Scalable, High-Fidelity All-Electronic Control of Trapped-Ion Qubits

C.M. Löschnauer,[1] J. Mosca Toba,[1] A.C. Hughes,[1] S.A. King,[1] M.A. Weber,[1] R. Srinivas,[1,2] R. Matt,[1] R. Nourshargh,[1] D.T.C. Allcock,[1,3] C.J. Ballance,[1,2] C. Matthiesen,[1] M. Malinowski,[1,*] and T.P. Harty[1]

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

[2] *Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom*

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

(Received 29 April 2025; revised 6 August 2025; accepted 18 September 2025; published 20 October 2025)

The central challenge of quantum computing is implementing high-fidelity quantum gates at scale. However, many existing approaches to qubit control suffer from a scale-performance trade-off, impeding progress towards the creation of useful devices. Here, we present a vision for an electronically controlled trapped-ion quantum computer that alleviates this bottleneck. Our architecture utilizes shared current-carrying traces and local tuning electrodes in a microfabricated chip to perform quantum gates with low noise and crosstalk regardless of device size. To verify our approach, we experimentally demonstrate low-noise site-selective single- and two-qubit gates in a seven-zone ion trap. We implement electronic single-qubit gates with fidelities $\geq$99.99912(8)% and demonstrate consistent performance with low crosstalk across the device. We also electronically generate two-qubit maximally entangled states with 99.97(1)% fidelity and long-term stable performance over continuous system operation. These state-of-the-art results highlight the potential of our all-electronic approach for coherent control of trapped-ion qubits in large-scale quantum computers.

DOI: 10.1103/h4wk-v31j

I. INTRODUCTION

The last three decades have brought about impressive demonstrations of key building blocks of quantum computers (QCs). Recent proof-of-principle experiments have shown that quantum gates on physical qubits can be implemented with high fidelity [1–3], and that quantum error correction can be used to improve error rates [4,5]. Scaling to larger system sizes has been demonstrated for some platforms [6–9], as have methods for scaling using quantum interconnects [10–13] and integrated control [14–17]. However, to build commercially valuable QCs, it is necessary to develop architectures that are *simultaneously* low noise and scalable. This requires co-optimizing the performance metrics—such as fidelity, connectivity, selectivity, and parallelism—with the scalability metrics—including manufacturability, qubit variability, footprint, signal delivery, and power consumption.

Chip-scale trapped-ion quantum computers (TIQCs)—with atomic qubits suspended above a control chip [18]—have the potential to satisfy these criteria by combining the low-noise performance of atomic qubits [19] with the scalability of chip-integrated electronics. While TIQCs are, by many metrics, the best-performing QC systems today [20], existing commercial systems are difficult to scale to useful sizes, as they rely on free-space light delivery to enable laser-based gates. Despite promising proof-of-principle demonstrations of laser-based gates with integrated photonics [15,21], significant performance, manufacturability, footprint, signal-delivery, and power-consumption challenges associated with integrated light generation, modulation, and distribution remain to be overcome [22–24].

An alternative method is to perform quantum gates using magnetic fields generated by chip-integrated traces [25]. This approach is immediately compatible with commercial microfabrication processes, opening the door to full device integration and scalability. Nevertheless, while previous work demonstrated high-quality single- and two-qubit laser-free gates [1,26,27], the possibility of using this approach to build universal, integrated, multi-zone TIQCs has remained unexplored so far.

In this paper, we propose a TIQC architecture that combines laser-free quantum gates with local electric-field

*Contact author: mm@oxionics.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

tuning for high-performance integrated control, and we then demonstrate key building blocks, namely all-electronic single- and two-qubit gates with state-of-the-art fidelities across a multi-zone microfabricated ion trap. We further show that electrical tuning allows for highly flexible, low-crosstalk coherent control. Our results showcase the potential of high-fidelity electronic qubit control for building large-scale QCs manufactured using established foundry processes.

II. ARCHITECTURE

Our architecture for all-electronic coherent control of trapped-ion qubits is illustrated in Fig. 1. This employs a chip ion trap [18], where voltages applied to electrodes on the top metal layer hold ion qubits tens of micrometers above the surface. Building on previous proposals and demonstrations of addressed laser-free quantum gates [26,28–30], two elements are incorporated into the chip to achieve site-selective control. First, a shared drive is used to deliver a control field globally to all qubits at once [28]. Second, local tuning electrodes are used to tune the control Hamiltonian on a per-zone basis, allowing for site-selective operations [31–33].

The shared drive is implemented by passing oscillating (ac) currents through traces defined inside the chip. These ac currents generate ac magnetic fields at all sites, which form the basis of single- and multi-qubit quantum gates [34,35]. To ensure efficient power and input-output use, traces generating magnetic fields in different zones are wired in series to a single source [26,36]. Local tuning is implemented by applying additional dc voltages to the top-layer electrodes. These voltages generate electric fields that adjust the ion positions; as the on-chip currents generate large magnetic-field gradients in the near field of the chip, even small changes in ion positions result in significant variations in the magnetic fields they experience [37,38]. The voltages can also generate electric-field curvatures that adjust ion motional frequencies. As shown experimentally in Sec. III, these degrees of freedom can be used to locally modulate the strength of single- and multi-qubit interactions.

The combination of the shared ac magnetic drive and local dc electric tuning allows coherent control with high fidelity, site selectivity, and parallelism. For example, to perform quantum operations globally (all qubits "active"), the local electrodes are used to fine-tune small differences in the coupling rates between qubits and the magnetic drive so as to make the interaction Hamiltonians identical for all zones. Site-selectivity can then be achieved by using the local electrodes to "hide" qubits in certain zones so as to make their interaction Hamiltonian close to the identity [29,39,40]. Finally, local electrodes can continuously tune the Hamiltonian between the "active" and "hidden" values to perform different operations in different zones in parallel.

Compared to existing architectures for TIQC which use local laser and/or magnetic drives [30,41], our architecture brings several performance and scalability benefits. First, all qubit tuning is achieved using dc voltages, which can be efficiently multiplexed using only standard chip-integrated electronics [42]. This alleviates chip manufacturability and wiring challenges, allowing the use of standard microfabrication and packaging processes.

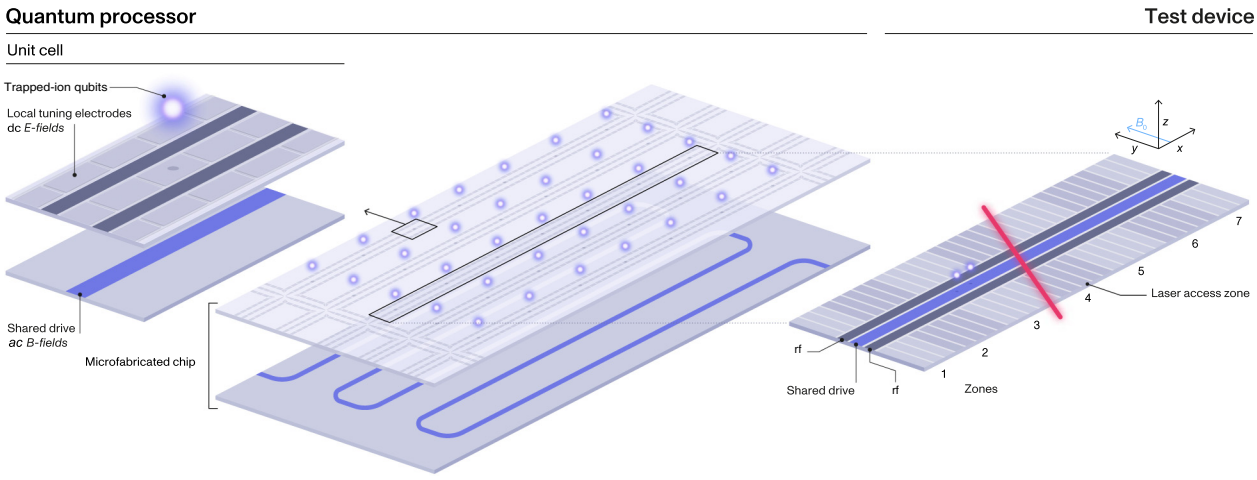


FIG. 1. Illustration of the all-electronic approach to the coherent control of trapped-ion qubits. The TIQC is a microfabricated chip (center) that contains a repeating 2D grid of unit cells (left), each composed of trapped-ion qubits stored near local tuning electrodes. A shared drive is routed to deliver ac magnetic fields to all unit cells at once. The strength of qubit coupling with the shared drive is controlled by applying dc electric fields to the tuning electrodes. Right: a linear seven-zone test device used for the experiments described in Sec. III. The tuning electrodes and a shared drive are placed on the same layer. The shared current-carrying trace is routed along $\hat{x}$, and the ions are trapped directly above at the height of $z \approx 40$ μm. Laser light is delivered to zone 4 for ion loading, state preparation, and measurement, but it is not used for coherent control. A static magnetic field $B_0 \approx 8.5$ mT is oriented along $\hat{y}$.

Furthermore, dc voltages can be distributed across a chip with negligible crosstalk and power dissipation, and the electric-field patterns from multiple electrodes can be well localized to specific trap zones. Finally, our method does not require new dedicated structures, instead repurposing the electrodes already present in the TIQC for qubit shuttling and stray-field compensation.

Second, all coherent control is achieved using electric and magnetic fields, which can be generated from electronically integrated, high-stability, low-cost sources with low phase and amplitude noise [43]. This is in contrast to laser-based gates, which typically rely on large, complex, free-space optical sources. Furthermore, while laser-based gates face stringent fidelity limits due to spontaneous emission [44,45], laser-free gates have a negligible fundamental noise floor [26,46], allowing significant fidelity improvements through hardware engineering alone. Finally, wiring the shared drive in series is highly power efficient, allowing the number of qubits to be increased while keeping the power density constant.

To verify the performance of our architecture, we have designed and built a seven-zone microfabricated ion trap with local dc electrodes and a shared ac current-carrying trace. In the next section, we demonstrate that all-electronic coherent control can be used to perform single- and two-qubit gates in a site-selective fashion. Furthermore, we show the low-noise capabilities of our technology by benchmarking single-qubit gates with 99.99916(7)% fidelity and by creating two-qubit maximally entangled states with 99.97(1)% fidelity—to our knowledge, this represents record-level performance for any QC platform to date. Finally, we discuss the path to scaling this performance to mid-scale QCs with thousands of qubits, taking into account the incoherent laser-based operations required for TIQCs, such as qubit state preparation and measurement (SPAM).

III. EXPERIMENTS

The architecture validation experiments are performed in a seven-zone single-layer linear microfabricated trap cooled to ≈5 K (Fig. 1, right). Each zone contains six to ten independent electrodes over ≈350 μm along x which act as local tuners, and a current-carrying trace routed in the center of the chip acts as the shared drive. The qubit states $|\uparrow\rangle$ and $|\downarrow\rangle$ are encoded in the $4S_{1/2}$ Zeeman sublevels of $^{40}Ca^{+}$ ions trapped 40 μm above the chip surface. The qubit frequency $\omega_0 \approx 2\pi \times 240$ MHz is set by a static magnetic field of $B_0 \approx 8.5$ mT, aligned in the plane of the chip perpendicular to the trap axis. Laser beams at 375, 397, 423, 729, 854, and 866 nm are aligned above the central zone to facilitate ion loading, ion cooling, and qubit initialization and measurement; in all other zones, all ion control is laser-free.

Each experimental shot begins with Doppler and sideband cooling of one (for single-qubit gate experiments) or two "probe" ions (for two-qubit gate experiments) to near the motional ground state [47]. The probe ions are then transported using time-dependent voltages on the dc electrodes to the center of the zone to be probed—the "target" zone. Next, electrode voltages in all other zones are set so as to also contain trapping wells in their respective centers, with the well properties calibrated to be identical to the target zone well but not containing any ions. These spectator wells allow us to investigate zone-to-zone crosstalk. Following that, local electrode dc voltages are adjusted to selectively activate interactions in some zones only, and a sequence of ac current pulses, near-resonant with the qubit transition, is applied to the shared trace to drive the desired operation. Subsequently, the probe ion(s) are shuttled back to the central zone, state $|\downarrow\rangle$ is shelved using a sequence of 729-nm pulses, and state-dependent fluorescence is collected on a complementary metal-oxide semiconductor camera to measure the qubit state [48]. In the experiments presented below, single-qubit gates are performed in all seven trap zones, while two-qubit gates are only performed in zones 2–6 due to a reduced tuning-electrode count in zones 1 and 7. The single-qubit gates are benchmarked on a stationary qubit, meaning that the qubit is not shuttled between "active" and "hidden" positions during the benchmarking, and we probe one zone at a time.

A. High-fidelity tunable single-qubit gates

Consider a qubit with frequency ω_0 in zone i [Fig. 2(a)]. Applying an oscillating current $I\cos(\omega_0 t + \phi)$ for duration t_1 through the shared trace results in a single-qubit rotation unitary $U_1 = \exp(-i\sigma_\phi\theta_1/2)$, where $\sigma_\phi = \cos(\phi)\sigma_x + \sin(\phi)\sigma_y$, $\sigma_{x/y}$ are the Pauli spin matrices, and the gate angle $\theta_1 = \Omega_1 t_1$, where the Rabi frequency Ω_1 depends on the magnitude and polarization of the magnetic field from the shared trace (see Appendix). In our geometry, Ω_1 is proportional to the z-polarized component of the magnetic field B_z. To achieve site-selective control, we leverage the fact that directly above the shared trace $B_z \approx 0$ while $\partial B_z/\partial y$ remains significant [Fig. 2(b)]. Thus, zone-selective ion translation along y results in zone-selective adjustment of θ_1.

To validate the method, we measure the Rabi frequency for a range of ion positions y in the central zone. The results are summarized in Fig. 2(b). At $I \approx 70$ mA, we find that we can tune the Rabi frequency between $\Omega_1 \approx 2\pi \times 17$ kHz and $\Omega_1 \approx 2\pi \times 330$ kHz by translating an ion from $y \approx -0.2$ μm (qubit "hidden") to $y \approx 3.5$ μm (qubit "active"). This translation is performed by tuning voltages on local electrodes by ≈0.1 V. We correct for the effect of the residual Rabi frequency at the "hidden" position (caused by ≈0.2° misalignment of the static magnetic

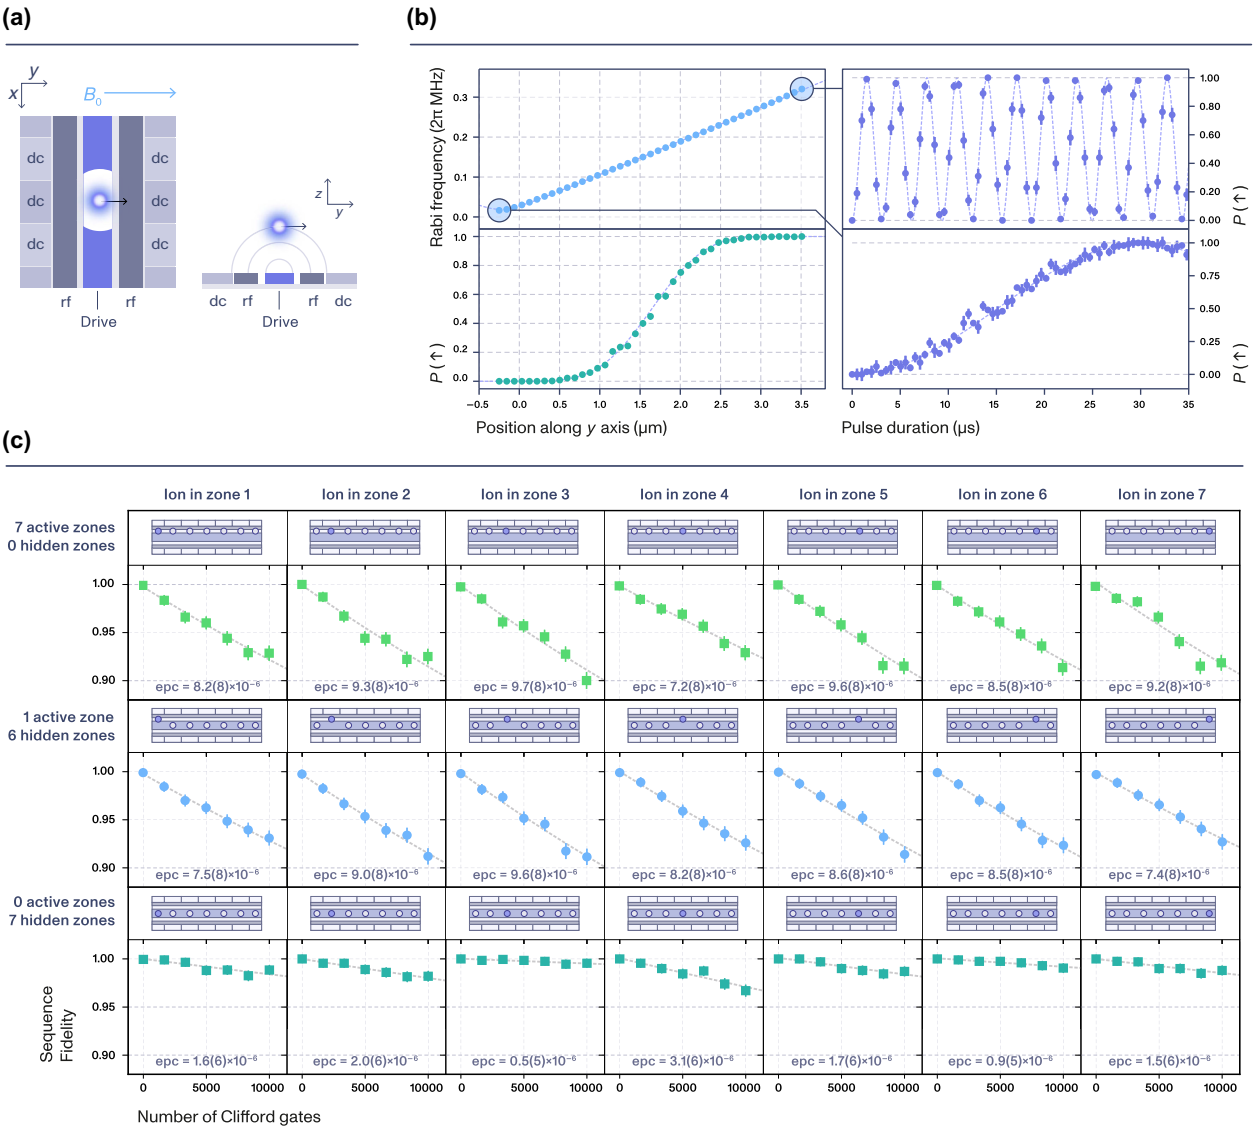


FIG. 2. Summary of single-qubit control experiments. (a) Single-qubit control is achieved by using dc electrodes to adjust qubit position along $\hat{y}$. Directly above the shared trace ($y \approx 0$) the ac magnetic field is oriented in the y direction. Since the static magnetic field is oriented along y as well, we find that the single-qubit Rabi frequency $\Omega_1 \propto B_z \approx 0$, i.e., the qubit is "hidden" (see Appendix). At the same time, $|\partial\Omega_1/\partial y| \propto |\partial B_z/\partial y| > 0$, thus translating the qubit along y causes interactions to be turned on. (b) Top left: experimental measurements of the single-qubit Rabi frequency Ω_1 vs y. Right: Rabi oscillations in the "active" (top) and "hidden" (bottom) positions. Bottom left: Spin-flip probability $P(\downarrow)$ vs y for an SK1 sequence calibrated to drive a π-pulse in the "active" position. (c) Results of single-qubit RB throughout the device. Each experiment uses one qubit to record the spin-flip probability $P(\downarrow)$ in a target zone in one of three configurations: all zones "active" (top), target zone "active" and other zones "hidden" (middle), and all zones "hidden" (bottom). Benchmarking is carried out in one target zone at a time. The dc trapping wells required for a given configuration are applied in all zones simultaneously, but only the respective target zone contains an ion. The zone configuration is shown above the data for each benchmarking experiment, where active and hidden zones are indicated by axially centered and displaced circles, respectively, and the filled circle indicates the target zone containing the qubit, while the other zones contain empty wells (empty circles). In all plots, error bars represent a $\pm 1\sigma$ confidence interval for the spin-flip probability. The exponential fits (dashed lines) are used to estimate the error per Clifford (epc) in each configuration; see Appendix for details.

field) by employing a Solovay-Kitaev-1 gate sequence, which uses three resonant pulses to improve robustness to static Rabi frequency errors [49]. As shown in Fig. 2(b) for a total gate duration of 8 μs, we can switch between $\theta_1 \approx 0$ (qubit "hidden") and $\theta_1 \approx \pi$ (qubit "active") with high contrast and a high degree of robustness against ion position errors.

We quantify the gate errors using single-qubit randomized benchmarking (RB) [50,51] with all zones set to "active", as shown in Fig. 2(c) (top) and described further

in the Appendix. Using all-electronic control, we achieve an average error of $8.8(8) \times 10^{-6}$ per Clifford, consistent across all seven zones and in agreement with independently measured noise in the signal delivery chain. This is, to our knowledge, the lowest single-qubit gate error ever recorded in a multi-zone quantum processor.

To estimate the crosstalk errors, we repeat RB with one zone set to be in the "active" configuration at a time [Fig. 2(c), middle]. The resulting error per Clifford of $8.4(7) \times 10^{-6}$ is unchanged within measurement uncertainty, indicating that setting other zones to "hidden" or "active" configuration changes target zone qubit gate errors by $\lessapprox 1 \times 10^{-6}$. A more dominant source of crosstalk is the residual magnetic interaction experienced by a qubit in a "hidden" zone due to the shared drive field. To estimate its magnitude, we run the RB sequence on the qubit with all zones set to "hidden" configuration [Fig. 2(c), bottom], finding a spin-flip probability of $1.6(8) \times 10^{-6}$ per Clifford, once again indicating very low crosstalk. We note that the remaining error on the qubit in a "hidden" zone is dominated by coherent effects and can be further suppressed with higher-order composite pulses [49].

B. High-fidelity tunable two-qubit gates

Consider a pair of qubits with frequency ω_0 and a motional mode frequency ω_m in zone i [Fig. 3(a)]. Applying a bichromatic oscillating current $I_{\pm}(t) = I\cos((\omega_0 \pm (\omega_m + \delta))t + \phi)$ for duration t_2 to the shared trace results in a two-qubit entangling unitary $U_2 = \exp(-i\sigma_\phi \otimes \sigma_\phi \theta_2/2)$ as long as $\delta = 2\pi n/t_2$, where n is an integer [25,52] (see Appendix). The entangling gate angle is then given by $\theta_2 = \Omega_2^2 t_2^2/(4n\pi)$, where the gate Rabi frequency Ω_2 depends on the magnitude of the magnetic-field gradient along the direction of ion motion [25].

To implement arbitrary quantum computation, it suffices to complement site-selective tunable single-qubit rotations with global maximally entangling two-qubit gates ($\theta_2 = \pi/2$). We demonstrate this capability by preparing maximally entangled states as follows. We optically pump two ions to initial state $|\uparrow\uparrow\rangle$ and set the radial rocking mode to be at a frequency $\omega_m \approx 2\pi \times 3.7$ MHz and aligned in the plane of the chip. We then apply a sequence of two bichromatic ac current pulses through the shared trace, each with $I \approx 0.7$ A and $\delta \approx 2\pi \times -16$ kHz, resulting in $\Omega_2 \approx 2\pi \times 6$ kHz. Each pulse has a duration of $t_2/2 \approx 60\ \mu$s, and we flip the phase difference between the currents by π between the pulses to improve robustness to fluctuations in ω_m [53]. This sequence ideally prepares a maximally entangled state $(|\uparrow\uparrow\rangle - i|\downarrow\downarrow\rangle)/\sqrt{2}$. We estimate the fidelity with that state by measuring the qubit populations at the end of the pulse, as well as parity contrast after an additional $\pi/2$ "analysis" pulse with a variable phase [54].

We begin the experiments by roughly tuning the two-qubit interactions throughout the trap. We match mode frequencies and orientations of all five two-qubit wells simultaneously via the local dc electrodes, resulting in the high-contrast parity oscillations shown in Fig. 3(c). No statistically meaningful zone-to-zone variations are observed. We then proceed to characterize the gate performance in zone 4 with high precision, carrying out fine parameter calibration and repeating the measurement sequence over $\approx$12 h to estimate the fidelity with a statistical uncertainty of $\leq 1 \times 10^{-4}$. The results are shown in Fig. 3(b). Averaging over all $\approx 10^6$ data points, we find the error of preparing a maximally entangled state of $3(1) \times 10^{-4}$, a record for any QC platform. The uncertainty corresponds to a 68% confidence interval estimated via bootstrapping [55], and the fidelity estimate has been corrected for the state-preparation error of $1.4(2) \times 10^{-4}$ and the measurement error of $3.0(2) \times 10^{-4}$ which were independently monitored throughout the data-acquisition period (see Appendix for details).

To benchmark the baseline passive stability of the system, we continued the data acquisition for a further 48 h, bringing the total number of experimental runs to $\approx 5 \times 10^6$. We find that, beyond $t \approx 12$ h, the uncertainty in the error deviates from the $t^{-1/2}$ scaling expected for purely statistical fluctuations [Fig. 3(d), inset], indicating a fidelity drift. Nonetheless, even when averaged over the full $t = 60$ h of data acquisition, we record an error of $4.7(7) \times 10^{-4}$, which is state of the art for any QC platform. The average fidelity drift rate is at the level of $\approx 1 \times 10^{-4}$ per day over the full data run. During that period, only the motional frequency ω_m was regularly calibrated using an automated routine, with all other system parameters, including the qubit frequency, fixed at the initial calibration value.

Based on independent memory benchmarking experiments [19,56], we estimate that the majority of the remaining infidelity can be accounted for by qubit frequency fluctuations. At the same time, analysis of population errors in the experiments above, as well as independent measurements of mode heating rates (<1 phonon per second on the gate motional mode), allow us to upper-bound the combined error of motional mode fluctuations and heating to $<1 \times 10^{-4}$. These results suggest that the entangling-operation error rate can be reduced to $<1 \times 10^{-4}$ by employing standard techniques to protect from qubit dephasing, such as dynamical decoupling [27,57–59] and/or field-insensitive encodings [60]. Simultaneously, the high level of stability in the fidelity measurements indicates that the errors caused by parameter drifts can also be kept to $<1 \times 10^{-4}$ by adding only sparse calibration routines, e.g., one additional automated calibration per gate parameter per day. Together, these results show a direct path to reliable and robust quantum control with two-qubit gate error rates $<1 \times 10^{-4}$.

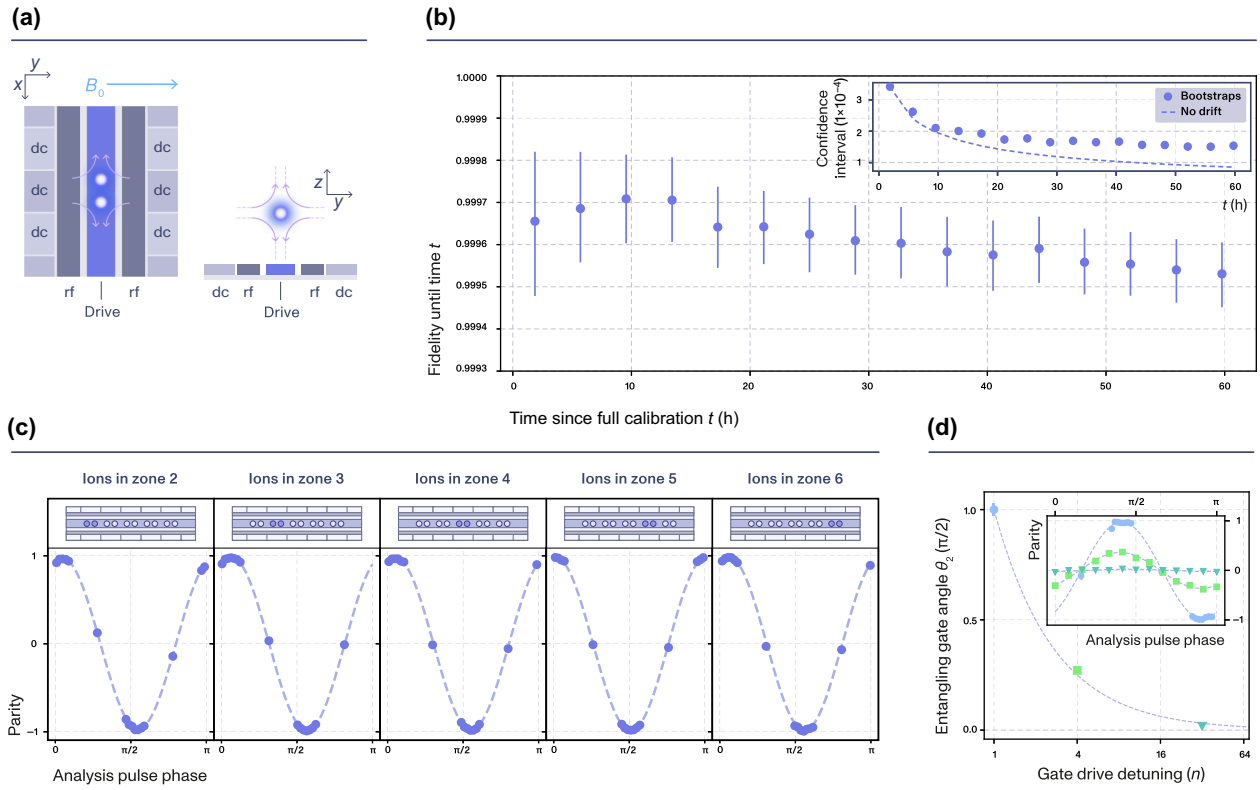


FIG. 3. Summary of two-qubit control experiments. (a) Two-qubit Mølmer-Sørensen interactions are implemented by driving currents through the shared trace. A state-dependent force on the in-plane radial mode is generated using the magnetic-field gradient $\partial B_z/\partial y$, and the fact that $B_z \approx 0$ is used to minimize off-resonant qubit frequency shifts. Quadrupole potentials generated by the dc electrodes are used to tune the motional mode frequencies and orientations, adjusting the effective interaction strength. (b) The entangled-state fidelity in zone 4 following full system calibration. We record a fidelity of 0.9997(1) over the first 12 h and a fidelity of 0.9995(1) over the full 60 h of data acquisition. The $\pm 1\sigma$ confidence intervals (CIs) are calculated by bootstrapping with 10 000 resamples. After 12 h, the CI deviates visibly from the $1/\sqrt{t}$ scaling (inset), indicating drifts in the underlying process. (c) Qubit parity oscillations in five zones (not corrected for state-preparation and measurement errors), demonstrating the ability to prepare maximally entangled states $\theta_2 \approx \pi/2$ throughout the device. Parity is measured in one target zone at a time. The required dc trapping wells are applied in all zones simultaneously, but only the respective target zone contains a pair of ions. Diagrams above the data display the trap well configuration and occupation, where filled circles indicate the occupied target zones. (d) Experimental measurements of the entangling angle θ_2 when the mode frequency is electrically adjusted to $n \times \delta$, where δ corresponds to a detuning of a maximally entangling operation. The inset shows parity oscillations used to infer the value of θ_2.

Finally, we note that while the experiments above focused on the maximally entangling operations, the method can be extended to implement tunable-angle two-qubit operations. We demonstrate this experimentally by electrically adjusting the potential experienced by the ions to offset the motional frequency by $n \times \delta$ from the shared drive. We then apply the same current pulse as above through the shared trace and measure the contrast of qubit parity oscillations to estimate the entangling angle θ_2. In Fig. 3(d) we show that this allows for the entangling angle to be adjusted in discrete steps as $\theta_2 = \pi/(2n)$. We note that further continuous control over θ_2 is possible by adjusting the mode orientation and that the two-qubit interactions can be fully switched off ($\theta_2 = 0$) by splitting the ion crystal before the shared drive is turned on. Such tunable-angle two-qubit gates could be used for further error reduction through more efficient unitary decomposition [61–63].

IV. DISCUSSION

The experiments presented in Sec. III demonstrate that the all-electronic coherent control architecture discussed in Sec. II allows for low-noise, site-selective multi-zone quantum control of trapped-ion qubits. Here, we outline how these coherent-control primitives can be combined with already demonstrated ion-trap and microfabrication technology to build next-generation high-performance QCs with many thousands of qubits. These mid-scale QCs are expected to provide early computational value [64], while simultaneously serving as architectural proof-points for large-scale fault-tolerant QCs.

We envision the all-electronic control architecture to be at the core of the TIQC built on a microfabricated multi-layer silicon chip [65]. Following the quantum-charge-coupled-device approach [37], the unit cells will be arranged on a 2D grid as in Fig. 1, with junctions placed throughout the device to facilitate all-to-all connectivity via ion shuttling [66]. To minimize qubit frequency variations across the chip, the device will be placed in a highly homogeneous magnetic field, and quantum information will be encoded into atomic states with reduced magnetic-field sensitivity [60].

While coherent qubit control is all electronic, qubit SPAM requires laser light, which can be delivered from off-chip sources to the ions using chip-integrated photonics [15,16,67]. We propose implementing all SPAM using shared laser drives only, delivering light to all unit cells simultaneously, and once again using local electrodes to fine-tune the operations by moving ions in and out of laser beams. This alleviates the need to build optical systems for large-scale delivery of individually modulated laser beams, which has historically been a major roadblock to scaling laser-based TIQCs [68–72]. Furthermore, SPAM requires orders of magnitude less optical power than high-fidelity coherent optical operations [73], relaxing constraints on optical losses and beam quality. Thus, our photonic delivery network can be built out of purely passive integrated components (waveguides, splitters, and grating couplers) [74] and using only processes and materials that are well known to the semiconductor industry [21,75].

We envision that, using all-electronic coherent control primitives, ever larger QCs can be built by repeating more and more unit cells. Assuming a zone spacing of $\approx 350\ \mu m$ as in Sec. III, a 10 000-qubit device requires a chip area of $\approx 3.5 \times 3.5\ cm^2$. While building an ion-trap chip of this size undoubtedly poses large engineering challenges, we note that similar chip sizes are already used for current camera sensors and field-programmable gate array chips. Power dissipation from the shared trace will be minimized by placing it on a dedicated low-loss metal layer and keeping the chip at cryogenic temperatures. For example, a shared trace of 50-μm width, 2-μm thickness, and 2×10^{-10}-Ωm resistivity [76] carrying a current of 1 A would lead to power dissipation of 0.35 mW per cell, allowing a chip with tens of thousands of qubits to be operated in an off-the-shelf 4-K cryostat, and with further order-of-magnitude gains possible by reducing the ion-electrode distance.

Brute-force device scaling could be limited by the practical wiring limits of purely passive devices. However, as argued in recent theoretical [42] and experimental [77] work, most control electrodes can be co-wired across the device, such that a modest number ($\approx$50–100) of analog voltage inputs suffices to achieve all-to-all connectivity regardless of device size. Furthermore, due to their low bandwidth, the local tuning electrodes can be efficiently multiplexed using only simple, cryogenically compatible, low-power chip-scale electronics [42]. This allows hundreds of electrodes to be controlled from a single analog input. By employing these two methods, a few hundred input lines—a typical number of analog signals in today's TIQCs—can already control devices with tens of thousands of qubits. Avoiding stray ac fields caused by capacitive and inductive coupling for such a co-wired multiplexed device will be essential, requiring the careful routing, shielding, and filtering of signal lines customary for high-frequency circuits.

Practical QC manufacturing also requires careful design tolerancing and efficient device calibration, as cell-to-cell variations can have catastrophic effects on device performance and yield [78,79]. Compared to many QC modalities, the all-electronic coherent-qubit-control TIQC architecture poses minimal new demands on the chip design, tolerancing, and manufacturing processes. In contrast to solid-state qubits, trapped-ion qubits leverage the physical separation between the qubit and the chip to suppress the detrimental effects of surface and material disorder on gate performance [80]. The experiments presented in Sec. III allow us to upper-bound the two-qubit gate error from surface-induced motional dephasing to less than 1×10^{-4}, and we expect further gains can be achieved by employing coherent control [53] and surface-treatment techniques [81–84].

Given that TIQCs allow for low surface noise and high qubit frequency uniformity, the primary goal of QC tolerancing and calibration is to equalize the control-field parameters in different unit cells. Notably, the on-chip electronic control elements—the shared trace and the local electrodes—have typical minimum size scales of tens of micrometers, several orders of magnitude larger than the resolution of modern semiconductor processes. Thus, large-scale highly uniform fabrication of multilayer qubit control structures is relatively straightforward. While integrated photonics requires submicrometer dielectric features [85], their uniformity needs can be considerably relaxed by carrying out dissipative operation in the saturated steady-state regime [86–88]. Finally, the qubit coupling inhomogeneities caused by residual chip manufacturing and design imperfections can be shimmed out using local tuning electrodes. The combination of low calibration complexity, high parameter stability (Sec. III), and the ability to parallelize calibrations in different zones will allow continuous TIQC operation with minimal downtime.

In summary, we have presented a vision for TIQCs using all-electronic coherent control to alleviate the scale and performance challenges of today's QCs. We have experimentally tested the fundamental coherent-control primitives of our architecture—site-selective single- and two-qubit gates—and verified that both can be implemented with state-of-the-art fidelities. We have discussed the prospects for scaling these systems up and argued

that next-generation devices with thousands of qubits can be fabricated using already-established microfabrication technology. This shows a path to increasing QC size without sacrificing performance, enabling mid- and large-scale QCs that can solve problems of commercial and scientific importance.

ACKNOWLEDGMENTS

We thank the entire Oxford Ionics team for their contributions to this work. We also thank Daniel Slichter and Mario Gely for their helpful comments on an earlier version of the manuscript.

All authors are employees and stock/option holders of Oxford Ionics Ltd.

The experiments were carried out by A.C.H., C.M.L., J.M.T., M.A.W., R.S., and S.A.K. R.J.N. led the development of the experimental system and infrastructure, to which the whole Oxford Ionics team contributed, in particular C.M. and S.A.K. The control architecture was developed by C.J.B., C.M., D.T.C.A., M.M., R.S., and T.P.H. Data analysis and interpretation were done by A.C.H., C.M., C.M.L., J.M.T., R.N., R.S., and M.M. The manuscript was written by M.M. with significant contributions from C.J.B. and C.M. and input from all authors.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: METHODS

1. Single-qubit gate methods

a. Single-qubit gate mechanism

Consider a shared trace placed along the $\hat{x}$ direction as in Sec. III. An oscillating current $I\cos(\omega_0 t+\phi)$ applied to the shared trace generates a magnetic field $\vec{B}\cos(\omega_0 t+\phi)$ at the ion location. For simplicity, we assume that the wire is very long along the x direction, such that the oscillating magnetic field $\vec{B}=B_y\hat{y}+B_z\hat{z}$ is polarized in the y–z plane.

For a σ-polarized qubit placed in a static magnetic field oriented along $\hat{y}$ such as the one discussed in Sec. III, the Hamiltonian can then be written as $H=H_0+H_1$, where the bare qubit Hamiltonian $H_0=\frac{1}{2}\hbar\omega_0\sigma_z$ and the interaction Hamiltonian $H_1=RB_z\sigma_x\cos(\omega_0 t+\phi)$. The coefficient $R=g_s\mu_B/2$ describes the strength of the magnetic dipole moment between the qubit states. In deriving H_1, we have neglected the contribution of the term proportional to B_y. This component causes a modulation of the qubit energy, giving rise to additional resonances. In our experiments, the driving field is far detuned from these resonances and we therefore omit them in all of the following. Transforming H into the interaction picture with respect to H_0 and taking the rotating-wave approximation to drop all oscillating terms results in a time-independent Hamiltonian $H'=\frac{1}{2}\hbar\Omega_1\sigma_\phi$, where $\Omega_1=RB_z/\hbar$ and $\sigma_\phi=\sigma_x\cos\phi+\sigma_y\sin\phi$. Integrating H' from $t=0$ to $t=t_1$ results in a single-qubit rotation unitary as in the main text.

For the geometry described in Fig. 2, we expect that at $y=0$ (i.e., directly above the shared drive) $B_z=0$ and $\partial B_z/\partial z=0$. Thus, in the absence of any imperfections, the single-qubit Rabi frequency $\Omega_1=0$ at $y=0$ for any value of z. At the same time, $\partial B_z/\partial y\neq 0$, which allows Ω_1 to be tuned by translating the ion along y. The experimental data in Fig. 2(b) differ from this simplified model in two ways. First, the Rabi-frequency minimum is shifted from $y_0=0$ to $y_0\approx -0.2$ μm. We believe this to be caused by a small tilt ($\approx 0.3^\circ$) of the static magnetic field in the y–z plane. Second, the Rabi frequency at $y=y_0$ does not drop to zero, and we instead record a minimum $\Omega_1\approx 2\pi\times 17$ kHz when applying current $I\approx 70$ mA. We believe this to be caused by a small tilt ($\approx 0.2^\circ$) of the static magnetic field in the x–y plane. We expect both of these effects can be compensated to a high degree using additional magnetic-field coils.

b. Single-qubit operation benchmarking

The quality of single-qubit operations is measured using single-qubit Clifford RB, which is implemented by decomposing Clifford gates onto single-qubit gates with θ_1 in the set $\{\pi/2,\pi\}$ and ϕ selected from the set $\{0,\pi/2,\pi,3\pi/2\}$. Each single-qubit gate, in turn, is implemented as a sequence of three square pulses: the target rotation followed by a pair of rotations with $\theta_1=2\pi$ and phase $\phi'=\phi\pm\arccos(-\theta_0/4\pi)$, which implements a Solovay-Kitaev-1 (SK1) composite pulse [49]. The longest SK1 sequence ($\theta_1=\pi$) takes 8 μs, including 7.5 μs of coherent evolution and 0.5 μs of dead time. All single-qubit RB includes Pauli randomization to eliminate possible bias from asymmetric readout errors.

We infer the error per Clifford gate ϵ by fitting the measured sequence fidelity (equivalent to the spin-flip probability) to the expected dependence on sequence length m, $F(m)=a(1-2\epsilon)^m+1/2$, where a includes the effects of SPAM errors. The uncertainty on the error per Clifford gate is extracted from the least-squares fit as the square root of the variance of ϵ.

2. Two-qubit gate methods

a. Two-qubit gate mechanism

The two-qubit gate mechanism can be analyzed by extending the discussion in Sec. A 1 to include ion motion and spin-motion coupling. Consider a qubit at frequency ω_0 placed directly above a narrow and long shared trace, coupled to a single radial mode of motion at frequency ω_m oriented along $\hat{y}$. A bichromatic oscillating current $I\cos([\omega_0\pm(\omega_m+\delta)]t)$ generates a bichromatic magnetic

field $\vec{B}\cos([\omega_0 \pm (\omega_m + \delta)]t)$ at the ion location. Due to the trace geometry, we find that $B_z \approx 0$ and $\partial B_y/\partial y \approx 0$, allowing us to approximate the system Hamiltonian as $H = H_0 + H_2$, where the bare Hamiltonian for spin and motion is given by $H_0 = \frac{1}{2}\hbar\omega_0\sigma_z + \hbar\omega_m\left(\hat{a}^\dagger\hat{a} + 1/2\right)$ and the spin-motion coupling Hamiltonian is given by

$$H_2 = R\frac{\partial B_z}{\partial y}f(t)y_0\sigma_x\left(\hat{a}^\dagger + \hat{a}\right),$$

where $\hat{a}$ is the motion-lowering operator, $y_0 = \sqrt{\hbar/(2m\omega_m)}$ is the size of the ground-state ion wave packet, with m denoting ion mass, and

$$f(t) = \cos\left[(\omega_0 + \omega_m + \delta)t\right] + \cos\left[(\omega_0 - \omega_m - \delta)t\right].$$

Once again transforming H into the interaction picture with respect to H_0, and performing the rotating-wave approximation to drop all terms oscillating at frequencies above δ, we arrive at $H' = \frac{1}{2}\hbar\Omega_2\sigma_x(\hat{a}^\dagger e^{i\delta t} + \hat{a}e^{-i\delta t})$ with $\Omega_2 = R(\partial B_z/\partial y)y_0/\hbar$, which is a standard Hamiltonian for a state-dependent force that can generate two-qubit Mølmer-Sørensen unitaries as discussed in the main text.

b. SPAM error estimation

The estimate of entangled-state fidelity includes a correction for SPAM errors. We first describe how we characterize the SPAM errors.

We denote the probability of ion i being prepared in $|\downarrow\rangle$ instead of $|\uparrow\rangle$ as $\epsilon_{\mathrm{sp},i}$. The average per-ion state-preparation error $\epsilon_{\mathrm{sp}} = \left(\epsilon_{\mathrm{sp},1} + \epsilon_{\mathrm{sp},2}\right)/2$ is measured by applying the same experimental sequence as used for entangled-state generation, but with no current applied to the shared drive, and recording the average probability $P(\downarrow)$ to find an ion in $|\downarrow\rangle$ instead of $|\uparrow\rangle$. The duration of fluorescence detection is selected to ensure negligible error ($\ll 1 \times 10^{-4}$) from $|\uparrow\rangle$ and $|\downarrow\rangle$ photon count overlap, and the Rabi frequency of the shelving pulse at 729 nm is selected to ensure negligible probability of off-resonant excitation of $|\uparrow\rangle$. Thus, we attribute all measurements of $|\downarrow\rangle$ to imperfect state preparation, and we set $\epsilon_{\mathrm{sp}} = P(\downarrow)$. Over the full 60 h of data acquisition, we record a stable state-preparation error of $\epsilon_{\mathrm{sp}} = 1.4(2) \times 10^{-4}$.

For perfect readout, an ion in state $|\downarrow\rangle$ is recorded as "dark," while an ion in state $|\uparrow\rangle$ is recorded as "bright." We denote the probability of ion i in $|\uparrow\rangle$ being erroneously recorded as "dark" as $\epsilon_{\mathrm{d,i}}$ and the probability of an ion in $|\downarrow\rangle$ being erroneously recorded as "bright" as $\epsilon_{\mathrm{b,i}}$. We also denote the average per-ion readout error as $\bar{\epsilon}_{\mathrm{r}} = (\epsilon_{\mathrm{b}} + \epsilon_{\mathrm{d}})/2$. As $\epsilon_{\mathrm{b}} \ll 1 \times 10^{-4}$, it can be neglected in the analysis. To estimate $\epsilon_{\mathrm{d}} = \left(\epsilon_{\mathrm{d},1} + \epsilon_{\mathrm{d},2}\right)/2$, we apply the same experimental sequence as used for entangled-state generation, but with the two-qubit entangling pulse replaced with a single-qubit π rotation which ideally prepares $|\downarrow\downarrow\rangle$. We then record the probability $P(\downarrow\uparrow + \uparrow\downarrow)$ of measuring one ion bright, attributing all such events to imperfect state preparation and measurement, i.e., $P(\downarrow\uparrow + \uparrow\downarrow) = \epsilon_{\mathrm{sp},1} + \epsilon_{\mathrm{sp},2} + \epsilon_{\mathrm{d},1} + \epsilon_{\mathrm{d},2} = 2\left(\epsilon_{\mathrm{sp}} + \epsilon_{\mathrm{d}}\right)$. Using the estimate of ϵ_{sp} presented earlier, we thus infer $\epsilon_{\mathrm{d}} = 6.0(5) \times 10^{-4}$ and an average readout error of $\bar{\epsilon}_{\mathrm{r}} = 3.0(2) \times 10^{-4}$, stable throughout the data acquisition period. The readout error is consistent with the expected probability of spontaneous emission from the shelve state during the readout window.

c. Fidelity estimate of two-qubit operation

The quality of two-qubit operations is measured by the fidelity with which they transform the initial pure state $|\uparrow\uparrow\rangle$ into a maximally entangled state of the form $\left(|\uparrow\uparrow\rangle + e^{i\phi_0}|\downarrow\downarrow\rangle\right)/\sqrt{2}$.

We estimate the entangled-state fidelity F by performing partial-state tomography [89], combining the population $P = P(\downarrow\downarrow) + P(\uparrow\uparrow)$ and the amplitude C of parity oscillations $\Pi(\phi) = P(\uparrow\uparrow) + P(\downarrow\downarrow) - P(\uparrow\downarrow) - P(\downarrow\uparrow)$ to give $F = (P + C)/2$. Parity oscillations are induced by applying a $\pi/2$ analysis pulse with a variable phase ϕ to the state after the entanglement operation, and we recover C from a maximum-likelihood estimate of the sinusoidal $\Pi(\phi)$ dependence using a binomial fitter.

State-preparation errors ϵ_{sp} mean we do not start in a pure state $|\uparrow\uparrow\rangle$, hence creating a state different from $\left(|\uparrow\uparrow\rangle + e^{i\phi_0}|\downarrow\downarrow\rangle\right)/\sqrt{2}$ during the entanglement operation, while state-readout errors ϵ_{r} lead to population measurements that do not faithfully reflect the true populations. We consider the impact of both types of errors on our measurement of the fidelity estimate $F' = (P' + C')/2$ independently, and we apply corrections to recover an estimate of the true fidelity F. Assuming uncorrelated SPAM errors and $F \approx 1$, we detail below that $F = F' + 2\epsilon_{\mathrm{sp}} + 3\bar{\epsilon}_{\mathrm{r}}$ up to the leading order in ϵ_{sp} and ϵ_{r}, which is used to obtain the errors quoted in the main text.

We discuss corrections due to readout errors first. Given knowledge of the readout errors, which we assume to be independent between ions, we can calculate the conditional probabilities of observing "dark" or "bright" given the actual state of the qubit. This allows us to convert the measured readout populations to true underlying qubit populations. Taking the maximally entangled target state as a starting point (i.e., not considering state-preparation errors), keeping only terms of first order in readout error and taking into account that a high-fidelity gate will result in populations that are very close to their ideal values, we can relate expected measurements of P' and C' to the true quantities: $P_{\mathrm{corr,\,r}} = P' + 2\bar{\epsilon}_{\mathrm{r}}$ and $C_{\mathrm{corr,\,r}} = C' + 4\bar{\epsilon}_{\mathrm{r}}$. Hence, the readout correction to fidelity is given by $F_{\mathrm{corr,\,r}} = F' + 3\bar{\epsilon}_{\mathrm{r}}$.

Next, we consider corrections due to state-preparation errors. Imperfect state preparation amounts to applying

the gate operation over several different input states rather than just the target input state $|\uparrow\uparrow\rangle$. Given a known state-preparation error, we can infer the probability distribution of actual input states and determine the output states and corresponding populations produced by the entangling operation (and a subsequent $\pi/2$-pulse, if applicable). This allows us to relate the measured P' and C' to the corrected quantities P and C. The latter quantities should be observed if the input state was always exactly $|\uparrow\uparrow\rangle$. Keeping only terms of first order in state-preparation error and taking into account that a high-fidelity gate will result in populations that are very close to their ideal values, we find that $P_{\text{corr, sp}} = P' + 2\epsilon_{\text{sp}}$ and $C_{\text{corr, sp}} = C' + 2\epsilon_{\text{sp}}$. Therefore, the state-preparation correction to fidelity can be written as $F_{\text{corr, sp}} = F' + 2\epsilon_{\text{sp}}$.

Combining readout and state-preparation errors, we arrive at our estimate of the true Bell-state fidelity $F = F' + 2\epsilon_{\text{sp}} + 3\bar{\epsilon}_{\text{r}}$, where $\epsilon_{\text{sp}} = 1.4(2) \times 10^{-4}$ and $\bar{\epsilon}_{\text{r}} = 3.0(2) \times 10^{-4}$. To obtain the confidence intervals for the fidelity estimate, we bootstrap [55] the raw population, parity, and SPAM data with 10 000 resamples that are run through the same analysis procedure.

d. Two-qubit gate fidelity limitations

As we are using a Zeeman qubit for the present demonstration, the qubit coherence is affected by ambient magnetic-field fluctuations, which directly map to qubit frequency fluctuations (dc Zeeman shifts), and fluctuating ac Zeeman shifts, for instance due to instability of the bichromatic gate tones. In order to estimate the impact of qubit frequency stability on the fidelity of the two-qubit operations we probe qubit decoherence over the duration of a single gate using memory benchmarking and Ramsey spectroscopy. These measurements show a loss of contrast of a few times 1×10^{-4}, which suggest our two-qubit gate error is dominated by qubit dephasing. We can also estimate potential improvement gained from dynamically decoupling the qubit during the gate by applying continuous dynamical qubit decoupling during Ramsey spectroscopy, which points to a reduction of our gate error to the 1×10^{-4} level.

[1] 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).

[2] 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).

[3] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver, High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler, Phys. Rev. X **13**, 031035 (2023).

[4] R. Acharya, *et al.*, Google Quantum AI, Suppressing quantum errors by scaling a surface code logical qubit, Nature **614**, 676 (2023).

[5] M. P. da Silva, *et al.*, Demonstration of logical qubits and repeated error correction with better-than-physical error rates, ArXiv:2404.02280.

[6] A. D. King, *et al.*, Quantum critical dynamics in a 5000-qubit programmable spin glass, Nature **617**, 61 (2023).

[7] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, Nature **618**, 500 (2023).

[8] D. Bluvstein, *et al.*, Logical quantum processor based on reconfigurable atom arrays, Nature **626**, 58 (2024).

[9] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, Nature (2025).

[10] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Entanglement of single-atom quantum bits at a distance, Nature **449**, 68 (2007).

[11] 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).

[12] Y. P. Zhong, H.-S. Chang, K. J. Satzinger, M.-H. Chou, A. Bienfait, C. R. Conner, É. Dumur, J. Grebel, G. A. Peairs, R. G. Povey, D. I. Schuster, and A. N. Cleland, Violating Bell's inequality with remotely connected superconducting qubits, Nat. Phys. **15**, 741 (2019).

[13] 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).

[14] 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).

[15] 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).

[16] 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).

[17] A. M. J. Zwerver, *et al.*, Qubits made by advanced semiconductor manufacturing, Nat. Electron. **5**, 184 (2022).

[18] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland, Surface-electrode architecture for ion-trap quantum information processing, Quantum Info. Comput. **5**, 419 (2005).

[19] 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).

[20] S. A. Moses, *et al.*, A race-track trapped-ion quantum processor, Phys. Rev. X **13**, 041052 (2023).

[21] K. K. Mehta, C. D. Bruzewicz, R. McConnell, R. J. Ram, J. M. Sage, and J. Chiaverini, Integrated optical addressing of an ion qubit, Nat. Nanotechnol. **11**, 1066 (2016).

[22] C. Mordini, A. R. Vasquez, Y. Motohashi, M. Müller, M. Malinowski, C. Zhang, K. K. Mehta, D. Kienzler, and J. P. Home, Multi-zone trapped-ion qubit control in an integrated photonics QCCD device, Phys. Rev. X **15**, 011040 (2025).

[23] J. Kwon, W. J. Setzer, M. Gehl, N. Karl, J. Van Der Wall, R. Law, M. G. Blain, D. Stick, and H. J. McGuinness, Multi-site integrated optical addressing of trapped ions, Nat. Commun. **15**, 3709 (2024).

[24] C. W. Hogle, D. Dominguez, M. Dong, A. Leenheer, H. J. McGuinness, B. P. Ruzic, M. Eichenfield, and D. Stick, High-fidelity trapped-ion qubit operations with scalable photonic modulators, npj Quantum Inf. **9**, 74 (2023).

[25] 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).

[26] 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).

[27] M. A. Weber, M. F. Gely, R. K. Hanley, T. P. Harty, A. D. Leu, C. M. Löschnauer, D. P. Nadlinger, and D. M. Lucas, Robust and fast microwave-driven quantum logic for trapped-ion qubits, Phys. Rev. A **110**, L010601 (2024).

[28] D. Leibfried, E. Knill, C. Ospelkaus, and D. J. Wineland, Transport quantum logic gates for trapped ions, Phys. Rev. A **76**, 032324 (2007).

[29] 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).

[30] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, Blueprint for a microwave trapped-ion quantum computer, Sci. Adv. **3**, e1601540 (2017).

[31] 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).

[32] 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, Phys. Rev. Lett. **131**, 020601 (2023).

[33] R. T. Sutherland, R. Srinivas, and D. T. C. Allcock, Individual addressing of trapped ion qubits with geometric phase gates, Phys. Rev. A **107**, 032604 (2023).

[34] U. Warring, C. Ospelkaus, Y. Colombe, K. R. Brown, J. M. Amini, M. Carsjens, D. Leibfried, and D. J. Wineland, Techniques for microwave near-field quantum control of trapped ions, Phys. Rev. A **87**, 013437 (2013).

[35] R. T. Sutherland, R. Srinivas, S. C. Burd, D. Leibfried, A. C. Wilson, D. J. Wineland, D. T. C. Allcock, D. H. Slichter, and S. B. Libby, Versatile laser-free trapped-ion entangling gates, New J. Phys. **21**, 033033 (2019).

[36] L. Enthoven, M. Babaie, and F. Sebastiano, Optimizing the electrical interface for large-scale color-center quantum processors, IEEE Trans. Quantum Eng. **5**, 5500317 (2024).

[37] 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).

[38] F. Mintert and C. Wunderlich, Ion-trap quantum logic using long-wavelength radiation, Phys. Rev. Lett. **87**, 257904 (2001).

[39] 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, Phys. Rev. Lett. **131**, 120601 (2023).

[40] N. K. Lysne, J. F. Niedermeyer, A. C. Wilson, D. H. Slichter, and D. Leibfried, Individual addressing and state readout of trapped ions utilizing rf micromotion, ArXiv:2402.05857.

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

[42] M. Malinowski, D. Allcock, and C. Ballance, How to wire a 1000-qubit trapped-ion quantum computer, PRX Quantum **4**, 040313 (2023).

[43] H. Ball, W. D. Oliver, and M. J. Biercuk, The role of master clock stability in quantum information processing, npj Quantum Inf. **2**, 16033 (2016).

[44] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland, Errors in trapped-ion quantum gates due to spontaneous photon scattering, Phys. Rev. A **75**, 042329 (2007).

[45] I. D. Moore, W. C. Campbell, E. R. Hudson, M. J. Boguslawski, D. J. Wineland, and D. T. C. Allcock, Photon scattering errors during stimulated Raman transitions in trapped-ion qubits, Phys. Rev. A **107**, 032413 (2023).

[46] R. T. Sutherland, Q. Yu, K. M. Beck, and H. Häffner, One- and two-qubit gate infidelities due to motional errors in trapped ions and electrons, Phys. Rev. A **105**, 022437 (2022).

[47] H. J. Metcalf and P. van der Straten, Laser cooling and trapping of atoms, JOSA B **20**, 887 (2003).

[48] A. H. Burrell, D. J. Szwer, S. C. Webster, and D. M. Lucas, Scalable simultaneous multiqubit readout with 99.99% single-shot fidelity, Phys. Rev. A **81**, 040302 (2010).

[49] J. T. Merrill and K. R. Brown, Progress in compensating pulse sequences for quantum computation, in *Quantum Information and Computation for Chemistry*, edited by S. A. Rice, A. R. Dinner, and S. Kais (2024).

[50] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and

D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).
[51] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and robust randomized benchmarking of quantum processes, Phys. Rev. Lett. **106**, 180504 (2011).
[52] A. Sørensen and K. Mølmer, Quantum computation with ions in thermal motion, Phys. Rev. Lett. **82**, 1971 (1999).
[53] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe, Coherent error suppression in multiqubit entangling gates, Phys. Rev. Lett. **109**, 020503 (2012).
[54] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland, Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate, Nature **422**, 412 (2003).
[55] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap* (CRC Press, New York, 1994).
[56] P. J. J. O'Malley, *et al.*, Qubit metrology of ultralow phase noise using randomized benchmarking, Phys. Rev. Appl. **3**, 044009 (2015).
[57] L. Viola and S. Lloyd, Dynamical suppression of decoherence in two-state quantum systems, Phys. Rev. A **58**, 2733 (1998).
[58] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker, Robust trapped-ion quantum logic gates by continuous dynamical decoupling, Phys. Rev. A **85**, 040302 (2012).
[59] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas, High-fidelity trapped-ion quantum logic using near-field microwaves, Phys. Rev. Lett. **117**, 140501 (2016).
[60] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland, Long-lived qubit memory using atomic ions, Phys. Rev. Lett. **95**, 060502 (2005).
[61] A. Parra-Rodriguez, P. Lougovski, L. Lamata, E. Solano, and M. Sanz, Digital-analog quantum computation, Phys. Rev. A **101**, 022305 (2020).
[62] B. Foxen, *et al.*, Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms, Phys. Rev. Lett. **125**, 120504 (2020).
[63] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, and A. Wallraff, Improving the performance of deep quantum optimization algorithms with continuous gate sets, PRX Quantum **1**, 020304 (2020).
[64] L. Clinton, T. Cubitt, B. Flynn, F. M. Gambetta, J. Klassen, A. Montanaro, S. Piddock, R. A. Santos, and E. Sheridan, Towards near-term quantum simulation of materials, Nat. Commun. **15**, 211 (2024).
[65] 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).
[66] 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 a radio-frequency Paul trap, Phys. Rev. Lett. **130**, 173202 (2023).
[67] 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).
[68] 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).
[69] C.-Y. Shih, S. Motlakunta, N. Kotibhaskar, M. Sajjan, R. Hablützel, and R. Islam, Reprogrammable and high-precision holographic optical addressing of trapped ions for scalable quantum control, npj Quantum Inf. **7**, 57 (2021).
[70] 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).
[71] A. Binai-Motlagh, M. Day, N. Videnov, N. Greenberg, C. Senko, and R. Islam, A guided light system for agile individual addressing of Ba^+ qubits with 10^{-1} level intensity crosstalk, ArXiv:2302.14711.
[72] A. S. Sotirova, B. Sun, J. D. Leppard, A. Wang, M. Wang, A. Vazquez-Brennan, D. P. Nadlinger, S. Moser, A. Jesacher, C. He, F. Pokorny, M. J. Booth, and C. J. Ballance, Low cross-talk optical addressing of trapped-ion qubits using a novel integrated photonic chip, ArXiv:2310.13419.
[73] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland, Microwave quantum logic gates for trapped ions, Nature **476**, 181 (2011).
[74] K. K. Mehta and R. J. Ram, Precise and diffraction-limited waveguide-to-free-space focusing gratings, Sci. Rep. **7**, 2019 (2017).
[75] G. N. West, W. Loh, D. Kharas, C. Sorace-Agaskar, K. K. Mehta, J. Sage, J. Chiaverini, and R. J. Ram, Low-loss integrated photonics for the blue and ultraviolet regime, APL Photonics **4**, 026101 (2019).
[76] J. R. Sambles, K. C. Elsom, and G. Sharp-Dent, The effect of sample thickness on the resistivity of aluminium, J. Phys. F: Metal Phys. **11**, 1075 (1981).
[77] R. D. Delaney, L. R. Sletten, M. J. Cich, B. Estey, M. Fabrikant, D. Hayes, I. M. Hoffman, J. Hostetter, C. Langer, S. A. Moses, A. R. Perry, T. A. Peterson, A. Schaffer, C. Volin, G. Vittorini, and W. C. Burton, Scalable multi-species ion transport in a grid based surface-electrode trap, ArXiv:2403.00756.
[78] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, E. Magesan, J. A. Smolin, J.-B. Yau, V. P. Adiga, M. Sandberg, M. Brink, J. M. Chow, and J. S. Orcutt, Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors, npj Quantum Inf. **7**, 129 (2021).
[79] C. Berke, E. Varvelis, S. Trebst, A. Altland, and D. P. DiVincenzo, Transmon platform for quantum computing challenged by chaotic fluctuations, Nat. Commun. **13**, 2495 (2022).
[80] N. P. de Leon, K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. S. Palmer, N. Samarth,

S. Sangtawesin, and D. W. Steuerman, Materials challenges and opportunities for quantum computing hardware, Science **372**, eabb2823 (2021).

[81] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas, Reduction of heating rate in a microfabricated ion trap by pulsed-laser cleaning, New J. Phys. **13**, 123023 (2011).

[82] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland, 100-fold reduction of electric-field noise in an ion trap cleaned with in situ argon-ion-beam bombardment, Phys. Rev. Lett. **109**, 103001 (2012).

[83] K. R. Brown, J. Chiaverini, J. M. Sage, and H. Häffner, Materials challenges for trapped-ion quantum computers, Nat. Rev. Mater. **6**, 892 (2021).

[84] M. Berlin-Udi, C. Matthiesen, P. N. T. Lloyd, A. M. Alonso, C. Noel, B. Saarel, C. A. Orme, C.-E. Kim, A. J. Nelson, K. G. Ray, V. Lordi, and H. Häffner, Changes in electric field noise due to thermal transformation of a surface ion trap, Phys. Rev. B **106**, 035409 (2022).

[85] G. J. Beck, J. P. Home, and K. K. Mehta, Grating design methodology for tailored free-space beam-forming, J. Lightwave Technol. **42**, 4939 (2024).

[86] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, A. S. Sørensen, D. Leibfried, and D. J. Wineland, Dissipative production of a maximally entangled steady state of two quantum bits, Nature **504**, 415 (2013).

[87] D. C. Cole, S. D. Erickson, G. Zarantonello, K. P. Horn, P.-Y. Hou, J. J. Wu, D. H. Slichter, F. Reiter, C. P. Koch, and D. Leibfried, Resource-efficient dissipative entanglement of two trapped-ion qubits, Phys. Rev. Lett. **128**, 080502 (2022).

[88] M. Malinowski, C. Zhang, V. Negnevitsky, I. Rojkov, F. Reiter, T.-L. Nguyen, M. Stadler, D. Kienzler, K. K. Mehta, and J. P. Home, Generation of a maximally entangled state using collective optical pumping, Phys. Rev. Lett. **128**, 080503 (2022).

[89] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, and D. J. Wineland, Experimental entanglement of four particles, Nature **404**, 256 (2000).