Information-Theoretic Limits of Control

Hugo Touchette* and Seth Lloyd†
d'Arbeloff Laboratory for Information Systems and Technology, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
(Received 27 May 1999; revised manuscript received 16 September 1999)

Fundamental limits on the controllability of physical systems are discussed in the light of information theory. It is shown that the second law of thermodynamics, when generalized to include information, sets absolute limits to the minimum amount of dissipation required by open-loop control. In addition, an information-theoretic analysis of control systems shows feedback control to be a zero sum game: each bit of information gathered from a dynamical system by a control device can serve to decrease the entropy of that system by at most one bit additional to the reduction of entropy attainable without such information. Consequences for the control of discrete state systems and chaotic maps are discussed.

PACS numbers: 05.45.Gg, 05.20.−y, 89.70.+c

Information and uncertainty represent complementary aspects of control. Open-loop control methods attempt to reduce our uncertainty about system variables such as position or velocity, thereby increasing our information about the actual values of those variables. Closed-loop methods obtain information about system variables and use that information to decrease our uncertainty about the values of those variables. Although the literature in control theory implicitly recognizes the importance of information in the control process, information is rarely regarded as the central quantity of interest [1]. In this Letter we address explicitly the role of information and uncertainty in control processes by presenting a novel formalism for analyzing these quantities using techniques of statistical mechanics and information theory. Specifically, based on a recent proposal by Lloyd and Slotine [2], we formulate a general model of control and investigate it using entropylike quantities. This allows us to make mathematically precise each part of the intuitive statement that in a control process information must constantly be acquired, processed, and used to constrain or maintain the trajectory of a system. Along this line, we prove several limiting results relating the ability of a control device to reduce the entropy of an arbitrary system in the cases where (i) such a controller acts independently of the state of the system (open-loop control) and (ii) the control action is influenced by some information gathered from the system (closed-loop control). These results not only combine concepts of dynamical entropy and information in a unified picture, but also prove to be fundamental in that they represent the ultimate physical limitations faced by any control systems.

The basic framework of our present study is the following. We assign to the physical plant $\mathcal{X}$ that we want to control a random variable X representing its state vector (of arbitrary dimension) and whose value x is drawn according to a probability distribution $p(x)$. Physically, this probabilistic or ensemble picture may account for interactions with an unknown environment, noisy inputs, or unmodeled dynamics; it can also be related to a deterministic sensitivity to some parameters which make the system effectively stochastic. The recourse to a statistical approach then allows the treatment of both the unexpectedness of the control conditions and the dynamical stochastic features as two faces of a single notion: *uncertainty*.

As is well known, a suitable measure quantifying uncertainty is entropy [3,4]. For a classical system with a discrete set of states with probability mass function $p(x)$, it is expressed as

$$H(X) \equiv -\sum_x p(x) \log p(x) \tag{1}$$

(all logarithms are assumed to the base 2 and the entropy is measured in bits). Similar expressions also exist for continuous state systems (fine-grained entropy), quantum systems (von Neumann entropy), and coarse-grained systems obtained by discretization of continuous densities in the phase space [5]. In all cases, entropy offers a precise measure of disorderliness or missing information by characterizing the minimum amount of resources (bits) required to encode unambiguously the ensemble describing the system [4–6]. As for the time evolution of these entropies, we know that the fine-grained (or von Neumann) entropy remains constant under volume-preserving (unitary) evolution, a property closely related to the fact that only one-to-one mappings of states, i.e., *reversible* transformations preserving information, are exempt of dissipation [7]. Coarse-grained entropies, on the other hand, usually increase in time even in the absence of noise due to the finite nature of the partition used in the coarse graining which, in effect, induces a "randomization" of the motion [8].

In this context, we now address the problem of how a control device can be used to reduce the entropy of a system or to immunize it from sources of entropy, in particular, those associated with noise, motion instabilities, incomplete specification of states, and initial conditions. Although the problem of controlling a system requires more than limiting its entropy, the ability to limit entropy is a prerequisite to control. Indeed, the fact that a control

 0031-9007/00/84(6)/1156(4)$15.00

process is able to localize a system in definite stable states or trajectories simply means that the system can be constrained to evolve into states of low entropy.

To illustrate, in its most simple way, how the entropy of a system can be affected by external systems, let us consider a basic model consisting of a system $\mathcal{X}$ coupled to an environment $\mathcal{E}$. For simplicity, we assume that the states of $\mathcal{X}$ form a discrete set evolving over discrete time intervals Δt [9]. The initial state is again distributed according to $p(x)$, and the effect of the environment is taken into account by introducing a perturbed conditional distribution $p(x' \mid e)$, where x' is a value of the state later in time and e, a particular *realization* of the stochastic perturbation appearing with probability $p(e)$. For each value e, we assume that $\mathcal{X}$ undergoes a unique evolution, referred to here as a *subdynamics*, taken to be entropy conserving in analog to the Hamiltonian time evolution for a continuous physical system:

$$H(X' \mid e) \equiv -\sum_{x'} p(x' \mid e) \log p(x' \mid e) = H(X). \quad (2)$$

After the time transition $X \to X'$, the distribution $p(x')$ is obtained by tracing out the variables of the environment and is used to calculate the change of the entropy $H(X') = H(X) + \Delta H$. From the concavity property of entropy, it can be easily shown that $\Delta H \geq 0$, with equality if and only if (iff) the state $\mathcal{E}$ is perfectly specified, i.e., if a value e appears with probability one. In practice, however, the environment degrees of freedom are uncontrollable and the uncertainty associated with the environment coupling can be suppressed by "updating" our knowledge of X after the evolution.

One direct way to reveal that state is to imagine a measurement apparatus $\mathcal{A}$ coupled to $\mathcal{X}$ in such a way that the dynamics of the composed system $\mathcal{X} + \mathcal{E}$ is left unaffected. For this measurement scheme, the outcome of some discrete random variable A of the apparatus is described by a conditional probability matrix $p(a \mid x')$ and the marginal $p(a)$ from which we can derive $H(X' \mid A) \leq H(X')$ with equality iff A is independent of X [4]. In this last inequality we have used $H(X' \mid A) \equiv \sum_a H(X' \mid a) p(a)$, and $H(X' \mid a)$ given similarly as in Eq. (2). Now, upon the application of the measurement, one can define the reduction of entropy of the system conditionally on the outcome of A by $\Delta H_A = H(X' \mid A) - H(X)$, which, obviously, satisfies $\Delta H_A \leq \Delta H$, and $H(A) \geq \Delta H - \Delta H_A$. In other words, the decrease in the entropy of $\mathcal{X}$ conditioned on the state of $\mathcal{A}$ is compensated for by the increase in entropy of $\mathcal{A}$. This latter quantity represents information that $\mathcal{A}$ possesses about $\mathcal{X}$. Accordingly, the entropy of X given A plus the entropy of A is nondecreasing, which is an expression of the second law of thermodynamics as applied to interacting systems [10–12].

It must be stressed that the reduction of entropy of $\mathcal{X}$ discussed so far is conditional on the outcome of A. By assumption, $\mathcal{X}$ is not affected by $\mathcal{A}$; as a result, according to an observer who does not know this outcome, the entropy of $\mathcal{X}$ is unchanged. In order to reduce entropy for all observers unconditioned on the state of any external systems, a direct dynamical action on $\mathcal{X}$ must be established externally by a *controller* $\mathcal{C}$ whose influence on the system is represented by a set of control actions $x \xrightarrow{c} x'$ triggered by the controller's state c. Mathematically, these actions can be modeled by a probability transition matrix $p(x' \mid x, c)$ giving the probability that the system in state x goes to state x' given that the controller is in state c. The specific form of this *actuation* matrix depends on the subdynamics envisaged in the control process: some actions, for example, may correspond to dissipative (volume contracting) control strategies forcing several initial conditions to a common state, while others can model uncontrolled transitions perturbed by noise leading to entropy increasing actuation rules. Hence, the systems $\mathcal{X}$ and $\mathcal{C}$ need not in general model a closed Hamiltonian system; $\mathcal{X}$, as we already noted, can be an open system, i.e., one affected by external systems (e.g., environment) on which one has usually no control. Formally speaking, though, one can always embed any open-system evolution in a higher dimensional closed system whose dynamics mimics a Hamiltonian system. This can be done by supplementing an open system with a set of *ancillary* variables acting as an environment $\mathcal{E}$ in order to construct a global volume-preserving transition matrix such that, when the ancillary variables are traced out, the reduced transition matrix thus obtained reproduces the dynamics of the system $\mathcal{X} + \mathcal{C}$.

Note that these ancillary variables need not have any physical significance: they are only there for the purpose of simplifying the analysis of the evolution of the system. In particular, any control strategy must be independent of the choice of $\mathcal{E}$ which means, within our model, that the control of the system $\mathcal{X}$ can be assured only by the choice of the control variable C whereby we can force an ensemble of transitions leading the system to a net entropy change ΔH. Since the overall dynamics of the system, controller, and environment is Hamiltonian, Landauer's principle immediately implies that if the controller is initially uncorrelated with the system, so that $p(x, c) = p(x)p(c)$ for all x and c as in the case with open-loop control, a decrease in entropy ΔH for the system must be compensated for by an increase in entropy of at least ΔH for the controller and the environment [12]. Furthermore, using again the concavity property of H, it can be shown that the maximum decrease of entropy achieved by a particular subdynamics of control variable $\hat{c}$ is always *optimal* in the sense that no probabilistic choice of the control parameter can improve upon that decrease. Explicitly, we have the following theorem (we omit the proof which follows simply from the concavity property).

Theorem 1.—For open-loop control, the maximum value of $\Delta H_{\text{open}} \equiv H(X) - H(X')$ can always be attained for a *deterministic* choice of the control variable, i.e., with $p(\hat{c}) = 1$ and $p(c) = 0$ for all $c \neq \hat{c}$, where $\hat{c}$ is the value of the controller leading to $\max \Delta H_{\text{open}}$. Any nondeterministic choice of the control variables either achieves the maximum or yields a smaller value.

From the standpoint of the controller, one major drawback of acting independently of the state of the system is that no information other than that available from the state of $\mathcal{X}$ itself can provide a reasonable way to determine which subdynamics are optimal or accessible given the initial state. For this reason, open-loop control strategies usually fail to operate efficiently in the presence of noise because of their inability to react or be adjusted in time. In order to account for all the possible behaviors of a stochastic dynamical system, we have to use the information contained in its evolution by considering a *closed-loop* control scheme in which the state of the controller is allowed to be correlated to the initial state of $\mathcal{X}$. This correlation can be thought of as a measurement process described earlier that enables C to gather an amount of information given formally in Shannon's information theory [3,4] by the *mutual information* $I(X;C) \equiv H(X) + H(C) - H(X,C)$, where $H(X,C) \equiv -\sum_{x,c} p(x,c) \log p(x,c)$ is the *joint* entropy of X and C. Having defined these quantities, we are now in a position to state our main result which is that the maximum improvement that closed-loop can give over open-loop control is limited by the information obtained by the controller. More formally, we have the following theorem.

Theorem 2.—The amount of entropy ΔH_{closed} that can be extracted from any dynamical system by a closed-loop controller satisfies

$$\Delta H_{\text{closed}} \leq \Delta H_{\text{open}} + I(X;C), \tag{3}$$

where ΔH_{open} is the maximum entropy decrease that can be obtained by open-loop control and $I(X;C)$ is the mutual information gathered by the controller upon observation of the system state.

Proof.—We construct a closed system by supplementing an ancilla $\mathcal{E}$ to our previous system $\mathcal{X} + \mathcal{C}$. Moreover, let $\mathcal{C}$ and $\mathcal{E}$ be collectively denoted by $\mathcal{B}$ with state variable B. Since the complete system $\mathcal{X} + \mathcal{B}$ is closed, its entropy has to be conserved, and thus $H(X,B) = H(X',B')$. Defining the entropy changes of $\mathcal{X}$ and $\mathcal{B}$ by $\Delta H_X = H(X) - H(X')$ and $\Delta H_B = H(B') - H(B)$, respectively, and by using the definition of the mutual information, this condition of entropy conservation can also be rewritten in the form $\Delta H_X = \Delta H_B - I(X';B') + I(X;B)$ [12]. Now, let ΔH_{open} be the maximum amount of entropy decrease of $\mathcal{X}$ obtained in the open-loop case where $I(X;C) = I(X;B) = 0$ [by construction of $\mathcal{E}$, $I(X;E) = 0$]. From the conservation condition, we hence obtain $\max \Delta H_X = \Delta H_{\text{open}} + I(X;B)$, which is the desired upper bound for a feedback controller. □

The above results can be generalized to continuous state systems by using the definitions of entropy and mutual information involving integrals instead of sums over the states, as mentioned in the note [9]. In the continuous state version of theorem 2, it is worth noting that $I(X;C)$ is well defined and non-negative even when X and C are both continuous variables or when one is continuous and the other discrete. Also, in the limit $\Delta t \to 0$, Eq. (3) becomes a rate equation limiting the improvement obtainable by the addition of continuous feedback control.

To illustrate the two theorems, suppose that we control a system in a mixture of the states $\{0, 1\}$ using a controller restricted to use the following two actions:

$$\begin{cases} c = 0\colon x \to x' = x\,, \\ c = 1\colon x \to x' = \textsc{not}\, x \end{cases} \tag{4}$$

(in other words, the controller and the system behave like a so-called "controlled-NOT" gate). Since these actuation rules simply permute the state of $\mathcal{X}$, $H(X') \geq H(X)$ with equality if we use a deterministic control strategy or if $H(X) = H_{\max} = 1$ bit, in agreement with our first theorem. Thus, $\Delta H_{\text{open}} = 0$. However, by knowing the actual value of x [$H(X)$ bit of information] we can choose C to obtain $\Delta H_X = H(X)$, therefore achieving Eq. (3) with equality. Evidently, as implied by this equation, information is required here as a result of the nondissipative nature of the actuations and would not be needed if we were allowed to use dissipative subdynamics. Conversely, no open-loop controlled situation is possible if we confine the controller to use entropy-increasing actuations as, for instance, in the control of nonlinear systems using *chaotic* dynamics.

In order to demonstrate this last statement, let us consider the feedback control scheme proposed by Ott, Grebogi, and Yorke (OGY) [13] as applied to the logistic map,

$$x_{n+1} = r x_n (1 - x_n), \qquad x \in [0,1] \tag{5}$$

(the extension to more general systems naturally follows). The OGY method, specifically, consists of applying to Eq. (5) small perturbations $r \to r + \delta r_n$ according to $\delta r_n = -\gamma(x_n - x^*)$, whenever x_n falls into a region D in the vicinity of the target point x^*. The gain $\gamma > 0$ is chosen so as to ensure stability [14]. For the purpose of chaotic control, all the accessible control actions determined by the values of δr_n, and thereby by the coordinates $x_n \in D$, can be constrained to be entropy increasing for a proper choice of D, meaning that the Lyapunov exponent $\lambda(r)$ associated with any actuation indexed by r is such that $\lambda(r) > 0$ [15]. Physically, this implies that *almost* any initial uniform distribution for X covering an interval of size ε "expands" by a factor $2^{\lambda(r)}$ on average after one iteration of the map with parameter r [16–18]. Now, for an open-loop controller, it can readily be shown in that case that no control of the state x is possible; without knowing the position x_n, a controller merely acts as a perturbation to the system, and the optimal control strategy then consists of using the smallest Lyapunov exponent available so as to achieve $\Delta H_{\text{open}} = -\lambda_{\min} < 0$. Following theorem 2, it is thus necessary, in order to achieve a controlled situation $\Delta H_X \geq 0$, to have $I(X;C) \geq \lambda_{\min}$ using a measurement channel characterized by an information capacity [4] of at least $\lambda_{\min}$ bit per use.

In the controlled regime ($n \to \infty$), this means specifically that, if we want to localize the trajectory generated by Eq. (5) uniformly within an interval of size ε using a set of chaotic actuations, we need to measure x within an interval no larger than $\varepsilon 2^{-\lambda_{\min}}$. To understand this, note that an optimal measurement of $I(X;C) = \log a$ bits consists, for a uniform distribution $p(x)$ of size ε, of partitioning the interval ε into a subintervals of size ε/a. The controller under the partition then applies the same actuation $r^{(i)}$ for all the coordinates of the initial density lying in each of the subintervals i, therefore stretching them by a factor of $2^{\lambda(r^{(i)})}$. In the optimal case, all the subintervals are directed toward x^* using $\lambda_{\min}$ and the corresponding entropy change is thus

$$\Delta H_{\text{closed}} = \log \varepsilon - \log 2^{\lambda_{\min}} \varepsilon/a = -\lambda_{\min} + \log a \,, \tag{6}$$

which is consistent with Eq. (3) and yields the aforementioned value of a for $\Delta H_{\text{closed}} = 0$. Clearly, this value constitutes a lower bound for the OGY scheme since not all the subintervals are controlled with the same parameter r, a fact observed in numerical simulations [19].

In summary, we have introduced a formalism for studying control problems in which control units are analyzed as informational mechanisms. In this respect, a feedback controller functions analogously to a so-called Maxwell's demon [20]. In fact, when applied to microscopic systems, our results provide absolute limits to the ability of such a demon to convert heat to work by obtaining information [10,12]. Our main result showed that the amount of entropy that can be extracted from a dynamical system by a controller is upper bounded by the sum of the decrease of entropy achievable in open-loop control and the mutual information between the dynamical system and the controller created during an initial interaction. This upper bound sets a fundamental limit on the performance of any controllers whose designs are based on the possibilities to accede low entropy states. Hence, its practical implications can be investigated for the control of linear, nonlinear, and complex systems (discrete or continuous), as well as for the control of quantum systems. For this latter topic, our probabilistic approach seems particularly suitable for the study of quantum controllers.

The authors thank J.-J. E. Slotine for helpful discussions. This work has been partially supported by NSERC (Canada), and by a grant from the d'Arbeloff Laboratory for Information Systems and Technology, MIT.

*Electronic address: htouchet@mit.edu
†Electronic address: slloyd@mit.edu

[1] R.F. Stengel, *Optimal Control and Estimation* (Dover, New York, 1994).
[2] S. Lloyd and J.-J. E. Slotine, Int. J. Adapt. Cont. Sig. Proc. **10**, 499 (1996).
[3] C.E. Shannon and W. Weaver, *The Mathematical Theory of Communication* (University of Illinois Press, Urbana, 1963).
[4] T.M. Cover and J.A. Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
[5] A. Wehrl, Rev. Mod. Phys. **50**, 221 (1978).
[6] R. Shaw, Z. Naturforsch. **36A**, 80 (1981); J.D. Farmer, Z. Naturforsch. **37A**, 1304 (1982).
[7] R. Landauer, IBM J. Res. Dev. **3**, 183 (1961); Nature (London) **355**, 779 (1988).
[8] C. Beck and F. Schlögl, *Thermodynamics of Chaotic Systems* (Cambridge University Press, Cambridge, England, 1993).
[9] This assumption is used throughout the paper, although all our results apply equally to continuous variables evolving continuously in time. To apply the formulas given in the text to continuous variables, the sums must be changed to integrals, while the continuous time limit is obtained by taking $\Delta t \to 0$.
[10] C.H. Bennett, Sci. Am. **257**, No. 5, 108 (1987); W.H. Zurek, Nature (London) **341**, 119 (1989); C.M. Caves, in *Complexity, Entropy, and the Physics of Information,* edited by W.H. Zurek (Addison-Wesley, Reading, 1991), pp. 91–115.
[11] R. Schack and C.M. Caves, Phys. Rev. Lett. **71**, 525 (1993); Phys. Rev. E **53**, 3387 (1996).
[12] S. Lloyd, Phys. Rev. A **39**, 5378 (1989).
[13] E. Ott, C. Grebogi, and J.A. Yorke, Phys. Rev. Lett. **64**, 1196 (1990); T. Shinbrot, C. Grebogi, E. Ott, and J.A. Yorke, Nature (London) **363**, 411 (1993).
[14] H.G. Schuster, *Deterministic Chaos* (VCH, Weinheim, 1995), 3rd ed.
[15] The few negative values appearing in the chaotic domains of the Lyapunov spectrum $\lambda(r)$ can be effectively suppressed either by adding a small noise component to the dynamical system or by a coarse graining of the state space. J.P. Crutchfield, J.D. Farmer, and B.A. Huberman, Phys. Rep. **92**, 45 (1982).
[16] A.N. Kolmogorov, Dokl. Akad. Nauk SSSR **119**, 861 (1958); **124**, 754 (1959).
[17] Ya.G. Sinai, Dokl. Akad. Nauk SSSR **124**, 768 (1959); **125**, 1200 (1959).
[18] V. Latora and M. Baranger, Phys. Rev. Lett. **82**, 520 (1999).
[19] H. Touchette and S. Lloyd (to be published).
[20] See *Maxwell's Demon, Entropy, Information, Computing,* edited by H.S. Leff and A.F. Rex (Princeton University Press, Princeton, New Jersey, 1990) for a review and references.