Multiple attractors, long chaotic transients, and failure in small-world networks of excitable neurons

Hermann Riecke
Department of Engineering Sciences and Applied Mathematics, Northwestern University, Evanston, Illinois 60208

Alex Roxin
Computational Neuroscience, Departament de Tecnologia, Universitat Pompeu Fabra, 08003 Barcelona, Spain

Santiago Madruga
Max-Planck-Institute for Physics of Complex Systems, D-01187 Dresden, Germany

Sara A. Solla
Department of Physiology and Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208

We study the dynamical states of a small-world network of recurrently coupled excitable neurons, through both numerical and analytical methods. The dynamics of this system depend mostly on both the number of long-range connections or "shortcuts", and the delay associated with neuronal interactions. We find that persistent activity emerges at low density of shortcuts, and that the system undergoes a transition to failure as their density reaches a critical value. The state of persistent activity below this transition consists of multiple stable periodic attractors, whose number increases at least as fast as the number of neurons in the network. At large shortcut density and for long enough delays the network dynamics exhibit exceedingly long chaotic transients, whose failure times follow a stretched exponential distribution. We show that this functional form arises for the ensemble-averaged activity if the failure time for each individual network realization is exponentially distributed. © *2007 American Institute of Physics.* [DOI: 10.1063/1.2743611]

Many systems in nature can be described as a network of interconnected nodes. Networks in a growing list of systems, from social and ecological webs to the neural anatomy of simple organisms, have been shown to exhibit complex topological features that distinguish them from both ordered lattices and purely random networks. Beyond the investigation of the structural and geometrical properties of such networks, a new class of question arises when dynamical degrees of freedom are placed at their nodes. As the investigation of such dynamical complex networks proceeds, it has become increasingly clear that the network architecture can significantly influence the dynamics of the system. An understanding of emergent dynamics on complex networks requires investigating the interplay between the intrinsic dynamics of the node elements and the connectivity of the network in which they are embedded. In order to address some of these questions in a specific scenario of relevance to the dynamical states of neural ensembles, we study here the collective behavior of excitable model neurons in a network with small-world topology. The small-world network has local lattice order, but includes a number of randomly placed connections that may provide connectivity shortcuts. This topology bears a schematic resemblance to the connectivity of the cerebral cortex, in which neurons are most strongly coupled to nearby cells within $50-100$ μm, but also make projections to cells millimeters away. We find that the dynamics of this small-world network of excitable neurons depend mostly on both the density of shortcuts and the delay associated with neuronal projections. In the regime of low shortcut density, the system exhibits persistent activity in the form of propagating waves, which annihilate upon collision and are spawned anew via the reinjection of activity through shortcut connections. As the density of shortcuts reaches a critical value, the system undergoes a transition to failure. The critical shortcut density results from matching the time associated with a recurrent path through the network to an intrinsic recovery time of the individual neurons. Furthermore, if the delay associated with neuronal interactions is sufficiently long, activity re-emerges above the critical density of shortcuts. The activity in this regime exhibits long, chaotic transients composed of noisy, large-amplitude population bursts. A numerical investigation of the interplay between network topology and interaction delays in this regime reveals a mechanism that underlies the observed stretched-exponential distribution of failure times for the chaotic network activity.

I. INTRODUCTION

It has been widely recognized that the connectivity of a network of active elements has a profound impact on its function. Substantial effort has therefore been devoted to the characterization of network connectivity,[1,2] leading to the identification of various measures that are significant in de-

termining the geometrical properties of the system. Particularly relevant among them are the average and maximal length of the minimal paths that connect two arbitrary nodes in the network, the clustering coefficient, which characterizes the propensity of all neighbors of a given node to be connected to each other, and the distribution for the degree, defined as the number of links that emanate from a node.

A large number of networks with complex topology fall into the class of small-world networks, characterized by short average path length and high clustering coefficient. A simple realization of such a small-world network consists of a regular lattice supplemented by a number of randomly placed connections that tend to provide connectivity shortcuts. This topology schematically resembles the connectivity of the cerebral cortex, in which neurons are most strongly coupled to nearby cells within $50-100$ μm, but also project to distant cells, which can be millimeters away. Small-world properties of neuronal networks have been found in a comprehensive morphological characterization of *in vitro* two-dimensional networks,[3] and in a 1:1 network model of the rat dentate gyrus that incorporates known data about cell types, cell-specific connectivity, and axonal branch length.[4]

The dynamics of elements coupled through the connectivity of a complex network has been studied in detail for the case of oscillatory elements. The emphasis has been in determining the role of network topology in the ability to achieve synchronization in a system of coupled oscillators. The existence of long-range connections, which reduce the effective size of the network, has been found to substantially enhance synchronizability.[5,6] At the same time, the heterogeneity of the degree distribution found in many complex networks limits the ability of the oscillators to synchronize. Stable synchronization emerges from a balance between these two competing aspects of complex network topology.[7]

Excitable elements are the components of another important class of dynamical systems, whose emergent behavior is not characterized by synchronization. Locally coupled networks of excitable elements exhibit traveling waves (e.g., Refs. 8–10). If these waves annihilate upon collision, as is typically the case, persistent activity usually requires either an external drive or spontaneous excitation by noise. In models for neural systems driven by noise,[11–13] networks with nonlocal connections between their elements exhibit a tendency towards relatively ordered oscillations in the population activity. The spatial structure of such noise-induced waves becomes less coherent with an increase in the fraction of nonlocal connections.[14] The combination of local connectivity with a small number of nonlocal connections allows a time-periodic localized external input to entrain the whole system much faster than in a purely local network; at the same time, the oscillations are much more coherent than in a truly random network.[15,16] Oscillatory activity at the population level has also been observed in small-world networks of binary McCulloch-Pitts neurons connected through both excitatory and inhibitory synapses.[17]

The dynamics of networks of excitable elements depend very strongly on both the range of the coupling and on the length of the refractory period relative to the time scales associated with propagation. In Ref. 18, the authors consider both scale free and random networks of three-state excitable neurons, and find that short refractory periods enable the propagation of sustained activity through the activation of short loops. These short loops are frequently found in many types of complex and random networks. In the noiseless case, they can only be activated if the initial conditions capture the broken symmetry associated with a choice of direction for the propagation of activity.[18] A related model of three-state excitable elements on a small-world network has been proposed to study the propagation of infectious diseases. In this model, the introduction of nonlocal connections was found to induce a transition to a state with coherent population oscillations.[19]

Most studies of the dynamics of complex networks have assumed that the interaction between node elements is bidirectional. This is a very reasonable assumption in the context of epidemic models,[19] and it would also apply to ensembles of neurons connected via gap junctions.[12] In the absence of noise, activity that is initiated through the excitation of individual neurons in an otherwise quiescent state results in wave fronts that propagate symmetrically in both available directions; if both local and nonlocal connections are bidirectional, all wave fronts are generated as symmetric pairs that annihilate upon collision. In such a system, persistent activity can only arise from initial conditions that suitably break the symmetry between the two possible directions for the propagation of activity.

In the cortical neural ensembles that motivate this work, the coupling between neurons is predominantly not bidirectional; most connections involve chemical synapses that transmit information from the presynaptic axon to the postsynaptic dendrite. In this scenario, it is more appropriate to consider directed networks with unidirectional connections. In previous work,[20] we have investigated networks in which the local connections are bidirectional, based on the assumption that the probability for reciprocal axo-dendritic connections is quite high for neurons in close proximity, while the nonlocal connections that provide shortcuts are unidirectional. This network is a modification of the by now classical small-world network,[21] in that the added nonlocal connections are directed. We found that just a few unidirectional shortcuts suffice for sustaining persistent activity, even when activity arises from localized excitations that do not break the symmetry between the two possible directions of propagation. However, as the density of shortcuts is increased, an increasing number of network configurations support only a brief burst of population activity after which the activity dies out. When the speed associated with the propagation of activity is low, this failure of network activity was found to be delayed and to occur only after many cycles of chaotic population bursts. The simplicity of this model allowed for an analytic description of the failure transition and for detailed numerical analysis.

The properties of this simple model[20] provide important insight into the phenomena found in simulations of more elaborate models motivated by specific biological systems. In Ref. 22, the connection between the topology of a neural network and its tendency towards epileptic seizures has been studied and related to the degree of recurrent connectivity in

different parts of the hippocampus. The origin of bursting behavior was addressed in Ref. 23. The common view on bursting behavior is that it arises when fast spiking drives a slow process, typically associated with slow kinetics, that in turn can shut off the spiking activity. However, no such slow kinetics are needed if the network displays small-world connectivity.[23] Both the seizing activity described in Ref. 22 and the bursting activity described in Ref. 23 find a common interpretation in the mechanisms that underlie the failure transition found in the simple model analyzed in our earlier work.[20] A rapid spread of activity followed by persistent oscillations has also been observed in recent chemical experiments based on the Belousov-Zhabotinsky reaction; in this system, unidirectional shortcuts were implemented through the photosensitive properties of the reaction.[24,25]

Here we build on our previous results, and present a detailed characterization of the persistent states and the dependence of their properties on the density of shortcuts; we also investigate the long chaotic transients and provide an explanation for the stretched exponential that characterizes their eventual failure. In Sec. II we define the model: a network of excitable integrate-and-fire neurons coupled via excitatory pulses in a small-world topology with unidirectional shortcuts. In Sec. III we discuss the persistent states and the crossover from persistent activity to failure for the case of rapidly propagating waves. In Sec. IV we analyze the exceedingly long chaotic transients in the regime of slowly propagating waves. In the concluding Sec. V we discuss our results in light of other work on neural networks with small-world topology.[4,22,23]

II. NEURAL MODEL AND NETWORK CONNECTIVITY

We consider a one-dimensional network: a ring of N identical integrate-and-fire neurons. The state of the ith neuron is described by its membrane potential V_i, which is dynamically controlled by the spiking activity of the neurons that project onto it,

$$\tau \frac{dV_i}{dt} = -V_i + RI_{\text{ext}} + g_{\text{syn}} \sum_{j=1}^{N} w_{ij} \delta(t - t_j - \tau_D). \quad (1)$$

In order to fully specify the dynamics of integrate-and-fire neurons, this set of N equations needs to be supplemented with a condition for spiking whenever the membrane potential reaches a specified threshold, and with a reset condition for the membrane potential after the emission of a spike,

$$V(t^+) = V_{\text{res}} \quad \text{whenever } V(t^-) = V_{\text{th}}. \quad (2)$$

In Eq. (1), τ is the membrane time constant, g_{syn} is the synaptic strength measuring the change in membrane potential due to each incoming spike, $w_{ij} = 1$ or 0 indicates the presence or absence of a synaptic connection from neuron j to neuron i, t_j is the time at which neuron j fires a spike, I_{ext} is an external current, and R is the membrane resistance. The effective delay τ_D in the neuronal interaction includes both the time for the spike or action potential to propagate along the axon and the time needed for initiating the triggered action potential. When this latter time dominates over the axonal delay, the dependence of τ_D on the physical distance

between the presynaptic and the postsynaptic neurons can be neglected. Postsynaptic currents due to synaptic activation are considered instantaneous and are therefore modeled as delta functions. The process of spike emission is described in Eq. (2): whenever the membrane potential of a neuron reaches the threshold value V_{th}, a spike is emitted and the membrane potential is reset to the value V_{res}. Without loss of generality, we write these two equations in terms of dimensionless quantities by measuring V_i and g_{syn} with respect to V_{res}, setting $V_{\text{res}} = 0$ and $V_{\text{th}} = 1$, and rescaling time by τ. (In Ref. 20, we used $\tau = 10$.) Also, we replace RI_{ext} by the steady-state voltage V_∞; this is the asymptotic value that V_i would reach in the absence of synaptic input. We restrict the model to the case of excitable rather than spontaneously oscillating neurons by setting $V_\infty < V_{\text{th}} = 1$, and consider only initial conditions such that at most a few neurons are triggered to spike while the rest of the network is in a quiescent state. In the absence of noise, this initial condition implies that neurons can only fire at times that are integer multiples of τ_D. Since the dynamical evolution of the membrane potentials can be integrated exactly between subsequent spikes, the time step for the numerical computations is taken to be $\Delta t = \tau_D$.

Cortical neurons often receive not only local input from nearby neurons but also input from some distant neurons through long-range projections. We mimic this heterogeneous connectivity through an extremely simplified network architecture: each neuron is bidirectionally connected to its $2k$ nearest neighbors, i.e., $w_{ij} = 1$ for $|i - j| \leq k$, $j \neq i$, and unidirectionally connected to pN randomly chosen neurons. The parameter p thus indicates the density of additional unidirectional connections, as a fraction of the total number N of neurons.

The dynamics that arise from Eqs. (1) and (2) in the case of purely local connectivity ($p = 0$) and in the absence of noise depend only on the interplay between the strength g_{syn} of the synapses, the number $2k$ of local connections per neuron, and the delay τ_D. The dependence on g_{syn} and τ_D is most easily illustrated for first-nearest-neighbor connectivity, $k = 1$. In this case, if the presynaptic input is weak enough to satisfy $g_{\text{syn}} + V_\infty < V_{\text{th}} = 1$, this presynaptic input is insufficient to cause a spike and the activity is not propagated. In contrast, if the presynaptic input is strong enough to satisfy $g_{\text{syn}} + V_\infty > V_{\text{th}} = 1$, it results in a propagating wave of speed $1/\tau_D$. After spike emission, the voltage is reset to $V_{\text{res}} = 0$, and the neuron is only ready to fire again after it has recovered to the extent that an input of magnitude g_{syn} is sufficient to trigger another spike. This recovery time is given by

$$T_R = \ln\left(\frac{V_\infty}{V_\infty + g_{\text{syn}} - 1}\right). \quad (3)$$

Note that T_R is not the intrinsic refractory period of the neuron, since this integrate-and-fire neuron can fire at arbitrarily large frequencies for sufficiently strong input g_{syn}.

Due to the bidirectionality of the local connections, the firing of each neuron not only triggers a spike in the neuron ahead of it in the direction of wave propagation, but also gives an input to the neuron behind it, which thus receives an input at a time $2\tau_D$ after its own firing. If $T_R > 2\tau_D$, this input is not sufficient to trigger a new spike, and the activity
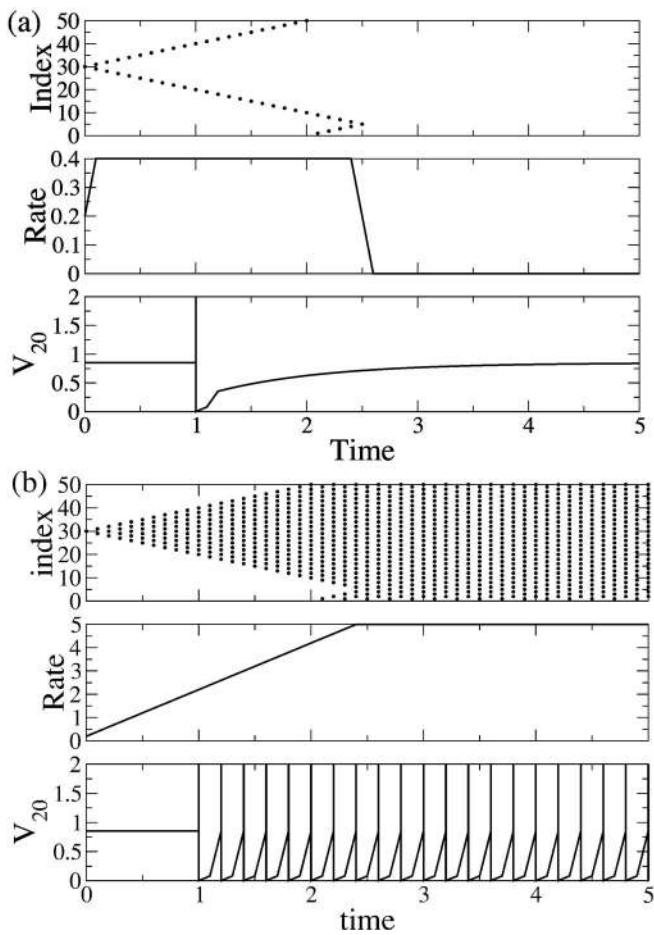
FIG. 1. Dynamics in a regular network of 50 neurons with first-nearest-neighbor coupling. Parameters are $V_\infty=0.85$ and $\tau_D=0.1$. (a) A weak synaptic coupling $g_{syn}=0.2$ results in a propagating wave. The two wave fronts meet and annihilate at time 2.5. Top: Raster plot showing spike times for each neuron. Middle: Population firing rate. Bottom: Membrane potential of neuron number 20. (b) A strong synaptic coupling $g_{syn}=1.0$ results in a wave that entrains all neurons in its path. The final state of the network consists of two synchronous groups of neurons firing out of phase with each other. Top, middle, and bottom panels as in (a). Note that the maximum possible firing rate is $1/\tau_D=10$.

propagates away from the site of initiation as a wave to which each neuron contributes exactly one spike, see Fig. 1(a). However, if $T_R \leq 2\tau_D$, the wave front entrains all the neurons in its wake, eventually leading to synchronized activity of the whole network. In the absence of autapses, i.e., for $w_{ii}=0$, the network breaks up into two synchronous groups of neurons that fire out of phase with one another, see Fig. 1(b). In the general case, a neuron receives inputs from k neighbors as the wave approaches; the input coming from a neuron at a distance n is discounted by the factor $e^{-(n-1)\tau_D}$. A propagating wave can thus be sustained if $g_{syn}e^{\tau_D}\sum_{n=1}^{k}e^{-n\tau_D} + V_\infty > V_{th}=1$.

For simplicity, we will focus on the case of first-nearest-neighbor coupling in the regime in which waves of excitation propagate but do not entrain activity in their wake. This choice implies $k=1$ and constrains the allowable values of V_∞ and g_{syn}; we use $V_\infty=0.85$ and $g_{syn}=0.2$ unless otherwise noted. In this regime, the collision of two waves leads to their mutual annihilation and, after having fired in a propa-
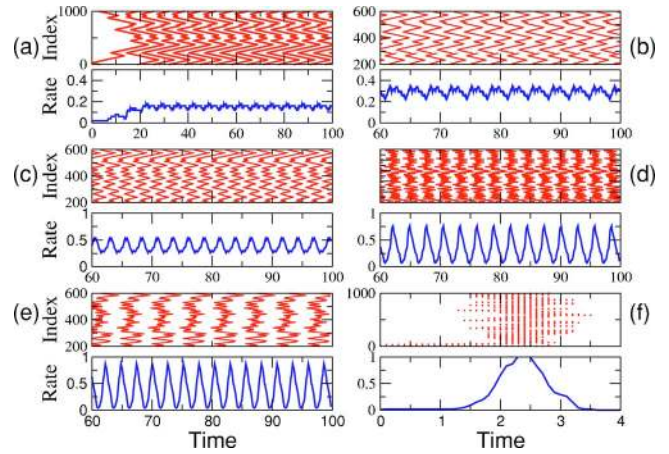

FIG. 2. Examples of network dynamics for $\tau_D=0.1$ and for different values of the density p of shortcuts. The values $p=0.01$, 0.05, 0.10, 0.15, 0.20, 0.25 correspond to panels (a)–(f), respectively. Each panel shows both spiking activity and population firing rate. Note different temporal scale on panel (f).

gating wave, a neuron can be triggered to fire by a single input of size g_{syn} after a time

$$T_R^{(1)} = \ln\left(\frac{V_\infty - g_{syn}e^{2\tau_D}}{V_\infty + g_{syn} - 1}\right). \tag{4}$$

This calculation includes the input received from the neuron ahead in the wave, at a time $2\tau_D$ after spiking.

The model specified by these assumptions together with Eqs. (1) and (2) is used here as a simple model for the generation and propagation of waves of activity in cortical tissue. As discussed in the following sections, the incorporation of random connections qualitatively alters the dynamics of the network and sustains a rich variety of spatiotemporal patterns.

III. THE ORDERED REGIME: ATTRACTORS AND FAILURE

The dynamics of the model depend on several parameters. Once the input current $V_\infty=0.85$ and the synaptic strength $g_{syn}=0.2$ have been fixed, the dynamics arising from Eqs. (1) and (2) are determined by the remaining two parameters: the fraction p of randomly placed shortcuts and the delay τ_D associated with the neural interaction.

The dynamics for $p \neq 0$ differ qualitatively from those for $p=0$. The presence of shortcuts allows the waves of excitation to be reinjected into portions of the network which have been previously excited. This process of reinjection may lead to persistent network activity, as shown in Fig. 2 for several values of the density p. As the waves spread outward from the initial site of activation, they encounter shortcut connections that inject activity elsewhere in the network. As for the $p=0$ case shown in Fig. 1(a), wave fronts that meet annihilate. After some time, the activity settles into a stable pattern in which the rates of wave generation and annihilation are balanced.

When averaged over time and across network configurations at a fixed value of the shortcut density p, the firing rate of these persistent states increases rapidly with p, and saturates around $p \sim 0.1$ [Fig. 3(a)]. This saturation is a conse-
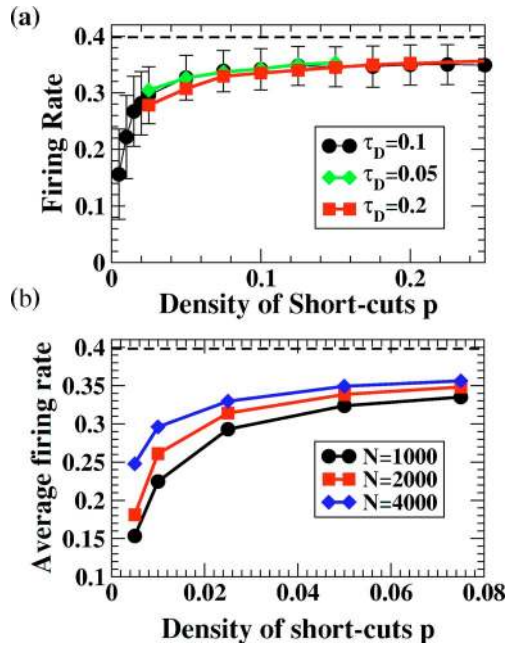
FIG. 3. Population firing rate, averaged over 1000 configurations for p < 0.2 and over 5000 configurations for p > 0.2. The maximal value of the firing rate, $1/T_R^{(1)}$, is indicated in both panels by a dashed horizontal line. (a) Firing rate as a function of the density of shortcuts p for different values of the delay τ_D for a network of size N=1000. Bars give the standard deviation of the firing rate across configurations for τ_D=0.1. (b) Firing rate as a function of the density of shortcuts p for different values of the system size N with τ_D=0.1.

FIG. 4. Dynamics in a network of N=2000 neurons with a density of short-cuts p=0.1. The same network configuration and the same initial activation is used in both panels. The only difference is in the delay time: τ_D=0.05 for the top panel and τ_D=0.10 for the bottom panel. Note that more waves travel through the system when the delay is longer. Circles indicate a new wave spawned for τ_D=0.10 but not for τ_D=0.05.

quence of the neuron's finite recovery period $T_R^{(1)}$ of Eq. (4); its inverse is the maximal firing rate, indicated as a dashed line in both panels of Fig. 3. As shown in this figure, the firing rate comes very close to this maximal value. As the system size is increased, saturation is reached at smaller values of p [Fig. 3(b)].

The firing rate is essentially the inverse of the time between successive waves passing through a specific node in the ring of neurons. Thus, one may expect that decreasing the wave speed by increasing τ_D would reduce the mean firing rate. However, this is not the case. As shown in Fig. 3(a), the firing rate is quite insensitive to wave speed. The reason for this effect is illustrated in Fig. 4. The two panels show the spatiotemporal pattern of activity for the same network configuration and the same initial activation, but for different delay times τ_D. At the larger delay time, additional waves are excited through the shortcuts. An example of the appearance of such a new wave can be seen at t=380τ_D, a time indicated by a circle in both panels of Fig. 4. At this time, a new wave is spawned by neuron 1071 for τ_D=0.1 but not for τ_D =0.05. These additional waves increase the firing rate in a manner that may even overcompensate for the reduced wave speed, as demonstrated by the higher density of waves for τ_D=0.1 at times beyond t=800τ_D.

Only a few shortcuts are present for low p, and many pathways leading to persistent activity consist only of large closed loops, which result in low firing rates. As p increases, the typical loop size decreases, and network configurations with only large loops and correspondingly low firing rates become increasingly unlikely. This effect is illustrated in the
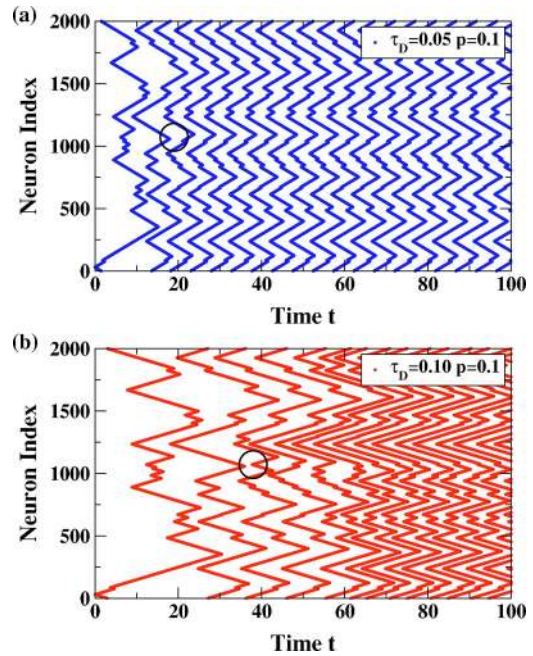
probability density function for the firing rate, shown in Fig. 5. As p increases, the distribution is shifted towards larger firing rates and it narrows substantially, reflecting a saturation close to the maximal firing rate set by the recovery time $T_R^{(1)}$. This maximal firing rate is marked by a dashed vertical line in Fig. 5.

For small values of p, the state of persistent activity settles into a periodic pattern whose oscillations increase in amplitude with increasing p. This effect is clearly illustrated in panels (a)–(e) of Fig. 2. A quantitative description of this effect is shown in Fig. 6, in which the standard deviation of the firing rate, averaged over a large number of network configurations that exhibit persistent activity for a given value of p, is used to characterize the amplitude of the oscillations. An additional feature of this regime, also apparent from panels (a)–(e) of Fig. 2, is that the time for these oscil-
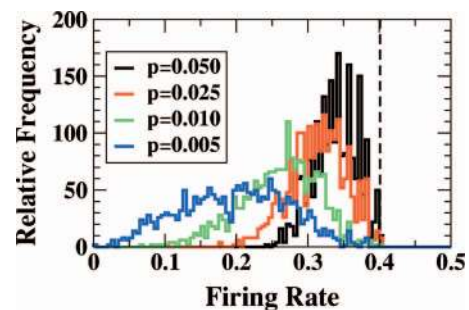

FIG. 5. (Color) Probability distribution of firing rates for several values of p, for networks of N=2000 neurons with τ_D=0.1. As p increases, the distributions narrow in width as they shift towards the maximal firing rate $1/T_R^{(1)}$, indicated by a dashed vertical line.
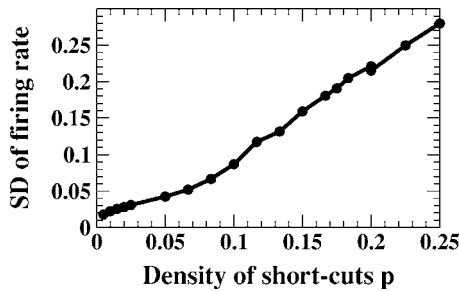
FIG. 6. Amplitude of oscillatory activity, measured through the standard deviation of the firing rate, as an increasing function of the shortcut density p. The standard deviation is obtained as an average over 5000 network configurations that exhibit persistent activity for $\tau_D=0.1$; $N=1000$.

lations to become established after the excitation of a single neuron decreases with increasing p. This property is in agreement with the results of Refs. 15 and 16, which found that, when excited through a small cluster of driven oscillatory neurons, small-world networks of either Hodgkin-Huxley or FitzHugh-Nagumo neurons are entrained much more quickly than regular networks (with no shortcuts) of the same types of neurons.

For larger values of p, the activity patterns can be quite complicated. In this regime, as in the small p regime, all neurons get excited during an oscillation cycle of network activity. However, in this regime not all neurons and not all connections between them are necessary for the persistence of activity, as illustrated in Fig. 7. The full raster plot of network activity shown in the top panel depicts the spikes from all neurons (black dots). This plot also identifies the spikes from those neurons that are essential for sustainability (red dots). These neurons provide a pathway for recurrence,[26] and they are found as follows. At an arbitrary time once the steady state has been reached, all neurons that fire at that time step are labeled. A backwards search is then performed for preceding ancestors of these labeled neurons,
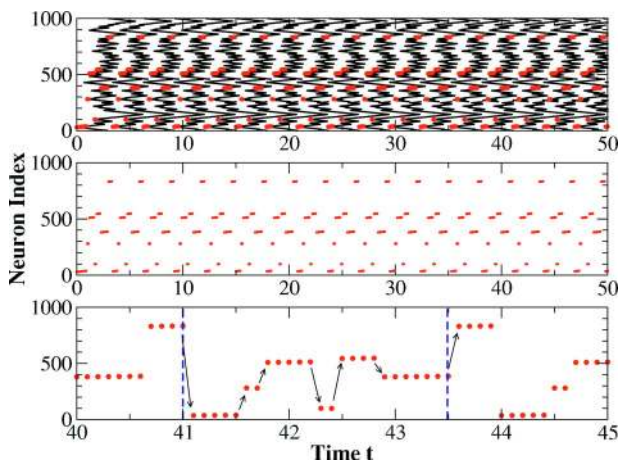

FIG. 7. Spiking activity in a network of $N=1000$ neurons with a density of shortcuts $p=0.1$ is shown in the top panel for $\tau_D=0.1$. The activity of neurons that are part of the backbone pathway is indicated in red. The middle panel shows the spiking activity within the backbone. The enlargement shown in the bottom panel displays periodic activity, with a period longer than the recovery time $T_R^{(1)}=2.494$ (delimited by dashed vertical lines). See text for details.

i.e., those presynaptic neurons that fired one delay τ_D ago and triggered the activity of the labeled neurons. These one-step ancestors are labeled in turn, and the process is iterated backwards in time until $t=0$ is reached. The pattern of labeled neurons quickly converges to a small subset, as shown in the middle panel of Fig. 7. Only the neurons in this reduced subset contribute to the persistence of the pattern; the remaining neurons could be cut out of the network without destroying the persistent activity. The pattern of activity along this backbone, shown in the bottom panel of Fig. 7, is periodic in time. The period is only slightly longer than the recovery time $T_R^{(1)}$, indicated by the dashed vertical lines (note the expanded time scale in this panel). Such backbone pathways have also been identified in an experimental study of the excitable Belousov-Zhabotinsky reaction;[25] in this photosensitive system, unidirectional shortcuts were implemented through local optical excitation.

A striking feature of this regime is that a network configuration capable of sustaining persistent activity displays an extraordinarily large number of different attractors. To assess the number of coexisting attractors we focus on the regime of low shortcut density, where all solutions are periodic. Each solution can be characterized by its period, its mean firing rate, and the standard deviation of the firing rate. A labeling of attractors based only on these three measures is likely to underestimate the total number of attractors; however, it suffices to find a large number of them, as shown in Fig. 8. This figure illustrates attractor multiplicity for a randomly chosen network configuration of $N=1000$ neurons with $p=0.05$ and $\tau_D=0.1$. Each of the 471 distinct stable patterns of persistent activity identified in Fig. 8(a) arises from an initial condition in which only one neuron is activated. As shown in Fig. 8(b), which shows the number of initial conditions that lead to a solution with a given period, many of these stable patterns have the same period (note the logarithmic scale). In contrast, the standard deviation of the firing rate [Fig. 8(c)] shows great variability, reflecting different temporal evolutions of the firing rate within a period.

A large number of attractors is typical for these networks. Figure 9 shows the number of attractors as a function of network size; black circles represent data points obtained as averages over 20 different network configurations of a fixed size. Most attractors have quite small basins of attractions: within the restricted set of initial conditions in which only a single neuron is excited, most attractors can be reached from only one such initial condition (red squares). The overall number of attractors increases roughly linearly with system size for large values of N, as does the number of attractors with different periods (blue diamonds), but the latter is about an order of magnitude smaller than the total number of attractors. The identification of attractors requires that networks achieve a steady state of persistent activity, but since the duration of the transients grows with system size, the computation time grows faster than N^2. This precludes us from simulating significantly larger system sizes than shown in Fig. 9; such data would be necessary to reliably estimate scaling relations between network size and number of attractors. (For $N=4000$ the computation takes over 2 weeks on a desktop PC.) Preliminary computations with more general
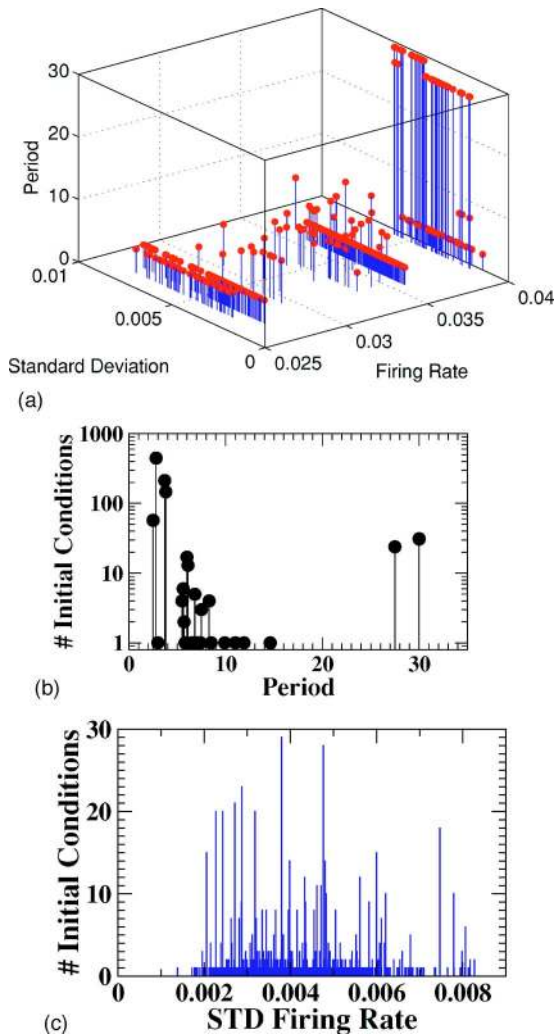
(a)

(b)

(c)

FIG. 8. Large number of attractors for a specific network configuration, with $N=1000$ neurons, $p=0.05$, and $\tau_D=0.1$. (a) The mean firing rate, its standard deviation, and the period of all 471 distinct attractors evoked by initial conditions in which only one neuron is activated. (b) Number of initial conditions that lead to an attractor with a given period; many attractors, although distinct, have the same period (note the logarithmic vertical scale). (c) Number of initial conditions that lead to an attractor with a given standard deviation of the firing rate.
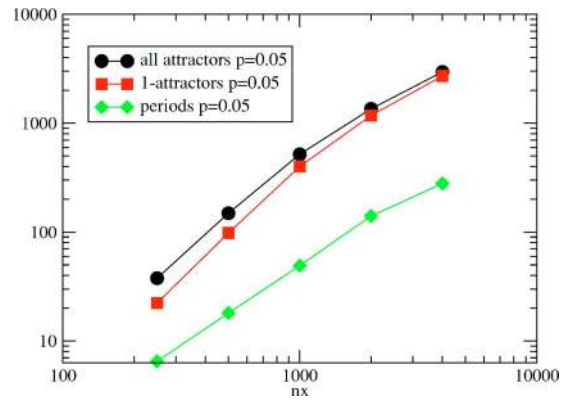

FIG. 9. Number of attractors: all attractors (circles), attractors with basin of attraction of size one (squares), and attractors with different periods (diamonds; note different vertical scale on the right) as a function of network size N.

initial conditions reveal many more attractors than those shown in Fig. 8. Thus, while the restricted initial conditions that give rise to Fig. 9 suggest an almost linear increase in the number of attractors with system size, the full number of attractors may grow substantially faster.

At this point, the origin for this exceedingly large number of different attractors is not clear. In all-to-all coupled oscillator systems, factorially large numbers of attractors are due to the permutation symmetry associated with the global coupling.[27] The small-world networks investigated here do not possess such symmetry. In this case, the large number of attractors found for a specific network configuration is likely to be due to a combination of geometrical and dynamical causes: the coexistence of many backbones that can independently support sustained activity, and the variability in membrane potentials that can result in distinct but not too different firing patterns. While we find transitions between different attractors in the presence of finite-amplitude

noise,[26] we have not investigated whether this system exhibits the extreme noise sensitivity found in the case of all-to-all coupled oscillators, where it is due to the crowding of large numbers of attractors.[27]

As the density p of shortcuts increases, the distance traveled by the waves before encountering the entrance to a shortcut decreases. Consequently, waves of excitation spread throughout the network more rapidly, as shown in the progression of spatiotemporal patterns of Figs. 2(a)–2(e) for increasing values of p. If p becomes too large, activity spreads too fast and it quickly dies away, as shown in Fig. 2(f). Since the shortcuts are randomly placed, different network configurations will exhibit different dynamics. Thus, while the overall likelihood of persistent activity decreases with increasing p, the actual network dynamics depend on the particular network configuration.

The mechanism that leads to the extinction of network activity is easily elucidated.[20] Once a neuron has emitted a spike, its membrane potential is reset to a fixed value, chosen here to be $V_{res}=0$. While the membrane potential of the neuron recovers towards its resting value V_∞, activity spreads through the network, eventually finding its way back. Once this occurs, the neuron receives synaptic input equal to g_{syn}. This input will be sufficient to trigger a spike only if this neuron has recovered sufficiently. It is therefore clear that as the number of shortcuts is increased and activity spreads more rapidly, the network is less likely to sustain persistent activity. Whether or not this mechanism of premature return will lead to the extinction of activity in a given network depends on its particular configuration. Many different network configurations for a given value of p have been simulated in order to measure the fraction that fail to sustain persistent activity. This probability of failure is shown as a function of p for different values of the system size N in the left inset of Fig. 10. In agreement with our intuitive argument, the probability that a network drawn at random fails to sustain activity increases with increasing p. In fact, there is a sharp crossover from a low p regime characterized by the ability to sustain persistent activity into a large p regime in which activity will always fail. The transition between these
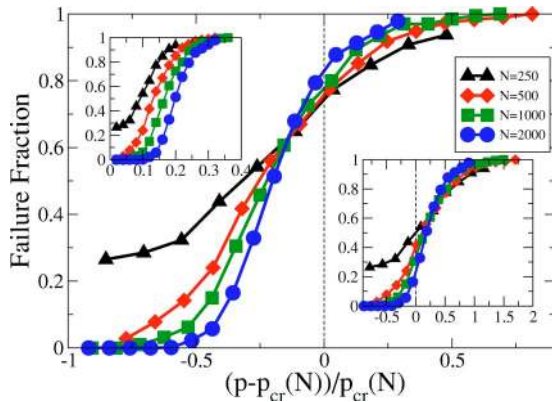
FIG. 10. Probability of failure as a function of shortcut density for different values of the system size N. The density of shortcuts is scaled by p_{cr} from Eq. (8). Left inset: unscaled data. Right inset: data scaled by p_{cr} from Eq. (7).

two regimes occurs at higher values of p as the system size N increases.

As discussed in Ref. 20 this transition can be captured in a mean-field approximation in which the return time is assumed to be identical for all neurons. In this approximation, the maximum return time T_A needed for the activity to traverse the entire network can be expressed as a function of only p and N. Setting this time to be equal to the recovery time yields an upper bound for the critical density of shortcuts at which a transition from persistent activity to failure occurs,

$$T_A(p_{cr}) = T_R^{(1)}. \qquad (5)$$

An approximate form for T_A is easily derived. Assume an initial condition in which a single neuron fires at time $t=0$. Given a density p of shortcuts, the entrance to a shortcut will be typically encountered after $1/p$ neurons have fired, which occurs after a time $\tau_D/2p$; the factor of 2 is due to the two wave fronts that emerge from the initially activated neuron and propagate in opposite directions. Due to the activity injected through the shortcut, four wave fronts are now propagating through the system; $2/p$ neurons will fire during the subsequent time interval of duration $\tau_D/2p$, at the end of which a new shortcut entrance will typically be found and two more wave fronts will be generated. The process is iterated, with $2^{k-1}/p$ neurons firing during the kth cycle. It takes n cycles to ensure that all neurons have fired, with n such that

$$\sum_{k=0}^{n-1} 2^k = pN, \qquad (6)$$

which leads to a total time

$$T_A(p) \equiv n\frac{\tau_D}{2p} = \tau_D \frac{\ln(1+pN)}{2p\ln 2}. \qquad (7)$$

This is a purely geometric result for the time it takes for activity to traverse the entire extent of the network; this time is related trivially to the largest distance in the network. The geometric mean-field properties of small-world networks have been analyzed by Newmann, Moore, and Watts,[28] who used a continuum limit to calculate the fraction of a small-world network that is covered by starting at a single point and extending outwards a distance r in both directions; this is equivalent to following the spread of waves of activation during a time $r\tau_D$ for the neural network considered here. The calculation in Ref. 28 takes into account two effects that were omitted from their earlier calculation[29] and from our derivation, Eq. (7). First, as the network is covered, a shortcut might lead to a part of the network that has already been traced over; this contribution should not be counted. In the neural scenario, this is equivalent to an attempt at injecting activity into a neuron that has already fired but not yet recovered to the point where it can fire again. Such a shortcut does not contribute to sustained activity. Second, when two covering fronts meet, they stop and no longer contribute. In the neural network, this corresponds to activity wave fronts that meet and annihilate. The incorporation of these two additional mechanisms leads to a two-component model that correctly describes both the covered fraction of the network and the number of fronts.[28] The result, when applied to the neural network of Eq. (1), yields

$$\sqrt{\left(1+\frac{4}{pN}\right)}\tanh\left[\sqrt{\left(1+\frac{4}{pN}\right)}\frac{pT_A(p)}{2\tau_D}\right] = 1. \qquad (8)$$

The identification $T_A(p_{cr}) = T_R^{(1)}$ then yields the mean-field estimate $p_{cr}^{(MFT)}$ for the density at the failure transition.

The main panel in Fig. 10 shows the failure rates as a function of shortcut density, with p rescaled by the critical density $p_{cr}^{(MFT)}$ that follows from Eq. (8). All rescaled curves intersect at a common value of p, which defines the transition point at p_{cr}. The mean-field theory yields an upper bound, $p_{cr} < p_{cr}^{(MFT)}$. It is interesting to note that a similar rescaling using the estimate for p_{cr} that results from Eq. (7) also produces a family of rescaled curves that intersect at a common value of p, as shown in the right inset of Fig. 10. While $p_{cr}^{(MFT)}$ overestimates the true p_{cr} at which the rescaled curves cross, the estimate based on Eq. (7) underestimates the true p_{cr}. In an earlier report[20] we incorrectly used T_R instead of $T_R^{(1)}$ as the recovery time in Eq. (8). The corresponding rescaled curves intersected at zero. We cannot tell whether this agreement with the true p_{cr} was just a coincidence or the consequence of a subtle cancellation between factors leading to overestimation or underestimation. In any case, it is worth emphasizing that all three methods, although only approximate, succeed in rescaling the curves in the left inset of Fig. 10 so as to obtain a unique crossing. This is due to the fact that these phenomenological approaches result in estimations for the critical density p_{cr} with very similar dependencies on system size N; these estimates seem to provide a good approximation to the true dependence on system size.

The failure transition occurs at the value of p for which the geometrical quantity $T_A(p)$ equals $T_R^{(1)}$. This recovery time is thus the central quantity that determines persistence or failure for a given shortcut configuration. It is worth pointing out that this recovery time is not the same as the absolute refractory period T_r. During the refractory period, neurons are inhibited from receiving synaptic input, while the mem-
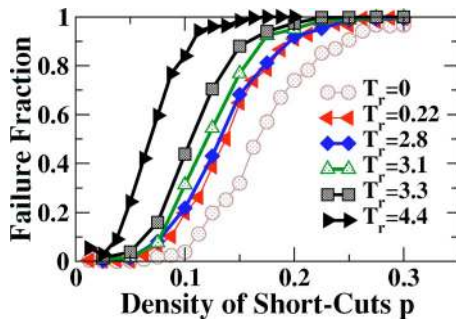
FIG. 11. Dependence of failure probability on shortcut density for different values of the refractory period T_r. Shown are averages over 500 network configurations, with $N=1000$, $\tau_D=0.1$, and $T_R=2.83$.

FIG. 12. Dependence of failure probability on shortcut density for different values of the delay $\tau_D=0.06, 0.08, 0.10, 0.12, 0.14, 0.16$. Shown are averages over 2000 network configurations for $N=1000$. Note the nonmonotonic character of the curves for large enough delay. In the large p regime, failure probability is estimated within a finite time $T^*=100$.

brane potential still relaxes towards its resting value. While T_r is an intrinsic property of individual neurons, $T_R^{(1)}$ depends strongly on the strength of the synaptic coupling. The curves in Fig. 11 show the fraction of network configurations that fail to sustain persistent activity for different values of the refractory period T_r. For $T_r < 2\tau_D$, the input that neurons in the wake of a propagating wave receive at a time $2\tau_D$ after their own firing is unaffected, and the relevant recovery time is $T_R^{(1)}$ [cf. Eq. (4)]. The corresponding curves are as those for $T_r=0$, shown in Fig. 10. The input received at $2\tau_D$ after firing is suppressed for $T_r > 2\tau_D$. In this low p regime, the persistent states do not depend on neurons receiving additional inputs before the one that triggers a spike, and the failure transition is controlled by T_R [cf. Eq. (3)]. The failure transition is thus independent of T_r as long as $2\tau_D < T_r < T_R$. It is only for $T_r > T_R$ that the refractory period suppresses relevant input to the neurons and affects the failure transition. This is illustrated in Fig. 11, where the refractory period T_r is seen to have little effect on persistent activity when it takes values between $2\tau_D=0.20$ and $T_R=2.83$.

In our analysis of persistent activity and transition to failure, we have also considered the possibility of an upper bound in the length of allowed shortcuts. We summarize our results without showing the corresponding numerical data: the behavior of the system is qualitatively unchanged as long as this upper bound exceeds a threshold value; below threshold, the network is essentially only locally coupled and the failure probability rapidly approaches one. Other modifications to the distribution of shortcut lengths might include allowing for a nonuniform distribution. It is unclear how a nonuniform distribution of shortcut lengths would affect the dynamics. However, it has been shown that one-dimensional networks maintain a small-world structure if the distribution of shortcut lengths is power-law with power <2.[30] We assume that the regime we have discussed in this section will also be present in such a scenario.

IV. THE DISORDERED REGIME: CHAOTIC TRANSIENTS IN SLOW WAVES

For small values of the density of shortcuts p, a small-world network of integrate-and-fire neurons is quite likely to sustain persistent activity. As discussed in the preceding section, the spatiotemporal pattern of network activity in this regime is high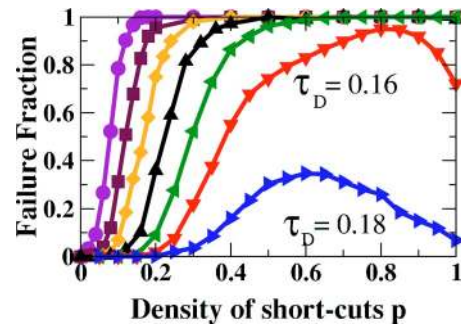ly regular and most often periodic, despite the complex, heterogeneous topology of the network itself. As p increases, a transition takes place: more and more network configurations exhibit activity which peaks and then shuts down. Interestingly, for large enough τ_D there is an additional change in network dynamics as p increases. As shown in Fig. 12 for $\tau_D=0.16$ and 0.18, the failure probability initially increases with p, as for low values of τ_D, but it then turns downwards again as p increases further. To understand this reentrant phenomenon, we first analyze the spatiotemporal dynamics characteristic of this seemingly persistent activity at large p.

Examples of network dynamics for slow waves ($\tau_D=0.16$) are shown in Fig. 13. For values of p below or near the theoretical transition to failure [Figs. 13(a)–13(c)], the activity is similar to that shown for fast waves ($\tau_D=0.10$) in Fig. 2. However, in the reentrant regime [Fig. 13(d)], the activity is chaotic [cf. Fig. 15(a) below] and the population firing rate exhibits irregular peaks that reflect near-synchronous activity involving a large fraction of the network.

A detailed, quantitative analysis for $\tau_D=0.18$ shows that the change in behavior occurs already before the maximum of the failure probability curve. For fast waves, corresponding to small τ_D, the amplitude of the oscillations in the population activity, as measured by the standard deviation of the firing rate, was found to increase monotonically with p (cf. Fig. 6). However, for slow waves, corresponding to large τ_D, this amplitude of oscillations is nonmonotonic and decreases over the range $0.4 \leq p \leq 0.6$ (Fig. 14). It is instructive to compute the spectral entropy of the population firing rate,

$$\mathcal{S} = -\sum_{\omega} P(\omega) \ln P(\omega), \qquad (9)$$

which measures the number of significant peaks in the power spectrum $P(\omega)$. The spectral entropy exhibits a significant increase over the same range $0.4 \leq p \leq 0.6$ [Fig. 15(b)], indicating an increase in the complexity of the dynamics. The variability of the spectral entropy across network configurations with the same density of shortcuts exhibits a broad maximum in the same range of p values, and reaches very small nonzero values (about 0.06) in the strongly chaotic regime. The detailed evolution towards chaotic dynamics,
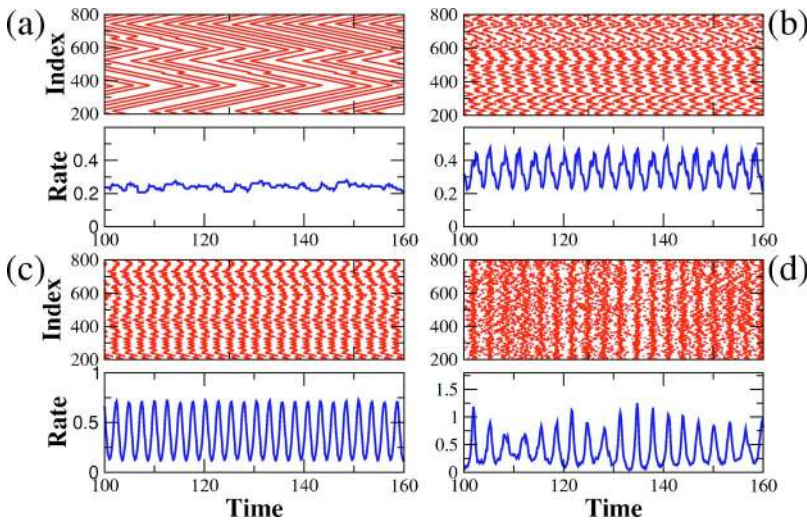
FIG. 13. Examples of network dynamics for $\tau_D=0.16$ and for different values of the density p of shortcuts. The values $p=0.01$, 0.2, 0.4, and 1.0 correspond to panels (a)–(d), respectively. Each panel shows both spiking activity and population firing rate. The reentrant activity in (d) is noisy and exhibits synchronized population bursts.

which depends on the specific configuration of shortcuts, is not investigated in further detail in this work.

What underlies the emergence of reentrant activity with increasing p? The answer lies in the interplay between network topology and the delay τ_D. The mean-field model that provides a description of the failure transition, Eq. (5), is based on the assumption that the maximal firing frequency of each neuron is limited by the recovery time $T_R^{(1)}$. However, a small-world network constructed by adding shortcuts allows for neurons to receive more than one incoming shortcut. The probability of such configurations is small at low p; in this regime, the fraction of neurons with two incoming shortcuts is given by $s_2 \sim p^2/2$,[26] which is indeed negligibly small for $p < 1$. (We assume that the number of neurons with more than two shortcut inputs can be ignored and we compute the most likely rather than the expected value of the number of neurons with two shortcut inputs.) Mean-field results that ignore such configurations describe the failure transition provided it occurs at $p_{cr}<1$. However, as p approaches 1, i.e., for $|p-1| \ll 1$, the fraction of neurons with two incoming shortcuts becomes $s_2 = 1 - \sqrt{2}/2 + (p-1)/2$, which is nearly 0.3 for $p=1$. In this regime, a significant fraction of neurons is likely to receive several synaptic inputs during one cycle of network activity. Such neurons would not be constrained by the recovery time $T_R^{(1)}$, but would rather be primed to fire

earlier, potentially allowing the activity to persist where it otherwise would fail. A neuron that has received n inputs at times t_r, $r=1,\ldots,n$, since its last firing has a recovery time $T_R^{(n)}$ given by

$$T_R^{(n)}(t_1,\ldots,t_n) = \ln\left(\frac{V_\infty - g_{\text{syn}}\Sigma_{r=1}^n e^{t_r}}{V_\infty + g_{\text{syn}} - 1}\right). \tag{10}$$

Note that Eq. (10) reduces to Eq. (3) for $n=0$ and to Eq. (4) for $n=1$ (with $t_1=2\tau_D$). In general, $T_R^{(n)}$ depends on the specific firing times of the n neurons that provide inputs through shortcuts; these times depend in turn on the details of activity propagation in each specific network configuration. However, since integrate-and-fire neurons become increasingly sensitive to their input as time passes after their firing, the value of $T_R^{(n)}$ in Eq. (10) is bounded below by $T_{R,\text{min}}^{(n)}$, which

FIG. 15. Temporal complexity of activity patterns. (a) Representative power spectra for $p=0.2$ ($\mathcal{S}=4.6$) and $p=1.0$ ($\mathcal{S}=7.4$). (b) Spectral entropy $\mathcal{S}$ (mean and standard deviation across 200 configurations) for $\tau_D=0.18$ and $N=1000$, over 15 000 time steps.
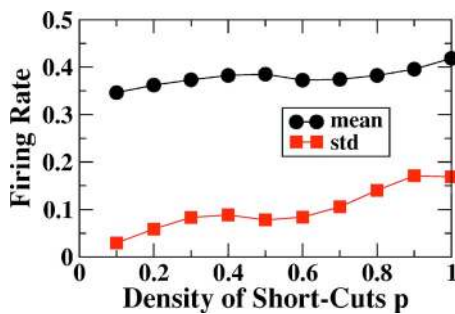

FIG. 14. Average population activity (mean firing rate) and amplitude of its oscillations (standard deviation of the firing rate) as a function of the short-cut density p. Note the nonmonotonic behavior of these curves. Data obtained as an average over those configurations (out of 200) for which the activity persists for at least 15 000 steps for $\tau_D=0.18$ and $N=1000$.
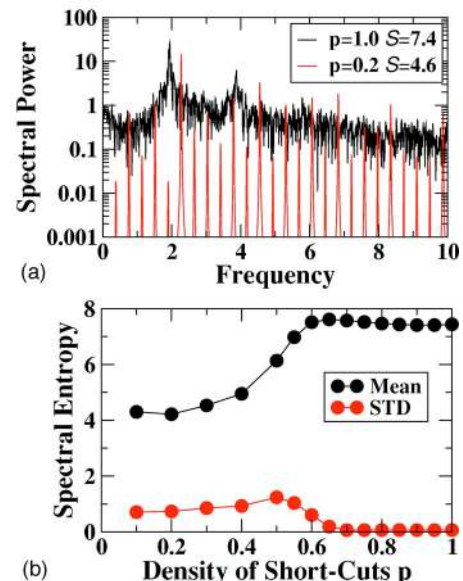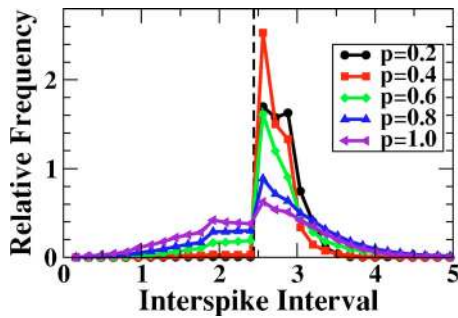
FIG. 16. Distribution of interspike intervals (ISI) for τ_D=0.16 and different values of the shortcut density p. For low p, all allowed ISI are above the recovery period $T_R^{(1)}$ (dashed vertical line). For $p \gtrsim 0.6$, multiple inputs result in ISI below $T_R^{(1)}$.
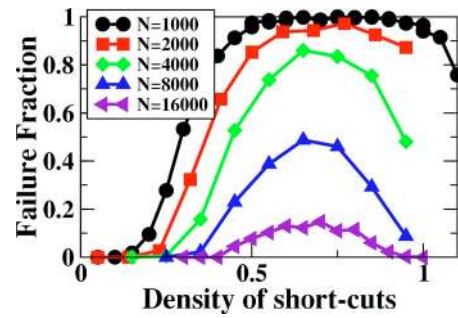

FIG. 17. Failure probability as a function of shortcut density p for τ_D =0.14 and different values of system size N. Shown are averages over 400 configurations; failure probability is estimated within a finite time T^*=28. Note the prevalence of prolonged activity in larger systems.

occurs when all n inputs coincide at $T_{R,\mathrm{min}}^{(n)}$ itself. This lower bound is given by

$$T_{R,\mathrm{min}}^{(n)}(t_i = T_{R,\mathrm{min}}^{(n)}) = \ln\left(\frac{V_\infty}{V_\infty + ng_{\mathrm{syn}} - 1}\right). \quad (11)$$

For small enough p, the interspike intervals (ISI) of almost all neurons in almost all network configurations are bounded below by $T_R^{(1)}$; this property has allowed us to calculate the time for activity to spread throughout the whole network using a purely geometric approach. For higher values of p, there may be a subset of neurons with shorter allowable ISI. However, many neurons will still receive only one input per cycle, and the frequency of their spiking activity should reflect this fact. The distribution function for the ISI shown in Fig. 16 supports this argument. Fast spiking activity with ISI$<T_R^{(1)}$ occurs appreciably only for $p \gtrsim 0.6$, and becomes both more common and faster with increasing shortcut density. Our earlier analysis[20] showed that the spikes with ISI$>T_R^{(1)}$ occur in population bursts, with no such spikes in between bursts. In the absence of other spikes, the activity would die out during these intervals between bursts. However, the fast spiking neurons that receive multiple inputs via multiple incoming shortcuts are primed to carry over the activity during the intervals between bursts; their fast spiking sustains activity during the time needed for the slow spiking neurons to recover.

Long delays τ_D contribute in several ways to bridging the periods of low activity between bursts. For larger τ_D, the failure transition is shifted towards larger shortcut densities, thus significantly enhancing the number of neurons that receive multiple shortcut inputs. At the same time, in order to bridge the time between the return time T_A and the recovery time $T_R^{(1)}$ of the slow spiking neurons, fewer fast spiking neurons are needed if the delay is longer. Moreover, the exponential recovery of integrate-and-fire neurons towards their resting potential implies that later inputs have a stronger impact on the recovery period than earlier ones [cf. Eqs. (10) and (11)]. With increasing τ_D, all inputs are shifted to later times relative to the most recent spike of the postsynaptic neuron; this shift significantly reduces the recovery time. However, an increased delay τ_D is not necessary to establish a regime of prolonged activity. As shown in Fig. 17 for fixed delay τ_D=1.4, increasing the system size from N=1000 to

N=16 000 shifts the failure transition to sufficiently large p that the number of neurons with multiple inputs is sufficient to bridge the gap between bursts of slow spiking activity, even for this shorter delay.

For low values of p, the spatiotemporal dynamics are most often periodic. In those cases the dynamics can truly be called persistent. For $p \gtrsim 1$, the chaotic nature of the dynamics precludes such a clear assessment; in fact, failure is possible even after very long times. In this regime, prolonged activity relies on bridging the quiescent period between bursts of slow spiking activity through fast spiking neurons that receive multiple shortcut inputs; it is necessary that these shortcuts are actually activated at suitable times during the quiescent part of the cycle. Thus, while in one cycle the activity during the burst may have excited such a pathway, the different activity pattern in the next burst may fail to do so; the activity could then die out. Indeed, we find that the prolonged activity characteristic of large shortcut densities eventually fails for essentially all configurations. Examples of such long-lived transients are shown in Fig. 18, where population firing rates are shown for a fixed network configuration of shortcuts at increasing values of τ_D.

The value of τ_D has a strong influence on the duration of the transient. As shown in Fig. 18, the overall trend is for the lifetime of the transient activity to increase with increasing delay (note the change in temporal scale from panel to panel). However, the actual dependence on τ_D is more subtle,
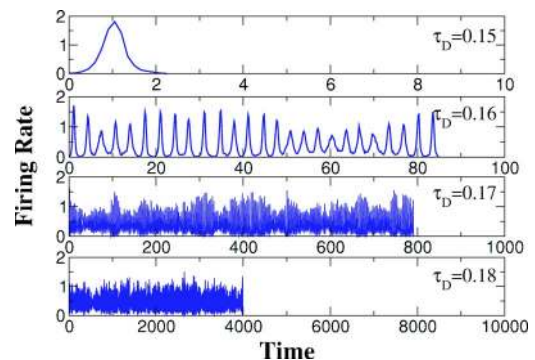

FIG. 18. Population firing rate for four values of the delay τ_D, with p=1. The same network configuration and the same initial activation is used in all four panels. Failure tends to occur later as τ_D increases; note the change in temporal scale from panel to panel.
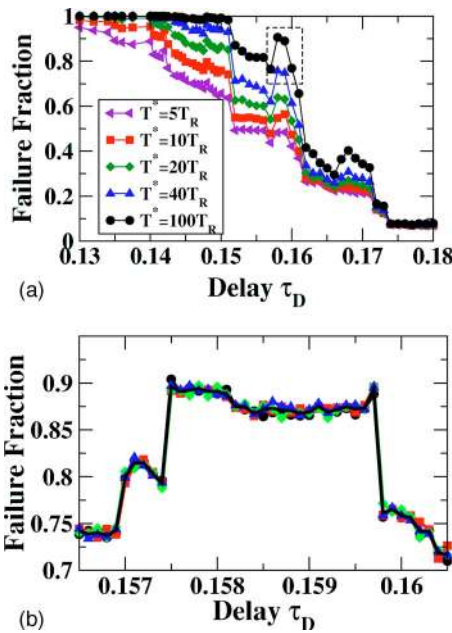
FIG. 19. (a) Failure rates for different values of the final time $T^*=5T_R$, $10T_R$, $20T_R$, $40T_R$, and $100T_R$. Shown are averages over 2000 configurations with $N=1000$ and $p=1$. (b) Enlarged view of the data within the box in (a) for $T^*=100T_R$ illustrates the fine structure in the failure rate. Shown are four runs based on averages over 8000 configurations each; the black line is the average over the four runs.

as shown in Fig. 19. The fraction of network configurations for which the prolonged activity fails before a specified time T^* is reached exhibits a surprising degree of structure in its dependence on τ_D. Most surprising is the finding that an increase in τ_D does not always decrease the probability of failure, but can in fact enhance it. These changes can occur over very small intervals in τ_D, as shown in the expanded window of Fig. 19(b). This fine structure is reminiscent of resonances, although these are more like antiresonances; the values of τ_D within these windows are in some sense optimal for escaping from prolonged activity. While details of the mechanism underlying this structure are not yet understood, it is clear that the dependence on τ_D reflects the significance of the ratio τ_D/T_R. This effect is illustrated in Fig. 20, which shows the dependence of the failure probability on τ_D for two network sizes: $N=500$ and $N=1000$. While the probability of failure is higher overall for the smaller network, the location of the "resonant" windows in τ_D does not show
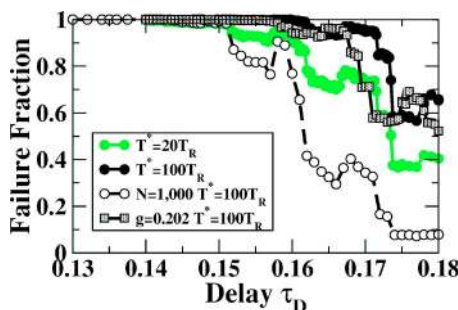

FIG. 20. Failure rates averaged over 8000 network configurations with $N=500$ and $p=1$. Data are shown for both $g_{syn}=0.200$ (circles) and $g_{syn}=0.202$ (squares). Data for $N=1000$ (cf. Fig. 19) shown for comparison.
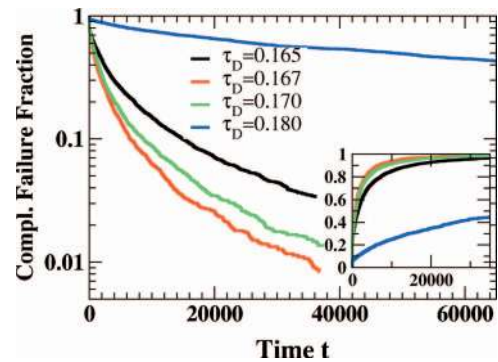

FIG. 21. (Color) Complementary failure fraction $1-\mathcal{F}$ as a function of time for different values of the delay τ_D. Note the logarithmic scale on the vertical axis. Inset: Failure fraction $\mathcal{F}$ for the same data; the vertical scale is now linear.

much dependence on system size. However, even a small reduction in recovery time from $T_R=2.83$ to $T_R=2.79$, due to an increase in synaptic strength from $g_{syn}=0.200$ to $g_{syn}=0.202$, results in a significant shift of the "resonant" windows towards lower values of τ_D.

To assess whether any significant fraction of the network configurations sustains truly persistent activity, we consider the failure fraction $\mathcal{F}$ as a function of the final time T^*, so as to extrapolate to $T^* \to \infty$. The complementary fraction $1-\mathcal{F}(T^*)$ of network configurations that sustain persistent activity up to a time T^* is shown for different values of the delay τ_D in Fig. 21. As anticipated from the nonmonotonicity in Fig. 19, the fraction of failing network configurations is largest for $\tau_D=0.167$; a very rapid drop in failure rate occurs from $\tau_D=0.17$ to $\tau_D=0.18$. When considered as a function of time for fixed τ_D, the behavior of the curves in Fig. 21 indicates that the decay is not exponential.

To obtain an approximate analytic form for the failure fraction $\mathcal{F}(T^*)$, let us consider a specific network configuration for a given value of p. The duration of the activity until failure will then depend on the initial condition. In this numerical experiment we choose a specific network configuration and consider 2000 different initial conditions; to reduce the computational effort these simulations are done for a smaller system with $N=200$. The initial conditions are random and given by $V_i=V_0+\xi_i$, $i=1\ldots N$, with $V_0=0.85$ and ξ_i drawn from a uniform distribution in the interval $[-0.5, +0.5]$. A failure time is determined for the activity triggered by each initial condition. The resulting distribution of failure times exhibits exponential behavior of the form $\beta e^{-t/T}$ for large times; this allows the extraction of a characteristic failure time T associated with this network configuration. An exponential distribution of failure times suggests that the chaotic dynamics effectively lead to a fixed probability for the activity to die out after each population burst.

The characteristic failure time T is then computed for many network configurations of the same size and same shortcut density, leading to a distribution $\rho(T)$ of characteristic failure times, as illustrated in Fig. 22. The large-T behavior of this distribution is well fit by an exponential decay,
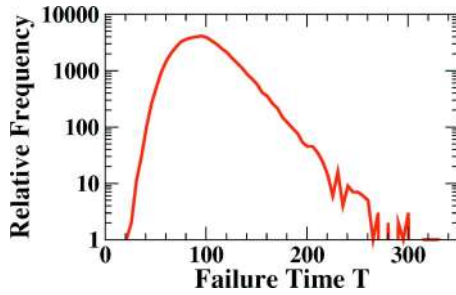
FIG. 22. Distribution of characteristic failure times T, from 50 000 network configurations with $N=200$, $p=1$, and $\tau_D=0.16$. Note the logarithmic scale on the vertical axis.

$$\lim_{T \to \infty} \rho(T) \propto e^{-\alpha T}. \tag{12}$$

The distribution $\rho(T)$ of characteristic failure times determines the average number of failures expected to have occurred by time T^*,

$$\mathcal{F}(T^*) = \int_0^\infty dT \rho(T)(1 - \beta e^{-T^*/T}). \tag{13}$$

Inserting the asymptotic behavior identified in Eq. (12) into Eq. (13) and assuming that the prefactor β does not depend on the characteristic failure time T yields the expected failure rate for large t,

$$\mathcal{F}(t) = 1 - 2\beta\sqrt{\alpha t} K_1(2\sqrt{\alpha t}), \tag{14}$$

where $K_1(x)$ is the first-order modified Bessel function of the second kind. The asymptotic expansion of this Bessel function for large arguments leads to

$$\mathcal{F}(t) \sim 1 - \beta\sqrt{\pi}(\alpha t)^{1/4} e^{-2(\alpha t)^{1/2}}, \tag{15}$$

which displays stretched exponential behavior. The analytic result of Eq. (14) provides a good fit to the time dependence of the failure probability, as shown in Fig. 23. For $\tau_D=0.18$, the fit with $\alpha=0.67\times10^{-5}$ and $\beta=0.91$ is good over essentially the whole time range. For $\tau_D=0.165$ (inset), the fit is not quite as good; the curvature of the analytic function seems to be smaller than that of the data. A fit to the $\tau_D=0.165$ data for $t \geq 3400$ yields $\alpha=0.12\times10^{-3}$ and $\beta=0.64$. This fit (blue dotted line) underestimates the data for long

times. A fit with $\alpha=0.10\times10^{-3}$ and $\beta=0.54$ (red dashed line) reduces the underestimation at long times, but is not as good for smaller times. In any case, the analytic results of Eqs. (14) and (15) indicate the existence of very long transients of prolonged activity. This activity will always eventually fail; $\mathcal{F}(t) \to 1$ as $t \to \infty$.

V. CONCLUSION

In this paper we have used a minimal model to study the influence of network topology on the dynamics of coupled excitable elements. The network consists of a ring of locally connected elements; the connectivity is enhanced through random shortcuts that connect arbitrarily distant elements. Since our goal is to capture gross features of cortical connectivity, we assume that these shortcuts provide only unidirectional connections. This is in contrast to the bidirectional shortcuts that are appropriate for modeling epidemic propagation[19] or regular diffusive processes.

The dynamics of the system exhibit three distinct regimes, depending on the density of shortcuts and the speed of the waves that propagate through the network. For low but nonzero density of shortcuts, activity persists for essentially all network configurations when triggered by the initial excitation of a single neuron. This activity is predominantly periodic, and the mean firing rate of these persistent states shows only little dependence on the wave speed or the density of shortcuts once $p \gtrsim 0.05$. This firing rate is quite close to the maximal firing rate allowed by the recovery period of the neurons.

The recovery period is not to be confused with an absolute, intrinsic refractory period; rather, it is the time from one spike until the membrane potential has recovered to a value such that a single synaptic input of specified strength will suffice to trigger a new spike. This recovery period can be much longer than the absolute refractory period; this phenomenon has been observed in neurons that exhibit a slow after-hyperpolarization that underlies the slow oscillations (<1 Hz) observed *in vivo* in cat[31] and in cortical slices of ferret.[32] There, the recovery period induced by after-hyperpolarization can last as long as a few seconds. Of relevance to the dynamics investigated here is the time associated with the propagation speed of such slow oscillations over the whole network; specifically, the dependence of this propagation time on network connectivity. This dependence has been studied in cortical models[33,34] that do not incorporate shortcuts but use other mechanisms to control and vary network connectivity: either a variable spatial width of the Gaussian distribution that controls the probability that two neurons are connected,[34] or a trimodal probability distribution that captures a type of patchy connectivity in the cortex.[33] As expected, the speed of activity propagation was found to increase with increased connectivity; it was conjectured that this connectivity dependence underlies the large difference in propagation times observed for slow waves in olfactory cortex and neocortex.[34]

For low density of shortcuts, the persistent activity triggered by a localized excitation is periodic. However, the overall dynamical behavior of the system can be quite com-
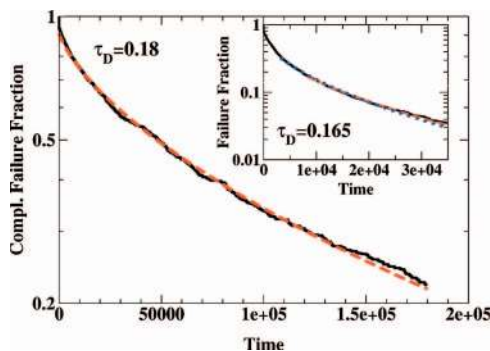

FIG. 23. (Color) Complementary failure fraction $1-\mathcal{F}$ as a function of time for $p=1$ and $\tau_D=0.18$ (inset for $\tau_D=0.165$). Dashed red and dotted blue lines are fits to Eq. (14). Note the logarithmic scale on the vertical axis.

plex due to the coexistence of a large number of stable solutions for a given network configuration. Whether the number of attractors grows as fast with system size as it does for globally coupled oscillator networks is not known at this point. In the case of global coupling, the permutation symmetry leads to a factorial growth of the number of attractors and to attractor crowding;[27] however, this permutation symmetry is broken in the small-world topology. So far as we have been able to investigate the dynamics of increasingly large networks, we have not been able to reach saturation of the number of attractors evoked through different initial conditions. The mechanism underlying this large number of attractors is not yet apparent; however, it is clear that noise will induce switching between these different attractors.[26]

As the density of shortcuts is increased, the number of network configurations that can sustain persistent activity decreases, until persistent activity is essentially no longer possible. For fast waves, the transients after a localized excitation consist of a single population burst followed by activity extinction. For slow waves, the crossover to complete failure occurs at larger shortcut densities than for fast waves. In addition, slow waves can exhibit exceedingly long transients that comprise thousands of population bursts. This activity eventually fails; the failure times for different network configurations at a fixed density of shortcuts follow an exponential distribution, which leads to a stretched exponential distribution for the expected probability of failure as a function of time. The mechanism responsible for these long transients is crucially dependent on the existence of pathways that bridge the quiescent periods between population bursts. These pathways are supported by the topology of the small-world network; this mechanism is thus expected to differ from the one relevant for dilute random networks of pulse-coupled oscillators.[35]

Another quantity of interest is the fraction of network configurations that has already failed by a specified time T^*. This quantity exhibits an overall decreasing trend with decreasing wave speed, or increasing recovery time τ_D. In addition to this decreasing trend, this quantity exhibits an intricate fine structure that includes sharp, resonance-like increases of the failure fraction with decreasing wave speeds. A naive argument supports the decreasing trend, but not the sharp resonances; a decrease in wave speed is likely to allow for additional shorter loops to contribute to the activity, and thus enhance the chances for persistence. However, the existence of sharp increases in failure probability indicates that the activation of one such loop can block the propagation of activity through previously active loops, and thus induce failure. While it is clear that such a switching between loops can occur with increasing τ_D, the mechanism that underlies these increased failures is not yet fully understood.

The crossover to failure, which can be understood in detail based on an analytic mean-field result for the effective size of these idealized small-world networks, provides the basis for understanding a number of recent studies of related but more complex neural network models.[4,22,23]

In Ref. 22, the connection between network connectivity and epilepsy in the hippocampus was investigated by considering small-world networks of three different types of neu-

rons: noisy and leaky integrate-and-fire neurons, stochastic Hodgkin-Huxley cells, and Poisson spike-train cells. As parameters such as synaptic strength, number of synapses per neuron, and proportion of local versus shortcut connections, were varied, the network displayed dynamical behaviors described as "normal," "seizing," and "bursting." For low shortcut density, noise-driven activity was found to remain at a low level; this pattern of activity was associated with normal behavior. With increasing shortcut density, the level of activity triggered by a noise-driven event strongly increased, due to the recruitment of a vastly larger number of neurons. This pattern of activity was likened to seizing behavior. According to our analysis, this regime corresponds to patterns of connectivity that support persistent activity. Yet further increases in the shortcut density were found to induce bursting dynamics, characterized by irregular bursts involving a large fraction of all neurons, separated by quiescent periods. This pattern of activity corresponds in our analysis to network configurations associated with failure, for which each noise-triggered event leads to a population burst that brings essentially all neurons into their recovery period.

In Ref. 23, the role of network connectivity in sustaining population bursting activity was investigated by considering a small-world network of Morris-Lecar neurons. Activity was initiated through a localized set of pacemaker neurons. It was found that in the presence of shortcuts, network activity builds up over several driven cycles into bursts that involve a large fraction of neurons firing within a small time window. Both the time needed to build up such bursts and the time between them were found to decrease with increasing shortcut density. The appearance of repeated bursts is related to the failing configurations discussed here. As expected from our analysis, bursting behavior was supplanted by persistent activity when the wave speed was reduced (cf. increasing τ_D in Fig. 12). The slow buildup towards bursting activity observed in Ref. 23 appears to be specific to the Morris-Lecar neurons used in their model.

In Ref. 4, a 1:20 functional model of the rat dentate gyrus was constructed to investigate the functional consequences of two types of changes in network architecture that take place concurrently during epileptogenesis: loss of hillar mossy cells and sprouting of granular cell axons. The removal of mossy cells resulted in a massive reduction in the total number of connections, but this effect was to some extent compensated by an increase in local connectivity due to spatially restricted sprouting of granular cell axons. Surprisingly, as long as mossy cells were not almost fully eliminated, they continued to provide intermediate and long-range connections onto granule cells, thus preserving the short network diameter characteristic of small-world connectivity. The hyperexcitability that arises from these structural changes corresponds in our analysis to the sustained propagation of fast waves in the presence of a small density of shortcuts.

[1] R. Albert and A. L. Barabasi, Rev. Mod. Phys. **74**, 47 (2002).
[2] M. E. Newman, SIAM Rev. **45**, 167 (2003).
[3] O. Shefi, I. Golding, R. Segev, E. Ben-Jacob, and A. Ayali, Phys. Rev. E **66**, 021905 (2002).
[4] J. Dyhrfjeld-Johnsen, V. Santhakumar, R. J. Morgan, R. Huerta, L. Tsimring, and I. Soltesz, J. Neurophysiol. **97**, 1566 (2007).
[5] M. Barahona and L. M. Pecora, Phys. Rev. Lett. **89**, 054101 (2002).
[6] H. Hong, M. Y. Choi, and B. J. Kim, Phys. Rev. E **65**, 026139 (2002).
[7] T. Nishikawa, A. E. Motter, Y.-C. Lai, and F. C. Hoppensteadt, Phys. Rev. Lett. **91**, 014101 (2003).
[8] P. C. Bressloff, J. Math. Biol. **40**, 169 (2000).
[9] P. C. Bressloff, Physica D **155**, 83 (2001).
[10] D. Golomb and G. B. Ermentrout, Phys. Rev. Lett. **86**, 4179 (2001).
[11] R. D. Traub, D. Schmitz, J. G. Jefferys, and A. Draguhn, Neuroscience **92**, 407 (1999).
[12] T. J. Lewis and J. Rinzel, Network Comput. Neural Syst. **11**, 299 (2000).
[13] T. J. Lewis and J. Rinzel, Neurocomputing **38**, 763 (2001).
[14] M. Perc, New J. Phys. **7**, 252 (2005).
[15] L. F. Lago-Fernández, R. Huerta, F. Corbacho, and J. A. Sigüenza, Phys. Rev. Lett. **84**, 2758 (2000).
[16] L. F. Lago-Fernández, F. J. Corbacho, and R. Huerta, Neural Networks **14**, 687 (2001).
[17] D. R. Paula, A. D. Araujo, J. S. Andrade, H. J. Herrmann, and J. A. Gallas, Phys. Rev. E **74**, 017102 (2006).
[18] A. R. Carvunis, M. Latapy, A. Lesne, C. Magnien, and L. Pezard, Physica A **367**, 595 (2006).
[19] M. Kuperman and G. Abramson, Phys. Rev. Lett. **86**, 2909 (2001).
[20] A. Roxin, H. Riecke, and S. A. Solla, Phys. Rev. Lett. **92**, 198101 (2004).
[21] D. J. Watts and S. H. Strogatz, Nature (London) **393**, 440 (1998).
[22] T. I. Netoff, R. Clewley, S. Arno, T. Keck, and J. A. White, J. Neurosci. **24**, 8075 (2004).
[23] J. Shao, T. Tsao, and R. Butera, Neural Comput. **18**, 2029 (2006).
[24] M. Tinsley, J. X. Cui, F. V. Chirila, A. Taylor, S. Zhong, and K. Showalter, Phys. Rev. Lett. **95**, 038306 (2005).
[25] A. J. Steele, M. Tinsley, and K. Showalter, Chaos **16**, 015110 (2006).
[26] A. Roxin, Ph.D. thesis, Northwestern University (2003).
[27] K. Wiesenfeld and P. Hadley, Phys. Rev. Lett. **62**, 1335 (1989).
[28] M. E. Newman, C. Moore, and D. J. Watts, Phys. Rev. Lett. **84**, 3201 (2000).
[29] M. E. Newman and D. J. Watts, Phys. Rev. E **60**, 7332 (1999).
[30] T. Petermann and P. de Los Rios, Phys. Rev. E **73**, 026114 (2006).
[31] F. Amzica and M. Steriade, J. Neurophysiol. **73**, 20 (1995).
[32] M. V. Sanchez-Vives and D. A. McCormick, Nat. Neurosci. **3**, 1027 (2000).
[33] A. Compte, M. V. Sanchez-Vives, D. A. Mccormick, and X. J. Wang, J. Neurophysiol. **89**, 2707 (2003).
[34] M. Sanchez-Vives and A. Compte, Lect. Notes Comput. Sci. **3561**, 133 (2005).
[35] A. Zumdieck, M. Timme, T. Geisel, and F. Wolf, Phys. Rev. Lett. **93**, 244103 (2004).