Electrons in artificial atoms

R. C. Ashoori

Progress in semiconductor technology has enabled the fabrication of structures so small that they can contain just one mobile electron. By varying controllably the number of electrons in these 'artificial atoms' and measuring the energy required to add successive electrons, one can conduct atomic physics experiments in a regime that is inaccessible to experiments on real atoms.

THE puzzle of atomic spectra was a prime motivation for the development of quantum mechanics. Niels Bohr unravelled the mystery by determining that the wavelike nature of electrons allowed them to occupy only discrete orbits within an atom, with well defined energies. Starting about 10 years ago, advances in semiconductor technology allowed the fabrication of structures so small that their discrete quantum level structure was resolvable. In the past few years, powerful new spectroscopic probes have revealed a wealth of new physics in these 'artificial atoms'.

Essentially, artificial atoms are small boxes about 100 nm on a side, contained in a semiconductor, and holding a number of electrons that may be varied at will. As in real atoms, electrons are attracted to a central location. In a natural atom, this central location is a positively charged nucleus; in an artificial atom, electrons are typically trapped in a bowl-like parabolic potential well in which electrons tend to fall in towards the bottom of the bowl. One can consider the artificial atom as a tiny laboratory in which quantum mechanics and the effects of electron–electron interactions can be studied.

The spacing between atoms in a semiconductor crystal is typically about 0.3–0.4 nm. In artificial atoms, electrons are confined in structures about 100 nm in diameter. Thus, an artificial atom in a crystal comprises many real atoms. The quantum-mechanical theory of solids explains why the electrons do not get trapped on the real atoms of the crystal and instead only sense the potential well of the artificial atom. This theory dictates that some electrons in a crystal behave as free electrons, albeit with a different mass[1]. For example, electrons in the semiconductor gallium arsenide appear to carry a mass that is only 7% of the mass of free electrons.

The physical characteristics of an artificial atom (sometimes called a 'quantum dot') differ considerably from that of a natural atom for one important reason: artificial atoms are typically much larger than real atoms. The electron orbits do not simply scale with size. Imagine an atom containing many electrons whose size is continuously variable. As it is made larger, the Coulomb energy arising from the repulsion between electrons orbiting around the nucleus decreases because the average spacing between electrons increases. However, there is another energy scale in the problem: the separation in energy of the different orbits of electrons in the artificial atom. As the atomic size increases, the differences in the orbital energies decrease faster than the Coulomb energy. It follows that in a large atom, the effects of electron–electron interactions are relatively more important than in a small atom.

It may not be long before these distinctions affect modern electronics. The present view of small electronic devices depicts them as controlling the motion of small and classical 'seas' of a few thousand electrons. As devices shrink, this view no longer holds. In fact, it is already possible to create electronic devices small enough to have device characteristics sensitive to the motions of single electrons within them, even at room temperature[2,3]. Electronic devices may no longer be seen as small seas of electrons but instead as very large 'atoms'. Building devices at this size scale that have reproducible and desired electronic properties will be an immense challenge. This Review describes structures in which researchers have now succeeded in precisely controlling the number of contained electrons, down to as few as one electron. It focuses on the unusual physical features of these large atoms unveiled by these studies. At low temperatures, electrons fall into distinct quantum-mechanical energy levels of these atoms, and the large Coulomb energy exerts a profound influence. Transitions never observed in the spectra of 'natural' atoms are readily seen for artificial ones. The electrons in artificial atoms can also act lethargically, reluctant to displace themselves to make way for another electron, even on a timescale of milliseconds. Most strangely, the large Coulomb repulsion can make it seem that electrons attract one another.

Measurements of the orbital energies of electrons in an artificial atom thus yield a wealth of data about a new physical regime. A large number of experimental probes have been used. Some techniques are electrical, such as measuring the current through a single dot[4–10], or the capacitance of arrays[11–13] and single quantum dots[14]. Others have employed infrared and optical spectroscopies on arrays of dots[15–17] and on individual dots[18].

Construction and measurement

I will focus here on results from two newly developed techniques[5,14,19] which have particularly affected our understanding of quantum dots. These methods are unique in allowing extraordinarily high-resolution spectroscopy of single quantum dots. The energy resolution of these spectroscopies is limited solely by the sample temperature. Both types of experiments have been performed on artificial atoms constructed inside a gallium arsenide semiconductor and at very low temperatures (0.05–0.30 K).

Figure 1*a* shows schematically the type of sample used in one of these experiments. The quantum dot is located between two capacitor plates. It is close enough to one of the plates to allow single electrons to tunnel (or hop) between the artificial atom and the nearby plate. Tunnelling is a quantum-mechanical process that allows electrons to pass through barriers that would be impenetrable classically. The artificial atom is far enough from the other capacitor plate to prohibit tunnelling to this plate.

How is such a sample actually realized? Utilizing modern semiconductor technology, it is possible to grow gallium arsenide (GaAs) semiconductor crystals one atomic layer at a time[20]. Moreover, the material composition can be changed in successive layers. Some of the gallium in a layer can be replaced with aluminum to create aluminum gallium arsenide (AlGaAs). AlGaAs acts as an insulator, whereas electrons move freely in GaAs. By sandwiching a 10-nm thickness of GaAs between AlGaAs layers, one confines electrons in a GaAs 'quantum well'. This quantum well is so thin that at low temperatures, only the lowest quantum energy state of the well is occupied by electrons. The electrons have no freedom to move in the direction perpendicular to the well, but may only move laterally in it. By

placing electrostatic gates on the surface of the wafer, we can laterally confine this two-dimensional electron gas and create a quantum dot.

Figure 1*b* demonstrates how the schematic picture shown in Fig. 1*a* is realized in an experiment. A crystal wafer is grown layer by layer, starting from the bottom of the diagram. The first layer is a silicon-doped GaAs layer. The silicon doping makes this layer metallic, and it acts as the bottom capacitor plate of Fig. 1*a*. Then a thin (10-nm) AlGaAs barrier layer is grown. This barrier is thin enough that electrons can leak through it. Above that, there is a GaAs quantum well, and grown on top of the well is a thick (not leaky) AlGaAs barrier. On top of the crystal wafer, chromium is deposited to form the top capacitor plate, which I will refer to as the 'gate'. Additional sample processing[21,22] on the sample surface is used to create a gate which laterally confines electrons in the quantum well below, creating a quantum dot.

Electric fields can be created by applying a voltage between the plates of the capacitor. If the top plate (gate) is made positive compared to the bottom one, electrons from the bottom plate will be attracted in the direction of the top plate, towards the artificial atom. Single electrons can thus be coaxed to tunnel into the dot or expelled from it through the application of voltages on the top plate.

After each electron is added to the dot, additional gate voltage is usually needed to coax another electron onto the dot. As I discuss below, this happens partly because after each electron is added to the dot, the electron charge in the dot is larger, and other electrons are repelled from entering. There are therefore specific gate voltages at which an electron is added to the dot, and it is this spectrum of voltages (the electron addition spectrum) which has been particularly useful in understanding quantum dots.

The motion of the single electrons into or out of the dot can be detected using a simple physical principle. When a single electron tunnels into the artificial atom, it moves closer to the top plate of the capacitor, and electrons in the top plate tend to be pushed away from the plate; that is, some charge is induced on the top plate. In these samples, the amount of charge induced is about half of an electron's charge. A specialized transistor enables detection of this charge and thereby allows measurement of the gate voltages at which single electrons were added to the artificial atom. A small a.c. voltage of frequency around 100 kHz is added to the d.c. gate voltage. When the gate voltage is adjusted to a voltage at which an electron can be added to the dot, the a.c. voltage causes the electron to tunnel back and forth between the dot and the bottom plate. At these gate voltages, charge appears synchronously with the a.c. voltage. A synchronous detector then registers a signal only at these voltages, yielding a peak in the response. The scheme measures the capacitance signal due to a single electron and is known as single-electron capacitance spectroscopy (SECS).

Results from this method, taken on a sample cooled to 0.3 K, are shown in Fig. 1*c*. The first peak on the left of the graph corresponds to the first electron in the quantum dot. Subsequent peaks correspond to additional electrons added to the quantum dot. The peak structure is completely reproducible as one scans the gate voltage up and down, and the widths of the peaks directly reflect the temperature of the sample.

The other type of experiment is performed on the type of sample displayed in Fig. 2*a*. Like the SECS method, this experiment also measures the gate voltages at which electrons are added to the dot. The detection method is however, completely different. Two current leads are in such close proximity to the dot that electrons may tunnel between the leads and the dot. A small voltage difference is applied between the left and right leads, and the experiment consists of monitoring the current through the quantum dot as the gate voltage is varied. I will refer to this type of experiment known as 'gated transport spectroscopy' (GTS).

For a current to flow, single electrons must pass through the quantum dot. This can only happen under special conditions. Just as in the SECS experiments, at particular gate voltages, it becomes energetically favourable for one more electron to be added to the

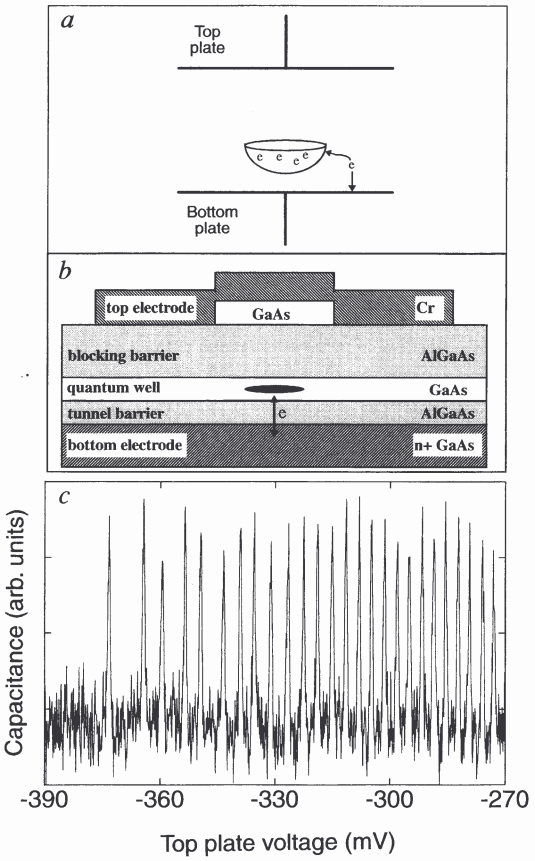

FIG. 1 *a*, Schematic diagram of an artificial atom located between two capacitor plates. The artificial atom is actually two-dimensional; the bowl-like shape is to represent the force tending to move electrons to the centre of the atom. *b*, Diagram of the sample used in single-electron capacitance spectroscopy (SECS) experiments in a crystal grown using molecular-beam epitaxy. The artificial atom is the black disk in the quantum well. *c*, Capacitance of the sample containing the artificial atom as a function of the top plate (gate) voltage. The first peak on the left represents the first electron entering the artificial atom.

dot. If the gate voltage is varied above or below these values, the number of electrons on the dot is stable and differs by one. However, when the gate voltage is set precisely at the gate voltage needed for an electron to be added to the dot, the number of electrons on the dot may fluctuate by one. This fluctuation can occur by one electron tunnelling onto the dot from the left lead and later an electron tunnelling off the dot to the right lead. This process gives rise to a detectable current flowing through the dot. In the experiment, a current peak therefore appears as the gate voltage is swept through the position where one more electron is added to the dot. These peaks are clearly demonstrated in the gate voltage scan of Fig. 2*b*. Unlike the SECS experiment, the first peak on the left corresponds not to the first electron in the quantum dot but to perhaps the thirtieth. To have electrical conduction through the dot, it must be close enough to the right and left leads for tunnelling to occur to both. For GTS experiments, it has as yet proven practically impossible to create a dot containing fewer than around ten electrons[23,24], for which appreciable electron tunnelling occurs to both electrodes. Most GTS experiments start with 25 or more electrons in the dot[25–27].

As with the SECS experiment, the peak positions reflect the energy required to add each successive electron to the dot. The peaks have a non-zero width because there is a range of gate voltages over which fluctuations in the electron number on the dot can occur. Once again, this range is directly related to the temperature of the sample.

Interpreting the spectra

In both Figs 1*c* and 2*b*, electrons can then be counted one by one as they move into the artificial atom. A special feature of the SECS experiment is that it even allows precise knowledge of the number of electrons in the dot by allowing observation of the first electron in the dot.

The gate voltage scales in Figs 1*c* and 2*b* can be directly understood in terms of the energy required to add each successive electron to the quantum dot. There are essentially two reasons why it takes extra energy to add each additional electron to the dot[12,19]. First, the electrons in the dot make it more difficult for other electrons to enter the dot. As more electrons are added, this repulsion increases and it takes more energy (charging energy) to add another electron. Second, the Pauli exclusion principle requires electrons to be in different quantum levels in the dot. Additional electrons must be promoted into higher energy levels (quantum level spacing). In the types of GaAs quantum dots discussed in this Review, the charging energies are typically a factor of 5 greater than the mean quantum level spacing.

The application of a gate voltage balances these energy requirements so that additional electrons can be added to the dot. Electrons can lower their energies by entering the quantum dot and thereby be closer to the attractive gate[28]. Owing to the simple geometry of these structures, this energy change varies in linear proportion to the gate voltage. Multiplying the gate voltage scales in Figs 1*c* or 2*b* by a geometric factor converts this scale into an energy scale for the quantum dot. For the case of Fig. 1*c*, this geometric factor is 0.5. The spacings between the peaks in both figures then directly reflect the additional energies required to add successive electrons to the quantum dot. The SECS and GTS techniques thereby yield direct energy-level spectroscopy of quantum dots.

A model for the artificial atom

To understand the behaviour of the dot containing more than one electron, we need to examine in some detail the quantum energy levels of a quantum dot containing one electron. In most experiments, an electron is tightly bound in a quantum well in the z direction and has no freedom to move in this direction. I will then consider the electron to be bound laterally in the quantum well in the x–y plane by a potential $1/2m\omega_0^2 r^2$. Here, $r^2 = x^2 + y^2$ measures the distance from the centre of the quantum dot. The simple mechanical analogue of this system is a ball in a bowl whose cross-section is a parabola, and ω_0 is the frequency that a classical particle would exhibit while oscillating in this bowl. This case of a circularly symmetric parabolic potential approximates well with the actual case in experiments and turns out to be rather simple to solve[29].

In solving the Schrödinger equation for the parabolic potential, one obtains quantum level energies given by the simple formula:

$$E_{nl} = \hbar\omega_0(2n + |l| + 1) \qquad (1)$$

where $\hbar$ is Planck's constant divided by 2π. There are two quantum numbers, n and l, corresponding to the fact that the electron moves in two dimensions. n is a positive integer which corresponds to the number of nodes in the wavefunction as one moves radially out from the dot centre, and $2|l|$ gives the number of nodes seen in moving circumferentially about the dot centre. l may be any integer, positive or negative. States of positive l correspond to electron wavefunctions moving anticlockwise with time as viewed looking down from the positive z axis, negative l states move in the opposite direction. These correspond to a series of circular orbits, with states of larger $|l|$ orbiting farther away from the dot centre. Figure 3*a* displays a picture of the orbits for a particular value of n. The mean radius of the different l states increases as $\sqrt{l}$.

Apart from having the quantum numbers n and l discussed above, electrons also have a spin quantum number, s, which can take values of $\pm 1/2$. An electron with $s = +1/2$ has a spin which points along the z axis perpendicular to the plane of the quantum dot (spin up); an electron with $s = -1/2$ has a spin which points against it (spin down). In zero magnitude field, electrons with either spin state have precisely the same energy.

It is useful to examine the energy levels determined by equation (1). Plotting the energy levels as a function of l, we obtain the results of Fig. 3*b*. The lines in the figure connect states that have the same value of n. The bottom 'V' is for states with $n = 0$. Each state can hold two electrons, one with electron spin up and one with spin down. At low temperatures, non-interacting electrons will simply fall like marbles into the lowest available states. Each of the circles shown in Fig. 3*b* can hold two electrons; the filled circles represent filled states in a quantum dot containing 30 electrons. Note that for this case of zero magnetic field, there are typically several states for the same value of the energy; that is, the states are degenerate.

If a magnetic field, B, pointed in the z direction is applied, the energies and orbits of the levels change somewhat. Classically, in the absence of any confining potential, the magnetic field effectively confines electrons, causing them to move in circles in the clockwise direction with a frequency $\omega_c = eB/mc$. The quantum-mechanical energy levels are now given by:

$$E_{nl} = [l\hbar\omega_c/2 + \hbar\sqrt{(\omega_c/2)^2 + \omega_0^2}(2n + |l| + 1)] \qquad (2)$$

The energy has changed in two basic respects. There is now a first term that depends on l and not just the absolute value of l. Therefore, electron states with negative l values have lower energies than those with positive l values. Because the different l states represent wavefunctions that rotate in time about the dot centre, there is a magnetic moment associated with these states. This is the reason for the first term in equation (2). Positive l states have magnetic moments pointing opposite to the applied magnetic field, giving them higher energy than the negative l states whose moments point along the field. Finally, the second term in equation (2) is somewhat different from equation (1). The energies are higher, and this is because the magnetic field has enhanced the electron confinement in the dot.

The magnetic field also creates an energy difference between the spin-up and spin-down electrons in a state. The electron spins have an associated magnetic moment, and so in an externally

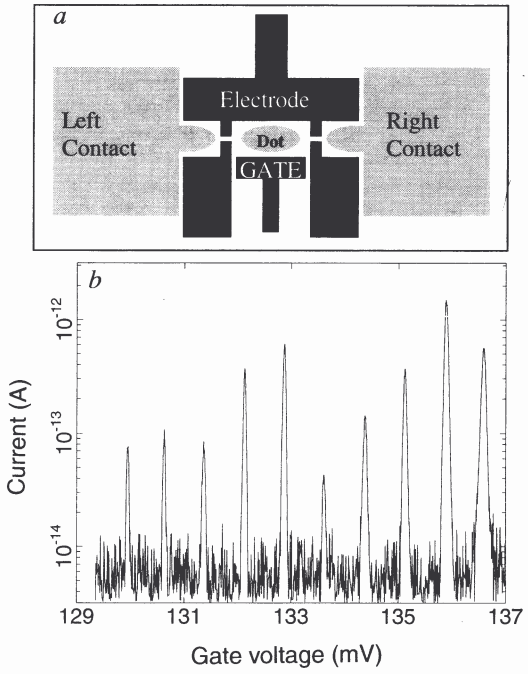

FIG. 2 *a*, Schematic diagram of samples used in gated transport experiments. *b*, Current through an artificial sample as a function of gate bias. The first peak on the left is thought to represent the addition of the thirtieth electron to the artificial atom. (Courtesy O. Klein)

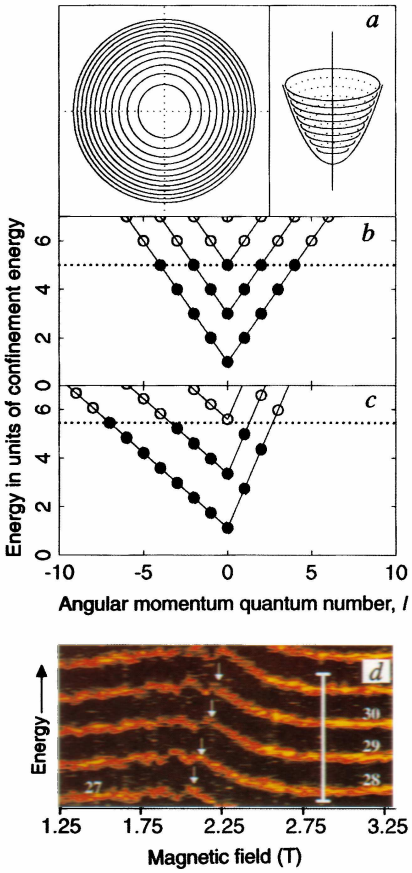

FIG. 3 *a* Left, orbits (as seen looking down on the two-dimensional artificial atom) for different values of magnitude of angular momentum $|l|$ in the artificial atom. The radius of each orbit increases as the square root of $|l|$. Right, orbits as seen in the bowl-like potential of the artificial atom. *b*, Energies of the different quantum levels as a function of angular momentum quantum number $|l|$ in zero magnetic field. The different 'Vs' correspond to different values of the radial quantum number, *n*. *c*, As *b*, but in the presence of a magnetic field. *d*, Colour-coded image from a SECS experiment. The red and yellow traces display the energy required to add each successive electron to the dot. To the right of the arrows, all electrons are in the lowest Landau level. The traces are numbered corresponding to the number of electrons contained in the artificial atom. The vertical bar corresponds to an energy of 5 meV.

applied magnetic field the electrons act rather like bar magnets, tending to align their moments with the external field (Zeeman effect). The energy difference between spin-up and spin-down states is often small compared to the other energy splittings in the quantum dot. To simplify the discussion, I first consider this 'spin splitting' here to be zero.

Figure 3*c* displays the quantum levels after the application of a magnetic field. Essentially two things have happened due to the field. The 'Vs' have rotated to the left, and the energy separation between the Vs has increased. If we again think of electrons as marbles filling slots, the marbles then fall from states on the right arm of the Vs to the left arms. Moreover, because the spacing between the Vs increases with increased field, the marbles fall from the upper to the lower Vs. In a magnetic field, we refer to each of the Vs as Landau levels. At very high fields, all of the marbles are in the left arm of the lowest V. In this case, all of the electrons are in the 'lowest Landau level'.

The peak positions in the SECS or the GTS experiment simply track the energy of the highest-energy electron in the quantum dot. This electron may be in a different quantum level depending on the magnetic field, and each of these levels has a different evolution in a magnetic field. Therefore, the peak position is expected to zig and zag as the highest-energy electron moves from state to state. In the absence of spin splitting, each of the energy levels depicted in Fig. 3*b* and *c* holds two electrons of precisely the same energy. In this case, one would expect the two electrons to undergo identical level shifts (zigs and zags) as the magnetic field strength is varied. In reality, each level of Fig. 3*c* is split into two different energies (two circles displaced vertically, one for spin up and the other for spin down) by the magnetic field. At 2 tesla (2 T), the amount of this splitting is only about 0.05 meV, or approximately 0.5 K in temperature units. This value is rather small compared to the energy differences between orbitals in a quantum dot which are typically around 1–2 meV.

Both GTS[25] and SECS[21] measurements demonstrate the zigzag behaviour. In the SECS experiments, one creates plots such as the one shown in Fig. 3*d*. This is a compendium of many data sets such as the one shown in Fig. 1*c*, now taken at varying values of the magnetic field strength. In Fig. 3*d*, the horizontal axis represents magnetic field strength and the vertical axis represents the voltage across the capacitor (or energy). The capacitance is plotted on a colour scale with yellow, red and black representing the highest, intermediate and lowest capacitances, respectively. Capacitance peaks are therefore signified by yellow and red traces. The traces clearly display a zigzag behaviour that ceases abruptly, just as we would expect from the simple-minded model of Fig. 3*c*. The zigzags stop at the positions of the arrows in Fig. 3*d*, the magnetic field strength beyond which all of the electrons are in the lowest Landau level.

Subtle differences exist between the observed traces and the results expected from the model of Fig. 3*b* and *c*. Because the spin splitting is very small, the model predicts pairs of two successive electron traces (such as for electrons 27 and 28 or electrons 29 and 30 in Fig. 3*d*) to have nearly identical zigzag features. This is not the case; no pairs of nearly identical traces are seen[21]. Also, the model predicts zigzags to occur well before the magnetic field value at which all electrons fall into the lowest Landau level; the data only display zigzags just below this field value. What is responsible for these and other deviations from the ideal model behaviour? One issue is most suspect: the simple model neglects the fact that the electrons interact with each other.

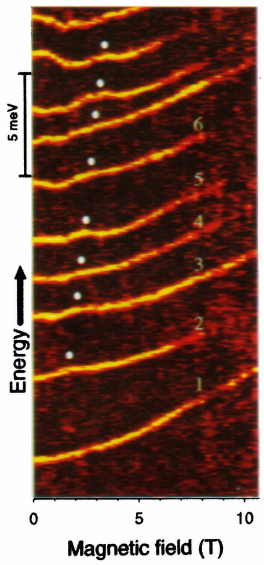

FIG. 4 *a*, Colour-coded image from a SECS experiment for the first ten electrons in a quantum dot. The lowest red and yellow trace displays the energy required to add the first electron into the dot. The white circles indicate a kink which is thought to be the field at which all of the electronic spins align with the magnetic field.

Adding one to ten electrons

There is another important change in the quantum states with the application of magnetic field. States of all l values shrink in radius as the magnetic field is increased. For high magnetic fields (when ω_c is larger than ω_0), the radius of a particular l state shrinks as $1/\sqrt{B}$. All of the circles shown in Fig. 3*a* tend to converge to the dot's centre.

Let us now examine the bottom trace of Fig. 4. This trace shows SECS results for the magnetic-field dependence of the energy of the quantum dot containing only one electron. Note that this energy increases as the magnetic field strength is increased owing to the enhanced confinement of the electron by the magnetic field. This behaviour of the one-electron energy turns out to be easily predictable using equation (2). This electron stays in the lowest energy state of the quantum dot, with $n = 0$ and $l = 0$. A fit to this trace using equation (2) determines that the diameter of the first electron's wavefunction is about 40 nm.

For more than one electron in the dot, interactions between electrons make understanding the spectra substantially more complicated and rich than one might conclude from the discussion above[30]. This is seen in the basic problem of just two electrons in a quantum dot.

The second trace from the bottom in Fig. 4 is for two electrons in the artificial atom. Notice that this trace appears qualitatively different from that of the first electron. Rather than smoothly moving up in energy, the two-electron energy shows a very clear kink at a magnetic field strength of 1.5 T. The origin of this kink, like the zigzags seen in Fig. 3*d*, is a shuffling of electrons between quantum states in the dot. In this case, it is the flipping of the spin of one of the two electrons within the quantum dot.

First, I consider two electrons in the dot with no Coulomb repulsion between them. Because of the Pauli exclusion principle, the two electrons must exist in quantum states that have at least one quantum number which is not the same. In zero magnetic field, the electron energy does not depend on the spin direction. These two electrons thus fall into the lowest available energy states which have quantum numbers $n = 0$, $l = 0$, and $s = \pm 1/2$. Such a state with oppositely pointing spins is known as a singlet state. At sufficiently high field, the electron which in zero field had its spin magnetic moment pointing against the applied field will, owing to the Zeeman effect, flip its spin so that the moment points along the field. When this happens, both electrons will have their spins pointing the same way, and they will have the same value of s. In this 'triplet' state, one of the electrons moves from an $l = 0$ state to an $l = 1$ state. Thus this singlet to triplet crossing arises when the Zeeman energy exceeds the energy to promote an electron to the $l = 1$ state. For the quantum dot of Fig. 4 in the absence of Coulomb repulsion, one would expect this crossing to occur at 25 T. In reality, the crossing is observed to occur at around 1.5 T. Why is this so?

The simple model of electrons as bar magnets clearly fails. When we include the Coulomb interaction between the electrons, we see that *it* actually drives the singlet to triplet crossing. In fact, the singlet–triplet crossing is predicted to occur even in the absence of a Zeeman interaction. The basic physics is simple, and I will approximate the actual case by considering one-electron wavefunctions. In reality, the wavefunctions change somewhat when the Coulomb interaction is added to the problem, and the electron motions are actually correlated in time to keep the electrons widely separated.

When the two electrons both have the same spatial quantum numbers (that is, $n = 0$ and $l = 0$ for both electrons), the energy of repulsion between the two electrons is large because the electron wavefunctions are identical. In contrast, when one of the two electrons moves into an $l = -1$ state, the electrons are on average farther apart, decreasing the Coulomb energy. There is a competition between the propensity for electrons to move towards the centre of the dot where the confinement potential is low and the tendency for electrons to repel each other. As the magnetic field strength is increased in the two-electron dot, states of larger $|l|$ move closer to the dot centre. When the radius of the $l = -1$ state becomes sufficiently small, the balance tips in favour of one electron moving from the dot centre out to the $l = -1$ state. To satisfy the Pauli exclusion principle, the wavefunction describing the two-electron system must change sign when the coordinates of the two electrons are interchanged. This turns out to mean that an electron must flip its spin direction as it moves to the $l = -1$ state. The two-electron problem has been solved exactly[31,32], and the predicted singlet to triplet crossing falls close to that observed in Fig. 4.

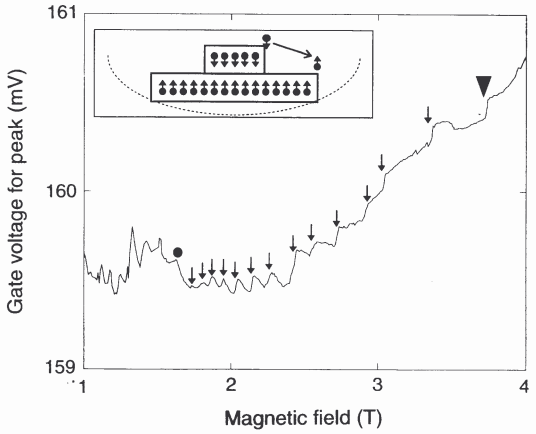

FIG. 5 Peak positions from a gated transport spectroscopy (GTS) experiment for a quantum dot containing about 30 electrons. The filled black circle marks the magnetic field value at which all of the electrons fall into the lowest orbital Landau level. The small arrows demarcate the spin flips of individual electrons in the dot. The last arrow on the right marks the field at which the electrons in the dot become spin-polarized. The black triangle marks the field values at which 'hole bunches' are created within the dot (adapted from ref. 34). Inset, schematic diagram of the spin-flipping process.

What about the dot containing more electrons? The spectra can no longer be predicted exactly. There exist several methods for deriving approximate spectra, the most accurate of which is known as 'exact diagonalization'[26,33–35]. This method is especially computer-intensive, and calculations can only be made for the case of a few electrons in the dot. In exact diagonalization, one typically approximates the many-body electron wavefunctions using combinations of a finite number of single-electron wavefunctions. These methods have been used to model the data in Fig. 4, and it is possible to identify the 'kinks' in these traces with predicted transitions. The main series of kinks, indicated by white circles in each of the traces, is thought to be the magnetic field value of the transition for flipping the last electron moment in line with the magnetic field[35]. Once again, these transitions are largely driven by the Coulomb interaction between electrons and not by the interaction of the electron moments with the external field.

Many-electron interactions

For yet more electrons, there is again strong evidence that transitions in the dot are driven by interactions between the electrons. Figure 5 displays peak positions as a function of magnetic field for conductance measurements in a GTS experiment at a temperature of about 0.08 K. This dot contains around 30 electrons. The filled circle indicates the magnetic field value, as marked by the arrows in Fig. 3*d*, where the last electron falls into the left arm of the lowest V of Fig. 3*c*. These lower-temperature measurements reveal something else. Beyond the filled circle, a new series of kinks develops in the spectra. As all of the electrons are in the lowest possible spatial energy states, these kinks must be due to spin-flips. As with the two-electron case above, the Zeeman interaction is simply too weak to cause these energy level crossings.

How do the interactions produce the observed level crossings? Let us start with the case of having all of the electrons in the lowest Landau level, and each orbital l state occupied by two electrons. With the application of stronger magnetic field, these orbits (shown in Fig. 3*a*) tend to converge towards the centre of the dot. If no electrons were to change their orbital states, in the limit of an infinite magnetic field all of the electrons would be compressed to a point at the centre of the dot. Because electrons repel each other, this obviously cannot happen. To lower the electron density at the centre of the dot, spin-down electrons are transferred to spin-up sites at the dot's edge. This process is shown schematically in Fig. 5 inset. Even though there is only a small energy difference between spin-down and spin-up states for non-interacting electrons, the spin-down states are depopulated because there exist lower-energy states at the edge of the dot.

How does one model this situation mathematically? The simplest scheme is to account for the electron repulsion by creating a simple 'mean-field' potential which each electron feels as a result of repulsion from all of the other electrons in their average positions. This potential is added to the confinement potential in the dot, and then the electrons are allowed to rearrange into different orbits to take this potential into account. As the electrons have moved, the problem needs to be solved again for the new electron distribution. The iterations continue until the electron distribution stops changing.

This 'self-consistent' approach[25] does indeed produce the type of spin-flipping behaviour suggested by the data on Fig. 5. However, it fails to explain this data in detail. The spin flips in the data occur more frequently at lower fields, and less frequently at higher fields. The self-consistent model predicts a much weaker dependence on field than demonstrated in observations. What is missing here?

Electrons attract?

The self-consistent model treats the electrons as though they are independent entities whose only interaction with other electrons is through the Coulomb repulsion and through the fact that individual electrons must occupy different single-particle orbits. In reality, the interactions are much more subtle. The electron motions inside the dot are actually correlated. For instance, in the dot containing two electrons, the correlations tend to place the electrons on opposite sides of the dot. Some of this happens automatically because electrons avoid each other as a result of the Pauli exclusion principle (the fact that the many-electron wavefunction must be antisymmetric under exchange of two electrons), and this phenomenon is known as the exchange interaction. The next most sophisticated calculation scheme beyond the self-consistent method is the 'Hartree–Fock' (HF) technique, which approximates the electron repulsion using the averaged fields of all of the electrons in their orbits. However, HF does take proper account of the exchange interaction. An HF calculation, when compared with the results of Fig. 5, indeed predicts nearly the same spin-flipping behaviour[36].

The exchange interaction leads to a remarkable phenomenon in quantum dots. HF calculations demonstrate an apparent short-range *attraction* between electrons. Considering wavefunctions for individual electrons, the exchange interaction tends to cause electrons to fill adjacent l states without leaving any gaps. This happens because the exchange interaction is operative when single-electron wavefunctions have some spatial overlap; when the wavefunctions overlap, the many-electron wavefunction must still go to zero when the two electron positions approach each other. This phenomenon means that on average electrons in adjacent orbits are actually kept further apart (for example, at any time, they are on opposite sides of the quantum dot) than are electrons in some more widely separated orbits.

The exchange interaction produces an interesting dilemma for electrons in a quantum dot in high magnetic field. Consider the situation where the process described in Fig. 5 inset has continued as the magnetic field has been increased, so that now all of the electrons in the dot have their spins pointing up. As the magnetic field is increased further, the orbits continue to converge to the centre of the dot. Once again, we are faced with the problem of all the electrons condensing to a point in the limit of high fields. One solution would be for electrons to leave 'holes' (empty l states) in the electron puddle—orbits in the middle of the puddle which are unoccupied—so that the electrons are not spaced too closely together. However, the exchange interaction makes the creation of holes energetically very costly.

HF calculations suggest the following scenario. To save exchange energy, it is wise to keep any holes in a bunch rather than dispersing the holes throughout the dot. Figure 6*b* depicts the predicted behaviour. The left-hand figure shows the electron droplet with the magnetic field set just above the point at which all of the electrons become spin-polarized. The right-hand figure shows the droplet at a field just beyond a critical value at which a bunch of holes is introduced into the droplet. By putting the holes into a bunch, most of the electrons retain their nearest neighbours and hence the short-range 'attraction' is satisfied. At the same time, the size of the droplet has expanded, and the energy due to electron repulsion has been reduced. The data set displayed in Fig. 5 and the SECS results shown in Fig. 6*a* both demonstrate evidence of bunches of holes being admitted into the electron droplet at fields beyond the fields required for spin polarization of the droplet. Figure 6*a* also shows evidence of subsequent bunches being admitted as the field is increased further.

The variations in brightness of the traces in Fig. 6*a* arise from differences in the probability for electron tunnelling between the artificial atom and the nearby metal[21,22]. The dim and bright regions of the traces result respectively from low and high tunnelling probabilities. One may interpret the tunnelling

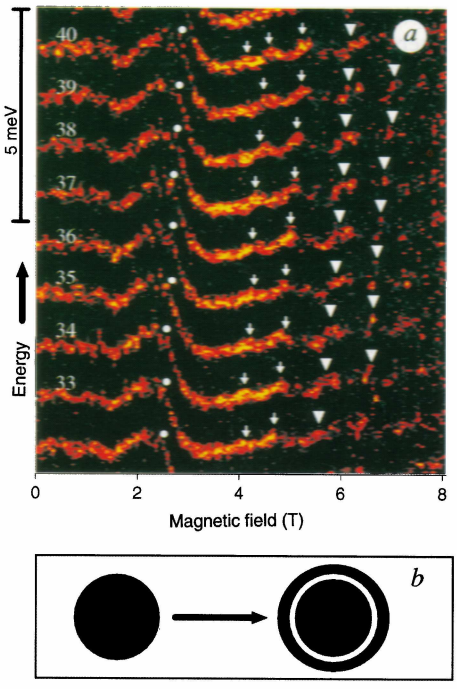

FIG. 6 *a*, Colour-coded image from a SECS experiment. The numbers labelling each trace signify the number of electrons in the artificial atom. The white circle indicates the magnetic field value at which all of the electrons fall into the lowest Landau level. The arrows point to the spin flips of individual electrons in the dot. The last arrow on the right marks the field at which the electrons in the dot become spin-polarized. The triangles indicate the field values at which 'hole bunches' are created within the dot. *b*, Schematic diagram of the hole-bunch creation process. The black regions depict areas of the dot which are populated with electrons.

probability as the willingness of the electrons of the artificial atom to make room for another electron. The tunnelling probabilities are enhanced at the positions where hole bunching occurs. They are instead diminished for magnetic fields at which HF calculations predict no holes in the interior of the artificial atom, such as beyond the rightmost arrow in Fig. 6*a* (at around 4.5 T) when the electrons in the artificial atom become completely spin-polarized.

Although the HF calculation captures the essence of the hole-bunch creation, it ignores the effect of electron correlation on the Coulomb repulsion. Exact diagonalization and other[37–39] calculations attempt to take these correlations into account. These suggest an even richer behaviour in the spectra as we move to higher magnetic fields and more detailed measurements.

Future research

The study of artificial atoms is a widely expanding research area. At present, the prospect of understanding electron correlations in a simple system has been a driving force for much of the theoretical work. All of the measurements discussed in this Review describe the spectroscopy of the lowest available energy levels of the artificial atoms. Measurements of excited states [26,40,41] constitute another class of measurements which may also yield significant information on the effects of correlations[42,43]

An exciting new field of research is emerging. Artificial 'molecules' consisting of two or more closely spaced artificial atoms are now being created[44,45] and studied theoretically[46,47]. These are typically structures similar to those of Fig. 2*a*, but which contain more than two or more dots. In some sense, these specimens are the first step toward artificial solids made of large arrays of linked dots[48]. Many of the key characteristics of real solids and real molecules, such as their magnetic and optical properties, and whether they are insulators or metals, depend largely on the strength of the couplings between the atoms. Unlike natural molecules, the strength of interactions between the neighbouring artificial atoms in artificial molecules can be varied at will (tuned) by varying the voltages on electrodes. The exciting feature here is that researchers may learn to engineer desired properties into artificial solids simply by tuning electron voltages that control the coupling between artificial atoms.

With the creation of the artificial atom, the ultimate limit of small-sized electronics is being achieved. So far, artificial atoms have allowed us to test the effects of interactions between electrons in an unprecedented way. An open question remains the practicality of using the ability to manipulate and trap single electrons for electronic devices. A device similar to that of Fig. 2*a* and known as the single-electron transistor (SET) has demonstrated spectacular sensitivity for sensing of electrical charge[49], and there have been suggestions for using either SETs[50] or artificial atoms themselves[51] as computing devices. Until recently, all single-electron devices functioned only at extremely low temperatures. Now, two novel structures[2,3] for single-electron transistors have been produced in silicon, unlike the gallium arsenide structures discussed in this Review. They both clearly display room-temperature operation. The smaller a single-electron device can be produced, the larger the energy for single-electron charging, and the higher the operation temperature. It turns out that advanced silicon processing, chiefly the ability to produce controllably very thin silicon oxide layers, facilitates the production of extremely small artificial atoms.

Despite these recent improvements in the operating temperature of the devices, two important problems remain. First, it has not yet proven feasible to create devices with reproducible thresholds for adding single electrons. Present digital logic systems require sharply defined thresholds for 'on' and 'off' states of a transistor. Second, single-electron devices are inherently low-current structures, and there is a 'fan-out' problem: unlike conventional transistors, the outputs of single-electron devices cannot rapidly switch many other devices. Researchers are now learning to circumvent both of these problems. For instance, schemes have been devised for memory circuits that combine conventional transistors and single-electron devices[52]. The method is insensitive to variations in device thresholds, and hybrid designs utilizing conventional transistors can defeat the fan-out problem. Because single-electron devices are so small, huge numbers of them can be packed in the space between conventional transistors on a chip. Such an implementation might yield nearly a thousand-fold increase in the storage density of a memory chip.

The work described in this Review has shown that researchers can now routinely measure extremely small amounts of electrical charge. In fact, it is now possible to measure a charge as small as 1/10,000 of an electron charge. Theory predicts that this number might reach as low as 1/1,000,000 of an electron charge[53]. Such a small charge might be induced on a single-electron device by a small motion of a single, distant charge. We are just now learning to make use of this extreme sensitivity. □

R. C. Ashoori is at the Department of Physics, 13-2053, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA.

1. Ziman, J. M. *Principles in the Theory of Solids* (Cambridge Univ. Press, 1972).
2. Yano, K. *et al. IEEE Int. Elect. Dev. Meeting* in *IEEE International Electron Devices Meeting* 541–544 (IEEE, New York, 1993).
3. Takahashi, Y. *et al. Electron. Lett.* **31,** 136–137 (1995).
4. Reed, M. A. *et al. Phys. Rev. Lett.* **60,** 535–537 (1988).
5. Meirav, U., Kastner, M. A. & Wind, S. J. *Phys. Rev. Lett.* **65,** 771–774 (1990).
6. Su, B., Goldman, V. I. & Cunningham, J. E. *Science* **255,** 313–315 (1992).
7. Main, P. C. *et al. Physica* B**189,** 125–134 (1993).
8. Guéret, P., Blanc, N., Germann, R. & Rothuizen, H. *Phys. Rev. Lett.* **68,** 1896–1899 (1992).
9. van der Vaart, N. C. *et al. Phys. Rev. Lett.* **73,** 320–323 (1994).
10. Schmidt, T. *et al. Phys. Rev.* B**51,** 5570–5573 (1995).
11. Hansen, W. *et al. Phys. Rev. Lett.* **62,** 2168–2171 (1989).
12. Silsbee, R. H. & Ashoori, R. C. *Phys. Rev. Lett.* **64,** 1991 (1990).
13. Ashoori, R. C., Silsbee, R. H., Pfeiffer, L. N. & West, K. W. in *Nanostructures and Mesoscopic Systems* (eds Reed, M. & Kirk, W.) 323–334 (Academic, San Diego, 1992).
14. Ashoori, R. C. *et al. Phys. Rev. Lett.* **68,** 3088–3091 (1992).
15. Sikorski, C. & Merkt, U. *Phys. Rev. Lett.* **62,** 2164–2167 (1989).
16. Meurer, B., Heitmann, D. & Ploog, K. *Phys. Rev. Lett.* **68,** 1371–1374 (1992).
17. Bawendi, M. G., Steigerwald, M. L. & Brus, L. E. *A. Rev. phys. Chem.* **41,** 477–496 (1990).
18. Brunner, K. *et al. Phys. Rev. Lett.* **69,** 3216–3219 (1992).
19. McEuen, P. L. *et al. Phys. Rev. Lett.* **66,** 1926–1929 (1991).
20. Cho, A. Y. *J. Cryst. Growth* **111,** 1–13 (1991).
21. Ashoori, R. C. *et al. Phys. Rev. Lett.* **71,** 613–616 (1993).
22. Ashoori, R. C. *et al. Surf. Sci.* **305,** 558–565 (1994).
23. Weis, J., Haug, R. J., von Klitzing, K. & Ploog, K. *Phys. Rev.* B**46,** 12837–12840 (1992).
24. Borsosfoldi, Z. *et al. Appl. Phys. Lett.* **66,** 3666–3668 (1995).
25. McEuen, P. L. *et al. Phys. Rev.* B**45,** 11419–11422 (1992).
26. Johnson, A. T. *et al. Phys. Rev. Lett.* **69,** 1592–1595 (1992).
27. Staring, A. A. M. *et al. Phys. Rev.* B**46,** 12869–12872 (1992).
28. Kastner, M. A. *Physics Today* **46,** 24–31 (1993).
28. Darwin, C. G. *Proc. Camb. phil. Soc. Math. phys. Sci.* **27,** 86–90 (1930).
30. Maksym, P. A. & Chakraborty, T. *Phys. Rev. Lett.* **65,** 108–111 (1990).
31. Wagner, M., Merkt, U. & Chaplik, A. V. *Phys. Rev.* B**45,** 1951–1954 (1992).
32. Pfannkuche, D., Gudmundsson, V. & Maksym, P. A. *Phys. Rev.* B**47,** 2244–2250 (1993).
33. MacDonald, A. H. & Johnson, M. D. *Phys. Rev. Lett.* **70,** 3107–3110 (1993).
34. Hawrylak, P. *Phys. Rev. Lett.* **71,** 3347–3350 (1993).
35. Palacios, J. J., Martin-Moreno, L., Chiappe, G., Louis, E. & Tejedor, C. *Phys. Rev.* B**50,** 5760–5763 (1994).
36. Klein, O. *et al. Phys. Rev. Lett.* **74,** 785–788 (1995).
37. Jain, J. K. & Kawamura, T. *Europhys. Lett.* **29,** 321–326 (1995).
38. Macucci, M., Hess, K. & Iafrate, G. J. *Phys. Rev.* B**48,** 17354–17363 (1993).
39. Ferconi, M. & Vignale, G. *Phys. Rev.* B**50,** 14722–14725 (1994).
40. Foxman, E. B. *et al. Phys. Rev.* B**47,** 10020–10023 (1993).
41. Weis, J., Haug, R. J., v. Klitzing, K. & Ploog, K. *Phys. Rev. Lett.* **71,** 4019–4022 (1993).
42. Meir, Y., Wingreen, N. S. & Lee, P. A. *Phys. Rev. Lett.* **70,** 2601–2604 (1993).
43. Oaknin, J. H., Martín-Moreno, L., Palacios, J. J. & Tejedor, C. *Phys. Rev. Lett.* **74,** 5120–5123 (1995).
44. van der Vaart, N. C. *et al. Phys. Rev. Lett.* **74,** 4702–4705 (1995).
45. Waugh, S. R., Westervelt, R. M., Campman, K. & Gossard, A. C. *Phys. Rev. Lett.* **75,** 705–708 (1995).
46. Palacios, J. J. & Hawrylak, P. *Phys. Rev.* B**51,** 1769–1777 (1995).
47. Matveev, K. A., Glazman, L. I. & Baranger, H. U. *Phys. Rev. Lett.* (submitted).
48. Stafford, C. A. & Das Sarma, S. *Phys. Rev. Lett.* **72,** 3590–3593 (1994).
49. Averin, D. V. & Likharev, K. K. in *Mesoscopic Phenomena in Solids* (eds Altshuler, B. L., Lee, P. A. & Webb, R. A.) 173–271 (North-Holland, Amsterdam, 1991).
50. Likharev, K. K. *IBM J. Res. Dev.* **32,** 1444–1158 (1988).
51. Tougaw, P. D. & Lent, C. S. *J. appl. Phys.* **75,** 1818–1825 (1994).
52. Likharev, K. K. & Korotkov, A. N. in *ISDRS'95* (IEEE, New York, in the press).
53. Averin, D. V. & Likharev, K. K. in *Single Charge Tunneling* (eds Grabert, H. & Devoret, M. H.) 311–332 (Plenum, New York, 1992).

ACKNOWLEDGEMENTS. I am grateful to my collaborators at AT&T Bell Laboratories together with whom I developed SECS measurements during the time I spent there: H. Stormer, J. Weiner, L. Pfeiffer, K. Baldwin and K. West. I thank M. Bawendi and M. Kastner for their comments on this manuscript. The work described here on SECS measurements is supported primarily by the US Office of Naval Research and the David and Lucille Packard Foundation.