Ion-trap quantum information processing: experimental status

D. Kielpinski

Centre for Quantum Dynamics, Griffith University, Nathan QLD 4111, Australia

Atomic ions trapped in ultra-high vacuum form an especially well-understood and useful physical system for quantum information processing. They provide excellent shielding of quantum information from environmental noise, while strong, well-controlled laser interactions readily provide quantum logic gates. A number of basic quantum information protocols have been demonstrated with trapped ions. Much current work aims at the construction of large-scale ion-trap quantum computers using complex microfabricated trap arrays. Several groups are also actively pursuing quantum interfacing of trapped ions with photons.

Contents

I. Introduction	1
II. Ion traps for QIP	3
III. Small ion-trap quantum registers	5
A. Readout and initialisation of ion qubits	5
B. Single-qubit operations	7
C. Coherent control of ion motion	8
D. Universal quantum logic and QIP protocols	10
IV. Errors and error handling in ion-trap QIP	11
A. QIP error handling	11
B. Diagnostics	12
C. Physical errors	13
V. Toward large-scale ion trap quantum computing	15
A. The QCCD architecture	15
B. Implementing the QCCD	17
C. Quantum interfacing of ions and photons	17
VI. Conclusion	18
Acknowledgments	19
References	19

I. INTRODUCTION

Quantum mechanics offers algorithms for efficient factorization of large numbers [1] and efficient searching of large databases [2], two problems that appear insoluble in classical computing. Because much modern cryptography relies on the difficulty of factoring large numbers, a large-scale quantum computer could have a large impact on many areas of technology, Internet commerce being only one example. At the same time, the delicate and demanding nature of QIP, with every quantum accounted for, requires a more subtle technology than that used to construct a classical computer. The search for physical systems supporting QIP has ranged far and wide, across optics, atomic physics, and condensed-matter physics [3, 4, 5]. Several physical implementations, namely linear optics, trapped ions, and superconducting electronic circuits, demonstrate the essential ingredients of QIP, including initialization to a known quantum state, efficient readout of the quantum state, long qubit coherence time, and universal quantum logic. In particular, small quantum computers have already been constructed with trapped ions [6], and a number of basic QIP algorithms [7], quantum memory schemes [8, 9], and communication protocols [10] have been demonstrated. In this review, we discuss the experimental status and prospects of ion-trap QIP, referring to the theory of QIP and to other physical QIP implementations only for context. This is not to underestimate the excellent work in these other

areas, but only to keep the review at a manageable length. For a general overview of QIP, the reader is encouraged to consult Nielsen and Chuang’s essential text [3], and for a recent review, [5].

Most quantum information processing devices are made up of two-level systems, “qubits,” where each qubit is analogous to a single bit in a classical computer. The quantum state of the device encodes information, and an appropriate unitary evolution of the state of the register can perform a computing task. In our case, a qubit corresponds to a trapped ion, with the two qubit states being two electronic energy levels of the ion. Laser cooling of several trapped ions causes the ions to form a Coulomb crystal, in which the ions are held in the equilibrium positions given by the combination of the trapping force and their mutual Coulomb repulsion. The techniques of optical pumping and electron shelving, well known for decades, provide efficient initialization and readout. Trapped-ion qubits can have extremely long memory times, at least 20 seconds for optical [11] and 10 minutes for microwave [12] transitions from the ground state. The qubit energy levels are often identical to those used for the long-lived transitions of atomic clocks, and techniques originally developed for atomic frequency metrology have proved invaluable in the development of ion-trap QIP.

Universal quantum computation requires both single-qubit operations and a nontrivial two-qubit logic gate. The key to two-ion logic is the collective motion of the trapped ions, which can be cooled to the quantum ground state and can store quantum information for hundreds of milliseconds [13, 14]. Laser coupling of the ion spin states to this shared quantum degree of freedom allows engineering of indirect spin-spin interactions and thus two-qubit logic gates. These techniques have enabled the demonstration of many basic QIP protocols with trapped ions [7, 10, 15, 16, 17, 18, 19].

Imperfect control of the quantum information processor leads to *decoherence*, the loss of information through correlation with an uncontrolled environment. In ion-trap QIP, decoherence arises mainly from environmental electromagnetic noise, from technical noise in the logic laser, and from spontaneous emission from excited electronics levels. A decoherence-free subspace (DFS) encoding has proved resistant to the main source of environmental noise [8]. The complementary technique of quantum error correction has recently been demonstrated in principle [9], but repeated error correction, which offers the prospect of unlimited coherence time, is still only a dream. Operation of a large-scale quantum computer will probably require both quantum error correction and DFS encoding.

No physical quantum information processor has yet demonstrated computational power anywhere near that of present-day desktop computers. A viable QIP implementation must therefore be *scalable*. In other words, the present-day techniques for QIP in a given physical system must remain useful for much larger versions of the same system without consuming a disproportionate amount of physical or computational resources. The most widely considered roadmap for large-scale ion-trap QIP builds on the successes of small ion-trap QIP devices by proposing a modular architecture of interconnected quantum registers [13, 20]. This roadmap is now being realised by investigators around the world and steps along the way have already proven crucial for demonstrations of several basic QIP protocols.

The growth of interest in QIP since the discovery of Shor’s factoring algorithm has fostered work on QIP implementations in a wide variety of atomic, optical, and condensed-matter systems (see [5] for a recent review). There is broad consensus on the requirements for QIP: efficient initialisation and detection of qubit states, long qubit coherence times, and a set of one- and two-qubit gates that enable access to the entire Hilbert space of the quantum register. Among QIP implementations, only trapped ions, superconducting circuits, and single-photon linear optics currently meet these criteria. Trapped ions make nearly ideal small quantum registers, but significant technical work is still required for large-scale quantum computing. It is easier to see how to fabricate a large number of superconducting qubits, but decoherence times are still very short. Single photons again have long coherence time, but no deterministic two-qubit gate has yet been demonstrated, so long chains of gates cannot be reliably implemented. While other physical systems may soon achieve the criteria for QIP, it is encouraging that large-scale

QIP already appears feasible (if challenging) in these three cases.¹

II. ION TRAPS FOR QIP

Almost all experiments in ion-trap QIP so far have used radio-frequency (RF) traps to confine ions under ultra-high vacuum. In these traps, one applies large RF voltages to an electrode structure made of conducting material in order to create a quadrupole electric field with a minimum in free space. For RF voltages of a few hundred volts and an appropriate frequency Ω_{RF} in the MHz range, the RF field induces a deep (few eV) ponderomotive potential that confines ions harmonically at the field minimum with an oscillator frequency $\nu \ll \Omega_{\text{RF}}$ [21]. In the QIP context, the residual ion motion at the RF drive frequency, or “micromotion,” can generally be neglected once the effects of stray electric fields have been cancelled. The low pressure ($\sim 10^{-14}$ bar) in the vacuum chamber means that collisions of residual gas atoms with the ions are infrequent, on the order of one every 100 seconds. The ions are then well isolated from environmental perturbations, enabling the precise quantum state control needed for QIP experiments.

One popular method for producing ions uses a resistively heated oven to generate an atomic beam of the species to be ionised. The atomic beam passes through the trapping region and intersects a beam of low-energy (< 100 eV) electrons emitted from a resistively heated thoriated tungsten filament. Electron-impact ionisation produces ions in the trapping region, and the strong, pseudo-conservative trapping force is sufficient for confining the hot ion sample, which can have a temperature of several thousand kelvin. Some experiments have replaced the oven-generated atomic beam by laser ablation of a solid target containing the species to be ionised [22, 23]. In many cases one can replace electron-impact ionisation by photoionisation [24], which offers the advantage of much higher ionisation efficiency and isotope selectivity. The latter property is particularly useful for species like Yb^+ [25], for which several isotopes have high natural abundance.

Linear RF traps are preferred for QIP experiments because they nearly eliminate micromotion along one motional axis. A simple electrode structure for a linear trap [26, 27] is shown in Fig. 1. Essentially the trap is a quadrupole mass filter plugged at the ends with static potentials. To operate the trap, one applies RF voltage to the rods 1 and 3 of Fig. 1, while rods 2 and 4 are held at RF ground. The induced ponderomotive potential confines the ions to the RF nodal line, which lies along the $\hat{z}$ axis. The electrodes on the trap axis are held at a positive DC voltage relative to the rod electrodes, pushing the (positive) ions toward the center of the trap. For motional amplitudes characteristic of laser-cooled ions, the resulting trap potential is harmonic in all three directions, with trap frequencies up to 50 MHz for light ions such as Be^+ [28]. Because the RF electric fields at the electrode surfaces are so high, electrical breakdown limits any attempt to radically increase the ion motional frequency [29].

In traditional ion-trap experiments, the trap electrodes are fabricated from bulk metal and the electrode assembly is clamped into a custom jig, usually of ceramic. Since the most widely accepted path to large-scale QIP with ion traps requires the fabrication of large arrays of interconnected traps [13, 20], recent experiments have moved toward microfabricated electrodes and quasi-planar (multilayer or single-layer) geometries. The first microfabricated traps used laser-machined substrates, but traps based on gallium arsenide [30] and silicon [31] have now been fabricated by MEMS techniques. The design and fabrication of trap arrays has been greatly simplified by the introduction of the “surface” trap geometry [32]. In this design all electrodes lie in a plane, and a trap array of great complexity can be fabricated on an insulating substrate by a single photolithographic step. Surface traps using a printed-circuit-board substrate [33] allow rapid prototyping of designs for trap arrays. Silicon-based surface traps offer the possibility of very large scale integration of such a trap array with CMOS control circuitry [34].

¹ A warning to novices: the intense experimental competition in this field has led some investigators to make exaggerated claims on the basis of their results, and these claims have sometimes appeared in reputable journals. I do not believe that this paper cites any work of that kind.

FIG. 1: Electrode structure of a linear RF trap. High-voltage RF applied to the rod electrodes provides ponderomotive confinement of ions to the trap axis, while the endcaps are held at DC potential for axial confinement. GND: electrical ground. The electrodes marked RF have equal RF voltages at all times. The DC electrodes are held at RF ground.

FIG. 2: A crystal of eight Yb^+ ions in the trap at Griffith University, Brisbane, Australia. The ions are cooled by a 369.5 nm diode laser [38] and are viewed by fluorescence imaging on a CCD camera.

When multiple ions are present in the trap, one must consider the Coulomb repulsion between ions as well as the ions' interaction with the harmonic trapping potential. If the ions are sufficiently cold, the classical equilibrium positions of the ions are given by minimizing the potential energy. For sufficiently weak axial confinement, the equilibrium positions all lie on the trap axis $x = y = 0$, so that the ions line up in a string, as shown in Fig. 2. Minimizing the combined trap and Coulomb potential gives the equilibrium positions along the axis [35, 36]. In the case of a two-ion string, the ion-ion distance is given by $(e^2/(2\pi\epsilon_0 m\nu_z^2))^{1/3}$, where m is the ion mass and ν_z is the axial oscillation frequency of a single ion. In current experiments the separation between ions is on the order of 3 to 30 μm , with smaller separation for lower mass and for higher trap frequency. Hence it is experimentally difficult to maintain the positioning of a focused laser beam to the tolerance required for individual laser addressing of the trapped ions. While most current ion-trap QIP experiments use readout and quantum logic schemes that are designed to avoid the requirement of individual laser addressing, the Innsbruck group has had remarkable success in ion QIP using individual addressing ([37], and see below).

After laser cooling, the residual motion of the ions is very small, so we can linearize the total potential about the equilibrium positions. The resulting harmonic oscillations constitute normal modes of the ion crystal. Solving the relevant eigenvalue equation gives the normal mode frequencies and eigenvectors (for numerical values see [36]). For any number of ions, the lowest-frequency mode is always the center-of-mass (COM) mode, in which the ion string moves as a unit, with no relative motion between the ions. The Coulomb interaction then has no effect on the dynamics of the COM mode, so the COM frequency is just equal to the single-ion trap frequency ν_z . Because of the symmetry of the ion string about $z = 0$, the ions' relative amplitudes of motion in a given mode are either symmetric or antisymmetric about the center of the string.

FIG. 3: Atomic energy levels for a trapped-ion qubit. The qubit transition between $|\downarrow\rangle$ and $|\uparrow\rangle$ is shown in green. To detect the qubit state, one uses a laser to repeatedly drive the transition shown in orange, causing fluorescence from the $|\downarrow\rangle$ state but not from the $|\uparrow\rangle$ state. The detection laser also cools the ions to the crystalline state. One reinitialises the qubit by driving a transition from $|\uparrow\rangle$ to an auxiliary energy level that decays into $|\downarrow\rangle$. The auxiliary level can be the upper state of the detection transition or some other excited state. The purple arrow schematically indicates this reinitialisation process.

III. SMALL ION-TRAP QUANTUM REGISTERS

Current experiments in trapped-ion QIP rely on laser excitation of electronic transitions for all QIP operations. Fig. 3 shows a typical set of atomic energy levels for an ion that is being used as a qubit. There are two relevant transitions out of the electronic ground state $|\downarrow\rangle$. One of these is a strong optical transition whose upper state $|e\rangle$ decays rapidly by spontaneous emission. This transition is used for initialisation and detection of the qubit. The other is a coherent optical or RF transition whose upper state $|\uparrow\rangle$ is extremely long-lived. The states $|\downarrow\rangle$, $|\uparrow\rangle$ form the logical basis, corresponding to the role of 0 and 1 in a classical computer. In analogy to the spin-1/2 system, one often refers to the logical degree of freedom as the “spin”. Coherent rotations of this spin correspond to single-qubit logic gates. Of course, all real ions have much more complex level structures than shown in Fig. 3, and so there are many variations and complications of this basic scheme depending on the exact level structure of the ion being used, but in almost all cases, the overall goal of these schemes is to reproduce the level structure of Fig. 3.

A. Readout and initialisation of ion qubits

Quantum computation requires the preparation of the computational register in a well-defined input state at the beginning of the computation and the efficient readout of the state of the register at the end of the computation. These are accomplished with the assistance of the strongly-allowed $|\downarrow\rangle \rightarrow |e\rangle$ transition of Fig. 3, referred to as the readout transition. By choosing the correct ion species with an appropriate energy level structure, one arranges that the readout laser only induces fluorescence from $|\downarrow\rangle$, while the laser is far off resonance for $|\uparrow\rangle$. Then a high fluorescence rate indicates $|\downarrow\rangle$, and a low one, $|\uparrow\rangle$. This technique for high-efficiency internal state discrimination has been used in studies of trapped ions for decades [39, 40, 41].² The fluorescence of the $|\downarrow\rangle$ state distinguishes it from the $|\uparrow\rangle$ state with $> 98\%$ efficiency per ion in each repetition of the experiment. The fluorescence has a total power of picowatts per ion and is emitted nearly isotropically, so current experiments in ion QIP use complex multi-element objective lenses to achieve large numerical aperture (NA) and high light collection efficiency. For current optical imaging systems, with f-number $\sim f/1$, the readout time can be as short as a few hundred microseconds [43, 44]. A recent experiment reports readout error rates as low as 1.5×10^{-4} [45]. Optical pumping on the same transition

² Recently another efficient detection method has been demonstrated when the energy levels are not appropriate for the electron shelving technique [42].

FIG. 4: Statistics of detected photons for a single ion prepared in a) the $|\uparrow\rangle$ state and b) the $|\downarrow\rangle$ state. Statistics were collected over 1000 repetitions of the experiment. Figure courtesy NIST Ion Storage Group.

initialises the register with a residual error that is much smaller than the detection error and requires only a few microseconds. These excellent readout and initialisation properties are key ingredients in the success of ion-trap quantum computing.

Qubit readout requires us to convert an analog quantity (number of photons measured) to a digital quantity (qubit state), so a knowledge of the photon statistics is crucial for high readout efficiency. Fig. 4 shows a histogram of the photon statistics for a single ${}^9\text{Be}^+$ ion prepared in $|\uparrow\rangle$, and another histogram for an ion prepared in $|\downarrow\rangle$. While the $|\downarrow\rangle$ histogram is Poissonian, the $|\uparrow\rangle$ distribution has a long, non-Poissonian tail. The tail arises because one generally prolongs the detection period until off-resonant repumping of the $|\uparrow\rangle$ state to $|\downarrow\rangle$ begins to cause photon scattering as well [44, 46]. The two histograms are readily distinguished, and by counting each readout of > 1 photons as $|\downarrow\rangle$ and all other readouts as $|\uparrow\rangle$, one can easily discriminate $|\downarrow\rangle$ from $|\uparrow\rangle$ in a single repetition of the experiment with 98% efficiency. Recent work has investigated the adaptive readout of an ion qubit over several repetitions of an experiment, reaching detection errors as low as 10^{-4} , with a requirement of seven repetitions (on average) per readout [47].

In experiments with one ion, one might detect the ion in $|\downarrow\rangle$ on some repetitions and in $|\uparrow\rangle$ on others, for instance by preparing the ion in a superposition state. Fig. 5 shows the photon statistics for the state $|\downarrow\rangle + |\uparrow\rangle$. In each repetition of this experiment, the ion is projected into either the state $|\downarrow\rangle$ or the state $|\uparrow\rangle$, in accordance with the quantum measurement postulate. Thus the probability distribution of photon counts is a linear combination of that found for $|\downarrow\rangle$ (Fig. 4a) and that found for $|\uparrow\rangle$ (Fig. 4b), a direct verification of “wavefunction collapse” into an eigenstate of the measurement basis.

For multiple ions, the simplest readout method is to collect the fluorescence from the entire ion string at once, making no attempt to spatially resolve the ions [6]. In this case one reads out $N_{|\downarrow\rangle}$, the number of ions in $|\downarrow\rangle$, rather than independently reading out each qubit. The photon statistics over many repetitions of the experiment is compared with theoretical distributions that depend on the mean number of photons collected per ion in $|\downarrow\rangle$, the mean number of photons due to background light, and the effect of off-resonant repumping [6, 46]. This method can resolve the relative probability of the outcomes $N_{|\downarrow\rangle}$ for four ions with 98% accuracy over 1000 repetitions and has been used successfully in six-ion experiments [48].

Even if the ions are not spatially resolved during detection, individual laser addressing of the ions can provide individual readout, though at the cost of considerable technical difficulty. Only the Innsbruck experiment routinely achieves individual addressing [37], and only at the laser wavelength used for quantum logic operations, rather than that used for detection. In this technique [16], one performs a SWAP operation between $|\downarrow\rangle$ and an auxiliary Zeeman sublevel of the metastable state $|\uparrow\rangle$ on all the ions except the one to be read out. Since all Zeeman sublevels of $|\uparrow\rangle$ are equally dark to the detection laser, fluorescence will be detected from the ion string only if the ion to be read out is in the state $|\downarrow\rangle$. Successive SWAP operations and detection periods then read out the ions individually.

FIG. 5: Statistics of detected photons for a single ion prepared in the superposition state $|\downarrow\rangle + |\uparrow\rangle$. Statistics were collected over 1000 repetitions of the experiment. The solid line is the best fit to the theoretical count distribution; the integral under the peak near zero photons is approximately equal to that under the peak near 15 photons, indicating nearly equal populations in the two states. Figure courtesy NIST Ion Storage Group.

B. Single-qubit operations

Universal quantum logic requires a set of single-qubit operations that are adequately realized by coherent coupling of the ion internal states [3]. A source of coherent radiation can drive the appropriate coupling if its frequency matches the energy splitting of the qubit states. Qubits using the hyperfine levels of the ground state (“hyperfine qubits”) are relatively easy to manipulate using commercial sources of microwave radiation, which easily achieve frequency stability on the 10 Hz scale. On the other hand, logic on qubits that use long-lived optical transitions (“optical qubits”) requires an ultrastable laser (less than 1 kHz linewidth) that can be tuned to the qubit transition wavelength [49]. Such lasers present considerable technical challenges and are rarely found outside laboratories devoted to optical frequency metrology. However, the underlying physics of the atomic interaction with the radiation field is the same for both hyperfine and optical qubits. The same interaction again underlies nuclear magnetic resonance, where the relevant basis states are often the spin states of a proton in a magnetic field, thus the term “spin” for the qubit state.

In manipulating an optical qubit, one attempts to achieve the ideal case of highly monochromatic radiation in resonance with an ion transition of very narrow linewidth. In this regime the radiative interaction drives Rabi flopping between the ion internal states (spin states) [13, 50] according to the Hamiltonian³

$$\mathbf{H}_{\text{Rabi}} = 2\hbar\Omega \left[e^{i(-\delta t + \vec{k} \cdot \vec{x} + \phi_L)} \sigma_+ + \text{h.c.} \right] \quad (1)$$

where σ_+ is a Pauli operator on the qubit state and the applied radiation has frequency $\omega_{\downarrow\uparrow} + \delta$, wavevector $\vec{k}$, Rabi frequency Ω , and phase ϕ_L at the position of the ion, $\vec{x}$. For an ion initially prepared in $|\downarrow\rangle$, the final state $\Psi(t)$ after an interaction time t with $\delta = 0$ is given by $\Psi(t) = \cos(\Omega t)|\downarrow\rangle - ie^{i\phi_L} \sin(\Omega t)|\uparrow\rangle$. By varying the duration of the applied radiation, one observes a sinusoidal oscillation of the ion fluorescence with period π/Ω . A radiation pulse that induces a total Rabi angle $\Omega t = \pi$ converts $|\downarrow\rangle$ to $|\uparrow\rangle$ and vice versa. Such a pulse is called a π pulse, and one likewise speaks of $\pi/2$ pulses and so forth. Experiments with optical qubits typically require laser intensities on the

³ We use the conventions of [13] in defining Ω and ϕ_L .

order of 100 W cm^{-2} to reach Rabi frequencies of hundreds of kHz, so the laser only needs to have a power of a few mW.

As will be seen below, it is convenient to use lasers to manipulate hyperfine as well as optical qubits. Usually one drives two-photon stimulated Raman transitions between hyperfine levels, with the Raman laser beams far detuned from any allowed transition so that no spontaneous photon emission can occur. On two-photon resonance, adiabatic elimination of the excited state gives an effective Hamiltonian that has the same form as the Rabi Hamiltonian, Eq. (1) [13, 29]. To simplify the expression for the Raman Rabi frequency Ω_R , we consider the case where only one atomic excited state contributes to the two-photon transition amplitude. The frequency difference between the Raman beams is much less than the detuning Δ of the Raman laser beams from the excited state. Writing the resonant Rabi frequencies of the two Raman laser beams as Ω_1 and Ω_2 , we find $\Omega_R \equiv 2\Omega_1\Omega_2/\Delta$. By conservation of momentum, we find that $\vec{k}$ in Eq. (1) is replaced by the difference $\vec{k}_1 - \vec{k}_2$ between the Raman beam wavevectors. Likewise ϕ_L is replaced by the difference $\varphi_1 - \varphi_2$ between the phases φ_1, φ_2 of the Raman beam electric fields, evaluated at the position of the ion. The ideal Rabi flopping behavior of optical and hyperfine qubits is identical, but in practice laser frequency noise affects the single-photon coupling of optical qubits much more than it affects stimulated Raman coupling of hyperfine qubits, because any common-mode frequency noise of the Raman laser beams cancels out. At a typical excited state detuning of $\approx 100 \text{ GHz}$, a laser intensity of about 100 W cm^{-2} is again sufficient for Rabi frequencies of hundreds of kHz.

C. Coherent control of ion motion

So far we have described initialization of our quantum register, operations on single qubits, and detection of the register state. However, one more ingredient is needed to perform universal quantum logic: a gate that entangles two particles [51, 52, 53]. In all deterministic two-ion gates that have been experimentally demonstrated, one indirectly creates a qubit-qubit interaction by transferring quantum information through a motional mode of the ion crystal. After Doppler cooling, the residual thermal excitation of the ions' quantized motion still has a mean occupation number $\bar{n} \approx 5 - 100$ for typical axial trap frequencies of 0.1 - 1 MHz. High-fidelity quantum logic operations require further cooling of the ion motion to nearly $\bar{n} = 0$, which is generally accomplished by resolved-sideband cooling [54, 55, 56], although other methods have been demonstrated [57].

Resolved-sideband cooling operates on the qubit transition of Fig. 3, rather than the detection transition. Because the qubit transition is so long-lived, all the motional modes of the ion crystal have frequencies much larger than the transition linewidth. In this case, the ion motion along the laser beam gives rise to sideband transitions offset from the $|\downarrow\rangle - |\uparrow\rangle$ transition frequency $\omega_{\downarrow\uparrow}$ by multiples of the trap frequencies, as can be seen from Eq. (1). The number of observed sidebands increases with ion temperature, as expected classically. In the following, we need only consider the first-order “red” and “blue” sidebands of each motional mode, which occur at frequencies $\omega_{\downarrow\uparrow} - \nu$ and $\omega_{\downarrow\uparrow} + \nu$ for a motional mode of frequency ν .

Red sideband excitation, combined with the initialisation process (see Fig. 3), provides a convenient mechanism for cooling to the ground state. When the ion crystal makes a red sideband transition, it loses one phonon in the corresponding motional mode. Initialisation heats the crystal by an amount corresponding to the recoil energy of a single photon, but the recoil energy is much smaller than the energy splitting of motional levels, so the motional mode readily reaches the quantum ground state. Single ions have been cooled to the ground state of motion along the trap axis with residual thermal excitation $\langle n \rangle$ of 0.001 quanta [49]; in other words, the probability of finding the ion outside the motional ground state was 0.1%. Three-ion crystals have been cooled to the axial ground state with 99% efficiency [10] by successive cooling on the three motional modes.

To describe coherent coupling between an ion spin and a motional mode cooled to the ground state, we return to the Rabi Hamiltonian Eq. (1), now considering $\vec{x}$ as the quantum position operator of the ion. If the laser wavevector is parallel to the trap axis, only the sidebands for motion along $\hat{z}$ are observed; many experiments use this geometry and we assume it for simplicity. It is easy to quantize the normal modes of the ion motion, since each mode is just

FIG. 6: Experimentally measured Rabi oscillations on the blue sideband as a function of laser pulse duration, for the initial state $|\downarrow\rangle|n=0\rangle$. After [58].

a simple harmonic oscillator. Writing the operator for small displacements of the i th ion as $\mathbf{x}_i$ and the conjugate momentum as $\mathbf{p}_i$, we define the annihilation operator $\mathbf{a}_k$ for the k th mode of N ions in the usual way:

$$\mathbf{a}_k = \sqrt{\frac{2\pi Nm\nu_k}{2}} \sum_{i=1}^N v_k^{(i)} \left(\mathbf{x}_i + \frac{i}{2\pi Nm\nu_k} \mathbf{p}_i \right) \quad (2)$$

where ν_k is the mode frequency, $\vec{v}_k$ is the normalized eigenvector of ion motional amplitudes for the mode, and m is the ion mass. The dynamics is readily expressed in terms of the Lamb-Dicke parameter $\eta \equiv k_z z_0$, where $z_0 = (2m_{\text{ion}}\omega_z)^{-1/2}$ is the zero-point wavepacket spread of a single ion along the trap axis and $\vec{k}$ is the laser wavevector. In the multi-ion case, generalised Lamb-Dicke parameters apply to the various normal modes of the ion crystal; in particular, the center-of-mass mode for N ions always has $\eta_{\text{COM}} = \eta/\sqrt{N}$. For low-error quantum logic, one must arrange that $\eta_k \ll 1$ and $\langle n_k \rangle \eta_k \ll 1$ for all axial motional modes, where the modes are labeled by k and n_k is the number operator of the i th mode.

The logic lasers in ion-trap QIP experiments are either focused onto one ion at a time, or they illuminate all ions equally. In the former case, the Rabi Hamiltonian Eq. (1) near a spin-motion resonance can be approximated as [29]

$$\mathbf{H} = e^{i\phi_L} \Omega \sigma_{j,+} + \text{h.c.} \quad \delta = \omega_{\downarrow\uparrow} \quad \text{carrier} \quad (3)$$

$$= \eta_k e^{i\phi_L} \Omega \sigma_{j,+} \mathbf{a}_k + \text{h.c.} \quad \delta = \omega_{\downarrow\uparrow} - \nu_k \quad \text{red sideband} \quad (4)$$

$$= \eta_k e^{i\phi_L} \Omega \sigma_{j,+} \mathbf{a}_k^\dagger + \text{h.c.} \quad \delta = \omega_{\downarrow\uparrow} + \nu_k \quad \text{blue sideband} \quad (5)$$

where $\mathbf{a}_k$ is the annihilation operator of the k th normal mode and $\vec{\sigma}_j$ is the Pauli spin operator of the j th ion. For the case of equal illumination, the individual spin operator σ_j is replaced by the collective spin operator $\vec{\mathbf{J}} \equiv \sum_j e^{i\phi_{L,j}} \vec{\sigma}_j$. The collective spin is just the coherent sum of the individual ion spin operators $\vec{\sigma}_j$.

From Eq. (5), we see that coherent driving on a motional sideband induces Rabi flopping dynamics between the collective spin state and a particular motional mode. The sideband Hamiltonian is that of the Jaynes-Cummings model, familiar from quantum optics. By driving on the blue sideband of a single ion, one coherently couples the states $|\downarrow\rangle|n\rangle$ and $|\uparrow\rangle|n+1\rangle$, as demonstrated experimentally in [58] (see Fig. 6). A $\pi/2$ pulse on the blue sideband transforms $|\downarrow\rangle|0\rangle$ to the superposition state $|\downarrow\rangle|0\rangle + |\uparrow\rangle|1\rangle$. (Hereafter we omit normalisation factors for wavefunctions.) This state exhibits entanglement between spin and motion and the $\pi/2$ pulse can be considered as a quantum controlled-NOT gate between the spin qubit and a motional qubit consisting of the states $\{|n=0\rangle, |1\rangle\}$ [59]. Many interesting QIP tasks can be performed using only the spin/motion quantum states of a single ion, but these have been thoroughly described elsewhere [60]. In this review, we concentrate on the use of the spin/motion coupling as a means to create spin/spin logic gates.

In dealing with hyperfine qubits, one must remember that a Raman transition has an effective wavevector $\vec{\Delta k} = \vec{k}_1 - \vec{k}_2$ that is the difference of two optical wavevectors. The effective wavelength for a Raman transition is usually in the optical range, although the transition frequency is in the microwave range. The effective wavevector $\vec{\Delta k}$ only becomes small if the Raman laser beams are nearly copropagating. A common geometry for operations on hyperfine qubits uses Raman beams with an angular separation of 90° and with $\vec{\Delta k}$ parallel to the trap axis $\hat{z}$. Then

the Raman process only couples to motion along the $\hat{z}$ axis, leading to the simplification of the sideband spectrum noted above.

D. Universal quantum logic and QIP protocols

A wide variety of two-ion gates can be implemented using the collective motion, from a controlled-NOT [61, 62] to a controlled phase gate [63, 64], or any of a class of other two-qubit gates [65, 66, 67]. These gates are deterministic, i.e., they execute a desired unitary operation at a desired moment (apart from technical limitations). Although deterministic gates are not absolutely necessary for large-scale QIP [68], access to deterministic gates is a key advantage of ion-trap quantum computing in the near term and drastically reduces the physical resources required to construct a large ion-trap QIP device.

Here is a simple example of a quantum gate that uses the shared motion to entangle two ions, along the lines of Cirac and Zoller's original proposal [61]. This example requires the presence of an additional qubit transition from $|\downarrow\rangle$ to an auxiliary state $|e\rangle$. Suppose we use a $\pi/2$ carrier pulse on both ions to prepare the state $(|\downarrow_1\rangle + |\uparrow_1\rangle)(|\downarrow_2\rangle + |\uparrow_2\rangle)|n_k = 0\rangle$, where n_k is the number operator on the k th motional mode. The gate then proceeds as follows:

1. One performs a π pulse on the red sideband of ion 1. The red sideband has no effect on $|\downarrow_1\rangle|0\rangle$, so one obtains the state $(|\downarrow_1\rangle|0\rangle + |\downarrow_1\rangle|1\rangle)(|\downarrow_2\rangle + |\uparrow_2\rangle)$.
2. One performs a 2π pulse on the red sideband of ion 2, but on the auxiliary $|\downarrow\rangle - |e\rangle$ transition. Only the $|\downarrow_2\rangle|1\rangle$ state is affected by this pulse, and it is transformed to $-|\downarrow_2\rangle|1\rangle$ by the action of H_{Rabi} (Eq. 1). The resulting state is $|\downarrow_1\downarrow_2\rangle|0\rangle + |\downarrow_1\uparrow_2\rangle|0\rangle - |\downarrow_1\downarrow_2\rangle|1\rangle + |\downarrow_1\uparrow_2\rangle|1\rangle$.
3. One performs another π pulse on the red sideband of ion 1. Again $|\downarrow_1\rangle|0\rangle$ is unaffected, so the state becomes $|\downarrow_1\rangle(|\downarrow_2\rangle + |\uparrow_2\rangle) + |\uparrow_1\rangle(|\downarrow_2\rangle - |\uparrow_2\rangle)$. We can see that this state is entangled by reversing the $\pi/2$ pulse on the carrier, but only on ion 2; then we obtain $|\downarrow_1\uparrow_2\rangle - |\uparrow_1\downarrow_2\rangle$, the Bohm-EPR state [69, 70].

A more widely used class of two-qubit gates is based on the geometric phase (Berry's phase) associated with transporting the two-ion state around a closed path in the phase space of the motional mode used for logic [63, 71]. For instance, simultaneous application of the sideband Hamiltonians of form $\mathbf{J}_x \mathbf{a}_k$, $\mathbf{J}_x \mathbf{a}_k^\dagger$ produces a unitary evolution driven by the commutator $\mathbf{J}_x^2$. Such gates have been used in experiments to realize the entangling operators $\sigma_x \sigma_x$ [66] and $\sigma_z \sigma_z$ [64]. They do not require individual laser addressing of the ions and so they are widely used in QIP experiments. They are also relatively insensitive to the motional state, an important advantage for robust quantum computing experiments because of both imperfect initialisation and ion heating (see Sec. ??). In some cases one can generalise the two-qubit gate to provide a direct multiqubit interaction [71, 72], and up to six ions have been entangled in this way [48]. In a more algorithmic approach, entangled states of up to eight ions have been created using a sequence of controlled-NOT operations [73].

By combining vibrational multiqubit gates with the detection and single-qubit logic techniques outlined above, it is currently possible to construct an ion-trap quantum computer with up to four qubits and capable of implementing tens of logic gates in a computation. Many quantum computing algorithms have been demonstrated using trapped ions, including the Deutsch-Josza algorithm [15], the Grover search algorithm [18], and the semiclassical Fourier transform [7], which is the engine of Shor's efficient factoring algorithm [1]. Likewise, several quantum communication protocols have been implemented, among them teleportation [10, 16], dense coding [17], and entanglement purification [19].

These experiments meet the most stringent requirements for QIP with small numbers of qubits. Conceptually, the system is nearly ideal. Every qubit and all relevant motional modes are initialised with low error. Quantum gates with small errors can be performed on demand, and the long qubit coherence time allows the application of a complex sequence of gates. The final state is easy to read out with low error. The quantum states generated in these experiments clearly show entanglement, ruling out classical explanations for the behavior of the quantum register.

The QIP protocols that have been demonstrated with trapped ions by no means exhaust the theoretical proposals that can be realised in this system. Many fascinating phenomena of quantum mechanics remain to be elucidated with small ion-trap quantum registers, and alternate approaches to QIP, such as one-way quantum computing [74], can also be explored with trapped ions. Quantum simulations of physical systems, especially in quantum field theory, are an attractive near-term goal for ion-trap QIP. At the same time, it is generally admitted that a new technical approach is needed to scale up ion-trap quantum registers to hundreds of ions and thousands of logic gates. Large ion crystals exhibit a dense sideband spectrum, and unwanted sideband excitations degrade logic gate performance unacceptably for more than ~ 10 ions. Most experimental efforts toward large-scale ion QIP now use an array of interconnected ion traps, in a scheme discussed below (Sec. V A).

IV. ERRORS AND ERROR HANDLING IN ION-TRAP QIP

Large-scale QIP requires long-term storage of quantum information and low-error quantum gates. Errors in classical digital computing are relatively easy to detect and correct by simultaneously processing many copies of the same information. In principle, each logical bit in a classical CPU can be stored in the position state of a single electron, but the CPU actually uses a macroscopic number of electrons, all of which occupy the same state, to store each bit. Implicitly, classical computing relies on the availability of many copies of the same state. However, the destructive nature of quantum measurement makes it impossible to reliably copy a *single* instance of a quantum state [75], requiring a radically different approach to error handling. The sensitivity of QIP to physical errors is currently addressed by 1) purely QIP-based error-handling protocols that do not depend on the qubit implementation, 2) careful engineering of qubit encoding and gate operations to take advantage of the qubit physics, and 3) brute-force reduction of sources of technical noise.

In order to describe the effect of errors mathematically, the density matrix ρ supplants the wavefunction as the description of the quantum state of the system, so as to include both quantum coherence and classical (incoherent) randomness on an equal footing [76]. In a rigorous formalism, the Hamiltonians and unitary operators of fully coherent quantum dynamics are replaced by superoperators [77]. However, one can often model the effect of a noise source on QIP more simply by assuming that a classical control parameter, like laser power or trap frequency, is varying in a random fashion, and taking an appropriate ensemble average over the realisations of the random cases [13].

A. QIP error handling

QIP-based error handling, and especially quantum error correction (QEC) [78, 79], is widely believed to be crucial for the operation of large-scale quantum computers. Without error handling, QIP would consist simply of controlling the quantum evolution of a system so that the initial state encoded a problem, and the final state, a solution. However, complex physical systems usually display chaotic dynamics, i.e., an exponential sensitivity to errors in the initial conditions, so that the maximum length of an accurate computation scales as the log of the error rate. In both classical computing and QIP, frequent and repeated error correction allows accurate computation to proceed indefinitely. This property, called fault tolerance, requires gate and memory errors to remain below certain threshold values; otherwise the process of repeated error correction induces more errors than it repairs. For particular error-handling architectures with a high computational overhead, theoretical estimates of the threshold error can range as high as a few percent [80], but gate errors below 10^{-3} appear computationally desirable, though technically challenging. A basic QEC protocol has been demonstrated with trapped ions, the detection and correction of an ion spin-flip [9], but the gate operations involved in QEC did not reach fault-tolerance, so repeated QEC was not possible. Fault-tolerance necessarily requires an attack on the physical sources of decoherence (Sec. IV C).

Decoherence-free subspace (DFS) encoding, a useful complement to QEC, arises in the context of collective decoherence processes, i.e., those processes which have the same effect on each qubit. One of the most prominent decoherence mechanisms in ion-trap QIP is the collective dephasing caused by fluctuating magnetic fields. Since the ion string is very small compared to the spatial wavelength of these magnetic fields, the interaction Hamiltonian is just proportional to $\mathbf{J}_z$. All states of the quantum register with the same number of spins in $|\uparrow\rangle$ are degenerate with each other. Such a degenerate subspace is called a decoherence-free subspace (DFS) [81, 82, 83]. Any superposition of states in this DFS is protected from the collective dephasing. The method has been demonstrated to improve memory lifetime in ion-trap QIP under ambient conditions [8]. The $\mathbf{J}_x^2$ multiqubit gate, available in ion-trap QIP, enables universal quantum computation on qubits encoded in the dephasing DFS [20]. The DFS concept can be generalised to protection against collective amplitude noise [81, 83, 84], but this kind of noise does not figure prominently in current ion-trap QIP experiments.

B. Diagnostics

To apply any error-reduction method effectively, we must clearly understand the methods of quantum state characterisation and their application in ion-trap QIP. A simple diagnostic for single-qubit logic gates is simply to observe the Rabi oscillations of the initial state $|\downarrow\rangle$ under increasing laser pulse duration (see Fig. 6). The slow decay of the oscillation amplitude permits an estimate of the maximum feasible number of single-qubit logic gates, while the form of the decay envelope can help to discriminate between various sources of technical noise [13, 29]. To sample memory errors independently from gate errors, one implements Ramsey interferometry [85] by applying a carrier $\pi/2$ pulse, waiting a time T , and applying another carrier $\pi/2$ pulse with a variable phase. The probability of finding $|\downarrow\rangle$ oscillates sinusoidally with the phase, and the contrast of the oscillations as a function of T measures the overall error of the Ramsey sequence. Generally T is much longer than the $\pi/2$ pulse duration, so the loss of contrast with increasing T can be attributed purely to memory error.

Quantum state tomography is the general QIP method for characterisation of quantum states [3]. In this method, one measures the quantum state in several noncommuting measurement bases in order to reconstruct the density matrix. For a single qubit, one can obtain a complete reconstruction by measuring the spin state along σ_x , σ_y , and σ_z . In ion-trap experiments, one implements the measurement along σ_x by performing a $\pi/2$ pulse with $\phi = 0$ and subsequently measuring in the usual σ_z basis; a similar procedure gives the measurement along σ_y . Density matrices for experimentally produced two-qubit states can be readily obtained in this manner [86]. Quantum process tomography extends the tomographic idea to reconstruction of quantum gates. The superoperator corresponding to an experimental application of a gate is fully described by the action of the gate on the qubit basis states. For a two-ion gate, one performs quantum state tomography of the final state produced from $|\downarrow\downarrow\rangle$, then for that produced from $|\downarrow\uparrow\rangle$, and so on. A second tomographic reconstruction from the final-state density matrices gives the gate superoperator. The diagnostic use of process tomography has yielded experimental improvements for two-ion entangling gate implementations and also enables comparisons between various gate schemes [87].

Unfortunately, quantum state tomography is resource-intensive. Measuring the state populations along σ_z to 1% accuracy requires on the order of 10^4 experimental repetitions, even for quantum-limited detection. Tomography of a general N -particle state requires measurement along $\mathcal{O}(n^2)$ noncommuting measurement bases. For eight ions, this poses a significant challenge [73], as illustrated in Fig. 7. Quantum process tomography requires even more resources, scaling as $\mathcal{O}(n^4)$ for an n -qubit gate. Although theoretical efforts are underway to simplify tomography of large quantum systems, many other diagnostic methods currently supplement tomographic analysis in ion-trap QIP.

Multi-qubit states often lend themselves to a partial, but highly informative, characterisation by specially designed measurements. A classic benchmark for verifying “quantumness” of a two-qubit state is the violation of a Bell inequality [88]. Bell inequalities give upper bounds on classical two-particle correlation functions, but quantum mechanics predicts that higher correlation values should be possible, and indeed they are observed in experiments on two ions [43]. Because the assumptions required to derive a Bell inequality are very general, experimental tests of Bell inequalities show that the counterintuitive features of quantum mechanics are essential for a correct physical theory.

FIG. 7: Partial data from a tomographic analysis of eight ion qubits prepared in the W-state $|\uparrow\downarrow\ldots\downarrow\rangle + |\downarrow\uparrow\ldots\downarrow\rangle + \ldots + |\downarrow\downarrow\ldots\uparrow\rangle$ [73]. Only the absolute value of the density matrix elements is shown. Corresponding data on the phase of the density matrix elements (not shown) completes the tomography dataset. Figure courtesy of R. Blatt and H. Häffner, U. Innsbruck.

In some cases, a special measurement is scalable in that it remains useful for entangled states involving more and more qubits. For instance, the state $|\downarrow\downarrow\ldots\downarrow\rangle + |\uparrow\uparrow\ldots\uparrow\rangle$ has been recently produced for up to six ions [48]. Ideally, the density matrix has only four nonzero elements: $\rho_{\downarrow\downarrow\ldots\downarrow\downarrow\ldots\downarrow\downarrow}$, $\rho_{\uparrow\uparrow\ldots\uparrow\uparrow\ldots\uparrow\uparrow}$, $\rho_{\downarrow\downarrow\ldots\downarrow\uparrow\uparrow\ldots\uparrow\uparrow}$, and $\rho_{\uparrow\uparrow\ldots\uparrow\downarrow\downarrow\ldots\downarrow\downarrow} = \rho_{\downarrow\downarrow\ldots\downarrow\uparrow\uparrow\ldots\uparrow\uparrow}^*$. The diagonal population elements $\rho_{\downarrow\downarrow\ldots\downarrow\downarrow\ldots\downarrow\downarrow}$, $\rho_{\uparrow\uparrow\ldots\uparrow\uparrow\ldots\uparrow\uparrow}$ are readily measured by the usual detection methods. To measure the experimental value of the coherence $\rho_{\downarrow\downarrow\ldots\downarrow\uparrow\uparrow\ldots\uparrow\uparrow}$, one applies a carrier $\pi/2$ pulse simultaneously to all ions with a variable phase ϕ . One then detects the number $N_{\downarrow}$ of ions in $|\downarrow\rangle$ and computes the parity $\Pi(\phi) \equiv (-1)^{N_{\downarrow}}$. Assuming an ideal carrier pulse, one finds

$$\Pi(\phi) = 2|\rho_{\downarrow\downarrow\ldots\downarrow\uparrow\uparrow\ldots\uparrow\uparrow}| \cos N\phi \quad (6)$$

so one extracts the desired coherence element of the density matrix by measuring the parity as a function of phase. The deviation of $|\rho_{\downarrow\downarrow\ldots\downarrow\uparrow\uparrow\ldots\uparrow\uparrow}|$ from its ideal value of 0.5 provides useful bounds on a wide variety of error processes. One can also place limits on other sources of gate error by Fourier analysis of $\Pi(\phi)$.

C. Physical errors

The physical error sources during quantum logic operations are quite different from the sources of “memory” errors incurred during the times when an ion is sitting idle. In general, logic errors dominate, and their sources are quite different for optical than for hyperfine qubits, while the sources of memory errors are common to both kinds of qubit. The minimum gate error reported in the refereed literature is 3%, obtained with a geometric phase gate operating on a Raman transition between hyperfine states [64]. For a hyperfine qubit, spontaneous emission from the excited state in the Raman process dominates the gate error [64]. Spontaneous emission error can be reduced to an arbitrarily low value by increasing the excited-state detuning, at the cost of increased Raman laser power [89]. Theoretical calculations for a number of ion species show that laser intensities of perhaps 10^5 W cm^{-2} and detunings of 10 THz are required for error rates on the order of 10^{-3} [90]. For an optical qubit, the dominant source of gate error is usually laser frequency noise [62], which is purely technical in origin. Even so, an ultrastable laser linewidth on the order of 100 Hz can contribute as much as 10% error [86]. A recent preprint from the Innsbruck group gives an error of 7×10^{-3} , obtained by a geometric phase gate operating on an optical qubit [91], not impossibly far from the target error rate of 10^{-4} commonly assumed for large-scale ion-trap QIP architectures [34, 92].

Although heating of the ion motion is not a major error source during a single gate operation, recoiling the ions would require spontaneous scattering of laser light and might easily destroy the quantum coherence. In current QIP

FIG. 8: Heating rate measurements for ion traps of different sizes d , expressed in units of motional quanta per second. The measurements were taken over the last ten years by researchers around the world, using a variety of macroscopic and microfabricated traps, and follow a general d^{-4} power law. The large blue dots represent data taken with a single trap with movable electrodes, and clearly rule out the d^{-2} scaling expected for Johnson (thermal) noise [95]. The large scatter of the data is only partly understood. Figure courtesy of C. Monroe, U. of Maryland.

experiments, the thermal excitation continues to accumulate and can degrade the two-qubit logic gates toward the end of a long gate sequence, so significant efforts have been expended to understand and eliminate ion heating. Unfortunately, the exact mechanism of ion heating remains elusive, although some features of the experimental data can be understood phenomenologically. The heating almost certainly arises from fluctuating electric fields that are nearly spatially uniform over the extent of a few-ion crystal [93, 94]. However, the observed heating rates indicate fluctuating fields that are much stronger than those expected from thermal noise [14], and the inferred spectral density of electric field noise scales approximately as $1/f$, rather than the constant spectral density of thermal noise. As shown in Fig. 8, the heating rate scales with trap size d as approximately d^{-4} ; this result has been rigorously confirmed using a trap with movable electrodes [95]. These observations are consistent with the idea that the heating is driven by the movement of small patches of charge on the trap electrodes, perhaps caused by adsorbed gas [13, 14]. It also appears that contamination of the electrodes by the atomic beam used for loading can substantially increase the heating rate [14, 96, 97], perhaps accounting for some of the scatter of the data in Fig. 8. Recent results show that the ion heating is suppressed by several orders of magnitude at cryogenic temperatures [98, 99], though the thermal noise limit has not yet been observed.

Memory errors are well understood in trapped-ion QIP. The dephasing of a two-level atomic system is the fundamental limitation to the precision of atomic clocks. The efforts of the metrology community over the last few decades have already uncovered, measured, and found ways to minimise the sources of decoherence that apply to an ion-trap quantum memory, and these methods are rapidly being adapted for trapped-ion QIP. The error rates for memory implementations depend on the precise choice of ion energy levels for representing the qubit states. A fundamental limit to the memory lifetime of an optical qubit is given by the radiative decay rate, of order 1 second in current experiments, while hyperfine qubits have radiative lifetimes comparable to the age of the universe.

The dominant source of memory error in current experiments is generally dephasing by ambient magnetic field noise, which shifts the qubit energy levels through the Zeeman effect. Even when the quantum information is encoded in a DFS, the residual error caused by fluctuating magnetic field *gradients* limits memory time [11]. Magnetic field dephasing can be suppressed by several orders of magnitude by choosing a qubit transition with vanishing first-order Zeeman shift, a fact well known by the metrology community and explicitly demonstrated in the QIP context [100]. A two-ion gate compatible with such a “clock” transition has been demonstrated [101] and has been found to be resistant to laser phase fluctuations [67]. A DFS made with clock-state qubits might offer entangled states with lifetimes far in excess of the current record of 20 seconds [11] and on to the 10 minute coherence times demonstrated for ion-trap atomic clocks [12].

V. TOWARD LARGE-SCALE ION TRAP QUANTUM COMPUTING

A. The QCCD architecture

As we have seen, one can use a small number of trapped ions to construct a quantum register. However, manipulating a large number of ions in a single trap presents immense technical difficulties, and scaling arguments suggest that a single trap can only support computations on tens of ions [13, 102, 103]. To build up a large-scale quantum computer, Wineland and co-workers proposed a “quantum charge-coupled device” (QCCD) architecture consisting of a large number of interconnected ion traps [13, 20]. By adjusting the operating voltages of these traps, one can confine a few ions in each trap or shuttle ions from trap to trap.⁴ In any particular trap, one can manipulate a few ions using the methods already demonstrated, while the connections between traps allow communication between sets of ions. Since both the speed of quantum logic gates [104] and the shuttling speed are limited by the trap strength, shuttling ions between memory and interaction regions can consume an acceptably small fraction of a clock cycle. The QCCD architecture is not only a dream: the NIST group used small QCCD devices to demonstrate many of the protocols discussed in Sec. III D [7, 10, 17, 19].

Figure 9 shows a schematic of an ion-trap array used in the QCCD architecture. Trapped ions storing quantum information are held in the memory regions. To perform a logic gate, one moves the relevant ions into an interaction region by applying appropriate voltages to the electrode segments. In the interaction region, the ions are held close together, enabling the Coulomb coupling necessary for entangling gates. Lasers are focused through the interaction region to drive the gates. Detection can occur in the interaction regions or in separate parts of the trap array. The QCCD array can be considered as a network of single-trap QIP devices, so that the architecture is essentially modular. The techniques used for quantum logic are those already demonstrated for single-trap quantum registers, so many of the problems of a large-scale QCCD device, for instance the problems of fabricating large trap arrays and of addressing the many interaction regions with laser beams, are easily understood in terms of classical physics and engineering [34].

One can create the trapping and transport potentials needed for the QCCD using a combination of radiofrequency (RF) and quasistatic electric fields. Figure 9 shows a conceptual picture, including only the electrodes that support the quasistatic fields. By varying the voltages on these electrodes, we can perform ion transport and confinement along the local trap axis, which lies along the arrows in Figure 9. Two more layers of electrodes lie above and below the static electrodes. Applying RF voltage to the outer layers confines the ions transverse to the local trap axis, just as for the standard linear trap. This geometry allows stable transport of the ions around T- and X-junctions, so one can build complex, multiply connected trap structures.

The first experiments toward the QCCD scheme used a trap with six axial electrode segments, giving two trapping regions separated by 1.2 mm. Single ions were reliably transported between the trapping regions in $\sim 20 \mu\text{s}$, with little loss of qubit coherence and motional heating much less than 0.1 quantum per round trip [97]. Mastering the more complex task of splitting an ion crystal with low heating [10] enabled the NIST group to demonstrate some of the three- and four-qubit QIP protocols discussed in Sec. V A. The extension of these linear trap arrays to T- and X-junctions is challenging because residual RF fields cause unwanted trapping potentials along the ion transport direction. Transport around a T-junction has been achieved with less than 2% error, at the cost of a massive heating of the ion motion to a few thousand kelvin [105]. In principle, these residual RF fields can be eliminated by careful trap design. Recent work indicates that transport speed can be increased by an order of magnitude using careful shaping of the voltage waveforms that drive the trap electrodes [106], but it remains to be seen whether this technique leads to substantial heating from the quantum ground state.

⁴ The QCCD architecture is named in analogy to the ubiquitous charge-coupled device (CCD) camera, which uses shuttling of electrons for image readout.

FIG. 9: Schematic of the QCCD architecture for large-scale ion-trap QIP, after [20]. Ions (dots) are confined to the local trap axis by RF potentials and are shuttled between memory, interaction, and detection regions by quasistatic potentials applied to the DC electrodes (shaded boxes). For simplicity, only the DC electrodes are shown; the RF electrodes can be implemented in a variety of geometries (see text).

Even with the best trap design, it is hard to imagine that the ions will remain in the motional ground state indefinitely, especially under transport. Although laser cooling of the qubit ions cannot occur during QIP, *sympathetic cooling* of the qubit ions by another ion species has been experimentally demonstrated [107, 108]. Confining both species in the interaction region allows use of the cooling species as a heat sink, with the Coulomb interaction providing energy transfer from the qubit ions. No decoherence need occur if the laser wavelengths relevant to the cooling species are sufficiently far detuned from the transitions of the qubit species.

Additional decoherence mechanisms for the qubit states also arise from ion transport. For instance, the spatial variations of the magnetic field strength along the transport path shift the qubit energy levels through the Zeeman effect, so that, e.g., $|\downarrow\rangle + |\uparrow\rangle \rightarrow |\downarrow\rangle + e^{i\alpha}|\uparrow\rangle$ with a phase α depending on the transport path. Current experiments compensate for these phases using spin-echo refocusing, i.e., by swapping the states $|\downarrow\rangle$ and $|\uparrow\rangle$ in an appropriate way during ion transport.

FIG. 10: Integration of a trap array with a microfabricated array of phase Fresnel lenses. The lenses collimate ion fluorescence from the individual detection zones for reimaging onto detectors or optical fibers.

B. Implementing the QCCD

The QCCD architecture requires large-scale system integration of widely disparate electronic and optical techniques in an ultra-high vacuum environment. Microfabrication of trap arrays has been discussed above in Sec. II. Low-noise, rapid ion transport demands parallel delivery of precisely shaped arbitrary voltage waveforms to all trap electrodes, a need currently met by computer-controlled digital-to-analog converters [105, 106]. However, this technology is difficult to scale past a few tens of electrodes and integration of classical control circuitry with the trap array appears essential in the long run [34]. Such integration would ideally enable programming of ion transport sequences by simple digital commands from an external computer. A recently demonstrated silicon-based surface trap [109], fabricated using a CMOS-compatible process, represents a first step toward integration of classical circuitry with the QCCD.

Highly parallel quantum logic and initialisation operations require that an array of laser beams be delivered to the trap sites through switching or repositioning. Currently, laser addressing is achieved by changing the deflection angle of an acousto-optic modulator [37]. In recent work, a micromirror array fabricated by a microelectromechanical systems (MEMS) process was used to steer a near-infrared laser beam, allowing laser addressing on a 5×5 grid of spacing $8\mu\text{m}$ with a switching time of $\sim 10\mu\text{s}$ [110]. Similar MEMS micromirrors have been fabricated with high reflectivity at the UV wavelengths relevant to ion QIP [111]. These devices operate outside the UHV chamber, allowing wide latitude in the choice of fabrication technology.

Methods for highly parallel detection of ion fluorescence have not been explored much. The objective lenses used in current experiments have apertures $\sim 50\text{ mm}$ and a useful field of view of $< 1\text{ mm}$, so only one or two detection regions can be currently realised. The semiconductor industry uses more complex multi-element lenses with field of view $> 10\text{ mm}$ for lithographic exposures, but these lenses are the result of a massive engineering effort, concentrated at a few specific wavelengths, that would be difficult to replicate in the QIP community. Alternatively, one can realise parallel detection using an array of light collection regions, one for each detection site. The author and co-workers have proposed using a microfabricated array of Fresnel lenses for this purpose [112], as shown in Fig. 10. In [112] a single phase Fresnel lens with $f = 3\text{ mm}$ was designed for the 369.5 nm Yb^+ transition and was fabricated by lithography and etching of fused silica, in a scalable and UHV-compatible process. From optical tests, we predict a fluorescence collection efficiency of 4% for this lens, several times higher than current multi-element lens systems [113] and nearly sufficient for fault-tolerant QIP [92].

C. Quantum interfacing of ions and photons

Since trapped ions are essentially stationary, quantum communication between ion-trap QIP devices will require the coherent transfer of quantum information between ion qubits and an altogether different physical qubit implementation that is better suited to long-distance travel. Single photons are a natural candidate for this quantum interface, in view of the well-developed QIP techniques in linear optics, and since ion QIP is already heavily reliant

on optical technology. Optical quantum communication among ion-trap QIP devices may prove a useful route to large-scale quantum computing [114]. A near-term application for ion-trap QIP is the construction of a quantum repeater for optical quantum cryptography over long distances [115].

So far, the most complex ion-photon networking tasks have been performed by probabilistic detection of ion fluorescence. This method relies on the entanglement between the final state of a spontaneously emitting ion and the polarisation of the emitted photon, first demonstrated in [116]. If the single photons emitted in fluorescence from two widely separated ions are made to interfere at a beamsplitter, a subsequent coincident detection of the photons projects the ions into an entangled state. Remote entanglement of two ions in independent traps separated by a meter has been achieved [113] and shown to violate a Bell inequality [117].

Figure 11 illustrates the protocol for ion-photon entanglement used in [116]. A single ion with $^2S_{1/2}$ ground state and nuclear spin $1/2$ is initially prepared in the $|F, m_f\rangle = |1, 0\rangle$ hyperfine level. Laser excitation to the $^2P_{3/2}|F' = 2, m_{F'} = 1\rangle$ state is followed by spontaneous emission of a single photon. If the photon has π polarisation, the final atomic state must be $|F, m_F\rangle = |1, 1\rangle$, while for σ^+ polarisation the final state must be $|1, 0\rangle$. In the plane perpendicular to the quantisation axis, an emitted π (σ^+) photon is vertically (horizontally) polarised. The statistics of ion and photon measurement results under single-qubit rotations reveal the nonclassical correlations. This protocol is readily extended to remote ion-ion entanglement through the nonclassical interference of emitted photons at a beamsplitter [113].

Several experiments have taken steps toward deterministic coupling between ion and photon quantum states by using high-finesse resonators to collect a large fraction of the photons spontaneously emitted by a single ion. The subwavelength confinement of the ion allows fine control over the ion-photon coupling [118] and selective excitation of the vibrational sidebands that the ion imposes on the cavity mode [119]. Ion-cavity systems have been used to demonstrate stable probabilistic single-photon generation over tens of minutes [120], illustrating their suitability as robust building blocks for QIP. However, in current experiments the single-ion cooperativity parameter $C = g^2/(2\kappa\gamma)$ is on the order of one, so the coherence of the ion-photon coupling is low. Here g is the Rabi frequency of the ion when a single photon is present in the cavity, κ is the cavity decay rate, and γ is the decay rate of the ion transition. Proximity effects pose a major obstacle to increasing C . The short ($\lesssim 100\mu\text{m}$) resonators used for neutral-atom cavity experiments achieve high C largely because of their small mode volume. Since the single-photon energy remains constant, the electric field of a single photon grows, and so does g . This path is not available in ion-trap experiments, making high C solely reliant on high mirror reflectivity.

VI. CONCLUSION

All physical implementations of QIP are currently in their infancy. None has yet clearly demonstrated the capacity for fault-tolerant quantum computing or the possibility of outcompeting classical computers in any way. Nevertheless, ion-trap experiments fulfill the requirements for effective small-scale QIP. Ion qubits can be initialised and read out on demand, and the deterministic quantum logic gates afforded by the motional modes enable universal quantum logic. The low error in these basic QIP operations has allowed successful tests of many simple QIP protocols.

Trapped ions appear to be an attractive system for large-scale quantum computing under the QCCD architecture. Experimental progress toward the QCCD has already advanced the state of the art in small-scale QIP. The ion-trap community is successfully addressing the technical problems of scaling, namely trap array fabrication, fast laser beam steering, and parallel detection. Demonstrations of quantum interfacing of ions with photons promise optical quantum communication between QCCD devices. With these advantages, trapped ions will likely occupy a prominent role in QIP for many years to come.

FIG. 11: Protocol for generating probabilistic ion-photon entanglement used in [116]. Laser excitation is followed by spontaneous emission of a single photon whose polarisation is correlated with the final atomic state.

Acknowledgments

I thank Rainer Blatt, Wolfgang Lange, Chris Monroe, and David Wineland for helpful discussions. This work was supported under Australian Research Council grants DP0773354 (Kielpinski) and FF0458313 (Wiseman) and by the US Air Force under grant FA4869-08-1-4005.

-
- [1] P. W. Shor, *Proc. 35th Ann. Symp. Found. Comp. Sci.* (IEEE Computer Society, Los Alamitos, CA, 1994), p. 116.
 - [2] L. K. Grover, *Phys. Rev. Lett.* **79**, 325 (1997).
 - [3] M. A. Nielsen and I. A. Chuang, *Quantum Information and Quantum Computation* (Cambridge Univ. Press, Cambridge, UK, 2000).
 - [4] R. Clark, ed., *Experimental Implementation of Quantum Computation '01* (Rinton, Princeton, NJ, 2001).
 - [5] T. P. Spiller, W. J. Munro, S. D. Barrett, and P. Kok, *Contemp. Phys.* **46**, 407 (2005).
 - [6] D. Kielpinski, *J. Opt. B* **5**, R121 (2003).
 - [7] J. Chiaverini, J. Britton, D. Leibfried, E. Knill, M. D. Barrett, R. B. Blakestad, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, et al., *Science* **308**, 997 (2005).
 - [8] D. Kielpinski, V. Meyer, M. A. Rowe, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland, *Science* **291**, 1013 (2001).
 - [9] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, et al., *Nature* **432**, 602 (2004).
 - [10] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri,

- et al., *Nature* **429**, 737 (2004).
- [11] H. Häffner, F. Schmidt-Kaler, W. Hänsel, C. F. Roos, T. Körber, M. Chwalla, M. Riebe, J. Benhelm, U. D. Rapol, C. Becher, et al., *Appl. Phys. B* **81**, 151 (2005).
 - [12] J. J. Bollinger, D. J. Heinzen, W. M. Itano, S. L. Gilbert, and D. J. Wineland, *IEEE Trans. Instr. Meas.* **40**, 126 (1991).
 - [13] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, *J. Res. NIST* **103**, 259 (1998).
 - [14] Q. A. Turchette et al., *Phys. Rev. A* **61**, 063418 (2000).
 - [15] S. Gulde, M. Riebe, G. P. T. Lancaster, C. Becher, J. Eschner, H. Häffner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt, *Nature* **421**, 48 (2003).
 - [16] M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, et al., *Nature* **429**, 734 (2004).
 - [17] T. Schaetz, M. D. Barrett, D. Leibfried, J. Chiaverini, J. Britton, W. Itano, J. D. Jost, C. Langer, and D. J. Wineland, *Phys. Rev. Lett.* **93**, 040505 (2004).
 - [18] K.-A. Brickman, P. C. Haljan, P. J. Lee, M. Acton, L. Deslauriers, and C. Monroe, *Phys. Rev. A* **72**, 050306 (2005).
 - [19] R. Reichle, D. Leibfried, E. Knill, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, *Nature* **443**, 838 (2006).
 - [20] D. Kielpinski, C. Monroe, and D. J. Wineland, *Nature* **417**, 709 (2002).
 - [21] P. K. Ghosh, *Ion Traps* (Clarendon, Oxford, 1995).
 - [22] R. J. Hendricks, D. M. Grant, P. F. Herskind, A. Dantan, and M. Drewsen, *Appl. Phys. B* **88**, 507 (2007).
 - [23] D. R. Leibbrandt, R. J. Clark, J. Labaziewicz, P. Antohi, W. Bakr, K. R. Brown, and I. L. Chuang, *Phys. Rev. A* **76**, 055403 (2007).
 - [24] N. Kjærgaard, L. Hornekær, A. Thommesen, Z. Videsen, and M. Drewsen, *Appl. Phys. B* **71**, 207 (2000).
 - [25] C. Balzer, A. Braun, T. Hannemann, C. Paape, M. Ettler, W. Neuhauser, and C. Wunderlich, *Phys. Rev. A* **73**, 041407 (2006).
 - [26] J. D. Prestage, G. J. Dick, and L. Maleki, *J. Appl. Phys.* **66**, 1013 (1989).
 - [27] M. G. Raizen, J. M. Gilligan, J. C. Bergquist, W. M. Itano, and D. J. Wineland, *J. Mod. Opt.* **39**, 233 (1992).
 - [28] S. R. Jefferts, C. Monroe, E. W. Bell, and D. J. Wineland, *Phys. Rev. A* **51**, 3112 (1995).
 - [29] D. Kielpinski, Ph.D. thesis, University of Colorado (2001), available at <http://www.boulder.nist.gov/timefreq/ion/qucomp/papers.htm>.
 - [30] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab, and C. Monroe, *Nature Phys.* **2**, 36 (2006).
 - [31] D. J. Wineland et al., *Proc. 17th Intl. Conf. on Laser Spectroscopy* (Academic Press, 2006).
 - [32] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland, *Quant. Info. Comp.* **5**, 419 (2005).
 - [33] K. R. Brown, R. J. Clark, J. Labaziewicz, P. Richerme, D. R. Leibbrandt, and I. L. Chuang, *Phys. Rev. A* **75**, 015401 (2007).
 - [34] J. Kim et al., *Quant. Info. Comp.* **5**, 515 (2005).
 - [35] A. Steane, *Appl. Phys. B* **64**, 623 (1997).
 - [36] D. F. V. James, *Appl. Phys. B* **66**, 191 (1998).
 - [37] H. C. Nägerl, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler, and R. Blatt, *Phys. Rev. A* **60**, 145 (1999).
 - [38] D. Kielpinski, M. Cetina, J. A. Cox, and F. X. Kärtner, *Opt. Lett.* **31**, 757 (2006).
 - [39] H. G. Dehmelt, *Bull. Am. Phys. Soc* **20**, 60 (1975).
 - [40] D. J. Wineland, J. C. Bergquist, W. M. Itano, and R. E. Drullinger, *Opt. Lett.* **5**, 245 (1980).
 - [41] T. Sauter, W. Neuhauser, R. Blatt, and P. E. Toschek, *Phys. Rev. Lett.* **57**, 1696 (1986).
 - [42] M. J. McDonnell, J.-P. Stacey, S. C. Webster, J. P. Home, A. Ramos, D. Lucas, D. N. Stacey, and A. Steane, *Phys. Rev. Lett.* **93**, 153601 (2004).
 - [43] M. A. Rowe, D. Kielpinski, V. Meyer, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland, *Nature* **409**, 791 (2001).
 - [44] M. Acton, K.-A. Brickman, P. C. Haljan, P. J. Lee, L. Deslauriers, and C. Monroe, *Quant. Info. Comp.* **6**, 465 (2006).
 - [45] A. Myerson, D. Szwed, S. Webster, D. Allcock, M. Curtis, G. Imreh, J. Sherman, D. Stacey, A. Steane, and D. Lucas (2007), [arXiv:0802.1684](https://arxiv.org/abs/0802.1684).
 - [46] B. E. King, Ph.D. thesis, University of Colorado (1999), available at <http://jilawww.colorado.edu/www/sro/thesis/king/index.html>.
 - [47] D. B. Hume, T. Rosenband, and D. J. Wineland, *Phys. Rev. Lett.* **99**, 120502 (2007).
 - [48] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, et al., *Nature* **438**, 639 (2005).
 - [49] C. Roos, T. Zeiger, H. Rohde, H. C. Nägerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt, *Phys. Rev. Lett.* **83**, 4713 (1999).
 - [50] L. Allen and J. H. Eberly, *Optical Resonance and Two-Level Atoms* (Dover, New York, 1975).
 - [51] D. Deutsch, A. Barenco, and A. Ekert, *Proc. Roy. Soc. London A* **449**, 669 (1995).
 - [52] S. Lloyd, *Phys. Rev. Lett.* **75**, 346 (1995).
 - [53] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, *Phys. Rev. A* **52**, 3457 (1995).
 - [54] D. J. Wineland and W. M. Itano, *Phys. Rev. A* **20**, 1521 (1979).
 - [55] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, *Phys. Rev. Lett.* **62**, 403 (1989).

- [56] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, Phys. Rev. Lett. **75**, 4011 (1995).
- [57] C. F. Roos, D. Leibfried, A. Mundt, F. Schmidt-Kaler, J. Eschner, and R. Blatt, Phys. Rev. Lett. **85**, 5547 (2000).
- [58] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland, Phys. Rev. Lett. **76**, 1796 (1996).
- [59] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland, Phys. Rev. Lett. **75**, 4714 (1995).
- [60] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, Rev. Mod. Phys. **75**, 281 (2003).
- [61] J. I. Cirac and P. Zoller, Phys. Rev. Lett. **74**, 4091 (1995).
- [62] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, Nature **422**, 408 (2003).
- [63] G. J. Milburn, S. Schneider, and D. F. V. James, Fortschr. Phys. **48**, 801 (2000).
- [64] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, et al., Nature **422**, 412 (2003).
- [65] A. Sørensen and K. Mølmer, Phys. Rev. Lett. **82**, 1971 (1999).
- [66] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, et al., Nature **404**, 256 (2000).
- [67] P. C. Haljan, K.-A. Brickman, L. Deslauriers, P. J. Lee, and C. Monroe, Phys. Rev. Lett. **94**, 153602 (2005).
- [68] E. Knill, R. Laflamme, and G. J. Milburn, Nature **409**, 46 (2001).
- [69] A. Einstein, B. Podolsky, and N. Rosen, Phys. Rev. **47**, 777 (1935).
- [70] D. Bohm, *Quantum Theory* (Dover, 1989).
- [71] X. Wang, A. Sørensen, and K. Mølmer, Phys. Rev. Lett. **86**, 3907 (2001).
- [72] K. Mølmer and A. Sørensen, Phys. Rev. Lett. **82**, 1835 (1999).
- [73] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, et al., Nature **438**, 643 (2005).
- [74] R. Raussendorf and H.-J. Briegel, Phys. Rev. Lett. **86**, 5188 (2001).
- [75] W. K. Wootters and W. H. Zurek, Nature **299**, 802 (1982).
- [76] J. J. Sakurai, *Modern Quantum Mechanics* (Addison-Wesley, Reading, MA, 1994), revised ed.
- [77] D. F. Walls and G. J. Milburn, *Quantum Optics* (Springer, New York, 2008), 3rd ed.
- [78] P. W. Shor, Phys. Rev. A **52**, R2493 (1995).
- [79] A. M. Steane, Phys. Rev. Lett. **77**, 793 (1996).
- [80] E. Knill, Nature **434**, 39 (2005).
- [81] P. Zanardi and M. Rasetti, Phys. Rev. Lett. **79**, 3306 (1997).
- [82] L.-M. Duan and G.-C. Guo, Phys. Rev. Lett. **79**, 1953 (1997).
- [83] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Phys. Rev. Lett. **81**, 2594 (1998).
- [84] L. M. Duan and G. C. Guo, Phys. Rev. A **57**, 737 (1998).
- [85] N. F. Ramsey, *Molecular Beams* (Oxford Univ. Press, London, 1963).
- [86] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Häffner, W. Hänsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt, Phys. Rev. Lett. **92**, 220402 (2004).
- [87] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt, Phys. Rev. Lett. **97**, 220407 (2006).
- [88] J. S. Bell, Physics (N.Y.) **1**, 195 (1965).
- [89] R. Ozeri et al., Phys. Rev. Lett. **95**, 030403 (2005).
- [90] R. Ozeri et al., Phys. Rev. A **75**, 042329 (2007).
- [91] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt (2008), arXiv:0803.2798.
- [92] A. M. Steane, Quant. Info. Comp. **7**, 171 (2007).
- [93] D. F. V. James, Phys. Rev. Lett. **81**, 317 (1998).
- [94] B. E. King et al., Phys. Rev. Lett. **81**, 1525 (1998).
- [95] L. Deslauriers, P. C. Haljan, P. J. Lee, K.-A. Brickman, B. B. Blinov, M. J. Madsen, and C. Monroe, Phys. Rev. A **70**, 043408 (2004).
- [96] R. G. DeVoe and C. Kurtsiefer, Phys. Rev. A **65**, 063407 (2002).
- [97] M. A. Rowe, A. Ben-Kish, B. DeMarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenković, et al., Quant. Info. Comp. **2**, 257 (2002).
- [98] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe, Phys. Rev. Lett. **97**, 103007 (2006).
- [99] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang, Phys. Rev. Lett. **100**, 013001 (2008).
- [100] C. Langer et al., Phys. Rev. Lett. **95**, 060502 (2005).
- [101] P. C. Haljan, P. J. Lee, K.-A. Brickmann, M. Acton, L. Deslauriers, and C. Monroe, Phys. Rev. A **72**, 062316 (2005).
- [102] R. J. Hughes et al., Phys. Rev. Lett. **77**, 3240 (1996).
- [103] D. G. Enzer, *Experimental Implementation of Quantum Computation '01* (Rinton, Princeton, NJ, 2001), p. 99.
- [104] A. Steane, C. F. Roos, D. Stevens, A. Mundt, D. Leibfried, F. Schmidt-Kaler, and R. Blatt, Phys. Rev. A **62**, 042305 (2000).
- [105] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk, Appl. Phys. Lett. **88**, 034101 (2006).
- [106] G. Huber, T. Deuschle, W. Schnitzler, R. Reichle, K. Singer, and F. Schmidt-Kaler, New J. Phys. **10**, 013004 (2008).
- [107] B. B. Blinov, L. Deslauriers, P. Lee, M. J. Madsen, R. Miller, and C. Monroe, Phys. Rev. A **65**, 040304 (2002).
- [108] M. D. Barrett, B. DeMarco, T. Schaetz, V. Meyer, D. Leibfried, J. Britton, J. Chiaverini, W. M. Itano, B. Jelenković,

- J. D. Jost, et al., Phys. Rev. A **68**, 042302 (2003).
- [109] J. Britton et al. (2006), quant-ph/0605170.
 - [110] C. Knoernschild, C. Kim, B. Liu, F. P. Lu, and J. Kim, Opt. Lett. **33**, 273 (2008).
 - [111] C. Kim, C. Knoernschild, B. Liu, and J. Kim, IEEE J. Sel. Top. Quant. Elec. **13**, 322 (2007).
 - [112] E. W. Streed, B. G. Norton, J. J. Chapman, and D. Kielpinski, arXiv:0805.2437.
 - [113] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Nature **449**, 68 (2007).
 - [114] L.-M. Duan, B. B. Blinov, D. L. Moehring, and C. Monroe, Quant. Info. Comp. **4**, 165 (2004).
 - [115] D. L. Moehring, M. J. Madsen, K. C. Younge, J. R. N. Kohn, P. Maunz, L.-M. Duan, C. Monroe, and B. B. Blinov, J. Opt. Soc. Am. B **24**, 300 (2007).
 - [116] B. B. Blinov, D. L. Moehring, L.-M. Duan, and C. Monroe, Nature **428**, 153 (2004).
 - [117] D. N. Matsukevich, P. Maunz, D. L. Moehring, S. Olmschenk, and C. Monroe, Phys. Rev. Lett. **100**, 150404 (2008).
 - [118] G. R. Guthöhrlein, M. Keller, K. Hayasaka, W. Lange, and H. Walther, Nature **414**, 49 (2001).
 - [119] A. B. Mundt, A. Kreuter, C. Becher, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt, Phys. Rev. Lett. **89**, 103001 (2002).
 - [120] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther, Nature **431**, 1075 (2004).