



PAPER

Universal control of ion qubits in a scalable microfabricated planar trap

OPEN ACCESS

RECEIVED

17 September 2015

REVISED

7 December 2015

ACCEPTED FOR PUBLICATION

29 December 2015

PUBLISHED

19 February 2016

C D Herold, S D Fallek, J T Merrill, A M Meier, K R Brown, C E Volin and J M Amini

Georgia Tech Research Institute, Atlanta, GA 30332, USA

E-mail: Creston.Herold@gtri.gatech.edu**Keywords:** quantum computing, quantum control, trapped ions, surface-electrode trapSupplementary material for this article is available [online](#)

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

We demonstrate universal quantum control over chains of ions in a surface-electrode ion trap, including all the fundamental operations necessary to perform algorithms in a one-dimensional, nearest-neighbor quantum computing architecture. We realize both single-qubit operations and nearest-neighbor entangling gates with Raman laser beams, and we interleave the two gate types. We report average single-qubit gate fidelities as high as 0.970(1) for two-, three-, and four-ion chains, characterized with randomized benchmarking. We generate Bell states between the nearest-neighbor pairs of a three-ion chain, with fidelity up to 0.84(2). We combine one- and two-qubit gates to perform quantum process tomography of a CNOT gate in a two-ion chain, and we report an overall fidelity of 0.76(3).

1. Introduction

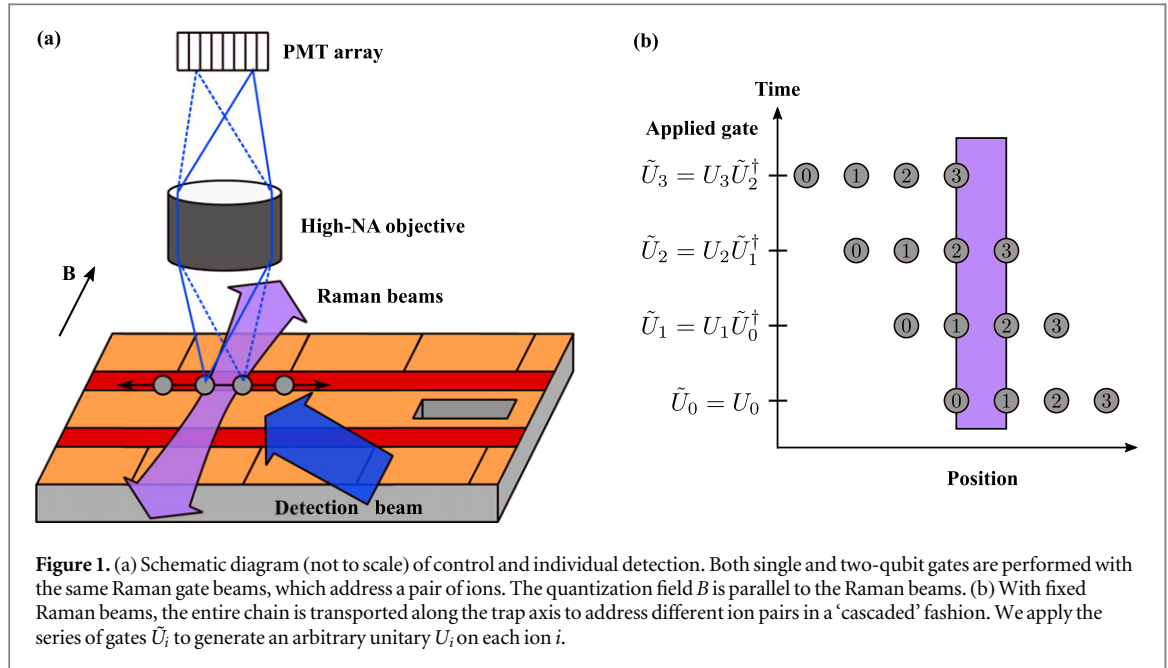
The size of the Hilbert space describing a quantum computer grows exponentially with the number of qubits. Despite this, any quantum computation can be implemented to arbitrary accuracy by composing elements from a set of one- and two-qubit operations [1, 2]. Any such operation set is called a universal gate set. Universal gate sets have been demonstrated with NMR [3], superconducting [4, 5], neutral atom [6], and trapped ion [7, 8] qubits.

Trapped ions are attractive for building an extensible quantum system due to their long coherence times and ability to be entangled via photons and phonons [9, 10]. Ion trap quantum computer architectures could use transport [11] or photonic interconnects [12] to provide connectivity between disparate regions of the quantum computer. These architectures can benefit from the fabrication precision and the scalability of microfabricated surface-electrode ion traps [13]. While the components of a universal gate set have been shown in microfabricated traps (see e.g. [14–18]), a complete universal gate set has not yet been demonstrated in a single system.

In this work, we demonstrate a universal gate set on two- and three-ion chains in a surface-electrode trap. Our architecture is based on transporting the chain sequentially through a single gate location to provide ion addressing; the gate beams are pulsed on only when the chain is stationary. We perform both single-qubit operations and nearest-neighbor entangling gates. We measure single-qubit gate fidelities for chains of two to four ions using randomized benchmarking, Bell state fidelities in the nearest-neighbor pairs of a three-ion chain, and process fidelity of a CNOT gate in a two-ion chain. The process fidelity measurement demonstrates our ability to interleave one- and two-qubit operations in the same experiment. All features are scalable to longer ion chains.

2. System overview

The experimental system is shown schematically in figure 1(a). We trap chains of $^{171}\text{Yb}^+$ ions 60 μm above the surface-electrode ion trap described in [19]. This trap features through-chip vias instead of wirebonds. The



improved optical access provided by this feature allows for tight focusing of the Raman gate beams, which propagate parallel to the surface of the trap. Qubits are encoded in the hyperfine clock states of $^{171}\text{Yb}^+$ with $|0\rangle \equiv |^2S_{1/2}, F=0, m_F=0\rangle$ and $|1\rangle \equiv |F=1, m_F=0\rangle$. Details on trapping, state preparation, and readout of $^{171}\text{Yb}^+$ ions can be found in [20].

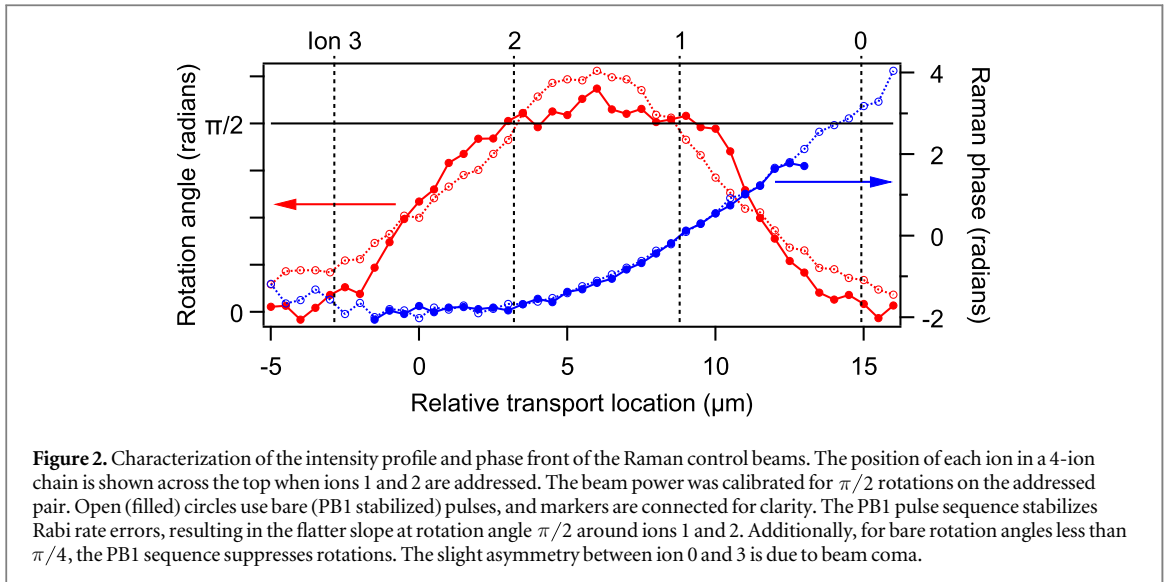
We load individual ions and merge each ion into a harmonic well. In this way, we deterministically build chains of a desired number of ions. During state detection, each ion in the chain is imaged onto a separate channel of a multi-channel photomultiplier tube (PMT) [21]. This allows us to detect the state of each ion in the chain independently. Despite efforts to reduce crosstalk by using only every fourth PMT channel, we observe 5% crosstalk between ions due to coma in the imaging system. This crosstalk is calibrated and accounted for with a joint determination of which ions were bright in a given experiment set. With 400 μs of fluorescence collection, our state detection fidelity is 0.98(1).

A single pair of Raman beams provides both one- and two-qubit gates. The beams are derived from a mode-locked tripled YAG laser, in a configuration similar to that described in [22, 23]. The circularly polarized beams propagate perpendicular to the trap axis and parallel to the quantizing magnetic field. In this configuration, the gate beams do not couple to the axial motional modes or the $|F=1, m_F=\pm 1\rangle$ Zeeman levels. Furthermore, the trapping potentials are designed so that one set of radial secular modes is rotated normal to the beam direction. This limits coupling of the Raman beams to only a single set of motional sideband transitions along with the carrier, reducing the effect of off-resonant shifts and removing the dependence of the carrier transition on the temperature of these modes. After Doppler cooling all modes, we perform sideband cooling with the Raman beams only on the radial mode used for gates, resulting in sub-quanta occupation of excited motional states.

Mølmer–Sørensen (MS) gates between pairs of neighboring ions provide the entangling gates for this system. The Raman beam in one direction (consisting of the red- and blue-detuned beams¹) is focused to a waist ($1/e^2$ half-width of the intensity profile) of roughly 7 μm while the opposing Raman beam has a waist of 13.5 μm . The average power in these beams is on the order of 10 and 100 mW, respectively. The $1/e$ half-width of the Raman interaction Rabi rate (see figure 2) is approximately equal to our ion spacing of about 6 μm , permitting pairwise addressing of ions within a chain as in figure 1. The Raman beams remain static throughout the course of the experiment; we transport the ion chain to bring different pairs of ions into the beam.

Still addressing ions pairwise, we employ the cascading transport scheme shown in figure 1(b) to implement an arbitrary set of one-qubit unitary operations $\{U_i\}$, where i is the ion index. The ion chain is first transported such that the $i=0, 1$ ion pair (i.e. the 01 pair) is illuminated by the laser. The unitary $\tilde{U}_0 = U_0$ is performed, leaving the endmost ion in the desired state. Then the ions are transported to address the subsequent pair. The operation $\tilde{U}_1 = U_1 \tilde{U}_0^\dagger$ results in the net operation $\tilde{U}_1 \tilde{U}_0 = U_1 U_0^\dagger U_0 = U_1$ on the second ion, accounting for the

¹ These beams are generated with a single acousto-optic modulator, the output of which is imaged onto the trapping region with several telescopes to ensure the two beams overlap at their foci.



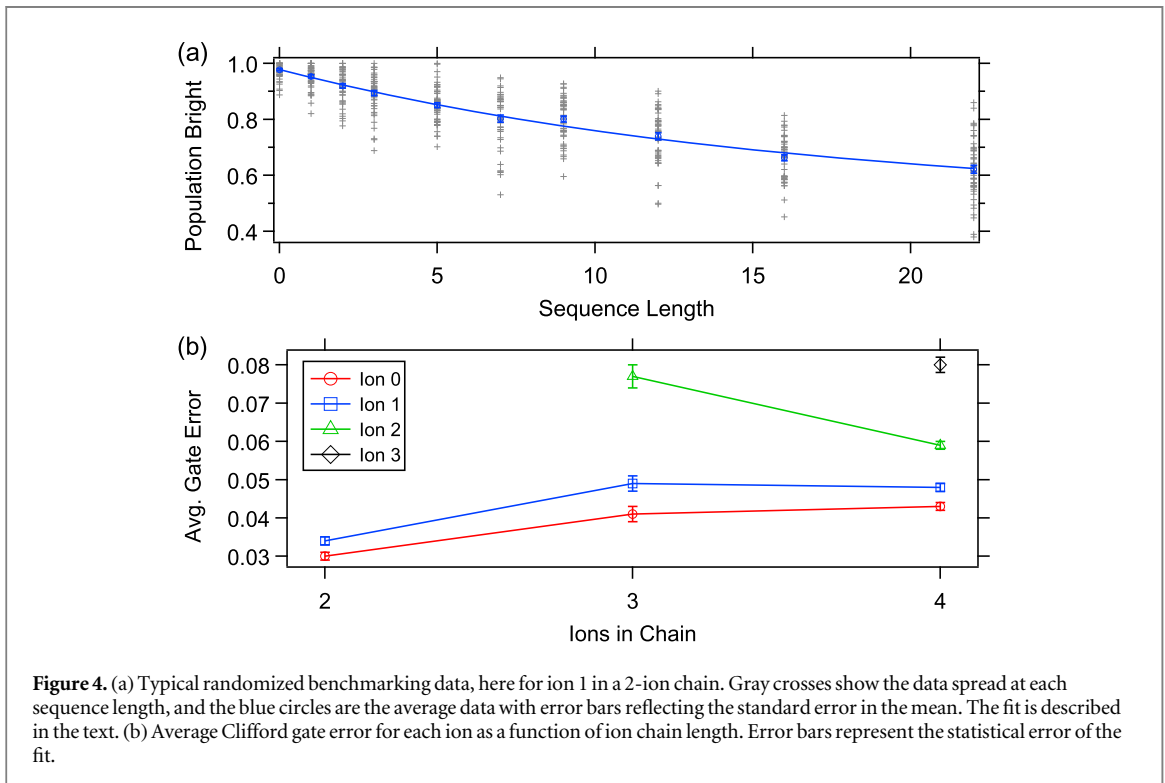
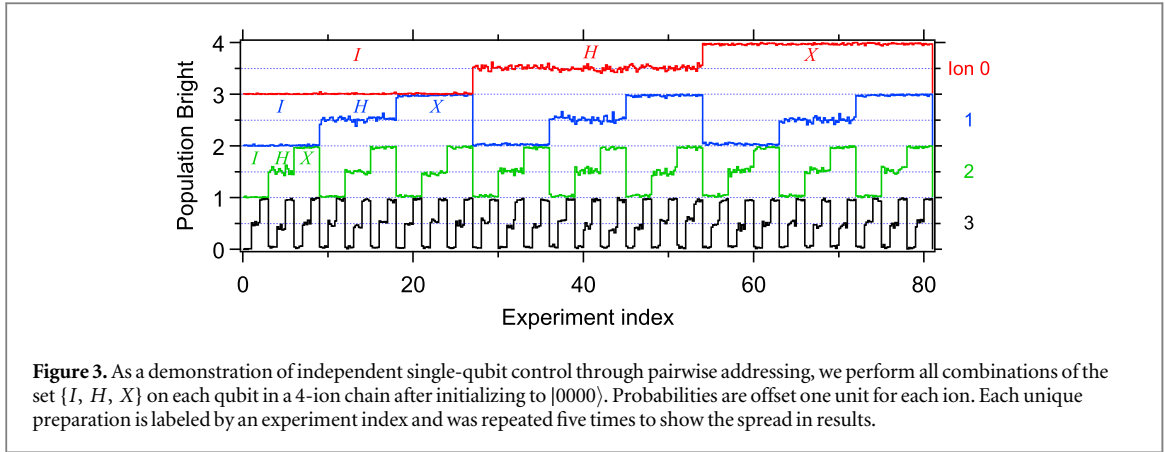
effect of the first operation. This procedure is repeated sequentially along the chain until all ions are targeted, resulting in the desired gate on each ion. Universal control is achieved in this transport architecture by interleaving such cascaded single-qubit gates with pairwise MS gates.

In this work, ions are only transported over very short distances corresponding to the spacing between ions or the length of the chain. Each transport step is performed adiabatically in $100 \mu\text{s}$. We varied the transport duration between $100 \mu\text{s}$ and 1 ms but could not see any effect on gate fidelity at the present error levels due to heating. In addition to shuttling the ions through the gate beams, we adjust the strength of the axial harmonic well to optimize the ion spacing for each type of gate. For two-qubit operations, where a strong coupling between the ions' motion is desired, the axial well is stronger than for single qubit operations, where crosstalk to unaddressed ions needs to be minimized. For Doppler and sideband cooling, the typical axial trap frequency is $2\pi \times 0.5 \text{ MHz}$, and the radial frequencies are $2\pi \times (1.8, 2.1) \text{ MHz}$. The lower radial frequency corresponds to the mode addressed for all gates. We reduce the axial frequency by a factor of $\sqrt{2}$ from the cooling trap frequencies to spread the ions out for single-qubit gates, and we increase it by $\sqrt{1.5}$ times to further split the radial modes for two-qubit operations. Similarly, for detection the ion spacing is optimized to match the spacing of the photomultiplier array.

3. Single-qubit control

As described above, all single-qubit control is implemented via pairwise addressing. For each ion pair, we calibrate the ion chain position to balance the carrier transition Rabi rate on the two addressed ions. We then calibrate the gate beam intensity to give a $\pi/2$ rotation in $4 \mu\text{s}$. These calibrations are stabilized against Rabi rate errors using PB1 compensating pulse sequences [24–26], which are composed of the calibrated $\pi/2$ rotations. Including hardware programming overhead, a PB1 composite $\pi/2$ operation is performed in $120 \mu\text{s}$. Arbitrary single-qubit unitaries are constructed from these composite $\pi/2$ rotations about different axes in the equatorial plane of the Bloch sphere. The angle the axis of rotation makes with the x -axis is set by the phase of the Raman beams and is controlled via an acousto-optic modulator. We developed a compiler (see appendix A) to efficiently implement arbitrary unitaries in the fewest $\pi/2$ rotations.

Implementation of an arbitrary set of unitaries across the chain under actual experimental conditions is more complicated than the idealized scheme presented in figure 1(b). First, wavefront curvature across the gate beams causes ions to experience a position-dependent phase of the Raman interaction (see figure 2), resulting in a spatially varying $\pi/2$ rotation axis. Second, the pairwise addressing of our gate beams is imperfect, and ions adjacent to the addressed pair are observed to flop at reduced but significant rates. To a large degree, this crosstalk is mitigated by our choice of compensating pulse sequence since PB1, a so-called ‘passband’ sequence, also suppresses small excitations (figure 2). Due to comatic aberrations, the crosstalk is asymmetric, and it is a reasonable approximation to treat the right-hand side of our gate beam as crosstalk-free when PB1 is used. Therefore, we only calibrate and correct for the small rotation angle of left-side crosstalk. This asymmetry dictates the direction in which we cascade single-qubit operations. As an example of independent single-qubit



control, we prepare each ion in a four ion chain in all 3^4 combinations of $\{|0\rangle, |1\rangle, |+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}\}$ (figure 3).

To characterize the single-qubit gate fidelity, we employ randomized benchmarking [27]. Starting with all ions in state $|0\rangle$, a randomly chosen Clifford gate is performed on each ion, cascading the gates across the chain as described above. Each Clifford operation requires on average 400 μs of Raman pulses and is followed by 100 μs transport between gate locations. Subsequent passes build a random sequence of Clifford gates of length L on each ion. The final pass consists of an inverting gate calculated to take each ion to $|1\rangle$. The results for one ion in a two-ion chain are shown in figure 4(a) for varying sequence lengths. We performed benchmarking on two- to four-ion chains and fit the average gate fidelity \bar{F}_{RB} for each ion, following [28], to

$$\bar{F}_{\text{RB}} = \frac{1}{2} + \frac{1}{2}(1 - 2\varepsilon_m)(1 - 2\varepsilon_g)^L, \quad (1)$$

where ε_g is the average gate error and ε_m includes error from state preparation and measurement as well as error in the inverting gate. The average ε_m for fits on all chain lengths is 0.025(2). We also performed a three parameter fit as in [29], where the constant offset is allowed to vary from $\frac{1}{2}$, however these fits gave non-physical results with negative ε_m . Because the measurement imbalance in our apparatus is small, the two-parameter fit with offset fixed at $\frac{1}{2}$ is justified; a typical reduced- χ^2 for the fits is 1.6.

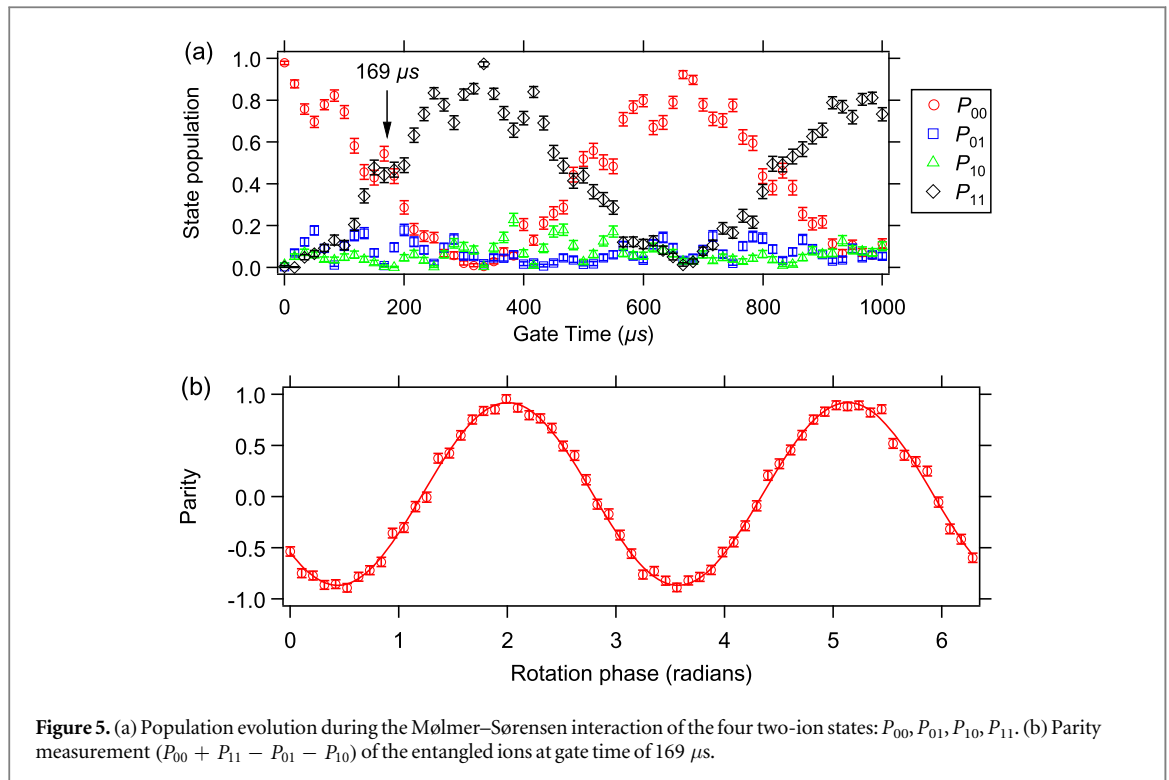


Figure 5. (a) Population evolution during the Mølmer–Sørensen interaction of the four two-ion states: P_{00} , P_{01} , P_{10} , P_{11} . (b) Parity measurement ($P_{00} + P_{11} - P_{01} - P_{10}$) of the entangled ions at gate time of 169 μs .

As can be seen in figure 4(b), we measure average single-qubit gate fidelity as high as 0.970(1) for two-ion chains. However, the average gate error increases with ion chain length and position within the chain in the cascade direction. Longer chains require more calibrations and more operations, which lead to more time for dephasing between successive operations on a given qubit. As a result, one expects gate error to increase with chain length.

That the end ions, i.e. ions 2 and 3, have a higher gate error can be explained by imperfect crosstalk correction. As we noted in the discussion of figure 2, the wings of our comatic gate beams cause small-rotation crosstalk on ions adjacent to the addressed pair. The PB1 sequence suppresses the average value of this crosstalk to the point where it is negligible on the left-hand (low ion number) side of the beam; on the right-hand side, we calibrate and correct for the remaining crosstalk in our cascaded gate sequences. Therefore, ions 0 and 1, unlike ions 2 and 3, are not subject to this crosstalk. Despite our careful calibration, we suspect that the increased slope due to PB1 (figure 2) may be *increasing* our sensitivity to crosstalk fluctuations on the right-hand (high ion number) side of the beam. Improving the passive stability of our gate beams should increase the overall single-qubit gate fidelity. We note that the average gate error for ion 2 decreased from a three-ion to four-ion chain. This is most likely due to relatively poorer calibration; systematic errors are not included in the error bars.

4. Two-qubit gates

4.1. Two-ion entanglement

To complete our universal gate set, we require an entangling gate. Here we employ the MS gate on a pair of ions [30]. During MS gates, we strengthen our axial confinement to increase the frequency splitting of our radial modes in comparison to single-qubit gates. This reduces motional excitation of undesired modes, yielding higher fidelity. State evolution of a typical gate is shown in figure 5(a) where ion motion has been excited and de-excited twice per gate, i.e. two ‘loops’ per gate, resulting in a 169 μs gate time. This is substantially faster than our PB1 compensated implementation of single-qubit unitaries.

To verify entanglement following the MS gate, Bell state fidelity estimation is performed using the ‘parity flopping’ technique [31]. The resulting parity curve is shown in figure 5(b). Bell state fidelity is measured to be $F_{\text{Bell}} = 0.93(2)$. We expect this fidelity is largely limited by imperfect overlap of the red and blue sideband beams during the MS interaction; a redesigned optics system will remedy this issue. This MS gate, along with the individual qubit rotations in section 3, comprise our universal gate set for a two-ion chain. In section 5, we demonstrate combined one- and two-qubit gates to produce a CNOT gate.

Table 1. Parity measurement for pairwise entanglement in a three-ion chain. P_{00} and P_{11} are the measured populations for the state where the unaddressed ion is in $|0\rangle$, i.e. for an MS gate on the 12-pair, the states $|000\rangle$ and $|011\rangle$. Similarly, the parity amplitude is extracted with the unaddressed ion in state $|0\rangle$. P_1 is a measure of the population build-up on the unaddressed ion, where we trace over the addressed ions.

Ion pair	P_{00}	P_{11}	Parity amp.	P_1	F_{Bell}
No echo 01	0.32(1)	0.45(1)	0.66(1)	0.18(1)	0.72(1)
Echo 01	0.37(1)	0.46(1)	0.62(1)	0.10(1)	0.72(2)
No echo 12	0.47(1)	0.41(1)	0.81(1)	0.094(7)	0.84(2)
Echo 12	0.41(1)	0.47(1)	0.76(1)	0.023(3)	0.82(2)

4.2. Pairwise entangling gates on three ions

Universal control of ion chains requires the ability to entangle any pair of ions. Within our architecture, we entangle neighboring ions in pairs. However, this entanglement can be transferred to non-neighboring pairs through SWAP gates, which can be composed from CNOT gates. In this section, we characterize entangling gates on pairs of neighboring ions within a three-ion chain, the last component of a universal gate set for three ions.

To estimate the MS gate fidelity, we measure F_{Bell} as in the previous section. All ions are prepared in $|0\rangle$, the two targeted ions are shifted into the laser beam, then an MS gate and analysis operation are performed. Parity is calculated using only the state populations where the unaddressed ion remains in the desired $|0\rangle$ state, and the resulting fidelities are summarized for each targeted pair in table 1.

We observe moderate build-up of undesired population in the $|1\rangle$ state on the unaddressed ion during the MS gate; this is reported as P_1 in table 1. Due to the spatial asymmetry of our gate beam (shown in the rotation angle of figure 2), the MS crosstalk on the 01 pair is worse than the 12 pair. In appendix B, we describe an ‘echoing’ sequence which decouples a subset of ions from a general MS interaction. For an MS gate on two ions in a three-ion chain, we perform the MS gate with two motional loops. The echo is inserted after each loop and flips the spin of the untargeted ion (or equivalently the spins of the targeted ions) with a Y -gate. Early results show that the echo sequence reduces P_1 ; however, the Bell-state fidelity is not improved. This is likely due to the error introduced by the additional single-qubit operations used in the echo. We expect that increased single-qubit gate fidelity should make the echo more effective and allow us to perform high-fidelity MS gates on any adjacent pair of ions in longer chains.

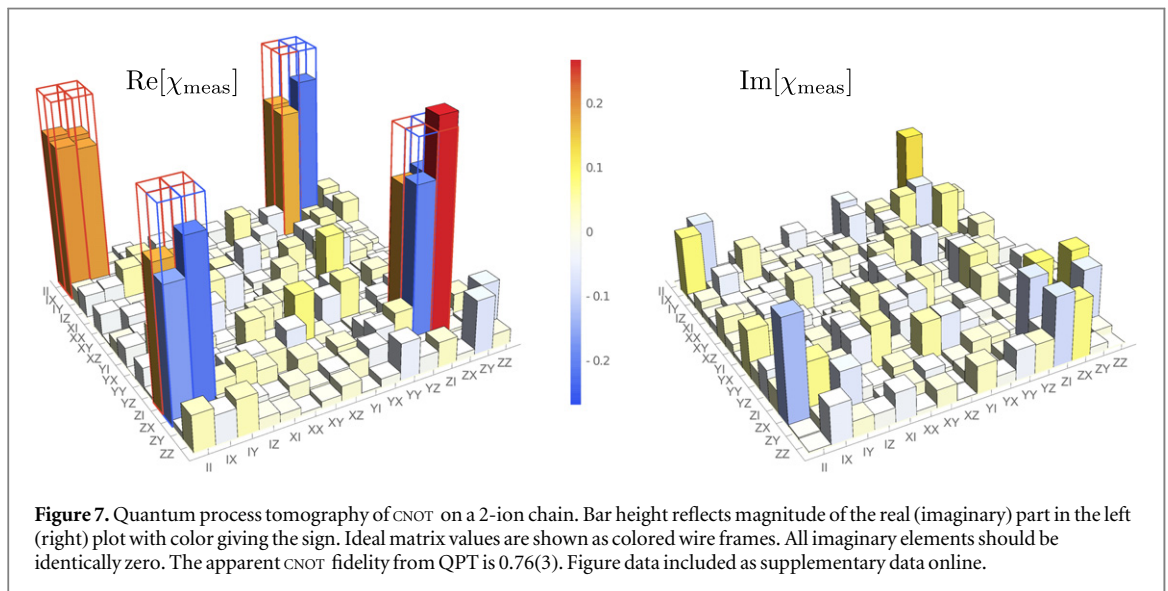
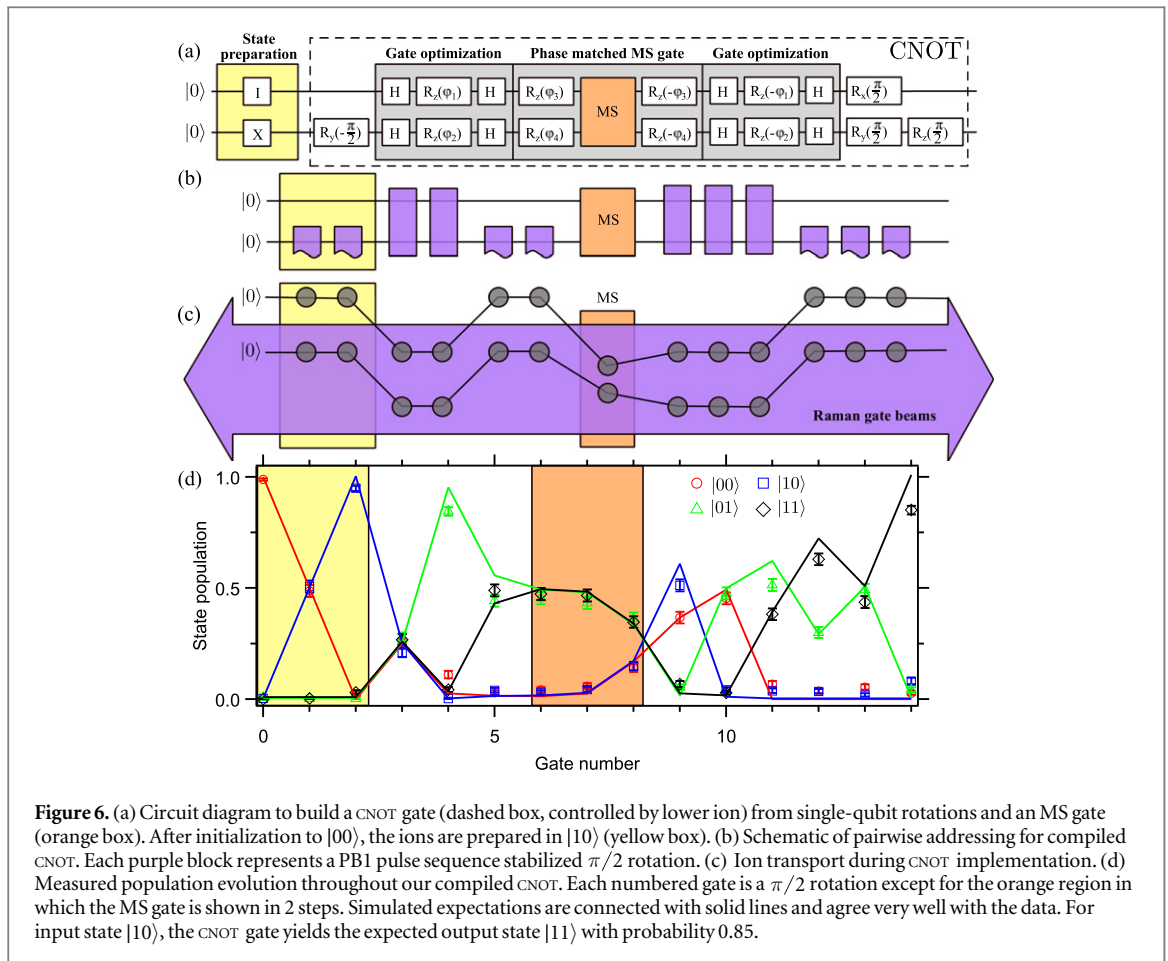
5. CNOT characterization

For a demonstration of a small circuit using our universal gate set, we perform the CNOT gate on two ions. The CNOT is built from an MS gate wrapped with a particular set of single qubit rotations, as shown in figure 6(a). In figures 6(b) and (c) we represent how the CNOT circuit fits into the cascading architecture discussed in section 2. Single-qubit rotations are performed in two cascading passes, one before the MS gate and one after. As in section 4, axial confinement is increased for the MS gate and then subsequently relaxed for single-qubit operations, adding an additional $50 \mu\text{s}$ for each. We utilize four steps of transport ($100 \mu\text{s}$ each, as before), excluding the transport needed for state preparation and detection.

Successful combination of one- and two-qubit operations requires accounting for the Raman phase difference due to increased axial confinement during MS gates. We add z -rotations (R_z) before and after the MS gate (with rotation angles φ_3 and φ_4) to ensure it is in phase with the single-qubit rotations. In addition, we add the optimization operations shown in figure 6(a). These gates do not affect the outcome of the ideal circuit, but they allow us to minimize the number of physical operations in the compiled single-qubit cascades by varying φ_1 and φ_2 . Further discussion of this procedure is provided in appendix A.

As a diagnostic tool, we measure the population evolution throughout the implementation of the CNOT (markers in figure 6(d)). To measure the populations at a given point within the CNOT sequence, we trigger an override on the gate beams, shutting down all subsequent operations. This freezes the populations, which are read out at the end of the CNOT sequence. When combined with on-the-fly simulation of ideal rotations (points connected by solid lines), such a scan has proved a useful tool for diagnosing calibration errors. We have also built an optimization routine into our simulation tool to find the best parameter calibrations to match observed populations. Thus, from a small set of carefully chosen gate scans like figure 6(d), we can quickly improve overall fidelity.

To systematically characterize our CNOT gate, we perform quantum process tomography (QPT) [32]. We prepare and analyze each ion in all combinations of the set $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}_y = (|0\rangle - i|1\rangle)/\sqrt{2}$ resulting in 256 measurements. The preparation and analysis gates are *not* compiled into the single-qubit gates in the



CNOT itself. Rather, two additional passes of single-qubit operations are used, similar to the state preparation (yellow box) in figure 6. Each measurement is repeated 160 times and the maximum likelihood probability to be in $|00\rangle$ is used to generate the measured process matrix χ_{meas} , shown in figure 7. The fidelity of our operation with the ideal CNOT process, χ_{ideal} is measured as $F = \text{Tr}(\chi_{\text{ideal}}\chi_{\text{meas}}) = 0.76(3)$, where we have reported the mean and standard deviation of 10^4 parametric bootstrap resamplings of our measurements. From the resampled χ_{meas} we estimate a typical matrix element error bar to be 0.02. There are only 16 non-zero matrix elements of χ_{ideal} , all in the real part. Nearly all other measured matrix elements are consistent with zero within three standard deviations.

Using a simple uncorrelated error model to account for four single-qubit gates (from randomized benchmarking fidelity) and one MS gate (from Bell state fidelity), we expect an approximate fidelity of $(\bar{F}_{\text{RB}})^4 F_{\text{Bell}} = (0.97)^4 \times 0.93 = 0.82$. Yet, single-qubit errors in preparation and analysis also reduce QPT fidelity estimates. With the same QPT procedure, we measure a two-qubit identity operation fidelity of 0.95, leading us to estimate a QPT fidelity of 0.78 in good agreement with our measurement.

6. Outlook

We have demonstrated a universal gate set for two- and three-ion chains for the first time in a surface-electrode ion trap. Universality is realized through a novel control scheme which leverages the precision transport capability of our trap. This scheme is extensible to longer chains. Importantly, we demonstrated universal control in a *single* experiment by implementing a circuit consisting of single-qubit rotations and an MS gate to perform a CNOT gate. Single-qubit gates on four-ion chains were characterized as well as echoed MS gates on pairs within a three-ion chain. Minor system improvements currently underway should improve fidelities to levels where basic quantum algorithms are within reach.

Our results bolster the promise of scalable ion trap quantum computing. Through the integration of junctions, microfabricated ion traps have already been extended to two-dimensional architectures. Integrated optics are a demonstrated path to a scalable photonic interface with ions in large trap arrays. We have shown that microfabricated traps are also compatible with high-fidelity universal quantum control of ion chains.

Acknowledgments

This material is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) under US Army Research Office (ARO) contract W911NF1010231. All statements of fact, opinion, or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of IARPA, the ODNI, or the US Government.

Appendix A. Compiler

In our control scheme, we must implement arbitrary unitary operations on individual qubits for universal control. Even for algorithms designed around a discrete set of single-qubit gates, accounting for the crosstalk shown in figure 2 would require unitary operations outside that set. Here, we describe how we construct arbitrary unitaries from our calibrated, stabilized, physical operations. Additionally, methods to minimize the number of physical rotations required are detailed.

As discussed in section 3, by adjusting the phase of our Raman gate laser, we can perform rotations about any axis in the equatorial plane of the Bloch sphere. Defining the angle that axis makes with the x -axis as φ , our physical operations $R_\varphi(\pi/2)$ are composite $\pi/2$ rotations stabilized with the PB1 sequence. While we have constructed various analytic solutions to perform arbitrary unitaries in a fixed number of axis angles φ_j , direct numerical search has proved more efficient. In a test of 1000 randomly generated unitaries, we always find a solution in three or four rotations, with an average of 3.25.

Our compiler takes a set of requested circuit-level gates for the ion chain and computes a target unitary U_i for each ion i . First, U_0 is performed on the 01 ion pair, using the rotations calculated from our numerical search. While this leaves ion 0 in the desired state, the physical operations act in a z -rotated basis due to the phase front curvature of the Raman gate beams (figure 2). Each physical rotation on ion 1 is also through an angle of $\pi/2$; the rotation due to the crosstalk on ion 2 is much smaller. Neglecting crosstalk beyond next-nearest neighbors, the resulting physical operations are summarized in step 1 of table A1, where the net physical unitary $\tilde{U}_i^{(j)}$ denotes the physical effect on ion j of a unitary targeting ion i . In each step, the unitary $\tilde{U}_i^{(j=i)}$ is performed, while the compiler stores a unitary corresponding to the calibrated crosstalk on $j = [i + 1, i + 2]$. In step 2, we then need to perform the product of the targeted unitary U_1 and the inverse of the stored unitary from the previous step. Thus, in each subsequent step, we numerically search for physical rotations whose net effect is to undo the crosstalk from previous operations and perform the targeted unitary U_i .

To correct crosstalk, we store a unitary corresponding to the history for each ion throughout the algorithm. However, we can take further advantage of this storage to reduce the number of applied physical $R_\varphi(\pi/2)$, thereby improving overall fidelity. This is implemented in two ways. First, subsequent single-qubit gates are concatenated for each ion, whenever possible. If they do not need to be performed, we append the inverse operation to the stored unitaries. At the next physical operation cascade, the inverse history is composed with

Table A1. In contrast to the idealized pairwise control shown in figure 1(b), here we show how our compiler must account for crosstalk.

Step	Addressed				
	pair	Ion 0	Ion 1	Ion 2	Ion 3
1	01	$\tilde{U}_0^{(0)} = U_0$	$\tilde{U}_0^{(1)}$	$\tilde{U}_0^{(2)}$	—
2	12	—	$\tilde{U}_1^{(1)} = U_1(\tilde{U}_0^{(1)})^\dagger$	$\tilde{U}_1^{(2)}$	$\tilde{U}_1^{(3)}$
3	23	—	—	$\tilde{U}_2^{(2)} = U_2(\tilde{U}_1^{(2)}\tilde{U}_0^{(2)})^\dagger$	$\tilde{U}_2^{(3)}$
4	34	—	—	—	$\tilde{U}_3^{(3)} = U_3(\tilde{U}_2^{(3)}\tilde{U}_1^{(3)})^\dagger$
Net gate:		U_0	U_1	U_2	U_3

any additional requested gate, allowing us to perform multiple single qubit gates and crosstalk corrections in a single cascade of operations. Second, we often perform requested unitaries up to a stored R_z . An arbitrary unitary can almost always be performed in only two $R_\varphi(\pi/2)$ rotations, up to a missing R_z , which is nearly half that required for an exact unitary. This has the effect of transforming an ion's relative frame and can simply be appended to the stored unitary. Subsequent operations on this ion need only shift the Raman gate phase to account for the frame shift. In fact, without global operations, there is no need for a common definition of zero phase. Instead we define a phase zero for each ion by the Raman beam phase front when that ion is targeted.

In some instances, performing unitaries up to a phase still produces an exact result. For example, state preparation on the fiducial state with every ion in $|0\rangle$ is the same with an R_z inserted:

$$U_{\text{exact}}|0\rangle = (U_{\text{exact}}R_z^\dagger)R_z|0\rangle = U'R_z|0\rangle = U'|0\rangle \quad (\text{A.1})$$

since z-rotations do not change $|0\rangle$. Similarly, operations immediately prior to detection can be performed up to an R_z without effecting the probability to be in $|0\rangle$ or $|1\rangle$.

Prior to performing an MS gate, the stored unitary of all targeted ions must be physically undone as arbitrary single-qubit operations do not generally commute with the MS gate. Additionally, the relative frames of targeted ions must be rotated to match the MS gate beams. It would seem we are forced to perform exact single-qubit unitary operations around all MS gates. However, these requirements can be softened. If we diagonalize the MS interaction with appropriately rotated Hadamard gates on each ion, R_z operations will commute and can be tracked to the other side. This is the origin of the 'gate optimization' additions to figure 6(a) and results in only two $R_\varphi(\pi/2)$ prior to the MS gate and three after it. Exact unitaries would require six to eight $R_\varphi(\pi/2)$ operations.

In summary, our compiler stores a unitary history for operations on each ion. This enables us to track crosstalk corrections, unperformed single-qubit gates, and unperformed R_z rotations. We are able to program algorithms directly from a quantum circuit, and the compiler produces a compact set of operations which limit the number of physical rotations actually performed.

Appendix B. MS echo

We describe an 'echoing' scheme that enables pairwise entanglement despite the imperfect pairwise addressing of our MS gate beams. The intuitive motivation is that flipping the state of an ion in the MS basis (i.e. $|+\rangle_x \leftrightarrow |-\rangle_x$) halfway through the gate will cause the net effect of the MS interaction on this ion to be identity. In fact, we show that the scheme is exact (within the approximations outlined below) and could be used to entangle any m ions in an N ion chain with 2^{N-m} single-qubit flips on the untargeted ions; these 'echo' operations must be equally distributed throughout the MS gate, requiring it to be performed in 2^{N-m} equal steps.

Following [30], we write the interaction picture Hamiltonian, \tilde{H}_{int} , for two Raman beams with symmetric detuning δ from the carrier transition, assuming identical intensity profiles for the red- and blue-detuned beams. Assuming the detuning is close to one sideband frequency ν , we ignore all other modes as well as terms oscillating at $\nu + \delta$. In the Lamb–Dicke approximation, \tilde{H}_{int} is given by

$$\tilde{H}_{\text{int}} \approx \sqrt{2}\Omega\eta\tilde{J}_x [x \cos(\nu - \delta)t + p \sin(\nu - \delta)t], \quad (\text{B.1})$$

where Ω and η are the maximum Rabi rate and Lamb–Dicke factor, respectively, for any ion in the chain, x is the position operator and p the momentum operator (defined in [30]), t is time, and we have defined

$$\tilde{J}_x = \sum_{i=1}^N c_i X_i \quad \text{and} \quad c_i = \frac{1}{2} \frac{\Omega_i \eta_i}{\Omega \eta}, \quad (\text{B.2})$$

where X_i is the Pauli X operator acting on the i th ion. The Rabi rate Ω_i and Lamb–Dicke factor η_i are different for each ion i in the N -ion chain. Equation (B.1) also neglects off-resonant carrier transitions, which is a valid approximation provided $\Omega \ll \delta$ [30]. This Hamiltonian is a generalization of equation (6) in [30], and the results for the propagator U follow directly with a modified spin-dependence, \tilde{J}_x . If all ions had the same Rabi rate and Lamb–Dicke factor, then all c_i would be $\frac{1}{2}$ and \tilde{J}_x would reduce to the total angular momentum operator J_x .

The authors of [30] describe how the propagator resulting from a Hamiltonian of the form (B.1) causes eigenstates of J_x to traverse a circular path in phase space. For a propagation time $2\pi n/(\nu - \delta)$ with $n = 1, 2, 3, \dots$, the phase space trajectory returns to the origin, and the ions are left in the original motional state. At these times, the propagator $U = \exp(-iA\tilde{J}_x^2)$, where A is proportional to the geometric area enclosed by the loop in phase space. In our general case, the proportionality factor \tilde{J}_x^2 is now a sum over pairs of X operators, weighted by the relative intensity and Lamb–Dicke factor (c_i , defined in (B.2)). Simplifying

$$\tilde{J}_x^2 = \sum_{i=1}^N \sum_{j=1}^N c_i c_j X_i X_j = \sum_{i=1}^N c_i^2 + 2 \sum_{i=1}^N \sum_{j>i}^N c_i c_j X_i X_j, \quad (\text{B.3})$$

where the first term on the right-hand side is a global phase, independent of the internal states of the ions; we ignore this global phase and redefine

$$U \equiv \exp\left(-2iA \sum_{i=1}^N \sum_{j>i}^N c_i c_j X_i X_j\right). \quad (\text{B.4})$$

We define an ‘echo’ operation

$$E_k = \exp\left(-i\frac{\pi}{2} Y_k\right) = -iY_k \quad (\text{B.5})$$

which swaps x -eigenstates of the k th ion through a π rotation about the y -axis. Unitary transformations of the propagator are built from the E_k :

$$U_{k_1 k_2 \dots k_L} = \left(\prod_{l=1}^L E_{k_l}^\dagger\right) U \left(\prod_{l=1}^L E_{k_l}\right), \quad (\text{B.6})$$

$$= \left(\prod_l E_{k_l}^\dagger\right) \exp\left(-2iA \sum_{i=1}^N \sum_{j>i}^N c_i c_j X_i X_j\right) \left(\prod_l E_{k_l}\right), \quad (\text{B.7})$$

$$= \exp\left[-2iA \sum_{i=1}^N \sum_{j>i}^N c_i c_j \left(\prod_l E_{k_l}^\dagger\right) X_i X_j \left(\prod_l E_{k_l}\right)\right], \quad (\text{B.8})$$

$$= \exp\left\{-2iA \sum_{i=1}^N \sum_{j>i}^N c_i c_j \left[\prod_l (-1)^{\delta_{k_i i} + \delta_{k_j j}}\right] X_i X_j\right\}, \quad (\text{B.9})$$

where (B.8) can be shown by expanding the operator exponential with a Taylor series, and (B.9) uses $X_i Y_k = (-1)^{\delta_{ik}} Y_k X_i$.

We assert that a subspace D within an N -ion chain can be decoupled so that only the ions not in D participate in the net MS interaction through products of these $U_{k_1 \dots}$ with $k_l \in D$. For example, the following sequence will decouple the k th ion in a N -ion chain

$$U_k U = \exp\left\{-2iA \sum_{i=1}^N \sum_{j>i}^N c_i c_j [1 + (-1)^{\delta_{ki} + \delta_{kj}}] X_i X_j\right\}. \quad (\text{B.10})$$

Since $j \neq i$, k will only equal i or j or neither. Therefore, if i or j is equal to k , the term in (B.10) in square braces vanishes; conversely, the terms where $k \neq i, j$ survive, exactly as desired for an MS gate on all ions but the k th.

Generally, a sequence to decouple an d -qubit subspace D can be constructed from the product of the $U_{k_1 k_2 \dots}$ where the ion labels k_i are all combinations of 0 to d indices in D . Table B1 summarizes possible decoupling sequences up to $d = 4$. The number of such operators is given by a sum over binomial coefficients:

$$\sum_{i=0}^d \binom{d}{i} = 2^d. \quad (\text{B.11})$$

For an MS gate on m ions out of an N ion chain there are $d = N - m$ ions in D , requiring the MS gate be performed in 2^{N-m} equal steps as initially asserted. Finally, we note that since the transformed U operators all

Table B1. Sequences to decouple up to $d = 4$ ions. The transformed U operators all commute. In practice, there may be an optimal order in which to perform these operations.

d	Possible sequence
1	$U_k U$
2	$U_{kl} U_k U_l U$
3	$U_{klm} U_{kl} U_{km} U_{lm} U_k U_l U_m U$
4	$U_{klmn} U_{klm} U_{kln} U_{kmm} U_{lmn} U_{ki} U_{km} U_{kn} U_{lm} U_{ln} U_k U_l U_m U_n U$

commute, they can always be rearranged to give a decoupling sequence with a single E_k between each MS interaction. Generally, this implies the ability to optimize the echo sequence for any single-qubit control scheme.

In section 4.2, we demonstrate the most rudimentary echoing scheme. By decoupling one ion within a three-ion chain, we generate pairwise entanglement despite imperfect pairwise addressing. Within a longer chain, we note one could shorten the required echo sequence by neglecting to echo ions which have a near-zero interaction (c_i) with the gate beams.

While this echoing technique can decouple any arbitrarily connected subset of ions, we have only considered a single vibrational mode under the assumption of near-detuning. Spectator modes, which tend to lead to residual motional entanglement, may be a key error source. In the future, we anticipate using modulated sequences, like those described in [33], to alleviate this issue. Such sequences could also be interleaved with echo pulses to arbitrarily decouple selected ions with high fidelity.

References

- [1] Barenco A, Bennett C, Cleve R, DiVincenzo D, Margolus N, Shor P, Sleator T, Smolin J and Weinfurter H 1995 *Phys. Rev. A* **52** 3457
- [2] Leibfried D *et al* 2003 *Nature* **422** 412
- [3] Vandersypen L M, Steffen M, Breyta G, Yannoni C S, Sherwood M H and Chuang I L 2001 *Nature* **414** 883
- [4] Bialczak R C *et al* 2010 *Nat. Phys.* **6** 409
- [5] Chow J M *et al* 2012 *Phys. Rev. Lett.* **109** 060501
- [6] Maller K M, Lichtman M T, Xia T, Sun Y, Piotrowicz M J, Carr A W, Isenhower L and Saffman M 2015 *Phys. Rev. A* **92** 022336
- [7] Schmidt-Kaler F *et al* 2003 *Appl. Phys. B* **77** 789
- [8] Hanneke D, Home J P, Jost J D, Amini J M, Leibfried D and Wineland D J 2009 *Nat. Phys.* **6** 13
- [9] Kim J 2014 *Physics* **7** 119
- [10] Hucul D, Inlek I V, Vittorini G, Crocker C, Debnath S, Clark S M and Monroe C 2015 *Nat. Phys.* **11** 37
- [11] Kielpinski D, Monroe C and Wineland D J 2002 *Nature* **417** 709
- [12] Monroe C and Kim J 2013 *Science* **339** 1164
- [13] Seidelin S *et al* 2006 *Phys. Rev. Lett.* **96** 253003
- [14] Wang S X, Labaziewicz J, Ge Y, Shewmon R and Chuang I L 2010 *Phys. Rev. A* **81** 062332
- [15] Ospelkaus C, Warring U, Colombe Y, Brown K R, Amini J M, Leibfried D and Wineland D J 2011 *Nature* **476** 181
- [16] Warring U, Ospelkaus C, Colombe Y, Jördens R, Leibfried D and Wineland D J 2013 *Phys. Rev. Lett.* **110** 173002
- [17] Wilson A C, Colombe Y, Brown K R, Knill E, Leibfried D and Wineland D J 2014 *Nature* **512** 57
- [18] Mount E, Kabytayev C, Crain S, Harper R, Baek S Y, Vrijsen G, Flammia S, Brown K R, Maunz P and Kim J 2015 *Phys. Rev. A* **92** 060301
- [19] Guise N D *et al* 2015 *J. Appl. Phys.* **117** 174901
- [20] Olmschenk S, Younge K C, Moehring D L, Matsukevich D N, Maunz P and Monroe C 2007 *Phys. Rev. A* **76** 052314
- [21] Manning T A 2014 Quantum information processing with trapped ion chains *PhD Thesis* University of Maryland, College Park
- [22] Islam R *et al* 2014 *Opt. Lett.* **39** 3238
- [23] Hayes D, Matsukevich D N, Maunz P, Hucul D, Quraishi Q, Olmschenk S, Campbell W, Mizrahi J, Senko C and Monroe C 2010 *Phys. Rev. Lett.* **104** 140501
- [24] Wimperis S 1994 *J. Magn. Reson. A* **109** 221
- [25] Brown K R, Harrow A W and Chuang I L 2004 *Phys. Rev. A* **70** 052318
- [26] Brown K R, Harrow A W and Chuang I L 2005 *Phys. Rev. A* **72** 039905
- [27] Knill E, Leibfried D, Reichle R, Britton J, Blakestad R, Jost J, Langer C, Ozeri R, Seidelin S and Wineland D 2008 *Phys. Rev. A* **77** 012307
- [28] Brown K R, Wilson A C, Colombe Y, Ospelkaus C, Meier A M, Knill E, Leibfried D and Wineland D J 2011 *Phys. Rev. A* **84** 030303
- [29] Magesan E *et al* 2012 *Phys. Rev. Lett.* **109** 080505
- [30] Sørensen A and Mølmer K 2000 *Phys. Rev. A* **62** 022311
- [31] Sackett C *et al* 2000 *Nature* **404** 256
- [32] Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [33] Choi T, Debnath S, Manning T, Figgatt C, Gong Z X, Duan L M and Monroe C 2014 *Phys. Rev. Lett.* **112** 190502