Highly Optimized Tolerance: Robustness and Design in Complex Systems

J.M. Carlson
Department of Physics, University of California, Santa Barbara, CA 93106

John Doyle
Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125
(December 3, 1999)

Highly optimized tolerance (HOT) is a mechanism that relates evolving structure to power laws in interconnected systems. HOT systems arise, e.g., in biology and engineering, where design and evolution create complex systems sharing common features, including (1) high efficiency, performance, and robustness to designed-for uncertainties, (2) hypersensitivity to design flaws and unanticipated perturbations, (3) nongeneric, specialized, structured configurations, and (4) power laws. We study the impact of modifying percolation by incorporating increasing levels of design and find that even small amounts of design lead to HOT states, changing the behavior of the model completely.

PACS numbers: 05.40.+j, 64.60.Ht, 64.60.Lx, 87.22.As, 89.20.+a

Recently we introduced a mechanism for power laws in complex systems, which we refer to as *highly optimized tolerance* (HOT) [1]. The name is intended to reflect systems designed for high performance in an uncertain environment and operated at densities well above a standard critical point. In this letter we focus particular attention on properties associated with varying the level of design, roughly analogous to varying the rate of mutation in biological organisms or varying the extent to which alternative prototypes are tested prior to implementation of an incremental change in man made systems. Through design and evolution, HOT systems achieve *rare* structured states which are *robust* to perturbations they were designed to handle, yet *fragile* to unexpected perturbations and design flaws. Familiar examples include communication and transportation systems which are regularly modified to maintain high density, reliable throughput for increasing levels of user demand. However, as the sophistication of these systems is increased, engineers encounter a series of tradeoffs between greater productivity and the possibility of catastrophic failure. Such robustness tradeoffs are central properties of the complex systems which arise in biology and engineering. They also distinguish HOT states from the generic ensembles typically studied in statistical physics under the widely popularized scenarios of the "edge of chaos" (EOC) [2] and self-organized criticality (SOC) [3]. We present our results in the context of percolation. We find that even algorithms which are restricted to local explorations of a small fraction of the state space can lead to dramatic departures from a conventional critical phenomenon.

To illustrate HOT, we construct the simplest possible example. We begin with two-dimensional site percolation [4] on an $N \times N$ square lattice. We focus on spatial features, ignoring most aspects of the dynamics. We study the development of a designed configuration (see Fig. 1) as we incrementally increase the density by occupying sites one at a time. We make a loose analogy with forest fires, where occupied sites correspond to trees, and risk is associated

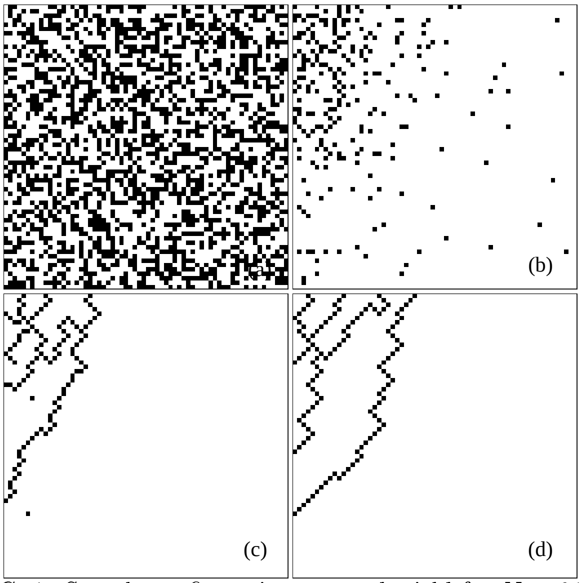

FIG. 1. Sample configurations at peak yield for $N = 64$ and varying values of the design parameter D, (a) the random case $D = 1$, (b) $D = 2$, (c) $D = N$, and (d) $D = N^2$. Unoccupied sites are black, and occupied sites are white.

with fires. A key quantity is the yield Y, defined to be the average density of trees left unburnt in a configuration after a single spark hits. If a spark hits an unoccupied site, nothing burns. When the spark hits an occupied site the fire burns every site in the connected cluster c of nearest neighbor occupied sites, where c is the number of sites lost relative to the system size. At each density we calculate Y for an ensemble of configurations, subject to a probability distribution $P(i, j)$ of sparks. We also obtain a distribution of the loss in individual fires $f(c)$ (the cumulative distribution of events of size greater than or equal to c is $F(c)$), such that $Y = \rho - < f >$, where $< f >$ is computed with respect to both the ensemble of configurations and the distribution $P(i, j)$. We increment the density by adding one occupied site to the previous configuration in its original state before a spark lands. Different levels of design are distinguished by different algorithms for choosing the next site for addition based on yield, reminiscent of systems in biology and engineering where incremental changes occur through natural selection or design modifications favoring higher yields. As we increase the density we obtain a yield curve $Y(\rho)$. At a given density the maximum possible yield is trivially bounded by ρ (no loss). In the thermodynamic limit, whenever $Y(\rho)$ falls below the maximum yield the mean event size $< f >$ is of the same order as the size of the system.

We first consider random percolation with no design. Sequences of configurations are generated by randomly occupying sites. At a given density all possible configurations are equally likely, and sites are independently occupied with probability $p = \rho$, and vacant with probability $1 - p$. By translation invariance, in the thermodynamic limit results for the random case are independent of the distribution of sparks $P(i, j)$. For finite systems edge effects modify the distribution.

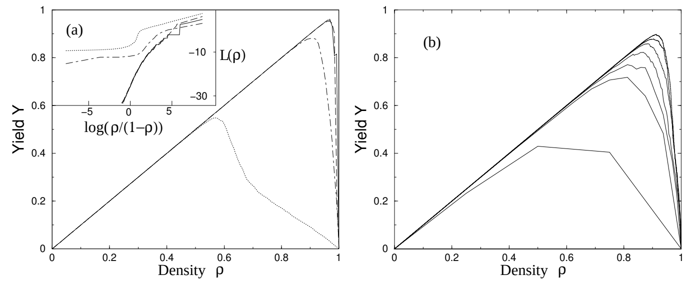

FIG. 2. Yield vs. Density $Y(\rho)$: (a) for design parameters $D = 1$ (dotted curve), 2 (dot-dashed), N (long dashed) and N^2 (solid) with $N = 64$, and (b) for $D = 2$ and $N = 2, 2^2, ...2^7$ running from the bottom to top curve. The results have been averaged over 100 runs.The inset to (a) illustrates corresponding loss functions $L = \log(< f > /(1- < f >))$, on a scale which more clearly differentiates between the curves.

The yield curve is depicted by the lowest curve in Fig. 2a. At low densities the results coincide with the maximum yield. Near $\rho = p_c \approx 0.6$ there is a crossover, and $Y(\rho)$ begins to decrease monotonically with ρ, approaching zero at high density. The crossover becomes sharp as $N \to \infty$ and is an immediate consequence of the percolation transition, marking the emergence of an infinite cluster at $p = p_c$. In the thermodynamic limit only events involving the infinite cluster result in a macroscopic loss and $Y(\rho) = \rho - {P_\infty}^2(p)$. Here $P_\infty(p)$ is the percolation order parameter, i.e., the probability a given site is in the infinite cluster. A typical random configuration at peak yield is illustrated in Fig. 1a. The fractal appearance of the clusters is a key signature of criticality. The distribution of fires $F(c)$ is asymptotically a power law, illustrated for $N = 64$ in Fig. 3a.

The goal of design is to push the yield towards the upper bound for densities which exceed the critical point. This requires selecting nongeneric (measure zero) configurations, which we refer to as *tolerant states*. We define HOT states to be those which specifically optimize yield in the presence of a constraint (see Figs. 1b-d). As the level of design increases, the connected clusters become increasingly regular in shape and separated by well defined barriers consisting of closed contours of unoccupied sites. A HOT state corresponds to a forest which is densely planted to maximize the timber yield, with fire breaks arranged to minimize the spread of damage.

To determine HOT states we specify constraints on the optimization, defined here in terms of a modification of the random percolation model which incorporates design. At each increment in density, we generate a set of test configurations. A test configuration at density $\rho = (n + 1)/N^2$ is generated from the previous configuration at density $\rho = n/N^2$ (where n is the number of occupied sites) by adding a grain to a previously unoccupied site. In the limiting case, all possible test configurations are explored. Then, in a manner reminiscent of evolution by natural selection, the next configuration in the sequence is taken to be the test configuration which produces the highest yield upon averaging over $P(i, j)$. If there is degeneracy, one of the highest yield configurations is selected at random. We define the design parameter to be the number D of test configurations which are initially examined. Thus D lies between $D = 1$ (random percolation) and $D = N^2$ (all possible configurations associated with the addition of one occupied site are tested).

While the choice of $P(i, j)$ is irrelevant at large N for random configurations, for designed states knowledge of $P(i, j)$ can be exploited to produce better designs. For our numerical examples we use:

$$\begin{aligned} P(i, j) &= P(i)P(j) \\ P(x) &\propto 2^{-[(m_x + (x/N))/\sigma_x]^2} \end{aligned} \tag{1}$$

where $m_i = 1$, $\sigma_i = 0.4$, $m_j = 0.5$ and $\sigma_j = 0.2$. In Fig. 1 the maximum value of $P(i, j)$ coincides with the upper left hand corner $i = j = 1$, while the minimum value is in the lower right corner $i = j = N$. We choose the tail of a Gaussian to dramatize that power laws emerge through design even when the external distribution is far from a power law. We choose an asymmetric distribution to lift all degeneracies in the maximally designed case $D = N^2$. This leads to a deterministic sequence of uniquely determined locally optimal designs through the full range of densities.

For $D = N^2$, initially the low $P(i, j)$ sites are filled, in what begins as a checkerboard pattern. At higher densities, the upper left corner remains a checkerboard, while regions near the lower right corner form a cellular pattern. The cells consist of compact regions enclosed by well defined barriers which are one lattice site wide. As the density increases further, cells begin to merge. When a barrier between neighboring cells is broken, the entire barrier is filled in before other sites are occupied. Fig. 1d illustrates the configuration at the maximum yield point.

As shown in Fig. 3a, the distribution of events $F(c)$ exhibits a power law tail when $P(i, j)$ is given by Eq. (1). This is not the case for all $P(i, j)$. If $P(i, j)$ is uniform, the system evolves towards cells of equal size, while if the position of the next spark is known, there is no fire. However, we obtain power laws for a broad class of $P(i,j)$, including Gaussian, exponential, and Cauchy.

There are several noteworthy distinctions between the distributions obtained for the cases $D=1$ (random) and $D=N^2$. Asymptotically both are characterized as power laws $F(c)\sim c^{-\alpha}$. However, for the random case asymptotically the cumulative distribution is very flat ($\alpha\approx 0.055$), while for the designed case $\alpha\approx 1$. Furthermore, averaged over the ensemble of random configurations, $F(c)$ is smooth and continuous, while for the designed case the non-cumulative distribution $f(c)$ consists of a discrete set of points (marked by *'s), associated with the particular set of optimal cells obtained for the deterministic algorithm. The sizes of consecutive cellular regions c increases with increasing event size, at a rate which scales inversely with the rate of decay of the tails of $P(i,j)$. This increase in size balances the decay of $P(i,j)$ to produce a power law $F(c)$ in the HOT state.

Interestingly, in the tolerant regime our algorithm produces power law tails for a range of densities below the maximum yield, and without ever passing through a state that resembles the (fractal) critical state. This is illustrated in Fig. 3b where we plot the discrete cumulative event size distributions $F(c)$ along with a least squares fit to the power law for equal density increments when $D=N^2$. Clearly for very low densities, the event size distribution does not extend to large events because there are so few occupied sites. However, even in that regime, over the limited range of sizes which are present, the distribution is roughly a power law. At higher densities, the power law extends to large events for a broad range of densities. In this regime, the key distinction between the distributions is a relative suppression of the amplitude of the tails for lower densities. At moderate densities, the largest domain is essentially formed, while the neighborhood of the upper left hand corner retains the checkerboard pattern, contributing a sharp spike in $F(c)$ at $c=1/N^2$. Thus backing away from the maximum yield point (a strategy commonly employed in engineering), the distribution retains heavy tails, but the overall rate of large events decreases. The fact that power laws are not a special feature associated with a single density is in sharp contrast to a traditional critical phenomena.

Next we turn to the case of intermediate values of the design parameter, $1<D<N^2$, so that only a subset of the possible test configurations is considered at each density increment. A sequence of maximum yield configurations for increasing values of D is illustrated in Fig. 1. Already for $D=2$, there is a clear preference for occupying sites in the lower right corner of the lattice, while the boundaries between regions are poorly formed compared to the deterministic case. However, for $D=N$, the cellular structure of the HOT state is clear, with only a few "flaws," corresponding to unoccupied sites in the middle of regions rather than on a boundary. Because the algorithm is stochastic when $D<N^2$ the specific cellular pattern corresponding to the maximum yield point for a given realization depends on the history and differs from that obtained in the deterministic design.

Yield curves for $D=1,2,N$, and N^2 are included in Fig. 2a. On average, the peak yield increases monotonically, and the drop off beyond the maximum yield point becomes increasingly sharp with increasing D. We observe instances of higher yields for specific configurations in cases with less design since the algorithm is locally, rather than globally, optimized. Furthermore, even for relatively small values of D the maximum yield is quite close to that obtained for the deterministic case.

We define a loss function $L=\log(<f>/(1-<f>))$, which is plotted as function of $\log(\rho/(1-\rho))$ in the inset to Fig. 2a. Compared to $Y(\rho)$, these curves more clearly reflect the differences for different values of D at low and high densities. The loss is greatest in the random case for the full range of ρ, including low densities where there is no discernible distinction in Y because $<f>$ is not macroscopic. Interestingly, both Y and L nearly coincide for $D=N$ and N^2 (the curve for $D=N^2$ is more jagged because the deterministic case is not smoothed by an ensemble average over the selection of test sites). While at most densities $D=2$ has greater losses compared to $D=N$ or N^2, the curves cross at a density beyond the maximum yield points, indicating that the lower level of design $D=2$ is more robust in the regime which is dominated by large failure events.

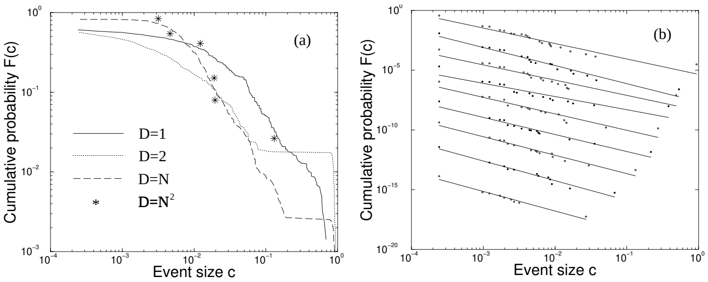

FIG. 3. Cumulative distributions of events $F(c)$: (a) at peak yield for $D=1,2,N$ and N^2 with $N=64$, and (b) for $D=N^2$, and $N=64$ at equal density increments of 0.1, ranging at $\rho=0.1$ (bottom curve) to $\rho=0.9$ (top curve).

The corresponding distributions of events are illustrated in Fig. 3a. As in the case $D=N^2$, the distributions $F(c)$ are discrete power laws when averaged over $P(i,j)$ for each realization of the stochastic designs. As illustrated in Fig. 3a, for $D=2$ and N when we average over the ensemble of configurations the distribution becomes smooth for smaller values of c, with a discrete jump separating the largest events. Here the rapid growth in the separations between points in the discrete $F(c)$ is sufficient to prevent smoothing of the range of largest events in the ensemble average. Comparing different values of $D>1$, we see that the relative weight of the largest event is increasingly suppressed as the level of design is increased from $N=2$ to N to N^2, indicative of a sequence of improved designs.

As N increases (refining the grid so that the range of $P(i,j)$ is held constant), the maximum yield associated with each value of D increases systematically towards $Y = 1$. This is illustrated for $D = 2$ in Fig. 3b. Even the case $D = 2$ corresponds to a diverging number of design choices as $N \to \infty$. As a consequence, the relative cost in yield of occupying less favorable sites is reduced for larger N.

Because random states dominate the set of all possible configurations, properties of the measure zero subset of HOT states are ignored in conventional statistical physics. However, in the limit of large N we achieve high yield HOT states for a variety of (measure zero) design scenarios. This is especially true in the percolation model where barriers can be designed at cost which vanishes in the limit of large N: the cost of a barrier scales like N compared to the yield of a region which scales like N^2. As a result even regular grids lead to high yields for large N, regardless of $P(i,j)$. A regular grid design does not optimize the placement of the grid lines for a finite system unless $P(i,j)$ is uniform.

The special design features which the HOT states develop through the local evolutionary algorithm are high concentrations of occupied sites in regions of low $P(i,j)$ and cellular structures separated by barriers. These features lead to robust, high yield states at densities well above the random critical point. For a given density the expected event sizes associated with HOT states are much smaller than those of random configurations. In Fig. 3a the random case exhibits the flattest distribution with events of the largest average size, in spite of the fact that it corresponds to the lowest density. However, the design features of HOT states also lead to new sensitivities which increase with increasing levels of design. For example, HOT states are sensitive to changes in $P(i,j)$. If the system is designed for a Gaussian $P(i,j)$ (e.g. Eq. (1)) but then subject to a uniform distribution of hits, then the distribution of events increases with the size of the region: $f(c) \sim c$, resulting in the most dramatic change in the distribution of events when $D = N^2$. In contrast, the random critical state is essentially independent of $P(i,j)$. Furthermore, HOT states are highly sensitive to design flaws. If a vacant site on the barrier between neighboring regions is filled then an event triggered on either side penetrates the barrier joining surrounding regions. In contrast, in random configurations small changes do not alter the distribution of events.

Percolation seems to be the simplest template for introducing HOT states and contrasting their properties with criticality. Even the random case can be viewed as a very primitive design with density as the *only* design parameter. In this case, the critical point coincides with the maximum yield, making this a natural alternative to SOC whereby primitive systems might evolve to criticality. More importantly, adding even modest levels of additional design, such as choosing one of two random sites at each density, moves yields well above the random critical point.

The complexity which arises in EOC/SOC, as well as more generally in statistical physics, is a generic property. Features associated with sets of measure zero in configuration space are disregarded. This traditional approach leaves out some of the most fundamental properties of the complex systems which arise in biology and engineering, and which tend to have all the features of the HOT state. Here only sets of measure zero are retained as systems are refined through design or evolution towards states of increasing complexity. As our simple percolation model suggests, such designs need not be deterministic or globally optimal, but can result from stochastic local and incremental evolution that characterizes natural selection as well as most engineering innovation.

In advanced systems, designed features are so dominant and pervasive that they are easily taken for granted. While generic complexity emerges from a featureless substrate, the complexity in designed systems often leads to apparently simple, predictable, robust behavior. As a result, designed complexity becomes increasingly hidden, so that its role in determining the sensitivities of the system tends to be underestimated by nonexperts, even those scientifically trained. Complex systems in engineering and biology are dominated by robustness tradeoffs, which result in both high performance and new sensitivities to perturbations the system was not designed to handle. Domain specific models of such systems are often complicated, with a great deal of built-in structure included to reproduce the extreme robustness and sensitivities of these systems. Thus while the forest fire analogy in this paper is useful pedagogically, our models have limited direct relevance to real forest fires. Nevertheless, rare, structured, HOT states must be the building blocks of any theory of complex systems which is of relevance to biology and engineering.

ACKNOWLEDGMENTS

This work was supported by the David and Lucile Packard Foundation, NSF Grants No. DMR-9510394 and DMR-9813752, and a DOD MURI grant for "Mathematical Infrastructure for Robust Virtual Engineering."

[1] J. M. Carlson and John Doyle, *Phys. Rev. E* **60**, 1412 (1999)
[2] R. Lewin, *Complexity - Life at the Edge of Chaos,*(Macmillan, New York, 1992); S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, (Oxford Univ. Press, New York, 1993).
[3] P. Bak, C. Tang, K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987); P. Bak, *How Nature Works: The Science of Self-Organized Criticality*, (Copernicus, New York, 1996).
[4] D. Stauffer, *Introduction to Percolation Theory,* (Taylor, London 1985).