J Comput Neurosci (2014) 36:469–481
DOI 10.1007/s10827-013-0481-5

How well do mean field theories of spiking quadratic-integrate-and-fire networks work in realistic parameter regimes?

Agnieszka Grabska-Barwińska · Peter E. Latham

Received: 16 May 2013 / Revised: 9 August 2013 / Accepted: 11 September 2013 / Published online: 5 October 2013

Abstract We use mean field techniques to compute the distribution of excitatory and inhibitory firing rates in large networks of randomly connected spiking quadratic integrate and fire neurons. These techniques are based on the assumption that activity is asynchronous and Poisson. For most parameter settings these assumptions are strongly violated; nevertheless, so long as the networks are not too synchronous, we find good agreement between mean field prediction and network simulations. Thus, much of the intuition developed for randomly connected networks in the asynchronous regime applies to mildly synchronous networks.

Keywords Recurrent network · Synchronization · Quadratic integrate and fire neuron · Theta neuron · Random networks · Mean field theory

1 Introduction

A long term goal in computational neuroscience is to understand the relationship between network parameters–especially connectivity–and network behavior. This relationship has been studied extensively in randomly connected networks of excitatory and inhibitory neurons (Amit and Brunel 1997a, b; van Vreeswijk and Sompolinsky 1998; Brunel 2000; Latham et al. 2000a; Hansel and Mato 2001; Lerchner et al 2006a, b; Renart et al. 2010; Hertz 2010), for which the following picture has emerged: randomly connected networks can operate in a relatively small variety of regimes, with the regime characterized mainly by the degree and type of synchrony (Brunel 2000; Hansel and Mato 2001). Of these, only the completely asynchronous regime (in which the cross-correlograms are vanishingly small) is well understood quantitatively. This regime, however, is hard to access in realistic networks: parameters have to be carefully adjusted to reduce synchrony among neurons, and it is next to impossible to eliminate synchrony altogether. Therefore, here we ask: how well does analysis designed to work in the asynchronous regime apply to synchronous networks? The answer, not surprisingly, depends on the degree of synchrony. Somewhat more surprisingly, even when networks are reasonably highly synchronous, the asynchronous analysis makes accurate quantitative predictions of the mean firing rates of the excitatory and inhibitory populations, and makes good qualitative predictions of the distribution of firing rates. Thus, even in the relatively synchronous regime, we can rely on these models to provide intuition about the dynamics of randomly connected excitatory and inhibitory networks.

Our analysis is based on the quadratic integrate and fire neuron (Ermentrout and Kopell 1986; Ermentrout 1996; Gutkin and Ermentrout 1998; Brunel and Latham 2003), chosen because it provides a very good description of the spiking dynamics of type I neurons at low firing rate (Ermentrout and Kopell 1986), and because there is a reasonably accurate analytic expression for the firing rate of these neurons as a function of synaptic drive (Brunel and Latham 2003). The analytic expression is available for

Action Editor: Brent Doiron

A. Grabska-Barwińska (✉) · P. E. Latham
Gatsby Computational Neuroscience Unit,
University College London,
London, UK
e-mail: agnieszka@gatsby.ucl.ac.uk

P. E. Latham
e-mail: pel@gatsby.ucl.ac.uk

essentially arbitrary synaptic time constant; we took advantage of this to test a range of synaptic time constants, and so broaden the validity of our conclusions. Although we focus on the quadratic integrate and fire neuron, our approach can be extended, at least in principle, to any single neuron model. However, in most cases–especially those with conductance based synapses–approximate schemes must be used to characterize the relationship between synaptic drive and firing rate (Shriki et al. 2003).

Consistent with our goal of testing the validity of the asynchronous assumption, when carrying out our analysis we assume that spike trains emitted by any two neurons are uncorrelated. With this assumption, the network equilibria are fully described by a set of algebraic equations that relate the firing rate of each neuron to the firing rates of all other neurons in the network. We solve them using a mean field approach based on the self-consistent signal to noise analysis of Shiino and Fukai (1992, 1993), and compare this solution to simulations in a regime in which our assumptions are violated, and the neurons are not asynchronous.

2 Mean field analysis of a network of quadratic integrate and fire neurons

Our goal is to compute the distribution of firing rates in a recurrently connected network of excitatory and inhibitory neurons. Our starting point is a set of equations describing the time evolution of the membrane potential of each neuron in the network. We then proceeded in two steps. First, we reduce the time-dependent membrane potential equations to a set of algebraic firing rate equations. Second, we solve them using mean field techniques.

The first step depends critically on both the single neuron model and the synaptic coupling. For the former we use the quadratic integrate and fire neuron (Ermentrout and Kopell 1986; Ermentrout 1996; Gutkin and Ermentrout 1998). For the latter we use current-based synapses, and assume that each spike produces an instantaneous rise in membrane potential followed by an exponential decay (Koch 1998). In the limit that the network is large–the limit of interest here–this input is reasonably well approximated by filtered white noise. Thus, to compute the firing rate of our model neurons we simply need to compute the firing rate of a quadratic integrate and fire neuron receiving filtered white noise. Fortunately, the firing rate of such a neuron (or at least an approximation to it) has been computed as a function of the mean and variance of the fluctuating input (Brunel and Latham 2003).

The remainder of this section proceeds as follows: in Section 2.1 we write down the equations describing the single neuron dynamics, and provide an expression for the firing rate; in Section 2.2 we write down the full network equations; and in Section 2.3 we derive the mean field equations. Then, in Section 3, we compare the predictions of our mean field model to numerical simulations of the network equations.

2.1 Single neuron dynamics

Using V and h to denote the membrane potential of a neuron and its synaptic drive, respectively, the single neuron dynamics of a quadratic integrate and fire neuron receiving fluctuating input can be written

$$\frac{\tau_m}{V_{th}-V_r}\frac{dV}{dt} = \frac{(V-\overline{V})^2}{(V_{th}-V_r)^2} + \mu + h(t) \tag{2.1a}$$

$$\tau_s \frac{dh}{dt} = -h + \sigma \tau_m^{1/2}\xi(t). \tag{2.1b}$$

Here τ_m and τ_s are the membrane and synaptic time constants, respectively, V_r, V_{th} and μ are parameters that set the neuron firing rate in the absence of the synaptic drive ($h(t) = 0$), $\overline{V}$ is the voltage midway between the resting membrane potential, V_r, and the threshold, V_{th},

$$\overline{V} \equiv \frac{V_r + V_{th}}{2}, \tag{2.2}$$

$\xi(t)$ corresponds to Gaussian white noise,

$$\overline{\xi(t)\xi(t')} = \delta(t-t'), \tag{2.3}$$

and σ sets the overall level of the noise. Here and in what follows, an overline indicates an average over time. Because of the quadratic dependence on V, the voltage can reach $+\infty$ in finite time; when that happens, a spike is emitted, and the voltage is reset to $-\infty$. (To handle the infinities in our numerical simulations, we change to angular variables; see Appendix B).

Poisson spikes at a sufficiently high firing rate produce synaptic drive that corresponds approximately to white noise (Walsh 1981; Tuckwell 1988). Thus, a neuron embedded in a network in which each neuron receives a large number of inputs, as is the case in our networks, with spike statistics that are approximately Poisson would receive synaptic drive that looks like white noise (as in Eq. (2.1b)). If we simply assume that neurons are Poisson, we can compute both the mean drive, μ, and the fluctuations in the drive, σ^2, to any particular post-synaptic neuron as a function of the firing rates of its pre-synaptic neurons. Then, if we could compute the firing rate of the postsynaptic neuron as a function of μ and σ^2, we could derive a set of algebraic equations whose solution tells us the firing rate of every neuron in the network. Unfortunately, it is not, as far as we know, possible to compute the single neuron firing rate exactly. However, an approximate expression for the firing

rate exists (Brunel and Latham 2003). That approximate rate, denoted $\nu_{QIF}(\mu, \sigma^2)$, is given by

$$\nu_{QIF}(\mu, \sigma^2) = \frac{\nu_{0S} + (\tau_s/\tau_m)^2 \nu_{0L} \rho_{2S}/\rho_{2L}}{1 + (\tau_s/\tau_m)\rho_{2S} + (\tau_s/\tau_m)^2 \rho_{2S}/\rho_{2L}}. \tag{2.4}$$

The various quantities that appear in this expression depend on μ and σ^2 (a dependence that is suppressed for clarity) via

$$\nu_{0S} = \frac{1}{\pi\tau_m}\left[\int_{-\infty}^{\infty} \frac{d\xi}{\pi^{1/2}} \exp\left[-\mu\xi^2 - \sigma^4\xi^6/48\right]\right]^{-1} \tag{2.5a}$$

$$\rho_{2S} = \pi\sigma^2 \frac{\tau_m \nu_{0S}}{2} \int_{-\infty}^{\infty} \frac{d\xi}{\pi^{1/2}} \xi^2 \exp\left[-\mu\xi^2 - \sigma^4\xi^6/48\right] \tag{2.5b}$$

$$\nu_{0L} = \frac{\mu^{1/2}}{\pi\tau_m} \tag{2.5c}$$

$$\rho_{2L} = \frac{\sigma^2}{16\mu^2}. \tag{2.5d}$$

With a small amount of algebra, Eq. (2.4) reduces exactly to Eq. (5.2) of Brunel and Latham (2003).

To determine how well the approximate expression given in Eq. (2.4) captures the true firing rate, we performed simulations with a range of μ and σ^2. The results are shown in Fig. 1. Agreement is best when τ_s is small; deviations were less than 1 Hz and, except at very low firing rates, relative deviations were a few percent. Agreement got worse as τ_s increased, with deviations up to 5 Hz for $\tau_s = 100$ ms. However, as can be seen from the solid lines in Fig. 1, our networks mainly operate where the approximation is good. As a result, even the largest inaccuracies ($\tau_s = 100$ ms) have a minor effect on firing rate distributions. Thus, we do not expect the fact that we have an approximate firing rate to have much effect on the accuracy of our mean field theory.

2.2 Network equations

We now turn to a network of excitatory (E) and inhibitory (I) neurons, which, in addition to recurrent interactions, receives excitatory input from an external population (X); see Fig. 2. The network equations are very similar to Eq. (2.1a); the main difference (besides an explosion of subscripts) is that the white noise term in Eq. (2.1b) is replaced by synaptic drive and an offset,

$$\frac{\tau_m}{V_{th} - V_r}\frac{dV_{Li}}{dt} = \frac{(V_{Li} - \overline{V})^2}{(V_{th} - V_r)^2} + \mu_L + h_{Li}(t) \tag{2.6a}$$

$$\tau_s \frac{dh_{Li}}{dt} = -h_{Li} + \delta\mu_{Li} + \sum_{M=E,I,X} \frac{\tau_m}{K_M^{1/2}} \sum_{j=1}^{N_M} \sum_l J_{LM}^{ij} \delta\left(t - t_{Mj}^l\right). \tag{2.6b}$$

Here J_{LM}^{ij} is the connection strength from neuron j of type M to neuron i of type L (note that J_{LI}^{ij} is negative), N_M is the number of neurons of type M, K_M is the average number of connections from neurons of type M, t_{Mj}^l is the time of the l^{th} spike emitted by neuron j of type M, and $\delta(\cdot)$ is the Dirac δ-function. The external neurons are taken to be Poisson with constant firing rate ν_X. The factor $K_M^{-1/2}$ that appears in Eq. (2.6b) ensures that for sufficiently large networks, the

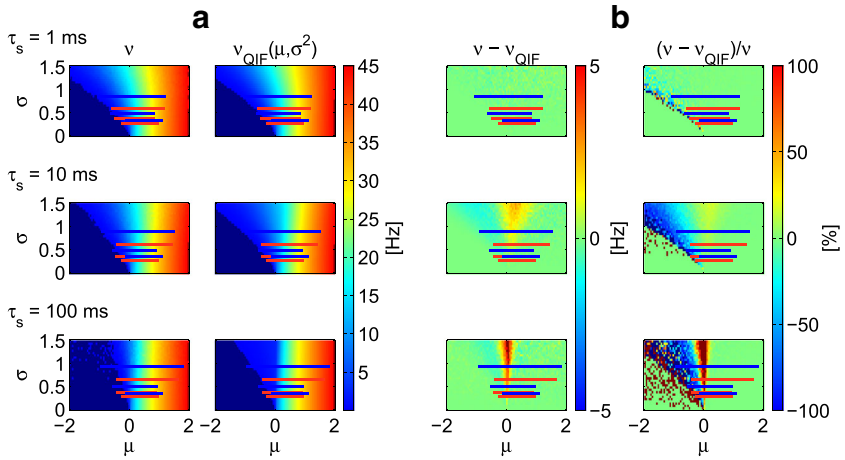

Fig. 1 Comparison of mean field predictions and simulated firing rates for a quadratic integrate and fire neuron. **a** Firing rate versus μ and σ for $\tau_m = 10$ ms and $\tau_s = 1, 10$ and 100 ms. The left panel is from 10000 seconds of simulation of Eq. (2.1b); the right panel from Eq. (2.4). **b** Absolute (*left panel*) and relative (*right panel*) error between the simulations and analytic expression. The relative error is thresholded at ±100 %, to prevent exposure of the least interesting errors, which were due to division by nearly zero firing rates. Superimposed on all images are thick horizontal lines indicating the working range of our simulations–every line indicates the range of inputs to an excitatory (*red*) or inhibitory (*blue*) population ($h_L \pm 3\Delta_{h_L}$, see Eq. (2.13b)). The three pairs of lines correspond to the three networks we tested: the disconnected network (Fig. 5), exhibiting the lowest σ; the "default" network (Fig. 6), exhibiting intermediate σ; and the strongly connected network (Fig. 7), exhibiting the highest σ

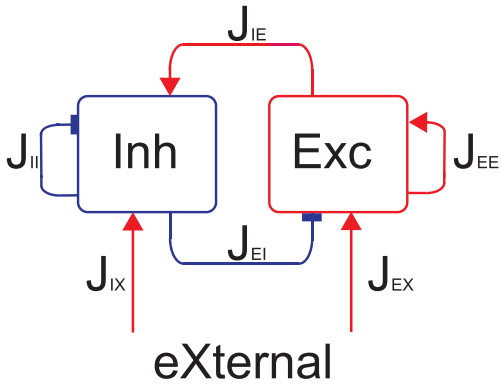

Fig. 2 Network architecture. *Red* indicates excitatory connections, *blue* indicates inhibitory connections. As in the main text, *E*, *I* and *X* correspond to excitatory, inhibitory and external connections, respectively

distribution of firing rates is independent of network size (van Vreeswijk and Sompolinsky 1998).

We take the connectivity matrix to be random and sparse,

$$J_{LM}^{ij} = \begin{cases} J_{LM}\left(1 + \Delta\, \zeta_{LM}^{ij}\right) & \text{with probability } \epsilon \\ 0 & \text{with probability } 1-\epsilon \end{cases} \tag{2.7}$$

where Δ is the standard deviation of the nonzero synaptic weights, and the ζ_{LM}^{ij} are a set of uncorrelated, zero mean, unit variance random variables, chosen so that the randomness does not cause the connection strength to change sign; i.e., chosen so that $1 + \Delta\, \zeta_{LM}^{ij} \geq 0$. Note that the connection probability, ϵ, is independent of neuron type, so $K_M = \epsilon N_M$. Finally, we let $\delta\mu_{Li}$ be a Gaussian random variable with variance $\Delta_{\mu_L}^2$,

$$\delta\mu_{Li} \sim \mathcal{N}\left(0, \Delta_{\mu_L}^2\right). \tag{2.8}$$

A list of parameters for our default network is given in Table 1. For some sets of simulations (Figs. 5, 6, 7 and 8), we varied connections strengths; for others (Fig. 9), we varied network size.

2.3 Mean field analysis

The first step in analyzing the network equations is to break the synaptic drive in Eq. (2.6b) into time-independent and temporally fluctuating pieces; the former contributes to the mean synaptic drive (μ in Eq. (2.1a)), the latter to the fluctuations in the synaptic drive (σ in Eq. (2.1b)). This gives us

$$h_{Li}(t) = \overline{h}_{Li} + \delta h_{Li}(t) \tag{2.9}$$

where, recall, an overline represents a time average. Averaging Eq. (2.6b) over time, and noting that $\overline{dh_{Li}/dt} = 0$, we see that

$$\overline{h}_{Li} = \delta\mu_{Li} + \sum_{M=I,E,X} \frac{\tau_m}{K_M^{1/2}} \sum_{j,l} J_{LM}^{ij}\, \overline{\delta\left(t - t_{Mj}^l\right)}. \tag{2.10}$$

Computing the time average of the δ-functions is straightforward,

$$\overline{\delta\left(t - t_{Mj}^l\right)} = \lim_{T\to\infty} \frac{1}{T} \int_0^T dt \sum_l \delta\left(t - t_{Mj}^l\right) = \nu_{Mj} \tag{2.11}$$

Table 1 Default parameters used in the mean field calculations and simulations

Parameter	Value	Description
τ_m	10 ms	Membrane time constant
τ_s	1, 10 or 100 ms	Synaptic time constant
μ_E, μ_I	-0.25	Mean synaptic drive
$\Delta\mu_E$, $\Delta\mu_I$	0.2	Standard deviation of synaptic drive
$\epsilon\ (=K/N)$	0.1	Connection probability
Δ	0.2	Standard deviation of non-zero connection strengths
N_E	16,000	Number of excitatory neurons
N_I	4,000	Number of inhibitory neurons
N_X	2,000	Number of external neurons
ν_X	15 Hz	Mean firing rate of external neurons
J_{EE}, J_{EI}, J_{EX}	0.25, -0.6, 1.2	Mean synaptic weights onto excitatory neurons
J_{IE}, J_{II}, J_{IX}	0.35, -0.9, 1.5	Mean synaptic weights onto inhibitory neurons

Typical EPSPs for the default network range from 0.07 to 0.10 mV, and typical IPSPs from −0.35 to −0.52 mV. For our model, average PSP size from a cell of type L to a cell of type M is approximately equal to $(V_{th} - V_r)(\tau_s/\tau_m)^{\tau_m/(\tau_m-\tau_s)} J_{LM}/K^{1/2}$; see Eq. (14) of (Latham et al. 2000a) with $\mathcal{E}_j$ replaced by V_{th}, W_{ij} replaced by $J_{LM}/K^{1/2}$, and r_j set to 1. For some sets of simulations, we changed the recurrent connections strengths (Figs. 5, 6, 7 and 8) and the number of neurons (Fig. 9)

where ν_{Mj} is the firing rate of neuron j of type M. The second equality follows from the fact that the integral over time counts spikes, and by definition firing rate is the number of spikes divided by time. Inserting this expression into Eq. (2.10) leads to

$$\overline{h}_{Li} = \delta\mu_{Li} + \sum_{M=E,I,X} \frac{\tau_m}{K_M^{1/2}} \sum_j J_{LM}^{ij}\, \nu_{Mj}. \tag{2.12}$$

The dependence on index appears nontrivial. To deal with it analytically, we make one of our main mean field approximations, which is that the term $\sum_j J_{LM}^{ij}\, \nu_{Mj}$ can be treated as a Gaussian random variable with respect to index i. According to the central limit theorem, this approximation is valid if the terms in the sum, $J_{LM}^{ij}\, \nu_{Mj}$, are sufficiently weakly correlated. We assume that they are; the extent to which our results are consistent with simulations is a partial measure of the validity of this approximation (it is only a partial measure because we make other approximations).

With the Gaussian approximation, Eq. (2.12) becomes

$$\overline{h}_{Li} = h_L + \Delta_{h_L}\eta_{Li} \tag{2.13a}$$

$$h_L \equiv \sum_{M=I,E,X} K_M^{1/2} J_{LM} \tau_m \nu_M \tag{2.13b}$$

where ν_M is the average firing rate of population M (see Eq. (2.15a) below) and the η_{Li} are zero mean, unit variance Gaussian random variables with respect to index i. We show in Appendix A that the total variance, $\Delta_{h_L}^2$, can be expressed in terms of network parameters as

$$\Delta_{h_L}^2 = \Delta_{\mu_L}^2 + \sum_{M=E,I,X} J_{LM}^2 \left(1 + \Delta^2 - \epsilon\right) \tau_m^2 \nu_M^2. \tag{2.14}$$

The population averaged firing rate, ν_M, and the second moment of the firing rate, ν_M^2, have natural definitions,

$$\nu_M \equiv N_M^{-1} \sum_j \nu_{Mj} \tag{2.15a}$$

$$\nu_M^2 \equiv N_M^{-1} \sum_j \nu_{Mj}^2\,. \tag{2.15b}$$

We now turn to the second term in Eq. (2.9), $\delta h_{Li}(t)$. Inserting Eq. (2.9) into (2.6b) and using Eq. (2.12) for $\overline{h}_{Li}$, we find that $\delta h_{Li}(t)$ evolves according to

$$\tau_s \frac{d\delta h_{Li}}{dt} + \delta h_{Li} = \sum_{M=E,I,X} \frac{\tau_m}{K_M^{1/2}} \sum_j J_{LM}^{ij} \times \left(\sum_l \delta\left(t - t_{Mj}^l\right) - \nu_{Mj} \right). \tag{2.16}$$

To solve this equation, we need the temporal statistics of the the right hand side. Consistent with our mean field approximation, we assume that it is a Gaussian process, so all we need is its covariance (by construction the time average is zero). To compute that, we make two approximations: the neurons are independent, and they fire with Poisson statistics. Neither of these are totally accurate; spike times across different neurons are correlated, and quadratic integrate and fire neurons (like all realistic neurons) exhibit a refractory period. Fortunately, though, the refractory period produces relatively small errors in firing rates–so long as the rates aren't too high, the error is on the order of 10 % (Deger et al. 2012). And, because of almost complete cancellation between excitatory and inhibitory synaptic drives, neurons are only weakly correlated (Renart et al. 2010; Hertz 2010). Thus, while our mean field analysis won't perfectly describe the network, it should not be far off. Note that the alternative, computing the covariance structure self-consistently, is hard, and typically requires network simulations (Lerchner et al. 2006a, b).

With the independent Poisson assumption, the terms on the right hand side of Eq. (2.16) consist of sums of Poisson processes, each of which is δ-correlated. Thus, the sums are δ-correlated, and, after a small amount of algebra (carried out in Appendix A), we find that δh_{Li} evolves according to

$$\tau_s \frac{d\delta h_{Li}}{dt} = -\delta h_{Li} + \sigma_L \tau_m^{1/2}\, \xi_{Li}(t) \tag{2.17}$$

where $\xi_{Li}(t)$ is δ-correlated white noise (see Eq. (2.3)), and σ_L^2 is given by

$$\sigma_L^2 = \sum_{M=E,I,X} J_{LM}^2 \left(1 + \Delta^2\right) \tau_m \nu_M \tag{2.18}$$

(see Eq. (A.9)).

We can now rewrite Eq. (2.6b) as a stochastic differential equation. Using Eq. (2.9) for h_{Li}, Eq. (2.13a) for $\overline{h}_{Li}$ and Eq. (2.17) for the time evolution of $\delta h_{Li}(t)$, Eq. (2.6b) becomes

$$\frac{\tau_m}{V_{th} - V_r} \frac{dV_{Li}}{dt} = \frac{(V_{Li} - \overline{V})^2}{(V_{th} - V_r)^2} + \mu_L + h_L + \Delta_{h_L}\eta_{Li} + \delta h_{Li}(t) \tag{2.19a}$$

$$\tau_s \frac{d\delta h_{Li}}{dt} = -\delta h_{Li} + \sigma_L \tau_m^{1/2}\, \xi_{Li}(t)\,. \tag{2.19b}$$

These equations are identical in form to the single neuron dynamics given in Eq. (2.1b). Thus, we can use Eq. (2.4) to write down the firing rate of any particular neuron,

$$\nu_{Li} = \nu_{QIF}\left(\mu_L + h_L + \Delta_{h_L}\eta_{Li}, \sigma_L^2\right). \tag{2.20}$$

Equation (2.20) gives us a set of equations for the firing rates of the neurons. As such, it can give us the distribution of firing rates, but it cannot tell us which neuron has which rate. However, for a randomly connected network the distribution is all we need, since there is nothing to distinguish one neuron from another.

Our approach to finding the firing rate distribution is illustrated schematically in Fig. 3. For a given network, the mean synaptic drive to the population of neurons is modeled as a Gaussian distribution (bottom panel), with mean $\mu_L + h_L$

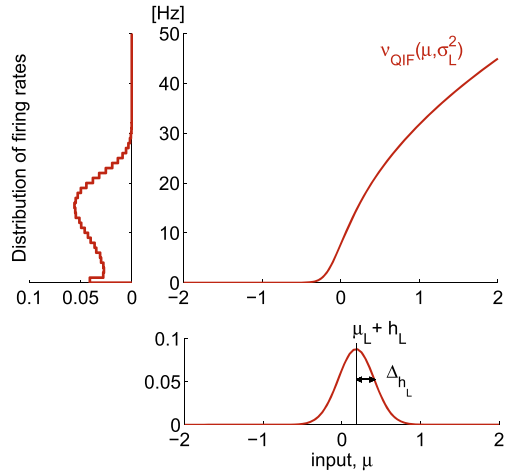

Fig. 3 Computing the distribution of firing rates from the distribution of neuronal inputs. Bottom: (Gaussian) probability distribution of mean input in a network of neurons. Middle: Firing rate of our quadratic integrate and fire neuron, $\nu_{QIF}(\mu,\sigma^2)$ (Eq. (2.4)) as a function of input, μ, with σ^2 fixed at σ_L^2. Left: The probability distribution over firing rate, $p(\nu)$, derived by mapping $p(\mu)$ through the nonlinearity $\nu = \nu_{QIF}\left(\mu,\sigma_L^2\right)$. The resulting distribution of firing rates, here binned at 1 Hz, is non-Gaussian, The parameters μ_L, h_L, Δ_{h_L} and σ_L^2 were set to the the mean field values for the excitatory neurons of the default network; the resulting firing rate distribution is thus identical to the top left plot in Fig. 6

and variance $\Delta_{h_L}^2$. Assume for the moment that we knew h_L and $\Delta_{h_L}^2$, as well as the variance associated with the temporal fluctuations, σ_L^2. We could, then, translate any particular mean input to a firing rate via $\nu(\mu) = \nu_{QIF}\left(\mu,\sigma_L^2\right)$; more importantly, we could translate the distribution of means (bottom panel in Fig. 3) to the distribution of firing rates (left panel).

We don't, though, know h_L, $\Delta_{h_L}^2$ and σ_L^2, as they depend on the firing rate distribution. Fortunately, this dependence is only via the first two moments: h_L and σ_L^2 depend on the first moments via Eqs. (2.13b) and (2.18), respectively, and $\Delta_{h_L}^2$ depends on the second moments via Eq. (2.14). Thus, ν_E, ν_I, ν_E^2, and ν_I^2, which constitute our order parameters, fully determine h_L, $\Delta_{h_L}^2$ and σ_L^2, and so they fully determine the distribution of firing rates. To determine the values of ν_E, ν_I, ν_E^2, and ν_I^2, we simply average over index. Fortunately, the only dependence on index in Eq. (2.20) is through η_{Li}, which is a zero mean, unit variance Gaussian random variable. In the large N limit, we may, therefore, replace averages over indices by integrals over continuous Gaussian variables; this leads to

$$\nu_L^k = \int d\eta \frac{e^{-\eta^2/2}}{\sqrt{2\pi}} \left[\nu_{QIF}\left(\mu_L + h_L + \Delta_{h_L}\eta, \sigma_L^2\right)\right]^k \qquad (2.21)$$

where k is either 1 or 2; $k = 1$ captures the first moment and $k = 2$ captures the second. Once we know the moments of the firing rate distribution, the second step–computing the entire distribution–amounts to computing (numerically) the integral

$$p(\nu_L) = \int d\eta \frac{e^{-\eta^2/2}}{\sqrt{2\pi}} \delta\left(\nu_L - \nu_{QIF}\left(\mu_L + h_L + \Delta_{h_L}\eta, \sigma_L^2\right)\right). \qquad (2.22)$$

Equation (2.21) constitutes our mean field equations for the network; once solved, Eq. (2.22) gives us the distribution of firing rates. In the next section we compare the mean field predictions of the the firing rate distributions with simulations. To help visualize the operating regime, we also show nullclines in average firing rate space. These are constructed as follows. First we solve (numerically) Eq. (2.21) with $k = 2$; that is, we solve for ν_E^2 and ν_I^2 in terms of ν_E and ν_I. Once we do that we are left with mean field equations for only the first moments, ν_E and ν_I. These correspond to Eq. (2.21) with $k = 1$ and L set to either E or I. Because we know how ν_E^2 and ν_I^2 depend on ν_E and ν_I, we can express Δ_{h_L} (the only term that depends on the second moments of the firing rates) in terms of ν_E and ν_I. The two resulting equations for ν_E and ν_I represent curves in $\nu_E - \nu_I$ space, those curves are the excitatory ($L = E$) and inhibitory ($L = I$) nullclines; see Wilson and Cowan (1972) and Latham et al. (2000a) for details on how they are constructed. For strongly coupled networks operating in the balanced regime–presumably the regime of interest for the brain–they should intersect where the slope of the excitatory nullcline is positive (van Vreeswijk and Sompolinsky 1998).

3 Numerical results

When deriving our mean field equations, we assumed uncorrelated and Poisson spikes, constant firing rates, and white noise synaptic drive, and we used an approximate expression for the firing rate. To determine the effect of these assumptions and approximations, we performed network simulations. For all simulations we integrated Eq. (2.6a), with the added condition that a spike was emitted when the voltage reached $+\infty$, at which point it was reset to $-\infty$. To avoid numerical issues with the infinities, we made the change of variables $V = \overline{V} + (V_{th} - V_r)\tan(\theta/2)$; see Appendix B. In these variables, a spike is emitted when θ passes through π. In addition, to speed up the simulations, we replaced the external input ($M = X$ in Eq. (2.6b)) with white noise; see Appendix C.

We typically performed simulations with a set of three "default" networks. These networks differed only in their synaptic time constants, which were either 1, 10 or 100 ms; in all cases the membrane time constant was 10 ms (see Table 1). Figure 4 shows activity for the default network

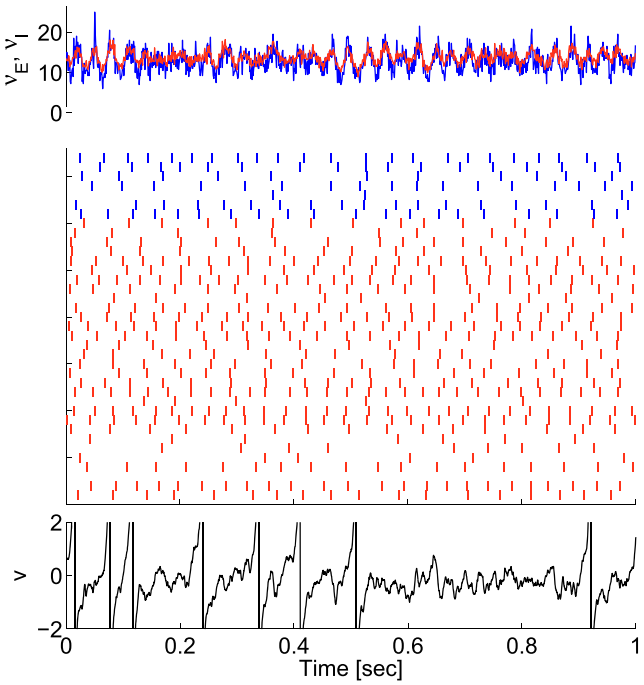

Fig. 4 Activity of the default network. Top: Population averaged firing rate of excitatory (*red*) and inhibitory (*blue*) neurons. Note that the network was not perfectly asynchronous (oscillations are visible). Middle: Spike rasters from sample neurons. There are four times as many excitatory neurons as inhibitory ones, consistent with our network architecture. Bottom: Membrane potential of a sample neuron

with $\tau_s = 1$ ms. The top plot shows the time-dependent population averaged firing rate for the excitatory and inhibitory neurons, with firing rate computed in 1 ms bins; the center plot shows spike rasters from a subset of the neurons (with blue for inhibitory neurons and red for excitatory ones); and the bottom plot shows the membrane potential of a sample neuron; it resembles spike trains recorded *in vivo*.

As an initial test of our mean field predictions, we considered networks of non-interacting neurons driven by external input approximated by white noise (see Eq. C.1). These simulations test the approximate expression for the firing rate given in Eq. (2.4). The results are shown in Fig. 5, where we plot the predicted and observed distributions of firing rates for the excitatory (left column) and inhibitory (center column) neurons for a range of synaptic time constants. The predicted firing rate distributions (thick lines in the left and center panels) are close to the simulated ones, as are the predicted and simulated mean firing rates (vertical lines). In the right column, we plot the excitatory and inhibitory nullclines (the solutions to Eq. (2.21) with $k = 1$ and $L = E$ and I, respectively) along with a 100 ms trajectory of the population averaged excitatory and inhibitory firing rates (binned at 1 ms). Consistent with the fact that the populations are decoupled, the excitatory and inhibitory nullclines are vertical and horizontal, respectively, and the excitatory and inhibitory firing rates are uncorrelated. Overall, the close match between the predictions and simulations indicates that the approximate firing rate (Fig. 1) will not be a limiting factor in the accuracy of our mean field models.

Next we tested mean field predictions in a coupled network – the default network given in Table 1. The results are shown in Fig. 6. Unlike in our uncoupled networks, the excitatory and inhibitory populations now interact, and the interactions are strong enough that the firing rates are

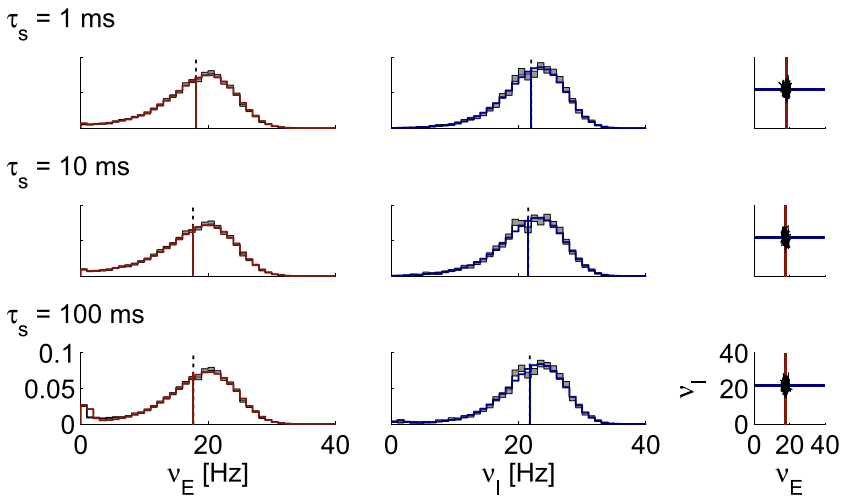

Fig. 5 Histograms of firing rates in two populations of neurons in a disconnected network. Parameters are from the default network, Table 1, except that $J_{LM} = 0$. Firing rates derived from the simulation are represented by the shaded region (population average ± 1 standard deviation, $n_E = 16000$, $n_I = 4000$); these were computed from 5 s of simulations and binned at 1 Hz. The gray area indicates the count error, $n_k^{1/2}/N_L$ where n_k is the number of neurons with firing rates falling into the k-th bin. The solid lines are the histograms predicted by the mean field theory. Vertical lines represent ν_E, ν_I; the dashed line is estimated from the simulations and the solid line from mean field theory. The right column shows the excitatory and inhibitory nullclines, along with a 100 ms trajectory. Because the populations are uncoupled, the nullclines are vertical and horizontal and the firing rates are independent

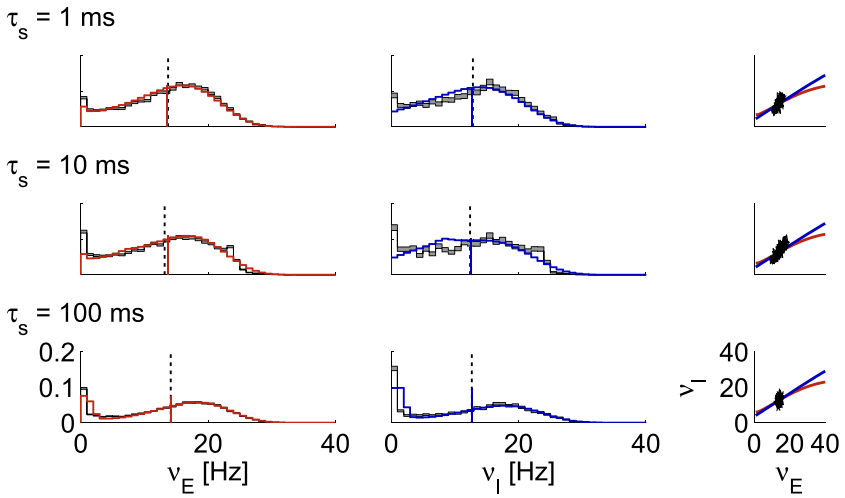

Fig. 6 Histograms of firing rates in two populations of neurons in a coupled network. Same as Fig. 5, except that the default network, Table 1, is used with connections intact. Because the excitatory and inhibitory populations are now coupled, the nullclines are no longer vertical and horizontal, and, at least for $\tau_s = 1$ ms and $\tau_s = 10$ ms, the excitatory and inhibitory firing rates are correlated. Not surprisingly, the mean field predictions do not match the simulations as well as in Fig. 5. For $\tau_s = 100$ ms, however, excitatory and inhibitory rates are uncorrelated; consequently, the mean field theory and simulations match better. This is a general trend in our data (see Figs. 7 and 8)

correlated (right column). Nevertheless, the theoretical and simulated firing rate histograms match reasonably well, with the best matches at the shortest ($\tau_s = 1$ ms, first row), and the longest ($\tau_s = 100$ ms, third row) synaptic time constants. The largest mismatch between theory and simulation happens at $\tau_s = 10$ ms, and mainly for the inhibitory neurons. The mismatch is not so surprising given the average trajectory of firing rates (right column), which–in violation of the asynchronous assumption–exhibit strong synchronization. For all networks, the theoretical prediction for the average firing rate is almost identical to the value we get from simulations. This is true even for $\tau_s = 10$ ms, where the network exhibits strong synchronization. This is a hint that, at least when it comes to mean firing rates, our mean field theory is very robust.

How do our mean field equations hold up when the neurons become synchronized? To address this question, we increased the connection strengths, a manipulation that tends to make the network more synchronous. In particular, we doubled all recurrent connections; that is, we increased J_{LE} and J_{LI} by a factor of two compared to what we used for the default network. As can be see in Fig. 7, when $\tau_s = 1$ and 10 ms, the networks became synchronized, as indicated by the strong correlations in excitatory and inhibitory firing rates (right panels). Not surprisingly, the theoretical and simulated firing rate distributions are now very different (about as different as they were for $\tau_s = 10$ ms in Fig. 6, which was also fairly synchronized). Note, though, that again the predicted and simulated population averaged firing rates are very similar. For the long time constant network, $\tau_s = 100$ ms, there was very little synchronization (the excitatory and inhibitory rates were weakly correlated). Consistent with this, there is a good match between theory and simulations.

To quantify how synchronization affects the accuracy of our mean field model, we randomly varied the strengths of the recurrent connections and plotted the match between theory and simulations versus degree of synchrony. Specifically, we let $J_{KL} \to J_{KL}(1 + 0.2\eta_{KL})$, where η_{KL} is zero mean, unit variance Gaussian random variable. Similar to Brunel and Hakim (1999), network synchrony, denoted S, was defined to be the maximum cross-covariance between the instantaneous population averaged excitatory and inhibitory firing rates, normalized by the population averaged firing rates,

$$S \equiv \max_{\tau} \int_0^T \frac{dt}{T} \frac{(\nu_E(t) - \nu_E)(\nu_I(t+\tau) - \nu_I)}{\nu_E \nu_I}. \tag{3.1}$$

Note that we have slightly abused notation: $\nu_E(t)$ and $\nu_I(t)$ are the instantaneous population averaged excitatory and inhibitory firing rates, whereas ν_E and ν_I are the population averaged firing rates with an additional average over time. We compute the above integral by discretizing time into 1 ms bins.

In Fig. 8 we plot error versus synchrony (left column), with colors corresponding to different synaptic time constants (black for $\tau_s = 1$ ms, red for $\tau_s = 10$ ms, yellow for $\tau_s = 100$ ms). As predicted, the smaller the synchrony, the better the match of theory to simulations. For large synaptic time constants, all networks exhibited negligible synchrony, and the theory worked very well. In the right column we plot error versus firing rate. This plot shows two things: we explored a relatively large range of firing rates, and firing rate alone is not a good predictor of error.

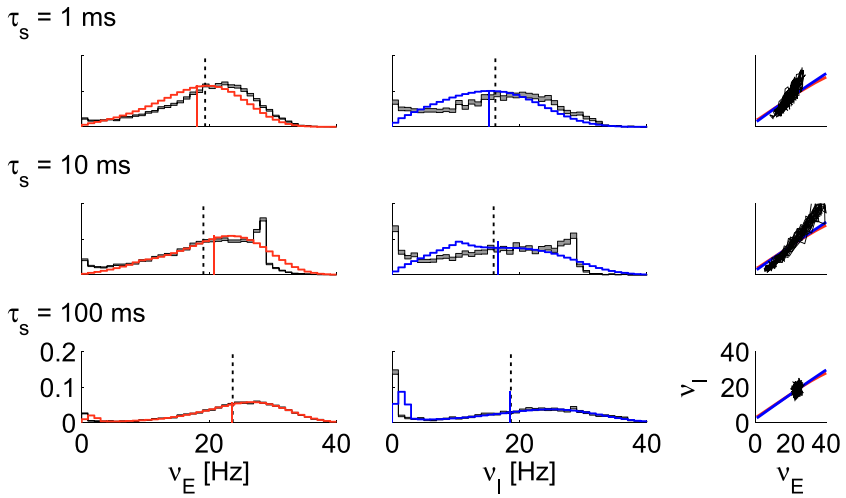

Fig. 7 Histograms of firing rates in a strongly coupled network. Same as Fig. 6 except that the recurrent connections, J_{KE} and J_{KI}, are increased by a factor of two. As a result, network activity becomes even more synchronous, at least at $\tau_s = 1$ ms and $\tau_s = 10$ ms. Not surprisingly, the match between the firing rate distribution computed from our mean field theory and from the simulations is not as good as it was in Figs. 5 and 6, although the population averaged firing rates are not far off. Note, though, that the network is still desynchronized for the longest time constants $\tau_s = 100$ ms. Here the mean field model does good job predicting the firing rate distribution, except at low firing rate, where the mismatch is explained by the inaccuracy of our approximation of firing rate model (see Fig. 1)

Finally, we investigate the effect of changing the size of the network. We varied network size from 500 to 60,000 neurons, keeping all other parameters fixed to those of the default network (see Table 1). In Fig. 9 we plot, in the top panel, the degree of synchrony, S (Eq. (3.1)), versus network size. Up to networks of around 4,000 neurons, synchrony dropped steadily and consistently for the three synaptic time constants tested (1, 10 and 100 ms). However, above 4,000 neurons, synchrony rose for the intermediate synaptic time constant, $\tau_s = 10$ ms. The initial decrease in synchrony is probably associated with the decrease in fluctuations that comes with larger networks, as fluctuations tend to drive oscillations around the fixed point. The subsequent rise for larger networks is harder to explain. However, based on

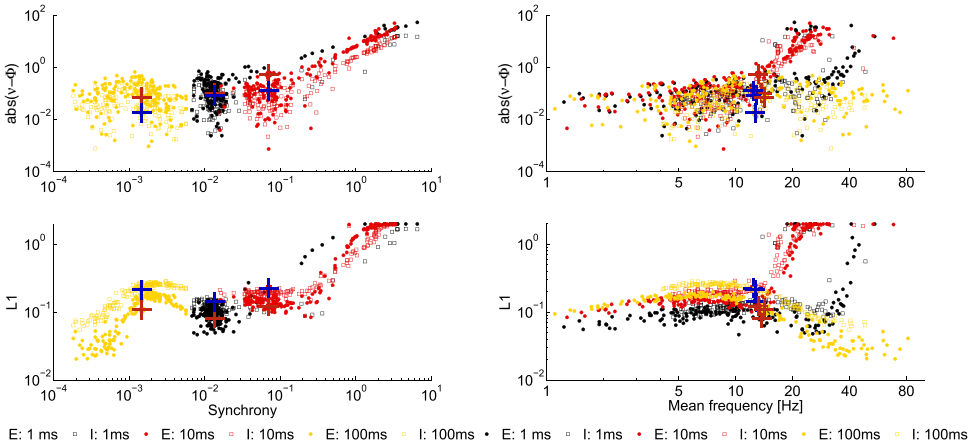

Fig. 8 The match between mean field predictions and simulation gets worse as synchrony increases. Mean field theory prediction error for excitatory (*dots*) and inhibitory (*squares*) populations. Colors indicate different synaptic time constants, and the pluses come from the default network (Fig. 6), with red corresponding to excitatory neurons, and blue to inhibitory ones. The synchrony measure plotted on the x-axis is the peak covariance between $\nu_E(t)$ and $\nu_I(t)$ (normalized by firing rates; see Eq. (3.1)). Top left: absolute difference between the simulated and predicted population averaged firing rates (ν_E and ν_I). Bottom left: sum over bins of the absolute differences between the simulated and predicted firing rate probability distributions, theory vs experiment (L1 norm; maximum value = 2). Top right: absolute difference between the simulated and predicted population averaged firing rates (ν_E and ν_I), this time versus firing rate. Bottom right: L1 norm versus firing rate. Note that the degree of synchrony does a good job predicting the accuracy of the mean field model, while the firing rate by itself does not

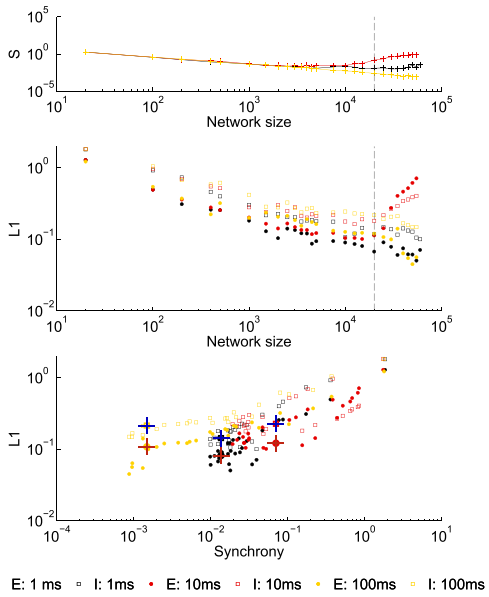

Fig. 9 Effect of network size on the match between mean-field and simulations. Default network, except that we varied the number of neurons from 500 to 60,000 (with, as usual, a factor of four more excitatory neurons than inhibitory ones). The default size (20,000) is marked with a dashed line. Color conventions are the same as in Fig. 8. Top: as network size increases, the degree of synchrony drops for $\tau_s = 1$ and 100 ms, but not for $\tau_s = 10$ ms; see text for an explanation. Middle: The quality of our prediction improves for larger networks, unless their activity synchronizes (as for $\tau_s = 10$ ms, see above). Bottom: network sizes are pooled and the L1 norm of the firing rate distribution (as in Fig. 8) is plotted versus synchrony. Not surprisingly, the higher the synchrony, the worse the prediction. Pluses come from the default network (Fig. 6)

previous work (Brunel and Hakim 1999; Rappel and Karma 1996), we suspect it's because fluctuations are a two-edges sword. On the one hand, they drive oscillations; on the other hand, they act as a noise source which tends to decorrelate neurons. At the intermediate synaptic time constant, $\tau_s = 10$ ms, the asymptotic network state is probably an oscillatory one, and those oscillations were being masked by the fluctuations associated with finite size effects. As the size of the network, and thus the size of the fluctuations, dropped, the oscillations were uncovered.

In the middle and bottom panels of Fig. 9 we plot the L1 norm of the firing rate distributions (described in the caption of Fig. 8). For synaptic time constants, τ_s, of 1 and 100 ms, the L1 norm decreases (implying our mean field predictions are better) as the number of neurons increases (middle panel). When $\tau_s = 10$ ms, however, our mean field predictions get worse at larger network sizes; this is consistent with the increase in synchrony seen in the top panel. When the L1 norm is plotted versus synchrony (bottom panel), the story is simpler: the larger the synchrony, the larger the L1 norm, and the worse our mean field predictions. This plot suggest that the degree of synchrony has a larger effect on our mean field predictions than the network size. Similar, though slightly more noisy, results are achieved for the mean population activity (not shown).

4 Discussion

Using what are by now relatively standard mean field methods (Shiino and Fukai 1992, 1993; Amit and Brunel 1997a, b; van Vreeswijk and Sompolinsky 1998; Roudi and Latham 2007), we computed the distribution of excitatory and inhibitory firing rates in large networks of recurrently connected spiking neurons. Our main result is that we could assume that activity was asynchronous and Poisson (assumptions that are clearly violated), and still get relatively good agreement with network simulations–so long as the network is not too synchronous; see Fig. 8. This indicates that much of the intuition developed for these kinds of networks–see in particular the seminal work of van Vreeswijk and Sompolinsky (1998)–applies even in the mildly synchronous regime. It should be noted, though, that all our analysis was based on the quadratic integrate and fire neuron with current based synapses. Whether our results apply to other single neuron and synaptic models is an open–and, we believe, interesting–question.

In our analysis we used current-based quadratic integrate and fire neurons, chosen both because they are a good model of type I neurons (Ermentrout and Kopell 1986) and because there is a good approximate expression for their firing rate given colored noise input (Brunel and Latham 2003), which in turn is a reasonably good approximation to synaptic input (Walsh 1981; Tuckwell 1988). While an equivalent analytic treatment would be difficult for conductance-based models and more realistic neurons, we do not expect any major surprises: the nullclines, which ultimately govern the range of network behavior, should be similar (Latham 2002), and, as we have shown, even for relatively synchronous firing the mean field model still provides a good description, at least at the level of population averaged firing rates.

Although the random connectivity used here is a major idealization, it is important for two reasons. First, it forms the substrate upon which computations are built, and has been shown to play a major role in determining exactly how those computations are carried out (Latham et al. 2000a, b; Salinas 2003; Latham and Nirenberg 2004; Roudi and Latham 2007). Second, it leads naturally to the next question: would mean field theory apply to networks with structured connectivity, which are just as prone to oscillations as randomly connected ones? Mean field theory has been applied to structured networks in a limited number of cases, (Amit and Brunel 1997a, b; Latham and Nirenberg 2004; Roudi and Latham 2007), but a thorough understanding of such networks awaits development.

Acknowledgments We thank Nicolas Brunel for helping initiate the project and for critical reading of the manuscript. We thank Peter Dayan for productive discussions. P.E.L. and A.G-B. were supported by the Gatsby Charitable Foundation. We also acknowledge the hospitality of the Kavli Institute for Theoretical Physics, where a portion of this work was performed.

Conflict of interests The authors declare that they have no conflict of interest.

Appendix A: Statistics of the synaptic drive

In the main text we approximated $\overline{h}_{Li}$ as a Gaussian random variable with respect to index, i, and the right hand side of Eq. (2.16) as Gaussian white noise. With this approximation, all we need are the variance of $\overline{h}_{Li}$ and the covariance of the right hand side of Eq. (2.16). Here we compute those quantities.

We start with the variance of $\overline{h}_{Li}$, Eq. (2.12). To isolate the index-independent and index-dependent terms, we write

$$J_{LM}^{ij} = \epsilon J_{LM} + \delta J_{LM}^{ij} \tag{A.1}$$

where ϵJ_{LM} is the population averaged value of J_{LM}^{ij} (see Eq. (2.7)) and $\delta J_{LM}^{ij} \equiv J_{LM}^{ij} - \epsilon J_{LM}$ represents the index-dependent fluctuations around that average (sometimes referred to as the quenched noise). Making this substitution, using Eq. (2.15a) for the mean firing rate, and recalling that $\epsilon = K_M/N_M$, Eq. (2.12) becomes

$$\overline{h}_{Li} = h_L + \delta\mu_{Li} + \sum_{M,j} \frac{\tau_m}{K_M^{1/2}} \delta J_{LM}^{ij} \nu_{Mj} \tag{A.2}$$

where h_L is given in Eq. (2.13b) and the sum is over $M = E$, I and X. The last term in this expression is the sum of a large number of variables. The weights inside the sum are truly random, so if the firing rates and the weights are sufficiently weakly correlated, this sum is a Gaussian random variable with respect to index, i. Here we assume they are, although this is clearly an approximation: the firing rates, ν_{Mj}, are functions of the connection strengths, and so the variables inside the sum are not quite independent. However, in practice this is a good approximation, especially if ϵ (which is a measure of the sparseness of the connectivity; see Eq. (2.7)) is small, something that tends to reduce correlations. Given this approximation, and the fact that, by construction, the mean is zero, all we need is the variance. This variance (plus the variance of $\delta\mu_{Li}$, which, by construction, is $\Delta_{\mu_L}^2$ (see Eq. (2.8)), is given by

$$\Delta_{h_L}^2 = \Delta_{\mu_L}^2 + \sum_{M,M',j,j'} \frac{\tau_m^2}{(K_M K_{M'})^{1/2}} \nu_{Mj} \nu_{M'j'} \frac{1}{N_L} \sum_i \delta J_{LM}^{ij} \delta J_{LM'}^{ij'} \tag{A.3}$$

where, as in Eq. (2.13a), we use $\Delta_{h_L}^2$ for the total variance. When $j \neq j'$ or $M \neq M'$, in the large K limit the sum is approximately zero; when $j = j'$ and $M = M'$, the sum over i is just the variance of J_{LM}^{ij}. Thus, using Eq. (2.7) for the variance of J_{LM}^{ij}, Eq. (A.3) becomes, after a small amount of algebra,

$$\Delta_{h_L}^2 = \Delta_{\mu_L}^2 + \sum_M J_{LM}^2 \left(1 + \Delta^2 - \epsilon\right) \tau_m^2 \nu_M^2 \tag{A.4}$$

where ν_M^2 is the second moment of the firing rate (Eq. (2.15b)).

We next compute the covariance of the right hand side of Eq. (2.16). Using $C_{LL'}^{ii'}(\tau)$ to denote the covariance between neuron i of type L and neuron i' of type L' at times separated by τ, we have

$$C_{LL'}^{ii'}(\tau) = \sum_{M,M',j,j'} \frac{\tau_m^2}{(K_M K_{M'})^{1/2}} J_{LM}^{ij} J_{L'M'}^{i'j'} \times \left\langle \sum_{l,l'} \left[\delta\left(t - t_j^l\right) - \nu_{Mj}\right] \left[\delta\left(t + \tau - t_{j'}^l\right) - \nu_{M'j'}\right] \right\rangle. \tag{A.5}$$

The angle brackets represent an average over the distribution of spike times. Real neurons have a nontrivial correlational structure; if nothing else, there is a refractory period. However, we ignore that and make the approximation that the neurons are Poisson. In that case, as shown by (Rice 1954), and as is relatively easy to derive, the average over the distribution of spikes yields

$$\left\langle \sum_{l,l'} \left[\delta\left(t - t_j^l\right) - \nu_{Mj}\right] \left[\delta\left(t + \tau - t_{j'}^{l'}\right) - \nu_{M'j'}\right] \right\rangle = \nu_{Mj}\delta(\tau)\delta_{jj'}\delta_{MM'} \tag{A.6}$$

where δ_{ij} is the Kronecker delta ($\delta_{ij} = 1$ if $i = j$ and 0 otherwise). Thus, Eq. (A.5) becomes

$$C_{LL'}^{ii'}(\tau) = \delta(\tau) \sum_{M,j} \frac{\tau_m^2}{K_M} J_{LM}^{ij} J_{L'M}^{i'j} \nu_{Mj}. \tag{A.7}$$

Assuming, as usual, that the connections strengths are approximately independent of the firing rates, we may average the connection strengths and firing rates separately. Using Eq. (2.7) for the distribution of connection strengths, we have

$$\begin{aligned} C_{LL'}^{ii'}(\tau) &= \delta(\tau)\tau_m^2 \sum_M \frac{\nu_M}{K_M} \sum_j J_{LM}^{ij} J_{L'M}^{i'j} \\ &= \delta(\tau)\tau_m^2 \sum_M J_{LM}^2 \left[\epsilon(1 - \delta_{ii'}\delta_{LL'}) + \left(1 + \Delta^2\right)\delta_{ii'}\delta_{LL'}\right] \nu_M. \end{aligned} \tag{A.8}$$

An important observation is that $C_{LL'}^{ii'}(\tau)$ is nonzero even when $i \neq i'$ and/or $L \neq L'$. Thus, the driving terms for different neurons are correlated; this in turn implies that spike times are correlated across neurons. This would seem

to imply that our independence approximation is badly violated. However, as shown by (Renart et al. 2010; Hertz 2010), for balanced networks operating in the asynchronous regime, correlations between excitatory and inhibitory neurons largely cancel, leaving the mean correlation on the order of $1/N$. Thus, in large networks the independence approximation tends to work relatively well. This means we can focus on the autocorrelation, C^{ii}_{LL}, which is somewhat simpler than the full covariance,

$$C^{ii}_{LL}(\tau) = \delta(\tau)\tau_m^2 \sum_M J^2_{LM}\left(1+\Delta^2\right)\nu_M. \tag{A.9}$$

This expression leads to Eqs. (2.17) and (2.18).

Appendix B: Transforming from the quadratic integrate and fire neuron to the θ-neuron

For quadratic integrate and fire neurons, action potentials are emitted when the voltage reaches $+\infty$, at which point the voltage is reset to $-\infty$. Integrating to infinity, however, poses a problem numerically. To get around this, we make the change of variables

$$V_{Li} = \overline{V} + (V_{th} - V_r)\tan(\theta_{Li}/2). \tag{B.1}$$

This moves the points at $V_{Li} = \pm\infty$ to $\theta_{Li} = \pm\pi$, and also removes the singularities at $\pm\infty$. Inserting this into Eq. (2.6a) we see that θ_{Li} evolves according to

$$\tau_m \frac{d\theta_{Li}}{dt} = (1-\cos\theta_{Li}) + (1+\cos\theta_{Li})(\mu_L + \mu_{Li} + h_{Li}). \tag{B.2}$$

A spike is emitted when $\theta_{Li} = \pi$, at which point it is reset to $-\pi$.

Appendix C: White noise approximation to external input

To speed up the simulations, we use Gaussian white noise instead of actual spike trains for the external input (the term with $M = X$ in Eq. (2.6b)). To do that, we make the replacement

$$\frac{\tau_m}{K_X^{1/2}} \sum_{j,l} J^{ij}_{LX}\,\delta\left(t - t^l_{Xj}\right) \rightarrow$$
$$J_{LX}\tau_m\overline{\nu}_X\left(K_X^{1/2} + (1+\Delta^2-\epsilon)^{1/2}\eta_{LXi} + \left[\frac{1+\Delta^2}{\tau_m\overline{\nu}_X}\right]^{1/2}\xi_{LXi}(t)\right) \tag{C.1}$$

where η_{LXi} is a zero mean, unit variance Gaussian random variable with respect to index, i, $\xi_{LXi}(t)$ is Gaussian white noise, and we assumed that all the external neurons have the same firing rate ν_X (which allowed us to replace $\left(\nu_X^2\right)^{1/2}$ with ν_X); see Eqs. (2.13b), (2.14) and (2.18).

References

Amit, D., & Brunel, N. (1997a). Dynamics of a recurrent network of spiking neurons before and following learning. *Network*, *8*, 373–404.

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

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, *8*(3), 183–208.

Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, *11*(7), 1621–1671.

Brunel, N., & Latham, P. (2003). Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Computation*, *15*, 2281–2306.

Deger, M., Helias, M., Boucsein, C., Rotter, S. (2012). Statistical properties of superimposed stationary spike trains. *Journal of Computational Neuroscience*, *32*(3), 443–463.

Ermentrout, B. (1996). Type i membranes, phase resetting curves, and synchrony. *Neural Computation*, *8*, 979–1001.

Ermentrout, B., & Kopell, N. (1986). Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM Journal on Applied Mathematics*, *46*, 233–253.

Gutkin, B., & Ermentrout, B. (1998). Dynamics of membrane exci tability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics. *Neural Computation*, *10*, 1047–1065.

Hansel, D., & Mato, G. (2001). Existence and stability of persistent states in large neuronal networks. *Biophysical Reviews and Letters*, *86*, 4175–4178.

Hertz, J. (2010). Cross-correlations in high-conductance states of a model cortical network. *Neural Computation*, *22*(2), 427–447.

Koch, C. (1998). *Biophysics of computation: information processing in single neurons (Computational Neuroscience)*, 1st edn. Oxford University.

Latham, P. (2002). Associative memory in realistic neuronal networks. *Advances in neural information processing systems* (Vol. 14). Cambridge: MIT.

Latham, P., & Nirenberg, S. (2004). Computing and stability in cortical networks. *Neural Computation*, *16*, 1385–1412.

Latham, P., Richmond, B., Nelson, P., Nirenberg, S. (2000a). Intrinsic dynamics in neuronal networks. I. theory. *Journal of Neurophysiology*, *83*, 808–827.

Latham, P., Richmond, B., Nirenberg, S., Nelson, P. (2000b). Intrinsic dynamics in neuronal networks. II. experiment. *Journal of Neurophysiology*, *83*, 828–835.

Lerchner, A., Sterner, G., Hertz, J., Ahmadi, M. (2006a). Mean field theory for a balanced hypercolumn model of orientation selectivity in primary visual cortex. *Network*, *17*(2), 131–150.

Lerchner, A., Ursta, C., Hertz, J., Ahmadi, M., Ruffiot, P., Enemark, S. (2006b). Response variability in balanced cortical networks. *Neural Computation*, *18*(3), 634–659.

Rappel, W.J., & Karma, A. (1996). Noise-induced coherence in neural networks. *Physical Review Letters*, *77*(15), 3256–3259.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., Harris, K.D. (2010). The asynchronous state in cortical circuits. *Science*, *327*(5965), 587–590.

Rice, S. (1954). Mathematical analysis of random noise. In *Selected papers on noise and stochastic processes* (pp. 130–294). Dover.

Roudi, Y., & Latham, P. (2007). A balanced memory network. *PLoS Computational Biology*, *3*, 679–1700.

Salinas, E. (2003). Background synaptic activity as a switch between dynamical states in a network. *Neural Computation*, *15*, 1439–1475.

Shiino, M., & Fukai, T. (1992). Self-consistent signal-to-noise analysis and its application to analogue neural networks with asymmetric connections. *Journal of Physics A*, *25*, L375–L381.

Shiino, M., & Fukai, T. (1993). Self-consistent signal-to-noise analysis of the statistical behavior of analog neural networks and enhancement of the storage capacity. *Physical Review E*, *48*, 867–897.

Shriki, O., Hansel, D., Sompolinsky, H. (2003). Rate models for conductance-based cortical neuronal networks. *Neural Computation*, *15*, 1809–1841.

Tuckwell, H. (1988). *Introduction to theoretical neurobiology* (Vol. 2) Cambridge: Cambridge University.

van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput*, *10*, 1321–1371.

Walsh, J. (1981). A stochastic model of neural response. *Advances in Applied Probability*, *13*, 231–281.

Wilson, H., & Cowan, J. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*, 1–24.