Information is Physical

Rolf Landauer
IBM Thomas J. Watson Research Center
PO Box 218, Yorktown Heights, NY 10598

Abstract

This paper is related to, but not equivalent to, the keynote talk given at the workshop on Physics of Computation. This published version concentrates more on matters of detail, and is not intended as a survey.

About fifty years ago the attempt to move toward modern computers, on a laboratory basis, commenced. About forty years ago the first "large scale" electronic computers, made commercially, were distributed. Another decade after that we learned that throwing away information incurs an inevitable energy cost. About two decades ago it became clear that computers, in principle, could use 1:1 mappings; information did not have to be discarded. And ten years ago the proceedings of the first session, resembling this one, appeared in print. As someone who was a full time participant during four of those five decades, I am both gratified and disturbed. Immensely pleased at how much has happened in computer technology, and pleased at how this field has developed. Twenty years ago a meeting of this sort would have been difficult to foresee. It would have been difficult to image the rich and diverse paths of development, involving many people, and many viewpoints. I am disturbed because some truths have made very slow headway. Science is full of fashion. Some subjects achieve instant recognition; other notions creep forward slowly and stumblingly. Much of what is contained in this volume has come along a slow and difficult path.

My talk at the actual workshop in Dallas started with an attempt to list the variety of questions which are encountered when viewing the interface of physics and information handling at a fundamental level. I then went on to emphasize the part that I have tried to understand. Information is physical and not an abstract entity. Information is inevitably tied to physical degrees of freedom through a charge, a spin, a hole in a punched card, or similar devices. That ties information to the laws of physics and ties it to the parts available in the real universe. What can we learn from that? In this written version I will not, as in the talk, attempt to sketch what we have learned in this field in a systematic way. Instead, I will concentrate on a few special cases where simple concepts deserve broader exposure.

Reversible computation, invented by Charles Bennett, is illustrated in Fig. 1a. The left-hand end of a horizontal chain is the initial state. Motion to the right yields forward steps through a sequence of states represented by successive circles. Different letters correspond to different initial states. For each state (except, perhaps, the initial state) there is a unique predecessor state. That is in contrast to Fig. 1b, where two distinguishable paths merge into one; information is lost. As Bennett showed, the event illustrated in Fig. 1b can be avoided. Motion in the space of Fig. 1 takes place in a noisy and dissipative system. Forward motion in Fig. 1a can best be understood by an analogy with electronic motion in a lattice. A small forward force will assure a predictable drift velocity, even though locally the motion may look very diffusive. At a junction, as shown in Fig. 1b, there is more phase space to the left of the junction than to its right. If the system were left in equilibrium, it would be found preferentially on the left. To overcome this we must apply a larger bias force near the junction, and dissipate $kTln2$. The above argument is only one of a number leading to the same conclusion, which is still challenged, on occasion. At least two of the preliminary position papers submitted to this

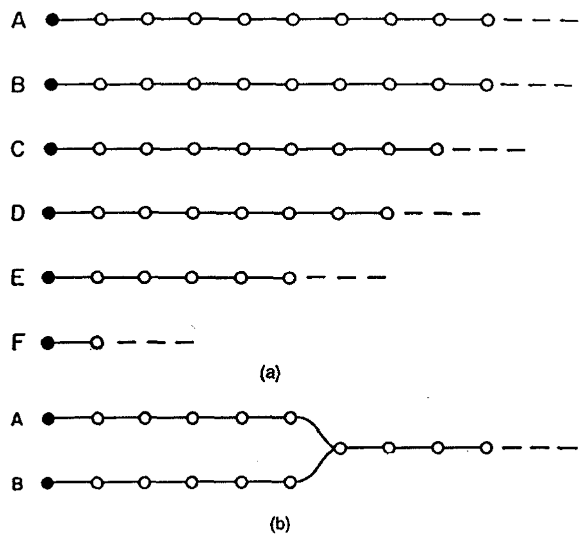

Figure 1:

0-8186-3420-0/93 $03.00

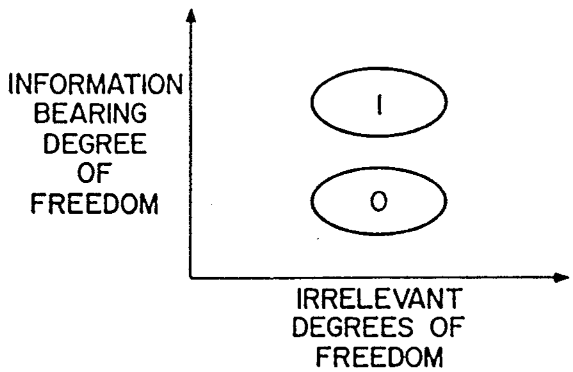

Figure 2:

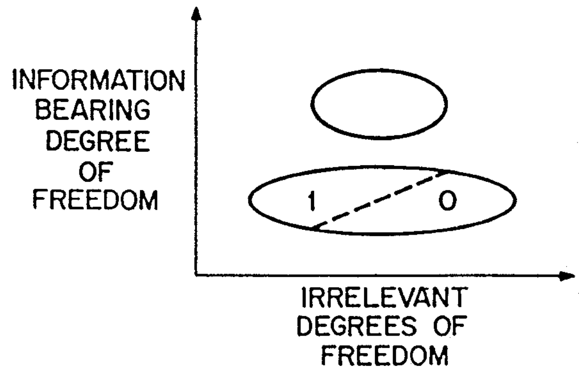

Figure 3:

workshop stated such a challenge. (I have no way of knowing whether they will appear unchanged in this collection.) Let me, therefore, repeat another argument, presented in its most detailed version in Ref. [1]. Consider a typical logical process which discards information, e.g., a logic variable which is reset to *0*, regardless of its initial state. Fig. 2 shows, symbolically, the phase space of the complete computer considered as a closed system, with its own power source. The erasure process we are considering must map the *1* space down into the *0* space. Now, in a closed conservative system phase space cannot be compressed, hence the reduction in the vertical spread must be compensated by a lateral phase space expansion, i.e., a heating of the horizontal irrelevant degrees of freedom, typically thermal lattice vibrations. Indeed, we are involved here in a process which is similar to adiabatic magnetization (i.e., the inverse of adiabatic demagnetization), and we can expect the same entropy increase to be passed to the thermal background as in adiabatic magnetization, i.e., $kln2$ per erasure process. At this point, it becomes worthwhile to be a little more detailed. Fig. 3 shows the end-result of the erasure process in which the original *1* and *0* spaces have both been mapped into the vertical range originally occupied by the *0*. This is, however, rather like the isothermal compression of a gas in a cylinder into half its original volume. The entropy of the gas has been reduced and the surroundings have been heated, but the process is not irreversible, the gas can subsequently be expanded again. Similarly as long as *1* and *0* occupy distinct phase space regions, as shown in Fig. 3, the mapping is reversible. *The real irreversibility comes from the fact that the 1 and 0 spaces will subsequently be treated alike and will eventually diffuse into each other.*

I will not, in this brief note, allude to all the supplemental cautions. For example, we can reset a bit with less energy expenditure if *1* and *0* are not equally probable. Unfortunately, when the publications in a field span several decades, there is a tendency by late arrivals to assume that a discussion like this repeats all that is known and understood.

Reversible computation was, initially, an invention intended to answer conceptual questions, and not as a serious technological proposal. Eventually, Likharev showed that Josephson junction circuits could be used to accomplish reversible computation. At this symposium three papers suggested how field effect transistors can be used to approximate reversible computation. Not, necessarily, to avoid an energy loss of order kT, but to substantially reduce the existing dissipation per switching event. Only time will tell whether the penalties incurred by these proposals are worthwhile. Nevertheless, the proposals are most welcome because they serve to underline the role of reversible computation in dissipative and noisy systems. The widespread concern with idealized conservative Hamiltonian systems, classical or quantum mechanical, has tended to displace the concern with *real* dissipative systems. The semiconductor proposals also serve as a reminder that reversible computers can be clocked, and need not move diffusively [2].

Quite aside from questions about the realizability of dissipationless systems, invoked by many authors in this field, there are questions about their desirability. These questions arise from the fact that actual systems cannot be expected to have the exact Hamiltonian needed for the desired evolution, in time, of the information bearing degrees of freedom. Real systems have noise and friction, and clearly have manufacturing defects, i.e., have Hamiltonians which deviate from the ideal. In our further discussion we emphasize this deviation.

Consider the quantum mechanical Hamiltonian $H_o + H_e$, where H_o is the desired Hamiltonian, and H_e the error or deviation. With time, H_e causes more and more of the undesired states to be brought into the computational state. Eventually, the computation will go totally off track. To restore a computation toward its desired evolution, i.e., apply error correction, we have to throw away information about the source of the error. And, as we know, throwing away information requires energy dissipation; it cannot be done within a Hamiltonian framework. In the classical case we also have to face manufacturing errors. To the extent that they only affect the motion *along* the tracks shown in Fig. 1a, they do not induce errors. All the proposals for reversible computation in dissipative classical system involve explicitly, or implicitly, large forces or walls which prevent jumping between the tracks of Fig. 1a. Can we not do that equally well in the quantum case? Unfortunately, the quantum mechanical computer proposals specify a formal Hamiltonian which causes bits to interact and evolve as we would like them to do. As a consequence of the lack of description of physical apparatus, we do not understand what restrictions on the nature of H_e really require, and whether it is plausible to prevent track jumping in Fig. 1a.

The error in the Hamiltonian causes a second, and perhaps more fundamental, problem. Let us return to the electron motion analogy for Fig. 1a. In a periodic lattice, electrons can propagate freely, quantum mechanically. If there are random defects, however, then we face what the condensed matter physicist calls *localization*. Transmission through a section of disordered lattice diminishes exponentially with its length. Similarly, in the case of a computation if, for example, the energy depends on the exact bit pattern, we can expect an exponentially diminishing probability of completing a long computation. Unfortunately, the theoreticians active in the physics of computation tend to be unfamiliar with localization, and have simply ignored this point, despite my own repeated allusions to it. The existence of localization does not guarantee that there is an insurmountable problem. Let me point to one possible direction for further thought. In the quantized Hall effect electrons can move along the edge of a sample in the presence of defects, and without reflection back towards their source. Perhaps, there is a way to apply this to computational trajectories. But it is up to the advocates of quantum mechanical computation to face up to localization, and to find a way to circumvent it.

The growing understanding of reversible computation eventually led to a reexamination of two older related questions, which were supposedly settled. The energy cost of classical measurement, as expressed in Maxwell's famous demon paradox, was settled definitively. That piece of progress has received adequate notice in the literature and in the record of this workshop. The other question which needed clarification relates to the cost of sending a bit from one location to another. The progress in recent years in that direction, however, has received very little public visibility.

The concern with quantum mechanical computation leads to a closely related question. Is there a limit to the number of bits that can be stored in a given volume with a given energy content? Papers have appeared in the literature which declare such limits. I do not believe that such limits exist, at least not if based only on quantum mechanics and statistical mechanics. The key point: It is best to store bits in bistable wells. In analogy to the study of energy flow in thermodynamic cycles, or study of the signal capacity of a noisy channel, we do not include the rest energy of the apparatus or its manufacturing energy. That can be amortized over many uses and we only count the energy above the ground state, in each case. Fig. 4a shows a particle in a bistable well and Fig. 4b shows the two lowest wave functions for that potential. Information can be stored, conceivably, by putting the particle in one of these two states. More likely, by putting it in the left hand or right hand well and utilizing the slow tunneling rate, in the presence of a high barrier, to preserve information. In either case, we note that the energy elevation above the ground state is of order ΔE. That energy splitting can be made as small as we wish, by choosing a sufficiently impenetrable barrier. Similarly the bistable wells can be as narrow as we wish; there is no density limit. Of course, if we turn to particle physics and cosmology we may be able to discover constraints on the nature of such wells.

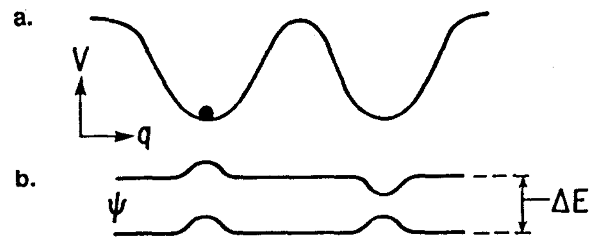

Figure 4:

Note that we have discussed the energy elevation, above the ground state, needed to store information. That is not necessarily an energy dissipation.

This author has had a two-fold role in the development of this field. The material above relates to my interest in the ultimate physical limits of the information handling process. That field has, to a large entent, concerned itself with setting aside supposed limits which arose from casual speculation. My other role: A critic, with some technological background, who pointed to flaws in excessively optimistic proposals, put forth by physical scientists. I have provided such discussion in Ref. [3]. Closely related material has been stated by my colleague R.W. Keyes, Ref. [4].

My colleagues who need contract support for their work do not always welcome such cautious messages. But we cannot hope to build a successful advanced technology effort on excessive promises. Let me extend the critical messages, here, in a direction not specifically addressed in the cited references. There have been widespread casual suggestions that periodic arrays of quantum dots, of resonant tunneling structures, or of other small entities of recent interest to the physics community, can lead to cellular automata. There was an ONR workshop in London in 1990, *Workshop on Applications of Quantum-Coupled Devices to Cellular Automata*. A relatively detailed proposal has been provided in Ref. [5]. Now a cellular automaton (c.a.) requires a well-defined interaction between a set of neighboring elements, and no interactions *at all* between more widely separated elements. It is hard to see how that can be achieved, except by wires, or their close equivalent, such as optical fibers. Furthermore, we need an external clock to cause the c.a. to advance a step, simultaneously over the whole plane. That, however, just conceivably, could come from an oscillating magnetic field or from a pulsed light source. Once we use wired connections between neighboring elements, and for the clock, why not just use transistor circuitry in each cell? That can certainly do the job.

Cellular automata which can perform universal computation are not likely to be of the simple "voting" kind, where the setting of an element depends only on whether the number of its neighbors in the *1* state exceeds some specified minimum. The physical interaction needs to be more complicated than that.

Cellular automata, despite their appeal to physicists in modeling a range of physical phenomena, are not really very suitable general purpose computers as stated in Ref. [6]. Indeed, many investigators who exhibit cellular automata, whose evolution mimics some physical law, do not even tell us how to use these same c.a. rules for the initial program loading! D. Ferry [7] has expressed a similar assessment about c.a.

References

While the actual talk listed some of the key contributions to this field, the citations that follow have a very different purpose. They contain very specific secondary details.

[1] R. Landauer, "Stability and Instability in Information Processing and in Steady State Dissipative Systems", *Physik 1971.* Stuttgart: B.G. Teubner, 1971, p. 286.

[2] R. Landauer and M. Büttiker, "Drift and Diffusion in Reversible Computation", *Physica Scripta T9*, pp.155-164, 1985.

[3] R. Landauer, "Advanced Technology and Truth in Advertising", *Physica A*, vol. 168, pp.75-87, 1990; "Can We Switch by Control of Quantum Mechanical Transmission?", *Phys. Today*, vol. 42, pp.119-121, Oct. 1989; "Nanostructure Physics: Fashion or Depth", M.A. Reed and W.P. Kirk, eds., Nanostructure Physics and Fabrication. New York: Academic Press, 1989, pp.17-30.

[4] R.W. Keyes, "The Future of Solid State Electronics", *Phys. Today*, vol. 45, pp.42-48, Aug. 1992; "Physics of Digital Devices", *Rev. Mod. Phys.*, vol. 61, pp.279-287, 1989; "What Makes a Good Computer Device?", *Science*, vol. 230, pp. 138-144, Oct. 1985.

[5] C.S. Lent and W. Porod, "Bistable Saturation in Coupled Quantum Dots for Cellular Automata", submitted to *Appl. Phys. Lett.*

[6] R. Landauer, "Reversible Computation", O.G. Folberth and C. Hackl eds., Der Informationsbergriff in Technik und Wissenschaft. München: R. Oldenbourg, 1986, p. 139-158.

[7] D.K. Ferry, "Quantum Magnetotransport in Lateral Surface Superlattices", to be published in Prog. in Quant. Elect..