Spectroscopy and quantum optics with trapped ions

Richard Thompson

Blackett Laboratory, Imperial College,
Prince Consort Rd., London SW7 2BZ, UK

Email: r.c.thompson@ic.ac.uk

Abstract

In these lecture notes, I describe the use of ion traps in experimental investigations of spectroscopy and quantum optics. Ion traps are well suited to this type of investigation because of the well-controlled conditions under which ions are held in traps and because they are well isolated from the environment. The notes start with an account of the way that ion traps work, concentrating on the radiofrequency or Paul trap. The techniques of laser cooling in ion traps are then discussed. Laser cooling is important because it allows new experiments to be performed which otherwise would not be possible. In particular, it allows one to work with single atomic particles and in some circumstances to cool them even to the ground state of the vibrational motion. The rest of the notes deal with various experimental studies undertaken in spectroscopy and quantum optics with trapped ions, including observations of quantum jumps; cavity quantum electrodynamics; nonclassical states; quantum logic gates; the quantum Zeno effect; and frequency standards and fundamental constants. These notes do not attempt to give a full account of the theory of these phenomena, but rather to give an idea of the special characteristics of ion traps, of the very wide range of investigations that have been undertaken with them, and of the potential they hold for future investigations in spectroscopy, quantum optics and other areas of physics.

1 Introduction

Ion traps have now been in use for about forty years. In that time they have moved from being a spin-off from work with linear mass analysers to constituting a whole research field, with applications in mass spectrometry, measurements of fundamental constants, atomic physics, spectroscopy, quantum optics, frequency standards, cavity quantum electrodynamics and quantum computing. For many of these applications (though by no means all) it was the advent of laser cooling which allowed the technique to develop so effectively. Laser cooling allows signals to be observed even from single atomic ions and enables scientists to prepare ions in specific internal states and well-defined states of motion in the trap.

These lectures will concentrate on the use of ion traps for studies in spectroscopy and quantum optics. There are other applications of ion traps, particularly in mass spectrometry of atomic and molecular ions and in studies with trapped electrons. There are also many similarities between work with neutral atoms and that with atomic ions (see for example [1, 2]).

For a general review of work with trapped ions see, for example, the review by Thompson [3] or the book by Ghosh [4] and for applications in fundamental studies in physics see Horvath *et al* [5].

The lecture notes are organised as follows. First, we deal with the theory of ion traps, concentrating mainly on the radiofrequency or Paul trap (section 2). Then we look at how laser cooling is applied to ions in traps (section 3). The rest of the lecture notes is taken up with various applications of trapped ions: quantum jumps (section 4); cavity quantum electrodynamics (section 5); nonclassical states (section 6); quantum logic gates (section 7) ; the quantum Zeno effect (section 8); and frequency standards and fundamental constants (section 9).

2 Basic theory of ion traps

2.1 The Penning trap

Penning and Paul traps use the same basic set of electrodes, shown in figure 1a. This consists of a *ring* electrode, shaped like the inner part of a doughnut, and two *endcap* electrodes, which are similar to hemispheres. Actually these electrodes follow the shapes of the equipotential surfaces corresponding to a pure quadrupole potential of the form:

$$\phi(r,z) = \frac{U_0}{R_0^2}(r^2 - 2z^2) \tag{1}$$

where $r^2 = x^2 + y^2$, $R_0^2 = r_0^2 + 2z_0^2$ is a geometrical constant depending on the size of the trap (with r_0 and z_0 defined in figure 1a) and U_0 is the potential of the ring with respect to the endcaps. U_0 is made negative in order to generate a potential well in the axial (z) direction. This traps positively charged ions in one dimension. However, this potential cannot trap in three dimensions on its own, but needs some additional feature to give three-dimensional stability. In the Penning trap [6] this comes from an additional static axial magnetic field. Then an ion attracted out towards the ring electrode is forced into a combination of circular orbits in the radial (x, y) plane. The two resulting motions (at the *modified cyclotron* frequency (ω_c') and the lower *magnetron* frequency (ω_m) can perhaps best be seen by transforming into the frame rotating at half the cyclotron frequency (given by $\omega_c = eB/m$), in which the magnetic field is effectively cancelled out [7]. In this frame the two radial motions are seen to be a rotation in the positive or negative sense around the centre of the trap at a frequency ω_1 given by

$$\omega_1 = \sqrt{\omega_c^2/4 - \omega_z^2/2} \tag{2}$$

where ω_z is the axial oscillation frequency given by

$$\omega_z^2 = 4e(-U_0)/mR_0^2. \tag{3}$$

U_0 is negative for trapping positive ions. The two radial frequencies are then found by transforming back into the laboratory frame and are found to be

$$\omega_c' = \omega_c/2 + \omega_1 \tag{4}$$

and

$$\omega_m = \omega_c/2 - \omega_1. \tag{5}$$

The Penning trap is a static trap but it has the disadvantage that the magnetron motion is unstable, meaning that the total energy associated with the magnetron motion is negative. Therefore in the presence of collisions there will always be a tendency for the magnetron orbit to increase in size, necessitating work under ultra-high vacuum (UHV) conditions (say less than 10^{-9} mbar).

2.2 The radiofrequency or Paul trap

The Paul trap [8] uses an alternative method for generating an effective three-dimensional trapping potential. In this case it is a *dynamic* trap. It works by using a potential which has an AC component, that is, we use a potential given by $U_0 + V_0 \cos \Omega t$. For half of the cycle this gives a potential which is stable (i.e. has a minimum) in the z direction but is unstable (i.e. has a maximum) in the radial plane. For the other half of the cycle it is unstable in the z direction but is stable in the radial plane. The magic of the Paul trap is that the parameters can be chosen such that overall the motion can be stable in *both* directions. To see this, we need first to write down the equation of motion of the ion (say for the z motion):

$$\ddot{z} = (U_0 + V_0 \cos \Omega t)(4z)(e/mR_0^2). \tag{6}$$

This can be written in a standard form called a Mathieu equation:

$$\frac{d^2 z}{d\tau^2} + (a_z - 2q_z \cos 2\tau)z = 0 \tag{7}$$

where $\tau = \Omega t/2$ and a_z and q_z are given by

$$a_z = -\frac{16eU_0}{m\Omega^2 R_0^2} \tag{8}$$

$$q_z = \frac{8eV_0}{m\Omega^2 R_0^2}. \tag{9}$$

This equation has stable solutions for certain ranges of values of a_z and q_z [9]. The final motion is quite complicated in general, but consists of a fast driven motion (the *micromotion*) at the applied frequency (Ω) and also motion at lower frequencies (the *macromotion* or *secular motion*). For small values of a_z and q_z we can represent the slow motion as resulting from an effective potential set up by the driven motion. This effective potential arises because the driven motion takes place in an inhomogeneous field so that the average force on the ion does not cancel to zero over a complete cycle of the driving field. The effective potential energy of the ion (also called a *pseudopotential*) is given by

$$V_{eff} = \frac{1}{8}m(a_z + \frac{1}{2}q_z^2)\Omega^2 z^2 \tag{10}$$

so that the oscillation frequency in this potential is

$$\omega_z = \frac{1}{2}\beta_z\Omega \tag{11}$$

where

$$\beta_z^2 = a_z + \frac{1}{2}q_z^2. \tag{12}$$

Similarly we can do the same analysis for the radial motion, and here we find that the values of a_r and q_r are $(-1/2)$ times the corresponding values for the axial motion. Therefore the stability conditions for the radial motion are different from those for the axial motion. The final situation is summarised in the *stability diagram* shown in figure 2. This shows that there are regions in the (a_z, q_z) plane where both motions are simultaneously stable. Although there are several such stability regions, in fact work has only been performed in the first and largest stability region, close to the origin. It can be seen that this has q_z ranging up to a value of approaching unity, whereas a_z is typically much less, and is often set at zero (i.e. there is no applied static potential).

Since the a and q parameters depend on the charge to mass ratio of the ions, the Paul trap has the potential to be used as a mass-selective device. Indeed, the Paul trap was originally developed in Paul's group in Bonn in the early 1950's out of work on linear mass filters [8]. If the trap parameters are chosen such that a=0, then the trap is stable for all masses down to a particular value (determined by the maximum value of q_z, which is around 0.9). On the other hand, the parameters can be chosen to put the operating point in the corner of the stability diagram and then the trap will be stable only for a very narrow range of masses. In this way it will act as a selective device, only trapping in a narrow range of charge-to-mass ratio. Furthermore, since the oscillation frequencies in the trap are mass-dependent, it is possible to selectively excite the motion of one species of ion in isolation. These motionally excited ions can then be detected using time-of-flight methods, for example. Commercial devices are available for use in mass spectrometry applications where the mass-selective nature of the Paul trap is exploited [10].

A useful concept for an ion trap is the trap depth, i.e. the potential energy which an ion must have in order to escape from the confining potential. In a Paul trap this is the magnitude of the pseudopotential at the edge of the electrode. In the axial direction this has the value D_z where $eD_z = \frac{1}{2}m\omega_z^2 z_0^2$. Typically this has the value of a few electron volts.

Since the oscillation frequency in the trap depends on (V/R_0^2), the applied voltage has to rise as the square of the size of the trap if the oscillation frequency is to be kept constant. If a high oscillation frequency is required (especially if the Lamb-Dicke regime is to be achieved - see section 3.3) then the applied voltages will have to be high and the trap also needs to be small, so that prohibitively large applied RF voltages can be avoided. This has led to the design and construction of miniature Paul traps (see section 2.4).

2.3 Motion in the Paul trap

The motion of ions in the Paul trap can become quite complicated as there are in general three components to the motion: first, the driving field at angular frequency Ω; second the radial secular motion at frequency ω_r (there are in fact two independent components to this at the same frequency which can be considered as the x and y amplitudes or equivalently as the clockwise and anticlockwise amplitudes); and third the axial secular motion at frequency ω_z. In a perfect trap the motion in a given direction (say the z direction) can be written (for low values of a_z and q_z) as:

$$z(t) = A\cos\omega_z t[1 + \frac{1}{2}q_z\cos\Omega t]. \tag{13}$$

This shows that it is not possible to eliminate the micromotion (at angular frequency Ω) unless the ion is stationary (A=0). This is because it is the motion at Ω that sets up the pseudopotential in which the ion is moving. Thus a single ion can in principle reside at the centre of the trap without moving. However, if there are more ions, they repel each other because of the Coulomb interaction so that even if they have no secular motion, they do not reside at the centre of the trap. This means that they always have some residual micromotion, which limits the minimum value of their mean kinetic energy, i.e. the minimum temperature to which they can be cooled.

In a real trap there are always small perturbations (e.g. contact potentials on the electrodes) that give rise to extra static fields which offset the equilibrium position of the ions from the centre of the trap. This leads to an extra term in the motion at frequency Ω, unrelated to the secular motion. The result of this is that even a single ion will not be free of the micromotion when it is at the equilibrium position. In many experiments, where a single ion is required to be stationary at trap centre, measures have to be taken to cancel out any of these stray fields. This is generally done by minimising the amplitude of the residual micromotion using extra electrodes to apply additional DC fields to the ion.

2.4 Practical Paul traps

Many ion traps correspond quite closely to the diagram of the electrodes presented here (figure 1a). However, for many experiments it is necessary to confine ions very tightly, with high secular oscillation frequencies. In this case miniature traps are constructed, which are difficult to manufacture with electrodes curved accurately in three dimensions. Therefore the electrodes are generally constructed with simpler shapes, designed such that the leading term in the potential is the desired quadrupole term with all other higher-order terms being sufficiently small that they do not affect the potential appreciably, especially near the centre of the trap. Examples of traps with simplified electrode designs can be found in several papers, for example [11, 12, 13]. These traps are all of the order of 1 mm or less in diameter. Such traps are designed mainly for use with one ion (or at most a few ions). The applied frequency is typically tens of MHz and the voltage may need to be several hundred volts, giving secular frequencies of a few MHz. Some of these miniature traps are designed as a high quality-factor quarter-wave RF resonator giving a large voltage amplification factor, thus eliminating the need to generate a large RF voltage externally [14]. Larger traps designed for trapping large numbers of ions generally do not need to create such a steep trapping potential, and therefore they can be run with lower frequencies (hundreds of kHz) and lower voltages (of the order of 100-200 V) [15]. Such traps may be 10-40 mm in diameter. In all cases the secular frequencies are typically 0.1 to 0.2 times the applied frequency (i.e. β is typically 0.2 - 0.4).

2.5 Linear RF trap

A development of the Paul trap is the linear RF trap. This is made from four rods and an RF potential is applied between the two opposite pairs of rods. This makes a two-dimensional trap (see figure 1b), and some means has to be provided to keep the ions in at both ends as otherwise they will escape. This can be done, for example, by adding endcap rods at both ends on which a DC potential is placed [16]. The reason for the interest in the linear trap is that instead of having a single point where an ion is free of micromotion, there is now a line that has the same property. Therefore it is possible to have a number of ions in a linear trap, *all* of which are free of micromotion. A number of applications are seen for this approach, but especially for frequency standards (where a larger signal-to-noise ratio can be obtained for a given maximum second-order Doppler shift: see section 9) and quantum computing, where several ions are required to be free of micromotion simultaneously (see section 7).

An alternative way to get a similar result is to have a very long set of electrodes which are bent round to form a ring trap [17]. Then no axial confinement is necessary as the ions are allowed to go round the whole circumference of the ring trap. The ring trap has been used for studies of crystallisation of ions at low temperatures [18]. In this experiment it was possible to form an ion crystal of up to one million ions.

2.6 Operation of ion traps

There is not enough space here to go into details of the operation of ion traps but we need to discuss the subject briefly before moving on. More details can be found in reviews or in the book by Ghosh [4].

First, we consider the loading of traps. This is generally achieved using a weak atomic beam crossed with an electron beam. The atomic beam often comes from a small oven which is heated electrically, and the electron beam comes from an electron gun or just a small filament that is heated so that it emits electrons which are then accelerated towards the electrodes. In a Penning trap the electrons have to travel along the magnetic field lines into the trap but in a Paul trap there is no such restriction on the location of the filament. The atoms will generally have thermal energies, whereas the electrons need to have sufficient energy to ionise the atoms (generally 10-20 eV for the creation of singly-charged ions). Once created, the ions will be attracted towards the centre of the trap and will generally have total energies of several eV (before cooling). It is always necessary for the base pressure in the vacuum system to to be low (say $< 10^{-9}$ mbar) for successful operation of traps. However, in the case of a Paul trap, a low atomic weight buffer gas may be introduced at low pressure for cooling purposes (see section 3.1).

Then we need to be able to detect ions. For large clouds, electronic detection techniques are available, which are often similar to those used for the detection of electrons in traps. There are also destructive means of detection which can be used, based, for example, on time-of-flight techniques for ions ejected from the trap. However, for small clouds of ions and continuous, non-destructive observation, there is no real alternative to optical means of detection. This is based on the detection of light from a laser beam which has been resonantly scattered by ions in the trap. This is very sensitive (see section 4.1) and does not perturb the ion cloud strongly. It is also highly selective and can even distinguish between isotopes of the same element. If laser cooling is being used (see the next section) then we simply need to detect the scattered laser cooling light, and this can give us information on the progress of the laser cooling: as the laser cooling proceeds, the scattered fluorescence signal generally increases strongly. It is necessary to use a sensitive detector which can detect single photons, such as a photomultiplier or a high sensitivity imaging detector.

3 Laser cooling of trapped ions

3.1 The need for laser cooling

When ions are loaded into ion traps they have energies which are typically of the order of a few electron volts, much larger than normal thermal energies. This is because they are not generally created at the centre of the trap, so they gain energy as they fall towards the trap centre in the trap potential. The hot ions will have large amplitude motions and it will be possible for them to be ejected from the trap. In fact, in the Paul trap the energy tied up in the micromotion can also be coupled into the secular motion of the ions through ion-ion collisions [19]. This can also lead to loss of ions from the trap.

In order to localise the ions near the centre of the trap and to increase the density of the ion cloud, it is necessary to cool them. By this we mean a reduction in the mean kinetic energy of the ions, so the concept of temperature can still be applied even in the case of a single ion. For the magnetron motion in a Penning trap, the *total* energy associated with this motion actually increases as cooling takes place, because the potential energy of this motion, which is dominant, is negative.

Cooling the ions has several other advantages. First, it reduces the chances of ions being lost through collisions. Second, it reduces the Doppler effect (to all orders) so that higher resolution can be obtained in spectroscopic measurements and higher fluorescence signals from trapped ions can be obtained. Third, it also increases the interaction time between ions and any other radiation with which the ions are interacting (e.g. a laser beam or microwave radiation) thus reducing transit time broadening effects.

Cooling may be accomplished by several different means. For large clouds of ions in a Paul trap (but not in a Penning trap because of the unstable nature of the magnetron motion) a buffer gas can be used to reduce the temperature to something approaching room temperature [20]. This can be very effective and it is very simple to do: it just involves letting buffer gas (generally helium) into the vacuum system at a low pressure (typically 10^{-6} mbar). The buffer gas needs to be light and unreactive in order to be effective, so helium is usually the best choice.

However, buffer gas cooling can at best bring the temperature down to room temperature and generally it is not this low. In order to take the temperature lower it is necessary to use laser cooling. This can be applied in all cases except large clouds of ions in a Paul trap, where the RF heating prevents laser cooling being effective.

It should be pointed out that many successful experiments have been performed with buffer-gas cooled ions in traps, in particular spectroscopic experiments of various sorts (e.g. precision microwave spectroscopy of Ba^+ [21], Doppler-free optical spectroscopy of Th^+ [22]). For an example involving frequency standards, see section 9.3.

3.2 Doppler cooling

Laser cooling can be applied to trapped ions or to neutral atoms [2]. For ions the most important process is Doppler cooling which was first suggested for ions by Wineland and Dehmelt in 1975 [23]. Advanced techniques such as polarisation gradient cooling are not generally applicable in ion traps due to the applied fields (but see [24]). There have been just one or two demonstrations of this sort of laser cooling in an ion trap [25]. For the lowest temperatures for single trapped ions, Raman cooling is used (see section 3.4).

Doppler cooling works in ion traps in a very similar way to neutral atoms, but there are some changes (see [26] for a full treatment). In particular, since the ions are bound by electromagnetic forces and move in closed orbits, they only need one laser beam to cool all degrees of freedom, so long as the beam is inclined at an angle to all the axes of the trap so that it has a component along all three directions. This is a big simplification compared to the case of neutral atoms, where up to six laser beams are required.

There are a number of different limits which need to be treated slightly differently, depending on the relative values of the relevant frequencies in the problem. If the ion oscillation frequencies (ω_z, ω_r) are less than the linewidth (γ) of the laser cooling transition (generally the resonance transition of the ion) (i.e. $\omega_z, \omega_r \ll \gamma$), then this is called the weak binding limit and this is where laser cooling works in a very similar manner to the way in which it works for neutral atoms. This is the usual case for trapped ions. We also assume here that the recoil energy, R, given by

$$R = \hbar^2 k^2 / 2m \tag{14}$$

is much less than $\hbar\gamma$ (the recoil energy is the kinetic energy of an atom having a momentum equal to that of one photon). For optimal laser cooling the laser is tuned to $(\gamma/2)$ below the centre of the transition and in this case the limiting temperature is the same as for neutral atoms:

$$E_{min} = \frac{1}{2} k_B T_{min} \approx \frac{1}{4} \hbar\gamma. \tag{15}$$

where E_{min} refers to the kinetic energy of the ion. This gives a minimum temperature $T_{min} \approx \frac{1}{2}\hbar\gamma/k_B$ which for magnesium is of the order of 1 mK.

If, on the other hand, the oscillation frequencies are larger than the linewidth ($\omega_z, \omega_r \gg \gamma$) then we have the tight binding limit and this is best thought of in terms of the vibrational quantum numbers of the ion in the trap potential well. In this case the Doppler effect gives rise to a carrier at the resonance frequency with a set of sidebands spaced at the trap oscillation frequency (see section 3.3). Now the optimal cooling is obtained by tuning the laser to the position of the first sideband below resonance. Each time a photon is absorbed we now lose on average one vibrational quantum. This is referred to as *sideband cooling*. In this case the limiting mean vibrational quantum number, $\langle n_{min} \rangle$, is given by:

$$\langle n_{min} \rangle = \frac{5}{16} \frac{\gamma^2}{\Omega^2} \ll 1. \tag{16}$$

In the weak binding limit we can think of a photon as being absorbed at a specific time and place in the motion of the ion. In the strong binding case, the absorption is best thought of as being a transition from one state of vibrational motion to another, so the absorption is in some sense spread over the entire orbit.

For ions in Penning traps the cooling process is complicated by the fact that the laser beam needs to be offset from the centre of the trap in order to cool both radial motions (magnetron and cyclotron) at the same time. The beam needs to be offset to the side where the magnetron motion causes the ions to travel in the same direction as the laser beam [27]. This is related to the unstable nature of the magnetron motion. As a result of this, the temperature associated with the magnetron motion cannot be reduced as much as that associated with the modified cyclotron motion. The magnetron temperature is typically more like 1K rather than the Doppler limit of typically 1 mK. The cooling can be improved by making use of two laser beams rather than one [28].

In recent work in our group at Imperial College, we have calculated cooling rates for both the magnetron motion and for the cyclotron motion, as a function of the two main parameters in the problem: the laser frequency detuning and the offset of the laser beam from trap centre [29]. These calculations confirm that in order to cool both motions simultaneously the laser beam needs to be offset as stated above and the detuning needs to be to the low-frequency side of the transition. The calculated cooling rates for the two motions are very different, with the magnetron cooling rate always being much smaller (by a factor of 10 to 100) than that of the cyclotron motion. This arises because the cooling for the magnetron motion comes from a balance of two opposing effects which are nearly equal in magnitude, so the overall cooling effect is small (and can even be negative if care is not taken in the choice of parameters).

In recent experiments [30] we have been able to measure these cooling rates for each component of the ion motion individually for the first time, and we were able to confirm that the magnetron cooling rate is very small compared to the cyclotron one. These experiments treat the ion cloud as a driven classical oscillator, damped by the laser cooling, and they measure the damping rate from the plot of the phase difference between the driving force (an oscillating potential applied to the trap electrodes) and the driven motion of the cloud. As with any driven, damped oscillator, this phase difference changes by π as the driving frequency goes through the resonance frequency, and the width of the curve is proportional to the damping rate.

3.3 Lamb-Dicke regime

For ions bound in the simple harmonic potential of an ion trap, the radiated light can be thought of as coming from a frequency-modulated source, as a result of the oscillatory motion of the ions. Due to this motion, the ions do not radiate with the conventional Gaussian Doppler line shape, which is a continuous function of frequency, but rather they radiate at the unmodulated carrier frequency with sidebands spaced at integer multiples of the trap oscillation frequency each side of the carrier. The amplitudes of the various sidebands are determined by the distribution of amplitudes of the ion oscillation. In the case where $\omega \ll \gamma$ these sidebands merge into each other and we retrieve the normal Gaussian distribution. If, however, we have $\omega \gg \gamma$ then the sidebands are well resolved.

There is a special case when the amplitude of the motion (generally for a single ion) is so small that the amplitudes of the various sidebands become small enough to be ignored. This is referred to as the *Lamb-Dicke regime* [31]. In this case the motion has an amplitude which is much less than $(\lambda/2\pi)$ where λ is the wavelength of the radiated light. This corresponds to the phase of the radiated light deviating from that of the unmodulated phase by an amount which is always much less than one radian. We have given a classical description of this effect here but a similar result follows from a full quantum mechanical treatment. The effect is the same as that referred to as *Dicke narrowing* in the context of collisions in atomic gases.

The reason that this is important is that now all the radiation (or equivalently, all the absorption) is concentrated in the carrier, which is free of the first-order Doppler effect. Thus, the first order Doppler effect can be overcome, not by selecting particular atoms, but by forcing all the radiation to be at a frequency which is unaffected by the first order Doppler effect. In fact the second order Doppler effect (relativistic time dilation) is still present but this is of course also reduced dramatically because of the reduction in the average kinetic energy of the ion. The fractional second order Doppler effect, $\delta\nu/\nu$, is given by $\gamma - 1$ where

$$\gamma - 1 = E_{kin}/mc^2 = 3kT/2mc^2 \tag{17}$$

For laser cooled ions this can be reduced to one part in 10^{15} or so (at a temperature of the order of 1K).

The first demonstration of the achievement of the Lamb-Dicke regime for an optical transition was by Bergquist *et al* [32] using mercury ions. They cooled a mercury ion initially using light tuned to the resonance line at 194 nm and then probed it at 282 nm on a narrow transition to a metastable level. Two sidebands are seen on each side of the carrier, corresponding to the Doppler temperature of a few mK (the vibrational frequency of the ion is about 1.5 MHz). Then the ion is cooled by sideband cooling on the 282 nm transition itself (this has a lower Doppler limit as it is a narrow transition). This is achieved by tuning a narrow band laser to the first lower (red) sideband corresponding to the $\Delta n = -1$ transition. Then we are left with just one sideband on either side of the carrier, and the amplitudes of the sidebands become asymmetrical. This is because in the quantum mechanical treatment of the problem it is clear that we can go from $n = 1$ (or above) in the ground state to either $n+1$ or $n-1$ in the excited atomic state, but from $n = 0$ we can only go to $n+1$ as a vibrational quantum number of $n = -1$ is not possible. Thus the asymmetry of the sidebands tells us how close the ion is to the lowest vibrational level ($n = 0$). In this case an equivalent temperature of about 50μK is achieved, though temperature is not a particularly useful concept in this situation. The best way to specify the state of the system is in terms of the mean vibrational quantum number, $\langle n \rangle$, which had a value of about 0.05 in this experiment. This means that the system is in its ground state 95 % of the time. In this case a simple calculation of the temperature from the mean energy (ignoring the zero-point energy) gives a false result:

$$\langle E \rangle \approx \langle n \rangle \hbar\omega = k_B T \tag{18}$$

is not appropriate if $\langle n \rangle < 1$. Instead we must use the full Boltzmann distribution to find the relation between $\langle E \rangle$, $\langle n \rangle$ and T:

$$\langle E \rangle = \langle n \rangle \hbar\omega = \frac{\hbar\omega}{\exp\left(\hbar\omega/k_B T\right) - 1} \tag{19}$$

which, for $k_B T \ll \hbar\omega$, becomes

$$\langle E \rangle = \langle n \rangle \hbar\omega = \hbar\omega \exp\left(-\hbar\omega/k_B T\right). \tag{20}$$

The *Lamb-Dicke parameter*, η, gives an idea of how large the ground state wavefunction is compared to the wavelength of the radiation, and the Lamb-Dicke regime cannot be achieved unless $\eta \ll 1$. It is defined by

$$\eta = k x_0 = k\sqrt{\hbar/2m\omega} \tag{21}$$

where $k = 2\pi/\lambda$ and $x_0 = \sqrt{\hbar/2m\omega}$ is the spread of the ground state wavefunction in the potential well, which has oscillation frequency ω. In the above experiment, η has a value of approximately 0.09.

3.4 Raman cooling

The above work is unusual in that standard Doppler cooling is used to achieve a very low temperature, and the Lamb-Dicke limit, by making use of a very narrow transition. One alternative route to very low temperatures in an ion trap is to use Raman cooling. Here we rely on a *driven* Raman transition between two specific energy levels of the ion in the trap, using a pair of relatively intense laser beams whose difference frequency corresponds to the difference in energy between the levels. The intermediate virtual level is chosen to lie close to an excited state of the ion so that the rate of the Raman transition is enhanced.

Raman cooling has been used by the NIST group in Boulder to cool single Be ions to the lowest vibrational state in the trap [33]. It works in the following manner. First, the ion is cooled using Doppler cooling to the region of a few mK (the Doppler limit for Be is 0.5 mK, corresponding to a mean vibrational quantum number of about 0.5 in this trap, which has a vibrational frequency of 11.2 MHz (x direction). Then the Raman beams are used to transfer the ion from the electronic ground state (with vibrational quantum number n) to a different hyperfine level (with vibrational quantum number $n-1$). The Raman pulse is timed to be a π pulse - i.e. to transfer the ion with close to 100% probability. The ion is then excited resonantly to a state from which it decays back spontaneously to the ground state, probably remaining in the vibrational state $n-1$ (so long as $\eta \ll 1$; in fact the value is roughly 0.2 in this experiment). Thus the overall effect of this cycle is to reduce n by 1. The cycle is then repeated several times, and this should reduce n to zero with high probability. In fact Wineland's group are able to achieve an average value of n of about 0.014 in one dimension. This means that the ground state is occupied at least 98 % of the time. With cooling applied in all three dimensions, this figure becomes 92 %. Thus the ion can effectively be prepared in the ground vibrational state of the system. This forms the basis of many other experiments (see sections 6 and 7).

3.5 Sympathetic cooling

There is only a small number of ion species to which laser cooling can be successfully applied, due to the need to have a closed laser cooling cycle. This means that the laser excites the ion on a strongly-allowed transition to an excited state from which it can decay *only* back to the ground state, or to at most one or two other states from which ions can be fed back into the cooling cycle with extra lasers. Additionally, all the required laser wavelengths must be achievable with continuous and tunable lasers such as diode lasers, dye lasers or titanium sapphire lasers. These requirements restrict the number of available ion species to about ten, and exclude all molecules, where the level structure is far too complicated to find a simple laser cooling scheme.

However, one possibility is to trap two species of ion in the same trap at the same time, and then to laser cool one of them and to allow these laser cooled ions to cool the other ions in the trap by collisional processes. This is similar to buffer-gas cooling except that the buffer gas in this case is laser cooled and is also held in the trap. Surprisingly, perhaps, this method works very well and has been used in several experiments (e.g. [34]; see also section 9.3).

One interesting possibility is to apply these techniques to *molecular* ions in a trap. This could prove a good way in which to prepare molecular ions at very low temperatures, with not only the external motion of the molecules cooled but also, if the conditions can be optimised, with the *internal* motion (i.e. vibration and rotation) at low temperatures. At Imperial College, we have been testing the effect of sympathetic cooling by introducing various ion species into a trap with laser-cooled magnesium present [35]. We have detected the presence of these ions by driving their oscillation frequencies in the trap with additional oscillating potentials on the trap electrodes. At resonance these ions heat up, and this causes the magnesium ions to heat up as well, resulting in a change in the level of fluorescence. Thus we can detect the presence of different ions in the trap. We can also measure the temperature of the magnesium ions roughly by measuring their excitation spectrum from a scan of the laser frequency through the atomic resonance. From simulations we have been able to show that all the ions in the trap are expected to be very nearly in thermal equilibrium with each other, so therefore a measurement of the temperature of the magnesium ions can be used to infer the temperature of the molecular ions as well. This experiment was able to show that the temperature remains very low, so long as there are not too many of the molecular ions present [35].

3.6 Ion crystals

If the temperature of an ion cloud is reduced far enough, the ions will cease to behave like a gas, but will fall into a regular crystal-like structure. This is an interesting phenomenon, particularly as the typical inter-ion spacing in such a crystal is of the order of 10 μm, some ten thousand times bigger than the typical spacing in a conventional crystal.

The formation of an ion crystal is determined by the value of Γ, which is given by

$$\Gamma = \frac{e^2}{4\pi\epsilon_0 akT} \tag{22}$$

where a, the Wigner-Seitz radius, defined by $(4/3)\pi a^3 n = 1$, is a measure of the inter-ion spacing. When Γ is greater than 2, we expect liquid-like behaviour, and when it is greater than 170 (for an infinite plasma), we expect crystallisation to take place.

Many experiments have been performed on this topic and crystals have been observed and photographed in many different types of traps, including Penning, RF, linear and ring traps (see, for example, [3]). Recently, Bragg scattering has been observed from an ion crystal in a Penning trap. One particular point of interest is the formation and melting of the crystals, and there has been much discussion of whether the melting process is an order-chaos transition. Ion crystals have a particularly important area of application in the use of trapped ions for quantum gates and quantum computers - see section 7.

4 Quantum jumps

4.1 Single ion techniques

We have already discussed laser cooling of single ions, but at this stage it is worth thinking about how it is possible to work at all with single ions in ion traps. In particular, we consider how to prepare such an ion in a trap and what signal levels may then be expected.

The first experiments with single ions started only a few years after laser cooling was first introduced for trapped ions. The first laser cooling was reported in 1978 [11, 36] and a photograph of a single ion was published in 1980 [37], with a spectrum of a single laser cooled ion the next year [38]. Preparation of a single ion in a trap is generally a case of turning the intensity of the atomic beam and the electron beam down to very low values so that on average only very few ions are loaded on each loading attempt. If the average number of ions loaded is less than one, then if a signal is seen, it is likely to be from a single ion. The technique of quantum jumps (see section 4.2) can then be used to confirm unambiguously whether it is really a single ion or not.

An alternative is to load a small number of ions and then to attempt to lose all the ions except one. In some experiments it is possible to see the signal level reduce in definite steps as each ion is lost. This may be achieved through charge-exchange collisions with background gas or by taking the trap parameters to the edge of the stability region so that the trap is only just stable for a single ion and not for a cloud of ions.

It is worth considering how much optical signal can be seen with a single ion in a trap, because it is surprisingly large. The point is that a single cold ion can remain in the laser beam all the time and can also remain in resonance with the laser all the time, because it is no longer affected by the Doppler effect once it is cold. Then the ion can scatter of the order of $A/2$ photons per second, where A is the Einstein A-coefficient of the resonance transition, if the transition is saturated (this generally requires of the order of 100 μW or less in a focussed laser beam). Since A is typically 10^8 per second, we can collect of the order of 10^6 photons with a typical solid angle for collection, given the presence of the electrodes which generally limit the available solid angle. Of these, we may be able to detect 10^5 photons per second, with typical quantum efficiencies of photomultipliers and optical filter efficiencies. Thus in a real experiment we may expect to see 10 to 100 thousand counts per second detected from a single ion. This has indeed been seen in several different experiments [39, 40, 41] (this represents an overall efficiency of the detection system of the order of one part in a thousand). Indeed, for those experiments where the scattered light is in the visible region of the spectrum, experimenters have been able to see the ion with the naked eye [41]. Unfortunately, for many ions the resonance transition is in the UV region (since excitation energies of ions are generally larger than those for atoms, which are usually in the visible), so for these experiments one has to rely on other detectors than the eye. Many images of single ions or small crystals of ions in Paul or linear RF traps have now been published.

4.2 Setup for quantum jumps

It is very hard to observe directly a quantum mechanical system jumping between two levels, as the overall detection efficiency for the photon emitted in the transition is typically only 10^{-3} or less. However, if there are 3 levels in a suitable arrangement, the jumps on one transition can be used to turn the resonance fluorescence on another transition on and off. This is the basis for the experimental observation of quantum jumps in real time.

Quantum jumps therefore require an atomic system which has at least three levels arranged in a V-configuration with one ground state coupled to two excited states, one of which is a short-lived resonance level and the other of which is a long-lived metastable level. When the ion is in the ground state, it can be excited into the resonance level by the laser cooling radiation and it will then scatter resonance fluorescence at a high rate determined by the A-value for the transition. However, if the ion is excited in some manner into the metastable level (for instance, by an off-resonant transition, a low-probability decay from the resonance level, or a direct excitation from the ground state by a second laser) then it will no longer be able to scatter resonance light, and so the observed fluorescence signal will drop to zero. At some stage the ion will decay back to the ground state and then the process will start again. The times of switching between the *on* and *off* states are randomly distributed, giving a so-called *random telegraph signal*.

The theory of this process was first described by Cook and Kimble in 1985 [42]. The *on* and *off* times both have an exponential distribution (rather like radioactive decay) with mean values determined by the rate of excitation into the metastable state and the rate of decay out of it. These mean times can be several seconds (eg [41]) or they may be as short as 10-20 ms [39], depending on the atomic system and the various rates involved. The signal observed is of course quite different from what one observes if a larger number of ions is present in the trap, and before the first experimental observations were made, it was a matter of some debate as to whether one would really observe the random telegraph signal. This illustrates the way that quantum mechanics predicts manifestly different behaviour for a single atomic particle compared to that seen with an ensemble, and highlights the importance and significance of experiments with single ions.

4.3 Observations of quantum jumps

The first observations of quantum jumps were made in 1986 in several different experiments with Ba^+ and Hg^+ ions. All these early experiments were performed in Paul traps, mostly miniature traps optimised for the trapping of single particles. The Ba experiment is particularly interesting as the main fluorescing transition is in the blue and this means that the ion can be observed undergoing quantum jumps in real time with the naked eye [41]. In this experiment two lasers are needed: one to drive this blue transition from the $^2S_{1/2}$ ground state to the $^2P_{1/2}$ resonance level, and one to recycle ions that have decayed from this level to a metastable $^2D_{3/2}$ level (this transition is in the red). The quantum jumps take place to a second metastable level, $^2D_{5/2}$, via an off-resonant excitation involving the $^2P_{3/2}$ state. The mean on and off times are of the order of several seconds, so the characteristic quantum jump signal can be clearly seen (see figure 3). Another interesting feature of this experiment is that cooperative quantum jumps were seen in runs where 2 or 3 ions were trapped at the same time. Of course, a certain number of chance coincidences would be expected to be seen, but the observation in this experiment seems to indicate that these were more frequent than expected.

Since these early observations, many other experiments have been performed, and these are well summarised in reviews by Blatt and Zoller [43] and Cook [44]. The statistics of the quantum jumps have been verified in different systems and we now have a very good understanding of the way that single particles behave in these conditions. Since normal quantum mechanics deals with ensembles, it is not always the most appropriate way to treat the behaviour of single particles and new treatments of quantum mechanics which are more suited to this situation have been developed. See [39] and [45] for a discussion of these approaches to quantum mechanics.

Quantum jumps may be observed even in a two level system if a magnetic field is used to split the ground state into its Zeeman components. In this way it has been possible to observe quantum jumps in Mg ions [39, 40]. Then one of the ground state Zeeman levels acts as the metastable level and in the absence of laser radiation its lifetime is essentially infinite. The experiment is ideally suited to a Penning trap as the magnetic field is already there and has the right order of magnitude. In this experiment one laser (tuned to a strong transition out of one of the ground state Zeeman levels) can be used to serve several purposes: it acts as the laser source for laser cooling; it is used to produce a fluorescence signal for detection of the ions; it drives the ions into the metastable state; and it drives them out of the metastable state back into the ground state. The spontaneous Raman transitions which are driven by the laser take place at a rate which is determined by the laser intensity and the detunings of the various levels, which is itself proportional to the strength of the magnetic field. In the Imperial College experiment the mean *off* time was of the order of 15 ms and theory shows that the ratio of *on* to *off* times is (for a perfect set-up) 16:1. A plot of quantum jumps from this experiment is shown in figure 4.

4.4 Shelving technique

Quantum jumps are not merely an interesting quantum mechanical phenomenon. They have use in the construction of new frequency standards using ion traps. Here the problem is that in order to obtain the greatest benefit from the use of trapped ions in frequency standards one would like to use just a single ion, but then the signal level is going to be very low, especially on a narrow transition such as would be useful for a standard. The solution is to use the ion as an *atomic amplifier* to increase the observed signal, in the manner first suggested by Dehmelt [46]. He points out that when an ion absorbs a photon taking it to a metastable level this results in the loss of many photons on the resonance transition which would otherwise have been detected. Thus although one is unable to detect the single absorbed photon with high probability, one can detect the fact that it has been absorbed with very nearly unit probability by recording the absence of resonance fluorescence of the ion. This is called *electron shelving* and is used in all single ion frequency standards to obtain a high detection efficiency for transitions to the metastable state (see section 9.4).

5 Cavity quantum electrodynamics

5.1 Advantages of ion traps

The subject of quantum electrodynamics in a cavity (cavity QED) is a large and important one which is reviewed elsewhere (see, for example, [47]). Here we deal only with the special advantages which are offered by ion traps for experiments in cavity QED, with a brief account of the types of experiment in progress at present in this area.

The main advantage offered by the use of ion traps is the ability to have a definite number of ions in the cavity, and for this number to remain constant as a function of time. In other schemes, the atoms come from an atomic beam and so there is always a *distribution* of atom number. Only the mean number can be controlled, and not the exact number of atoms at any time. Also the atoms fly through the cavity, giving a limited interaction time of the order of 100 μs. Although this has advantages in terms of gaining information about the state of the atom on leaving the cavity (the Rydberg states which are commonly employed can be detected with almost unity efficiency), it means that the number of atoms in the cavity is not constant, which is a distinct disadvantage.

With an ion trap we can be sure that there is zero, one, two etc ions in the cavity and this will not change with time, as the ions are confined by the trap potential. The number of ions present can be verified using quantum jumps or by calibrating the average fluorescence level observed per ion.

It is also possible to locate a single trapped ion accurately at the position of the mode waist in the optical cavity. This allows the use of a concentric optical cavity, which has a very small beam waist. In other experiments this is not possible as the position of atoms in an atomic beam cannot be controlled well enough, so confocal cavities are then used, which are very sensitive to aberrations.

Against these advantages has to be set the inconvenience of using an ion trap: the electrodes get in the way of the cavity and it is therefore difficult to make a cavity suitably small around the position of the trapped ion(s). Also since most ion resonance wavelengths are in the UV, whereas those for neutral atoms are generally in the visible, we may have the added inconvenience of having to generate narrow-band, tunable laser radiation in the UV. This generally involves frequency-doubling techniques which yield much lower laser powers than are available in the visible. Finally, the ultra-high reflectance coatings needed for the construction of ultra-high finesse optical cavities are much harder to manufacture for UV wavelengths than for visible wavelengths, due to the increased absorption of the dielectrics used for the coatings. There are some ions which have visible transitions (e.g. Ba^+ and Ca^+) but there the atomic structure is such that more than one laser wavelength is required.

Overall, therefore, although in many ways it is desirable to perform cavity QED experiments with trapped ions, the experiments are technically very demanding and for this reason progress to date has been slow.

5.2 Experiments in progress

The experiments being set up at present are looking either for the modification of the spontaneous emission rate of an ion located at the centre of a cavity or for evidence of the *single atom laser* where a single ion generates coherent optical radiation in a cavity [48].

An excited atom in a cavity has its spontaneous decay rate modified because the rate at which radiation goes into the optical cavity is either enhanced (if the cavity is resonant with the frequency of the radiation emitted by the ion) or suppressed (if the cavity resonance frequency is well away from that of the ion). Thus if the rate of decay of excited ions in the cavity is measured, this should be dependent on the tuning of the cavity resonance. The experiment is difficult for several reasons, but in particular because the emission rate is only modified for that part of the radiation which goes into the cavity. Radiation which is outside the solid angle subtended by the cavity mode at the position of the ion is not affected by the position of the cavity resonance frequency. Thus the maximum modification to the spontaneous emission rate is limited by the solid angle subtended by the cavity mode, which is generally small. If the cavity is not resonant with the atomic frequency the reduction in spontaneous emission rate will therefore be quite small, but if it is resonant the enhancement can still be quite large.

There are several different ways of observing the effect of the cavity on the spontaneous emission rate of a transition. One way of detecting this for a weak transition down from a metastable state is by looking for a change in the rate of quantum jumps when the cavity is brought into resonance with the transition. Alternatively, it is possible to observe the light emitted directly into the cavity mode (through a partially transmitting cavity mirror) or out of the cavity. A different approach would be to measure the change in the natural linewidth of the transition. These different approaches are suited to different experimental conditions and experiments are currently being set up in a few laboratories around the world looking for these effects.

The single atom laser is a particularly interesting experiment in cavity QED with a single trapped ion [48]. Under the right experimental conditions, a single ion can be made to lase, and the properties of this system are quite unusual, and different from those of the micromaser, for instance, where single atoms fly through the apparatus one at a time. This system has been studied theoretically by several authors but its practical implementation is difficult due to constraints on the atomic structure and practical difficulties in the construction of the cavity. A concrete proposal for its implementation from Walther's group at Garching makes specific predictions for the Ca^+ ion [48]. This work shows that two lasing thresholds are expected: one where the laser turns on and another where it turns off again. For certain parameters, the light is expected to show sub-Poissonian statistics.

6 Nonclassical states

6.1 Introduction

Quantum mechanics allows a much greater range of possible states of a system compared to those allowed by classical mechanics, as is well known. This is because a quantum mechanical system can exist in a *superposition* of different eigenstates, whereas a classical system has to be in a definite state at any time. One consequence of this is that it is possible to construct quantum mechanical superpositions that have noise levels which are lower than the standard quantum mechanical limits set by the uncertainty principle. In fact, in these so-called *squeezed states* the noise associated with one type of measurement is decreased while that associated with another is increased, so there is no violation of the uncertainty principle overall [49]. Furthermore, some quantum mechanical states have properties associated with them that cannot be realised in classical systems (for example, sub-Poissonian statistics of a light field).

All these features of quantum mechanical systems have been of great interest to physicists over the years as they probe the very foundations of quantum mechanics and shed light on the inherent statistical nature of the theory. There has therefore been a great deal of effort which has gone into experimental demonstrations of nonclassical states (especially of light) in order to see whether quantum mechanics correctly predicts these often counter-intuitive phenomena. With ion traps these experiments fall into two areas: demonstration that the resonance fluorescence emitted by a single ion is sub-Poissonian and antibunched; and generation of non-classical states of motion in the vibrational state of a single trapped ion. These experiments are discussed in the following sections.

6.2 Nonclassical light

It is well known that the light scattered by a single atomic particle on resonance has non-classical properties. This is because (in a simplified semiclassical view) after the particle has emitted a photon in decaying from the excited state to the ground state there has to be a delay while the particle is excited again before it can emit another photon. This means that there is a very low probability of two photons being emitted together: there will on average be a delay of the order of the spontaneous emission lifetime before the second photon is emitted. This is in contrast to classical light where two photons may be emitted together. The end result of this is that the resonance fluorescence is both *antibunched* and *sub-Poissonian*. Antibunched light has a low probability of two photons arriving at the detector close to each other in time, while sub-Poissonian light has a variance of photon counts which is less than that associated with normal counting processes having Poisson statistics (with the variance equal to the mean). These two properties are related but distinct.

Walther's group at Garching has made demonstrations of the non-classical nature of resonance fluorescence from a single ion. In the latest experiment [50] they study a single magnesium ion in a miniature end-cap trap [13]. They can choose to measure the particle nature of the light scattered by the ion or its wave nature. By counting photons and recording the statistics of the counts they are measuring the particle nature of the light, and this shows that the light is indeed antibunched, as expected [50]. This is done by measuring $g^{(2)}(\tau)$, the intensity correlation function, which is the normalised probability of detecting a photon an interval τ after one photon has been detected. For a single ion, $g^{(2)}(\tau)$ is predicted to approach zero as τ approaches zero, and this is confirmed in the experiment (an earlier experiment by the same group shows plots of $g^{(2)}(\tau)$ for 1, 2 and 3 ions and also demonstrates that the light is sub-Poissonian [51]).

However, by measuring instead a beat signal obtained by mixing the resonance fluorescence from the ion with the laser light driving the transition, they are also able to see the spectrum of the resonance fluorescence, which probes the wave nature of the fluorescence. As expected, this spectrum is very narrow (remember the laser linewidth is eliminated by this heterodyne technique) and the measured linewidth of the elastically scattered light is less than 6 Hz (see figure 5). This experiment is technically very difficult as the signal level available with only a single ion is very low.

6.3 Nonclassical states of motion

Nonclassical properties of simple harmonic oscillators (SHOs) in quantum mechanics are generally treated in the context of modes of the radiation field, as radiation modes are a straightforward example of a quantum mechanical simple harmonic oscillator with the position and momentum becoming the electric field components of the mode [49]. Several different types of radiation can be defined: the vacuum state, where the mode is not excited at all; Fock states (or number states), which are states with a definite number of excitations (photons) in the mode; coherent states (with an amplitude and phase which are as definite as the uncertainty principle will allow - these are the states generated by lasers); thermal states (an incoherent superposition of number states such as is generated by a classical thermal source); and squeezed states (where the uncertainty in one quadrature (e.g. phase) is reduced at the expense of increased noise in another (e.g. amplitude)). All these different types of state of the radiation field (and more!) have been generated in experiments on beams of light.

However, with a trapped ion we have available a system which also matches the requirements of a quantum mechanical simple harmonic oscillator if we concentrate on the vibrational state of the ion. This is directly a single oscillator mode, well isolated from the environment and, with the Raman techniques developed for laser cooling, it is possible to manipulate the state of the system at will. The NIST group has, in a remarkable series of experiments, demonstrated the main non-classical properties of quantum mechanical SHOs using a single Be^+ ion in a miniature RF trap [52].

The experiment builds on the initial preparation of the $n = 0$ state of the vibrational motion using the Raman techniques discussed earlier (section 3.4). Once the ion is in the $n = 0$ state, it can then be manipulated using similar techniques. For example, a Raman pulse tuned to the $\delta n = +1$ (blue) sideband will place the ion in the $n = 1$ state (in the other hyperfine level). If it is returned to the ground electronic state on the carrier, we have generated a Fock state with $n = 1$ and if it is returned on the (lower) sideband corresponding to $\delta n = +1$ again, we generate the Fock state with $n = 2$. In this way they are able to generate any Fock state with high purity and by plotting the probability of making a Raman transition as a function of the length of a subsequent Raman pulse on the blue sideband, they can directly see the Rabi oscillations between these two states. This Rabi frequency is itself a function of n, as expected for the theory of this system:

$$\Omega_{n,n+1} = \sqrt{n+1}\eta\Omega \tag{23}$$

where Ω is the Raman coupling parameter. For this experiment, they find that η, the Lamb-Dicke parameter, has the value 0.202. A coherent state can be generated in several ways, for example by applying a short coherent drive to the ion in the trap using a weak electric field applied across the endcaps. This is of course a forced (or driven) SHO, and the result is a motion with a definite amplitude and phase, limited in uncertainty only by the uncertainty principle, as for any quantum mechanical system. This state can also be expressed as a coherent sum over Fock states of the oscillator, as was verified in the experiment, by the observation of collapses and revivals in the ground state occupation probability after a subsequent Raman pulse was applied.

The NIST group were able to generate a thermal state in simple manner too. All they did was to laser cool the ion *without* using the Raman cooling at all. The state generated by Doppler cooling is a thermal state with a temperature corresponding to the Doppler limit. In their case this gave an average value of $\langle n \rangle$ equal to 1.3 ± 0.1. A squeezed state was also generated, this time by driving the $n = 0$ state with Raman beams having a frequency difference corresponding to $\delta n = 2$. They were able to achieve a squeeze parameter β of 40 ± 10 in this manner, corresponding to a value of $\langle n \rangle$ of roughly 7.1.

Finally, the NIST group were able to generate a *Schrödinger cat state* using the vibrational state of the ion [53]. A Schrödinger cat state is a superposition of *macroscopically* different states of the system (in the original conception it was of course a superposition of "alive" and "dead" states of the cat). In this realisation a complicated series of operations is used to put the ion into a superposition of two electronic states, each having a large amplitude of motion but with a phase shift ϕ between them. We then have in effect two wavepackets in the trap oscillating with different phases. This can be observed by mixing the two states with another Raman pulse. Then these two wavepackets can overlap and interfere. The experiment consists of measuring these interferences as a function of the phase difference ϕ, and the results are striking (see figure 6). The measurements show that the wavepackets are roughly 7 nm wide, and in the largest amplitude motion state generated, they have a maximum separation of 80 nm. In this way they really have been able to create a superposition of macroscopically different states of a single atomic particle in the trap.

7 Ion trap logic gates

7.1 Quantum computing

As discussed earlier, quantum mechanical systems have the unique property that they can be prepared in a quantum superposition of different states. This is not the same as a statistical mixture of states, for which the system is in one state or another, but it is not known which until a measurement is made - only the probability of it being in either state is known. If a system is in a quantum superposition, it is in some sense in both states at the same time. However, when a measurement is made, the system will be projected into one state or the other with probabilities determined by the coefficients of the different states.

One application of this idea is that of *quantum computing*. This is a large subject on its own, and is treated, for example, by Steane [54]. Put simply, the advantage of making a computer out of quantum mechanical logic gates rather than classical ones is that they can then operate on input states which are quantum superpositions of different states, and therefore in effect perform many calculations at the same time, giving an output which is itself a superposition of different output states. However, this does not always help because any measurement of the output state can only yield one piece of information and the rest is lost. The trick is therefore to find problems which can be expressed in such a way that a single measurement of the output *can* be employed to give useful information.

There appear to be a few problems which fall into this category. One is the factorisation of very large numbers, which is an extremely difficult problem on classical computers, taking a very long time to complete. There exist algorithms (see [54]) which can find the factors of large numbers in a much smaller number of steps with a quantum computer than with a classical one. Another possible area of application is in algorithms for the sorting of lists.

However, there are many problems to be solved. One is that clean, well-isolated quantum systems have to be found which can be used to realise a quantum computer. These systems have to be able to exist in a quantum mechanical superposition of states with sufficiently low decoherence rates that many operations can be performed on the states while preserving the coherence. If the coherence is lost, then the quantum superposition turns into a statistical mixture of states, and this loses all the advantages of having a quantum mechanical system. For practical systems this problem appeared to be insurmountable but recently the development of *quantum error correction codes* has meant that quantum computers can be made that correct for errors as they are generated through the sophisticated use of redundancy in the operations [54].

In a quantum computer the information is held in the form of *qubits*. A qubit is a single quantum mechanical piece of information, analogous to a bit in a classical computer, except that now a qubit may exist in any superposition of 1 and 0 rather than just one or the other. In terms of the density matrix of the system, this is represented by the presence of off-diagonal elements in the density matrix. The qubits are used as inputs to *quantum gates*, which operate on one or more qubits to give output states. These gates are similar to the gates in classical computers and complex operations are built up out of a small number of different types of elementary gates. One of the most important of these is the *controlled-NOT* or C-NOT gate. This reverses the state of the *target* qubit (i.e. $|1\rangle \rightarrow |0\rangle$ and $|0\rangle \rightarrow |1\rangle$) if and only if the value of the *control* qubit is $|1\rangle$.

It appears that there are two main ways of realising a quantum computer at present. One is using nuclear magnetic resonance (NMR) techniques on a solid sample. Elementary operations have been performed with this type of system [54], but there is a difficulty in extending the size of such a quantum computer because of signal to noise problems with large numbers of qubits. Also note that this is not strictly a *single* quantum system, but rather a very large ensemble of such systems. The other way of realising a quantum computer is using a string of laser cooled ions in a linear RF trap [55]. This type of trap has the distinct advantage that several ions in a string can all be in the Lamb-Dicke regime, and their positions are therefore fixed, so that they can be individually addressed by laser beams. There may also be possibilities for quantum computing using cavity QED [47].

7.2 Single ion gates

The first suggestion that a quantum computer could be made from a string of ions in an RF linear trap was made by Cirac and Zoller [56]. The suggestion was that each qubit could be represented by the electronic state of an ion in the trap, and that coupling between them could be performed using the centre of mass motion of the string of ions. A procedure was given for the construction of a C-NOT gate using this idea. This used the useful feature that a 2π pulse of radiation on a transition between two states leaves the system in the same state as it started in, except that the *sign* of the wavefunction is reversed. Normally this has no effect, but in this scheme it transforms a superposition of the form $|0\rangle + |1\rangle$ to $|0\rangle - |1\rangle$ where the radiation only couples state $|1\rangle$ to a third state. This is used very ingeniously together with other operations in the gate to transform the target qubit if the control bit is set but to leave it alone if the control bit is not set.

Only one group has managed so far to build a gate using trapped ions [57]. In fact, although based loosely on the ideas in [56], this gate works with a single ion. The two qubits are represented by the electronic state of the ion (in fact, two of the Zeeman-split ground state hyperfine levels of a Be^+ ion - the target qubit) and the first two vibrational states of the ion in the trap potential well (the control qubit). A third atomic level is also needed at some point. With this system, the NIST group have been able to demonstrate that a C-NOT gate can be made to work with high efficiency. The results are shown in figure 7. In this figure, which is like a truth table for the gate, the black bars represent the target qubit while the white bars represent the control qubit. The state of the control bit is unchanged by the gate, as expected (at least 90% of the time) and if the control qubit is zero the target qubit is also unchanged. However, when the control qubit is one, the target qubit is reversed by the gate with high efficiency.

Although the efficiency of this gate is nothing like as good as it would need to be in order to be a part of a quantum computer, it constitutes a convincing demonstration that quantum gates can be built with trapped ions and that the experimental problems can be overcome.

7.3 Future prospects

In their initial experiment, the NIST group also looked at the decoherence rates in the system [57]. They found the rate to be a few kHz, to be compared to the switching speed of 20 kHz. The decoherence rates need to be made lower in order to improve the fidelity of the gate. The problems to be dealt with are discussed in great detail by Wineland *et al* [58]. This paper shows that there are severe problems to be dealt with, and several unknowns, but that the prospects are good for demonstrating real computation with quantum mechanical systems in the next few years.

The NIST group has since begun to work on two-ion systems. A crystal of two ions in a trap has two modes of oscillation along the axis of the crystal: a centre of mass (COM) mode where the ions move together and a stretching mode where the ions move in opposite directions. Here it is of course necessary to cool both modes of motion in order to achieve the Lamb-Dicke regime for the two ions. The NIST group has managed to cool a two-ion crystal to $\langle n_{COM} \rangle = 0.11$ and $\langle n_{stretch} \rangle = 0.01$ [59]. This is an important first step towards the construction of a gate with two ions, which is itself a significant step towards the construction of many-ion gates.

The NIST group have made important measurements of the heating rates in a two-ion crystal. They find that the heating rate for the stretching mode is much lower than that for the centre of mass mode, and suggest that this mode would be the better one to use for quantum gates [59]. The measured rates in this experiment were $\delta\langle n \rangle / \delta t = 20$ ms^{-1} (COM mode) and < 0.18 ms^{-1} (stretch mode).

In future work it is likely that logic gates with more than one ion will soon be constructed but the use of such gates for real computation is a long way away. The most recent work reported by the NIST group is the deterministic entanglement of two trapped ions [60], which represents another important step. At present the most important aspect is the demonstration that the ideas which have been developed for computation with quantum mechanical systems can indeed be realised in practice. Other groups around the world are also working in this area, including Los Alamos [61] and Innsbruck [62].

8 Quantum Zeno Effect

8.1 Original proposal

The quantum Zeno effect was first introduced by Misra and Sudarshan in 1977 [63]. These authors pointed out that, in a quantum mechanical system that decays from one state to another, a measurement taking place during the decay collapses the wave function into either the initial state or the final state. Now in the early stages of the decay the probability of such a measurement resulting in the particle ending up in the final (decayed) state rises *quadratically* with time, even though for longer times the decay proceeds exponentially (which can be approximated as a *linear* decay for short times). Since the effect of a measurement is to destroy all the coherences in the system, in effect it restarts the decay process. Now if the decay is continually put back onto the quadratic part of the curve, the end result is that the overall rate of decay is slowed down from the rate expected from the exponential curve.

This was termed the *quantum Zeno effect* by analogy with the classic paradox due to Zeno about an arrow in flight. The point is that if we make many measurements on a decaying system (such that they probe this early quadratic region) we can slow down the decay. In the limit of continuous observation, we can expect the decay to slow to zero. This is a classic example of the measurement of a quantum mechanical system having an effect on the subsequent behaviour of the system.

The trouble with the observation of this phenomenon for real systems is that the time period over which the decay is quadratic is extremely short and inaccessible experimentally. The length of this period is related to the bandwidth of the available states in the decay, which is extremely large for a normal decay.

8.2 Experimental observation

Itano *et al* [64] performed an experiment in 1990 to observe the quantum Zeno effect. In their experiment they worked on a *driven* radiofrequency transition between hyperfine Zeeman sublevels in a cloud of Be ions in a Penning trap. In this way they avoided the problem of the short length of time available when the probability of a transition rises quadratically. In this case the length of the quadratic period depends on the rate at which the transition is driven, which is under experimental control. Of course the effect is now not quite the same as in the original papers, but the principle is very similar.

Itano *et al* used a π-pulse of radiofrequency radiation to drive the transition, so in the absence of any measurements the probability of the transition taking place is nearly equal to unity. However, if the coherent transition is interrupted by short measurements, they show that the overall transition probability is reduced in both theory and experiment [64]. Here a measurement consists of a brief laser pulse which is tuned to a transition out of *one* of the two hyperfine Zeeman levels involved and could in principle be used to detect whether the ions were in that state (by the observation of a photon) or the other one (by the failure to observe a photon). The laser pulse has to be sufficiently long and intense to ensure that an ion in that state has a high probability of being excited by the laser.

The wave function at time t is given by

$$|\psi(t)\rangle = \cos(\Omega t/2)|i\rangle - i\sin(\Omega t/2)|f\rangle \tag{24}$$

where Ω is the Rabi frequency, and i and f are the initial and final states. For a π-pulse, $\Omega t = \pi$, and this transfers all the population from i to f. If the π-pulse is interrupted at a time $t_1 = \pi/n\Omega$, then we find that the probabilities of the ion being found in the two states are:

$$P_i(t_1) = \cos^2(\pi/2n) \tag{25}$$

$$P_f(t_1) = \sin^2(\pi/2n). \tag{26}$$

At this point the decay restarts and the calculation can be repeated for each of the $(n-1)$ remaining measurement pulses. In practice the state of the ions is only determined in this experiment at the end of the π-pulse, when the fraction of ions that have made the transition is determined by measuring the amount of fluorescence on a cycling transition which yields a signal of many photons per ion in the first few ms of irradiation. The results obtained in this experiment are shown in figure 8.

8.3 Detailed theoretical models

Itano *et al* used a simple wavefunction collapse model to predict the fraction of ions to make the transition as a function of the number of measurement pulses, as discussed above (with some modifications) and the experiment verified these predictions to within experimental error. They concluded that this wavefunction collapse model gave a good description of the process.

However, many theorists were unhappy with this approach, and the publication of these experimental results gave rise to much theoretical work which aimed to give a more detailed and more rigorous description of what was going on in this experiment. A good treatment of this problem is given by Frerichs and Schenzle [65] (see also [66]). They treat the system of 3-level ion, radiofrequency radiation and laser radiation as a single system using a Bloch equation approach. This gives predictions for how intense the laser pulse has to be for it to count as a measurement. This was something that had to be introduced in an *ad hoc* manner in the treatment of Itano *et al.* The Bloch equation approach shows that the effect of the laser pulse is to destroy the coherences (i.e. the off-diagonal elements of the density matrix) which are built up by the radiofrequency radiation. This is the mechanism for the collapse of the wavefunction in this case.

8.4 Proposed experiment

At Imperial College we are in the final stages of building an experiment to test the predictions of these theoretical models. We will perform an experiment which is similar to that of Itano *et al* [64] but we need to perform the experiment on a single ion so that we can have a closer control of the experimental conditions. This will clearly reduce the signal level and the experiment will therefore need to run in a stable manner for a long time to collect data. We will study the variation of the effect of the measurement pulses with the strength of the pulses in an attempt to verify the theoretical predictions for what constitutes a measurement. We will also look at variations in the experiment which should allow us to introduce some decay by coupling to an excited state of the ion in the manner proposed by Plenio *et al* [67]. This will take the experiment closer to the scheme originally proposed by Misra and Sudarshan [63]. Another possibility is to perform a quasi-continuous measurement, as proposed in [68]. The trap for this experiment has been constructed and tested and the experiment is now in the process of being put together. We have already been able to work with single ions in the trap, and we have been able to drive the microwave transition under high resolution. We have also been able to observe a π-pulse of microwaves on this transition. For more details, see [69].

9 Frequency standards and fundamental constants

9.1 Background to frequency standards

Standards are needed for all of the physical units in use. In earlier times, length standards were based on physical artifacts such as a standard metre but this became difficult because the metre could not be reproduced at will at different locations and its accuracy was limited. Therefore this was replaced by an optical wavelength standard based on a highly reproducible optical transition in krypton gas. The metre was then defined to be equal to a certain number of wavelengths of this radiation. This was reproducible to about a part in 10^9 but in the last 15 years even this became inadequate. However, instead of then defining a new length standard based on a well-characterised laser wavelength with better reproducibility and accuracy, it was decided to base the unit of length on that for time and frequency, which was already reproducible to something like a part in 10^{13}. This was done by defining the velocity of light to be a fixed number (exactly 299 792 458 m/s). Then an accurate measurement of the optical frequency of a stable laser also determines its wavelength to the same precision.

The unit of time (and frequency) is defined using the caesium atomic clock, which runs on a microwave transition in caesium atoms at a frequency of around 9 GHz. The frequency of this transition is defined to have a certain fixed value and if the clock is run under certain well-defined conditions the frequency is guaranteed to have this value. The clock can then be used to generate a standard frequency against which other frequencies or time intervals can be calibrated. Standard laser frequencies can then be used as practical wavelength (or optical frequency) standards once their frequency has been determined in a measurement using a so-called frequency chain starting from the Cs clock.

In recent years the current generation of atomic clocks have been pushed to their limits of reproducibility and accuracy, and there is a need for new and more accurate clocks. The pressure for these improvements comes from pure science (e.g. the measurement of fundamental constants and studies in astronomy) and also from technology such as navigation.

9.2 Advantages of ion traps

Ion traps have been seen as having potential for ultra-high resolution spectroscopy and frequency standards for many years (see, for example, [23]). There are several advantages offered by the use of ion traps for frequency standards. First, the ions are located in a very well isolated environment with no collisions with foreign gas molecules or with walls, so it is possible to realise their inherent transition frequencies with high precision and reproducibility. The transitions are not broadened or shifted by external perturbations to a very good approximation.

Second, because ions can be cooled in traps, we can drastically reduce the broadening and shift due to the Doppler effect. This is unlike any other Doppler-free techniques, where only the first-order effect is eliminated. The second order effect (relativistic time dilation) is only reduced if there is a genuine reduction in the kinetic energy of the particles. The use of the Lamb-Dicke regime (section 3.3) is a further help here as the carrier frequency is completely free of the first-order Doppler effect. Although in a conventional RF trap it is only possible to cool a single ion to the Lamb-Dicke regime, an advantage of the linear RF trap is that several ions can be cooled to a string where all of them are free of micromotion and can therefore be in the Lamb-Dicke regime simultaneously. This allows an increase in signal to noise ratio over the use of a single ion.

Third, the interaction time between the ions and the radiation (either laser light or microwaves) can be made very long with trapped ions, and this enables much narrower linewidths to be obtained, as a result of the reduced transit time broadening. As an illustration, note that in a caesium atomic clock the time of flight of an atom through the atomic beam apparatus is of the order of 10 ms, and this sets a limit on the maximum interaction time. The corresponding transit time linewidth is of the order of 100 Hz. In the caesium clock the centre of the transition has to be found to about 1 mHz, which is one part in 10^5 of the linewidth. In an ion trap, the interaction time can be as long as 500 s, giving a linewidth of the order of 1 mHz. Thus even if the centre of this transition can only be found to an accuracy of say 1 % of the linewidth (for a similar transition frequency), this represents a significant improvement over the caesium clock, and systematic errors will be much less of a problem in this case.

There has therefore been much interest in the construction of new frequency standards using ion traps, and this is the driving force behind many of the laboratories involved in ion trap research. Much progress has been made in the development of standards both in the microwave region of the spectrum (especially Hg^+, 40 GHz and Be^+, 300 MHz) and the optical region (e.g. Ba^+, 2 μm and Hg^+, 282 nm). There is a significant prospect that the next generation of frequency standards will be based on transitions observed in trapped ions. It is desirable for these standards to be based on transitions with as high a frequency as possible because then for a given linewidth $\delta\nu$ (limited in general by the interaction time) the *fractional* linewidth will be low, leading to a high quality factor (Q), defined by

$$Q = \frac{\nu_0}{\delta\nu} \tag{27}$$

where ν_0 is the frequency of the transition. The higher Q is, the better the reproducibility and stability of the frequency standard is likely to be.

9.3 Microwave standards

Microwave frequency standards are typically in the region of 1-40 GHz, corresponding to ground state hyperfine transitions in ions having an alkali-like atomic structure (perhaps in the presence of a magnetic field) and the wavelength of these transitions lie in the range 1-30 cm. Therefore the ions are automatically within the Lamb-Dicke regime so long as the size of the ion cloud is of the order of a few mm or less. Therefore the need for laser cooling is not as strong as in the case of optical transitions (see below). However, in order to eliminate the second-order Doppler effect it is still advantageous to employ laser cooling. This is difficult for a large cloud of ions in an RF trap as RF heating prevents low temperatures being achieved. However, in a Penning trap there is no problem in cooling a large cloud of ions.

Mercury is a favourable candidate as it has a large mass, leading to low velocities, and it has a large ground state hyperfine structure splitting in the isotope $^{199}Hg^+$ (roughly 40.5 GHz), leading to a good fractional frequency stability. In fact prototype frequency standards using mercury were being actively worked on very early in the days of ion traps [70]. Experiments have recently been performed by the NIST group [71] and by the JPL group [72]. The NIST experiment [71] uses a linear ion trap with a string of a few ions which are laser cooled using radiation at 193 nm (eight ions were used in this experiment). It achieves a linewidth of 250 mHz (using a Ramsey-type interrogation scheme with an interrogation time of 1.8 s). The most recent results with this experiment report a Ramsey interrogation time of 100 s and, when run as a standard, the stability is reported to be about 4×10^{-15} over a 2 hour period [73]. This trap is run at cryogenic temperatures to keep the effects of collisions with background gas as small as possible. The other experiment [72] uses ions which are not laser cooled, but the signal is high as there is a large number of ions. The fractional second order Doppler shift is significant (of the order of 10^{-12}) but is kept at a constant value by keeping the experimental conditions constant (especially the number of ions in the trap, which is maintained to approximately 0.1 %). The linear trap is able to achieve a smaller second order Doppler shift than a standard RF trap for a fixed number of ions. This experiment achieves a stability of 2×10^{-15} over a period of 24 000 s and is claimed to have an inherent stability of at least 10^{-15}. The linewidth of the microwave resonance is 160 mHz. Most recent results with this experiment are reported in [74].

Beryllium is another element which has been worked with over a long period. The NIST group has recently run their Be^+ experiment as a prototype standard with very impressive performance figures. This experiment uses the technique of *sympathetic cooling* [34] where two different ionic species are held in the same trap at the same time (here

Mg^+ and Be^+). The Mg^+ ions are laser cooled using light at 280 nm from a frequency-doubled dye laser. Through collisions, these ions sympathetically cool the Be^+ ions, and so the temperature of the Be^+ ions is reduced to very low values without them being directly laser cooled. This allows very long interaction times to be used when probing the clock transition, since the ions are not able to heat up during that period, as they otherwise would. As a result, in this experiment they are able to use an interaction time of typically 100 s, giving a linewidth of 5 mHz, and the standard has a stability of 3 $\times 10^{-14}$ over a period of 10 000 s [75]. The final limits on the accuracy of this standard are to do with the background pressure in the vacuum system. Figure 9 shows a lineshape of the microwave resonance used as a frequency standard in Be^+.

These various examples show that ion trap microwave standards are already at a similar level of performance to that of the caesium standard, and are still capable of improvement. Meanwhile, the Cs atomic clock has itself been improving due to the use of the techniques of optical pumping and laser cooling, so there is now a healthy competition between these different approaches to the production of new standards.

9.4 Optical standards

For an optical standard one requires an atomic structure which is very similar to that used in quantum jump experiments, and the standard arrangement is a three-level system with the resonance level and a metastable level both coupled to the ground state by a strong and a weak transition respectively. The lifetime of the metastable level has to be of the order of seconds so that the linewidth of the so-called *clock transition* down to the ground state (which is obviously limited by the natural linewidth of the metastable level) can be of the order of 1 Hz. It is also necessary to have an ultra-stable laser available with a comparable linewidth in order to take advantage of this low natural linewidth, and several laboratories around the world are working towards this level of stabilisation.

Detection is performed by observing the *absence* of many resonance transition photons when the ion is excited into the metastable level by a photon on the weak clock transition. Thus even with a single ion it is possible to detect nearly every transition made on the clock transition due to the use of this electron shelving technique (section 4.4). The two lasers are applied at different times in order to avoid light shifts and broadening of the clock transition. Most of these experiments are performed with a single ion because of the need to achieve the Lamb-Dicke regime and the need to avoid RF heating effects. Thus a miniature RF trap with strong confinement and a single ion is the most common arrangement. Such experiments have been performed in Hg^+, Ba^+, Ca^+, Sr^+ and Yb^+. Most of these experiments have not progressed as far as the microwave standards experiments, due to the technical difficulties involved: tight confinement, strong laser cooling, ultra-stable laser sources etc. However, a linewidth of less than 100 Hz on an optical transition in Hg^+ has been observed by the NIST group [76], giving a Q-value of roughly 10^{13}, with a potential reproducibility of one part in 10^{15} or better. The most recent results from this work are reported in [77].

An extreme example of a potentially very narrow transition is one in Yb^+ which has been studied by the group at NPL in the UK. This transition at 467 nm goes from the $S_{1/2}$ ground state of the ion to a very long lived $F_{7/2}$ state which has a lifetime of several years. The NPL group measured this transition in a single laser cooled ion in a miniature RF trap using a quantum jump technique [78] (see figure 10). The laser with which they probed the ion had a linewidth of 350 kHz, and this is reflected in the linewidth of the observed transition, but the potential linewidth of this transition is very small indeed. The difficulty, of course, is obtaining a sufficient signal from such a weak transition to steer the frequency of an ultra-stable laser source. The Yb^+ ion also has other transitions of interest for optical frequency standards and is being studied by several groups around the world.

Another important ion for this work is indium, as there is a transition in In^+ which has excellent properties for high reproducibility and absence of systematic effects. This is because the clock transition is a $J = 0$ to $J = 0$ transition. However, the laser wavelengths required are difficult to generate and the transition rate on the cooling transition is rather low, leading to difficulties with cooling. This ion is discussed, for example, in [79].

Because the frequency of an optical transition is so much higher than that of a microwave transition, it is likely that in the end the transition used to define the second will be an optical one rather than a microwave one because the potential for high stability is so much better. However, it is likely to be a while before such standards are generally available and microwave ion trap standards are likely to be in use before then. There are of course other systems which may be used as well as ions in traps (particularly laser cooled atoms) but the work with trapped ions is very promising and has already come a long way towards the realisation of the ultimate frequency standard.

9.5 Fundamental constants

Precision determinations of fundamental constants depend on highly accurate and reproducible experiments, so it is not surprising that there are many experiments with trapped ions which are able to improve on earlier values of some of the fundamental constants. In particular, the most accurately known fundamental constant, the g-value of the electron, was determined in a Penning trap (see for example [80]). This experiment uses a comparison of the cyclotron frequency to the spin-flip frequency to measure $(g - 2)$. Similar types of experiments can be used to measure g for hydrogen-like ions [81], muons [82] and electrons bound in ions [83].

Traps are also used for precision mass measurements, such as the proton mass, which can now be determined to 2 parts in 10^{11}, with a potential accuracy of 5 parts in 10^{12} [84]. These measurements may be linked to a possible future redefinition of the mass scale using an atomic mass rather than an artefact as at present. Less accurate measurements (around a part in 10^7) are routinely used for radioactive isotopes, for the measurement of nuclear binding energies [85].

Finally, tests of fundamental physical theories, such as CPT, are possible in traps, as reported in [86]. These tests rely on the accurate measurement of various frequencies in traps, such as cyclotron frequencies, in different systems and at different times.

For many purposes, ion traps are ideal for these types of measurement, because of the good isolation of the trapped particles from various perturbations. We are certain to see much progress in these areas as the techniques continue to advance.

10 Conclusion

These lectures have attempted to demonstrate the wide range of problems in spectroscopy and quantum optics to which ions in traps can be applied. It has not been possible to be comprehensive, and there are inevitably some areas which have had to be left out. However, I hope that it has been made clear that trapped ions have been used for many critical experiments in spectroscopy and have also made many contributions to our understanding of quantum optics in the last twenty years, since laser cooling of trapped ions was first demonstrated. Ion traps offer a uniquely well-controlled and isolated environment for small numbers of atomic particles and the applications, as we have seen, are numerous. The development of new frequency standards is a particularly important area of application of traps, and this drives much of the work with trapped ions. I have every confidence that in the next twenty years they will continue to have a dramatic impact in the areas of physics discussed here and beyond.

Acknowledgements

I wish to thank my colleagues at Imperial College for their help in reading these notes, and the Engineering and Physical Sciences Research Council for their support of our work at Imperial College.

References

[1] L Orozco, in "New Perspectives in Quantum Mechanics", S Hacyan, R Jauregui and R Lopez-Pena, eds, *AIP Conf. Proc* **464**, 27 (1999).

[2] S L Rolston, in "New Perspectives in Quantum Mechanics", S Hacyan, R Jauregui and R Lopez-Pena, eds, *AIP Conf. Proc* **464**, 91 (1999).

[3] R C Thompson, in *Adv At Mol Opt Phys* **31**, 63-136 (1993).

[4] P K Ghosh, *Ion traps*, Oxford University Press (1995).

[5] G Zs K Horvath, R C Thompson and P L Knight, *Contemporary Physics* **38**, 25-48 (1997).

[6] F M Penning, *Physica* **3**, 873-94 (1936).

[7] R C Thompson and D C Wilson, *Z Phys D* **42**, 271-7 (1997).

[8] E Fischer *Z. Phys* **156** 1-26 (1959).

[9] N W McClachlan *Theory and application of Mathieu functions*, Oxford University Press (1947).

[10] R E March and R J Hughes *Quadrupole storage mass spectrometry*, Wiley Interscience, New York (1989).

[11] W Neuhauser, M Hohenstatt, P E Toschek and H Dehmelt, *Phys Rev Lett* **41**, 233-6 (1978).

[12] E C Beaty, *J Appl Phys* **61**, 2118-122 (1987).

[13] C A Schrama, E Peik, W W Smith and H Walther, *Opt Comm* **101**, 21 (1993).

[14] e.g. S R Jefferts, C Monroe, A S Barton and D J Wineland, *IEEE Trans Instrum Meas* **44**, 148-50 (1995).

[15] e.g. M D McGuire, R Petsch and G Werth, *Phys Rev A* **17**, 1999-2004 (1978).

[16] J D Prestage, G J Dick and L Maleki, *J Appl Phys* **66**, 1013-7 (1989).

[17] I Waki, S Kassner, G Birkl and H Walther, *Phys Rev Lett* **68**, 2007-2010 (1992).

[18] G Birkl, S Kassner and H Walther, *Nature* **357**, 310-3 (1992).

[19] R Blümel, C Kappler, W Quint and H Walther, *Phys Rev A*, **40**, 808-23 (1989).

[20] e.g. F Plumelle, M Desaintfuscien, J L Duchene and C Audoin, *Opt Comm*, **34**, 71-6 (1980).

[21] G Werth, *Phys Scr* **T22**, 191 (1988).

[22] W Kalber, G Meisel, J Rink, R C Thompson, J Mod Opt **39**, 335 (1992).

[23] D J Wineland and H Dehmelt *Bull Am Phys Soc* **20**, 637 (1975).

[24] D J Wineland, J Dalibard and C Cohen-Tannoudji *J Opt Soc Am* **9**, 32-42 (1992).

[25] G Birkl, J A Yeazell, R Rückerl and H Walther, *Europhys Lett* **27**, 197 (1994).

[26] W M Itano and D J Wineland, *Phys Rev A* **25**, 35-54 (1982)

[27] W M Itano, L R Brewer, D J Larson and D J Wineland, *Phys Rev A* **38**, 5698-5706 (1988).

[28] D J Heinzen, J J Bollinger, F L Moore, W M Itano and D J Wineland, *Phys Rev Lett* **66**, 2080-3 (1991).

[29] R C Thompson and J Papadimitriou, in preparation (1999).

[30] M A van Eijkelenborg, M E M Storkey, D M Segal and R C Thompson, *Opt Comm* **159**, 169 (1999).

[31] R H Dicke *Phys Rev* **89**, 472-3 (1953).

[32] J C Bergquist, W M Itano and D J Wineland, *Phys Rev A* **36**, 428-30 (1987).

[33] C Monroe, D M Meekhof, B E King, S R Jefferts, W M itano, D J Wineland and P Gould, *Phys Rev Lett* **75**, 4011-4 (1995).

[34] D J Larson, J C Bergquist, J J Bollinger, W M Itano and D J Wineland, *Phys Rev Lett* **57**, 70-3 (1986).

[35] M A van Eijkelenborg, M E M Storkey, D M Segal and R C Thompson, submitted to *Phys Rev A* (1999).

[36] D J Wineland, R E Drullinger and F L Walls, *Phys Rev Lett* **40**, 1639-41 (1978).

[37] W Neuhauser, M Hohenstatt, P E Toschek and H Dehmelt, *Phys Rev A* **22**, 1137-40 (1980).

[38] D J Wineland and W M Itano, *Phys Lett* **82A**, 75-8 (1981).

[39] N Gisin, P L Knight, I C Percival, R C Thompson and D C Wilson, *J Mod Opt* **40**, 1663-71 (1993).

[40] R G Hulet, D J Wineland, J C Bergquist and W M Itano, *Phys Rev A* **37**, 4544-7 (1988).

[41] Th Sauter, R Blatt, W Neuhauser and P E Toschek, *Opt Comm* **60**, 287-92 (1986).

[42] R J Cook and H J Kimble, *Phys Rev Lett* **54**, 1023-6 (1985).

[43] R Blatt and P Zoller *Eur J Phys* **9**, 250-6 (1988).

[44] R J Cook in *Progress in Optics XXVIII* (E Wolf, ed) Elsevier, pp361-416 (1990).

[45] B M Garraway and P L Knight, *Phys Rev A* **49**, 1266-74 (1994).

[46] H Dehmelt, *IEEE Trans Instrum Meas* **IM-31**, 83-7 (1982).

[47] S Haroche, in "New Perspectives in Quantum Mechanics", S Hacyan, R Jauregui and R Lopez-Pena, eds, *AIP Conf. Proc* **464**, 45 (1999).

[48] G M Meyer, H-J Briegel and H Walther, *Europhys Lett* **37**, 317-22 (1997).

[49] M C Teich and B E A Saleh, *Quantum Opt* **1**, 153-91 (1989).

[50] J T Hoeffges, H W Baldauf, T Eichler, S R Helmfrid and H Walther, *Opt Comm* **133**, 170-4 (1997).

[51] F Diedrich and H Walther, *Phys Rev Lett* **59**, 2931-4 (1987).

[52] D M Meekhof, C Monroe, B E King, W M itano and D J Wineland, *Phys Rev Lett* **76**, 1796-9 (1996).

[53] C Monroe, D M Meekhof, B E King and D J Wineland, *Science* **272**, 1131 (1996).

[54] A Steane, *Rep Prog Phys* **61**, 117-73 (1998).

[55] A Steane, *Appl Phys* **B 64**, 623-42 (1997).

[56] J I Cirac and P Zoller, *Phys Rev Lett* **74**, 4091 4 (1995).

[57] C Monroe, D M Meekhof, B E King, W M Itano and D J Wineland, *Phys Rev Lett* **75**, 4714-7 (1995).

[58] D J Wineland, C Monroe, W M Itano, D Liebfried, B King and D M Meekhof, *Journal of Research of the National Institute of Standards and Technology* **103**, 259 (1998).

[59] B E King, C J Wyatt, Q A Turchette, D Liebfried, W M Itano, C Monroe and D J Wineland, *Phys Rev Lett*, in press (1998)(quant-ph/9803023).

[60] Q A Turchette, C S Wood, B E King, C J Myatt, D Liebfried, W M Itano, C Monroe and D J Wineland, preprint (1998)(quant-ph/9806012).

[61] R J Hughes *et al*, *Fort der Physik - Prog of Phys* **46**, 329-61 (1998).

[62] H C Nagerl, W Bechter, J Eschner, F SchmidtKaler and R Blatt, *Appl Phys B* **66**, 603-608 (1998).

[63] B Misra and E C G Sudarshan, *J Math Phys* **18**, 756-63 (1977).

[64] W M Itano, D J Heinzen, J J Bollinger and D J Wineland, *Phys Rev A* **41**, 2295-300 (1990).

[65] V Frerichs and A Schenzle, *Phys Rev A* **44**, 1962-8 (1991).

[66] A Schenzle, *Contemp Phys* **37**, 303-20 (1996).

[67] M Plenio, P L Knight and R C Thompson, *Opt Comm* **123**, 278-86 (1996).

[68] A Beige and G C Hegerfeldt, *J Phys A* **30**, 1323-34 (1997).

[69] R C Thompson, J-L Hernandez-Pozos, J Hoeffges, D M Segal and J R Vincent in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 388 (1999)

[70] M Jardino, M Desaintfuscien, R Barrilet, J Viennet, P Petit and C Audoin, *Appl Phys* **24**, 107-12 (1981).

[71] M G Raizen, J M Gilligan, J C Bergquist, W M Itano and D J Wineland, *J Mod Opt* **39**, 233-42 (1992).

[72] J D Prestage, R L Tjoelker, G J Dick and L Maleki, *J Mod Opt* **39**, 221-232 (1992).

[73] D J Berkeland, J D Miller, J C Bergquist, W M Itano and D J Wineland, *Phys Rev Lett* **80**, 2089-92 (1998).

[74] J D Prestage, R L Tjoelker and L Maleki, in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 357 (1999).

[75] J J Bollinger, D J Heinzen, W M Itano, S L Gilbert and D J Wineland, *IEEE Trans Instrum and Meas*, **40**, 126-8 (1991).

[76] D J Wineland, J C Bergquist, J J Bollinger, W M Itano, D J Heinzen, S L Gilbert, C H Manney and M G Raizen, *IEEE Trans Ultrasonics, Ferroelectrics and Freq Control*, **37**, 515-23 (1990).

[77] B C Young *et al*, in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 337 (1999).

[78] M Roberts, P Taylor, G P Barwood, P Gill, H A Klein, W R C Rowley, *Phys Rev Lett* **78** 1876-9 (1997).

[79] W Nagourney, J Torgerson and H Dehmelt, in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 343 (1999).

[80] R K Mittleman, I I Ioannou and H G Dehmelt, in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 13 (1999).

[81] M Diederich *et al*, in "Trapped Charged Particles and Fundamental Physics", D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 43 (1999).

[82] A Grossmann *et al*, in “Trapped Charged Particles and Fundamental Physics”, D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 52 (1999).

[83] G Marx, G Tommasco and G Werth, in “Trapped Charged Particles and Fundamental Physics”, D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 57 (1999).

[84] R S Van Dyck Jr, D L Farnham, S L Zafonte and P B Schwinberg, in “Trapped Charged Particles and Fundamental Physics”, D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 101 (1999).

[85] G Bollen *et al*, in “Trapped Charged Particles and Fundamental Physics”, D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 111 (1999).

[86] R Bluhm, V A Kostelecky and N Russell, in “Trapped Charged Particles and Fundamental Physics”, D H E Dubin and D Schneider, eds, *AIP Conf Proc* **457**, 133 and 138 (1999) .

Figure captions

Figure 1 Electrodes of ion traps: (a) Paul and Penning traps; (b) linear RF trap.

Figure 2 Stability diagram of a Paul trap. The motion is stable in both axial and radial directions within the shaded area. See the text for the definitions of a ($= a_z$) and q ($= q_z$).

Figure 3 Quantum jumps in a single Ba^+ ion (from [41]). When the ion is in the *off* state there is still a small fluorescence signal due to scattered light off the electrodes of the miniature Paul trap.

Figure 4 Quantum jumps in a single Mg^+ ion (from [39]). This experiment was performed in a Penning trap at a magnetic field of about 1 T. The fluorescence photons were counted in bins of 1 ms duration.

Figure 5 Resonance fluorescence heterodyne spectrum of a single Mg^+ ion in a miniature end-cap trap (from [50]). As expected, the elastically scattered component has a linewidth limited only by the resolution of the detection electronics.

Figure 6 Realisation of a Schrödinger cat state using a single Be^+ ion (from [53]). The plots show the interference between the two macroscopically distinguishable components of the ion motion in the coherent superposition, as a function of the phase difference ϕ between them.

Figure 7 Demonstration of a controlled-NOT (C-NOT or CN) quantum gate using a single trapped ion (see the text for details) (from [57]).

Figure 8 Results of the quantum Zeno experiment (from [64]). Here n is the number of times the π-pulse of microwave radiation is interrupted by the brief measurement pulses. The transition probability is only determined at the end of the π-pulse.

Figure 9 Be^+ microwave resonance at 303 MHz measured on a cloud of ions in a Penning trap (from [75]). The interaction time here was of the order of 500 s.

Figure 10 Spectrum of the 467 nm transition in a single Yb^+ ion (from [78]). The resolution is limited by the laser linewidth, which is 350 kHz. The laser frequency was changed in steps of 250 kHz between measurements.