Journal of Statistical Physics, Vol. 54, Nos. 5/6, 1989

A Physicist's Sandbox

Kurt Wiesenfeld,[1] Chao Tang,[2] and Per Bak[2]

Received April 28, 1988

We discuss some recent results suggesting that certain spatially extended dynamical systems naturally evolve toward a state characterized by domains of all length scales. The analogy with second-order phase transitions has prompted the name "self-organized criticality"; specific results are available for cellular automaton models, which can be thought of as caricatures of a sandpile undergoing avalances. The potential generality of the results stems from the very simple (nonlinear) diffusion dynamics governing the system.

KEY WORDS: Self-organization; external noise; criticality.

1. INTRODUCTION

This paper concerns the study of spatially extended physical systems. A recent resurgence of interest in this subject stems from the success of the field of nonlinear dynamics, whose rapid progress in describing temporal complexity has spawned vigorous effort.[3]

The present work takes a rather particular point of view—namely, that substantial progress toward understanding the dynamical behavior of spatially extended systems can be achieved via the methods of "modern" dynamical systems theory. This is not intended to sound radical; however, one should recognize that (so far, at least) dynamical systems theory has not led to advances of this sort. Relative to the strides made for temporally complex behavior, progress on spatiotemporal complexity has been maddeningly slow. We are still very much in the "look and see" phase, with no organizing principles to parallel, for example, "routes to chaos," phase

---

[1] Physics Department, Georgia Institute of Technology, Atlanta, Georgia 30332.
[2] Physics Department, Brookhaven National Laboratory, Upton, New York 11973.
[3] In addition to the papers in the present volume, see ref. 1.

0022-4715/89/0300-1441$06.00/0

space reconstruction of attractors, etc. Moreover, powerful tools such as the center manifold reduction seem to play a limited role here.

The main stumbling block is this: the interaction of very many degrees of freedom plays an essential role in spatiotemporal dynamical systems. (Indeed, in those instances when we can say something significant about spatially nontrivial systems, it is because the reduction to a few degrees of freedom is possible.) Therefore, the simple and crucial geometrical insights of few-degree-of-freedom models are insufficient. The utility of carrying over those concepts most useful in the analysis of low-dimensional phase space—e.g., measuring spectra of Liapunov exponents—has thus far proved inconclusive at best. In effect, we have no corresponding geometrical insight into inherently many-degree-of-freedom dynamical systems.

As a first step toward a general understanding, one wishes to uncover examples of spatiotemporal behavior which, though complex, can nevertheless be *recognized* and *characterized.* Here, we present such an example, which is a cellular automaton representing a nonlinear diffusion process. Remarkably, the complexity of its dynamics can be understood and (statistically) quantified.[2–4] The results are very suggestive, insofar as the behavior is reminiscent of the *critical point* of thermodynamic systems undergoing a second-order phase transition. Thus, there is significant structure on all length scales simultaneously, which is characterized by power law distribution functions. The analogy with critical phenomena is conceptually useful, but there is an important difference: in our automaton, the dynamics *naturally evolves* toward the critical point; no tuning of some external parameter is necessary. Consequently, we have called this behavior "self-organized criticality."[2]

The contents of this paper are as follows. In the next section we define the "sandpile automaton" and review its observed behavior. We turn to a discussion of the underlying structure of the automaton's dynamics in Section 3 and suggest why this sort of "grainy" diffusion process might be widespread among spatially extended systems. The temporal fingerprint of self-organized criticality is a power spectrum $S(f) \propto f^{-\phi}$; Section 4 briefly discusses this mechanism of generating "flicker noise" in light of some other proposed mechanisms which are also based solely on general ideas from nonlinear dynamics. Section 5 proposes avenues ripe for further study of the phenomenon of self-organized criticality, including the possibility of laboratory verification. Finally, we close with a description of a demonstration illustrating these ideas that can be performed in the comfort of one's own home.

2. SANDPILE AUTOMATON

In this section, we describe simulations for the "sandpile automaton," which is a type of nonlinear diffusion equation.[2] Although the name is partly for fun, the underlying issues are usefully thought of in terms of the dynamics of sand, for which we all have a certain intuition.

Imagine building up a sandpile by lightly sprinkling sand a few grains at a time (Fig. 1). Eventually, a mound is formed, which is somewhat rough in its shape, varying from spot to spot, but has some mean slope. Of course, this slope cannot be too great, for then sand would flow. (The maximum such slope is called the "angle of repose.") The addition of more sand is unpredictible: it may simply nestle where it is sprinkled, or it may start a small rearrangement of sand, or it may on occasion trigger a more far-reaching collapse of the sandpile—in other words, an avalanche. The question we want to answer is: what is the distribution of sandslides induced by the addition of sand, once the system has reached statistical steady state?

We begin with a one-dimensional model of this dynamics (Fig. 2). We divide space up into an array of cells; the sand is allowed to pile up in a

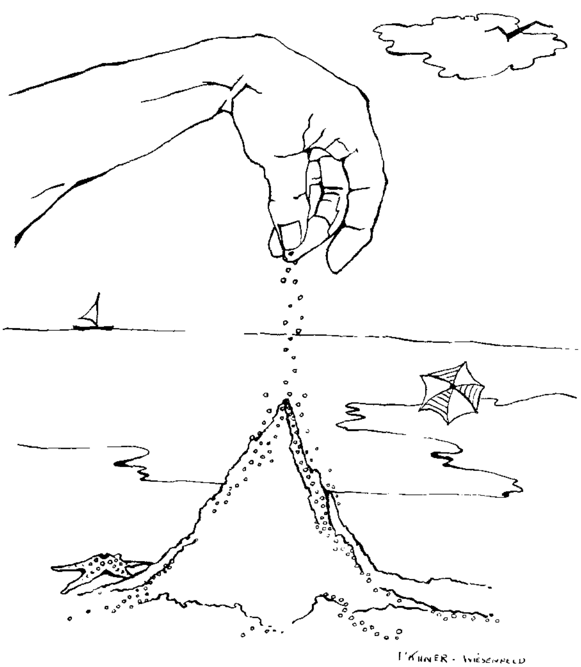

Fig. 1. The formation of a sandpile: artist's conception.

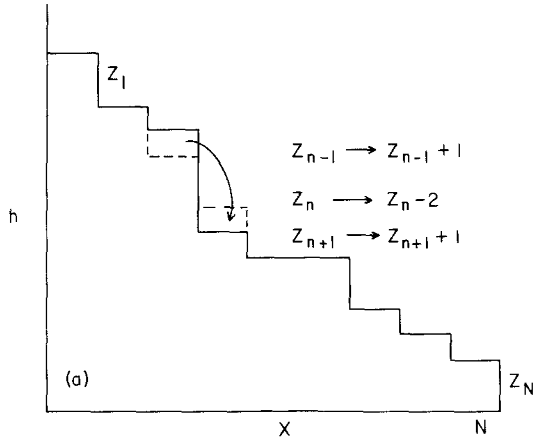

Fig. 2. One-dimensional "sandpile" automaton. The state of the system is specified by an array of integers representing the height difference between neighboring plateaus.

column above each cell, the height taking on only integer values. The relevant dynamical variable is the local slope z_n, which is taken to be the height difference of the columns to the left and the right of the point n.

There are two operations. The first is the addition of one unit of sand: from the picture, this corresponds to a change in two neighboring slope variables:

$$\begin{aligned} z_n &\to z_n + 1 \\ z_{n-1} &\to z_{n-1} - 1 \end{aligned} \tag{1}$$

The other operation is the relaxation of the sandpile in case the slope locally exceeds the threshold z^*, causing one unit of sand to slide downhill. As depicted in Fig. 2, this corresponds to the change

$$\begin{aligned} z_n &\to z_n - 2 \\ z_{n+1} &\to z_{n+1} + 1 \\ z_{n-1} &\to z_{n-1} - 1 \end{aligned} \tag{2}$$

provided $z_n > z^*$, otherwise there is no change. If this looks like a (discrete) diffusion operator, that is because it is, as described further in the next section. The threshold condition is what makes it a *nonlinear* diffusion dynamics.

Finally, we need to specify boundary conditions. We put a wall at one end (so that sand never flows out the left side),

$$z_0 = 0$$

and an edge at the other (so that tumbling sand simply "disappears" at the right side):

$$z_N \to z_N - 1; \qquad z_{N-1} \to z_{N-1} + 1$$

when $z_N > z^*$. This choice of boundary conditions is but one possibility; the net effect is to create an overall flow of sand from left to right.

These rules define the dynamics of the cellular automaton, starting from any initial configuration $\{z_n\}$. If any of the slopes z_n are above threshold, then the system is allowed to relax via successive application of rule (2), until $z_n \leqslant z^*$ for all n. Then, another unit of sand is added, say at a randomly chosen position, and then relaxation is applied (if necessary), and so on.

The results of such a simulation are interesting; in fact, this 1D case is simple enough that one can solve the dynamics exactly. For any initial configuration, the sandpile evolves to its *minimally stable state*, defined by $z_n = z^*$ for all n. The effect of adding sand to this state is to cause this unit to slide from site to site, until it falls off the edge at $n = N$, leaving the pile at its original state. We call this the minimally stable state, since it is the only configuration in which addition of sand necessarily causes a sandslide. We may say that this state is at once sensitive and yet robust with respect to external noise: sensitive, because a local perturbation propagates globally throughout the system; robust, because the configuration $\{z_n\}$ is ultimately unaffected by the perturbation.

The emergence of the minimally stable state in one spatial dimension is curious; however, the *dynamics* of this state is rather bland. Far more interesting is the situation in two or more spatial dimensions, as we now discuss.

The automaton rules (1), (2) can be generalized to more than one spatial dimension. For example, in 2D, addition of sand corresponds to

$$\begin{aligned} z(x-1, y) &\to z(x-1, y) - 1 \\ z(x, y-1) &\to z(x, y-1) - 1 \\ z(x, y) &\to z(x, y) + 2 \end{aligned} \tag{3}$$

and relaxation to

$$\begin{aligned} z(x, y) &\to z(x, y) - 4 \\ z(x, y \pm 1) &\to z(x, y \pm 1) + 1 \\ z(x \pm 1, y) &\to z(x \pm 1, y) + 1 \end{aligned} \tag{4}$$

provided $z(x, y) > z^*$, where we have the square array (x, y) for $x, y = 1, ..., N$. In fact, the correspondence between the sandpile picture and these diffusion rules is not as straightforward as in the 1D case; after all, the slope is no longer a scalar field, so that the threshold condition is properly one on the gradient of the height function. Although a somewhat contorted correspondence can be concocted,[2] it is best to consider this as being the simplest discrete scalar diffusion dynamics in two dimensions. (Ultimately, our interest is not in sandpiles *per se*, nor in the behavior of any particular cellular automaton; rather, we are looking for behavior that might prove typical of some large class of dynamical systems with spatial degrees of freedom. In this sense, one may view these rules as analogous to "Ising-type" models studied in statistical mechanics.)

Some results are summarized in Figs. 3–5, as we now describe.

The most important observation is that the minimally stable state—with $z(x, y) = z^*$ for all (x, y)—is no longer robust in spatial dimensions greater that one. A perturbation of this state triggers a catastrophic avalanche extending throughout the system. Instead, the system settles down to a statistically stationary state, such that *on average* for each unit of sand added one unit falls off the boundary. This stationary state consists of locally connected *domains*, such that disturbance at a single site induces a chain reaction affecting all cells within this domain. The fascinating point is this: one observes that, at any instant of time, *domains of all sizes exist*

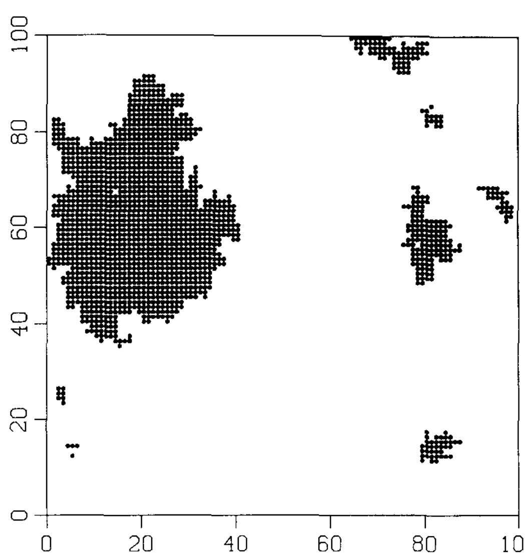

Fig. 3. Typical domain structure of the self-organized critical state. Each cluster is triggered by a single perturbation at one of the sites of the 100 × 100 array.

simultaneously. Figure 3 illustrates this qualitative structure, showing the set of sites triggered by adding a single unit at each of several different sites.

On a more quantitative note, one can ask for the distribution function $D(s)$ of domains of size s. The answer is that the distribution obeys a power law, reminiscent of the domain structure of systems tuned to the point of a second-order phase transition. Figure 4 shows a typical distribution function for an array of modest size, 50×50 sites, built up "from scratch," i.e., with initial conditions $z(x, y) = 0$ at all sites. One can see that this fits a pretty respectable power law for two decades, up to an area of about $s = 300$; the falloff at large s is a finite-size effect. We believe that the dynamics has driven the system just to the point where structure on all length scales is sustained, so that a simulation on a very large array will produce a power law distribution function up to correspondingly large domain sizes.

The self-organized critical state is also acheived in simulations in three spatial dimensions. Figure 5 shows the results of a simulation for a 3D array, again showing a power law distribution, but with a different exponent than in the 2D case.

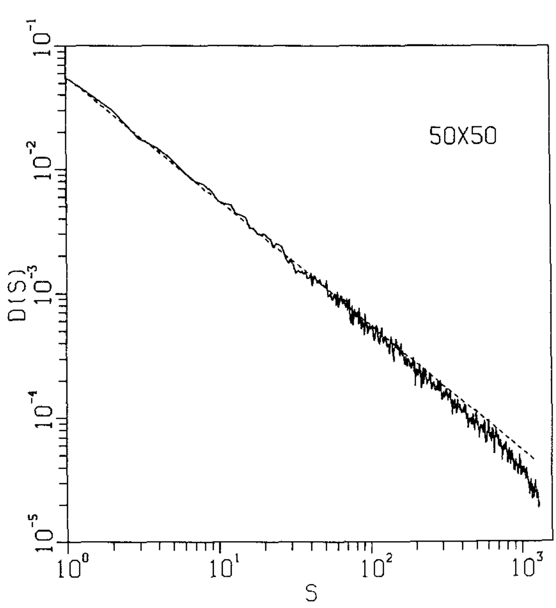

Fig. 4. Domain size distribution for the sandpile built from scratch, 50×50 array. Once the automaton has reached its statistical stationary state, 100,000 units of sand have been added (one at a time) to induce this same number of avalanches. The data have been coarse grained, and fit a slope of -1.0.

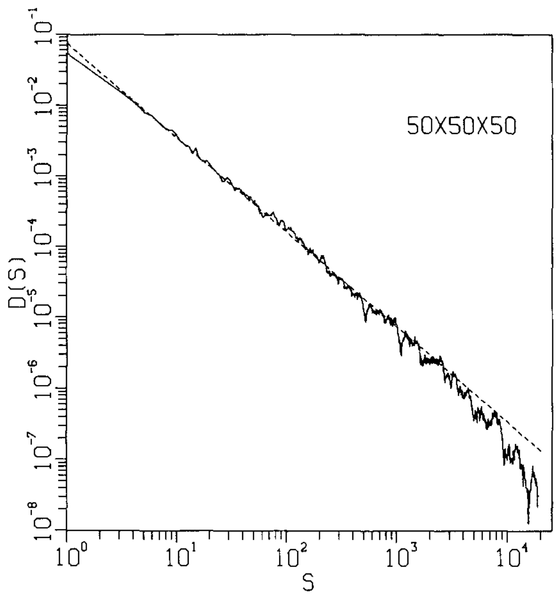

Fig. 5. Domain size distribution for a $50 \times 50 \times 50$ array; 50,000 units of sand have been added after reaching the stationary state. The data have been coarse grained, and fit a slope of -1.33.

The fact that this (statistical) dynamical critical state is self-organized is a crucial point; the situation stands in contrast to the power law distributions observed in equilibrium thermodynamic systems poised at a second-order phase transition point. Somehow, the dynamics naturally carries the sandpile system to its critical point, i.e., the critical state is an *attractor* of the dynamics.

To see that this state of affairs is not peculiar to one set of rules, it is important that variations of the model do not alter the basic features. We have tested a number of such variations (though naturally an infinite number of other rules remain). Without belaboring the point, we have found that the appearance of a self-organized critical point is a robust property of this nonlinear diffusion process. For example, one can start from different initial configurations; one can change the boundary conditions; one can alter the perturbation rule for "adding sand"; and one can introduce quenched disorder by removing (at random) certain bonds in the square lattice (so that sand cannot flow along certain bonds). For all these variations, we have observed power law distribution functions[2] for various quantities such as the avalanche distribution function $D(s)$.

A detailed description of many of these results can be found in ref. 2.

3. DISCUSSION: GOING BEYOND THE SANDPAIL

Naturally, the motivation for studying the above automaton goes beyond the particular issue of avalanches in sandpiles. In this section we want to indicate the general diffusion dynamics underlying the model, and explain why it is this represents a potentially important class of systems. The main point is this: *this picture represents the coarse-grained phase space dynamics—both spatially and temporally—of spatially extended (dissipative) dynamical systems.*

The first step is to rewrite the automaton rules (1)–(4) in a more familiar dynamical form. Let $\Psi(n, t)$ be some scalar function of discrete space n and discrete time t. The usual diffusion equation becomes, allowing for a nonconstant diffusion parameter K and a random noise ξ,

$$\Psi(n, t+1) - \Psi(n, t) = \mathscr{D}[K\mathscr{D}[\Psi]] + \xi(n, t) \tag{5}$$

where $\mathscr{D}$ is the discrete spatial derivative, e.g., in one dimension $\mathscr{D}[\Psi(n, t)] = \Psi(n+1) - \Psi(n)$. (Note that in our automaton ξ is nonnegative.) The diffusion "parameter" K is really a *functional* with threshold behavior, so that it is zero if the slope $\mathscr{D}[\Psi]$ is below some critical value, and a constant otherwise. Operating on both sides by the linear operator $\mathscr{D}$, one finds that Eq. (5) becomes an equation for $z = \mathscr{D}[\Psi]$, to wit

$$z(n, t+1) - z(n, t) = \mathscr{D}^2[K(z)z] + \mathscr{D}[\xi(n, t)] \tag{6}$$

and this is precisely the automaton dynamics (1), (2) in 1D and (3), (4) in 2D; it is easily implemented in higher dimensions, of course.

The next question concerns the appropriateness of a threshold-type behavior for the diffusion parameter. For the sandpile problem this seems quite natural, but there is a more general context in which this behavior is relevant. Imagine a phase space filled with very many attractors. (For example, each attractor may correspond to a particular time-independent spatial configuration of the sandpile.) Then the "microscopic" time-continuous dynamics finds its way to some attractor on a short time scale; small perturbations have no effect, since the system relaxes back to the local attractor. Only on a longer time scale will a sufficiently large perturbation cause the system to hop into some other basin of attraction, and thus into a different long-lived configuration. If we only pay attention to the coarse-grained dynamical evolution, then the system obeys a diffusion equation of the sort simulated here. On this time scale, and in the presence of noise, each attractor is viewed as being a metastable state of the dynamics.

The existence of huge numbers of metastable states is known to be a common feature in spatially disordered many-body systems such as spin glasses. However, it *also* seems to be a common feature of more general dissipative dynamical systems possessing many degrees of freedom. For example, evidence for huge numbers of coexisting attractors are observed in coupled maps, arrays of ordinary differential equations, and in partial differential equations. This phenomenon has been termed "attractor proliferation," and may prove central in understanding the dynamics of many spatially extended systems.[5] In trying to track down the long-term statistical behavior of such a system in the presence of noise, one is naturally led to dynamics typified by the sandpile automaton.

The picture of phase space trajectories in an attractor-filled space leads immediately to the concept of "dynamical selection of minimally stable states."[6,7] As long as the total set of attractors fits into a compact phase space volume Γ, one can ask for the expected final state if the system's initial conditions are far from equilibrium (i.e., outside the attracting volume Γ). Heuristically (see Fig. 6), in the absence of monstrous basins of

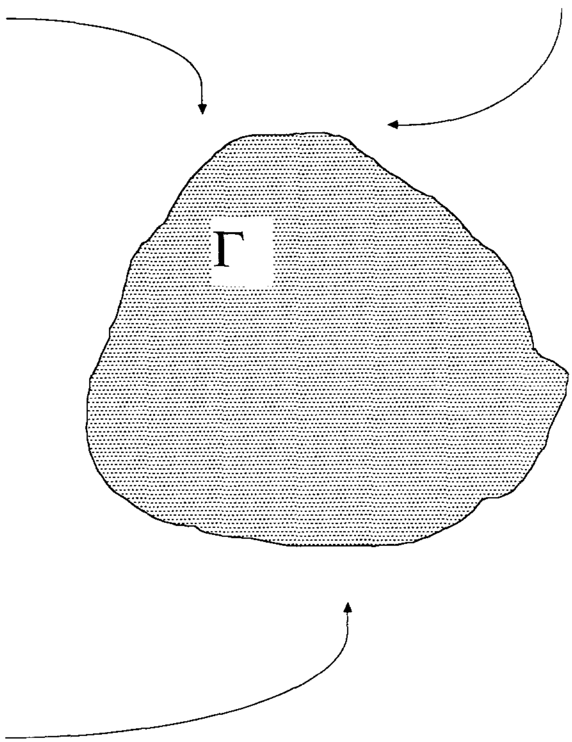

Fig. 6. Heuristic demonstration of selection of minimally stable attractors. A very large number of attractors occupies a compact phase space volume Γ. Trajectories outside Γ are likely to be "trapped" by those attractors near the boundary of Γ.

attraction, the deterministic trajectory is likely to get trapped by some attractor near the surface of Γ. (Of course, the appeareance of peculiar "fingered" basins for interior attractors could alter this simple picture.)

What is missing from the above picture is the effect of *spatial dimensionality*. This is a feature not made explicit in phase space plots; nevertheless, we expect dimensionality to play a significant role based on physical grounds. What we have at this stage is little more than the simulations reported above: in 1D the minimally stable state is achieved, but in higher dimensions the minimally stable state is overly sensitive to noise, and the system rather achieves a statistical steady state—visiting some set of metastable states. The ensuing steady-state fluctuations give rise to power law distribution functions, reminiscent of thermodynamic systems tuned to their second-order phase transition points. The facinating thing is that the present system *naturally* evolves toward this critical state without tuning some external parameter. In effect, the dynamics provides some feedback to the evolving domain structure, and the feedback drives the system to its statistically stationary state.

4. TEMPORAL BEHAVIOR AND 1/*f* NOISE

The results of our simulations show that the system naturally evolves toward a steady state which is characterized by spatial domains of all sizes. The size distribution of these minimally stable domains fits a power law. We can also inquire as to the temporal behavior of the self-organized critical state. It turns out that the power spectrum also exhibits power law behavior, $S(f) \sim f^{-\phi}$, with ϕ between 0 and 2; in other words, the temporal signature of self-organized criticality is "$1/f$ noise." (Indeed, a mean field theory of the sandpile yield an exponent of $\phi = 1$.)[3,4]

We should emphasize that the power law behavior of $S(f)$ is neither the goal nor the main result of these models; rather, it is by-product of the self-organized critical state, being a direct consequence of the spatial scaling structure. Nevertheless, one can hardly say the words "one over f" without drawing a crowd; and of course, it is one of the great mental distractions of physics. For this reason, we devote this section to a brief discussion of this aspect of the automaton behavior, including some comments concerning the possible connection to some other general dynamical $1/f$ mechanisms recently reported.

Consider the automaton in its steady state. The existence of domains of all sizes immediately implies temporal fluctuations on all time scales, even with input perturbations lasting only a single time step. This is easy to understand, since the domains (by virtue of their minimal stability) have the property that a disturbance at a single site spreads throughout the domain: the larger the domain, the longer lived the disturbance.

This expectation is borne out in simulations.[2] Suppose one upsets a single site, and measures the lifetime of the ensuing avalanche. After the disturbance dies out, one upsets some other site, and so on. (There is no need to "reset" the configuration, since the critical state is perforce a self-maintaining, dynamic balance.) Figure 7 shows the resulting distribution of lifetimes; it follows a tolerable power law. The response to a continual "white noise" input perturbation—with each kick localized to a single site for a single time step—will be a superposition of the individual responses provided the perturbations are sufficiently dilute. (We do not know what happens if the input noise is so strong that a single domain is externally excited many places at once, but the simple-minded superposition picture may well break down in this regime.) Figures 8 and 9 show a representative time series and power spectrum, respectively.

The idea that a distribution of relaxation times, operating simultaneously and independently, can lead to $1/f$ noise is an old one.[8] The usual van der Ziel argument based on linear relaxation processes has to be modified in its details in the present context,[2] but the basic idea is nevertheless the same. Again, the interesting new feature is that the dynamics naturally carries the spatially extended system into the scaling

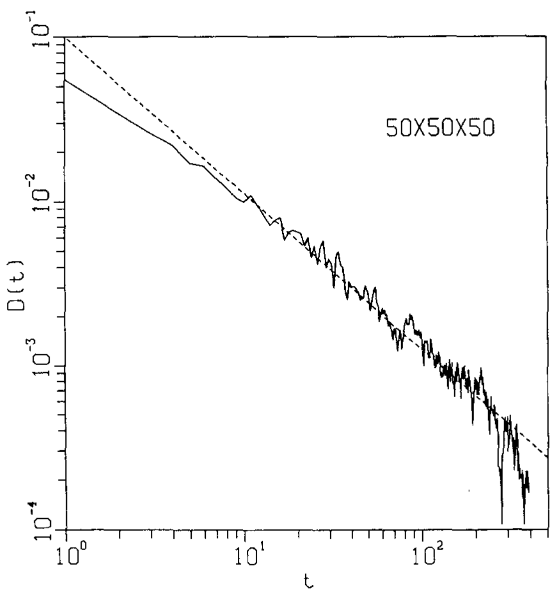

Fig. 7. Distribution of lifetimes for a $50 \times 50 \times 50$ array; the dashed line has slope -0.95.

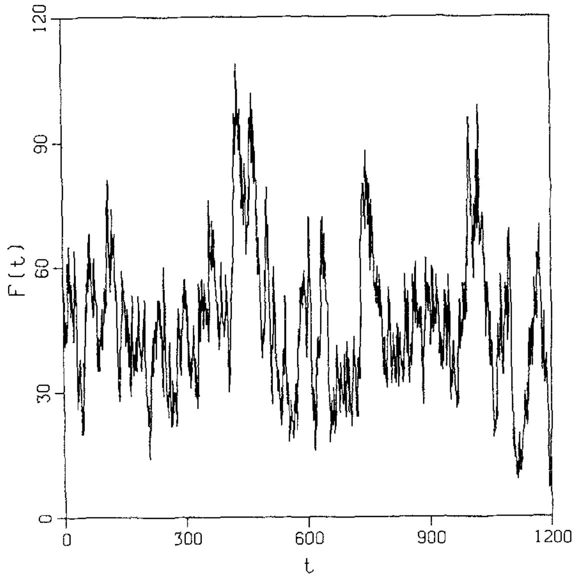

Fig. 8. Time series generated by superposing randomly the response represented by individually perturbed domains, for a $20 \times 20 \times 20$ array. Note the fluctuations on a wide range of time scales.

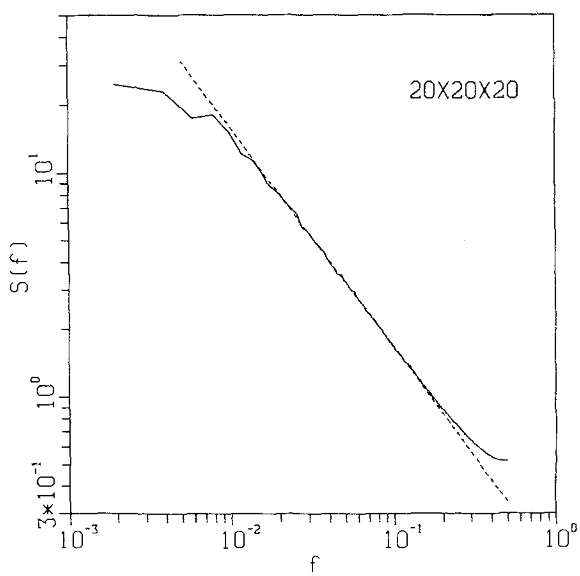

Fig. 9. Power spectrum for the time series of Fig. 8. The dotted line has slope -0.98. The crossover to white noise at very low frequencies is a finite-size effect.

regime, without the need to specify *a priori* any distribution of relaxation times.

The appeal of such a "generic" dynamical picture lies in the purported "generic" nature of $1/f$ noise.[9] The sheer breadth of phenomena which (we are told) show $1/f$ power spectra is mystifying: the light from quasars, the intensity of sunspots, the flow of sand in an hourglass, the water level of the Nile river, the membrane potential of a lobster neuron, and even the traffic on a Japanese highway, have been reported as displaying $1/f$ fluctuations.

Now, *if* one is ready to regard the wide occurrence of $1/f$ noise as being due to some common underlying mechanism, then one must search for some explanation almost entirely independent of physical details. ("Almost," since not *all* systems display $1/f$ noise.) For this reason, the field of nonlinear dynamics seems a good place to search for some such "universal" mechanism. We would like to discuss the present mechanism very briefly in the context of three other nonlinear dynamical mechanisms that have recently appeared in the literature.

Perhaps the simplest dynamical systems model yielding $1/f$ noise is the one-dimensional iterative map of the interval studied by Manneville.[10] This represents a *special case* of the generic intermittency threshold, in which the tangent point coincides with the endpoint of the interval. Manneville showed that the ensuing chaotic dynamics has a power spectrum that varies as $1/f(\ln f)^2$ at low frequencies; later, Miracky *et al.* modified the map to get a precise $f^{-\phi}$, $\phi = 1$, dependence.[11] At first blush, this looks quite promising, since we know that simple iterative maps are good models for a wide variety of physical systems. Unfortunately, this mechanism has the failing that it requires precise tuning of some external parameter. That is, the problem is simply shifted to explaining why systems typically might be sitting at this special parameter value. As we have emphasized, the sandpile automaton automatically evolves toward the critical state (thus the term "self-organized").

Keeler and Farmer discussed precisely this point in their work on an array of coupled logistic maps,[12] and showed that their system displayed robust intermittency—that is, intermittency that persists over a whole parameter interval. Their simulations showed that the low-frequency macroscopic behavior of the array fits a $f^{-\phi}$ power law, with a relatively large ϕ (a value of $\phi \sim 2.5$ is quoted in their paper, though this exponent apparently varies with the external parameter value). Consequently, this model has the same basic virtues as our sandpile: are there any differences? Perhaps the most fundamental difference is the *spatial* behavior of these systems: Keeler and Farmer show that the logistic array displays "a stable natural wavelength," whereas the sandpile automaton has no characteristic length scale. Thus, an experiment which tests the correlations between temporal and spatial fluctuations could decide between these two alternatives.

Finally, Huberman and Kerszberg showed how relaxation in hierarchical structures can lead to anomalous low-frequency noise spectra.[13] Under certain circumstances—that is, for certain hierarchical architectures—$1/f$ noise can result. The basic idea is that diffusion between different levels in the hierarchy generates a hierarchy of time scales. It is difficult to see how our sandpile's critical state—which seems to be well described by a collection of independent domains—corresponds to some "ultrametric topology." On the other hand, it is possible that at some more abstract level, the underlying dynamics may be expressible in these terms. Nevertheless, at present it seems simpler to regard the sandpile's $1/f$ fluctuations in terms of the old-fashioned picture of van der Ziel's grabbag of independent relaxation times.

5. FUTURE VISTAS

Our current understanding of the sandpile automaton leads us to suspect that certain features may be "generic" to spatially extended dynamical systems. The main goal, then, is to extract which features these are, and to abstract those organizing principles responsible for them. As we now outline, several avenues suggest themselves for future study; our understanding of spatiotemporal dynamics in general and self-organized criticality in particular is sufficiently immature that any such further work would be valuable.

The question arises as to whether the "critical exponents" measured for the automata are in any sense "universal," i.e., do their precise values depend on the details of the governing dynamical equations (in which case they are not universal) rather than being determined only by gross distinctions such as the spatial dimensionality and symmetry (in which case there are universality classes)? Our present understanding leads us to expect power law distribution functions, but not necessarily universal values for the exponents; however, as further rules are tested there may well emerge a pattern to the observed values, which would suggest an underlying structure waiting to be discovered. One approach to this problem is to develop a "mean field" theory to establish "classical" exponents; this goal has very recently been accomplished.[3,4]

A second line of development is to find a still simpler model that is in some sense more solvable, so that one need not rely solely on numerical simulations. For example, one can remove the inherent randomness of the present model by always adding sand at the center of the sandpile, yielding the so-called "central seeding model."[14] The behavior of this automaton is quite fascinating: the dynamics is purely periodic in time, and the pile passes through its minimally stable state during this cycle (followed, of course, by a catastrophic avalanche). Nevertheless, the distribution of avalanche sizes still follows a power law. The hope is that this model is sufficiently regular to allow for a detailed understanding of its dynamical behavior.

Along these lines, one can try to find a "minimal" model that displays the essential features, for example, an automaton which has only two states per cell (as does the elementary cellular automaton most often encountered in the technical literature and on T-shirts). Alternatively, one may revert to one spatial dimension with interactions extending beyond nearest neighbors. (Some preliminary studies report that it is possible to have interactions that destabilize the minimally stable configuration, again leading to power law distribution functions.[15])

It is also of interest to establish which features are somehow remnants of the gross truncations inherent to the cellular automaton discretized space, time, and state variables. For example, essential to the evolving picture of self-organized criticality is the presence of a huge number of stable or metastable states: can this be a property of smooth dynamical systems typified by ordinary and partial differential equations? The answer to this seems to be "yes"; in fact, the existence of huge numbers of phase space attractors seems to be a fairly common property of spatially coupled arrays of nonlinear dynamical systems. The overall ramifications of this wealth of attractors are unclear, but the emergence of self-organized critical structures is at least a possible outcome, and potentially a generic result.

Finally, we come to the question of experimental realization of these ideas. In a sense, the present system is better adapted to direct experimental testing than other recent spatiotemporal models (such as coupled logistic maps, or one-dimensional reaction-diffusion equations), in large part due to the global predictions of the statistically stationary state and the correlations between spatial and temporal fluctuations. In particular, the sandpile picture can be taken literally (though we recommend that anyone doing this does so with a grain of salt), and examined directly in the laboratory.[16] However, other systems are available which possess a threshold-type nonlinear relaxation reminiscent of the present automaton dynamics. For example, it has been suggested that magnetic flux creep in type II superconductors fits this picture.[17] The basic idea is that impurities can pin magnetic flux inside the bulk material, allowing the superconductor to maintain a far-from-equilibrium stationary state; relaxation proceeds only if the local vortex gradient exceeds a critical threshold value (see, e.g., ref. 18). In fact, flux is expelled in a series of "jumps," as flux bundles of varying magnitudes Φ become depinned. It is possible to measure the distribution $D(\Phi)\,d\Phi$ of these jumps, as the system relaxes due to thermal activation; the signature of self-organized criticality would be that $D(\Phi)$ is proportional to some power of Φ.[4]

Beyond such specifically tailored systems, our suspicion is that self-organized criticality is a consequence of only some general dynamical circumstances, and as such may be ubiquitous. In any event, further work is required to bring about any deeper understanding of the connection between the sandpile automaton and particular physical realizations.

6. A HOME EXPERIMENT

As we have repeatedly emphasized, the ideas suggested by the dynamics of our nonlinear diffusion automaton are of interest mainly because of their potential generality. The concepts of domains possessing local minimal stability, pieced together to form a globally self-organized critical state; the power law spatial distribution functions; and the accompanying "flicker noise" temporal fluctuation spectrum are the main features we suggest may be found in many spatially extended nonequilibrium systems. From experience in both critical phenomena and low-dimensional dynamical systems, the possibility of genericity and universality lead us to take seriously the results we have presented, with tongue in cheek, using the language of "sandpiles."

Having said this, however, it remains true that real sandpiles offer an extraordinarily charming demonstration, and in the comfort of one's own home. Indeed, some sophisticated versions of sandpile-type experiments have been reported recently (using both glass microspheres and irregularly shaped aluminum oxide grains)[(16)]; however, we now describe a poor man's version, aimed more toward leisure-time observation than toward precision measurement.

To demonstrate self-organized criticality, one needs a shoebox and a cup or two of sand; salt or pepper will do in a pinch. The sand should be gathered up as steeply as possible into one corner of the box. One can try to directly mimic the sandpile automaton by sprinkling additional sand onto the peak; however, a nice alternative is to *very lightly* dampen the sand with water. The angle of repose (i.e., the threshold slope) is larger for wet sand, so as the water evaporates one observes a sequence of slides—some very small, others quite large—occurring at random sites. The evaporation process can be sped up by placing the box on a warm surface, or under direct sunlight.

[4] Actually, pure thermal activation would lead to a gradual but inexorable degradation toward the zero flux equilibrium state. To maintain critically indefinitely, one needs a mechanism to (randomly) add flux to compensate this leakage.

ACKNOWLEDGMENTS

We would like to thank Elaine Wiesenfeld for creating both Fig. 1 and one of the authors. We also thank Mac Beasley for suggesting the connection between these ideas and flux creep in type II superconductors, and Leo Kadanoff for discussing his group's results prior to publication.

REFERENCES

1. Spatio-temporal Coherence and Chaos in Physical Systems, *Physica* **23D**:1–482 (1987).
2. P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**:381 (1987); *Phys. Rev. A* **38**:364 (1988).
3. C. Tang and P. Bak, *Phys. Rev. Lett.* **60**:2347 (1988).
4. C. Tang and P. Bak, *J. Stat. Phys.* **51**:797 (1988).
5. K. Wiesenfeld and P. Hadley, unpublished.
6. S. N. Coppersmith and P. B. Littlewood, *Phys. Rev. B* **36**:311 (1987).
7. C. Tang, K. Wiesenfeld, P. Bak, S. N. Coppersmith, and P. B. Littlewood, *Phys. Rev. Lett.* **58**:1161 (1987).
8. A. van der Ziel, *Physica* **16**:359 (1950).
9. W. H. Press, *Commun. Mod. Phys. C* **7**:103 (1978).
10. P. Manneville, *J. Phys.* (Paris) **41**:1235 (1980).
11. R. F. Miracky, M. Devoret, and J. Clarke, *Phys. Rev. A* **31**:2509 (1985).
12. J. D. Keeler and J. D. Farmer, *Physica* **23D**:413 (1987).
13. B. A. Huberman and M. Kerszberg, *J. Phys. A* **18**:L331 (1985).
14. K. Wiesenfeld and B. McNamara, unpublished.
15. L. Wu, S. Zhou, C. Liu, L. P. Kadanoff, and S. R. Nagel, *Bull. Am. Phys. Soc.* **33**:534 (1988); L. P. Kadanoff, private communication.
16. M. Jaeger, C. Liu, and S. R. Nagel, *Bull. Am. Phys. Soc.* **33**:534 (1988); *Phys. Rev. Lett.* **62**:40 (1989).
17. M. R. Beasley, private communication.
18. R. P. Huebener, *Magnetic Flux Structures in Superconductors* (Springer-Verlag, Berlin, 1979).