PNAS
Proceedings of the National Academy of Sciences of the United States of America www.pnas.org

Persistence of Nonoptimal Strategies

H. A. Ceccatto, and B. A. Huberman

PNAS 1989;86;3443-3446
doi:10.1073/pnas.86.10.3443

Proc. Natl. Acad. Sci. USA
Vol. 86, pp. 3443–3446, May 1989
Applied Mathematics

Persistence of nonoptimal strategies

H. A. CECCATTO* AND B. A. HUBERMAN†

Department of Applied Physics, Stanford University, Stanford, CA 94305; and Xerox Palo Alto Research Center, Palo Alto, CA 94304

Communicated by Leonid Hurwicz, February 17, 1989

ABSTRACT Metastable configurations in open computational systems with local minima in their optimality functions are shown to be very long lived, which makes them effectively stable. When rare transitions to the global optimum do occur, they happen extremely fast, in analogy to models of punctuated evolution in biology. These results are obtained by introducing a thermodynamic-like formalism that allows for a simple analysis of nonlinear game dynamics in computational ecosystems.

Open computational systems provide an important example of distributed computation that is both concurrent and asynchronous (for a collection of articles dealing with the subject, see ref. 1). Their slow emergence on the computing scene creates a number of interesting problems, which range from their programmability to understanding their global behavior in terms of local properties. In addition, the ability of processes to spawn others in remote computers and servers of the system, their competition for resources, and the lack of global controls makes for a community of computational agents that is reminiscent of biological and social organizations. These computational ecosystems are made up of processes that interact with each other by passing messages that embody arbitrary computations, making their interactions effectively infinite in their range.

The existence of computational ecologies brings to mind the spontaneous appearance of organized behavior in biological and social systems (2‡, 4), where agents can engage in cooperating strategies while working on the solution of particular problems. Such a phenomenon was found to emerge in the proverbial "prisoner's dilemma" when two-agent interactions with memory were allowed. Recently, spontaneous organization has also been shown to exist in open computational systems when agents can choose among many possible strategies while collaborating in the solution of computational tasks (5). Their dynamic behavior in the presence of delays and nonlinearities, however, while leading to stable attractors, excludes at times the possibility of having evolutionary stable strategies (ESS).

There are by now a number of distributed computational systems that exhibit many of the above characteristics and that offer increased performance when compared with traditional operating systems. ENTERPRISE, for example, is a market-like scheduler where independent processes or agents are allocated at run time among remote idle processors through a bidding mechanism (6). The system has been shown to provide substantial improvements over standard schedulers, even in the face of large delays and inaccurate estimates of processing times. Another, more sophisticated proposal includes algorithms that allow for both processor scheduling as an auction process and for distributed garbage collection through which unreferenced loops can be collected across trust boundaries (7).

An interesting problem in open systems is posed by the appearance of sudden changes in the nature of the network in which they are embedded. One may envision a computational ecology where agents have reached a fixed point consisting of mixed strategies, which are used in the solution of a complex and lengthy computational problem. The rapid availability of newer and/or faster computers or resources to the system can then introduce a different optimality criterion, because a new strategy mix may now increase the agent's performance in the solution of the problem. If the system is adaptable, one would expect it to move from the once nearly optimal strategy mix to the new one in the course of time.

In many situations of interest, however, the use of better computational resources in a network becomes economically feasible only when many processes have access to them. This introduces an interesting dilemma, since local changes in the previous strategy mix lower the performance of the system, thus preventing the appearance of spontaneous adaptive behavior. One may then ask the following question: given a bistable situation in a computational system, with one having a higher overall payoff than the other, how does the overall system relax towards the globally optimal one?

This paper presents a general solution to this problem. It does so by introducing a formalism analogous to statistical thermodynamics. By constructing a simple optimality function whose minima give the local fixed points of the system, we explicitly calculate the time evolution of a system that contains several basins of attraction for the game dynamics. We then show that, under fairly general conditions, the time it takes for a system to cross over from a local fixed point that is not optimal to a global one that is can grow exponentially with the number of agents in the system. This implies that metastable configurations in the adaptive landscape become effectively stable for large numbers of agents. When such a crossover does occur, however, it happens extremely fast (logarithmically in the number of agents), giving rise to a phenomenon analogous to punctuated equilibria in biology (8). A corollary of these results is that open systems with metastable strategies cannot spontaneously adapt to changing constraints, thereby necessitating the introduction of coordinating agents in order to do so. Finally, we comment on the applicability of these results to biological and social organizations.

Consider a collection of N computational agents capable of asynchronously choosing among several possible strategies, which, for the sake of simplicity in the presentation, we take to be two.§ We will denote by $G_1(n_1)$ and $G_2(n_2)$ the respective density-dependent payoffs, where n_i is the number of agents engaged in strategy i. To what extent agents choose given strategies depends on their perceived payoffs because, although there are many situations when cooperation between processes increases the speed at which a given problem is solved (thus increasing the payoff), crowding of resources by

Abbreviation: ESS, evolutionary stable strategies.
‡A more recent account of evolutionary game theory is presented in ref. 3.
§The formalism can be easily generalized to many more strategies, as shown in ref. 5.

many agents can reduce the payoff. A simple, but not necessarily unique, way of modeling competition between cooperation and crowding is by introducing a density-dependent payoff function with the form

$$G(n) = \left(P + C\frac{n}{N}\right)\left(R - \Gamma\frac{n}{N}\right), \quad [1]$$

where P is the payoff accrued in the absence of cooperation between agents and C/N is the extra benefit to be accrued due to cooperation. R characterizes the total capacity of a finite computational resource, and Γ/N is the amount by which the payoff decreases as each additional agent chooses to use the strategy.

By analogy with thermodynamics, we introduce an optimality function $\Omega(\beta, \mu)$, formally equivalent to a free energy, with β a continuous parameter ($0 < \beta < \infty$) playing the same role as that of an inverse temperature, which determines the amount of imperfect knowledge that individual processes have when making a choice among strategies. Perfect knowledge implies $\beta = \infty$, whereas maximal uncertainty is denoted by $\beta = 0$. The variable $\mu = 2n/N - 1$ is the deviation from an even choice among the two strategies. When all agents engage in one particular strategy $n = 0, N$ and $\mu = +1, -1$, whereas the symmetric case corresponds to half of the agents choosing each strategy (i.e., $n = N/2$ and $\mu = 0$). The minima of the function Ω determine the locally stable strategies of the system. Thus, if we call g the derivative of Ω with respect to μ (i.e., $g = d\Omega/d\mu$), the fixed points of the dynamics are determined by the condition

$$g(\beta, \mu) = 0, \quad [2]$$

with $g'(\beta, \mu) > 0$. In order to make use of this formalism, we need an explicit form for the dependence of the function $\Omega(\beta, \mu)$ in terms of the payoffs of the system and the uncertainty factor β. Let us denote by $\eta_1(\beta, \mu)$ and $\eta_2(\beta, \mu)$ the probabilities that an agent perceives strategy 1 to be better than strategy 2 and vice versa. In terms of the corresponding payoffs G_1 and G_2, one can write them as

$$\eta_1(\beta, \mu) = D \exp\{\beta[G_1(\mu) - G_2(\mu)]\} \quad [3a]$$

and

$$\eta_2(\beta, \mu) = D \exp\{\beta[G_2(\mu) - G_1(\mu)]\}, \quad [3b]$$

where D is a small proportionality constant.¶

With these definitions, the connection between the function g and the payoff function G is given by

$$g(\beta, \mu) = (1 + \mu)\eta_2(\beta, \mu) - (1 - \mu)\eta_1(\beta, \mu). \quad [4]$$

In order to demonstrate that the zeroes of the g function do indeed generate the stable equilibria of the system, we need to consider the dynamics of the system, which is governed by an equation that determines how the probability of having n agents engaged in strategy 1 at time t, $P_t(n)$, evolves in time. It is given by the master equation

$$P_{t+1}(n) = \sum_{m=n-1}^{n+1} P_t(m|n)T(m|n), \quad [5]$$

¶This is a plausible assumption for a wide class of payoff functions, and it gives the same results as those obtained with an error function in ref. 5. Notice that in the case where the ηs satisfy the normalization condition $\eta_1 + \eta_2 = 1$ the corresponding equations take the form $\eta_{1,2} = ½\{1 \pm \tanh \beta[G_1(\mu) - G_2(\mu)]\}$.

where the transition matrices, $T(m|n)$, which determine the rate at which agents switch from one strategy to the other (i.e., from strategy m to strategy n), are given by ref. 5:

$$T_\beta(n|n+1) = \frac{\alpha}{N}(N - n)\eta_1(\beta, \mu);$$

$$T_\beta(n+1|n) = \frac{\alpha}{N}(n + 1)\eta_2(\beta, \mu). \quad [6]$$

The prefactor α/N, which denotes the probability per unit time that an agent decides to evaluate the possible strategies, plays a similar role to that of an attempt frequency in the rate theory of thermodynamic processes.

We will now show that the equilibrium probability distribution, $P_e(\mu)$, defined by having $P_{t+1} = P_t = P_e$, is given by

$$P_e(\mu) = C \exp[-N\Omega(\beta, \mu)], \quad [7]$$

with Ω the optimality function defined above [$\Omega = \int_{-1}^{\mu} g(\beta, \mu)d\mu$] and C a normalization constant. It can be easily established that the condition $dP_e(\mu)/d\mu = 0$ is equivalent to Eq. **2**, with the function g determined by Eq. **4**. In other words, the probability distribution has sharp maxima at the zeroes of $g(\beta, \mu)$. This can be seen from the stationary solution of the master equation, Eq. **5**, which verifies the equality

$$P_e(n + 1)T_\beta(n + 1|n) = P_e(n)T_\beta(n|n + 1). \quad [8]$$

The solution of Eq. **8**, together with Eq. **6** is, up to a normalization factor, given by

$$P_e(n) \propto \begin{bmatrix} N \\ n \end{bmatrix} \prod_{r=1}^{n} \frac{\eta_1(\beta, (r-1)/N)}{\eta_2(\beta, (r-1)/N)}, \quad [9]$$

which, by taking logarithms and using Stirling's approximation, can be written (to leading order in N) in the same form as Eq. **7**. This completes our proof.

From these results it follows that we have converted the problem of finding the long time solution of the dynamical problem to that of finding the minima of Ω, a much simpler procedure that directly determines the ESS. To establish how the system flows towards the actual minima of Ω, we notice that the time evolution of the fraction of agents in a given strategy mix is given by ref. 5,

$$\frac{d\mu}{dt} = -\alpha g(\beta, \mu), \quad [10]$$

so that if $g(\mu)$ has only one zero (i.e., Ω has only a global minimum), the system will relax exponentially to that value of μ, with relaxation time $1/\alpha$. Thus, in adaptive systems, changes in the overall constraints of the system, which are reflected in a shift in the value of the Ω minimum, are followed exponentially fast. If, on the other hand, $g(\mu)$ has several zeroes, for short times the system will relax to the closest zero of g for which $g'(\mu) > 0$. At longer times, however, the system will reach the global minimum of Ω but with a slower relaxation rate, which we calculate below.

The usefulness of this formalism becomes apparent when considering problems with complicated tradeoffs between cooperation and resource utilization. Consider for the sake of simplicity a heterogeneous system with two types of resources, each of which generates a payoff G of the form given in Eq. **1**. If the density dependence of G_1 on μ is weaker than that of G_2 on μ, one can obtain a situation with more than one possible fixed point, as shown in Fig. 1. The associated optimality function is depicted in Fig. 1*b*, whereas $g(\mu)$ is shown in Fig. 1*c*.

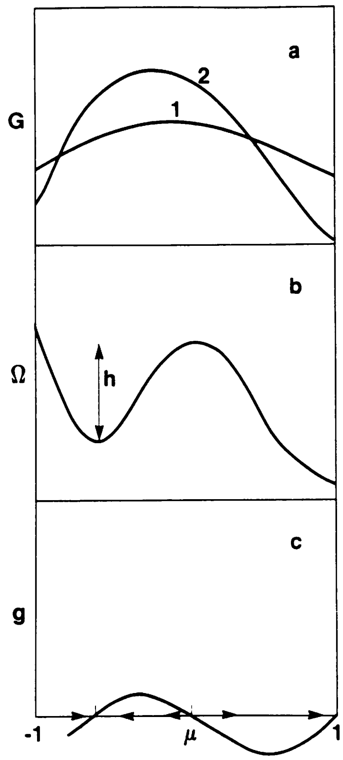

FIG. 1. (*a*) Payoffs for a heterogeneous open system with two types of resources, inducing different density dependence. (*b*) The associated optimality function. (*c*) Derivative of optimality function and associated dynamical flows.

Given any initial strategy mix, the system will quickly settle into the nearby minimum, which may or may not be the optimal one. The interesting case occurs when, due to changes in either constraints or initial characteristics, a given locally stable strategy is not an absolute minimum, for it then becomes relevant to ask how it spontaneously evolves into the global one.

We first note that in the absence of imperfect knowledge the system would always stay in the relative minimum and thus not adapt, for small excursions away from it reduce the local payoffs. It is only in the case of imperfect knowledge (i.e., wrong evaluation of the payoffs) that many agents can change from one strategy to the other. This is because, in evaluating the number of agents in the other strategy, imperfect knowledge amounts to the assumption that many have already moved. We therefore need to compute the time that it takes for a large number of them to change strategies in such a way so as to drive the system toward the new optimality minimum.

Such a calculation, analogous to particle decay in a bistable potential or to phase nucleation in thermodynamics, has been performed many times in the past, and the answer can be found in textbooks (see, for example, ref. 9; for an earlier calculation in the context of Ising models with infinite range interactions, see ref. 10). For a system with very long-range interactions, as is the case of agents in a computational ecology, the time, t, it takes for N agents to cross over from a metastable ESS to the optimal one is given by

$$t = \text{constant } e^{(\beta N h)}, \quad [11]$$

with h the height of the barrier, as shown in Fig. 1*b*. Therefore, for any large system with a finite amount of imperfect knowledge, this time becomes infinite, making the metastable state effectively stable. When large fluctuations do appear, however, the transit time from one minimum to the other is proportional to the logarithm of the number of agents, making this transition effectively instantaneous when compared to the actual time to produce the large fluctuation.

Such a process is clearly illustrated in Fig. 2, where we depict the results of a simulation by Kephart *et al.* (14) of a system with 10 agents capable of engaging in two strategies and with an overall payoff function of the form shown in Fig.

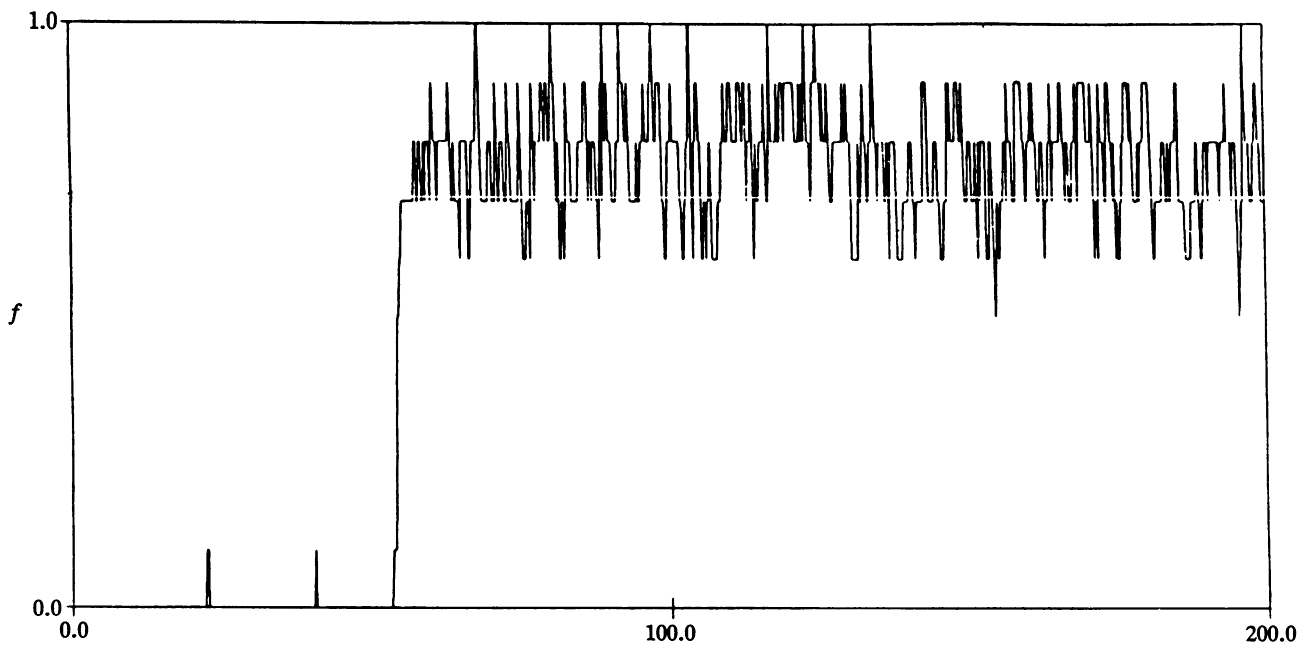

FIG. 2. Computer simulation of the evolution of a system with 10 agents asynchronously updating their choices among two strategies. The ordinate depicts the measured fraction of agents engaged in a given strategy mix for a system with the payoffs of Fig. 1. The x axis corresponds to time. The initial configuration has all agents in a pure strategy with a local minimum in the optimality function. Notice the small fluctuations before the sharp transition, due to the small size of the system.

1*a*. With all agents initially engaged in a strategy corresponding to the nonglobal minimum, the simulation monitored the fraction of agents, f, engaging in a given strategy mix as a function time. As can be seen, the system remained in the original configuration (i.e., $f = 0$) for a long time and made a very sharp transition to the optimal strategy mix ($f = 0.85$) in spite of its small size.

This scenario, which is very similar to that of punctuated equilibria in evolution (11–13), has a number of interesting implications. A most important one is that a large collection of computational agents in an open system will not spontaneously generate adaptive behavior when the introduction of novel constraints produces metastable configurations. In such situations, a global agent has to exist in order to (*i*) become aware of the advantage produced by another fixed point and (*ii*) to induce a coordinated action whereby processes simultaneously change their ESS. Without it, the evolution of open computational systems would be characterized by a number of unproductive strategy mixes with very few and rare collective transitions into more adapted ones. This lack of adaptability could pose serious disadvantages when compared to systems that either evolve with the help of global agents or are in close proximity to more optimal ESS.

When considering the implications of these results for biological or social organizations, an important point to notice is the effective range of the interactions among agents. For short-range ones, the time to escape a metastable ESS is much shorter than the one calculated above. This is because short-range interactions render the population with which an agent communicates effectively smaller, thus leading to a significant decrease in the value of N entering the exponential of Eq. **11**. For long-range communications, however, the effects of a local change in the strategy mix are broadcasted to all agents of the system, thereby making the effective radius of the nucleating droplet infinite and leading to the effects we discussed above. To the extent that a large community of social agents can have long-range interactions, and if no other mechanisms exist to lower optimality barriers, one would expect it to display nonadaptive behavior in the absence of global processes. This would not be so, however, for systems that are modularly organized into smaller and nearly independent units. In this case, an effective dynamical hierarchy appears as agents inside the units communicate with each other on shorter time scales than those involving agents belonging to different units. This leads in turn to effective cluster sizes small enough to spontaneously nucleate the global ESS at faster rates than the ones calculated in this paper.

Finally, we expect that similar metastable effects exist in complex dynamical systems with more complicated fixed points, such as oscillatory states or chaotic orbits, with the attractors playing the role of the local minima in the optimality function. In these cases, the crossover from metastable local attractors to globally stable ones will be long enough so as to make the system appear effectively frozen in given dynamical configurations.

We thank J. Kephart for providing us with the results of his simulation. This work was partially supported by ONR Contract N00014—82-0699. H.A.C. was also supported by a fellowship from Consejo Nacional de Investigaciones Cientificas y Tecnicas, Argentina.

1. Huberman, B. A., ed. (1988) *The Ecology of Computation* (North-Holland, Amsterdam).
2. Maynard-Smith, J. & Price, G. R. (1973) *Nature (London)* **246,** 15–18.
3. Maynard-Smith, J. (1986) *Physica* **22D,** 43–49.
4. Axelrod, R. & Hamilton, W. D. (1981) *Science* **211,** 1390–1396.
5. Huberman, B. A. & Hogg, T. (1988) in *The Ecology of Computation*, ed. Huberman, B. A. (Elsevier/North-Holland, Amsterdam), pp. 77–115.
6. Malone, T., Fikes, R. E., Grant, K. R. & Howard, M. T. (1988) in *The Ecology of Computation*, ed. Huberman, B. A. (Elsevier/North-Holland, Amsterdam), pp. 177–206.
7. Miller, M. & Drexler, E. (1988) in *The Ecology of Computation*, ed. Huberman, B. A. (Elsevier/North-Holland, Amsterdam), pp. 133–176.
8. Eldredge, N. & Gould, S. J. (1972) in *Models in Paleobiology*, ed. Schopf, T. J. M. (Freeman, San Francisco), pp. 82–115.
9. van Kampen, N. G. (1981) *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam).
10. Griffiths, R. B., Weng, C. Y. & Langer, J. S. (1966) *Phys. Rev.* **149,** 301–305.
11. Barton, N. H. & Charlesworth, B. (1984) *Annu. Rev. Ecol. Syst.* **15,** 133–164.
12. Lande, R. (1985) *Proc. Natl. Acad. Sci. USA* **82,** 7641–7645.
13. Newman, C. M., Cohen, J. E. & Kipnis, C. (1985) *Nature (London)* **315,** 400–401.
14. Kephart, J., Hogg, T. & Huberman, B. A. (1989) *Distributed Artificial Intelligence*, ed. Huhns, M. N. (Morgan Kaufmann, Los Altos, CA), Vol. 2, pp. 96–102.