nature neuroscience

Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons

Srdjan Ostojic

Asynchronous activity in balanced networks of excitatory and inhibitory neurons is believed to constitute the primary medium for the propagation and transformation of information in the neocortex. Here we show that an unstructured, sparsely connected network of model spiking neurons can display two fundamentally different types of asynchronous activity that imply vastly different computational properties. For weak synaptic couplings, the network at rest is in the well-studied asynchronous state, in which individual neurons fire irregularly at constant rates. In this state, an external input leads to a highly redundant response of different neurons that favors information transmission but hinders more complex computations. For strong couplings, we find that the network at rest displays rich internal dynamics, in which the firing rates of individual neurons fluctuate strongly in time and across neurons. In this regime, the internal dynamics interact with incoming stimuli to provide a substrate for complex information processing and learning.

Neural activity in the neocortex is highly irregular. Understanding how such irregular activity is generated and what function it serves have been two major goals of theoretical neuroscience.

The mechanistic origin of irregular activity has been proposed to lie in a tight balance between excitatory and inhibitory synaptic inputs[1–5]. Such a balance leads to a highly fluctuating net input current, the mean of which is below threshold, so that action potentials are generated by fluctuations. Experimental measurements *in vivo* have corroborated this picture[6,7]. Theoretical studies have moreover shown that networks of neurons can self-organize to reach such a balanced state[2,8] in which different neurons asynchronously emit action potentials[9]. Neural activity in this state is chaotic in the sense that slight changes in initial conditions lead to drastically different patterns of spike times[2,10,11]. As a result, networks of deterministic neurons in the asynchronous state display activity that looks random, and that can be described by a firing rate at which action potentials are emitted stochastically.

Networks of excitatory and inhibitory neurons in the asynchronous, balanced state are believed to serve as basic computational units in the cortex, each unit network encoding a single variable in its population firing rate[1,5]. In that picture, the computational capacity of a unit network is limited, the activity of all neurons being essentially redundant. Interesting computations are generated by coupling several such units together[12,13], and these are often investigated by representing the activity of each unit network solely by its population firing rate[14]. A seminal study[15] showed that networks consisting of randomly connected rate units display a transition from an inactive state to a highly heterogeneous, chaotic state if the synaptic coupling is strong. Recent works have found that in this heterogenous regime, rate networks possess high computational capabilities[16–19] because of rich internal dynamics[20] and the existence of an exponentially large number of fixed points (possible attractors) that provide the substrate for complex nonlinear computations[21,22]. In contrast, short memory timescales seem to limit the computational capacity of spiking networks exhibiting spike-time chaos[23].

The heterogeneous chaotic state in networks of rate units is distinct from spike-time chaos in an asynchronous network of spiking neurons: the latter denotes chaos at a microscopic level that generates variability of spike times even at a constant firing rate[11], while the former denotes macroscopic chaos in which the firing rates of a large number of populations strongly fluctuate[24]. Whether and how complex dynamics corresponding to macroscopic, rate chaos appear in a network of spiking neurons is an unresolved question, and, as a consequence, it is not clear how the computational capacity of a network of rate units can be exploited by a network of spiking neurons.

Here we reexamine the dynamics and the computational capacity of an unstructured network of sparsely connected excitatory and inhibitory spiking neurons[9]. As the strength of synaptic coupling increases, we find that the network displays a transition from classical asynchronous activity, in which all neurons fire at a constant rate, to heterogeneous asynchronous activity, in which each neuron fires at a different, time-varying rate. We show that this transition is analogous to the transition found in networks of rate units[15], each spiking neuron playing the role of a rate unit. This finding implies that a single, randomly coupled network of spiking neurons possesses much richer computational properties than suspected so far. In particular, the classical asynchronous state reliably propagates signals but has limited computational capacities, while the heterogeneous asynchronous state instead propagates signals less accurately but provides an efficient substrate for computations such as the categorization of temporal inputs.

Group for Neural Theory, Laboratoire de Neurosciences Cognitives, INSERM U960, École Normale Supérieure, Paris, France. Correspondence should be addressed to S.O. (srdjan.ostojic@ens.fr).

Received 27 October 2013; accepted 23 January 2014; published online 23 February 2014; doi:10.1038/nn.3658

RESULTS

The classical asynchronous state

We consider the prototypical network of excitatory and inhibitory leaky integrate-and-fire (LIF) neurons studied in ref. 9 (see Online Methods). The simplified features of this network lead to a small number of free parameters and allow a detailed understanding of the dynamics. In absence of time-varying inputs, such a network displays asynchronous spontaneous activity if inhibition dominates excitation[3,9]: although the dynamics are fully deterministic, each neuron fires irregularly, its activity resembling a Poisson process. This asynchronous state is an emergent, self-sustaining state of the network[2,8] that can be understood by examining the dynamics of a typical neuron: the irregular firing of its presynaptic afferents leads to highly fluctuating inputs to the neuron; if inhibition is strong enough to counterbalance excitation, the total input is on average subthreshold and the action potentials are generated by fluctuations, so that this typical neuron itself also fires irregularly. Irregular inputs therefore lead to irregular outputs, and owing to the recurrent nature of connections in the network the asynchronous state sustains itself.

A more formal analysis combined with some approximations leads to a mathematical description of the asynchronous state[3,9,25] (see Online Methods). This description predicts that, in the asynchronous state, the equilibrium firing rate ν_0 and synaptic couplings are quantitatively related through a self-consistent equation

$$\nu_0 = F\left(\mu(\nu_0), \sigma^2(\nu_0)\right) \qquad (1)$$

where F is the current-to-rate transfer function of an isolated neuron receiving a white noise input[3,26] (Online Methods equation (1)). The mean μ and variance σ^2 of the total input to a typical neuron in the network depend on the presynaptic firing rates and the synaptic couplings (see Online Methods equations (5) and (6)). The self-sustaining nature of the asynchronous state is therefore reflected in the appearance of the equilibrium firing rate on both right and left hand side of equation (1).

A comparison with direct simulations of network dynamics shows that equation (1) accurately predicts the firing rate for moderate values of the overall synaptic coupling (**Fig. 1a**). However if this coupling is increased to large values, the actual firing rates in the network diverge from the predictions. We will show that this divergence is due to a fundamental change in the nature of the dynamics. A key aspect of the classical asynchronous state described by equation (1) is that neurons fire at a rate that is constant in time. This seems natural, as the input to the network is constant in time. If, however, the synaptic interactions are strong enough, a qualitatively different type of asynchronous state emerges, in which the firing rates strongly fluctuate in time and across neurons although the input is constant and uniform.

Heterogeneous instability of classical asynchronous state

Depending on the values of the synaptic couplings and external inputs, the classical asynchronous state can become unsustainable on the part of the network. Its stability can be assessed by allowing the firing rates to deviate from the equilibrium value predicted by equation (1). If the dynamics amplify such a perturbation, the asynchronous state is unstable, and other types of activity emerge, such as runaway activity[3], multistability[3] or oscillations[9,25]. Previous studies considered only homogenous perturbations, meaning that the allowed deviation from the equilibrium firing rate was identical for all neurons belonging to a given subpopulation. Here we adopt a different approach and allow the rate of each neuron to deviate from the equilibrium firing rate by a different amount. This is equivalent to interpreting each neuron as a stochastic process with its own, potentially time-varying firing rate. We found that the network dynamics can amplify such a heterogeneous perturbation and lead to a new collective state.

The response of the network dynamics to slow heterogeneous perturbations is specified by a stability matrix (see Online Methods equation (12)). This matrix is random, owing to the random assignment of the connections, and its eigenvalues can be determined using random matrix theory[27,28]. It possesses an eigenvalue associated with the unit eigenvector, which corresponds to a homogeneous perturbation. This eigenvalue is always negative in the inhibition-dominated regime studied here[3]. The other eigenvalues correspond to heterogeneous perturbations and, in the limit of a large number of neurons, are densely distributed within a circle in the complex plane centered on 0 (**Fig. 1b**). The radius λ_{max} of that circle depends on network parameters (see Online Methods equation (16)). As the synaptic coupling J increases (all other parameters being kept fixed), the value of λ_{max} increases, and it exceeds unity at a critical value J_c (**Fig. 1b**). For couplings larger than J_c, the network possesses eigenvalues larger than unity, and the stability analysis predicts that heterogeneous perturbations are amplified by the dynamics and destabilize the asynchronous state.

The critical coupling J_c closely corresponds to the synaptic coupling at which the mean firing rates in simulations start diverging from

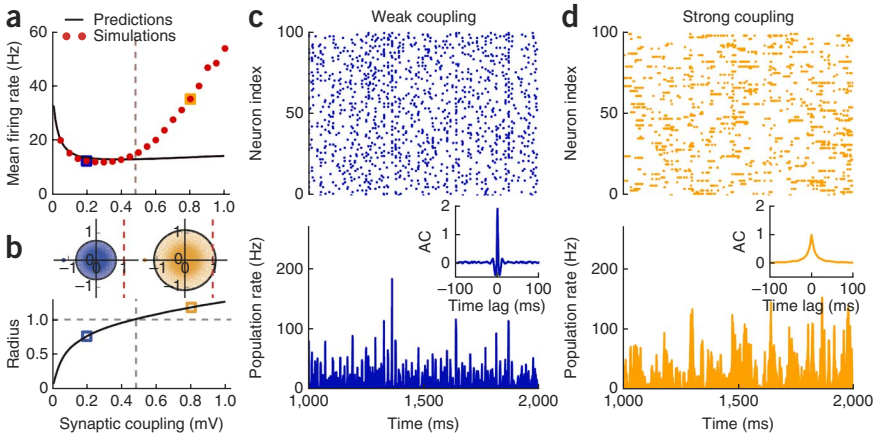

Figure 1 Instability of the asynchronous state in a network of LIF neurons. (**a**) Mean firing rate of the neurons in the network as function of the overall synaptic coupling J. Solid line, analytical prediction for the classical asynchronous state (equation (1)); dots, results of numerical simulations. (**b**) Linear stability of the classical asynchronous state. Top, eigenvalues of the stability matrix (equation (12)) for $J = 0.2$ mV (left) and for $J = 0.8$ mV (right). Bottom, radius λ_{max} of the linear spectrum as function of synaptic coupling J. The value of λ_{max} crosses unity at a critical coupling J_c, above which the asynchronous state is unstable. (**c**) Network activity in the asynchronous state ($J = 0.2$ mV; blue square in **a**). Top, rastergram for 100 neurons in the network; bottom, average population activity in 1-ms bins; inset, autocorrelation of the population activity. (**d**) Network activity above the instability ($J = 0.8$ mV; orange square in **a**). Same quantities as in **c**. Parameter values: constant input $\mu_0 = 24$ mV, relative inhibition strength $g = 5$, number of connections per neuron $C = 1{,}000$, number of neurons $N = 10{,}000$.

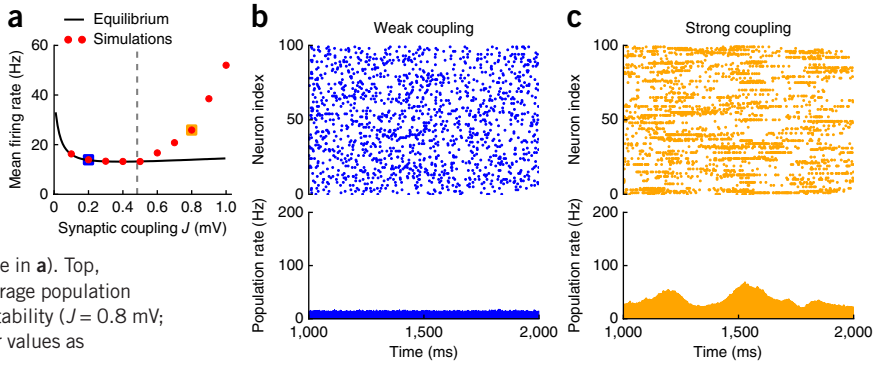

Figure 2 Instability of the asynchronous state in a network of Poisson neurons. The activity of each neuron is fully specified by a temporally varying firing rate. Action potentials are generated stochastically from the underlying rate but do not influence the dynamics. (**a**) Mean firing rate of the neurons in the network as function of synaptic coupling J. Solid line, predictions for the firing rate in the equilibrium state (equation (1)); dots, results of numerical simulations. (**b**) Network activity in the equilibrium state ($J = 0.2$ mV; blue square in **a**). Top, rastergram for 100 neurons in the network; bottom, average population activity in 1-ms bins. (**c**) Network activity above the instability ($J = 0.8$ mV; orange square in **a**). Same quantities as in **b**. Parameter values as in **Figure 1**.

the predictions of equation (1) (**Fig. 1a**,**b**). Rastergrams of simulated spiking activity reveal that above the transition the activity becomes structured (**Fig. 1c**,**d**): neurons tend to fire bursts of spikes closely separated in time, but the timing and size of these bursts varies strongly in time and across neurons. The population-averaged activity looks, however, similar on both sides of the instability (**Fig. 1c**,**d**). It fluctuates strongly on short timescales owing to fast synaptic transmission, which results in a strong central peak in the autocorrelation of the population activity. Independently of the synaptic coupling, if the number of neurons N in the network is increased, the amplitude of the autocorrelation function is reduced proportionally to $1/N$, indicating that the zero-lag synchrony is not a fundamental feature of the dynamics but merely an effect attributable to finite size. The network activity is therefore asynchronous on both sides of the transition.

In summary, combining numerical simulations with a mathematical analysis revealed that as the synaptic coupling increases, the classical asynchronous state loses stability and a different type of asynchronous state emerges. On the population-averaged level, the two states are difficult to distinguish, but we will show that, at the level of individual neurons, they correspond to markedly different dynamics. In the following, we will call the state reached at strong synaptic couplings the heterogeneous asynchronous state, in contrast to the classical, homogeneous asynchronous state that develops at moderate couplings.

Heterogeneous instability in a network of Poisson neurons

To further examine the nature of the heterogeneous asynchronous state that appears at strong synaptic couplings, we introduce a simplified network in which the connectivity matrix is unchanged but each LIF neuron is replaced by a Poisson neuron. The activity of each Poisson neuron is fully specified by its individual time-varying firing rate, and the neurons interact solely through their rates. The dynamics of the firing rate ν_i of ith neuron are given by

$$\tau \frac{d\nu_i}{dt} = -\nu_i + F\left(\mu_i(\nu), \sigma_i^2(\nu)\right) \qquad (2)$$

Here τ is the time constant of the dynamics (see Online Methods), F is the LIF current-to-rate transfer function and $\mu_i(\nu)$ and $\sigma_i^2(\nu)$ are, respectively, the mean and variance of the total input to neuron i (see Online Methods equations (9) and (10)). Spikes can be generated stochastically at the rate prescribed for each neuron but play no role in the dynamics.

In the Poisson network, each neuron is represented by a standard rate-model[14] that describes accurately the slow-timescale dynamics of integrate-and-fire neurons in the asynchronous state[29,30]. We start by showing that the simplified Poisson network displays the same heterogeneous instability as the LIF network, although it lacks the short-timescale dynamics present in the LIF network.

By construction, the Poisson network exhibits an equilibrium state analogous to the classical asynchronous state in the LIF network. In that equilibrium state, the firing rates are constant in time, equal for all neurons and identical to the firing rates in the classical asynchronous state of the LIF network given by equation (1) (see **Fig. 2a**, $J < 0.5$). As expected, the rastergrams generated by the Poisson network lack the short-timescale dynamics present in the LIF network (**Fig. 2b**).

The stability of the equilibrium state of the Poisson network to slow heterogeneous perturbations is specified by the eigenvalues of the same stability matrix as the LIF network (Online Methods equation (12)). The Poisson network therefore exhibits an instability identical to that of the LIF network, at the same critical value of synaptic coupling (**Fig. 1b**). Simulations of the Poisson network show that above the critical coupling, the mean firing rates strongly deviate from the firing rates predicted for the equilibrium state (**Fig. 2a**, $J > 0.5$ mV), in a manner similar to that of the LIF network. For strong couplings, although the rastergram generated by the Poisson network (**Fig. 2c**, top) again lacks the short-timescale dynamics of the LIF network, it clearly displays the same characteristic structured activity as observed for identical parameters in the LIF network (**Fig. 1c**, top).

Fluctuating firing rates in the heterogeneous state

We have shown that a network of Poisson neurons exhibits the same transition at strong synaptic couplings as the network of LIF neurons. We now exploit this analogy to gain insight into the nature of the transition.

In the Poisson network, the dynamics are directly specified in terms of time-varying firing rates of individual neurons. Direct simulations show that, for low synaptic couplings, the rates of all neurons are equal and constant in time, as expected at equilibrium (**Fig. 3a**, top left); above the transition, the firing rates instead strongly fluctuate in time and across neurons (**Fig. 3a**, top right), similarly to their behavior in the chaotic state found in networks of rate units[15,31]. The structured activity seen in the rastergrams above the transition corresponds to high instantaneous firing rates that lead to bursts of closely spaced spikes. Moreover, the steep increase of mean firing rates with synaptic coupling (**Fig. 2a**) is explained by strong fluctuations in the firing rates, combined with the constraint that these fluctuations should be positive.

In contrast to those in the Poisson network, instantaneous firing rates in the LIF network are not inherent properties of individual neurons. One can nevertheless approximate the spiking activity of each neuron by a Poisson process with an underlying time-dependent

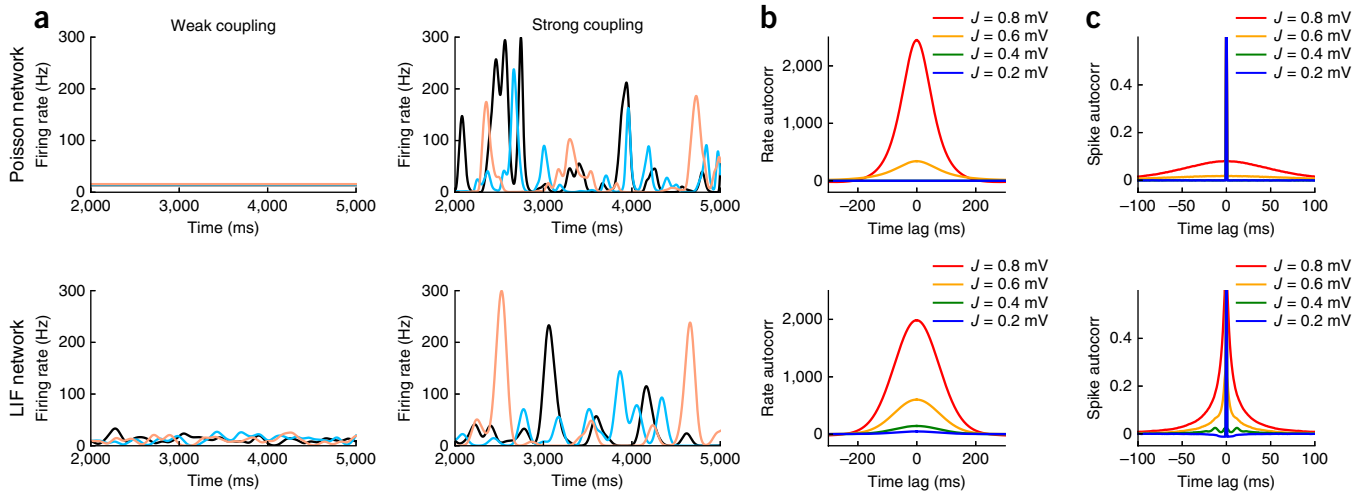

Figure 3 Strong synaptic couplings lead to highly fluctuating instantaneous firing rates of individual neurons. Top, Poisson network; bottom, LIF network. (**a**) Firing rates for three example neurons. Left, low synaptic coupling ($J = 0.2$ mV); right, strong synaptic coupling ($J = 0.8$ mV). For LIF neurons, we estimated the instantaneous firing rates by convolving the spike trains with a 50-ms-wide Gaussian filter. (**b**) Autocorrelation (autocorr) function of instantaneous firing rates, averaged over neurons in the network. (**c**) Autocorrelation function of spike trains, averaged over neurons in the network. All parameters as in **Figures 1** and **2**.

firing rate that we estimated by convolving spike trains with a 50-ms-wide Gaussian filter. The resulting firing rates display a clear similarity with the rates of neurons in the Poisson network (**Fig. 3a**, bottom): for weak couplings, they fluctuate weakly around the equilibrium values; for strong couplings, they fluctuate strongly in time and across neurons.

In both Poisson and LIF networks, above the transition the autocorrelation functions of instantaneous firing rates become broad and their magnitudes increase with synaptic coupling (**Fig. 3b**), indicating that the firing rate dynamics are possibly chaotic[15,31]. The close resemblance between the Poisson and LIF networks is consistent with the fact that the rate dynamics in the Poisson network capture the slow-timescale dynamics of the LIF network. Indeed, the firing rates in the LIF network were estimated by low-pass-filtering spike trains and therefore do not contain short-timescale dynamics. Spike-train autocorrelation functions in the LIF network display a short-timescale structure absent in the Poisson network (**Fig. 3c**) but nevertheless display the characteristic broad profile above the transition and show amplitudes that increase with synaptic strength.

In summary, a comparison of the LIF network with a simplified network of Poisson neurons reveals that the transition from classical to heterogeneous asynchrony is fundamentally a transition between a resting state in which the firing rates of all neurons are constant in time to a resting state in which the firing rates vary strongly in time and across neurons.

Computational properties in the two asynchronous states

The two types of asynchronous activity found in the resting state imply vastly different computational properties. To illustrate these properties, we examined the dynamics of the network in response to two different time-dependent inputs, as seen from an external readout neuron (**Fig. 4**). If such a neuron receives action potentials through NMDA-based synapses, which have a decay timescale of about 100 ms, its input effectively consists of instantaneous firing rates of different neurons similar to those displayed in **Figure 3a**.

The central property of the classical asynchronous state is that, at rest, the network is at an equilibrium point, meaning that the firing rates of all neurons are approximately constant. The main effect of a slow temporally varying input is to modulate the equilibrium firing rate, the activity of individual neurons at each time point being close to the instantaneous equilibrium rate. As a result, the firing rates of all individual neurons are similar to each other and close to the average activity of the population (**Fig. 4b**, left).

In contrast, in the heterogeneous asynchronous state, the network at rest is not at a stable functioning point but exhibits strong internal dynamics in which different neurons fire at varying rates. A temporal input interacts with the internal state of the network to produce complex dynamics in which the firing rates vary greatly across neurons (**Fig. 4b**, right). In particular, the output of each neuron bears little resemblance to the population average.

If the task of the network is to relay the temporal input, then this task is performed more efficiently in the classical asynchronous state than in the heterogeneous state. Indeed, the high redundancy of the activity in the classical asynchronous state favors information transmission[1,5]: sampling the activity of a small subset of neurons in the network provides a good estimate of the population average (**Fig. 4c**). In the heterogeneous state, in contrast, the error on the estimate for the same number of sampled neurons is 5 times larger (**Fig. 4e**), meaning that 25 times as many neurons are needed to obtain the same accuracy (the error scales inversely to the square root of the number of neurons).

While the redundancy in the classical asynchronous state favors information transmission, it hinders more complex transformations of inputs. In principle, the firing rate of every neuron is a potential degree of freedom in the dynamics of the network, and the response to an input can be represented as a trajectory in a high-dimensional space where every dimension represents the activity of one neuron[32,33]. The fact that the response of all neurons is highly similar implies that the dynamics effectively explore only one dimension of this huge space. A principal components analysis applied to the output of the network generated by two inputs shown in **Figure 4a** indeed confirmed that 95% of the variance is explained by the projection on a single dimension (**Fig. 4d**). In the heterogeneous state, in contrast, the dynamics explore many more dimensions (**Fig. 4f**), the effective number of dimensions depending on the number and nature of the inputs.

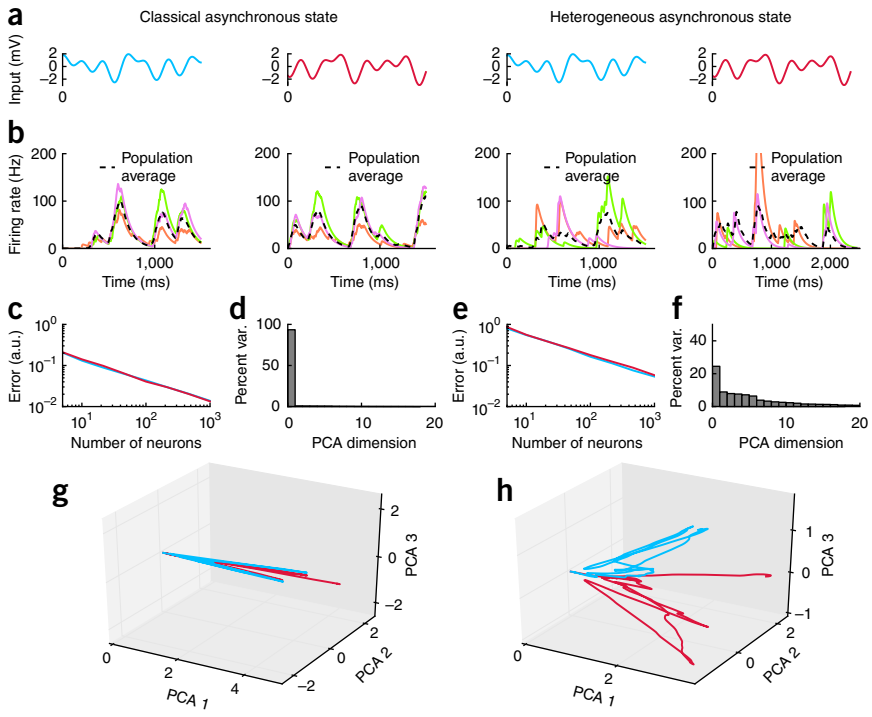

Figure 4 Temporal inputs are processed differently in the two types of resting asynchronous activity. A subset of neurons in the network receives an input current that varies in time. We examine the resulting output of the network as seen by a neuron that reads out the activity of the network through NMDA-like synapses. Left, the network at rest is in the classical asynchronous state; right, the network at rest is in the heterogeneous asynchronous state (same parameters as in **Fig. 1c**,**d** respectively). (**a**) Two different temporally varying inputs given to the network. (**b**) Firing rates in response to the two inputs, as seen by the readout unit. Colored traces, three example neurons; dashed black line, average response of the network. (**c**,**e**) Difference between the mean population response and the average response obtained by pooling a subset of neurons from the network, as function of the size of the pooled subset. (**d**,**f**) Dimensionality of the network response, quantified by the percentage of variance (var.) explained by successive dimensions in a principal component analysis (PCA). (**g**,**h**) Projection on the first three PCA components of the network response to the two inputs shown in **a**, for the classical (**g**) and heterogeneous (**h**) asynchronous states.

The internal dynamics in the heterogeneous asynchronous state transform the one-dimensional temporal input into a high-dimensional output, thereby providing the substrate for complex spatiotemporal computations[20]. The readout unit can, for instance, easily classify the two temporal inputs shown in **Figure 4a** into two separate classes by performing linear discrimination. Indeed a projection of the activity on the first three principal components shows that the responses to the two inputs lie in two separate parts of the activity space (**Fig. 4h**). By adjusting the readout weights in a perceptron-like manner, the readout unit can selectively respond to the input of one kind but not of the other. In contrast, in the classical asynchronous state, as the activity is essentially one-dimensional, the trajectories greatly overlap and cannot be distinguished by a linear readout (**Fig. 4g**). This is an illustration of the general principle that projecting the activity nonlinearly in a high number of dimensions, as is the case in the heterogeneous state, provides the substrate for complex computations[34,35].

Strong inhibition favors classical asynchronous activity

We now return to the conditions under which heterogeneous asynchronous activity appears in the absence of inputs. So far we have mainly explored the effect of the overall synaptic strength, but the location of the transition depends also on the other two main parameters in the network: the strength of inhibition relative to excitation and the constant offset current.

The value of the ratio between excitation and inhibition strongly influences the critical coupling above which heterogeneous activity appears (**Fig. 5a**). Weaker inhibition reduces the critical coupling and favors heterogeneous activity (but note that the critical coupling remains nonzero for minimal values of relative inhibition), while stronger inhibition favors classical asynchrony. In a range of parameter values, the transition between classical and heterogeneous activity can therefore also be achieved by keeping the overall coupling constant and varying the ratio between inhibition and excitation.

The location of the transition also depends on the constant offset current that sets the resting potential in the neurons. Increasing this current decreases the critical coupling (**Fig. 5b**), so that the transition from classical to heterogeneous asynchrony can be induced by increasing this current without varying the synaptic strength. The range of synaptic couplings for which this is possible is, however, small (**Fig. 5b**), and the strength of the transition is weaker than when synaptic couplings are increased.

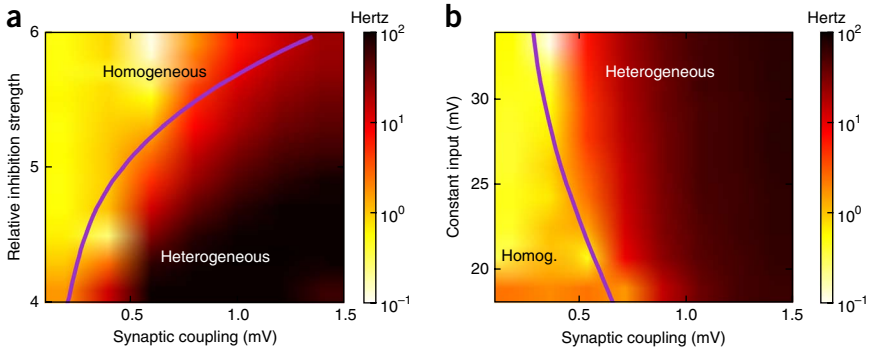

Figure 5 Influence of network parameters on asynchronous activity. (**a**) Nature of the asynchronous state as a function of overall synaptic coupling J and relative inhibition strength g. (**b**) Nature of the asynchronous state as function of overall synaptic coupling J and constant input current μ_0. The color plots display, on a logarithmic scale, the deviation of the mean firing rates from the firing rates predicted by equation (1) for the classical asynchronous state. Dark shades correspond to strong deviations, which are the signature of the heterogeneous asynchronous state (**Fig. 1a**). Light shades correspond to weak deviations and therefore to the homogeneous (homog.) asynchronous state. The purple lines display the analytical predictions for the transition.

DISCUSSION

We have shown that an unstructured, sparsely connected network of spiking neurons can display two different types of asynchronous activity at rest. For weak overall synaptic couplings and/or strong inhibition, the network is in the well-known asynchronous state, in which individual neurons fire irregularly at rates that are constant in time. If the overall synaptic couplings are strong and/or inhibition is just strong enough to balance excitation, a new type of resting state emerges. In that state the neurons still fire irregularly and asynchronously, but the firing rates of individual neurons fluctuate strongly in time and across neurons. We therefore call this new state the heterogeneous asynchronous state.

The two regimes of spontaneous asynchronous activity have different computational properties, as seen in their responses to temporally varying inputs. In the classical asynchronous state, the responses of different neurons are highly redundant, which favors a reliable transmission of information but limits the capacity of the network to perform nonlinear computations on the stimuli. In the heterogeneous asynchronous state, the responses of different neurons to the input instead strongly vary. This variability in the population degrades the transmission of information but provides a rich substrate for a nonlinear processing of the stimuli, as performed, for instance, in decision-making and categorization. Several experimental studies have recently reported high variability in the spontaneous state and multidimensional responses to stimuli similar to those found in the heterogeneous asynchronous regime[20,32,33,36–38].

Relation to previous studies

The heterogeneous asynchronous state described here is closely related to the chaotic state found in networks of rate units[15]. Whether and how such a state occurs in networks of spiking neurons has so far been unclear. Theoretical studies of spiking networks considered the instantaneous firing rate to be a population-averaged quantity (see, for example, refs. 3,9,25,39–41), rather than a property of an individual neuron, and typically examined networks consisting of a small number of populations. While these studies uncovered various instabilities of the classical asynchronous state (see, for example, refs. 9,12,42), they did not differentiate between two different types of asynchronous activity corresponding to the two states described in ref. 15.

The main conceptual novelty of the present work is to interpret the instantaneous firing rate as an individual property of each spiking neuron in the network. This change of viewpoint naturally leads to an instability of the classical asynchronous state mathematically analogous to the chaotic instability in rate networks[15]. Some qualitative aspects of the transition described here differ, however, from that in ref. 15 owing to differences between models. In the network considered here, excitation and inhibition are segregated and firing rates are constrained to be positive, in contrast to those in ref. 15. As a consequence, mean firing rates are nonzero in the equilibrium state and increase strongly with synaptic coupling above the transition, which is not the case in ref. 15.

Why haven't previous numerical studies of spiking networks identified the heterogenous asynchronous state found here? The bursts of activity characteristic of that state (**Fig. 1a**) have in fact been observed before (see, for example, ref. 43), but they have not been related to fluctuating firing rates and to the instability found in ref. 15. In another study[44], the authors showed that adding clustering to the connectivity of the network leads to dynamics that share some features with the heterogeneous asynchronous state found here in a fully random network. These dynamics, however, originate in the bistability of the different clusters in the network rather than in an instability similar to that in ref. 15. Most other numerical studies of unstructured spiking networks have kept the overall synaptic coupling constant, whereas the heterogeneous state emerges most strongly by varying that parameter. A notable exception is ref. 45, where the authors varied the synaptic strength but seem not to have observed the transition reported here. A possible explanation is that they used strong relative inhibition that favors classical asynchrony. Another possible explanation is that the synaptic delays used in that study were shorter than the refractory period, in which case the overall coupling is artifactually reduced and the transition to the heterogeneous state inhibited (see Online Methods).

Generality of the findings

The evidence for two types of asynchronous state presented here is based on quantitative comparisons between numerical simulations and mathematical analyses of network dynamics. To allow mathematical tractability, we considered a network that includes several simplifying features and incorporates only the minimal biophysical constraints. The existence of two distinct types of asynchronous activity is, however, much more general and does not rely on the simplified features of the network or the validity of the approximations used in the mathematical analysis (discussed below). Indeed, the long-timescale dynamics of a large class of neural models can be described by rate models[29,30]. As shown in ref. 15, the only requirement for the appearance of heterogeneous activity is to randomly couple many nonlinear rate units.

The ability to predict the location of the transition in parameter space, however, depends on the features of the network and the accuracy of the approximations. Here we have studied a highly uniform network in which each neuron receives exactly the same number of incoming connections, so that the mean firing rates of all neurons are identical. Relaxing this assumption leads to a distribution of mean firing rates but does not impair mathematical tractability[25,46]. Another important assumption is that the total input to a typical neuron can be described by Gaussian white noise. Increasing the synaptic coupling leads to the progressive breakdown of this assumption, but that effect alone generates only small deviations in the firing rate[47]. Another issue is that the firing statistics of integrate-and-fire neurons in response to white noise inputs in general deviate from the Poisson process[48], which in turn implies nonvanishing correlation times in the synaptic input. In the subthreshold regime and for low firing rates, these correlation times are in general shorter than the long timescale of rate dynamics, but they can become comparable in the strongly hyperpolarizing regime (see, for example, the deviation between simulations and predictions for strong relative inhibition in **Fig. 5a**).

Computations and learning in the heterogenous state

In the heterogeneous asynchronous state, the interaction between the external input and the internal dynamics of the network leads to a high-dimensional output that can be exploited by a readout unit to perform complex computations[20,34,35]. The internal state of the network, however, strongly fluctuates in time, so that repeated presentations of the same external stimulus might lead to very different high-dimensional responses. Indeed, the dynamics are sensitive to the initial conditions of the network, so that any external noise impairs robust learning.

Previous studies have examined the computational properties of the heterogeneous state in networks of rate units[16,17,19], which suffer from the same sensitivity to initial conditions and noise. These studies

have found that the robustness of the network response can be increased without impairing its computational capabilities. One way of improving the robustness is to provide strong inputs that quench the temporal but not the spatial variability in the network[31,49]. This effect is exploited in learning schemes that add a feedback loop from the readout unit to the network[16,17]. Another possibility is to perform plastic changes on the level of the synapses internal to the network, in such a way that some innate trajectories of the dynamics are stabilized[19]. Transposing these learning schemes to the setting of networks of spiking neurons studied here is an important direction for future studies.

METHODS

Methods and any associated references are available in the online version of the paper.

ACKNOWLEDGMENTS

The author is grateful to N. Brunel, V. Hakim, D. Martí and M. Stern for discussions and feedback on the manuscript. This work was supported by France's Agence Nationale de la Recherche ANR-11-0001-02 PSL* and ANR-10-LABX-0087.

COMPETING FINANCIAL INTERESTS

The author declares no competing financial interests.

1. Shadlen, M.N. & Newsome, W.T. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **4**, 569–579 (1994).
2. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
3. Amit, D.J. & Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **7**, 237–252 (1997).
4. Troyer, T.W. & Miller, K.D. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput.* **9**, 971–983 (1997).
5. Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
6. Shu, Y., Hasenstaub, A. & McCormick, D.A. Turning on and off recurrent balanced cortical activity. *Nature* **423**, 288–293 (2003).
7. London, M., Roth, A., Beeren, L., Häusser, M. & Latham, P.E. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* **466**, 123–127 (2010).
8. Renart, A. *et al.* The asynchronous state in cortical circuits. *Science* **327**, 587–590 (2010).
9. Brunel, N. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* **8**, 183–208 (2000).
10. Monteforte, M. & Wolf, F. Dynamical entropy production in spiking neuron networks in the balanced state. *Phys. Rev. Lett.* **105**, 268104 (2010).
11. Monteforte, M. & Wolf, F. Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Phys. Rev. X* **2**, 041007 (2012).
12. Brunel, N. & Wang, X.-J. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* **11**, 63–85 (2001).
13. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
14. Dayan, P. & Abbott, L.F. *Theoretical Neuroscience* (MIT Press, 2001).
15. Sompolinsky, H., Crisanti, A. & Sommers, H. Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259–262 (1988).
16. Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
17. Sussillo, D. & Abbott, L.F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).
18. Toyoizumi, T. & Abbott, L.F. Beyond the edge of chaos: amplification and temporal integration by recurrent networks in the chaotic regime. *Phys. Rev. E* **84**, 051908 (2011).
19. Laje, R. & Buonomano, D.V. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
20. Buonomano, D.V. & Maass, W. State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
21. Wainrib, G. & Touboul, J. Topological and dynamical complexity of random neural networks. *Phys. Rev. Lett.* **110**, 118101 (2013).
22. Sussillo, D. & Barak, O. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Comput.* **25**, 626–649 (2013).
23. Wallace, E., Maei, H.R. & Latham, P.E.L. Randomly connected networks have short temporal memory. *Neural Comput.* **25**, 1408–1439 (2013).
24. Churchland, M.M. & Abbott, L.F. Two layers of neural variability. *Nat. Neurosci.* **15**, 1472–1474 (2012).
25. Brunel, N. & Hakim, V. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.* **11**, 1621–1671 (1999).
26. Amit, D.J. & Tsodyks, M.V. Quantitative study of attractor neural network retrieving at low spike rates. 1. Substrate spikes, rates and neuronal gain. *Network* **2**, 259–273 (1991).
27. Rajan, K. & Abbott, L.F. Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.* **97**, 188104 (2006).
28. Timme, M., Wolf, F. & Geisel, T. Topological speed limits to network synchronization. *Phys. Rev. Lett.* **92**, 074101 (2004).
29. Ostojic, S. & Brunel, N. From spiking neuron models to linear-nonlinear models. *PLoS Comput. Biol.* **7**, e1001056 (2011).
30. Schaffer, E.S., Ostojic, S. & Abbott, L.F.A. Complex-valued firing-rate model that approximates the dynamics of spiking networks. *PLoS Comput. Biol.* **9**, e1003301 (2013).
31. Rajan, K., Abbott, L.F. & Sompolinsky, H. Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E* **82**, 011903 (2010).
32. Churchland, M.M., Yu, B.M., Sahani, M. & Shenoy, K.V. Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* **17**, 609–618 (2007).
33. Rabinovich, M., Huerta, R. & Laurent, G. Neuroscience. Transient dynamics for neural processing. *Science* **321**, 48–50 (2008).
34. Rigotti, M., Rubin, D.B.D., Wang, X.-J. & Fusi, S. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front. Comput. Neurosci.* **4**, 24 (2010).
35. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
36. Mazor, O. & Laurent, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
37. Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I. & Shenoy, K.V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
38. Churchland, M.M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* **13**, 369–378 (2010).
39. Ginzburg, I. & Sompolinsky, H. Theory of correlations in stochastic neural networks. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **50**, 3171–3191 (1994).
40. Gerstner, W. Time structure of the activity in neural network models. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **51**, 738–758 (1995).
41. Gerstner, W. Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. *Neural Comput.* **12**, 43–89 (2000).
42. Wong, K.F. & Wang, X.-J. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
43. Rangan, A.V. & Young, L.S. Emergent dynamics in a model of visual cortex. *J. Comput. Neurosci.* **35**, 155–167 (2013).
44. Litwin-Kumar, A. & Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* **15**, 1498–1505 (2012).
45. Tetzlaff, T., Helias, M., Einevoll, G.T. & Diesmann, M. Decorrelation of neural-network activity by inhibitory feedback. *PLoS Comp. Biol.* **8**, e1002596 (2012).
46. Roxin, A., Brunel, N., Hansel, D., Mongillo, G. & van Vreeswijk, C. On the distribution of firing rates in networks of cortical neurons. *J. Neurosci.* **31**, 16217–16226 (2011).
47. Helias, M., Deger, M., Rotter, S. & Diesmann, M. Instantaneous non-linear processing by pulse-coupled threshold units. *PLoS Comp. Biol.* **6**, e1000929 (2010).
48. Ostojic, S. Interspike interval distributions of spiking neurons driven by fluctuating inputs. *J. Neurophysiol.* **106**, 361–373 (2011).
49. Rajan, K., Abbott, L.F. & Sompolinsky, H. Inferring stimulus selectivity from the spatial structure of neural network dynamics. in Advances in Neural Information Processing Systems 23 (eds. Lafferty, J., Williams, C.K.I., Shawe-Taylor, J., Zemel, R. & Culotta A.) 1975–1983 (Curran, 2010).

ONLINE METHODS

Network of spiking neurons. The network considered here is identical to the network studied in ref. 9. It consists of N leaky integrate-and-fire neurons, with $N = 10{,}000$ in the simulations. A fraction $f = 0.8$ of these neurons are excitatory, the remaining inhibitory. The membrane potential dynamics of the ith neuron are given by

$$\tau_m \frac{dV_i}{dt} = -V_i + \mu_0 + RI_i(t) + \mu_{ext}(t) \quad (3)$$

where $\tau_m = 20$ ms is the membrane time constant, μ_0 is a constant offset current, RI_i is the total synaptic input from within the network and $\mu_{ext}(t)$ is a time-dependent external current. When the membrane potential crosses the threshold $V_{th} = 20$ mV, an action potential is emitted and the membrane potential is reset to the value $V_r = 10$ mV. The dynamics resume after a refractory period of $\tau_r = 0.5$ ms. The total synaptic input to the ith neuron is

$$RI_i(t) = \tau_m \sum_j J_{ij} \sum_k \delta(t - t_k^{(j)} - \Delta) \quad (4)$$

where J_{ij} is the amplitude of the postsynaptic potential evoked in neuron i by an action potential occurring in neuron j, $\Delta = 0.55$ ms is the synaptic delay and δ is the delta function. Note that with the above implementation of refractoriness, it is important to have synaptic delays longer than the refractory period. If the delays are shorter, spikes that reach a neuron while it is refractory do not have any effect, and the overall coupling is effectively reduced. This is an artifact of the combination of instantaneous synapses and the implementation of the refractory period.

Each neuron receives C incoming connections, C being much smaller than the total number of neurons N, so that the connectivity is sparse. We kept the value of C fixed at 1,000. A fraction $f = 0.8$ of these connections are excitatory, the remaining inhibitory. All excitatory synapses have the same strength J, and all inhibitory synapses have the same strength $-gJ$, where g quantifies the relative strength of inhibition compared to excitation. As a result, the network is highly uniform, all neurons being statistically equivalent. Each neuron receives, in addition, a constant external input μ_0 that is equivalent to setting the resting potential. The values of J, g and μ_0 were systematically varied and are indicated in the figure legends. In contrast with the implementation by Brunel (2000)[9], there is no external noise, so the network is fully deterministic.

The external input $\mu_{ext}(t)$ is zero in all figures except **Figure 4**. In that figure, $\mu_{ext}(t)$ is nonzero for 2,000 randomly chosen excitatory neurons and consists of a superposition of sinusoids shown in **Figure 4a**.

Mathematical analysis of network dynamics. To mathematically analyze the network dynamics, we used mean field theory[3,9,25]. Mean field theory describes the regime in which each neuron receives a large number of inputs at every time step, each input being in isolation too small to generate an action potential. In such a situation, the total synaptic input to each neuron can be approximated by a Gaussian white noise process that is independent across neurons. Assuming that each neuron emits spikes as a Poisson process of constant rate ν_0 (identical for all neurons, as all neurons receive exactly the same number of incoming connections), the mean μ and s.d. σ of the equivalent white noise input are given[3] by

$$\mu = \tau_m \nu_0 CJ\left(f - (1-f)g\right) + \mu_0 \quad (5)$$

$$\sigma^2 = \tau_m \nu_0 CJ^2\left(f + (1-f)g^2\right) \quad (6)$$

The mean μ of the input vanishes when the relative strength g of inhibition is equal to $f/(1-f) = 4$. This corresponds to perfect balance between excitation and inhibition. Here we concentrate on the inhibition-dominated regime $g > 4$.

The equilibrium firing rate ν_0 in the asynchronous state is given by the self-consistent equation

$$\nu_0 = F\left(\mu(\nu_0), \sigma^2(\nu_0)\right) \quad (7)$$

where F is the current-to-rate transfer function of the LIF neuron receiving a white noise input[3,26]

$$F(\mu,\sigma^2) = \left[\tau_r + 2\tau_m \int_{\frac{V_r-\mu}{\sigma}}^{\frac{V_{th}-\mu}{\sigma}} du e^{u^2} \int_{-\infty}^{u} dv e^{-v^2}\right]^{-1} \quad (8)$$

To assess the stability of the equilibrium state, we allow the firing rate of every neuron to deviate by a different amount from the equilibrium firing rate, and determine whether the network dynamics suppress or amplify this perturbation. For such a heterogeneous perturbation, the total input to a given neuron can still be approximated by a Gaussian white noise, but now the mean and s.d. of this process vary across neurons. Writing the firing rate of the jth neuron at time t as $\nu_0 + \nu_j(t)$, the mean μ_i and s.d. σ_i of the input to neuron i read

$$\mu_i(t) = \mu + \tau_m \sum_j J_{ij} \nu_j(t) \quad (9)$$

$$\sigma^2{}_i(t) = \sigma^2 + \tau_m \sum_j J_{ij}^2 \nu_j(t) \quad (10)$$

The stability of the asynchronous state to slow (zero frequency), heterogeneous perturbations is determined by linearizing equation (1) for each neuron in the network. The linearized dynamics for the firing rate of neuron i read

$$\nu_i(t + dt) = \sum_j G_{ij} \nu_j(t) \quad (11)$$

where G is the stability matrix given by

$$G_{ij} = \tau_m (J_{ij}\gamma_\mu + J_{ij}^2 \gamma_{\sigma^2}) \quad (12)$$

Here γ_μ and γ_{σ^2} are the gains of the current-to-rate transfer function with respect to a variation in mean and variance of the input:

$$\gamma_\mu = \left.\frac{\partial F}{\partial \mu}\right|_{\mu(\nu_0),\sigma^2(\nu_0)} \quad (13)$$

$$\gamma_{\sigma^2} = \left.\frac{\partial F}{\partial \sigma^2}\right|_{\mu(\nu_0),\sigma^2(\nu_0)} \quad (14)$$

Equation (11) indicates that the heterogeneous perturbation is amplified if the matrix G possesses eigenvalues larger than unity. The matrix G is an asymmetric, sparse random matrix. The distribution of its elements varies from column to column, depending on whether the column corresponds to inputs from an excitatory or an inhibitory neuron. Eigenvalue distributions for matrices of this type have been studied in refs. 27,28, and we exploit here the results of these studies. The matrix G possesses an eigenvalue

$$\tau_m \left(CJ(f - (1-f)g)\gamma_\mu + CJ^2(f + (1-f)g^2)\gamma_{\sigma^2}\right)$$

associated with the unit eigenvector, which corresponds to a homogeneous perturbation. This eigenvalue is always negative in the inhibition-dominated regime studied here[3] because the value of γ_{σ^2} strongly decreases with J so that the positive contribution proportional to J^2 becomes negligible.

The other eigenvalues correspond to heterogeneous perturbations and for $N \to \infty$ are densely distributed within a circle in the complex plane centered at the origin. For a matrix with independently distributed elements, the radius λ_{max} is given by[27]:

$$\lambda_{max} = \sqrt{N(f\sigma_E^2 + (1-f)\sigma_I^2)} \quad (15)$$

where σ_E^2 and σ_I^2 are the variances of the elements in the excitatory and inhibitory columns.

Here the distribution of elements is not independent because of the constraint that every neuron receives exactly the same number of inputs. This constraint can, however, be neglected for large N, so that the distribution of elements in each column is approximately independent and binomial: each excitatory synapse has

a strength J with probability C/N, and each inhibitory synapse has a strength $-gJ$ with probability C/N. In the limit $N\to\infty$, this yields

$$\lambda_{\max} = \tau_m J\sqrt{C}\sqrt{f(\gamma_\mu + J\gamma_{\sigma^2})^2 + (1-f)g^2(-\gamma_\mu + gJ\gamma_{\sigma^2})^2} \quad (16)$$

Note that $\lambda_{\max}$ depends on the synaptic coupling explicitly, but also implicitly through γ_μ and γ_{σ^2}.

Poisson network. In the Poisson network, each neuron is described by a firing rate that fully specifies its activity. The dynamics of the firing rate of neuron i are given by

$$\tau\frac{dv_i}{dt} = -v_i + F\left(\mu_i(t),\sigma_i^2(t)\right) \quad (17)$$

where F is the current-to-rate transfer function of an LIF neuron receiving white noise, and, as in equations (9) and (10),

$$\mu_i(t) = \mu + \tau_m\sum_j J_{ij}v_j(t) \quad (18)$$

$$\sigma^2{}_i(t) = \sigma^2 + \tau_m\sum_j J_{ij}^2 v_j(t) \quad (19)$$

The effective timescale τ in the rate model can be related to the slowest dynamical timescale present in the corresponding network of integrate-and-fire neurons. Its value depends on the parameters of the network[29,30] and is typically shorter than the membrane timescale. For simplicity, we treat it as an arbitrary constant and set its value to the membrane timescale τ_m.

At equilibrium, the firing rate in the Poisson network is given by equation (1). Linearizing the dynamics around equilibrium leads to

$$v_i(t+dt) = \sum_j G_{ij}v_j(t) \quad (20)$$

with the matrix G given in equation (12), so that the linear stability analysis in the Poisson network is identical to that in the LIF network.

Simulations and data analysis. The mean firing rates ν_0 in **Figures 1a** and **2a** were computed by simulating $T = 20$ s of network dynamics.

The average population activity $n(t)$ in **Figures 1c**,**d** and **2c**,**d** were determined the number of spikes in the full network per 1-ms bin. The autocorrelation of the population activity is defined as

$$A_{\text{pop}}(\tau) = \frac{1}{\nu_0^2 T}\sum_{t=1}^{T}(n(t+\tau)-\nu_0)(n(t)-\nu_0) \quad (21)$$

In **Figure 3**, the instantaneous firing rate $v_i(t)$ of each LIF neuron was computed by convolving its spike train with a 50-ms Gaussian filter. The average autocorrelation of firing rates is defined as

$$A_{\text{rate}}(\tau) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(v_i(t+\tau)-\nu_0)(v_i(t)-\nu_0) \quad (22)$$

The average autocorrelation of spike trains $n_i(t)$ (computed in 1-ms bins) is defined as

$$A_{\text{spikes}}(\tau) = \frac{1}{\nu_0 NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(n_i(t+\tau)-\nu_0)(n_i(t)-\nu_0) \quad (23)$$

The firing rates in **Figure 4** were computed by convolving the spike trains with $s(t) = \exp(-t/\tau_r) - \exp(-t/\tau_d)$ for $t > 0$, where $\tau_r = 1$ ms and $\tau_d = 100$ ms are, respectively, the rise and decay timescales of an NMDA-based synapse. The principal component analysis in **Figure 4** was determined by diagonalizing the covariance matrix of the firing rates obtained in response to the successive presentation of the two stimuli.

doi:10.1038/nn.3658