Appl. Phys. B 64, 623–642 (1997)

Applied Physics B Lasers and Optics

Invited paper

The ion trap quantum information processor

A. Steane

Department of Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford, OX1 3PU, England
(E-mail: a.steane@physics.ox.ac.uk)

Received: 16 October 1996/Revised version: 2 April 1997

Abstract. An introductory review of the linear ion trap is given, with particular regard to its use for quantum information processing. The discussion aims to bring together ideas from information theory and experimental ion trapping, to provide a resource to workers unfamiliar with one or the other of these subjects. It is shown that information theory provides valuable concepts for the experimental use of ion traps, especially error correction, and conversely the ion trap provides a valuable link between information theory and physics, with attendant physical insights. Example parameters are given for the case of calcium ions. Passive stabilisation will allow about 200 computing operations on 10 ions; with error correction this can be greatly extended.

PACS:42.50.Vk; 89.70.+c; 32.80.Pj

This paper is a contribution to the rapidly developing field of quantum information theory and experiment. Quantum information is an interdisciplinary subject, in which computer scientists and other experts in the theory of classical information and computing are not necessarily familiar with quantum mechanics, and physicists and other experts in quantum theory are not necessarily familiar with information theory. Furthermore, whereas the field has enjoyed a rich theoretical treatment, there is a lack of an experimental basis to underpin the ideas. This is especially significant to the issue of error correction, or more generally any stabilisation of a quantum computer, which is among the most important unresolved issues in this field. The aim of this paper is to offer an aid to people from different sides of the subject to understand issues in the other. That is to say, the ideas of quantum information and computing will be introduced to experimental physicists, and a particular physical system that might implement quantum computing will be described in detail for the benefit of theoreticians. I hope to give sufficient information to form more or less a 'blueprint' for the type of quantum information processor currently achievable in the lab, highlighting the various experimental problems involved. The discussion is like a review in that it brings together the work of other authors rather than provides much original material. However, an exhaustive review of the wide range of subjects involved is not intended, and as a result it will not be possible to do justice to the efforts of the many people who brought the experimental and theoretical programmes to their present state of accomplishment.

The plan of the article is as follows. In Sect. 1 the concepts of quantum information processing are introduced. In Sect. 2 the general requirements for realising an experimental processor, making use of a 'universal' set of quantum logic gates, are described. In Sect. 3 the linear ion trap is considered as a system in which these ideas can be applied. A physical process by which quantum logic gates may be applied in an ion trap is described in detail. Limitations on the size of the processor (number of quantum bits) and speed of operation ('switching rate') are discussed. In Sects. 4 and 5 the main experimental techniques required to realise the ion trap processor in the lab are discussed; these are laser cooling of the ions, and low-noise generation of the correct direct current (dc) and radio frequency (rf) voltages for the trap electrodes, as well as a good choice of electrode design. In Sect. 6 we begin to establish definite values for the experimental parameters, by considering specific candidate ions to which the methods can be applied. Example values are given for the singly-charged calcium ion. In Sect. 7 experimental limitations such as unwanted heating of the ion motion are discussed. This leads to an estimate for the maximum number of unitary operations (quantum gates) that could be carried out in the processor before the coherence of the system is destroyed. It is found that for an example case of around 10 ions, a few hundred operations represents a severe experimental challenge. The use of quantum error correction to enhance the performance is then discussed. This should allow great increases in the number of operations, while preserving coherent evolution. The conclusion outlines the most important avenues for future investigation.

1 Quantum information and computing

Quantum information theory is concerned with understanding the properties of quantum mechanics from an information theoretic point of view. This turns out to be a very fruitful approach, and leads naturally to the idea of information pro-

cessing or computing, so that one poses the question "what are the possibilities for, and the limitations of, information processing in a physical system governed by the laws of quantum mechanics?" A good deal of theoretical insight into this question has been gained; introductory discussions may be found in [3, 4, 8, 24]. For instance, it is possible to identify a small set of 'building blocks'; if they could be realised and many of them combined, a 'universal quantum computer' could be constructed. The computer is 'universal' in the sense that it could simulate, by its computations, the action of any other computer, and so is more or less equal to or better than any other computer [1]. The phrase 'more or less equal' has a technical definition which will be elaborated in Sect. 3.3. A specific set of such building blocks is a set of two-state systems (think of a line of spins), and a simple unitary interaction that can be applied at will to any chosen small set of these two-state systems [2, 3]. In this context it is useful to describe the interaction in terms of its propagator $U = \exp(\mathrm{i}H\Delta t/\hbar)$ rather than its Hamiltonian H. Here Δt is some finite interval of time (one 'clock period' in computing terminology) at the end of which the propagator has had just the effect desired on the computer. After this time the interaction H falls to zero (is turned off). Such a propagator is referred to as a 'gate', by analogy with a logic gate in a classical computer.

For quantum information processing, these requirements may be summed up as the need for a system ('quantum computer' or QC) with a Hilbert space of sufficient number, D, of dimensions, over which you have complete experimental control. That is, you can tell your system to go from any of its states to any other, without uncontrollable error processes such as relaxation and decoherence. Also, one must be able to prepare the initial state and measure the final state.

It is usual to consider a Hilbert space whose number of dimensions is a power of 2, i.e. $D = 2^K$, in which case we say we have a system of K quantum bits or 'qubits'. Examples of qubits are the spin state of an electron (2 orthogonal states and so a single qubit) the polarization state of a photon (a single qubit), and the internal state of an atom having two energy levels of total spin 1 and 2 (8 states and so 3 qubits). Although these are all equivalent from the point of view of the properties of Hilbert space, they are very different from the point of view of experimental implementation. The use of the word 'qubit' rather than 'two-state system' emphasizes this equivalence between otherwise very different quantum systems. In fact, the idea of a qubit has further significance, since it can be shown [5, 6] that the essential properties of any quantum state of any system can be transposed (by interactions allowed by the laws of physics) into the properties of a finite set of qubits and back again [7]. The important point is that the average number of qubits required to do this is equal to the von Neumann entropy of the initial state ("quantum noiseless coding theorem", also referred to as "quantum data compression" [8]). Therefore the qubit gives a measure of *information* content in quantum systems, and is thus the correct quantum equivalent of the classical bit.

With the invention of a new word for the quantum two-state system accepted, there is justified resistance to the adoption of the terms 'computer' and 'computing' to describe the larger quantum systems with which we are concerned. This is because it is an open question whether a true quantum 'computer' could ever function, since once the physical system has sufficient degrees of freedom to be meaningfully called a 'computer', the large-scale interference necessary for parallel quantum computing (see below) may always be destroyed in practice, owing to the sensitivity of such interference to decoherence. For this reason, the more modest term 'information processor' is used here as much as possible. The 'processing' might consist of quite simple manipulations, such as allowing one qubit to interact with another, followed by a measurement of the state of the second qubit. Even such a simple operation has a practical use, since it can be used for error detection at the receiving end of a quantum communication channel, leading to the possibility of secure quantum key distribution for cryptography [9–11].

Decoherence and dissipation in quantum mechanics is a subject in its own right, and has been discussed since the birth of quantum theory. Recent reviews and references may be found in [12–14]. Its impact on quantum computers in particular has been considered [15–17], and will be taken into account in Sect. 7.

Let us temporarily neglect decoherence and all sources of experimental imprecision, and suppose that we are able to prepare a quantum system, of, say, two hundred qubits, and drive it through any prescribed evolution. Such an experimental apparatus would be able to address many interesting questions in physics, but let us concentrate on questions of information processing. Can such a system perform information processing in new (and hopefully powerful) ways? That is, in ways which are ruled out for *any* computer based on standard binary logic bits 0 and 1, rather than quantum states, i.e. qubits? The answer to this question is *yes*. Exactly what is the essence of quantum, as opposed to classical, information processing is still not fully clear, but it appears to depend on quantum superposition, entanglement, and interference. It would make this paper too long to devote much space to this lengthy subject; the reader is referred to [3, 4, 24]. However, a few words are in order.

A classical computer is given an input, which we can imagine without loss of generality to be a binary number x. An algorithm A prescribes various operations to be performed, resulting in an output $A(x)$. In the course of the computation, the computer uses an internal memory in which intermediate results are stored, typically requiring more memory space than would be needed to store just x or $A(x)$ alone.

Similarly, a quantum computer is prepared in an input state $|x\rangle$, and driven towards an output state $|A(x)\rangle$ (the notation adopted here is slightly over-simplified, but this will not affect the point I wish to make). A feature of a quantum computer is that in the course of the computation, superposition states such as $|y\rangle + |z\rangle$ are involved. It is significant that here the quantum state 'stores', after a certain fashion, *both* values y and z, *without* using any more physical resources than those required to store either y or z alone. However, the same could be said for the idea of using classical fields for information storage (for example in a hologram): they also can support superposition. The quantum computer has a further subtlety. With 200 qubits, there are in total 2^{200} dimensions in Hilbert space. This means that a 200-qubit computer can store, in a quantum superposition of mutually orthogonal states, 2^{200} different numbers. The important point is that this is a *huge* effective storage space: it would be quite impossible to manipulate holograms or classical fields with this number of elements, or indeed to build an electronic computer with this size of memory (even with one atom per memory bit, there are not sufficient atoms

in the whole of the Earth!). The next point is that a single operator U applied to a system in such a superposition will evolve all elements of the superposition *simultaneously*, since $U(|y\rangle+|z\rangle)= U|y\rangle+U|z\rangle$. These properties together are referred to as *quantum parallelism*. They hint at the possibility of a huge speed-up in execution time for well-designed quantum algorithms, compared to anything possible on a classical computer. It is significant also that such exponentially large superpositions can be produced in the quantum computer's state by means of a managable number (here, 200) of rotations of the states of individual qubits, so the large superposition is not bought at the cost of large numbers of operations to produce it.

There is a drawback, however. The final result of an algorithm is a single number, not 2^{200} different numbers, so we require our quantum algorithm to be able to bring together its huge number of intermediate results. This relies on quantum interference. The difficulty of implementing quantum computation comes down to the fact that we rely on an interference among a huge number of different states, but such interferences are very sensitive to experimental imprecision.

If a computational task can be framed in such a way as to take advantage of quantum parallelism, and produce an output that depends on a quantum interference between the exponentially large number of intermediate results, then a great speed-up in execution time is obtained, compared to any computer that cannot use such methods (i.e. any classical computer). A computational task that is particularly amenable to this approach is that of finding the period of a function that is simple to evaluate, but whose period is very long, and cannot be deduced by any quicker means than evaluating the function on many inputs. It was by reducing the task of factorising a large integer to this form that Shor showed that an ideal quantum computer can solve the important factorisation problem [18].

Algorithms like that of Shor represent an important insight into the nature of quantum mechanics, but it must be remarked that the existence of such an algorithm does not in itself imply that a quantum computer capable of running it can be built, since in the discussion above we temporarily neglected experimental imprecision and decoherence. An interesting point thus emerges: does nature actually prevent the realisation of efficient quantum algorithms like Shor's, not directly, but via the 'back door' of decoherence? If one chooses any system that might support quantum computation, and makes reasonable estimates of the rate of decoherence, through thermal effects, spontaneous emission, or experimental imprecision, a rough calculation will show that Shor's algorithm cannot succeed on interesting cases, i.e. factorising large numbers ($> 10^{100}$). This has been emphasised by Haroche and Raimond [19]. However, a two-fold attack is underway on this problem. First, one may search for simpler quantum algorithms, by which even a small quantum computer, of, say, 25 qubits, might run important computations. This search is so far unfruitful, but it is interesting to note that 25 qubits is currently the limit on the size of a quantum system that can be thoroughly simulated by classical computers. Second, one may seek ways to make the quantum computer more robust. Here there has been considerable success, based on the idea of *quantum error correction* [65–67, 70–72]. This is a new concept, which will be discussed in Sect. 7.1.

2 Minimum requirements for quantum information processing

It can be shown [20] that to produce arbitrary unitary transformations of the state of a set of qubits, which is what one wants for information processing, it is sufficient to be able to produce arbitrary rotations in Hilbert space of any individual qubit, i.e. the propagator

$$\exp(-\mathrm{i}\boldsymbol{\theta}\cdot\boldsymbol{\sigma}/2)= \begin{pmatrix} \cos(\theta/2) & -\mathrm{e}^{-\mathrm{i}\phi}\sin(\theta/2) \\ \mathrm{e}^{\mathrm{i}\phi}\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}, \tag{1}$$

and to be able to carry out the 'controlled-rotation' operation $\mathrm{CROT} = |00\rangle\langle 00|+|01\rangle\langle 01|+|10\rangle\langle 10|-|11\rangle\langle 11|$ between any pair of qubits. The notation used here is standard, the kets $|0\rangle$ and $|1\rangle$ refer to two orthogonal states of a qubit. This basis is referred to as the 'computational basis', since this aids in designing useful algorithms for the QC. From a physical point of view, it is useful to take the computational basis to be the ground and excited eigenstates of the Hamiltonian of the relevent two-level system, though this is by no means required and any basis will serve. States such as $|01\rangle$ are product states $|0\rangle = |0\rangle\otimes|\,1\rangle$ where the first ket refers to one qubit, and the second to another. For our purposes the qubits will always be distinguishable so we do not need to worry about the symmetry of the states (with respect to exchange of particles) and any related quantum statistics.

As mentioned previously, an operation like CROT is a propagator acting on the state of a pair of qubits. In matrix form it is written

$$U_{\mathrm{CROT}} = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & -1 \end{pmatrix} \tag{2}$$

in the basis $|00\rangle$,$|01\rangle$,$|10\rangle$,$|11\rangle$, where matrix elements that are zero have not been written. The appellation 'controlled rotation' comes from the fact that if the first qubit is in the state $|0\rangle$, CROT has no effect, whereas if the first qubit is in the state $|1\rangle$, CROT rotates the state of the second by the Pauli σ_z operator.

The two operators just described form a universal set, which means that any possible unitary transformation can be carried out on a set of qubits by repeated use of these operators or 'quantum gates', applied to different qubits [20]. Another commonly considered quantum gate is the 'controlled not' or 'exclusive or' (XOR) gate

$$U_{\mathrm{XOR}} = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 01 & \\ & & 10 & \end{pmatrix}; \tag{3}$$

see also (16). This gate has no effect if the first qubit is in the state $|0\rangle$, but applies a NOT operation (σ_x Pauli spin operator) to the second qubit if the first is in the state $|1\rangle$. In the computational basis, this means that the state of the second qubit becomes the XOR of the two input qubit values. We have introduced CROT before XOR in this discussion, going against standard practice, because we shall see later that CROT is easier to implement in an ion trap.

It should be emphasised that this model in terms of quantum gates operating on quantum bits is by no means the only way to think about quantum computation, but is the way that is best understood at present, and is certainly very powerful. The next most studied model is currently that of cellular automata [21], but the field is still young and further models will no doubt be developed in the future.

A further simplification of the physical construction of a quantum computer is as follows. Instead of seeking a means to carry out CROT between any pair of qubits directly, it is sufficient to have one special qubit that can undergo CROT with any of the others. This special qubit acts as a one-bit 'bus' to carry quantum information around the computer, making repeated use of the SWAP operation $|00\rangle\langle 00|+|10\rangle\langle 01|+|01\rangle\langle 10|+|11\rangle\langle 11|$. To carry out CROT between any pair of qubits x and y, one makes use of the bus bit B as follows: $\mathrm{CROT}\,(x,y) = \mathrm{SWAP}\,(B,x)\cdot\mathrm{CROT}\,(B,y)\cdot\mathrm{SWAP}\,(B,x)$. The operation SWAP can be built out of three XORs with the order of the bits alternating: $\mathrm{SWAP}\,(B,x) = \mathrm{XOR}\,(B,x)\cdot\mathrm{XOR}\,(x,B)\cdot\mathrm{XOR}\,(B,x)$. However, in practice this construction is unnecessarily complicated, since SWAP can be applied more or less directly in most physical implementations.

The use of a bus bit makes the physical construction of a quantum information processor much simpler, and indeed most current proposals use this concept. However, it has the major disadvantage that more than one gate (acting on different sets of qubits) cannot be carried out simultaneously (i.e. in parallel), except single qubit rotations. If we accept this limitation, the minimum requirement for our processor is arbitrary rotations of any single qubit, plus CROT and SWAP between the bus qubit and any of the others. This is the minimum set of 'computing operations', in the sense that arbitrary transformations can be carried out by means of this small set. However, this establishes neither that arbitrary transformations can be carried out *efficiently*, nor that they can be carried out *without uncorrectable errors*, both of which are important additional considerations for a computer. We will return to these issues in Sects. 3.3 and 7.1.

A further ingredient for quantum information processing is that the result of the process – here the final state of the quantum system – must be able to be measured without errors. A basis is chosen (typically the eigenbasis of the system Hamiltonian) and a measurement of all the qubits is carried out in this basis.

To make a modest processor (a few qubits) the easiest approach is probably to use single particles with several internal degrees of freedom. Examples are a spin $J = 2^{K-1}-1/2$ in a magnetic field (say $J = 7/2$ giving $2J+1 = 8$ dimensions and therefore $K = 3$ qubits); a molecule or confined particle with 2^K accessible vibrational states ('accessible' in this context means the experimenter can cause computing operations among the states at will). This approach will be interesting in the short term. However, it is difficult to imagine it being extended in the longer term to enable the realisation of a really interesting processor with hundreds of qubits. Also, it is not clear how to apply arbitrary operations to a single particle (spin, molecule) with an evenly spaced ladder of energy levels, owing to level degeneracies in the interaction picture.

There are now several proposed physical systems that might one day make a quantum computer [22–26]. We will concentrate on the system of a line of ions in an ion trap, since it appears to be the most promising at present. However, developments in solid-state physics may overtake us, and one should bear this in mind. It is not easy to couple the quantum information out of an ion trap system (i.e. in the form of qubits, not classical measurements), which is important for quantum communication. In this regard the approach based on strong coupling between an atom and a cavity mode appears more useful, since there a bit of quantum information could in principle be transferred into the polarisation state of a photon, which then leaves the system in a chosen direction (a 'flying qubit') [27]. However, such ideas could be applied to trapped ions, making a form of hybrid processor, so the ion trap system remains an interesting candidate even for quantum communication purposes.

3 Ion trap method

For reviews and references on ion trapping, see for example [28–32]. The ion trap system that interests us uses a line of N trapped ions. Each ion has two stable or metastable states, for example two hyperfine components of the electronic ground state (which usually requires an odd isotope), or two Zeeman sublevels of the ground state, separated by applying a magnetic field. The ground state and a metastable electronic exited state (e.g. a D state for ions of alkaline earth elements) might also be used, but this is a poor choice since the laser linewidth and frequency, as well as most of the mirrors etc. on the optical bench, will have to be very precisely controlled for such an approach to work. There have been optimistic estimates of the computational abilities of an ion trap processor, based on the use of such optical transitions, but one should beware of the lack of realism in such estimates. This will be discussed more fully once we have seen exactly how the system is intended to operate.

There are N laser beam pairs, each interacting with one of the ions, (or a single beam which can be directed at will to any chosen ion), see Fig. 1. Each ion provides one qubit, the two-dimensional Hilbert space being spanned by two of the ion's internal energy eigenstates. A further $(N+1)$th qubit acts as a 'bus' enabling the crucial CROT operations. This qubit is the vibrational motion of the whole ion string in the trap potential. This motion must be quantised, in other words the ion cloud temperature must be reduced well below the 'quantum limit' defined by the axial vibrational frequency in the ion trap:

$$k_{\mathrm{B}}T \ll \hbar\omega_{\mathrm{z}}\,. \tag{4}$$

The first major experimental challenge (after making a trap and catching your ions) is to cool the ions down to this quantum regime. Note that the quantum regime for the trapped motion of the ion is *not* related to the "Lamb–Dicke" regime, which will be considered below. In brief, it will be shown that one wants to operate well into the quantum regime, but on the border of the Lamb–Dicke regime.

So far the quantum regime has been achieved only for a single ion of either mercury in two dimensions [36] or beryllium in three dimensions [37]. Both experiments used optical sideband cooling in the resolved-sideband (tight-trapping) limit. This and other possible cooling techniques will be discussed. Traps for neutral atoms have also attained the motional ground state, most spectacularly in the case of Bose Einstein condensation [39], but also in optical lattices [40]. These systems

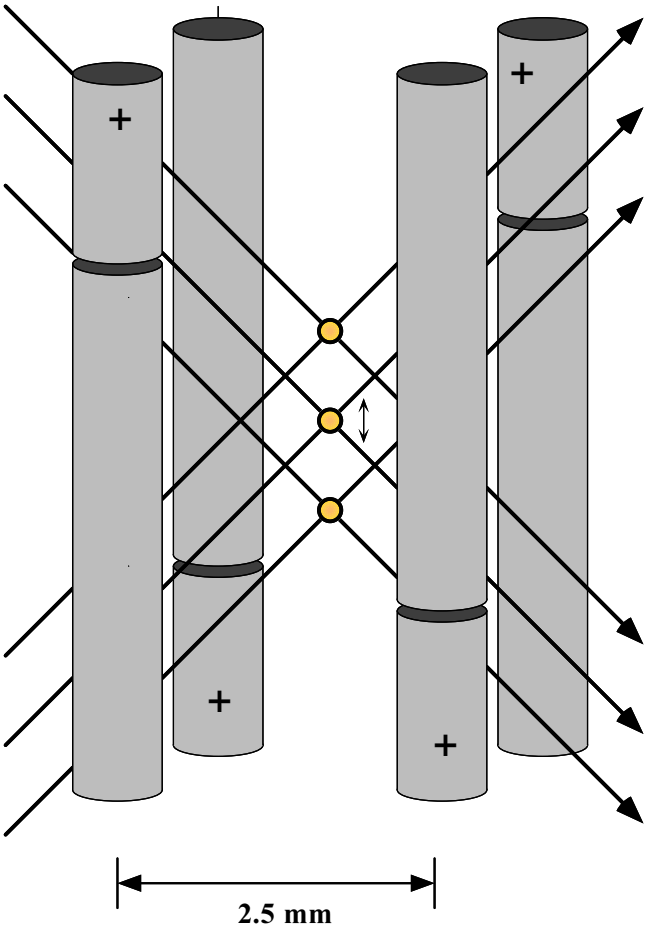

Fig. 1. Experimental arrangement. A line of three ions sits between cylindrical electrodes, here seen sideways on. Pairs of laser beams excite Raman transitions, which impart momentum changes to the ions along the axial direction of the trap. The double-ended arrow indicates the direction of the resulting oscillations, it can be regarded as a pictorial representation of the fourth 'qubit' in the system (see text, Sect. 3.2). The electrodes are split in order to allow a constant voltage to be applied between their ends, so that an axial potential minimum occurs in the region where the long electrode segments overlap. Radial confinement is provided by alternating voltages (see text, Sect. 5)

do not (at present) provide full control of individual atoms and interactions between pairs, so we will not discuss them. However, they lend further weight to the impression that it is in atomic physics and quantum optics, rather than solid-state devices, that quantum information processing will be most fruitful in the immediate future.

To get to the quantum regime, it appears to be neccessary to use a Paul rather than Penning trap, since rf technology allows tighter confinement than does high-magnetic-field technology. Therefore only the Paul trap (rf trap) will be considered from now on, although we may permit ourselves to add a magnetic field if we wish, for some other reason such as to enhance the stability or split the Zeeman levels. In any case, tighter confinement enables a faster 'switching-time' for quantum gates such as CROT, so, as a general rule, tight traps are the best option, though there are some qualifications to this rule, which are discussed in Sect. 3.4.

Note that if several ions are in a three-dimensional rf trap of standard geometry (with the rf voltage between end caps and a ring), then matters are complicated since no more than one ion can be at the centre of the trap potential. Away from the centre, ions undergo rf micromotion and this causes heating if there is more than one ion, owing to collisions (Coulomb repulsion) which force the micromotion out of quadrature with the rf field. To avoid this, one must use a linear or ring geometry. The confinement along the axis is then either due to a static field from end cap electrodes (linear case), or to repulsion between ions combined with their confinement to a ring shape. In this case, only radial micromotion is present, but this vanishes for all the ions if they lie along the axis at the centre of the radial potential, so rf heating is avoided. The ring case must imply a small micromotion tangential to the ring, since the tangential and radial confinement can not be completely decoupled, but as far as I know this has not yet been found to be a problem.

3.1 Average motion

We will model a row of N ions in a trap as a system of N point charges in a harmonic potential well of tight radial confinement, i.e. $\omega_x,\omega_y \gg \omega_z$; see Fig. 2. The oscillation frequencies ω_x,ω_y, and ω_z are parameters that will be obtained from the electrode geometry and potentials in Sect. 5. The total Hamiltonian is

$$H = \sum_{i=1}^{N} \frac{1}{2} M \left(\omega_x^2 \hat{X}_i^2 + \omega_y^2 \hat{Y}_i^2 + \omega_z^2 \hat{Z}_i^2 + \frac{\hat{\boldsymbol{P}}_i^2}{M^2} \right) + \sum_{i=1}^{N} \sum_{j>i} \frac{e^2}{4\pi\varepsilon_0 \, |\hat{\boldsymbol{R}}_i - \hat{\boldsymbol{R}}_j|} . \tag{5}$$

For $\omega_x,\omega_y \gg \omega_z$ and at low temperatures, the ions all lie along the z axis, so we can take $|\hat{\boldsymbol{R}}_i - \hat{\boldsymbol{R}}_j| \simeq |\hat{Z}_i - \hat{Z}_j|$, and the radial and axial motion can be separated. The axial motion interests us, so the problem is one-dimensional. A length scale is given by

$$z_{\mathrm{s}} = \left(\frac{e^2}{4\pi\varepsilon_0 M \omega_z^2} \right)^{1/3} , \tag{6}$$

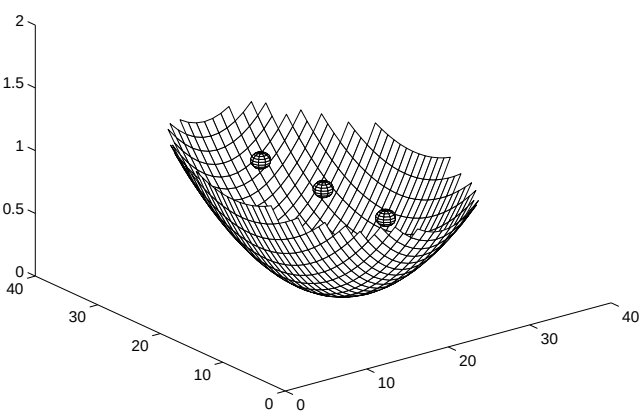

Fig. 2. Schematic illustration showing an anisotropic harmonic potential with the positions of three trapped ions indicated

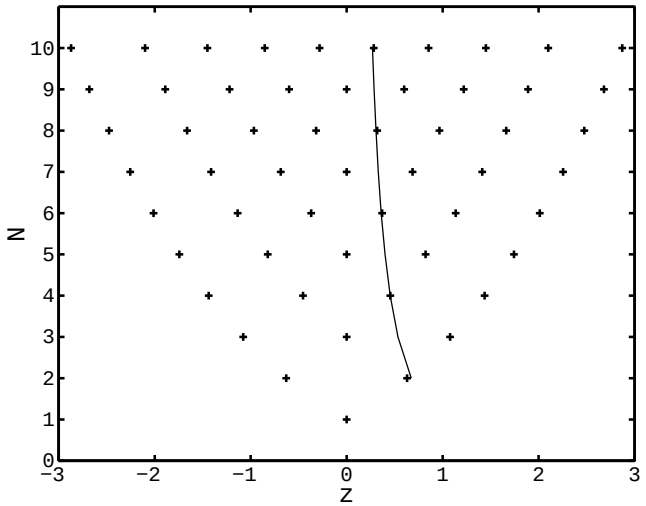

Fig. 3. Equilibrium positions for a line of point charges in a quadratic potential, as a function of the number of charges N. The positions are in units of z_0 defined by (6). The curve is Eq. (9)

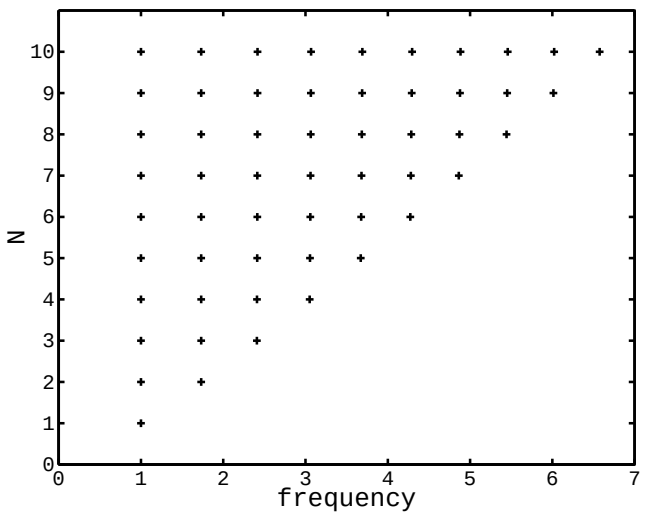

Fig. 4. Normal mode frequencies for a line of point charges in a quadratic potential, as a function of the number of charges N. The frequencies are in units of ω_z

which is of the order of the separation between the ions (typically 10 to 100 μm). Solving the classical equations of motion (i.e. the operators $\hat{Z}$, $\hat{P}_z$ become classical variables z, p_z), one obtains the equilibrium positions shown in Fig. 3. With more than two trapped ions, the outer ions tend to push the inner ones closer together, so the ion positions depend on N [see (9)]. Remarkably, however, the frequencies of the first two normal modes of oscillation about these equilibrium positions are independent of N (for small oscillations) [25], and those of higher modes are nearly independent of N. The frequencies of the first two modes are ω_z and $\sqrt{3}\omega_z$, and those of higher modes are given approximately by the list $\{1,\ \sqrt{3},\ \sqrt{29/5},\ 3.051,\ 3.671,\ 4.272,\ 4.864,\ 5.443,\ 6.013,\ 6.576\}$, which gives the frequency of the highest mode, in units of ω_z, for $N = 1$ to 10. The near independence of N of the mode frequencies is illustrated by Fig. 4.

The lowest mode of oscillation corresponds to harmonic motion of the centre of mass of the ion string. In this mode, all the ions move to and fro together. It is important that the frequency of this mode is significantly different from that of any other mode, since this means that experimentally one can excite the centre-of-mass mode without exciting any of the others.

We can now proceed directly to a quantum mechanical treatment, simply by treating the centre-of-mass coordinate z_{cm} as a harmonic oscillator. The classical result that the centre-of-mass normal mode has frequency ω_z remains valid even though the ion wavefunctions may now overlap, since all the internal interactions among the ions cancel when one calculates the centre-of-mass motion. Since we have an oscillator of mass NM and frequency ω_z, the energy eigenfunctions are

$$\psi_n(z_{\mathrm{cm}}) = \left(\frac{NM\omega_z}{\pi\hbar 2^{2n}(n!)^2}\right)^{1/4} H_n\left(z_{\mathrm{cm}}\sqrt{NM\omega_z/\hbar}\right) \times \mathrm{e}^{-NM\omega_z z_{\mathrm{cm}}^2/2\hbar}. \tag{7}$$

The spatial extent of the Gaussian ground state probability distribution is indicated by its standard deviation

$$\Delta z_{\mathrm{cm}} = \sqrt{\hbar/2NM\omega_z}. \tag{8}$$

Since we want a different laser beam to be able to address each of the ions, we require Δz_{cm} to be small compared to the separation between ions. The closest ions are those at the centre of the line. A numerical solution of (5) yields the following formula for the separation of the central ions:

$$\Delta z_{\mathrm{min}} \simeq 2.0\, z_{\mathrm{s}} N^{-0.57}. \tag{9}$$

This formula is plotted for $N \leq 10$ in Fig. 3. An approximate analytical treatment for $N \gg 1$ does not predict a power-law dependence of Δz_{min} on N, but rather $\Delta z_{\mathrm{min}} \propto z_{\mathrm{s}}(\log(N)/N^2)^{1/3}$ [42]. However, (9) is more accurate for $N < 10$ and remains accurate for the range of N that interests us (up to, say, $N = 1000$). Setting $\Delta z_{\mathrm{cm}} \ll \Delta z_{\mathrm{min}}$ yields

$$\frac{\omega_z}{M} \ll \frac{32N^{1.86}}{\hbar^3}\left(\frac{e^2}{4\pi\varepsilon_0}\right)^2 \simeq 2.4\times 10^{21}\, N^{1.86}\,\mathrm{Hz/u}, \tag{10}$$

where u is the atomic mass unit 1.66057×10^{-27} kg. This condition is easily fulfilled in practice, with ω_z no greater than a GHz, and M between 9 and 200 u. Therefore it is legitimate to picture the ions as strung out in a line, each sitting in a small wavepacket centred at its classical equilibrium position, not overlapping the others. Note that (10) does not guarantee that the ions are sufficiently separated to be addressed by different laser beams, only that their wavefunctions do not overlap.

In the above, it was assumed that the radial confinement was sufficient to cause the ions to lie along the z axis, rather than form a zigzag or helix about it. The onset of such zigzag modes has been studied numerically [41] and analytically [42]. They occur when the ions approach sufficiently closely that the local potential minimum at the position of an ion on the z axis becomes a saddle point. For a string of ions uniformly spaced by Δz (which is not the case in our harmonic trap), the transition from a line to a zigzag occurs when

$\omega_r^2 \simeq 4.2072(z_s/\Delta z)^3\omega_z^2$ [43], where we have taken the case $\omega_x = \omega_y \equiv \omega_r$. Setting $\Delta z = \Delta z_{\min}$ leads to the condition

$$\frac{\omega_r}{\omega_z} > 0.73N^{0.86} \tag{11}$$

for the prevention of zigzig modes. For $N \gg 1$, an approximate analytic treatment yields the condition [42]

$$\frac{\omega_r}{\omega_z} > 0.77\frac{N}{\sqrt{\log N}}. \tag{12}$$

These numerical and approximate analytic formulae are within 10% agreement for $3 < N < 2000$.

3.2 Principle of operation

The principle of operation of an ion trap 'information processor' was described by Cirac and Zoller [25], and the most important elements of such a system were first realised in the laboratory by Monroe et al. [44]. Whereas the transition operators given by Cirac and Zoller were calculated for standing-wave excitation of allowed single-photon transitions, experimentally Monroe et al. employed travelling-wave excitation of two-photon Raman transitions (cf Figs. 1 and 8). The basic form of the operators is independent of the type of excitation used, however. The method may be understood by reference to Fig. 5, which shows the relevant energy levels for one of the ions in the trap. We consider three of the ion's internal energy eigenstates $|F_1,M_1\rangle$, $|F_2,M_2\rangle$ and $|F_{\rm aux},M_{\rm aux}\rangle$, and various excitations of the centre-of-mass motion. The ion's internal energy levels are separated in frequency by ω_0 and $\omega_{\rm aux}$

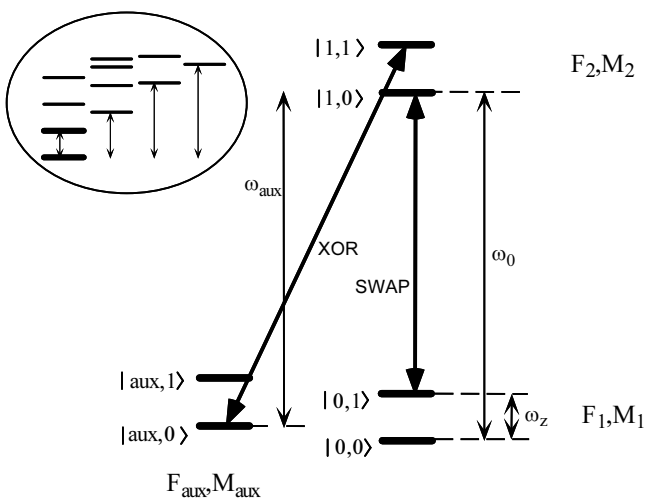

Fig. 5. Energy levels and transitions in a single ion significant for information processing with a line of trapped ions. The labels F, M indicate different internal states of the ion. Each internal state has an associated set of vibrational levels for each of the vibrational modes. Here, just the ground and first levels of the lowest mode (spacing ω_z), are shown in the main diagram, and the insert shows the further low-lying vibrational levels whose excitation we wish to avoid. The full arrows indicate transitions at frequencies $\omega_0 - \omega_z$ and $\omega_{\rm aux} + \omega_z$, which are used in the SWAP$_{(-i)}$ and CROT operations described in the text. Note that radiation at a given frequency couples not only the levels at the two ends of the relevant arrow on the diagram, but also other pairs of levels with the same difference of vibrational quantum number. The figure shows $\omega_{\rm aux}$ to be of the same order as ω_0, because this is what typically occurs when alkali-like ions are used. The vibrational frequency is smaller, $\omega_0 \simeq 1000\omega_z$ as indicated on Fig. 5. Note that all these levels are low-lying, separated from the ground state only by hyperfine and Zeeman interactions (see Fig. 8), so their natural lifetime against spontaneous emission of rf photons is essentially infinite. Figure 5 shows the lowest-lying excitations of the second, third, and fourth normal modes as well as the first, to act as a reminder of the location of the closest extraneous levels whose excitation we wish to avoid. The energy eigenstates of the vibrational motion may be written as $|n_1,n_2,n_3,\ldots\rangle$, where the n_i are the excitations of the various normal modes. Only the ground state $|0,0,0,\ldots\rangle$ and first excited state of the centre of mass $|1,0,0,\ldots\rangle$ will be involved in the operations we wish to invoke. This centre-of-mass vibrational degree of freedom is often referred to somewhat loosely as a 'phonon'. The 'computational basis' consists of the states

$$\begin{aligned}
|0,0\rangle &\equiv |F_1,M_1\rangle \otimes |0,0,0,\ldots\rangle\,,\\
|0,1\rangle &\equiv |F_1,M_1\rangle \otimes |1,0,0,\ldots\rangle\,,\\
|1,0\rangle &\equiv |F_2,M_2\rangle \otimes |0,0,0,\ldots\rangle\,,\\
|1,1\rangle &\equiv |F_2,M_2\rangle \otimes |1,0,0,\ldots\rangle\,.
\end{aligned} \tag{13}$$

It will now be shown how to carry out CROT between any single ion's internal state and the bus (phonon) bit, then how to carry out arbitrary rotations of the internal state of an ion, then how to carry out SWAP between any ion and the bus bit. From the discussion in Sect. 2, these three operations form a universal set and so allow arbitrary transformations of the stored qubits in the processor.

The auxillary states $|{\rm aux},i\rangle \equiv |F_{\rm aux},M_{\rm aux}\rangle \otimes |i,0,0,\ldots\rangle$ ($i = 0,1$) are available as a kind of 'shelf' by means of which useful state-selective transformations can be carried out among the computational basis states. If one applies radiation at the frequency $\omega_{\rm aux} + \omega_z$, then inspection of Fig. 5 will reveal that only transitions between $|1,1\rangle$ and $|{\rm aux},0\rangle$ will take place (if we assume that unwanted levels such as $|1,2\rangle$ are unoccupied).[1] If one applies a 2π pulse at this frequency, then the state $|1,1\rangle$ is rotated through 2π radians, and therefore simply changes sign. In the computational basis, the effect is equal to that of the CROT operator described in Sect. 2, see (2).

A 2π pulse at frequency $\omega_0 - \omega_{\rm aux} - \omega_z$ also produces a controlled rotation, only now the minus sign appears on the second element down the diagonal of the unitary matrix, rather than on the fourth, causing a sign change of the component $|0,1\rangle$ rather than of $|1,1\rangle$. This case will be called C$_\neg$ROT, the negation symbol $\neg$ referring to the fact that here the second qubit is rotated if the first is in the state $|0\rangle$ rather than $|1\rangle$.

To rotate an ion's internal state without affecting the centre-of-mass motion, one applies radiation of frequency ω_0. If such radiation has phase ϕ with respect to some defined origin of phase, and duration sufficient to make a $p\pi$ pulse, then the effect in the computational basis is

$$V^p(\phi) \equiv \begin{pmatrix} \cos(p\pi/2) & -{\rm i}{\rm e}^{-{\rm i}\phi}\sin(p\pi/2) \\ -{\rm i}{\rm e}^{{\rm i}\phi}\sin(p\pi/2) & \cos(p\pi/2) \end{pmatrix}_{\rm ion} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}_{\rm cm}, \tag{14}$$

[1] Cirac and Zoller originally proposed to produce the selective effect of this CROT operation by means of a chosen laser polarisation rather than frequency. However, frequencies can be experimentally discriminated more precisely than polarisations, which explains why Monroe et al. chose to use a frequency-selective rather than polarisation-selective method.

where we have followed the notation of [25], but used p instead of k to avoid confusion with the wave vector. Note that to apply such rotations succesfully, it is necessary to have the phase of the radiation under experimental control. That means under control at the position of the ion, not just in some stable reference cavity. This constitutes a severe experimental constraint, which makes computational basis states separated by radio frequencies highly advantageous compared to states separated by optical frequencies. On the other hand, in order to have the right phase experimentally, note that one need not worry about the continuous precession at frequency ω_0 caused by the internal Hamiltanian of each ion. The laser field keeps step with this precession, as becomes obvious when one uses the interaction picture, which we have done implicitly in writing (14). A possible problem arises when different ions have different internal energies, owing to residual electric and magnetic fields in the apparatus. However, this particular problem has a fairly simple experimental solution: since each ion is addressed by a different laser beam pair, one can independently tune the laser beams driving each ion, by means of acousto-optic modulators in the beam paths. Such modulators will be required in any case to allow the laser beam intensities to be switched.

The centre-of-mass motion acts as the 'bus' qubit described in Sect. 2. To carry out $\mathrm{XOR}(B,x)$, between the 'bus' and the internal state of a single trapped ion, Monroe et al. applied first a $\pi/2$ pulse at frequency ω_0,

$$V^{1/2}(-\pi/2) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}_{\text{ion}} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}_{\text{cm}}, \tag{15}$$

followed by CROT as described in the paragraph after (13), followed by a second $\pi/2$ pulse at ω_0 with phase displaced by π with respect to the first,[2] i.e. $V^{1/2}(\pi/2)$. A straightforward calculation shows that this sequence produces exactly

$$\mathrm{XOR}(\mathrm{cm},\,\mathrm{ion}) = \begin{pmatrix} 1000 \\ 0001 \\ 0010 \\ 0100 \end{pmatrix}, \tag{16}$$

By symmetry, to obtain $\mathrm{XOR}(\mathrm{ion},\,\mathrm{cm})$, one might imagine using a similar sequence, but with the $\pi/2$ pulses applied at frequency ω_z so as to affect the vibrational state without affecting the internal state. However, the vibrational degree of freedom is not really a two-level system, so this will not work (indeed, it will cause unwanted multiple excitations of the vibrational motion). To perform CROT, we made use of a transition at frequency $\omega_{\text{aux}}+\omega_z$. Note that this relied on the fact that there was no population in the state $|\text{aux},1\rangle$ (which would have become coupled to $|1,2\rangle$, which is outside the computational Hilbert space). This illustrates the general method by which the vibrational state is influenced: one uses radiation at a frequency offset from an internal resonance of the ion by ω_z, thus coupling levels of vibrational quantum numbers differing by 1. To avoid coupling higher-lying vibrational states, one of the possible initial states must be unoccupied when such a transition is invoked.

[2] In fact Monroe et al. state that they used $V^{1/2}(\pi/2)$ for the first pulse, and $V^{1/2}(-\pi/2)$ for the third, producing XOR with an additional rotation of the centre-of-mass state.

The transition at frequency $\omega_0-\omega_z$ is indicated on Fig. 5. A moment's reflection allows one to convince oneself that as long as there is no population in the $|1,1\rangle$ state (nor in extraneous states such as $|0,2\rangle$), application of this radiation will cause only transitions between $|1,0\rangle$ and $|0,1\rangle$, and hence a SWAP operation is available between the bus qubit and any other. Applying a $p\pi$ pulse at phase ϕ and frequency $\omega_0-\omega_z$, we obtain the operation

$$U^p(\phi) \equiv \begin{pmatrix} 1 & 0 & 0 & \times \\ 0 & \cos(p\pi/2) & -\mathrm{i}e^{-\mathrm{i}\phi}\sin(p\pi/2) & \times \\ 0 & -\mathrm{i}e^{\mathrm{i}\phi}\sin(p\pi/2) & \cos(p\pi/2) & \times \\ 0 & 0 & 0 & \times \end{pmatrix}, \tag{17}$$

where the final column of crosses indicates that an initial state $|1,1\rangle$ is carried out of the computational basis by $U^p(\phi)$. The case $p=1$, that is a π pulse, produces a SWAP operation with an additional $-\mathrm{i}$ phase factor, which we will write $U^1(0) = \mathrm{SWAP}_{(\mathrm{i})}$. Applying $U^1(0)$ to ion x, followed by $\mathrm{C}_{\neg}$ROT to ion y (i.e. using the frequency $\omega_0-\omega_{\text{aux}}-\omega_z$), followed by $U^1(0)$ once again to ion x, has the effect of a CROT operation between x and y. That is, $\mathrm{CROT}(x,y) = \mathrm{SWAP}_{(\mathrm{i})}(B,x)\cdot \mathrm{C}_{\neg}\mathrm{ROT}(B,y)\cdot \mathrm{SWAP}_{(\mathrm{i})}(B,x)$, as long as the initial state of x and the bus is not $|1,1\rangle$. To apply the method, one uses the bus as a 'work bit' that is arranged always to return to state $|0\rangle$ before operations such as $U^p(\phi)$ are applied, so the quantum information processing can go forward without problem.[3]

So far we have described operations on the ion trap by means of $p\pi$ pulses. A complimentary technique is that of adiabatic passage, in which a quantum system is guided from one state to another by a strongly perturbing Hamiltonian applied slowly. For example, instead of swapping one ion's internal state with the bus qubit, and then swapping the bus with another ion, one could swap the internal state of two ions 'via' the bus but without ever exciting the first vibrational level. The details are described for a related system in [26]. This method has experimental advantages in being insensitive to features such as the timing and interaction strength. However, it can be more sensitive to off-resonant coupling to other levels, compared with the pulse technique, and this consideration led Monroe et al. to favour rf pulses. Both $p\pi$ pulses and adiabatic passage will probably have their uses in a practical QC (whether based on an ion trap or some other system).

The laser pulses described provide the universal set of 'quantum logic gates' for the linear ion trap. To complete the operation of our processor or QC, we require that the final state of the quantum information processor can be measured with high accuracy. This is possible for trapped ions by means of the 'electron-shelving' or 'quantum-jumps' technique [28, 30, 32–34]. That is, one may measure whether a given ion is in state $|F_1,M_1\rangle$ or $|F_2,M_2\rangle$ by illuminating it with radiation resonant with a transition from $|F_1,M_1\rangle$ to some high-lying level, whose linewidth is small enough so that transitions from $|F_2,M_2\rangle$ are not excited. If fluorescence is produced (which may be detected with high efficiency), the ion state has collapsed to $|F_1,M_1\rangle$, if none is produced, the ion state has collapsed to $|F_2,M_2\rangle$. This method requires that spontaneous decay from the high-lying level to $|F_2,M_2\rangle$ is

[3] Indeed, the bus may even be measured at those times when it should be in the ground state $|0\rangle$, producing a slight stabilisation or error detection; see Sect. 7.1. In the ion trap, however, one can only thus measure the vibrational state by first swapping it with the internal state of a prepared ion.

forbidden, which is possible in this case through the angular momentum selection rules for electric dipole radiation. Note that typically a quantum process may need to be prepared, run, and measured several times in order to gather more information about the processor's final state than is available from a single measurement of all the qubits. This is related to the ideas of quantum state tomography, which have recently been demonstrated in an ion trap experiment (see [35] and references therein).

3.3 Efficient gate sequences

It was shown in the previous section that CROT can be applied to any pair of qubits, and arbitrary rotations of single qubits can be carried out. Hence, as explained in Sect. 2, any arbitrary sequence of unitary transformations of the quantum processor can be brought about. However, the most efficient methods will not blindly adopt a simple repetition of CROTs and rotations to solve any problem. There may be much more efficient methods, by using other possible pulse sequences. Cirac and Zoller emphasize this by demonstrating how to apply a $\mathrm{C}^n\mathrm{ROT}$ operation, in which the σ_z operator is applied to one ion's internal state only if n other ions are in the state $|1\rangle$, using a number of pulses equal to $2(n-2)+3$. This is efficient in that the number rises only linearly with n, and the multiplying factor is small (i.e. 2 rather than 48 as in [45]).

Efficiency in computer science has a rigorous definition. Without the details, the essential point is that if the number of elementary computational steps (here, quantum gates) required to complete an algorithm rises exponentially with the size of the input to the algorithm, then the algorithm is inefficient. The definitions can be made rigorous, which we will not attempt to do, but essentially each algorithm addresses not one instance of a problem, such as to "find the square of 2357", but a whole class of problems, such as, "given an integer x, find its square". The 'size of the input' to the algorithm is measured by the amount of information required to specify x, which is the number of digits in the binary expression of x, i.e. $\log_2(x)$. A computation is inefficient if the number of steps is exponential in $\log(x)$, i.e. is proportional to x. Similarly, a quantum gate involving n qubits is inefficient if the number of physical operations, such as laser pulses, required to implement it is exponential in n (e.g. increases as 2^n). The strict definition of the universal computer mentioned in Sect. 1 also involves this efficiency aspect: when a universal computer simulates the action of another, the number of operations in the simulation algorithm must not rise exponentially with the amount of information required to define the simulated computer.

Although we emphasised in Sect. 2 that a small set of gates is 'universal' in that all unitary transformations can be composed by them, this does *not* necessarily imply that they can be used to build the particular transformations we may want in an efficient way. In this sense, the word 'universal' is misleading.

So far, networks of quantum gates have been designed for the most part without regard to the exact physical process that might underlie them. However, in such an approach it is not obvious which gates to call 'elementary', since a physical system like the ion trap may be particularly amenable to some transformations. We have already seen an example in the SWAP gate in the ion trap, which can be carried out without recourse to a sequence of XOR gates. This implies that a thorough understanding of a particular system like the ion trap may lead to progress in finding efficient networks. The important insight in Cirac and Zoller's construction of the $\mathrm{C}^n\mathrm{ROT}$ gate is that the method makes use of π pulses at frequency $\omega_{\mathrm{aux}}+\omega_z$. In other words, during the implementation of this gate the ions are deliberately carried out of the computational Hilbert space. Alternatively, one could regard the 'shelf' level $|F_{\mathrm{aux}},M_{\mathrm{aux}}\rangle \otimes |\mathrm{i},0,0,\ldots\rangle$ as within the computational Hilbert space, in which case we have more than one qubit available per ion. Later, in Sect. 7.1, we will consider using vibrational modes in addition to the lowest one in order to have more than one 'bus' qubit.

3.4 Switching rate

The previous section showed how the ion trap information processor worked, by invoking radiation of prescribed frequency and duration in the form of $p\pi$ pulses. The 'switching rate' of the processor is limited by the duration of these pulses.

Let Ω be the Rabi frequency for resonant excitation of the internal transition at frequency ω_0 for a free ion. This will be determined by the linestrength of the transition and the laser power available. For a two-level atom one has $\Omega^2 = 6\pi\Gamma I/\hbar c k^3$ where Γ is the linewidth of the transition, I is the intensity of the travelling wave exciting the transition, and k is the wavevector. When we consider excitations of the internal state alone of an ion in a trap, i.e. $\Delta n = 0$ where n is the vibrational quantum number, this 'free ion' Rabi frequency still applies. However, when changes in the vibrational state are involved, i.e. transitions at frequency $\omega_0 \pm \omega_z$ producing $\Delta n = \pm 1$, an additional scaling factor $\Delta z_{\mathrm{cm}} k_z$ appears, where Δz_{cm} is the extent of the ground-state vibrational wave function given in (8), and $k_z = k\cos(\theta)$ is the wavevector component along the z direction. Using (8), we have

$$\Delta z_{\mathrm{cm}} k_z = \left(\frac{\hbar k^2 \cos^2(\theta)}{2NM\omega_z}\right)^{1/2} \equiv \frac{\eta}{\sqrt{N}}, \tag{18}$$

where η is the Lamb–Dicke parameter for a single trapped ion. In the case of weak excitation, the effective Rabi frequency for the vibrational-state-changing transitions is $\eta\Omega/\sqrt{N}$, a result that can be interpreted as arising from conservation of momentum. The factor $\sqrt{N}$ appears because the whole ion string moves *en masse* and therefore has an effective mass NM (Mössbauer effect). The Lamb–Dicke parameter can also be written in terms of the recoil energy (energy of recoil of an ion after emission of a single photon)

$$E_{\mathrm{R}} \equiv \frac{(\hbar k)^2}{2M}, \tag{19}$$

giving $\eta \equiv \cos(\theta)(E_{\mathrm{R}}/\hbar\omega_z)^{1/2}$.

We can now obtain a measure of the switching rate R by taking it as the inverse of the time to bring about a 2π pulse on a vibrational-state-changing transition, i.e.

$$R \simeq \frac{\eta\Omega}{2\pi\sqrt{N}} \quad (\eta \leq \sqrt{N}). \tag{20}$$

Outside the Lamb–Dicke limit[4] (i.e. for $\eta > \sqrt{N}$) the ion–radiation interaction is more or less equivalent to that of a free ion, so the factor $\eta/\sqrt{N}$ no longer applies; it is replaced by a factor less than or equal to 1.

It was remarked in the previous section that to maintain phase control between (and during) computing operations, there is a strong advantage in having the transition frequencies ω_0, ω_{aux} in the rf to microwave rather than optical region of the electromagnetic spectrum. However, if the relevant transitions are driven directly by microwave radiation, with a frequency of the order of the vibrational frequency ω_z, then the Lamb–Dicke parameter is extremely small (of order $(\hbar\omega_z/2Mc^2)^{1/2}$), so vibrational-state-changing transitions are almost impossible to drive. One way to avoid this would be to make the trap extremely weak, but this has the disadvantage of making the system sensitive to perturbations and lowering the switching rate. Instead, it is better to drive the microwave transitions by Raman scattering at optical frequencies. This combines the advantage of a large photon momentum and hence strong driving of vibrational-state-changing transitions, with the possibility of accurate phase control since only the phase *difference* between the pair of laser beams driving a Raman transition need be accurately controlled. The Raman technique was adopted for these reasons by Monroe et al. [44]. The same reasoning leads to the advantage of Raman scattering for precise laser manipulation of free atoms [47, 48]. A clear theoretical analysis is provided by [49].

The maximum switching rate is dictated by the three frequencies Ω, ω_z, and $E_R/\hbar$ in a subtle way. If only low laser power is available, $\Omega \ll \omega_z$, then the Rabi frequency limits the switching rate and the best choice for ω_z is that which makes $\eta \sim \sqrt{N}$, i.e.

$$\eta^2 \equiv \frac{\cos^2(\theta) E_R}{\hbar\omega_z} \simeq N\,. \tag{21}$$

Therefore the recoil energy, given by the choice of ion and transition, dictates the choice of trap strength, for a given number of ions. Typical recoil energies for an ion are in the region $E_R \sim 2\pi\hbar \times (10\text{–}200)$ kHz, and traps with this degree of confinement are now standard. In this situation, increasing the number of ions does not affect the switching rate, but reduces the required trap confinement, making the system more sensitive to perturbations.

If higher Rabi frequencies are available, one's intuition suggests that ω_z becomes the limit on the switching rate, since Ω must be less than ω_z or the power broadening will no longer allow the different vibrational levels to be discriminated. However, at high ω_z one has $\eta \ll 1$ (Lamb–Dicke regime) so the switching rate on $\Delta n = \pm 1$ transitions cannot reach ω_z if $\Omega < \omega_z$. Placing the *ad hoc* limit $\Omega < \omega_z/10$ in (20), one obtains

$$R < \frac{1}{20\pi}\left(\frac{E_R\omega_z}{\hbar N}\right)^{1/2}. \tag{22}$$

The switching rate is thus limited by the geometric average of ω_z and $E_R/\hbar$, and the processor slows down when more ions are involved. For example, to achieve a switching at the recoil frequency, i.e. $R = E_R/2\pi\hbar$, with $N = 10$ ions, (22) implies $\omega_z = 1000E_R/\hbar$ and $\Omega = 100E_R/\hbar$. To keep the ions in a straight line, (11) requires $\omega_r > 5300E_R/\hbar$ which is very hard to achieve experimentally.

There is another problem with increasing ω_z in order to increase R. When ω_z is large, $\eta \ll \sqrt{N}$, so the transitions that do not change the vibrational state ($\Delta n = 0$) are much more strongly driven by the laser than those that do [$\Delta n = \pm 1$, equation (20)]. This increases the unwanted off-resonant driving of $\Delta n = 0$ transitions when $\Delta n = \pm 1$ transitions are invoked to perform quantum gates between an ion and the phonon 'bus'. Cirac et al. [25, 57] have emphasized the possibility of using standing-wave rather than travelling-wave excitation to avoid this problem, since $\Delta n = 0$ transitions are suppressed if the ion is positioned in the node of a standing wave. However, it may not be technically feasible to achieve this for more than a few of the ions.

In principle it should be possible to run an ion trap processor at rates of order ω_z by relaxing the condition $\Omega < \omega_z$ and allowing off-resonant transitions, but the simple analysis given in Sect. 3.2 is then no longer valid. One can no longer use a two-level model for each transition of the ion/centre-of-mass system. The ac (alternative current) Stark effect (light shift) will be all-important, and different computational basis states will become mixed by the ion–light interaction. The optical Bloch equations remain solvable (numerically if not analytically), and a detailed analysis should still enable useful elementary computing operations to be identified. Such an analysis is a possible avenue for future work [46].

4 Cooling

To make the quantum information processor described in the previous sections, the main initial requirements of an experimental system are cooling to the quantum regime (4), and confinement to the border of the Lamb–Dicke regime (21). The ions must be separated by at least several times the laser wavelength [see (9)], but this is automatically the case, for small numbers of trapped ions, since with current technology the ions are always separated by many times the width of their vibrational ground state wavefunction [inequality (10)], which is itself approximately equal to the laser wavelength, given that the Lamb–Dicke parameter is of order 1.

Surveys of cooling methods in ion traps are given in [31, 32]. For cooling to the quantum regime, there are two possible approaches. Either one may cool to the ground state of a tight trap having $\eta \ll 1$, then adiabatically open the trap to $\eta \sim 1$, or one may apply cooling to a trap already at $\eta \sim 1$. The advantage of the former approach is that one does not require cooling below the recoil limit $k_B T_R = E_R$. The advantage of the latter is that strong confinement is not necessary, but attaining the quantum regime with $\eta \sim 1$ requires sub-recoil cooling.

Cooling to the quantum regime has so far been demonstrated for trapped ions by means of sideband cooling in the resolved-sideband limit [36, 37]. This is described in Sect. 4.1 below. However, it may be interesting to pursue other approaches, as discussed in Sect. 4.2 and Sect. 4.3.

The physics of sideband cooling is very closely related to that involved in information processing in the ion trap. This is no coincidence, and a similar link will probably be found in all

[4] The Lamb–Dicke limit is the condition that the ions' motion is confined to a region small compared to a wavelength of the radiation under consideration. Since in the present case the ion string is in the ground state of the trap (quantum regime), the Lamb–Dicke limit corresponds to small η.

physical implementations of quantum information processing. The relationship is sufficiently close that one may say that once the goal of laser cooling to the motional ground state is achieved in any given experimental ion trap, a primitive form of quantum information processing can proceed immediately, since all the required experimental components will be in place. Conversely, quantum error correction (see Sect. 7.1) is a special type of 'cooling'.

4.1 Sideband cooling

Sideband cooling is the name for the simplest type of laser cooling of a confined ion. The name comes from how the photon-scattering process looks in the resolved-sideband limit (see below). In the case of free ions, the corresponding process is radiation pressure or Doppler cooling.

There are several significant frequencies or energies. First, we have the vibrational frequency in the ion trap potential, ω_z. Next, we have the radiative width of the transition used to do the cooling, Γ. Either a single photon transition is used, in which case Γ is its natural width (or possibly its broadened width if another laser is used to broaden a very narrow level as in [36]), or a stimulated Raman transition is used, in which case Γ is some combination of the inverse of the duration of the Raman pulses, and the time for optical pumping out of one of the states linked by the Raman transition. The physics in the Raman case and single-photon case is very similar. The Raman method is a way of providing a very narrow transition when one is not already available. It also combines the advantages of precise frequency control (in the rf regime) with large photon recoil (optical regime), which permits fast cooling, for the same reason that the switching rate for information processing is faster (Sect. 3.4). One could instead use an rf or microwave transition, but then the cooling would be a lot slower and may not compete well enough with heating processes.

Laser cooling of atoms is often done quite happily by using strong, resonant transitions. Indeed, such transitions are eagerly sought out. Why the talk of narrow transitions in the previous paragraph? It is because simple Doppler cooling leads to the well-known Doppler cooling limit $k_B T_D \simeq \hbar\Gamma/2$, when the recoil energy is small compared to $\hbar\Gamma$ (this applies in a trap as well as to free atoms). However, we want to get to the quantum limit (4), so we require

$$\omega_z \gg \Gamma. \tag{23}$$

This equation is a further constraint on the performance of the trap. It says the cooling transition must be narrow enough, or the trap confinement tight enough, to resolve the motional sidebands in the Lamb–Dicke spectrum.

In the resolved sideband limit, radiation pressure cooling is called sideband cooling. A nice way of understanding it is to consider it as a form of optical pumping towards the state of lowest vibrational quantum number [31]; see Fig. 6. Note that the recoil after spontaneous emission produces heating. The average change in the vibrational energy per spontaneous emission is equal to the recoil energy $\hbar\omega_z\langle\Delta n\rangle = E_R$ (a particularly clear derivation of this fact may be found in [51]). For a single trapped ion illuminated by low-intensity light, the cooling is governed by the following equation [51, 52]:

$$\frac{\mathrm{d}}{\mathrm{d}t}\langle H\rangle = \frac{I\sigma_0}{\hbar\omega_L}\sum_n P_n \sum_f (E_f - E_n + E_R)\left|\langle\psi_n|\,\mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{R}}|\psi_f\rangle\right|^2 \times g\left(\omega_L - (E_f - E_n)/\hbar\right), \tag{24}$$

where I is the intensity of the incident radiation (a single travelling wave), σ_0 is the resonant photon scattering cross section ($\sigma_0 = 2\pi\lambda^2 = (2\pi)^3/k^2$ for a two-level atom), $\hbar\omega_L = \hbar ck$ is the laser photon energy, P_n is the occupation probability of the nth energy level of the vibrational motion, of energy $E_n = \hbar\omega_z(n+1/2)$ and wavefunction ψ_n [see (7)], and $g(\omega)$ is the lineshape function. For a two-level atom,

$$g(\omega) = \frac{\Gamma^2/4}{(\omega-\omega_0)^2 + \Gamma^2/4}. \tag{25}$$

The quantity $\mathrm{d}\langle H\rangle/\mathrm{d}t$ is the rate of change of the mean total energy of the ion, averaged over an absorption/spontaneous emission cycle. Since the ion's internal energy is left unchanged, this is the rate of change of the mean kinetic energy. Equation (24) has a simple physical interpretation as a sum of energy changes associated with radiative transitions up and down the ladder of vibrational energy levels. At the lowest attainable temperature, $\mathrm{d}\langle H\rangle/\mathrm{d}t = 0$ and one possible solution

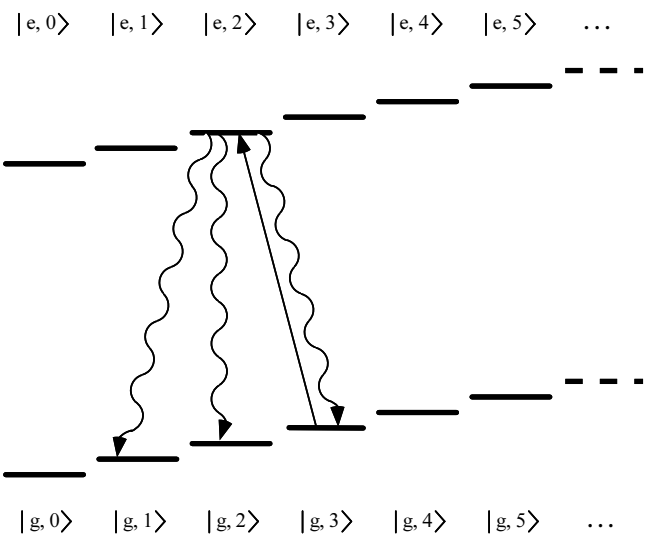

Fig. 6. Sideband cooling. A laser excites transitions in which the vibrational quantum number of a confined ion falls by 1 (or a higher integer). Spontaneous transitions bring the ion's internal state back to the ground state, with the vibrational quantum number changing by ± 1 or 0. On average the vibrational quantum number is reduced, until the vibrational ground state is reached. The internal ground and excited states are $|g\rangle$ and $|e\rangle$, and the figure shows the different vibrational levels spread out horizontally for clarity. When both $|g\rangle$ and $|e\rangle$ are long lived (for example they may be the computational basis states: cf Figures 5 and 8), the $|g,n\rangle \rightarrow |e,n-1\rangle$ transition is driven by a π pulse, ($U^1(0)$ operator), and the spontaneous transition is a Raman transition via an unstable excited state (optical pumping). Note that such experimental techniques are identical to those required for information processing and error correction. To cool a crystal of several ions, it is sufficient for the laser to interact with only one ion since the Coulomb coupling between ions causes rapid thermalisation of their motional state. However, the coupling between different normal modes is weaker, so these may need to be cooled separately by tuning the laser to the various normal mode sideband frequencies

of (24) is the thermal distribution

$$P_n = (1-s)s^n\,, \tag{26}$$

where s is the Boltzmann factor $s = \exp(-\hbar\omega_z/k_{\rm B}T)$, and the probability distribution has been normalised.

At sufficiently low temperatures, all but the lowest energy levels can be ignored in (24). By using $\boldsymbol{k}\cdot\boldsymbol{R} = \eta(\hat{a}^\dagger + \hat{a})$, where $\hat{a}|\psi_n\rangle = \sqrt{n}|\psi_{n-1}\rangle$, and expanding in powers of the Lamb–Dicke parameter, it is a simple matter to obtain

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\langle H\rangle \simeq I\sigma_0\frac{E_{\rm R}}{\hbar\omega_{\rm L}}\{\langle n\rangle\,[g(\omega_{\rm L}-\omega_z)-g(\omega_{\rm L}+\omega_z)]\\ +[g(\omega_{\rm L})+g(\omega_{\rm L}-\omega_z)]\}\,,\end{aligned} \tag{27}$$

where $\langle n\rangle = \sum nP_n$ is the ion's mean vibrational quantum number. Now assume $\Gamma \ll \omega_z$ [inequality (23)] and let the incident radiation be tuned to the first sideband below resonance, $\omega_{\rm L} = \omega_0 - \omega_z$, then the cooling limit $\mathrm{d}\langle H\rangle/\mathrm{d}t = 0$ leads to a mean vibrational quantum number [52, 53]:

$$\langle n\rangle \simeq \frac{5\Gamma^2}{16\omega_z^2}\,. \tag{28}$$

Note that since $\langle n\rangle$ is proportional to $(\Gamma/\omega_z)^2$, the experimental constraint (23) will ensure achievement of the quantum limit $\langle n\rangle \ll 1$. This also justifies our ignoring higher energy levels in deducing (28).

The above derivation assumed a single direction of propagation for the cooling laser, which will result only in cooling along one direction, so our calculation has been one dimensional. Taking into account the fact that spontaneous photons are emitted into all directions, we find they do not heat any given dimension quite as much as we assumed, and the factor 5 in (28) is replaced by $(1+4\alpha)$ where $\alpha \simeq 2/5$ depends on the dipole radiation emission pattern [54]. However, this corresponds to an experiment in which the motion in the other dimensions is heated, which we wish to avoid. To cool all three dimensions, one can either introduce three laser beams, or use a single beam propagating at an oblique angle to all the principle axes of the trapping potential, and tune it separately to resonance with the three sideband frequencies $\omega_{\rm L} - \omega_{x,y,z}$. For this one must have all three frequencies distinct, i.e. $\omega_x \neq \omega_y$.

It is commonly imagined that sideband cooling is not possible if the recoil energy is greater than the phonon energy $\hbar\omega_z$, since then the cooling that results from photon absorption is undone by the recoil from photon emission, and $\mathrm{d}\langle H\rangle/\mathrm{d}t > 0$. However, one can always tune to the next lower sideband, $\omega_{\rm L} = \omega_0 - 2\omega_z$, and good cooling is regained, as a thorough analysis of (24) will show. Therefore it is not necessary to be well into the Lamb–Dicke regime in order to attain the quantum limit by sideband cooling.

Note also that both (27) and (28) are significant in order to find the minimum temperature one will obtain in the lab. This is because there will always be heating mechanisms present, such as a coupling between the stored ions and thermal voltages in the electrodes (see Sect. 7), so it is the cooling *rate*, (27), not just the minimum possible temperature, which is important. This has been emphasised by Eschner et al. [38] who propose a subtle variation on sideband cooling. Their method uses repeated measurements of the motional state by looking for resonance fluorescence on a sideband. The measurement probes the high-energy part of the distribution of the ion's population amongst the quantum states of the trap. The observation of no fluorescence, i.e. a null detection, implies that the ion's state has collapsed onto the relatively lower energy part of its initial vibrational distribution. Further such measurements provide opportunities for further cooling. On any given application of this method, the measurement may heat or cool the ion, but cooling is more likely and one knows when it has occurred. This idea is reminiscent of forced evaporative cooling (see Sect. 4.3), only here it is applied to the probability distribution of a single confined particle rather than an ensemble.

4.2 Sisyphus cooling

The constraint (23) means that sideband cooling will either be slow and therefore not compete well with heating processes, or will require the use of Raman transitions. We can avoid $\Gamma \ll \omega_z$ and nevertheless use laser cooling to get close to the quantum regime, by the use of 'Sisyphus' cooling [55, 56]. This makes use of optical pumping and optical dipole forces (forces associated with a position-dependent ac Stark shift of the atomic energy levels) in a laser standing wave, on an atom with at least three internal states. When the dipole force is caused by a position-dependent polarisation of the standing wave, the cooling is referred to as 'polarisation gradient cooling'. Theoretical analyses [55–57] have so far concluded that the lowest temperatures attainable by this method correspond to a mean vibrational quantum number $\langle n\rangle \simeq 1$, i.e. just on the border of the regime we require. However, the cooling rate is important as well as the theoretical minimum temperature, and for this reason Sisyphus cooling may be attractive for cooling a whole string of ions [58], as required for the information processor, since it is relatively fast. This would form the precooling, which usually is necessary in any case before sideband cooling, or something similar, can be applied to get well into the quantum regime.

4.3 Statistical mechanical cooling methods

So far, all the cooling techniques described have been based on laser cooling. However, for trapped neutral atoms the technique of forced evaporative cooling has been shown to be extremely powerful, enabling the temperature in a weakly interacting atomic vapour to be brought well into the quantum regime of a trap, which for a cloud of bosons leads to Bose Einstein condensation [39].

In forced evaporative cooling, one starts with a large number of trapped particles in thermal equilibrium. Those of higher energy are forced to leave the trap, and those remaining rethermalise towards a lower equilibrium temperature. The technique relies on an ability to remove selectively particles of higher than average energy. One way to do this is to reduce the depth of the trap, allowing the faster particles to fly out. Clearly this approach will work only if the thermal energy is located more in some particles than in others, which is true for a gas of weakly interacting particles, but not for a crystalised system such as a cold string of trapped ions. However, evaporation may be useful in an ion trap as a first stage of cooling, to bring about crystalisation. Also, it is conceivable that a Bose condensate of neutral atoms may one day be sufficiently easy

to produce in the vicinity of an ion trap that it may be used as a cold reservoir to cool the ions through collisions. The use of one species to cool another is referred to as 'sympathetic cooling'.

5 rf requirements

We now turn to the design of the ion trap itself. The electrode structure of the trap consists of a two-dimensional rf quadrupole plus an axial static potential. Concentrating on the two-dimensional quadrupole, consider first the most simple case, in which the point in the centre of the electrode structure remains at zero potential, and we omit any axial confinement. The potential on one pair of diagonally opposed electrodes is $(U-V\cos\Omega_V t)/2$, and that on the other pair has equal magnitude and opposite sign to this. Here Ω_V is the frequency of the applied voltage, the subscript is necessary to distinguish it from the Rabi frequency of a driven atomic transition introduced in previous sections. The potential as a function of position in the x-y plane is $\phi(x,y,t) = (U-V\cos\Omega_V t)(x^2-y^2)/2r_0^2$ where r_0 is a measure of the electrode separation.[5] For the case of cylindrical electrodes, r_0 is the distance from the axis to the surface of the electrodes [59]. The trapping effect in the radial direction is stable as long as Ω_V is not too small, and is strong as long as Ω_V is not too large. This may be parametrised in terms of the standard parameters

$$a = \frac{4eU}{Mr_0^2\Omega_V^2}, \tag{29}$$

$$q = \frac{2eV}{Mr_0^2\Omega_V^2}, \tag{30}$$

where e is the charge on a trapped ion. For present purposes, a zero dc potential difference $U = 0$ may be used, so $a = 0$. The radial confinement is then stable as long as q is less than about 0.9 [28, 32]. The radial micromotion has a velocity amplitude of $q\Omega_V\varrho/2$ for an ion at average distance ϱ from the z axis. The average motion on a time scale slow compared to $1/\Omega_V$, the so-called secular motion, can be modelled in terms of the pseudopotential $\frac{1}{2}M\omega_r^2(x^2+y^2)/e$ with radial vibrational frequency

$$\omega_r = \sqrt{a^2+q^2/2}\,\frac{\Omega_V}{2} = \frac{q\Omega_V}{2\sqrt{2}} \quad (a=0). \tag{31}$$

Choosing $q = 1/\sqrt{2}$ so as to be comfortably in the zone of stability of the trap, we obtain $\omega_r = \Omega_V/4$. From this the Lamb–Dicke parameter for the radial confinement is obtained as

$$\eta_r = \left(\frac{2\sqrt{2}\,E_R k^2 r_0^2}{eV}\right)^{1/4}, \tag{32}$$

where k is the wavevector and E_R the recoil energy as defined in (19), and we have neglected the $\cos(\theta)$ term for simplicity. The significance of (32) is that, for a given ion and wavevector, the Lamb–Dicke parameter of the radial confinement is dictated primarily by the choice of electrode size (r_0) and rf voltage amplitude V. The required rf frequency Ω_V is dictated by V/r_0^2 through (30) and the stability condition $q \simeq 1/\sqrt{2}$.

Recall from the discussion of the switching rate, Sect. 3.4, that we want the Lamb–Dicke parameter for the *axial* motion to be around 1, assuming there is only a small number of ions in the trap. We also want the ions to adopt the shape of a linear string, so the radial confinement must be tighter than the axial confinement [see (11), (12)]. Taken together these two considerations imply that the Lamb–Dicke parameter for the *radial* motion should be much less than 1.

Let us now add to the linear trap an axial dc potential, so that the ions are confined in all three dimensions, and with no axial micromotion. The most obvious way to do this is to add positively charged electrodes to either end of the linear trap, but this introduces a difficulty in correctly balancing the rf potential so that there is no residual axial rf component. An ingenious way around this is to split the linear electrodes of the radial quadrupole field and impose a potential difference between their two ends, as described in [59]; see Fig. 1. In either case, the dc potential near the centre of the trap will take the form of a harmonic saddle-point potential

$$\phi_{dc}(x,y,z) = \frac{U_z}{z_0^2}\left[z^2 - \tfrac{1}{2}\left(x^2+y^2\right)\right], \tag{33}$$

where U_z is the potential on each electrode, and z_0 is a parameter that is a measure of the electrode separation (its exact value depending on the geometry). From this equation we obtain the vibrational frequency for the axial harmonic motion of a trapped ion:

$$\omega_z = \sqrt{\frac{2eU_z}{Mz_0^2}}. \tag{34}$$

This is also the frequency of the lowest mode of vibration of a string of trapped ions (centre-of-mass mode), as discussed in Sect. 3.1. The Lamb–Dicke parameter of the axial confinement is

$$\eta_z = \left(\frac{E_R k^2 z_0^2}{4eU_z}\right)^{1/4}. \tag{35}$$

Owing to Earnshaw's theorem, it is impossible to apply an axial dc potential without influencing the radial confinement. The dc potential $\phi_{dc}(x,y,z)$ has the effect of expelling the ions in the radial direction. In the presence of both static axial and fluctuating radial electric potentials, the secular (i.e. slow) radial motion is still harmonic, but the vibrational frequency is no longer ω_r but [32, 59]

$$\omega_r' = \left(\omega_r^2 - \tfrac{1}{2}\omega_z^2\right)^{1/2}. \tag{36}$$

However, as long as $\omega_r \gg \omega_z$, which is the case we are interested in, then $\omega_r' \simeq \omega_r$, so the previous discussion of the radial confinement remains approximately valid, and in particular the stability condition $q <\sim 0.9$ is not greatly changed. The depth of the trap (and hence the ease of catching ions) is given approximately by the smaller of eU_z and $eV/11$.

[5] In practice it is advantageous to avoid exact cylindrical symmetry in order to have all three vibrational frequencies distinct, but this will unnecessarily complicate the present discussion.

Table 1. List of candidate ions for information processing. Only singly-charged ions are considered, although some ions of higher charge may also be interesting. For each element, only the most abundant isotope, and those having non-zero nuclear spin are shown. Unstable isotopes are not shown, although most elements in the list (all but Mg and Li) have further isotopes of half-life longer than one week. Thalium and indium are omitted since their ground states have $J = 0$ and so lack hyperfine structure. The hyperfine splittings are for the ground state in all but helium-like lithium; they are taken from G. Werth in [30], and from [76]. The S–D wavelength is only shown when the D level lies below the P level. The recoil energy is based on the S–P wavelength. For Li the S,P,D labels do not apply; the transitions are from the metastable triplet state. Note that the fine structure splitting in the excited state (not shown) is also relevant to the Raman transition rate (see text)

Element, isotopes	Natural abundance (%)	Nuclear spin ($\hbar$)	Hyperfine splitting (GHz)	λ S–P (nm)	λ S–D (nm)	Recoil energy (kHz)
Be 9	100	3/2	1.25001767	313		226
Mg 24	79	0		280		106
25	10	5/2	1.7887631			
Ca 40	97	0		397	730	30
43	0.14	7/2	3.25560829			
87	7	9/2	5.00236835			
Sr 88	83	0		422	674	12.7
135	6.6	3/2	7.18334024			
137	11	3/2	8.03774167			
Ba 138	72	0		493	1760	5.94
199	17	1/2	40.507348			
201	13	3/2	30.16			
Hg 202	30	0		194	282	26.6
171	14	1/2	12.6428121			
173	16	5/2	10.4917202			
Yb 174	32	0		369	411	8.42
6	7.5	1	3.0018			
Li* 7	92.5	3/2	11.8900	539		94.7

6 Candidate ions

Table 1 gives a list of ions suitable for information processing. The list consists of ions whose electronic structure is sufficiently simple to allow laser cooling without the need for too many different laser frequencies. The list is not intended to be exhaustive, but contains most ions that have been laser cooled in the laboratory.

For information processing, a large recoil energy is attractive from the point of view of allowing a faster switching rate (22), but makes the Lamb–Dicke regime harder to achieve [see (32), (35)]. The choice of rf rather than optical transitions for information processing appears so advantageous as to be forced upon us. Since we require at least three long-lived low-lying states of the ion (the states $|0\rangle$, $|1\rangle$ and $|\mathrm{aux}\rangle$), this implies that the existence of hyperfine structure (i.e. a non-zero nuclear spin isotope), although it complicates the cooling process, may be advantageous. Indeed, for alkali-like ions (such as singly charged ions from group 2 of the periodic table), the electronic ground state is $J = 1/2$, so if the nuclear spin is zero there are only two long-lived internal states (the Zeeman components $|J,M\rangle = |1/2,\pm 1/2\rangle$), which is not sufficient. Having said this, we are not necessarily forced to choose $|\mathrm{aux}\rangle$ to be a third internal state of the ion (i.e. the choice implied by Fig. 5). One could make use of the second normal mode of oscillation of the ion string instead, choosing $|\mathrm{aux},0\rangle = |F_1,M_1\rangle \otimes |0,1,0,\ldots\rangle$ [cf equation (13)]. Recently, Monroe et al. [46] have shown that the need for an auxilliary level can be avoided altogether.

The other major consideration is the difficulty in generating the light required for cooling and information processing.

Examining Fig. 7 and Table 1, we see that ^{9}Be is an attractive choice, in that it allows the fastest switching rate and it requires only one laser wavelength for cooling, and the hyperfine splitting frequency of 1.25 GHz is accesible to electro-optic modulators. However, the wavelength of 313 nm requires the use of a dye laser (frequency doubled) which is disadvantageous. The next most promising candidate appears to be ^{43}Ca. It requires two laser wavelengths for cooling, 397 (or 393) nm and 866 (or 850) nm (Fig. 8), but both can be produced by diode lasers (one frequency doubled), which makes this ion very attractive (strontium has similar advantages). Diode lasers can be made very stable in both frequency and power. If more laser power is needed than is possible with diode lasers, then a titanium–sapphire laser can be used, which is also advantageous compared with

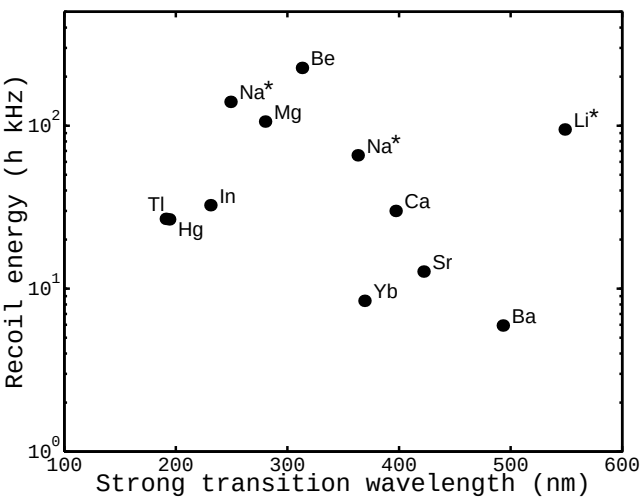

Fig. 7. The recoil energies and main (typically S–P) transition wavelengths for ions that may be amenable to quantum information processing (cf Table 1). A high recoil energy is advantageous for a high switching rate, but tends to be associated with a short wavelength. A rough rule is that the shorter the wavelength, the more complicated and therefore less stable is the laser system. The starred symbols are singly ionised ions from group 1, in which a metastable manifold is used for computing, making them unattractive in the long term

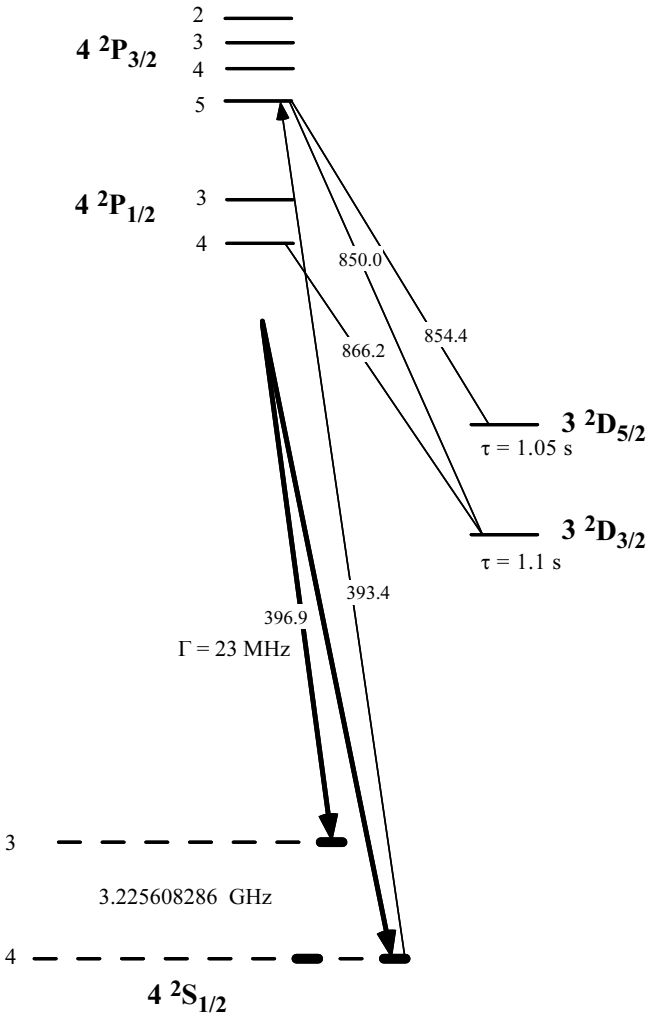

Fig. 8. Low-lying energy levels of the $^{43}Ca^{+}$ ion, of electronic structure $1s^2 2s^2 2p^6 3s^2 3p^6 nl$. The hyperfine structure of the D states is omitted to keep the diagram uncluttered. The other levels are labelled by the total angular momentum quantum number F (nuclear spin $I = 7/2$). The Zeeman sublevels are shown for the ground state hyperfine manifold, and a possible choice of computational and auxilliary levels is shown by the thickened sublevels. An example Raman transition is shown, for use both in sideband cooling and for the operation of quantum gates. The $3D$ levels are shown because they are involved in the state preparation (including cooling) and measurement, both of which are an integral part of a complete quantum 'computation'. The level separations and lifetimes are taken from references [62]

dye lasers. The hyperfine splitting of 3.26 GHz is accessible to electro-optic modulators, though less easily than the smaller splitting in beryllium. The obvious difficulty in working with ^{43}Ca is that it is a rare isotope, having a natural abundance of only 0.14% or 1 part in 700, making an isotopically enriched sample that much more expensive. However, one could carry out preliminary experiments using the 97% abundant ^{40}Ca.

For a group-2 ion, the internal states required for information processing, discussed in Sect. 3.2 and illustrated in Fig. 5, will be taken from the ground state hyperfine manifold. For ^{43}Ca, for example, one might take $|F_1 = 4, M_1 = 4\rangle$, $|F_2 = 3, M_2 = 3\rangle$ and $|F_{\text{aux}} = 4, M_{\text{aux}} = 2\rangle$. This choice is highlighted in Fig. 8. The degeneracy between the first and auxilliary levels is lifted by an imposed magnetic field of order 0.1 mT.

6.1 *Example: the $^{43}Ca^{+}$ ion*

To estimate laser power requirements, we will calculate the intensity required to saturate the $4S_{1/2}$–$4P_{3/2}$ transition in Ca^{+} (for laser cooling purposes) and that required for Raman transitions in the ground state via a quasi-resonance with this transition (for information processing purposes). With a two-level model for the allowed electric dipole transition, the saturation intensity (defined as the intensity giving a Rabi frequency equal to the FWHM (full width, half maximum) linewidth Γ divided by $\sqrt{2}$) is $I_S = 4\pi^2\hbar c\Gamma/6\lambda^3 = 48\,\text{mW/cm}^2$ (using $\Gamma = 2\pi \times 23\,\text{MHz}$, $\lambda = 397\,\text{nm}$). To initiate laser cooling, this intensity must be available in a laser beam wide enough to intersect a significant proportion of a 'hot' ion's trajectory in the trap. With a beam diameter of 1 mm, the required laser power is of order 0.5 mW, which is a large overestimate in practice.

Raman transitions from $|0\rangle$ to $|1\rangle$ via a near-resonance with an excited state $|e\rangle$ can be modelled as transitions in an effective two-level system, in which the effective Rabi frequency of the Raman transition is

$$\Omega_{\text{eff}} = \frac{\Omega_0 \Omega_1}{2\Delta}\,, \tag{37}$$

where Ω_0 and Ω_1 are the Rabi frequencies of the single-photon transitions from levels $|0\rangle$ and $|1\rangle$ to $|e\rangle$, and $\Delta \gg \Omega_0, \Omega_1$ is the detuning from resonance of both of these transitions. If level $|e\rangle$ decays only to levels $|0\rangle$ and $|1\rangle$, the single photon Rabi frequencies can be obtained from the laser intensity I and the linewidth Γ of the excited state, leading to $\Omega_{\text{eff}} \simeq I\Gamma^2/8I_S\Delta$. During a Raman transition, the average population of the excited state $|e\rangle$ is $\sim \Omega_0^2/4\Delta^2$, and to produce, for example, a 2π pulse, the pulse duration is $2\pi/\Omega_{\text{eff}}$. Therefore the probability of an unwanted spontaneous emission process such an operation is

$$p_{\text{em}} = \frac{\pi\Gamma}{\Delta}\,. \tag{38}$$

An interesting possibility, which has not yet been tried in an ion trap, is to use the Argon ion laser line at 488 nm to drive Raman transitions. In this case, we have $\Delta \simeq 6 \times 10^6\Gamma$, so $p_{\text{em}} \simeq 5 \times 10^{-7}$, allowing a million computing operations before spontaneous emission is a problem. The laser intensity required to obtain $\hbar\Omega_{\text{eff}} = E_R$, if we assume equation (37) applies, is then

$$I = 8I_S \frac{\Delta}{\Gamma}\frac{E_R}{\hbar\Gamma} \simeq 2.9 \times 10^9\,\text{W/m}^2\,. \tag{39}$$

To address a single ion, the laser is focussed to a tight spot of diameter $\sim 10\,\mu\text{m}$, so the required power is modest, of order 0.3 W. Although the main features of this argument are correct, in fact the situation is more complicated since the ion is not really the three-level system implied by (37). A complete analysis must take account of all the fine and hyperfine structure.

Let the fine structure splitting in the excited state be ΔE_{fs}, and let the hyperfine structure in the excited state be of order ΔE_{hfs}. To calculate the Raman transition's effective Rabi frequency Ω_{eff}, we must replace (37) by a sum over the transition amplitudes for all available transition routes. Destructive interference between these amplitudes can result in a much reduced

effective Rabi frequency, as described in [60]. Another way of understanding this is to note that if the detuning $\Delta \gg \Delta E_{\rm hfs}/\hbar$, then the nuclear and electron angular momenta I and J may be considered decoupled, to a first approximation, and the electric dipole transition matrix element only couples to J, so only $\Delta M_I = 0$ transitions are allowed. A similar argument applies to the LS coupling when $\Delta \gg \Delta E_{\rm fs}/\hbar$, in which case only $\Delta M_S = 0$ transitions are allowed. Unfortunately, this means that for $L = 0$ ground states, which is the usual case, Raman transitions from one ground state hyperfine level to another cannot be driven if the detuning Δ is increased much beyond $\Delta E_{\rm fs}/\hbar$ [61]. We conclude that a large fine structure is advantageous. For calcium, $\Delta E_{\rm fs}/\hbar \simeq 2\pi \times 7700\,{\rm GHz}$, and choosing Δ equal to this splitting gives $p_{\rm em} \simeq 10^{-5}$.

To confine Ca to the Lamb–Dicke regime of the 393 nm radiation, we require $\omega_z = E_{\rm R}/\hbar \simeq 2\pi \times 29\,{\rm kHz}$. This is a reasonable choice for information processing if we bear in mind the remarks made in Sect. 3.4 about off-resonant transitions. Choosing an axial electrode separation of ~ 10 mm, the voltage required on the axial electrodes is $U_z \sim 0.74$ volts. This is a low value compared to that typically used in ion traps. The fact that Lamb–Dicke confinement is achieved even with such a relatively weak trap comes from the property that the ions are cooled to the ground state of the axial motion, which here corresponds to a temperature small compared to the recoil limit $E_{\rm R}/k_{\rm B} \simeq 1.4\,\mu{\rm K}$. The low value of U_z shows that contact potentials will certainly be a problem, and one must be able to compensate them by seperately controlling the voltage on each electrode. To make the radial confinement ten times stronger than this axial confinement, we require an alternating voltage on the radial quadrupole electrodes of frequency $\Omega_V \simeq 2\pi \times 1.2\,{\rm MHz}$ [see (31)] and amplitude $V \simeq 9$ volts [see (30)], assuming a distance of ~ 1 mm from the axis to the radial electrode surfaces. In practice there is no great difficulty in generating voltages up to several hundred volts at frequencies in the region of tens of MHz. Since there is no reason not to have strong radial confinement, and some advantages, such as reduced collisional heating and loss, these larger parameter values should be used.

7 Performance limitations

Having begun in Sect. 3.2 with an idealised treatment, in which we assumed operations could be carried out in an ion trap with arbitrary precision, we were more realistic in Sect. 6. It now remains to discuss the limitations on the performance of the ion trap system for information processing purposes.

Two important figures of merit for a quantum information processor are the number of stored qubits, which so far in this paper has been the number of trapped ions N, and the number Q of elementary operations that can be carried out before dissipation or decoherence causes a significant loss of quantum information. To a first approximation, we may quantify dissipation or decoherence by a simple rate $\Gamma_{\rm d}$, in which case $Q = R/\Gamma_{\rm d}$, where the switching rate R is given in Sect. 3.4. If we model decoherence as if each ion were independently coupled to a thermal reservoir, leading to a phase decoherence rate γ for any individual ion, then we must take $\Gamma_{\rm d} = N\gamma$ since the quantum computation is likely to produce entangled states in which the off-diagonal elements of the density matrix decay at this enhanced rate [12, 13, 16, 17]. The origin of the factor N here is quite simple to understand. In a state such as $a|000\rangle + b|111\rangle$ the phase relationship between a and b is lost if *any one* of the three qubits dephases, so the rate for dephasing of the joint state is three times the single-qubit rate if the qubits dephase independently. It is just like in classical probability where if something can happen in N equally likely different ways, it is N times more likely to happen. When throwing a dice, we are three times more likely to get an even number than to get a one. However, decoherence of a many-ion state in an ion trap is not yet sufficiently well understood to tell whether such a model applies [63]. Two possible thermal reservoirs affecting the ion trap are electrical resistance in the electrodes and thermal radiation.

The major problems in an ion trap are spontaneous transitions in the vibrational motion, i.e. heating (a random walk up and down the ladder of vibrational energy levels), thermal radiation (driving internal rf transitions in the ions), and experimental instabilities such as in the laser beam power, rf voltages and mechanical vibrations, and fluctuating external magnetic fields [44]. The instabilities contribute to the heating and also imply that a laser pulse is never of exactly the right frequency and duration to produce the intended quantum gate. A 'decoherence rate' of a few kHz was quoted by Monroe et al. [44], consistent with the heating rate of 1000 vibrational quanta per second quoted in their earlier work [37]. With the switching rate of order 20 kHz, they obtained $Q \simeq 10$ with $N = 1$. A heating rate of 6 quanta per second was reported by Diedrich et al. [36].

A useful model of the motion of a trapped ion is a series LC circuit shunted by the capacitance of the trap electrodes [63, 64]. The inductance in the model is given by $l \simeq Mz_0^2/Ne^2$ where z_0 is of the order of the axial electrode separation. A resistance r is due to losses in the electrodes and other conductors in the circuit. This resistance both damps and heats the ionic motion with time constant l/r, leading to a heating rate in vibrational quanta per second [63]:

$$\Gamma_{\rm heat} \simeq \frac{r}{l}\frac{k_{\rm B}T}{\hbar\omega_z} \simeq \frac{rNe^2k_{\rm B}T}{Mz_0^2\hbar\omega_z}\,. \tag{40}$$

For example, substituting the parameters from Sect. 6.1, and using $r = 0.1$ ohm, $T = 300\,{\rm K}$, we obtain $\Gamma_{\rm heat} = 5\,{\rm s}^{-1}$. It should be borne in mind that one can only consider (40) to apply once other sources of electrical noise, such as rf pickup, have been reduced sufficiently, so one cannot hope to improve the performance merely by increasing the electrode separation z_0 and voltage U_z.

It is a simple matter to combine (20), (34), (37), and (40) in order to obtain $Q = R/\Gamma_{\rm heat}$ as a function of the experimental parameters. However, this does not bring much insight and it is better to think in terms of the switching rate and decoherence rate. If we take $N = 10$ ions with $\omega_z = E_{\rm R}/\hbar = 2\pi \times 29\,{\rm kHz}$, as suggested in the previous section, the switching rate is about 1 kHz, and $Q \sim 200$ where we use the value just quoted for $\Gamma_{\rm heat}$. These parameters indicate what will probably be achievable in the next few years.

It is not hard to show that the influence of spontaneous emission of photons by the ions in the trap is much less important than the severe experimental problems just mentioned. Spontaneous emission takes place during the application of a laser pulse, because of the unavoidable weak excitation of an excited state of the relevant ion, as noted in the previous

section. (It was already remarked that spontaneous emission between laser pulses is negligible, owing to the adoption of the ground-state hyperfine manifold for computing.) With the probability $p_{\rm em}$ from (38), the number of operations that can be carried out before spontaneous emission plays a significant role is $Q \simeq 1/p_{\rm em}$, which can be of the order of 10^5, as remarked after (38).

The conclusion is that for the moment the limitations of the ion trap are associated with the vibrational degrees of freedom, and with experimental instabilities. It is here that experimental and theoretical work must concentrate if progress is to be made. It remains misleading at present to talk of quantum 'computations' taking place in the lab.

7.1 Error correction

Although it is important to build an information processor with as much precision and stability as possible, in the longer term the aim of significant computations is almost certainly unrealisable without something that goes beyond such 'passive' stabilisation. The inherent instability of quantum computing was stressed in Sect. 1. It was initially thought that anything like active stabilisation of a quantum computer would be impossible, since it would rely on a means of monitoring the quantum state of the computer, which would irreversibly destroy the computation. However, the union of information theory with quantum mechanics has lead to another powerful concept, that of quantum error correction [65–67]. The essential idea is that K qubits of quantum information in the quantum computer can be stored ('encoded') in a carefully chosen way among $N > K$ two-state systems. The simplest case is to store a single qubit in three two-state systems, as follows:

$$\begin{aligned} |0_{\rm L}\rangle &\equiv (|000\rangle + |011\rangle + |101\rangle + |110\rangle)/2 \\ |1_{\rm L}\rangle &\equiv (|111\rangle + |100\rangle + |010\rangle + |001\rangle)/2 \end{aligned} \tag{41}$$

All we have done here is written down two orthogonal states, calling them $|0_{\rm L}\rangle$ and $|1_{\rm L}\rangle$. With two states we have a single *logical* qubit, $K = 1$, but it is stored physically in $N = 3$ separate ions. A general state of the logical qubit is just $a|0_{\rm L}\rangle + b|1_{\rm L}\rangle$. The logical qubit inhabits a two-dimensional subspace of the total Hilbert space of eight dimensions.

The computation is carried out in the specially chosen 2^K-dimensional subspace of the total Hilbert space (2^N dimensions) of the enlarged computer. For example, to apply the logical state rotation $|1_{\rm L}\rangle\langle 0_{\rm L}| - |0_{\rm L}\rangle\langle 1_{\rm L}|$, we must apply $V^1(-\pi/2) = |1\rangle\langle 0| - |0\rangle\langle 1|$ seperately to each ion.

The encoding is chosen so that the most likely errors cause the computer's state to go out of the special subspace (whereas computing operations keep the state within the special subspace). One can detect such departures, and force the computation back on track, *without* corrupting the stored quantum information, by making well-chosen joint measurements on the three qubits. The exact corrective procedure is deduced from the error correction methods which are a central part of classical information theory. This is a rather subtle and beautiful link between classical information theory and quantum mechanics. Unfortunately it would take too long to describe it fully here. A thorough discussion which does not assume much prior knowledge is given in [66]. The above 'simplest case', (41), is fully analysed there.

The encoding (41) is designed to reduce phase noise, i.e. a noise source that causes the state of each ion to precess a random amount, so that $|0_{\rm L}\rangle \to |000\rangle + {\rm e}^{{\rm i}\varepsilon(\theta_2+\theta_3)}|011\rangle + {\rm e}^{{\rm i}\varepsilon(\theta_1+\theta_3)}|101\rangle + {\rm e}^{{\rm i}\varepsilon(\theta_1+\theta_2)}|110\rangle$ where the θ_i are random phases and $\varepsilon \ll 1$ is a small parameter indicating the level of phase noise. A similar expression holds for $|1_{\rm L}\rangle$. The error term is linear in ε, since ${\rm e}^{{\rm i}\varepsilon\theta} \simeq 1 + {\rm i}\varepsilon\theta$. To understand the correction method, re-write the quantum state in terms of the new basis $|\bar{0}\rangle \equiv (|0\rangle + |1\rangle)/\sqrt{2}$, $|\bar{1}\rangle \equiv (|0\rangle - |1\rangle)/\sqrt{2}$. Then one finds $|0_{\rm L}\rangle = |\bar{0}\bar{0}\bar{0}\rangle + |\bar{1}\bar{1}\bar{1}\rangle$, $|1_{\rm L}\rangle = |\bar{0}\bar{0}\bar{0}\rangle - |\bar{1}\bar{1}\bar{1}\rangle$. Now, phase noise introduces erroneous terms such as $|\bar{1}\bar{0}\bar{0}\rangle$ into the quantum state. We can detect the presence of such terms, without corrupting the logical qubit, by making measurements that ask not 'what is the state of the first ion?' but rather 'are the first two ions in the same state?'. Such a measurement is performed by making use of a fourth ion, which is used as a 'check bit'. It is prepared in the state $|\bar{0}\rangle$, and then undergoes an XOR operation with the first ion, then the second. Its state is finally measured in the basis $|\bar{0}\rangle$, $|\bar{1}\rangle$, which in practice is done by first rotating the state by $V^{1/2}(-\pi/2)$, (15), then measuring in the computational basis $|0\rangle$, $|1\rangle$. The other significant measurement in this case is 'are the first and third ions in the same state?', for which another check bit ion may be used, or else the first one re-used.

After the check bit measurements have been completed, one has gathered some information about the noise in the system, but none about the original logical state (i.e. the values of the coefficients a and b). In addition, the measurements force the system into a state that is either noise-free, or related to the noise-free state by a rotation $\sigma_z \equiv |0\rangle\langle 0| - |1\rangle\langle 1|$ of one ion. In the latter case, the measurement results indicate which ion must be rotated, so the state can be corrected by applying $V^1(0)V^1(\pi/2) = -{\rm i}\sigma_z$ to the relevant ion.

The above sequence of logical operations and measurements is not guaranteed to correct the state, but yields in this case a final state in which the noise terms are second order rather than first order in ε [66]. Thus, we have gained a less noisy final state as long as ε is small. Even this simplest case is quite impressive, suppressing the noise by two orders of magnitude when $\varepsilon \simeq 0.01$. However, more advanced encodings achieve an even more powerful stabilisation without greatly increasing the complexity of the corrective procedure. In addition, the more general methods are not restricted to correction of phase noise only. A general quantum error correcting code (QECC) can be paramatrised in terms of the number K of information qubits 'encoded', the number N of physical two-state systems used, and the degree of noise suppression achieved. A QECC is called a "t-error-correcting" code if it can be used to restore the encoded state after up to t of the N systems have had arbitrary errors (i.e. state changes and entanglement of any type and size). In practice, noise will typically cause small errors in all the systems, rather than large errors in a few. If the erroneous terms in the density matrix are of order ε before correction, then a t-error-correcting code successfully corrects all terms up to order t in ε, so the noise is reduced to $O(\varepsilon^{t+1})$ [65–67]. The power of the method arises from the following, at first sight astonishing, result: the 'scale-up' N/K required to implement QEC remains bounded as $t \to \infty$.

Quantum error correction can usefully be compared and contrasted with the more simple 'watchdog' or 'quantum Zeno effect' idea which preceded it [68]. If a system is repeatedly measured, it will repeatedly collapse onto the measurement basis. This can be used to suppress any tendency of a system to precess away from the initially measured state, and such a suppression is called the watchdog or quantum Zeno effect. In the present context, the precession is caused by some 'error' Hamiltonian H_e. The Zeno effect occurs if many measurements can be applied within a time t sufficiently small that $1-|\langle\phi|\exp(-iH_e t/\hbar)|\phi\rangle|^2 \propto t^2$. In such a case, if m measurements are made during a time interval Δt, the probability the system is found in a state other than its initial state $|\phi\rangle$ is proportional to $m(\Delta t/m)^2$, which tends to zero as $m \to \infty$. For our present purposes, we do not wish to preserve the quantum computer in a particular state, but rather in a particular part of Hilbert space, which can be achieved with a generalised version of the Zeno effect [69]. There is thus some similarity with quantum error correction, only now if the system leaves the right part of Hilbert space, there is no corrective procedure. As a result the stabilisation is much less powerful than that of QEC, and is almost certainly useless for a quantum computer. This is because errors in a quantum computer will occur primarily during the action of the logic gates, and it is unlikely that the measurement necessary for a Zeno effect in the special sub-Hilbert space could be made repeatedly during the action of a gate.

By contrast with the Zeno effect, quantum error correction allows a finite error term in the system's density matrix to accumulate, and corrects it afterwards. This is particularly important to the operation of quantum gates. For example, a gate between two qubits involves a four-dimensional logical Hilbert space. To allow error correction, we must ensure that at no point is the whole action of this gate concentrated into a four-dimensional physical Hilbert space. This can be done as follows. Suppose each qubit in the QC is encoded into two physical two-state systems. A gate $U(a,b)$ between two such encoded qubits a,b can then be applied in four steps $U(a,b) = u(a_2,b_1)\cdot u(a_1,b_2)\cdot u(a_2,b_2)\cdot u(a_1,b_1)$, where the operators u are gates between a pair of two-state systems, and $\{a_1,a_2\},\{b_1,b_2\}$ are the sets of two-state systems storing qubits a and b. The important point is that error correction can be applied *between* the four u operations. At no stage is any quantum information stored in a physical Hilbert space only just large enough to hold it, neither is any gate U carried out in a single step. The proper combination of these features so as to allow stabilisation has been dubbed 'fault tolerance' [72–74]. It is found that overall stability can be achieved even when every operation is noisy, including those involved in the corrective procedure. This is under active investigation.

Quantum error correction was initially discussed with a general model of error processes in the quantum computer, in order to show that almost any imaginable error process might in principle be corrected by such techniques [65, 66, 70, 71]. Subsequently, a technique specifically adapted to the vibrational noise in an ion trap was proposed [75]. In this proposal, the Hilbert space is enlarged by making use of four internal states in each ion to store each qubit in the quantum computation. This enables a two-qubit gate to be carried out in four steps as outlined above. Next, we require a method of detecting and correcting the most likely errors in the vibrational state. This is done by using the first ($n=0$) and fourth ($n=3$) vibrational states, $|0,0,0,\ldots\rangle$ and $|3,0,0,\ldots\rangle$, instead of the first two, to store the 'bus' qubit. The vibrational quantum number n is measured whenever it should be 0, by swapping the phonon state with the state of additional ions introduced for the purpose, and probing them. If n is found to be 1, then corrective measures are applied based on the assumption that a single jump upwards from $n=0$ occured, the details are given in [75]. If n is found to be 2 or 4, then corrective measures are applied based on the assumption that a single jump down or up from $n=3$ occured. If n is found to be 0 as it should be, a corrective measure is still required to allow for the difference between such conditional evolution and the unitary evolution without jumps. This procedure enables single jumps up or down the ladder of vibrational levels to be corrected. Since these will be the most likely errors (at a sufficient degree of isolation from the environment), the effect overall is to stabilise the QC. In this case the figure of merit Q is roughly squared (where Q counts the possible number of logical gates U, not subgates u)—a remarkable enhancement.

The above procedure makes allowance for the fact that errors cause the vibrational state to explore a Hilbert space of more than two dimensions, so in the language of quantum information, the bus size is larger than a single qubit, though the bus is still only used to store a single logical qubit. The bus could be made larger still by using higher excitations of the fundamental normal mode, or by using higher-order normal modes. This should allow more powerful error correction, and hence further increases in Q. The basic theory of error correction gives hope that such increases in Q can be dramatic [65, 66, 72].

Error correction should not be regarded as a device merely of interest to quantum computers. Rather, it is a powerful method of enforcing coherent evolution on a quantum system that would otherwise be dissipative. Such a capability may be useful for quite general situations in which stability is important, for example in low-noise electronic circuits, and frequency standards. This may prove to be an area in which quantum information theory has provided a useful tool for other branches of physics.

8 Conclusion

Let us summarise the main avenues for future work involving the ion trap quantum information processor.

One of the basic aims of quantum information theory is to link abstract ideas on the nature of information with the laws of physics. The ion trap provides a means of establishing this link in a complete and concrete way. This will set the theory on a more firm basis.

An important task in the theory of quantum computation is that of identifying efficient multiple-qubit quantum gates. So far, it has been assumed that the efficient gates are those that can be divided into a set of sufficiently few two-qubit gates. However, a system like the ion trap may allow particular unitary transformations to be carried out efficiently without dividing them into many two-qubit operations. An investigation of this should be fruitful.

The principle of operation of the ion trap that we have described made use of various approximations whose influence on long quantum computations has yet to be analysed. For example, any given laser pulse on an ion will involve off-resonant stimulation of transitions other than the specific transition the pulse is designed to drive. Such effects may be unwanted, but their influence is unitary and accurately predictable. It would be interesting to investigate whether these effects can be taken into account in designing the sequence of laser pulses, so that they do not need to be corrected, or whether we are forced to regard them as errors.

Quantum computation will certainly require error correction if it is ever to be useful for computational purposes. The ion trap provides a guide to the specific type of error correction likely to be required in the future. The basic tools of error correction are now fairly well understood, but there is much work to be done in bringing them to bear on the ion trap. In addition, these ideas may offer significant advantages for other uses of the ion trap, such as frequency and mass standards; this should be explored.

Error correction only works once the level of noise in the trap is brought sufficiently low by careful construction and isolation. Experimental ion trap systems must be made much more stable than they are at present before they can take advantage of error correction of multiple errors among many qubits. This is not just a question of technology, but also of a better understanding of the noise processes, especially the influence of electrical noise in the electrodes providing axial confinement. The most immediate experimental challenge is to cool a many-ion crystal to the motional ground state.

Acknowledgements. I would like to acknowledge helpful conversations with R. Thompson, D. Segal, D. Stacey and especially J. Brochard. C. Monroe and D. Wineland provided very useful comments on the manuscript, and D. James helped in finding atomic data. I am supported by the Royal Society.

References

1. A.M. Turing: Proc. Lond. Math. Soc. Ser. 2 **42**, 230 (1936); see also Proc. Lond. Math. Soc. Ser. 2 **43**, 544 (1936)
2. D. Deutsch: Proc. R. Soc. Lond. A **400**, 97–117 (1985)
3. A. Ekert: in *Atomic Physics 14* (ed. by D.J. Wineland, C.E. Wieman, S.J. Smith) (AIP Press, New York 1995), p450. Also A. Ekert, R. Jozsa: Rev. Mod. Phys, **68**, 733 (1996)
4. A. Barenco: Contemp. Phys. **37**, 345 (1996)
5. R. Jozsa, B. Schumacher: J. Mod. Optics **41**, 2343 (1994)
6. B. Schumacher: Phys. Rev. A **51**, 2738–2747 (1995)
7. R. Cleve, D. DiVincenzo: Phys. Rev. A **54**, 2636 (1996)
8. C.H. Bennett: Phys. Today **48**, (October) 24–30 (1995)
9. D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, A. Sanpera: Phys. Rev. Lett. **77**, 2818 (1996)
10. C.H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. Smolin, W.K. Wootters: Phys. Rev. Lett. **76**, 722–725 (1995)
11. C.H. Bennett, D.P. DiVincenzo, J.A. Smolin, W.K. Wootters: Phys. Rev. A **54**, 3825 (1996)
12. H. Carmichael: *An open systems approach to quantum optics*, Lecture Notes in Physics, Vol.18 (Springer, Berlin 1993)
13. A.O. Calderia, A.J. Leggett: Phys. Rev. A **31**, 1059 (1985)
14. W.H. Zurek: Prog. Theor. Phys. **89**, 281–312 (1993); see also W.H. Zurek: Phys. Today, October, 36 (1991)
15. R. Landauer: Philos. Trans. R. Soc. London Ser. A. **353**, 367 (1995)
16. W.G. Unruh: Phys. Rev. A **51**, 992–997 (1995)
17. G.M. Palma, K-A. Suominen, A.K. Ekert: Proc. Roy. Soc. Lond. A **452**, 567–584 (1996)
18. P.W. Shor: In *Proc. 35th Annual Symp. on Foundations of Computer Science*, Santa Fe 1994, IEEE Computer Society Press; revised version quant-ph/9508027
19. S. Haroche, J.-M. Raimond: Phys. Today, August, 51–52, 1996.
20. D.P. DiVincenzo: Phys. Rev. A **51**, 1015–1022 (1995); A. Barenco: Proc. R. Soc. Lond. A **449**, 679–683 (1995); D. Deutsch, A. Barenco, A. Ekert: Proc. R. Soc. Lond. A **449**, 669–677 (1995)
21. N. Margolus: In *Complexity, Entropy and the Physics of Information, Santa Fe Institute Studies in the Sciences of Complexity,* Vol VIII, ed. by W.H. Zurek (Addison-Wesley, Reading, Ma 1990) p.273
22. S. Lloyd: Science **261**, 1569 (1993); Science **263**, 695 (1994)
23. G.P. Berman, G.D. Doolen, D.D. Holm, V.I. Tsifrinovich: Phys. Lett. **193**, 444–450 (1994)
24. D. DiVincenzo: Science **270**, 255 (1995)
25. J.I. Cirac, P. Zoller: Phys. Rev. Lett. **74**, 4091–4094 (1995)
26. T. Pellizzari, S.A. Gardiner, J.I. Cirac, P. Zoller: Phys. Rev. Lett. **75**, 3788 (1995)
27. Q.A. Turchette, C.J. Hood, W. Lange, H. Mabushi, H.J. Kimble: Phys. Rev. Lett. **75**, 4710 (1995)
28. R.C. Thompson: Adv. At. Mol. Opt. Phys. **31**, 63–136 (1993)
29. G. Werth: Contemp. Phys. **26**, 241–256; R. Blatt, P. Gill, R.C. Thompson: J. Mod. Opt **39**, 193–220 (1992)
30. "Trapped charged particles and related fundamental physics", special issue of Physica Scripta T **59**, proceedings of Nobel Symposium, ed. by I. Bergström, C. Carlberg, R. Schuch (1994)
31. W.M. Itano, J.C. Bergquist, J.J. Bollinger, D.J. Wineland: Physica Scripta T **59**, 106 (1995)
32. P.K. Ghosh: *Ion Traps* (Clarendon Press, Oxford 1995)
33. W. Nagourney, J. Sandberg, H. Dehmelt: Phys. Rev. Lett **56**, 2797 (1986); Th. Sauter, W. Neuhauser, R. Blatt, P.E. Toschek: Phys. Rev. Lett **57**, 1696 (1986); J.C. Bergquist, R.G. Hulet, W.M. Itano, D.J. Wineland: Phys. Rev. Lett **57**, 1699 (1986)
34. R. Blatt, P. Zoller: Eur. J. Phys. **9**, 250 (1988); R.J. Cook: in *Progress in Optics XXVIII*, ed by E. Wolf (Elsevier, Amsterdam 1990), p.361
35. D. Liebfried, D.M. Meekhof, B.E. King, C. Monroe, W.M. Itano, D.J. Wineland: Phys. Rev. Lett. **77**, 4281 (1996)
36. F. Diedrich, J.C. Bergquist, W.M. Itano, D.J. Wineland: Phys. Rev. Lett. **62**, 403 (1989)
37. C. Monroe, D.M. Meekhof, B.E. King, S.R. Jefferts, W.M. Itano, D.J. Wineland, P. Gould: Phys. Rev. Lett. **75**, 4011 (1995)
38. J. Eschner, B. Appasamy, P.E. Toschek: Phys. Rev. Lett. **74**, 2435 (1995)
39. M.H. Anderson, J.R. Ensher, M.R. Matthews, C.E. Wieman, E.A. Cornell: Science **269**, 198 (1995); K.B. Davis, M.-O. Mewes, M.R. Andrews, N.J. van Druten, D.S. Durfee, D.M. Kurn, W. Ketterle: Phys. Rev. Lett. **75**, 3969 (1995)
40. P. Verkerk, B. Lounis, C. Salomon, C. Cohen-Tannoudji, J.-Y. Courtois, G. Grynberg: Phys. Rev. Lett. **68**, 3861 (1992); P.S. Jessen, C. Gerz, P.D. Lett, W.D. Phillips, S.L. Rolston, R. J.C. Spreeuw, C.I. Westbrook: Phys. Rev. Lett **69**, 49 (1992); G. Grynberg, B. Lounis, P. Verkerk, J.-Y. Courtois, C. Salomon: Phys. Rev. Lett. **70**, 2249 (1993)
41. J.P. Schiffer: Phys. Rev. Lett. **70**, 818–821 (1993)
42. D.H.E. Dubin: Phys. Rev. Lett. **71**, 2753–2756 (1993)
43. H. Totsuji, J.-L. Barrat: Phys. Rev. Lett. **60**, 2484 (1988)
44. C. Monroe, D.M. Meekhof, B.E. King, W.M. Itano, D.J. Wineland: Phys. Rev. Lett. **75**, 4714 (1995)
45. A. Barenco, C.H. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J.A. Smolin, H. Weinfurter: Phys. Rev. A **52**, 3457 (1995)
46. C. Monroe, D. Leibfried, B.E. King, D.M. Meekhof, W.M. Itano, D.J. Wineland: submitted to Phys. Rev. A.
47. M. Kasevitch, D.S. Weiss, E. Riis, K. Moler, S. Kasapi, S. Chu: Phys. Rev. Lett. **66**, 2297 (1991)
48. J. Reichel, O. Morice, G.M. Tino, C. Salomon: Europhys. Lett. **28**, 477–482 (1994)
49. K. Moler, D.S. Weiss, M. Kasevitch, S. Chu: Phys. Rev. A **45**, 342 (1992)
50. J.I. Cirac, R. Blatt, A.S. Parkins, P. Zoller: Phys. Rev. Lett. **70**, 762 (1993)
51. C. Cohen-Tannoudji: *Cours de physique atomique et moleculaire, annee scolaire 1985–1986*, "Refroidissement des ions piégés", Collège de France lecture notes, unpublished
52. D.J. Wineland, W.M. Itano: Phys. Rev. A **20**, 1521 (1979)
53. W. Neuhauser, M. Hohenstatt, P. Toschek, H. Dehmelt: Phys. Rev. Lett. **41**, 233 (1978)
54. S. Stenholm: Rev. Mod. Phys. **58**, 699 (1986)
55. D.J. Wineland, J. Dalibard, C. Cohen-Tannoudji: J. Opt. Soc. Am. B **9**, 32–41 (1992)
56. J.I. Cirac, R. Blatt, P. Zoller, W.D. Phillips: Phys. Rev. A **46**, 2668 (1992)
57. J.I. Cirac, R. Blatt, A.S. Parkins, P. Zoller: Phys. Rev. A **48**, 1434 (1993)

58. G. Birkl, J.A. Yeazell, R. Rückerl, H. Walther: Europhys. Lett. **27**, 197–202 (1994)
59. M.G. Raizen, J.M. Gilligan, J.C. Bergquist, W.M. Itano, D.J. Wineland: Phys. Rev. A **45**, 6493 (1992)
60. R.A. Cline, J.D. Miller, M.R. Matthews, D.J. Heinzen: Opt. Lett. **19**, 207 (1994)
61. I thank C. Monroe for pointing this out to me.
62. F. Plumelle, M. Desaintfuscien, M. Houssin: IEEE Trans. Instr. Meas. **42**, 162 (1993); F. Arbes, M. Benzing, T. Gudjons, F. Kurth, G. Werth: Z. Phys. D **29**, 159–161 (1994); F. Arbes, M. Benzing, Th. Gudjons, F. Kurth, G. Werth: Z. Phys. D **31**, 27–30 (1994)
63. D.J. Wineland, J.J. Bollinger, W.M. Itano, D.J. Heinzen: Phys. Rev. A **50**, 67–88 (1994)
64. D.J. Wineland, H.G. Dehmelt: J. Appl. Phys. **46**, 919 (1975)
65. A.R. Calderbank, P.W. Shor: Phys. Rev. A, **54**, 1098 (1996)
66. A.M. Steane: Proc. Roy. Soc. Lond. A **452**, 2551 (1996)
67. E. Knill, R. Laflamme: Phys. Rev. A **55**, 900 (1997)
68. L.A. Khalfin: Zh. Eksp. Teor. Fiz. **33**, 1371 (1957) [Sov. Phys. JETP **6**, 1053 (1958)]; R.G. Winter: Phys. Rev. **123**, 1503 (1961); B. Misra, E. C.G. Sudarshan: J. Math. Phys. **18**, 756 (1977)
69. A. Berthiaume, D. Deutsch, R. Jozsa: In *Proceedings of the Workshop on Physics and Computation, PhysComp 94*, Los Alamitos; IEEE Computer Society Press pp.60-62
70. P.W. Shor: Phys. Rev. A **52**, R2493 (1995)
71. A.M. Steane: Phys. Rev. Lett. **77**, 793 (1996)
72. P. Shor: In *Proc. 37th Symp. on Foundations of Computer Science*, to be published (Preprint quant-ph/9605011)
73. D.P. DiVincenzo, P.W. Shor: Phys. Rev. Lett. **77**, 3260 (1996)
74. A.M. Steane: Phys. Rev. Lett. , (1997)
75. J.I. Cirac, T. Pellizzari, P. Zoller: Science **273**, 1207 (1996)
76. H. Landolt, R. Börnstein, A. Eucken: *Landolt-Börnstein Zahlenwerte und Funktionen, (Atom und Molekularphysik),* (Springer-Verlag, Berlin 1952); J. Kowalski et al.: Hyperfine Interactions **15**, 159 (1983)