Entropy Cost of Information

Paul N. Fahn *
Department of Computer Science and
Information Systems Laboratory
Stanford University
Stanford, CA 94305

Abstract

An entropy analysis of Szilard's one-molecule Maxwell's demon suggests a general theory of the entropy cost of information. The entropy of the demon increases due to the decoupling of the molecule from the measurement information. In general, neither measurement nor erasure is fundamentally a thermodynamically costly operation; however, the decorrelation of the system from the information must always increase entropy in the system-with-information. This causes a net entropy increase in the universe unless, as in the Szilard demon, the information is used to decrease entropy elsewhere before the correlation is lost. Thus information is thermodynamically costly precisely to the extent that it is not used to obtain work from the measured system.

1 Introduction

James Clerk Maxwell introduced his famous demon in 1871 as a thought experiment to explore the second law of thermodynamics, which asserts the nondecrease of entropy. In the 123 years since its introduction, the study of Maxwell's demon has come to include questions of the thermodynamic significance of measurement and information. In particular, the one-molecule demon proposed by Szilard in 1929 [5], brought information into the heart of the debate, where it has remained ever since. Analyses by Brillouin [3] and Bennett [2] have used the demon to explore the thermodynamic costs of information processing; they believed measurement and erasure, respectively, to be thermodynamically costly events in the sense that they cause entropy increases. This paper uses a new entropy analysis of the Szilard demon as a starting point for a theory which sees decorrelation between information and system as the key entropy-producing event.

*fahn@cs.stanford.edu

2 The Szilard Engine

The following description of the one-molecule Maxwell's demon is Bennett's [2] slight modification of Szilard's original presentation [5]. We consider further modifications in Section 5.1. The terms "engine" and "demon" are used interchangeably.

The engine has two main parts: a cylinder and a memory device. The cylinder contains one molecule and has movable pistons at either end; the pistons can move inwards but not outwards past their initial positions. The cylinder is in thermal equilibrium with its environment at temperature T. The memory device has three (macroscopic) states: '0', 'L', and 'R'. The exact nature of the memory is immaterial; we will only assume (for now) that each macroscopic state of the memory corresponds to an equal number of microscopic physical states of the memory device.

0. At the start of the process, the pistons are pulled out so that the volume of the cylinder is at its maximum. The memory is in state '0'.

The steps of the engine's operation are as follows; each step is, by assumption, performed without any inefficiency. A diagram of the engine after each step, for one particular cycle of operation, is shown in the left column of Figure 2.

A. Insertion of Partition. A thin partition is quickly inserted into the cylinder, dividing it into left and right halves of equal volume.

B. Measurement. A measurement is made to determine whether the molecule is trapped in the left or right half of the cylinder, and the memory is set to state 'L' or state 'R' accordingly. No assumptions about the measurement method are made.

C. Compression. The piston on the side of the cylinder not containing the molecule is pushed in until it touches the partition.

D. Removal of Partition. The partition is quickly removed.

E. Expansion. The one-molecule gas expands by

0-8186-6715-X/94 $4.00

colliding with the piston that has been pushed in at the compression step. The energy lost by the molecule is replaced by heat from the environment at temperature T. The expansion stops when the original volume of the cylinder is attained.

F. Erasure. The memory is set to state '0'. Thus the measurement information, either 'L' or 'R', is erased.

At the end of the erasure step, the demon's cycle of operation is complete. The cylinder is at the same volume, the memory is in the same state, and the molecule is at the same temperature, as at the start.

3 Entropy

How should the entropy of a system be determined? We insist on the following properties.

1. All the standard thermodynamic identities involving entropy should hold.

2. Entropy should be a function of the thermodynamic state of the system. Thus the total entropy change of the engine over the complete cycle must be zero, since the engine is in the same thermodynamic state at the end as at the start.

3. Entropy should not depend on whether someone possesses information about the system, or the kind or amount of information possessed. In other words, entropy must be completely "objective" and should neither depend on nor refer to what an observer may know or not know about the system in question.

The following formulation of entropy meets all the above criteria and will be used to analyze the operation of the Szilard Engine.

Let R be a region of the phase space which characterizes a physical system in equilibrium. The entropy of this region R is proportional to the natural logarithm of the volume of R:

$$S = k \log V(R) \tag{1}$$

where the constant of proportion k is Boltzmann's constant. The volume function $V(R)$ should be thought of as counting the states in R; these states are uniformly distributed in phase space.

Now consider an experimental procedure to prepare a thermodynamic system. After the experiment is performed a single time, the system will be in some microscopic state, corresponding to a single point in phase space (no measurement of the microscopic state is implied). Other points in phase space will correspond to the state of the system on subsequent repetitions of the experiment. Eventually, after many repetitions of the procedure, the corresponding points in phase space will uniformly fill a region R of phase space; then $S = k \log V(R)$ is the entropy of the thermodynamic system.

When analyzing the entropy of the demon at various points in its cycle, we must think of the experiment being repeated numerous times, and not restrict our attention to a single execution of the operation cycle. In particular, it must be kept in mind that the molecule will be trapped in both halves of the cylinder, rather than in just a single half, as would be the case if we considered just a single execution of the demon's cycle. The latter view of entropy does not satisfy the desired properties listed above and leads to results that contradict the second law.

For example, what is the entropy of the molecule after the partition is inserted compared to before the insertion? As we repeat the experiment over and over, we note that since the molecule is trapped sometimes in the left and sometimes in the right half of the cylinder, all positions within the entire cylinder are possible states for the molecule after the partition has been inserted, just as they were before the partition was inserted. Thus the entropy has not changed (the partition is considered to be arbitrarily thin). If we had, on the other hand, considered only a single cycle execution, we would say that the molecule occupies a region of phase space half as large as the region occupied before insertion, and that the entropy decreased. But this would be incorrect: it would imply that we could decrease entropy simply by sliding the partition in, and would constitute a violation of the second law.

4 Entropy Analysis of the Demon

We now proceed to examine the entropy increases and decreases that occur at each step of the demon's operation.

The state of the demon can be specified by two parameters, the horizontal position of the molecule and the state of the memory. The other state parameters will have no effect on the entropy and thus do not need to be explicitly considered in the phase space. For example, the velocity distribution of the molecule never changes, since it is always at temperature T; velocity thus plays no role in the entropy changes. Therefore it suffices to use a two-dimensional rather than a seven-dimensional diagram of phase space.

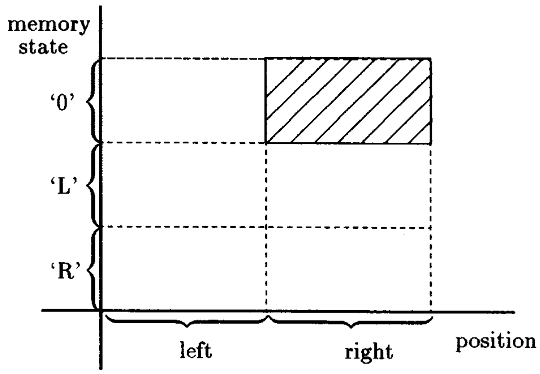

Figure 1: Phase Space of the Szilard Engine

Figure 1 shows the phase space of the demon. The horizontal axis marks positions of the molecule; the vertical axis marks microscopic states of the memory. At any instant, the microscopic state of the engine is represented by a point in this diagram. A macroscopic state of the engine, specifying only the macroscopic state of the memory ('0','L', or 'R') and which half of the cylinder contains the molecule, is a set of microscopic states of the engine and corresponds to a rectangular region in the phase space diagram. For example, the shaded region in Figure 1 corresponds to the memory in state '0' while the molecule is in the right half of the cylinder.

We proceed to examine the entropy of the engine in equilibrium after each step; the corresponding region in phase space is shown in the right column of Figure 2.

0. Initial configuration. Before the process begins, the molecule can occupy any position within the cylinder and the memory is in state '0'. Let ν denote the volume of the corresponding region of phase space. Then the entropy is $S_0 = k \log \nu$.

1. After insertion of the partition. The molecule can still occupy any position (as we consider all the available states over many repetitions of the demon's cycle) and the memory is still in state '0'. Thus the entropy is unchanged and $S_1 = S_0$.

2. After the measurement. The molecule may occupy either the left or the right half, and the memory may be in either state 'L' or state 'R'. However, these configurations are correlated so that only two macroscopic states occur: 1) the molecule is in the left half of the cylinder and the memory reads 'L', or 2) the molecule is in the right half and the memory reads 'R'. The other combinations cannot occur. Looking at the phase space diagram, it is clear that the system now occupies a different region of phase space than it did just before the measurement, but the volume of that region is unchanged. Thus the entropy does not change due to the measurement. $S_2 = S_0$.

3. After the compression. Pushing in the cylinder has no effect on either the position of the molecule or the state of the memory. Thus the macroscopic state (region in phase space) is unchanged, and so is entropy. $S_3 = S_0$.

4. After removal of the partition. Removal of the partition has no effect on either the position of the molecule or the state of the memory; entropy is unchanged. $S_4 = S_0$.

5. After the expansion of the gas. After the gas has expanded to fill the entire volume, the molecule can have any position and the memory can be in either state 'L' or state 'R' *in any combination*, unlike the case in phases 2 through 4. For example, the molecule can be in the left half of the cylinder while the memory is in state 'R'; this is a configuration that was not available previously. Thus the number of available states for the demon has increased, i.e., the entropy of the demon has increased. Looking at the phase space diagram, it is clear that the volume in phase space has exactly doubled:

$$S_5 = k \log 2\nu = k \log \nu + k \log 2 = S_0 + k \log 2$$

and the entropy change is exactly

$$\Delta S_{\text{demon}} = k \log 2 = 1 \text{ bit.} \qquad (2)$$

Another entropy change occurs at this step as well. As the gas expands, it loses energy, which is replaced by an influx of heat from the demon's environment at temperature T, thereby decreasing the entropy of the environment. The expansion step under consideration is a case of isothermal expansion, one of the basic processes in thermodynamics (it is, for example, one of the steps of the Carnot engine). As any textbook of thermodynamics shows, the entropy change is $dS = dQ/T$, where dQ represents the energy lost by the heat reservoir. The amount of entropy gained by the expanding system is the same as that lost by the reservoir. From Equation 2, we can then conclude that the amount of heat transferred is $kT \log 2$ and the entropy change in the demon's environment is

$$\Delta S_{\text{env}} = -k \log 2 = -1 \text{ bit.}$$

Note that the total entropy change in the universe due to the expansion is

$$\Delta S_{\text{expand}} = S_{\text{demon}} + S_{\text{env}} = k \log 2 - k \log 2 = 0. \quad (3)$$

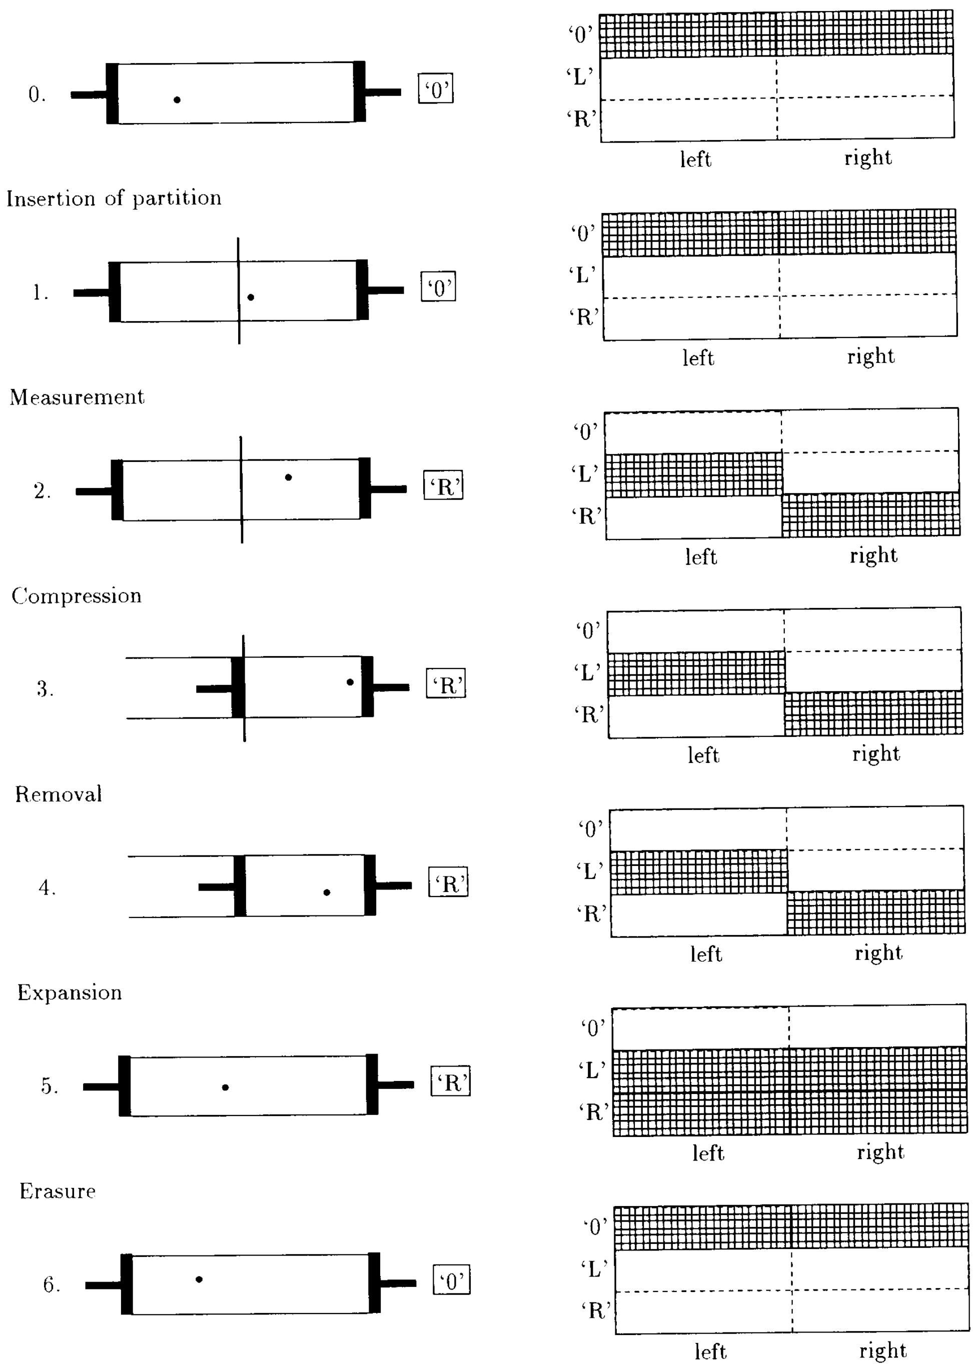

Figure 2: The Szilard Engine. The diagram shows the state of the demon and the phase space after each step. The phase space refers to many repetitions of the process, not to the specific example shown on the left.

6. After the erasure of the memory. Erasure of the memory does not affect the possible position of the molecule. But now the memory can be in only one state, state '0'. A glance at the phase space diagram makes it clear that the demon occupies the same region of phase space that it did at the start, and, as we expected, $S_6 = S_0 = k \log \nu$. But this volume is half of the phase space volume at phase 5, i.e., $S_6 = S_5 - k \log 2$ and the entropy of the demon has decreased by $\Delta S_{\text{demon}} = -1$ bit.

The entropy of the demon's environment also changes due to the memory erasure. Landauer [4] and Bennett [1] have argued that the environment incurs an increase of $\Delta S_{\text{env}} = k \log 2 = 1$ bit; they use the second law to infer this entropy increase. In the current case, however, one is analyzing a Maxwell's demon at least partly in order to decide whether or not the second law is valid; therefore, it would be inappropriate to assume its validity within the analysis itself. However, the same conclusion is easily reached by assuming instead that the microscopic laws of physics are time-reversible (in the sense of invertible functions, not in the sense of reversible thermodynamic processes). Then the degrees of freedom of any closed system must be conserved, which implies that the degrees of freedom of the demon's environment must increase when the degrees of freedom of the demon itself decrease due to the erasure. (It is possible that the alternative assumptions of the second law and time-reversible microscopic physical laws are in fact equivalent.)

The total entropy change in the universe due to the erasure is

$$\Delta S_{\text{erase}} = S_{\text{demon}} + S_{\text{env}} = -k \log 2 + k \log 2 = 0. \quad (4)$$

Erasure entails a transfer of entropy from the demon to the environment.

The entropy changes for the entire cycle of operation are summarized in Table 1.

Note that if we had considered the entropies of the molecule and the memory separately, and then added them together, we would have gotten erroneous results. The separate entropies can be obtained by considering the volumes of the support sets along the horizontal and vertical axes in the phase space diagrams. But this method would not take into account the correlations between the molecule position and the memory, and as a result would not reveal the entropy increase that occurs when more states are made available to the combined system during the expansion step.

5 Discussion

Looking at Table 1, we first note that entropy changes only occur during the expansion and erasure steps, and not during the measurement step (keep in mind that we have assumed errorless measurement).

Second, we note that the sum of the entropy changes in each row and in both columns is zero. That is, there is no net change in entropy at any step, and neither the demon nor its environment suffer a net change over the course of the operation cycle; happily, the second law is not violated. It thus appears at first glance that no step can justifiably be called thermodynamically costly in the context of the Szilard Engine. However, by looking at the operations in more general contexts, we conclude below that decorrelation between information and system, and not erasure or measurement, is a thermodynamically costly event in a fundamental sense. We note that Bennett, although using a similar phase space for the demon [1], reaches very different conclusions about which operations are "thermodynamically costly".

In the model of a tri-state memory that we have described, erasure of 1 bit of memory decreases the entropy of the memory by $k \log 2$ (1 bit) and increases the entropy of the memory's surroundings, or environment by the same amount; this is true regardless of whether the erasure is part of some cyclic process or if it is simply an isolated operation. Both the entropy decrease and the entropy increase are due to the erasure, which, in effect, causes a transfer of entropy from the memory to its environment. Furthermore, even these entropy changes are not truly fundamental to the act of erasing 1 bit of memory. Depending on how the memory is implemented, the entropy changes due to erasure may vary both in magnitude and sign, or may not occur at all; this is discussed further in Section 5.1. It would thus appear that erasure, when considered in full generality, should not be called "thermodynamically costly".

How about the expansion step of the demon? In the case of the Szilard Engine, this step also entails merely a transfer of entropy, rather than a net increase. But, considering the expansion step more carefully, the real cause of the entropy increase in the demon is the evolution of the system in such a way that the correlation between the state of the memory and the position of the molecule is lost. The correlation restricts which states of the molecule can occur with which states of the memory. The expansion eliminates the correlation and lifts the restriction: the expansion allows the demon to occupy states, such as memory 'R' and molecule position left, that previously had not been

step	demon	environment	total net change
A. Insertion	0	0	0
B. Measurement	0	0	0
C. Compression	0	0	0
D. Removal	0	0	0
E. Expansion	+1	−1	0
F. Erasure	−1	+1	0
total	0	0	0

Table 1: Entropy changes at each step of the Szilard Engine (demon), in bits. 1 bit $= k \log 2$.

accessible. None of this involves the work and the heat transfer that occur at the same time as the expansion. That is, the work and the heat transfer are not intrinsic to the entropy increase of the demon; the same entropy increase would have occurred, for the same reasons, without any work being done or any heat transferred. Thus the expansion step is just a specific case of a more general "decorrelation" operation. Unlike erasure, only an entropy increase — and not a decrease — is fundamental to a decorrelation operation. For this reason, it seems justified to refer to decorrelation as a thermodynamically costly operation.

5.1 Modified Szilard Engines

These points are further illustrated by considering modified versions of Szilard's Engine.

First we slightly change our model for the demon's memory: instead of the reference state '0' having the same size (in phase space) as the 'L' and 'R' states, suppose that the '0' state is twice as large. This is a natural model for the memory as it would occur, for example, if we implemented the memory using two bits of standard semiconductor memory found in today's computers; one can think of the '00' and '11' states mapping to 'L' and 'R' respectively and the '01' and '10' states both mapping to our reference state '0'.

The operation of the engine proceeds as before, and is illustrated in Figure 3. The entropy of the engine now decreases by $k \log 2$ (1 bit) during the measurement step, since the measurement reduces the size of the available phase space. At the same time, there is an offsetting increase in the entropy of the environment, for the same reasons that an increase occurred during erasure in the original memory model. The entropy changes during the decorrelation (expansion) step are the same as before. But now there is no entropy change at all during the erasure step; the volume of phase space is the same before as after the erasure, and thus the erasure can be performed reversibly, with no consequent entropy changes either within the demon or in its environment. The entropy changes for this demon are presented in Table 2.

One can in fact generalize this argument by considering the volume A in phase space of the '0' memory state to be an arbitrary parameter of the Szilard Engine. It is not too hard to see that as A varies, the entropy changes in the measurement and erasure steps will vary in tandem so that the sum of the changes in those two steps is constant. That is, the measurement and erasure steps together share a net entropy change, with each step's portion determined by the parameter A. Thus we see that *the entropy changes due to measurement and erasure vary in an arbitrary way and therefore are not fundamental; the entropy change due to decorrelation does not vary arbitrarily and therefore is fundamental.*

Another modification to the demon shows that decorrelation steps can create net entropy increases in the universe, rather than just entropy transfers. Consider the following modification to the Szilard Engine (see Figure 4). The cylinder is as before, except that there are no pistons at the ends. The partition is inserted, the measurement is made (and the memory set accordingly), then the partition is removed, and finally the memory is erased. This is a cyclic process and is basically identical to the Szilard Engine except that the removal and expansion steps have been combined so that the gas does no work and the entropy of the environment is not decreased at any step (no heat transfer takes place).

The entropy analysis for this demon is similar to the previous cases; the entropy changes are shown in Table 3. As always, entropy is measured when the system comes to equilibrium after each step. First note that

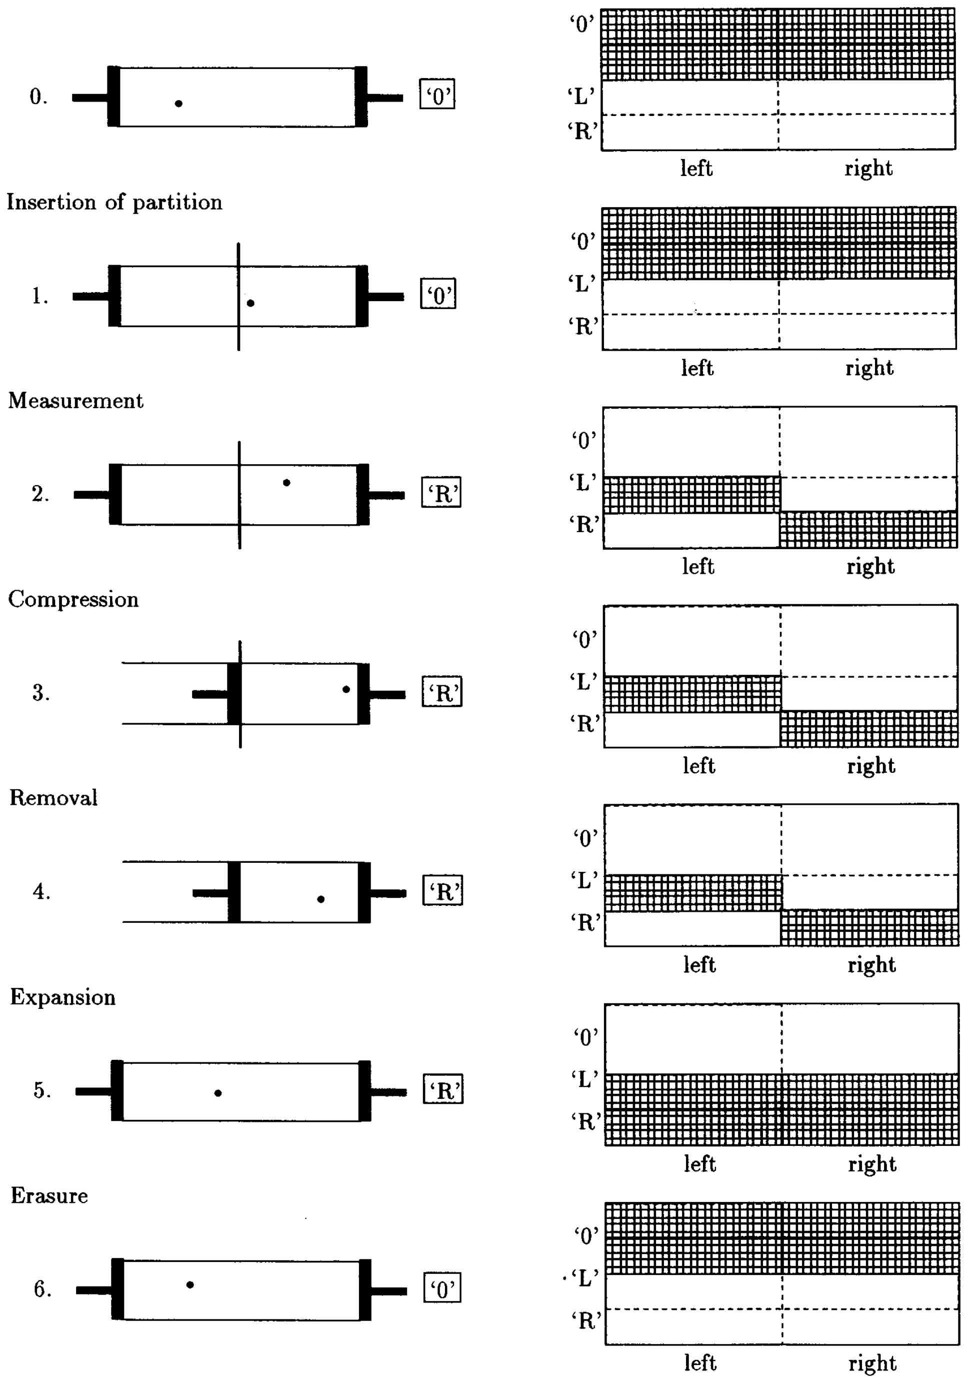

Figure 3: Szilard Engine with memory implemented using two elementary bits. The entropy changes occurring at the measurement and erasure steps have changed but the changes at the expansion steps are the same.

step	demon	environment	total net change
A. Insertion	0	0	0
B. Measurement	−1	+1	0
C. Compression	0	0	0
D. Removal	0	0	0
E. Expansion	+1	−1	0
F. Erasure	0	0	0
total	0	0	0

Table 2: Entropy changes at each step of the Szilard Engine with modified memory implementation, in bits.

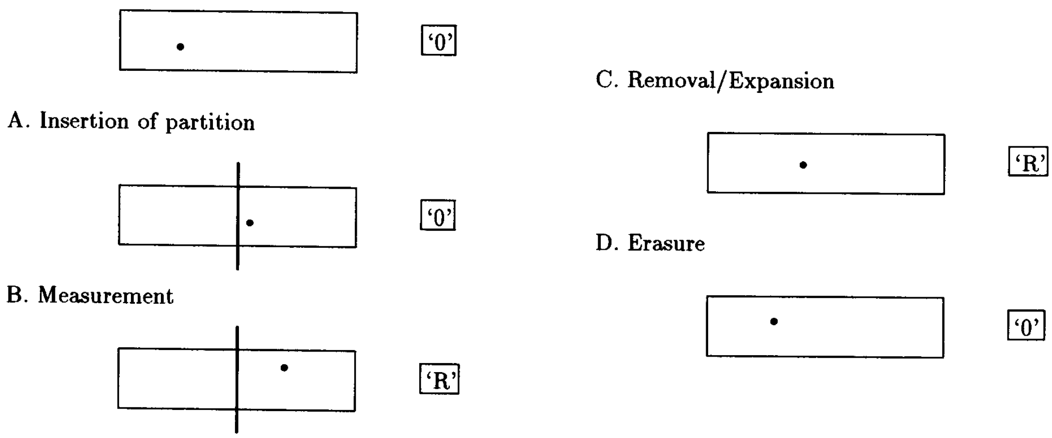

Figure 4: Another version of the Szilard Engine, without any work done. Unlike the original engine, a net increase in entropy results from the operation of this version.

not all the rows and columns add to zero. The column showing the entropy changes in the environment shows an increase of 1 bit of entropy; that is, running this demon for one cycle of operation increases the entropy of the universe by 1 bit.

The row for the removal/expansion step, which we can now call the decorrelation step, has a net entropy increase of 1 bit. The net entropy changes due to the other operations is still zero. In other words, the decorrelation causes a net entropy increase in the universe, unlike all the other steps in this process, including measurement and erasure. Once again, it seems justified to refer to decorrelation as a thermodynamically costly operation.

It is instructive to consider a further modification: eliminating the memory altogether. The partition is inserted, and then removed. Clearly there is no entropy change at all in this cyclic process. We conclude that the decorrelation which causes an increase of entropy is inherently an information-sensitive event: without the information, there's no decorrelation and there's no entropy increase.

In the special case of the Szilard Engine, the decorrelation step did not cause a net increase in entropy. Why not? Because in the same step, the gas did work on its environment, causing an entropy decrease that offset the increase due to the decorrelation. The information embodied in the state of the memory was used to obtain work from the gas; the information could be used in this way only while it was still correlated with the state of the gas.

6 Theory of Entropy and Information

The above results suggest a more general theory about the entropy costs associated with information. We summarize the results described above in terms

step	demon	environment	total net change
A. Insertion	0	0	0
B. Measurement	0	0	0
C. Removal/Expansion	+1	0	+1
D. Erasure	−1	+1	0
total	0	+1	+1

Table 3: Entropy changes for the demon that does no work.

of such a theory, without reference to any Maxwell's demon.

Suppose there is a physical system and a physical memory device whose state can record the results of measurements made on the system. To avoid unpleasant paradoxes, entropy must be calculated for the joint system comprising both the memory and the physical system; we will call this joint system the metasystem, to avoid confusion.

The results of a measurement are recorded in a memory; information can be thought of as the record of a measurement. The measurement creates a correlation between the memory and the physical system. An entropy change may occur in the metasystem due to the measurement, of sign and magnitude determined by the specifics of the case. Of course, an entropy decrease will be accompanied by an entropy increase of equal or greater magnitude in the metasystem's environment.

Decorrelation occurs when the physical system evolves in such a way that its state is no longer correlated with the information, i.e. with the state of the memory. That is, the physical system reaches an equilibrium in which its state and the state of the memory are independent. There may be processes where decorrelation never occurs; this would be the case if the information were sufficient to predict the exact state of the physical system at all times subsequent to the measurement.

Decorrelation causes an increase of entropy in the metasystem. The amount of increase depends on the details of the process and system, but an entropy increase must always occur, since decorrelation always increases the number of accessible states. In some cases, the information can be used to induce the physical system to do work on its environment, which may cause a decrease in the environment's entropy. By the second law of thermodynamics, the amount of this entropy decrease cannot exceed the amount of the entropy increase due to the decorrelation. This suggests the following principle:

> *The entropy cost of information is the degree to which the information is not used to obtain work from the system. This cost is assessed when the correlation between information and system is lost.*

The term 'work' is used generically to include any entropy-reducing process.

Erasure of information, like measurement, may produce either an increase or a decrease in the entropy of the memory; if there is a decrease, there will also be an offsetting increase in the entropy of the memory's environment.

If the measurement is noisy or subject to error, then the entropy of the resulting state will be larger than if there had been no error. This is equivalent to the case of an errorless measurement followed by a step of partial decorrelation, which causes an entropy increase; this partial decorrelation is combined into the measurement.

Of course, these theories are not complete and call for further research. They also raise other questions regarding information, measurement and entropy in a wider variety of systems and processes. For example, one would like an explicit analysis of measurement in a quantum system and in other cases where the measurement information and the physical system are related in a more complicated way than in the case of the Szilard Engine. One would also like a more comprehensive model of erasure, a theory for the energy required to erase, and a closer examination of the idea that erasure entails a loss of information.

Acknowledgements

Many thanks to Tom Cover, Don Kimber, and Juan Parrondo for interesting and fruitful discussions.

References

[1] Charles H. Bennett. The thermodynamics of computation — a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.

[2] Charles H. Bennett. Demons, engines and the second law. *Scientific American*, 257:108–116, 1987.

[3] Leon Brillouin. *Science and Information Theory*. Academic Press, New York, 2nd edition, 1962.

[4] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.

[5] Leo Szilard. On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. *Behavioral Science*, 9:301–310, 1964. Originally published in *Zeitschrift für Physik*, 53:840–856, 1929.