11 Trapped Ions and Atoms

Experiments with single particles were considered impossible well into the 20^{th} century [225]. This was made possible by early work in the group of Hans Dehmelt, who created traps for charged particles and used them for measurements on single electrons, protons and ions [226].

When the first suggestions were made for specific physical systems as qubits, trapped atomic ions were among them[178]. Atomic ions have some attractive properties for use as qubits: qubits can be defined in ways that make decoherence very slow while simultaneously allowing for readout with high efficiency. To avoid perturbing these ideal properties, it is best to isolate the ions in space. This can be achieved with electromagnetic traps, which arrange electric and magnetic fields in such a way as to create a potential minimum for the ion at a predetermined point in space.

For neutral atoms, which have no charge or static dipole moments, the forces from electromagnetic fields are much weaker, so trapping them is even more difficult. This became possible only with the development of high-intensity lasers, which interact with the transition dipole moments between stationary states.

Apart from trapping of neutral atoms, lasers have a range of other applications, including

- Generating gate operations
- Reading out the results
- Initializing the qubits
- Cooling the motional degrees of freedom
- Trapping neutral atoms.

11.1 Trapping ions

11.1.1 Ions, traps and light

Earnshaw's[1] theorem states that static electromagnetic fields cannot trap a charge in a stable static position[2]. However, using a combination of static and alternating electromagnetic fields makes it possible to confine ions in an effective (time-averaged) potential[228].

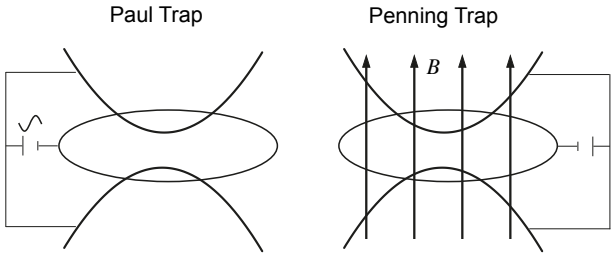

Figure 11.1: Two classical ion traps.

Figure 11.1 shows schematically the geometries used in the two traditional traps, the Paul[3] and Penning[4] traps. Both consist of an axially symmetric set of electrodes. The electrodes on the symmetry axis have the same potential, while the ring has the opposite polarity. The resulting field is roughly that of a quadrupole, where the field vanishes at the center and increases in all directions.

In the case of the Paul trap, the voltage on the electrodes varies sinusoidally: The electrodes

[1] Samuel Earnshaw (1805 - 1888)
[2] In the purely electrostatic case the existence of a minimum of the electrostatic potential in a charge-free region would violate Gauss' law. See [227] for a discussion of Earnshaw's theorem in a modern context.
[3] Wolfgang Paul (1913 – 1993)
[4] Frans Michel Penning (1894 – 1953)

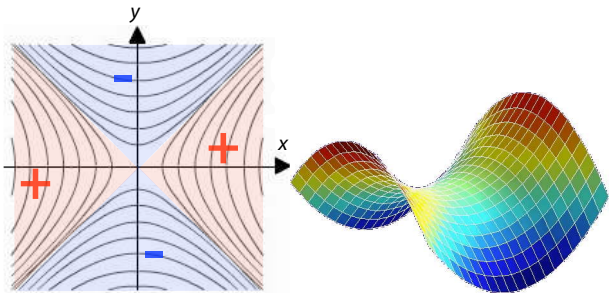

Figure 11.2: Quadrupole potential.

generate a potential

$$\Phi(x, y, t) = (U - V\,\cos(\omega t))\frac{x^2 - y^2}{2r_0^2},$$

where x points towards the ring electrode and y toward the polar end caps. The ion is therefore alternately attracted to the polar end caps or to the ring electrode. If the parameters U, V and ω of the potential are chosen correctly, the average potential over a cycle corresponds to a net force that pushes the ion towards the center of the trap. In the exact center, the field is zero and any deviation results in a net restoring force.

The Penning trap uses the same electrodes, but the electric field is static: it is repulsive for the end caps. The ions are prevented from reaching the ring electrode by a longitudinal magnetic field, which forces the ions onto circular orbits.

11.1.2 Linear traps

The basic Paul and Penning traps are well suited for storing single ions. While it is also possible to store several ions, their arrangement is difficult to control and the Coulomb-repulsion between them forces them out of the field-free center, so they are subject to the time-dependent trapping field. For quantum computing applications, it is desirable to trap larger numbers of ions and to arrange them in such a way that they are easier to address by laser beams. This is achieved by extending the traps in one dimension into so-called linear traps.

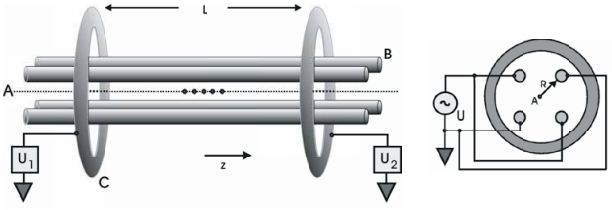

Figure 11.3: Linear quadrupole trap.

Figure 11.3 shows the geometry of a linear Paul trap, which consists of four parallel rods that generate a quadrupole potential in the plane perpendicular to them. The quadrupole potential is alternated at a radio-frequency, and the time-averaged effect on the ions confines them to the symmetry axis of the trap, while they are free to move along this axis. A static potential applied to the end rings surrounding the rods prevents the ions from escaping along the axis. The resulting effective potential (averaged over an rf cycle) can be written as

$$V = \omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2,$$

where ω_α, $\alpha = x, y, z$ are the vibrational frequencies along the three orthogonal axes. By design, one has $\omega_x = \omega_y \gg \omega_z$, i.e., strong confinement perpendicular to the axis and weak confinement parallel to the axis.

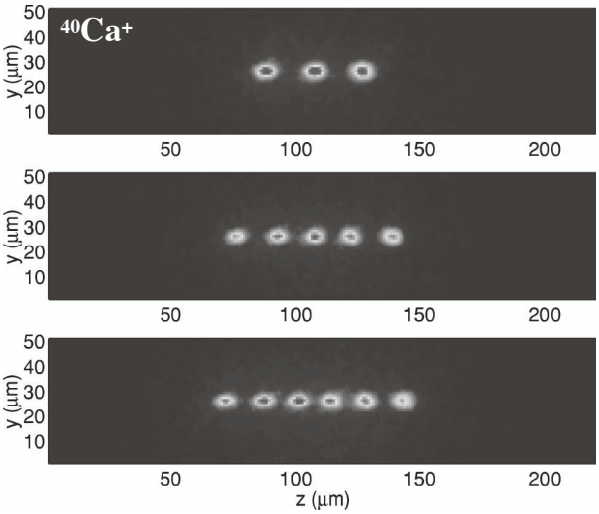

Figure 11.4: Strings of ions in linear traps.

Ions that are placed in such a trap will there-

fore preferentially order along the axis, as shown in figure 11.4. The distance between the ions is determined by the equilibrium between the confining potential $\omega_z^2 z^2$ and the Coulomb repulsion between the ions.

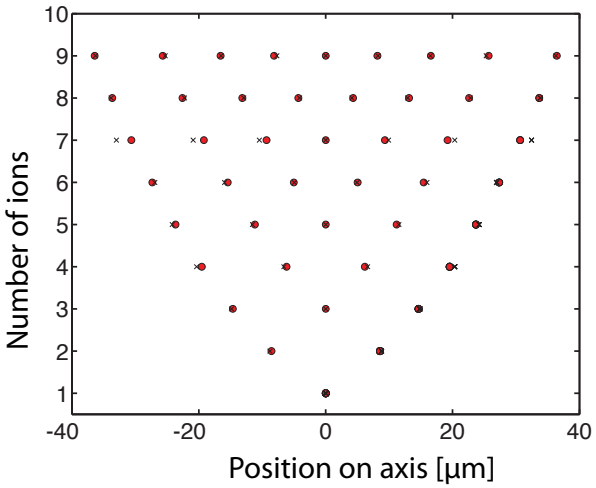

Figure 11.5: Position of multiple trapped ions in linear traps.

Figure 11.5 shows the positions of the ions along the axis as a function of the number of trapped ions. In general, the ions have finite thermal energy, moving around their equilibrium positions. Thes motions can be expanded in eigenmodes, as discussed in section 11.2.2.

This type of trap has two important advantages for quantum computing applications: it allows one to assemble many ions in a linear chain where they can be addressed by laser beams and the equilibrium position of the ions (on the symmetry axis) is field-free. This is in contrast to the conventional Paul trap where the Coulomb repulsion between the ions pushes them away from the field-free point. As a result, two or more ions in a Paul trap perform a micro-motion driven by the rf potential. In the linear Paul trap, the field-free region is a line where a large number of ions can remain in zero field and therefore at rest.

11.2 Interaction with light

The interaction of light with atomic ions is essential for building a quantum computer on the basis of trapped ions: it is used for initializing, gating, readout, and for controlling the motional degrees of freedom. We therefore discuss here some of the basics of the interaction between light and atomic ions. Most of it is formally equivalent to the NMR case, except that the interaction is that of an electric dipole with an electric field and that the frequencies are six orders of magnitude higher (10^{14} Hz rather than 10^8 Hz).

11.2.1 Optical transitions

When light couples to atomic ions, the electric field of the optical wave couples to the atomic electric dipole moment:

$$\mathcal{H}_e = -\vec{E} \cdot \vec{\mu}_e,$$

where $\vec{E}$ is the electric field and $\vec{\mu}_e$ the atomic electric dipole moment. For the purpose of quantum information processing applications, it is important to distinguish between "allowed" and "forbidden" optical transitions. In the first case, the matrix element of the electric dipole moment operator for the transition is of the order of 10^{-29} C m. The order of magnitude of these values is determined by the product

$$\begin{aligned} e\,a_0 &= 1,6 \cdot 10^{-19}\mathrm{C} \cdot 5,3 \cdot 10^{-11}\mathrm{m} \\ &\approx 0,85 \cdot 10^{-29}\mathrm{Cm}. \end{aligned}$$

In the case of "forbidden" transitions, it can be many orders of magnitude smaller. Atoms do not have a permanent electric dipole moment. This implies that the diagonal matrix elements vanish,

$$\langle \Psi | \vec{\mu}_e | \Psi \rangle = 0.$$

The size of the electric dipole moment determines not only the strength of the interaction with the laser field and thus the ease with which

the ion can be optically excited, it also determines the lifetime of the electronically excited states. According to Einstein's theory of absorption and emission, the spontaneous emission rate is proportional to the square of the matrix element. States that have an optically allowed transition to a lower lying state therefore decay on a timescale of $\approx 10^{-8}$ s. Since the decay destroys any stored information, this makes such states unsuitable for use in quantum computers.

While an atom has an infinite number of energy levels, it is often sufficient to consider a pair of states to discuss, e.g., the interaction with light. Writing $|g\rangle$ for the state with the lower energy (usually the ground state) and $|e\rangle$ for the higher state, the relevant Hamiltonian can then be written as

$$\mathcal{H}_{2LS} = -\omega_0 \mathbf{S}_z - 2\omega_1 \cos(\omega t)\mathbf{S}_x.$$

Here, $\hbar\omega_0 = \mathcal{E}_e - \mathcal{E}_g$ is the energy difference between the ground and excited state and $2\omega_1 \cos(\omega t)$ is the coupling between the laser field (with frequency ω) and the atomic dipole moment. The operators $\mathbf{S}_x$ and $\mathbf{S}_z$ are pseudo-spin-1/2 operators as defined in section 4.3.1.

If the Hamiltonian is written in this way, the analogy to the real spin-1/2 system discussed in Chapter 10, is obvious. This allows us to treat two-level transitions as virtual spins-1/2 [47]. In the interaction representation with respect to the laser frequency, the coordinate system "rotates" at the laser frequency ω around the z-axis of the virtual spin. Neglecting the counter-rotating component at frequency 2ω, we get the effective Hamiltonian

$$\mathcal{H}^r_{2LS} = -(\omega_0 - \omega)\mathbf{S}_z - \omega_1 \mathbf{S}_x, \tag{11.1}$$

in close analogy to the rotating frame representation of magnetic resonance (see section 5.4). In optics, this is known as the *rotating wave approximation*.

11.2.2 Motional effects

When an atom is not at rest, its transition frequency is shifted through the Doppler effect:

$$\omega_0^D = \omega_0 + \vec{k} \cdot \vec{v},$$

where $\vec{k}$ is the wave vector of the laser field and $\vec{v}$ the atomic velocity. In free atoms, the velocity can have arbitrary values, with the probability of a specific velocity determined by the Boltzmann distribution. The optical spectra of ensembles of atoms are therefore broadened and/or shifted according to their motional state.

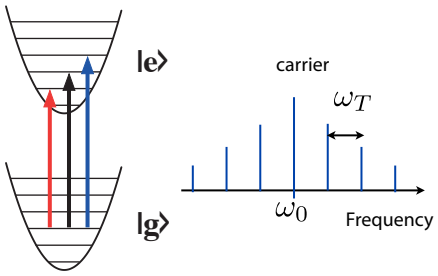

Figure 11.6: Energy levels of the trapped atom (left) and the resulting spectrum (right).

In trapped ions, the motional energy is quantized. Depending on the trap potential, the motional states can often be approximated by a collection of harmonic oscillators. The lowest mode is always the center of mass motion of the full system, in analogy to the motion of atoms in a crystal. A change of the fundamental vibrational mode can be compared to the Mössbauer effect, where the recoil from the photon is shared between all atoms in the crystal.

Harmonic oscillator motion does not shift the frequency by arbitrary amounts, but creates sidebands that are separated from the carrier frequency ω_0 by the harmonic oscillator frequency. As shown in Figure 11.6, the trap motion creates a set of sidebands whose frequencies can be written as $\omega_n = \omega_0 + n\omega_T$, where $-\infty < n < \infty$ is the order of the sideband and ω_T is the trap frequency.

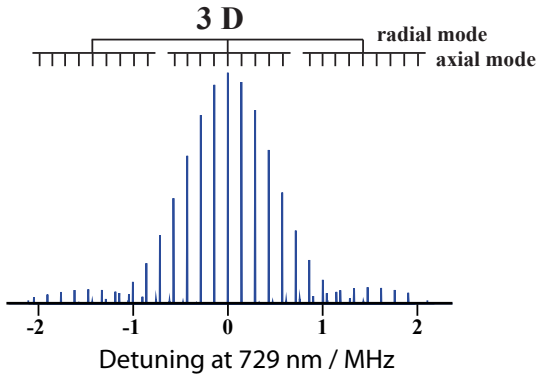

Figure 11.7: Sideband pattern for axial and radial modes.

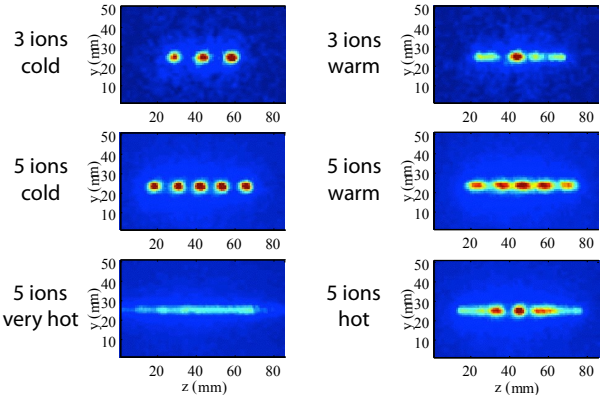

Figure 11.9: Thermal excitation of multiple trapped ions.

Since every motional degree of freedom creates such a sideband pattern, the resulting spectrum can contain a large number of resonance lines.

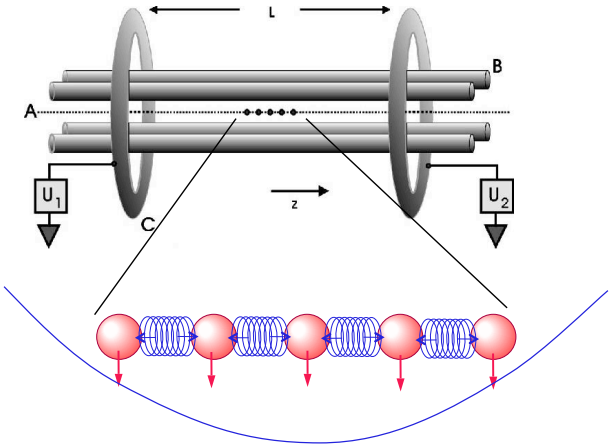

Figure 11.8: Coupled motions of multiple trapped ions.

In the case of many trapped ions (see fig. 11.8), their motion is coupled through the Coulomb-repulsion between them. The number of possible modes increases with the number of trapped ions - in close analogy to phononic modes in crystals. If we consider only axial motion, the number of modes is equal to the number of ions.

Trapped ions are normally not in the motional ground state, but carry thermal excitations. Depending on the amount of thermal excitation, the position of the ions can have significant uncertainty. Most implementations of trapped ion quantum computers require that the ions be in or close to the motional ground state.

In all techniques suggested to date for quantum computing with trapped ions, the spatial coordinates of the qubit ions play an important role either as a qubit or as a variable used for coupling different qubits. If the spatial degrees of freedom are used in the computation, the motional state of the ion must be well controlled and initialized to a specific state, which is usually the motional ground state. The ions must therefore be cooled into their ground state as a part of the initialization process [229].

11.2.3 Force from lasers

The technique to bring them into the ground state is laser cooling, which was developed in the 1980's [230, 231, 232, 233, 234, 235]. It relies on the transfer of linear momentum from photons to atoms during an absorption (and emission) process. Suitable arrangements allow one to use this momentum transfer to create extremely strong forces that push the atoms in the direction of the laser beam. Adjusting the experimental parameters properly, these forces can be conservative (i.e., they form a potential) or they can be dissipative friction forces. Conservative forces are useful for logical gate operations, while frictional forces are useful for initialization and cooling.

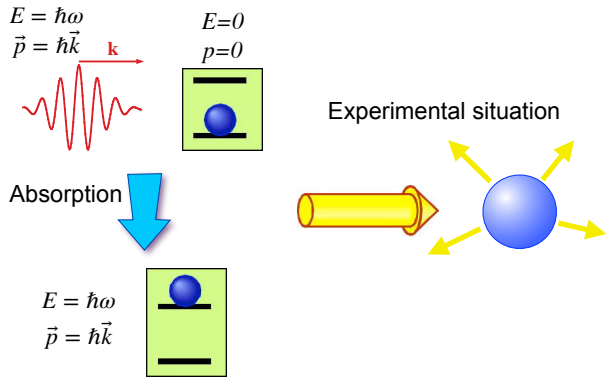

Figure 11.10: Photon momentum as the source of mechanical effects of light.

The origin of these mechanical effects of light can be traced to the momentum $\hbar k$ that every photon carries. As shown in Figure 11.10, the photon momentum is transferred to the atom whenever a photon is absorbed. During the subsequent spontaneous emission process, the recoil of the photon emission also contributes to the mechanical effects of the light on the atom. However, the emission is, in contrast to the absorption process, not directed. The average effect of all emission processes therefore vanishes.

The momentum change due to the transfer of a single photon momentum is relatively small; it corresponds to a change in the atomic velocity of a few cm/s - much less than the thermal velocity of typically several 100 m/s. As an example, we calculate the momentum transferred by a single photon at a wavelength of 589 nm, a prominent wavelength in the spectrum of Na:

$$\begin{aligned}\Delta p &= \frac{h}{\lambda} = \frac{6.626 \cdot 10^{-34}\text{Js}}{589 \cdot 10^{-9}\text{m}} \\ &= 1.125 \cdot 10^{-27}\frac{\text{m kg}}{\text{s}}.\end{aligned}$$

Given the mass $m_{\text{Na}} = 3.818 \cdot 10^{-26}$ kg of the sodium atom, this corresponds to a change in its velocity of

$$\Delta v = \frac{\Delta p}{m_{\text{Na}}} = 2.95\frac{\text{cm}}{\text{s}}.$$

This many order of magnitude smaller than the thermal velocity of an atom, which is of the order of several hundred m/s. This estimate was first made by Einstein in 1917 [236] and verified experimentally by Frisch 1933 [237] with a classical light source. Since the atoms scattered less than three photons in his experiment, the effect was very small. Frisch therefore collimated an atomic beam such that its velocity component in the direction of the photonic momentum was reduced to less than 1 m/s.

The situation changes if an allowed atomic transition is excited by a laser, that supplies a large number of photons. The atom can then absorb and re-emit the photon within a few nanoseconds (16 ns for Na). It can therefore scatter almost 10^8 photons per second, and the momentum transferred by them adds up to a force

$$\begin{aligned}F &= \frac{\Delta p}{\tau} = \frac{1.125 \cdot 10^{-27}\frac{\text{m kg}}{\text{s}}}{16\text{ ns}} \\ &= 7.03 \cdot 10^{-20}\text{N},\end{aligned}$$

corresponding to an acceleration of

$$\begin{aligned}a &= \frac{F}{m_{\text{Na}}} = \frac{7.03 \cdot 10^{-20}\text{N}}{3.82 \cdot 10^{-26}\text{kg}} \\ &= 1.84 \cdot 10^{6}\frac{\text{m}}{\text{s}^2} = 188\,000\,g.\end{aligned}$$

This implies that an atom arriving with the velocity of a jet plane can be stopped over a distance of a few centimeters.

11.2.4 Laser cooling

In the case of trapped ions, the situation may also be discussed in terms of resolved motional sidebands. Cooling is then achieved by irradiating the lower-frequency sidebands, as shown in Figure 11.11. Since the photon energy is smaller than the atomic excitation energy, the difference in energy is supplied by the atom from its kinetic energy: While the atom is excited, the quanum number of the motion is reduced by one unit.

In reality, the laser drives not only the $|g,3\rangle \leftrightarrow |e,2\rangle$ transition, but all $|g,n\rangle \leftrightarrow |e,n-1\rangle$ transitions for $n > 0$. For each absorption event, the

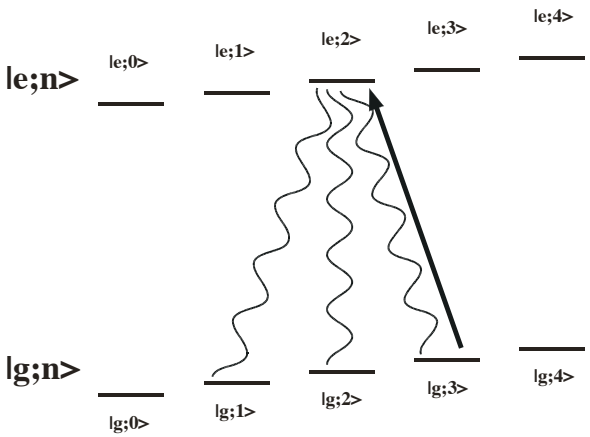

Figure 11.11: Schematics of sideband cooling for a single degree of freedom.

vibrational quantum number is reduced by one unit, since the photon energy is smaller than the energy difference of the two internal states. The emission process occurs with roughly equal probabilities into the different ground states, thus not affecting the average vibrational energy. The only state that is not coupled to the laser is the $|g, 0\rangle$ state, since no transition with a frequency below the carrier originates from this state. As a result, all atoms eventually are driven into this state in the absence of heating mechanisms.

11.3 Quantum information processing with trapped ions

Cold trapped ions were among the first candidates for qubits (see, e.g., [238]), but it took several years of intense experimental work to realize this potential [239].

11.3.1 Metastable qubits

Since the atomic ions stored in traps have a large number of states, there are many distinct possibilities for defining qubits. However, most states are not suitable for storing information. The most important reason is that spontaneous decay times through allowed transitions are of the order of a few nanoseconds. This violates the requirement of long decoherence times. Accordingly, both states of the qubits must either be sublevels of the electronic ground state or metastable states, i.e., states where all transitions to lower lying states are "forbidden".

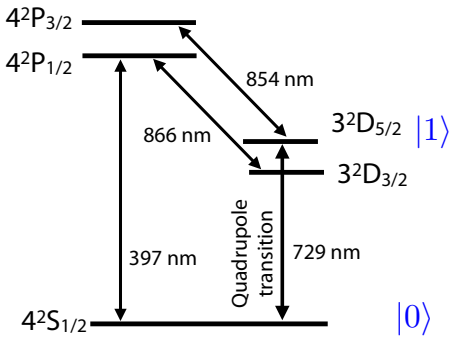

Figure 11.12: Possible qubit implementation using a metastable state in Ca^+.

A typical example of a a qubit implementation is the Ca^+ ion [240]. Figure 11.12 shows its level scheme. In its ground state [Ar](4s), the single valence electron is in the 4s orbital, which is abbreviated by the term symbol $4^2S_{1/2}$. If the electron is excited into a 3d orbital, it has angular momentum $L = 2$, and can only decay to the ground state by emitting two quanta of angular momentum. These quadrupole transitions are "forbidden" in the dipole approximation, resulting in long lifetimes of the excited state. Nägerl *et al.* [241] therefore suggested using the transition between the $4^2S_{1/2}$ ground state and the $3^2D_{5/2}$ excited state as a qubit.

Apart from the computational basis states, ions have many other states that cannot be completely omitted. In particular, the $3^2D_{3/2}$ state is important, since it can be populated and also has

a long lifetime. To bring it back into the qubit system, the 866 nm transition to the $4^2P_{1/2}$ state can be driven with an additional laser. From there, the ions quickly decay to the ground state.

11.3.2 Hyperfine qubits

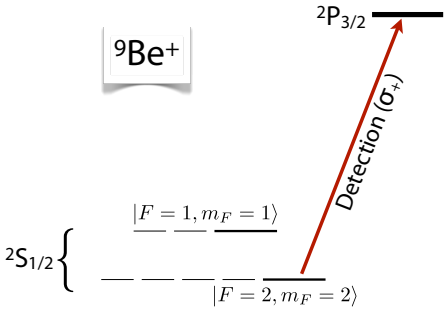

Figure 11.13: Possible qubit implementation using two hyperfine states of $^9Be^+$.

The second common choice is to encode the quantum information in sublevels of the electronic ground state [242, 243]. Figure 11.13 shows as an example the possible encoding of a qubit in the hyperfine levels of the electronic ground state of Be^+. The two qubit states correspond to the $|F=2, m_F=2\rangle$ and $|F=1, m_F=1\rangle$ hyperfine states. Since the spontaneous transition rate between ground states is very small, the lifetime is again long compared to all relevant timescales.

The transitions from the two ground state hyperfine levels to the electronically excited state $^2P_{1/2}$ are sufficiently well resolved to allow one to optically distinguish whether the ion is in the $|2,2\rangle$ or $|1,1\rangle$ state. This fulfills the most important requirement for the readout process.

The initialization of the qubits must bring the ion into a specific internal state as well as into the motional ground state. While the laser cooling for the initialization of the external state was described above, the initialization of the internal state can be achieved by optical pumping. The principle of optical pumping is very similar to sideband cooling: a laser drives the system in such a way that only the desired state of the ion

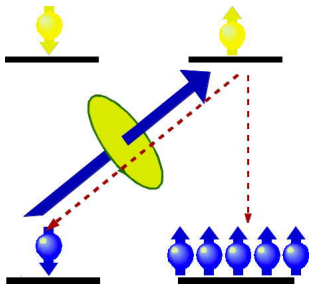

Figure 11.14: Initialization of a hyperfine qubit by optical pumping.

does not couple to the laser, while ions in other states can absorb light, become excited and return to an arbitrary sublevel of the ground state. These absorption / emission cycles are repeated until the ion falls into the state that does not couple. Given enough time, all ions will therefore assemble in the uncoupled state. In this case, the dissipative process that is required for the initialization step is spontaneous emission.

11.3.3 Single-qubit gates

The way to generate (pseudo-)spin rotations that correspond to single qubit gates depends on the specific choice of the qubit states. If the two states encoding the qubit are connected by an optical transition, it is possible to apply laser pulses that have the same effect as RF pulses acting on spin qubits. The corresponding Hamiltonian (11.1) has the same structure as that of a spin-1/2. Since the spatial separation of the ions is typically of the order of 10 optical wavelengths, it is possible to use tightly focused laser beams aimed at individual ions to separately address the qubits [11]. While the optical transitions used for such qubits must be "forbidden", the tightly focused laser beams that are required for addressing qubits individually provide sufficiently high Rabi frequencies for efficient excitation.

If the qubit is defined by two hyperfine states that are connected by a magnetic dipole transition, the situation is even more directly related to magnetic resonance. In this case, the tran-

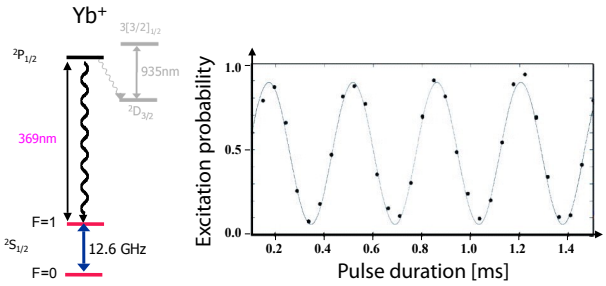

Figure 11.15: Effect of a microwave pulse on a hyperfine qubit.

sition between the two qubit states is a magnetic dipole transition, which can be driven by microwave fields [244]. As shown in figure 11.15, this results in the usual Rabi oscillations. Since the wavelength of microwave radiation is large compared to the distance between the ions, microwaves will interact with all qubits simultaneously. Addressing of individual qubits therefore requires a magnetic field gradient to separate the transition frequencies of the ions.

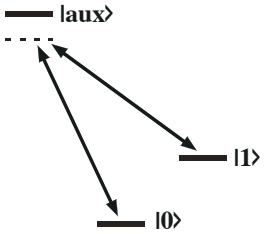

Figure 11.16: Raman excitation of a hyperfine qubit.

The second possibility for driving transitions of hyperfine qubits is to use Raman laser pulses. For this purpose, one uses two laser fields [245], whose frequency difference matches the energy level separation of the two qubit states, as shown in figure 11.16. The laser frequency is close to a transition to an auxiliary state. Choosing an appropriate set of parameters (frequencies, field strengths), it is possible to generate laser pulses that effectively drive the transition between the two qubit states, with negligible excitation of the auxiliary state [242]. In many cases, one uses a frequency-swept set of laser pulses, which can make the operation more robust.

11.3.4 Two-qubit gates

Two-qubit gates that can form the basis of a universal quantum computer, require, in addition to the single-qubit operations, an interaction between qubits. In the case of trapped ions, the main interaction is the Coulomb repulsion between neighboring ions, which are separated by a few micrometers in typical traps. In the qubits described so far, the Coulomb interaction affects all qubit states in the same way and it is therefore not suitable for driving a gate operation directly. Nevertheless, it can be utilized for two-qubit operations in different ways, depending on the qubit implementation.

The Coulomb repulsion between the ions couples their motional degrees of freedom. As in a solid, the motion of ions in a trap is best described in terms of eigenmodes that involve all ions. This quantized motion is often involved in quantum information processing.

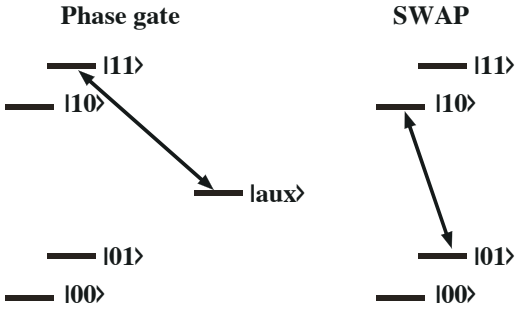

Figure 11.17: Selective laser pulse to generate a phase shift of state $|11\rangle$ (left) and a SWAP operation (right).

We first discuss a two-qubit gate that uses the internal degrees of freedom of a $^9Be^+$ ion as the target qubit and the harmonic oscillator motion as the control qubit of a CNOT gate [242]. Figure 11.17 shows two examples of simple two-qubit gates that can be realized by such a scheme. The notation $|\alpha\beta\rangle$ refers to the internal state α and the motional state β.

In the first example, resonant radiation that couples only the state $|11\rangle$ to an auxiliary state executes a 2π pulse. As in any two-level system,

the system consisting of $|11\rangle$ and $|\text{aux}\rangle$ acquires a phase $e^{i\pi} = -1$ by the pulse. Since the other states are not affected, the overall effect of the pulse on the computational basis states is

$$P_4 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

This phase gate can be combined with two $\pi/2$ pulses into a CNOT operation [242].

Another important two-qubit gate, the SWAP operation, can be generated by a π pulse to the transition $|01\rangle \leftrightarrow |10\rangle$ (see Figure 11.17). In one of the early implementations, where one qubit was represented by an ion and the other by its motional states, this corresponded to a pulse on the red sideband.

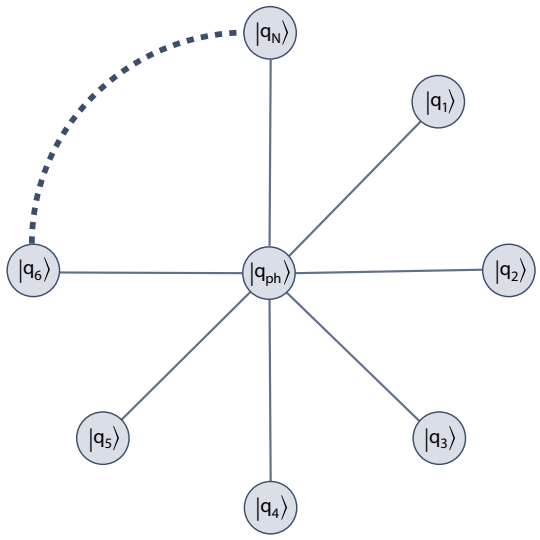

Figure 11.18: The motional degree of freedom $|q_{ph}\rangle$ acts as a bus qubit that connects to all computational qubits $|q_i\rangle$.

While motional degrees of freedom are not ideal as actual qubits, they appear to be useful for executing two-qubit gates between ions. As shown in figure 11.18, the motion couples to every computational qubit. At can therefore serve as an intermediary between two computational qubits. Their function is similar to that in a bus of a classical computer, which connects all parts of the system with each other. In the case of trapped ions, every ion is coupled to the eigenmodes of the oscillators. This possibility has been explored widely [246] [178].

A two-qubit gate between ions j and k can therefore be executed by first swapping the information from ion j into the oscillator mode, executing the two-qubit gate between oscillator and ion k, as described above, and subsequently swapping the information from the oscillator back to ion j. Since the harmonic oscillator motion involves all ions, this procedure works for any pair of ions, irrespective of their distance.

11.3.5 Readout

One of the important advantages of trapped ion quantum computers is the possibility of optically reading out the result with a very high selectivity and success probability. A photon from a laser focused to an ion and tuned to an allowed optical transition is absorbed with almost 50% probability and the photon is re-emitted after typically 10 ns. If the photon is collected and sent to an appropriate detector, such as an avalanche photo diode, it can be detected with a probability of $> 90\%$, depending on the wavelength.

Since the collection efficiency of the detection system is only a few percent, however, this is still not sufficient for high-fidelity readout. It is therefore necessary to repeat the absorption-emission process several thousand times to obtain an unambiguous signature of the state of the qubit. These repetitions must be performed without changing the state of the qubit. This can be achieved if the laser frequency is tuned to an optical cycling transition from the state that is to be detected, focuses it on the ion to be measured, and detects the fluorescence emitted.

The term “cycling transition” means that the state to which the ion is excited can only fall back to the particular ground state from which it was excited. Figure 11.19 shows an example of such a cycling transition between an electronic state with total angular momentum $F = 1$

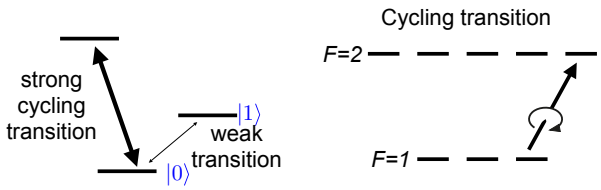

Figure 11.19: Optical readout of a single qubit: the left-hand part shows the relevant states and transitions, the right-hand part an example of a cycling transition.

and a second state with $F = 2$. If circularly polarized light couples to the $|F = 1, m_F = 1\rangle$ ground state, it excites the atom into the $|F = 2, m_F = 2\rangle$ excited state. The selection rule $\Delta m_F = \pm 1$ does not allow for transitions to any other state than the $|F = 1, m_F = 1\rangle$ ground state.

For suitable transitions, up to 10^8 photons per second can be scattered. If the detection system has a 1% collection efficiency, this yields a very reliable decision whether the ion is in the particular state or not.

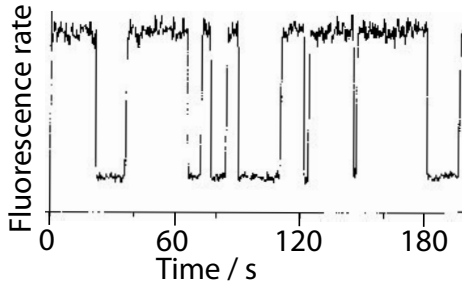

Figure 11.20: Fluorescence of a single Ba ion. The quantum jumps indicate changes of the internal quantum state of the ion [247].

Figure 11.20 shows an example for an observed signal [247]: when the single Ba ion is in the observed state, it scatters approximately 2200 photons per second; the background rate is less than 500 photons per second. As shown in the example data, the fluorescence level is an excellent indicator if the ion is in the state that is being measured. The sudden drops in the fluorescence level indicate that the ion jumps into a different state, which is not coupled to the transition being irradiated. These transitions are referred to as "quantum jumps".

The detection scheme sketched here only provides a measure of the atom being in state $|0\rangle$; a similar measurement of state $|1\rangle$ is only possible if that state is also part of a cycling transition. The complementary measurement of the atom being in state $|1\rangle$ can be achieved in different ways. The first possibility is to take the absence of a result for the state $|0\rangle$ measurement as a measurement of the atom being in state $|1\rangle$. This is possible since the system (under ideal conditions) *must* be either in state $|0\rangle$ or state $|1\rangle$. A second possibility is to perform first the measurement of state $|0\rangle$ and then apply a logical NOT operation and a second measurement of state $|0\rangle$. Since the NOT operation interchanges the two states, a subsequent measurement of the state $|0\rangle$ is logically equivalent to a measurement of state $|1\rangle$ before the NOT operation.

11.4 Experimental implementations

11.4.1 Systems

One of the most popular ions for quantum information studies is the Ca^+ ion [241, 134]. Figure 11.12 shows a possible identification of qubit states. For laser cooling, excitation of resonance fluorescence and optical pumping of the ground state, different transitions are used. The experiment therefore requires laser sources at the wavelengths 397 nm, 866 nm, and 854 nm. If the E2 transition between the ground state and the metastable $D_{5/2}$ state is used as the qubit, a fourth laser with a wavelength of 729 nm is required. Its frequency stability must be better than 1 kHz.

The long lifetimes make hyperfine ground states very attractive for quantum information process-

ing applications. Examples for such systems are the $^{171}Yb^{+}$[244] and $^{9}Be^{+}$ ions [245].

The linear Paul trap was mostly used for quantum information processing, but some variants are also being tested. Tight confinement of the ions is advantageous as it increases the separation between the vibrational levels and therefore facilitates cooling into the motional ground state. In addition, the vibrational frequencies are involved in the logical operations. Accordingly higher vibrational frequencies imply faster clocks.

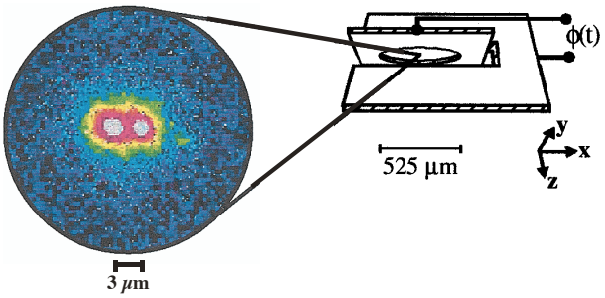

Figure 11.21: Two ions in a small elliptical trap [229].

Tight confinement can be achieved mainly by miniaturization of the traps. For the example shown in Figure 11.21, the smallest trapping frequency is 8.6 MHz [229]. However, miniaturization is not without difficulties: it increases, e.g., the effect of uncontrolled surface charges in the trap and it makes addressing of the ions more difficult.

11.4.2 Initial results

The earliest quantum logic operation was reported by the group of Wineland [242]. They used a $^{9}Be^{+}$ ion where one of the qubits was a pair of internal states, two hyperfine sublevels of the electronic ground state, the $|F = 2, m_F = 2\rangle$ and $|F = 1, m_F = 1\rangle$ states with an energy difference of 1.25 GHz. This qubit represented the target qubit. The control qubit was defined by the two lowest harmonic oscillator states, which were separated by 11 MHz. A sequence of three Raman pulses was used to implement a CNOT gate.

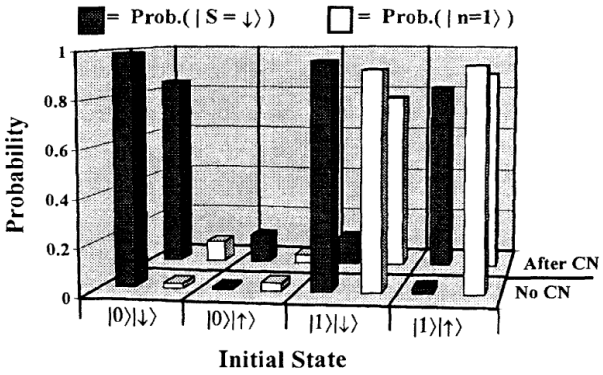

Figure 11.22: Experimental test of the CNOT gate on single $^{9}Be^{+}$ ion [242].

Figure 11.22 shows the populations of the four possible states of the system before (front row) and after (back row) the application of the CNOT gate. The black bars indicate the measured probabilities of finding the system in the internal ground state ↓, the white bars the probability of finding them in the motional ground state $n = 0$. These probabilities were obtained by collecting fluorescence while irradiating with a laser tuned to a cycling transition. The resulting photon count is much higher for the internal ground state.

To measure the motional state, they applied a laser pulse on the red sideband to SWAP the internal and external state qubits and performing another measurement on the cycling transition. The control qubit (the motional state), which is shown in white, does not change during the CNOT operation. The target qubit (internal state), shown in black, remains also roughly constant when the control qubit is in the $|0\rangle$ state (shown in the first two columns) but changes when the control is 1 (3^{rd} and 4^{th} column).

Other achievements with this system include cooling of two ions into the vibrational ground state and their entanglement [229, 245]. For this purpose the authors did not address the ions individually, but modified the effective Rabi

frequency through fine-tuning of their micromotion. The resulting state was not a singlet state (but close to it) and the scheme is not directly applicable to quantum computing.

Using Ca^+ ions in a linear trap, optical addressing of individual ions was demonstrated [248], and in a chain of three ions, coherent excitation [249].

11.4.3 Two-qubit operations

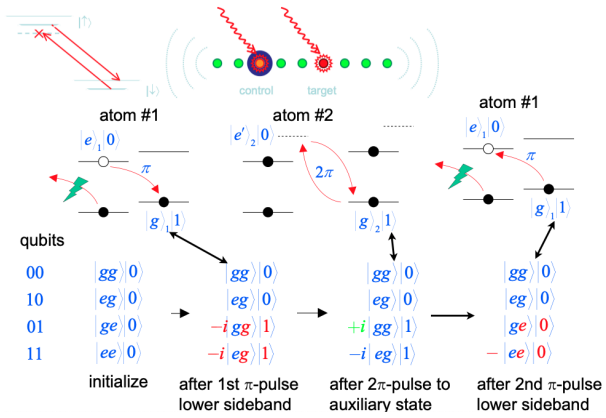

Figure 11.23: The 2-qubit Cirac-Zoller gate.

The two-qubit Cirac–Zoller gate[178] is a popular method for implementing 2-qubit gate operations. As shown in figure 11.23, it uses a sequence of 2 π and one 2π pulse selectively applied to the two targeted ions. The first and the last pulse are tuned to the lower sideband of the first atom, while the second pulse drives the transition between the state $|g,1\rangle$ state of ion 2 and an auxiliary state. Here, $|g\rangle$ refers to the internal state of ion 2 and $|1\rangle$ to the motional state of the system. The 2π pulse selectively inverts the state $|gg,1\rangle$. The full sequence implements a controlled phase gate on the 2-ion system, while the motional state is brought back to the ground state.

Figure 11.24 shows an early implementation of this scheme on two trapped Ca^+ ions [11]. The authors used a laser tuned to a blue-shifted sideband, where, in addition to the electronic transition of the given ion, the collective motion

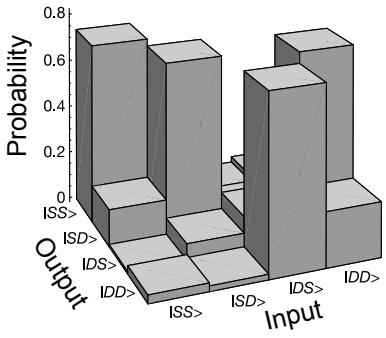

Figure 11.24: CNOT gate implemented on two trapped Ca^+ ions [11].

of the two ions was also excited. To obtain a CNOT gate, they combined the CPhase gate with single-qubit gates, which were realized by a laser beam whose frequency was resonant with the quadrupole transition and which was focused so tightly that it interacted only with a single ion. The final state was measured by exciting the S–P transition of the trapped ions and measuring the fluorescence. Since the ions can only be excited when they are in the S state, high fluorescence counts are indicative of the qubit being in the $|0\rangle$ state.

A two-qubit gate has also been implemented on two trapped beryllium ions by Leibfried *et al.* [12]. They used two hyperfine states of the electronic ground state to store the quantum information. In this experiment, the motion of the ions was excited by two counter-propagating laser beams, whose frequencies differed by 6.1 MHz. As a result, the ions experience a time-dependent effective potential that resonantly excites the oscillatory motion in the trap. The parameters of the excitation were chosen such that the ions were not directly excited, but instead their quantum states were transported around a closed loop in parameter space. As shown by Berry [250], the parameters of such a circuit can be chosen in a way that the transported states acquire a net phase. Leibfried *et al.* used this procedure to implement a phase gate on their system. Since the laser beams interact with both ions, additional lasers will be required for generating specific single-qubit gates in this system.

In a similar system, a Grover-type search was implemented [251].

The number of trapped and controlled ions has increases significantly over the years. Using internal as well as motional degrees of freedom, Monz et al were able to entangle up to 14 qubits [252] and later up to 50 qubits [253].

11.4.4 Challenges

One of the biggest problems of ion traps is that the ions, as charged particles, are relatively sensitive to stray fields in the vicinity. These fields can adversely affect the motion of the ions and, if they are time dependent, they heat the ions. Typical heating times are of the order of 1 ms [229] for two ions in a trap. With increasing numbers of ions, heating rates tend to increase so that not only the number of particles that couple to these stray fields, but also the number of degrees of freedom that can be driven, increases.

Some of the gate operations require that the atoms be cooled into the motional ground state. This can not be achieved by simple Doppler cooling but requires, in addition, sideband cooling. Depending on the type of trap and gate operation, the cooling procedure has to be applied to all (or only specific) motional modes.

Like all other implementations of quantum computers, ion traps will have to demonstrate that they can perform a sufficiently large number of gate operations. As the number of ions in a trap increases, it becomes more and more difficult to control the ions. In particular, the trap frequency (i.e. the confinement) decreases, while the number of motional modes increases and heating effects become more effective. It appears thus unlikely that individual traps will be able to accept a sufficiently large number (i.e., hundreds) of ions.

Several solutions to this problem have been proposed, such as storing the ions in multiple traps. It has been suggested [255] that it should be possible to couple these separate traps through photons, thus creating an arbitrarily large quantum

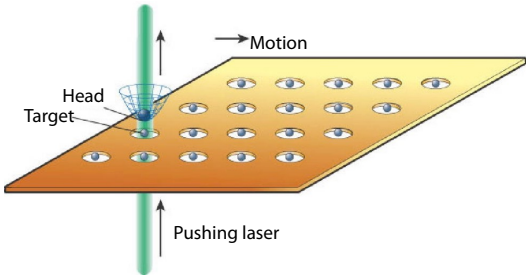

Figure 11.25: A proposed system of individual microtraps [254].

register with a linear overhead. As a first step towards this goal, quantum interference between two remote trapped $^{174}Yb^{+}$ ions was reported [256].

11.4.5 Scalability

While these demonstration experiments were done on a small number of ions, proposals exist how the number of ions could be scaled up, particularly by integrating the trap electrons on a chip [169]. Operation of such a micro-fabricated trap was first demonstrated for a single ion [257].

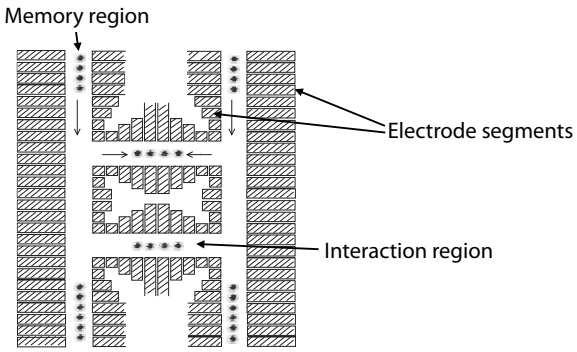

Figure 11.26: Proposed architecture for a large-scale ion-trap quantum computer [169].

The current generation of trapped-ion quantum computers relies on microfabricated traps that can trap, store and manipulate many individual atomic ions. An early proposal was the quantum charge-coupled device (QCCD), a microfabricated array of electrodes that can trap the

ions and shift them around between “interaction-” and "memory-”regions [169]. This allows one, e.g., to apply 2-qubit gate operations on arbitrary pairs of ion by bringing them close to each other. One of the current frontrunners in the race towards ion-based quantum computers is the company Quantronics. It uses a QCCD-type architecture [258]. A similar path is chosen by some European groups [257].

Addressing of qubits by lasers must be achieved in the far-field diffraction-limited regime, where the separation between the ions must be large compared to an optical wavelength. This requirement sets a lower limit on the distance between the ions and therefore on the strength of the axial confinement potential. Since this potential also determines the vibrational frequency that enters the clock speed, it is obvious that ion traps cannot be operated with arbitrary speed.

In the case of hyperfine qubits, it is not necessary to use lasers for driving gate operations. Instead, one can use microwave pulses, provided a magnetic field gradient makes the ions distinguishable in frequency space[244]. The gradient also induces spin-state dependent interactions between the ions, which can be used for 2-qubit gates.

11.5 Neutral atoms

Neutral atoms can also be used as qubits [259]. Compared to trapped ions, they offer potentially lower decoherence rates, since their interactions with the environment are weaker. For the same reason, neutral atoms are more difficult to trap, store and manipulate.

11.5.1 Potential and force

The first prerequisite for using neutral atoms as qubits is a means to control their position and velocity. Since electrostatic forces cannot be used, one has to resort to electromagnetic waves that interact with the induced dipole moment of the atoms and / or to magnetic fields that interact with the static magnetic dipole moment of the atoms.

For quantum computing applications, the main tool for generating mechanical forces acting on atoms are laser beams. As discussed in section 11.2.3, the interaction with a laser field can create a highly viscous medium that reduces the velocity of ions and atoms and results in a velocity distribution corresponding to a low effective temperature. This effect is used also for neutral atoms. However, since the absorption-emission cycles involved in laser cooling destroy any information stored in the atoms, the cooling laser must be switched off for the actual computation. To trap the particles, one uses a different regime where the effective forces are conservative and can be described by an effective potential. This type of interaction can be understood by considering the potential energy surface generated by the laser field. Starting from the classical expression for the energy of an electric dipole μ_e in an electric field E,

$$U = -\vec{E} \cdot \vec{\mu}_e,$$

we calculate the force acting on the atom as

$$F = -\vec{\nabla} U = \vec{\nabla}(\vec{E} \cdot \vec{\mu}_e).$$

In the absence of saturation, the induced dipole moment $\vec{\mu}_e$ increases linearly with the strength of the field, $\vec{\mu}_e \propto \vec{E}$, and the effective potential U is thus proportional to the square of the field strength.

The sign of the potential depends on the difference between the laser frequency and the atomic transition frequency: For a red-detuned laser (i.e. laser frequency smaller than the transition frequency), the atomic dipole oscillates in phase with the laser field, $\vec{\mu}_e \cdot \vec{E} > 0$ and the energy becomes negative. In this case, the atom is pulled into the region of maximal field, where its potential energy is minimal. In the case of a blue-detuned laser, the atomic dipole is out of phase with respect to the field, the energy becomes positive and the atom is pushed out of the high-field region.

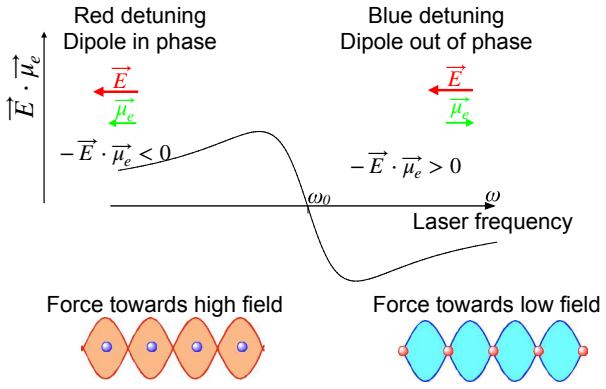

Figure 11.27: Effect of laser-detuning with respect to the optical resonance frequency: for a red-detuned laser, the atomic dipole oscillates in phase with the laser field and the atoms are pulled into the high field region. For a blue-detuned laser, the induced dipole is out of phase and the atoms are pushed out of the areas of high laser intensity.

11.5.2 Traps

A simple example of a laser-based trap for neutral particles is a tightly focused laser beam. Such traps were initially used for the manipulation of neutral atoms, but also for macroscopic particles [260, 233, 261]. The depth of the trap (i.e. the maximum kinetic energy that a particle can have without escaping from the trap) is determined by the laser intensity and the detuning of the laser frequency from the atomic resonance: the strength of the induced dipole moment decreases linearly with the frequency difference. While it would therefore be advantageous to tune the laser close to the resonance, this would also cause absorption. In the context of quantum information processing, however, absorption of light from the trapping laser must be avoided since this would cause decoherence. One therefore uses a large detuning and high laser intensity.

This type of trap can be extended to arrays of tightly focused laser beams [262, 263], as shown in figure 11.28. To make such a scheme scalable, the

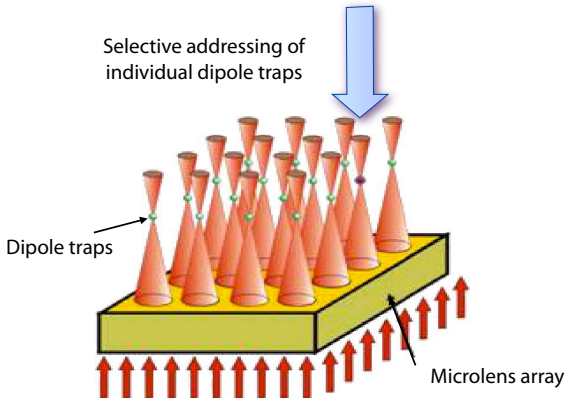

Figure 11.28: Array of dipole traps generated by focusing a laser beam with an array of micro-lenses.

different foci can be generated by micro-lenses, either in one or two dimensions [264]. This also allows for parallel manipulation of many qubits.

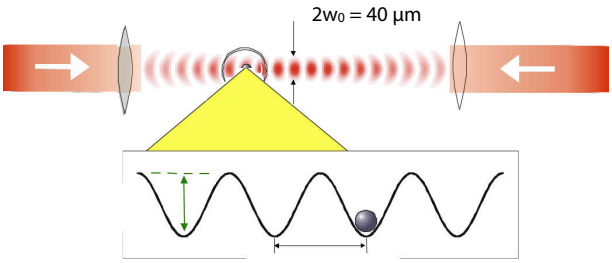

Figure 11.29: Dipole trap generated by two counter-propagating laser beams.

Another geometry for trapping neutral atoms is a standing wave generated by superimposing two counter-propagating laser beams, which are red-detuned with respect to the transition frequency. As shown in figure 11.29, the atoms are then drawn towards the antinodes of the beam, where the resulting intensity is maximal and the potential energy minimal. These energy minima form a linear sequence, separated by half of the laser wavelength.

The principle can be extended to two or three dimensions: two (or three) pairs of counterpropagating beams in orthogonal dimensions generate a 2D array of stable trap sites. Many parameters of the resulting optical lattices, such as the

dimensionality, form, depth and position can be precisely controlled through the geometry, polarization and intensity of the laser beams generating the lattice.

11.5.3 Motional control

A trap can be filled, e.g., by placing it directly in a cloud of cold atoms. The depth of a trap is of the order of 100 mK, with oscillator frequencies around 100 kHz. Using Raman cooling techniques, up to 80% of the atoms can be put into the motional ground state, which corresponds to temperatures in the μK range.

Such a process will always generate a random filling of the different minima of the trap, which is not compatible with the requirements of quantum information processing. If the traps are small enough, the interaction between the atoms reduces the probability of filling the trap with more than one atom [265].

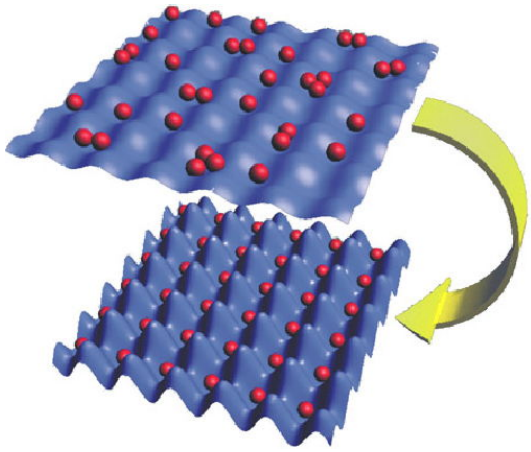

Figure 11.30: Array of dipole traps generated by two orthogonal pairs of counter-propagating beams showing random filling and the Mott insulator state.

In an array of such traps, a parameter range exists, in which an ensemble of cold atoms will preferentially occupy every microtrap with a single atom [266]. This regime is called a Mott insulator state (see also chapter 11.6) and may be used to create quantum registers if the separation between the microtraps is in a range suitable for individually addressing the qubits.

If the separation between the microtraps is smaller than the required distance between the qubits or if the fluctuations of the populations are too large, active control of the populations is required. For a linear trap, this was demonstrated experimentally [267, 268] by combining the trapping laser with a second standing wave trap, at a right angle to the quantum register.

For these experiments, it is necessary to shift the trap potential. This can be achieved in one dimension by shifting the phase of the counter-propagating laser beams that form the standing wave [269]. Shifting the phase of one beam by ϕ shifts the standing-wave pattern, i.e. the trap potential, by a distance

$$\delta = \frac{\lambda}{2}\frac{\phi}{2\pi},$$

where λ is the laser wavelength. If the phase shift is time-dependent, i.e. realized by a frequency shift of one of the laser beams, this results in a linear motion of the trap potential. Acceleration can be implemented by a frequency chirp.

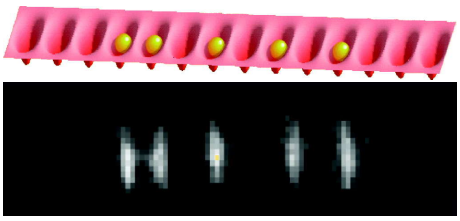

Figure 11.31: Schematic representation of an atomic conveyor belt and image of atoms in a conveyor belt.[270]

On the basis of such phase shifts, it is possible to implement a wide variety of time-dependent potentials, including “atomic conveyor belts”, such as the one represented in figure 11.31.

Laser-optical traps are in general state-selective: depending on the internal state of the atom, the interaction can be strong or weak, and the atoms can be confined to or expelled from regions of

high laser intensity. This effect must be taken into account when gate operations are applied that change the internal state of the atoms. The state-selectivity can also be used, to specifically manipulate subsets of the atoms that are in specific internal states.

11.5.4 Gate operations

Single qubit gate operations can be performed on neutral atoms in much the same way as on atomic ions, using laser pulses with appropriate wavelength and polarization. The qubit states will typically be hyperfine sublevels of the electronic ground state. Transitions between them can be excited either by Raman laser pulses or by microwave pulses [271, 272].

For two-qubit operations, a state-dependent interaction must be present between the different qubits. In contrast to trapped ions, neutral atoms do not experience Coulomb forces. They can, however, interact by electric or magnetic dipole couplings, which depend on their internal state. This interaction is of course significantly weaker than the Coulomb interaction, and has a shorter range. This shorter range can also be beneficial, since it reduces unwanted long-range interactions, and since these interactions are state-dependent, it is possible to turn them on and off in a controlled manner.

The strength of the interaction between qubits can be controlled, e.g., via the distance between them. This requires independent trapping of two subensembles, e.g. by using two dipole traps with different polarizations, such that one trap represents the dominant interaction for atoms in one internal state, while the orthogonal polarization dominates for the other state [273]. The two potentials have the same periodicity but are displaced with respect to each other, and the displacement may be controlled. Thus, atoms in different spin states may be brought into contact with each other in a well-defined way and for a well-defined time. If the time-dependence of the interaction is properly adjusted, the overall effect of such a controlled collision between cold atoms generates an entangling quantum logical gate operation [274, 275].

The strength of the interaction and therefore the speed of the gate operations can be increased significantly if the atoms are briefly promoted to a highly excited Rydberg state [276]. For suitable Rydberg states, interactions have been observed over distances up to 10 µm. Another possibility for controlling the coupling between atoms is to bring them into optical resonators. If two atoms interact with the same resonator mode, they experience an indirect interaction with each other [277].

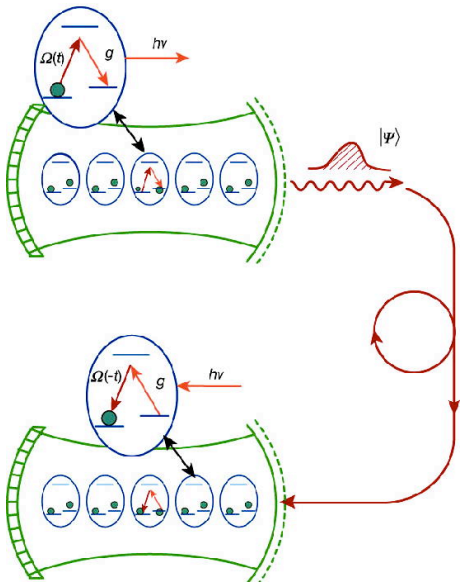

Figure 11.32: Couplings between atoms can be generated by placing them in an optical resonator and by connecting optical resonators through optical fibers.

Alternatively, the atoms can be put into separate optical resonators, which are then coupled to each other, e.g. through an optical fiber [255]. These schemes involve some degree of population of electronically excited states. Accordingly, a major challenge for their implementation is the need to avoid spontaneous emission, which destroys the quantum coherence in these systems.

11.6 Interacting atoms in optical lattices

Some interesting applications of trapped atomic ensembles are found in the field of quantum simulations. Most of the problems that were considered are inspired by condensed matter physics. In ideal crystals as well as in multidimensional trapped atom experiments, the system has periodic boundary conditions. In a pioneering experiment, Greiner et al [266], following a theoretical suggestion of Jaksch et al. [278], observed a "quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms", as the title of the paper says. The paper demonstrated for the first time that it is possible to construct physical realizations of theoretical models for condensed-matter systems with adjustable values of the model parameters and without many of the complications present in real-world condensed-matter systems. A lot of activity followed, both experimental and theoretical. The phenomena and the tools employed in their study have been discussed in a number of reviews of varying level, perspective, and length [279, 280, 281, 282, 283].

Atoms in optical lattices thus are excellent illustrations of Feynman's [14] concept of using quantum systems to simulate other quantum systems. It also seems possible to construct gates, initialize qubits and perform other operations essential for the implementation of universal quantum computation. A significant drawback, however, remains the lack of individual adressability of the atoms stored in an optical lattice.

11.6.1 Particles in a periodic potential

For a discussion of the Hubbard model, we need a few basic notions from the theory of crystalline condensed matter. In ordinary solids the electrons move in a periodic potential generated by the ion charges. The interaction between electrons is often neglected or treated implicitly in some form of effective-field approximation unless it is absolutely necessary to proceed otherwise. The interatomic distances in a solid are of the order of the atomic radius, which leads to overlap between electronic wavefunctions of neighboring atoms, to chemical bonding, and, under suitable conditions, to the ability of electrons to move around in the crystal. In an "artificial solid" of atoms in an optical lattice this is different: the lattice constant (the distance between neighboring potential wells) is given by the light wavelength, $\lambda_L \approx 10^{-6}$m, much larger than the typical size of an atom, about 10^{-10}m. All short-range variations in the interatomic potential can thus be neglected when we discuss effects of the interatomic interactions in an optical lattice. Furthermore, the electrons in a real solid are inevitably spin-1/2 fermions, whereas an optical lattice can be populated with either bosonic or fermionic atoms.

Before discussing interatomic interaction effects, however, we have to understand the behavior of a single atom (or, equivalently, of a number of noninteracting atoms) in a D-dimensionally periodic potential of the form

$$V(\vec{r}) = V_0 \sum_{i=1}^{D} \sin^2 k_L r_i. \tag{11.2}$$

The potential is generated by superposing D standing-wave laser beams of wavenumber $k_L = 2\pi/\lambda_L$ in orthogonal directions, thus creating a simple cubic (or square, or one-dimensional) lattice of potential minima with lattice constant $a = \lambda_L/2$. The potential strength V_0 is given by the intensity of the laser beam. As the intensity varies across the beam, (for example in a Gaussian shape with maximum intensity in the center of the beam), V_0 should be considered weakly position-dependent. In addition, V_0 may also vary along the beam due to focusing effects. The spatial variation of V_0, though necessary to keep the atoms from moving out of sight, will be neglected.

The motion of a free particle is completely characterized by the momentum $\vec{p} = \hbar\vec{k}$, and the

particle's wave function is a plane wave

$$\psi_{\vec{k}}(\vec{r}) = \frac{1}{\sqrt{\Omega}} e^{i\vec{k}\cdot\vec{r}}, \tag{11.3}$$

where Ω is a normalization volume. The energy of the particle is

$$\varepsilon_{\vec{k}} = \frac{\hbar^2 k^2}{2m}. \tag{11.4}$$

In an optical lattice the laser frequency is tuned to be roughly (but not precisely) equal to a transition frequency of the atom; the emission or absorption of a photon will thus be accompanied by a kinetic energy of recoil

$$E_R = \frac{\hbar^2 k_L^2}{2m}. \tag{11.5}$$

A periodic potential acting on a free particle changes both wave function and energy. Instead of a plane wave (11.3) the wave function becomes a modulated plane wave, or Bloch function,

$$\psi_{\vec{k}}(\vec{r}) = u_{\vec{k}}(\vec{r}) e^{i\vec{k}\cdot\vec{r}}, \tag{11.6}$$

where the function $u_{\vec{k}}(\vec{r})$ (the "Bloch factor") has the periodicity of the potential $V(\vec{r})$ (11.2). The dispersion relation (energy-momentum relation) $\varepsilon_{\vec{k}}$ (11.4) changes gradually as the potential strength is slowly increased. For very weak potential $\varepsilon_{\vec{k}}$ essentially keeps its free-particle form (11.4); however, it turns out to be convenient to write it in the form

$$\varepsilon_{\vec{k},\vec{g}} = \frac{\hbar^2}{2m}(\vec{k}+\vec{g})^2, \tag{11.7}$$

where $\vec{g}$ is a vector of the reciprocal lattice, defined by requiring that the plane waves $e^{i\vec{g}\cdot\vec{r}}$ have the periodicity of the lattice (or the potential $V(\vec{r})$). In the simple cubic case, all components of every $\vec{g}$ are integer multiples of $\frac{2\pi}{a} = \frac{4\pi}{\lambda_L}$. In the dispersion (11.7) $\vec{k}$ is then restricted to the first Brillouin zone, the region of reciprocal space ($\vec{k}$ space) closer to $\vec{g} = \vec{0}$ than to any other $\vec{g}$. In the simple cubic case that is the cube $[-\pi/a, \pi/a]^3$.

In $\varepsilon_{\vec{k},\vec{g}}$ then $\vec{g}$ classifies different branches of the dispersion relation, leading to different energy bands. The lowest energy band obviously is the one with $\vec{g} = \vec{0}$. At the Brillouin zone boundary, for example at $\vec{k} = (\pi/a, 0, 0)$, the $\vec{g} = \vec{0}$ and $\vec{g} = (-2\pi/a, 0, 0)$ bands are degenerate. Degenerate perturbation theory shows that the degeneracy is lifted by a weak potential V_0. The formerly degenerate energy levels are pushed away in opposite directions and an energy gap is created. The set of energy bands and gaps is called the band structure, and is an important means in understanding the behavior of crystalline condensed matter; compare, for example [284, 285] for details. For sufficiently weak potential strength V_0 the maximum energy of the lowest band is of the order of

$$\frac{\hbar^2}{2m}\left(\frac{\pi}{a}\right)^2 = \frac{\hbar^2}{2m}\left(\frac{2\pi}{\lambda_L}\right)^2 = \frac{\hbar^2 k_L^2}{2m} = E_R. \tag{11.8}$$

The bandwidth of the lowest energy band for "nearly free" particles is thus equal to the recoil energy (11.5).

The Hamiltonian for non-interacting fermions or bosons in an optical lattice can now be written in the occupation number (or "second quantization") formalism:

$$\mathbf{H} = \sum_{\vec{k},\vec{g}} \varepsilon_{\vec{k},\vec{g}} c^\dagger_{\vec{k},\vec{g}} c_{\vec{k},\vec{g}}\,, \tag{11.9}$$

where $c^\dagger_{\vec{k},\vec{g}}$ ($c_{\vec{k},\vec{g}}$) is the creation (annihilation) operator for a particle in a single-particle energy eigenstate with quantum numbers $\vec{k}$ and $\vec{g}$. The operator $n_{\vec{k},\vec{g}} = c^\dagger_{\vec{k},\vec{g}} c_{\vec{k},\vec{g}}$ is the occupation number operator for that eigenstate. (Note that the Bloch functions $\psi_{\vec{k}}$ and the Bloch factors $u_{\vec{k}}$ in (11.6) should also bear $\vec{g}$ or some other appropriate band index as $\vec{k}$ is restricted to the first Brillouin zone.)

11.6.2 Interactions between atoms

Let us now discuss the effects of interatomic interactions in an optical lattice. Since typical interatomic distances are comparable to λ_L and thus very large compared to atom sizes, the interatomic potential essentially only acts through

its long-range part which can be parametrized in terms of the scattering length a_S:

$$V_{\text{int}}(\vec{r}-\vec{r}') = \frac{4\pi\hbar^2}{m} a_S \delta(\vec{r}-\vec{r}'). \tag{11.10}$$

In order to treat this short-range potential, it is convenient to use a basis set of localized single-particle wave functions instead of the Bloch functions (11.6) extending over the whole "crystal". This basis set is given by the Wannier functions [284, 285]

$$w_n(\vec{r}-\vec{l}) = \frac{1}{\sqrt{N}} \sum_{\vec{k}} e^{-i\vec{k}\cdot\vec{l}} \psi_{n\vec{k}}(\vec{r}). \tag{11.11}$$

Here, N is the number of lattice sites $\vec{l}$ in the system (equal to the number of $\vec{k}$ vectors in the first Brillouin zone) and n is an index indicating the band of interest. (Labeling bands by the reciprocal lattice vectors $\vec{g}$ is convenient only close to the free-particle case.) Each Wannier function is centered around a lattice site $\vec{l}$, and decays with growing distance from $\vec{l}$. If the potential wells of the optical lattice are very deep and well separated from each other (that is, for large V_0) the Wannier functions for the lowest energy bands are the lowest energy eigenstates of a single potential well (similar to orbitals of an isolated atom), and the Bloch functions are linear combinations of atomic orbitals.

Denoting the creation and annihilation operators for particles in a Wannier state by $c^\dagger_{\vec{l}n}$ and $c_{\vec{l}n}$, respectively, the Hamiltonian of non-interacting particles may be written as

$$\mathbf{H} = \sum_{\vec{l},\vec{l}'} t_{\vec{l}-\vec{l}'n} c^\dagger_{\vec{l}n} c_{\vec{l}'n}\,, \tag{11.12}$$

where the "hopping elements" $t_{\vec{l}-\vec{l}'n}$ are given by

$$t_{\vec{l}-\vec{l}'n} = \frac{1}{N} \sum_{\vec{k}} \varepsilon_{\vec{k}n} e^{i\vec{k}\cdot(\vec{l}-\vec{l}')}. \tag{11.13}$$

Indices labeling spin or other internal degrees of freedom have been suppressed. The formal expression (11.13) for the hopping elements can be rewritten in terms of the Wannier functions at lattice sites $\vec{l}$ and $\vec{l}'$. For the large-V_0 case the Wannier functions are well localized within each potential minimum and the hopping elements will be negligible except for $\vec{l}$ and $\vec{l}'$ nearest neighbors. If the Wannier functions are isotropic, all non-vanishing hopping elements will have the same value, which we call $-t_n$. The non-interacting particles then are described by the energy bands

$$\varepsilon_{\vec{k}n} = -2t_n \sum_{i=1}^{D} \cos k_i a. \tag{11.14}$$

11.6.3 The Hubbard model

In the same spirit, the interaction between atoms via the potential (11.10) can be discussed. The interaction term of the Hamiltonian then contains a sum of terms, each with four electron creation and annihilation operators and an integral involving Wannier functions located at four lattice sites. From the localization properties of the Wannier functions it is then clear that the dominant term is the one where all Wannier functions are located at the same lattice site. Neglecting all other terms, the interaction is given by the single value

$$U = \frac{4\pi\hbar^2}{m} a_S \int d^3r\, |w_n(\vec{r})|^4. \tag{11.15}$$

Since we focus on a single band, the band index n can be omitted. The total Hamiltonian depends on the statistics of the atoms involved. If the atoms are spin-half fermions (with internal quantum number $\sigma = \uparrow, \downarrow$) we obtain the Hubbard model [286, 287, 288] in its original form

$$\mathbf{H} = -t \sum_{\vec{l},\vec{l}',\sigma} c^\dagger_{\vec{l}\sigma} c_{\vec{l}'\sigma} + U \sum_{\vec{l}} n_{\vec{l}\uparrow} n_{\vec{l}\downarrow} \tag{11.16}$$

($\vec{l},\vec{l}'$ nearest neighbors), where $n_{\vec{l}\sigma} := c^\dagger_{\vec{l}\sigma} c_{\vec{l}\sigma}$ is the number operator. The model (11.16) and its many extensions are popular in solid-state physics for modeling the electron corre-

lation effects believed to be important in magnetism, metal-insulator transitions, and high-temperature superconductivity. If the atoms are bosons (all in the same internal atomic state) the Hamiltonian is known as the bosonic Hubbard model

$$\mathbf{H} = -t \sum_{\vec{l},\vec{l}'} c^{\dagger}_{\vec{l}} c_{\vec{l}'} + \frac{U}{2} \sum_{\vec{l}} n_{\vec{l}}(n_{\vec{l}} - 1), \quad (11.17)$$

originally [289] employed to describe superfluids in porous media or granular superconductors.

In contrast to the situation in real-world condensed matter systems modeled by Hubbard-type Hamiltonians, the parameter values t and U for atoms in an optical lattice can be easily tuned by varying the laser field strength V_0 in (11.2). If V_0 increases, the potential wells get steeper and narrower and the wave functions get compressed, so that U (11.15) increases. Approximating the potential well by a D-dimensional paraboloid and the wave function by the appropriate oscillator ground state, one obtains $U \sim V_0^{D/2}$.

→ **Problem**

By the same mechanism the overlap between wave functions in neighboring potential wells will decrease as V_o grows, and hence the nearest-neighbor hopping amplitude t will decrease exponentially. The ground-state properties of the Hubbard model depend only on the ratio U/t and thus on the laser field strength V_0.

11.6.4 Experimental simulation: setup

In the experiment of Greiner et al. [266] ultracold bosonic ^{87}Rb atoms were trapped in an optical lattice produced by a laser with $\lambda_L = 852$nm. The recoil energy (11.8) then is $E_R \approx k_B \cdot 0.15\mu$K and the trapping potential was $V_0 \lesssim 22E_R$. Non-interacting free ($V_0 = 0$) bosons at zero temperature will condense into the lowest plane-wave state (11.3), with $\vec{k} = \vec{0}$. The situation does not change decisively if a weak interatomic interaction and a weak lattice potential are present: the state remains a macroscopically coherent (superfluid) many-boson state.

With growing strength of the lattice potential V_0, the single-particle states become progressively localized and the repulsive interaction U dominates the Hamiltonian (11.17) more and more. If the total number of particles is small enough, every lattice site (or potential well) contains at most one atom. For strong enough V_0 each atom will be strongly localized in one potential well and will not enjoy enough overlap to its neighbors to develop long-range phase coherence of the wave function. If every potential well contains exactly one atom[5], further atoms can only be added at the price of an excitation energy U per atom. The same energy gap also prevents the formation of doubly occupied sites (and accompanying vacancies) which would be needed to achieve particle transport. The resulting state is obviously incompressible and, thinking in terms of the original (electronic) Hubbard model, insulating. Therefore it is known as the Mott (-Hubbard) insulator state.

The transition between the superfluid and Mott insulating states was demonstrated in a time-of-flight experiment [266]. For small or moderate potential strength V_0 all bosons convene in the lowest-energy extended Bloch state in a coherent manner. As the Bloch state is a periodically modulated plane wave (11.6) it contains Fourier components with different $\vec{g}$ values, with $\vec{g} = \vec{0}$ dominating as V_0 goes to zero. In the experiment the optical potential is switched off suddenly and the atoms are allowed to expand freely. The Fourier components with different $\vec{g}$ separate spatially according to their different propagation speeds. After a fixed expansion time an absorption image is taken. The absorption images of Figure 11.33 therefore map directly the distribution of the atoms in reciprocal space.

11.6.5 Results

For free atoms, ($V_0 = 0$) the lowest-energy Bloch state is an unmodulated $\vec{k} = \vec{0}$ plane wave. The corresponding absorption image in Figure

[5] The situation is similar for any other integer number of atoms per site.

11.33,a therefore shows a single spot in the center. As V_0 grows, additional Fourier components enter the lowest Bloch state and become visible in the absorption image, which develops into a two-dimensional projection of the reciprocal lattice (Figure 11.33 b, c, and d). However, upon further growth of V_0 the reciprocal lattice spots fade away again, their intensity being soaked up by a big central blob. That blob (Figure 11.33 g and h) is witness to the fact that the state has evolved into one where each potential minimum houses one atom, with no phase relation between neighboring atoms and hence no interference visible. The large extent of the central absorption spot in reciprocal space reflects the localization of each atom in real space. This shows clearly the change in the nature of the ground state as the ratio U/t is varied. The expected change in the nature of the low-energy excitation spectrum from gapless in the superfluid state to gapped in the Mott insulating state could also be observed [266] by analyzing tunneling between neighboring potential wells which are energetically displaced with respect to each other in an applied external field.

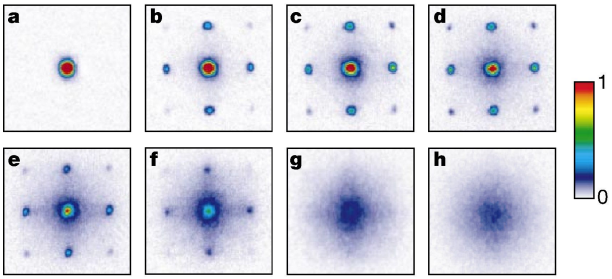

Figure 11.33: Absorption images reflecting the Fourier component structure (momentum distribution) of the many-particle wave function of ^{87}Rb atoms in an atomic lattice of strength V_0. Images were obtained after switching off the lattice potential suddenly and allowing atoms to expand for 15 ms. Potential strengths V_0 in units of the recoil energy E_R are a:0, b:3, c:7, d:10, e:13, f:14, g:16, and h:20 [266].

Fermionic atoms in optical lattices have also been studied. Köhl et al. [290] stored a large number of ^{40}K atoms in a three-dimensional simple cubic lattice and obtained absorption images after switching off the potential and allowing the atoms to expand ballistically. As the Pauli principle strictly forbids double occupation of single-particle energy levels, the fermionic ^{40}K atoms fill up the available Bloch states (11.6) up to the Fermi energy. The surface in $\vec{k}$ space which separates occupied states (at low energy) from empty states (at high energy) is called the Fermi surface. For a simple cubic optical lattice potential of the form (11.2) the Fermi surface is spherical at low density, but less so at higher density. For a completely filled band the Fermi surface is equal to the boundary of the first Brillouin zone. The energy gap to the next higher band then is the minimum energy for a single-particle excitation. For electrons in a solid that situation corresponds to an insulator (or a semiconductor, if the energy gap is small enough). In contrast to the interaction-induced Mott insulator discussed above, the present case is termed band insulator and obviously does not rely on interaction effects. In the optical-lattice experiment on ^{40}K the shape of the Fermi surface could be measured for various particle densities. Also, employing the magnetic-field dependence of the scattering between two spin species (Feshbach resonance), interaction effects like the transfer of atoms into higher bands could be observed.

These pioneering experiments on bosonic and fermionic atoms in optical lattices show that quantum simulation of correlated many-body systems may soon be within reach of experimental possibilities. These exciting prospects have led to a very large number of proposals for correlation effects in many-body systems that could be studied with atoms in optical lattices, see the review by Lewenstein et al. [283].

Problems

For sufficiently large potential strength V_0 the optical lattice potential (11.2) can be approximated by a harmonic oscillator potential

$$V_{\text{osc}} = \frac{m\omega^2}{2}\vec{r}^2,$$

where $\vec{r}$ is the D-dimensional vector of displacement from the potential minimum.

a) Calculate the Hubbard interaction U (11.15), approximating the Wannier function $w_n(\vec{r})$ by the normalized oscillator ground-state wave function

$$\phi_0(\vec{r}) = \pi^{-D/4} a^{-D/2} \exp -\frac{1}{2}\left(\frac{\vec{r}}{a}\right)^2,$$

where $a = \sqrt{\frac{\hbar}{m\omega}}$ is the characteristic length of the quantum harmonic oscillator. Show that U grows as $V_0^{D/2}$.

b) The nearest-neighbor hopping amplitude t can be approximated by the overlap (the integral of the product of the wave functions) between the ground-state wave functions in neighboring potential wells. Calculate t and determine its dependence on the parameters of the optical lattice potential. Show that t decreases as V_0 grows.