

Event-driven simulations of a plastic, spiking neural network

Chun-Chung Chen^{1,2} and David Jasnow²

¹*Physics Division, National Center for Theoretical Sciences, Hsinchu, Taiwan 300, Republic of China*

²*Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA*

(Received 29 March 2011; published 7 September 2011)

We consider a fully connected network of leaky integrate-and-fire neurons with spike-timing-dependent plasticity. The plasticity is controlled by a parameter representing the expected weight of a synapse between neurons that are firing randomly with the same mean frequency. For low values of the plasticity parameter, the activities of the system are dominated by noise, while large values of the plasticity parameter lead to self-sustaining activity in the network. We perform event-driven simulations on finite-size networks with up to 128 neurons to find the stationary synaptic weight conformations for different values of the plasticity parameter. In both the low- and high-activity regimes, the synaptic weights are narrowly distributed around the plasticity parameter value consistent with the predictions of mean-field theory. However, the distribution broadens in the transition region between the two regimes, representing emergent network structures. Using a pseudophysical approach for visualization, we show that the emergent structures are of “path” or “hub” type, observed at different values of the plasticity parameter in the transition region.

DOI: [10.1103/PhysRevE.84.031908](https://doi.org/10.1103/PhysRevE.84.031908)

PACS number(s): 87.18.Sn, 87.19.lj, 87.19.lw

I. INTRODUCTION

Neurons can form plastic networks through connecting synapses with weights that are changed dynamically by the neural activity in the form of neural spikes. It has been established that the precise timing of these spikes determines whether and how the synaptic weights will be increased (potentiated) or decreased (depressed) [1–4]. With better understanding of spike-timing-dependent plasticity (STDP), it becomes increasingly important to find out its implications on the underlying network structure and, consequently, on the neural activity itself. Theoretical models of a plastic neural network typically consist of three components: neural dynamics, synaptic transmission, and network plasticity [5,6]. As has been shown previously [6], a simple self-consistent mean-field scheme can be constructed when these three ingredients of a plastic neural network are given. Reducing the neural and synaptic dynamics to a determination of response functions characterizing a mean-field medium, the network plasticity is left governed by one-dimensional random-walk dynamics. Such a mean-field scheme was applied [6] to a network of leaky integrate-and-fire neurons [7–9], coupled through a jump-and-decay synaptic conductance dynamics [10], and under STDP rules [2], controlled by a *plasticity parameter* w^* representing the expected value of synaptic weights when firings of the neurons are purely driven by noise. While we do not expect that there is a broad range of cell types with different plasticity parameter values but otherwise similar physiological characteristics, different values of w^* can correspond to different stages of a (quasistatically) growing and developing network. The mean-field theory (MFT) predicts a first-order phase transition and hysteresis from a low w^* regime of noise-driven activity to a high w^* regime of self-sustaining activity. It also predicts a narrow synaptic weight distribution as long as the overall rate of change of synaptic strength is low. In the current study, we perform intensive event-driven simulations [11] on the same model network to compare with the predictions of MFT; the simulations also reveal emergent

network dynamics and structures that are not captured by the MFT.

In Sec. II that follows, we briefly describe the simulated model of a plastic neural network, which was elaborated in our MFT study [6]. We then present in Sec. III the method and algorithm of the event-driven simulations used to explore the dynamics of the model. The results of the simulations for static uniform networks are reported in Sec. IV, where we present the mean activity levels for networks of different sizes as functions of synaptic input. These activity levels agree well with the MFT predictions in the stable phases of low and high activity but show a gradual increase in a transition region ($w_1^* < w^* < w_2^*$, from the onset to the persistence of threshold firing events) instead of the sharp first-order jump predicted by MFT. We also analyze the sporadic patterns of self-sustaining threshold firings in the transition region to identify two types of threshold firing events. In Sec. V, we present simulation results with high resolution in the plasticity parameter w^* for a plastic network with $N = 32$ neurons that has reached a stationary state of plasticity. The change in the synaptic-weight conformation of the network in the transition region manifests itself in the synaptic-weight distribution, which is seen to broaden twice, along with a bimodal elevation of the average firing activity compared to that of a static network. To characterize the emergent synaptic-weight structure of the network in the transition region, we employ a pseudophysical approach for visualization in Sec. VI to generate a two-dimensional (2D) layout of the network. Such an approach reveals a *path* or *loop* conformation near the lower end of the transition region ($w^* \sim w_1$) for network sizes up to $N \approx 32$ and a *hub* conformation near the high end ($w^* \sim w_2$) for $N \approx 24$ and greater. We then conclude and summarize our findings in Sec. VII.

II. MODEL

While our method is applicable to other combinations of models of neurons, synapses, and plasticity, we follow the

same choices made in our mean-field approach [6], which we describe briefly below.

In the integrate-and-fire model for the neurons, the state of a neuron i is described by a membrane potential V_i , which follows the differential equation of a leaky integrator [12]

$$\tau_m \frac{dV_i}{dt} = V_0 - V_i + G_i(\mathcal{R} - V_i), \quad (1)$$

where τ_m is the leak time for the membrane charge, V_0 is the resting potential when the neuron is in the quiescent state, and \mathcal{R} is the reversal potential for the ion channels on the synapses. The total synaptic conductance G_i for the neuron is given by the sum

$$G_i \equiv \sum_j w_{j,i} Y_j \quad (2)$$

over all presynaptic neurons, j . The synaptic weights $w_{j,i}$ define the network, and the same active transmitter fraction Y_j is assumed for all efferent synapses of neuron j . In addition to the continuous dynamics (1), a neuron fires when its membrane potential reaches a threshold value, V_{th} . Then, its membrane potential drops immediately to a reset value, V_r . The action potential of the integrate-and-fire model is assumed to be instantaneous and is not modeled explicitly. The spike train produced by the neuron i is defined as the function

$$S_i(t) \equiv \sum_n \delta(t - t_{i,n}), \quad (3)$$

where $t_{i,n}$ is the time when the neuron i fires for the n th time.

The fraction Y_j of the active transmitters is described by the Tsodyks-Uziel-Markram (TUM) model [10] of neural transmission, where the transmitters are distributed in three states: “active,” with the fraction Y ; “inactive,” with the fraction Z ; and “ready-to-release,” with the fraction X . For efferent synapses of a presynaptic neuron j , these fractions follow the dynamics [10]

$$\begin{aligned} \frac{dX_j}{dt} &= \frac{Z_j}{\tau_R} - u S_j X_j, \\ \frac{dY_j}{dt} &= -\frac{Y_j}{\tau_D} + u S_j X_j, \\ \frac{dZ_j}{dt} &= \frac{Y_j}{\tau_D} - \frac{Z_j}{\tau_R}, \end{aligned} \quad (4)$$

where τ_D is the decay time of active transmitters to the inactive state, τ_R is the recovery time for the inactive transmitters to the ready-to-release state, and u is the fraction of ready-to-release transmitters that is released to the active state by each presynaptic spike. With the conservation rule

$$X_j + Y_j + Z_j = 1, \quad (5)$$

there are two independent variables per presynaptic neuron. Consistent with the TUM dynamics, the values of the factors multiplying S_j at the discontinuities are to be evaluated immediately before the discontinuities.

While the integrate-and-fire and TUM dynamics are both deterministic, we model the stochasticity of the network with additional noise-driven firing events following Poisson statistics with the frequency λ_N for each neuron. The noise-driven firings are treated the same way as threshold firings;

TABLE I. Values of parameters used in calculations.

Integrate and fire		TUM model	
Resting potential V_0 :	−55 mV	Decay time τ_D :	20 ms
Leak time τ_m :	20 ms	Recovery time τ_R :	200 ms
Firing threshold V_{th} :	−54 mV	Release fraction u :	0.5
Reset potential V_r :	−80 mV	Noise frequency λ_N :	1 Hz
Reversal potential \mathcal{R} :	0 mV	Plasticity rate r :	0.01

that is, the membrane potentials are brought instantaneously to the reset value V_r and the firing times are included in the spike trains (3) of the transmitter dynamics (4).

To minimize the computational cost, the active transmitter fractions Y_j double as the exponentially decaying “window functions” in our version of the plasticity rules suggested by van Rossum *et al.* [2]. The synaptic weights are taken to follow the dynamical equation

$$\frac{dw_{j,i}}{dt} = \Delta Y_j S_i - r w_{j,i} Y_i S_j, \quad (6)$$

where Δ is the parameter for additive potentiation and r is the parameter for multiplicative depression. We define the plasticity parameter as the ratio

$$w^* \equiv \frac{\Delta}{r}, \quad (7)$$

which is equal to the expectation value of the synaptic weight when we have the symmetry $\langle Y_j S_i \rangle = \langle Y_i S_j \rangle$, e.g., when S_i and S_j are both Poisson spike trains of the same frequency. Fixing w^* leaves the depression factor r as an overall control parameter for the *rate of plasticity*.

It is common for a model of a biological system to carry a large number of empirical parameters. Instead of exploring all possible ranges of these parameters, we fix them with physiologically plausible values that are commonly found in the literature. Unless otherwise stated, the values of the parameters are as listed in Table I.

III. SIMULATION METHOD

The disparity between the time scales of spiking activities and neural plasticity presents a significant challenge to computer simulations of plastic neural networks. Particularly for STDP, where precise timing is crucial in determining synaptic changes, we cannot alleviate the computation requirement through use of larger integration time steps. However, for continuous dynamics that can be solved analytically, one can improve the efficiency of computation through an *event-driven* approach similar to that described by Brette [9], which gives machine precision timing for the spikes and requires a limited amount of computation upon each spike production.

In the event-driven approach, instead of calculating the state of the system for a fixed increment of time, one calculates the time for the next discontinuous event. This approach is feasible when the continuous dynamics of the system is solvable so that the time for the next discontinuous event can be evaluated efficiently. For the leaky integrate-and-fire model considered,

an analytical form can be written down for the trajectory of the normalized membrane potential [9]

$$v_i \equiv \frac{V_i - V_0}{\mathcal{R} - V_0} = x e^x \left[E_1(x) - E_1(G) + \frac{v}{G e^G} \right], \quad (8)$$

where E_1 is the exponential integral function of the first kind, $x \equiv G e^{-t/\tau_m}$, and $v = v_i(0)$, $G = G_i(0)$ are the initial states of the neuron. For the TUM dynamics, the transmitter fractions follow the trajectories

$$Y_i = e^{-t/\tau_D} Y, \quad (9)$$

$$Z_i = \left[Z + \frac{\tau_R}{\tau_R - \tau_D} Y \right] e^{-t/\tau_R} - \frac{\tau_R}{\tau_R - \tau_D} Y_i,$$

where $Y = Y_i(0)$ and $Z = Z_i(0)$ are their initial values. With the trajectory (8), one can solve for the time of the next threshold crossing $V_i \rightarrow V_{th}$ through, e.g., a root solver. The computation time for a root solver to find the position of the crossing depends only logarithmically on the precision required, compared to the linear increase for the fixed-time-step approach. Further improvement of the computational efficiency can be achieved by, e.g., precomputing a lookup table for the time-to-threshold crossing given the current state of a neuron and interpolating when the desired value falls between precomputed points. For the simulated network with Poisson noise, the algorithm is as outlined below:

Step 1. For each neuron i in the system, calculate the time-to-fire t_{tf}^i from the continuous, deterministic dynamics of the model.

Step 2. Draw the time to its next noise-triggered firing t_N^i from an exponential distribution with the expectation value λ_N^{-1} for each neuron i .

Step 3. Find the minimum of the times t_{tf}^i and t_N^i among all neurons to be the time to the next discrete event Δt of the system.

Step 4. Advance the simulated clock by Δt and update the state of the system using Eqs. (8) and (9).

Step 5. Fire the selected neuron i by resetting its membrane potential to $V_i = V_R$ and increasing the affected active transmitter fraction Y_i by $u(1 - Y_i - Z_i)$.

Step 6. Repeat this process from step 1.

For the current model, the simulated dynamics boils down to the evaluation of the time-to-fire $t_{\text{tf}}(V, G)$ given the initial membrane potential V and total synaptic conductance G as well as the evaluation of the trajectories (8) and (9). We note that instead of the common approach of defining the model with the set of differential equations (1) and (4), it is equally valid and computationally preferable to define the continuous dynamics of the model with the trajectories (8) and (9) for the neurons and synapses.

IV. STATIC NETWORK

Before incorporating plasticity, we first perform event-driven simulations on fully connected networks of up to 128 neurons with uniform, fixed synaptic weights w . This exercise will prepare us for the inclusion of plasticity and also reveal the emergence of “structure” in the network. The average activity levels for different system sizes are shown in Fig. 1. compared with predictions from MFT [6]. As expected, the




FIG. 1. (Color online) Mean activity level vs total synaptic input (Kw) for fully connected uniform networks (node degree $K = N - 1$) with static synaptic weight (w). Dashed lines are predictions from mean-field theory without fluctuation corrections (long dashes) and with the inclusion of shot-noise-like fluctuations (short dashes) in total synaptic conductance for a network with $K = 63$ afferent synapses per neuron.

average activity levels of the simulated network coincide with mean-field predictions in the stable phases of low and high activity. However, in the transition region between the two phases, the activity levels from network simulations increase continuously instead of exhibiting the jumps and hysteresis predicted by mean-field theory. We note that the hysteresis in MFT [6] comes about when there are two stable states in the system: a quiet state with only noise-triggered firings that are insufficient to ignite the entire system, and an active state in which firings are system wide and self-sustaining. However, for a finite-size network having sufficient fluctuations and running for a sufficient time, the system can make transitions between the two states, leading to a single-valued mean activity level of the system showing no hysteresis. Comparing networks of different sizes, the rise of the average firing frequency in the transition region is steeper for larger networks since fluctuations in the total synaptic conductance *decrease* with an *increase* in the number of afferent synapses [6] rendering mean-field theory a better approximation.

The firing pattern of the $N = 64$ network is shown in Fig. 2 for four different values of synaptic input Kw as marked in Fig. 1. With computer simulations, we are able to distinguish firings due to threshold crossings of the membrane potentials (marked with dark vertical line segments) from noise-triggered firings (marked with light vertical line segments). We can thus define a system to be active when there is any neuron with a membrane potential set to cross the firing threshold without the help of further noise events. Operationally, this is when the time-to-fire t_{tf}^i calculated in step 1 of the event-driven algorithm outlined previously is finite for any of the neurons. The switching between the quiet and active states of the network are evident on the strips B and C. This leads to clusters of threshold firings during which the system remains active. In the specific model we considered, there are two types of threshold firings clusters that can be identified from the polygraphs. The first represents *bursting*, which are brief clusters (up to a few hundred milliseconds) of threshold firings of generally similar durations. These can be seen for all threshold activities on strip B and some on strip C as marked

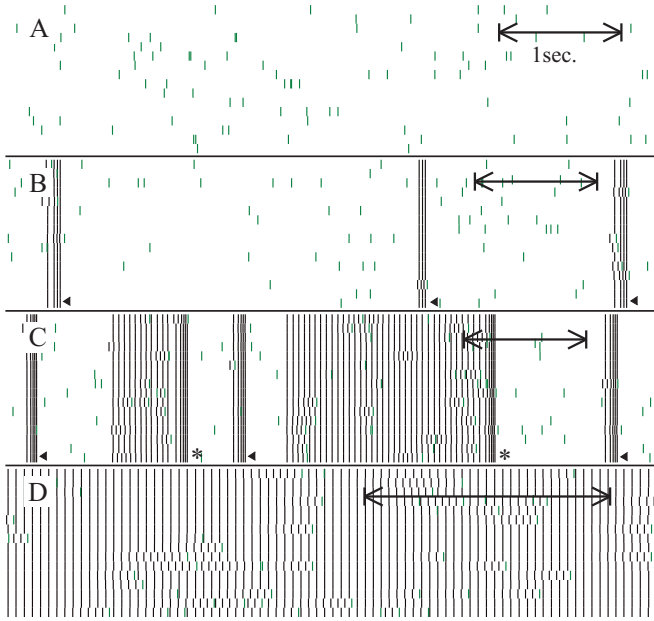


FIG. 2. Firing polygraphs of the $N = 64$ network at four different values of synaptic strength w labeled in Fig. 1. For each value, only activities of 16 randomly chosen neurons are shown, with each occupying $1/16$ of the height of the corresponding strip. The light and dark vertical line segments mark the firings due to noise and threshold crossing, respectively. The time on each strip of recording runs toward the left with a scale indicated by the 1 s time interval on each strip. Symbols (\blacktriangleleft and $*$) mark clusters of threshold firings as detailed in text.

by the symbol “ \blacktriangleleft ” in Fig. 2. The second type of threshold firing activity is *intermittently persisting* activity, which can last for seconds, are of a wider range of durations and are seen in two time segments marked by the symbol “ $*$ ” on strip C of Fig. 2. The two types of clusters of threshold firings can also be identified from the duration distribution of self-sustained threshold firings shown in Fig. 3. The distinction between the two kinds of activities can be attributed to the short-term depression caused by the inactive state of synaptic transmitters in the TUM model: When threshold firings of the network are just ignited after a quiet period, most transmitters are in the ready-to-release state and activities can propagate easily leading to a somewhat higher firing rate of the neurons in the beginning. However, as most transmitters are deposited into the inactive state during repeated firings, the firing frequency is reduced. At the low-activity end of the transition region, the reduction of firing frequency typically continues all the way to a cessation of any threshold firings, leading to the short bursting events of similar duration. Near the high-activity end of the transition region, as the transmitters become inactive, the reduced firing frequency can settle to a stable value that can last for seconds before all threshold firings stop. From the simulation results such as what are shown in Fig. 3, the stopping of the threshold firings appears to follow Poisson statistics; that is, it can be characterized by a stopping rate τ^{-1} . The stopping time or *duration* distribution of sustained threshold firings decays exponentially for large duration with the time constant τ that increases with w and suggests a divergence at some $w = w_p$. For the $N = 32$

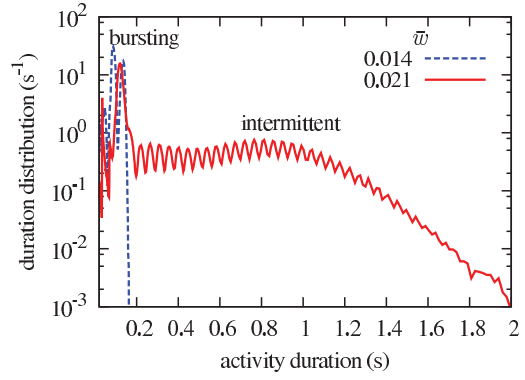


FIG. 3. (Color online) Time-weighted duration distribution of sustaining threshold firing activities. The system is considered active when any of the neurons is set to cross the firing threshold without further noise events. The duration is measured for each episode of the system staying continuously active. The oscillations in the distribution come from the periodic nature of persistent firing activities. Values of the synaptic weight \bar{w} (0.014, dashed line and 0.021, solid line) correspond to the locations B and C, respectively, in Figs. 1 and 2.

network, we estimate $w_p \approx 0.0487$ and that the “stopping-time constant” diverges roughly with the scaling $\tau \sim (w_p - w)^{-2}$ as suggested in Fig. 4. We note that the scaling form of the stopping time remains the same when we simplify the TUM dynamics in our model by removing the inactive state (data not included here), suggesting at least some degree of universality. These results suggest that the divergence of the stopping-time constant τ signals that the system enters the phase of truly self-sustaining activity.

The Poisson statistics, as seen in the exponential tail of our *duration distribution* for large durations is in contrast with the critical avalanches observed by Beggs and Plenz in neocortical slices [13]. Furthermore, instead of having avalanches with a distribution of sizes, during the active periods (as seen in Fig. 2) of our network, the threshold firings are generally system wide. Nonetheless, the apparent power-law divergence of the stopping time τ as the system enters the active phase is suggestive of criticality for such a transition. In a




FIG. 4. (Color online) Scaling of stopping-time constant estimated from the tail of active duration distribution in Fig. 3 for an $N = 32$ static uniform network. The best estimate of $w_p = 0.0487$ for the diverging point (center curve) is shown along with two slight deviations. The dashed line is to show the estimated slope of -2 for comparison.

uniform system of constant synaptic weights as considered here, the critical point would require a tuning of the synaptic weights to their critical value $w = w_p$. However, for a more realistic model with a smooth distribution of local synaptic weights, near the onset of a self-sustaining phase, a power-law dependence of the stopping time could be a reasonable expectation [14]. Alternatively, in a biological network, tuning might result from functional necessity to have the system operate near a critical point [15]. It is not unimaginable that the power-law size distribution of avalanches in neocortical slices might also be related to an underlying critical point, which could correspond to the onset of, say, system-wide bursting instead of self-sustaining activity. Such possibilities warrant much further investigation that likely would require significant changes to the network considered here (for example, making it heterogeneous and not fully connected) that are beyond the scope of the current study.

V. DYNAMIC (PLASTIC) NETWORK

For the current study, we focus on the stationary states of the system under the plasticity dynamics. A uniform network with synaptic weight $w_{j,i} = w^*$ for all synapses is used as the initial configuration, and simulations are conducted at each data point for at least 2^{31} ms (≈ 25 days) of simulated time to reach the stationary state. We concentrate our calculations on a fully connected network of $N = 32$ neurons with a high resolution in the plasticity parameter w^* . For each w^* , all measurements are averaged over an ensemble of 32 or more independent runs of different random sequences. The results of the calculations are summarized in Fig. 5, where we see that the firing frequency of a stationary *plastic* network coincides with that of a uniform *static* network ($w = w^*$ for all synapses) in the low-activity, noise-dominated regime as well as in the persistently active regime. In the transition region between the noise-driven and self-sustaining regimes, the firing rate of the network is enhanced by the plasticity, and while the behavior is not universal, with TUM synapses the amount of enhancement exhibits an interesting bimodal shape.

In the insets of Fig. 5, we show the synaptic weight distribution, also averaged over the same ensemble, at each marked point along the activity-plasticity curve. In the two stable regimes, noise dominated and persistently active, the synaptic weights have a Gaussian distribution with a narrow width (within about 5% of w^*) as predicted by MFT [6]. However, coincident with the bimodal enhancement of firing frequency, the synaptic weight distribution of the stationary plastic network shows dramatic changes in the transition region along the activity-plasticity curve: Increasing the plasticity parameter w^* from the noise-dominated low-activity regime, we see a discontinuous jump in the firing rate and sudden broadening of the weight distribution with tiered side peaks at $w^* \approx 0.017$. The enhanced firing and broadened distribution ease off after the jump, until these features become insignificant around $w^* = 0.04$. At this point, with a slight increase of w^* , two disconnected side peaks pop out in the synaptic weight distribution. However, such a splitting in the weight distribution is only accompanied by a gradual enhancement of average firing rate of the plastic network. The two side peaks eventually merge back to the main peak with further




FIG. 5. (Color online) Activity-plasticity (parameter) plot for a fully connected plastic network of $N = 32$ neurons compared with the activity level of a static network of the same size with uniform weight $w = w^*$. The plastic network is allowed to reach a stationary state by running for about 25 days of simulated time. All results are averaged over an ensemble of 32 or more runs of different random sequences. Insets of the plot show the normalized synaptic weight distributions on a log-log scale with a range of 2 (4) decades on the horizontal (vertical) axis at different points marked on the activity-plasticity curve.

increases of w^* . This eventually returns the network to a uniform conformation with a narrow Gaussian synaptic-weight distribution before the plasticity parameter reaches $w^* \approx 0.1$.

Instead of a simple increase of the Gaussian width, the observed broadening of the synaptic-weight distribution in the transition region just summarized (see Fig. 5) comes with structured side peaks, which signify emergent network conformations that we try to decipher and visualize with a pseudophysical approach in Sec. VI. For the $N = 32$ and smaller networks, we are able to obtain a stationary state from the simulation run that is insensitive to the initial configuration at each data point (value of w^*). However, we do see runaway synaptic weights under the plasticity rules used for the $N = 48$ and $N = 64$ networks in narrow, isolated intervals at $w^* \approx 0.0268$ and $w^* \approx 0.0194$, respectively, near the high-activity end of the transition region before the runs (2^{31} ms, ~ 25 days, of simulated time) are complete. For these instances, the diverging synaptic weights and firing frequencies of the driven neurons slow the simulation down to a near standstill, and we cannot always maintain a stationary system as we vary the plasticity parameter across the runaway point. We note that such runaways are commonly seen in models with Hebbian plasticity and are typically “cured” with cutoffs in the range of synaptic weights [16–19]. We do not need such a device for the $N = 32$ network results presented. For larger networks, when we do see such “isolated” runaways, we regard them as pathological [20,21] and exclude them from the ensembles.

VI. NETWORK “LAYOUT”

Characterizing the structure of nonuniform networks is an active field of research [22–25]. Here we adopt a more intuitive and visual approach to extract any emergent structures of the stationary plastic network obtained above. Under our model setup, such structures reside in the framework of a fully connected network of uniform connectivity, represented by the main peak in the synaptic weight distribution. In this section,

we attempt to uncover the topological meaning of the side peaks seen for the distribution in the transition region between the two stable phases of the system.

In order to visualize the networks such that any emerging features can be better revealed, we introduce a pseudophysical system of fictitious interacting particles in two dimensions representing the neurons, recognizing, of course, that the actual all-to-all network is not strictly two dimensional. The nodes (pseudoparticles) are imagined to be located on the plane, confined to the region near an arbitrary origin by an isotropic harmonic potential. The nodes interact with a (fictitious) pairwise repulsive force that is determined by the synaptic weights between the neurons. To have the pseudoparticles representing neurons that are more strongly connected stay close to one another, we choose a repulsive pairwise force

$$\vec{F}_{i,j} = \frac{\hat{e}_{j,i}}{w_{i,j}^2 + w_{j,i}^2} \quad (10)$$

to act on particle i due to j with $\hat{e}_{j,i} = \vec{r}_{j,i}/|\vec{r}_{j,i}|$ being the unit vector along the $\vec{r}_{j,i} = \vec{r}_i - \vec{r}_j$ direction (in the two-dimensional pseudospace). In fitting the entire layout of the pseudoparticles (representing neurons) to a fixed display area, the strength k of the confining harmonic potential is adjusted so that, in an equilibrium arrangement of the particles, the radial coordinate \vec{r} of the particle farthest from the origin of the potential has magnitude equal to 1. Such a pseudophysical system of neurons is relaxed to reach a stable arrangement of neurons in which the net force on each pseudoparticle (neuron) is zero. The relaxation is accomplished through a deterministic, overdamped dynamics in which the velocity of each node is proportional to the net force it experiences,

$$\frac{d\vec{r}_i}{dt} = -k\vec{r}_i + \sum_j \vec{F}_{i,j}. \quad (11)$$

The representative layouts for the marked data points in Fig. 5 are shown in Fig. 6, where we only show a fraction of the synapses with strongest weights along with the pseudopositions to reduce obscurity. We note that the stable layouts reached are not unique for the system, and it is possible for the pseudoparticle system to be trapped in a metastable configuration in which strongly connected neurons are hindered by other neurons and fail to get close to each other during the relaxation. We do not attempt to further relax the layout (by, e.g., perturbing the positions of the neurons) after a stable configuration is reached during the relaxation process. (We note that for larger systems where the “hindrance” can be more severe, one may consider a three- or higher-dimensional pseudophysical system to help alleviate the problem.)

As revealed in layouts of Fig. 6, the jump in the firing rate near $w^* = 0.017$ is accompanied by the formation of a path conformation in the network where synapses of stronger synaptic weights connect a sequence of neurons and facilitate the propagation of activities along the path. Such a path can form a closed loop near the jump but become less likely to do so as w^* is increased. The path or loop structure is stable in the system such that once it forms, it will generally remain to the end of our simulation run at a fixed w^* value. As we only present results of simulations from uniform initial networks, there is no hysteresis loop in Fig. 5. However, the jump in the firing rate signals a first-order phase transition. While we have not attempted a systematic study of hysteresis effects with larger systems, we do see some degree of hysteresis in smaller ($N \leq 16$) networks: If we use a loop conformation obtained in the stationary state of the network at a w^* value above the jump as the initial condition for the simulation of the network at a w^* value below the jump, the loop conformation can sometimes persist over the course of the entire simulation. With a path or loop conformation of the network, the tiered side peaks of the synaptic weight distribution in this area come from the

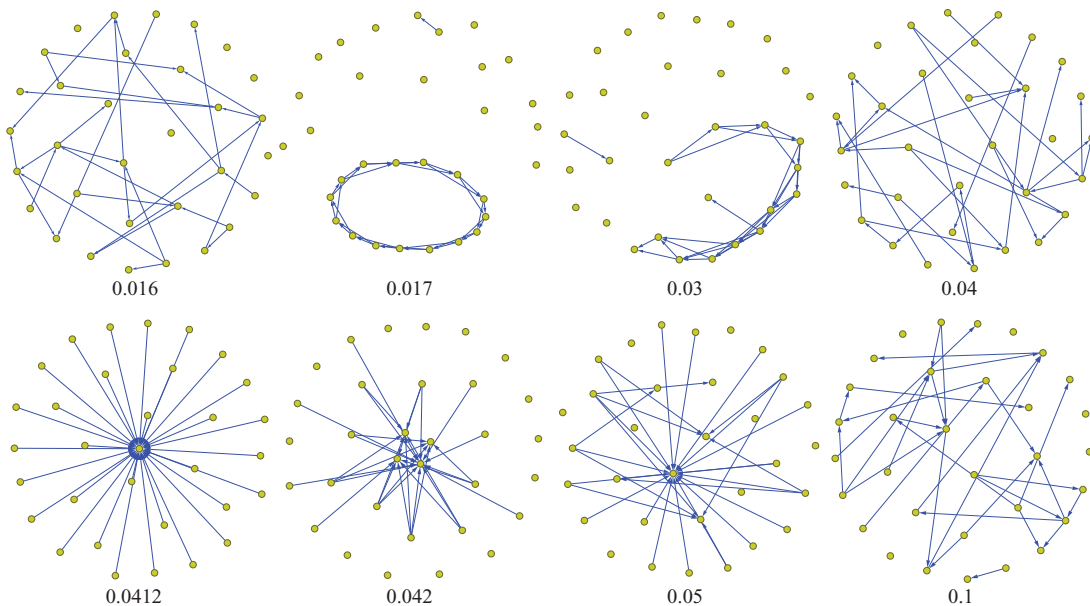


FIG. 6. (Color online) Typical resultant layouts of neurons (represented by the dots) following the pseudophysical approach described in the text at the points marked on the activity-connectivity curve in Fig. 5. The number labeled below each layout is the corresponding value of the plasticity parameter w^* . Only the strongest 31 synaptic connections (lines with arrows) are shown in each layout.

skip-over synapses along the path, i.e., synapses that are along the direction of the path but bypass some of the neurons.

Eventually, with the increase of w^* , the path falls apart and the network returns to a uniform conformation around $w^* = 0.04$. The emergence of disjoint side peaks in the synaptic weight distribution with a slight increase of w^* from 0.04 as seen in Fig. 5 is accompanied by the formation of a *sink hub* structure in the network where the high-activity level of the hub neuron and its strong afferent synapses reinforce each other because of the positive-feedback nature of the synaptic plasticity. This type of hub structure is first seen at $w^* = 0.0411$ in our simulations and can go in and out of existence several times during the course of our simulation run at a fixed w^* . This partially explains why we only see a gradual increase of the average activity level when a hub starts to form. The sink hub structure becomes more and more stable as w^* increases until multiple hubs appear in the network around $w^* = 0.0414$. The network returns to a uniform conformation through the increasing number of less and less significant hubs. Finally, at sufficiently large w^* the network has reached a homogeneous, high-activity state.

We can now try to understand the formation of “path” and “hub” conformations in different segments of the transition region separating homogeneous, stable low- and high-activity states of the network. Near the low-activity end of the transition region, the sequence of neurons along the path represents a firing order in a bout of threshold firings that is amplified and preserved by the STDP as the path of stronger synaptic weights. Especially when such a path forms a loop, threshold firing activity can cycle through all neurons in the loop no matter which one of them is triggered by the noise. However, near the high-activity regime, each neuron generally fires multiple times during an episode of sustaining threshold firings with a high frequency. The firing order of neurons loses its meaning because there are many ways of pairing up their spikes. In this area of the transition region, the main distinguishable feature of a neuron is simply its firing rate. In fact, the expectation value of a synaptic weight w for a synapse in this range can be uniquely determined from the firing rates of its pre- and postsynaptic neurons when the firing rates of the neurons are high enough ($\lambda > \tau_D^{-1}, \tau_m^{-1}$) so that the actual timing of spikes becomes unimportant. Under such a condition, the stationary synaptic weight under the plasticity dynamics (6) for, say, the $j \rightarrow i$ synapse, is given by

$$w_{j,i} = w^* \frac{\bar{Y}_j \lambda_i}{\bar{Y}_i \lambda_j}, \quad (12)$$

where $\bar{Y}_i \equiv \langle Y_i \rangle$ and $\lambda_i \equiv \langle S_i \rangle$ are, respectively, the average active transmitter fraction and mean firing rate of the neuron i . With Eq. (12), the condition

$$w_{i,j} w_{j,k} w_{k,i} = w^{*3} \quad (13)$$

should hold for any neurons $i \neq j \neq k$ in the network [25]. And, indeed, we have verified that the condition (13) is well satisfied by the hub conformations obtained in the simulations, but is apparently violated by the path conformations (data analysis not presented here).

While it is tempting to correlate the layouts resulting from our pseudophysical arrangement with configura-

tions of physical neurons, we must caution that one can only view such a pseudophysical approach as an arbitrary and nonunique way of revealing network structures. Furthermore, the structures found by the pseudophysical approach do not represent any direct physical information contained in the parameters of the neuron network. Nonetheless, it helps to provide an intuitive “structural” understanding of the network conformation as we have shown above.

VII. CONCLUSIONS

The event-driven algorithm presented in Sec. III improves the accuracy and efficiency of the simulations allowing us to gain an intensive view of the phase space of a stationary plastic neural network. With a large number of parameters typical for a model of a biological system, we fix all but one, the plasticity parameter, with physiologically plausible values for our investigation. Also, to minimize conceptual complications, we study only the stationary properties of fully connected networks driven by uniform, noncorrelated Poisson noise. Our results show the network develops interesting structures in the transition region between the stable phases corresponding to low- and high-activity regimes. While our fully connected network with a limited number of neurons is certainly inadequate to address the dynamics of a real brain, it can be a reasonable starting point for cultured networks consisting of hundreds of neurons with virtually “all-to-all” interactions [13,26–29]. Our finding suggests that emergent structures in these networks are more likely to be seen when it is firing intermittently in the transition region perhaps during its development from a weakly coupled system of neurons to a more strongly connected network.

The characterization of network connectivity is an active field of research with established quantifications such as the “clustering coefficient” [30] and “modularity” [24]. However, in the current study, we adopt a simple visualization approach to gain a more intuitive view of the network structure. The resultant identification of the path and hub conformations explains, or more accurately, coincides, with the appearance of structured side peaks in the synaptic-weight distribution and the elevated firing rate of the neurons. While these path and hub conformations are the simplest forms that can appear in a connected network, the present work demonstrates that they can arise naturally in transition regions of a neural network under STDP and driven only by noise. The connection-weight conformation (i.e., the distribution and correlations of synaptic weights) of a naturally occurring network is often strongly influenced by activity during its formation and maturation. However, studies of network dynamics typically focus on behavior of the network under given, static network topology or weight conformation. Advances in the understanding of the network plasticity are setting the stage for addressing how the observed conformation of these networks can come about. Our study of a pure and simple network is just an initial step along this direction, and it leaves open questions of how variations of biological details that are fixed in our model and different network topology, geometry, or sparseness might affect the emergent structures and alter the dynamical behavior of the network.

Furthermore, while in this work we have only considered a stationary network driven by uncorrelated noise, arguably the ultimate “goal” of a plastic neural network is *learning* to serve specific or varying functional requirements. Recent studies have shown that STDP is effective in detection of phase-of-firing coded information for an individual postsynaptic neuron [31], especially with a systematic oscillatory background field [32], or in the storage of different timing-based firing patterns presented to a network [33]. Such applications are possible through the reinforcement of the timing information by the STDP. From our analysis in Sec. VI, the hub conformations of the stationary network synaptic weights are consistent with a “conservative” vector field (see [25]) that can be derived from the firing rates of the pre- and postsynaptic neurons. On the other hand, such “conservation” is violated in path conformations where the firing order of neurons on the path are seen to be *preserved* by the STDP. This preservation of firing order suggests that we can expect such timing-based learning to be effective when path conformations are important, that is, near the low-activity end of the transition region shown in Fig. 5. It is among the goals of our future work to subject the system to more meaningful, both timing- and rate-based inputs

to explore the changes in the stationary structures as well as the transient dynamics of the resultant network.

Finally, there remain pronounced and puzzling features observed in the current study that require further elucidation. For example, while both types of conformations (hubs and paths) appear abruptly as the plasticity parameter w^* is varied, the average firing rate increases with a discontinuous jump accompanying the appearance of the path conformation, but continuously for the hub. Also, the system does not “morph” from the path conformation to the hub conformation directly, but instead, returns to a uniform network before developing the hub structure. This feature could suggest a symmetry or “topological” difference between the two types of conformations that prevents a direct transition from one conformation to the other. We hope that further investigation of such observations can lead to a better understanding of plastic neural networks in general.

ACKNOWLEDGMENT

This research was supported in part by Computational Resources on PittGrid (www.pittgrid.pitt.edu).

-
- [1] L. F. Abbott and S. B. Nelson, *Nat. Neurosci.* **3**, 1178 (2000).
 [2] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano, *J. Neurosci.* **20**, 8812 (2000).
 [3] P. D. Roberts and C. C. Bell, *Biol. Cybern.* **87**, 392 (2002).
 [4] Y. Dan and M.-M. Poo, *Neuron* **44**, 23 (2004).
 [5] A. Morrison, A. Aertsen, and M. Diesmann, *Neural Comput.* **19**, 1437 (2007).
 [6] C.-C. Chen and D. Jasnow, *Phys. Rev. E* **81**, 011907 (2010).
 [7] W. Gerstner and W. M. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity* (Cambridge University Press, Cambridge, 2002).
 [8] A. Burkitt, *Biol. Cybern.* **95**, 97 (2006).
 [9] R. Brette, *Neural Comput.* **18**, 2004 (2006).
 [10] M. Tsodyks, A. Uziel, and H. Markram, *J. Neurosci.* **20**, RC50 (2000).
 [11] R. Brette, M. Rudolph, T. Carnevale, M. Hines, D. Beeman, J. Bower, M. Diesmann, A. Morrison, P. Goodman, F. Harris, M. Zirpe, T. Natschläger, D. Pecevski, B. Ermentrout, M. Djurfeldt, A. Lansner, O. Rochel, T. Vieville, E. Muller, A. Davison, S. E. Boustani, and A. Destexhe, *J. Comput. Neurosci.* **23**, 349 (2007).
 [12] P. Dayan and L. F. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (Massachusetts Institute of Technology Press, Cambridge, MA, 2001).
 [13] J. M. Beggs and D. Plenz, *J. Neurosci.* **23**, 11167 (2003).
 [14] We note that a simple average over an ensemble of *uniform* networks with a probability distribution $p(w) = \text{const.}$ for $0 < w < w_p$ results in a stopping-time distribution $P(\tau) \sim \tau^{-3/2}$ provided $\tau \sim (w_p - w)^{-2}$ as estimated in Fig. 4. This will hold for more general smooth functions $p(w)$ near w_p , but it is not clear that this ensemble will represent a real network with a distribution of synaptic strengths w .
 [15] R. Legenstein and W. Maass, *Neural Networks* **20**, 323 (2007).
 [16] A. Amarasingham and W. B. Levy, *Neural Comput.* **10**, 25 (1998).
 [17] S. Song, K. D. Miller, and L. F. Abbott, *Nat. Neurosci.* **3**, 919 (2000).
 [18] J. Rubin, D. D. Lee, and H. Sompolinsky, *Phys. Rev. Lett.* **86**, 364 (2001).
 [19] T. Toyozumi, J.-P. Pfister, K. Aihara, and W. Gerstner, *Neural Comput.* **19**, 639 (2007).
 [20] M. E. Hasselmo, *Neural Networks* **7**, 13 (1994).
 [21] A. Greenstein-Messica and E. Ruppin, *Neural Comput.* **10**, 451 (2011).
 [22] E. Bullmore and O. Sporns, *Nat. Rev. Neurosci.* **10**, 186 (2009).
 [23] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, *Phys. Rep.* **424**, 175 (2006).
 [24] M. E. J. Newman, *Proc. Natl. Acad. Sci. USA* **103**, 8577 (2006).
 [25] The loop product $w_{i,j}w_{j,k}w_{k,i}$ is an attempt to generalize the curl operator, $\nabla \times$, in metric spaces to a network. The condition (13) is then equivalent to the curl-free condition of the “vector field” $v_{i,j} \equiv \ln(w_{i,j}/w^*)$, which seems to hold in the hub conformations but not the loop or path conformations for the stationary plastic networks.
 [26] G.-Q. Bi and M.-M. Poo, *J. Neurosci.* **18**, 10464 (1998).
 [27] R. Segev, Y. Shapira, M. Benveniste, and E. Ben-Jacob, *Phys. Rev. E* **64**, 011920 (2001).
 [28] O. Shefi, I. Golding, R. Segev, E. Ben-Jacob, and A. Ayali, *Phys. Rev. E* **66**, 021905 (2002).
 [29] P.-Y. Lai, L. C. Jia, and C. K. Chan, *Phys. Rev. E* **73**, 051906 (2006).
 [30] A. Barrat, M. Barthelemy, R. Pastor-Satorras, and A. Vespignani, *Proc. Natl. Acad. Sci. USA* **101**, 3747 (2004).
 [31] T. Masquelier, R. Guyonneau, and S. J. Thorpe, *PLoS ONE* **3**, e1377 (2008).
 [32] T. Masquelier, E. Hugues, G. Deco, and S. J. Thorpe, *J. Neurosci.* **29**, 13484 (2009).
 [33] S. Scarpetta, A. de Candia, and F. Giacco, *Frontiers in Synaptic Neuroscience* **2**, 32 (2010).