Design degrees of freedom and mechanisms for complexity

David Reynolds and J. M. Carlson
Department of Physics, University of California, Santa Barbara, California 93106

John Doyle
Control and Dynamical Systems, California Institute of Technology, Pasadena, California 91125
(Received 15 October 2001; published 15 July 2002)

We develop a discrete spectrum of percolation forest fire models characterized by increasing *design degrees of freedom* (DDOF's). The DDOF's are tuned to optimize the yield of trees after a single spark. In the limit of a single DDOF, the model is tuned to the critical density. Additional DDOF's allow for increasingly refined spatial patterns, associated with the cellular structures seen in highly optimized tolerance (HOT). The spectrum of models provides a clear illustration of the contrast between criticality and HOT, as well as a concrete quantitative example of how a sequence of robustness tradeoffs naturally arises when increasingly complex systems are developed through additional layers of design. Such tradeoffs are familiar in engineering and biology and are a central aspect of the complex systems that can be characterized as HOT.

DOI: 10.1103/PhysRevE.66.016108 PACS number(s): 05.65.+b, 64.60.Ht, 05.45.−a

I. INTRODUCTION

Complexity and design are common concepts that describe many aspects of our everyday experiences, yet both concepts have generated much controversy in the scientific literature. Design is typically associated with a planned or deliberate selection based on performance goals, as in high technology and engineering [1]. However, design also plays an important role in management and policy for both social and natural systems. Evolution by natural selection can be thought of as nature's mechanism for design, whereby high-fitness organisms are selected over poor performers, without requiring the intervention of a deliberate designer [2]. While it is widely recognized that chance, evolution, and design all play roles in complex technological, social, and biological systems, their relative roles remain controversial. Even for engineering systems, where design is deliberate and performance goals are relatively straightforward to quantify, it is difficult to assess the extent to which designs are close to optimal, since even the simplest mathematical descriptions are often computationally intractable. A further challenge is to describe the role and extent of design in a way that is clear and systematic, yet does not rely on detailed knowledge of the design process.

The relationship between design and complexity has similar ambiguities and controversies. One popular notion of complexity involves systems that exhibit variable features spanning a range of scales in time and space. Broadly speaking, there have been two extreme theoretical points of departure from this notion of complexity, which are differentiated in terms of their stance regarding the role of design. At one extreme is the point of view that complexity ''emerges'' in systems that are otherwise internally homogeneous and simple [3]. This view underlies the theory of self-organized criticality (SOC) [4–6], and the edge of chaos [7] and suggests that large-scale structure arises naturally and at no apparent cost through collective fluctuations in systems with generic interactions between individual agents. In this scenario structure is associated with bifurcation points and critical phase transitions. ''Self-organized'' describes systems that are dynamically attracted to such transition states.

In the alternative point of view, complexity is associated with intricately designed or highly evolved systems. This notion of complexity is familiar in biology and engineering, and plays an important role in the extensive literature developed in the artificial life (ALife) community. ALife studies how increasingly complex computer organisms evolve through competition [8]. More recently, a theoretical framework referred to as highly optimized tolerance (HOT) was introduced [9–11], which emphasizes the role of robustness to uncertainties in the environment as a driving force towards increasing complexity in biological evolution and engineering design. It is the HOT framework on which this paper is based.

The suggestion that robust design is the primary mechanism for complexity is motivated by the observation that for most biological and technological systems, the vast majority of components are present for robustness rather than for basic functionality of the organism or machine [12]. Designed systems are typically internally extremely heterogeneous and complex, and this complexity is introduced to create simple, reliable, robust external behavior, despite uncertainty in component parts and in the environment. In designed systems, internal complexity is used to minimize external complexity. This deviates from emergent complexity, which emphasizes how internally simple systems can yield externally complex behavior. Nonetheless, due to the broad spectrum of environmental uncertainties, and the robustness and performance tradeoffs that are inherent in evolution and design, systems with ''designed complexity'' exhibit external structure and failure modes over a broad range of scales, albeit at higher densities and throughput than those associated with the ''emergent'' systems. As we have described elsewhere in more detail [9–13], emergent and designed complexity are extremely different, and we have argued that the latter is much more relevant to biological, social, and technological systems where design and evolution amidst environmental uncertainty plays a role.

The aim of this paper is to determine how the characteristics of designed systems change as the resolution of the design is varied. Two questions then are central to this line of research. Can systems be classified in terms of their relative complexity? Is it possible to quantify the amount of design? For real systems, this defines an inverse problem. The goal would be to determine the extent to which a system deviates from a generic, random configuration, or some other null hypothesis, by identifying a series of performance goals (inequalities) that measure these deviations, along with tunable parameters, which can be used to narrow the set of acceptable configurations from the generic, random starting point, towards selected states that satisfy the performance goals. However, we are not yet at a point where we can solve the inverse problem for real systems. Instead, we propose a forward design problem in the context of a lattice model and introduce a single, simple framework for optimization of yield (the performance goal) with respect to a specified number of *design degrees of freedom* (DDOF's), which represent the tunable parameters. Varying the number of DDOF's allows us to interpolate between systems with minimal design, and those that are highly designed.

We investigate the consequences of increasing DDOF's in the context of the percolation forest fire model [14], which is a variant of site percolation [15] on the two-dimensional square lattice. A spark impacts a standard percolation configuration on an individual site, and burns the associated connected cluster. This model has been studied previously in the context of SOC (Refs. [14] and [16]) and HOT [9–13,17,18]. By creating a spectrum of models that interpolates between minimal and maximal design limits we can contrast the features associated with criticality and HOT, and the tradeoffs that lead to incremental shifts from structure that is associated with fluctuations and emergence to that associated with deliberate, optimized layouts. Most importantly, we hope to create a model system in which design and necessity have clear meanings, albeit severely abstracted from their meaning in engineering and biology. For our model system, the notions of emergence, complexity, and order have natural interpretations, which will allow us to clarify some of the potential ambiguities.

For the full spectrum of models, we assume that the lattices are chosen to maximize yield in an external environment. The measure of yield, and the external environment remain fixed as we vary the number of tunable parameters. When the DDOF's are optimized for the yield of ''trees'' after a single ''spark,'' we obtain random configurations at the critical density in the limit of a single DDOF, and a fully optimized HOT configuration, consisting of compact connected clusters separated by well defined linear firebreaks, in the limit of infinite DDOF's. Interestingly, in our formulation the intermediate cases do not exhibit a gradual crossover, in the sense of, say, a smooth variation of the density, or a gradual shift in the power law or the typical size or shape of events. Instead, the design variables, which correspond to local densities in different regions in space, always converge to either critical or unit density. This occurs even when the number of DDOF's is as low as two. As the DDOF's increase, different portions of the lattice shift from a density that approaches the critical density to unit density. In the limit of large numbers of DDOF's near critical regions form increasingly narrow linear barriers that bound compact unit density regions characteristic of HOT.

Studying the system as a function of the number of DDOF's provides a concrete, quantitative measure of the structured sensitivity that is a central feature of the *robust, yet fragile* HOT systems. As the number of DDOF's increases, the system becomes increasingly robust to common perturbations: the average yield increases, as does the ability of the layout to resolve features in the distribution of sparks. The system also becomes increasingly fragile to rare events, changes in the distribution of sparks, and flaws in the design. This occurs because the increasing density puts more of the system at risk for potential failure, and narrower barriers imply increasing sensitivity to design flaws.

The remainder of this paper is organized as follows. In Sec. II we describe the percolation forest fire model, and introduce our scheme for increasing the number of design degrees of freedom.

In Sec. III we describe a sequence of numerical results for increasing DDOF's. While optimization of a system with just a few design parameters can be computed brute force, constrained optimizations are required to extend our results to large numbers of DDOF's.

In Sec. IV we describe the results of analytical calculations for a spatially uniform distribution of sparks. HOT does not lead to power laws in this case. Instead, with large numbers of DDOF's the system breaks up into equal size HOT regions of unit density, separated by linear barriers approaching the critical density.

We conclude in Sec. V with a summary of our results. We also discuss alternative schemes of varying the number of DDOF's, and the role of DDOF's in more realistic settings.

II. THE PERCOLATION FOREST FIRE MODEL WITH DESIGN DEGREES OF FREEDOM

The percolation forest fire model consists of a two-dimensional $N \times N$ lattice. Each site is either occupied by a tree or is vacant, and each contiguous set of nearest neighbor occupied sites defines a connected cluster. The forest is subject to external perturbations, represented by sparks. When a spark hits a vacant site on the lattice nothing happens. When a spark hits an occupied site it burns all the trees in the connected cluster associated with the site.

The impact site (i,j) for the spark is drawn from a probability distribution $P(i,j)$. If $P(i,j)$ is uniform, then each site is equally likely to be hit. If $P(i,j)=\delta(i-i_0)\,\delta(j-j_0)$ then site (i_0,j_0) is hit with absolute certainty. All other distributions lie between these two extremes, and represent the more realistic case of variable risk, where ignitions are common in some regions and rare in others.

In random percolation [15], the state of the system is fully characterized by the density ρ. Individual sites are independently occupied with probability ρ, and vacant with probability $1-\rho$. Properties of the system are determined by ensemble averages in which all configurations at density ρ are taken to be equally likely. In the thermodynamic limit trans-

lational invariance of the ensemble renders the choice of $P(i,j)$ completely irrelevant, and the density is *a priori* the only tunable parameter. However, when additional tunable parameters (defined below) allow for more detailed resolution of the spatial layout of the lattice, then the form of $P(i,j)$ plays a key role in determining the optimal state.

We define yield Y to be the number of trees remaining after a single spark,

$$Y=\rho-\langle l\rangle. \tag{1}$$

Here ρ is the density before the spark and $\langle l\rangle$ is the average loss due to the fire, computed over the distribution of sparks $P(i,j)$ as well as the configurations in the ensemble. We optimize yield as a function of the tunable parameters, given a distribution of sparks $P(i,j)$.

The number of DDOF's is a count of the parameters that are deliberately tuned to select the configuration or ensemble of configurations that define the state of the system. For random percolation, by construction there is only one DDOF—the density ρ. Once ρ is set the state of the system is defined by the ensemble of all possible configurations at density ρ. In the opposite extreme, if we specifically choose whether each site individually is occupied or vacant on an $N\times N$ lattice, then we have N^2 DDOF's, which diverges in the limit $N\rightarrow\infty$. For large lattices, computing the globally optimal configuration with respect to N^2 DDOF's rapidly becomes numerically intractable, requiring consideration of 2^{N^2} candidate lattices in order to select the best configuration. Constraints may be imposed on the optimization to restrict either the number of DDOF's, or the search space for optimal configurations, or both.

Previously a variety of constrained optimization schemes were considered in the context of HOT, which preserve N^2 DDOF's on an $N\times N$ lattice, but restrict the search for optimal configurations. One example, which we modify for finite DDOF's in Sec. III D, corresponds to a local incremental algorithm for increasing the density. This algorithm consists of a local optimization in configuration space. Sites are occupied one at a time, always choosing the next site to occupy in order to maximize yield for the incremental change in density [10]. The result is a sequence of unique configurations [when $P(i,j)$ is chosen to avoid degeneracies] of increasing density, which trace over the full interval $\rho\in[0,1]$. In the case of degeneracies (i.e., two choices of the next site to occupy produce the same yield), a site is selected randomly from the candidates, and the process continues until the lattice is fully occupied. Collecting the complete set of configurations that arises when the random choices associated with degeneracies are taken into account results in an ensemble of possible configurations at each density. With or without degeneracies, configurations selected by the local incremental algorithm correspond to a set of measure zero in the space of possible configurations at a given density. The ensemble defines a yield curve, $Y(\rho)$, which has a maximum at some $\rho=\rho_{\max}$. Beyond this value there is a sharp drop in yield. In the thermodynamic limit, $\rho_{\max}$ approaches unity. Figure 1 contrasts an example of the maximum yield point obtained by this algorithm with a typical random configuration at the critical density for a 64×64 lattice. For a more complete description of this model see Ref. [10].

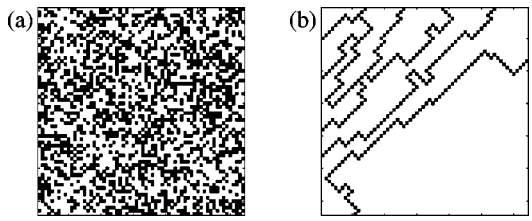

FIG. 1. Sample configurations for the (a) critical and (b) HOT percolation forest fire models on a 64×64 lattice. White sites are occupied, and black sites are vacant. The critical configuration represents a sample from the ensemble of randomly generated configurations at the critical density $\rho_c\approx 0.59$. The HOT configuration is obtained for the (constrained) local, incremental optimization scheme. HOT configurations have much higher densities and are stylized for the distribution of sparks $P(i,j)$ given in Eq. (2), which is sharply peaked in the upper left corner of the lattice, where the vacancies are most concentrated, and form linear fire breaks.

In the local incremental algorithm, there are of order N^2 DDOF's (since sites are individually assigned to be occupied or vacant), but the number of configurations that are searched through to locate a maximum are far less ($<N^4$) than the 2^{N^2} required for a global search. In spite of the fact that the local incremental algorithm samples fewer configurations, both the brute force global optimization and the local incremental algorithm lead to highly designed systems, with many common properties. Both lead to power law distributions of fire sizes for a broad class of $P(i,j)$, densities, and yields that approach unity in the limit of large system sizes, and compact connected clusters of trees separated by linear barriers. In both cases the number of design degrees of freedom diverges as $N\rightarrow\infty$.

Finally, we consider intermediate numbers of DDOF's. There are many ways to interpolate between one and infinite DDOF's. We consider a particular choice below, which is convenient and tractable. We believe the basic trends produced by our scheme are characteristic of the general problem of incrementing DDOF's in percolation (this is discussed in more detail in the conclusion), although in the future we plan to explore alternatives, especially those relevant to specific applications.

Our scheme is based on subdividing the $N\times N$ lattice into equal square cells (see Fig. 2). This defines an $M\times M$ lattice, with each cell containing $n^2=(N/M)^2$ sites. Individual cells are characterized by a density ρ_{IJ}, where (I, J) defines the cell coordinate on the $M\times M$ design lattice.

The (I, J) sites are a coarse graining of the (i, j) coordinate system describing the underlying lattice of vacant and occupied sites. Our scheme is loosely analogous to an inversion of the traditional real space renormalization employed in statistical mechanics to analyze critical phenomena [19]. However, rather than starting with a configuration on the underlying lattice, and rescaling to obtain an equivalent density for the coarse grained version ultimately culminating in a fixed point, here we individually prescribe (and eventually optimize) a density for each design cell of the $M\times M$ lattice, which determines the ensemble of allowed configurations on

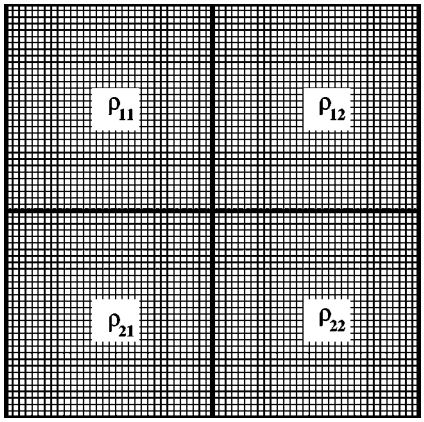

FIG. 2. Definition of the $M\times M$ design lattice (here $M=2$), which is superimposed on the underlying $N\times N$ percolation lattice of vacant and occupied sites. Each cell in the design lattice is characterized by a density ρ_{IJ} and all configurations at density ρ_{IJ} are equally likely in the cell. The set of $\{\rho_{IJ}\}$ with $I, J\in[1,M]$ define the M^2 DDOF's.

the portion of the underlying lattice enclosed in the cell. When ρ_{IJ} takes a particular value in a cell on the $M\times M$ design lattice, then all microscopic configurations associated with the n^2 primitive sites on the underlying $N\times N$ lattice described by that density are equally likely in the cell. In other words, sites on the microscopic lattice within the cell are independently occupied with probability ρ_{IJ}, and properties associated with the cell at that density are defined by the ensemble average, just as in random percolation, but restricted to the subspace of the underlying lattice that is defined by the (I, J) cell. It is possible to take $N\to\infty$, yet keep M finite. In that case, the individual cells on the design lattice are described by the thermodynamic limit of standard random percolation at the prescribed cell density.

III. NUMERICAL RESULTS

In this section we present the results of a sequence of numerical computations, which illustrate the effects of increasing design degrees of freedom. We begin with brute force calculations on finite lattices and limited DDOF's. Our observations lead to insights that simplify calculations in the limit of an infinite underlying system. For small numbers of DDOF's we compute globally optimal solutions. For large numbers of DDOF's we introduce a modified version of the local incremental algorithm on the design lattice.

For an $M\times M$ design lattice superimposed on an $N\times N$ underlying lattice, the sequence of computations is summarized as follows:

(A) $M=1$. We begin by reviewing the case of a single DDOF, where optimization of the yield leads to criticality.

(B) *N finite, M small*. This corresponds to the case of a finite underlying lattice and just a few DDOF's. In this case we compute the optimal subregion densities of the $M\times M$ lattice by averaging over random sample configurations at different sublattice densities. Interestingly, this calculation shows that as N gets large, the subregion densities converge to either unit density or a density that approaches the critical density ρ_c. (The "critical" regions are bounded away from ρ_c, for technical reasons we will discuss this later.)

(C) $N\to\infty$, *M small*. We use the fact that ρ_{IJ} converges to $\rho_c-\epsilon$ (ϵ infinitesimal) or unity to simplify our search of the state space in the thermodynamic limit. This allows brute force calculation of the globally optimal configuration for $M\leqslant 5$ on an $M\times M$ coarse grained lattice. We begin to observe the design lattice breaking up into compact domains of unit density, separated by uncrossable barriers of density $\rho_c-\epsilon$.

(D) $N\to\infty$, *M large*. To extend our results to larger design lattices, we modify the local incremental algorithm considered previously in the context of finite lattices, to the case of the infinite underlying lattice, but a finite design lattice. Here cellular patterns similar to those previously associated with HOT clearly emerge on the design lattice.

We have verified that the qualitative results presented here hold for a range of spark distributions, including Gaussian, exponential, and Cauchy. However, for consistency in this paper, unless otherwise specified, we present our numerical results for the same exponential distribution, scaled in a manner that allows us to directly compare results for different M and N. Specifically, we define the coordinate system so that the origin lies in the upper left-hand corner of the lattice, and that x increases moving horizontally towards the right, and y increases moving vertically downward from that point. We scale the discrete underlying lattice into the unit square, so that $x=i/N$, $i=1,\ldots N$, and $y=j/N$, $j=1,\ldots N$, and take

$$P(x,y)=A\exp\{-[(x/\eta_x)+(y/\eta_y)]\}, \tag{2}$$

where η_x and η_y are characteristic lengths for the decay of the distribution along the x and y axes, respectively. We take $\eta_x=\frac{1}{10}$, and $\eta_y=\frac{1}{4}$, deliberately breaking the symmetry of the distribution to avoid degeneracies. The constant A is determined by normalization. To compute the probability of a spark hitting a particular site (i, j) on the discrete lattice, we integrate the continuous distribution over the square corresponding to the site,

$$P(i,j)=\int_{(j-1)/N}^{j/N}\int_{(i-1)/N}^{i/N}P(x,y)\,dx\,dy. \tag{3}$$

To compute the total probability of hitting a cell (I, J) in the design lattice, we compute the corresponding integral over the area of the design cell,

$$P(I,J)=\int_{(J-1)/M}^{J/M}\int_{(I-1)/M}^{I/M}P(x,y)\,dx\,dy. \tag{4}$$

Specific designed configurations depend on the details of $P(x,y)$ up to the resolution of the $M\times M$ design grid.

A. Criticality—the optimal solution for a single DDOF

We begin with a single design degree of freedom, $M=1$. The state of the system is characterized by the density ρ,

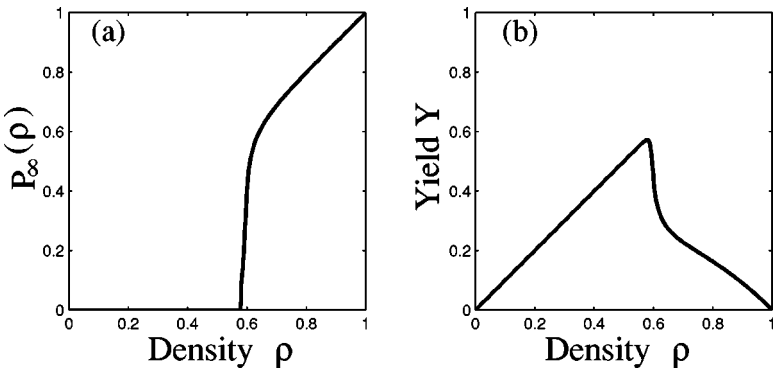

FIG. 3. The onset of percolation and the maximum yield point coincide for a single design degree of freedom. (a) illustrates the percolation probability $\mathcal{P}_\infty(\rho)$ that is simply related to yield $Y=\rho-\mathcal{P}_\infty(\rho)^2$ illustrated in (b).

and all configurations of density ρ are weighted equally. In the limit of large N, the choice of $P(x,y)$ is irrelevant because the ensemble is *a priori* translation invariant. The results for $M=1$ follow directly from known numerical results for site percolation on a square lattice, which we summarize below.

The model exhibits a continuous phase transition at density $\rho_c \approx 0.592$, which is associated with the emergence of an infinite connected cluster. In the limit $N\to\infty$, for $\rho<\rho_c$ there is no infinite cluster. For $\rho>\rho_c$ an infinite cluster exists somewhere on the lattice with probability one. At $\rho=\rho_c$ the probability of an infinite cluster lies between zero and unity, and depends on the shape of the lattice. For a square-shaped lattice the probability of crossing between opposite sides was derived using conformal field theory, which yields an exact probability of 1/2 [20,21]. It was later confirmed numerically [22]. More general expressions are also derived for rectangular regions, where the crossing probability is a function of the aspect ratio.

At low densities, the lattice is sparsely populated, and breaks up into isolated clusters of a range of sizes, cutting off sharply at a characteristic size, defining the correlation length ξ. Clusters of size greater than ξ are extremely rare. Below the critical density ξ is finite and independent of the system size as long as the system is large enough. This implies that at low densities there is on average zero macroscopic loss associated with a fire ignited by a single spark. In other words, the average number of sites lost in a fire remains bounded, of order ξ^2 or less, and does not scale with the number of occupied sites ρN^2. While some rare configurations do contain macroscopic connected clusters at low densities (e.g., configurations in which all sites are connected), they are sufficiently unlikely within the ensemble of all possible layouts that they do not contribute any statistical weight to the loss in the thermodynamic limit.

Approaching the critical density ξ diverges, $\xi\sim(\rho_c-\rho)^{-\nu}$, with $\nu=4/3$. At ρ_c the system exhibits clusters that extend to the size of the system. At criticality, the largest clusters are fractal and the number of sites scales as N^{d_f}, where $d_f=91/48<2$. Since the number of sites overall scales as $\rho_c N^2$, even at ρ_c the macroscopic loss in density is zero, although there is a nonzero probability of crossing.

At densities greater than or equal to the critical density, there is an infinite cluster and the probability that any given site is on the infinite cluster defines the percolation probability $\mathcal{P}_\infty(\rho)$, illustrated schematically in Fig. 3(a). Below ρ_c, $\mathcal{P}_\infty(\rho\leqslant\rho_c)=0$ since there is no infinite cluster. At ρ_c, $\mathcal{P}_\infty(\rho\leqslant\rho_c)=0$ because even if an infinite cluster exists, it is fractal. At density 1, the system is fully occupied, corresponding to a single systemwide cluster, so $\mathcal{P}_\infty(1)=1$. As the density decreases from unity back towards ρ_c, there is a single macroscopic percolation cluster along with many separate isolated islands that are microscopic in size. The microscopic islands become exponentially rare above a characteristic size $\xi'\sim(\rho-\rho_c)^{-\nu'}$, with $\nu'=\nu=\frac{4}{3}$ (analogous to the microscopic islands described above for densities $\rho<\rho_c$). As the density decreases, the macroscopic percolation cluster becomes increasingly sparse, and the characteristic size ξ' of the isolated islands increases. However, the only time a spark leads to a macroscopic decrease in the density is when the percolation cluster is hit. This occurs with probability $\mathcal{P}_\infty(\rho)$.

Consequently, the yield can be simply expressed in terms of the percolation probability

$$Y=\mathcal{P}_\infty(\rho)[\rho-\mathcal{P}_\infty(\rho)]+[1-\mathcal{P}_\infty(\rho)]\rho=\rho-\mathcal{P}_\infty(\rho)^2. \tag{5}$$

In the first equality, the first term corresponds to the probability of hitting the infinite cluster, in which case the density that remains after the hit is the initial density minus the density associated with the infinite cluster. The second term corresponds to the probability of missing the infinite cluster, in which case the full density is recovered. For $\rho\leqslant\rho_c$, $P_\infty(\rho)=0$, and $Y=\rho$ is monotonically increasing in ρ [see Fig. 3(b)]. For $\rho>\rho_c$, $Y(\rho)$ is a monotonically decreasing function of ρ. While the density is increasing, the increasing size of the infinite cluster leads to increasing average losses. Of course, not all configurations have large losses. HOT configurations are designed for small losses up to unit density. However, typical configurations have large losses, and dominate the random ensemble.

Therefore, for $M=1$ the critical density ρ_c maximizes yield. At ρ_c, $P_\infty(\rho_c)=0$. The infinite cluster (if it exists) is a system spanning only sparse fractal object of microscopic density. Additionally, the distribution of cluster sizes is a power law, reflecting the fractal self-similarity of the critical state, with a finite-size scaling cutoff determined by the system size. The power law in the cluster size distribution leads to a power law in the fire size distribution (the fire size distribution is based on the occurrence of fires, where the prob-

ability of hitting a cluster is proportional to its area) which written cumulatively takes the form

$$F(l) \sim l^{-\alpha} g(\xi(\rho)/N). \tag{6}$$

Here $F(l)$ is the cumulative probability of a fire that results in the loss of greater than or equal to l sites. For site percolation on a two-dimensional square lattice the exponent $\alpha \approx 0.05$. Note that this distribution is very flat, and differs from the noncumulative density describing the cluster size distribution by roughly two in the exponent (going from clusters to fires increments the exponent $-\alpha$ by roughly one, and going from noncumulative to cumulative distributions also increments the exponent by one). The function $g(\xi/N)$ determines the finite-size scaling cutoff. As $N \to \infty$ this implies that the loss in density of the characteristic large fires becomes vanishingly small, scaling as N^{d_f}/N^2.

In summary, maximizing yield $Y(\rho)$ as a function of density for a single DDOF leads to $\rho = \rho_c$, and the fire characteristics are the common features associated with systems at a critical phase transition. At this density the system on average sustains no net loss. Nonetheless, the distribution of fire sizes extends up to a size the scales with the size of the system. The fires are sparse, and fractal in shape, and their characteristics are uncorrelated with the spatial distribution of ignitions.

While optimization for yield with a single DDOF leads to the critical state, it is important to note that the arguments that lead to criticality here and in SOC are slightly different. SOC is based on an implicit dynamical argument, which balances the infinitesimal rate of ignitions and slow but steady growth of trees. SOC seeks a fixed point that is a statistically steady state of the dynamics in a system that exhibits a separation of times scales. Thus SOC involves tuning rates, while criticality involves tuning densities. While in some cases, tuning rates may seem more appealing, both involve the same number of DDOF's. Even in this case some mysteries remain. The SOC forest fire model does not correspond to a critical system in the usual sense of equilibrium statistical mechanics, and exhibits scaling properties different than ordinary percolation [14,23]. In contrast, our model explicitly invokes optimization, so that deliberate feedback or evolutionary selection pressure is the underlying mechanism for selecting the state, even in the limit of a single tunable parameter. While yield (i.e., mean productivity) is a natural candidate for fitness, alternative optimization functions based, e.g., on some linear combination of the variance and loss could be defined in a manner that may lead to optimal behavior away from criticality.

B. Explicit optimization on finite lattices with few DDOF's

Next we consider cases with $M > 1$, but small enough to allow explicit computation of the optimal solution for finite lattices. Specifically, we compute yield as a function of the M^2 design cell densities ρ_{IJ} by generating a random sampling of configurations in which we independently vary the ρ_{IJ}. From this we determine the optimal yield configuration as a function of the M^2 cell densities. For small enough M and N we can compute enough random configurations to obtain smooth curves and convincingly locate the maximum. We subsequently increase N for fixed M in order to extrapolate our results to the thermodynamic limit of the underlying lattice, while keeping the number of DDOF's fixed.

Something interesting happens even for the smallest number of design cells. As N becomes large, each of the densities ρ_{IJ} either converges rapidly to unity or more gradually towards ρ_c. This was seen previously in the mean field limit [18]. The results presented here provide evidence that similar results hold on finite-dimensional lattices.

Below we illustrate numerical results for the case $M=2$, which divides the underlying lattice into four $N/2 \times N/2$ design cells. Similar results are obtained for even smaller numbers of DDOF's, e.g., when the $N \times N$ lattice is divided into two $N/2 \times N$ design cells (we will return to this case after discussing $M=2$). Here we use the first case that preserves symmetric design cells for our most detailed discussion for simplicity and consistency of notation and figures throughout the paper.

We considered a sequence of lattice sizes $N=16$, 32, 64, 128 of the underlying lattice. Each lattice is divided into four equal cells, defining a 2×2 square design lattice. We compute the yield as a function of the four independent densities $Y(\rho_{11},\rho_{12},\rho_{21},\rho_{22})$, by averaging over 100 randomly generated configurations of each density in each design cell.

Our numerical algorithm for generating these configurations is as follows. For each of the 100 members of the ensemble we use to compute each average, we assign a random number $z(i,j)$ in the interval $z(i,j) \in [0,1]$ to each site (i,j) of the underlying lattice. The random configuration in design cell (I,J) corresponding to density ρ_{IJ} is obtained by generating the configuration in the design cell where sites are occupied when $z(i,j) \leqslant \rho_{IJ}$ and vacant when $z(i,j) > \rho_{IJ}$. We accumulate statistics for yield as a function of the four design parameters by independently incrementing the densities in small steps and then averaging over the 100 different realizations of the random numbers $z(i,j)$.

In Fig. 4 for $N=64$ we illustrate cross sections of Y expressed individually as a function of each the four densities ρ_{IJ} along slices of the five-dimensional space $Y(\rho_{11},\rho_{12},\rho_{21},\rho_{22})$, which pass through the absolute maximum value of Y [e.g., we plot $Y(\rho_{11})$ for fixed values of ρ_{12}, ρ_{21}, and ρ_{22} coinciding with the maximum]. Note that for each plot, the maximum of $Y(\rho_{IJ})$ occurs for $\rho_{IJ} \approx \rho_c$ or $\rho_{IJ}=1$. In this particular case, the maximum value of Y is $Y=0.7493$ and is obtained for $\rho_{11}=0.4515$, $\rho_{12}=\rho_{21}=\rho_{22}=1$.

The unit density maxima at $\rho_{IJ}=1$ are associated with increasing values of $Y(\rho_{IJ})$ at the endpoint of the interval defining possible values, and are well defined even for finite N because the discrete underlying lattice plays a minor role when the cell is fully occupied. In contrast, the position of the maximum that occurs in the upper left cell for the density ρ_{11} near ρ_c does depend on the system size. In Fig. 5 we illustrate the results of increasing the size of the underlying lattice, which illustrates that the position of the maximum (i.e., the value of ρ_{11} where $\partial Y/\partial \rho_{11}=0$) is converging towards the critical density ρ_c.

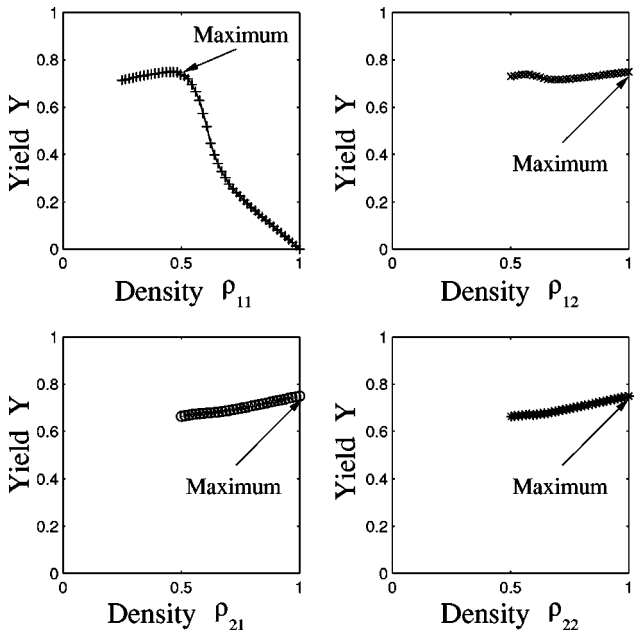

FIG. 4. Cross sections of the yield through the maximum plotted as a function of one of the four design cell densities: $Y(\rho_{11},\rho_{12},\rho_{21},\rho_{22})$, the other three remaining fixed. The results are illustrated for $N=64$.

Sample solutions obtained for different distributions of sparks are illustrated in Fig. 6. The optimal solutions (top row) and the corresponding spark distributions $P(x,y)$ (bottom row) are shown. Figure 6(a) corresponds to the asymmetric exponential in Eq. (2). Here the optimal solution corresponds to the upper left cell at the critical density, and the other three cells are at unit density, as described in Fig. 4. Figure 6(b) represents another case in which $P(x,y)$ is extremely sharply peaked. In this case, the distribution is Gaussian, and we obtain a solution of the same form as for case (a). As the sharpness of $P(x,y)$ is reduced, other solutions are observed. Figure 6(c) represents a case where the distribution is exponential in y and uniform in x, and the solution corresponds to near critical density regions in the upper half of the lattice, and unit density regions at the bot-

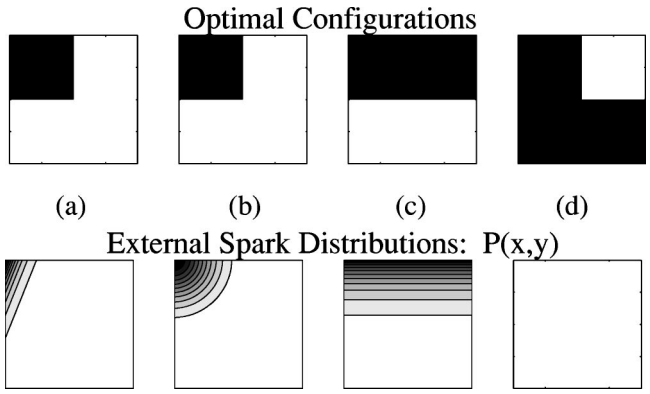

FIG. 6. Sample optimal configurations of the design lattice for four different $P(x,y)$. Black cells correspond to density near ρ_c, while white cells are at unit density. Below the configurations we illustrate the corresponding spark distributions in gray scale, ranging from higher values (black) to lower values (white). The distributions are (a) the asymmetric exponential in Eq. (2), (b) a Gaussian, (c) an exponential that depends only on y, and decays moving downward in the lattice from the peak value taken at the top, and (d) a uniform distribution.

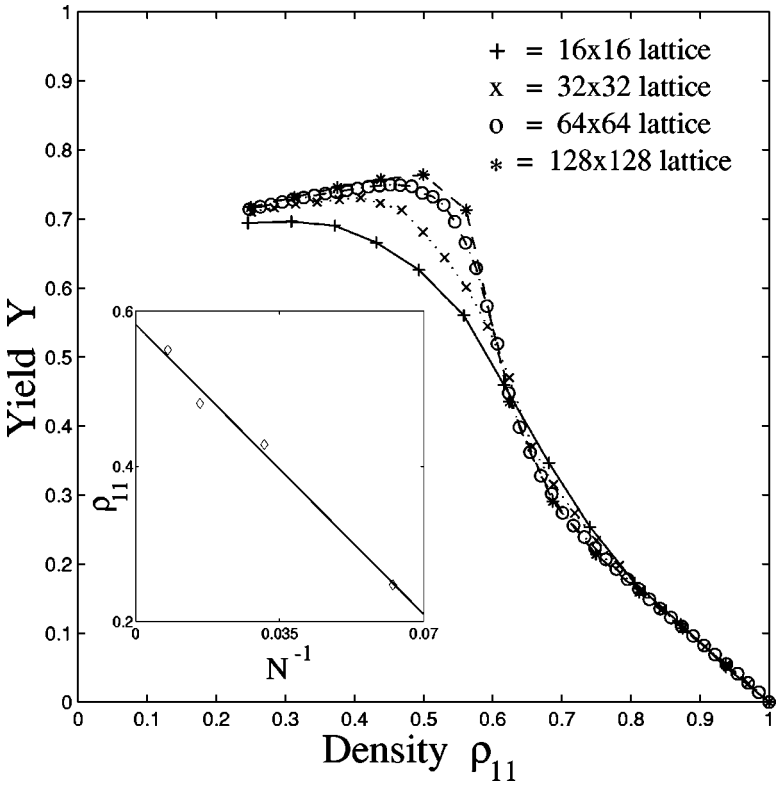

FIG. 5. Lattice size dependence of the results presented in Fig. 4 illustrates that the value of ρ_{11} at the maximum of $Y(\rho_{11},\rho_{12},\rho_{21},\rho_{22})$ converges to ρ_c as the system size increases. Results are shown for lattice sizes $N=16,32,64,128$, and illustrate that the density ρ_{11} that maximizes Y is converging towards the critical density $\rho_c \approx 0.592$ with increasing system size. The main figure illustrates the cross sections analogous to those presented in Fig. 4, but with increasing system sizes, and the inset illustrates the value of ρ_{11} at the maximum of Y as a function of inverse system size.

tom. Figure 6(d) represents the case of a uniform distribution of hits. Here the solution corresponds to a single unit density region, with the remainder of the lattice near the density ρ_c. Because the distribution is uniform, the placement of the unit density cell is arbitrary, leading to four degenerate optimal solutions. In general, in the limit $N\to\infty$ the optimal solution for the design lattice can only depend on the distribution of sparks $P(x,y)$ up to the resolution of the design lattice. This follows from the fact that each cell is described by the translation invariant ensemble of configurations characterized only by the density, which washes out all structures associated with $P(x,y)$ beyond the cumulative probability $P(I,J)$ [Eq. (4)] that the region is hit. This essentially trivial result is consistent with our finite lattice simulations, where the edge effects associated with finite grids on the underlying lattice, which are relatively small to begin with, become increasingly insignificant in the limit of large N.

Similar results are also obtained for both lesser and greater resolution M of the design lattice. At the cost of introducing asymmetric cells, we can reduce the number of DDOF's to two. In this case we again find that the cell densities converge to ρ_c or unity. For the case of a uniform distribution of sparks $P(x,y)$ the optimal solution corresponds to the critical density in each cell: $\rho_1=\rho_2=\rho_c$. This solution is maintained until a point where there is sufficient asymmetry in $P(x,y)$. At this point the cell that is less likely to be hit, say cell two, has optimal density 1, while the more likely cell, cell one, trivially remains at the critical density. The amount of asymmetry required is a function of the critical density ρ_c, and is thus dependent on the underlying percolation model. Letting $R(1)=P(1,1)\equiv p$ be the cumulative probability of hitting region 1, and $R(2)=P(1,2)\equiv 1-p$, in the limit of large system sizes we can compute the yield as a function of the second design cell density, $Y(\rho_c,\rho_2)$,

$$Y(\rho_c,\rho_2)=\rho_c+p\rho_2+(1-p)[\rho_2-\mathcal{P}_\infty(\rho_2)]. \qquad (7)$$

When the asymmetry is such that $p>\rho_c$, then the optimal solution for cell two shifts from ρ_c to unit density.

For larger values of M our computations are limited by the numerical intensity of computing the optimal value of Y as a function of M^2 design parameters. We have gone as high as $M=3$, which requires a reduction in the number of configurations we randomly sample to compute the averages. A more efficient method of computing optimal solutions begins with the assumption $N=\infty$ and makes use of the observation that the cell densities converge towards ρ_c or 1. This is presented in the next two sections.

C. Global optimization for an infinite underlying lattice and few DDOF's

In this section, we obtain globally optimal solutions by brute force optimization of the $M\times M$ design lattice layout, subject to the assumption that individual cells are near criticality or at unit density. Both criticality and unit density coincide with fixed points of the renormalization group in the standard statistical mechanical formulation of percolation. However, properties of these two fixed points are extremely different. Indeed, the special loss and crossing properties associated with cell densities approaching ρ_c and at unity, which hold (following standard results in percolation theory) in the limit $N\to\infty$ simplify our subsequent calculations enormously. The description of cells at unit density is trivial: if a spark strikes the cell all sites are lost, as are all sites in any other unit-density cell that is connected to the sparked cell by a contiguous path of nearest neighbor unit-density cells. When a spark strikes within a cell at or below density ρ_c, there is no macroscopic loss of density. The probability that fire will propagate into a neighboring cell from a hit in a critical cell is also zero, because the infinite cluster (if it exists) is sufficiently sparse that the chance of hitting it are negligible.

The only subtleties arise when considering the possibility of fires propagating from one unit-density cell to another through a region of density at or near ρ_c. First consider a spark that ignites a unit-density cell, which is separated from another unit-density cell by an intermediate cell that must be crossed from top to bottom or left to right to obtain a connection. If the intermediate cell has density greater than ρ_c, the crossing probability is unity. If the intermediate cell has density exactly ρ_c, there is a finite probability of crossing, which depends on the shape of the intermediate region. As previously stated, if the contiguous intermediate ρ_c region is square, the crossing probability is exactly $\frac{1}{2}$ [20,21]. However, if the density is $\rho_c-\epsilon$, where ϵ can be taken infinitesimally small, the probability drops to zero. Thus for $M>2$ DDOF's (where the issue of intermediate cells becomes important), the ''critical'' regions will remain bounded away from criticality (at infinitesimal cost in density, but substantial gain in yield) to avoid connecting unit-density domains.

The final case we must consider is that of next nearest neighbor unit-density cells. That is, consider two unit-density cells that share a common corner, but are otherwise separated by cells at or near the critical density. At and above the critical density, the probability of a path connecting the unit-density cells through the intermediate critical cell is one. Above the critical density this follows from the fact that connected paths between adjacent sides of a cell are more likely than crossings between opposite sides, because of the many finite paths connecting adjacent sides. The existence of many finite paths also insures that for densities below criticality, the probability that the corner connected unit-density cells are joined through the intermediate cell remains finite. Unlike the probability of a crossing between opposite sides of a cell, which is strictly zero below the critical density, the probability of a crossing between adjacent sides is nonzero for all $\rho>0$. For example, occupation of the single site at the corner of the cell (which occurs with probability ρ) is sufficient. In fact, numerically we find that the probability of a crossing between adjacent sides of a square cell monotonically increases with density, and continuously approaches unity at $\rho=\rho_c$. In particular, we find that when $\rho=\rho_c-\epsilon$, the probability of a crossing between adjacent sides remains essentially unity $[1-O(\epsilon^\chi)$ from some $0<\chi<1]$. Thus, corner connections are equivalent to nearest neighbor connections on the design lattice in the limit of large N. In our finite lattice simulations (for design lattices up to 3×3), we ob-

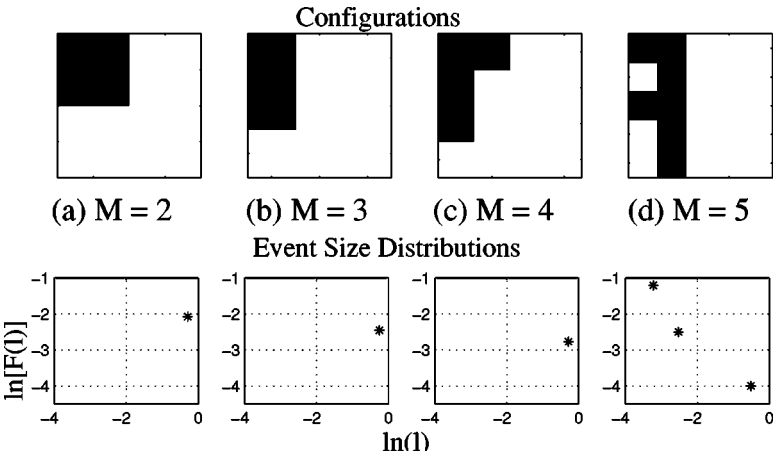

FIG. 7. Global optimization of the design lattice for $M \leqslant 5$. Solutions are obtained by considering all possible design lattice configurations, with $\rho_{IJ} \in \{\rho_c - \epsilon, 1\}$. For $M=1$ the optimal solution is the critical density. The top row illustrates the optimal configurations for $M=2,...,5$ (black corresponds to a density approaching ρ_c, and white corresponds to unit density), and the bottom row illustrates the macroscopic portion of the corresponding event size distributions for each case.

serve that corner connections are always excluded in optimal solutions, for $P(I,J)$ sufficiently smooth [24]. For example, with $M=2$ and a uniform distribution of hits, the optimal solution has only one unit-density cell (rather than two that are diagonally connected) as illustrated in Fig. 6(d). See Appendix A for a further discussion of corner connectivity.

When we assume an infinite underlying lattice, we are not *a priori* restricted to considering site percolation on the two-dimensional square lattice as the underlying statistical model. Site percolation on the two-dimensional square lattice fixes the value of ρ_c at 0.592. While the particular layout of optimal solutions does depend on ρ_c, our qualitative results are independent of the specific value. For consistency we retain this choice. The disadvantage is that site percolation on the two-dimensional square lattice has not been solved rigorously. If we had chosen an underlying model for which percolation results were mathematically rigorous, many of our results for adding DDOF's would be mathematically precise as well.

In the limit $N \to \infty$, we use the following results from percolation theory and the previous subsection to calculate both the yield and the event size distributions.

1. Propagation

(1) Cells at $\rho_c - \epsilon$ experience no macroscopic loss in density in a fire, and fires do not propagate macroscopic distances across the cell.

(2) Fires do not propagate from left to right or from top to bottom across cells at density $\rho_c - \epsilon$. The $\rho_c - \epsilon$ cells thus act effectively as fire breaks for vertical and horizontal propagation.

(3) Fires will propagate between adjacent edges of cells with density $\rho_c - \epsilon$ (see Appendix A). This implies a corner connection between cells at unity density is effectively the same as a shared edge.

(4) Cells at unit density experience total loss when a spark hits the cell, or when fires propagate into the cell from nearest (edge connected) or next nearest (corner connected) neighbor cells at unit density.

2. Event sizes

(5) The cluster size distribution in cells at density $\rho_c - \epsilon$ is identical to that of site percolation in the neighborhood of the critical density. Each cell at density $\rho_c - \epsilon$ contributes cumulative weight to the event size distribution given by the total probability $P(I,J)$ [Eq. (4)] of hitting within the cell. With probability $(1-\rho_c)P(I,J)+O(\epsilon)$ it hit a vacant site within the cell, and with probability $\rho_c P(I,J)+O(\epsilon)$ an occupied site is hit. The critical events span a range of sizes which scales as $F(l) \sim l^{-\alpha}$ with $\alpha \sim 0.05$, but which all scale to zero loss in the limit $N \to \infty$ [see item (7)].

(6) The cluster size distribution associated with cells at unit density is determined by their connectivity, each with statistical weight determined by the cumulative hit probability summed over the area of connected regions (as defined above).

(7) The overall event size distribution thus has two contributions: one from cells at density approaching ρ_c and the other from cells at unit density. These scale differently in the limit $N \to \infty$. Specifically, the events in the $\rho_c - \epsilon$ regions are infinitesimal, with the largest events scaling as N^{d_f}, with $d_f = \frac{91}{48}$, compared to the density, which scales as N^2, and distributed according to a power law $F(\sim l) \sim l^{-\alpha}$, where $\alpha \approx 0.05$. The events associated with the unit-density cells consist of a discrete set of sizes, with macroscopic loss.

Next we compute the globally optimal configuration by explicitly considering all of the 2^{M^2} possible configurations, and picking the one with highest yield. Because the number of configurations increases extremely rapidly with M, this brute force global optimization rapidly slows down and becomes computationally intractable for $M>5$.

Our results for $M \leqslant 5$ are illustrated in Fig. 7 for a fixed distribution of sparks $P(x,y)$ [Eq. (2)]. Black signifies density $\rho_c - \epsilon$ in the design lattice, and white signifies density 1. With $M=1$ (not shown) criticality is the optimal solution, and the event size distribution is that of criticality. For $M=2$ we recover the pattern of critical and unit-density cells that the corresponding discrete lattice simulations was converging towards with increasing N [Fig. 6(a)]. Specifically, the cell in the upper left corner is near the critical density, while the remaining cells are at unit density. A single event is illustrated in the corresponding event size distribution—the large event that occurs if any of the sites within one of the three unit-density cells is hit. Since the size of all of the critical events in the $\rho_c - \epsilon$ cell scale to zero as $N \to \infty$, we

cannot illustrate them on the same logarithmic graph.

As M increases, increasing resolution of the design grid results in increasingly refined patterns. The first aspect of this, which emerges for $M=3,4$, is increased resolution of the spark distribution $P(x,y)$, through more refined placement of the $\rho_c-\epsilon$ barriers. This placement is subject to irregularities imposed by the combination of the continuous distribution of sparks and the finite design grid. For both $M=3,4$ there is still a single region at unit density. However, it is clear from the comparison of the event sizes and probabilities for $M=2,3,4$ that with increasing M the size of the large event is staying roughly constant, while the probability of hitting the region is decreasing.

For larger M ($M=5$), we begin to see multiple (here three) unit-density regions, separated by barriers formed by nearest neighbor connected paths of design cells at density $\rho_c-\epsilon$. The width of the barriers corresponds to the width of a single design cell, and they are concentrated near the upper left-hand corner of the lattice that has the highest probability of sparks. They are also concentrated towards the left edge of the lattice, more than the top edge, because the asymmetry of the exponential distribution implies the probability of sparks falls off more rapidly along the horizontal axis, compared to the vertical axis. As a result, smaller events occur in regions that are more likely to be sparked, while large events occur in regions where sparks are rare. Although the data are still extremely sparse, the relative probabilities of fires in three unit-density regions are consistent with power law statistics produced for much larger lattices (see Fig. 9).

As M increases the patterns become increasingly reminiscent of the HOT configurations obtained previously through constrained optimization on finite lattices. In those studies, individual sites were chosen to be occupied or vacant, and connected barriers of vacant sites defined firebreaks separating compact connected clusters of occupied sites. The difference here is that our finite lattice is the design lattice, which is superimposed on an infinite underlying lattice. Barriers and occupied sites correspond to design cells at density $\rho_c-\epsilon$ and unity, respectively. There are a few technical differences between our calculations here and the analogous optimizations on finite lattices: on the finite lattice fires do not spread between next nearest neighbors (i.e., corner connections), and calculations of the yield weight barrier sites as zero (vacant) rather than $\rho_c-\epsilon$ (critical). By taking these differences into account, algorithms for global optimization of small lattices are easily modified for the small finite design lattices superimposed on infinite underlying lattices considered in this section. This suggests the next step in our numerical optimization, which is to modify our constrained optimizations for finite lattices, to develop analogous constrained optimizations for a finite design lattice superimposed on an infinite underlying lattice.

D. Local optimization for an infinite underlying lattice and many DDOF's

In this section we generalize the local incremental algorithm discussed in Sec. II and [10] for obtaining HOT states on finite lattices to the case of a finite design lattice superimposed on an infinite underlying lattice. Rules (1)–(7) defined in Sec. III C for calculating yields and event size distributions continue to hold. The only difference is that here we restrict the search space for optimal solutions. Thus we are not guaranteed or even likely to converge to the globally optimal solution for a given value of M. Nonetheless, as in our previous analysis of finite lattices [9,10], the general properties of increasingly designed states are independent of the specific constrained optimization performed.

The algorithm is defined as follows: We begin with each cell at density $\rho_c-\epsilon$. We incrementally increase the density to unity by converting cells to unit density one at a time. The choice of which cell to convert at each incremental step is determined by testing all possible remaining choices (i.e., cells that are still at density $\rho_c-\epsilon$) to determine which cell, if converted, leads to the highest yield configuration on the lattice, given that one additional cell must be converted. That best cell is then converted, and the procedure continues until the lattice is fully occupied. In the case of degeneracies (two or more choices produce identical outcomes), one of the choices is selected at random.

For each M, this procedure generates a discrete curve of yield vs density (Fig. 8), with density increments of $(1-\rho_c)/M^2$ as additional sites on the lattice are occupied. The configuration associated with the maximum value of Y (marked by the arrow in Fig. 8) is the optimal configuration for the search. These are illustrated in Fig. 9 for increasing values of M. The search is local in the sense that the configuration associated with each increment in density is based on the configuration at the previous increment. Compared to the brute force global optimization in the previous section, where all of the possible 2^{M^2} configurations are considered as candidates for the optimal configuration, here we search over a restricted space of less than M^4 possibilities. This allows us to consider much larger values of M than are accessible in the brute force global optimization.

Here we see that for small values of M ($M=1,2,3$) the results agree with results obtained by global optimization in Fig. 7. However, for more than two DDOF's this need not generally be the case. For example, for the case of three DDOF's, the local incremental algorithm will always convert the cell that is least likely to be hit from $\rho_c-\epsilon$ to unit density first. However, for some $P(x,y)$ the globally optimal solution will not have that site at unit density. Nonetheless, the qualitative features obtained from the local incremental algorithm are similar to those obtained by global optimization. Here we again see that increasing values of M lead to increasingly refined patterns composed of compact connected clusters at unit density, separated by barriers one design cell wide at density near ρ_c. Furthermore, because we access larger values of M, we begin to deduce the emergence of systematic trends associated with increasing design, which we verify analytically for the case of a uniform distribution of sparks in Sec. IV. These are deduced from Figs. 9 and 10, where Fig. 10 illustrates the peak values of Y obtained from searches analogous to that illustrated in Fig. 8. We conclude this section with a summary of our observations.

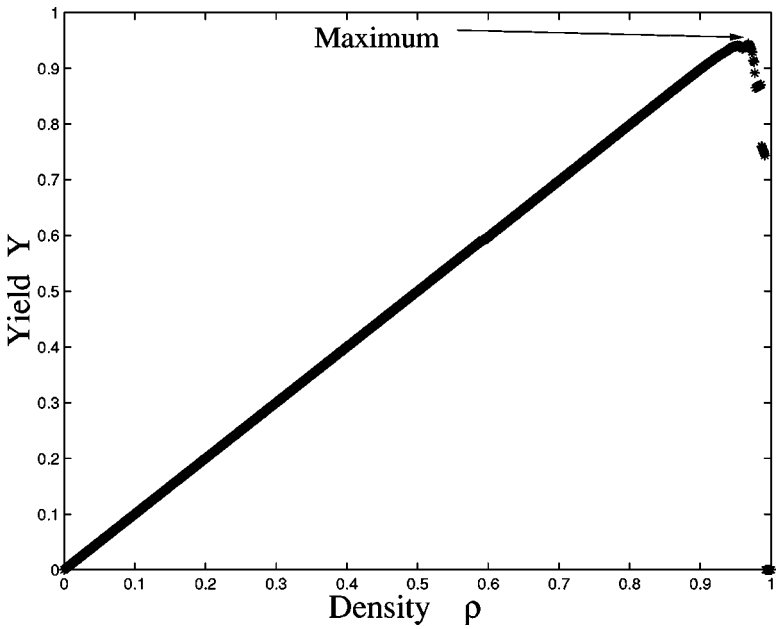

FIG. 8. Search for an optimal design configuration for $M=64$ using the local, incremental algorithm. Design cells are converted from density approaching ρ_c to unit density one at a time, in a manner that maximizes Y at each increment. This results in the dense trace of black *'s (which appear as a solid line for most of the curve) illustrated in the figure. The optimal design configuration [Fig. 9(f)] corresponds to the maximum Y for this search.

(1) Increasing M leads to increasing densities and increasing yields. This is illustrated in Fig. 10, where we plot the yield as a function of density for the locally optimized lattices for successive numbers of DDOF's. The one exception to this trend in our numerical data is seen in comparing $M=4$ and $M=8$, where the yield of the $M=4$ solution is slightly greater than for $M=8$. This inversion does not disrupt the overall trend, and arises due to the finite size of the design lattice.

(2) Increasing M typically leads to decreasing average loss. In Fig. 10 the loss is measured by the vertical drop of the optimal yield from the diagonal line (zero loss). The critical solution obtained for $M=1$ has zero loss. However, the higher-density configurations obtained for $M>1$ have nonzero loss associated with the unit-density cells. This is apparent in our numerical results, which show a drop from the diagonal line for $M=2$. As M increases, the drop tends to decrease (even though the density is increasing, indicating more converted regions). As in item (1) there are exceptions due to finite design grid effects. Nevertheless, the overall trend is illustrated by the fact that the slope of the curve drawn through the optimal yield points is steeper than that of the diagonal one.

(3) There are several aspects of items (1) and (2) that are also apparent from the event size distributions (Fig. 9). The increasing density with increasing M [item (1)] implies that in the event size distribution statistical weight from the microscopic critical portion of the distribution (not shown) is shifted to the macroscopic portion, which we refer to as the HOT tail. This is apparent in comparing the cumulative weight in the HOT tail [deduced from the value of $F(l)$ associated with the left-most data point], which increases with M. Increasing DDOF's also leads to increasingly refined patterns, which adds breadth to tail in the event size distribution. Figure 9 illustrates the event size distribution of the optimal configuration for different values of M. The fact that the average loss decreases with increasing M [item (2)] is another way of saying that the mean size in the event distribution decreases with increasing M.

(4) Finally, increasing density and spatial resolution of the pattern, which is associated with higher yields for higher DDOF's simultaneously introduces new sensitivities, reflecting the robust yet fragile nature of designed systems. Higher densities imply increased fragility to changes in the distribution of sparks, and flaws in the design pattern, to which the critical configuration is entirely insensitive.

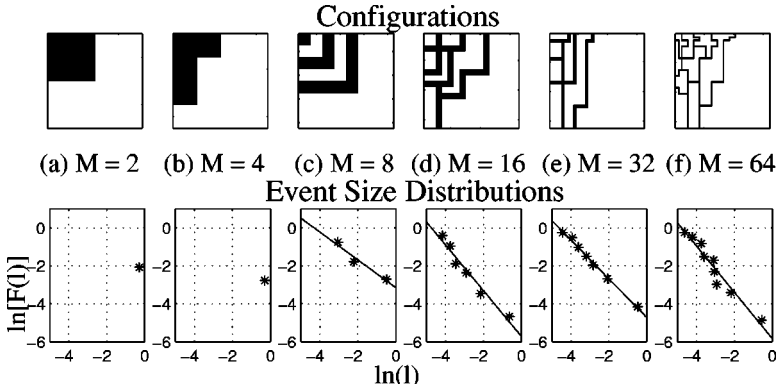

FIG. 9. Optimal configurations and event size distributions for the design lattice obtained using the local incremental algorithm for $M=2,4,8,16,32,64$ and $P(x,y)$ from Eq. (2).

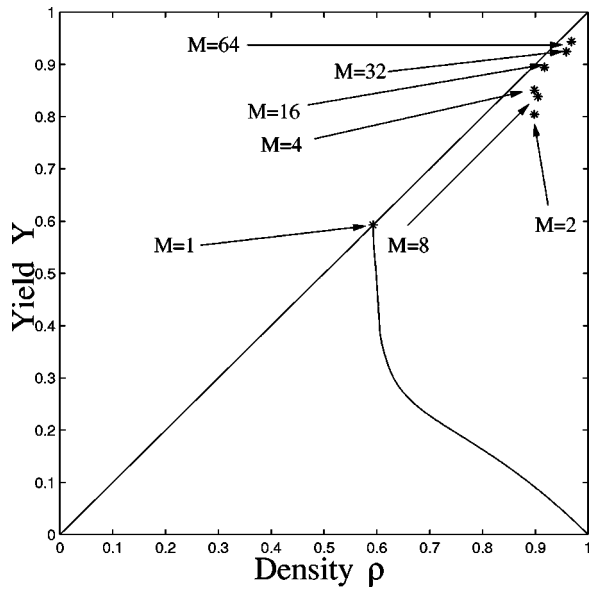

FIG. 10. Yield vs density for the optimal configurations on the design lattice obtained using the local incremental algorithm for $M=1,2,4,8,16,32,64$ and $P(x,y)$ from Eq. (2). For each M the maximum is obtained from a run in which the density is incrementally increased as in Fig. 8. Here for clarity, for all but $M=1$, we omit the full scan over densities, and retain only the peak value.

IV. ANALYTICAL RESULTS FOR A UNIFORM DISTRIBUTION OF SPARKS

The symmetries associated with a uniform distribution of sparks make optimization of the yield in the percolation forest fire model with M^2 DDOF's more tractable analytically. As in the previous section, we focus on the case of an infinite underlying lattice. In this limit, the optimal design cell densities are either approaching ρ_c or at unit density, and rules (1)–(7) of Sec. III B can be used to determine the propagation properties between cells and the event size distributions.

Optimal solutions typically consist of compact unit-density regions of equal size, surrounded by near critical barriers. Unit-density regions will tend to be of equal size, because for uniform sparks, the area of a unit-density region defines both the hit probability of the region and the loss. If one unit-density region is larger than another, then the larger cell will both be more likely to be hit, and will also suffer a greater loss, causing an increase in the average loss relative to that which would be obtained if the regions were of equal size. The unit density regions will tend to be square, surrounded by critical barriers one design cell in width in order to maximize density. Such a configuration minimizes the number of critical density cells required to isolate the region.

Deviations from this typical layout are associated with packing constraints that are encountered in fitting the optimal-size unit-density regions into a design lattice of finite M. This results in some spread in the size distribution of the regions to accommodate the edges, which is of negligible importance as $M\rightarrow\infty$. When the finite size of the design grid is not incompatible with the optimal region size, all regions are square and of equal size even for finite M.

For a uniform distribution of sparks, the microscopic critical events still exhibit power laws. However, we no longer obtain power law distributions for the macroscopic events, since they are all of equal size. Fully designed HOT configurations on finite lattices do not exhibit power laws for uniform distributions of sparks either. In that case, the cells are also square and of equal size, though the optimal barriers are diagonal relative to the underlying lattice, which minimizes the number of vacancies (since fires do not propagate between corner connected sites on the underlying lattice).

The calculations in this section provide a quantitative, analytical illustration of the features associated with adding DDOF's that emerged as numerical trends in the previous section. Namely, increasing DDOF's leads to increasing densities, decreasing losses, increasing yields, and a shift in statistical weight in the event size distribution from critical events towards the HOT tail (now a family of events, all of equal size). We begin by considering small M, where just a few DDOF's produce macroscopic increases in yield. This is followed by asymptotic analysis of the large M limit, where we show yield approaches unity, and determine the characteristic event size of the unit-density regions.

A. Exact solutions for small numbers of DDOF's

When M is sufficiently small it is possible to exactly calculate the optimal configuration by explicitly considering relatively few choices. The uniform distribution of sparks simplifies the problem by introducing many degeneracies in the yield for different spatial patterns of the cell densities.

The yield can be written as

$$Y=\frac{1}{M^2}\left(\sum_{\{I,J\}}\rho_{IJ}-\sum_{\mathcal{R}_k}P(\mathcal{R}_k)A(\mathcal{R}_k)\right). \tag{8}$$

The first term on the left-hand side is the total density, written as a sum over the densities of the design cells $\{I, J\}$, and the second term is the average loss. The loss comes from each of the unit-density design cells, which form a set of connected regions $\{\mathcal{R}_k\}$ (the edge and corner connected clusters) on the design lattice. Here $P(\mathcal{R}_k)$ is the probability of hitting region $\mathcal{R}_k$, and $A(\mathcal{R}_k)$ is the corresponding area (i.e., the number of design cells in region $\mathcal{R}_k$).

We obtain the following results for small M:

$M=1$. As previously stated, the optimal solution for a single design cell corresponds to the critical density.

$M=2$. For a 2×2 design grid, we consider the yields associated with zero to four design cells set to unit density, with the remaining cells at the critical density. Since edge and corner connections of unit-density cells on the design lattice are equivalent, there is no distinction between different arrangements of unit-density cells once the number is fixed. Letting Y_s denote the yield for s design cells at unit density, by explicit calculation we obtain is $Y_0=0.592$, $Y_1=0.632$, $Y_2=0.546$, $Y_3=0.336$, and $Y_4=0.0$. Thus for $M=2$ the optimal solution corresponds to one design cell at unit density, and the three remaining cells at density ρ_c. This is exactly what we found numerically on finite lattices [Fig. 6(d)].

TABLE I. Analytical results for the optimal yield, density, and average loss for a uniform distribution of sparks, with $N\to\infty$ and $M\times M$ design lattice, with unit-density regions of size m.

M	m	Yield	Density	Loss
1	0	0.592	0.592	0
2	1	0.632	0.694	0.062
3	1	0.724	0.773	0.049
large M	$(1-\rho_c)^{1/3}M^{2/3}$	$1-3(1-\rho_c)^{2/3}M^{-2/3}$	$1-2(1-\rho_c)^{2/3}M^{-2/3}$	$(1-\rho_c)^{2/3}M^{-2/3}$
$M\to\infty$	∞, but $m/M\to 0$	1	1	0

$M=3$. For a 3×3 design grid a larger number of candidate configurations must be taken into account. Degeneracies introduced by the uniform distribution of sparks again simplifies the search considerably, so that we need only consider 13 distinct configurations. After some algebra, the optimal solution is found to correspond to unit-density regions in each of the four corner design cells, with the remaining five cells forming a plus sign at density approaching ρ_c. The yield in this configuration is $Y=0.724$.

Determining optimal solutions for larger values of M becomes increasingly tedious. However, even $M\leqslant 3$ suggests a trend that we expect to continue for large M. Namely, the optimal state of the design lattice breaks up into unit-density regions (in these cases corresponding to a single design cell) surrounded by critical barriers one design cell wide. For increasing M, we obtain increasing densities, increasing yields, and decreasing average losses. These results are summarized in Table I. As for nonuniform $P(x,y)$, we again find that very few DDOF's produce macroscopic increases in yield.

For small M design lattices, the critical density barriers occupy a larger portion of the lattice than the unit-density regions, even though they are a single design cell wide. However, as M increases, the design cell size decreases, and the critical barriers shrink in width, occupying less and less area. Eventually the optimal number of design cells in each unit-density region becomes greater than one, and will continue to increase with increasing M. Since the barriers are a single design cell wide, as $M\to\infty$ it is possible to have an infinite number of barriers, which nonetheless occupy a vanishing fraction of the lattice, leading to the asymptotic results of unit yield and density in the limit $M\to\infty$, discussed in the next section.

B. Asymptotic solution for many DDOF's

When the number of DDOF's is large, the optimal solution for a uniform distribution of sparks divides the $M\times M$ design lattice into $m\times m$ square unit-density regions of equal size, divided by critical boundaries of density $\rho_c-\epsilon$. The objective of this section is to calculate the optimal size m of the unit-density regions for fixed M, in the limit $M\to\infty$. Asymptotically, we can ignore corrections associated with packing the optimal size regions onto a finite design grid. These terms correspond to small adjustments of m away from the optimal value that does not alter the scaling.

The solution is represented schematically in Fig. 11. The $M\times M$ design grid is broken up into $(m+1)\times(m+1)$ regular repeat units. Each of these consists of a square unit-density region of m^2 design cells, and linear perimeter of $(2m+1)$ cells at density $\rho_c-\epsilon$. Neglecting terms of $O(\epsilon)$ the yield that corresponds to this configuration is given by

$$Y=\frac{m^2}{(m+1)^2}+\frac{2m+1}{(m+1)^2}\rho_c-\frac{m^4}{M^2(m+1)^2}. \tag{9}$$

The first two terms on the left-hand side correspond to the density of the $\rho=1$, and $\rho_c-\epsilon$ cells, respectively. The last term is the average loss associated with the unit-density regions.

It is a straightforward exercise to optimize m for fixed M. We obtain

$$m\approx(1-\rho_c)^{1/3}M^{2/3}. \tag{10}$$

The key steps are outlined in Appendix B. From this solution we can obtain asymptotic results for yield Y,

$$Y\approx 1-3(1-\rho_c)^{2/3}M^{-2/3}, \tag{11}$$

the density

$$\rho\approx 1-2(1-\rho_c)^{2/3}M^{-2/3}, \tag{12}$$

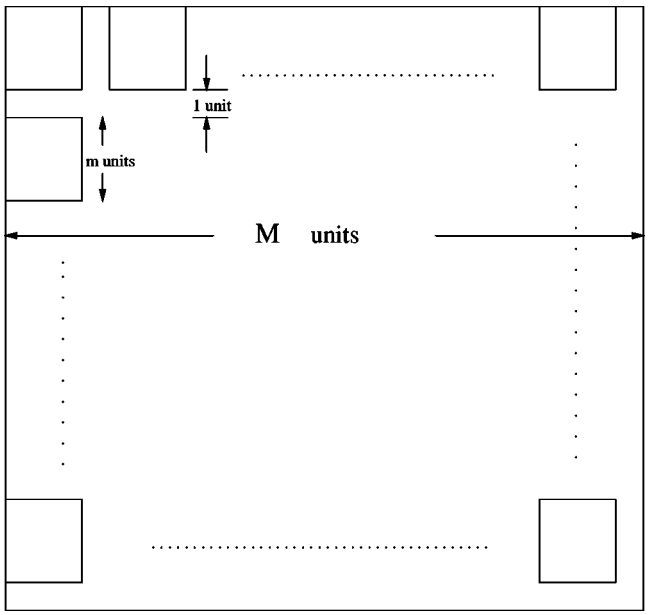

FIG. 11. Schematic diagram of the optimal layout for a uniform distribution of sparks.

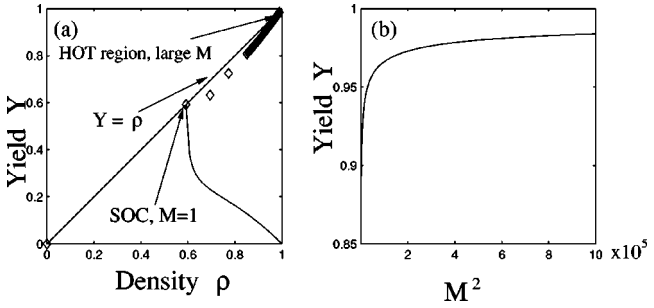

FIG. 12. Figure (a) illustrates yield vs density for the optimal configurations on the design lattice obtained using the uniform $P(x,y)$ solution for $M=1$, 2, 3 and $15 \leqslant M \leqslant 1000$, while (b) illustrates the yield as the number of DDOF's are increased. Here we plot the leading asymptotic results, and neglect small corrections due to fitting.

and the average loss

$$\langle l \rangle_{\text{ave}} \approx (1-\rho_c)^{2/3} M^{-2/3}, \tag{13}$$

which are included in Table I. Numerical results plotting the yield as a function of density appear in Fig. 12(a) for large values of M.

For finite M, corrections associated with possible misfits of the optimal $m \times m$ squares on the finite design grid can be calculated, and leads to some rectangular regions when the fit is not perfect. A perfect fit occurs when M, or $M-1$ is an integer multiple of the optimal $(m+1)$. Here $M-1$ correspond to cases where the $m \times m$ squares fit perfectly, but the last critical barrier is removed from the boundary row. However, as $M \to \infty$ corrections are negligible, and do not change the leading asymptotic results given above.

As with small M these results illustrate the general trends associated with increasing DDOF's, which are consistent with the numerical results of the previous section. Namely, increasing DDOF's lead to increasing densities and yields, and decreasing losses. Results for small $M=1$, 2, 3 (first three diamonds, in order of increasing Y) and large M (increasing M corresponds to increasing Y) are depicted graphically in Fig. 12. As $M \to \infty$ the yield approaches the maximum value of unity. The fact that the slope of Y vs. ρ in Fig. 12 is steeper than the diagonal defined by $Y=\rho$ illustrates that the average loss decreases with increasing ρ.

Similar to our numerical results for nonuniform $P(x,y)$, the distribution of events breaks up into microscopic critical events associated with $\rho_c - \epsilon$ barriers and events involving macroscopic losses associated with the unit-density regions. In this case the critical events have the same power law statistics. Increasing DDOF's again results in increasing weight associated with the unit-density regions in the overall distribution of events. However, for a uniform distribution of sparks the statistics of the ''HOT tail'' are not described by a power law, and instead correspond to a single size, or in some cases a few sizes when packing corrections are taken into account. The characteristic size of unit-density regions decreases with increasing M and becomes microscopic in the limit $M \to \infty$. However, because they correspond to a diverging number of unit-density design cells, they are still infinitely large relative to the largest (fractal) events in the critical regions, which lie within the barrier network.

TABLE II. Contrasting properties of criticality/SOC vs HOT.

Characteristic	Criticality and SOC	HOT
Density	Low	High
Internal configuration	Generic, random	Highly structured
Loss	Infinitesimal	Macroscopic
Events	Fractal	Compact
Statistics	Power laws	Power laws
Origin of power laws	Internal fluctuations, criticality	Optimization in a variable environment
Robustness	Generic	Robust, yet fragile
Increased resolution	Does not matter	New structure, and sensitivities

V. CONCLUSION

At a time when the study of complex system plays an increasing role in science, particularly in interdisciplinary endeavors, developing a more quantitative measure of such general concepts as complexity and design is important because it leads to a more precise and common vocabulary that can be applied to different systems. The scientific field of complex systems aims to link simple models and general principles that arise in physics, mathematics, and engineering to a wide range of real and genuinely complicated applications that span many disciplines. In order to strengthen these links it is useful to examine how the amount of design included in simple models may affect the nature of the complexity that is observed. In simple models basic concepts can be investigated in detail, albeit in an abstract context. While it is difficult if not impossible to imagine quantifying precisely the amount of design in everyday complex systems such as ecosystems or the Internet, it is a common engineering and policy task to evaluate how tradeoffs associated with alterations in design and added complexity may change the performance of a system for both better and worse.

Here we introduced a method for incrementing the number of tunable design degrees of freedom in percolation forest fire models. This allows us to interpolate between low and high design limits, corresponding to critical and HOT states, respectively. The intermediate design states do not represent a smooth transition in the underlying configuration of the lattice, but rather a shift in the statistical weight associated with near critical regions towards unit-density regions that form HOT lattice configurations. The contrasting characteristics of criticality and HOT are listed in Table II.

Construction of a design lattice of cells superimposed on an underlying lattice simplifies our analysis, given that percolation is *a priori* a lattice problem. Our method is reminiscent of an inverted real space renormalization, but serves to highlight the fact that highly designed lattices would not be renormalizable in the traditional sense. Rescaling the density of a configuration, without specifically taking into account

the barriers, would wash out the key features responsible for robustness.

One might argue that our definition of DDOF's in terms of a design lattice has biased our results towards solutions that exhibit separate HOT and critical features segregated into different regions in space. This is certainly a reasonable concern, and worthy of further investigation. Our construction is artificial, and presumes the ability to define sharp boundaries between regions of different density. However, we expect similar results would hold even if the borders between design cells were more diffuse. The fact that $\rho_c - \epsilon$ corresponds to the highest density at which there is zero loss and zero connectivity suggests that criticality may play a special role in the initial formation of barriers in any implementation of increasing DDOF's. Our design lattice *a priori* restricts the near critical barriers to specific design cells. An alternative scheme might expand the density $\rho(x,y)$ as a polynomial in x and y, with additional DDOF's associated with added terms in the expansion. For a single DDOF, only the density is tuned, again resulting in criticality. The next step would be associated with introducing a gradient. While such an expansion forces a smooth variation in the density, it is the spatial placement of the critical density that plays the central role in determining the solution, and serves as a boundary between a lower-density region in which there is no net loss, and a high-density region in which loss is macroscopic. For large numbers of DDOF's in this scheme, critical barriers could be resolved with increasing sharpness and spatial placement, just as in our design lattice.

In terms of yield, the specific value of the density in the barrier regions becomes less and less important with increasing DDOF's, and of vanishing importance in the limit of infinite DDOF's. The key function of barriers is to isolate unit-density regions. For large DDOF's, the contribution to density from the increasingly narrow barriers becomes negligible, so that the barriers could be any density, even zero, as long as their density is below ρ_c. Because fluctuations in the barrier density around ρ_c can induce connectivity, finite lattice simulations exhibit barrier densities that are sufficiently below ρ_c to make connections extremely rare. Thus the role of criticality may be an artifice of taking small DDOF's, combined with a thermodynamic limit on the underlying lattice. When we take the thermodynamic limit of the underlying lattice, the function describing connectivity becomes completely deterministic, except at ρ_c, where the probability of connection between opposite sides of a barrier depends on the shape of the barrier, and lies somewhere between zero and unity [20,21]. The fact that our barriers are tuned to ϵ below the critical density is a testament to the fact that we are optimizing for yield, with no risk of fluctuations, which results in barriers tuned to the maximum density to prevent connectivity. Clearly, were we to introduce the possibility of density fluctuations and other stochastic effects, these configurations would be highly sensitive, and a more conservative barrier density and barrier width that is more robust to fluctuations would win out.

Additional methods for increasing the number of DDOF's may be motivated by specific applications, and could involve additional mechanisms for the spread of cascading failure beyond the nearest neighbor connectivity represented in percolation. However, it is our belief that most natural and man-made complex systems lie near the high design limit described by HOT.

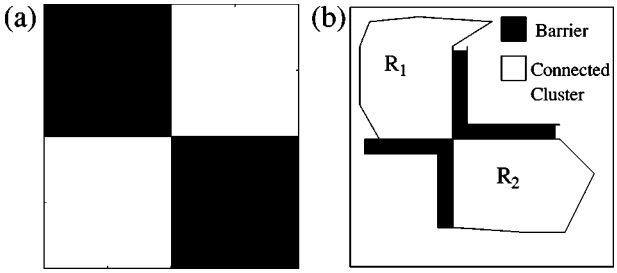

FIG. 13. Corner connectivity. Figures (a) and (b) illustrate unit-density regions that are corner connected (next nearest neighbors). While $M=2$ in (a), M is essentially arbitrary in (b). It is presumed in the text that both are subject to a uniform $P(I,J)$ and that the unit-density cells are macroscopic.

High-yield lattices correspond to an extremely small subset of all possible configurations, for which most single-site perturbations are neutral in terms of yield. While the temptation is to say that these high-yield lattices are robustly high yield, the correct statement is that they are extremely *robust, yet fragile*, because of their extreme sensitivity to a few rare events.

ACKNOWLEDGMENTS

This work was supported by the David and Lucile Packard Foundation, NSF Grant No. DMR-9813752, and EPRI/DoD.

APPENDIX A: CORNER CONNECTIVITY

In this appendix we describe why a corner connection between next nearest neighbor unit-density cells makes propagation of fires between the unit-density cells sufficiently likely that such configurations can be considered equivalent to a shared edge. The net effect is to prevent diagonal barriers on the design lattice that are commonly observed when the configuration is optimized on the underlying lattice. This distinction is readily seen by comparing the HOT configuration obtained using the local incremental algorithm in Fig. 1(b) with the corresponding result for the local incremental algorithm applied to the design lattice in Fig. 9(f), which has only vertical and horizontal barriers. The results in this appendix fall short of a rigorous proof, but a combination of numerical and analytical results strongly suggest that when the density in the intermediate cell is $\rho_c - \epsilon$, the chances of a crossing between adjacent edges in the cell is $1 - O(\epsilon^{\chi})$, where $0 < \chi < 1$. Thus as ϵ approaches zero, fires will spread almost surely between corner connected unit-density cells as we assume in Sec. III C, propagation rules (3)–(4).

Figure 13 illustrates two examples of corner connections on the design lattice. Figure 13(a) represents the case for $M=2$, while Fig. 13(b) represents the case for some much larger value of M. In the latter case, the individual design

cells are not drawn, and the $\rho_c-\epsilon$ barriers correspond to vertically and horizontally connected lines of cells, one design cell wide. The case of relatively large unit-density regions, compared to narrow $\rho_c-\epsilon$ barriers corresponds to large values of M, where the cost in density of adding a few more $\rho_c-\epsilon$ cells to prevent a corner connection is negligibly small. However, even in the case of small M we see corner connections prohibited numerically [see, e.g., Fig. 6(c), where corner connections are excluded for uniform hits].

In both the small and large M cases shown in Fig. 13, the two white unit-density regions share a common corner. If a spark strikes one of the unit-density regions, what is the probability the fire spreads through the intermediate region to the other unit-density cell? If the intermediate cell has unit density, the corner connected cells are *a priori* part of the same connected cluster. If the cell has density ρ_c or $\rho_c-\epsilon$ we have to do a little more work to show that this is still almost surely the case.

The probability of crossing between adjacent sides of the cell at densities near criticality is well studied [20–22]. In particular, an exact formula for crossing probability between segments of the boundary of a simply connected compact region was originally derived in Ref. [20] and is called Cardy's formula. It follows from this that if the intermediate cell has density ρ_c or greater, the crossing probability is unity, due to the existence of an infinite cluster. To prevent top to bottom crossing, the barrier cells must be at or below density $\rho=\rho_c-\epsilon$.

To determine the general behavior of this diagonal crossing probability P_{dc} for arbitrary densities, we consider a finite underlying $N\times N$ lattice, in the limit of large N. In the absence of an infinite cluster ($\rho<\rho_c$), the relevant terms come from paths of finite length. Unlike crossings between opposite sides, there are many finite paths connecting adjacent sides. For example, the shortest path involves the single corner-most site, which is occupied with probability ρ. Thus $P_{\rm dc}>\rho$, $\forall\rho$. The next shortest path involves the three occupied sites that connects the adjacent sides but leaves the corner site unoccupied. Including the shortest and second shortest path gives us a lower bound $P_{\rm dc}>\rho+\rho^3(1-\rho)$. While this process is far from elegant, it can be extended for an arbitrary number of finite paths to obtain a lower bound of high accuracy. We find that the bound is a continuous function of ρ, increasing monotonically approaching ρ_c. As N increases, we find that $P_{\rm dc}(\rho_c)$ approaches unity (in agreement with Cardy's formula), and obtain a family of curves that converge towards the limiting form illustrated in Fig. 14. By inspection, as $\rho=\rho_c-\epsilon$, approaches ρ_c from the left, $P_{\rm dc}$ approaches unity. However, it loses analyticity at $\rho=\rho_c$ (as is standard for order parameters). By inspection, at $\rho=\rho_c-\epsilon$, $P_{\rm dc}=1-A\epsilon^{\chi}$ to leading order where A is some constant and $0<\chi<1$.

Finally, we note that decreasing the density in the barrier cells further (which decreases the probability of a diagonal connection) is not favored in calculations that optimize yield. Instead, higher yields are obtained when the barriers are at density $\rho_c-\epsilon$, and corner connections are treated as equivalent to edge connections. The worst case scenario is for small M, because the cost in density of additional barrier cells is much greater than for large M.

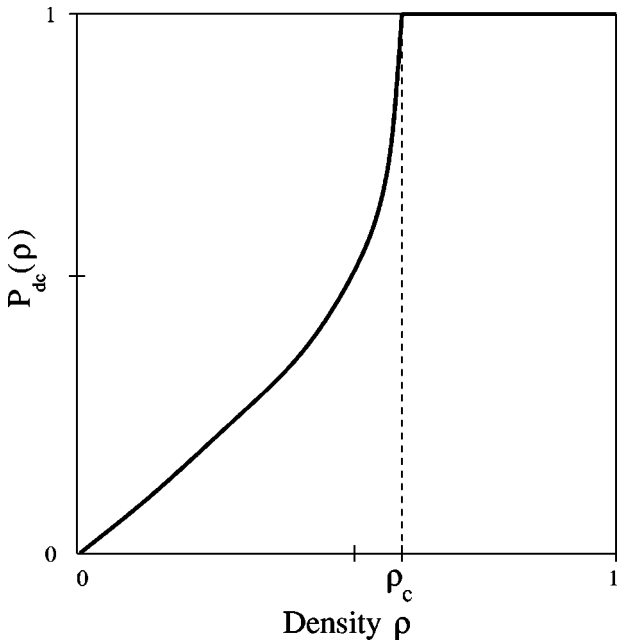

FIG. 14. The probability for crossing between adjacent sides as a function of the density $P_{\rm dc}(\rho)$ is similar to that of $\mathcal{P}_\infty$ in that it behaves like an order parameter. When $\rho=\rho_c-\epsilon$, the diagonal crossing probability has the leading order form $P_{\rm dc}=1-A\epsilon^{\chi}$.

In fact, for the case $M=2$, depicted in Fig. 13(a), any density in the intermediate cells leads to a yield that is below the yield obtained when all cells are at the critical density. To see this, let the total probability that the unit cells are connected be denoted by p. From the previous arguments, $p=2P_{\rm dc}-P_{\rm dc}^2$ and $p\to 1-A^2\epsilon^{2\chi}$ as $\rho\to\rho_c-\epsilon$. In the case of a uniform $P(I,J)$, we obtain the following expression for the yield Y (assuming the intermediate cells are at density $\rho<\rho_c$):

$$Y=\tfrac{1}{8}(1-p)^2+\tfrac{1}{4}+\tfrac{1}{2}\rho. \tag{14}$$

But we know that $1-p<1-\rho$, which implies the following:

$$Y<Y_1=\tfrac{1}{8}(1-\rho)^2+\tfrac{1}{4}+\tfrac{1}{2}\rho. \tag{15}$$

From this estimate, we note that $dY_1/d\rho>0$ on the interval of interest $0<\rho<\rho_c$. Thus, from this we find that $Y_1<Y_{1\,\max}=Y_1(\rho=\rho_c)\approx 0.5668$. However, since $Y<Y_1, Y<Y_{1\,\max}\approx 0.5668<\rho_c$. Therefore, the yield for the configuration in Fig. 13(a) is less than the critical density. Thus, such a configuration is not even a candidate optimal configuration.

For large M [Fig. 13(b)], suppose that the connected clusters R_1 and R_2 are composed of l_1 and l_2 unit design cells, respectively. Without loss of generality, let us assume that $l_1\geq l_2$. Consider the contributions to the yield that arise due to the presence of a corner connection between regions R_1 and R_2, and denote it by $Y_{\rm cc}$ (cc=corner connected). We want to then compare $Y_{\rm cc}$ to the yield $Y_{\rm ncc}$ (ncc=no corner connection) obtained if we change a single unit-density cell in the corner of R_1, to prevent a corner connection. All we then need to do is show that $Y_{\rm ncc}-Y_{\rm cc}>0$. For $Y_{\rm cc}$ we obtain

$$Y_{\rm cc}=\frac{1}{M^2}\left(l_1+l_2+\frac{2l_1l_2}{M^2}(1-p)^2-\frac{(l_1+l_2)^2}{M^2}\right), \quad (16)$$

while for $Y_{\rm ncc}$ we obtain

$$Y_{\rm ncc}=\frac{1}{M^2}\left(l_1+l_2-(1-\rho_c)-\frac{(l_1-1)^2+l_2^2}{M^2}\right). \quad (17)$$

Thus combining Eqs. (16) and (17) we obtain

$$Y_{\rm ncc}-Y_{\rm cc}=\frac{1}{M^2}\left(-(1-\rho_c)+\frac{2l_1l_2}{M^2}[1-(1-p)^2]-\frac{2l_1-1}{M^2}\right). \quad (18)$$

In the large M limit (see Sec. IV B) $l_1, l_2 \sim M^{4/3}$. Consequently, $Y_{\rm ncc}-Y_{\rm cc}>0$ since $l_1l_2/M^2 \sim M^{2/3}$ and clearly dominates Eq. (18). In this case as well, we observe that configurations with corner connections do not produce the highest yields and thus are not optimal configurations.

APPENDIX B: ASYMPTOTIC CALCULATION OF THE OPTIMAL REGION SIZE FOR A UNIFORM DISTRIBUTION OF SPARKS

In this appendix we sketch the key algebraic steps in the calculation of the optimal size m^2 of the unit-density regions for M^2 DDOF's. In particular, we optimize Eq. (9) as a function of m with fixed M, to obtain the optimal m in Eq. (10) to leading order in M, as $M\to\infty$.

Let $\gamma=m/M$. In the limit $M\to\infty$, γ approaches a continuous variable, reflecting the fact that the problem of packing the optimal regions into the design lattice, which constrains our solution for finite M, becomes a negligible problem in the asymptotic limit.

Rewriting Eq. (9) in terms of γ, we obtain

$$Y=\frac{\gamma^2}{\left(\gamma+\frac{1}{M}\right)^2}+\frac{2\gamma+\frac{1}{M}}{M\left(\gamma+\frac{1}{M}\right)^2}\rho_c-\frac{\gamma^4}{\left(\gamma+\frac{1}{M}\right)^2}. \quad (19)$$

In the asymptotic limit, where γ is continuous, the optimal yield satisfies

$$\frac{dY}{d\gamma}=0. \quad (20)$$

Equation (20) applied to Eq. (19) yields

$$m[m^3+2m-M^2(1-\rho_c)]=0, \quad (21)$$

where the solution $m=0$ corresponds to a minimum yield (for $M>0$) at the critical density. In the limit of large M, the linear term in m in the parenthesis can be ignored, and we recover Eqs. (10), from which it is straightforward to deduce Eqs. (11)–(13).

We can also obtain these results by using a scaling argument. Assuming

$$m=AM^{\delta} \quad (22)$$

to lowest order in M^{-1}, we obtain the yield

$$Y=1-2A^{-1}(1-\rho_c)M^{-\delta}-A^2M^{2\delta-2}. \quad (23)$$

Assuming this scaling for m holds asymptotically in M, we seek a solution for δ that sustains the maximum Y for increasing values of M. Note that the exponent δ enters the last two terms on the left-hand side with opposite sign. The scaling of these two terms balances when $\delta=\frac{2}{3}$, which corresponds to the optimal solution. To see this, consider $\delta>\frac{2}{3}$. While the second term is relatively smaller, the exponent in the final term $2\delta-2>-\frac{2}{3}$ is larger, leading to a more rapid decay in yield. Similarly, if $\delta<\frac{2}{3}$ the second to last term dominates, and produces smaller yields. Since the optimal m maximizes yield, we conclude $\delta=\frac{2}{3}$. Additionally, we can fix A by $\partial Y/\partial A=0$ and consequently recover Eqs. (10)–(13).

[1] R. Pool, *Beyond Engineering* (Oxford University Press, Oxford, New York, 1997).

[2] J. T. Bonner, *The Evolution of Complexity by Means of Natural Selection* (Princeton University Press, Princeton, New Jersey, 1988).

[3] G. A. Cohen, D. Pines, and D. Meltzer, *Complexity: Metaphors, Models, and Reality* (Perseus, Cambridge, MA, 1999).

[4] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. Lett. **59**, 381 (1987).

[5] See, e.g., P. Bak, *How Nature Works: The Science of Self-Organized Criticality* (Copernicus, New York, 1996).

[6] D. L. Turcotte, Rep. Prog. Phys. **62**, 1377 (1999).

[7] S. A. Kauffman, *Investigations* (Oxford University Press, Oxford, New York, 2000).

[8] See, e.g., C. Adami, *Introduction to Artificial Life* (Springer-Verlag, New York, 1998), and references therein.

[9] J. M. Carlson and J. Doyle, Phys. Rev. E **60**, 1412 (1999).

[10] J. M. Carlson and J. Doyle, Phys. Rev. Lett. **84**, 2529 (2000).

[11] J. Doyle and J. M. Carlson, Phys. Rev. Lett. **84**, 5656 (2000).

[12] J. M. Carlson and J. Doyle, Proc. Natl. Acad. Sci., U.S.A. **99**, 2538 (2002).

[13] T. Zhou, J. M. Carlson, and J. Doyle, Proc. Natl. Acad. Sci., U.S.A. **99**, 2049 (2002).

[14] P. Bak, K. Chen, and C. Tang, Phys. Rev. A **147**, 290 (1990); B. Drossel and F. Schwabl, Phys. Rev. Lett. **69**, 1629 (1992).

[15] D. Stauffer, *Introduction to Percolation Theory* (Taylor, London, 1985).

[16] B. D. Malamud, G. Morein, and D. L. Turcotte, Science **281**, 1840 (1998).
[17] T. Zhou and J. M. Carlson, Phys. Rev. E **62**, 3197 (2000).
[18] C. Robert, J. Carlson, and J. Doyle, Phys. Rev. E **63**, 056122 (2001).
[19] S. K. Ma, *Modern Theory of Critical Phenomena* (Benjamin, New York, 1976).
[20] J. Cardy, J. Phys. A **25**, L201-6 (1992).
[21] J. Cardy, e-print math-ph/0103018 v2.
[22] R. M. Ziff, Phys. Rev. Lett. **69**, 2670 (1992).
[23] K. Chen and P. Bak, Phys. Rev. E **62**, 1613 (2000).
[24] Corner connections are excluded in optimal solutions as long as $P(I,J)$ is sufficiently smooth. This applies for the exponential $P(x,y)$ we consider here, and other continuous functions. However, $P(I,J)$'s can be constructed for which corner connections are allowed. For example, when $P(I,J)=0$ in alternating cells in a checkerboard pattern, cells with $P(I,J)=0$ will have unit density. Because they are never hit, there is no risk of propagation through a corner connection. We thank the referee for suggesting this example.