Available online at www.sciencedirect.com
ScienceDirect
Journal of Theoretical Biology 252 (2008) 213–220
Journal of Theoretical Biology
www.elsevier.com/locate/yjtbi

Thermodynamics of natural selection III: Landauer's principle in computation and chemistry

Eric Smith*

Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

Received 28 June 2007; received in revised form 4 February 2008; accepted 7 February 2008
Available online 16 February 2008

Abstract

This is the third in a series of three papers devoted to energy flow and entropy changes in chemical and biological processes, and their relations to the thermodynamics of computation. The previous two papers have developed reversible chemical transformations as idealizations for studying physiology and natural selection, and derived bounds from the second law of thermodynamics, between information gain in an ensemble and the chemical work required to produce it. This paper concerns the explicit mapping of chemistry to computation, and particularly the Landauer decomposition of irreversible computations, in which reversible logical operations generating no heat are separated from heat-generating erasure steps which are logically irreversible but thermodynamically reversible. The Landauer arrangement of computation is shown to produce the same entropy-flow diagram as that of the chemical Carnot cycles used in the second paper of the series to idealize physiological cycles. The specific application of computation to data compression and error-correcting encoding also makes possible a Landauer analysis of the somewhat different problem of optimal molecular recognition, which has been considered as an information theory problem. It is shown here that bounds on maximum sequence discrimination from the enthalpy of complex formation, although derived from the same logical model as the Shannon theorem for channel capacity, arise from exactly the opposite model for erasure.

Keywords: Carnot cycle; Landauer's principle; Thermodynamics; Natural selection; Molecular recognition

1. Introduction

Two previous papers (Smith, 2008a, b) have developed a general relation between energy flow and entropy transfer in chemical and biological processes. Paper I (Smith, 2008a) argued that separation of timescales and isolation of subsystems by catalysts makes many aspects of metabolism and natural selection engine-like, so that bounds on their performance can sensibly be computed from idealized limits of reversible transformation. Paper II (Smith, 2008b) then showed that the chemical equivalents of Carnot cycles (Fermi, 1956) provide a natural basis into which to decompose the kinds of composite reversible transformations which arise in metabolism. These cycles define precise measures of chemical work and relate them both to internal energy and to entropies of formation and distribution.

The purpose for the current paper is to relate the foregoing energy/entropy constraints from chemistry to the classical theory and thermodynamics of computation. It is plausible that such a relation should exist because Shannon's theory of communication (Shannon and Weaver, 1949) identifies entropy reductions as measures of information. Indeed, all of statistical mechanics (including biochemistry) may be interpreted as an information-based approach to the study of incompletely determined systems (Jaynes, 2003). Computation, in turn, is the process by which information is "produced", in a manner to be made precise below. Thus it is tempting to think of natural selection, via metabolism, as an essentially computational process that instills the information particular to biomatter in the course of incorporating material from the growth medium into living structures. Anticipating this connection, the relation between energy and entropy was

*Tel.: +1 505 946 2764.
E-mail address: desmith@santafe.edu

0022-5193/$ - see front matter
doi:10.1016/j.jtbi.2008.02.013

proposed in the first two papers as a relation between *work* W and *information* $\mathscr{I}$

$$\mathrm{d}W = k_B T\,\mathrm{d}\mathscr{I}, \quad (1)$$

in which k_B is Boltzmann's constant and T is the temperature of the heat bath with which the processes of interest are in thermal equilibrium.

Eq. (1) is the particular expression of the second law of thermodynamics that characterizes the energetic requirements of computation (Landauer, 1961; Bennett, 1982). A key assumption in the computation literature is that the degrees of freedom carrying data are physically identical, in which case $\mathscr{I}$ represents a reduction in entropy as input data are converted to output "answers" by the computer. (The main result of Paper II (Eq. (22)) was that if the degrees of freedom are not physically identical, the form of Eq. (1) can be kept if $\mathscr{I}$ is a Kullback–Leibler divergence (Cover and Thomas, 1991) of the distribution of probabilities from a Gibbs equilibrium. Alternatively, $\mathscr{I}$ may be defined as an entropy change if enthalpies of formation are added to externally supplied work to define W.)

Somewhat more than the re-naming of energies and entropies in the second law is required, however, to motivate referring to natural selection as a computational process. The two distinctive features of computation are that a computer can exchange entropy between data streams and a heat bath when both share a single temperature (unlike a classical heat engine), and that an ideally efficient computer combines *logical irreversibility* with *thermodynamic reversibility*. Paper II demonstrated that chemistry shares the first property, by virtue of one or more pairs of dual state variables (μ, N), added to the minimal (p, V, T) state variables of classical heat engines. Here N is the number of some species of particle and μ its chemical potential, and p, V, and T are pressure, volume, and temperature, respectively. This paper will show that the essential state diagram of "chemical Carnot cycles" (Fig. 3 of Paper II) is also the state diagram corresponding to the classical Landauer (1961) decomposition of computation.

The value of a computational interpretation of chemistry comes from the way the Landauer decomposition separates logical operations which can be performed without generating any heat (Keyes and Landauer, 1970; Bennett, 1973, 1982) from processes of *erasure* needed to reclaim computational resources, which inherently convert data entropy and work to heat. The careful use of this decomposition makes it possible not only to associate specific transformations of a chemical Carnot cycle with natural steps in computation, but also to unravel some of the complex relations between communication-based notions of information, and diverse interpretations of Eq. (1) that arise in biology.

In the original reliable-communication approach to information (Shannon and Weaver, 1949; Cover and Thomas, 1991), thermodynamics was not directly of interest, but the problem of error-correcting encoding was, for it determined the asymmetry between the roles of signal and noise. Encoding and decoding are computational processes, and by decomposing them into distinct logic and erasure stages, energetic costs may be assigned to the handling of signal or noise in different communication problems. One biological phenomenon that has been treated explicitly with the communications concept of *channel capacity* is optimal molecular recognition (Schneider 1991a, b, 1997). It was shown in Paper II that this problem is consistent with Eq. (1) if the enthalpy of "priming" the recognition system, rather than an external source of work, defines the $\mathrm{d}W$ responsible for the entropy reduction that occurs when proteins bound to random sequences migrate reliably to their specific target sequences. This paper will further refine that result, showing that the application of Shannon's theorem to define a "machine capacity" for recognition (Schneider, 1991a) arises from a novel computational model of decoding. Whereas in the original problem of reliable telephone communication, decoders were assumed to erase channel noise and output the decoded signal, in molecular recognition the binding protein "erases" the entropy corresponding to the *signal* in Shannon's formula for channel capacity, leaving the noise which is un-erasable because it arises from thermal fluctuations. The clarity afforded by the Landauer decomposition of reversible decoding allows us to identify the source of the energy cost of recognition (Schneider, 1991b), consistent both with Eq. (1) and with the derivation of the Shannon limit on the entropy of recognition in Schneider (1991a).

The value of using chemistry as an instance of computation is to clarify the role of ensembles in the joint representation of data and heat degrees of freedom. Rissanen (1989) has correctly argued that any consistent assignment of thermodynamic cost to an algorithm must be computed from an ensemble in which data and heat degrees of freedom are jointly stochastic, but objections to this claim remain in the literature (Earman and Jorton, 1998, 1999). In chemistry there is no temptation to regard the information in variables (μ, N) as privileged with respect to that in "thermal" variables (p, V, T). The demonstration that the "Landauer cycle", as we have termed it (Smith, 2008b) for chemical work, is also the cycle for computation will show why logical irreversibility does not imply thermodynamic irreversibility when the heat bath as well as the computer is made part of the state space.

The remainder of the paper is organized as follows: Section 2 reviews the standard Landauer decomposition of logically irreversible (but thermodynamically reversible) computation, and makes precise the sense in which a computation "produces" information as it converts an input data stream to an output data stream. Here it is demonstrated that the chemical version of the Landauer cycle (Smith, 2008b) is also the computational version corresponding to the Rissanen ensemble. Section 3 applies the Landauer decomposition to the problems of data compression and error-correcting encoding and decoding, making contact between thermodynamics and Shannon's channel capacity expression. Here the problem of molecular recognition is analyzed, demonstrating the consistency of the computational interpretation with both this domain-specific process and the very different concerns of general physiology and selection. The computational and second-law properties of such "self-powering" decoders as binding proteins is further clarified by mapping sequence selection to the thermodynamically similar process of first-order phase transition. Finally, results are summarized in Section 4.

2. Logic and thermodynamics of computation

The idea that *an instance of* computation carried out on physical hardware must obey the second law of thermodynamics is uncontroversial. Landauer's (1961) principle is the stronger assertion that *an algorithm for* computation has an intrinsic energy cost, related by the second law to the information it produces. The stronger assertion is needed to overcome the problem of Maxwell demons (Brillouin, 2004), which have been of particular interest as models for the chemical processes that instill information in biomatter.

Since universal computation can be constructed from a small number of Boolean logical operations (Hopcroft and Ullman, 1979), the essential elements in the Landauer argument are that the logic of these operations can be implemented on reversible hardware (Bennett, 1973), and that any reduction in the number of output variables compared to input variables for an irreversible logic gate can be achieved by some standard model of erasure. Both reversible logic and standard erasure may be obtained in many ways. The convenient assumption to make for "standard" models of erasure is that the bits to be erased are first maximally compressed, so that the heat rejected upon erasure may be computed directly from the number of bits erased. In this case the second law implies that the rejected heat is independent of the model of erasure used. Since maximally compressed bits are by definition independent and identically distributed (IID), they also do not depend on the logic gate from which they arise. Therefore it is sufficient to demonstrate the Landauer decomposition for any logical operation, once it has been shown that reversible gates exist to implement a complete Boolean algebra. This section will review the standard Landauer decomposition, taking reversible logic as given and using the Szilard (1929) single-particle gas as a model for erasure.

2.1. The Landauer decomposition for elementary operations

The key to capturing the Rissanen (1989) ensemble view, and to making contact with the chemical Landauer cycle, is to include data input and output streams as ensembles, and to be explicit about the requirements for reversible exchange of data between these streams and the computer itself. The data streams and the computational algorithm will first be represented in general form here, and then reduced to the particular case of "exclusive-or" (XOR) to provide notation and graphical representation for an example.

Represent the input state of a computer with a random variable X, whose values x make up a statistical ensemble. Similarly, let the output state correspond to a random variable Y whose values y make up another ensemble. Regard successive samples of X or Y as coming from input and output *data streams* which act as reservoirs. Since the Landauer decomposition works by separating reversible logic from idealized erasure, we must include the coupling between the data streams and computer in this decomposition. The simplest assumption is that the transfer of data from the stream to the computer input or output state is performed reversibly—corresponding to the thermodynamic treatment of adiabatic coupling to reservoirs in Paper II—which then requires that the inputs start and outputs end in canonical (not random) states. If the output y is a deterministic function $y = f(x)$, the ensemble entropy of y cannot exceed that of x (Cover and Thomas, 1991). A computer that takes data repeatedly from a stream of instances of X and produces a stream of instances Y thus converts the joint data stream $Y \times X$ at any time to a joint data stream of equal or lower entropy with each successive act of computation. It is this entropy reduction that we will relate to an obligatory heat flow.

It is convenient to illustrate the subsequent computation for a logical operation that directly maps most-compressed inputs to most-compressed outputs, after which the generalization to all of Boolean logic follows simply by the elaboration of a small number of cases. Suppose therefore that X comprises two binary digits $X \equiv (X_1, X_2)$, each drawn IID, with probabilities $(0.5, 0.5)$ to take values $(0, 1)$. The computation on X is XOR, denoted $y = x_1 \oplus x_2$, if the output y takes value 0 when the two digits X_1 and X_2 are the same, and value 1 when they are different. The task is to separate XOR into (reversible) transformation and (irreversible) erasure stages. The reversible stage takes the two input values (x_1, x_2) and canonical states $(0, 0)$ of two output registers, and transforms the output to (y, x_1) while transforming the input to canonical values $(0, 0)$. This is the step that can be done without work input or heat production, in a variety of models (Bennett, 1982). However, doing so requires definite states of the initial Y and final X registers, which is the same constraint as that required for reversible exchange of variables with the data streams.

Only the data value y will be transferred to the output data stream, so in order to restore the computer to a state capable of processing another input, the value x_1 in the second position must be erased. As output x_1 is also IID $\sim(0.5, 0.5)$, by the assumed distribution of inputs, it is already a maximally compressed variable, meaning that erasure will remove the maximal possible information per reset event. One standard model for erasure uses Szilard's single-particle ideal gas (Szilard, 1929). A particle on the left-hand side of a two-chamber box is taken to represent value 0 and a particle in the right-hand side to represent value 1. Erasure takes place by removal of a barrier between the state-0 and state-1 sides of the box, followed by compression with a piston to reset the particle in canonical state 0 while the box is in contact with a thermal reservoir, after which the wall is re-inserted. The entropy that must be rejected to the reservoir (in thermodynamic units) is $k_B \log 2$, where log stands for the natural logarithm (in logical units, ''1 bit'') for an IID Boolean variable, and if the reservoir is at temperature T the energy required from the piston is $k_B T \log 2$. Energy lost to the reservoir is heat because the reservoir degrees of freedom are not explicitly controlled by the computer apparatus, while energy input from the piston is work because the piston position is a (mechanical) boundary parameter. The input–output map and essential steps in erasure are illustrated in Fig. 1.

2.2. The distinct informations

When the possibility of observing various data values is added to this description of a computational step, multiple distinct entropy reductions can be considered, which admit the interpretation of various measures of information. The model provides a precise distinction of their meanings, which readily generalizes to more complex computations and also maps onto the description of natural selection.

The entropy of the computer's ensemble of states is 2 bits—equal to the entropy of the input data—in all stages of computation prior to the compression marked ''reset'' in Fig. 1. If the value of the input or output registers (whichever currently holds the data) is observed at any stage prior to the step labeled ''erase'', the entropy of the resulting distribution is reduced to zero, a gain of 2 bits of information about the complete input data value. Note that the removal of the wall in the transition labeled ''erase'' in Fig. 1 does not reduce the actual entropy of the computer's state, but it does remove a prescription for how we could measure the output registers so as to determine the input data value. In this sense it is not Jaynes's (1957, 2003) emphasis on what one ''knows'' that determines the entropy, but rather what one *controls* physically.[1]

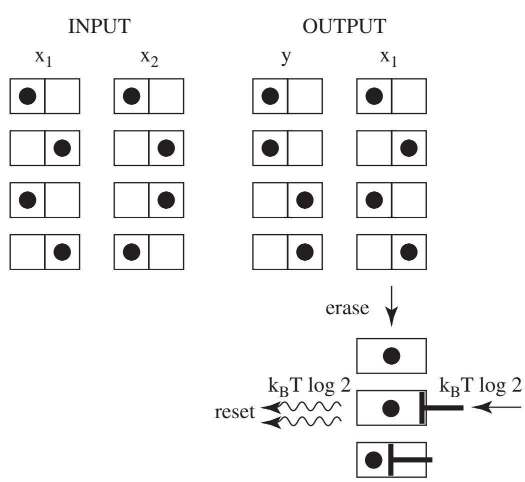

Fig. 1. States and stages in a computer that computes XOR. Two input variables (X_1, X_2) are represented in the leftmost two columns as Szilard boxes, with the particle in the left-hand chamber of a box representing state 0 and in the right-hand chamber representing state 1. Four equally probable input values are arrayed vertically. Two output variables are represented in the rightmost two columns, after the reversible transformation from inputs to outputs. Not only XOR, but the identity function on X_1, is computed at this stage. To compute only XOR, we reset the second output variable, by removing the partition and then using a piston to compress the particle into a canonical 0 state, after which the divider is replaced (not shown).

The ''reset'' transition does change the ensemble of the computer's states, leaving the x_1-valued bit of the output in a canonical state 0. The entropy of the computer state is thereby reduced by $k_B \log 2$ (1 bit), while heat $k_B T \log 2$ is rejected to the reservoir. This entropy reduction therefore provides the definition of $\mathrm{d}\mathscr{I}$ in Eq. (1), and because the bits have been assumed identical, the only work $\mathrm{d}W$ is the mechanical work drawn from the piston. (The more general case, in which an enthalpy of reaction can be added to the formal work variable to define the effective $\mathrm{d}W$, is developed in Paper II, and will be revisited in the computational analysis of molecular recognition below.) The Landauer claim is that while the Szilard single-atom gas realizes the bound (1) from the second law of thermodynamics, the bound is itself a thermodynamic identity binding on any model of logic with erasure.

The computer state after ''reset'' has entropy 1 bit, and the only possible observation remaining is of Y. Observing Y again reduces the entropy of the computer state to zero, yielding 1 bit of information. This ''information in the output of the computation'' equals the original uncertainty about the parity of the input X not removed by XOR, which is neutral with respect to parity.

It is the data entropy converted to heat by Landauer's principle—$S(X) - S(Y)$ (also equal to 1 bit in the example)—which corresponds to the entropy reduction when natural selection acts to create a biased distribution from a uniform equilibrium distribution. It can be understood as the entropy of all input degrees of freedom with respect to which the computation of XOR is *not neutral*—in other words, XOR converts $x_1 = 1 \rightarrow 0$ while leaving $x_1 = 0$ unchanged. To the extent that the calculation XOR is defined by the degrees of freedom on which it is not neutral, the entropy reduction $S(X) - S(Y)$ could be described as information ''about'' which calculation is performed. The entropy in the output of the computation is related to the input entropy and the heat Q rejected to the reservoir by

$$S(X) = S(Y) + Q/T, \tag{2}$$

in cases like these that saturate the Landauer bound.

[1]The following is pedantic, but has emerged as a confusion in reviews of earlier versions of this paper, and therefore seems necessary: ''Knowledge'' of data values, as it is being used for analysis in this section, corresponds to real reductions in entropy only if it is realized by partitioning the ensembles from which data are drawn. These partitionings are not part of the computer's function, so the reset is the only computational function resulting in entropy transfer.

2.3. Logical versus thermal reversibility

Objections to this kind of statistical treatment of the data stream together with the heat bath (Earman and Jorton, 1998, 1999) arise because particular data values are of interest from a computational point of view, while those of the bath are not. Ordinarily in the modeling of a Hamiltonian dynamical system, a choice is made either to treat all degrees of freedom microscopically and reversibly (even for systems where statistical mechanics would normally be used, D'Souza and Margolus, 1999), or to treat the values of all non-state variables as stochastic. Analysis of the physical substrate of deterministic computation happens to be a domain that falls in between, where determinism of the computation is a property of a subset of the microscopic degrees of freedom, while the average work and heat properties of an optimal computer depend only on a notion of thermodynamic reversibility, in which entropy is preserved through a transformation but the mapping of microscopic configurations need not be.

The erasure/reset step in computation is *logically* irreversible, but *thermodynamically* reversible. The physical operations of the computer during this pair of transformations can certainly be reversed: we could add energy and entropy from the heat bath to a canonically set 1-bit variable while extracting the piston to remove the energy as work, and then insert a wall to obtain a standard randomized variable with one bit of entropy. Combining this introduced variable with Y and inverting the reversible XOR, we would obtain a data stream with the same entropy as X, though not the same sequence of data values.

For conventional thermodynamic examples, including the above analysis of metabolism, the failure of microscopic reversibility is unremarkable. The fact that deterministic computation is also not reversible should not obscure the fact that the path in the space of state variables, shown in Fig. 2, is essentially the same as that of Fig. 3 of Paper II (Smith, 2008b).

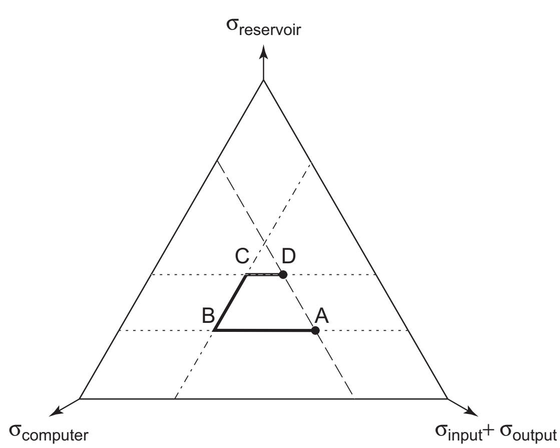

Fig. 2. The canonical Landauer path corresponding to the computation model of Fig. 3 of Paper II (Smith, 2008b). Arc AB increases entropy in the computer by accepting data from the input, and arc CD sends data to the output. Only the erasure/reset arc BC is canonically assumed to be performed in contact with the thermal reservoir. There is no arc corresponding to DA' of Fig. 3 of Paper II, because we have ignored possible conditions for the computer to be in "equilibrium" with the different input and output data streams. On real hardware (such as a chemical system) the Landauer cycle of Paper II may be more realistic.

3. Error correction, molecular recognition, and phase transition

Eq. (2) identifies the information that a computational algorithm imposes on its inputs X with the subset of the entropy in X that the algorithm recognizes and rejects as heat Q. The residual entropy in the output Y is potential information about X not imposed by the algorithm, or equivalently, entropy not recognized and eliminated by any of the algorithm's erasure steps. Both data compression and error correction are instances of imposing information that fall within this representation of computation. Two applications of error correction, opposite in the roles played by erasure, are Shannon's original problem of reliable signal transmission (Shannon and Weaver, 1949; Cover and Thomas, 1991), and Schneider's (1991a, b) theory of optimal molecular recognition. Relating erasure in these processes to the classical thermodynamics of phase transition will show how the "self-powering cycles" considered in Paper II, Section 2.3 can be realized.

3.1. Compression and error-correcting codes

Data are compressible if the bits on which the input X is carried are capable of greater variation than exists in the actual input ensemble. In the task of maximal lossless compression, the actual variation in X defines the length of the desired output representation Y, and the possible but unrealized variation in X counts the number of distinct *encodings* by which the compressed output Y could be represented in X. An algorithm corresponding to the correct encoding can reversibly decode X into Y, or encode Y as X. Thus lossless compression is an instance of reversible computation.

Lossy compression occurs when the entropy of X exceeds that of the output Y under a decoding algorithm, so that some bits of the input are regarded as *noise*, which may vary freely with respect to the output bits in Y. In the classical problem of reliable signal transmission (Shannon and Weaver, 1949; Cover and Thomas, 1991), an input is first encoded with some redundancy, then mixed with noise during transmission, and finally projected by a many-to-one map (an *irreversible* decoder) to recover an approximation to the input.

The Landauer decomposition to illustrate the energetics of the classical channel is shown in Fig. 3. The encoding of input x onto an expanded representation $c(x)$ with the same entropy may be performed by a reversible computer. Transmission mixes noise z with $c(x)$ (with "mixing" defined by the channel model) to produce strings $c(x) \oplus z$ with entropy $S(x) + S(z)$, if $c(x)$ and z are independent. For finite code blocks and realistic noise models, a given output

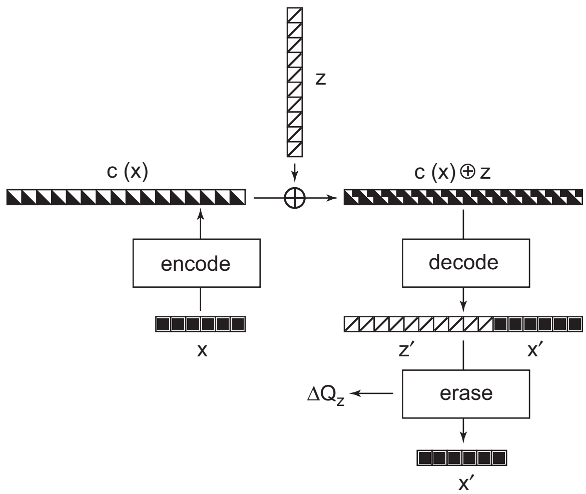

Fig. 3. Computation diagram for conventional channel error correction. x is a random variable holding a string of maximally compressed input bits (1-bit entropy per bit signal, indicated by filled shading). $c(x)$ is the data expanded with an error-correcting code, representing the same entropy with more bits (partly shaded). z is a compressed representation of environmental noise, mixed ($\oplus$ in the diagram) by the transmission channel with the encoded signal to yield a string $c(x) \oplus z$ which, under optimal encoding, now also contains 1 bit of entropy per (expanded) signal bit. Decoding may be implemented as a reversible logical operation, whose function is to rearrange the received bits so that estimates x' of the original data bits occupy the leading places, yielding by implication a representation z' of the inferred channel noise in the remaining bits. If encoding is optimal, the representation z' is itself maximally compressed. For conventional channel encoding, the noise bits are not of interest and may be reset by an ideal erasure with rejection of heat $\Delta Q_z = k_B T S_z$ to the environment, where S_z is the entropy in units of the natural logarithm (nats) of the noise ensemble z'. In the long-signal limit with optimal encoding, x' differs from the input x, hence z' from z, only with vanishing probability.

$c(x) \oplus z$ can be produced in more than one way, but it is possible to define a canonical decomposition $c(x') \oplus z'$ as the one most likely given the characteristics of x and z. The usual irreversible decoder may then be taken as composed of a reversible computer—which produces the estimate x' for the signal, and by implication an estimate z' for the noise—and a classical eraser of the bits carrying z'.

For a discrete, memoryless channel, *Shannon's theorem* shows (Cover and Thomas, 1991, p. 184) that the information capacity C of the channel, per symbol, is the maximum over encodings c of the mutual information $I(c(x); c(x) \oplus z)$ between the symbols in the channel input and output. In the particular case of the Gaussian channel, which is a good model for many thermodynamic processes, the capacity per symbol in units of the natural logarithm (nats) is given by (Cover and Thomas, 1991, p. 241)

$$C = \frac{1}{2} \log\left(\frac{P+N}{N}\right). \tag{3}$$

Here P and N were originally used to denote power levels of electrical signals, but for the general Gaussian channel they are whatever measures of variance determine the range of permitted input symbols in $c(x)$ and z, respectively.

The geometric interpretation of the Gaussian channel is that $\sqrt{P+N}$ is the radius of a "ball" in the space of signals $\{c(x) \oplus z\}$, while $\sqrt{N}$ is the radius of the smaller ball created by the noise ensemble $\{z\}$ about any single codeword $c(x)$. The ratio of the two radii is the number (per symbol) of distinct codewords that can be placed in the symbol space by an optimal code, without overlapping. The distinctive feature of the Gaussian channel as a model for the Shannon theorem is that the noise degrees of freedom can be erased in a manner that is neutral with respect to the signal set $\{x\}$. Such erasure rejects heat $\Delta Q_z = k_B T S_z$, where S_z is the entropy of the inferred noise ensemble $\{z'\}$. The entropy of $\{z'\}$ may equal the entropy of $\{z\}$ even for unreliable transmission, but for an optimal code in the limit of infinite code block size, the logical assignments of z' also converge to those of z.

3.2. Molecular recognition

Without changing the transmission or encoding models, *but reversing the role of erasure*, Schneider (1991a, b) has applied the same theory of optimal encoding to the problem of sequence selection during molecular recognition. In this problem, the non-covalent binding interface between a macromolecule (for definiteness, suppose it is a protein) and its substrate implements the mixing function of a channel. The "data" x are drawn from the ensemble of protein-sequence complexes created by the initial binding of the protein to a random location on the substrate. z is phase space randomness of the degrees of freedom in the protein-sequence complex created by thermal noise. The binding interface is a Gaussian channel, which adds energies from binding affinity to energies from thermal activation. Schneider's observation is that by exploiting coordinated fluctuation in the position and momentum coordinates of multiple non-covalent binding degrees of freedom, a protein can implement an encoding $c(x)$ that distinguishes more sequences than could be distinguished by simply summing the binding affinities independently. In this model, P and N represent energies, rather than powers, with P derived from the "priming" energy of the protein and $N = k_B T/2$ per degree of freedom arising from thermal fluctuations.

For the recognition problem, the thermal "noise" cannot be erased because it is in equilibrium with the surrounding temperature bath. Rather, the objective is to remove the entropy of the "data", which corresponds to projecting an ensemble of random initial binding sites into an ensemble in which the protein is bound only to the desired target sequence. The channel capacity per Hamiltonian degree of freedom[2] remains that of Eq. (3) (Schneider, 1991a), but the enthalpy dissipated as heat in the course of recognition

[2]Meaning that position and momentum degrees of freedom for a single dimension of spatial motion are counted independently.

is now the priming energy $P \rightarrow \Delta Q_x = k_B T S_x$, and S_x is the entropy of the largest ensemble $\{x\}$ of sequences from which a unique target site can be selected by a given protein (Schneider, 1991b).

3.3. Erasure and phase transition

The association of Landauer's erasure with the relaxation of a protein as it converges on a target sequence is physically natural. Erasure is the resetting of bits from arbitrary to fixed values. Convergence of a (primed) bound protein to the target sequence results from the successive restriction of physical degrees of freedom as the priming energy is lost and the resulting accessible bound sequences more and more closely approximate the target sequence.[3]

The way in which, during molecular recognition, enthalpy can be used to reject the entropy of the "signal" in Shannon's theorem, but not the entropy of the "noise", extends to a universal property of first-order phase transitions. In a first-order phase transition, a material moved along its coexistence curve from the fully disordered to the fully ordered phase dissipates a *latent heat* (Kittel and Kroemer, 1980), for which the energy is drawn from enthalpy of the disordered configuration. Transformation along the coexistence curve corresponds to the case of optimal molecular recognition, in that both are thermodynamically reversible; if they are performed in a single direction, it is assumed that an outside entity decides that direction.

The latent heat on the coexistence curve can of course be released in an irreversible transformation that differs arbitrarily little from full passage across the curve, if the material is prepared in a metastable disordered state and the boundary conditions are then taken into a regime which favors the ordered state. The enthalpy of the metastable primed state is maintained, in the language of Schneider (1991a), by "frustration", which corresponds to the statistical unlikelihood in a phase transition of spontaneous molecular alignment consistent with the low-energy *order parameter* (Kittel and Kroemer, 1980). Note that, while it is acceptable to refer to the free energy reduction of the primed protein in the recognition problem, it is not correct in the limit of optimal recognition to refer to the whole system as irreversible. In that limit, the entropy of the protein-sequence ensemble must be included as an addition to the protein-internal entropy, and the resulting free energy is neutral with respect to the transition, as the changes from energy and entropy exchanges with the heat bath balance.

Both first-order phase transitions and molecular recognizers are self-powering chemical entropy rejectors, in the language of Paper II, Section 2.3. Their algorithms correspond to the identification of an order parameter or a target sequence, respectively. Their heat is generated entirely from enthalpy without requirement for external work. In some cases, as in that of recognition, we may recognize the Gibbs free energy of a subsystem as being reduced by the transformation, but the intuition gained from this representation only becomes applicable if the transformation is at least weakly irreversible. Optimal recognition is an alternative to the cyclic processes underlying the reversible evolution studied in Paper II, which permits study of the chemistry/computation relation in a non-cyclic context.

[3]A similar map could be made for protein folding, wherein a succession of local folds is accompanied by rejection of heat as the protein approaches the targeted complete fold.

4. Summary and conclusions

This sequence of three papers has attempted to decompose an elementary constraint (1) from the second law of thermodynamics into the kinds of chemical and computational processes by which it is realized. The fundamental difference between processes involving chemical work, and the classical analysis of heat engines, is that chemical work can be used to reduce the entropy of a system when the system and its environment share a single temperature. While the classical analysis which decomposes reversible thermal transformations into a basis of Carnot cycles may still be used, the cycles require the addition of an intrinsically "chemical" pair of dual intensive/extensive state variables, in order to take non-trivial forms. It is with the introduction of this new pair, which plays the thermodynamic role of a particle number N and its dual chemical potential μ, that reversible transformations first become eligible for interpretation as fixed-temperature computations.

The problem of bounding the physical requirements of computation has been given a general-form analysis by the Landauer decomposition, whose consequences have been the main topic of this paper. This decomposition recognizes that the *logic* of computation is not itself irreversible; only the processes of erasure that reclaim computer resources for re-use are logically irreversible, and in the limit of ideally efficient computation even these remain thermodynamically reversible. The realization that chemical ensembles provide a definition of information independent of particular logical operations, and that the chemical Carnot cycle (Fig. 3 of Paper II) is also the model for idealized cyclic computation (Fig. 2), adds further intuition to an argument of Rissanen (1989) for ensemble treatment of data: The algorithm for a general-purpose computer is only independent of its data if it is implemented without prior constraint on the data instances to which it is to be applied. Therefore, any energetic properties intrinsic to the algorithm must be *defined* from ensembles of data, and there is no reason of principle to distinguish this ensemble from the one containing thermal degrees of freedom which are uncontrolled for the same reason. The naïve paradox of conjoint logical irreversibility and thermal reversibility appears also in the problem of molecular recognition, where sequence identification is irreversible if expressed in terms of the free energy of the relaxing protein, but reversible when the ensemble entropy of binding sites is also counted.

With these associations in place, it is possible to say precisely in what respect chemical processes in metabolism or natural selection correspond to computation, and in some cases even to relate these to particular applications of computation such as data compression or error correction. An advantage of the Landauer separation of logic from erasure is that different combinations of these may be understood in common terms. Thus, the thermodynamics of Shannon's reliable communication differ as completely as possible from the thermodynamics of optimal molecular recognition (the two processes erase the opposite entropies in the expression for channel capacity), yet both are consistent with the second law, and they share an expression for capacity because they are based on the same computational logic.

Acknowledgment

I thank Insight Venture Partners for support of this work, and grant FIBR: The Emergence of Life: From Geochemistry to the Genetic Code Award 0526747, National Science Foundation for support of a collaboration within which it was performed. Thanks to Harold Morowitz and Shelley Copley for perspectives on metabolism and catalysis, and to Duncan Foley and David Krakauer for discussion and references about Landauer's principle. I especially thank David and an anonymous referee, for support of this work and for very helpful comments and criticisms.

References

Bennett, C.H., 1973. Logical reversibility of computation. IBM J. Res. Dev. 17, 525–532.

Bennett, C.H., 1982. The thermodynamics of computation—a review. Int. J. Theor. Phys. 21, 905–940.

Brillouin, L., 2004. Science and Information Theory, second ed. Dover Phoenix Editions, Mineola, NY.

Cover, T.M., Thomas, J.A., 1991. Elements of Information Theory. Wiley, New York.

D'Souza, R.M., Margolus, N.H., 1999. Thermodynamically reversible generalization of diffusion limited aggregation. Phys. Rev. E 60, 264–274.

Earman, J., Jorton, J.D., 1998. Exorcist XIV: the Wrath of Maxwell's Deamon, part I: from Maxwell to Szilard. Stud. Hist. Philos. Mod. Phys. 29, 435–471.

Earman, J., Jorton, J.D., 1999. Exorcist XIV: the Wrath of Maxwell's Deamon, part II: from Szilard to Landauer and beyond. Stud. Hist. Philos. Mod. Phys. 30, 1–40.

Fermi, E., 1956. Thermodynamics. Dover, New York.

Hopcroft, J.E., Ullman, J.D., 1979. Introduction to Automata Theory, Languages, and Computation. Addison-Wesley, Reading, MA.

Jaynes, E.T., 1957. Information theory and statistical mechanics. Phys. Rev. 106, 620–630 reprinted in Rosenkrantz, R.D. (Ed.). Jaynes, E.T.: Papers on Probability, Statistics and Statistical Physics. D. Reidel, Dordrecht, Holland, 1983.

Jaynes, E.T., 2003. Probability Theory: The Logic of Science. Cambridge University of Press, New York.

Keyes, R.W., Landauer, R., 1970. Minimal energy dissipation in logic. IBM J. Res. Dev. 14, 152–157.

Kittel, C., Kroemer, H., 1980. Thermal Physics, second ed. Freeman, New York.

Landauer, R., 1961. Irreversibility and heat generation in the computing process. IBM J. Res. Dev. 3, 183–191.

Rissanen, J., 1989. Stochastic Complexity in Statistical Inquiry. World Scientific, Teaneck, NJ.

Schneider, T.D., 1991a. Theory of molecular machines I: channel capacity of molecular machines. J. Theor. Biol. 148, 83–123.

Schneider, T.D., 1991b. Theory of molecular machines II: energy dissipation from molecular machines. J. Theor. Biol. 148, 125–137.

Schneider, T.D., 1997. Information content of individual genetic sequences. J. Theor. Biol. 189, 427–441.

Shannon, C.E., Weaver, W., 1949. The Mathematical Theory of Communication. University of Illinois Press, Urbana, IL.

Smith, E., 2008a. Thermodynamics of natural selection I: Energy flow and the limits on the organization. J. Theor. Biol., in press, doi:10.1016/j.jtbi.2008.02.010.

Smith, E., 2008b. Thermodynamics of natural selection II: Chemical Carnot cycles. J. Theor. Biol., in press, doi:10.1016/j.jtbi.2008.02.008.

Szilard, L., 1929. Uber die entropieverminderung in einem thermodynamischen system bei eingriffen intelligenter wesen. Zeit. Phys. 53, 840–856.