modulated qubit drives (e.g., laser beams) to individual trap regions. Ion transport is only leveraged in a limited way, e.g., to move an ion away from a laser beam to switch off the interaction. However, as systems grow and off-chip local drive modulation becomes impractical, transport-assisted gates become a powerful tool, as they allow for local control with only a small number of global qubit drives. In other words, transport-assisted gates reduce the problem of

*mm@oxionics.com

implementing quantum gates at scale to the problem of ion transport at scale.

However, the ability to execute ion transport in large-scale systems is hindered by the wiring challenge. Fast low-heating ion transport requires precise dynamical control of voltages on many electrodes [15,16]. For example, QCCD architecture demonstrations from Pino *et al.* [17] and Kaushal *et al.* [18] have both used about ten electrodes per qubit [19], each wired to a separate DAC outside of the vacuum system (see Fig. 1). This "standard approach" is likely not scalable even to moderate system sizes. For example, a 1000-qubit chip would require approximately 10 000 input lines—more than today's cutting-edge CPUs and at a bleeding edge of packaging feasibility. And while electrode cowiring, i.e., permanently connecting multiple electrodes to the same voltage source, has been proposed [20,21] and used [22,23] to reduce the number of inputs per electrode, existing cowiring approaches only allow for limited functionality, such as qubit storage or shuttling between remote zones.

One previously proposed solution to the wiring challenge is to form an integrated "quantum processing unit" (QPU), which combines an ion-trap chip with the DACs. The integration could be either monolithic [24] or achieved by packaging together several independently fabricated chips [25]. However, DAC integration comes with major challenges:

(1) *Power dissipation*. For example, Stuart *et al.* [24] have developed compact cryogenic (4 K) DACs with a power consumption of approximately 30 mW per channel, or approximately 300 W for 1000 qubits. While this can be optimized, DAC power dissipation presents a significant challenge, especially in cryogenic environments, which are beneficial for reducing noise in trapped-ion systems.

(2) *Data bandwidth*. Fully flexible DAC control requires streaming large volumes of data to the QPU, typically approximately 50 Mbit/s per DAC [18,26]. For a 1000-qubit QPU, this corresponds to approximately 500 Gbit/s of data flow between the DAC and the control system. Designing an appropriate interface would not be trivial and in practice might require integration of further digital electronics, e.g., to store the waveforms [27]. This increases the complexity of QPU design and fabrication and may limit the flexibility of operations.

(3) *Footprint*. An integrated filtered DAC will typically require a much larger chip area than the electrode itself. For example, a single 1-nF planar filter capacitor would use an area of approximately 500 μm $\times$ 500 μm in a standard CMOS process [28]. The DAC block itself may also require considerable space—e.g., 130 μm $\times$ 270 μm in Ref. [24]. Developing low-noise voltage sources with areas comparable to ion-trap electrodes (approximately 100 μm $\times$ 100 μm or less) is thus a challenge in itself and might require advanced techniques such as wafer stacking and trench capacitors [29–31].

Because of the problems highlighted above, current approaches are insufficient to address the wiring challenge, even for intermediate-scale QPUs with approximately 1000 qubits.

In this paper, we present an architecture called WISE (Wiring using Integrated Switching Electronics), which addresses the challenge of wiring such intermediate-scale trapped-ion quantum computers. Our method relies on simple integrated electronics that entail minimal power dissipation, low data bandwidth, and small footprint, alleviating all the major challenges of DAC integration. In WISE, all complex high-footprint high-power electronics are placed off chip, allowing for large-scale control without compromising performance.

The paper is structured as follows. In Sec. II, we give an overview of the electrode wiring architecture and ion-transport methods. We then describe how WISE can perform all the key operations of a large-scale QCCD architecture: arbitrary qubit reconfiguration (Sec. III) and transport-assisted gates (Sec. IV). Subsequently, in Sec. V, we discuss the hardware implementation in more detail, demonstrating that WISE is indeed compatible with

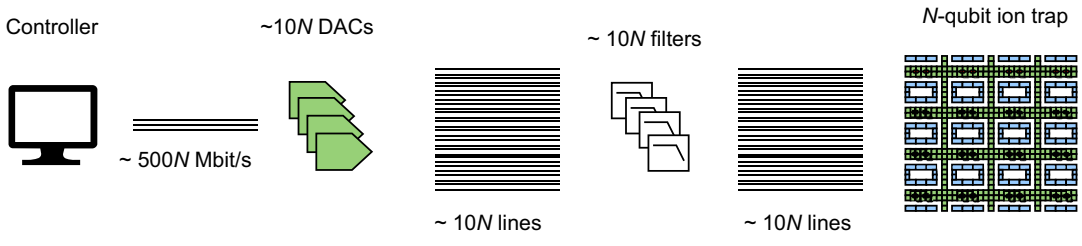

FIG. 1. An illustration of the standard approach to the electrical wiring of trapped-ion quantum computers. Qubit transport in an N-qubit ion trap is achieved by using approximately $10N$ electrodes. Each electrode is wired through an individual filter to an individual DAC. The DAC output waveforms are set through a digital interface, with typical data flow rates of approximately 50 Mbit/s per DAC.

ion-trap fabrication and operation constraints. Finally, we combine all those insights in Sec. VI, and describe how to build a 1000-qubit trapped-ion quantum computer capable of running arbitrary quantum circuits with all-to-all connectivity but using only approximately 200 input lines.

II. WISE METHOD

In a QCCD device, ion positions are controlled by voltages applied to trap electrodes. These electrodes serve two different purposes. The first is *dynamic*: to deliver time-varying waveforms which execute the desired transport primitives, such as shuttling [32,33], merging, splitting [34,35], or crystal rotations [36]. The second is *quasistatic*: to compensate stray electric fields generated by, e.g., local charges and differences in work functions [37, 38]. In WISE, each electrode is explicitly assigned to be either dynamic or quasistatic [39,40]. We then employ two techniques:

(1) *Dynamic electrode parallelization*. Dynamic electrodes are cowired to a fixed number of DACs, assigned through integrated switches.
(2) *Quasistatic electrode demultiplexing*. Quasistatic "shim" electrodes are controlled through a small number of DACs through integrated demultiplexers in a "sample-and-hold" fashion.

In WISE, instead of DAC integration, we primarily use switch integration. Unlike DACs, switches require very small data input rates and can be operated with negligible power dissipation. Furthermore, they are relatively simple structures, made of a small number of transistors and inverters. Thus, they can be readily integrated into an ion-trap fabrication process [24] and have been demonstrated to be cryogenically compatible, also in the context of quantum computing [41,42]. The resulting high-level wiring architecture is shown in Fig. 2.

A. Dynamic electrode parallelization

Consider an ion trap composed of multiple identical zones, each holding the same number of qubits. The main observation behind dynamic electrode parallelization is that, as long as the zone-to-zone electric field crosstalk is sufficiently small, the same transport operation can be executed in multiple areas of the chip at the same time by applying the same voltage waveforms to multiple electrodes. Furthermore, instead of permanently cowiring electrodes, a switch network can be used to dynamically select which voltage waveform is applied to which electrode. Thus, for example, parallel splits can be executed by connecting one set of DACs, delivering one set of split waveforms, to all the zones where we want to implement a split, and disconnecting it from all the zones where we do not want to implement a split.

Due to dynamic electrode parallelization, instead of using one DAC per electrode, it suffices to use one DAC per waveform. Furthermore—as we argue in Sec. III E—while it might be optimal to execute different transport operations in different zones at the same time (e.g., a split in some zones and a merge in some other zones), it is nonetheless efficient to only perform only one transport operation at any given time (e.g., a split in some zones, followed by a merge in some other zones). Due to those insights, a modest number of DACs suffices to execute arbitrary transport sequences.

Figure 2 (top row) illustrates the dynamic electrode wiring that leverages those insights. Outside the QPU, a fixed number of DACs are used to output qubit transport waveforms, while in the QPU, switches are used to select which electrode connects to which DAC. In subsequent sections, we show how dynamic electrode parallelization allows for efficient qubit reconfiguration in one-dimensional (1D) and two-dimensional (2D) ion traps (Sec. III) and why approximately 100 DACs suffice for the purpose, regardless of qubit number N. In Sec. IV A, we discuss how dynamic electrode parallelization can be used to perform transport-assisted quantum gates. Finally, Sec. V A describes the hardware implementation of the dynamic electrode switch network.

B. Quasistatic electrode demultiplexing

In WISE, all electrodes that are not dynamic are quasistatic, meaning that they are held at a constant voltage during any given qubit reconfiguration and during any given quantum gate [43]. The main insight is that holding a fixed voltage on electrodes can be less resource intensive than applying time-dependent waveforms, because electrodes can maintain their set voltage even while disconnected from DACs. Thus, a single DAC can be demultiplexed to control multiple electrodes. In this mode of operation, the DAC output is connected to shim electrodes one by one, charging them to the necessary voltage before disconnecting. While the DAC is disconnected, the electrode is connected to ground through an integrated capacitor, which holds the dc voltage while acting as an rf shunt [24].

Since quantum gates and transport operations are very sensitive to electric field noise [44,45], multiplexer operation can lead to additional errors, especially if the multiplexing frequency or the switching frequency is near the motional frequency of the ions (approximately 1–10 MHz). To alleviate that source of noise, we operate in a "sample-and-hold" fashion, where all the electrodes are charged first, and the sensitive operations are performed with the multiplexer turned off.

Figure 2 (bottom row) illustrates the shim electrode wiring in the WISE architecture. Outside the QPU, a small number of DACs cycle through shim-voltage set points,

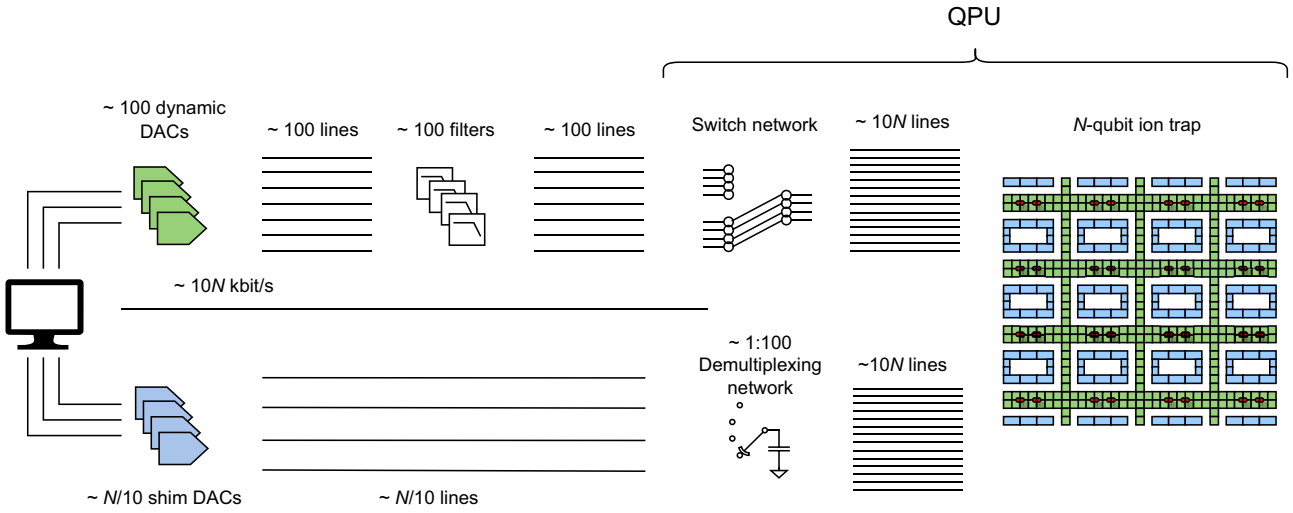

FIG. 2. A high-level illustration of the WISE architecture. We form a quantum processing unit (QPU) by combining an ion trap with a switch-demultiplexing network. In this way, all the high-density input/output (I/O) is confined to the QPU and the interface between the QPU and the outside world requires significantly fewer connections. We divide ion-trap electrodes into dynamic ones (green, approximately ten dynamic electrodes per qubits) and quasistatic ones, i.e., "shims" (blue, approximately ten shim electrodes per qubit). The dynamic electrodes are controlled by approximately 100 DACs, regardless of the system size, due to an integrated switch network. The number of shim DACs is reduced from one DAC per electrode in the standard approach to one DAC per approximately 100 electrodes in our approach, due to an integrated demultiplexing network. The high-speed digital interface is only required between the controllers and the DACs, while digital communication between the QPU and the controllers requires only small signal bandwidth and can be done serially.

while in the QPU, demultiplexers are used to connect the DACs to on-chip capacitors and electrodes. In subsequent sections, we describe the role of shim electrodes in qubit reconfiguration (Sec. III F) and in transport-assisted gates Sec. IV B. Finally, Sec. V B describes the hardware implementation of the shim-demultiplexing network, and motivates the choice of approximately 100 shim electrodes per DAC.

III. QUBIT RECONFIGURATION

In this section, we show how dynamic electrode parallelization can be used for arbitrary and efficient reconfiguration in the QCCD architecture. This is structured as follows. First, we describe how to perform a "switchable swap" of two qubits in a linear (1D) array. Second, we show how to use switchable swaps to perform arbitrary reconfiguration of N qubits in a linear array. Third, we extend the method to construct arbitrary reconfiguration of $N = m \times n$ qubits in a regular 2D array, where every qubit can be swapped with any of its four neighbors (i.e., one qubit per junction). Fourth, we extend the construction to a realistic N-qubit QCCD device, consisting of two-qubit chains held in a 2D array with k qubits per junction. Finally, we argue that the method is practical by calculating the expected run time and error rate of a worst-case reconfiguration.

A. Switchable swap

A switchable swap refers to a procedure where two neighboring qubits are physically swapped conditioned on the settings of on-chip switches.

Consider two qubits (1,2) held in separate linear zones (1,2), as illustrated in Fig. 3(a). Each zone contains $N_{\mathrm{de/lz}} \sim 10$ dynamic electrodes. We implement a switchable swap as follows. If switches (s_1, s_2) are set to $(1, 1)$, we connect the dynamic electrodes in both zones to $2 \times N_{\mathrm{de/lz}} \sim 20$ DACs, which play waveforms that result in a "swap sequence," such as shown in Fig. 3(b), which reorders the qubits. On the other hand, if the switches are set to $(0, 0)$, we connect the electrodes to $2 \times N_{\mathrm{de/lz}} \sim 20$ DACs, which execute a "stay-still" sequence, keeping the qubits in place. This allows the implementation of a switchable swap using $4 \times N_{\mathrm{de/lz}} \sim 40$ DACs [46] and 2 bits of information [47].

B. 1D reconfiguration

We now enlarge the processor to be a linear repeating array with zones $i = 1, 2, \ldots, N$, as illustrated in Fig. 4. Every second zone is connected to the same fixed set of DACs; i.e., if $s_i = s_{i+2}$, then every dynamic electrode in zone $i + 2$ executes the same waveform as the corresponding dynamic electrode in zone i. Thus, by playing the same waveforms as in Sec. III A and sending an N-bit word $s = (s_1, s_2, \ldots, s_N)$ to the QPU, we can implement a

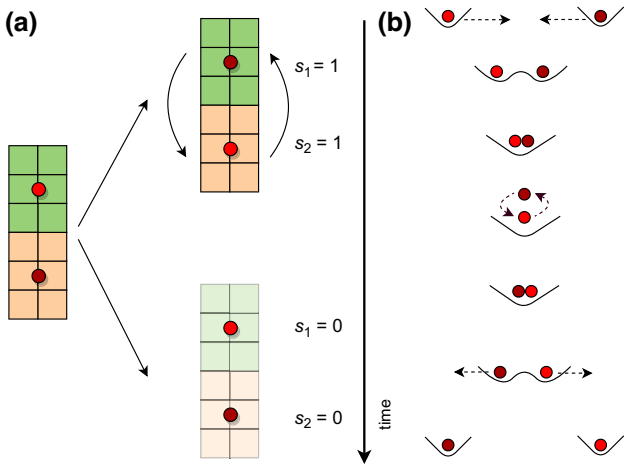

FIG. 3. (a) Basic switchable swaps. Two qubits (red) are placed in neighboring linear zones (1,2), shown in green and orange, respectively. Each zone contains $N_{\mathrm{de/lz}}{\sim}10$ dynamic electrodes (the figure shows $N_{\mathrm{de/lz}} = 6$ for simplicity), as well as rf and shim electrodes (not shown). When switches s_1 and s_2 are set to $(s_1, s_2) = (1, 1)$, the zones are "active" and the qubits undergo a swap. If, however, $(s_1, s_2) = (0, 0)$, the zones are "inactive" (drawn as partially transparent) and the qubits remain in their original locations. (b) The 1D swap sequence. Two qubits (red) in neighboring linear zones are swapped by bringing them together (shuttling) and merging their potential wells, followed by a crystal rotation, well split, and a shuttle step. At each step, the potential experienced by each ion is the net potential of all neighboring dynamic, shim, and rf electrodes.

switchable swap of qubits $(i, i+1)$ for every odd i in parallel. We call this step an "odd swap." Similarly, an "even swap"—a switchable swap of qubits $(i, i+1)$ for every even i in parallel—can be accomplished by playing a different "swap waveform" and sending another N-bit word s. The implementation of an arbitrary odd or even swap in a N-qubit array requires $4 \times N_{\mathrm{de/lz}}{\sim}40$ DACs, the same as a switchable swap, and N bits of information.

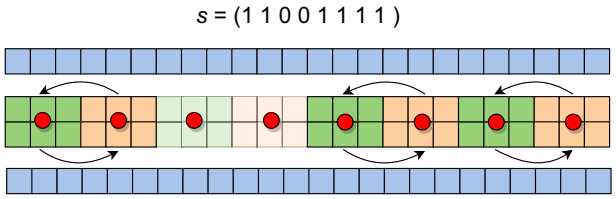

FIG. 4. An odd swap in a linear 1D array of $N = 8$ qubits with $N_{\mathrm{de/lz}} = 6$ dynamic electrodes per zone. In this example, the bit-select word is $s = (1, 1, 0, 0, 1, 1, 1, 1)$. Thus, zones 1, 2, 5, 6, 7, and 8 are active (their qubits undergo swaps), while zones 3 and 4 remain inactive (their qubits remain stationary). In addition to dynamic electrodes (shown in green and orange for odd and even zones, respectively), rf electrodes are shown in white and shim electrodes in blue.

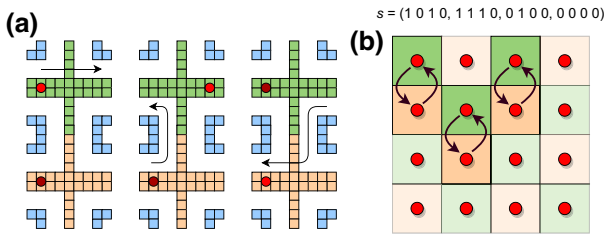

FIG. 5. (a) The 2D qubit-swap sequence. Two qubits in neighboring junctions can be swapped without crystal rotations as a sequence of three shuttling steps. (b) The 2D qubit swap. During a 2D swap, zone (i, j) is active (solid color) if and only if $s_{i,j} = 1$, in which case its qubit undergoes a swap operation. In the above image, the ion trap is a regular 4×4 array and the bit-select word $s = (s_{1,1}, s_{1,2}, \ldots, s_{2,1} \ldots, s_{4,4})$ has length 16. This particular swap step is an odd vertical swap, i.e., zones (i, j) and $(i, j+1)$ undergo a swap if and only if $i + j$ is even and $(s_{i,j}, s_{i,j+1}) = (1, 1)$.

As is well known, these odd-even swap primitives suffice to perform arbitrary qubit reconfiguration in 1D using an algorithm known as "odd-even sort" [48]. Specifically, denote the current qubit configuration as $\vec{x} = (x_1, x_2, \ldots, x_N)$ and the target configuration as $\pi(\vec{x}) = (\pi(x_1), \pi(x_2), \ldots, \pi(x_N))$. In the first time step, qubits x_i and x_{i+1} are swapped if $\pi(x_i) > \pi(x_{i+1})$ for every odd i. In the second time step, qubits x_i and x_{i+1} are swapped if $\pi(x_i) > \pi(x_{i+1})$ for every even i. These steps are repeated until $\vec{x} = \pi(\vec{x})$. The odd-even sort is the time-optimal method of sorting qubits in 1D, with a worst-case run time of N time steps.

C. 2D regular array reconfiguration

Consider a rectangular 2D grid of qubits, where every qubit can be swapped with any of its four neighbors. This corresponds to an ion trap tiled with X junctions, with one qubit per junction [4]. In this case, qubit swap can still be performed using a swap sequence as in Fig. 3(b). Alternatively, it can be executed as a sequence of shuttling steps, as illustrated in Fig. 5(a). Regardless of the physical implementation, we can still consider a switchable qubit swap to be a logical primitive and the swap can be either horizontal (in a row) or vertical (in a column).

The 1D reconfiguration algorithm can be extended onto a regular 2D grid as follows [49–51]. Consider a grid of $N = m \times n$ qubits arranged in a 2D array of $m \times n$ junction zones. We enumerate zones as (i, j), where $i = 1, 2, \ldots, m$ and $j = 1, 2, \ldots, n$. A zone is considered "odd" if $i + j$ is odd and "even" otherwise, as shown in Fig. 5(b). We achieve arbitrary qubit reconfiguration as follows. First, we rearrange the qubits in every row in parallel such that, for every column, the target row of every qubit is different, which is always possible due to Hall's marriage theorem [51]. This rearrangement proceeds by horizontal

odd-even swap as outlined in Sec. III B and thus takes at most m time steps. Afterward, we rearrange the qubits in every column in parallel to place every qubit in the target row. This is executed by vertical odd-even swap and takes at most n time steps. Finally, we proceed with the final row-wise rearrangement, which takes at most m time steps. Thus, the WISE architecture allows for arbitrary permutation of $N = m \times n$ qubits in 2D in at most $2m + n$ time steps [52] and two types of zones. Assuming that each junction zone contains $N_{\mathrm{de/jz}} \sim 20$ dynamic electrodes [53], the reconfiguration requires $4 \times N_{\mathrm{de/jz}} \sim 80$ DACs, and N bits of information per time step.

D. Realistic 2D array reconfiguration

The regular 2D array of junction zones described in Sec. III C is not a preferred arrangement of zones for several reasons. First, junctions typically require a larger electrode count and footprint than linear segments. Thus, allocating one junction per qubit may be exceedingly costly and it will likely be preferable to operate with $k > 1$ qubits per junction. Second, quantum computing requires not just an array of individual qubits but an array of qubit chains to facilitate multiqubit gates. Thus, any reconfiguration method must consider how qubits come together in larger chains. Finally, a QPU might require specialized zones, such as "ion-loading zones" [54] or "qubit-readout zones" [55]. A practical reconfiguration method should thus allow for different zone types.

Fortunately, the methods presented above can be easily extended to such more realistic traps. As an example, consider an N-qubit 2D ion trap with $k = 6$ qubits per junction. The trap contains $N_{\mathrm{jz}} \approx N/k$ junction zones, which facilitate vertical swaps, and $N_{\mathrm{lz}} \approx N(1 - 1/k)$ linear zones, which facilitate horizontal swaps, and where quantum gates are implemented. An illustration of the trap and qubits is shown in Fig. 6.

In order to perform arbitrary reconfiguration using a dynamic electrode parallelization, we divide the trap into four zone types—junction odd, junction even, linear odd, and linear even—arranged as shown in Fig. 6. Since every odd zone only neighbors even zones, row-wise and columnwise odd-even sort operations remain unchanged. We assume once again that each linear zone contains $N_{\mathrm{de/lz}} \sim 10$ dynamic electrodes, while each junction zone contains $N_{\mathrm{de/jz}} \sim 20$ dynamic electrodes. The algorithm for sorting this more realistic trap of $m \times n$ qubits using parallel dynamic control is illustrated in Fig. 7. In an array of $N = m \times n$ qubits, each step can be executed by using $4 \times N_{\mathrm{de/lz}} + 4 \times N_{\mathrm{de/jz}} \sim 120$ dynamic DACs and N bits of information. The algorithm has an approximate run time of $2m + kn$ time steps. Similar techniques can be applied to handle longer chains, different numbers of qubits per junction, and specialized zones.

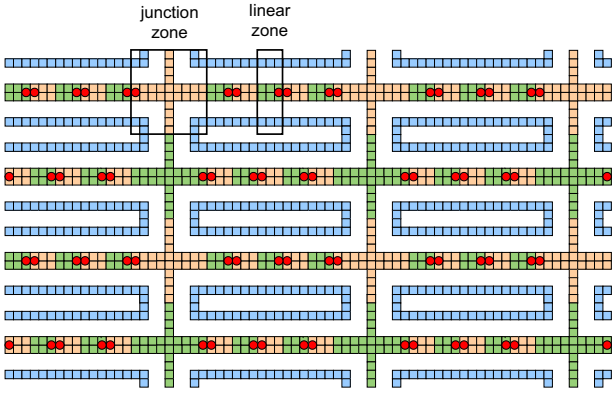

FIG. 6. A schematic illustration of a 2D ion trap with $k = 6$ qubits per junction. We distinguish four zone types: junction odd (orange), junction even (green), linear odd (green), and linear even (orange). Before and after reconfigurations, qubits are stored in the linear zones in small chains to allow for multiqubit quantum gates.

E. Performance

The total duration of qubit reconfiguration depends on the time it takes to execute a single step of parallel swaps. This, in turn, depends on the duration of other primitives, such as ion splitting, merging, rotation, and/or linear shuttling [in the case of a 1D swap in Fig. 3(b)] or junction shuttling [in the case of a 2D swap in Fig. 5(a)]. As an order of magnitude, we assume a qubit swap duration of $t_0 = 100\ \mu\mathrm{s}$, which is representative of the capabilities of today's small-scale quantum computers [57,58]. However, significantly faster swaps should become feasible in the future [59].

Furthermore, in a 2D array, the reconfiguration time depends on the number of zones $N = m \times n$, as well as on the number of qubits per junction k. Finally, the reconfiguration time crucially depends on the target permutation, with longer-range connectivity requiring more reconfiguration steps.

As a pessimistic estimate, we consider the worst-case duration of arbitrary qubit routing. For the algorithm in Sec. III D, it is easy to verify that, for any given (N, k), the reconfiguration time is minimized when $m = \sqrt{kN/2}$ and $n = \sqrt{2N/k}$. In that case, arbitrary reconfiguration of an N-qubit array requires approximately $\sqrt{8kN}$ swap steps.

With those assumptions, Fig. 8 (a) shows the maximum time necessary to perform arbitrary qubit routing in a realistic 2D array. We find that, with $k = 6$, we achieve an arbitrary reconfiguration of a 1000-qubit array in $t_r = 22$ ms. To put that in more meaningful terms, we compute the memory error associated with qubit reconfiguration, assuming the error model measured for $^{43}Ca^{+}$ clock qubits in Ref. [13]. The result is shown in Fig. 8 (b). We find that for $k = 6$ and $N = 1000$, we can expect a memory error (per qubit per reconfiguration) of 2×10^{-5}, significantly

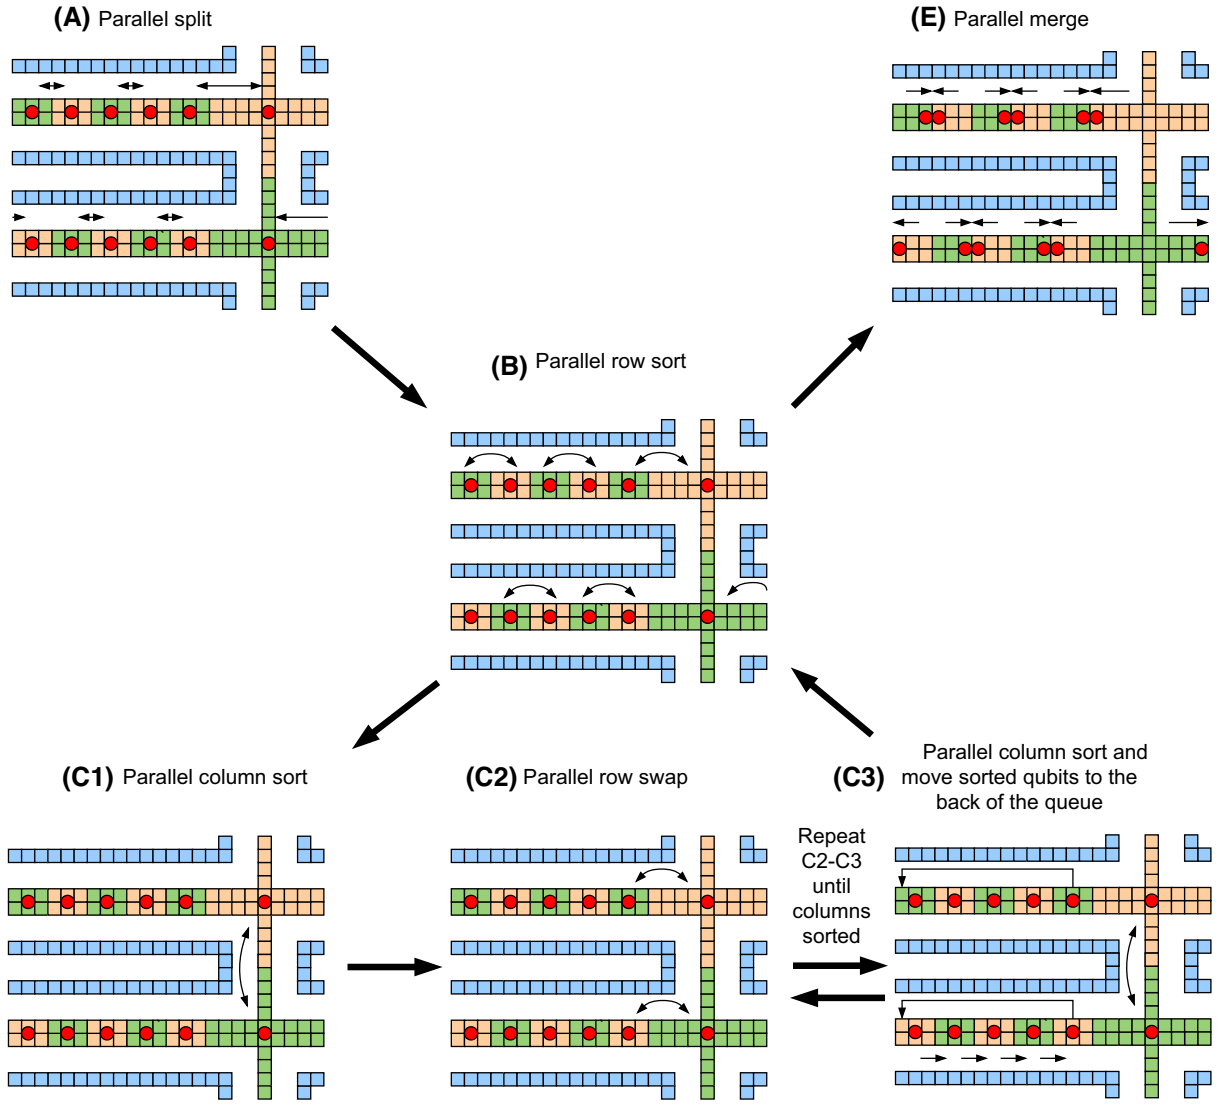

FIG. 7. The sort algorithm for a realistic 2D ion trap with $k = 6$ qubits per junction. (A) First, two-qubit chains in linear segments are split in parallel such as to place every qubit in a different zone. (B) Parallel row-wise odd-even sort is used exactly as if the array was regular, rearranging the qubits such that qubits in the same column have unique destination rows. This takes at most m time steps. (C) Parallel column odd-even sort is applied on every kth qubit in sequence. This takes $k \times n$ time steps [56]. (D) Parallel row-wise odd-even sort is used exactly as if the array was regular. (E) Finally, qubits are remerged into two-qubit chains between linear zones in parallel.

below two-qubit gate errors in any quantum computing platform today.

This demonstrates that, despite limited freedom of operation, parallel dynamic control is compatible with high-fidelity reconfiguration at scale. At the same time, as each step requires N bits of information to specify the switch settings, the overall required data flow rate between the controller and the chip is N/t_0, which evaluates to 10 Mbit/s for a 1000-qubit chip. It is easy to send this amount of data over a single serial line, offering a major improvement over the standard approach. We emphasize that the numbers in this section describe routing times necessary for arbitrary all-to-all connectivity. In practice, judicious qubit mapping and algorithm choice can significantly decrease the overheads associated with qubit reconfiguration. Furthermore, specific reconfigurations can be executed much faster using less generic routing algorithms.

F. Shim electrodes

The quasistatic shim electrodes play a passive role in qubit reconfiguration: they are charged to the desired values before reconfiguration begins and are disconnected during ion transport. The primary function of shim electrodes is to minimize the stray-field offsets between zones

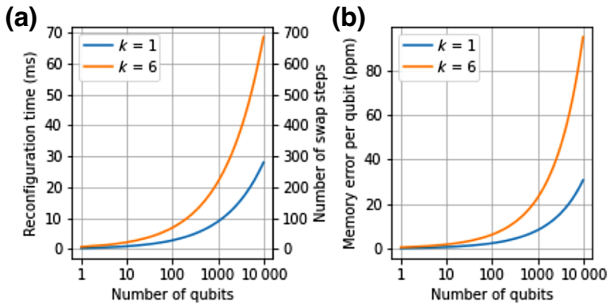

FIG. 8. The 2D qubit-array reconfiguration time and error. (a) The approximate worst-case time and number of time steps for reconfiguring a 2D array of qubits in the optimal configuration, assuming a swap time of $t_0 = 100\,\mu$s. (b) The associated memory error, assuming the decoherence error model in Ref. [13].

such that the same dynamic swap waveform leads to successful qubit swaps in different zones. Among all the transport primitives, ion splitting is typically most sensitive to the stray-field setting, requiring axial stray-field error of $E \lessapprox 5$ V/m [60]. However, larger stray-field errors can be tolerated as long they are radial. Finally, differences in stray fields can lead to heating and even ion loss during transport operations such as crystal rotations or junction transport. However, these can typically tolerate errors as large as $E \sim 100$ V/m [58].

The second role of shim electrodes is to passively compensate for differences in dynamical electrode moments in different zones due to, e.g., fabrication imperfections. While quasistatic shims cannot null out dynamical errors at every point in time, they can be used for fine tuning the potentials along critical points in a transport trajectory, e.g., to prevent frequency crossings during crystal rotations [61].

Future experiments must verify the minimum number of necessary shim electrodes per zone. In subsequent sections, we conservatively assume that each zone contains $N_{\mathrm{se}/z} = 10$ shims [62]. We anticipate that improvements in control design may allow for as few as $N_{\mathrm{se}/z} \approx 3$ shims per zone, controlling 6° of freedom (e.g., three field terms and three curvature terms) at one point in space during a two-qubit swap.

IV. TRANSPORT-ASSISTED QUANTUM GATES

Consider an N-qubit trap, with a global spatially inhomogeneous qubit drive coupling to qubits in all zones. In order to perform transport-assisted quantum gates, we engineer a situation where, for every qubit i, the gate Rabi frequency Ω_i can be adjusted by adjusting the qubit position r_i. There are several different ways in which this can be implemented in hardware.

For example, one way to implement this scenario is by using integrated microwaves. A simplified case with a linear trap and a single conductor is shown in Fig. 9(a). A

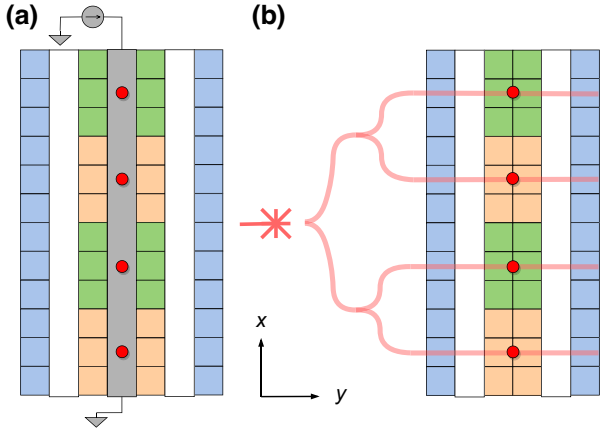

FIG. 9. (a) A schematic illustration of a linear ion trap supporting laser-free transport-assisted gates. Individual qubits (red) are held above a microwave conductor (gray) running along x across all zones. When the current is applied, all qubits experience a position-dependent magnetic field, allowing for transport-assisted gates. (b) A schematic illustration of a linear ion trap supporting laser-based transport-assisted gates. A laser beam is passively split into multiple channels, each delivered to a separate qubit in a separate zone. Once the laser is turned on, the qubits experience a position-dependent laser intensity, allowing for transport-assisted gates.

trace carrying current I along the trap axis x creates a magnetic field $B \propto I$, near-resonant with the qubit frequency, in every zone in parallel. This leads to qubit coupling with Rabi frequency $\Omega(y) \propto I \times y$, where y is the distance from the trap axis. Furthermore, it is possible to achieve $\Omega(y = 0) \approx 0$ by, e.g., exploiting polarization selection rules [63], or using multiple parallel microwave lines [64,65]. Thus, we obtain $\Omega(y) \approx \alpha I y$, where α is a constant. Therefore, single-qubit interactions can be switched off by placing the qubit at $y = 0$ and effectively modulated by moving the qubit along y [66,67].

Another way to implement this scenario is to use laser light, e.g., as illustrated in Fig. 9(b). A laser beam, near-resonant with the qubit transition, is coupled into a trap-integrated waveguide [68–70] and passively split between multiple paths of comparable intensity [71]. Inside each zone, a grating coupler is used to focus the light near the qubit location at $x = 0$ to a spot size w [72], which leads to position-dependent Rabi frequency $\Omega(x) \propto \Omega_0 \exp(-x^2/w^2)$. Thus, single-qubit interactions can be modulated by moving the qubit around $x = \pm w/\sqrt{2}$, and switched off by placing the qubit at $x \gg w$.

A. Dynamic electrode parallelization

We can leverage dynamic electrode parallelization for single-qubit gates, regardless of the implementation, as follows. We set up the dynamic waveforms such that, when $s_i = 1$, the qubit in zone i is moved to position r_1 with

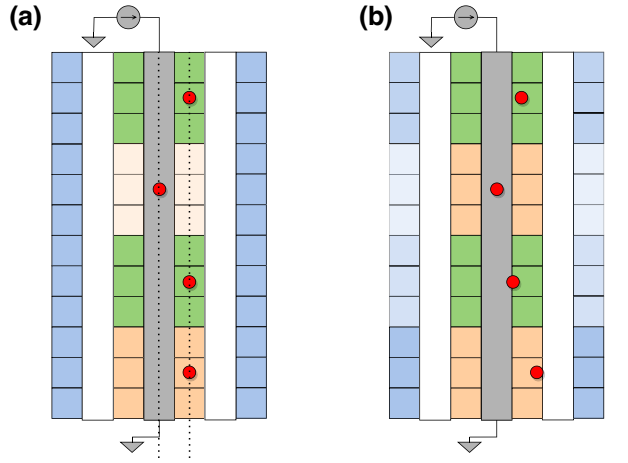

FIG. 10. An illustration of transport-assisted laser-free single-qubit gates using (a) dynamic electrode parallelization and (b) shim demultiplexing. In (a), a switch-select word s selects which zones are active ($s = 1$, solid color) and which are inactive ($s = 0$, semitransparent). Once the switches are set, a dynamical waveform moves the qubits in active zones to locations where they experience nonzero Rabi frequency Ω, while qubits in inactive zones are moved to locations where $\Omega_{1q} \approx 0$. Afterward, a current pulse is applied to the central conductor, executing the target operation. Finally, the qubits are returned to their original locations. In (b), the multiplexer is first turned on, setting shim electrodes in different zones to different values (shades of blue). This moves qubits in different locations to different zones, adjusting the Rabi frequency locally. Afterward, the multiplexer is turned off and a current pulse is applied to the central conductor, executing the target operations.

Rabi frequency Ω. On the other hand, if $s_i = 0$, the qubit is moved to position r_0, where $\Omega = 0$. An illustration of the scenario is shown in Fig. 10(a).

With the drive resonant with the qubit frequency, this generates a Hamiltonian $H_\phi = \sum s_i \Omega \sigma_\phi^{(i)}$, where $\sigma_\phi = \cos\phi\sigma_x + \sin\phi\sigma_y$, ϕ is the global drive phase, and Ω is the global drive strength. If the drive is detuned from the qubit transition frequencies by $\delta \gg \Omega$, we instead obtain a Hamiltonian $H_z = \sum_i s_i \sigma_z \Omega^2/(2\delta)$ [73].

Due to unavoidable experimental imperfections, it is impossible to ensure that $\Omega = 0$ when $s_i = 0$. Thus qubits in zones with $s_i = 0$ will always experience some residual Rabi frequency $\Omega_\varepsilon = \varepsilon \times \Omega$. This, if uncorrected, results in approximately ε^2 infidelity per qubit per gate, which can be a challenge for transport-assisted gates. We mitigate this error using a two-step approach. First, we make use of any knowledge of ε to coherently undo undesired rotations with subsequent operations. Second, we use composite pulse schemes to cancel out any residual but unknown systematic errors. For example, using the SK1 pulse sequence [74], one can achieve single-qubit gate infidelities $< 10^{-5}$ while addressing imperfections as large as $\varepsilon \approx 0.1$, which is readily achievable in practice [75–77]

In addition to single-qubit gates, dynamic electrode parallelization allows us to switch two-qubit interactions on and off. When the global drive is turned on to generate a state-dependent force, we can write the effective entangling interaction between qubits i and $i+1$ as $H = \Omega\sigma_\phi^{(i)}\sigma_\phi^{(i+1)}$. This interaction only occurs if the qubits are in the same potential well [78]. Therefore, we set up dynamical waveforms such that, if $(s_i, s_{i+1}) = (1,1)$, qubits $(i, i+1)$ are merged into the same well prior to the two-qubit gate, while if $(s_i, s_{i+1}) = (0,0)$, they remain in separate zones. Thus, we generate a Hamiltonian $H_2 = \sum_i s_{2i-1}\Omega\sigma_\phi^{(2i-1)}\sigma_\phi^{(2i)}$.

Hamiltonians H_ϕ and H_z implement single-qubit gates on arbitrary subsets of qubits in parallel, while H_2 implements two-qubit gates on arbitrary subsets of qubits in parallel. Together, they can be used to implement a universal primitive gate set [79,80] and thus universal quantum computation. Explicitly, a quantum circuit is decomposed onto N_g different basis gates and implemented by using cycles of N_g time steps, each implementing one of N_g gate layers.

Finally, we note that while the local operation phase is set by the global drive phase, it can vary zone to zone, e.g., when the total length of the current conductor becomes comparable to the wavelength (in the microwave implementation), or due to length offsets between integrated waveguides routed to different zones (in the laser implementation). However, these phase variations can be accounted for and compensated using local rotations, e.g., by employing phase-insensitive operation blocks [81]

B. Shim demultiplexing

While dynamic electrode parallelization can be readily applied to circuits with a small number of basis gates N_g, it is prohibitively costly for circuits with $N_g \gg 1$. This creates a challenge in the NISQ regime, where the ability to implement a continuous variable-angle gate set is beneficial [82–84]. Shim demultiplexing can be used to overcome this limitation of dynamic electrode parallelization by allowing continuous local control of Ω. This leads to the ability to perform different single-qubit gates in parallel.

In this approach, instead of moving the qubits between two fixed positions r_1 and r_0, we adjust r continuously using demultiplexed shim electrodes, as illustrated in Fig. 10(b). After the shims are set, we turn off the demultiplexer and execute the operation by applying the global qubit drive. The shims are then set to new values before the next gate layer or transport layer.

Thus, shim demultiplexing allows us to implement a Hamiltonian $H_\phi = \sum_i \Omega_i \sigma_\phi^{(i)}$, where Ω_i is locally adjustable. This executes parallel rotations with locally adjustable angles across the processor using only a global drive. Combined with global two-qubit gates implemented

as before, we obtain a powerful toolbox for NISQ algorithms and beyond.

There are several drawbacks associated with transport-assisted gates via shim demultiplexing. First, as electrode voltages have to be adjusted before every gate layer, demultiplexing slows down the effective clock rate of the quantum computer. Second, the lack of local phase control restricts the use of composite pulses to reduce systematic errors. However, local phase control can be restored using demultiplexed integrated in-phase and quadrature (IQ) mixers, implementing the control scheme proposed in Ref. [85].

V. PHYSICAL IMPLEMENTATION

Ion-trap fabrication and operation are fundamentally compatible with electronic integration and CMOS processes. However, as discussed in Sec. I, care must be taken with regard to issues of power dissipation, footprint, and data flow. Furthermore, low-complexity electronics are preferred given the relative immaturity of cryogenic CMOS design tools. Finally, care must be taken to avoid introducing additional error sources, such as rf pickup or stray fields.

In this section, we sketch out the structures that need to be integrated into the QPU to implement the WISE architecture. We discuss the design and main performance metrics of the dynamic electrode switch network (Sec. V A) and the shim-demultiplexing network (Sec. V B). We discuss the main error sources and, in Sec. V C, summarize why the physical implementation of the WISE architecture indeed alleviates all the main issues of integrated DACs.

A. Dynamic electrode switch network

1. Switch implementation

The basic building block of the dynamic electrode switch network is shown in Fig. 11. The bit-select word $s = (s_1, \ldots, s_N)$ is loaded via a serial line into a parallel shift register. Afterward, the entry s_i is used to select whether the electrode $j = 1, 2, \ldots, N_{\mathrm{de}/z}$ in zone i of type t (odd or even) is connected to DAC $V_{t,1,j}$ or $V_{t,0,j}$. The switch itself is built using two transmission gates, each implemented using an NMOS-PMOS transistor pair and an inverter. The DAC outputs are only filtered off chip.

2. Switch network

While in principle one switchable electrode per zone suffices to implement a switchable swap (Appendix B) the most flexible wiring is illustrated in Fig. 12, where all $N_{\mathrm{de}/z}$ dynamic electrodes in every one of N zones are connected to switches. This allows for maximum flexibility in the waveform design, as every action can have a completely separate custom waveform set. As all the switches in zone

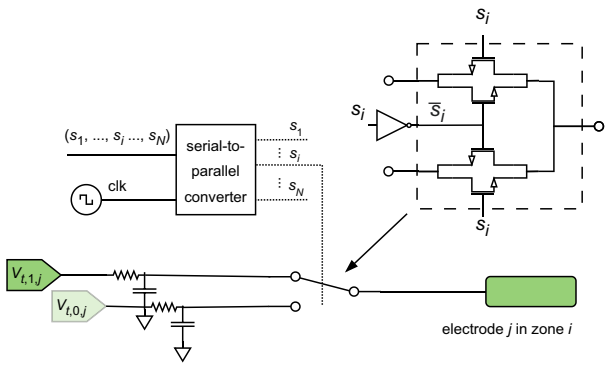

FIG. 11. The basic implementation of a dynamic electrode switch network. The bit-select word s is loaded into a serial-to-parallel converter, which connects to a single-pole double-throw switch, selecting whether the electrode is active ($s = 1$, solid) or inactive ($s = 0$, transparent). The switch is implemented using two transmission gates, each implemented as a pair of transistors.

i are controlled by the same digital line s_i, the length of the bit-select word s is N.

Since each electrode is always connected to at least one DAC, there is no requirement to integrate capacitors, as off-chip capacitors suffice to provide a dc-electrode rf ground. If necessary, residual rf pickup on dynamic electrodes can be reduced by optimizing transmission gate resistance.

3. Operation timing

The switch network is operated as follows. Consider a single operation step, either a swap (duration t_0) or a transport-assisted gate (duration t_{tag}). While the DACs are playing the required waveforms, the word s of length N bits describing the next transport step is loaded through a serial link into a parallel register. At the end of the transport step, the DAC voltages are set such that $V_{t,0} = V_{t,1}$, i.e., each electrode is at the same voltage regardless of the switch setting. Then, data from the parallel register are transferred onto the switches and the next transport step begins. Note that since $V_{t,0} = V_{t,1}$, setting the switches does not affect electrode voltages and thus does not lead to motional excitation of the ions. Thus, assuming that $t_{\mathrm{tag}} > t_0$, as long as the interface supports a data rate of more than N/t_0, the switch network does not bottleneck the device performance. With $t_0 = 100\ \mu$s as in Sec. III E, we find that the required data flow rate is 10 N kbit/s. Thus, a single serial line with a bandwidth of 50 Mbit/s suffices for $N \approx 5000$ qubits.

B. Shim demultiplexing

1. Demultiplexing network

A high-level illustration of the shim-demultiplexing network is shown in Fig. 13(a). The QPU consists of multiple

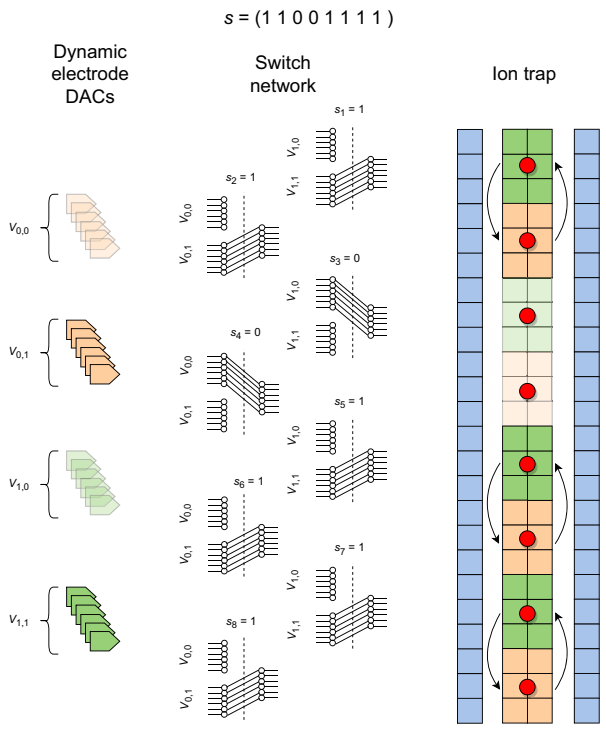

FIG. 12. The implementation of parallel dynamic control. The linear 1D trap (right) is undergoing an odd swap. Each zone consists of dynamic electrodes, shown in green for odd zones and in orange for even zones. Each dynamic electrode is connected through an individual single-pole double-throw switch to one of two DACs. All switches for zone i are controlled by the same bit s_i. Electrodes in odd zones are connected either to the DAC group $V_{1,0}$ (if $s_i = 0$) or to $V_{1,1}$ (if $s_i = 1$). Likewise, electrodes in even zones are connected either to $V_{0,0}$ (if $s_i = 0$) or to $V_{0,1}$ (if $s_i = 1$). In the above image, the electrode color matches the color of the DAC to which it is currently connected. In this (toy) example trap, there are $N_{\mathrm{de/lz}} = 6$ dynamic electrodes per linear zone and hence $N_{\mathrm{de/lz}} = 6$ DACs per group, for a total of $N_{\mathrm{dDAC}} = 4 \times N_{\mathrm{de/lz}} = 24$ dynamic electrode DACs.

analog $1 : M$ demultiplexers that serve to connect the shim DAC to the shim electrodes one at a time. A single on-off signal clocks all the demultiplexers on the QPU. Each shim electrode is connected to an integrated capacitor after the demultiplexer.

Figure 13(b) illustrates the demultiplexer circuit in more detail. On every clock cycle, on-chip digital logic increments an L-bit register x by one, starting from $x = 0$. The output $x = (x_1, x_2, \ldots, x_L)$ of the register is sent to a digital demultiplexer with $M = 2^L$ outputs. The circuit sets the output $y_x = 1$, while the rest of its outputs are set to 0. Demultiplexer outputs y_x act as switch control signals, connecting the shim DAC to electrode x if and only if $y_x = 1$. The individual switches are implemented as transmission gates as before.

After the switch is closed, we wait for time t_{ec}, allowing the electrode to charge, and then advance the clock by one cycle. The procedure is repeated in total M times, charging N_{se} shim electrodes in time $t_{\mathrm{sc}} = t_{\mathrm{ec}} \times M$ using N_{se}/M DACs and a single clock line. When $x = 0$ and all the electrodes are charged, we pause the clock signal and execute the subsequent operation layer with shim electrodes floating.

2. Integrated capacitors

Operation with floating electrodes requires integrated capacitors to provide an rf ground. To make the design practical, capacitances must be sufficiently small to avoid footprint bottlenecks highlighted for integrated DACs.

When disconnected, the voltage on each shim electrode oscillates with amplitude $V_{s,\mathrm{rf}} = V_{\mathrm{rf}} \times C_{s,\mathrm{rf}}/C$, where V_{rf} is the voltage on the rf electrode, $C_{s,\mathrm{rf}}$ is a parasitic capacitance between the rf and the shim electrode, and C is the capacitance of the integrated capacitor. This leads to additional micromotion, especially if there are variations of $C_{s,\mathrm{rf}}/C$ between electrodes.

Consider zones of area A_z, each containing $N_{\mathrm{se}/z}$ shim electrodes, each connected to a capacitor of area A_c. As long as $N_{\mathrm{se}/z} \times A_c < A_z$, the capacitors do not represent a significant bottleneck on the device size. Thin-film planar capacitors allow for $C/A_C \sim 3$ fF/μm^2 [29]. Thus, a capacitor with dimensions $A_c = 100\ \mu\mathrm{m} \times 100\ \mu\mathrm{m}$ provides a capacitance of $C \sim 30$ pF. Assuming that $A_z = 400\ \mu\mathrm{m} \times 400\ \mu\mathrm{m}$, we can connect $N_s \approx 16$ shims per zone to such capacitors without increasing the device footprint. Assuming that $V_{\mathrm{rf}} = 100$ V and $C_{s,\mathrm{rf}} = 1$ fF [37], we expect an rf pickup of $V_{s,\mathrm{rf}} = 3$ mV. This should be sufficient for high-fidelity quantum operations, especially since the pickup can be very uniform with integrated capacitors [86]. If necessary, the rf pickup can be further reduced by using larger values of C, e.g., by increasing the device footprint, increasing the number of capacitor layers, or using vertical trench capacitors [29,87].

3. Timing and I/O count

The shim-demultiplexing architecture allows us considerable freedom in trading off the shim charging speed, network complexity, performance, and the number of inputs. Increasing the multiplexing order M decreases the number of shim DACs but increases the shim charging time t_{sc}, slowing down the effective clock speed of the quantum computer. While decreasing the electrode charging time t_{ec} directly improves the clock speed, this must be traded off against the additional design and fabrication complexity of high-speed electronics. In addition, faster electrode charging can lead to additional motional excitation of ions, necessitating subsequent recooling. Finally, depending on the number of primitive gates N_g, single-qubit operations will either be executed using dynamic electrode parallelization (which does not require frequent

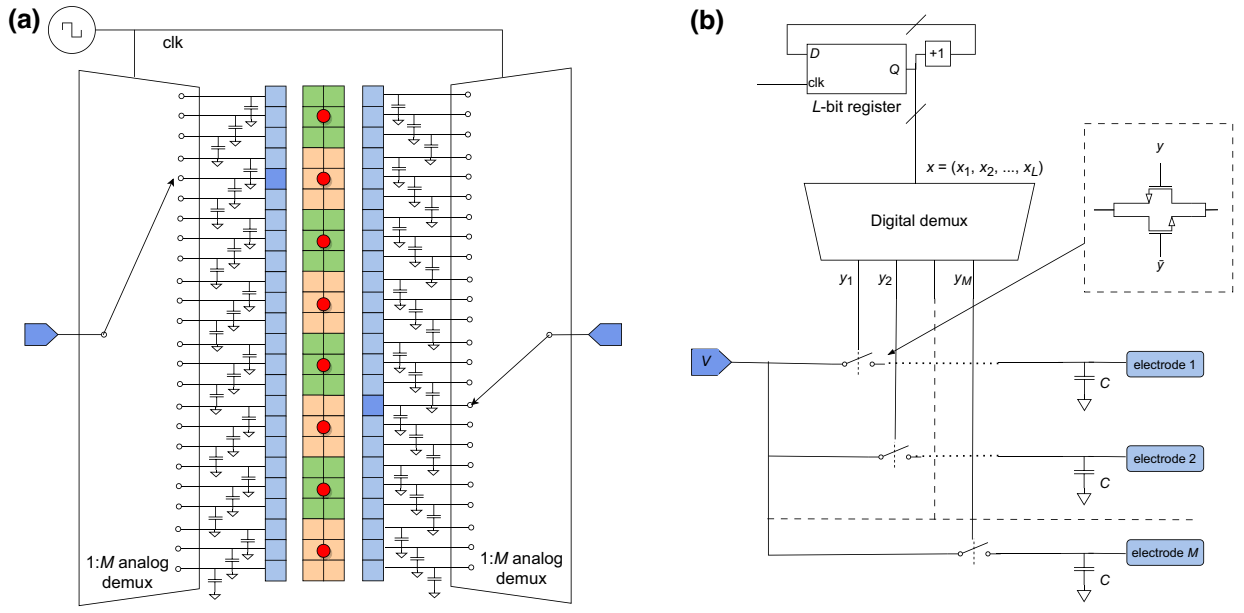

FIG. 13. (a) The basic architecture for shim demultiplexing. Each shim DAC (blue) is connected to one of M possible shim electrodes (blue) via an analog demultiplexer. An input clock controls all the demultiplexers on the chip, selecting whether they are turned on or off. Each shim electrode is connected to ground through an integrated capacitor. (b) Details of the analog-demultiplexer implementation. An L-bit register (one for the whole chip) advances x by one every clock cycle. The value of x is sent through a digital demultiplexer and controls a transmission-gate switch, selecting which of the $M = 2^L - 1$ electrodes is connected to the DAC. When $x = 0$, all the electrodes are disconnected and the qubit reconfiguration and/or quantum gates can proceed.

shim recharging) or shim demultiplexing (which requires shim recharging between every gate layer).

As an example, consider a $N = 1000$-qubit chip with $N_{\mathrm{se}/z} = 10$ shims per zone, reconfiguration time $t_r = 22$ ms, and $N_l = 10$ transport-assisted gate layers between reconfiguration steps, each requiring all shim electrodes to be recharged. Using a charging time of $t_{\mathrm{ec}} = 3$ μs and $M = 128$, the shim charging time is $t_{\mathrm{sc}} = M \times t_{\mathrm{ec}} = 384$ μs. Thus, shim charging takes $N_l \times t_{\mathrm{sc}}/t_r \approx 20\%$ of system time—significant, but not a run-time bottleneck. At the same time, the total of $N_{\mathrm{se}} = N_{\mathrm{se}/z} \times N \sim 10,000$ shim electrodes will require $N_{\mathrm{se}}/M \sim 80$ shim-voltage inputs, which is easily achievable in today's systems. Finally, the charging time of $t_{\mathrm{ec}} = 3$ μs is long enough to not introduce significant electronics design challenges and to limit motional excitation of the ions.

4. Voltage errors

In our architecture, shim-voltage errors can arise on top of errors from the DAC. The three main sources are charge injection, electrode discharging, and the photoelectric effect.

Charge injection refers to an additional voltage that appears once the switch is opened due to the redistribution of charge accumulated on the control transistors. This results in a fractional voltage shift of approximately C_t/C, where C_t is the transistor gate-source capacitance. Assuming a shim capacitance $C = 30$ pF, a transistor capacitance $C_t = 30$ fF, and a typical shim stray field $E_s = 200$ V/m [87], the charge-injection effect is at the level of $E_s \times C_t/C \sim 0.2$ V/m. The effect is expected to be systematic and thus can be compensated by calibration. Furthermore, the magnitude of charge injection can be reduced by reducing C_t through advanced fabrication processes.

Electrode discharging means that electrode voltage V decays over time once the switch is opened. The main loss channels are due to current flow through the capacitor and the transistor. We will assume the combined decay time constant of $\tau = 180$ s, based on results reported in Ref. [42] for cryogenic transistors and thin-film alumina capacitors of comparable values.

To evaluate the effect of electrode discharging on quantum operations, consider once again a $N = 1000$ qubit chip with reconfiguration time $t_r = 22$ ms (as in Sec. III E) and two-qubit gate time $t_{2q} = 100$ μs. Assuming once again that $E_s = 200$ V/m, we find a shim field drift of $E_s \times (1 - e^{-t_r/\tau}) \approx 24$ mV/m during reconfiguration. As the transport operation most sensitive to stray-field offsets—qubit split—can tolerate axial field offsets of approximately 5 V/m [60], this level of drift is fully compatible with a low-error low-heating reconfiguration. The most stringent requirement on electrode voltage stability is that stray-field drifts can lead to gate-mode frequency drifts,

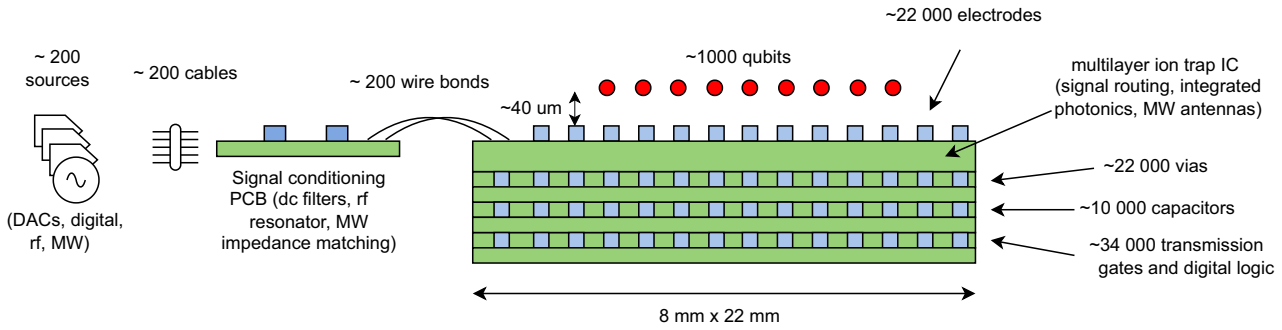

FIG. 14. An illustration of the possible electrical wiring of a 1000-qubit chip. The QPU (right) combines an ion-trap integrated circuit with integrated capacitors and switches and requires a footprint of 8 mm × 22 mm (excluding interconnects). The QPU is controlled using approximately 200 electrical inputs, delivered from approximately 200 off-chip sources via wire bonds.

causing two-qubit gate errors. During a two-qubit gate, we expect a stray-field drift $E_s \times (1 - e^{-t_{2q}/\tau}) \approx 0.1$ mV/m. Assuming a gate-mode anharmonicity of $df/dE = 1$ kHz per V/m [88], this corresponds to a motional frequency drift of $df = 0.1$ Hz, resulting in a negligible two-qubit gate error of approximately $df^2 \times t_{2q}^2 \sim 10^{-10}$ [89]. Thus, we expect that electrode discharging will not bottleneck either qubit reconfiguration or quantum gate fidelity [90].

Due to the photoelectric effect, the voltage held on the shim electrode—and hence the corresponding stray field—is perturbed by laser light. This effect is most pronounced for short-wavelength light (blue and ultraviolet) but can be meaningful for longer-wavelength radiation as well [91]. The impact of the photoelectric effect can be minimized by carefully managing laser scatter and electrode contamination, as well as by using laser-free gates instead of laser-based gates.

C. Comparison to integrated DACs

In a nutshell, our solution replaces the problem of integrating filtered DACs with the problem of integrated switches and small capacitors. Thus, it is only a reasonable solution if the latter issue is easier than the former. Fortunately, as highlighted earlier in the section, our architecture alleviates all the main concerns of integrated DACs:

(1) *Power dissipation*. Unlike a DAC, a transmission-gate switch results in negligible additional power consumption.
(2) *Data bandwidth*. Unlike a DAC, a transmission-gate switch requires only a single bit of information per operation. Thus, the data-streaming rate is reduced to $1/t_0 = 10$ kbit/s per zone, allowing for a large number of switches with a single serial line.
(3) *Footprint*. Unlike a DAC, a transmission-gate switch requires a footprint only a few times larger than that of a single transistor. The transistor size, in turn, depends on the required voltage level, as well as the details of the fabrication process. Assuming a transmission gate footprint of $A_{\text{tg}} \approx 50\ \mu\text{m} \times 50\ \mu\text{m}$, which is realistic for ± 10 V logic levels [92], we find that the switching networks introduce little to no overheads on the device footprint (Appendix A). Furthermore, all large filter capacitors (with typical values of approximately 1 nF) are placed off chip in our architecture. While WISE architecture requires integrated capacitors to eliminate rf pickup on shim electrodes, those can be small ($C \sim 30$ pF) and their footprint manageable ($A_C \approx 100\ \mu\text{m} \times 100\ \mu\text{m}$).

All in all, integrating switches at high density is considerably more practical than integrating DACs at high density.

VI. WIRING A 1000-QUBIT CHIP

Let us now put all the numbers from the previous sections together to discuss the wiring of a 1000-qubit trapped-ion quantum computer. Table I in Appendix A shows the main system parameters, collected from the previous sections. In this section, we summarize the most important quantities. The result is illustrated in Fig. 14.

Our device is a 2D array of $m \times n = 54 \times 19 = 1026$ zones, consisting of $N_{\text{lz}} = 855$ linear zones and $N_{\text{jz}} = 171$ junction zones. The device holds $N = 1026$ qubits ($k = 6$ qubits per junction) at a height of $h = 40$ um above the surface, in a region of $x \times y = 22 \times 8$ mm (total area $A_t = 164$ mm^2), using $N_{\text{de}} = 11\,970$ dynamic electrodes and $N_{\text{se}} = 10\,260$ shim electrodes. Each shim electrode is connected to one of the $C = 30$ pF thin-film planar capacitors, which occupy a total area of $A_{t,c} = 103$ mm^2 on a buried layer. Another buried layer contains the switching network—a total of $N_{\text{tg}} = 34\,200$ transmission gates, with an area of 86 mm^2—as well as the necessary digital logic.

After the switching network, the dynamic electrodes are connected to $N_{\text{dDAC}} = 120$ off-chip DACs. The shim electrodes are controlled by $N_{\text{sDAC}} = 81$ off-chip DACs via 1 :

TABLE I. System-parameter table for a 1000-qubit chip.

Quantity	Symbol	Formula	Value
Number of qubits = number of zones	N		1026
Number of qubits per junction	k		6
Zone count	$m \times n$	$\approx \sqrt{kN/2} \times \sqrt{2N/k}$	54×19
Number of junction zones	N_{jz}	$n \times m/k$	171
Number of linear zones	N_{lz}	$N - N_{\mathrm{jz}}$	855
Number of dynamic electrodes per linear zone	$N_{\mathrm{de/lz}}$		10
Number of dynamic electrodes per junction zone	$N_{\mathrm{de/jz}}$		20
Number of shim electrodes per zone	$N_{\mathrm{se}/z}$		10
Number of dynamic electrodes	N_{de}	$N_{\mathrm{de/lz}} \times N_{\mathrm{lz}} + N_{\mathrm{de/jz}} \times N_{\mathrm{jz}}$	11 970
Number of shim electrodes	N_{se}	$N_{\mathrm{se}/z} \times N$	10 260
Number of electrodes	N_e	$N_{\mathrm{de}} + N_{\mathrm{se}}$	22 230
Ion height	h		40 μm
Average zone size	$x_z \times y_z$	$10h \times 10h$	400 μm × 400 μm
Chip size (active region)	$x \times y$	$(m \times x_z) \times (n \times y_z)$	22 mm × 8 mm
Total area (active region)	A_t	$x \times y$	164 mm^2
On-chip capacitance density	C/A_C		3 $\mathrm{fF}/\mu\mathrm{m}^2$
On-chip capacitor size	A_C		100 μm × 100 μm
Total area (capacitors)	$A_{t,c}$	$A_C \times N_{\mathrm{se}}$	103 mm^2
Shim capacitance to ground (when floating)	C	$C/A_C \times A_C$	30 pF
Radio-frequency voltage	V_{rf}		100 V
Shim-to-rf capacitance	$C_{s,\mathrm{rf}}$		1 fF
Radio-frequency voltage on shim electrodes	$V_{s,\mathrm{rf}}$	$V_{\mathrm{rf}} \times C_{s,\mathrm{rf}}/C$	3 mV
Number of dynamic switch settings	N_{set}	0 or 1, i.e., stay or swap	2
Number of dynamic DACs	N_{dDAC}	$2 \times N_{\mathrm{set}} \times (N_{\mathrm{de/lz}} + N_{\mathrm{de/jz}})$	120
Multiplexing order	M		128
Number of shim DACs	N_{sDAC}	$N_{\mathrm{se}}/(M-1)$	81
Length of switch-select word	N_{sw}	$\log_2(N_{\mathrm{set}}^N)$	1026 bits
Serial link data rate	dN_s/dt		50 Mbit/s
Switch-select time	t_{ss}	$N_{\mathrm{sw}}/dN_s/dt$	21 μs
Number of transport-assisted gates per dynamic electrode	$N_{\mathrm{tag/de}}$		2
Number of transport-assisted gates per shim electrode	$N_{\mathrm{tag/se}}$		1
Number of transport-assisted gates	N_{tag}	$N_{\mathrm{tag/de}} \times N_{\mathrm{de}} + N_{\mathrm{tag/se}} \times N_{\mathrm{se}}$	34 200
Transport-assisted-gate size	A_{tag}		50 μm × 50 μm
Total area (transport-assisted gates)	$A_{t,\mathrm{tag}}$	$N_{\mathrm{tag}} \times A_{\mathrm{tag}}$	86 mm^2
Charging time per shim electrode	t_{ec}		3 μs
Total shim charging time	t_{sc}	$t_{\mathrm{ec}} \times M$	384 μs
Qubit swap time	t_0		100 μs
Maximum qubit-reconfiguration time	t_r	$2 \times m + k \times n$	22 ms
Transistor capacitance	C_t		30 fF
Typical shim field	E_s		200 V/m
Charge-injection field systematic	dE_c	$\sim E_s \times C_t/C$	∼0.2 V/m
Shim discharging time constant	τ		180 s
Shim field drift during reconfiguration		$E_s \times (1 - e^{-t_r/\tau})$	25 mV/m
Shim field drift during two-qubit gate pulse		$E_s \times (1 - e^{-t_{2q}/\tau})$	0.1 mV/m
Single-qubit gate pulse time	t_{1q}		1 μs
Two-qubit gate pulse time	t_{2q}		100 μs
Transport-assisted single-qubit gate time	$t_{\mathrm{tag},1q}$	$t_{\mathrm{sc}} + t_{1q}$	385 μs
Transport-assisted two-qubit gate time	$t_{\mathrm{tag},2q}$	$t_{\mathrm{sc}} + t_{2q}$	484 μs
Maximum system speed		$1/t_{\mathrm{tag},1q}$	2600 layers per second
Minimum system speed		$1/(t_{\mathrm{tag},2q} + t_r)$	40 layers per second

128 on-chip analog demultiplexers. In addition, approximately ten electrical connections are required to deliver the digital control signals (with a data-streaming rate of 50 Mbit/s) and power the digital logic, as well as to provide global trapping and qubit drive fields. In total, we use approximately 210 electrical connections between the

control electronics and the QPU, comparable to present-day ion-trap setups. Additional interconnects will be used for optical wiring (for laser delivery and readout) but the discussion of these is beyond the scope of this work.

Quantum operations are performed as interleaved layers of qubit reconfiguration and transport-assisted gates, with arbitrary-angle single-qubit rotations implemented by shim demultiplexing. The reconfiguration takes between $t_0 = 100$ μs (for simple nearest-neighbor routing) and $t_r = 22$ ms (for arbitrary large-scale routing). Quantum gates are enabled by global qubit drives, which implement single-qubit gates in $t_{1q} = 1$ μs and two-qubit gates in $t_{2q} = 100$ μs. However, the duration of each transport-assisted gate layer is dominated by the shim charging time of $t_{sc} = 384$ μs. Depending on the operation, the overall circuit execution speed is therefore between $1/(t_{1q} + t_{sc}) = 2600$ circuit layers per second (for sequential single-qubit gates in identical qubit configuration) and $1/(t_r + t_{2q} + t_{sc}) = 44$ circuit layers per second (for sequential two-qubit gates in maximally different qubit configurations).

VII. CONCLUSIONS

We have shown how integrating simple switching electronics into the QPU allows for a 1000-qubit trapped-ion quantum computer to be operated with only approximately 200 DACs. The QPU allows for transport-assisted gates, with low errors and effective all-to-all connectivity.

Today's trapped-ion systems already routinely use approximately 200 DACs per QPU [93]. Furthermore, today's ion traps are routinely made on silicon substrates [94], with multiple experiments demonstrating compatibility with substrate-integrated passive [95] and active [96,97] electronics, as well as CMOS processes [98,99]. Therefore, the proposed design shows that we can increase the scale of trapped-ion quantum computers from today's $N \sim 10$–30 qubits to $N \sim 1000$ qubits without fundamental roadblocks in electronic wiring. This opens the path to useful NISQ-scale quantum computation with trapped ions.

A. Challenges and open questions

First and foremost, we must verify the feasibility of performing the same transport operations in different trap zones using fixed dynamical waveforms and local shims. This may require optimizing chip fabrication to minimize zone-to-zone variations of electrode moments and to shield stray charges. Second, processor zones and waveforms must be carefully designed to minimize zone-to-zone crosstalk, such that a transport operation in one zone succeeds regardless of the waveform being applied to other zones. If necessary, spacing between zones can be increased without increasing the number of dynamic DACs count using "bucket-brigade" or "conveyor-belt" type shuttling [20,21,100–102]. Third, the rf pickup on the dynamic electrodes caused by finite switch impedance may affect the transport operations in undesirable ways, necessitating the development of low-loss switches or the integration of additional capacitors.

The most important verification of shim demultiplexing is to confirm the stability of voltage offsets caused by charge injection. If problematic, on-chip switches can be miniaturized to reduce transistor capacitance; hence the stored charge. This can, however, lead to increased current leakage, in turn shortening the electrode discharging time. The second assumption that requires verification is that high-fidelity quantum gates can be performed with floating shim electrodes. Finally, experiments should establish realistic time scales for electrode charging, in turn guiding the decision of how many shim electrodes can be connected to a single DAC.

Future experiments must also verify the feasibility of high-fidelity transport-assisted gates. Despite proof-of-principle demonstrations, all highest-fidelity gates in trapped-ion systems thus far have been executed with externally modulated localized drives [103–105]. Further work is needed to demonstrate that precise control over local potentials can be used to overcome zone-to-zone variations, enabling high-fidelity parallel gates. Furthermore, fast and robust calibration procedures must be developed to stabilize local stray-field drifts in larger-scale systems [38].

B. Larger-scale systems

While this paper has focused on wiring a 1000-qubit chip, it is interesting to ask if WISE can be applied to larger-scale systems. For example, could a 10^6-qubit chip be wired in the same fashion?

On the face of it, dynamic electrode parallelization can be used to control arbitrarily many qubits from a fixed number of DACs. In practice, this will be limited by, e.g., the parasitic capacitances of the switches and electrodes and the resulting DAC load swings. Another relevant effect is that the rf pickup on $N_{de/dDAC}$ dynamic electrodes connected to one DAC increases in proportion to $N_{de/dDAC}$, affecting the design of DAC filters and, in turn, the transport speed.

On the other hand, shim demultiplexing always requires the number of shim DACs N_{sDAC} and QPU I/O lines to grow in proportion to the qubit number N (for example, we used $N_{sDAC} \sim N/10$ in earlier sections). Scaling the method to $N \gg 1000$ requires decreasing the proportionality factor, which can be achieved by (a) increasing the multiplexing order above $M = 128$ shims per DAC, and (b) decreasing the number of shim electrodes per zone to $N_{se/z} < 10$. While the former increases total shim charging time t_{sc}, this is not an issue if the gates are implemented using dynamic electrode multiplexing and if the

shim discharging time scales are long. If necessary, t_{sc} can be reduced by decreasing the charging time per electrode t_{ec} or by sequentially charging additional on-chip capacitors during quantum operations and then connecting them in parallel to the shim electrodes.

All in all, while WISE can be applied to systems with $N > 1000$ qubits, the architecture has its limitations, which will bound the maximum number of qubits it can efficiently support to some number N_{max}. Regardless of its precise value (which is difficult to estimate), an $N > N_{max}$-qubit QPU can be then constructed as repeating units of N_{max} qubits. While DAC and control system integration become necessary for large enough N, WISE architecture can nonetheless be employed to significantly reduce the number of DACs per qubit compared to the standard approach, making the power dissipation and footprint issues much more tangible.

To summarize, the WISE architecture opens the door to building trapped-ion quantum computers that are 1–2 orders of magnitude larger than today's systems. Looking ahead to even larger devices, there are a number of architectural decisions that will significantly impact system wiring. For example, can we reduce the system resource cost if we specialize the WISE architecture for implementing a specific fault-tolerant quantum error-correction code and, if so, which quantum error correction code is optimal [106,107]? What is the optimal chip size and connectivity for encoding logical qubits [108–110]? Does the code benefit from long-range connectivity and, if so, of what kind [111–113]? Regardless of the exact large-scale design, we believe that the methods behind WISE architecture—after thorough refinement in NISQ-scale devices—can serve as its backbone.

ACKNOWLEDGMENTS

We are grateful to the entire team at Oxford Ionics for their support and helpful discussions.

APPENDIX A: 1000-QUBIT SYSTEM PARAMETER TABLE

A detailed list of parameters of a 1000-qubit system and their derivation is shown in Table I.

APPENDIX B: PARALLEL DYNAMIC CONTROL WITH ONE SWITCH PER ZONE

The implementation presented in the main text assumes that every dynamic electrode is connected to a separate switch. This is a conservative assumption and it should, in fact, be possible to only connect some dynamic electrodes to switches while leaving the remainder of the dynamic electrodes permanently cowired. In this appendix, we present an example implementation of that idea, where every zone features only one switchable electrode.

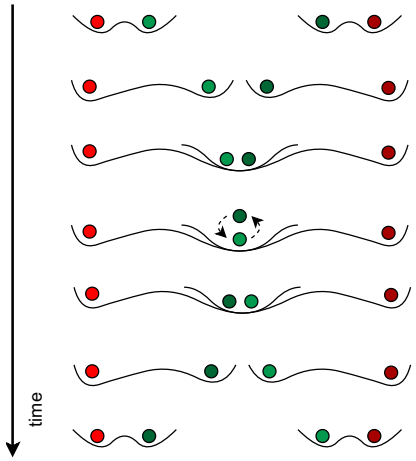

FIG. 15. An extended 1D swap waveform. If there are four qubits (two green and two red) at $t = 0$, the green qubits undergo a swap, while the red qubits remain stationary. Alternatively, we can think of it as a two-qubit swap waveform: if there are two qubits in green locations at $t = 0$, they undergo a swap, while if the qubits are initially at red locations, they remain stationary.

Consider an extension of the 1D swap waveform between zones $(1, 2)$ shown in Fig. 15. In it, green qubits undergo a swap but red qubits remain stationary. An alternative way to use this waveform is as follows. Suppose that there are only two qubits—one in zone 1 and one in zone 2—and that there is a switchable "push field" that can be used to push qubits in each zone either left or right. Specifically, if $s_i = 1$, the qubit in zone i is pushed into the green location, while if $s_i = 0$, it is pushed into the red location. After the switch is set, the waveform is executed. Thus, if $(s_1, s_2) = (1, 1)$, the qubits undergo a swap, while if $(s_1, s_2) = (0, 0)$, they remain in the original order.

This swap waveform can be physically implemented as illustrated in Fig. 16. We use dynamic electrodes cowired to the same DACs to create a double-well potential (the first step of the extended transport waveform) in each zone. At the same time, for each zone, we can select whether the ion is pushed into the red or green well at $t = 0$ by applying either positive voltage $+V$ or negative voltage $-V$ to one electrode, as illustrated in Fig. 16. Afterward, the dynamic DACs play the remainder of the extended transport waveform. This sequence allows us to locally and digitally select whether or not the ion undergoes a swap. Thus, one switch per zone suffices to implement parallel dynamic control. The same method can be used to condition the execution of a 2D swap waveform.

In the basic implementation, the $\pm V$ rails connect to separate DACs and are initially set to the same voltage, V_0. Once each electrode is connected to one rail, the voltage on the $+V$ rail is increased to $V_0 + V$ and the voltage on the $-V$ rail is reduced to $V_0 - V$. This ramping can be done smoothly to avoid undesired motional excitation,

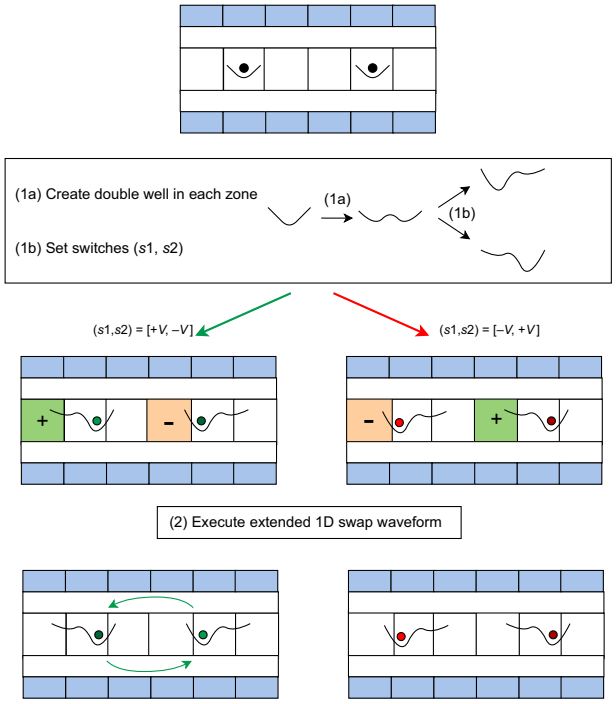

FIG. 16. The single-switch implementation of parallel digital control. Initially, two qubits (black) are trapped in two separate zones using two separate waveforms. The initial wells are evolved into double wells in every zone (1a), while local switches s_i are used to select if the qubit in zone i is pushed into the left or the right well (1b), thus selecting if the qubits will undergo a swap once the extended 1D waveform is applied. Steps (1a) and (1b) can be executed simultaneously or sequentially, in any order.

with DAC outputs filtered off chip. After all electrodes are set to either $V_0 + V$ or $V_0 - V$, the swap waveform is played and the rails return to V_0. This implements one cycle of parallel swaps.

Alternatively, parallel switching can be implemented by purely digital means. In this method, the $\pm V$ rails are permanently set to $V_0 \pm V$. This can considerably simplify circuit design and make it compatible with standard digital CMOS. However, a digital implementation may require capacitors after the switches to smooth out the switching impulses. Alternatively, qubits can be displaced with minimum motional excitation and without filtering by a sequence of switching events using bang-bang methods [20,41].

[1] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, Rev. Mod. Phys. **94**, 015004 (2022).

[2] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, Appl. Phys. Rev. **6**, 021314 (2019).

[3] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland, Complete methods set for scalable ion trap quantum information processing, Science **325**, 1227 (2009).

[4] M. Webber, S. Herbert, S. Weidt, and W. K. Hensinger, Efficient qubit routing for a globally connected trapped ion quantum computer, Adv. Quantum Technol. **3**, 2000027 (2020).

[5] M. A. Rowe, D. Kielpinski, V. Meyer, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland, Experimental violation of a Bell's inequality with efficient detection, Nature **409**, 791 (2001).

[6] L. E. de Clercq, H.-Y. Lo, M. Marinelli, D. Nadlinger, R. Oswald, V. Negnevitsky, D. Kienzler, B. Keitch, and J. P. Home, Parallel Transport Quantum Logic Gates with Trapped Ions, Phys. Rev. Lett. **116**, 080502 (2016).

[7] C. M. Seck, A. M. Meier, J. T. Merrill, H. T. Hayden, B. C. Sawyer, C. E. Volin, and K. R. Brown, Single-ion addressing via trap potential modulation in global optical fields, New J. Phys. **22**, 053024 (2020).

[8] H. N. Tinkey, C. R. Clark, B. C. Sawyer, and K. R. Brown, Transport-Enabled Entangling Gate for Trapped Ions, Phys. Rev. Lett. **128**, 050502 (2022).

[9] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, Experimental issues in coherent quantum-state manipulation of trapped atomic ions, J. Res. Natl. Inst. Stand. Technol. **103**, 259 (1998).

[10] D. Kielpinski, C. Monroe, and D. J. Wineland, Architecture for a large-scale ion-trap quantum computer, Nature **417**, 709 (2002).

[11] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, Single ion qubit with estimated coherence time exceeding one hour, Nat. Commun. **12**, 233 (2021).

[12] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler, and U. G. Poschinger, A long-lived Zeeman trapped-ion qubit, Appl. Phys. B **122**, 254 (2016).

[13] M. A. Sepiol, A. C. Hughes, J. E. Tarlton, D. P. Nadlinger, T. G. Ballance, C. J. Ballance, T. P. Harty, A. M. Steane, J. F. Goodwin, and D. M. Lucas, Probing Qubit Memory Errors at the Part-per-Million Level, Phys. Rev. Lett. **123**, 110503 (2019).

[14] E. Pelofske, A. Bärtschi, and S. Eidenbenz, Quantum volume in practice: What users can expect from NISQ devices, IEEE Trans. Quantum Eng. **3**, 1 (2022).

[15] J. M. Amini, H. Uys, J. H. Wesenberg, S. Seidelin, J. Britton, J. J. Bollinger, D. Leibfried, C. Ospelkaus, A. P. VanDevender, and D. J. Wineland, Toward scalable ion traps for quantum information processing, New J. Phys. **12**, 033031 (2010).

[16] J. D. Sterk, H. Coakley, J. Goldberg, V. Hietala, J. Lechtenberg, H. McGuinness, D. McMurtrey, L. P. Parazzoli, J. Van Der Wall, and D. Stick, Closed-loop optimization of fast trapped-ion shuttling with sub-quanta excitation, npj Quantum Inf. **8**, 1 (2022).

[17] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig,

D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, Demonstration of the trapped-ion quantum-CCD computer architecture, Nature **592**, 209 (2021).

[18] V. Kaushal, B. Lekitsch, A. Stahl, J. Hilder, D. Pijn, C. Schmiegelow, A. Bermudez, M. Müller, F. Schmidt-Kaler, and U. Poschinger, Shuttling-based trapped-ion quantum information processing, ArXiv:1912.04712 (2019).

[19] For the sake of generality, we will use the term *qubit transport* rather than ion transport. This accommodates the possibility of encoding one qubit in a short ion chain or the possibility that qubit ions may be always transported in tandem with ancilla ions [59].

[20] J. Alonso, F. M. Leupold, B. C. Keitch, and J. P. Home, Quantum control of the motional states of trapped ions through fast switching of trapping potentials, New J. Phys. **15**, 023001 (2013).

[21] P. Holz, Ph.D. thesis, University of Innsbruck, ETH Zurich, 2019.

[22] M. C. Revelle, Phoenix and Peregrine ion traps, ArXiv:2009.02398 (2020).

[23] P. C. Holz, S. Auchter, G. Stocker, M. Valentini, K. Lakhmanskiy, C. Rössler, P. Stampfer, S. Sgouridis, E. Aschauer, Y. Colombe, and R. Blatt, 2D linear trap array for quantum information processing, Adv. Quantum Technol. **3**, 2000031 (2020).

[24] J. Stuart, R. Panock, C. D. Bruzewicz, J. A. Sedlacek, R. McConnell, I. L. Chuang, J. M. Sage, and J. Chiaverini, Chip-Integrated Voltage Sources for Control of Trapped Ions, Phys. Rev. Appl. **11**, 024010 (2019).

[25] N. D. Guise, S. D. Fallek, H. Hayden, C.-S. Pai, C. Volin, K. R. Brown, J. T. Merrill, A. W. Harter, J. M. Amini, L. M. Lust, K. Muldoon, D. Carlson, and J. Budach, In-vacuum active electronics for microfabricated ion traps, Rev. Sci. Instrum. **85**, 063101 (2014).

[26] Sinara Fastino wiki, https://github.com/sinara-hw/Fastino/wiki (2022).

[27] J. C. Bardin, E. Jeffrey, E. Lucero, T. Huang, S. Das, D. T. Sank, O. Naaman, A. E. Megrant, R. Barends, T. White, M. Giustina, K. J. Satzinger, K. Arya, P. Roushan, B. Chiaro, J. Kelly, Z. Chen, B. Burkett, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, R. Graff, P. Klimov, J. Mutus, M. J. McEwen, M. Neeley, C. J. Neill, C. Quintana, A. Vainsencher, H. Neven, and J. Martinis, Design and characterization of a 28-nm bulk-CMOS cryogenic quantum controller dissipating less than 2 mW at 3 K, IEEE J. Solid-State Circuits **54**, 3043 (2019).

[28] SkyWater SKY130 PDKdocumentation, https://skywater-pdk.readthedocs.io/en/main/rules/device-details.html.

[29] D. T. C. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki Jr., S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring, Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap, Appl. Phys. B **107**, 913 (2012).

[30] Z. D. Romaszko, S. Hong, M. Siegele, R. K. Puddy, F. R. Lebrun-Gallagher, S. Weidt, and W. K. Hensinger, Engineering of microfabricated ion traps and integration of advanced on-chip features, Nat. Rev. Phys. **2**, 285 (2020).

[31] M. G. Blain, R. Haltli, P. Maunz, C. D. Nordquist, M. Revelle, and D. Stick, Hybrid MEMS-CMOS ion traps for NISQ computing, Quantum Sci. Technol. **6**, 034011 (2021).

[32] 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, Phys. Rev. Lett. **109**, 080501 (2012).

[33] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland, Coherent Diabatic Ion Transport and Separation in a Multizone Trap Array, Phys. Rev. Lett. **109**, 080502 (2012).

[34] H. Kaufmann, T. Ruster, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger, Dynamics and control of fast ion crystal splitting in segmented Paul traps, New J. Phys. **16**, 073012 (2014).

[35] T. Ruster, C. Warschburger, H. Kaufmann, C. T. Schmiegelow, A. Walther, M. Hettrich, A. Pfister, V. Kaushal, F. Schmidt-Kaler, and U. G. Poschinger, Experimental realization of fast ion separation in segmented Paul traps, Phys. Rev. A **90**, 033410 (2014).

[36] H. Kaufmann, T. Ruster, C. T. Schmiegelow, M. A. Luda, V. Kaushal, J. Schulz, D. von Lindenfels, F. Schmidt-Kaler, and U. G. Poschinger, Fast ion swapping for quantum-information processing, Phys. Rev. A **95**, 052319 (2017).

[37] S. C. Doret, J. M. Amini, K. Wright, C. Volin, T. Killian, A. Ozakin, D. Denison, H. Hayden, C.-S. Pai, R. E. Slusher, and A. W. Harter, Controlling trapping potentials and stray electric fields in a microfabricated ion trap through design and compensation, New J. Phys. **14**, 073012 (2012).

[38] D. P. Nadlinger, P. Drmota, D. Main, B. C. Nichol, G. Araneda, R. Srinivas, L. J. Stephenson, C. J. Ballance, and D. M. Lucas, Micromotion minimisation by synchronous detection of parametrically excited motion, ArXiv:2107.00056 (2021).

[39] D. Kienzler, Doctoral thesis, University of Innsbruck, ETH Zurich, https://www.research-collection.ethz.ch/handle/20.500.11850/102459 (2015).

[40] C. Decaroli, Doctoral thesis, University of Innsbruck, ETH Zurich, https://www.research-collection.ethz.ch/handle/20.500.11850/505352 (2021).

[41] J. Alonso, F. M. Leupold, Z. U. Solèr, M. Fadel, M. Marinelli, B. C. Keitch, V. Negnevitsky, and J. P. Home, Generation of large coherent states by bang–bang control of a trapped-ion oscillator, Nat. Commun. **7**, 11243 (2016).

[42] Y. Xu, F. K. Unseld, A. Corna, A. M. J. Zwerver, A. Sammak, D. Brousse, N. Samkharadze, S. V. Amitonov, M. Veldhorst, G. Scappucci, R. Ishihara, and L. M. K. Vandersypen, On-chip integration of Si/SiGe-based quantum dots and switched-capacitor circuits, Appl. Phys. Lett. **117**, 144002 (2020).

[43] We say *quasi* to indicate that the voltage can be adjusted between steps; e.g., a different voltage set can be used for transport and for gates.

[44] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, Ion-trap measurements of electric-field noise near surfaces, Rev. Mod. Phys. **87**, 1419 (2015).

[45] I. Talukdar, D. J. Gorman, N. Daniilidis, P. Schindler, S. Ebadi, H. Kaufmann, T. Zhang, and H. Häffner, Implications of surface noise for the motional coherence of trapped ions, Phys. Rev. A **93**, 043415 (2016).

[46] In this particular example, the number of DACs could be reduced by half by simply leaving the electrodes unconnected whenever we want to keep qubits still. Further reduction in the number of DACs could be achieved by wiring the odd and even zones to the same DACs but with inverted mapping. However, we ignore these tricks in the rest of the paper for the sake of generality.

[47] Technically, as $(0,1)$ and $(1,0)$ are not valid control words, it suffices to send 1 bit here. However, controlling every switch individually will become important later, so we assume that setting one switch requires 1 bit of information throughout the paper.

[48] Wikipedia, Odd-even sort, https://en.wikipedia.org/wiki/Odd%E2%80%93even_sort (2022).

[49] N. Alon, F. R. K. Chung, and R. L. Graham, Routing permutations on graphs via matchings, SIAM J. Discrete Math. **7**, 513 (1994).

[50] A. M. Childs, E. Schoute, and C. M. Unsal, Circuit transformations for quantum architectures, ArXiv:1902.09102 (2019).

[51] A. Banerjee, X. Liang, and R. Tohid, Locality-aware qubit routing for the grid architecture, ArXiv:2203.11333 (2022).

[52] Changing the permutation order to column-row-column allows for arbitrary permutation in at most $2n+m$ steps, which is more efficient if $m > n$.

[53] C. Zhang, K. K. Mehta, and J. P. Home, Optimization and implementation of a surface-electrode ion trap junction, New J. Phys. **24**, 073030 (2022).

[54] X. Shi, S. L. Todaro, G. L. Mintzer, C. D. Bruzewicz, J. Chiaverini, and I. L. Chuang, Ablation loading of barium ions into a surface electrode trap, ArXiv:2303.02143 (2023).

[55] M. Ghadimi, V. Blūms, B. G. Norton, P. M. Fisher, S. C. Connell, J. M. Amini, C. Volin, H. Hayden, C.-S. Pai, D. Kielpinski, M. Lobino, and E. W. Streed, Scalable ion-photon quantum interface based on integrated diffractive mirrors, npj Quantum Inf. **3**, 1 (2017).

[56] Technically, an additional k time steps are needed for moving qubits in and out of junctions (step C2 in Fig. 7). Furthermore, if k is large enough, moving sorted qubits to the back of the queue (step C3 in Fig. 7) may occasionally hold up the parallel column sort. On the other hand, the algorithm can be slightly sped up by other means, such as initiating the sort of the next column while the previous one is not fully sorted. We leave these caveats for future work.

[57] A. Bermudez, X. Xu, R. Nigmatullin, J. O'Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller, Assessing the Progress of Trapped-Ion Processors Towards Fault-Tolerant Quantum Computation, Phys. Rev. X **7**, 041061 (2017).

[58] W. C. Burton, B. Estey, I. M. Hoffman, A. R. Perry, C. Volin, and G. Price, Transport of multispecies ion crystals through a junction in an RF Paul trap, ArXiv:2206.11888 (2022).

[59] A. M. Steane, How to build a 300 bit, 1 Giga-operation quantum computer, ArXiv:quant-ph/0412165 (2006).

[60] V. Negnevitsky, Doctoral thesis, University of Innsbruck, ETH Zurich, https://doi.org/10.3929/ethz-b-000295923 (2018).

[61] M. W. van Mourik, E. A. Martinez, L. Gerster, P. Hrmo, T. Monz, P. Schindler, and R. Blatt, Coherent rotations of qubits within a surface ion-trap quantum computer, Phys. Rev. A **102**, 022611 (2020).

[62] While placing $N_{se/z} \gg 10$ shims per zone increases the number of independent degrees of freedom, this has to be balanced against diminishing effectiveness of further subdividing the electrodes.

[63] M. A. Weber, C. Löschnauer, J. Wolf, M. F. Gely, R. K. Hanley, J. F. Goodwin, C. J. Ballance, T. P. Harty, and D. M. Lucas, Cryogenic ion trap system for high-fidelity near-field microwave-driven quantum logic, ArXiv:2207.11364 (2022).

[64] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland, Trapped-Ion Quantum Logic Gates Based on Oscillating Magnetic Fields, Phys. Rev. Lett. **101**, 090502 (2008).

[65] D. T. C. Allcock, T. P. Harty, C. J. Ballance, B. C. Keitch, N. M. Linke, D. N. Stacey, and D. M. Lucas, A microfabricated ion trap with integrated microwave circuitry, Appl. Phys. Lett. **102**, 044103 (2013).

[66] U. Warring, C. Ospelkaus, Y. Colombe, R. Jördens, D. Leibfried, and D. J. Wineland, Individual-Ion Addressing with Microwave Field Gradients, Phys. Rev. Lett. **110**, 173002 (2013).

[67] A. D. Leu, M. F. Gely, M. A. Weber, M. C. Smith, D. P. Nadlinger, and D. M. Lucas, Fast, high-fidelity addressed single-qubit gates using efficient composite pulse sequences, ArXiv:2305.06725 (2023).

[68] K. K. Mehta, C. Zhang, M. Malinowski, T.-L. Nguyen, M. Stadler, and J. P. Home, Integrated optical multi-ion quantum logic, Nature **586**, 533 (2020).

[69] R. J. Niffenegger, J. Stuart, C. Sorace-Agaskar, D. Kharas, S. Bramhavar, C. D. Bruzewicz, W. Loh, R. T. Maxson, R. McConnell, D. Reens, G. N. West, J. M. Sage, and J. Chiaverini, Integrated multi-wavelength control of an ion qubit, Nature **586**, 538 (2020).

[70] M. Ivory, W. J. Setzer, N. Karl, H. McGuinness, C. DeRose, M. Blain, D. Stick, M. Gehl, and L. P. Parazzoli, Integrated Optical Addressing of a Trapped Ytterbium Ion, Phys. Rev. X **11**, 041033 (2021).

[71] A. R. Vasquez, C. Mordini, C. Vernière, M. Stadler, M. Malinowski, C. Zhang, D. Kienzler, K. K. Mehta, and J. P. Home, Control of an Atomic Quadrupole Transition in a Phase-Stable Standing Wave, Phys. Rev. Lett. **130**, 133201 (2023).

[72] K. K. Mehta and R. J. Ram, Precise and diffraction-limited waveguide-to-free-space focusing gratings, Sci. Rep. **7**, 2019 (2017).

[73] H. Haeffner, C. F. Roos, and R. Blatt, Quantum computing with trapped ions, Phys. Rep. **469**, 155 (2008).

[74] J. T. Merrill and K. R. Brown, Progress in compensating pulse sequences for quantum computation, ArXiv:1203.6392 (2012).

[75] D. P. L. A. Craik, N. M. Linke, M. A. Sepiol, T. P. Harty, J. F. Goodwin, C. J. Ballance, D. N. Stacey, A. M. Steane, D.

M. Lucas, and D. T. C. Allcock, High-fidelity spatial and polarization addressing of Ca-43 qubits using near-field microwave control, Phys. Rev. A **95**, 022337 (2017).

[76] Y. Wang, S. Crain, C. Fang, B. Zhang, S. Huang, Q. Liang, P. H. Leung, K. R. Brown, and J. Kim, High-Fidelity Two-Qubit Gates Using a MEMS-Based Beam Steering System for Individual Qubit Addressing, Phys. Rev. Lett. **125**, 150505 (2020).

[77] I. Pogorelov, T. Feldker, Ch. D. Marciniak, L. Postler, G. Jacob, O. Krieglsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler, and T. Monz, Compact Ion-Trap Quantum Computing Demonstrator, PRX Quantum **2**, 020343 (2021).

[78] H only implements a global single-qubit rotation when applied to individual qubits, which can be corrected for in subsequent gates.

[79] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).

[80] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).

[81] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, High-Fidelity Universal Gate Set for $^{9}Be^{+}$ Ion Qubits, Phys. Rev. Lett. **117**, 060505 (2016).

[82] S. S. Bullock and I. L. Markov, Arbitrary two-qubit computation in 23 elementary gates, Phys. Rev. A **68**, 012318 (2003).

[83] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, Phys. Rev. A **69**, 032315 (2004).

[84] V. V. Shende, I. L. Markov, and S. S. Bullock, Minimal universal two-qubit controlled-NOT-based circuits, Phys. Rev. A **69**, 062321 (2004).

[85] R. Srinivas, C. M. Löschnauer, M. Malinowski, A. C. Hughes, R. Nourshargh, V. Negnevitsky, D. T. C. Allcock, S. A. King, C. Matthiesen, T. P. Harty, and C. J. Ballance, Coherent control of trapped ion qubits with localized electric fields, ArXiv:2210.16129 (2022).

[86] Uniform pickup on shim electrodes will not lead to appreciable uncompensatable micromotion, only to a shift of the rf null. Furthermore, this shift can be compensated for by adjusting the dimensions of the rf electrodes.

[87] N. D. Guise, S. D. Fallek, K. E. Stevens, K. R. Brown, C. Volin, A. W. Harter, J. M. Amini, R. E. Higashi, S. T. Lu, H. M. Chanhvongsak, T. A. Nguyen, M. S. Marcus, T. R. Ohnstein, and D. W. Youngner, Ball-grid array architecture for microfabricated ion traps, J. Appl. Phys. **117**, 174901 (2015).

[88] J. P. Home, D. Hanneke, J. D. Jost, D. Leibfried, and D. J. Wineland, Normal modes of trapped ions in the presence of anharmonic trap potentials, New J. Phys. **13**, 073026 (2011).

[89] M. Malinowski, Doctoral thesis, University of Innsbruck, ETH Zurich, https://doi.org/10.3929/ethz-b-000516613 (2021).

[90] Also note that it is in principle possible to keep track of these decaying fields and apply appropriate corrections to the drive fields.

[91] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt, Trapped-ion probing of light-induced charging effects on dielectrics, New J. Phys. **12**, 093035 (2010).

[92] J. M. Stuart, Doctoral thesis, Massachusetts Institute of Technology, https://dspace.mit.edu/handle/1721.1/142846 (2021).

[93] C. Decaroli, R. Matt, R. Oswald, C. Axline, M. Ernzer, J. Flannery, S. Ragg, and J. P. Home, Design, fabrication and characterization of a micro-fabricated stacked-wafer segmented ion trap with two X-junctions, Quantum Sci. Technol. **6**, 044001 (2021).

[94] D.-I. Cho, S. Hong, M. Lee, and T. Kim, A review of silicon microfabricated ion traps for quantum information processing, Micro Nano Syst. Lett. **3**, 2 (2015).

[95] P. Zhao, J.-P. Likforman, H. Y. Li, J. Tao, T. Henner, Y. D. Lim, W. W. Seit, C. S. Tan, and L. Guidoni, TSV-Integrated Surface Electrode Ion Trap for Scalable Quantum Information Processing, Appl. Phys. Lett. **118**, 124003 (2021).

[96] D. Reens, M. Collins, J. Ciampi, D. Kharas, B. F. Aull, K. Donlon, C. D. Bruzewicz, B. Felton, J. Stuart, R. J. Niffenegger, P. Rich, D. Braje, K. K. Ryu, J. Chiaverini, and R. McConnell, High-Fidelity Ion State Detection Using Trap-Integrated Avalanche Photodiodes, Phys. Rev. Lett. **129**, 100502 (2022).

[97] W. J. Setzer, M. Ivory, O. Slobodyan, J. W. Van Der Wall, L. P. Parazzoli, D. Stick, M. Gehl, M. Blain, R. R. Kay, and H. J. McGuinness, Fluorescence Detection of a Trapped Ion with a Monolithically Integrated Single-Photon-Counting Avalanche Diode, Appl. Phys. Lett. **119**, 154002 (2021).

[98] K. K. Mehta, A. M. Eltony, C. D. Bruzewicz, I. L. Chuang, R. J. Ram, J. M. Sage, and J. Chiaverini, Ion traps fabricated in a CMOS foundry, Appl. Phys. Lett. **105**, 044103 (2014).

[99] S. Auchter, C. Axline, C. Decaroli, M. Valentini, L. Purwin, R. Oswald, R. Matt, E. Aschauer, Y. Colombe, P. Holz, T. Monz, R. Blatt, P. Schindler, C. Rössler, and J. Home, Industrially microfabricated ion trap with 1 eV trap depth, Quantum Sci. Technol. **7**, 035015 (2022).

[100] F. Sangster and K. Teer, Bucket-brigade electronics: New possibilities for delay, time-axis conversion, and scanning, IEEE J. Solid-State Circuits **4**, 131 (1969).

[101] A. R. Mills, D. M. Zajac, M. J. Gullans, F. J. Schupp, T. M. Hazard, and J. R. Petta, Shuttling a single charge across a one-dimensional array of silicon quantum dots, Nat. Commun. **10**, 1063 (2019).

[102] I. Seidler, T. Struck, R. Xue, N. Focke, S. Trellenkamp, H. Bluhm, and L. R. Schreiber, Conveyor-mode single-electron shuttling in Si/SiGe for a scalable quantum computing architecture, npj Quantum Inf. **8**, 1 (2022).

[103] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, High-Fidelity Preparation, Gates, Memory and Readout of a Trapped-Ion Quantum Bit, Phys. Rev. Lett. **113**, 220501 (2014).

[104] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter, High-fidelity laser-free universal control of two trapped ion qubits, Nature **597**, 209 (2021).

[105] C. R. Clark, H. N. Tinkey, B. C. Sawyer, A. M. Meier, K. A. Burkhardt, C. M. Seck, C. M. Shappert, N. D. Guise, C. E. Volin, S. D. Fallek, H. T. Hayden, W. G. Rellergert, and K. R. Brown, High-Fidelity Bell-State Preparation with $^{40}Ca^{+}$ Optical Qubits, Phys. Rev. Lett. **127**, 130505 (2021).

[106] L. Z. Cohen, I. H. Kim, S. D. Bartlett, and B. J. Brown, Low-overhead fault-tolerant quantum computing using long-range connectivity, Sci. Adv. **8**, eabn1717 (2022).

[107] M. A. Tremblay, N. Delfosse, and M. E. Beverland, Constant-Overhead Quantum Error Correction with Thin Planar Connectivity, Phys. Rev. Lett. **129**, 050504 (2022).

[108] T. S. Metodi, D. D. Thaker, A. W. Cross, F. T. Chong, and I. L. Chuang, in *Proceedings of the 38th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO Vol. 38 (IEEE Computer Society, Washington, DC, 2005) p. 305.

[109] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects, Phys. Rev. A **89**, 022317 (2014).

[110] R. Nigmatullin, C. J. Ballance, N. de Beaudrap, and S. C. Benjamin, Minimally complex ion traps as modules for quantum communication and computing, New J. Phys. **18**, 103028 (2016).

[111] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network, Phys. Rev. Lett. **124**, 110501 (2020).

[112] S. Gao, J. A. Blackmore, W. J. Hughes, T. H. Doherty, and J. F. Goodwin, Optimisation of Scalable Ion-Cavity Interfaces for Quantum Photonic Networks, Phys. Rev. Appl. **19**, 014033 (2023).

[113] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, Entanglement of Trapped-Ion Qubits Separated by 230 Meters, Phys. Rev. Lett. **130**, 050803 (2023).