

December 22, 28, 2011

June 20, 2012

March 31, 2013

The End of the Thermodynamics of Computation: A No Go Result

John D. Norton

Department of History and Philosophy of Science

Center for Philosophy of Science

University of Pittsburgh

Pittsburgh PA 15260

<http://www.pitt.edu/~jdnorton>

jdnorton@pitt.edu

PSA 2012: Philosophy of Science Association Biennial Meeting

To appear in *Philosophy of Science*.

The thermodynamics of computation assumes that computational processes at the molecular level can be brought arbitrarily close to thermodynamic reversibility; and that thermodynamic entropy creation is unavoidable only in data erasure or the merging of computational paths, in accord with Landauer's principle. The no go result shows that fluctuations preclude completion of thermodynamically reversible processes. Completion can be achieved only by irreversible processes that create thermodynamic entropy in excess of the Landauer limit.

1. Introduction

Electronic computers degrade work to heat and the need for its removal sets a practical limit to their performance. The study of the thermodynamics of computation, surveyed in Bennett (1982), seeks the limits in principle to reduction of this dissipation. Since dissipation reduces with size, the most thermodynamically efficient computers are sought among those that use individual molecules, charges or magnetic dipoles as memory storage devices.

These molecular-scale processes are treated like macroscopic ones in one aspect: they can be brought arbitrarily close to the most efficient, non-dissipative processes, those that are

thermodynamically reversible. Their defining characteristic is that they are at equilibrium at every stage. They are brought slowly from start to finish by the successive nudges of miniscule disequilibria. It is assumed that the dissipative effects of these nudges can be made arbitrarily small by indefinitely extending the time allowed for the process to reach completion.

Some form of dissipation, however, is judged unavoidable. The controlling idea of the thermodynamics of computation is that the creation of thermodynamic entropy and the associated need to pass heat to the environment arise only with logically irreversible operations. These include the erasure of data and the merging of computational paths. The amount of thermodynamic entropy created is quantified by Landauer's principle. It asserts that at least $k \ln 2$ of thermodynamic entropy is created when one bit of data is erased. The result is an elegant account of the bounds to the thermodynamic efficiency of computation. They are independent of the physical implementation, but are set by the logical operations comprising the computation.

Alas, this image of a well-developed science is an illusion. The thermodynamics of computation is an underdeveloped muddle of vague plausibility arguments and misapplications of statistical physics. Earman and Norton (1998, 1999) track the science's history through the Maxwell demon problem and find it rife with circular reasoning and question begging. Norton (2005, 2011) urges that the arguments used to support Landauer's principle are fallacious and have never successfully advanced beyond flawed plausibility arguments. Erasure may reduce the range of possible values for data in a memory. But this reduction is not a compression of the accessible phase space of thermodynamic components that can be associated with a change of thermodynamic entropy. The volume of accessible phase space remains unchanged in erasure. Prior to erasure we may also be unsure as to the data stored and assign probabilities to the possibilities. That sort of probability, however, is not associated with a thermodynamic entropy.

Finally, Norton (2011) describes a "no go" result—that thermodynamically reversible processes at molecular scales are precluded from proceeding to completion by fluctuations. Individual computational steps can only be completed if they are sufficiently far from equilibrium to overcome fluctuations. As a result they create quantities of thermodynamic entropy in excess of those tracked by Landauer's principle. It follows that the lower limit to thermodynamic entropy creation is not set by the logical specification of the computation, but by the details of the particular physical implementation and the number of discrete steps it employs, whatever their function.

This paper will develop the no go result. It is motivated and stated in the next section and then illustrated in Section 3.

2. The No Go Result

2.1 A Preliminary Form

In a thermodynamically reversible process,¹ all component systems are in perfect equilibrium with one another at all stages. As result, they are impossible processes.² Nothing changes. Heat will not spontaneously pass from one body to another if they are at the same temperature. In ordinary thermodynamics, this awkwardness is overcome by introducing a slight disequilibrium. We minutely raise the temperature of the first body and let that minute temperature gradient drive the heat transfer, slowly. Because heat is now passing spontaneously from hot to cold, this is a dissipative process. The thermodynamic entropy created measures the amount of dissipation. For theoretical analyses, this entropy creation can be neglected since it can be made as small as we like by making the driving temperature difference appropriately small. The process will still go forward, but more slowly.

Matters are different when we allow for the molecular constitution of matter. For now the equilibrium of a thermodynamically reversible process is dynamic. If two bodies at the same temperature are in thermal contact, energy will spontaneously pass to and fro between them as energy fluctuations due to random, molecular-scale events. If we are to assure that heat passes from the one to the other, we must arrange for a disequilibrium that is sufficiently great to overcome the fluctuations.

¹ Typical erasure processes begin with a thermodynamically irreversible process in which the memory device is thermalized. For example, the wall dividing a two-chamber memory cell is raised so the molecule can access both chambers. The resulting uncontrolled, thermodynamically irreversible expansion creates the $k \ln 2$ of thermodynamic entropy tracked by Landauer's principle. As Norton (2005, Section 3.2) argues, a mistaken tradition misidentifies this thermalization as thermodynamically reversible since the replacing of the partition supposedly returns the original state of "random data."

² For an analysis of thermodynamically reversible processes, see Norton (forthcoming, §3).

Boltzmann's Principle, " $S = k \ln W$," that is, "entropy = $k \ln$ probability," measures the dissipation needed. An isolated system is to pass from state 1 with total thermodynamic entropy S_1 to state 2 with total entropy S_2 . The inverted principle tells us that, if the system can spontaneously move between the two states, then the probabilities P_1 and P_2 of the two states are related by

$$P_2/P_1 = \exp((S_2 - S_1)/k) \quad (1)$$

In macroscopic terms, negligible thermodynamic entropy creation is sufficient to drive processes to completion. If $S_2 - S_1 = 10k$, a macroscopically negligible amount, we find $P_2/P_1 = 22,026$, so that the final state 2 is strongly favored.

At the molecular level, these amounts of thermodynamic entropy are large. They exceed the entropy change of $k \ln 2 = 0.69k$ tracked by Landauer's principle. They must exceed it, for creation of merely $k \ln 2$ of entropy is insufficient to assure completion of a process. Then $P_2/P_1 = \exp(k \ln 2/k) = 2$. The process is only twice as likely to be in its final state 2 as in its initial state 1. This is a fatal result for the thermodynamics of computation. If we have any computing process with multiple steps operating at molecular scales, we must create thermodynamic entropy in each step if the process is to go forward, quite aside from any issues of logical irreversibility.

2.2 The Main Result

Boltzmann's Principle in the form (1) applies to isolated systems. In the thermodynamics of computation, the computing systems are treated as open systems, in equilibrium with a heat bath at the ambient temperature T . The main result arises when we adapt these considerations to such systems.

A computer is a system consisting of many interacting components, including memory cells, systems that read and write to the memory cells and other control components to implement the computer's program. At any moment, the combined system is in thermal equilibrium with the environment at temperature T . Hence, the system is canonically distributed over its phase space, according to the probability density

$$p(\mathbf{x}, \boldsymbol{\pi}) = \exp(-E(\mathbf{x}, \boldsymbol{\pi})/kT) / Z$$

where Z is the normalizing partition function and \mathbf{x} and $\boldsymbol{\pi}$ are multi-component generalized configuration and momentum coordinates.

Each computational step is carried out by a thermodynamically reversible process, whose stages are parameterized by λ . Fluctuations will carry the system spontaneously from one stage to another. As a result, the system is probabilistically distributed over the different stages. The probabilities are computed by Einstein's methods, as adapted by Tolman (1938, pp. 637-38), and conform to the probability density

$$p(\lambda) = \text{constant} \cdot Z(\lambda) \quad (2)$$

where $Z(\lambda)$ is given by

$$Z(\lambda) = \int_{\lambda} \exp(-E(\mathbf{x}, \boldsymbol{\pi})/kT) \, d\mathbf{x}d\boldsymbol{\pi}$$

This last integral extends over the volume of phase space accessible to the system when the process is at stage λ .

In the Einstein-Tolman analysis, each of these stages is given a thermodynamic description as if it were an equilibrium state, even though it may have arisen through a fluctuation. The canonically distributed system at stage λ is assigned a canonical free energy

$$F(\lambda) = -kT \ln Z(\lambda) \quad (3)$$

treating $Z(\lambda)$ as a partition function, where the free energy is defined as

$$F(\lambda) = E(\lambda) - TS(\lambda)$$

Here $E(\lambda)$ and $S(\lambda)$ are the mean energy and the thermodynamic entropy assigned to the system in stage λ . It now follows from (2) and (3) that

$$p(\lambda) = \text{constant} \cdot \exp(-F(\lambda)/kT)$$

and that the probability densities for the system fluctuating between stages λ_1 and λ_2 satisfy

$$p(\lambda_2)/p(\lambda_1) = \exp(-(F(\lambda_2) - F(\lambda_1))/kT) \quad (4)$$

The process is thermodynamically reversible. Hence it is in equilibrium at every stage. Equilibrium requires the vanishing of the generalized thermodynamic force $X(\lambda)$ acting on the system:³

³ At equilibrium, the total entropy S_{tot} of the system S_{sys} and the environment S_{env} is stationary.

Writing $d = \partial/\partial\lambda|_T$, that amounts to $0 = dS_{\text{tot}} = dS_{\text{sys}} + dS_{\text{env}}$. By supposition, the computer

$$X(\lambda) = - \partial/\partial\lambda|_T F(\lambda) = 0$$

Integrating over λ , we find that the free energy $F(\lambda)$ is constant over the stages of the process:

$$F(\lambda) = \text{constant} \quad F(\lambda_1) = F(\lambda_2) \quad (5)$$

From (4), we have that

$$p(\lambda) = \text{constant} \quad p(\lambda_1) = p(\lambda_2) \quad (6)$$

This last result (6) is the no go result. It precludes thermodynamically reversible processes proceeding as we expect.

Our default expectation is that these processes are in a quiescent equilibrium at every stage λ , perhaps with a slight disturbance due to fluctuations. We expect to bring the process from its initial to its final stage by minute disequilibrium nudges that advance the process arbitrarily slowly in the tiniest of steps. What (6) tells us is that fluctuations obliterate the quiescent equilibrium. If the system is in one stage λ at some moment, it is equally likely to be found at the next moment at any other stage. If we set up the process in its initial stage, it is as likely to leap by a fluctuation to the final stage as it is to stay where it is. If the process has arrived at the final stage, it is as likely to be flung by a fluctuation back to its initial stage, as it is to stay where it is. In a slogan, fluctuations obliterate thermodynamically reversible processes.

Fluctuations are temperature sensitive. Hence we might expect the confounding effects of fluctuations to be calmed and controlled by cooling the processes, perhaps even close to absolute zero. A review of the calculation above shows that the no go result (6) obtains no matter what the temperature, even if it close to absolute zero.⁴

system exchanges no work with the environment, but only heat in a thermodynamically reversible process. Hence $dS_{\text{env}} = dE_{\text{env}}/T = - dE_{\text{sys}}/T$, where the last equality follows from conservation of energy: $dE_{\text{env}} + dE_{\text{sys}} = 0$. Combining, we have $0 = dS_{\text{sys}} - dE_{\text{sys}}/T$. Hence the condition for equilibrium is $0 = d(E_{\text{sys}} - TS_{\text{sys}}) = -X_{\text{sys}}$.

⁴ Temperature does affect the free energy needed to override the fluctuations. We see below that a probabilistic favoring of 20:1 is achieved by a free energy reduction of $3kT$. This reduction diminishes as T decreases. However the thermodynamic entropy created remains at least $3k$, independent of the temperature.

2.2 What It Takes to Beat Fluctuations

If fluctuations obliterate thermodynamically reversible processes, how is it possible for these processes to figure in thermodynamic analysis at all? The answer is that the disequilibrium required to overcome fluctuations is negligible macroscopically. While the no go result applies to macroscopic systems, it is overcome by disequilibria too small to trouble us. However, at the molecular scale explored by the thermodynamics of computation, the situation is reversed. There, the disequilibria needed to overcome fluctuations dominate. Most importantly, it requires thermodynamic entropy creation in amounts that well exceed those tracked by Landauer's principle.

A few computations illustrate this answer. Relation (4) tells us that we can probabilistically favor the end stage λ_2 over the initial stage λ_1 if the end stage free energy $F(\lambda_2)$ is smaller than the initial stage free energy $F(\lambda_1)$. A decrease of $3kT$ is sufficient for a modest favoring in the ratio of 20:1, for then

$$p(\lambda_2)/p(\lambda_1) = \exp(-(-3kT)/kT) = \exp(3) = 20$$

The dissipation associated with the reduction in free energy $F(\lambda_2) - F(\lambda_1) = -3kT$ is a minimum increase in the thermodynamic entropy of⁵

$$\Delta S_{\text{tot}} = \Delta S_{\text{sys}} + \Delta S_{\text{env}} = 3k$$

where the change Δ is applied to the entropy of the universe as a whole S_{tot} , which is the sum of the system entropy S_{sys} and the environment entropy S_{env} . Even though this modest probabilistic favoring by no means assures completion of the process, the entropy creation of at least $3k$ is many times greater than the $k \ln 2 = 0.69k$ of entropy tracked by Landauer's principle in a single bit erasure.

⁵ To see this, use $F=E-TS$ to rewrite $F(\lambda_2) - F(\lambda_1) = -3kT$ as

$$S(\lambda_2) - S(\lambda_1) - (E(\lambda_2) - E(\lambda_1))/T = 3k$$

We have $\Delta S_{\text{sys}} = S(\lambda_2) - S(\lambda_1)$. By conservation of energy, $-(E(\lambda_2) - E(\lambda_1))$ is the energy gained by the environment. By supposition, this energy is passed by heat transfer only. In the least dissipative case of a thermodynamically reversible heat transfer that corresponds to the minimum increase of entropy $\Delta S_{\text{env}} = -(E(\lambda_2) - E(\lambda_1))/T$.

Since the ratio of probability densities grows exponentially with free energy differences in (4), further creation of thermodynamic entropy can bring probability density ratios that strongly favor completion of the process. For example, if we increase the free energy difference to $25kT$, then the end stage is strongly favored, for

$$p(\lambda_2)/p(\lambda_1) = \exp(-(-25kT)/kT) = \exp(25) = 7.2 \times 10^{10}.$$

In macroscopic terms, however, $25kT$ of free energy is negligible. This quantity, $25kT$, is the mean thermal energy of ten diatomic molecules, such as ten oxygen molecules. Hence, there is no obstacle to introducing a slight disequilibrium in a macroscopic system in order to nudge a thermodynamically reversible process to completion.

3. Illustrations of the No Go Result for a One-Molecule Gas

This no go result applies to all thermodynamically reversible processes in systems in thermal equilibrium with their environment. However its derivation and its statement as (6) is remote from its implementation in specific systems. It is helpful to illustrate how fluctuations obliterate a simple process described in the thermodynamics of computation, the thermodynamically reversible, isothermal expansion and compression of a one-molecule gas. The analysis of the last section provides the precise computation. Here I give simpler estimates of the disturbing effects of fluctuations.

3.1 Reversible, Isothermal Expansion and Compression

A monatomic one-molecule gas is confined to a vertically oriented cylinder and the gas pressure is contained by the weight of the piston. The process intended is a thermodynamically reversible, isothermal expansion or compression of the gas. Our expectation is that this process will proceed indefinitely slowly, with the weight of the piston maintained just minutely away from the equilibrium weight so that the expansion or compression is only just favored. As the piston is raised in an expansion, it draws work energy from the one-molecule gas; and this energy is restored to the one-molecule gas as heat from the environment. The gas exerts a pressure $P=kT/V$, for V the volume of the gas. Thus the work extracted in a doubling of the volume and thus also the heat passed to the gas is given by $\int_V^{2V} kT/V' dV' = kT \ln 2$. The thermodynamic entropy change in the gas is the familiar $k \ln 2$.

That is our expectation. It is confounded by fluctuations. Consider the piston first. It is a thermal system that is Boltzmann distributed over its height $h \geq 0$ above the piston floor according to

$$p(h) = (Mg/kT) \exp(-Mgh/kT)$$

where M is the piston mass. The mean of this distribution is kT/Mg and its standard deviation is also kT/Mg .

This latter number measures the extent of thermal fluctuations in the height of the piston. For a macroscopic piston, M will be very much larger than kT/g and the extent of fluctuations in height will be negligible. However in this case of a one-molecule gas, the piston must be very light if it is to be suspended at equilibrium by the pressure of the one-molecule gas. Hence its M is small and the fluctuations in height will be great. They can be estimated quantitatively as follows. The weight of the piston is Mg . The mean force exerted by the gas pressure is $(kT/V)A = kT/h$, where A is the area of the piston and h its height above the base of the cylinder, so that $V = Ah$. Setting these two forces equal as the condition for equilibrium, we recover the equilibrium height as⁶

$$h_{eq} = kT/Mg$$

Remarkably, this quantity h_{eq} is just the same as the mean height and standard deviation of the above distribution, both of which are also given by kT/Mg .

This extraordinary result can be expressed more picturesquely as follows. If we set up the piston so that its weight perfectly balances the mean pressure force of the one-molecule gas, it will not remain at the equilibrium height, but will fluctuate immediately through the entire volume of the gas. It will perhaps be suddenly flung skyward by a collision with molecule; and it may then fall precipitously between collisions. The intended process of a gentle, indefinitely slow expansion or contraction is lost completely behind the wild gyrations of the piston over the full volume of the one-molecule gas.

⁶ Hence the mean energy of height is $Mgh_{eq} = kT$. While this energy is associated with a single degree of freedom of the moving piston, it differs from the familiar equipartition mean energy per degree of freedom $(1/2)kT$, because the relevant term of the piston's Hamiltonian, Mgh , is linear in h and not quadratic, as the equipartition theorem assumes.

Similar results hold for heat transfer between the one-molecule gas and its environment. Since it is monatomic, the Boltzmann distribution of the gas energy E is

$$p(E) = 2(E/\pi)^{1/2} (kT)^{-3/2} \exp(-E/kT)$$

The mean of this distribution is the familiar equipartition energy $(3/2) kT$ and the standard deviation is $(3/2)^{1/2} kT = 1.225 kT$.⁷ Hence, simply by virtue of its contact with the environment at temperature T , the one-molecule gas energy will be swinging wildly through a range comparable in size to the total mean energy of the gas.

We had expected that we would track a quantity of heat $kT \ln 2 = 0.69 kT$ while the piston slowly and gently moves to halve or double the volume of the gas. What we find is that the piston is wildly and randomly flung to and fro through the entire volume of the gas, while the gas energy fluctuates similarly wildly over a range greater than the $0.69 kT$ of heat transfer we track. We had expected a process that proceeds calmly at arbitrarily slow speed from start to finish. Instead we find a chaos of wild gyrations with no discernible start or finish.

This is a rough analysis. To maintain the equilibrium of a thermodynamically reversible process would require that the weight Mg be adjusted as the volume V changes since the gas pressure will vary inversely with volume. Norton (2011, Section 7.5) replaces the uniform force field of gravity with another force field that varies with height in precisely the way needed to maintain mean quantities at equilibrium.

3.2 Generality

A one-molecule gas confined in a cylinder by a piston is fanciful and cannot be realized practically. It is, however, one of the most discussed examples in the thermodynamics of computation because it is easy to visualize. Its statistical and thermodynamic properties mimic those of more realistic systems with few degrees of freedom. We may model a memory device as a two-chambered cell with a single molecule trapped in one part. A more realistic

⁷ This and the earlier energy standard deviation can be computed most rapidly from Einstein's energy fluctuation theorem, which identifies the variance of the energy with $kT^2 d\langle E \rangle / dT$, where $\langle E \rangle$ is the mean energy. For the piston, $\langle E \rangle = kT$, so the variance is $(kT)^2 = (Mgh_{eq})^2$. For the monatomic gas, $\langle E \rangle = (3/2)kT$, so the variance is $(3/2)(kT)^2$. The standard deviation is the square root of the variance.

implementation of the memory device is a single electric charge trapped by a potential well in a solid state medium; or a magnetic dipole aligned into a specific orientation by a magnetic field.

The thermodynamic operations carried out on the one-molecule gas have analogs in the more realistic implementations. Mechanical variables such as volume and pressure are replaced by electric and magnetic correlates. The general results remain the same. If we halve the range of possible states of a memory device, we reduce its thermodynamic entropy by $k \ln 2$, just as we do when we halve the volume of a one-molecule gas. The large fluctuations exhibited by the one-molecule gas derive from its small number of degrees of freedom. Correspondingly, the more realistic implementations will exhibit similarly large fluctuations.

The two processes investigated were heating/cooling and expansion/contraction of the gas. These are instances of the two processes that appear in all thermodynamically reversible processes: heat transfer and exchange of generalized work energy. As a result, the analysis here has a quite broad scope. Consider thermodynamically reversible measurement, in which one device reads the state of another. For example, a magnetic dipole reads the state of a second dipole when the two slowly approach and align in a process that maintains equilibrium throughout. This detection or measurement process is a reversible compression of the phase space of the reader dipole and is thermodynamically analogous to compression of a one-molecule gas. As a result, this measurement process will be fatally disrupted by fluctuations. While a standard claim of the thermodynamics literature is that these measurements can be performed without dissipation, the no go result shows that dissipation is required if the fluctuations are to be overcome and the process driven to a correct reading.

References

- Bennett, Charles. 1982. "The Thermodynamics of Computation—A Review." *International Journal of Theoretical Physics*, **21**, 905-40; reprinted in Leff and Rex, 2003, Ch. 7.1.
- Earman, John and Norton, John D. 1998. "Exorcist XIV: The Wrath of Maxwell's Demon." Part I "From Maxwell to Szilard." *Studies in the History and Philosophy of Modern Physics*, **29**(1998), 435-471.
- Earman, John and Norton, John D. 1999. "Exorcist XIV: The Wrath of Maxwell's Demon." Part II: "From Szilard to Landauer and Beyond," *Studies in the History and Philosophy of Modern Physics*, **30**(1999), 1-40.

- Leff, Harvey S. and Rex, Andrew, eds. 2003. *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing*. Bristol and Philadelphia: Institute of Physics Publishing.
- Norton, John D. 2005. "Eaters of the lotus: Landauer's principle and the return of Maxwell's demon." *Studies in the History and Philosophy of Modern Physics*, **36**, 375–411.
- Norton, John D. 2011. "Waiting for Landauer." *Studies in History and Philosophy of Modern Physics*, **42**, 184–198.
- Norton, John D. Forthcoming. "Infinite Idealizations," Prepared for *Vienna Circle Institute Yearbook*. Springer: Dordrecht-Heidelberg-London-New York.
- Tolman, Richard C. 1938. *The Principles of Statistical Mechanics*. London: Oxford University Press.