

Experimental Test of Landauer's Principle at the Sub- $k_B T$ Level

Alexei O. Orlov, Craig S. Lent, Cameron C. Thorpe, Graham P. Boechler, and Gregory L. Snider*

Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556, U.S.A.

Received November 21, 2011; accepted January 6, 2012; published online June 20, 2012

Landauer's principle connects the logical reversibility of computational operations to physical reversibility and hence to energy dissipation, with important theoretical and practical consequences. We report the first experimental test of Landauer's principle. For logically reversible operations we measure energy dissipations much less than $k_B T \log 2$, while irreversible operations dissipate much more than $k_B T \log 2$. Measurements of a logically reversible operation on a bit with energy $30 k_B T$ yield an energy dissipation of $0.01 k_B T$. © 2012 The Japan Society of Applied Physics

1. Introduction

Is there a fundamental connection between computation and heat dissipation, between information and thermodynamics? Heat generation already limits the speed and packing density of modern microprocessors, so it is important to understand to what extent the problem is fundamental and therefore unavoidable, and to what extent a new kind of binary device, perhaps not a transistor, could ameliorate the problem.¹⁻³⁾

Computation appears to be an abstract mathematical process, mapping one set of input bits onto another set of output bits. Why would there be any fundamental connection between such a mapping and the type of microscopic motion we call heat? Two distinct notions of reversibility are at issue. *Logical* reversibility of a computation means that the value of the inputs can be inferred from the value of the outputs. *Physical* reversibility means that a movie of the physical computing machinery played backward would not violate the laws of physics, including importantly, the second law of thermodynamics. Some have argued that there is no connection between these different sorts of reversibility and that confusing the two is a basic category mistake.⁴⁻⁶⁾

Landauer claimed that since information in an actual computation is represented by physical systems, the laws of physics do indeed connect the logical processes being implemented with physical consequences like heat generation.^{7,8)} In particular, according to his argument, (a) there is a fundamental minimum heat dissipation of $k_B T \log 2$ required when a physical process implements a logically irreversible computation, i.e., when information is erased, and (b) no fundamental minimum energy need be dissipated for a physical process that implements a logically reversible computation, i.e., when information is preserved. Erasure here refers to the irretrievable transfer of physical state information into a thermodynamically large system, a heat bath, at temperature T . Bennett showed that in fact any logically irreversible computation could be embedded in a logically reversible computation,^{9,10)} making the result in principle completely general.

This relationship between physical and logical reversibility is now called Landauer's Principle (LP) and it provides a fundamental connection between information theoretic accounts of computation and the thermodynamic limits of physical computing systems. LP contradicts the view of Szilard,¹¹⁾ Brillouin,¹²⁾ and others,^{5,6)} that it is *measurement* that necessarily entails unavoidable heat dissipation. There is a "cost" for lowering the energy

dissipated, according to LP. To achieve lower and lower dissipation in the information-preserving case, the physical computation needs to be performed more and more gradually (quasi-adiabatically).¹³⁾ Of course, there is always some energy dissipation, and indeed for practical devices some dissipation is necessary for stability.¹⁴⁾ According to LP, however, if information is not erased, there is no fundamental lower limit to how much energy needs to be dissipated to the environment.

While LP can be said to enjoy fairly broad support, there is a substantial literature that either denies the validity of the principle outright,^{5,6,15)} or calls it into question on the grounds that existing "proofs" are fatally flawed, an argument given at length by Earman and Norton.¹⁶⁻¹⁹⁾ It is characteristic that the arguments for and against LP involve idealized systems, often of single-particle gases in chambers with frictionless pistons, partitions, and the like.²⁰⁾ Such systems can be, and often are, criticized for sweeping important details under the rug and ignoring possibly dissipative steps.

Heretofore, no direct experimental tests of LP at the $k_B T$ level have been conducted to help bring clarity to the fundamental issue.

We describe an experimental test of Landauer's Principle with an energy resolution that is substantially smaller than $k_B T$. Such a test is challenging because to reliably represent a bit in a thermal environment, the energy of the bit must be *larger* than $k_B T$, while the limiting dissipation we seek to measure is *less* than $k_B T$ and subject to thermal fluctuations. At room temperature $k_B T \log 2 = 0.69 k_B T = 2.8$ zJ. We compare the energy dissipated in the device when logically reversible and irreversible operations are performed.

2. Reversible and Irreversible Operations

We choose a system with three distinct states representing a binary "1", a binary "0", and a "null" state which holds no information. The null state is comparable to the space on a Turing machine tape. This approach avoids needless confusion that arises if one views erasure as resetting to a standard state, usually "0", which makes it difficult to distinguish erasure from bit-switching.

We consider the three processes shown schematically in Fig. 1. The first, illustrated by Fig. 1(a), is a copy operation in which the information in a system S is copied into the memory device M . Initially, S holds a "1" or "0" bit and M is in the null state. An interaction is then established which smoothly switches M into the same state as S . This operation is logically reversible and therefore can, according to LP, be implemented in a physically reversible process, dissipating

*E-mail address: snider.7@nd.edu

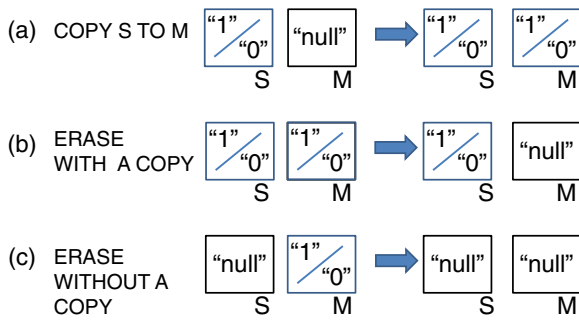


Fig. 1. (Color online) Schematic view of three binary switching operations. (a) The COPY operation results in the state of memory element M being switched to “1” or “0” to match the initial state of S. (b) ERASE WITH A COPY is a logically reversible operation which sets the memory element to the “null” state in the presence of a copy of the stored information. (c) The ERASE WITHOUT A COPY operation sets M to the “null” state with no copy of the bit present.

arbitrarily small amounts of energy (at the cost of speed). Bit information may be represented in physically different ways in the two systems. Note that any measurement is a COPY operation which copies the information about a physical system to the measurement device.

The second process is a logically reversible process which we call ERASE WITH A COPY, shown in Fig. 1(b). Here the role of system S is different—it holds a copy (a duplicate) of the information stored in the memory device M. The state of M is then smoothly switched to the null state. The presence of the copy of the information allows the dynamical path which M traverses as it moves from the active (“0” or “1”) state to the null state to depend on the content of the information. Of course the mere existence of a copy of the information does not necessarily mean that the information will be used effectively to switch the state of M to the null state in the least dissipative way; it simply makes it possible to do so. The system S can be as simple and automatic as a neighboring memory cell or as complex as the experimenter’s brain.

The third process is an irreversible ERASE WITHOUT A COPY process, as shown in Fig. 1(c). In this case, the system S holds no information about the content of M, and indeed it is assumed that no copy of the information stored in M exists anywhere else. The memory device is switched from this unknown active state to the null state. Erasure as a physical process must move the physical system M from either the “1” state or the “0” state to the null state. This many-to-one transition is at the heart of the (contentious) connection between physics and information. According to LP, this process must dissipate at least $k_B T \log 2$.

3. Experimental Realization

Each of these three processes is implemented experimentally using a switched voltage source to realize the system S and an RC circuit to realize the single bit memory element M.²¹⁾ Figure 2 illustrates the COPY experiment. The switch on the left is set to one of two voltage sources, each initially at ground (representing the null state), which then smoothly ramps V_{in} to a value of either $+V_0$ or $-V_0$, representing an input “1” or “0” bit respectively. The system S provides the input bit to be copied. The memory device M is composed of




Fig. 2. (Color online) The COPY operation performed experimentally. There is initially no charge stored on the capacitor in the memory element M, representing the “null” state. System S selects the bit value (0 or 1) to be stored by connecting the capacitor to one of two voltage ramps which smoothly increase to $\pm V_0$. The information in S is then copied to M as the capacitor is charged to $\pm Q_0$. If the switch in S is then moved to the hold position, the information is held in M.

a capacitor $C = 100$ pF, and series resistor $R = 1.1$ k Ω , as shown. A “1” (“0”) bit is represented in M by a charge of $+Q_0$ ($-Q_0$) and a corresponding voltage $+V_0$ ($-V_0$) on the capacitor where $V_0 = Q_0/C$. The bit energy is

$$E_{\text{bit}} = \frac{1}{2} \frac{Q_0^2}{C} = \frac{1}{2} C V_0^2. \quad (1)$$

Since the energies of the “0” and “1” states are identical this is a symmetric binary memory. The null state is represented by having no charge or voltage on the capacitor. A small amount of dissipation occurs in the series resistor R over the finite time T in which the voltage is ramped to $\pm V_0$. We measure the voltage $V_R(t) = V_{in}(t) - V_C(t)$ across the resistor and calculate the total energy

$$E_{\text{diss}} = \frac{1}{R} \int_0^T V_R^2(t) dt \quad (2)$$

dissipated during the process. If the switch S is then set to the “hold” position, the bit is simply maintained in M.

To implement ERASE WITH A COPY we start with a bit stored in memory system M. The switch S is first set to select the appropriate source at potential $\pm V_0$, depending on the value of the stored bit. This formally entails measuring the bit stored in M, which according to LP is equivalent to a reversible COPY operation, this time making a copy of the bit previously stored in M (the reverse of the COPY operation described above). Having this information enables the input to be set to match the present state of M. The voltage is then gradually ramped to 0 as shown schematically in Fig. 3. This gradually discharges the capacitor. Dissipation occurs in the resistor only because of the small but non-zero potential difference which develops across the resistor during the finite time of the transition. We measure the voltage across the resistor $V_R(t)$ as above and find the energy dissipated, E_{diss} using eq. (2) as the bit in M is erased. We have analyzed the loading of the measuring amplifier and shown that the effect is negligible. This process represents quasi-adiabatic erasure of the bit in M, in the presence of a copy—the copy is physically realized by the position of the switch in S.

Figure 4 shows V_{in} , V_C , and V_R , as first a “1” and then a “0” bit is transferred to M with a COPY operation and subsequently erased with an ERASE WITH A COPY operation. The time over which the voltage is ramped up or down is




Fig. 3. (Color online) The ERASE operation performed experimentally in two ways. Initially, the bit is represented by the charge $\pm Q_0$ held on the capacitor in memory device M, with corresponding voltage $V_C = \pm V_0$. For the ERASE WITH A COPY operation, the voltage is measured and the switch in S moved to the appropriate voltage source which starts at $\pm V_0$ and smoothly ramps to 0, with minimal dissipation in R. For the ERASE WITHOUT A COPY operation, the bit stored in M is unknown, so the switch S is moved to ground and the capacitor is discharged through R to ground, dissipating the full bit energy $(1/2)CV_0^2$.




Fig. 4. (Color) Measured voltages for reversible operations. (a) The input voltage ramps up for the COPY operation, then down for the ERASE WITH A COPY operation. (b) The measured capacitor voltage $V_C(t)$ (red) represents the signal that holds the bit. The blue curve shows V_R on the same scale for comparison. Neither is averaged. (c) The measured voltage across R, V_R , averaged over 3×10^6 samples.

$T = 640 \mu\text{s}$. Because the voltage source is changed slowly (compared to RC), the voltage across the resistor is always very small and is challenging to measure. At this scale the thermal environment produces considerable random fluctuations. The primary sources of noise are Johnson noise in the resistor, $S_{\text{Johnson}} = 4.3 \text{ nV}/\text{Hz}^{1/2}$, and noise in the amplifier used to measure the very small V_R , $S_{\text{amp}} = 9 \text{ nV}/\text{Hz}^{1/2}$. To adequately measure the voltage across the resistor and detect the increased power dissipation during switching, the measurement is averaged over many, N_{diss} , switching events. Note that no averaging is done for $V_C(t)$, the signal representing the bit itself; averaging is only necessary to detect the dissipation because it is so small. The bit energy must be substantially larger than $k_B T$ to be robust in a thermal environment.

The ERASE WITHOUT A COPY process is realized by first storing a bit on M, which is then treated as unknown. No measurement of the bit is made, so there is no copy of the information outside of M itself. As a result we cannot choose to connect M to the appropriate voltage source and ramp it down gradually. The only way to erase the bit is to connect M to ground directly, and let the capacitor discharge. The result is a substantial current flow and a correspondingly large energy is dissipated in the resistor, which is measured as above. The discharge happens too quickly for the acquisition system to fully capture all the energy dissipated, which must of course be the full bit energy $E_{\text{bit}} = (1/2)CV_0^2$.

4. Results

All experiments were performed at room temperature. For the COPY and ERASE WITH A COPY operations, bit energies of 3000 and $30 k_B T$, and switching times T of 64 , 256 , and $640 \mu\text{s}$ were used. The ERASE WITHOUT A COPY operation involves discharging the capacitor directly through the resistor and so happens very quickly. For that case we have $E_{\text{bit}} = 30 k_B T$ and we take for the switching time $T = 5 RC$. The results are shown graphically in Fig. 5 and tabulated in Table I.

In agreement with the predictions of LP, the logically reversible operations COPY and ERASE WITH A COPY show an energy dissipation which is significantly less than E_{bit} , and decreases as the switching time T is lengthened. For the specific case shown in Fig. 4, with a bit energy of $E_{\text{bit}} = 30 k_B T$, and a switching time of $640 \mu\text{s}$, the energy dissipated in each operation was only $E_{\text{diss}} \sim 0.01 k_B T = 41 \text{ yJ}$, significantly below $k_B T \log 2$. The only limitations to lowering the energy dissipated further are the lengthening switching time and the averaging necessary to measure the dissipation (averaging not required to measure the bit). Thus even at levels well below $k_B T$, the LP prediction is confirmed.

By contrast the logically irreversible ERASE WITHOUT A COPY function dissipates the full bit energy which is necessarily more than $k_B T$. Here we measure only $20.7 k_B T$ because of the difficulty of capturing the early rapid discharge.

5. Discussion

These experiments make clear that a key to physical reversibility in the ERASE WITH A COPY (or in any reversible function) is that the presence of the copy of the stored information allows the system to be biased in the state it is already in, which permits the gradual (adiabatic) shift to the null state. In this case the voltage source in S can be set to the same value as the voltage in M, provided there is a copy of the information available. This measurement is itself a COPY operation, which we have seen can be performed with arbitrarily small dissipation. If the bit in M is unknown, this cannot be done, and so an erasure procedure that works for an unknown bit must be implemented—connecting the signal to ground. Norton asks if one has proven that this is the best one can do.¹⁹ This concrete example drives home the point: how else could one do it? By what process could one gradually drain charge of an unknown sign off the capacitor, without first measuring the sign of the charge and thereby creating a copy of the bit? It seems apparent that no such process exists.

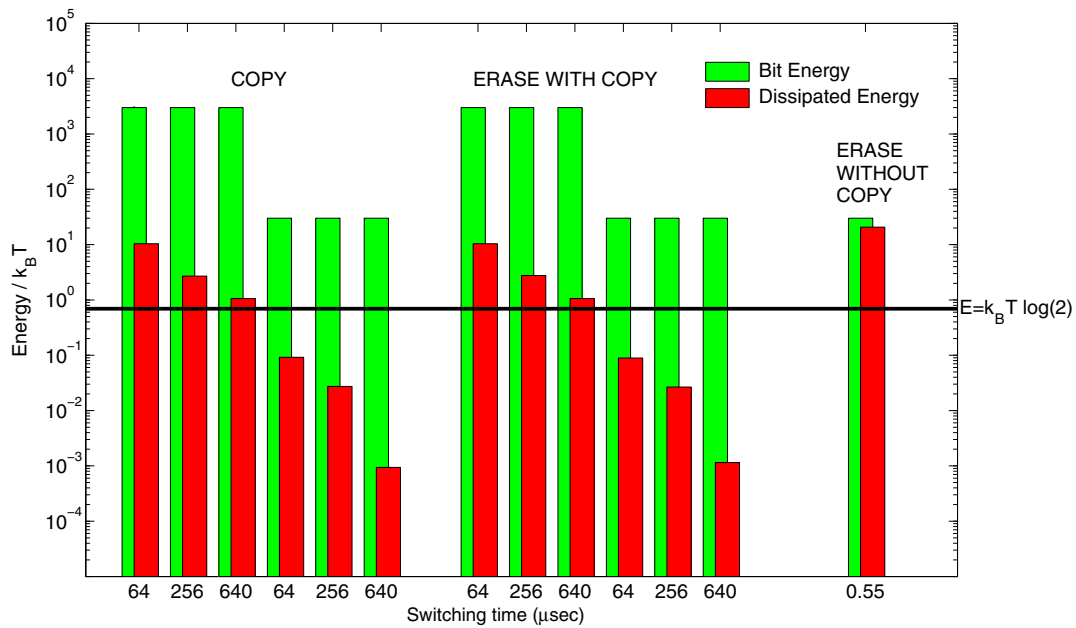


Fig. 5. (Color online) Bit energy and dissipated energy in performing reversible and irreversible functions. The measured energy dissipated when performing the COPY and ERASE WITH A COPY operations are shown for bit energies of 3000 and $30 k_B T$, and in each case for switching times of 64, 256, and 640 μs . For these logically reversible functions as the switching time increases, the measured energy dissipation decreases to much less than $k_B T \log 2$, shown with the black horizontal line. By contrast, the logically irreversible function ERASE WITHOUT A COPY, shown on the right, dissipates much more than $k_B T \log 2$.

Table I. Bit energy and dissipated energy for reversible and irreversible functions. Values are measured and tabulated for several switching times T and bit-storage voltages V_0 . For logically reversible functions COPY and ERASE WITH A COPY the dissipated energy decreases as the switching time increases, even below $k_B T \log 2$. For the logically irreversible function ERASE WITHOUT A COPY, dissipation is much greater than $k_B T \log 2$. Measurements of power dissipated in the resistor are averaged over N_{diss} switching events.

Function	$E_{\text{diss}}/k_B T$	$E_{\text{bit}}/k_B T$	T (μs)	V_0 (μV)	N_{diss}
COPY	10.4	3000	64	500	10^3
	2.70	3000	256	500	10^4
	1.06	3000	640	500	10^6
	0.0917	30	64	50	10^4
	0.0273	30	256	50	3×10^6
	0.00939	30	640	50	3×10^6
ERASE WITH A COPY	10.35	3000	64	500	10^3
	2.77	3000	256	500	10^4
	1.06	3000	640	500	10^6
	0.0893	30	64	50	10^4
	0.0266	30	256	50	3×10^6
	0.0115	30	640	50	3×10^6
ERASE WITHOUT A COPY	20.7	30	0.55	1000	1

It has been argued that fluctuations due to thermal noise from the environment would necessarily thwart attempts to either measure (COPY) or erase a bit without dissipating at least $k_B T \log 2$.^{4,5,15,19} It is true that thermal noise makes the measurement of very small amounts of energy dissipation difficult, so averaging was required to measure E_{diss} . But importantly, averaging is not necessary to measure the bit information itself. Consider the ERASE WITH A COPY experiment. Each individual member of the averaged ensemble

has different fluctuations in $V_R(t)$ resulting from thermal excitations, which are associated with extra energy flowing into, and out of, the resistor from the environment. These fluctuations do not, however, alter the energy budget of the process—extra energy transfers from the thermal environment to the system and then transfers back. The net result of thermal fluctuations does not increase or decrease the energy dissipated, which is here substantially less than $k_B T$. The voltage across an isolated resistor (disconnected from a circuit) has thermal Johnson noise, but the resistor does not therefore heat the room, even a little. Moreover these thermal fluctuations do not generate bit errors in this measurement, because the bit energy is substantially larger than $k_B T$; only during the initial part of writing the bit or the final part of erasing the bit does its energy become small. Each member of the COPY ensemble, for example, has a V_C which is initially at zero potential (with zero bit energy) and finally is at $\pm V_0$, regardless of the thermal fluctuations that occur during the charging process (which are visible in Fig. 4).

These results show that the reversible operation ERASE WITH A COPY can be implemented with arbitrarily small dissipation. The copy of the initial data, of course, still persists in system S and it is now subject to the same considerations as the original bit. It could be erased with very low dissipation if another copy is made, and then the same constraints would apply to that copy, etc. If it is eventually erased with no copy remaining, this will necessarily dissipate the bit energy, at least $k_B T \log 2$. In this experiment, the system S includes a conventional transistor-based signal generator and so erases the copy afterward dissipatively. This in no way invalidates the analysis of the system bit M. The same LP constraint which applies to M follows the bit until it is ultimately erased dissipatively or stored.

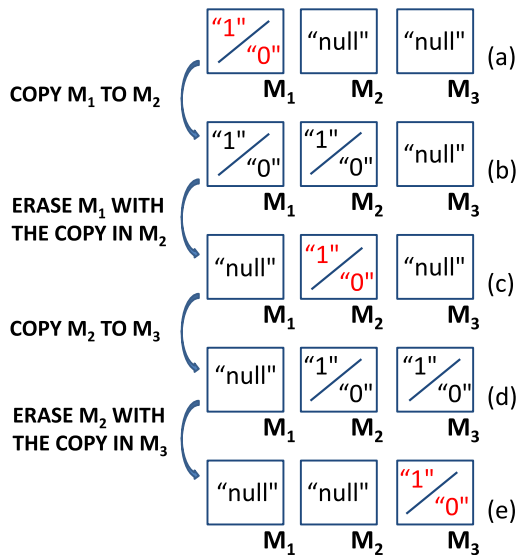


Fig. 6. (Color online) Repeated application of the COPY and ERASE WITH A COPY operations can be used to move information from place to place with minimal energy dissipation. There is no fundamental lower limit on the amount of energy that needs to be dissipated to transport a bit.

The COPY and ERASE WITH A COPY operations are of particular importance because they can potentially form the basis of on-chip information transmission. Figure 6 illustrates how a sequential application of these operations moves information from one place to another with minimal energy dissipation. The COPY operation makes a copy in M_2 of the information stored in M_1 . The copy in M_2 can then be used to erase M_1 with minimal dissipation. This is the same operation as shown in Fig. 1(b), but with the roles of the cells on left and right reversed. The process is then repeated, moving the information to M_3 and so on. In a linear array of bits this forms a shift register. Such an adiabatic shift register has been experimentally realized.^{22–24} Again, there is always some dissipation of energy, but it can be as little as desired at the cost of speed. This is a consequence of LP also pointed out (again with some controversy) by Landauer:²⁵ there is no fundamental lower limit to the amount of energy dissipation required to move information from place to place. Communication is in this sense fundamentally free. Present complementary metal oxide semiconductor (CMOS) circuits do not operate adiabatically in either computation or information transport around the chip, though they could,^{26–31} again at some cost in the speed of operation.

6. Conclusions

We have measured energy dissipation of significantly less than $k_B T \log 2$ when performing logically reversible functions COPY and ERASE WITH A COPY, but an irreversible function ERASE WITHOUT A COPY dissipates much more than

$k_B T \log 2$. In each case the signal energy stored in the bit was much greater than $k_B T \log 2$. The COPY procedure is equivalent to a measurement and shows definitively that the argument that measurement causes unavoidable energy dissipation is incorrect. Accounting for thermal fluctuations is important in measuring small signals in a thermal environment, but thermal fluctuations do not alter the basic energetics of computation, and do not necessarily produce bit errors. The first experimental test of Landauer’s Principle at the sub- $k_B T$ supports its validity.

Acknowledgments

Supported by NSF ECS-0725794, ECCS-0901659, and CHE-1124762.

- 1) N. G. Anderson: *IEEE Trans. Nanotechnol.* **7** (2008) 521.
- 2) K. Galatsis, A. Khitun, R. Ostroumov, K. L. Wang, W. R. Dichtel, E. Plummer, J. F. Stoddart, J. I. Zink, J. Y. Lee, Y.-H. Xie, and K. W. Kim: *IEEE Trans. Nanotechnol.* **8** (2009) 66.
- 3) J. J. Welsler, G. I. Bourianoff, V. V. Zhirnov, and R. K. Cavin: *J. Nanopart. Res.* **10** (2008) 1.
- 4) O. J. E. Maroney: *Stud. Hist. Philos. Mod. Phys.* **36** (2005) 355.
- 5) W. Porod, R. O. Grondin, D. K. Ferry, and G. Porod: *Phys. Rev. Lett.* **52** (1984) 232.
- 6) W. Porod, R. Grondin, D. K. Ferry, and G. Porod: *Phys. Rev. Lett.* **53** (1984) 1206.
- 7) R. Landauer: *IBM J. Res. Dev.* **5** (1961) 183.
- 8) R. W. Keyes and R. Landauer: *IBM J. Res. Dev.* **14** (1970) 152.
- 9) C. H. Bennett: *IBM J. Res. Dev.* **17** (1973) 525.
- 10) C. H. Bennett: *Int. J. Theor. Phys.* **21** (1982) 905.
- 11) L. Szilard: *Z. Phys.* **53** (1929) 840 [in German].
- 12) L. Brillouin: *J. Appl. Phys.* **22** (1951) 338.
- 13) J. Timler and C. S. Lent: *J. Appl. Phys.* **94** (2003) 1050.
- 14) W. H. Zurek: *Phys. Rev. Lett.* **53** (1984) 391.
- 15) R. K. Cavin, V. V. Zhirnov, J. A. Hutchby, and G. I. Bourianoff: *Fluctuation Noise Lett.* **5** (2005) C29.
- 16) J. Earman and J. D. Norton: *Stud. Hist. Philos. Mod. Phys., Part B* **29** (1998) 435.
- 17) J. Earman and J. D. Norton: *Stud. Hist. Philos. Mod. Phys., Part B* **30** (1999) 1.
- 18) J. D. Norton: *Stud. Hist. Philos. Mod. Phys., Part B* **36** (2005) 375.
- 19) J. D. Norton: *Stud. Hist. Philos. Mod. Phys.* **42** (2010) 184.
- 20) H. Leff and A. F. Rex: *Maxwell Demon 2: Entropy, Classical and Quantum Information, Computing* (IOP Publishing, Bristol, U.K., 2002).
- 21) G. P. Boechler, J. M. Whitney, C. S. Lent, A. O. Orlov, and G. L. Snider: *Appl. Phys. Lett.* **97** (2010) 103502.
- 22) A. O. Orlov, I. Amlani, G. H. Bernstein, C. S. Lent, and G. L. Snider: *Science* **277** (1997) 928.
- 23) A. O. Orlov, R. Kumamuru, R. Ramasubramaniam, C. S. Lent, G. H. Bernstein, and G. L. Snider: *Surf. Sci.* **532** (2003) 1193.
- 24) C. S. Lent, M. Liu, and Y. H. Lu: *Nanotechnology* **17** (2006) 4240.
- 25) R. Landauer: *Appl. Phys. Lett.* **51** (1987) 2056.
- 26) S. G. Younis and T. F. Knight: *Research on Integrated Systems* (MIT Press, Cambridge, MA, 1993) p. 234.
- 27) D. J. Frank and P. M. Solomon: *Proc. Int. Symp. Low Power Design (ISLPED 95)*, 1995, p. 197.
- 28) W. C. Athas, L. J. Svensson, J. G. Koller, N. Thartzanis, and E. Chou: *IEEE Trans. VLSI Syst.* **2** (1994) 398.
- 29) G. Dickinson and J. S. Denker: *IEEE J. Solid-State Circuits* **30** (1995) 311.
- 30) W. C. Athas: in *Low Power Design Methodologies*, ed. J. Rabaey and M. Pedram (Kluwer, Boston, MA, 1996) p. 63.
- 31) R. C. Merkle: *Nanotechnology* **4** (1993) 21.