

Computing with Atoms and Molecules

Prospects of Harnessing Quantum Mechanics for Faster Computers

Christopher Monroe and David Wineland
National Institute of Standards and Technology, Boulder, CO 80303

In 1965, Intel co-founder Gordon Moore predicted that the number of electronic components on a computer chip would double every year or two. “Moore’s Law” has been remarkably accurate even to this day, where the latest silicon processors now host some ten million transistors. This exponential growth in chip density has been the driving force behind the modern electronic age and has played a significant role in the booming world economy over the last couple decades. Will Moore’s Law continue indefinitely? Well, no. The problem is, adding components by simply expanding the size of chips will bog down their speed because of the time it takes for electronic signals to traverse the chip, so the only way to sustain this growth is to make transistors ever smaller. By the year 2010 or so, Moore’s Law predicts that a one-square-centimeter chip will have over one billion transistors, with each transistor approaching the size of a large molecule. At this point, significant further gains will require a fundamental change in the way we compute, for once we’ve entered the regime of individual molecules, atoms, and electrons, the laws of quantum mechanics reign.

In the early 1980s, Paul Benioff (Argonne National Laboratory) and Richard Feynman (California Institute of Technology) tinkered with the idea of quantum-mechanical computing elements such as single atoms. They showed that such tiny structures could, in principle, behave perfectly fine as electronic components. They even discussed using “quantum logic gates” largely following the laws of quantum mechanics. In 1985, David Deutsch (Oxford University) went a step further. By using the full arsenal of quantum mechanical rules, he proposed that the phenomenon of “quantum superposition” be harnessed to yield massively parallel computing – computing with multiple inputs at once in a single

device. Instead of miniaturizing chip components further, Deutsch posed a new way to get around the impending limit of Moore’s Law by taking advantage of different physical principles underlying these components.

So far, only a few algorithms such as number factoring have been discovered which can benefit from quantum parallelism, and it remains speculative whether quantum computers will ever replace the all-purpose computers in use today. However, interest in quantum information processing is now proliferating because its limits are not known, neither theoretically nor experimentally. There may well be larger classes of algorithms which benefit from quantum computers, and even though large-scale quantum computers have not yet been built, there do not seem to be any fundamental roadblocks. Thus, many consider quantum computers to be an intriguing possibility to circumvent the atomic limits to Moore’s Law.

QUBITS, QUANTUM PARALLELISM, and QUANTUM ALGORITHMS

Quantum mechanics is now hailed as one of the most successful theories in the history of science, as it has precisely predicted the behavior of the microscopic world from molecules and atoms to sub-nuclear phenomena. Although quantum mechanics underlies the whole of physics, there are aspects of it that appear highly non-intuitive since we are used to dealing with ordinary-life, macroscopic situations. At the heart of the matter is the concept of a quantum superposition, where a quantum object can exist in two states *at the same time*. The superposition “collapses” to one of its states only when the quantum object is measured by a macroscopic object, such as a

meter or a pair of eyes. This dichotomy bothered Albert Einstein, among others, who struggled to determine a theory which would unify quantum mechanics with the macroscopic world we live in.

Erwin Schrödinger illustrated this paradox vividly with his famous ‘Schrödinger’s Cat’ thought experiment: Imagine a cat isolated from the world in a cage with a single radioactive atom which has not yet decayed. The radiation of the atom is detected by a Geiger counter which triggers a hammer to smash a nearby vial of cyanide. The paradox arises when we consider what happens after a duration of one half-life of the radioactive atom. Here, quantum mechanics predicts that the atom is in a superposition of its original state and its decayed state. It follows then, that the cat must also be in a superposition of being both alive and dead! Dealing with this ridiculous situation usually involves erecting an artificial barrier between quantum and classical worlds – microscopic atoms and molecules are allowed to be in superpositions, but macroscopic objects like cats are not. This approach has largely allowed quantum mechanics to coexist peacefully with classical physics. But a quantum computer, hosting superpositions of information, is just a more humane version of Schrödinger’s Cat, and may push this artificial quantum/classical barrier closer to the macroscopic world.

The smallest piece of classical information is the binary digit, or bit, which can be either 0 or 1. The simplest quantum-mechanical unit of information is the quantum bit or ‘qubit.’ Qubits are able to store superpositions of 0 and 1, denoted by $\alpha|0\rangle + \beta|1\rangle$, where α and β are the weights of the superposition. In this notation, $|x\rangle$ signifies a quantum state and the + sign indicates a superposition. The states $|0\rangle$ and $|1\rangle$ may represent, for example, horizontal and vertical polarization of a single photon, or two particular energy levels within a single atom. The rules of quantum mechanics dictate that: (a) the evolution of the weights α and β is described by the Schrödinger wave equation, and (b) when the above quantum bit is measured, it yields either $|0\rangle$ or $|1\rangle$ with probabilities related to the weights α and β , respectively. The measurement of a quantum bit is much like flipping a coin – the results can only be described within the framework of probabilities.

The power of quantum computing is seen by considering a register of many qubits. As indicated in the sidebar, in general, N qubits can store a superposition of all 2^N binary numbers:

$$\gamma_1|000\dots 0\rangle + \gamma_2|000\dots 1\rangle + \dots + \gamma_{2^N}|111\dots 1\rangle.$$

Moreover, when a quantum computation is performed on this superposition, each piece gets processed in superposition. For example, quantum logic operations can shift all the qubits one position to the left, equivalent to multiplying the input by two. When the input state is in superposition, all inputs are simultaneously doubled with one turn of the crank (see Fig. 1a).

After this quantum parallel processing, the state of the qubits must ultimately be measured. Herein lies the difficulty in designing useful quantum computing algorithms. According to the laws of quantum mechanics, this measurement yields just one answer out of 2^N possibilities; worse still, there is no way of knowing which answer will appear! It seems quantum computers do not compute one-to-one functions (where each input results in a unique output as in the doubling algorithm above) any more efficiently than classical computers. The trick behind a useful quantum computer algorithm involves the phenomenon of quantum interference. Since the weights $\gamma_1, \gamma_2 \dots \gamma_{2^N}$ in the superposition (1) evolve according to a wave equation, they can be made to interfere with each other, as in any wavelike disturbance. Some weights interfere constructively, like the crests of an ocean wave, while others cancel, like when a valley meets a crest. In the end, the parallel inputs are processed with quantum logic gates so that almost all of the weights cancel, leaving only a very small number of answers, or even a single answer, as depicted in Fig. 1b. By measuring this answer (or repeating the computation a few times and recording the distribution of answers), information can be gained pertaining to all 2^N inputs.

In 1994, Peter Shor (AT&T Bell Labs) constructed a quantum algorithm to calculate the factors of a number M ; that is, numbers p and q whose product is M . He showed that a quantum computer is able to factor numbers exponentially faster than classical computers. This discovery led to a rebirth of interest in quantum computers, in part due to the importance of factoring for cryptography. The security of popular cryptosystems such as those used for internet commerce is derived from the *inability* to factor large numbers. More generally, Shor’s factoring algorithm also proved that quantum computers are indeed good for something, spurring physicists, mathematicians, and computer scientists to search for other algorithms amenable to quantum computing. In 1996, for example, Lov Grover (AT&T) proved that a quantum computer can search unsorted databases faster than any search conducted on a classical computer. Still, useful quantum algorithms are not plentiful, and it is unknown how many classes of problems will ultimately benefit from quantum computation. Nevertheless, quantum computing has

stimulated scientists to think about quantum mechanics in terms of information processing, and this may ultimately help unify the quantum and classical worlds, and perhaps resolve the paradox of Schrödinger's Cat.

QUANTUM COMPUTER HARDWARE

On the experimental side, what resources are needed to implement a useful quantum computer? In terms of the factoring algorithm, classical computers run out of steam when factoring numbers with a hundred-or-so digits (recently, a team of a few hundred computers took 5 months to factor a 155-digit number). A useful quantum factoring engine would therefore have to factor a number with at least 200 digits. This implies that it would need a few thousand qubits, and more than 10^9 quantum logic gate operations. Unfortunately, state-of-the-art technology hosts only a few qubits and can do only a few quantum logic operations. Quantum computing hardware is far behind the software, mainly because it is very difficult to maintain quantum-mechanical superpositions throughout the computation. Consider the following stringent (and apparently contradicting) hardware requirements for a quantum computer:

- (1) The qubits must be sufficiently shielded from the environment during the computation, as external influences behave like measurements and destroy quantum superpositions.
- (2) The qubits must interact strongly with each other in a controlled fashion to allow the formation of quantum logic gates and entangled superpositions.
- (3) The qubits must ultimately be measured through a controlled strong coupling to the environment represented by a measuring device.

There obviously must be a sufficient number of quantum bits available for a given application. Because this may involve thousands of qubits, it is not clear that a useful quantum computer will ever be built and, if it is built, what form it will take. One naturally looks to a condensed-matter or solid-state system because of the fantastic success of classical computers, but the requirements for quantum computers are significantly different and in some cases opposite. For example, classical computers rely on dissipation, or a strong coupling to the environment which ensures that a digital '1' remains a '1' and a '0' remains a '0' until we deliberately flip the bit. This "latching" feature allows classical computers to operate even in the presence of noise. In contrast, in a quantum computer, we want superpositions to be preserved until the final measurement; before this, any dissipation is bad! A variety of possible schemes are being investigated; at first they look quite different, but

they have many similarities. For brevity, we highlight schemes based in atomic physics currently considered to be the most attractive.

A MODEL QUANTUM COMPUTER

Two internal states in an atom can be used as qubit levels. Some states interact weakly with the environment (requirement 1) so that superpositions can be preserved for as long as hours; this is the same attribute that makes atoms good clocks. For instance, the qubit levels can be represented by the orientation of the intrinsic spin of the atom with respect to an external magnetic field, like the two states of a proton in a magnetic field (spin aligned or anti-aligned with the field). The intrinsic spin of an atom has an associated magnetic moment and magnetic field similar to that of a tiny bar magnet. Thus, these distinct energy states are much like those corresponding to different alignments of a bar magnet immersed in a magnetic field. Alternatively, spins embedded in a solid or liquid host could be used as qubits, but satisfying both requirements 1 and 2 can be difficult. Nuclear-magnetic-resonance (NMR) techniques can be exploited to derive qubits from spins within molecules. Requirement 1 can be satisfied by using spins which are sufficiently isolated from the environment, and requirement 2 can be satisfied by using the naturally occurring magnetic interactions between spins. Unfortunately, it is difficult to initialize and read out the quantum states in a molecule, and this system cannot readily be scaled to many quantum bits due to the finite number of spins on the molecule. Moreover, unless the molecular spins are cooled to near absolute zero, an exponential loss in signal strength prevents this system from being scaled beyond a few qubits.

In 1995, Ignacio Cirac and Peter Zoller (University of Innsbruck) suggested a way to construct a quantum computer with trapped atoms. As shown schematically in Fig. 2, the qubit states in this implementation are two spin states within an atomic ion. The ion qubits are confined in an electrode structure whose potentials provide a 3-dimensional harmonic trap, analogous to the 2-D well formed by a marble in a bowl. If the ions are cooled, each ion wants to rest at the bottom of the 3-D well but their Coulomb repulsion holds them apart; a balance between the trap and repulsion forces results in the ions forming a regular array. As shown in Figs. 2 and 3, the ions can be made to form a linear array by making the trap in the horizontal direction much weaker than the other two directions. In principle, there is no limit to the number of qubits which can be held in the trap. For typical conditions, the ions are

separated by a few micrometers and their spins are well-isolated from the environment and from each other, satisfying requirement 1. The qubits are coupled through their collective motion, satisfying requirement 2. At low temperature, the ions form a quasi-molecule and their motions are best viewed in terms of so-called normal modes. For example, for a trap containing two ions, the normal modes in the horizontal direction are the center-of-mass mode where the ions oscillate back and forth in unison and the “stretch” mode where they oscillate in opposition, much like a pair of pendulums connected by a spring. Using laser-cooling techniques, the motion in these modes can be nearly frozen out; they can be cooled to such a degree that the modes are put in their quantum-mechanical ground states. The ground state (labeled $|0\rangle$) and the first-excited state ($|1\rangle$) of motion of a selected mode themselves form a qubit. This qubit is special in that it is shared by all the ions; therefore, it can be regarded as a data-bus bit through which information can be transferred.

We now briefly discuss how to (a) prepare a qubit superposition state in a given ion, (b) transfer the information (the superposition state) of the selected ion onto the motional qubit, and finally (c) form a quantum logic gate between the motional qubit and another selected ion. We assume the internal state qubits are derived from a lower and upper energy orientation of the ion’s spin (labeled $|\downarrow\rangle$ and $|\uparrow\rangle$) with respect to a background uniform magnetic field, as outlined above. The ion’s spin state can be altered by applying an oscillating magnetic field. Not only can we flip its state from $|\downarrow\rangle$ to $|\uparrow\rangle$, but we can also make arbitrary superposition states $|\downarrow\rangle \rightarrow \alpha|\downarrow\rangle + \beta|\uparrow\rangle$ by applying the oscillating field for particular times. (These are the same types of manipulations performed on nuclear spins in MRI [magnetic resonance imaging].)

In addition to the uniform field, suppose we now superimpose a field whose magnitude depends on position; for simplicity we’ll assume the magnitude of the additional field is $+\Delta B$ at the ion’s left most position and $-\Delta B$ at its right-most position. When the ion oscillates back and forth, it now sees a field oscillating at the motional oscillation frequency with amplitude ΔB . When this frequency exactly matches the frequency associated with the difference in energy between the two spin states, energy is exchanged between the spin and the motion: $|\uparrow\rangle|0\rangle \rightarrow |\downarrow\rangle|1\rangle$. Moreover, if this operation is applied to a superposition state of the spin we can map an arbitrary state of the spin qubit to the motional data-bus qubit: $(\alpha|\downarrow\rangle + \beta|\uparrow\rangle)|0\rangle \rightarrow |\downarrow\rangle(\alpha|0\rangle + \beta|1\rangle)$. (The transition $|\downarrow\rangle|0\rangle \rightarrow |\uparrow\rangle|1\rangle$ cannot occur because it doesn’t conserve energy.) To localize the field gradient to a

particular ion, we use the field gradient associated with the fields of a laser beam focused onto the ion.

Very similar techniques are used to perform a logic operation between the motional qubit and another ion (selected by focusing the laser beam on this second ion). One kind of logic-gate which has been realized experimentally is similar to a classical exclusive-OR (XOR) gate; it is called a controlled-NOT (CN) gate. Its associated truth table, which describes the change in the state of the system before and after the operation, is:

$$\begin{aligned} |\downarrow\rangle|0\rangle &\rightarrow |\downarrow\rangle|0\rangle \\ |\uparrow\rangle|0\rangle &\rightarrow |\uparrow\rangle|0\rangle \\ |\downarrow\rangle|1\rangle &\rightarrow |\uparrow\rangle|1\rangle \\ |\uparrow\rangle|1\rangle &\rightarrow |\downarrow\rangle|1\rangle \end{aligned} \quad (2)$$

Here the motional qubit is called the control qubit; if it’s state is $|0\rangle$, the spin bit remains unchanged, if it is a $|1\rangle$, the spin bit flips. This truth table appears to be classical, but it can also act on superposition states. For example, we can transfer a superposition of the motional qubit into an entangled superposition between the motion and spin qubit:

$$|\downarrow\rangle(\alpha|0\rangle + \beta|1\rangle) \text{---(CN)} \rightarrow \alpha|\downarrow\rangle|0\rangle + \beta|\uparrow\rangle|1\rangle.$$

There are similar proposals for quantum logic on trapped neutral atoms. In this case, because the Coulomb interaction is absent, requirement 2 is satisfied by selectively moving the atoms very close together to realize coupling through a direct spin-spin interaction. In a related scheme called cavity-QED (cavity-quantum-electrodynamics), atoms communicate via photons that are confined in a common optical cavity or optically transferred between separate cavities. Some proposed condensed matter schemes for quantum computers are quite similar to these atomic-physics schemes. The qubits could be composed of electron or nuclear spins which are localized to particular sites in a solid host. If these sites are far enough apart, then the separate spins interact negligibly and superposition states can be preserved. If the sites are not too far apart, the electron clouds surrounding the sites can be distorted by the fields from a nearby surface electrode so that the neighboring electron spins interact with and couple to each other. This interaction could be used to realize a controlled-NOT gate.

Regardless of the physical system, an arbitrary quantum computation (e.g., generating the arbitrary superposition of Eq. (1)) can be constructed from a series of single-qubit manipulations and two-qubit controlled-NOT gates. These two gates form a

“universal” logic family, much like the AND and NOT gates in classical computing.

After the computation, we need an efficient way to read out the qubit states (requirement 3). In some condensed-matter schemes, this might be accomplished by coupling a site’s electron cloud to the gate of a single-electron transistor. Since the coupling can be made spin-dependent, it can be used to measure the site’s electron spin state. In atomic physics, a common technique for detection is to use state-dependent laser scattering. In the trapped-ion example, if the ion is in the $|\downarrow\rangle$ state it can be made to scatter many photons, while an ion in the $|\uparrow\rangle$ state scatters negligibly. This disparity can be so large that even if we detect only a small fraction of scattered photons we can still distinguish the $|\downarrow\rangle$ state from the $|\uparrow\rangle$ state with nearly 100% efficiency.

Realizing a computation composed of even a few hundred operations on 10 to 20 qubits will be a significant technical achievement. However, for quantum computation to be generally useful it must be scalable to much larger numbers. This is viewed as a long-term disadvantage of the NMR system as it will become very difficult to scale beyond approximately 10 qubits. Very large scale integration (VLSI) is a hallmark of modern classical computers; therefore, if the basic elements of quantum computation can be realized in a condensed-matter system, a host of existing methods might be employed to scale up to large size systems. The ion trap scheme appears to suffer in a similar way to the NMR scheme; specifically, as the number of ions in a trap increases, it will become more and more difficult to avoid coupling to unwanted motional resonances. However, this system could be multiplexed if we make arrays of traps comprising accumulators and storage cells, each containing only a few ions. Moving selected ions then transfers information. In condensed-matter systems, similar transfers might be accomplished by moving electrons between sites.

OUTLOOK

Currently, room-temperature molecular NMR has demonstrated simple quantum manipulations on a few qubits. However, this system is in a highly mixed state and entanglement is not manifest; this is related to why it is not scalable to large numbers. All the elements of the trapped-ion and cavity-QED schemes have been demonstrated on one and two atoms. For trapped ions, efforts to scale up to higher numbers of ions have been hampered by decoherence of the motional qubit, which is highly susceptible to external electric field fluctuations; however, it appears to be

straightforward to overcome these effects. A number of condensed-matter schemes look promising, but superpositions are very short-lived and quantum gates have not yet been demonstrated. From any perspective, a rather large “abyss” exists between theory and experiment as indicated in Fig. 4. Certainly future experiments will be able to handle more qubits and accomplish more coherent operations, but hopefully the gap will also be closed through future theoretical developments.

One recent theoretical breakthrough has been the idea of quantum error-correction. Here, classical error-correction schemes have been adapted to the quantum world. Errors, such as bit-flips caused by uncontrolled external field fluctuations or unwanted measurement by the environment, can be detected and corrected by measuring a subset of qubits and performing subsequent logic operations. Qubit superpositions can be maintained through these operations; a situation seemingly implausible in quantum mechanics because we are taught early on that measurement destroys all superpositions! In a long computation, such as factoring big numbers, employing error correction requires a fidelity significantly less than what would be required without it.

Independent of the future of quantum computing, the entire field of quantum-information science is still in its infancy. Certainly many issues are not understood and one can be optimistic that their resolution may lead to new insights into quantum mechanics, its relationship to information science, and its meaning in the physical world.

FURTHER READING

Colin P. Williams and Scott Clearwater, [Explorations in Quantum Computing](#), Springer-Verlag, New York (1998).

Seth Lloyd, “Quantum-Mechanical Computers,” *Scientific American*, October 1995, pp. 140-5.

David Deutsch and Artur Ekert, “Quantum Computation,” *Physics World*, March 1998, pp. 47-52.

SIDEBAR: SPOOKY COMPUTING

Quantum bits (qubits) can be prepared in superposition states of 0 and 1. Following quantum mechanical notation, a qubit superposition is expressed as $\alpha|0\rangle + \beta|1\rangle$, where α and β are the weights of the superposition. A measurement of this superposition is much like flipping a biased coin – the qubit yields a ‘0’ or ‘1’ with relative probabilities related to α and β , respectively.

With N qubits, we can construct the state

$$(\alpha_1|0\rangle_1 + \beta_1|1\rangle_1) (\alpha_2|0\rangle_2 + \beta_2|1\rangle_2) \cdots (\alpha_N|0\rangle_N + \beta_N|1\rangle_N), \quad (1s)$$

having $2N$ weights. But the most general state of N qubits is a superposition of all 2^N binary numbers:

$$\gamma_1|000\dots 0\rangle + \gamma_2|000\dots 1\rangle + \cdots + \gamma_{2^N}|111\dots 1\rangle, \quad (2s)$$

where the γ_i are the 2^N weights. The qubits are said to be *entangled* when the overall quantum state cannot be factored into individual qubit states as in Eq. (1s). Entangled qubits are correlated with each other when measured, even though each qubit is not in a definite state (0 or 1) before measurement. For example, each of the two qubits in the simple entangled state $|00\rangle + |11\rangle$ will always collapse to the same state when measured, yet neither qubit is in a well-defined state before the measurement. This implies that entangled qubits are somehow interconnected, even though there may not be any physical interaction between them. This remains true even when the qubits are separated by arbitrarily large distances. Albert Einstein described this situation as ‘spooky action-at-a-distance,’ and it remains one of the most mysterious features of quantum mechanics.

Quantum computing harnesses the implicit interconnections in entangled quantum states to speed up computations. To illustrate, consider the simple procedure of flipping the first qubit of an N -bit quantum computer. If the computer is prepared in the unentangled state of Eq. (1s), then this operation can always be reduced to swapping the weights α_1 and β_1 – a single operation. If the computer is instead prepared in the entangled state of Eq. (2s), then flipping just the first qubit is equivalent to swapping weights γ_1 with γ_{2^N-1+1} , γ_2 with γ_{2^N-1+2} , ... and γ_{2^N-1} with γ_{2^N} . This is a total of 2^{N-1} swap operations, exponentially more than in the unentangled case.

FIGURE CAPTIONS

Figure 1. Simplified evolution during a quantum algorithm on $N=3$ quantum bits. The inputs are prepared in superposition states of all $2^N = 8$ possible numbers (written in binary). The weights of the superposition are denoted by the greyscale, where black is a 100% weight and white is a 0% weight. (a) Quantum algorithm for simultaneously doubling all input numbers (Modulo 7), by shifting all qubits one position to the left and wrapping around the leftmost bit. The outputs are also in superposition, and a final measurement projects one answer at random. (b) Quantum algorithm involving wavelike interference of weights. Here, quantum logic gates cause the input superposition to interfere, ultimately canceling all of the weights except for one (101 in the figure) which can then be measured. For some algorithms, this lone answer (or the distribution of a few answers after repeated runs) can depend on the weights of all 2^N input states, leading to an exponential speedup over classical computers.

Figure 2. A Schematic representation of a model quantum computer based on trapped atomic ions. Electric potentials are applied to the horizontal cylindrical electrodes in order to provide a three-dimensional harmonic “trap.” The trap is relatively weak in the horizontal direction so that the ion equilibrium positions (at low temperature), determined by a balance between the trap and Coulomb repulsion forces, is a linear array. The ions’ internal qubit states can be formed by the direction of the magnetic moment in the ion which is parallel ($|\downarrow\rangle$ state of lower energy) or antiparallel ($|\uparrow\rangle$ state of higher energy) to an externally applied magnetic field. The two lowest energy states (denoted $|n=0\rangle$ and $|n=1\rangle$) of a selected mode of motion form an additional qubit. This qubit is shared among all the ions and can therefore provide a data bus. Logic operations are initiated by mapping the internal state superposition of one ion onto the shared motional qubit using a focused laser beam. A quantum logic gate is then performed between the motional qubit and a second ion by focusing the beam on that ion. To complete the logic operation between the two selected ions, the initial mapping operation is reversed.

Figure 3. Top: Ion trap electrode structure on a 10mm x 15mm ceramic substrate. The gold trap electrodes and leads (brownish color) have been evaporated onto the substrate, and associated circuit elements are visible on the surface of the structure. The trapping region is a horizontal slit which has been cut from the substrate, with an expanded view at the bottom of the figure. **Bottom:** Closeup of the ion trap region. Four beryllium ions which have been cooled to near their ground state of motion form a linear crystal in the trap; an image of the ions is superposed on the visible picture of the trap electrodes. Radio-frequency potentials of about 500 V at near 250 MHz have been applied between the upper and lower electrodes, and static potentials of several volts are applied between members of the segmented upper electrode. The trap apparatus is maintained in high vacuum (pressure $\cong 10^{-9}$ Pa, or about 10^{-14} of atmospheric pressure) to prevent heating and chemical reactions of the ions with the background gas.

Figure 4. The “Quantum-computing Abyss.” A large disparity exists between what is currently possible and what is required for large-scale problems like factorizing big numbers. Quantum error-correction techniques imply that by encoding two-state superpositions in entangled states composed of several qubits, the required number of error-free operations can be reduced from about 10^9 (factoring without error-correction) to about 10^4 or 10^5 . In several proposed implementations of a quantum computer this appears to be technically feasible.

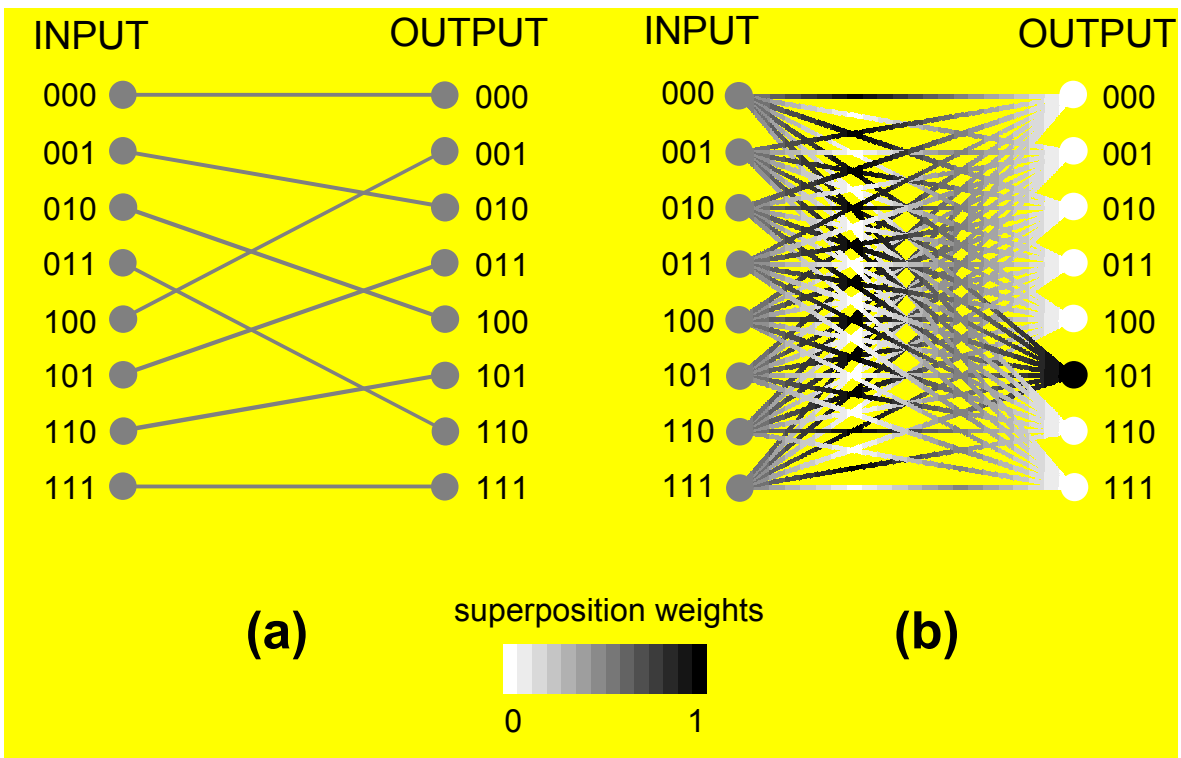


Figure 1

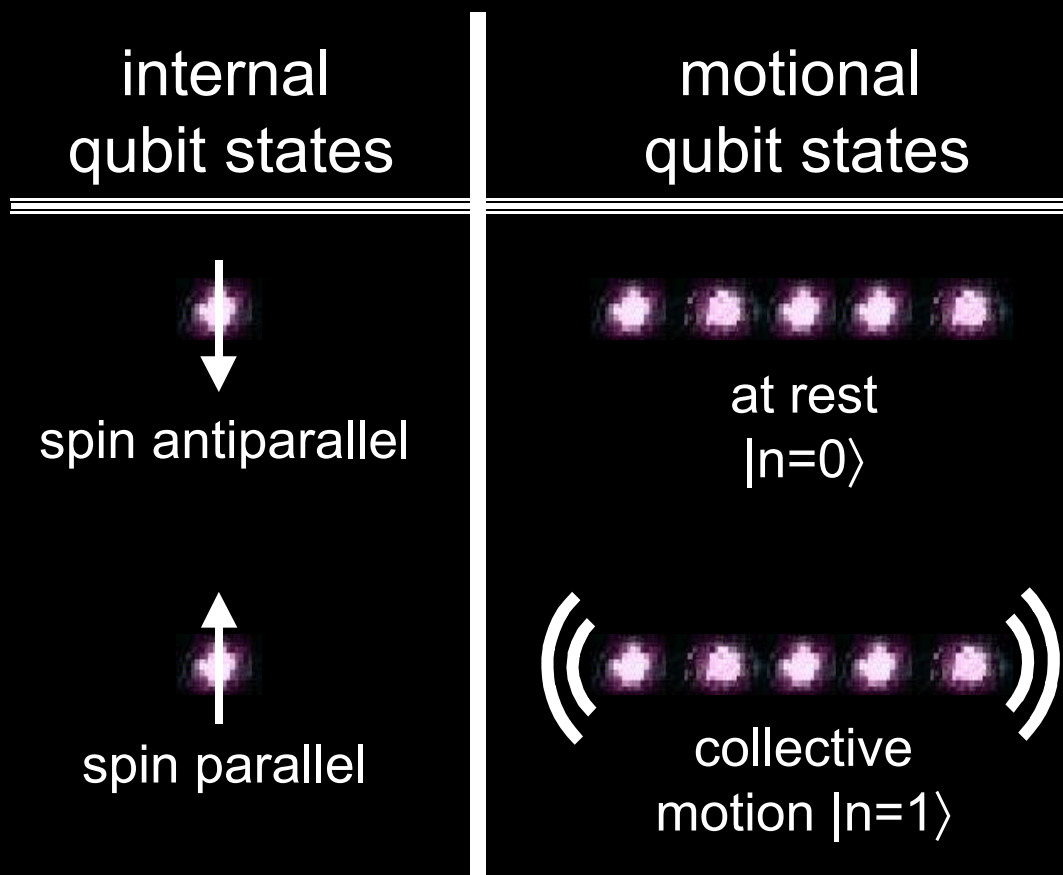



Figure 2




Figure 3

Quantum Computing Abyss

State-of-the-art
experiments

Requirements
in theory



quantum bits



operations




Figure 4