A Potentially Realizable Quantum Computer

Seth Lloyd

Complex Systems Group T-13
and *Center for Nonlinear Studies*
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Abstract: Arrays of weakly–coupled quantum systems can be made to compute by subjecting them to a sequence of electromagnetic pulses of well–defined frequency and length. Such pulsed arrays are true quantum computers: bits can be placed in superpositions of 0 and 1, logical operations take place coherently, and dissipation is required only for error correction. When operated at optical frequencies, the array could provide massively parallel computation at clock cycles of nano– to pico–seconds.

Introduction

As nanofabrication techniques allow the construction of devices of smaller and smaller size, problems of dissipation and error correction become more severe and quantum–mechanical effects inevitably come into play. Technological progress is beginning to make practical a question that was previously academic: What are the fundamental physical limits on computation? Landauer[1] pointed out that the only logical operations that necessarily require dissipation are irreversible ones, such as error correction, in which the logical input cannot be deduced from the output. This result led both to proposals for reversible, dissipationless logic devices,[2] and to Bennett's discovery[3] that computation could be carried out using reversible logic alone.[4] Subsequently, Benioff,[5] Deutsch[6] and Feynman[7] proposed computers in which bits, the fundamental 'quanta' of information, are registered by true quantum–mechanical quanta such as spins.[8–10]

These proposals for quantum–mechanical computers rely on 'designer Hamiltonians,' that are specially constructed to allow computation, and that do not necessarily correspond to any physical system. This report, in contrast, proposes a class of quantum computers, each step in the construction and operation of which lies within or close to the grasp of current technology. The proposed computers are composed of arrays of weakly coupled quantum systems that are subjected to a sequence of electromagnetic pulses of well-defined frequency and length. The systems that make up the arrays could be quantum dots in a semiconductor, for example, or nuclear spins in a crystal lattice, or localized electronic states in a polymer. If the technical steps required to control such a system, can be put together to make a working quantum computer, then the resulting machine, operating at optical frequencies, could provide massively parallel processing at clock rates of picoseconds to nanoseconds.

The idea of exploiting quantum effects to build molecular level computers is not new.[11–13] The proposal detailed here relies on inducing logic through the selective driving of resonances, a method used by Haddon and Stillinger[11] for inducing logic in molecules, and by Obermayer, Teich, and Mahler[14] for inducing a parallel logic in arrays of quantum dots. In this report it is shown that such a method extends to arrays of coupled quantum systems with generic weak interactions between them. Explicit techniques are detailed for performing computation and correcting errors. When operated in a quantum–mechanically coherent fashion, such computers are examples of logically reversible computers that dissipate less than $k_B T$ per logical operation; dissipation is only required for error correction. In fact, the systems described are true quantum computers in the sense of Deutsch:[6] bits can be placed in superpositions of 0 and 1, quantum uncertainty can be used to generate random numbers, and states can be created that exhibit purely quantum–mechanical correlations.[5−10]

How it works

For the purposes of exposition, consider a heteropolymer, $ABCABCABC\ldots$, in which each unit possesses an electron that has a long–lived excited state. For each unit, A, B or C, call the ground state 0, and the excited state 1. Since the excited states are long–lived, the transition frequencies ω_A, ω_B and ω_C between the ground and excited states are well-defined. In the absence of any interaction between the units, it is possible to drive transitions between the ground state of a given unit, B say, and the excited state by shining light at the resonant frequency ω_B on the polymer.[15−16] Let the light be in the form of a π pulse, so that $\hbar^{-1}\int \vec{\mu}_B\cdot\hat{e}\mathcal{E}(t)dt = \pi$, where μ_B is the induced dipole moment between the ground state and the excited state, $\hat{e}$ is the polarization vector for the light that drives the transition, and $\mathcal{E}(t)$ is the magnitude of the pulse envelope at time t. If the π pulse is long compared with $1/\omega_B$, so that its frequency is well–defined, and if the polymer is oriented, so that each induced dipole moment along the polymer has the same angle with respect to $\hat{e}$, then its effect is to take each B that is in the ground state and put it in the excited state, and to take each B in the excited state and put it in the ground state.

Now suppose that there are local interactions between the units of the polymer, given by interaction Hamiltonians H_{AB}, H_{BC}, H_{CA}. The interactions could arise, for example, from overlap between electron wave functions from unit to unit, from Van der Waals forces, from changes in the local structure of the polymer, or from induced or permanent multipole coupling. Almost any local interaction will do (additional measures must be taken to correct for the long range nature of permanent dipole coupling). Consider first the case in which these interaction Hamiltonians are diagonal in the original energy eigenstates for each unit (the effect of off–diagonal terms is considered below). The only effect of such interactions is to shift the energy levels of each unit as a function of the energy levels of its neighbors, so that the resonant frequency ω_B, for instance, takes on a value ω^B_{01} if the A on its left is in its ground state and the C on its right is in its first excited state. If the resonant frequencies for all transitions are different for different values of a unit's neighbors, then the transitions can be driven selectively: if a π pulse with frequency ω^B_{01} is applied to the polymer, then all the the B's with an $A = 0$ on the left and a $C = 1$ on the right will switch from 0 to 1 and from 1 to 0 are. If all transition frequencies are different, these are the only units that will switch. Each unit that undergoes a transition coherently emits or absorbs a photon of the given frequency: no dissipation takes place in the switching process.

Driving transitions selectively by the use of resonant π pulses induces a parallel logic on the states of the polymer: a particular resonant pulse updates the states of all units of a given type as a function of its previous state and the states of its neighbours. All units of the given type with the same values for their neighbours are updated in the same way. That is, applying a resonant pulse to the polymer effects the action of a cellular automaton rule on the states of units of the polymer.[14,17] In fact, it is easy to see that one can by the proper sequence of pulses realize *any* reversible cellular automaton rule that acts first on one type, then on another, then on another, etc. This result holds even when some units have more than one excited state.

The system is much more computationally powerful than a simple cellular automaton, since one can change the cellular automaton rule ¿from step to step by changing the sequence of pulses applied.[17] By selecting the sequence of pulses, one can make even the simplest of such systems perform any computation that one desires: pulse driven quantum computers are universal digital computers. The following sections present examples and results on how these computers can be programmed. Proofs of these results will be supplied elsewhere.[18]

Loading and unloading information

A simple sequence of pulses allows one to load information onto the polymer. There is one unit on the polymer that can be controlled independently — the unit on the end.[14] Simply by virtue of having only one neighbour, the unit on the end in general has different resonant frequencies from all other units of the same type. Suppose this unit is an A: the resonant frequencies $\omega_i^{A:end}$ for this unit are functions only of the state i of the B on its right. If these resonant frequences are different from the resonant frequencies ω_{ij}^{A} of the A's in the interior of the polymer, then one can switch the end unit from 0 to 1 on its own.

Suppose that all units are initially in their ground state. To load a 1 onto the polymer, apply a π pulse at frequency $\omega_0^{A:end}$. This pulse switches the end unit to 1. To move this 1 along the polymer, apply a π pulse with frequency ω_{10}^{B}. The only B that responds to this pulse is the first: it will switch to 1. Now apply a pulse with frequency $\omega_1^{A:end}$. This pulse switches the A on the end back to 0 (Figure 1). (This act of reversibly restoring a bit to zero using a copy of the bit is called 'uncopying,' and is typical of reversible computation schemes.[1–4]) A properly chosen sequence of pulses moves the 1 around the polymer at will; and a slightly more complicated sequence of pulses can be used to load any sequence desired onto the polymer.[18]

There are several ways to get information off the polymer. All involve a certain amount of redundancy, since detection efficiencies for single photons are not very good. The simplest way is to have many copies of the polymer. The same sequence of pulses will induce the same sequence of bits on each copy. To read a bit, one applies a sequence of pulses that moves it to the end. Then one applies two π pulses, with frequencies $\omega_{0,1}^{A:end}$. If either of these pulses is attenuated, then the bit on the end is a 1; if either is amplified, then the bit is a 0. This method has the disadvantage that all bits must be moved to the end of the polymer to be read.

If the light in the π pulses can be focussed to within a radius of a few wavelengths, information can also be read out in parallel, simply by copying the bit that is to be read out onto all or most units of the same type within a few wavelength neighbourhood, and then seeing whether π pulses aimed at that neighbourhood are attenuated or amplified. Other schemes that require less redundancy exist. For example, if the end unit has a fast decay mode (as described below in the section on dissipation), the signal for a bit being a 0 or 1 can be a photon of a different frequency than that of the switching pulse. Only a small numbers of such photons in a distinct frequency channel need be present to be detected with high accuracy.

Computation

Once information is loaded onto the polymer, a wide variety of schemes can be used to process it in a useful fashion. It is not difficult to find sequences of pulses that realize members of the following class of parallel processing computers.[18]

The polymer is divided up into sections of equal length. By choosing the proper sequence of pulses, and by properly formatting the input information, one can simulate the action of any desired reversible logic circuit on the information within each section. (Since every logical action described up until now is reversible, the entire circuit must be reversible: the logical operation induced by a sequence of π pulses can be reversed simply by applying the same sequence in reverse order.) The logic circuit realized is, of course, the same for each section, although the initial information on which the circuit acts can be different from section to section. A second sequence of pulses allows each section to exchange an arbitrary number of bits with the sections to its left and right. Input and output can be obtained from the sections on the end, as above, or ¿from each individual section using focussed light.

By choosing the proper section size and sequence of pulses, one can then realize a string of identical microprocessors of arbitrary reversible circuitry, each communicating with its neighbours (Figure 2). Such a device is obviously computationally universal, in the sense that one can embed in it the operation of a reversible universal Turing machine. A device with the parallel architecture described here, however, is likely to be considerably more useful than a Turing machine for performing actual computations. The number of pulses required to realize such a machine is proportional to the length of the wires, measured in terms of the number of units over which bits must be transported, and number of logic gates in one microprocessor.

Quantum computation

The resulting computer is not only a universal digital computer, but a true quantum computer. Bits can be placed in superpositions of 0 and 1 by the simple expedient of applying pulses at the proper resonant frequencies, but of length different from that required to fully switch the bit. For example, if in loading information on the polymer, as in the section above, instead of applying a π pulse, one applies a $\pi/2$ pulse of frequency $\omega_0^{A:end}$ of length T_1, the effect is to put the A unit on the end in the state, $1/\sqrt{2}\big(|0\rangle + e^{-i\phi_1}|1\rangle\big)$, where $\phi_1 = \pi/2 + \omega_0^{A:end}T_1$. Applied at a time T_2 later, a π pulse of frequency ω_{10}^B and length T_3 then puts the first two units in the state, $1/\sqrt{2}\big(|00\rangle + e^{-i\phi_2}|11\rangle\big)$, where $\phi_2 = 3\pi/2 + \omega_0^{A:end}(T_1 + T_2) + (\omega_1^{A:end} + \omega_{10}^B)T_3$. (Figure 3.)

In fact, by the proper sequence of pulses, it is possible not only to create any quantum state of N bits, but to effect any unitary transformation desired on those N bits.[18] The proposed device is not only a universal digital computer, but a universal quantum analog computer in the sense of Deutsch.[6] Of course, the number of steps over which coherent superpositions of computational states can be maintained is limited by interactions with the environment, which induce phase randomization and decoherence. Even if Deutsch's full program of quantum parallelism in computation cannot be realized for lengthy computations, however, the methods described can still be used to create and manipulate states that exhibit purely quantum–mechanical features, such as Einstein–Podolsky–Rosen correlations that violate Bell's inequalities.[19–21] In addition, by giving each bit a quantum 'twist' when loading it on the computer (for example, by applying a $\pi/2$ pulse or a $3\pi/2$ pulse at random), information could be encoded and stored in such a way that only the person who knows by how much each bit has been rotated could read the information. All others who try to read it will get no information, and will leave a signature of their attempt to read it in the process, by randomizing the states of the bits.[22]

Dissipation and error correction

Errors in switching and storing bits are inevitable. It is clear that without a method for error correction, the computer described here will not function.

Error correction is a logically irreversible process, and requires dissipation if errors are not to accumulate.[1] If in addition to a long–lived excited state, any of the units possesses an excited state that decays quickly to a long–lived state, this fast decay can be exploited to provide error correction. For example, each B could have an additional excited state, 2, that decays to the ground state, 0, in an amount of time short compared with the time in between pulses. Any B in a long–lived state, 1, e.g., can be restored to the ground state conditioned on the state of its neighbours by applying pulses with the resonant frequency $\omega^B_{ij}(12)$ of the transition between the states 1 and 2 given that its neighbours A and C are in the states i and j (Figure 4). The pulse need not have a well–defined length, provided that it is long enough to drive the transition efficiently.

If just one type of unit has a fast decay of the sort described, then one can realize not only any reversible cellular automaton rule that updates first one type of unit, then another, but any irreversible cellular automaton rule as well. The scheme described above that allows the construction of one–dimensional arrays of arbitrary parallel–processing reversible microprocessors then allows one to produce one–dimensional arrays of arbitrary irreversible microprocessors, each one of which can contain arbitrary error–correcting circuitry. Many error–correcting schemes are possible, using check sums and parity bits[23–24] multiplexing,[25] etc. A particularly simple and robust scheme is given in reference 18. For each logically irreversible operation accomplished, a photon is emitted incoherently to the environment. In contrast to the switching of bits using π pulses, in which photons are emitted and absorbed coherently, the switching of bits using fast decays is inherently dissipative. The amount of dissipation depends on what is done with the incoherently emitted photons. If the photon is absorbed and its energy thermalized, then considerably more than k_BT is dissipated; if the energy of the photon is put to work, dissipation can be brought down to close to k_BT.

Such a computer can function reliably in the face of a small error rate in principle. Error correction for the method of computation proposed here takes the place of gain and signal restoration in conventional circuits. Whether such a computer can actually be made to function reliably in the face of a finite error rate depends crucially on whether the error correction routine suffices to correct the number of errors generated in the course of the computational cycle, in between error correction cycles.

Suppose that the probability of error per unit per computational cycle is ϵ. Suppose that all bits come in $2k+1$ redundant copies, and that after each cycle, error correction is performed in parallel by having the copies vote amongst eachother as to their proper value, and all copies are restored to that value: there exist quick routines for performing this operation, that are insensitive to errors generated during their execution.[18] The error rate per cycle is reduced to $\eta \approx (4\epsilon)^k$ by this process. For a computation that uses b bits over c cycles to have a probability no greater than f for the failure of a single bit, we must have $b(1-\eta)^{bc} \geq b-f$, which implies that $\eta \leq 1/cb^2$. For example, suppose that the error rate per bit per cycle is a quarter percent, $\epsilon = .0025$. To have a computation involving 10^{12} bits over 10^{20} steps have a probability of less than 1% of getting a bit wrong requires that each bit have 47 redundant copies. Although such computers have much higher error rates and require much more error correction than conventional computers, because of their high bit density and massively parallel operation, error correction can be carried out without too great a sacrifice in space or time. Error correction is also the secret to making the computer perform when the intrinsic lifetime of the excited states is short.

Note that when a photon is emitted incoherently, the quantum coherence of the bit from which it was emitted, and of any other bits correlated with that bit, is destroyed. Incoherent processes and the generation of errors intrinsically limit the number of steps over which the computer can function in a purely quantum–mechanical fashion.

Errors

There are many potential sources of error in the operation of these pulsed quantum computers. The primary difficulties in the proposed scheme are the identification of a suitable material for the computer, and the delivery of effective π pulses. Microwave technology can give complete inversion with error rates of a fraction of a percent in NMR systems. Optical systems with long–lived localized states and small inhomogeneous broadening are difficult to find, but if they can be identified, the narrow bandwidths of such systems should allow for efficient population inversion. Different realizations of the computer suffer from different technical problems, such as the difficulty of insuring that all units of the same type in a polymer possess a common orientation with respect to the polarization of the applied pulse, or the difficulty of manufacturing quantum dots with both a sufficiently long lifetime and a sufficient degree of uniformity across an array. As noted above, a fraction of a percent error per bit per computational cycle can be tolerated; but a few percent is probably too much. Techniques such as pulse shaping and iterative excitation schemes enhance π pulse effectiveness and selectivity.[26−27]

In addition to the technological problem of supplying accurate π pulses, the following fundamental physical effects can cause substantial errors:

Effect of off–diagonal terms in interaction Hamiltonians. These terms have a number of effects. The simplest is to induce unwanted switching of individual units, with a probability of error per unit per pulse of $\left(\delta\omega_{off}/\omega\right)^2$ whenever a unit or its neighbour is switched. Here $\hbar\delta\omega_{off}$ is the characteristic size of the relevant off–diagonal term in the interaction Hamiltonian. Off–diagonal interactions also induce the propagation of excitons along the polymer: this process implies that a localized excited state has an intrinsic finite lifetime equal to the inverse of the bandwidth for the propagation of the exciton associated with that state.[28−29] For the polymer

$ABCABC\ldots$, the bandwidth associated with the propagation of an excited state of A can be calculated either by a decomposition in terms of Bloch states, or by perturbation theory, and is proportional to $\delta\omega_{off}^{AB}\delta\omega_{off}^{BC}\delta\omega_{off}^{CA}/(\omega_A-\omega_B)(\omega_A-\omega_C)$, where $\hbar\delta\omega_{off}^{AB}$, e.g., is the size of the term in H_{AB} that induces propagation of excitation from A to B. For a polymer of the form $12\ldots M12\ldots M\ldots$, the characteristic bandwidth goes as $\delta\omega_{off}^{M}/\Delta\omega^{M-1}$, where $\Delta\Omega$ is the typical size of the difference between the resonant frequencies of different types of units. For the computer to function, the inverse of the exciton propagation bandwidth must be much longer than the characteristic switching time. If the off–diagonal terms are of the same size as the on–diagonal terms, on average, then for the computer to function, the overall interaction between units must be weak, and M, the number of different kinds of units in the polymer, must be at least three. Small off–diagonal terms and a relatively large number of different types of units are essential for the successful operation of the computer.

Quantum–electrodynamic effects. The probability of spontaneous emission from a single unit is assumed to be small. In the absence of interactions, the spontaneous decay rate for a unit with resonant frequency ω is $4\omega^3\mu^2/3\hbar c^3$.[15] If the lifetime of an optical excited state is to be as long as microseconds, the induced dipole moment μ must be suppressed by symmetry considerations. Interactions between different units of the same type can give rise to quantum–electrodynamic effects such as super–radiance, and the coherent emission of a photon by one unit and coherent reabsorption by another.[26,30] Fortunately, the states that are being used for computation, in which each unit is in a well–defined excited or ground state, do not give enhanced probabilities for these processes. In the process of switching, however, and when bits are in superpositions of $|0\rangle$ and $|1\rangle$, super–radiant emission gives an enhancement of the spontaneous emission rate by a factor of n, where n is the number of units of the same type within a wavelength of the light used. Since the switching time is short compared to the lifetime, super–radiant emission is not a problem. (Though super–radiance can shorten the lifetime of quantum superpositions of logical states.)

Long range interactions. The dipole–dipole couplings of reference 14 fall off as $1/r^3$. For such a long range coupling, coherent switching will not work unless the shift in a unit's resonant frequency induced by nearby units of the same type is too small to throw the unit out of resonance. If the interactions are not nearest–neighbour, but still have a finite effective range, coherent switching can be maintained either by having a sufficiently large number of different types of units that units of the same sort fall out of the range of the interaction, or by encoding information and performing computation in such a way that no two units of the same type within the range of the interaction are ever switched at the same time.[18] In either case, the result of a resonant pulse is to realize a cellular automaton rule with a neighbourhood of radius larger than one.

None of the purely physical effects gives error rates that are insurmountable. If the π pulses are long compared to the inverse frequency shifts due to interaction, if the unperturbed resonant frequencies differ substantially between the different types of unit, and if the off–diagonal terms in the interaction Hamiltonians are small compared with the resonant frequencies and their differences, then this computing scheme will work in principle.

Although putting them together in a working package may prove difficult, precisely timed and shaped monochromatic laser pulses, well–oriented polymers, accurately fabricated semi–conductor arrays, and fast, sensitive photodetectors are all available in today's technology. Continuously tunable ti-sapphire lasers, or diode–pumped YAG lasers tuned by side–band modulation can currently supply frequency–stable picosecond pulses at nanosecond intervals with an integrated intensity that varies by a fraction of a percent. Currently available electro–optical shutters could be used to generate the proper pulse sequence at a nanosecond clock rate. Photodetectors equipped with photomultipliers and acoustic–optical filters can reliably detect tens of photons (or fewer) within a wavelength band a few nanometers wide.

A more difficult problem is finding a proper material for the computer. For operation in the optical regime, polymers or molecular crystals must admit localized, locally–interacting excited states with lifetimes not much shorter than microseconds. Although arrays of quantum dots created by X-ray lithography are not yet of sufficiently uniform quality, arrays of quantum dots and lines that have been created using interference techniques might be sufficiently uniform to realize the proposed scheme.

Numbers

The range of speed of operation of such a pulsed quantum computer within acceptable error rates is determined by the frequency of light used to drive transitions, and by the strength and character of the interactions between units. For square-wave pulses, the intrinsic probability of error per unit per pulse due to indiscriminate transition–driving is $\left(1/T\delta\omega_{on}\right)^2$, where T is the pulse length and $\delta\omega_{on}$ is resonant frequency shift induced by on-diagonal terms in the interaction Hamiltonian,[15] (this error can be reduced significantly by using shaped pulses[26]) while the probability of error per unit per pulse due to off–diagonal terms in the interaction Hamiltonians is $\left(\delta\omega_{off}/\omega\right)^2$. The decay of localized excitations due to exciton propagation gives a lifetime proportional to $\Delta\omega^{M-1}/\delta\omega_{off}^M$, where M is the number of different types of units.

Suppose that the excited states have transition frequencies corresponding to light in the visible range, say $\omega = 10^{15}$ sec^{-1}. In the absence of off–diagonal terms in the interaction Hamiltonians, the frequency shifts due to interaction do not need to be small compared to ω, and to obtain an intrinsic error rate of less than 10^{-6} per unit per pulse, using square wave pulses, the pulse length could be as short as 10^{-12} seconds, and as long as a few thousands of the intrinsic lifetimes of the excited states (assuming that a few thousand pulses are required for error correction). With shaped, hyperbolic secant pulses, the pulse length could be considerably shorter. The clock rate of such a computer could be varied to synchronize its input and output with conventional electronic devices. In the presence of off–diagonal terms of the same magnitude $\delta\omega_{off} \approx \delta\omega_{on} \sim \delta\omega$ as the on–diagonal terms, to obtain an intrinsic error rate of 10^{-6} per unit per pulse, one must have $\delta\omega = 10^{12}$ sec^{-1}, and a minimum pulse length of 10^{-9} seconds. If the computer has three different types of units, a typical exciton lifetime is on the order of 10^{-6} seconds, although spontaneous emission and other decay modes may shorten this.

If the units in the quantum computer are nuclear spins in an intense magnetic field, with dipole–dipole interactions, then the the pulses will have frequencies in the microwave or radiofrequency region, and the computers will have clock rates from microseconds to milliseconds.

Conclusion

Computers composed of arrays of pulsed, weakly–coupled quantum systems are physically feasible, and may be realizable with current technology. The units in the array could be quantum dots, nuclear spins, localized electronic states in a polymer, or any multistate quantum system that interacts locally with its neighbours, and that can be compelled to switch between states using resonant pulses of light. The exposition here has concentrated on one-dimensional systems with two or three states, but more dimensions and more states per unit provide higher densities of information storage and a wider range of possibilities for information processing, as long as the different transitions can still be driven discriminately.

The small size, high clock speeds and massively parallel operation of these pulsed quantum computers, if realized, would make them good devices for simulating large, homogeneous systems such as lattice gases or fluid flows. But such systems are capable of more than digital computation. When operated coherently, the devices described here are true quantum computers, combining digital and quantum analog capacities, and could be used to create and manipulate complicated many–bit quantum states. Many questions remain: What are the best physical realizations of such systems? (The answer may be different according to whether the devices are to be used for fast, parallel computing, or for generating novel quantum states.) How can they best be programmed? How can noise be suppressed and errors corrected? How can their peculiarly quantum features be exploited? What are the properties of higher dimensional arrays? The device proposed here, as with all devices in the next generation of nanoscale information processing, cannot be built and made to function without addressing fundamental questions in the physics of computation.

References

1. R. Landauer, *IBM J. Res. Develop.* **5**, 183–191 (1961).
2. K.K. Likharev, *Int. J. Theor. Phys.* **21**, 311-326 (1982).
3. C.H. Bennett, *IBM J. Res. Develop.* **17**, 525–532 (1973); *Int. J. Theor. Phys.* **21**, 905–940 (1982).
4. E. Fredkin, T. Toffoli, *Int. J. Theor. Phys.* **21**, 219-253 (1982).
5. P. Benioff, *J. Stat. Phys.* **22**, 563–591 (1980); *Phys. Rev. Lett.* **48**, 1581–1585 (1982); *J. Stat. Phys.* **29**, 515–546 (1982); *Ann. N.Y. Acad. Sci.* **480**, 475–486 (1986).
6. D. Deutsch, *Proc. Roy. Soc. Lond.* **A 400**, 97–117 (1985); *Proc. Roy. Soc. Lond.* **A 425**, 73–90 (1989).
7. R.P. Feynman, *Optics News* **11**, 11-20 (1985); *Found. Phys.* **16**, 507–531 (1986); *Int. J. Theor. Phys.* **21**, 467–488 (1982).
8. W.H. Zurek, *Phys. Rev. Lett.* **53**, 391–394 (1984).
9. A. Peres, *Phys. Rev.* **A 32**, 3266–3276 (1985).
10. N. Margolus, *Ann. N.Y. Acad. Sci.* **480**, 487–497 (1986); *Complexity, Entropy, and the Physics of Information, Santa Fe Institute Studies in the Sciences of Complexity* **VIII**, W.H. Zurek, ed., 273-288 (1991).
11. F.L. Carter, *Molecular Electronics* and *Molecular Electronics II*, Dekker, New York, 1982 and 1987.
12. D.K. Ferry, J.R. Barker, C. Jacobini, *Granular Nanoelectronics*, Plenum, New York, 1991. K.E. Drexler, *Nanosystems*, Wiley, New York, 1992. See also the now–defunct *Journal of Molecular Electronics*, and the proceedings of the *International Symposia of Molecular Electronics.*
13. J.J. Hopfield, J.V. Onuchic, D.N. Beratan, *Science* **241**, 817–820 (1988).
14. K. Obermayer, W.G. Teich, G. Mahler, *Phys. Rev.* **B 37**, 8096–8110 (1988). W.G. Teich, K. Obermayer, G. Mahler, *Phys. Rev.* **B 37**, 8111–8121 (1988). W.G. Teich, G. Mahler, *Phys. Rev.* **A 45** 3300, 1992.
15. L. Allen, J.H. Eberly, *Optical Resonance and Two-Level Atoms*, Wiley, New York 1975.
16. W.H. Louisell, *Quantum Statistical Properties of Radiation*, Wiley, New York 1973.
17. S. Wolfram, *Rev. Mod. Phys.* **55**, 601 (1983).
18. S. Lloyd, *Programming Pulsed Quantum Computers*, to be published.
19. A. Einstein, B. Podolsky, N. Rosen, *Phys. Rev.* **47**, 777 (1935).
20. G. Mahler, J.P. Paz, to be published.

21. D.M. Greenberger, M. Horne, A. Zeilinger, in *Bell's Theorem, Quantum Theory, and Conceptions of the Universe*, M. Kafatos, ed., Kluwer, Dordrecht, 1989. N.D. Mermin, *Phys. Today* **43**(6), 9 (1990).
22. C.H. Bennett, G. Brassard, S. Breidbart, S. Wiesner, *Advances in Cryptology: Proceedings of Crypto '82*, Plenum, New York, 267-275 (1982). S. Wiesner, *Sigact News* **15**(1), 78-88 (1983).
23. C.E. Shannon, W. Weaver, *The Mathematical Theory of Information*, University of Illinois Press, Urbana 1949.
24. R.W. Hamming, *Coding and Information Theory*, Prentice–Hall, Englewood Cliffs, 1986.
25. J. von Neumann, *Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components*, Lectures delivered at the California Institute of Technology, 1952.
26. W.S. Warren, *Science* **282**, 878-884 (1988). W.S. Warren, A.H. Zewail, *J. Chem. Phys.* **78**, 2279-2297, 2298-2311 (1983).
27. R. Tycko, A. Pines, J. Guckenheimer, *J. Chem. Phys.* **83**(6), 2775-2802 (1985).
28. R.S. Knox, *Theory of Excitons*, Academic Press, New York 1963.
29. P.W. Anderson, *Concepts in Solids*, Addison–Wesley, Redwood City 1963.
30. R.H. Dicke, *Phys. Rev.* **93**, 99 (1954).

Figure Captions:

Figure 1: A 1 is loaded onto the end of the polymer and moved over by a single unit. Solid lines represent actual energy levels, dotted lines represent unshifted energy levels, and lines with a dot represent occupied states. In a), a π pule of frequency $\omega_0^{A:end}$ has put the A on the end into an excited state. In b), a π pulse of frequency ω_{10}^{B} has put the B next to the end unit in an excited state, leaving all other B units in their ground states. In c), a π pulse of frequency $\omega_1^{A:end}$ has restored the end A reversibly to the ground state.

Figure 2: The architecture of the universal parallel computer. Input and output take place at the processor on the end. Each parallel processor has the same circuitry, realizing any desired logical function. After each processor cycle, an arbitrary number of bits are exchanged with the processors on the left and on the right. The two dimensional version of the computer has an analogous form.

Figure 3: The result of applying a $\pi/2$ pulse to the end unit, then copying the bit on the end unit to the next two bits. The resulting state has the three units on the end in the state, $1/\sqrt{2}(|000\rangle + e^{i\phi}|111\rangle$, an extended Einstein–Podolsky–Rosen–Bohm state.

Figure 4: Using dissipation to realize irreversible logical operations. Each B unit has an excited state that decays rapidly to the ground state. In a), a π pulse with frequency $\omega_{11}^{B}(12)$ takes each B unit that is in the first excited state, 1, and that has neighbours $A = 1, C = 1$, and puts it in the short–lived excited state 2. In b), the B unit in state 2 emits a photon and returns to the ground state. In this irreversible logical process, every B whose neighbours are 11 is restored to the ground state, regardless of its initial state. This type of logically irreversible process is crucial for error correction.