Journal of Statistical Physics, Vol. 62, Nos. 3/4, 1991

Algorithmic Treatment of the Spin-Echo Effect

Seth Lloyd[1] and Wojciech H. Zurek[2]

Received April 17, 1990; final October 19, 1990

We analyze the apparent increase in entropy in the course of the spin-echo effect using algorithmic information theory. We show that although the state of the spins quickly becomes algorithmically complex, then simple again during the echo, the overall complexity of spins together with the magnetic field grows slowly, as the logarithm of the elapsed time. This slow increase in complexity is reflected in an increased difficulty in taking advantage of the echo pulse. Our discussion illustrates the fundamental role of algorithmic information content in the formulation of statistical physics, including the second law of thermodynamics, from the viewpoint of the observer.

KEY WORDS: Spin-echo effect; algorithmic complexity; second law of thermodynamics.

1. INTRODUCTION

1.1. The Spin-Echo Effect

According to the second law of thermodynamics, disorder, once created, is almost impossible to destroy. The spin-echo effect[1,2] appears to contradict this dictum. In this effect, a large number of spins, initially aligned, precess in an inhomogeneous magnetic field until they are pointing every which way: the spins start out ordered and become disordered. If the spins are then subjected to a radiofrequency pulse of appropriately chosen frequency and duration, their subsequent evolution "undoes" the disorder, causing the spins to come back into alignment. In this paper we employ both conventional thermodynamics and algorithmic information theory to address

[1] California Institute of Technology, 452-48, Pasadena, California 91125.
[2] Los Alamos National Laboratory, Los Alamos, New Mexico 87545.

0022-4715/91/0200-0819$06.50/0

the question of just how much disorder is created in the spin-echo effect, and how effectively that disorder is banished. Use of the *physical entropy* given by the sum of the usual statistical Boltzmann–Gibbs–Shannon entropy and the algorithmic randomness (given by the size of the shortest description of the known aspects of the physical state of the system) has been recently proposed by one of us as essential to the consistent formulation of thermodynamics from the viewpoint of the observer.[3] This paper applies such ideas to a simple, experimentally accessible physical system.

The spin-echo effect was first demonstrated experimentally by Hahn.[1] We first describe the effect classically here, in its most easily visualizable form. We then give a complete quantum electrodynamic treatment to show that interactions between the spins and a quantized electromagnetic field leave the original "paradox"—the apparent contradiction with the second law—basically unaffected.

A macroscopic sample of matter such as glycerin is subjected to a strong, stable magnetic field B, pointing along the z axis. If $\mu B \gg k_B T$, where μ is the magnetic moment of the nuclei in the sample, then at equilibrium all of the spins except for a fraction $\sim e^{-\mu B/k_B T}$ will be pointing up along the z axis. The spins are now subjected to a pulse from a radiofrequency field B_1 directed along the x axis and oscillating at the spins' average Larmor frequency $\omega = 2\mu B/\hbar$. The pulse has length $\pi/2\omega_1$, where $\omega_1 = \mu B_1/\hbar$, which is just the amount of time required for the Larmor precession around the field of the pulse to rotate all the spins to a common orientation perpendicular both to the z axis and to the direction of the rf pulse in a corotating frame of reference. The spins now precess freely about the z axis.

Initially, all the spins are aligned, but local inhomogeneities in the magnetic field cause the spins to precess at slightly different rates. If the inhomogeneities are a fraction ζ of the total field B, then after $1/\zeta$ revolutions the spins will be pointing in all different directions. Inhomogeneities of 0.1% will cause the spins to become completely unaligned after 1000 revolutions, for example. As long as the coupling between different spins and between spins and the lattice is weak, this disorder is almost entirely due to the differing rates of Larmor precession of the spins.

Now, at any instant t between the start of the spins' precession at time zero and the spin–spin or spin–lattice relaxation time, a second rf pulse can be applied. This pulse is exactly twice the length of the first pulse, and causes the spins to precess by an angle π about the rotating magnetic field of the pulse. (In fact, as shown quantum mechanically in the next section, any length pulse will result in some spin echo: in the treatment here, we give the pulse lengths that give the most easily visualizable classical picture and the maximum spin echo.) In the frame that corotates with the pulse, the effect of this precession is to take the ith spin from an angle ϕ_i to an angle $-\phi_i$. That is, the second pulse conjugates the total phase that each spin has accumulated in the course of its Larmor precession.

The spins continue precessing, each in its own slightly different magnetic field. But now the precession of each spin *undoes* the phase that it accumulated before its phase was conjugated. At time t after the second pulse, the spins are all lined up again along a direction 180° away from their original orientation in the corotating frame. When the spins realign, they induce a signal in the rf coil—the *echo* of the signal that originally aligned the spins.

If the pulse that conjugates the phase of the spins arrives after the spin–spin or spin–lattice relaxation time, then the spins will fail to realign: in order to give an echo, the evolution of the spins following their conjugation must "undo" their entire previous evolution, including all significant interactions. No straightforward way is known to reverse the interaction between the spins and the lattice in which they sit, but a carefully selected sequence of rf pulses suffices to reverse not only the spins' evolution, but their interaction with the other spins as well.[2]

1.2. Entropy and Information

We begin with complete knowledge of the state of the spins. As the spins precess, we lose all knowledge about their individual orientations: the inhomogeneities in the field quickly cause the spins to reach a state of what appears to be maximum entropy. But when the second rf pulse is applied, the spins just as quickly return to a state of essentially zero entropy. Is the second law of thermodynamics being violated?

The answer is, of course, No. We now show how and why the second law is preserved. One possibility is that the entropy accumulated by the spins in the course of the evolution is somehow "dumped" into the quantum state of the electromagnetic pulse that effects the reversal. If such a transfer of entropy took place, one could maintain that the second law applies to the whole system (i.e., spins and field), but does not have to apply to the subsystems. In the full quantum electrodynamic treatment in the next section, we show that this is not the case: the entropy of the spins is not transferred to the field. The pulse acts essentially as a "mirror," reflecting the phases of the spins without recording them, rather than as a "photograph" in which an image of the phases is recorded.

A second possibility is that the entropy of the spins is not increasing: the spins are not actually disordered in an absolute sense, but only appear to be so. To judge this possibility, we need an intrinsic measure of disorder. The algorithmic complexity of the spins—the length of the shortest algorithm that can reproduce the spins' configuration—is just such a measure. In Section 4, we analyze the algorithmic complexity of the spins and show that the spins are in fact completely disordered at the time of reversal. The spins' algorithmic complexity does increase significantly. This conclusion could have been anticipated simply by pointing out that the collection of spins rotating at incommensurable frequencies is an ergodic dynamical system. Hence, in the course of its Poincaré cycle it must spend most of its time traversing "typical" (and, therefore, presumably "random") spin configurations which overwhelmingly contribute to its equilibrium entropy. Indeed, we shall show using conventional statistical mechanics that the spins' entropy increases significantly.

The third possibility, and correct answer, is that although the entropy of the spins does in fact first increase, then decrease in the course of the spin echo, the entropy of the total system—spins plus lattice plus magnetic field—remains almost constant for times significantly less than the spin–spin and spin–lattice relaxation times. The algorithmic complexity of the entire system evolves only slowly, increasing on average by the logarithm of total time elapsed. We are accustomed to think of entropy as an extensive quantity, so that the entropy of spins and lattice and electromagnetic field is equal to the entropy of the spins *plus* the entropy of the lattice *plus* the entropy of the electromagnetic field. For a system with strong correlations between its parts, however, entropy does not add up in this fashion.

The spins and magnetic field make up just such a highly correlated system: to know the orientation of a particular spin, one need only know the value of the magnetic field at the spin's site, together with the amount of time that the spin has been precessing since the first rf pulse oriented it along the y axis in the corotating frame. When such correlations exist, the entropy of the whole is significantly less than the sum of the entropies of the parts, and it is possible in principle to reduce the entropy of one component without increasing entropy elsewhere. The spin-echo apparatus uses the mutual information between magnetic field and spins to reduce the entropy of the spins without decreasing the entropy of field and spins combined.

The reader may protest that as a gas expands, collisions introduce correlations between the velocities of different molecules, and yet the entropy of these molecules is extensive (barring quantum symmetry and excluded-volume effects). Indeed, if one could follow Boltzmann's (probably apocryphal) instructions to Loschmidt concerning the velocities of the molecules and "reverse them," then the correlations between the molecules would allow them to return to a prior state of low entropy. The difference between the spins of the spin echo and the molecules of the gas is that for the spins there is a simply specifiable procedure to "reverse them," even in the presence of significant spin–spin interactions, whereas for the molecules in the gas, no such simply specifiable procedure exists. One could in principle put a mirror simultaneously in front of each molecule, thus reversing their velocities; but to specify the position and orientation of all the mirrors requires a vast amount of detailed information, which can in turn be translated into a *thermodynamic cost of reversal.*

This cost can be calculated by computing the accuracy to which the "mirrors" that reflect the particles have to be positioned in order to achieve a successful reversal[4]: after N collisions, to return the molecules of the hard-sphere gas to a volume in phase space V requires that the mirrors be positioned to within a volume $V/2^N$ in their own position–angle phase space. Given a cost of information storage of $k_B T \ln 2$ per bit (as developed in detail in ref. 4, and as explained in Section 3 of the present paper), one is led to associate a thermodynamic cost of $Nk_B T \ln 2$ with the actual implementation of the reversal.[3] This relatively large cost of reversal arises from the fact that the evolution of the hard-sphere gas amplifies errors exponentially in time. In contrast, the spin-echo system has a discrete spectrum, retains quantum mechanical coherence throughout its evolution, and hence cannot exhibit sensitivity to small changes in initial conditions. The lack of sensitivity to initial conditions in the spin-echo system is a necessary (but not sufficient) condition for the straightforward reversal of the spins' evolution.

In the final section we use algorithmic information theory to quantify the difficulty of getting work out of mechanical systems. In particular, just as it is possible to make the evolution of the spins undo itself and give the spin echo simply by conjugating the phases of the spins, it is possible in principle to reverse the evolution of any quantum mechanical system with a time-reversal-invariant Hamiltonian by conjugating the phases of its energy eigenstates, and to reverse the evolution of any classical integrable system by conjugating the phases of its angle variables. Why don't we recharge our batteries by reversing the dissipation of electrical energy? Why don't we walk up stairs backward, undoing the downward walk at no extra physical cost? While such feats are possible in principle, the theory of algorithms allows us to quantify just how impractical they are.

2. QUANTUM ELECTRODYNAMIC TREATMENT OF THE SPIN ECHO

Before calculating the increase and decrease in entropy of the spins over the course of the echo, we make sure that a paradox actually exists. After all, the entropy of water decreases when it freezes into ice, without causing any worries about the validity of the second law of thermodynamics: entropy is simply "pumped" from the water into its surroundings. We must show that the sum of the entropy of the parts of the spin-echo system actually decreases, and that entropy is not "pumped" from the spins into the electromagnetic field, as it is "pumped" from the position and momentum degrees of freedom of the hard-sphere gas into the position and angle degrees of freedom of the mirrors that reverse their trajectories.

Classically, in fact, the electromagnetic field is capable of registering arbitrarily small changes in position of the individual spins, and does indeed retain a "memory" of the orientation of the spins at the time of reversal. Quantum mechanically, although the spins do not radiate photons that determine their orientation at the time of reversal, they still absorb angular momentum from the photons in the rf field. Can one detect the orientation of the spins at the time of reversal by looking at changes in the photon population and angular momentum of the rf field, as one can detect the positions and velocities of the molecules in the hard-sphere gas at the time of reversal by looking at the orientations and changes in momentum of the mirrors that effect the reversal? We now give a detailed quantum electrodynamic treatment of the spin echo to show that although the spins absorb photons from the field, they leave no trace of their orientation at the time of reversal, and do not increase the entropy of the field.

The full Hamiltonian for a spin-1/2 particle subject to a magnetic field with strength $B_0 = \omega_0/\gamma$ along the z axis and to an rf field with frequency ω along the x axis is

$$\mathscr{H} = \hbar\omega a^\dagger a + \frac{\hbar\omega_0}{2}\sigma_z + \hbar\kappa(a + a^\dagger)\sigma_x \tag{2.1}$$

where κ is the effective coupling of the rf field to the spin within the sample, and we treat the magnetic field semiclassically.

Using the rotating-wave approximation (valid for weak coupling of the rf field), we have

$$\mathscr{H} = \hbar\omega a^\dagger a + \frac{\hbar\omega_0}{2}\sigma_z + \hbar\kappa(a^\dagger\sigma_- + a\sigma_+) \tag{2.2}$$

where $\sigma_\pm = (\sigma_x \pm i\sigma_y)/2$.

This Hamiltonian may be diagonalized by a trick of Jaynes and Cummings[5] to yield eigenstates

$$|\phi_+(n)\rangle = \cos\theta_n\,|n+1, \downarrow\rangle + \sin\theta_n\,|n, \uparrow\rangle \tag{2.3a}$$

$$|\phi_-(n)\rangle = -\sin\theta_n\,|n+1, \downarrow\rangle + \cos\theta_n\,|n, \uparrow\rangle \tag{2.3b}$$

where $|n, \uparrow\rangle$ is the state with n quanta in the rf field and the spin pointing up, $|n+1, \downarrow\rangle$ is the state with $n+1$ quanta in the field and the spin pointing down, and where

$$\tan\theta_n = \frac{\kappa(n+1)^{1/2}}{(\omega-\omega_0)/2+\lambda_n}$$
$$\lambda_n = \{[(\omega-\omega_0)/2]^2 + \kappa^2(n+1)\}^{1/2} \tag{2.4}$$

The $|\phi_\pm(n)\rangle$ are eigenstates of the Hamiltonian with eigenvalues $\hbar(\omega(n+1/2) \pm \lambda_n)$.

We can decompose the rf field number eigenstates in terms of $|\phi_\pm(n)\rangle$:

$$|n, \uparrow\rangle = \sin\theta_n\, |\phi_-(n)\rangle + \cos\theta_n\, |\phi_+(n)\rangle \tag{2.5a}$$

$$|n+1, \downarrow\rangle = \cos\theta_n\, |\phi_-(n)\rangle - \sin\theta_n\, |\phi_+(n)\rangle \tag{2.5b}$$

Now we can treat the quantum spin-echo effect. Initially, the spins are lined up along the x axis, and the rf field is off. The initial state of a typical spin is

$$|\psi(0)\rangle = \frac{1}{\sqrt{2}}(|0, \uparrow\rangle + |0, \downarrow\rangle) \tag{2.6}$$

The spin now precesses around the magnetic field along the z axis. At time t, the state is

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}}(e^{-i\omega_0 t/2}\, |0, \uparrow\rangle + e^{i\omega_0 t/2}\, |0, \downarrow\rangle) \tag{2.7}$$

At time t the rf field is turned on. Let us follow the evolution of a number eigenstate of the rf field coupled to the spin. We have

$$\begin{aligned} |\psi(t)\rangle &= \frac{1}{\sqrt{2}}(e^{-i\omega_0 t/2}\, |n, \uparrow\rangle + e^{i\omega_0 t/2}\, |n, \downarrow\rangle) \\ &= \frac{1}{\sqrt{2}}\{e^{-i\omega_0 t/2}[\sin\theta_n\, |\phi_+(n)\rangle + \cos\theta_n\, |\phi_-(n)\rangle] \\ &\quad + e^{i\omega_0 t/2}[\cos\theta_{n-1}\, |\phi_+(n-1)\rangle - \sin\theta_{n-1}\, |\phi_-(n-1)\rangle]\} \end{aligned} \tag{2.8}$$

After a time Δt, the spin evolves into a state

$$\begin{aligned} |\psi(t+\Delta t)\rangle = \frac{1}{\sqrt{2}}(&e^{-i\omega_0 t/2 - i(\omega(n+1/2)+\lambda_n)\Delta t}\sin\theta_n\, |\phi_+(n)\rangle \\ &+ e^{-i\omega_0 t/2 - i(\omega(n+1/2)-\lambda_n)\Delta t}\cos\theta_n\, |\phi_-(n)\rangle \\ &+ e^{i\omega_0 t/2 - i(\omega(n-1/2)-\lambda_{n-1})\Delta t}\cos\theta_{n-1}\, |\phi_+(n-1)\rangle \\ &- e^{i\omega_0 t/2 - i(\omega(n-1/2)-\lambda_{n-1})\Delta t}\sin\theta_{n-1}\, |\phi_-(n-1)\rangle) \end{aligned} \tag{2.9}$$

If n is large, so that $\kappa^2(n+1)/[(\omega-\omega_0)/2]^2 \gg 1$, then we have $\lambda_n \approx \kappa(n+1)^{1/2}$, $\theta_n \approx \pi/4$, and

$$\begin{aligned}|\psi(t+\Delta t)\rangle \approx \frac{1}{\sqrt{2}}(&-e^{-i\omega_0 t/2 - i\omega\Delta t/2} \sin\lambda_n \Delta t \, |n+1, \downarrow\rangle \\ &+ e^{-iw_0 t/2 - i\omega\Delta t/2} \cos\lambda_n \Delta t \, |n, \uparrow\rangle \\ &+ e^{i\omega_0 t/2 + i\omega\Delta t/2} \cos\lambda_{n-1} \Delta t \, |n, \downarrow\rangle \\ &- e^{i\omega_0 t/2 + i\omega\Delta t/2} \sin\lambda_{n-1} \Delta t \, |n-1, \uparrow\rangle) \end{aligned} \tag{2.10}$$

Following the experimental procedure for the maximum spin echo,[1] we apply the rf pulse for a time Δt such that $\lambda_n \Delta t \approx \lambda_{n-1} \Delta t \approx \pi/2$, so that

$$|\psi(t+\Delta t)\rangle = -\frac{1}{\sqrt{2}}(e^{-i\omega_0 t/2 - i\omega\Delta t/2} |n+1, \downarrow\rangle + e^{i\omega_0 t/2 + i\omega\Delta t/2} |n-1, \uparrow\rangle) \tag{2.11}$$

(If the rf pulse is applied for an arbitrary length of time, after the pulse the system is in a superposition of a state that gives an echo and a state that does not. That is, any length rf pulse such that $\Delta t \neq N\pi/\lambda_n$ will give some spin echo.) At time $t+\Delta t$ the rf field is turned off, and at time $2t+\Delta t$ we have

$$|\psi(2t+\Delta t)\rangle = -\frac{1}{\sqrt{2}}(e^{-i\omega\Delta t/2} |n+1, \downarrow\rangle + e^{i\omega\Delta t/2} |n-1, \uparrow\rangle) \tag{2.12}$$

$|\psi(2t+\Delta t)\rangle$ *is independent of* ω_0: the rf pulse puts all the spins in the same state at time $2t+\Delta t$, regardless of inhomogeneities in B_z that lead to different frequencies of precession. We have derived the desired result: the spins do not communicate their phase $\omega_0 t$ to the rf field at the time of reversal. We have also derived an unexpected result: $|\psi(2t+\Delta t)\rangle$ as calculated above for a pure number state does not give a spin echo! Each spin taken on its own is in a mixture at time $2t+\Delta t$: $\rho_{\text{spin}} = (\frac{1}{2})(|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|)$. If the rf field is in a number state, the spins fail to realign.

What has happened? Real-life rf fields are generally superpositions of number states, accurately approximated by a coherent state[6]: $|\psi_{\text{rf}}\rangle = e^{-|\alpha|^2/2} \sum_n (\alpha^n/n!) |n\rangle$. If the rf field is in a coherent state, we have

$$\begin{aligned}|\psi(2t+\Delta t)\rangle_{\text{coherent}} &= -\frac{1}{\sqrt{2}} e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{n!} (e^{-i\omega\Delta t/2} |n+1, \downarrow\rangle + e^{i\omega\Delta t/2} |n-1, \uparrow\rangle) \\ &= -\frac{1}{\sqrt{2}} e^{-|\alpha|^2/2} \sum_n \left(\frac{\alpha^{n-1}}{(n-1)!} e^{-i\omega\Delta t} |n, \downarrow\rangle + \frac{\alpha^{n+1}}{(n+1)!} e^{i\omega\Delta t} |n, \uparrow\rangle\right) \\ &\approx |\psi_{\text{rf}}\rangle \left(-\frac{1}{\sqrt{2}}\right) (e^{-i\omega\Delta t} |\downarrow\rangle + e^{i\omega\Delta t} |\uparrow\rangle) \end{aligned} \tag{2.13}$$

Both the rf field and the spins are essentially in pure states at time $2t+\Delta t$: the entropy of the spins at the time of reversal is not transferred to the field. Near the peak of the coherent state, $n \approx \alpha$, we have $\alpha^{n-1}/(n-1)! \approx \alpha^{n+1}/(n+1)!$, and the interference necessary to bring about the spin echo takes place.

If the rf field is in a coherent state, the spin echo takes place as usual. If the rf field is in a squeezed state sufficiently close to a number eigenstate, however, the $|\uparrow\rangle$ and $|\downarrow\rangle$ states of the spin cannot interfere and the spin echo does not occur. The situation is exactly analogous to the double-slit experiment: if the sheet in which the slits are cut is in a squeezed state sufficiently close to a momentum eigenstate, then one can discover which slit a photon has gone through by measuring the sheet's momentum after the photon has passed, and no interference pattern appears on the screen.[7]

The essential features of the spin-echo effect are preserved by the quantum mechanical treatment: the spins become realigned without transferring information to the field. From this point on, then, we will treat the spin echo classically.

3. ALGORITHMIC INFORMATION THEORY

The algorithmic information content $K_U(s)$ of a binary string s is given by the size, in bits, of the shortest program s_U^* that can be used to compute s on a universal computer U:

$$K_U(s) = |s_U^*| \tag{3.1}$$

The intuitive idea behind this definition is simple and quite powerful: the size of the message necessary to communicate certain binary strings is considerably shorter than their length. For instance, the binary expansion of π can be generated from a relatively concise algorithm. By contrast, there is only a small chance that such a simple algorithm will exist for a random sequence of 0's and 1's generated by coin flips.

Solomonoff,[8] Kolmogorov,[9-11] and Chaitin[12-15] have independently suggested how to capture this intuition by means of a rigorous formulation based on a theory of computation. When Eq. (3.1) is used to define algorithmic information content, the resulting formalism bears a strong resemblance to Shannon's information theory. The similarity can be further strengthened by insisting that the minimal programs be self-delimiting: that is, a computer given a minimal program as input should yield an output without being prompted by a special "end marker" that tells the computer to expect no more input. A self-delimiting program carries within it information about its own size that allows the computer to "decode" it without any additional data. Moreover, the sizes of self-delimiting programs satisfy the same Kraft inequality,[16]

$$\sum_{s_U^*} 2^{-K(s_U^*)} \leqslant 1 \tag{3.2}$$

that is so central to the theory of coding. The requirement that the minimal programs be self-delimiting can be met without loss of generality: any uniquely decodable code can be converted into a self-delimiting one without changing the size of the programs.

The physical significance of algorithmic information content derives from the fact that the states of a system can be represented by binary strings that give, for example, the locations of gas particles in phase space or the orientations of spins on a lattice. The regularity of the distributions of particles in phase space is reflected in the simplicity of a minimal algorithm that can reproduce their state (e.g., in the form of an appropriate plot). Indeed, it was pointed out by Bennett[17] that in the thermodynamic limit, the average algorithmic information content of the microstates in an ensemble $\mathscr{E}$,

$$\langle K(s_i) \rangle_{\mathscr{E}} = \sum_{s_i \in \mathscr{E}} p(s_i)\, K(s_i) \tag{3.3}$$

and the statistical (Shannon) entropy of the ensemble,

$$H(\mathscr{E}) = -\sum_{s_i \in \mathscr{E}} p(s_i) \log_2 p(s_i) \tag{3.4}$$

are almost identical:

$$H(\mathscr{E}) \leqslant \langle K(s_i) \rangle_{\mathscr{E}} < H(\mathscr{E}) + K(\mathscr{E}) \tag{3.5}$$

Here, $K(\mathscr{E})$ is the size of the description of the ensemble, which for thermodynamic (i.e., "characterized by a few macroscopic quantities") ensembles is negligible compared to the statistical entropy $H(\mathscr{E})$. Therefore,

$$\langle K(s_i) \rangle_{\mathscr{E}} \simeq H(\mathscr{E}) \tag{3.6}$$

for thermodynamic ensembles, and as Bennett[17] has concluded, one could use average algorithmic information content as a foundation for thermodynamics.

In addition to the algorithmic information content of a single string, one can define the conditional algorithmic information content $K(s \mid t)$ of a string s with respect to some other string t as the size of the minimum length program that needs to be supplied in addition to the "data" t in order to produce s on the output tape, Conditional algorithmic information satisfies the relation

$$K(s, t) = K(t) + K(s \mid t) + O(\log_2 K(s)) \tag{3.7}$$

Conditional algorithmic information content plays an important role in discussions of measurement and especially in consideration of the second law of thermodynamiçs from the viewpoint of a Maxwell's demon-like observer.[3,18] Measurement allows a demon to decrease the statistical entropy of the measured physical system. This decrease is compensated for by the increase in the size of the minimal description—the most concise form for the acquired data—that must be stored by the demon in its memory.

Following measurement, an amount of work

$$\Delta W^{+} = k_{\mathrm{B}} T \ln(2)[H(\rho) - H(\rho_i)] \tag{3.8}$$

can be extracted from the system by the slow isothermal "expansion" which transforms the postmeasurement density matrix ρ_i back into the premeasurement density matrix ρ. As this process can be performed cyclically, the second law may appear to be in danger. Fortunately, this is not the case. Even through the physical system was returned to the premeasurement state, the demon's memory was not. In order to violate the second law, the demon would have to restore its memory to the original, premeasurement "uncluttered" state. However, as suspected by Szilard,[19] pointed out by Landauer,[20,21] and emphasized in the context of Maxwell's demon by Bennett,[17,22] the erasure of useless information is expensive: its cost is $k_{\mathrm{B}} T \ln(2)$ of free energy per bit. The price for the demon of erasing the no longer useful record of the past measurement cannot be less than Boltzmann's constant times the algorithmic information $K_{\mathrm{D}}(i)$, where $K_{\mathrm{D}}(i)$ is the amount of algorithmic information, defined according to the demon's computational resources (hence the subscript), required to specify the result of the measurement. That is, the cost of erasure,

$$\Delta W^{-} = k_{\mathrm{B}} T \ln(2)\, K_{\mathrm{D}}(i) \tag{3.9}$$

must be subtracted from the gain of useful work, leaving the net gain

$$\Delta W = \Delta W^{+} - \Delta W^{-} = k_{\mathrm{B}} T \ln(2)[H(\rho) - H(\rho_i) - K_{\mathrm{D}}(i)] \tag{3.10}$$

However, on average[3,23]

$$\langle K_{\mathrm{D}}(i) \rangle \simeq H(\rho) - H(\rho_i) \tag{3.11}$$

Here,

$$\langle K_D(i) \rangle \equiv \sum_i p_i K_D(i) \tag{3.12}$$

where the $\{p_i\}$ are the probabilities for the complete set $\{i\}$ of mutually exclusive outcomes corresponding to density matrices $\{\rho_i\}$. Combining (3.10) and (3.11), we see that $\langle \Delta W \rangle \leqslant 0$.

This argument, given in more detail elsewhere,[3,18] demonstrates that the acquisition of information allows one to extract useful work from a physical system only when that system happens to be in a regular (that is, concisely describable) state: only then is the minimal cost of erasure, $k_B T \ln(2) K_D(i)$, less than the gain of useful work ΔW^+, Eq. (3.8). Even though the demon could "luck out" occasionally and discover that the system on which it has just performed a measurement turns out to be in a simple state, on the average it will have to erase as many bits as the entropy decrease effected by measurement, and dissipate at least as much free energy as the work gained.

It is therefore suggested to identify the sum

$$\mathscr{S}_D(\rho_i) = k_B \ln(2)[H(\rho_i) + K_D(i)] \tag{3.13}$$

as the true physical entropy, according to the demon. The key advantage of this definition arises from the fact that $\mathscr{S}$ is conserved in measurements on equilibrium ensembles.

The main lesson of this discussion is the realization that the disorder that is responsible for the second law need not be a measure of ignorance and need not be measured by probabilities—known disorder, quantified by the algorithmic randomness of individual states, $K_D(i)$, is equally costly from the thermodynamic point of view. The aim of the following sections is to show that this algorithmic viewpoint is also useful in discussions of the dynamical aspects of the second law.

4. ALGORITHMIC FORMULATION OF THE SPIN-ECHO EFFECT

To apply these algorithmic ideas to the spin-echo effect, we need a coarse graining of the state space. Even if we could measure the orientation of each spin individually, the angle of a given spin in the x–y plane can in general be resolved only to within an accuracy $\Delta\phi$. $\Delta\phi$ is determined by the power of our measuring devices and can take on any value. In our discussion, we take $\Delta\phi$ as given. The scale of the coarse graining in the other variables such as the time and the frequencies of the spins can then be expressed naturally in terms of $\Delta\phi$.

Let $K(\{\phi_i(t)\})$ be the algorithmic complexity of the angles of the spins $\{\phi_i(t)\}$ at time t relative to a coarse graining to an accuracy $\Delta\phi$: i.e.,

$$K(\{\phi_i(t)\}) \equiv K(\{\lfloor \phi_i(t)/\Delta\phi \rfloor\}) \tag{4.1}$$

where $\lfloor \phi_i(t)/\Delta\phi \rfloor$ is the greatest integer less than or equal to $\phi_i(t)/\Delta\phi$. Now $K(\{\phi_i(0)\}) \approx 0$, and $K(\{\phi_i(t)\}) \lesssim N \log_2(2\pi/\Delta\phi)$, since the angle of the spins is defined modulo 2π. Here ≈ 0 means $= 0 + O(1)$ and $\lesssim 0$ means $\leqslant 0 + O(1)$, where $O(1)$ is an additive constant that depends on the universal computer with respect to which the algorithmic complexity is defined. The discretization of angular space induces a discretization of time as well, since the state of the spins does not change by an amount $\Delta\phi$ in a time less than $\Delta t = \Delta\phi/\omega_0$, where ω_0 is the average frequency of Larmor precession of the spins (we assume that the spread of frequencies about the average frequency $\delta\omega$ is small compared to ω_0).

Initially the spins have algorithmic complexity ≈ 0, since all the spins are lined up along the y axis. At time t, we can specify the angles through which the spins have precessed either by giving them outright, or by specifying to an accuracy Δt the time since the spins were aligned, and the frequencies of the individual spins to an accuracy $\Delta\phi/t$. That is, we have

$$K(\{\phi_i(t)\}) \lesssim K(\{\omega_i\}, t) \tag{4.2}$$

where $K(\{\omega_i\}, t)$ is the algorithmic complexity of the integers $\{\lfloor \omega_i t/\Delta\phi \rfloor, \lfloor t/\Delta t \rfloor\}$.

We can estimate $K(\{\omega_i\}, t)$ as follows. Given ω_0 to within an accuracy $\Delta\phi/t$, the average amount of additional algorithmic information required to specify a given frequency to within an accuracy $\Delta\phi/t$ is $\log_2[\delta\omega/(\Delta\phi/t)]$. Our estimate is then

$$K(\{\omega_i\}, t) \sim K(\omega_0) + (N-1)\log_2\left(\frac{\delta\omega}{\Delta\phi/t}\right) + K(t) \tag{4.3}$$

where the factor $(N-1)$ appears rather than N because to specify the frequencies of the entire ensemble of N spins we need specify only the frequencies of $N-1$ spins together with their average frequency ω_0. Now $\Delta\phi = \omega_0 \Delta t$, and $K(t) \sim \log_2(t/\Delta t)$ for most t. Our estimate for the algorithmic complexity of the spins then becomes

$$K(\{\omega_i\}, t) \sim N \log_2\left(\frac{t}{\Delta t}\right) + (N-1)\log_2\left(\frac{\delta\omega}{\omega_0}\right) + K(\omega_0) \tag{4.4}$$

which is also our estimate for $K(\{\phi_i(t)\})$.

When $(t/\Delta t)/(\omega_0/\delta\omega) \geqslant 2\pi/\Delta\phi$, this estimate is greater than the upper bound for $K(\{\phi_i(t)\})$ obtained by listing the orientation of each of the spins explicitly, $N \log_2(2\pi/\Delta\phi)$. But this occurs if and only if $t\,\delta\omega \geqslant 2\pi$, which is just the point at which the spins are pointing every which way.

The algorithmic complexity of the state of the spins does indeed increase: their disorder is real, not just apparent. This disorder comes about because the state of the spins is becoming correlated with the state of the magnetic field in which the spins precess, and the magnetic field is itself disordered because of local inhomogeneities. After the rf pulse has conjugated the phases of the spins, their state starts to become decorrelated from the state of the magnetic field, and their algorithmic complexity eventually returns to its original low value. If the phase conjugation takes place at time t_r, then the algorithmic complexity of the spins at time $t_r + \tau$ is equal to their algorithmic complexity at time $t_r - \tau$. The spins rapidly become disordered, then just as rapidly become ordered again.

The overall algorithmic complexity of spins and magnetic field changes only slowly during this time, however. Because the magnetic field remains constant, and the state of the spins can be obtained simply from the state of the magnetic field and the elapsed time t, the overall algorithmic complexity grows by $\log_2(t/\Delta t)$ on average, as opposed to $N \log_2(t/\Delta t)$ for the algorithmic complexity of the spins alone. To see this slow growth, note that $K(\{\phi_i(t)\}, \{\omega_i\})$ depends on the accuracies $\Delta\omega_i$ to within which the frequencies ω_i can be determined. If the orientations of the spins can be determined to within an accuracy $\Delta\phi$, and if their evolution can be followed up to a time $t_{\max}$, then the ω_i must be known to within an accuracy $\Delta\omega_i = \Delta\phi/t_{\max}$. The coarse graining for the ω_i must be at least this detailed. Given such a coarse graining, we have $K_{\max}(\{\phi_i(t)\}, \{\omega_i\}) \approx K(\{\omega_i\}) + K(t)$. But $K(\{\omega_i\})$ is constant, and $K(t) \approx \log_2(t/\Delta t)$ for most t. The algorithmic complexity of spins and magnetic field together grows slowly by $\log_2(t/\Delta t)$, on average.

It is important to note that although $K(t) \approx \log_2(t/\Delta t)$ for most t, $K(t)$ often dips down much lower than this bound. In fact, no computable monotonically-increasing function grows slowly enough to be a lower bound for $K(t)$.[24] Nonetheless, for an arbitrarily selected t, the probability that $K(t)$ is l bits less than $\log_2(t/\Delta t)$ is less than 2^{-l}. For the great majority of times smaller than the recurrence time, therefore, our logarithmic approximation for $K(t)$ is accurate.

The slow growth of the overall algorithmic complexity, together with the quick growth of the algorithmic complexity of the spins alone, implies that the algorithmic mutual information between spins and magnetic field is growing rapidly:

$$\begin{aligned} K(\{\phi_i(t)\}: \{\omega_i\}) &\equiv K(\{\phi_i(t)\}) + K(\{\omega_i\}) - K(\{\phi_i(t)\}, \{\omega_i\}) \\ &\approx K(\{\phi_i(t)\}) - K(t) \\ &\approx N \log_2(t/\Delta t) \end{aligned}$$

The mutual information between the spins and magnetic field is almost equal to the algorithmic information of the spins themselves; the state of the magnetic field, together with the amount of time that has elapsed since the spins were aligned, suffices to determine $\{\phi_i(t)\}$.

The algorithmic picture of the spin-echo effect is clear: The spins do indeed become disordered and then become ordered again, but the overall disorder of spins and magnetic field only increases by a small amount. The structure of the magnetic field nonuniformities supplies the crucial "data" that allow the reversal to be accomplished at a relatively small price.

5. INFORMATION-THEORETIC TREATMENT OF THE SPIN ECHO

We now show that the normal statistical mechanical entropy of the spins does in fact first increase, then decrease in the course of the echo. Since the spins and magnetic field make up a Hamiltonian system (ignoring spin–spin and spin–lattice interaction), the overall statistical entropy of spins and field remains constant, as does the entropy of the magnetic field on its own. It is straightforward to calculate the statistical entropy of the spins at time t in terms of the entropy of the field.

Let $p(\omega_i)\,d\omega$ be the probability that the Larmor frequency of the ith spin lies in the range $[\omega_i, \omega_i + d\omega)$. Similarly, let $p_t(\phi_i)\,d\phi$ be the probability that ϕ_i lies within the range $[\phi_i, \phi_i + d\phi)$ at time t. Since $\phi_i(t) = \omega_i t$ at time t, we have $p_t(\phi_i)\,d\phi = (1/t)\,p(\phi_i/t)\,d\phi$. The entropies

$$H(\phi_i(t)) = -\int_0^{2\pi} p_t(\phi)\log_2 p_t(\phi/\Delta\phi)\,d\phi$$

$$H(\omega_i) = -\int_{-\infty}^{+\infty} p(\omega)\log_2 p(\omega/\Delta\omega)\,d\omega$$

are then simply related for times less than $2\pi/\delta\omega$:

$$H(\phi_i(t)) = H(\omega_i) + \log_2(\Delta\omega t/\Delta\phi) \tag{5.1}$$

Since we are using entropies over continuous probability distributions, we normalize so that distributions with spread $\Delta\phi$ and $\Delta\omega$ have entropy zero. Probability distributions with smaller spreads have negative entropies: however, since our coarse graining allows no accuracy greater than $\Delta\omega$, $\Delta\phi$, such negative entropies will not occur. In addition, since the orientations of the spins are only defined modulo 2π, whenever $H(\phi_i(t)) > \log_2(2\pi/\Delta\phi)$ we set the two equal.

At time $t = 0$ the statistical entropy of the spins is zero, and the entropy of spins and field taken together is just the entropy of the field. If the deviations of the Larmor frequencies about their mean are uncorrelated, then the statistical entropy of all the spins is just N times the entropy for a single spin, and the mutual information between spins and field as a function of time is

$$I(\{\phi_i(t)\}:\{\omega_i\}) = H(\{\phi_i(t)\}) + H(\{\omega_i\}) - H(\{\phi_i(t)\}, \{\omega_i\})$$
$$= H(\{\omega_i\}) + N \log_2 \left(\frac{\Delta\omega t}{\Delta\phi}\right) \tag{5.2}$$

where once again the right-hand side is set equal to zero when negative. The statistical entropy of the spins rises as $N \log_2 t$ and is equal to the mutual information between spins and magnetic field.

After the rf pulse has conjugated the phases of the spins, the entropy of the spins decreases just as rapidly: the natural evolution of the system "uses" the mutual information between spins and field to decrease the entropy of the spins. The amount of mutual information available is exactly enough to reduce the entropy of the spins to zero without requiring an increase in entropy elsewhere.[25,26]

6. STATISTICAL ENTROPY VS. ALGORITHMIC INFORMATION

The only difference between the algorithmic randomness of the spin-echo system and its statistical entropy is that while the entropy remains constant, the algorithmic complexity grows slowly by $\sim\log_2(t/\Delta t)$, the term that comes from specifying the amount of time that the system has evolved since the spins were aligned. Although initially a small term compared with the entropy of the spins or magnetic field at time t, this term becomes comparable with these entropies if one waits for a substantial fraction of the Poincaré recurrence time for the system.

For the macroscopic spin-echo system, the spin–spin and spin–lattice relaxation times are much smaller than the Poincaré time, and the spins will have become thermally randomized long before the term $\log_2(t/\Delta t)$ can amount to much. The term $\log_2(t/\Delta t)$ is comparable to the algorithmic complexity of the spins when $t \sim \Delta t 2^N$. For a collection of only a few spins, this term can make a difference after relatively short times. In addition, any integrable Hamiltonian system is analogous to the spin-echo system. The phase space of an integrable system can by a canonical transformation be decomposed into action variables—analogous to the frequencies of precession—and angle variables—analogous to the angles of the spins. One can give an algorithmic treatment for any integrable system that is wholly analogous to the algorithmic treatment of the spin-echo system. Since Hamiltonian systems are nondissipative, it is in principle feasible to follow accurately the evalution of such a system up to a substantial fraction of its Poincaré time.

The question then arises: Does the extra term $\log_2(t/\Delta t)$ make a difference in terms of the amount of energy that can be extracted from such a system as work? Or is the amount of energy that can be extracted constant over time, as the constancy of the fine-grained entropy seems to imply? We show below that the extra term does indeed make a difference in terms of the amount of energy that can be extracted from an integrable system as work: a device that is to extract this energy must devote more and more resources to the extraction as time goes on in order to extract the same amount of work. The effective *physical entropy*—the quantity that puts limits on how much work can be extracted from a system—rises slowly, as $\log_2(t/\Delta t)$. The algorithmic treatment of integrable systems thus gives a new contribution to thermodynamics.[3,18]

7. PHYSICAL ENTROPY

To see why extracting the same amount of work from an integrable system requires more and more resources as time goes on, we apply the idea of physical entropy developed in Section 3 to the problem of taking full advantage of the energy in the induction signal from the spin echo—say, for example, by storing it as electrical energy in a capacitor. We look at a "demon," a device that can capture and store this energy, and show that as time goes on, it must either use more and more memory to perform this task, or dissipate more and more energy.

The radiofrequency coil in which the spins sit is an LC circuit: the "demon" is simply a device that throws open a switch in the circuit between coil and capacitor after the charge driven by the pulse arrives, with the result that the capacitor receives the charge and stores the energy from the signal (Fig. 1a). Such a demon is capable of capturing and storing virtually all of the energy in the induction signal, but only if it throws the switch open at precisely the right time. If the switch is thrown too soon, no charge build up on the capacitor—too late, and all the charge leaks away. To store all the energy in the pulse, the demon must know the precise time of the pulse's arrival. To keep track of this time, the apparatus must devote $K(t) \sim \log_2(t/\Delta t)$ of memory space: here the amount of resources allocated to capturing the energy of the pulse, as measured by the amount of memory space required, rises as the logarithm of the time, on average. It is possible that the pulse can arrive at a large but algorithmically simple time, in which case the demon need only devote a small amount of resources to capturing the energy of the pulse. If the time of arrival is selected arbitrarily, however, the chances of $K(t)$ differing significantly from $\log_2(t/\Delta t)$ are small.

Suppose that the demon that is to capture the charge driven by the pulse does not keep track of when the pulse is to arrive: it can still capture this charge by using a trigger to open the switch after the voltage has risen to a predetermined level. Clearly, such a strategy will also capture most of the charge, and will not have to allocate resources to store the time of the pulse's arrival. The only sacrifice is that some of the energy in the pulse

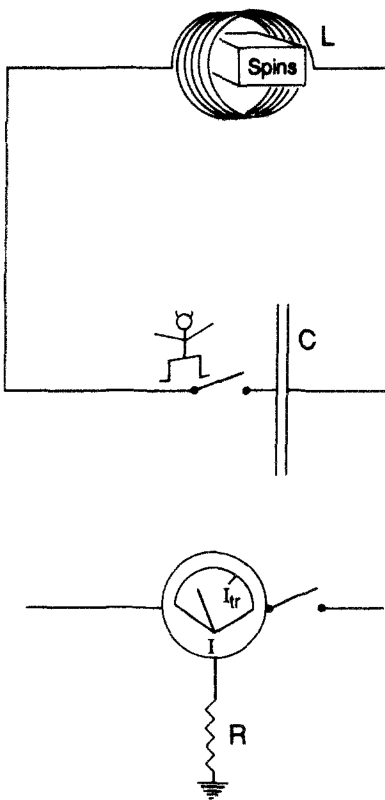

Fig. 1. (a) Spin-echo circuit with "demon." The spin-echo circuit is an LC circuit incorporating a radiofrequency coil with inductance L that contains the spins, and a capacitor C. To store the maximum amount of electrical energy from the induction pulse generated by the echo, the demon must open the switch when the charge on the capacitor reaches its maximum value. (b) Example of a "demon." The spin-echo circuit is connected to ground by a wire with a high resistance R. When the current flowing to ground surpasses a value I_{tr}, the switch is thrown open. I_{tr} is chosen so that when the switch is thrown, the maximum amount of charge is isolated on the capacitor.

must be diverted to the voltage detector. The requirement that the device not be triggered by a thermal fluctuation while waiting for the pulse then puts a minimum value on the energy that must be diverted.

Let the demon use a voltage detector of the most simple sort (Fig. 1b). Attach to the circuit a wire with a known resistance connected to ground. Since the wire's resistance is known, the amount of current flowing along the wire gives the voltage between circuit and ground; the amount of current can be registered by a magnetized needle whose deflection indicates the strength of the magnetic field in the neighborhood of the wire. When the needle's deflection exceeds a set value, it opens the switch, trapping charge on the capacitor. To minimize the dissipation caused by such a device, the resistance of the wire should be maximized, so that a smaller current within the wire reflects a larger voltage drop along the wire.

But the demon must be careful: if the resistance of the wire is too great, the triggering device may interpret the current from a thermal fluctuation as indicating the presence of the voltage pulse, and open the switch when no pulse is in fact present. If the device is triggered by an amount of energy E_{tr} in the electromagnetic field around the wire, then to ensure that the device is not triggered by a thermal fluctuation in the course of an interval of length t, we must have the probability of a false trigger $= e^{-E_{tr}/k_B T}(t/t_{fl}) \ll 1$, where T is the temperature of the wire, and t_{fl} is the amount of time it takes for a fluctuation of size $k_B T$ to arise on average. (The demon encounters a similar problem in balancing between minimizing dissipation and maximizing stored charge if it tries to trap the charge by inserting a diode in the circuit before the capacitor.[27])

Since the circuit is an LC circuit tuned to frequency ω, we have $t_{fl} \approx 2\pi/\omega = \Delta t/\Delta\phi$, and so we have

$$E_{tr} \gg k_B T\left(\ln\frac{t}{\Delta t} + \ln\frac{\Delta\phi}{2\pi}\right) \tag{7.1}$$

But E_{tr} is just the amount of energy dissipated in ensuring the capture of the energy of the pulse. The amount of dissipation, $E_{tr}/k_B T$, required to take advantage of the pulse grows as the logarithm of the amount of time in which the pulse can arrive.

We have sketched the two extreme alternatives: in the first, according to the demon, the pulse is in a definite state, arriving at a definite time, and the demon is able to take full advantage of the energy in the pulse, but must allocate $K \sim \log_2(t/\Delta t)$ of memory space to store the record of the pulse's arrival time. After the pulse has arrived, this record is of no further use to the demon, but can be erased only at the cost of $\sim k_B T \ln(2) \log_2(t/\Delta t)$ of dissipation. In the second alternative, the demon possesses no information about the time of arrival of the pulse. According to such a demon, the pulse has an entropy $\log_2(t/\Delta t)$. When the demon pins down the energy of the pulse to a definite state by storing it in the capacitor, the entropy $\log_2(t/\Delta t)$ must go somewhere: in the device described above, it goes to dissipation in the triggering mechanism. The total physical entropy $\mathscr{S} = H + K$ is the same for both scenarios, and rises as the logarithm of time.

There is also an intermediate alternative in which the demon has some information about when the pulse is to arrive, but that information is incomplete. For example, suppose that the demon knows the time t at which the pulse is to arrive only to an accuracy δt, $t \gg \delta t \gg \Delta t$. The amount of algorithmic information required to specify this inexact time of arrival is $\log_2(t/\delta t)$, on average. If the demon uses the triggering device described above to discover when the pulse actually arrives during the time interval δt, he must dissipate at least $k_B T \ln(\delta t/\Delta t)$. The total amount of thermodynamic resource required is thus at least

$$k_B T \ln(2) \log_2(t/\delta t) + k_B \ln(2) \log_2(\delta t/\Delta t) = k_B \ln(t/\Delta t) \tag{7.2}$$

as before.

There are two lessons to be learned here. First, to take full advantage of the energy inherent in a system, a device must have an algorithmically exact description of the form in which the energy is distributed. Second, since the algorithmic complexity of Hamiltonian systems increases slowly over time, more and more resources must be brought to bear in getting work out of such systems even though their fine-grained, statistical entropy remains constant.

8. SUMMARY

In the spin-echo effect, the disorder of the spins first increases, then decreases dramatically. This decrease of disorder does not violate the second law of thermodynamics, however, since the state of the spins is highly correlated with the state of the magnetic field, and entropy ceases to be an extensive variable. The overall fine-grained entropy of spins and magnetic field is constant. The algorithmic complexity of spins and field, by contrast, grows slowly by the logarithm of the time on average. The increase in algorithmic complexity is reflected in the increased number of resources that must be brought to bear to take full advantage of the spin echo.

In addition, a full quantum electrodynamic treatment of the spins and the rf field using the rotating-wave approximation reveals that the coherent nature of the electromagnetic field is essential for attaining the spin echo. If the rf field is in a number eigenstate, the echo will not take place. Squeezed states that can approximate number eigenstates at the frequencies required for the spin echo are not available at present, although optical squeezed states have been produced.[28] If sufficiently intense optical states close to number eigenstates could be produced, one could observe the quantum mechanical suppression of optical echoes.

ACKNOWLEDGMENTS

The authors would like to thank the Santa Fe Institute for its hospitality and support. This work was supported in part by the U.S. Department of Energy under contract DE-AC0381-ER40050.

REFERENCES

1. E. L. Hahn, *Phys. Rev.* **80**:580–594 (1950).
2. W.-K. Rhim, A. Pines, and J. S. Waugh, *Phys. Rev. B* **3**:684–696 (1971).
3. W. H. Zurek, *Phys. Rev. A* **40**:4731–4751 (1989).
4. W. H. Zurek, *Phys. Rev. Lett.* **53**:391–394 (1984).
5. E. T. Jaynes and F. W. Cummings, *Proc. I.R.E.* **51**:89–109 (1963).
6. R. J. Glauber, *Phys. Rev.* **131**:2766–2788 (1963).
7. W. K. Wootters and W. H. Zurek, *Phys. Rev. D* **19**:473–484 (1979).
8. R. J. Solomonoff, *Information Control* **7**:1–22 (1964).
9. A. N. Kolmogorov, *Inf. Transmission* **1**:3–11 (1965).
10. A. N. Kolmogorov, *IEEE Trans. Inf. Theor.* **14**:622–664 (1968).
11. A. N. Kolmogorov, *Usp. Mat. Nauk* **25**:602–613 (1970).
12. G. J. Chaitin, *J. ACM* **13**:547–569 (1966).
12. G. J. Chaitin, *Sci. Am.* **232**:47–52 (1975).
14. G. J. Chaitin, *IBM J. Res. Dev.* **21**:350–359 (1977).
15. G. J. Chaitin, *Algorithmic Information Theory* (Cambridge University Press, 1987).
16. R. W. Hamming, *Coding and Information Theory* (Prentice-Hall, Englewood Cliffs, New Jersey, 1986).
17. C. H. Bennett, *Int. J. Theor. Phys.* **21**:905–940 (1982).
18. W. H. Zurek, *Nature* **341**:119–124 (1989).
19. L. Szilard, *Z. Phys.* **53**:840–856 (1929); *Behav. Sci.* **9**:301–310 (1964).
20. R. Landauer, *IBM J. Res. Dev.* **5**:183–191 (1961).
21. R. Landauer, *Nature* **335**:779–784 (1988).
22. C. H. Bennett, *Sci. Am.* **257**:108 (1987).
23. C. M. Caves, in *Complexity, Entropy and the Physics of Information*, W. H. Zurek, ed. (Addison-Wesley, New York, 1990).
24. A. K. Zvonkin and L. A. Levin, *Uspekhi Mat. Nauk* **25**(6):85–128 (1970) [English translation in *Russ. Math. Surveys* **25**(6):83–124 (1970)].
25. J. Rothstein, *Am. J. Phys.* **24**:510–518 (1957).
26. S. Lloyd, *Phys. Rev. A* **39**:5378–5386 (1989).
27. R. P. Feynman, R. J. Leighton, and M. Sands, *The Feynman Lectures in Physics* (Addison-Wesley, Reading, Massachusetts, 1965).
28. M. C. Teich and B. E. A. Saleh, *Physics Today* **43**:26–34 (1990).