J Comput Neurosci (2006) 20:321–348
DOI 10.1007/s10827-006-7188-9

On the sensitive dependence on initial conditions of the dynamics of networks of spiking neurons

Arunava Banerjee

Received: 3 January 2005 / Revised: 19 December 2005 / Accepted: 4 January 2006 / Published online: 22 April 2006

Abstract We have previously formulated an abstract dynamical system for networks of spiking neurons and derived a formal result that identifies the criterion for its dynamics, without inputs, to be "sensitive to initial conditions". Since formal results are applicable only to the extent to which their assumptions are valid, we begin this article by demonstrating that the assumptions are indeed reasonable for a wide range of networks, particularly those that lack overarching structure. A notable aspect of the criterion is the finding that sensitivity does not necessarily arise from randomness of connectivity or of connection strengths, in networks. The criterion guides us to cases that decouple these aspects: we present two instructive examples of networks, one with random connectivity and connection strengths, yet whose dynamics is insensitive, and another with structured connectivity and connection strengths, yet whose dynamics is sensitive. We then argue based on the criterion and the gross electrophysiology of the cortex that the dynamics of cortical networks ought to be almost surely sensitive under conditions typically found there. We supplement this with two examples of networks modeling cortical columns with widely differing qualitative dynamics, yet with both exhibiting sensitive dependence. Next, we use the criterion to construct a network that undergoes bifurcation from sensitive dynamics to insensitive dynamics when the value of a control parameter is varied. Finally, we extend the formal result to networks driven by stationary input spike trains, deriving a superior criterion than previously reported.

Keywords Dynamical systems · Sensitive dependence · Spiking neurons

Action Editor: John Rinzel

A. Banerjee (✉)
Computer and Information Science and Engineering Department, University of Florida, Gainesville, FL 32611-6120,
e-mail: arunava@cise.ufl.edu

1. Introduction

It is well-documented that neurons in the cortex generate highly irregular spike trains that are generally not reproducible under repeated presentations of identical stimuli to the experimental subject (Tomko and Crapper, 1974; Burns and Webb, 1976; Tolhurst et al., 1983; Snowden et al., 1992; Britten et al., 1993). This has led to the widely held belief that the relevant information regarding a stimulus is contained in the time varying discharge rate of a cortical neuron and not in the particular spike train generated by it. There can however, be several mutually non-exclusive reasons for this phenomenon. First, the variability in the spike train generated by a neuron may simply be a reflection of the trial-to-trial variability in the input received by it from other neurons in the brain, a possibility that can not be ruled out even under the strictest of experimental regimen.

Second, the cortical neuron may be a highly noisy device, generating variable spike trains despite receiving identical inputs at its synapses. The extent to which this is the case has not entirely been resolved. On the one hand, current injection experiments mimicking post-synaptic potentials in a neuron have demonstrated that the spike generating mechanism of a neuron has low intrinsic noise, capable of producing spike trains that are reproducible to within a msec (Mainen and Sejnowski, 1995; Nowak et al., 1997). It has also been shown that spikes, once generated at the soma, rarely fail to arrive at the synapses on the axonal arbors of

a neuron (Cox et al., 2000). On the other hand, it has been demonstrated that synaptic transmission in neurons is highly unreliable with failure rates approaching 80% (Raastad et al., 1992; Rosenmund et al., 1993; Hessler et al., 1993). It is unclear whether this transmission failure is determined solely by thermal noise or is largely a deterministic process driven by, among other elements, the history of the pre- and post-synaptic neuron's spike activity.

Finally, the dynamics of cortical networks may be sensitive to initial conditions, where tiny fluctuations in spike timing caused by even the smallest amount of thermal noise result in successively larger fluctuations in the timing of subsequent spikes. It is this third possibility that is the subject of investigation of this article.

Sensitive dependence on initial conditions is not only an issue that is vital to the understanding of the dynamics of cortical networks, but is also a feature that can place significant constraints on the computational nature of such systems. First, if the dynamics of cortical networks is indeed sensitive to initial conditions, one can rule out the possibility that information is represented in the precise spatio-temporal pattern of spikes generated by the system, because such a representation could not be robust. We would then be required to refashion our search for the appropriate equivalence class of spike patterns that could robustly represent information. Second, and of particular relevance to experimental neuroscience, sensitive dependence on initial conditions can have a major impact on the topology of any potential attractors that might exist in the phase-space of cortical networks. Dynamical signatures that distinguish between attractors are likely to be more subtle if the attractors are chaotic as opposed to periodic/quasi-periodic. Insight into the topological characteristics of the attractors could therefore provide a principled basis for evaluating the efficacy of such widely used signatures as the average spike rate of neurons.

That sensitive dependence may play a role in the dynamics of cortical networks has long been suspected, both on experimental (Freeman and Skarda, 1985) as well as theoretical grounds (van Vreeswijk and Sompolinsky, 1996). Distinguishing chaotic dynamics experimentally from noisy periodic/quasi-periodic dynamics in a high dimensional dynamical system such as the cortex, is however known to be a notoriously difficult task (Grassberger and Procaccia, 1983; Osborne and Provenzale, 1989; Cazelles and Ferriere, 1992). Moreover, the theoretical result of (van Vreeswijk and Sompolinsky, 1996, 1998) is based on a highly simplified model of the neuron (namely, a binary stochastic unit) and network, and therefore, the extent to which the result applies to networks of neurons in the cortex remains unclear.

The vast majority of neurons in the brain communicate with one another using spikes. This core characteristic is often abstracted away in formal analyses of the dynamics of networks of neurons, where the inputs to and outputs from the neurons are modeled using continuous quantities representing variously, the firing rate or the instantaneous probability of spiking of a neuron (Amit and Brunel, 1997; Brunel and Hakim, 1999; Brunel, 2000; Latham et al., 2000; Seung et al., 2000). While these analyses have yielded important insights into global aspects of the dynamics of networks of neurons, the noted abstraction makes this approach unsuitable for the study of local characteristics of the particular spike trajectories that are generated by such systems. A second class of analyses, based on temporal positions of spikes in a Spike-response model (Gerstner and van Hemmen, 1992) of the neuron, does not suffer from this drawback. However, due to the lack of a full fledged representation of any arbitrary spike trajectory that could be generated by a network, the results in this case have been limited to phase-locked periodic solutions (Gerstner et al., 1996; Chow, 1998).

We have previously formulated an abstract dynamical system for networks of neurons where the spiking nature of the neuron assumes center-stage, and which has the capability of representing any spike trajectory generated by a network (Banerjee, 2001a). The system is based on a limited set of realistic assumptions and hence accommodates a wide range of neuronal models. We briefly review the abstract dynamical system in Section 2.1.

The strictly spike-based representation of the dynamics of a network of neurons lends itself to a precise determination of whether or not a spike trajectory generated by a given network is sensitive to initial conditions. Section 2.2 presents a formal description of this analysis. Briefly, the determination is based on a perturbation analysis of the spike trajectory. If the spikes on the spike trajectory at a given point in time, are perturbed, this will cause subsequent spikes on the trajectory to also be perturbed. Stated informally, the dynamics of the network is sensitive (respectively, insensitive) to initial conditions if these subsequent perturbations tend to grow progressively larger (respectively, smaller) with time. The first stage of the analysis shows how the perturbation in a newly generated spike can be computed as a function of the perturbations on those spikes that contributed to its generation. The second stage then cascades the perturbations on successively generated spikes to reveal the functional relationship between an initial set of perturbations on spikes, and a final set of perturbations on spikes generated in a distant future. Given any particular network of neurons and a corresponding spike trajectory, one can determine whether or not that trajectory is sensitive to initial conditions based on an inspection of this relationship.

This result, although an advance, is not particularly useful since we are seldom faced with the task of determining whether or not a spike trajectory generated by a particular network of neurons is sensitive. Our interest lies, instead, in the identification of those characteristics of generic networks of neurons that determine the sensitivity of their

spike trajectories. In order to address this more general question, we must first model appropriate aspects of the class of spike trajectories that are generated by such networks. In the final stage of the analysis presented in Section 2.2, the above noted cascade is modeled as a particular stationary stochastic process. Underlying the proof of the criterion for the dynamics of a network without inputs to be sensitive to initial conditions, as well as the criterion for it to be insensitive to initial conditions, in Banerjee (2001b), is this model.

Clearly, the result in Banerjee (2001b) is applicable only to the extent to which the stationary process succeeds in emulating the corresponding aspects of the real spike trajectories that are generated by networks of neurons. This issue is addressed in Section 2.3. First, the assumptions underlying the stationary stochastic process that models the cascade are described in detail. It is then demonstrated that these theoretical assumptions are indeed reasonable for a wide range of networks, particularly those that lack overarching structure.

A notable aspect of the criterion is the finding that sensitivity does not necessarily arise from randomness of connectivity or of connection strengths, in networks. The criterion guides us to cases that decouple these aspects. In Section 3 we present two instructive examples of networks, one with random connectivity and connection strengths, yet whose dynamics is insensitive to initial conditions, and another with structured connectivity and connection strengths, yet whose dynamics is sensitive to initial conditions. We then argue based on the criterion and features of the gross electrophysiology of the cortex that the dynamics of cortical networks, under normal operating conditions, ought to be almost surely sensitive to initial conditions. We supplement this with two additional examples of networks modeling cortical columns with widely differing qualitative dynamics, yet with both exhibiting sensitive dependence on initial conditions.

The criterion offers a succinct description of the conditions under which the dynamics of a network of spiking neurons should be almost surely sensitive, or almost surely insensitive, to initial conditions. Insight into the criterion enables us to construct networks whose dynamics is either sensitive or insensitive depending upon the value of a control parameter. In Section 4, we present example networks that undergo bifurcation from sensitive dynamics to insensitive dynamics as the control parameter is varied.

By considering networks of neurons that do not receive external inputs, we have been able to isolate the characteristic features of the intrinsic dynamics of the system from the confounding effects of inputs. Cortical networks are however, continually bombarded by inputs arriving from the thalamus and other cortical areas. The analysis of systems receiving external inputs suffers additional technical complications, such as the concept of sensitive dependence not being well-defined in all scenarios. In Section 5, we extend the formal result to the special case of networks receiving stationary input spike trains. We derive a criterion for the dynamics of such networks to be insensitive to initial conditions, and improve on the criterion reported in Banerjee (2001b) for the dynamics of such networks to be sensitive to initial conditions. Finally in Section 6, we discuss the ramifications of these results.

2. Sensitive dependence on initial conditions

We begin this section with a brief review of the abstract dynamical system that models recurrent networks of spiking neurons as formulated in Banerjee (2001a). Following this, we present the criterion for the dynamics of the system to be sensitive to initial conditions as derived in Banerjee (2001b), describing in greater detail the theoretical assumptions that underlie the formal result. We then investigate these theoretical assumptions and demonstrate that they are indeed reasonable for a wide range of network connectivities and connection strengths.

2.1. The abstract dynamical system

A neuron, at the highest level of abstraction, is a device that transforms multiple input sequences of spikes arriving at its various afferent (incoming) synapses into an output sequence of spikes on its axon. This transformation is realized by way of a quantity P, the membrane potential at the soma of the neuron. Post-synaptic potentials (PSPs) generated by the afferent spikes that have arrived at the various synapses of the neuron, and after-hyperpolarizing potentials (AHPs) generated by the efferent (outgoing) spikes that have departed the soma of the neuron, combine nonlinearly to generate P. The neuron emits a new spike whenever P, during its rising phase, crosses the threshold $\mathcal{T}$ of the neuron.

The formulation of the abstract dynamical system is based on the fundamental assumption that the membrane potential, P, can be derived from a bounded-time past history of all afferent and efferent spikes of a neuron. Spikes that have aged past this bound are considered to have negligible effect on the present value of P. This assumption is valid if the neuron is a finite precision device with fading memory.[1]

[1] We note models of the neuron in which PSP/AHPs decay exponentially fast to the resting level, such as the standard integrate-and-fire and spike-response models, belong to this category if one also assumes that the threshold is noisy, however small that noise might be. In such models, one can compute the present membrane potential from a record of the history of spikes from a bounded past. The above model is in fact, more general than has been described, since P does not necessarily have to be the membrane potential. It could be any quantity that denotes the instantaneous distance of the internal state of the neuron from the state that corresponds to being at threshold and generating a spike.

Since in any given system of neurons the afferent spikes into a particular neuron correspond to the efferent spikes of particular other neurons with appropriate axonal and synaptic delays, the state of a system of neurons can be specified by enumerating the temporal positions of all spikes generated by the neurons in the system over a bounded past. Such a record specifies the exact location of all spikes that are still situated on the axons of neurons, and via the functions P's specifies the current membrane potentials at the somas of the neurons. In effect, one avoids the technical difficulties that arise from attempting to represent such disparate quantities as membrane potentials and spikes under a common framework, by adopting a strictly spike based representation. Finally, since neurons can not generate successive spikes closer than their respective absolute refractory periods, there can only be finitely many spikes generated by each neuron during this bounded past.

Formally, let $i = 1, \ldots, \mathcal{S}$ denote the internal neurons of a system, and $i = \mathcal{S}+1, \ldots, \mathcal{R}$ the input neurons. Let Υ denote the duration of the bounded past as described above and r_i the absolute refractory period of neuron i. Then, neuron i will have generated at most $n_i = \lceil \Upsilon / r_i \rceil$ spikes during the past Υ time period. Let x_i^j $(1 \leq j \leq n_i)$ represent the time since the generation of a distinct spike by neuron i during the past Υ time period. The state of the system can then be specified by the $(\sum_{i=1}^{\mathcal{R}} n_i)$-tuple $\langle x_1^1, \ldots, x_1^{n_1}, \ldots, x_{\mathcal{R}}^1, \ldots, x_{\mathcal{R}}^{n_{\mathcal{R}}} \rangle \in [0, \Upsilon]^{\sum_{i=1}^{\mathcal{R}} n_i}$. It is conceivable that there will be times when neuron i will not have generated n_i spikes during the past Υ time period, n_i being merely an upper bound on the number of such spikes. At such times the remaining x_i^j's are set to 0. To maintain uniformity of vocabulary, x_i^j's with non-zero values are considered to correspond to live spikes and x_i^j's with value set at zero to dead spikes. Note that dead spikes are merely artifacts of our representation and do not physically exist. They are necessitated by two conflicting demands: the need to formulate a phase-space of fixed dimensionality, and the fact that neurons can generate variable number of spikes over any fixed interval of time. To elaborate using an example, a neuron all of whose spikes are currently dead is merely one that has not generated any spikes during the past Υ time period.

Each internal neuron i $(i = 1, \ldots, \mathcal{S})$ is assigned a C^∞ (i.e., smooth) membrane potential function $P_i(x_1^1, \ldots, x_1^{n_1}, \ldots, x_{\mathcal{R}}^1, \ldots, x_{\mathcal{R}}^{n_{\mathcal{R}}})$ that maps the present state description to the instantaneous potential at the soma of neuron i. The particular instantiation of the set of functions $P_i : [0, \Upsilon]^{\sum_{i=1}^{\mathcal{R}} n_i} \to \mathbb{R}$ determines both the electrophysiological properties of the neurons as well as their connectivity and connection strengths in the network. Figure 1(a) presents a schematic diagram of a system comprised of three internal neurons (depicted in black) and two input neurons (depicted in gray). The present state of the system is specified by the positions of the spikes (solid lines) in the shaded region at $t = 0$, and the state of the system at a future time T is specified by the positions of the spikes (solid lines) in the shaded region at $t = T$.

The dynamics of the system unfolds as follows. At any given moment in time all non-zero x_i^j's (live spikes) grow at a constant rate of 1. If the value of an x_i^j reaches Υ, it is reset to 0 (i.e., the live spike is turned dead). If input neuron i generates a spike, or if $P_i(\cdot) = \mathcal{T}$ and $dP_i/dt \geq 0$ for internal neuron i, exactly one of its x_i^j's set at 0 (a dead spike) is chosen and set to grow at a constant rate of 1 (i.e., turned live).

Whereas the $(\sum_{i=1}^{\mathcal{R}} n_i)$-tuple $\langle x_1^1, ..., x_1^{n_1}, ..., x_{\mathcal{R}}^1, ..., x_{\mathcal{R}}^{n_{\mathcal{R}}} \rangle \in [0, \Upsilon]^{\sum_{i=1}^{\mathcal{R}} n_i}$ does specify the state of the system, this representation is fraught with redundancy. First, for any given i, j, the two positions $x_i^j = 0$ and $x_i^j = \Upsilon$ ought to be identified with each other, both denoting dead spikes. Second, all permutations of $\langle x_i^1, \ldots, x_i^{n_i} \rangle$ for any given i ought to be identified with one another because the order in which spikes are assigned to the x_i^j's for any given neuron is inconsequential. These redundancies engender corresponding difficulties in the determination of trajectories in the phase-space. The first causes a trajectory to be discontinuous at the death of a spike, i.e., when an x_i^j is reset from Υ to 0. The second causes a trajectory to not be uniquely defined at the birth of a spike since any x_i^j set at 0 for that particular neuron can be chosen to grow at a constant rate of 1.

Since the objective of this article is to investigate the local neighborhood of trajectories and not the trajectories themselves, we confine the remainder of the review to certain notable aspects of the correctly formulated phase-space, and refer the interested reader to Banerjee (2001a, b) for a more comprehensive exposition.

It was shown in Banerjee (2001a) that the above redundancies can be eliminated by applying two successive transformations to the state description $\langle x_1^1, \ldots, x_1^{n_1}, \ldots x_{\mathcal{R}}^1, \ldots, x_{\mathcal{R}}^{n_{\mathcal{R}}} \rangle$. The first transformation, described informally, maps each x_i^j to a complex number on the unit circle. For each neuron i, the second transformation then maps the resultant n_i complex numbers to the coefficients of the complex polynomial whose roots are identically those complex numbers. The phase-space resulting from the transformations (denoted by $\prod_{i=1}^{\mathcal{R}} {}^i\mathbb{L}_{n_i}$) was shown to be a compact manifold with boundaries. It was also shown that for each membrane potential function $P_i(\cdot)$ there exists a corresponding C^∞ function $\widetilde{P}_i : \prod_{j=1}^{\mathcal{R}} {}^j\mathbb{L}_{n_j} \to \mathbb{R}$ on the transformed phase-space, and that for each i the subset of $\prod_{j=1}^{\mathcal{R}} {}^j\mathbb{L}_{n_j}$ that satisfies $\widetilde{P}_i(\cdot) = \mathcal{T}$ and $d\widetilde{P}_i(\cdot)/dt \geq 0$, is a closed subset of a C^∞ regular submanifold of codimension 1, which we denote by P_i^I.

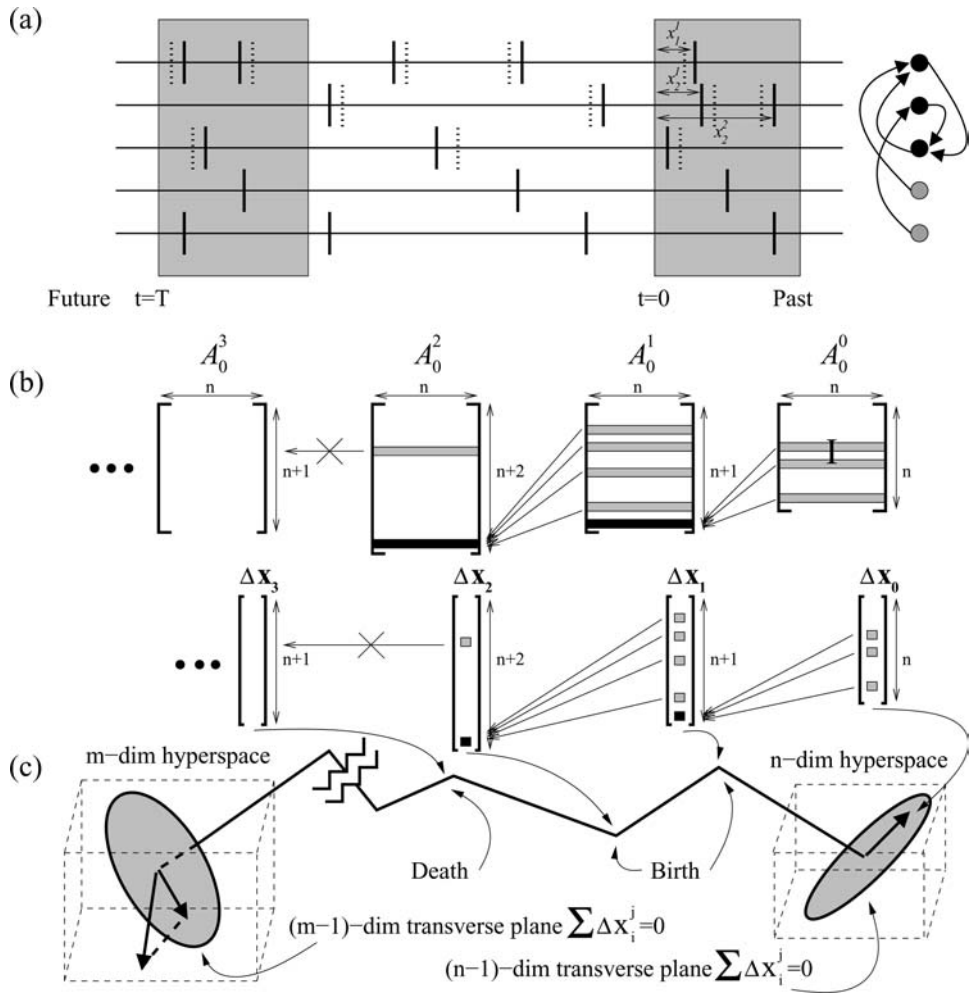

Fig. 1 (a) Schematic diagram of a system containing five neurons. Input neurons are depicted in gray and internal neurons in black. Directed edges denote the axons emanating from each neuron. Spikes are shown in solid lines. Gray boxes demarcate a bounded-time past history starting at time t. The temporal position of all spikes in the boxes specify the state of the system at $t = 0$ and $t = T$. The dynamics of the system may be visualized by translating the gray box at $t = 0$ to the left and noting the changes that take place within it. Perturbations on spikes are shown in dotted lines. Note that spikes generated by the input neurons are not perturbed. (b) The evolving perturbation vector $\Delta\mathbf{x_k}$ (lower section) and the corresponding perturbation matrix A_0^k (upper section). $\Delta\mathbf{x_k}$ gains a component according to Eq.(3) if the next event is a birth of a spike, and loses a component if the next event is a death of a spike. A_0^k correspondingly gains or loses a row. (c) Schematic diagram of a trajectory in the phase-space. The initial perturbation is constrained to lie transverse to the trajectory and the final perturbation is projected onto a transverse plane to the trajectory

The dichotomy between dead and live spikes described earlier introduces a corresponding structure on ${}^i\overline{\mathbb{L}}_{n_i}$. We denote by ${}^i\overline{\mathbb{L}}_{n_i}^j$ the subspace of ${}^i\overline{\mathbb{L}}_{n_i}$ that satisfies $\sigma_i \geq j$, where σ_i denotes the number of dead spikes in the state description of neuron i. In other words, ${}^i\overline{\mathbb{L}}_{n_i}^j$ denotes the subset of states of ${}^i\overline{\mathbb{L}}_{n_i}$ that have at least j dead spikes. It was shown in Banerjee (2001a) that for each j, ${}^i\overline{\mathbb{L}}_{n_i}^j$ is a C^∞ regular submanifold of ${}^i\overline{\mathbb{L}}_{n_i}$. Since, as noted earlier, dead spikes are merely artifacts of our representation, ${}^i\overline{\mathbb{L}}_{n_i}$ is assigned the topology generated by the family of all relatively open subsets of ${}^i\overline{\mathbb{L}}_{n_i}^j$, $\forall j \geq 0$.

The velocity field that generates the flow described earlier, can now be defined on the phase-space to complete the formulation of the abstract dynamical system. As shown in Banerjee (2001a), for systems without input neurons, a field $\mathcal{V}^1$ can be defined on the entire phase-space excepting the hypersurfaces P_i^I's, for when no neuron is on the verge of spiking and no spike is on the verge of death, and a field $\mathcal{V}^2$ can be defined on the P_i^I's for when one or more internal neurons are on the verge of spiking. Although no unique velocity field can be stipulated for systems with input neurons because the spike activity of the input neurons is not known a priori and is not determined by the state of the system, should an input trajectory be given, velocity fields for when one or more input neurons are also on the verge of spiking could be defined by appropriately modifying $\mathcal{V}^1$ and $\mathcal{V}^2$.

Since our objective is to investigate the local properties of trajectories through a perturbation analysis, the phase-space is additionally endowed with a Riemannian metric. The metric is chosen such that all flows corresponding to $\mathcal{V}^1$ on $\prod_{i=1}^{\mathcal{R}}({}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i}\setminus{}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i+1})$ (for all values of σ_i's) are measure preserving. We refer the interested reader to Banerjee, (2001b) for a detailed description of the Riemannian metric. It is sufficient for the purpose of this article to note that the metric likens the phase-space to the initial formulation at the local level, while preserving the global consistency afforded by the transformed representation.[2]

The advantage of imposing this particular metric lies in the fact that $\mathcal{V}^1$ is now a constant velocity field. As a result, the analysis of the local properties of a trajectory $\Psi_x(t)$ in the phase-space reduces to the analysis of a countably infinite sequence of discrete events each denoting the birth and/or death of one or more spikes. We shall hereafter refer to states in terms of their local representation $\langle x_1^1,\ldots,x_1^{n_1-\sigma_1},\ldots,x_{\mathcal{R}}^1,\ldots,x_{\mathcal{R}}^{n_{\mathcal{R}}-\sigma_{\mathcal{R}}}\rangle$. When it is necessary to define a local neighborhood on $\prod_{i=1}^{\mathcal{R}}{}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i}$ that also intersects with $\prod_{i=1}^{\mathcal{R}}{}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i^{-1}}$, as in the case of the analysis of the birth of a spike, we shall use $\langle x_1^1,\ldots,x_1^{n_1-(\sigma_1-1)},\ldots,x_{\mathcal{R}}^1,\ldots,x_{\mathcal{R}}^{n_{\mathcal{R}}-(\sigma_{\mathcal{R}}-1)}\rangle$ and set $x_i^1=0$ for all $i=1,\ldots,\mathcal{R}$, by convention. Finally, an infinitesimal perturbation on x_i^j will be denoted by Δx_i^j.

Before concluding this section, we note that even though the abstract dynamical system just described does have the capacity to model changing synaptic weights through changes in the functions P_i's, we do not consider such effects here. The non-stationary analysis necessary to incorporate synaptic weight update rules is beyond the scope of this article. However, we have conducted simulation experiments with the spike time dependent synaptic plasticity (STDP) rule (Markram et al., 1997; Bi and Poo, 1998; Froemke and Dan, 2002) and have found that the results do not dispute the conclusions of this article.

2.2. The criterion for sensitive dependence on initial conditions

We are now in a position to frame the question of sensitive dependence on initial conditions of trajectories in the phase-space of the abstract dynamical system. Consider the network of three internal neurons and two input neurons depicted in Fig. 1(a), initialized at the state described by the shaded region at $t=0$. As we progress in time, a new spike is generated by the second internal neuron. If the spikes of the internal neurons in the shaded region at $t=0$ were perturbed in time (dotted lines[3]), this would result in a perturbation on the new spike. This scenario would in turn repeat to produce further perturbations on future spikes generated by the internal neurons. Any initial set of perturbations would therefore propagate from spike to spike to produce a set of perturbations at any arbitrary future time $t=T$. Stated informally, the trajectory is sensitive (respectively, insensitive) to initial conditions if the successive perturbations tend to grow larger (respectively, smaller) as we progress in time.

We restrict the analysis in this section to systems that do not receive inputs, that is, to systems that do not contain input neurons. Systems containing input neurons are considered in Section 5. The analysis proceeds in stages. First, the impact of a birth of a spike and that of a death of a spike are each considered in isolation. These results are then used to deduce the cumulative impact of the sequence of births and deaths of spikes that comprise a trajectory.

Let $\langle\Delta x_1^2,\ldots,\Delta x_1^{n_1-(\sigma_1-1)},\ldots,\Delta x_{\mathcal{S}}^2,\ldots,\Delta x_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}\rangle$ be a perturbation on a trajectory (where the mapping to local coordinates is set such that $x_i^1=0$ for all $i=1,\ldots,\mathcal{S}$ and $x_i^j\neq 0$ for all $i=1,\ldots,\mathcal{S}$ and $j=2,\ldots,n_i-(\sigma_i-1)$) just prior to the birth of a spike (assigned without loss of generality to x_1^1) at neuron 1. Let the corresponding perturbation just past the birth be $\langle\Delta y_1^1,\Delta y_1^2,\ldots,\Delta y_1^{n_1-(\sigma_1-1)},\ldots,\Delta y_{\mathcal{S}}^2,\ldots,\Delta y_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}\rangle$. Consider the component Δy_1^1 on the new spike. Let $\langle 0,a_1^2,\ldots,a_1^{n_1-(\sigma_1-1)},\ldots,0,a_{\mathcal{S}}^2,\ldots,a_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}\rangle$ be the point on the trajectory such that

$$P_1\big(0,a_1^2,\ldots,a_1^{n_1-(\sigma_1-1)},\ldots,0,a_{\mathcal{S}}^2,\ldots,a_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}\big)=\mathcal{T}. \tag{1}$$

Δy_1^1 may then be computed by noting that on the perturbed trajectory

$$\begin{aligned}P_1\big(0,a_1^2&+\Delta x_1^2-\Delta y_1^1,\ldots,a_1^{n_1-(\sigma_1-1)}+\Delta x_1^{n_1-(\sigma_1-1)}\\&-\Delta y_1^1,\ldots,0,a_{\mathcal{S}}^2+\Delta x_{\mathcal{S}}^2-\Delta y_1^1,\ldots,a_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}\\&+\Delta x_{\mathcal{S}}^{n_{\mathcal{S}}-(\sigma_{\mathcal{S}}-1)}-\Delta y_1^1\big)=\mathcal{T}.\end{aligned} \tag{2}$$

The negative sign preceding Δy_1^1 relates to the fact that a positive Δy_1^1 corresponds to the new spike being generated earlier. A truncated Taylor expansion of (2) in conjunction with (1) then yields (3) below. All other components carry over as (4) below.

$$\Delta y_1^1=\sum_{i=1}^{\mathcal{S}}\sum_{j=2}^{n_i-(\sigma_i-1)}\big({}^1\alpha_i^j\times\Delta x_i^j\big)\quad\text{where}$$

[2] ${}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i}\setminus{}^i\overline{\mathbb{L}}_{n_i}^{\sigma_i+1}$) is mapped onto $(0,\Upsilon)^{n_i-\sigma_i}$ with its canonical basis deemed orthonormal.

[3] We consider the effect of a perturbation in the internal state of the system. Spikes generated by the input neurons are therefore not perturbed.

$$ {}^{1}\alpha_i^j = \frac{\partial P_1}{\partial x_i^j} \Bigg/ \left(\sum_{i=1}^{\mathcal{S}} \sum_{j=2}^{n_i-(\sigma_i-1)} \frac{\partial P_1}{\partial x_i^j} \right). \tag{3} $$

$$ \begin{aligned} \Delta y_i^j &= \Delta x_i^j \quad \text{for} \quad i = 1, \ldots, \mathcal{S}, \\ j &= 2, \ldots, n_i - (\sigma_i - 1). \end{aligned} \tag{4} $$

Since the $\partial P_1/\partial x_i^j$'s are evaluated at the instant of the generation of the new spike, i.e., when $P_1 = \mathcal{T}$ and $dP_1/dt \geq 0$, the denominator $\sum_{i,j}(\partial P_1/\partial x_i^j) = dP_1/dt \geq 0$. In general, we shall denote $(\partial P_l/\partial x_i^j)/\sum_{i,j}(\partial P_l/\partial x_i^j)$ by ${}^{l}\alpha_i^j$. Then, $\forall l \sum_{i,j} {}^{l}\alpha_i^j = 1$. Stated informally, the perturbation in a newly generated spike at neuron l can be represented as a weighted sum of the perturbations of those spikes in the state description that contribute to the generation of this new spike. The weight assigned to Δx_i^j is proportional to the value of $\partial P_l/\partial x_i^j$ at the instant of the generation of the new spike. The proportionality constant is set so that the weights sum to one.

Next, let $\langle \Delta x_1^1, \ldots, \Delta x_1^{n_1-\sigma_1}, \ldots, \Delta x_{\mathcal{S}}^1, \ldots, \Delta x_{\mathcal{S}}^{n_{\mathcal{S}}-\sigma_{\mathcal{S}}} \rangle$ be a perturbation on a trajectory (where the mapping to local coordinates is set such that $x_i^j \neq 0$ for all $i = 1, \ldots, \mathcal{S}$ and $j = 1, \ldots, n_i - \sigma_i$) just prior to the death of a spike (assigned without loss of generality to x_1^1) at neuron 1. Let $\langle \Delta y_1^1, \ldots, \Delta y_1^{n_1-\sigma_1}, \ldots, \Delta y_{\mathcal{S}}^1, \ldots, \Delta y_{\mathcal{S}}^{n_{\mathcal{S}}-\sigma_{\mathcal{S}}} \rangle$ be the corresponding perturbation just past the death. Then,

$$ \Delta y_1^1 = 0, \text{ and} \tag{5} $$

$$ \begin{aligned} \Delta y_i^j &= \Delta x_i^j \text{ for all } i = 1, ..., \mathcal{S},\ j = 1, ..., n_i - \sigma_i \text{ except} \\ i &= j = 1. \end{aligned} \tag{6} $$

The cumulative impact of the sequence of births and deaths of spikes that comprise a trajectory can now be computed as follows. Consider the section of the trajectory shown in Fig. 1(c) that corresponds to the evolution of the state of a network for a given interval of time. Let column vector $\Delta \mathbf{x_0}$ denote a perturbation on the trajectory at time $t = 0$ just prior to the first event (birth/death of a spike), and column vector $\Delta \mathbf{x_k}$ denote the corresponding perturbation just past the kth event (birth/death of a spike). As depicted in the lower section of Fig. 1(b), the perturbation vector gains and loses components with successive events, as specified by the equations derived above. Let A_0^k denote the matrix that embodies the propagation of the perturbation from event to event, such that $\Delta \mathbf{x_k} = A_0^k * \Delta \mathbf{x_0}$. Then, as depicted in the upper section of Fig. 1(b), the perturbation matrix A_0^k can be computed recursively as follows.

1. $A_0^0 = I$ (the $(n \times n)$ identity matrix, where n is the number of live spikes in the state at $t = 0$).
2. If the kth event corresponds to the birth of a spike at neuron l, then A_0^k is generated from A_0^{k-1} by identifying those rows r_i^j in A_0^{k-1} that correspond to spikes that contributed to the birth of the given spike (i.e., for each neuron i presynaptic to neuron l those spikes x_i^j that satisfy $\partial P_l/\partial x_i^j \neq 0$, as well as those spikes x_l^j of neuron l itself that satisfy $\partial P_l/\partial x_l^j \neq 0$), and introducing the new row $\sum_{i,j} {}^{l}\alpha_i^j r_i^j$ at an appropriate location into A_0^{k-1}.
3. If the kth event corresponds to the death of a spike, then A_0^k is generated from A_0^{k-1} by identifying that row r_i^j in A_0^{k-1} that corresponds to the given spike, and deleting it from A_0^{k-1}.

In effect, at any given stage, row r_i^j in A_0^k represents the component of the perturbation, Δx_i^j, of the corresponding spike x_i^j in the state description, as a function of the perturbation $\Delta \mathbf{x_0}$ on the initial set of spikes. The function is given by $\Delta x_i^j = r_i^j \cdot \Delta \mathbf{x_0}$, where $\cdot$ denotes the dot product.

Finally, we note that sensitive dependence on initial conditions is defined in terms of the divergence or convergence, not of states, but of trajectories. Hence, as depicted in Fig. 1(c), not only should the initial perturbation be constrained to lie on a plane transverse to the trajectory, but also, the rigid translational component of the final perturbation should be discarded to reveal the deviation transverse to the trajectory. Since the velocity of any point in the phase-space is of unit magnitude along each non-zero x_i^j, any transverse plane to the trajectory satisfies $\sum_{i,j} \Delta x_i^j = 0$.

Consider first an initial state description that has n live spikes. Let $\Delta \mathbf{x_0'}$ be any arbitrary $(n-1)$-dimensional perturbation vector. Then, $\Delta \mathbf{x_0} = C * \Delta \mathbf{x_0'}$ for the $n \times (n-1)$ matrix C described in (7), yields an n-dimensional perturbation vector $\Delta \mathbf{x_0}$ that lies on the transverse plane (i.e., satisfies $\sum_{i,j} \Delta x_i^j = 0$). In fact, any $n \times (n-1)$ matrix C of rank $(n-1)$, all of whose columns sum to 0, would suffice so long as the change in basis is accounted for. The C in (7) is based on the canonical basis.

Next, consider a final state description that has m live spikes. $\Delta \mathbf{x_k}$ is then an m-dimensional perturbation vector. Let $\Delta \mathbf{x_k'} = B * \Delta \mathbf{x_k}$ for the $(m-1) \times m$ matrix B described in (7). The operation corresponds to the deletion of the mean from each component in $\Delta \mathbf{x_k}$ followed by the discarding of a component, to yield $\Delta \mathbf{x_k'}$. The $(m-1)$-dimensional perturbation vector $\Delta \mathbf{x_k'}$ then lies on the transverse plane, based on the canonical basis.

For any given trajectory $\Psi_x(t)$ in the phase-space, $\Delta \mathbf{x_k'} = B * A_0^k * C * \Delta \mathbf{x_0'}$ for arbitrary value of k specifies the relation between an initial and a final perturbation, both of which lie on transverse sections of $\Psi_x(t)$. Assuming that A_0^k is an $(m \times n)$ matrix, the matrix of interest is then $B * A_0^k * C$, where B and C are the $(m-1) \times m$ and $n \times (n-1)$

matrices[4]:

$$B = \begin{bmatrix} 1-\frac{1}{m} & -\frac{1}{m} & \dots & -\frac{1}{m} & -\frac{1}{m} \\ -\frac{1}{m} & 1-\frac{1}{m} & \dots & -\frac{1}{m} & -\frac{1}{m} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ -\frac{1}{m} & -\frac{1}{m} & \dots & 1-\frac{1}{m} & -\frac{1}{m} \end{bmatrix}, \text{ and}$$

$$C = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ -1 & -1 & \dots & -1 \end{bmatrix}. \tag{7}$$

If $\lim_{k\to\infty} \|B * A_0^k * C\| = \infty$ (respectively, 0), then the trajectory is sensitive (respectively, insensitive) to initial conditions.[5] Given any particular network of neurons and a corresponding spike trajectory, one can therefore determine whether or not that spike trajectory is sensitive to initial conditions based on the above analysis.

This result, while not particularly useful in its own right, provides insight into how one may approach the broader question of which aspects of general networks of neurons play a crucial role in determining the sensitivity of their spike trajectories. The result demonstrates that it is the evolving matrix A_0^k, and not the spike trajectory as such, that lies at the heart of the matter. What is necessary to address the more general question is therefore an appropriate generative model for A_0^k. The upcoming discussion further elaborates on this distinction, following which the generative model for A_0^k is described.

The previous analysis envisions an investigator who has access to the dynamics of a particular system of neurons, from which she draws the data necessary to construct the evolving perturbation matrix A_0^k. At the birth of each new spike, she identifies those rows r_i^j's in the matrix that correspond to spikes that contributed to the generation of the new spike, and acquires their corresponding ${}^l\alpha_i^j$'s from the system. She then generates the new row $\sum_{i,j} {}^l\alpha_i^j r_i^j$ and inserts it into the matrix. At the death of each spike, she identifies the row r_i^j in the matrix that corresponds to that spike, and deletes it from the matrix.

Consider in contrast, a scenario where the investigator lacks access to the dynamics of the system of neurons, but is granted exactly the data necessary to construct the evolving perturbation matrix A_0^k. Here, at each step she is notified either of the birth or the death of a spike in the system. In the case of the birth of a spike, she is cited those rows in the matrix that were involved in the generation of the new spike, and is provided their corresponding ${}^l\alpha_i^j$'s. In the case of the death of a spike, she is simply cited the corresponding row in the matrix. Clearly, this is sufficient information for her to construct the evolving perturbation matrix A_0^k and draw conclusions regarding the sensitivity of the underlying trajectory. The investigator can, however, infer little else regarding the system of neurons and its dynamics. She can not even determine the number of neurons in the system, let alone their connectivity. In essence, the local neighborhood information that she receives regarding the trajectory could have been generated by the dynamics of any number of underlying systems of neurons. This observation can be stated formally as follows.

To every trajectory $\Psi_x(t)$ in a given phase-space, there corresponds a unique sequence of births (with contributing spikes and corresponding ${}^l\alpha_i^j$'s) and deaths of spikes that embodies its local neighborhood. Any two trajectories $\Psi_x(t)$ and $\Psi_y(t)$ that have identical such sequences are indistinguishable in their local neighborhoods, regardless of whether the trajectories lay in different sections of the same phase-space or in entirely different phase-spaces corresponding to distinct networks of neurons. Since the question of sensitivity depends on the nature of this sequence and not on the trajectory per se, the analysis can be based on such sequences rather than on trajectories.

The issue here is that there is no a priori guarantee that $\|B * A_0^k * C\|$ will converge for any arbitrary countably infinite sequence of births (with contributing spikes and corresponding ${}^l\alpha_i^j$'s) and deaths of spikes. However, for the class of local neighborhoods of trajectories whose corresponding sequences can be modeled by a stationary stochastic process, $\|B * A_0^k * C\|$ is guaranteed to converge to one of 0, 1, or ∞ for almost all (with respect to the stationary measure) trajectories owing to the subadditive ergodic theorem of Kingman (1973).

For the generative model, we shall therefore consider the stationary process specified by the "directly given" dynamical system $(\Omega, \mathcal{F}, \mu, T)$, defined as follows (see (Gray, 1988)). Briefly, each outcome in the sample space Ω shall instantiate an entire sequence of births (with contributing spikes and corresponding ${}^l\alpha_i^j$'s) and deaths of spikes, beginning with an initial number of live spikes. By assigning probabilities (according to the stationary measure μ) to sets of such outcomes (i.e., members of $\mathcal{F}$) we shall turn the generation of A_0^k (steps 1,

[4] Note that the dimensionality of B and C are determined by the dimensionality of A_0^k. Moreover, the values of the elements of B depend on the dimensionality of B. We shall continue to refer to these matrices as B and C although m and n might change depending on the context of their use.

[5] By virtue of the nature of these limiting values, all norms yield identical results. We use the Frobenius norm ($\|A\|_F = \sqrt{\text{Trace}(A^T * A)}$) in the analysis and simulations.

2, and 3 presented earlier) into a stationary stochastic process.

Formally, an outcome $\omega \in \Omega$ shall be an initial number of live spikes followed by a sequence of births (with contributing spikes and corresponding $\partial P_l/\partial x_i^j$'s) and deaths of spikes. Furthermore, we shall restrict Ω to those outcomes that have a bounded number (for appropriate choices of a lower and upper bound) of live spikes at all times. $\mathcal{F}$, the σ-algebra of subsets of Ω, shall be defined as the σ-algebra generated by the set of the following "basic events": each of the finitely many choices for an initial number of live spikes, each of the finitely many choices for the death of a spike at any given stage, and each of the finitely many choices for spikes to contribute to the generation of a new spike at any given stage, with each choice associated with a corresponding Borel σ-algebra for the ways that $\partial P_l/\partial x_i^j$'s can be assigned to those spikes. μ shall be a stationary probability measure, and $T : \Omega \to \Omega$ shall be the measure preserving (i.e., $\mu(T^{-1}(F)) = \mu(F)$ for all $F \in \mathcal{F}$) left shift transform.[6]

The sequence of random variables, $\log \|B * A_0^k(\omega) * C\|$ for $k = 0, 1, \ldots$, is then subadditive because $\|B * A_0^k * C\| = \|B * A_i^k * C * B * A_0^i * C\| \leq \|B * A_i^k * C\| \times \|B * A_0^i * C\|$. Here, A_0^i denotes the perturbation matrix for the segment, of the sequence specified by ω, beginning just prior to the first event and ending just past the ith event. Likewise, A_i^k denotes the perturbation matrix for the segment, of the sequence specified by ω, beginning just past the ith event and ending just past the kth event. Let $(\mathbf{1})$ denote the matrix all of whose elements are 1. The equality above follows from the fact that $C * B = I - \frac{1}{s}(\mathbf{1})$ where $C * B$ is an $(s \times s)$ matrix, and that $\forall l \sum_{i,j} {}^l\alpha_i^j = 1$. Consequently, based on the subadditive ergodic theorem,

$$Pr\left(\limsup_k \left\|B * A_0^k * C\right\| = \liminf_k \left\|B * A_0^k * C\right\|\right) = 1 \tag{8}$$

where $Pr(\cdot)$ denotes probability measure. We shall exploit (8) in the proof of the theorem in Section 5.

We assume hereafter that the trajectory under consideration lies in a non-trivial attractor (i.e., not the quiescent state) of a network of neurons that supports such a stationary measure. The following theorem was proved in Banerjee (2001b) with an additional set of assumptions regarding the stationary process. We consider these assumptions next. We do not repeat the proof here since we shall extend the theorem along similar lines in Section 5.

Theorem 1 *(Sensitivity without input). Let $\Psi_x(t)$ be a trajectory that is not drawn into the trivial fixed point of quiescence in a system that does not contain input neurons. Let $E(\sum_{i,j}({}^l\alpha_i^j)^2) = 1 + \delta < \infty$, where the expected value $E(\cdot)$ is taken over the set of all births of spikes in $\Psi_x(t)$. Then, if $\delta > \frac{2}{M_{low}}$ (respectively, $\delta < \frac{2}{M_{high}}$), $\Psi_x(t)$ is, with probability 1, sensitive (respectively, insensitive) to initial conditions. M_{low} and M_{high} denote the minimum and the maximum number of live spikes in $\Psi_x(t)$ across all time.*

2.3. The assumptions underlying the theorem and their validation

We demonstrated in the previous section that it is the nature of the sequence of births (with contributing spikes and corresponding ${}^l\alpha_i^j$'s) and deaths of spikes which embodies the local neighborhood of a trajectory, that determines whether or not the trajectory is sensitive to initial conditions. We subsequently described a generative model $(\Omega, \mathcal{F}, \mu, T)$ for such sequences and highlighted the fact that at a bare minimum, μ must be stationary in order for the notion of sensitivity to be well-defined almost everywhere. In this section, we further specify μ; the proof of the sensitivity theorem from the previous section is based on sequences generated by the following parameterized stationary process, with parameters M_{low}, M_{high},$\mathcal{D}_1$, P_{low}, P_{high},$\mathcal{D}_2$, and $\mathcal{D}_3$ reflecting the global aspects of the trajectory in the phase-space under consideration, as described below. How well the stationary process models the local neighborhoods of real spike trajectories is considered following this description. We show through numerical simulations of networks of spiking neurons that the assumptions underlying the process are indeed reasonable for a wide variety of network architectures and connection strengths.

The process is begun by choosing an integer n from a fixed distribution $\mathcal{D}_1$ over the range $[M_{low}, M_{high}]$, and constructing the $(n \times n)$ identity matrix $A_0^0 = I$. Parameters M_{low}, M_{high}, and $\mathcal{D}_1$ denote the minimum, the maximum, and the stationary distribution, of the number of live spikes in the trajectory under consideration across all time. At each step, the choice between the birth of a spike and the death of a spike is made from a stationary stochastic process, depending solely on the history of the number of live spikes in the system such that the stationary distribution of the resultant number of live spikes in the trajectory matches $\mathcal{D}_1$. The choice is therefore independent of the elements of A_0^{k-1}.

[6] To be more specific, if outcome $\omega_0 = n_0 D(a_1)D(a_2)B(b_1\rho_1 \ldots b_p\rho_p)D(a_3)\ldots$, then $T(\omega_0) = \omega_1$, where $\omega_1 = (n_0 - 1)D(a_2)B(b_1\rho_1 \ldots b_p\rho_p)D(a_3)\ldots$. In this representation of the outcomes, the leading positive integer n_0 is the initial number of live spikes in ω_0. $D(a_1)$ for $a_1 \in [1, n_0]$ denotes the death of the spike indexed by the positive integer a_1, and likewise for $D(a_2)$ and $D(a_3)$ for $a_2 \in [1, n_0 - 1]$ and $a_3 \in [1, n_0 - 1]$. $B(b_1\rho_1 \ldots b_p\rho_p)$ for $b_i \in [1, n_0 - 2]$ and $\rho_i \in \mathbb{R}$ denotes the birth of a spike with contributing spikes and their corresponding contributions specified by the positive integers b_i's (as with the a_i's) and the real numbers ρ_i's (the $\partial P_l/\partial x_i^j$'s).

Consider the state of the process just past the $(k-1)$th event. Let vector E_{k-1} be the expected value of the mean of the population of rows in A_0^{k-1}, i.e.,

$$E_{k-1} = E\left(\frac{1}{m}\sum_{i=1}^{m} v_i\right), \tag{9}$$

where $E(\cdot)$ denotes the expected value, m the number of rows in A_0^{k-1}, and v_i's the rows in A_0^{k-1}. Furthermore, let V_{k-1} and C_{k-1} be the scalar valued variance and covariance of the population of rows in A_0^{k-1}, defined as

$$V_{k-1} = E\left(\frac{1}{m}\sum_{i=1}^{m}(v_i - E_{k-1})\cdot(v_i - E_{k-1})\right), \text{ and} \tag{10}$$

$$C_{k-1} = E\left(\frac{1}{m(m-1)}\sum_{i=1,j=1,i\neq j}^{m,m}(v_i - E_{k-1})\cdot(v_j - E_{k-1})\right). \tag{11}$$

In the case of the birth of a spike, A_0^k is generated from A_0^{k-1} as follows. First, a $p \in [P_{low}, P_{high}]$ is chosen from a stationary process, the choice depending on the history of the number of live spikes in the system as well as the previous values of p, such that the stationary distribution of p matches $\mathcal{D}_2$. The choice is therefore independent of the elements of A_0^{k-1}. Parameters P_{low}, P_{high}, and $\mathcal{D}_2$ denote the minimum, the maximum, and the stationary distribution, of the number of spikes that can contribute to the generation of a spike in the trajectory across all time. Second, a random vector $\langle Y_1, \ldots, Y_p\rangle$ is chosen from a fixed exchangeable distribution $\mathcal{D}_3$, i.e., a distribution that satisfies $Pr(Y_1, \ldots, Y_p) = Pr(Y_{\sigma(1)}, \ldots, Y_{\sigma(p)})$ for any permutation σ[7]. $\mathcal{D}_3$ denotes the distribution of the $\partial P_l/\partial x_i^j$'s for the trajectory under consideration. Third, rows $v_1, \ldots, v_p$ are chosen from A_0^{k-1}, through a random sampling with replacement that is unbiased in the sense that the expected value, the variance, and the covariance of the samples match the population counterparts in A_0^{k-1}, i.e., for each sample v_i for $i = 1, ..., p$,

$$E(v_i) = E_{k-1}, \tag{12a}$$

$$V(v_i) = E((v_i - E_{k-1})\cdot(v_i - E_{k-1})) = V_{k-1}, \text{and} \tag{12b}$$

$$C(v_i, v_j)_{j\neq i} = E((v_i - E_{k-1})\cdot(v_j - E_{k-1}))_{j\neq i} = C_{k-1}, \tag{12c}$$

where v_j ranges over all the rows in A_0^{k-1} except the sample. Finally, the row $v_{\text{new}} = \sum_{i=1}^{p} y_i v_i$, where $y_i = Y_i / \sum_{i=1}^{p} Y_i$, is inserted into A_0^{k-1} to generate A_0^k. The random variables y_i's correspond to the ${}^l\alpha_i^j$'s at the birth of a spike. Note that the sampling with replacement makes it possible for any particular row to be chosen multiple times to contribute to the sum v_{new}. Physically, this would correspond to a scenario where a pre-synaptic neuron makes multiple synapses on a post-synaptic neuron, causing each pre-synaptic spike to contribute multiple PSPs towards the birth of a post-synaptic spike.

In the case of the death of a spike, A_0^k is generated from A_0^{k-1} as follows. A row v_{del} is chosen from A_0^{k-1}, through a random sampling that is unbiased in the sense that the expected value, the variance, and the covariance of the sample match the population counterparts in A_0^{k-1}, i.e.,

$$E(v_{\text{del}}) = E_{k-1}, \tag{13a}$$

$$V(v_{\text{del}}) = E((v_{\text{del}} - E_{k-1})\cdot(v_{\text{del}} - E_{k-1})) = V_{k-1}, \quad \text{and} \tag{13b}$$

$$C(v_{\text{del}}, v_j)_{j\neq \text{del}} = E((v_{\text{del}} - E_{k-1})\cdot(v_j - E_{k-1}))_{j\neq \text{del}} = C_{k-1}, \tag{13c}$$

where v_j ranges over all the rows in A_0^{k-1} except the sample. v_{del} is then deleted from A_0^{k-1} to generate A_0^k.

The assumptions underlying the process are therefore:

1 The choice between the birth and the death of a spike is *independent* of the elements of A_0^{k-1}.
2 In the case of the birth of a spike, p is chosen *independent* of the elements of A_0^{k-1}. Furthermore, $\langle y_1, \ldots, y_p\rangle$ is chosen *independent* of the rows $v_1, \ldots, v_p$ which are *unbiased* random samples with replacement of the rows in A_0^{k-1}.
3 In the case of the death of a spike, v_{del} is an *unbiased* random sample of the rows in A_0^{k-1}.

The crucial question to consider is whether these are reasonable assumptions to make for real trajectories that lie in non-trivial attractors (i.e., not the quiescent state) of networks of neurons. In other words, are the sequences of births (with contributing spikes and corresponding ${}^l\alpha_i^j$'s) and deaths of spikes generated by the attractor dynamics of networks of neurons modeled well, probabilistically speaking, by the outcomes of the stationary stochastic process?

We first note that the choice between the birth and the death of a spike at any given moment depends on the state of a system, i.e., on the location in the phase-space of the abstract dynamical system that models that particular network of neurons. Since this location information is not contained

[7] $Pr(\cdot)$ here denotes probability density. The proof in Banerjee (2001b) assumes i.i.d random variables $Y_1, \ldots, Y_p$. However, the proof holds for the more general case of exchangeable random variables without modification.

in the description of the local neighborhood of the trajectory as characterized by the elements of A_0^{k-1}, it is reasonable to assume that the choice between birth and death is independent of the elements of A_0^{k-1}. In the case of the birth of a spike, one can similarly argue that the choice of p, which corresponds to the number of spikes that contribute to the generation of the new spike, is likewise independent of the elements of A_0^{k-1}.

The remaining assumptions, that in the case of the birth of a spike $\langle y_1, \ldots, y_p \rangle$ is chosen independent of the $v_1, \ldots, v_p$'s which are in turn unbiased samples, and that in the case of the death of a spike v_{del} is an unbiased sample, are more complex and can not simply be argued for. To test the validity of these assumptions, we recorded the necessary data from simulations of a wide variety of network architectures and connection strengths. We report the results from these experiments in the remainder of this section.

All simulations were based on a model of the neuron akin to the Spike-response model (Gerstner and van Hemmen, 1992) in that the potential function $P(x_1^1, \ldots, x_1^{n_1}, \ldots, x_{\mathcal{R}}^1, \ldots, x_{\mathcal{R}}^{n_{\mathcal{R}}})$ of a neuron was modeled as $\sum_{i,j} {}_i P(x_i^j)$, that is, the sum of the excitatory and inhibitory PSPs induced by the arrival of spikes at the various synapses of the neuron, and the AHPs induced by the spikes that had departed the soma of the neuron. The benefit of this model lay in the ease with which the $\partial P_l / \partial x_i^j$'s for the various x_i^j's could be computed.[8] The time bound Υ defined earlier in this article was set at 500 msec. PSPs were modeled using the parameterized function (see (MacGregor and Lewis, 1977))

$$P(t) = \frac{Q}{d\sqrt{t}} e^{-\beta d^2/t} e^{-t/\tau}, \tag{14}$$

where Q denotes the connection strength, d denotes the distance (in dimensionless units) of the synapse from the soma, and β and τ control the rate of rise and fall of the PSP. AHPs were modeled using the function

$$P(t) = Re^{-t/\gamma}, \tag{15}$$

where R denotes the instantaneous fall in potential after a spike and γ controls its rate of recovery. Although the instantaneous fall of the AHP does not conform with the assumption of smoothness of P, the issue is of no significance since all functions were discretized in time for the purpose of simulation. As a result, the maximum gradient of the AHP was bounded from above and depended on the choice of the time step. The time step was chosen to be small enough so that spikes were not generated during the falling phase of the AHP.

Synaptic and axonal delays were combined and chosen from the realistic range [0.4, 0.9] msec. Four types of synapses, excitatory AMPA and NMDA and inhibitory GABA_A and GABA_B, were modeled. For AMPA synapses, β was set at 1.0, and τ was set at 20 msec when on excitatory and 10 msec when on inhibitory neurons. NMDA synapses were modeled with the simplifying assumption that they were free of Mg^{2+} blocking. In other words, the post-synaptic voltage dependence of NMDA synapses was not modeled. β for the NMDA synapses was set at 5.0 and τ at 80 msec. The resultant PSPs (for the same value of Q) were approximately a third smaller and had longer rise and fall times than those generated at the AMPA synapses. For GABA_A synapses, β was set at 1.1, and τ was set at 20 msec when on excitatory and 10 msec when on inhibitory neurons. The resultant PSPs had slightly longer rise times than those generated at the AMPA synapses. Finally, for GABA_B synapses β was set at 50.0 and τ at 100 msec. The parameter d that determined the distance of the synapse from the soma was chosen from the realistic range [1.0, 2.0]. Excitatory AMPA and NMDA synapses were considered to be co-located, whereas each inhibitory synapse was chosen to be either GABA_A or GABA_B. This choice was either made randomly or in a regular fashion (i.e., no GABA_B synapses) depending upon the architecture considered (see below). All AHPs were modeled to be identical, with R set at -1000.0 and γ at 1.2 msec. The threshold for all neurons was set at 1.0. The connection strengths, Q's, were either chosen randomly from the range [1.0, 10.0] or were set to fixed values depending upon the architecture considered. The Q's for GABA_A synapses on excitatory neurons were then increased to six times the magnitude. We found that this setting successfully thwarted runaway excitation in all the networks. Finally, the Q's at all synapses were scaled uniformly so that the average spike rate of the neurons in the network lay in a realistic range (i.e., did not exceed 20 Hz).

All experiments were conducted on networks composed of 1000 neurons. The networks did not contain any input neurons, and therefore all activity recorded corresponded to dynamics in non-quiescent attractors. Three broad categories of network architectures, with multiple sub-categories within each category, were investigated. The first was that of random networks, with 80% of the neurons chosen randomly to be excitatory and the rest inhibitory. Each neuron received inputs from 100 neurons chosen randomly. Two sub-categories were investigated—connection strengths, axonal and synaptic delays, as well as the distance of the synapses from the somas, were either chosen randomly from the ranges noted above or were set to constant values (delay of 0.8 msec, $d = 1.5$ for excitatory synapses and $d = 1.2$ for inhibitory synapses) across the entire network.

[8] The $\partial P_l / \partial x_i^j$ for each x_i^j could be computed independent of the value of the other x_i^j's.

The second category was that of clustered networks composed of 800 excitatory and 200 inhibitory neurons. Here, the neurons were grouped evenly into clusters with the connectivity within each cluster being all-to-all, and that between clusters being sparse. Several cluster sizes, ranging from 20 to 100, were investigated. As in the previous case, connection strengths, axonal and synaptic delays, as well as the distance of the synapses from the somas, were either chosen randomly from the ranges noted above or were set to constant values (delay of 0.8 msec, $d = 1.5$ for excitatory synapses and $d = 1.2$ for inhibitory synapses) across the entire network.

The third category was that of structured ring networks composed of 800 excitatory and 200 inhibitory neurons. Here, the excitatory and inhibitory neurons were interleaved evenly and placed around a ring. Each neuron received inputs from a sector of 100 contiguous neurons, the sector being centered on that neuron. Once again, connection strengths, axonal and synaptic delays, as well as the distance of the synapses from the somas, were either chosen randomly from the ranges noted above or were set to constant values (delay of 0.8 msec, $d = 1.5$ for excitatory synapses and $d = 1.0$ for inhibitory synapses) across the entire network.

Finally, multiple instantiations of each category/subcategory of network architectures were generated and their dynamics investigated. In each case, the network was probed for the existence of attractors besides the ever-present quiescent state attractor, by initializing it at various locations in the phase-space and recording the ensuing dynamics. In all cases where the initial state lay in the domain of attraction of the quiescent state, the dynamics of the network was found to converge rapidly to that state. In cases where the network generated sustained recurrent activity at moderate spike rates[9], suggesting dynamics in a non-trivial attractor situated in a realistic regime, we recorded the spike trajectory and gathered all information regarding the evolving perturbation matrix A_0^k.

At each birth of a spike, we recorded the rows, r_i^j's, in A_0^{k-1} that contributed to the generation of the new row, and their corresponding ${}^l\alpha_i^j$'s. We concurrently recorded the population mean, variance, and covariance of the rows in A_0^{k-1} (Eqs. (9)–(11)). At each death of a spike, we noted the discarded row r_i^j and concurrently recorded the population mean, variance and covariance of the rows in A_0^{k-1}. In order to assess how biased the sampling was at the birth of each spike, rather than evaluate each contributing row individually with respect to the noted population statistics, we first computed the sample mean, variance, and covariance of the set of contributing rows, and then evaluated these sample statistics with respect to the population statistics, across the sequence of successive births of spikes. This approach allowed us to assess the bias in the sampling, while evading the prospect of intractability that comparison to each individual contributing sample entailed. In order to assess how biased the sampling was at the death of each spike, we evaluated the discarded row with respect to the noted population statistics, across the sequence of successive deaths of spikes.

We generated joint probability histograms from the collected data for the sample statistic plotted against the corresponding population statistic. In all cases and for all the noted statistics, the probabilities were found to accumulate near the diagonal, implying that the assumption of unbiased sampling was indeed reasonable. We also found that for networks that lacked overarching structure, such as the random networks and the clustered networks with small cluster sizes, the probabilities were more closely concentrated on the diagonal. Subfigures (a) and (b) in Figs. 2–4 present, respectively, the joint distribution of the sample and population means for a fixed coordinate position of the r_i^j's (i.e., component ${}_pr_i^j$ of r_i^j, for a fixed coordinate position p) and the joint distribution of the sample and population variances, at the birth of spikes, for representative instances of the random, clustered, and ring networks, respectively. These examples were chosen to reflect the entire range of results obtained, from the best (the network in Fig. 2 that had all its parameters chosen randomly) to the worst (the network in Fig. 4 that had all its parameters set to constant values). The distributions for all the other statistics were similar.

In order to assess the dependence between the ${}^l\alpha_i^j$'s and the corresponding rows r_i^j's at the birth of spikes, we first decomposed each row r_i^j into its components, ${}_pr_i^j$'s, and then analyzed the dependence between the ${}^l\alpha_i^j$'s and the ${}_pr_i^j$'s, for fixed values of p. Note that ${}_pr_i^j$ represents the sensitivity of the spike x_i^j with respect to the pth spike in the initial state description. We generated the joint distribution $Pr({}^l\alpha_i^j, {}_pr_i^j)$ and compared it to the corresponding product of the marginals $Pr({}^l\alpha_i^j) * Pr({}_pr_i^j)$, for fixed values of p. Subfigures (c) and (d) in Figs. 2–4 present plots of $-\log(Pr({}^l\alpha_i^j, {}_pr_i^j))$ and $-\log(Pr({}^l\alpha_i^j) * Pr({}_pr_i^j))$, for fixed coordinate positions p's, from the noted instances of the random, clustered, and ring networks, respectively. The similarity of the plots attest to the near independence of the variables.

To further quantify the dependence between the ${}^l\alpha_i^j$'s and the components of the corresponding r_i^j's, we computed the mutual information between the two. Since the mutual information is the Kullback-Leibler divergence between $Pr({}^l\alpha_i^j, {}_pr_i^j)$ and $Pr({}^l\alpha_i^j) * Pr({}_pr_i^j)$, low values of mutual information would imply that the assumption of

[9] We found that this could be achieved in all the networks by simply scaling all Q's, as noted earlier.

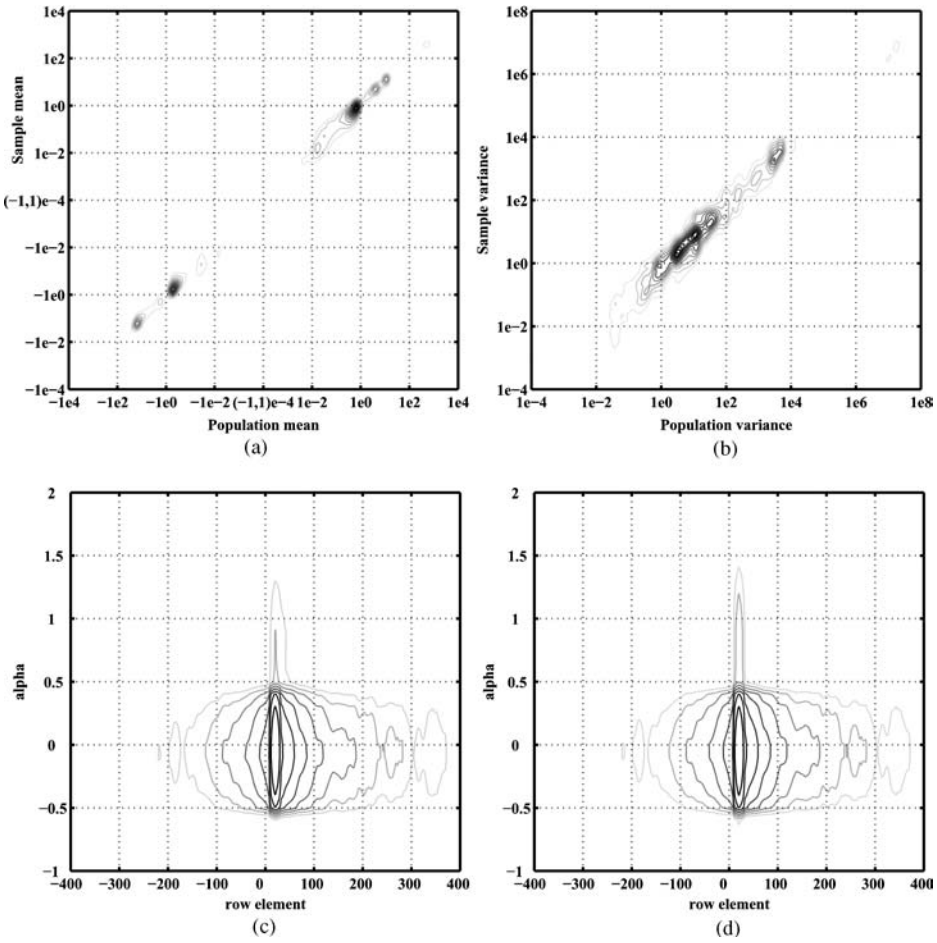

Fig. 2 Characteristic features of the data involved in the generation of A_0^k, distilled from the spike trajectory recorded from the instantiation of the random network. The probability histograms generated from the data are presented as contour plots. Darker contour lines represent higher values. (a) and (b): Distribution of the mean of a fixed coordinate position (subfigure (a)), and distribution of the variance (subfigure (b)), of the sample of rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against the corresponding mean and variance of all the rows in A_0^k. In both graphs, 15 contour lines span the range of probability values evenly. The lightest contour line therefore, denotes a value more than an order of magnitude smaller than the darkest contour line. (c) and (d): $-\log(Pr({}^l\alpha_i^j, {}_pr_i^j))$ (subfigure (c)) and $-\log(Pr({}^l\alpha_i^j) * Pr({}_pr_i^j))$ (subfigure(d)) for a fixed coordinate position p of the rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against their corresponding ${}^l\alpha_i^j$'s. In both graphs, 9 contour lines span the range $[10^{-2}, 10^{-8}]$ evenly in logarithmic scale

independence is reasonable. We found this to be the case in all our experiments. The values of the mutual information ranged from 0.0256 to 0.0503 bits for the random networks, from 0.0398 to 0.0782 bits for the clustered networks, and from 0.0590 to 0.1223 bits for the ring networks. By comparison, the entropy of the joint distributions (with the same discretization of the axes) ranged from 3.4626 to 7.5051 bits for the random networks, from 5.8434 to 9.6497 bits for the clustered networks, and from 3.8970 to 6.2328 bits for the ring networks.[10]

Based on all the above results, we can conclude that the assumptions underlying the stationary stochastic process (which in turn is the basis for the sensitivity theorem in Section 2.2) are indeed reasonable for a wide variety of network architectures and connection strengths.

3. Network architecture and sensitive dependence

An inspection of the formal criterion in the sensitivity theorem presented in Section 2.2 reveals that whether or not a trajectory in the phase-space of the abstract dynamical system (and hence, the spatio-temporal sequence of spikes generated by the corresponding network of neurons) is sensitive to initial conditions depends on the stationary distribution of the $\sum_{i,j}({}^l\alpha_i^j)^2$'s. The role that the connectivity and connection strengths of a network play in the determination of the sensitivity of the underlying spike trajectories is entirely indirect and through this quantity. In this section we present

[10] The entropy values do not make strict sense because they tend to infinity as the discretization is made finer. They were computed at the same discretization level as with the mutual information, and have been noted for mere comparison purposes.

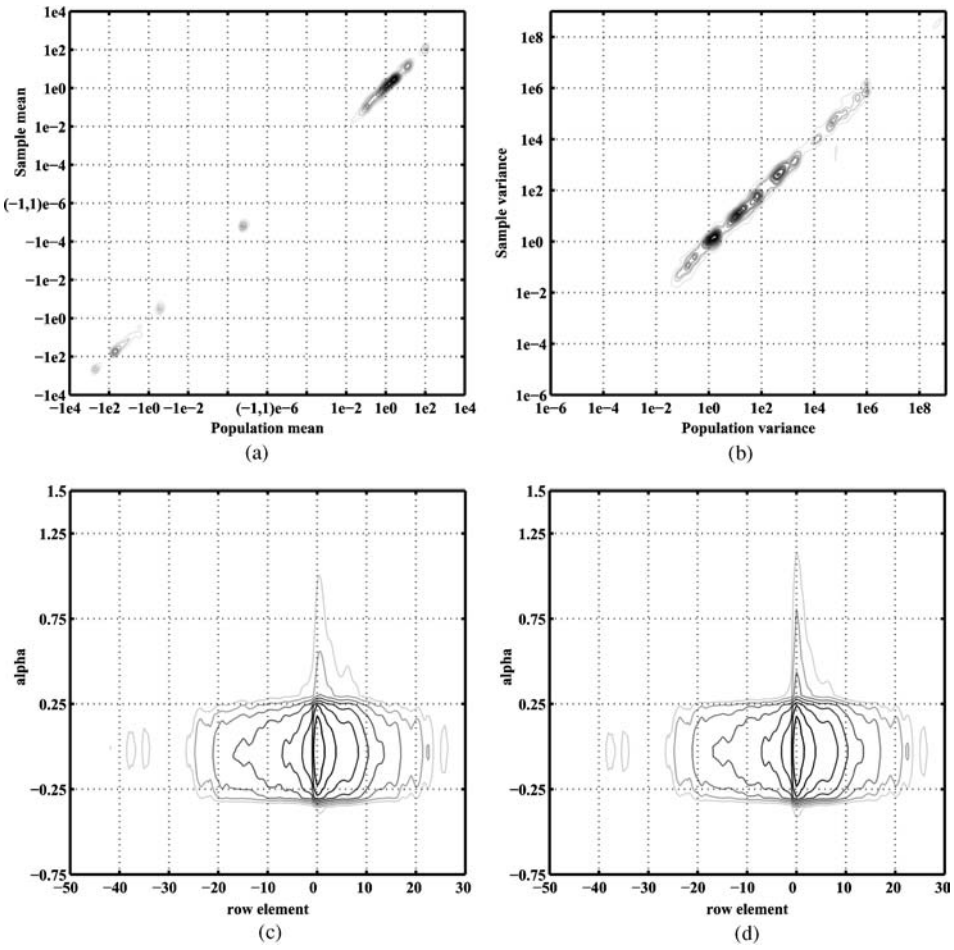

Fig. 3 Characteristic features of the data involved in the generation of A_0^k, distilled from the spike trajectory recorded from the instantiation of the clustered network. The probability histograms generated from the data are presented as contour plots. Darker contour lines represent higher values. (a) and (b): Distribution of the mean of a fixed coordinate position (subfigure (a)), and distribution of the variance (subfigure (b)), of the sample of rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against the corresponding mean and variance of all the rows in A_0^k. In both graphs, 15 contour lines span the range of probability values evenly. The lightest contour line therefore, denotes a value more than an order of magnitude smaller than the darkest contour line. (c) and (d): $-\log(Pr({}^l\alpha_i^j, {}_pr_i^j))$ (subfigure (c)) and $-\log(Pr({}^l\alpha_i^j) * Pr({}_pr_i^j))$ (subfigure(d)) for a fixed coordinate position p of the rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against their corresponding ${}^l\alpha_i^j$'s. In both graphs, 9 contour lines span the range $[10^{-2}, 10^{-8}]$ evenly in logarithmic scale

simulation results from instructive examples of networks that corroborate this observation.

Since the ${}^l\alpha_i^j$'s at the birth of a spike are constrained to satisfy $\sum_{i,j} {}^l\alpha_i^j = 1$, simple algebra demonstrates that $\sum_{i,j}({}^l\alpha_i^j)^2 < 1$ if all the ${}^l\alpha_i^j$'s are strictly positive. In addition, the value of $\sum_{i,j}({}^l\alpha_i^j)^2$ rises as the subset of the ${}^l\alpha_i^j$'s that are negative grows larger. If a spike is generated during the rising (respectively, falling) phase of the PSP/AHP of a contributing spike, then the ${}^l\alpha_i^j$ associated with that contributing spike is positive (respectively, negative). Since excitatory PSPs have short rising phases followed by prolonged falling phases, they are more likely to contribute negative ${}^l\alpha_i^j$'s. Conversely, since inhibitory PSPs have short falling phases followed by prolonged rising phases, they are more likely to contribute positive ${}^l\alpha_i^j$'s. Finally, since AHPs have instantaneous falling phases, their ${}^l\alpha_i^j$'s are strictly positive. These observations suggest that networks whose dynamics are dominated by inhibitory (respectively, excitatory) PSPs are less (respectively, more) likely to exhibit sensitive dependence on initial conditions, because at the instant of the generation of each new spike in the system a substantial portion of the ${}^l\alpha_i^j$'s are presumably positive (respectively, negative).

In particular, if all the neurons in a network are inhibitory, and in addition, all inhibitory PSPs have instantaneous falling phases followed by slow rises to the resting level as in the case of the standard integrate-and-fire model with synaptic inputs modeled as delta currents, one would be assured that $\sum_{i,j}({}^l\alpha_i^j)^2 < 1$ at every birth of a spike. The theorem would then predict that the dynamics of the network ought to be almost surely insensitive regardless of the connectivity and the connection strengths in the network.

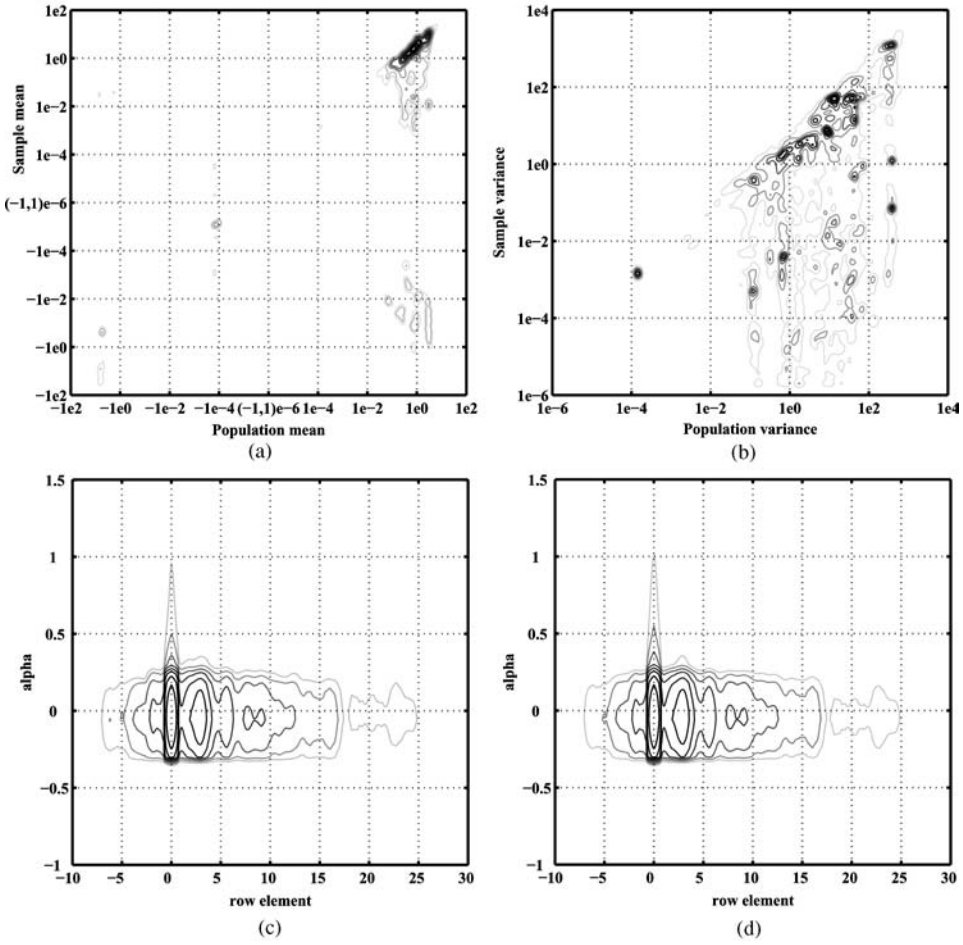

Fig. 4 Characteristic features of the data involved in the generation of A_0^k, distilled from the spike trajectory recorded from the instantiation of the highly structured ring network. The probability histograms generated from the data are presented as contour plots. Darker contour lines represent higher values. (a) and (b): Distribution of the mean of a fixed coordinate position (subfigure (a)), and distribution of the variance (subfigure (b)), of the sample of rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against the corresponding mean and variance of all the rows in A_0^k. In both graphs, 15 contour lines span the range of probability values evenly. The lightest contour line therefore, denotes a value more than an order of magnitude smaller than the darkest contour line. (c) and (d): $-\log(Pr({}^l\alpha_i^j, {}_pr_i^j))$ (subfigure (c)) and $-\log(Pr({}^l\alpha_i^j) * Pr({}_pr_i^j))$ (subfigure(d)) for a fixed coordinate position p of the rows, r_i^j's, drawn from A_0^k at the birth of each spike plotted against their corresponding ${}^l\alpha_i^j$'s. In both graphs, 9 contour lines span the range $[10^{-2}, 10^{-8}]$ evenly in logarithmic scale

In order to ascertain whether this prediction holds, we analyzed the dynamics of multiple instantiations of inhibitory networks with the previously described random, clustered, and structured ring architectures, and with connection strengths, axonal and synaptic delays, as well as the distance of the synapses from the somas either chosen randomly from the realistic ranges noted earlier or set to constant values across the entire network. In each case, we found that the trajectory was indeed insensitive to initial conditions. We present here the results from the most instructive case, that of a random inhibitory network with all its parameters chosen randomly, since it demonstrates that sensitivity does not necessarily arise from randomness of connectivity or connection strengths.

Figure 5 displays the results from the random inhibitory network. The system was composed of 1000 neurons with each neuron receiving synapses from 100 randomly chosen neurons. Axonal/synaptic delays were chosen randomly from a uniform distribution over [0.4, 0.9] msec. The inhibitory PSPs were modeled using

$$P(t) = -Qe^{-td/\tau} \quad \text{for } t \geq 0 \text{ , and, } 0 \text{ otherwise.} \tag{16}$$

Q, the strength of a synapse, was chosen randomly from a uniform distribution over [1.0, 6.0], d, the distance of a synapse from the soma, was chosen randomly from a uniform distribution over [1.0, 2.0], and τ was set at 20 msec. AHPs

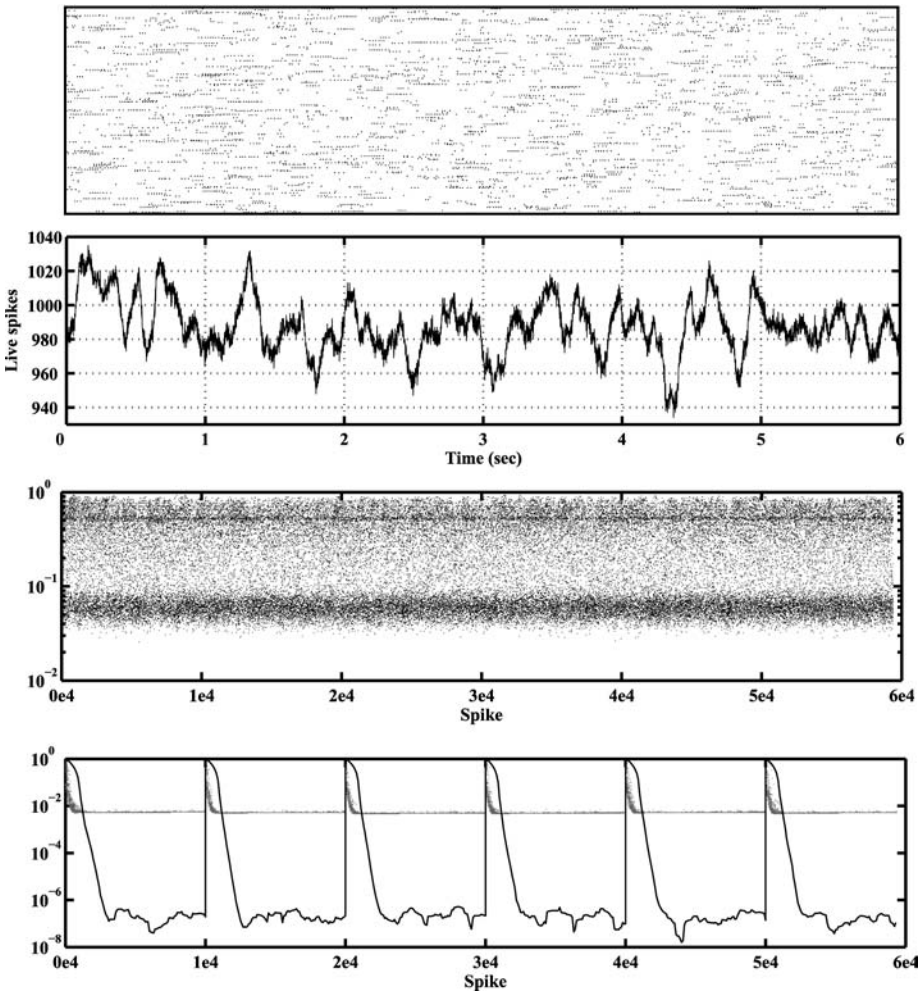

Fig. 5 Dynamics of the random inhibitory network. Panels (top) Spike raster plot of the first 250 neurons in the population for the first 2 of the 6 sec of simulation shown in the second panel. (second) Total number of inhibitory spikes in a window of $\Upsilon = 100$ msec for the 6 sec of simulation. (third) $\sum_{i,j}({}^l\alpha_i^j)^2$ for successively generated spikes during this period. The abscissa corresponds to successive spike generations and not time. There were 59379 spikes generated during the 6 sec period, and the mean $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 0.2649. (bottom) l^2-norm of the rows added to A_0^k for each such spike (gray) and $\|B * A_0^k * C\|_F$ normalized by the number of rows in $(B * A_0^k * C)$ (black). The abscissa, once again, corresponds to successive spikes and not time. A_0^k was reset to the identity matrix every 10000 spikes to reduce numerical error. Note that $\|B * A_0^k * C\|_F$ decayed toward 0 repeatedly, demonstrating that the trajectory was insensitive to initial conditions

were modeled using $-1000.0e^{-t/1.2}$, and the time bound Υ was set at 100 msec. Since the system lacked excitatory drive, the threshold of all the neurons was set at a constant negative value so that the network generated sustained recurrent activity. In the example shown in Fig. 5 this value was set at -24.0, causing the neurons in the network to spike at a rate of approximately 10 Hz.

The top panel in Fig. 5 displays the spike rasters of the first 250 of the population of 1000 neurons, for a duration of 2 sec. A visual inspection of the plot shows no discernible pattern to the spiking of the neurons, because, even if the trajectory were periodic, the period would likely be very large. The second panel displays the total number of live spikes at each time step for a period of 6 sec (the first 2 sec of which is reported in the top panel). Since Υ was set at 100 msec and there were approximately 1000 live spikes in the state description at each step, we can deduce that the neurons in the system spiked at approximately 10 Hz.

The third panel displays a raster plot in logarithmic scale of the value of $\sum_{i,j}({}^l\alpha_i^j)^2$ for each of the 59379 spikes that were generated during the 6 sec time period. Note that the value of $\sum_{i,j}({}^l\alpha_i^j)^2$ is strictly less than 1. The value of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 0.2649. The bottom panel displays various aspects of the perturbation matrix A_0^k for the spike trajectory, computed on-line during the simulation. Since the elements of A_0^k dropped to very small values as time progressed, A_0^k was reset to I (the identity matrix) after every 10000 spike generations to reduce numerical error. Displayed in gray is a raster plot (in logarithmic scale) of the l^2-norm of the successive rows introduced into A_0^k for the corresponding spikes generated in the system. It can be seen from the recursive definition of A_0^k in Section 2.2 that the elements of any row in A_0^k sum to 1. Therefore, the norm asymptotes eventually.[11] Displayed in black is $\|B * A_0^k * C\|_F$, normalized by the number of rows in $(B * A_0^k * C)$ to distinguish from the confounding effect of the size of the matrix. Note that $\|B * A_0^k * C\|_F$ decays toward 0 repeatedly, demonstrating that the trajectory was indeed insensitive to initial conditions.

[11] The l^2-norm of any row in A_0^k is bounded from below due to the noted constraint. The lower bound is reached when all elements of the row take the same value, $1/n$, where n is the number of columns in A_0^k.

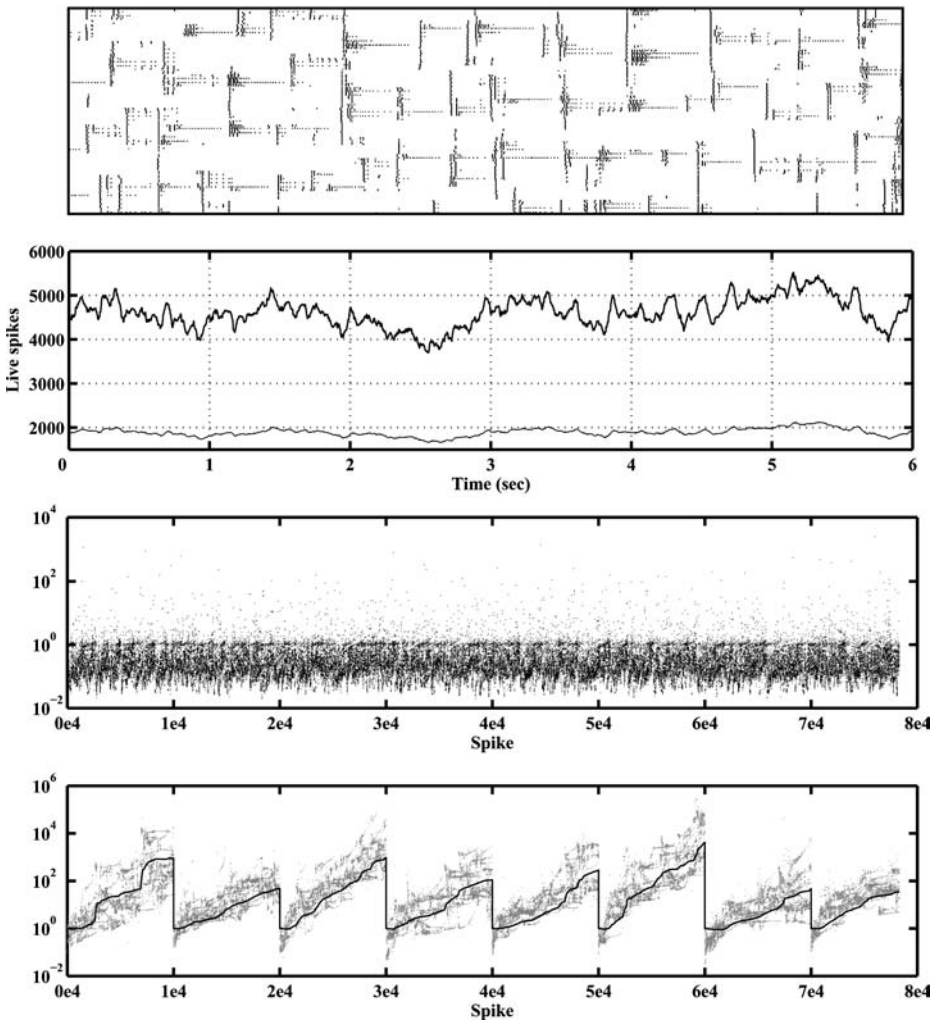

Fig. 6 Dynamics of the highly structured ring network. Panels (top) Spike raster plot of the first 250 neurons in the population for the first 2 of the 6 sec of simulation shown in the second panel. (second) Total number of excitatory (thick line) and inhibitory (thin line) spikes in a window of $\Upsilon = 500$ msec for the 6 sec of simulation. (third) $\sum_{i,j}({}^l\alpha_i^j)^2$ for successively generated spikes during this period. The abscissa corresponds to successive spike generations and not time. There were 78310 spikes generated during the 6 sec period, and the mean $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 45.1933. (bottom) l^2-norm of the rows added to A_0^k for each such spike (gray) and $\|B * A_0^k * C\|_F$ normalized by the number of rows in $(B * A_0^k * C)$ (black). The abscissa, once again, corresponds to successive spikes and not time. A_0^k was reset to the identity matrix every 10000 spikes to reduce numerical error. Note that $\|B * A_0^k * C\|_F$ grew repeatedly, demonstrating that the trajectory was sensitive to initial conditions

On the flip side, the theorem suggests that insensitive dynamics does not necessarily follow from structured connectivity and connection strengths in networks. To verify this conjecture, we ran detailed simulations of the dynamics of highly structured networks that the theorem predicted ought to be almost surely sensitive to initial conditions because their dynamics were dominated by excitatory PSPs. In all cases investigated, the trajectories were found to be sensitive to initial conditions. Figure 6 displays the results from an instantiation of the highly structured ring network that was described earlier. The network was composed of 800 excitatory and 200 inhibitory neurons that were interleaved evenly and placed around a ring. Each neuron received synapses from a sector of 100 contiguous neurons, the sector being centered on that neuron. As a result, any two neighboring neurons shared 96 of their 100 inputs. This had the effect of output spike trains from neighboring neurons being highly correlated, with sectors of neurons generating synchronized bursts of spikes. The axonal/synaptic delay for each neuron was set identically to 0.8 msec. PSPs were modeled using (14); all excitatory PSPs were set to be identical and modeled as $\{6.5/(1.5\sqrt{t})\}\{e^{-1.0*1.5^2/t}e^{-t/20} + e^{-4.0*1.5^2/t}e^{-t/80}\}$, and all inhibitory PSPs were set to be identical and modeled as $\{-39.0/(1.0\sqrt{t})\}e^{-1.0*1.0^2/t}e^{-t/20}$. All AHPs were modeled as $-1000.0e^{-t/1.2}$, and the time bound Υ was set at 500 msec. The threshold for the excitatory and the inhibitory neurons were set at 5.0 and 10.0 respectively, which led the excitatory (respectively, inhibitory) population to spike at a rate of approximately 11 (respectively, 20) Hz.

The top panel in Fig. 6 displays the spike rasters of a sector of 250 neurons from the population, for a duration of 2 sec. As noted earlier, the system generated synchronized bursts of spikes. The second panel displays the total number of live excitatory spikes (thick line) and inhibitory spikes (thin line) at each time step for a period of 6 sec (the first 2 sec of which is reported in the top panel). Since Υ was set at 500 msec and there were approximately 4400 live excitatory and 2000 live inhibitory spikes in the state description at each step, we can deduce that the excitatory (respectively,

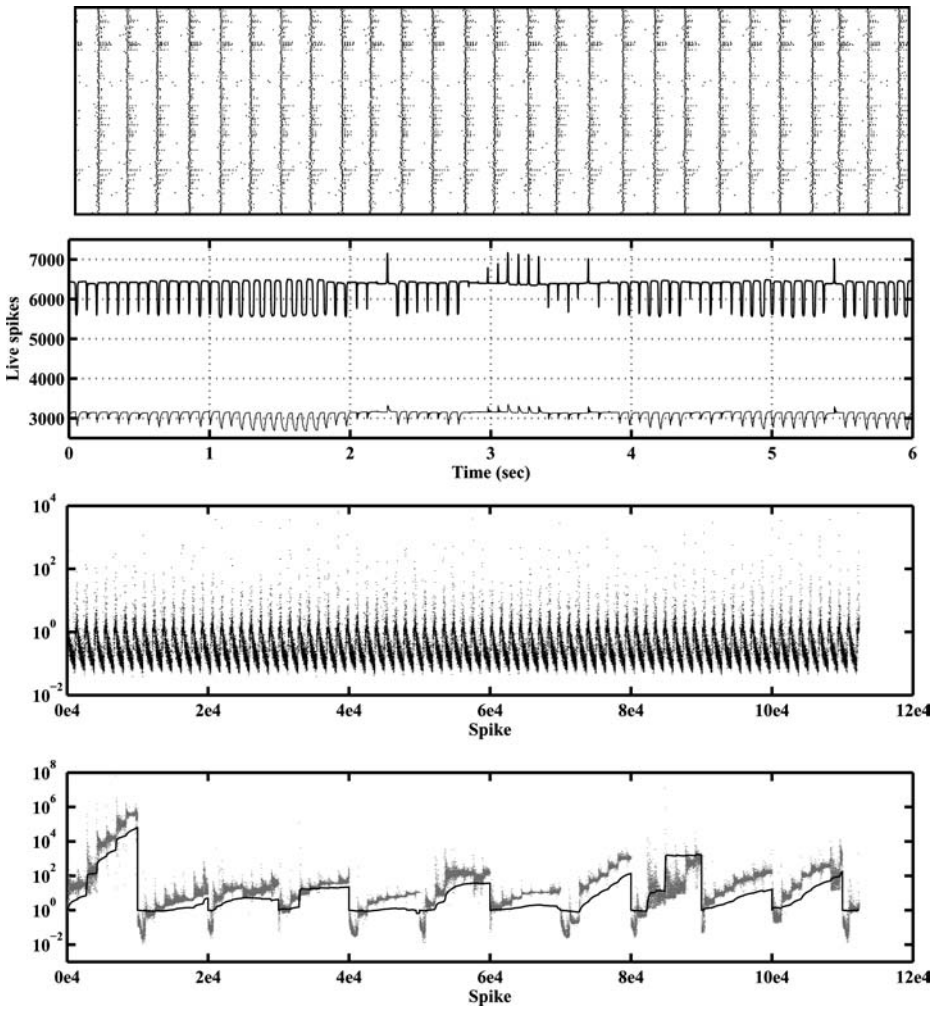

Fig. 7 Dynamics of the random bursty cortical network. Panels (top) Spike raster plot of the first 250 neurons in the population for the first 2 of the 6 sec of simulation shown in the second panel. (second) Total number of excitatory (thick line) and inhibitory (thin line) spikes in a window of $\Upsilon = 500$ msec for the 6 sec of simulation. (third) $\sum_{i,j}({}^l\alpha_i^j)^2$ for successively generated spikes during this period. The abscissa corresponds to successive spike generations and not time. There were 112374 spikes generated during the 6 sec period, and the mean $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 14.2163. (bottom) l^2-norm of the rows added to A_0^k for each such spike (gray) and $\|B * A_0^k * C\|_F$ normalized by the number of rows in $(B * A_0^k * C)$ (black). The abscissa, once again, corresponds to successive spikes and not time. A_0^k was reset to the identity matrix every 10000 spikes to reduce numerical error. Note that $\|B * A_0^k * C\|_F$ grew repeatedly, demonstrating that the trajectory was sensitive to initial conditions

inhibitory) neurons in the system spiked at approximately 11 (respectively, 20) Hz.

The third panel displays a raster plot in logarithmic scale of the value of $\sum_{i,j}({}^l\alpha_i^j)^2$ for each of the 78310 spikes that were generated during the 6 sec time period. Since some of the $\sum_{i,j}({}^l\alpha_i^j)^2$'s were rather large, we first generated a probability histogram from the data to check for a heavy-tail distribution. We found that the tail of the distribution decayed faster than a quadratic rate, indicating that $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was well-defined and less than ∞.[12] $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was then calculated from the data and was found to be 45.1933. The bottom panel is exactly as in Fig. 5. Since the elements of A_0^k grew to very large values as time progressed, A_0^k was reset to the identity matrix after every 10000 spike generations to reduce numerical error. As is clear from the figure, $\|B * A_0^k * C\|_F$ as well as the l^2-norm of the successive rows introduced into A_0^k for the corresponding spikes generated in the system, grew repeatedly, demonstrating that the trajectory was indeed sensitive to initial conditions.

The connectivity between neurons within cortical columns (intra-column as opposed to thalamic afferents), while having evolved to achieve specific functions, have been experimentally ascertained to fit random statistical distributions (Braitenberg and Schüz, 1991; Schüz, 1992). Approximately 80% of the neurons in cortical networks are excitatory and the rest inhibitory (Shepherd, 1998). This makes recurrent excitation the major theme of such networks, with excitatory PSPs dominating their dynamics. Under conditions where the neurons in the networks spike at low to moderate rates, fast acting AHPs (such as the one modeled in the simulations presented thus far) are substantially decayed by the time the neurons spike next. Consequently, at the birth of spikes, they contribute very small positive ${}^l\alpha_i^j$'s which have minimal impact on the value of $\sum_{i,j}({}^l\alpha_i^j)^2$. Although it is well known that slow acting AHPs can have a major impact on the global nature of the dynamics of networks of neurons, including the setting of the timescale for interburst intervals in such systems, their impact on the sensitivity of

[12] This test was performed on all the trajectories presented in this article. In each case, we found that the tail of the distribution decayed faster than a quadratic rate.

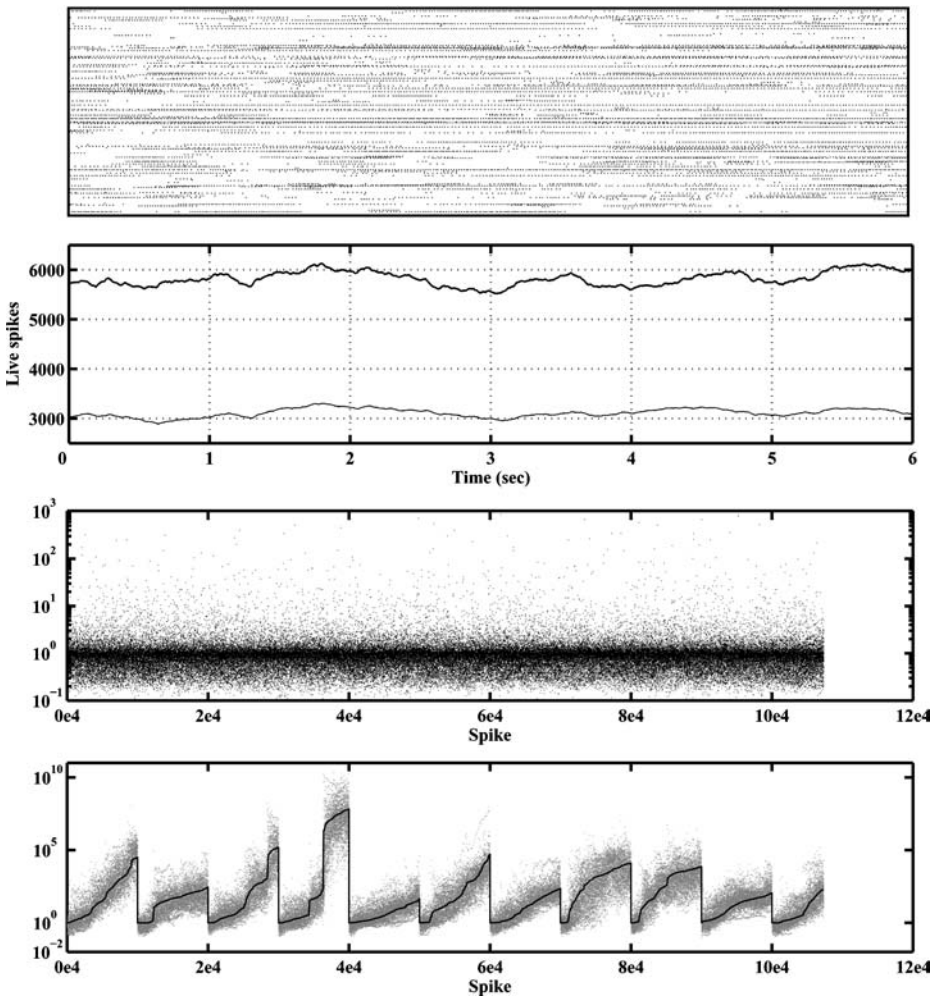

Fig. 8 Dynamics of the asynchronous cortical network. Panels (top) Spike raster plot of the first 250 neurons in the population for the first 2 of the 6 sec of simulation shown in the second panel. (second) Total number of excitatory (thick line) and inhibitory (thin line) spikes in a window of $\Upsilon = 500$ msec for the 6 sec of simulation. (third) $\sum_{i,j}({}^l\alpha_i^j)^2$ for successively generated spikes during this period. The abscissa corresponds to successive spike generations and not time. There were 107354 spikes generated during the 6 sec period, and the mean $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 7.4270. (bottom) l^2-norm of the rows added to A_0^k for each such spike (gray) and $\|B * A_0^k * C\|_F$ normalized by the number of rows in $(B * A_0^k * C)$ (black). The abscissa, once again, corresponds to successive spikes and not time. A_0^k was reset to the identity matrix every 10000 spikes to reduce numerical error. Note that $\|B * A_0^k * C\|_F$ grew repeatedly, demonstrating that the trajectory was sensitive to initial conditions

the underlying spike trajectories is again minimal. This is due to the fact that the $\partial P_l/\partial x_i^j$'s for slow acting AHPs are, by definition, small. They too, consequently, contribute very small positive ${}^l\alpha_i^j$'s at the birth of spikes that have minimal impact on the value of $\sum_{i,j}({}^l\alpha_i^j)^2$. The dominating effect of the excitatory PSPs in such networks is therefore likely to make $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ large, causing the dynamics of such systems to be almost surely sensitive to initial conditions. The following simulation results corroborate these observations.

We simulated several examples of cortical networks (80% excitatory and 20% inhibitory neurons) with random connectivity between neurons. The qualitative dynamics of the systems varied widely depending on the value of the parameters. However, in each case we found that the dynamics of the network was sensitive to initial conditions due to the reasons stated above. In Figs. 7 and 8 we present two examples of such networks. Both networks had excitatory AMPA and NMDA synapses, and inhibitory GABA_A and GABA_B synapses. The first network differed from the second in having weaker inhibitory connections and an additional slow acting AHP in excitatory neurons, that caused the dynamics of the network to be bursty. AHPs for the excitatory neurons in the first network were modeled as $-(1000.0e^{-t/1.2} + 20.0e^{-t/40})$ and that for the inhibitory neurons as $-1000.0e^{-t/1.2}$. All AHPs in the second network were modeled as $-1000.0e^{-t/1.2}$. The remaining parameters for the networks were as follows. Combined axonal/synaptic delays were chosen randomly from [0.4, 0.9] msec. PSPs were modeled using (14). The parameter d was chosen randomly from the range [1.0, 1.1] for all inhibitory synapses on excitatory neurons, and from [1.0, 2.0] for all other synapses. β and τ for the AMPA, NMDA, GABA_A and GABA_B synapses on excitatory and inhibitory neurons were set at the values specified earlier in Section 2.3. In the first network, Q for the AMPA synapses were chosen randomly from [1.0, 8.0] when on excitatory, and [1.2, 9.6] when on inhibitory neurons. Q for the NMDA synapses were chosen randomly from [1.0, 8.0]. Q for the GABA_A synapses were chosen randomly from [3.0, 24.0] when on excitatory, and [2.0, 16.0] when on inhibitory neurons. Finally, Q for the GABA_B

synapses were chosen randomly from [1.0, 8.0]. The threshold was set at 3.0 for excitatory and 20.0 for inhibitory neurons.

Apart from the AHPs, the second network differed from the first in two respects. First, the threshold for the inhibitory neurons was lowered to 6.0. Second, Q for the AMPA synapses were chosen randomly from [0.9, 9.0] when on excitatory, and [0.5, 5.0] when on inhibitory neurons. Q for the NMDA synapses were chosen randomly from [1.0, 10.0] when on excitatory, and [0.25, 2.5] when on inhibitory neurons. Q for the $GABA_A$ synapses were chosen randomly from [6.0, 60.0] when on excitatory, and [0.5, 5.0] when on inhibitory neurons. Finally, Q for the $GABA_B$ synapses were chosen randomly from [3.0, 30.0] when on excitatory, and [0.5, 5.0] when on inhibitory neurons.

As is clear from the top two panels of the respective figures, the qualitative dynamics of the two systems were significantly different. However, as the bottom panels in both figures demonstrate, the dynamics was sensitive to initial conditions in both cases. $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ was found to be 14.2163 and 7.4270 for the two systems respectively. It is of some importance to note here that although the dynamics of the bursty network was sensitive, it nevertheless displayed repeated patterns on a coarser time scale. This result demonstrates that sensitive dependence does not necessarily require asynchronous Poisson-like spiking behavior on the part of the neurons. The coarse pattern in the dynamics of the network was the outcome of the shape of the attractor in the phase-space.

The electrophysiology of the cortical neuron is substantially more complex than that expressed in the variant of the Spike-response model used in the above simulations. The reader may therefore wonder about the extent to which these results apply to real cortical networks in the brain. Although this is a legitimate concern, it is our belief that the results do extend to models of the neuron that are biophysically more accurate, due to the simplicity of the criterion $E(\sum_{i,j}({}^l\alpha_i^j)^2)$. Whereas the value of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ is likely to differ in real cortical networks, we do not foresee it dropping below 1 and turning the dynamics almost surely insensitive to initial conditions.

4. Bifurcation in spike dynamics

In the previous section, we presented instructive examples of networks that demonstrated that neither is insensitive dynamics a necessary consequence of structure, nor is sensitive dynamics a necessary consequence of the lack of structure, in the connectivity pattern and the connection strengths of a network. In this section we present a stronger result, that the sensitivity of the dynamics of a network can in fact be modulated independent of its connectivity and connection strengths. We present networks with fixed architecture and connection strengths, that undergo bifurcation from sensitive dynamics to insensitive dynamics as the value of a control parameter is varied.

We conjectured earlier that a network whose dynamics is dominated by excitatory PSPs is more likely to exhibit sensitive dependence since excitatory PSPs tend to contribute more negative ${}^l\alpha_i^j$'s that, in turn, raise the value of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$. The conjecture was based on the general profile of excitatory PSPs, i.e., on their having short rising phases followed by prolonged falling phases. If the duration of the rising phase of excitatory PSPs were to be made longer, arguments along similar lines would suggest that the value of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ would decrease, making the dynamics of the system less sensitive. If $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ were to drop below 1, the dynamics of the network would then become insensitive to initial conditions.

We simulated two classes of networks composed of 1000 neurons with 80% of the neurons chosen randomly to be excitatory and the rest inhibitory. In both classes, each neuron received synapses from 100 randomly chosen neurons, and axonal/synaptic delays were chosen randomly from [0.4, 0.9] msec. In the first class, all PSPs were modeled using the function

$$P(t) = \frac{Q}{d}\left(\frac{t}{\tau}\right)^{\beta} e^{-t/\tau}\frac{1}{\beta^{\beta} * e^{-\beta}}, \tag{17}$$

with τ set at 10 msec. Since $P(t)$ attained its maximum value at $t = \beta\tau$, the normalizing term $1/(\beta^{\beta} * e^{-\beta})$ assured that its maximum value was (Q/d) regardless of the value assigned to β. $\{Q, d\}$ were chosen randomly from {[6.0, 42.0], [1.0, 1.1]} for inhibitory synapses on excitatory neurons and {[1.0, 7.0], [1.0, 2.0]} for all other synapses. AHPs were modeled using $-1000.0e^{-t/5}$. The threshold for excitatory neurons was set at 3.0 and that for inhibitory neurons at 10.0. The simulations lasted 100 sec (5×10^5 steps at 0.2 msec per step). In the case of inhibitory synapses, β was fixed at 0.5, causing inhibitory PSPs to peak at 5.0 msec through the entire simulation. In the case of excitatory synapses, β was raised from 0.5 to 1.0 in increments of 0.05 every 10 sec. This caused excitatory PSPs to peak at progressively later times as the simulation progressed. In effect, β was the control parameter alluded to earlier.

The second class of networks was exactly like the first, except that instead of holding the peak value of excitatory PSPs constant, the area under the excitatory PSPs was held constant through the entire simulation. Stated formally, excitatory PSPs were modeled using the function

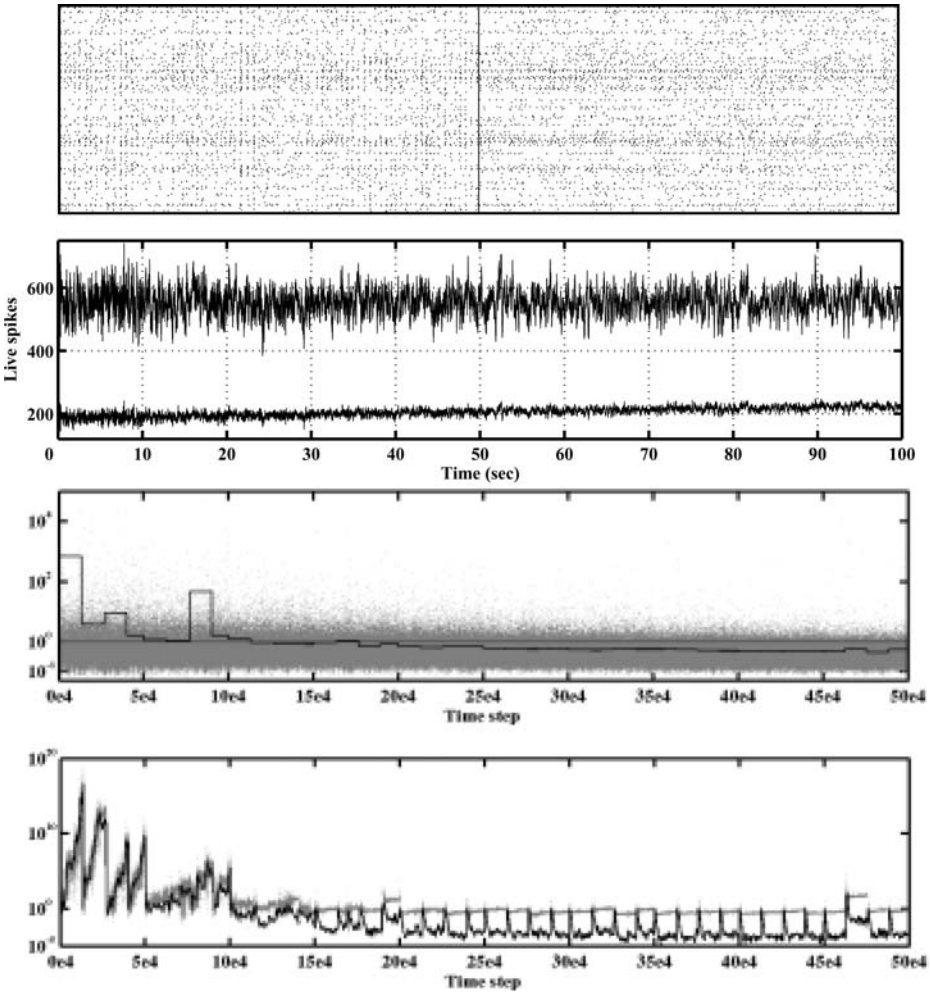

Fig. 9 Dynamics of the network displaying bifurcation as the value of the parameter β that controls the EPSP's time to peak, is increased. Panels (top) Spike raster plot of the first 250 neurons in the population for 2 sec at the beginning (left), and end (right) of the simulation shown in the second panel. (second) Total number of excitatory (upper) and inhibitory (lower) spikes in a window of $\Upsilon = 100$ msec for the entire 100 sec of simulation. (third) $\sum_{i,j}({}^{l}\alpha_i^j)^2$ for successively generated spikes at each time step (gray), and the mean $E(\sum_{i,j}({}^{l}\alpha_i^j)^2)$ computed in blocks that ended with A_0^k being reset (black). The abscissa corresponds to the time step (0.2 msec per step). For the sake of comparison a line at value 1 is also shown. (bottom) l^2-norm of the rows added to A_0^k for each such spike (gray) and $\|B * A_0^k * C\|_F$ normalized by the number of rows in $(B * A_0^k * C)$ (black). The abscissa, once again, corresponds to time step. A_0^k was reset to the identity matrix every 20000 spikes and at every 10 sec of simulation when β was changed. Note that $\|B * A_0^k * C\|_F$ grew repeatedly at the beginning of the simulation and decayed repeatedly at the end, demonstrating that the trajectory was sensitive at the beginning and insensitive at the end

$$P(t) = \frac{Q}{d}\left(\frac{t}{\tau}\right)^{\beta} e^{-t/\tau} \frac{1}{\int_0^{\Upsilon}(t/\tau)^{\beta}e^{-t/\tau}dt}. \tag{18}$$

The results from the simulations of the two classes were found to be very similar, and therefore, we only present results from the first class of networks.

Figure 9 displays the results from an instantiation of the first class of networks. The top panel displays the spike rasters of 250 neurons from the population for a duration of 2 sec, at the beginning of the simulation (left), and at the end of the simulation (right). As is clear from the figure, one can not distinguish between the sensitivity of the dynamics at the beginning and the insensitivity of the dynamics at the end, through visual inspection. The second panel displays the total number of live excitatory spikes (upper) and inhibitory spikes (lower) at each time step for the entire 100 sec simulation. Since Υ was set at 100 msec and there were approximately 560 live excitatory and 200 live inhibitory spikes in the state description at each step, we can deduce that the excitatory (respectively, inhibitory) neurons in the system spiked at approximately 7 (respectively, 10) Hz.

The third panel is like in earlier figures, except that the raster plot of the $\sum_{i,j}({}^{l}\alpha_i^j)^2$'s is shown in gray. The mean of the $\sum_{i,j}({}^{l}\alpha_i^j)^2$'s was computed for each interval that ended with either the perturbation matrix being reinitialized (after every 20000 spikes) or with β being changed (after every 10 sec of simulation). The value of the mean for each such interval is shown in black. For the sake of comparison, a line

at value 1 is also shown. The bottom panel is exactly as in earlier figures. Note that whereas $\|B * A_0^k * C\|_F$ repeatedly tended toward infinity at the beginning of the simulation, it repeatedly decayed toward 0 at the end, confirming that the trajectory was indeed sensitive at the beginning of the simulation and insensitive at the end. Moreover, this change occurred when the value of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ dropped below 1, as predicted by the theorem.

This result might seem inconsistent with our claim that the dynamics of cortical networks ought to be almost surely sensitive to initial conditions, given that the dynamics of the random network described above did become insensitive as the value of β was increased. This is however, not the case. In the above simulation, all inhibitory PSPs peaked at exactly 5 msec. In all simulations that we conducted where inhibitory PSPs peaked either uniformly at a time later than 5 msec, or randomly at times dependent on the position of the synapse on the neuron, we could not replicate the above bifurcation. In all such cases, $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ remained above 1 and the dynamics remained sensitive for all values of β tested; β was raised to values as large as 5.0, corresponding to excitatory PSPs peaking at 50 msec.

5. Systems with stationary inputs

So far we have considered systems without inputs, both in the simulations as well as in the formal analysis. We now extend the formal result to the special case of systems driven by stationary input spike trains. The trajectories under consideration now lie in the phase-space $\prod_{i=1}^{\mathcal{R}} {}^i\overline{\mathbb{L}}_{n_i}$, with $i = 1, \ldots, \mathcal{S}$ denoting the internal neurons of the system, and $i = \mathcal{S}+1, \ldots, \mathcal{R}$ the input neurons.

For the generative model of the local neighborhoods of such trajectories, we employ a variation of the "directly given" dynamical system $(\Omega, \mathcal{F}, \mu, T)$ defined in Section 2.2. An outcome $\omega \in \Omega$ is now an ordered pair of an initial number of live internal and live input spikes, followed by a sequence of births of internal spikes (with contributing internal and input spikes and corresponding $\partial P_l/\partial x_i^j$'s), deaths of internal spikes, births of input spikes, and deaths of input spikes. In essence, all spikes are now additionally labeled as internal or input spikes. The specification of the stationary measure μ also follows along the lines of Section 2.3, with two additional assumptions:

(i) The ratio of the number of internal spikes to the total (internal and input) spikes in successive state descriptions forms a stationary distribution with mean $\nu \in \mathbb{R}$. Moreover, this ratio stays close to ν at all times.

(ii) Input spikes and internal spikes are statistically indistinguishable from each other with regard to being chosen to contribute to the generation of internal spikes, as is also the case for the ${}^l\alpha_i^j$'s that are associated with them.

The difference between $A_0^k(\omega)$ in this stationary stochastic process and the previous, stems from this process having input spikes that are not perturbed. The sensitivity of the system is determined with respect to a perturbation in its internal state.

The process is begun, as before, by choosing an $n \in [M_{\text{low}}, M_{\text{high}}]$ from a fixed distribution, where M_{low} and M_{high} denote the minimum and the maximum number of total (internal and input) live spikes in the trajectory across all time. The $(n \times \nu n)$ matrix A_0^0 is then generated by taking the $(\nu n \times \nu n)$ identity matrix and padding it with $(n - \nu n)$ additional rows all of whose elements are zero. The first νn rows in A_0^0 correspond to the internal spikes in the system, and the remaining $(n - \nu n)$ rows correspond to the input spikes. Since input spikes are not perturbed, all elements in their corresponding rows are set to zeros.

We take an additional step to complete the initialization. Let C be the $(\nu n \times \nu n)$ matrix $I - \frac{1}{\nu n}(\mathbf{1})$, where I is the identity matrix, and $(\mathbf{1})$ is the matrix all of whose elements are 1. The matrix $\mathbb{A}_0^0 = A_0^0 * C$ is generated, and instead of applying the remainder of the stochastic process to A_0^0, we apply it to $\mathbb{A}_0^0$. Note that for any given νn-dimensional perturbation vector $\Delta\mathbf{x_0}$, $\mathbb{A}_0^0 * \Delta\mathbf{x_0}$ yields an n-dimensional perturbation vector, the first νn elements of which sum to zero (corresponding to perturbations on internal spikes satisfying $\sum_{i,j} \Delta x_i^j = 0$), and the remaining $(n - \nu n)$ elements of which are all zeros (corresponding to perturbations on input spikes).[13] Also, note that for the νn non-zero rows, v_i, in $\mathbb{A}_0^0$,

$$E\left(\frac{1}{\nu n}\sum_{i=1}^{\nu n} v_i\right)_{v_i \neq \langle 0 \ldots 0\rangle} = \langle 0, \ldots, 0\rangle. \tag{19}$$

The rest of the stochastic process is carried out as follows. At each step, the choice between the birth of an input spike, the death of an input spike, the birth of an internal spike, and the death of an internal spike, is made from a stationary stochastic process like the one described in Section 2.3, modified appropriately to take into account assumption (i) above. Consider the state of the process just past the $(k-1)$th event. We define E'_{k-1}, V'_{k-1}, and C'_{k-1} (analogous to E_{k-1}, V_{k-1}, and C_{k-1} in (9), (10), and (11)), with these new quantities computed over the population of the non-zero rows (i.e., rows that correspond to internal spikes) in $\mathbb{A}_0^{k-1}$. Formally,

[13] Strictly speaking, the last column of $\mathbb{A}_0^0$ should be discarded, and $\Delta\mathbf{x_0}$ should be a $(\nu n - 1)$-dimensional perturbation vector. The noted constraints on the elements of the n-dimensional perturbation vector $\mathbb{A}_0^0 * \Delta\mathbf{x_0}$, as well as (19), would still hold.

$$E'_{k-1} = E\left(\frac{1}{m}\sum_{i=1}^{m} v_i\right)_{v_i \neq \langle 0 \ldots 0\rangle}, \tag{20}$$

$$V'_{k-1} = E\left(\frac{1}{m}\sum_{i=1}^{m}(v_i - E'_{k-1})\cdot(v_i - E'_{k-1})\right)_{v_i \neq \langle 0 \ldots 0\rangle},$$

and (21)

$$C'_{k-1} = E\left(\frac{1}{m(m-1)}\sum_{i=1, j=1, i\neq j}^{m,m}(v_i - E'_{k-1}) \cdot (v_j - E'_{k-1})\right)_{v_i, v_j \neq \langle 0 \ldots 0\rangle}, \tag{22}$$

where m is now the number of non-zero rows in $\mathbb{A}_0^{k-1}$, and the v_i's, v_j's are the non-zero rows in $\mathbb{A}_0^{k-1}$.

In the case of the birth (respectively, death) of an input spike, $\mathbb{A}_0^k$ is generated from $\mathbb{A}_0^{k-1}$ by inserting (respectively, deleting) a row all of whose elements are zero.

In the case of the birth of an internal spike, $\mathbb{A}_0^k$ is generated from $\mathbb{A}_0^{k-1}$ as follows. The random number $p \in [P_{\text{low}}, P_{\text{high}}]$ and the random vector $\langle Y_1, \ldots, Y_p\rangle$ are chosen in the exact same manner as described in Section 2.3. Rows $v_1, \ldots, v_p$ are however, chosen from $\mathbb{A}_0^{k-1}$ through a slightly modified process. Each row v_i, for $i = 1, \ldots, p$, is chosen via a two step procedure. In the first step, v_i is randomly selected, with probability ν to be a non-zero row, and with probability $(1-\nu)$ to be a row all of whose elements are zero. If v_i is selected to be a non-zero row, it is then chosen in the second step from the non-zero rows of $\mathbb{A}_0^{k-1}$ through a random sampling with replacement that is unbiased (according to an appropriately modified version of (12), reflecting the restriction to non-zero rows, as in (20), (21), and (22)). Otherwise, v_i is set to $\langle 0, \ldots, 0\rangle$. Finally, the derived random variables $y_1, \ldots, y_p$, where $y_i = Y_i / \sum_{i=1}^{p} Y_i$, are computed, and $v_{\text{new}} = \sum_{i=1}^{p} y_i v_i$ is inserted at a random location into $\mathbb{A}_0^{k-1}$ to generate $\mathbb{A}_0^k$. Note that of the p rows chosen from $\mathbb{A}_0^{k-1}$, on average only νp rows have non-zero elements, and the rest are rows all of whose elements are zero.

In the case of the death of an internal spike, $\mathbb{A}_0^k$ is generated from $\mathbb{A}_0^{k-1}$ by choosing a row v_{del} from the non-zero rows of $\mathbb{A}_0^{k-1}$ through a random sampling that is unbiased (according to an appropriately modified version of (13), reflecting the restriction to non-zero rows, as in (20), (21), and (22)), and deleting it from $\mathbb{A}_0^{k-1}$.

In what follows, we first relate E'_{k-1}, V'_{k-1}, and C'_{k-1}, to E'_k, V'_k, and C'_k. We then use that relation to identify the constraint that causes $\lim_{k\to\infty} E(\|B * A_0^k * C\|_F^2)$ to either be ∞ or 0. Finally, we apply (8) to derive almost sure convergence of $\|B * A_0^k * C\|_F$ to the appropriate limits.

We begin by noting that the expectation in (20), (21), and (22) is taken over the stationary distribution on m, the number of non-zero rows in $\mathbb{A}_0^{k-1}$, or equivalently, the number of live internal spikes just past the $(k-1)$th event. For any fixed value of $k-1$, Ω can be partitioned into the pairwise disjoint sets ${}_{m,k-1}F \in \mathcal{F}$ for the finitely many choices of m, where ${}_{m,k-1}F$ denotes the set of outcomes in Ω that have exactly m live internal spikes just past the $(k-1)$th event. The expectation can then be resolved into

$$E'_{k-1} = \sum_{m=\nu M_{\text{low}}}^{\nu M_{\text{high}}} Pr(m) \times {}_{m,k-1}E'_{k-1} \tag{23}$$

(and likewise for V'_{k-1} and C'_{k-1}), where ${}_{m,k-1}E'_{k-1}$ is the corresponding quantity computed over ${}_{m,k-1}F$, and $Pr(m)$ is the probability assigned to ${}_{m,k-1}F$. Note that for ${}_{m,k-1}E'_{k-1}$, whereas the pair of subscripts on the left specifies the subset of Ω on which ${}_{m,k-1}E'_{k-1}$ is computed, the subscript on the right specifies the event just past which ${}_{m,k-1}E'_{k-1}$ is computed; they can be varied independently of each other.

In the following lemma, we restrict ourselves to the set of outcomes in any one such ${}_{m,k-1}F$. To ease readability, we have dropped the reference to ${}_{m,k-1}F$ in ${}_{m,k-1}E'_{k-1}$, ${}_{m,k-1}V'_{k-1}$, and ${}_{m,k-1}C'_{k-1}$, and refer to them as E'_{k-1}, V'_{k-1}, and C'_{k-1}, respectively. E'_k, V'_k and C'_k in the lemma refer to the corresponding quantities computed over the same set of outcomes, ${}_{m,k-1}F$, just past the kth event, i.e., ${}_{m,k-1}E'_k$, ${}_{m,k-1}V'_k$ and ${}_{m,k-1}C'_k$. It follows that

$$E'_k = \sum_{m=\nu M_{\text{low}}}^{\nu M_{\text{high}}} Pr(m) \times {}_{m,k-1}E'_k, \tag{24}$$

and likewise for V'_k and C'_k.

Lemma 1. *Let E'_{k-1}, V'_{k-1}, and C'_{k-1} be the above described quantities computed over ${}_{m,k-1}F \in \mathcal{F}$, the set of outcomes that have exactly m non-zero rows in $\mathbb{A}_0^{k-1}$, or equivalently, the set of outcomes that have exactly m live internal spikes just past the $(k-1)$th event.*

1. *If the kth event corresponds to the deletion of a non-zero row (v_{del}), or to the addition or deletion of a row all of whose elements are zero, then*

$$E'_k = E'_{k-1}, V'_k = V'_{k-1}, \text{ and } C'_k = C'_{k-1}.$$

2. *If the kth event corresponds to the addition of a non-zero row (v_{new}), then*

$$E'_k = \frac{m+\nu}{m+1} E'_{k-1}.$$

Moreover, if $pE(y_i^2) = (1+\delta) < \infty$ *for some* $\delta \in \mathbb{R}$ $(\delta > -1$ *necessarily), then*

$$V'_k = \left(\frac{m^2 + m\nu(1+\delta) - \nu^2\delta}{m(m+1)}\right) V'_{k-1} - \left(\frac{(m-1)\nu^2\delta}{m(m+1)}\right) C'_{k-1}, \text{ and}$$

$$C'_k = \left(\frac{2\nu}{m(m+1)}\right) V'_{k-1} + \left(\frac{(m+2\nu)(m-1)}{m(m+1)}\right) C'_{k-1}.$$

Proof:

1. The result is trivial since v_{del} is an unbiased sample, and the addition or deletion of a row all of whose elements are zero has no impact on the quantities in question.
   Before considering 2., we note that for the derived random variables $y_1, ..., y_p$,
   (a) $\forall i,\ pE(y_i) = 1$, and
   (b) $\forall i, j, i \neq j,\ pE(y_i^2) + p(p-1)E(y_i y_j) = 1$,
   because $Y_1, \ldots, Y_p$ are exchangeable random variables. Let $P(q) = \binom{p}{q}\nu^q(1-\nu)^{p-q}$, i.e., the probability of having q non-zero rows out of p draws.
2. $E(v_{\text{new}}) = \sum_{q=0}^{p} P(q)(E(\sum_{i=1}^{q} y_i v_i))$, where it is understood that the expected value $E(\cdot)$ is taken over the non-zero rows in question in $\mathbb{A}_0^{k-1}$.
   The y_i's being independent of the v_i's, and the v_i's being unbiased samples,

$$\begin{aligned} E(v_{new}) &= \sum_{q=0}^{p} P(q)\left(E(y_i)E\left(\sum_{i=1}^{q} v_i\right)\right) \\ &= \sum_{q=0}^{p} P(q)\left(\frac{1}{p} q E'_{k-1}\right) \\ &= \nu E'_{k-1}. \end{aligned}$$

Hence $E'_k = \dfrac{m+\nu}{m+1} E'_{k-1}$.

Moreover, since $E'_0 = \langle 0, \ldots, 0\rangle$ for the initial population of non-zero rows in $\mathbb{A}_0^0$ regardless of the value of m (see (19)), we can conclude that for all m, k, $E'_k = \langle 0, \ldots, 0\rangle$.

Since $E'_{k-1} = \langle 0, \ldots, 0\rangle$ for all values of m, the quantity (23) is also $\langle 0, \ldots, 0\rangle$ for all k. In the remainder of this proof, E'_{k-1} shall refer to the original quantity in (23).

$$\begin{aligned} V(v_{\text{new}}) &= E\left((v_{\text{new}} - E'_{k-1}) \cdot (v_{\text{new}} - E'_{k-1})\right) \\ &= E(v_{\text{new}} \cdot v_{\text{new}}) \\ &= \sum_{q=0}^{p} P(q)\left(E\left(\left(\sum_{i=1}^{q} y_i v_i\right) \cdot \left(\sum_{i=1}^{q} y_i v_i\right)\right)\right) \\ &= \sum_{q=0}^{p} P(q)\left(E\left(\sum_{i=1}^{q} y_i^2 (v_i \cdot v_i)\right) + E\left(\sum_{i=1, j=1, i\neq j}^{q,q} y_i y_j (v_i \cdot v_j)\right)\right). \end{aligned}$$

The y_i's being independent of the v_i's, and the v_i's being unbiased samples,

$$\begin{aligned} V(v_{\text{new}}) &= \sum_{q=0}^{p} P(q)\left(E\left(y_i^2\right)E\left(\sum_{i=1}^{q} v_i \cdot v_i\right) + E(y_i y_j)E\left(\sum_{i=1, j=1, i\neq j}^{q,q} v_i \cdot v_j\right)\right) \\ &= \frac{1+\delta}{p}(\nu p V'_{k-1}) + \frac{-\delta}{p(p-1)}\nu^2 p(p-1)\left(\frac{1}{m}V'_{k-1} + \frac{m-1}{m}C'_{k-1}\right) \\ &= \left((1+\delta)\nu - \frac{\delta\nu^2}{m}\right)V'_{k-1} - \delta\nu^2\frac{m-1}{m}C'_{k-1}. \end{aligned}$$

Note that the extra term in the solution to $E(\sum_{i=1, j=1, i\neq j}^{q,q} v_i \cdot v_j)$ is due to sampling with replacement.
Likewise,

$$\begin{aligned} C(v_{\text{new}}, v_i) &= E\left((v_{\text{new}} - E'_{k-1}) \cdot (v_i - E'_{k-1})\right) \\ &= E\left(v_{\text{new}} \cdot v_i\right) \\ &= \sum_{q=0}^{p} P(q)\left(E\left(\left(\sum_{j=1}^{q} y_j v_j\right) \cdot v_i\right)\right) \\ &= \nu\left(\frac{1}{m}V'_{k-1} + \frac{m-1}{m}C'_{k-1}\right). \end{aligned}$$

The result then follows from

$$V'_k = \frac{m}{m+1}V'_{k-1} + \frac{1}{m+1}V(v_{\text{new}}), \text{ and}$$

$$C'_k = \frac{m-1}{m+1}C'_{k-1} + \frac{2}{m+1}C(v_{\text{new}}, v_i).$$

□

In the above lemma, we computed ${}_{m,k-1}V'_k$ and ${}_{m,k-1}C'_k$ for four different scenarios, assuming hypothetically in each scenario that the kth event for all outcomes $\omega \in {}_{m,k-1}F$ is identical, being either the birth of an internal spike, the death of an internal spike, the birth of an input spike, or the death of an input spike. With the exception of the scenario corresponding to the birth of an internal spike, there was no change in value from ${}_{m,k-1}V'_{k-1}$ to ${}_{m,k-1}V'_k$, and from ${}_{m,k-1}C'_{k-1}$ to

${}_{m,k-1}C'_k$. Although the outcomes in ${}_{m,k-1}F$ will differ in their kth event, it follows from the lemma that any change in value from ${}_{m,k-1}V'_{k-1}$ to ${}_{m,k-1}V'_k$ and from ${}_{m,k-1}C'_{k-1}$ to ${}_{m,k-1}C'_k$, is determined by the probability assigned to the subset of ${}_{m,k-1}F$ whose kth event is the birth of an internal spike. For the entire sample space Ω, it follows from the counterparts to (23) and (24) for V' and C', that by appropriately constraining δ, we can enforce monotonic trends on fixed functions of V'_k, C'_k, as $k \to \infty$. Moreover, such constraints on δ can be computed by treating the kth event for all outcomes $\omega \in \Omega$ as the birth of an internal spike. The following theorem is based on these observations.

Theorem 2 *(Sensitivity with input) Let $\Psi_x(t)$ be a trajectory that is not drawn into the trivial fixed point in a system with input neurons. Let $E(\sum_{i,j}({}^l\alpha_i^j)^2) = 1 + \delta < \infty$ where the expected value $E(\cdot)$ is taken over the set of all births of internal spikes in $\Psi_x(t)$. Then, if*

$$\delta > \max\left(\frac{2 + O(1/M_{\text{low}})}{\nu} - 2, \frac{1}{1 - \nu O(1/M_{\text{low}})}\left(\frac{1}{\nu} - 1\right)\right)$$

(respectively, $\delta < \frac{1}{\nu} - 2$), $\Psi_x(t)$ is, with probability 1, sensitive (respectively, insensitive) to initial conditions. M_{low} denotes the minimum number of total (internal and input) live spikes in $\Psi_x(t)$ across all time.

Proof Let m denote the total number of rows (zero as well as non-zero) in $\mathbb{A}_0^k$. Then, by assumption, $m \in [M_{\text{low}}, M_{\text{high}}]$.

We first note that if $\lim_{k\to\infty} \|(I - \frac{1}{m}(\mathbf{1})) * \mathbb{A}_0^k\| = \infty$ (respectively, 0), then $\lim_{k\to\infty} \|B * A_0^k * C\| = \infty$ (respectively, 0). This follows from the fact that B is $I - \frac{1}{m}(\mathbf{1})$ without the last row, and that the rank of both $I - \frac{1}{m}(\mathbf{1})$ and B is $(m-1)$. We also note that of the k events, if k_b involve births of internal spikes, then by assumption, $k \to \infty$ implies $k_b \to \infty$ almost surely.

Let v_i's, v_j's denote the non-zero rows of $\mathbb{A}_0^k$. Then, abusing notation slightly for the sake of readability so that the square on the right hand side represents dot product with itself,

$$E\left(\left\|\left(I - \frac{1}{m}(\mathbf{1})\right) * \mathbb{A}_0^k\right\|_F^2\right) = E\left(\sum_{i=1}^{\nu m}\left(v_i - \left(\sum_{j=1}^{\nu m} v_j/m\right)\right)^2 + \sum_{i=\nu m+1}^{m}\left(\sum_{j=1}^{\nu m} v_j/m\right)^2\right)$$
$$= \nu\big((m-1)V'_k - (\nu m - 1)C'_k\big).$$

Therefore, since $C'_k \le V'_k$,

$$E\left(\left\|\left(I - \frac{1}{m}(\mathbf{1})\right) * \mathbb{A}_0^k\right\|_F^2\right) \le \nu m(V'_k - \nu C'_k).$$

Also, if $C'_k > 0$,

$$E\left(\left\|\left(I - \frac{1}{m}(\mathbf{1})\right) * \mathbb{A}_0^k\right\|_F^2\right) \ge \nu(m-1)(V'_k - \nu C'_k).$$

Else, if $C'_k \le 0$,

$$E\left(\left\|\left(I - \frac{1}{m}(\mathbf{1})\right) * \mathbb{A}_0^k\right\|_F^2\right) \ge \nu(m-1)V'_k.$$

Now, consider the right hand sides of the above inequalities. From Lemma 1 and the earlier observations,

$$\begin{aligned} V'_k - \nu C'_k &= \frac{(m-\nu)(m+2\nu+\delta\nu)}{m(m+1)}V'_{k-1} \\ &\quad - \frac{(m-1)(m+2\nu+\delta\nu)}{m(m+1)}\nu C'_{k-1} \\ &= \frac{(m-1)(m+2\nu+\delta\nu)}{m(m+1)}(V'_{k-1} - \nu C'_{k-1}) \\ &\quad + \frac{(1-\nu)(m+2\nu+\delta\nu)}{m(m+1)}V'_{k-1}. \end{aligned}$$

Since the last term of the final equation on the right hand side is always greater than 0,

$$V'_k - \nu C'_k \ge \frac{(m-1)(m+2\nu+\delta\nu)}{m(m+1)}(V'_{k-1} - \nu C'_{k-1}).$$

$$\text{If } \delta > \frac{2 + 2/(m-1)}{\nu} - 2 \text{ then } \frac{(m-1)(m+2\nu+\delta\nu)}{m(m+1)} > 1.$$

If $C'_k \le 0$ then from Lemma 1 and the earlier observations,

$$V'_k \ge \left(\frac{m^2 + m\nu(1+\delta) - \nu^2\delta}{m(m+1)}\right)V'_{k-1}.$$

$$\text{If } \delta > \frac{1}{1-\nu/m}\left(\frac{1}{\nu} - 1\right) \text{ then } \frac{m^2 + m\nu(1+\delta) - \nu^2\delta}{m(m+1)} < 1.$$

Conversely, from Lemma 1 and the earlier observations,

$$\begin{aligned} V'_k - \nu C'_k &= \frac{m(m+2\nu+\delta\nu)}{m(m+1)}(V'_{k-1} - \nu C'_{k-1}) \\ &\quad + \frac{-\nu(m+2\nu+\delta\nu)}{m(m+1)}(V'_{k-1} - C'_{k-1}). \end{aligned}$$

Since the last term of the equation on the right hand side is always lesser than 0,

$$V'_k - \nu C'_k \leq \frac{m(m+2\nu+\delta\nu)}{m(m+1)}(V'_{k-1} - \nu C'_{k-1}).$$

$$\text{If } \delta < \frac{1}{\nu} - 2 \text{ then } \frac{m(m+2\nu+\delta\nu)}{m(m+1)} < 1.$$

By monotonicity, and the fact that $k \to \infty$ implies $k_b \to \infty$ almost surely, if

$$\delta > \max\left(\frac{2+O(1/M_{low})}{\nu} - 2, \frac{1}{1-\nu O(1/M_{low})}\left(\frac{1}{\nu} - 1\right)\right)$$

$$\left(\text{respectively, } \delta < \frac{1}{\nu} - 2\right),$$

then $\lim_{k\to\infty} E(\|(I - \frac{1}{m}(\mathbf{1})) * \mathbb{A}_0^k\|_F^2) = \infty$ (respectively, 0), and as a consequence, $\lim_{k\to\infty} E(\|B * A_0^k * C\|_F^2) = \infty$ (respectively, 0).

Finally, we apply (8) to show almost sure convergence to the appropriate limits. Consider first the case where $\lim_{k\to\infty} E(\|B * A_0^k * C\|_F^2) = 0$. Since $\|B * A_0^k * C\|_F^2 \geq 0$ for all k, by Fatou's lemma, $E(\lim_k \inf \|B * A_0^k * C\|_F^2) \leq \lim_k \inf E(\|B * A_0^k * C\|_F^2)$. Consequently, noting that $\|B * A_0^k * C\|_F^2 \geq 0$, we get $E(\lim_k \inf \|B * A_0^k * C\|_F^2) = 0$, and the two together implies $Pr(\lim_k \inf \|B * A_0^k * C\|_F^2 > 0) = 0$. Applying (8) and noting once again that both $\|B * A_0^k * C\|_F^2 \geq 0$ and $\|B * A_0^k * C\|_F \geq 0$, we get $Pr(\lim_{k\to\infty} \|B * A_0^k * C\|_F = 0) = Pr(\lim_{k\to\infty} \|B * A_0^k * C\|_F^2 = 0) = 1$.

For the case where $\lim_{k\to\infty} E(\|B * A_0^k * C\|_F^2) = \infty$, we show that for all finite $c > 0$, $Pr(\lim_k \sup \|B * A_0^k * C\|_F < c) = Pr(\lim_k \sup \|B * A_0^k * C\|_F^2 < c^2) = 0$. We demonstrate this by showing that for any given $c > 0$ and $n_0 > 0$, there exists an $n > n_0$ such that $Pr(\cap_{k=n}^{\infty} \|B * A_0^k * C\|_F^2 < c^2) = 0$.

Since $\lim_{k\to\infty} E(\|B * A_0^k * C\|_F^2) = \infty$, given any $c > 0$ and $n_0 > 0$, there exists an $n > n_0$ such that $E(\|B * A_0^n * C\|_F^2) > c^2$. Hence, $Pr(\|B * A_0^n * C\|_F^2 < c^2) < 1$. Now consider just the set of outcomes, ω, that satisfy $\|B * A_0^n(\omega) * C\|_F^2 < c^2$. Since the stationary process has been constructed such that the choices made at each event are independent of the elements of A_0^k, arguments along similar lines show that there exists an $n' > n$ such that $Pr(\|B * A_0^{n'} * C\|_F^2 < c^2 \mid \|B * A_0^n * C\|_F^2 < c^2) < 1$. By repeatedly applying this argument, we get $Pr(\cap_{k=n}^{\infty} \|B * A_0^k * C\|_F^2 < c^2) = 0$. Finally, applying (8), we get $Pr(\lim_{k\to\infty} \|B * A_0^k * C\|_F^2 > c^2) = 1$ for all finite $c > 0$. □

We must point out that the bounds generated in this theorem, when restricted to the case without inputs, are weaker than those in the theorem in Section 2.2, as is clear from setting $\nu = 1$. For the reasonable value of $\nu = 0.5$, noting that M_{low} is a very large number, the criterion reduces to the dynamics being almost surely insensitive when $\delta < 0$ and almost surely sensitive when $\delta > 2$. These values are not significantly different from those deduced for systems without inputs. Based on the experimental values of $E(\sum_{i,j}({}^l\alpha_i^j)^2)$ derived in the previous sections, we reckon that the dynamics of cortical networks when driven by statistically similar input, such as that arriving from other cortical networks, ought to be almost surely sensitive to initial conditions.

6. Conclusions

Systems neuroscience has made significant strides in deciphering the principles governing the operation of those parts of the nervous system that are close to the sensory periphery of an organism (Jacobs and Werblin, 1998; Meister and Berry, 1999; Reinagel and Reid, 2000), to name but a few. In contrast, our understanding of the workings of those regions of the brain that are farther removed from the sensory and motor periphery, is rather lacking. Since multiple areas in the brain project onto such regions, the precise input to these regions is practically impossible to isolate and control. Furthermore, such inputs are received in the form of spike trains. Without a definite understanding of how, and what, information is coded in these spike trains, one can draw few conclusions regarding the nature of the information processing occurring in these regions. Given our limited comprehension of the nature of the inputs to and outputs from such regions of the brain, developing an in-depth formal understanding of the dynamical properties of networks of spiking neurons is a necessary first step towards deciphering the precise computational operations performed by those regions.

One of the core characteristics of any physical system, networks of neurons or otherwise, is whether or not its dynamics is sensitive to initial conditions, and if so, what causes it to be such. In this article, we have systematically explored this issue for systems of spiking neurons, by first identifying the formal assumptions underlying the sensitivity theorem in Banerjee (2001b), and subsequently demonstrating that the assumptions hold reasonably well for a wide variety of network architectures, particularly those that lack overarching structure.

A notable aspect of the criterion for sensitivity is the finding that it is not strongly linked to the pattern of connectivity and connection strengths of a network. In this article, we have provided ample evidence of this fact. A qualitative appraisal of the criterion guided the construction of networks with random architectures yet with dynamics that were insensitive to initial conditions, as well as networks with highly

structured architectures yet with dynamics that were sensitive to initial conditions. The networks thus constructed provide additional insights into sensitivity. For example, we found that asynchronous activity does not necessarily imply sensitive dynamics, as is apparent from an inspection of the spike raster plots from Figs. 5 and 8. Whereas the dynamics of the network in Fig. 5 is insensitive, the dynamics of the network in Fig. 8 is sensitive. We also found that synchronized activity does not necessarily imply insensitive dynamics. An examination of the spike raster plot from Fig. 7 demonstrates this. Moreover, synchronized activity may result from a variety of factors. Whereas the networks in Figs. 6 and 7 both exhibit synchronized activity, their root causes are different. The synchronized activity in Fig. 6 is the outcome of shared connectivity, whereas that in Fig. 7 is the result of slow inhibition. The clearest evidence of the fact that sensitivity of the dynamics of a network can be modulated independent of its connectivity and connection strengths was provided in our final example. The dynamics of the network underwent bifurcation as the rise time of excitatory PSPs was varied.

In all our simulations, including those not reported in this article, we have found that sensitive dependence on initial conditions is a robust feature of networks whose dynamics is dominated by excitatory PSPs, so long as the network operates in a regime where its constituent neurons spike at low to moderate rates. Since cortical networks satisfy both criteria, we have argued that their dynamics ought to be almost surely sensitive to initial conditions. We have presented two examples of model cortical networks with widely differing qualitative dynamics, yet with both displaying sensitive dependence on initial conditions.

Networks of neurons in the brain do not operate in isolation. They are incessantly bombarded by inputs arriving from the external environment as well as from other regions of the brain. We have extended the formal result in Banerjee (2001b) to the particular case of systems driven by stationary inputs, thus modeling a more realistic scenario.

The nature of the criterion for sensitivity indicates why it is important that spikes not be abstracted away in the analysis of the dynamics of systems of neurons. In any analysis that models spike trains using continuous valued rates or instantaneous probabilities of generating spikes, there can be no quantity corresponding to $E(\sum_{i,j}({}^{l}\alpha_i^j)^2)$. This is because not only does $E(\sum_{i,j}({}^{l}\alpha_i^j)^2)$ depend upon the precise conditions prevalent in the neurons at the instant of the generation of spikes, it is also evaluated over the birth of all spikes in a trajectory. However, as the theorem states, $E(\sum_{i,j}({}^{l}\alpha_i^j)^2)$ solely determines whether a trajectory is sensitive or insensitive to initial conditions. One is therefore led to conclude that there exist salient aspects of the dynamics of networks of spiking neurons that are, in principle, indiscernible in analyses that abstract away spike trains.

There remain many important questions to be addressed. Since sensitivity of dynamics would make any attractors in a given region of the phase-space chaotic, finding appropriate signatures that can distinguish between such attractors becomes a pressing issue. Although it is, in principle, possible for two or more attractors to have the same average spike-rates for the constituent neurons of a network, this has not yet been demonstrated in example model cortical networks. Finally, there is the issue of synaptic changes. How synaptic modification rules such as Hebbian long-term potentiation and depression (LTP/LTD), and spike-time dependent plasticity (STDP), affect the nature of the dynamics of systems of spiking neurons has only begun to be investigated.

Acknowledgment We are grateful to the anonymous referees for their critical reading of the manuscript. Their constructive comments have significantly improved this article.

References

Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cerebral Cortex 7: 237–252.

Banerjee A (2001) On the phase-space dynamics of systems of spiking neurons: I. model and experiments. Neural Computation 13: 161–193.

Banerjee A (2001) On the phase-space dynamics of systems of spiking neurons: II. formal analysis. Neural Computation 13: 195–225.

Bi G-q, Poo M-m (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. Journal of Neuroscience 18: 10464–10472.

Braitenberg V, Schüz A (1991) Anatomy of the Cortex: Statistics and Geometry. Springer-Verlag, Berlin Heidelberg New York.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Vision Neuroscience 10: 1157–1169.

Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. Neural Computation 11: 1621–1671.

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. Journal of Computational Neuroscience 8: 183–208.

Burns BD, Webb AC (1976) The spontaneous activity of neurones in the cat's cerebral cortex. Proceedings of the Royal Society of London, B, Biological Sciences 194: 211–233.

Cazelles B, Ferriere RH (1992) How predictable is chaos. Nature 355: 25–26.

Chow CC (1998) Phase-locking in weakly heterogeneous neuronal networks. Physica D 118: 343–370.

Cox CL, Denk W, Tank DW, Svoboda K (2000) Action potentials reliably invade axonal arbors of rat neocortical neurons. Proceedings of the National Academy of Sciences U.S.A. 97: 9724–9728.

Freeman WJ, Skarda C (1985) Spatial EEG patterns, nonlinear dynamics and perception: The non-Sherringtonian view. Brain Research Reviews 10: 147–175.

Froemke RC, Dan Y (2002) Spike timing dependent synaptic modification induced by natural spike trains. Nature 416: 433–437.

Gerstner W, van Hemmen JL (1992) Associative memory in a network of spiking neurons. Network 3: 139–164.

Gerstner W, van Hemmen JL, Cowan JD (1996) What matters in neuronal locking. Neural Computation 8: 1653–1676.

Grassberger P, Procaccia I (1983) Measuring the strangeness of strange attractors. Physica D 9: 189–208.

Gray RM (1988) Probability, Random Processes, and Ergodic Properties. Springer-Verlag, New York Berlin Heidelberg.

Hessler NA, Shirke AM, Malinow R (1993) The probability of transmitter release at a mammalian central synapse. Nature 366: 569–572.

Jacobs AL, Werblin FS (1998) Spatiotemporal patterns at the retinal output. Journal of Neurophysiology 80: 447–451.

Kingman JFC (1973) The ergodic theory of subadditive stochastic processes. Annals of Probability 1: 883–909.

Latham PE, Richmond BJ, Nelson PG, Nirenberg S (2000) Intrinsic dynamics in neuronal networks. I. theory. Journal of Neurophysiology 83: 808–827.

MacGregor RJ, Lewis ER (1977) Neural Modeling, Plenum Press, New York.

Mainen ZF, Sejnowski T (1995) Reliability of spike timing in neocortical neurons. Science 268: 1503–1507.

Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275: 213–215.

Meister M, Berry MJ (1999) The neural code of the retina. Neuron 22: 435–450.

Nowak LG, Sanchez-Vives MV, McCormick DA (1997) Influence of low and high frequency inputs on spike timing in visual cortical neurons. Cerebral Cortex 7: 487–501.

Osborne AR, Provenzale A (1989) Finite correlation dimension for stochastic systems with power-law spectra. Physica D 35: 357–381.

Raastad M, Storm JF, Andersen P (1992) Putative single quantum and single fibre excitatory post synaptic currents show similar amplitude range and variability in rat hippocampal slices. European Journal of Neuroscience 4: 113–117.

Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. Journal of Neuroscience 20:14: 5392–5400.

Rosenmund C, Clements JD, Westbrook GL (1993) Non-uniform probability of glutamate release at a hippocampal synapse. Science 262: 754–757.

Schüz A (1992) Randomness and constraints in the cortical neuropil, In *Information Processing in the Cortex* Aertsen, V, Braitenberg V (Eds.), Springer-Verlag, Berlin Heidelberg New York, pp. 3–21.

Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron 26: 259–271.

Shepherd GM (1998) The Synaptic Organization of the Brain, Oxford University Press, New York.

Snowden RJ, Treue S, Andersen RA (1992) The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. Experimental Brain Research 88: 389–400.

Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. Vision Research 23: 775–785.

Tomko G, Crapper D (1974) Neuronal variability: non-stationary responses to identical visual stimuli. Brain Research 79: 405–418.

van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274: 1724–1726.

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. Neural Computation 10: 1321–1372.