



Malleability of gamma rhythms enhances population-level correlations

Sonica Saraf¹ · Lai-Sang Young²

Received: 12 September 2019 / Revised: 23 December 2020 / Accepted: 27 January 2021 / Published online: 5 April 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

An important problem in systems neuroscience is to understand how information is communicated among brain regions, and it has been proposed that communication is mediated by neuronal oscillations, such as rhythms in the gamma band. We sought to investigate this idea by using a network model with two components, a source (sending) and a target (receiving) component, both built to resemble local populations in the cerebral cortex. To measure the effectiveness of communication, we used population-level correlations in spike times between the source and target. We found that after correcting for a response time that is independent of initial conditions, spike-time correlations between the source and target are significant, due in large measure to the alignment of firing events in their gamma rhythms. But, we also found that regular oscillations cannot produce the results observed in our model simulations of cortical neurons. Surprisingly, it is the *irregularity* of gamma rhythms, the absence of internal clocks, together with the *malleability* of these rhythms and their tendency to align with external pulses — features that are known to be present in gamma rhythms in the real cortex — that produced the results observed. These findings and the mechanistic explanations we offered are our primary results. Our secondary result is a mathematical relationship between correlations and the sizes of the samples used for their calculation. As improving technology enables recording simultaneously from increasing numbers of neurons, this relationship could be useful for interpreting results from experimental recordings.

Keywords Correlations · Gamma rhythms · Population dynamics · Synchrony · Time-delays

1 Introduction

The following two bodies of facts in neuroscience are well established. The first is that brain regions are highly interconnected and a significant part of signal transmission is via the spiking of neurons. These facts are well documented, e.g. in Binzegger et al. (2009); Felleman and Van Essen (1991); Sincich and Horton (2005). The second body

of facts, also well known, involves rhythms in the brain: rhythms in various frequency bands (δ , θ , α , β and γ) are ubiquitous (Buzsaki, 2011).

How exactly these rhythms impact signal transmission is an intriguing question: it is challenging to define precisely what it means to facilitate communication, let alone to measure or quantify it. It has been suggested that gamma rhythms are important for information transfer (Rodriguez et al., 1999; Fries, 2005; Fries, 2015; Pesaran et al., 2002; Womelsdorf et al., 2007; Sohal et al., 2009). It has also been documented that certain diseases and drugs, as well as heightened attention, learning, and memory formation are associated with altered gamma rhythms (Fries et al., 2001; Sederberg et al., 2003; Gonzalez-Burgos et al., 2010; Uhlhaas & Singer, 2010; McCarthy et al., 2012; Gonzalez-Burgos et al., 2015).

Motivated by these facts and ideas, this paper uses computational modeling to study the effect of gamma rhythms on signal transmission. We consider here two

Communicated by Jonathan D Victor and Alain Destexhe.

Action Editor: Jonathan D Victor and Alain Destexhe.

✉ Lai-Sang Young
lsy@cims.nyu.edu

¹ Center for Neural Science, New York University,
10003 New York, USA

² Courant Institute of Mathematical Sciences, New York
University, New York 10012, USA

local neuronal populations, one a source (sending) and the other a target (receiving), and our goal is to understand – on the mechanistic level – how gamma rhythms affect communication between them. The local populations are groups of a few hundred Excitatory and Inhibitory integrate-and-fire neurons, with connectivities similar to those in local circuits of the cerebral cortex. Population activity is driven by external input together with dynamical interaction among the neurons. We wanted our models to be realistic enough so that our results can be related to experiments. At the same time, the study is not customized to specific brain regions, in order to produce broadly relevant results.

As to how to quantify communication, we have elected to study correlations in spike firing on brief timescales of several milliseconds, a range that reflects gamma-band activity, and we chose to focus on *population-level correlations* between the source and target, meaning time correlations between the aggregated spike trains produced by the source and the target populations without regard to the participation of individual neurons. This metric is similar to that in Vogels and Abbot (2005), and was used in favor of correlations between pairs of neurons because we believe that it is a better reflection of communication between two populations. Though experimental capabilities are improving, we recognize that present technology does not yet permit the direct measurement of spike firing in entire populations, and most existing experimental results are for paired correlations (Roe & Ts'o, 1999; Nowak et al., 1999; Jia et al., 2013; Zandvakili & Kohn, 2015). To connect with existing results, we derived a formula relating correlations measured from different sample sizes.

There are other interpretations of “communication” besides correlated spiking: for example, some have interpreted Fries (2005) as proposing that higher firing rates in the target population was indicative of effective communication, and this study considered this possibility as well.

Our model populations needed to produce realistic gamma rhythms to enable us to investigate how rhythms impact communication. We have tried to emulate experimental data, which show that there are two aspects to gamma-band activity. One is its oscillatory behavior, and the other is the irregular, even episodic nature of the rhythm (Burns et al., 2011; Xing et al., 2012). It has been proposed in Fries (2005) and studied in theoretical works (Gielen et al., 2010; Börgers & Kopell, 2008; Battaglia et al., 2012; ter Wal & Tiesinga, 2017) that neuronal oscillations are responsible for effective communication, and we confirm these ideas using a more realistic model than was used in the aforementioned works. The main thrust of this paper, however, is that oscillatory behavior alone is insufficient: it is the *irregularity* in the oscillations that plays an important enabling role. Another feature, first pointed out in Chariker et al. (2018), is the *malleability* of the timing in

gamma events – meaning that a driving population can non-trivially influence the spike times of the receiving population, causing their firing events to align. This was found to be a very significant contributor to the robust correlation between connected populations.

Finally, in addition to identifying relevant phenomena, we also offer mechanistic explanations for the phenomena observed, seeking to shed light on cortical mechanisms. The proposals in this paper are based on experimental data, known theoretical and modeling results, and our own simulations. This emphasis on elucidating mechanisms is a distinguishing feature of the present paper.

2 Preliminaries I: Single-population models

In this section, we describe the models for the local populations used for the correlation studies to follow. We also document their gamma-band activity, which we will show in the sections to follow, serves to synchronize connected populations.

2.1 Model description

Local populations in this paper are intended to model local circuits in the cerebral cortex. Typically they consist of a few hundred to a thousand neurons, three quarters of which are Excitatory (E) and the rest Inhibitory (I). The neurons are randomly and homogeneously connected to one another according to certain specified probabilities, and they are modeled as integrate-and-fire neurons. Similar models were used in Chariker and Young (2015).

The local population models used for the simulations here are networks consisting of 300 E and 100 I-neurons. On average, each E-neuron is postsynaptic to about 80 other E-neurons and 50 I-neurons, and each I-neuron is postsynaptic to about 240 E-neurons and 50 other I-neurons. Connectivity between pairs of neurons are subject to variance, and are randomly drawn according to the means above. These connection probabilities, with E-to-E being more sparse than connections that involve I-neurons, are consistent with neuroanatomy (Chariker et al., 2016).

Leaving details to Supplementary Information, we outline the rules governing the dynamical interaction within the local population. The dynamics of individual neurons are described by the following conductance-based leaky integrate-and-fire (LIF) equation:

$$\dot{V} = -\frac{1}{\tau_{leak}} V - (V - V_E)g_E - (V - V_I)g_I \quad (1)$$

This equation describes the evolution of the membrane potential of a neuron, V , over time in milliseconds (ms). The membrane potential has been normalized so that $V = 0$

is the resting state, and when V reaches 1, the neuron fires an action potential following which its membrane potential is immediately reset to 0, where it remains in refractory for a couple of ms. Eq. (1) contains three constants: $\tau_{leak} = 20$ ms is the leak rate, and $V_E = \frac{14}{3}$, $V_I = -\frac{2}{3}$ are excitatory and inhibitory reversal potentials in normalized units (McLaughlin et al., 2000). These are accepted biophysical constants (Koch, 1999). The functions $g_E(t)$ and $g_I(t)$ are the excitatory and inhibitory conductances of the neuron in question. When an excitatory spike is received, $g_E(t)$ is temporarily elevated for a few ms; the same is true for $g_I(t)$ when a spike is received from an inhibitory neuron.

This system has 8 parameters. Four of them, $S^{QQ'}$, $Q, Q' \in \{E, I\}$, represent the synaptic coupling weights from neurons of type Q' to neurons of type Q . Two others, τ_E and τ_I , denote the rates at which the E and I-conductances, which elevate upon the arrival of a spike, decay to zero. The remaining two determine the amount of external drive fed into the system, described in the next paragraph.

In addition to the synaptic input received from within the local population, each neuron receives an excitatory external drive modeled as a Poisson point process. This drive is independent from neuron to neuron. The external

drive has two components: a *synaptic component*, which consists of E-spikes representing input from other regions of the brain with synaptic weight S^{QE} for neurons of type Q , and an “*ambient*” component with a smaller synaptic weight meant to represent all neurotransmitters not specifically modeled. The ambient component’s Poisson rate is assumed to be constant, whereas the synaptic component’s rates, λ_E and λ_I for E and I-neurons, are assumed to be low in background and to increase with drive.

In our simulations, the parameters above are chosen with guidance from realistic models of the visual cortex such as that in Chariker et al. (2016). Their exact values are unimportant for purposes of the present study, as long as they produce reasonable dynamics (including those depicted in Fig. 1). Since we are primarily interested in correlations between populations when the sending, also referred to as “source”, population is strongly driven, we use values of λ_E and λ_I that produce average E-firing rates of about 15 spikes/sec, which is consistent with the average stimulus-driven firing rates of local populations in realistic situations; see e.g. (Chariker et al. 2020).




Fig. 1 Spiking statistics in a strongly driven regime. **a** Interspike interval plots showing spike firing statistics of individual neurons. We randomly chose 20 E neurons and for each, calculated a histogram of time between each spike with 10 ms bins. We then averaged the 20 histograms. This was calculated using 8 seconds of simulation data. Simulations using a single neuron produced very similar outputs when sampled over a much longer time period. **b** Population statistics: Raster showing half a second of the simulation. Red dots repre-

sent excitatory spikes and blue are inhibitory spikes. A rhythm in the gamma band is clearly visible. **c** Summed spike plot for the E neurons of plot b, showing the fraction of the E-population spiking within each time bin of 4 ms. **d** PSD for the regime in panels a, b, and c. **e** Number of E spikes within 4 ms overlapping windows, sliding every 1 ms. Note the irregularity and occasional degradation of the spiking pattern

This completes our description of the local population model; details of the LIF equations and exact parameters used are given in Supplemental Information.

2.2 Relevant facts about gamma rhythms

We recall here some facts about the spiking patterns of the local population described in Sect. 2.1. The phenomena discussed in this subsection have been observed and documented in the real cortex (Gray & Singer, 1989; Henrie & Shapley, 2005; Cardin, 2016) and have been studied in other computational models (see Sect. 2.3). We recall these properties because they play important roles in the correlation studies to follow. We also wish to demonstrate that our model fairly realistically reproduces these known cortical phenomena.

Individual vs. collective spiking behaviors. In the real cortex, interspike intervals of individual neurons have long tail distributions that have been described as being exponential (Ostojic, 2011) or obeying power laws (Baddeley et al., 1997), and Fig. 1a shows that our model exhibits a similar behavior. These long tails show that individual neurons do not spike rhythmically. The resemblance to an exponential distribution indicates that the time intervals between spikes are close to being random with a fixed mean rate.

In contrast to the individual behavior depicted in Fig. 1a, we show in Fig. 1b rasters of the population over a time interval of 500 ms. Here one observes a tendency for the spikes to occur in clusters, leading to rises and falls in firing rates that produce a rhythm in the gamma band, a phenomenon well known to occur in many parts of the brain (Buzsaki, 2011; Cardin, 2016).

In our model, this rhythmic behavior was entirely self-organized: clustering of spikes was not implied by the LIF equation, and the presence of a rhythm was not apparent from ISI plots (Fig. 1a) or the PSDs of single neurons which do not display clear peaks or decay structures). Nor was it arranged through the wiring: connectivity among neurons in the local population was random and no grouping of neurons into clusters was programmed into the model. Lastly, gamma-band rhythms occurred for wide ranges of model parameters, and no specific engineering was needed to produce the rhythm shown, though different parameter choices led to rhythms with slightly different characteristics.

The production of a gamma-band rhythm is an example of an emergent phenomenon. A network phenomenon is called *emergent* if it does not originate from the laws governing individual network components, but occurs only as a result of the interaction among them. Here, the gamma-band rhythm is a product of the dynamical interaction among

neurons. Mechanisms for its generation will be discussed in more detail in Sect. 6.

A few important properties of gamma-band rhythms.

We highlight below a few properties of gamma rhythms that are important to recall for later on in this paper.

The clustering of spikes naturally leads to some amount of synchronization. We stress that in our model, as in the sensory areas of the real cortex (Henrie & Shapley, 2005; Xing et al., 2012), this synchronization is very partial and involves only a small fraction of the local population at any given time. In the summed spike plot of Fig. 1c, each bin in the histogram is 4 ms wide, and typically no more than 10% of the local population spikes within a 4 ms period. Occasionally the fraction is larger, but it seldom exceeds 20%. This is in contrast to the original picture of PING (Whittington et al., 2000; Börgers & Kopell, 2003), where the rhythm is produced by whole-population spikes.

A second important property of gamma-band activity is that it is broad-band (Henrie & Shapley, 2005; Jia et al., 2011). In Fig. 1d, we show the power spectral density (PSD) of our model. One can see that spectral power is distributed between 50–70 Hz, which is narrower than in the real cortex because it is a small, isolated local population. Still, the PSD is far from a delta function, which would be the case had this rhythm been periodic. In earlier studies such as Brunel and Wang (2003), the authors assumed this rhythm was periodic and sought to identify the period, but that is not how gamma rhythms are in the real brain.

A third property of gamma-band activity in the real cortex is that it is irregular and episodic in nature Chariker et al. (2018); Burns et al. (2011); Xing et al. (2012), and Fig. 1e shows that our model produces a rhythm much like that: not only is there variability in the heights of the peaks and in the distances between peaks, the rhythm also degrades from time to time (as it did around 8200 ms). It is not at all a rigidly oscillating signal.

We will show in Sect. 5 that the naturally irregular nature of these rhythms enables them to synchronize source-target populations more effectively than rigidly oscillatory signals can.

2.3 Comparison with existing models

Homogeneously connected neuronal populations have been studied by many theorists. Among the numerous papers on the topic, the following are two of the most influential sets of results.

Balanced state ideas, referring to the balancing of E- and I-currents in the limit where system size tends to infinity, were introduced in van Vreeswijk and Sompolinsky (1998).

While the infinite-size limit offers mathematical tractability, we have elected not to go that route because the real cortex does not support infinitely large, homogeneously connected groups of neurons.

Another much cited body of work is Brunel and Hakim (1999); Brunel (2000). In these papers, the authors considered homogeneously connected networks, also with system size N tending to infinity, with the additional assumption that

$$\frac{1}{N}(\# \text{ connections per neuron}) \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

Because this assumption implies that for N large, the sets of neurons presynaptic to two distinct neurons are likely disjoint, the authors proposed that neurons in the population are effectively driven by a mean field quantity plus independent Gaussian noise. We do not make the sparse connectivity assumption in our model because in the visual cortex (as in many other parts of the real cortex), connectivity is not sparse: I-neurons, in particular, are quite densely connected to all nearby neurons (Holmgren et al., 2003; Oswald and Reyes, 2011); E-to-E connectivity, while lower, is still at 10 – 15%. The fact that many neurons in the population have shared presynaptic neurons has dynamical consequences; shared inputs contribute to correlated behaviors.

Gamma-band rhythms have also been studied by a number of theorists, among them Börgers and Kopell (2003); Whittington et al. (2000); Brunel and Wang (2003), and others. These earlier models successfully captured the oscillatory aspect of the phenomenon but not the broad-band, episodic nature of gamma-band activity. Irregularity of the rhythm was highlighted in Rangan and Young (2013a, b) and studied further in Chariker and Young (2015) and Chariker et al. (2018). This characteristic of gamma rhythms is crucial for the main purposes of the present study. As mentioned at the end of Sect. 2.2, it is not just the oscillatory nature of the activity but also the irregularity of it that is essential in synchronizing source-target populations.

3 Preliminaries II: Correlations between two groups of neurons

This paper is about correlations in spiking activity on the population level. The present section is theoretical: a formal definition of population-level correlation is given in Sect. 3.1, and in Sect. 3.2 we compare this quantity to correlations between pairs of neurons, which have been studied in a number of theoretical and experimental papers.

3.1 Formal definitions

The notion of correlation of interest here is the correlation between instantaneous firing rates from two groups of neurons (also used in Vogels & Abbot (2005)). We do not distinguish between spikes fired by different neurons from within each group, lumping them all together into a single spike train representing the collective output from the group.

Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be two groups of neurons with sizes n_1 and n_2 respectively. To operationally define “instantaneous firing rate”, we have to fix a small time interval, which we take to be 4 ms. For the i th neuron in $\mathcal{G}_1$, we let $x_i([t, t + 4))$ denote the number of spikes fired by this neuron on the time interval $[t, t + 4)$, where t is in ms, and take

$$X_i(t) = \frac{1000}{4} x_i([t, t + 4))$$

to be our definition of instantaneous firing rate of neuron i in spikes/sec at time t .

We further fix a large time interval, $[0, T]$ for some integer T , and view X_i as a random variable defined on the probability space

$$\{t = 0, 1, 2, \dots, T - 4\}$$

with equal probability assigned to each sample point. The random variable Y_j representing instantaneous firing rates of neuron j in $\mathcal{G}_2$ is defined analogously.

The choice of the 4 ms interval above reflects the time scales we find relevant for our purposes: we are not especially interested in pinpointing the exact timing of spikes, and 4 ms is roughly the duration during which conductances in postsynaptic neurons are elevated. It is also short enough to reflect gamma-band activity, which dominates local-in-time firing patterns. The use of overlapping windows was to further smooth out the statistics.

Next, we define two random variables

$$X = \sum_{i=1}^{n_1} X_i, \quad Y = \sum_{j=1}^{n_2} Y_j$$

representing the firing rates of the two populations. The correlation between these two random variables, i.e.,

$$\rho(X, Y) = \frac{\mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y)}{\sqrt{\sigma^2(X)\sigma^2(Y)}} \quad (2)$$

where $\sigma^2(X)$ and $\sigma^2(Y)$ are the variances of the random variables X and Y respectively, is what we will refer to as the correlation in spiking activity between populations $\mathcal{G}_1$ and $\mathcal{G}_2$. Notice that the correlation between pairs of neurons is a special case of the definition above, with $n_1 = n_2 = 1$.

3.2 Population-level correlations versus paired correlations

We study correlations in spiking activity on the population level because we believe they more accurately describe the effectiveness of communication between local source and target populations. Bulk measurements on spiking activity have, however, been challenging to collect in the laboratory up until now, and many experimental results thus far have been for correlations between pairs of neurons (Roe & Ts'o, 1999; Nowak et al., 1999; Jia et al., 2013; Zandvakili & Kohn, 2015). A number of theoretical papers have also focused on single or averaged pairwise correlations between neurons within a single population (de la Rocha et al., 2007; Ostojic et al., 2009; Renart et al., 2010).

To clarify the difference between population-level correlation as defined in Eq. (2) and correlations between pairs of neurons, we let

$$\langle \rho(X_i, Y_j) \rangle := \frac{1}{n_1 n_2} \sum_{i,j} \rho(X_i, Y_j),$$

denote the mean of the correlations between pairs of neurons, one from $\mathcal{G}_1$ and the other from $\mathcal{G}_2$, averaged over all such pairs. The relation between $\rho(X, Y)$ and $\langle \rho(X_i, Y_j) \rangle$ is described in the following. To simplify the discussion, let us assume that $\text{Var}(X_i) = \sigma_1^2$ for all the neurons in $\mathcal{G}_1$; likewise $\text{Var}(Y_j) = \sigma_2^2$ for all the neurons in $\mathcal{G}_2$. Then

$$\langle \rho(X_i, Y_j) \rangle = \frac{1}{n_1 n_2} \frac{\sum_{i,j} \text{Cov}(X_i, Y_j)}{\sigma_1 \sigma_2}$$

where $\text{Cov}(\cdot, \cdot)$ is the covariance of the two random variables. As for $\rho(X, Y)$, we have

$$\rho(X, Y) = \frac{\sum_{i,j} \text{Cov}(X_i, Y_j)}{\sqrt{(n_1 \sigma_1^2 + C_1)(n_2 \sigma_2^2 + C_2)}} \quad (3)$$

where

$$C_1 = \sum_{1 \leq k, l \leq n_1, k \neq l} \text{Cov}(X_k, X_l),$$

and C_2 is defined analogously with Y_k, Y_l and n_2 in the place of X_k, X_l and n_1 .

From the formulas above, it is easy to glean the relation between $\rho(X, Y)$ and $\langle \rho(X_i, Y_j) \rangle$ in the following two situations:

(a) If both populations fire only whole-population spikes, i.e. $X_i = X_j$ for all i, j , and $Y_i = Y_j$ for all i, j , then one would expect $\rho(X, Y) = \langle \rho(X_i, Y_j) \rangle$, and the formulas above confirm this with $\sigma^2(X) = n_1^2 \sigma_1^2$, $\sigma^2(Y) = n_2^2 \sigma_2^2$.

(b) On the other hand, if the spike trains from distinct neurons within $\mathcal{G}_1$ (respectively $\mathcal{G}_2$) are entirely uncorrelated,

i.e. if $\text{Cov}(X_i, X_j) = 0$ for all $i \neq j$ and similarly for Y_i, Y_j , then $\sigma^2(X) = n_1 \sigma_1^2$, $\sigma^2(Y) = n_2 \sigma_2^2$, and

$$\rho(X, Y) = \frac{\sum_{i,j} \text{Cov}(X_i, Y_j)}{\sqrt{n_1 n_2} \sigma_1 \sigma_2}.$$

If one assumes that correlations between neurons within each population are nonnegatively correlated (which generally seems to be the case), then the two cases above represent the two ends of the spectrum of possible values for $\rho(X, Y)$:

$$\langle \rho(X_i, Y_j) \rangle \leq \rho(X, Y) \leq \sqrt{n_1 n_2} \langle \rho(X_i, Y_j) \rangle.$$

Notice from Eq. (3) that $\rho(X, Y)$ contains information about paired covariances *within* $\mathcal{G}_1$ and $\mathcal{G}_2$. The more synchronized each network is, the closer $\rho(X, Y)$ is to $\langle \rho(X_i, Y_j) \rangle$. At the other end of the spectrum, when spike firing within each network is close to being independent, $\rho(X, Y)$ can be larger than $\langle \rho(X_i, Y_j) \rangle$ by a factor comparable to network size. In a normal cortex, spike firing within local populations is not uncorrelated, but is also far from fully synchronized. Thus, the quantity $\rho(X, Y)$ generally lies somewhere between the two bounds above.

4 Time-adjusted correlations between source and target networks

We are now ready to proceed to the computational part of our study, to investigate population-level correlations between source and target networks. The setup for the rest of this paper is described in Sect. 4.1. This is followed by some preliminary observations. The main results are presented in Sect. 5.1.

4.1 A two-component feedforward network model

Given two networks, $\mathcal{N}_1$ and $\mathcal{N}_2$, of the type in Sect. 2.1, and a number $p \in (0, 1)$, we describe here how to construct a simple feedforward network with “connectivity” p , henceforth abbreviated as

$$\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2.$$

Here, $\mathcal{N}_1$ is the source (sending) network, $\mathcal{N}_2$ is the target (receiving) network. Connections from $\mathcal{N}_1$ to $\mathcal{N}_2$ are excitatory only, and excitatory input from $\mathcal{N}_1$ targets both E and I-neurons in $\mathcal{N}_2$. The meaning of the number p is as follows: Each neuron in $\mathcal{N}_2$ receives, on average, a fraction p of its total excitatory input from $\mathcal{N}_1$. “Total excitatory input” includes synaptic input from E-neurons within $\mathcal{N}_2$, the two

external sources described in Sect. 2.1, and feedforward input from $\mathcal{N}_1$.

The values of p of interest are generally ≤ 0.1 . Given p , we modify the source of external input to $\mathcal{N}_2$ as follows: Assuming that the total excitatory input to $\mathcal{N}_2$ is the same as that to $\mathcal{N}_1$, we first compute the amount of external synaptic input that corresponds to a fraction p of the total excitatory input. Let us refer to these amounts for E and I-neurons as $x_E = x_E(p)$ and $x_I = x_I(p)$, respectively. Then for each E (respectively I) neuron in $\mathcal{N}_2$, we remove x_E (resp. x_I) amount of synaptic input from its external drive and add connections from $\mathcal{N}_1$ whose spikes will replace the lost input. Presynaptic neurons from $\mathcal{N}_1$ are chosen randomly, and we assume an additional 1 ms transmission time for spikes from $\mathcal{N}_1$ to $\mathcal{N}_2$.

Further details of the construction of $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$ are given in [Supplementary Information](#).

4.2 Intrinsic optimal delays

Having built the network $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$, we now consider the correlation between the two E-populations. Using the notation of Sect. 3.1, we let $\mathcal{G}_1$ be the set of excitatory neurons in $\mathcal{N}_1$, and $\mathcal{G}_2$ the set of excitatory neurons in $\mathcal{N}_2$. Let X and Y be the random variables representing the instantaneous population firing rates of $\mathcal{G}_1$ and $\mathcal{G}_2$, and $\rho(X, Y)$ denote their correlation. To explore if the spiking events of the two populations would correlate better when measuring the spike times in the receiving network $\mathcal{G}_2$ with a delay, we considered also the quantities

$$\rho(X, Y_d) \quad \text{where} \quad Y_d(t) = Y(t + d)$$

for $d > 0$.

To locate the time delay that maximizes the correlations between $\mathcal{G}_1$ and $\mathcal{G}_2$ (if there is one), we computed $\rho(X, Y_d)$ for various values of d , at 0.5 increments from 0 to 30 ms with $p = .075$, i.e., when 7.5% of the excitatory drive in $\mathcal{N}_2$ comes from $\mathcal{N}_1$. (Values of p in the real cortex obviously vary, but $p \sim 0.075$ to 0.1 is thought to be fairly typical). Figure 2a shows the function $d \mapsto \rho(X, Y_d)$. The 10 graphs superimposed are for 10 different networks (constructed with the same connection probabilities) and 10 different sets of initial conditions. The results from the trials are remarkably similar, with the optimal time delays between 3.5 – 4 ms. In our simulations, convergence to $\rho(X, Y_d)$ for each d was very fast, allowing an approximate value of the optimal delay to emerge in less than a second, and the convergence to be complete in the next few seconds.

Note that the 3.5 – 4 ms optimal time delay observed is significantly longer than the 1 ms transmission time imposed on spikes from $\mathcal{N}_1$ to $\mathcal{N}_2$. This is because spikes from $\mathcal{N}_1$ do not immediately cause spikes in $\mathcal{N}_2$. They raise the excitatory conductance for neurons in $\mathcal{N}_2$, bringing their membrane potentials closer to threshold, and that increases their susceptibility to spike in the presence of more excitatory input.

Figure 2b shows the functions $X(t)$ and $Y(t + d)$ as functions of t for $d = 0$ (top) and $d = 3.5$ (bottom). In the top plot with no time delay, $\mathcal{N}_1$'s activity (red) peaks a little ahead of $\mathcal{N}_2$'s most of the time. In the bottom plot where the delay is optimal, we observe an excellent alignment of the gamma peaks produced by $\mathcal{N}_1$ and $\mathcal{N}_2$. Such alignments are not always present, however, because gamma rhythms degrade from time to time, as can be seen at the beginning of the time interval shown.




Fig. 2 Correlations with time delay in network activity for $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$, $p = 0.075$. **a** Computed are correlations between $X(t)$ and $Y(t + d)$ as functions of d (for 0.5 ms increments of d) for two networks connected as in Sect. 4.1. Each color represents a different trial using a different network with the same connection probabilities. The loca-

tions of the first peaks, which occur at approximately 3.5 – 4 ms for all of the trials, are taken to be the response time of $\mathcal{N}_2$ to $\mathcal{N}_1$. **b** shows the superimposed plots of $X(t)$ and $Y(t + d)$ for one of the trials in panel a, for $d = 0$ and 3.5 ms on the top and bottom, respectively; the graph of $X(t)$ is in red, and that of $Y(t + d)$ is in black

The existence of an intrinsic optimal time delay is an emergent phenomenon. We stress that this is a much stronger statement than the existence of a time shift that maximizes correlations for each initial condition. The notion of optimal time-delay we have observed is *independent of initial condition*.

The presence of a response time or phase-shift between post-synaptic and pre-synaptic firing has been observed experimentally (Fries, 2005; Jia et al., 2013; Bastos et al., 2015; Fries, 2015; Zandvakili & Kohn, 2015). In particular, Zandvakili and Kohn (2015) reported an increased probability of a V2 spike 3 ms after a V1 spike. Though this result was computed using pairs of neurons, it was in the ballpark of the optimal time delays we have found. We are unaware of in-depth theoretical studies of this issue prior to the present work.

4.3 Dependence on sample size

Although improvements in technology have enabled one to record from larger and larger neural samples and these capabilities are bound to improve further in the near future (see e.g. Steinmetz et al. (2018)), current recording technology cannot simultaneously record all spiking activity from local populations. An important question, therefore, is how many neurons are needed to obtain an accurate estimate of the population correlation and optimal time delays.

We studied this in the $\mathcal{N}_1 \rightarrow \mathcal{N}_2$ network using samples $\mathcal{G}_1$ and $\mathcal{G}_2$ consisting of N E-neurons from each layer, where N varied from 1 to the size of the full population. The correlations, ρ_N , for $p = 0.075$ computed with optimal time delays are shown in Fig. 3a, with 5 trials for each N . Ranges of the optimal time delays in the computation of ρ_N are shown in Fig. 3b.

Here, $N = 1$ corresponds to correlations between pairs of neurons. We see that the numbers are very small, about a tenth of correlations between entire E-populations. In addition to having small correlations, the trial-to-trial variances are large, having the same order of magnitude as the correlations themselves. These observations are consistent with individual neuron spikes being almost “random”, as depicted in Fig. 1a.

We also see that as N increases from a small value, ρ_N increases rapidly, starting to stabilize at about $N = 100$ and eventually asymptoting to the population value.

The effects of subsampling can be deduced from Eq. (3) as follows: Assuming that the covariances between neurons in $\mathcal{G}_1$ are identical, and the same is true for $\mathcal{G}_2$, we let c_1 and c_2 denote these paired covariances. We also assume that the covariances between any two neurons, one from $\mathcal{G}_1$ and one from $\mathcal{G}_2$, are identical and denote this value with c_{12} . Then it follows from Eq. (3) that

$$\rho_N = \frac{Nc_{12}}{\sqrt{(\sigma_1^2 + (N-1)c_1)(\sigma_2^2 + (N-1)c_2)}}. \quad (4)$$

As σ_i^2 is significantly larger than c_i (confirmed in data from our simulations), σ_i^2 is more dominant than $(N-1)c_i$ for small N , so the right side of Eq. (4) is roughly proportional to N for N small, as can be seen in the inset in Fig 3a. But c_1 and c_2 are not zero in part due to gamma-band activity within local populations (see Fig 1). This causes the expression on the right side of Eq. (4) to tend to a constant as N increases. The analysis above explains the shape of the plot in Fig 3a.

As an application to experimental neuroscience, Eq. (4) can be used to deduce the true values of population correlations from estimates of σ_i , c_i , and c_{12} .




Fig. 3 Correlations and optimal delays as functions of sample size N . **a** Correlations between two samples of size N from $\mathcal{N}_1$ and $\mathcal{N}_2$ in a $\mathcal{N}_1 \rightarrow \mathcal{N}_2$ network with $p = 0.075$. Results are based off of one, 8 second long simulation, with correlations computed for $N = 1, 10, 20, 30, \dots, 290$ randomly selected E-neurons. Shown are results of 5 trials (a different sample for each trial) for each N .

The population correlation using all 300 neurons is about 0.33 and is shown by the gray, dashed line. The inset shows more detail for smaller samples, with $N = 1, 3, 6, 9, \dots, 30$. **b** Maximum and minimum optimal time delays as a function of N , for the same simulation in panel a. The range of time delays was computed over 10 trials (using a different sample of neurons in each trial)

Another noteworthy – and somewhat surprising – finding is that when using as few as five neurons from each layer, the optimal time delays already stabilized and were around 3.5–4 ms, which was the optimal time delay for the full population found in Sect. 4.2; see Panel b in Fig. 3. Thus unlike population-level correlations, which are strongly dependent on sample size, optimal delay measurements stabilize quickly at small N and can already be reasonably captured with present recording technology.

5 Correlations: modeled neuronal populations vs periodic oscillations

Section 5.1 contains a more systematic study of correlations on the population level for the two-component feedforward networks constructed in Sect. 4.1. To explain the robust correlations observed, we examine an earlier proposal that firing rate oscillations in gamma rhythms may be responsible for synchronizing the two populations (Fries, 2005; Fries, 2015). In Sect. 5.2, we present results on correlations between source and target signals, both of which are assumed to be regular periodic oscillations, and compare them to our findings for the network $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$.

5.1 Correlations in our 2-component model: varying connectivity and gamma characteristics

To consolidate the results in Sect. 4.2, and to further analyze the relationships between correlation, network properties, and gamma characteristics, we carried out a systematic study consisting of three sets of simulations. All pertained to two-component networks of the form $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$. In each of the three sets, correlations were computed for $p = 0.05, 0.075, 0.10$ and 0.15 , a range thought to resemble connectivities in the real cortex. The local populations $\mathcal{N}_1$ and $\mathcal{N}_2$ were as described in Sect. 2.1. For all three sets, $\mathcal{N}_1$ and $\mathcal{N}_2$ had similar connectivities and firing rates, but gamma characteristics were varied.

The local populations with distinct gamma characteristics were of two flavors, to be referred to as “normal” and “synched”. “Normal” regimes are those depicted in the simulations in Sect. 2, so-called because the parameters there were chosen to produce gamma-band activity that emulates those in the sensory cortices (Henrie & Shapley, 2005). “Synched” regimes are, as the name suggests, more synchronized. Different circumstances, such as at the onset of a stimulus presentation, increased attention (Buzsaki, 2011), or the effects of drugs (e.g. anesthesia or ketamine) (McCarthy et al., 2012), may produce regimes that are more synchronized than our “normal” regimes. Our “synched” regimes have lower peak gamma frequencies, because when firing rates are maintained but more neurons participate in each spiking event – a definition of greater synchrony – there must be

fewer of such events, causing them to be further apart in time. Systematic ways to produce such regimes were carried out following ideas in Chariker et al. (2018) and parameters are given in Supplementary Information.

The three sets of simulations performed were for (i) normal driving normal, i.e., both $\mathcal{N}_1$ and $\mathcal{N}_2$ were normal, (ii) synched driving normal, i.e., $\mathcal{N}_1$ was synched and $\mathcal{N}_2$ was normal, and (iii) normal driving synched. The results of this subsection are summarized in Fig. 4. Correlations adjusted for optimal delay are shown in Fig. 4a; PSD and rasters for a synched regime are shown in Fig. 4b and should be compared to the corresponding plots for the normal regime shown in Fig. 1. Firing rates of the receiving population $\mathcal{N}_2$ are shown in Fig. 4c.

As can be seen from Fig. 4a, the correlation values computed in each of the three studies showed a steady rise as functions of connectivity p , and for each value of p , correlation values for independently drawn networks (using the same connection probabilities) and independently drawn initial conditions produced results that varied only mildly. The consistency seen in over 100 runs confirmed that these numbers are representative of the type of populations presented in Sect. 2.1 and the feedforward construction of Sect. 4.1. The existence of an optimal delay for each value of p independent of the network drawn and independent of initial condition was also confirmed.

We identify the following observations from Fig. 4a as being notable and requiring explanation:

- (a) The first point is the magnitudes of the correlations. That they increased with p in each study was expected, but values of correlations between 0.2 and 0.6 are quite large when $\mathcal{N}_1$ supplied such a small fraction of the excitatory current to neurons in $\mathcal{N}_2$.
- (b) The existence of optimal delays requires both an explanation and an interpretation.
- (c) Our third observation is that correlations are much higher when the source network is synchronized, i.e., synchronized systems entrain the spiking in regions downstream more effectively. They are also less susceptible to entrainment by source systems that are less synchronized.
- (d) The last item is about firing rates. Notice that for corresponding values of p , synchronized sources caused only a slight rise in firing rates in the target networks. They produced much stronger correlations, but not necessarily higher firing rates.

Some of the points above will be explained mechanistically in Sect. 6.2. As the oscillatory behavior of gamma-band activity clearly played a role in the results above (see Fig. 2b), we first have a look at the case of regular




Fig. 4 Correlations and firing rates as functions of connectivity and degree of synchrony. **a** Correlation of source and target networks as a function of connectivity. Shown are peak correlations as functions of p . The three different plots represent different source → target pairings. For the black line representing a normal → normal system, the colored dots are results from 5 trials using 5 different networks drawn with the same parameters. Note the strong correlations, as well as the low trial-to-trial variability. The values for the other lines were

also averaged over 5 trials. Ranges for optimal time-delays are written above each correlation value. Note that the synched source regime produced higher correlations, while the synched target regime lowered correlations. **b** Raster and PSD of a synched regime. **c** Mean firing rates of $\mathcal{N}_2$, the target network. x -axis is the percentage connectivity, 100p. Bar graphs show the average of 5 trials, 8 sec each. Note that for $p \leq 0.1$, firing rate did not increase appreciably with the increased synchrony of the source network

periodic oscillations to better understand the conceptual differences, if any.

5.2 Correlations between periodically oscillating systems

The robust correlations in Fig. 4a, which can be interpreted as effectiveness of communication between the two populations, together with item (c), the observation that synchronized networks better entrain populations downstream, are reminiscent of a body of work in the literature known by the name of *communication through coherence* (Fries, 2005; Fries, 2015).


In Fries (2005), the author proposed that communication between two neuronal groups mechanistically depended on coherence between them, and that the basis of this

coherence was neuronal oscillations. The author pointed to oscillatory synchronization in the source network, together with phase-locking between source and target groups, as being essential for effective communication. They placed a great deal of importance on the regularity of the oscillatory behavior and remarked that the absence of a reliable phase relation between the oscillations in the sending and receiving groups would be detrimental to communication.

In this subsection, we examine numerically the role of oscillatory behavior, as suggested in Fries (2005), as the sole mechanism for producing correlations, and compare the results to those in Sect. 5.1.

The following setup is considered. We assume, for simplicity, that before the source and target networks are connected, their oscillatory behaviors are represented by functions f_1 and f_2 respectively, where f_i has the form

Fig. 5 Correlations between periodically oscillating signals; see text for definitions of f_1 and $\tilde{f}_2$. **a** Three trials in a system in which the source and target have identical frequencies. Unadjusted correlations are dramatically different depending on the initial phase difference between them. **b** Correlation for a system where the source and target have different frequencies



$$f_i(t) = \sin(\omega_i t + \phi_i), \quad i = 1, 2.$$

In analogy with the notion of connectivity between $\mathcal{N}_1$ and $\mathcal{N}_2$, connectivity p here translates into modifying $f_2(t)$ so that it becomes

$$\tilde{f}_2(t) = pf_1(t) + (1 - p)f_2(t).$$

Correlations between $f_1(t)$ and $\tilde{f}_2(t)$ are then computed as before.

Figure 5a shows three simulations for which $\omega_1 = \omega_2$ for various values of ϕ_i at $p = 0.1$. Here we see that when f_1 and f_2 have compatible phases, such as when the two phases differ by 0.1π , oscillatory behavior can indeed be a powerful vehicle for promoting strong correlations, but the result is entirely dependent on the phase relation between the two oscillations at $t = 0$, and ϕ_i can be chosen to produce positive or negative correlations — if we do not incorporate a time delay into the computation of correlations. If we allow for time delays, then trivially the optimal delay is the shift that makes the two phases coincide at $t = 0$ and is entirely dependent on initial condition.

Figure 5b shows a simulation where $\omega_1 = 0.91 \times \omega_2$. The unadjusted correlation was computed to be 0.11. As can be seen from the plot, $\tilde{f}_2(t)$ advances in phase by about 9% each cycle, so that in about 11 cycles, the picture is essentially repeated. During these 11 cycles, the relative phases will vary from nearly coinciding, to almost anti-phase, and then return to coinciding. It is a mathematical fact that if we run the system for a long enough time, correlations will be the same, i.e. ~ 0.11 , independent of initial phases and independent of ω_1 and ω_2 as long as ω_1/ω_2 is irrational. If we run it for a short time, then correlations can vary depending on which stretch of the phase sequence we sample.

To summarize, for rigid oscillations phase-locking is impossible without the frequencies of the source and target being identical, and when they are identical, the phase relation without time adjustment can range from the two systems oscillating completely in phase or anti-phase (or anything in between). Time adjustments can align the phases but the amount of adjustment depends on the initial condition.

This picture differs substantially from what was observed in our simulations for $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$ as demonstrated in Sect. 5.1. For neuronal models, there is partial phase agreement (or alignment of gamma events) under a wide range of conditions: following a time adjustment that is intrinsic to the system, this alignment holds independently of initial conditions and without preconditions on the peak gamma frequencies of source and target. How this flexibility is achieved is the topic of the next section.

6 Mechanistic explanation for the robust correlations between source and target populations

That gamma-band rhythms are implicated in the relatively high correlations between source and target networks is abundantly clear; it was proposed in Fries (2005, 2015) and we have also seen it in our own simulations in Fig. 2b. We have also seen in Sect. 5.2 that oscillations alone cannot explain the phenomena observed. The aim of this section is to provide mechanistic explanations for the results reported in Sect. 5.1, and to do that, we need to first understand the mechanisms behind the generation of gamma-band activity in single populations, which we review in Sect. 6.1.

6.1 Mechanistic explanation for the generation of gamma-band rhythms in single populations

A widely accepted explanation for gamma-band activity proposed 20 years ago is PING (Whittington et al., 2000; Tiesinga et al., 2001; Börgers & Kopell, 2003; Börgers, 2017). The original picture of PING (of which there are variants e.g. “weak PING” (Börgers et al. (2005))), consists of a steady external drive together with E-to-I and I-to-E couplings within the local population: the external drive causes all of the E-neurons to spike essentially simultaneously; that causes all of the I-neurons to spike, which suppresses the entire population for a certain time period, until the external drive prevails, leading to an E-population spike and a repeat of the cycle.

PING produces population spikes that are periodic and highly regular. While this may be how gamma rhythms are in certain regions of the brain, we are primarily interested in the sensory cortices where experimental data do not support periodic population spikes.

The author of Brunel (2000) proposed to view gamma rhythms in terms of a Hopf bifurcation but that is for their reduced models, not the network. We do not know if the network regimes that motivated this reduction will support the gamma-band activity described in Sect. 2.2: to our knowledge, that has not been documented.

Another much cited theoretical paper on gamma rhythms is Brunel and Wang (2003), which had the theoretical focus of understanding the frequency understanding the frequency of the rhythm. The existence of a frequency is also inconsistent with sensory data, which showed that gamma rhythms are broad-band with wandering frequencies; see e.g. Xing et al. (2012). In any case, Brunel and Wang (2003) did not offer a mechanistic explanation.

A mechanistic depiction of neuronal behavior that lies behind gamma-band activity is the phenomenon of multiple firing events (MFEs) first proposed in Rangan and Young (2013a, b) and subsequently studied in Chariker and Young (2015). MFE refers to the precipitation of a spiking event initiated by the crossing of threshold by a few E-neurons due to external drive and/or normal current fluctuations. These first spikes lead to recurrent excitation, which may or may not cause other neurons to spike. When it leads to further spiking activity, both E- and I-neurons are activated and the firing event can last for 2 to 3 ms until activity is curtailed due to the suppressing effects of inhibition.

The idea that MFEs play an instrumental role in the generation of gamma-band rhythms was proposed in Rangan and Young (2013b) and Chariker and Young (2015), and re-examined against biological data in Chariker et al. (2018) using a previously constructed, realistic model of the visual cortex. The authors of Chariker et al. (2018) called this mechanism *recurrent-excitation-inhibition* (REI).

A mechanistic description of REI is as follows: When the crossing of threshold of some E-neurons leads to an MFE as described above, most of the neurons in the local population are left hyperpolarized in its aftermath because I-cells are activated along with E-cells, and they are quite densely connected. The decay of I-conductance and the depolarization of E-cells due to external input leads eventually to the next event. This explains the generation of a rhythm. The time constant for I-conductance decay in the LIF equations places the frequency of these spiking events in the gamma band.

A major difference between REI and PING is that in REI, when an MFE is precipitated, it can involve variable fractions of the population depending on the membrane potentials of the neurons postsynaptic to the ones that initiated the event, the E-I composition of the set of neurons activated, and the speed of I-neurons to activate and quell the barrage of spiking. Inter-event times are also variable: if many I-neurons are involved, I-conductance will build up, delaying the onset of the next event. Moreover, spiking events can degrade, as the initial spiking may not produce a substantial MFE. These types of variability lead to the irregular, episodic nature of gamma-band rhythms like those observed in the real cortex.

The single population models presented in Sect. 2.1 successfully captured this variability. While more analytically tractable models of single populations, such as those in van Vreeswijk and Sompolinsky (1998); Brunel (2000), may be appealing, we have not used them in our study because of a priori concerns that the diminished recurrent excitation and inhibition in these regimes – properties known to be implicated in the production of gamma-band activity – may impact the transfer of gamma rhythms between populations.

6.2 Explaining how gamma rhythms synchronize source and target populations

We now return to the two-component feedforward model $\mathcal{N}_1 \xrightarrow{P} \mathcal{N}_2$, and offer mechanistic explanations for phenomena (a) – (c) observed in Sect. 5.1 (restated below).

Our working hypothesis, following the proposal in Fries (2005), is that the strong correlation between $\mathcal{N}_1$ and $\mathcal{N}_2$ is in large measure derived from the synchronization of their gamma rhythms: the local-in-time firing rates of each population rise and fall with frequencies in the gamma band, and there is a tendency – considerably beyond pure chance – for the peaks of the firing rates in $\mathcal{N}_1$ to coincide with those in $\mathcal{N}_2$.

The hypothesis above leads immediately to a number of intriguing questions: (i) When unconnected, there is no reason for the gamma rhythms produced by $\mathcal{N}_1$ and $\mathcal{N}_2$ individually to be related in any way. How can a mere

7.5% connectivity cause the peaks to show such nontrivial alignment after a suitable time-delay adjustment? (ii) Why should there be a notion of intrinsic optimal time delay independent of initial condition? (iii) Why do more synchronized sources produce higher correlations, as in the “synched driving normal” case in Fig. 4a? Why is it that unlike rigid oscillations, two networks can produce high correlations even when they have incommensurate frequencies?

Below we propose answers to these questions. Our proposals are based on a combination of experimental data, known theoretical results, previous modeling work, and our own simulations.

With regard to question (i), observe first that though each local population generates a gamma rhythm through its internal dynamics in the sense that it produces waves of recurrent excitation followed by suppression, *there is no intrinsic timing associated with this rhythm*, i.e., there is no fixed clock to which the population must adhere. This is documented in the experimental literature; see Burns et al. (2011). By altering the external drive supplied to the local population, such as by increasing this drive during the depolarization phase, one can hasten the onset of the next MFE. Likewise, withdrawal of some of the drive can have the opposite effect. In other words, gamma rhythms are *malleable* – the timing of firing events are influenced by the input signals received by the local population.

As to why so low a connectivity from $\mathcal{N}_1$ can produce such a robust correlation, i.e., why spike times in $\mathcal{N}_2$ adapt readily to spiking from $\mathcal{N}_1$ when $\mathcal{N}_1$ provides only a small fraction of the total excitatory input to neurons in $\mathcal{N}_2$, the answer has to do with the fact that the excitatory and inhibitory currents received by a neuron are well balanced, not just when averaged over time as proposed in the well known theory of balanced states (van Vreeswijk & Sompolinsky, 1998), but also moment by moment; see the experimental results of Okun and Lampl (2008), the modeling results of Chariker et al. (2018), Fig. 4, and Joglekar et al. (2019); an example of this balance in our models is shown in Supplementary Information. Because of this tight balance, any temporary excess in excitation increases the possibility of producing a spike, making a cell especially sensitive to pulses from external sources. This phenomenon was first pointed out in Chariker et al. (2018); the authors described it as “the unreasonable effectiveness of external inputs”. *Strong sensitivity to external pulses contributes to the malleability of gamma rhythms.*

The discussion above takes us naturally to question (ii), which asks why there should be an intrinsic time delay. During a spiking event in $\mathcal{N}_1$, the synaptic input to neurons in $\mathcal{N}_2$ is increased, elevating the probability of a spiking event in $\mathcal{N}_2$ a few ms later. Not every spiking event in $\mathcal{N}_1$ will result in a spiking event in $\mathcal{N}_2$; it is just that the probability is increased. We propose that the intrinsic notion

of optimal delay we have observed is the statistical mean of the response time of neurons in $\mathcal{N}_2$, i.e., the average time that it takes for the membrane potentials of neurons in $\mathcal{N}_2$ to build up following an upsurge of synaptic input. This proposal is supported by the fact that typical optimal delay times are on the order of 3 – 4 ms, while the additional delay in transmission from source to target that we have imposed is 1 ms, and 2 – 3 ms is roughly the build-up time for MFEs.

For each fixed delay time d , the almost-sure convergence of correlations is in fact not surprising. It is a consequence of the ergodic theorem; our neuronal model with its Poisson drive is almost for certain ergodic (whereas rigid oscillations with identical frequencies are not). On the intuitive level, the value of d at which the maximum of $\rho(X, Y_d)$ is attained can be seen as the statistical mean of response times, but to make that more precise, one will have to define “an event”, or what constitutes “a response” to an event – clarification of these concepts we leave to future work.

We wish to point out that implicit to the idea of responses and response times is a presumed connection between action and reaction: a presumption that the target population adjusts its spiking patterns to those of the source – and this is part of what we mean by malleability. Indeed, when connectivity tended to zero, we found that correlations also became very small with no clear optimal time delays at which a peak occurred, even though both populations possessed similar spectrum profiles. In other words, the response of the target population to the source’s firing weakens as connectivity decreases.

In our simulations much of the convergence in correlations including the emergence of an optimal delay occurs within the first second. Given that the dynamical system has $\mathcal{O}(1000)$ state variables, this is very fast. Since correlations come from the alignment of gamma events, rapid convergence speaks to the flexibility of gamma dynamics in the target population, and how readily it adjusts its firing events to align them with those of the source.

Coming now to question (iii), there are two ways in which synchronized source networks behave differently with regard to producing correlations. The first is that synchronized regimes in $\mathcal{N}_1$ produce larger MFEs, meaning the number of E and I-neurons participating in each spiking event is, on average, larger; equivalently, the peaks in local-in-time firing rates are taller. Larger MFEs also produce stronger pushback by I-cells, leading to longer lulls between MFEs. More clearly defined firing events and more concentrated synaptic output during such events makes synchronized source networks especially effective at entraining the spikes of their targets. This leads to higher correlations.

But as explained in Sect. 5.1, synchronized networks (with the same firing rates as those in “normal” regimes) have lower peak frequencies. This implies that when a synched network drives a normal one, we are necessarily in a situation of




Fig. 6 Optimal-delay shifted spike count plots and rasters of the $\mathcal{N}_1 \xrightarrow{p} \mathcal{N}_2$ network with $p = 0.75$. Summed spikes of the source network is in red; summed spikes of the target network is in black. Both summed spikes and rasters are time shifted by the optimal delay to show the best alignment. **a** Normal driving normal system. Note the

incompatible frequencies. Indeed one can see in Fig. 6b, steady drifts from time to time in the relative phases between $\mathcal{N}_1$ and $\mathcal{N}_2$, much like those in Fig. 5b, but the patterns are not as rigid; they are interrupted by irregular behaviors and sometimes self-adjust.

Thus in the case of a synchronized network driving a normal one, the size of the correlation is influenced by two competing forces. Other things being equal, incompatible frequencies probably do lower correlations, but here it appears that the opposing forces prevail: The synchronizing power of larger MFEs in $\mathcal{N}_1$, aided by the malleability of the $\mathcal{N}_2$'s rhythm and its ability to reset occasionally to disengage from incompatible frequencies, leads to higher correlations than in the normal driving normal case in Fig. 4a.

This completes our answers to the three questions above. We mention in closing that unlike in systems with rigid oscillations, correlations attributable to the synchronization of gamma-band rhythms in source and target populations can never rise to anything close to 1, i.e., source and target populations can never be perfectly correlated. This is because gamma rhythms suffer inevitable degradations from time to time, a fact well documented in the experimental literature (Xing et al. (2012)). It is interesting that this feature both prevents correlations from becoming too high, and helps to keep them from being too low by offering the opportunity to reset when faced with incompatible initial conditions and/or frequencies. This together with the malleability of gamma rhythms as explained in this section is what produces the correlation values of 0.2 – 0.6 seen in our simulations.

7 Discussion

This is a theoretical paper studying communication between source (sending) and target (receiving) neuronal populations. We viewed the co-fluctuation in spiking activity between

matched frequencies of the source and target. **b** Synched driving normal system. Despite incompatible frequencies, the spiking events of the source and target match remarkably well, producing a correlation of 0.4 – 0.5 (Fig. 4a), which is significantly above those of rigid oscillatory systems with incompatible frequencies (Fig. 5b)

the two populations as an indicator of the effectiveness of communication, and proposed to capture this using a population-level correlation.

Our main results

Our primary results describe salient features of source-target correlations for which we offer mechanistic explanations. The amplification of population-level correlations through the malleability of gamma rhythms is one of our most exciting discoveries. *Our secondary results* relate population-level metrics to paired correlations and subsampling. We summarize below a few highlights of these two groups of results.

Three basic features of source-target correlations that we found can be summarized as follows. First, with connectivity between source and target networks as low as 7.5 – 10%, we found that correlations were robust and consistently within the range of 0.3 – 0.5 (for “normal driving normal”) after we adjusted for a delay in the spike times of the target network. Second, we identified an optimal delay insensitive to network details that maximizes correlations, and proposed to interpret it as the mean “response time” of the target to the source. A third observation is that synchronized networks are more effective than normal networks in entraining spike times downstream. These observations are recorded in Sect. 5.1.

The high population correlations are without a doubt mediated by gamma rhythms, as has been proposed in Fries (2005, 2015), i.e., these correlations are due in part to the alignment of the peaks and troughs of firing rates of the source and target populations. But rigid oscillations, which were used in Fries (2005, 2015) for illustration, cannot explain many of the phenomena observed (see Sect. 5.2), because they depend on exact frequency compatibility and alignment of initial phases

– correlations between real brain regions could not depend so delicately on such quantities!

What we found was that the irregular, even episodic, nature of gamma rhythms, together with their malleability, greatly enabled the alignment of gamma peaks between the source and target. Gamma rhythms in our models, as in the real brain, are irregular and naturally degrade from time to time, allowing the rhythms to reset. The malleability of the rhythm in the target network then allows it to align itself to that of the source network. This alignment is achievable with a relatively small amount of input from the source because of the tight moment-by-moment balance between excitatory and inhibitory currents that renders external pulses extremely effective. These ideas, which we regard as among the most important points of this paper, are explained in detail in Sect. 6.2.

Turning to our second group of results, we have advocated in this paper for the use of population-level metrics on correlations, while most of the results on correlations in the literature pertain to pairs of neurons. We also note that other metrics, such as mutual information, have been used in the literature (Dayan & Abbot, 2001; Grün & Rotter, 2010). We chose to use correlations to describe both response and response times because the correlation carries explicit information on the simultaneous rise and fall of local-in-time firing rates in the two populations. Quantitative relations between paired and population-level correlations are presented in Sect. 3.2, where it was revealed that hidden in the relationship is the degree of synchrony among neurons *within* the two networks. We discuss also in Sect. 4.2 the effects of subsampling, deriving a mathematical formula for the dependence of correlations on sample size.

Relation to existing literature

This work is related to a number of topics in the neuroscience literature. We discuss below three areas to which our results are most closely connected.

Closest to the present work is Fries (2005, 2015). Our study is along similar lines, but we used semi-realistic neuronal networks (something not done in Fries (2005, 2015)) and a more accommodating measure of coherence. To the degree that neuronal oscillations contribute to increased correlations, we agree with his findings; we find also using our network models that synchronized sources produce larger correlations (Sect. 5.1). It has also been suggested that communication in Fries (2005, (2015)) referred to measurable increases in firing rates of target regions. Our findings do not support the idea that higher correlations are necessarily accompanied by higher target firing rates (Fig 4c). Our main contribution to this topic is to point out that unlike rigidly oscillating systems where source and target are either phase-locked or

phase-incoherent, models with more realistic depictions of gamma activity reveal a more nuanced picture: coherence is never perfect but it is often significant, due to the absence of intrinsic clocks in gamma rhythms and the target population's tendency to align its gamma events with those in the source.

A second topic to which this work is intricately related is that of gamma-band activity, the properties of which have been well documented in the experimental literature (Gray & Singer, 1989; Henrie & Shapley, 2005; Jia et al., 2011; Burns et al., 2011; Xing et al., 2012). While the oscillatory behaviors of gamma rhythms are well known and have been hypothesized (though not confirmed) to be related to various neural phenomena (Pesaran et al., 2002; Sederberg et al., 2003; Buzsaki, 2011), the irregularity, not to mention malleability, of these rhythms had, up until recently, not been connected to known neural phenomena. Besides our own work here, the only other paper we know of that exploited the broad-band nature of gamma rhythms is Palmigiano et al. (2017). The authors of Palmigiano et al. (2017) pointed to the non-rigidity of gamma rhythms as a possible advantage for information routing in multi-component networks. There is some degree of overlap between their conclusion and ours. A major difference is that we have supported our findings with an in-depth analysis of neural mechanisms. We have offered novel mechanistic explanations for how the irregularity and malleability of gamma rhythms enhances communication.

A third topic impacted by this work is the relationship between population and individual neuron activity. In part due to the use of electrophysiology, much of neuroscience (both experiments and theory) has focused on properties of single neurons, which is of course important in its own right (Kawaguchi & Kubota, 1997; Cardin et al., 2007; Nowak et al., 2008). But neurons interact with one another, producing new emergent phenomena (Sect 2.2), and this collective behavior likely has greater influence on the dynamics downstream than individual neurons do (though there are exceptions). Indeed, the strong interest in correlations between pairs of neurons, both experimental and theoretical (Roe & Ts'o, 1999; Nowak et al., 1999; de la Rocha et al., 2007; Ostojic et al., 2009; Renart et al., 2010; Jia et al., 2013; Zandvakili & Kohn, 2015) i.e., the interest in how the behaviors of different neurons are related, is itself recognition from the neuroscience community of the importance of collective behavior. In this paper, we took this one step further, to promote the study of correlations between outputs of two populations. Under very mild assumptions, we were able to derive quantitative relations between correlations and sample size (Sect 4.2), a result that we hope will be relevant as improving technology enables experimentalists to simultaneously capture the spiking activity of larger and larger samples of local populations.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10827-021-00779-4>.

Acknowledgements We would like to thank Wolf Singer for valuable discussions, and Robert Shapley for providing helpful comments on the manuscript. We are grateful also to the reviewers, whose critical comments contributed to improvements in the paper.

Declarations

Conflicts of interest The authors declare no conflict of interest.

References

- Baddeley, R., Abbott, L. F., Booth, M. C. A., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London B*, pages 1775–1783.
- Bastos, A. M., Vezoli, J., & Fries, P. (2015). Communication through coherence with inter-areal delays. *Current Opinion in Neurobiology*, 31, 173–180.
- Battaglia, D., Witt, A., Wolf, F., and Geisel, T. (2012). Dynamic effective connectivity of inter-areal brain circuits. *PLOS Computational Biology*, 8.
- Binzegger, T., Douglas, R. J., & Martin, K. (2009). Topology and dynamics of the canonical circuit of cat v1. *Neural Networks*, 22, 1071–1078.
- Börgers, C. (2017). *An introduction to modeling neuronal dynamics*, Chap 30. Springer.
- Börgers, C., Epstein, S., & Kopell, N. J. (2005). Background gamma rhythmicity and attention in cortical local circuits: A computational study. *Proceedings of the National Academy of Sciences of the United States of America*, 102, 7002.
- Börgers, C., & Kopell, N. J. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Computation*, 15, 509–538.
- Börgers, C., & Kopell, N. J. (2008). Gamma oscillations and stimulus selection. *Neural Computation*, 20, 383–414.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8, 183–208.
- Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low ring rates. *Neural Computation*, 11, 1321–1671.
- Brunel, N., & Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? i. synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90, 415–430.
- Burns, S. P., Xing, D., & Shapley, R. M. (2011). Is gamma-band activity in the local field potential of v1 cortex a “clock” or filtered noise? *Journal of Neuroscience*, 31, 9658–9664.
- Buzsáki, G. (2011). *Rhythms of the brain*. Oxford University Press.
- Cardin, J. A. (2016). Snapshots of the brain in action: Local circuit operations through the lens of γ oscillations. *Journal of Neuroscience*, 36, 10496–10504.
- Cardin, J. A., Palmer, L. A., & Contreras, D. (2007). Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *Journal of Neuroscience*, 27, 10333–10344.
- Chariker, L., Shapley, R., & Young, L.-S. (2016). Orientation selectivity from very sparse lgn inputs in a comprehensive model of macaque v1 cortex. *The Journal of Neuroscience*, 36(49), 12368–12384.
- Chariker, L., Shapley, R., & Young, L.-S. (2018). Rhythm and synchrony in a cortical network model. *The Journal of Neuroscience*, 38(40), 8621–8634.
- Chariker, L., Shapley, R., & Young, L.-S. (2020). Contrast response in a comprehensive network model of macaque V1. *Journal of Vision*, 20(4):16. <https://doi.org/10.1167/jov.20.4.16>
- Chariker, L., & Young, L.-S. (2015). Emergent spike patterns in neuronal populations. *Journal of Computational Neuroscience*, 38, 203–220.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience*. MIT Press black.
- de la Rocha, J., Doiron, B., Shea-Brown, E., Josić, K., & Reyes, A. D. (2007). Correlation between neural spike trains increases with firing rate. *Nature*, 448, 802–806.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *TRENDS in Cognitive Sciences*, 9, 10496–10504.
- Fries, P. (2015). Rhythms for cognition: Communication through coherence. *Neuron*, 88, 220–235.
- Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291, 1560–1563.
- Gielen, S., Krupa, M., & Zeitler, M. (2010). Gamma oscillations as a mechanism for selective information transmission. *Biological Cybernetics*, 103, 151–165.
- Gonzalez-Burgos, G., Cho, R. Y., & Lewis, D. A. (2015). Alterations in cortical network oscillations and parvalbumin neurons in schizophrenia. *Biological Psychiatry*, 77, 1031–1040.
- Gonzalez-Burgos, G., Hashimoto, T., & Lewis, D. A. (2010). Alterations of cortical gaba neurons and network oscillations in schizophrenia. *Current Psychiatry Reports*, 12, 335–344.
- Gray, C. M., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 86, 1698–1702.
- Grün, S., & Rotter, S. (Eds.). (2010). *Analysis of Parallel Spike Trains*. Springer.
- Henrie, J. A., & Shapley, R. (2005). LFP power spectra in V1 cortex: The graded effect of stimulus contrast. *Journal of Neurophysiology*, 94, 479–490.
- Holmgren, C., Harkany, T., Svennenfors, B., & Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *The Journal of Physiology*, 551(1), 139–153.
- Jia, X., Smith, M. A., & Kohn, A. (2011). Stimulus selectivity and spatial coherence of gamma components of the local field potential. *The Journal of Neuroscience*, 31, 9390–9403.
- Jia, X., Tanabe, S., & Kohn, A. (2013). Gamma and the coordination of spiking activity in early visual cortex. *Neuron*, 77, 762–774.
- M. R. Joglekar, L. Chariker, R. Shapley, and L.-S. Young. *PLOS Computational Biology*, 15, 2019.
- Kawaguchi, Y., & Kubota, Y. (1997). Gabaergic cell subtypes and their synaptic connections in rat frontal cortex. *Cerebral Cortex*, 7, 476–486.
- Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press.
- McCarthy, M. N., Ching, S. N., Whittington, M. A., & Kopell, N. (2012). Dynamical changes in neurological diseases and anesthesia. *Current Opinion in Neurobiology*, 22, 693–703.
- McLaughlin, D., Shapley, R., Shelley, M., & Wielaard, D. J. (2000). A neuronal network model of sharpening and dynamics of orientation tuning in an input layer of macaque primary visual

- cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 8087–8092.
- Nowak, L. G., Munk, M. H., James, A. C., Girard, P., & Bullier, J. (1999). Cross-correlation study of the temporal interactions between areas v1 and v2 of the macaque monkey. *Journal of Neurophysiology*, 81, 1057–1074.
- Nowak, L. G., Sanchez-Vives, M. V., & McCormick, D. A. (2008). Lack of orientation and direction selectivity in a subgroup of fast-spiking inhibitory interneurons: Cellular and synaptic mechanisms and comparison with other electrophysiological cell types. *Cerebral Cortex*, 18(5), 1058–1078.
- Okun, M., & Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature Neuroscience*, 11, 535–537.
- Ostojic, S. (2011). Interspike interval distributions of spiking neurons driven by fluctuating inputs. *Journal of Neurophysiology*, pages 361–373.
- Ostojic, S., Brunel, N., & Hakim, V. (2009). How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. *Journal of Neuroscience*, 29, 10234–10253.
- Oswald, A.-M. M., & Reyes, A. D. (2011). Development of inhibitory timescales in auditory cortex. *Cerebral Cortex*, 21, 1351–1361.
- Palmigiano, A., Geisel, T., Wolf, F., & Battaglia, D. (2017). Flexible information routing by transient synchrony. *Nature Neuroscience*, 20(7), 1014–1022.
- Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., & Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience*, 5, 805–811.
- Rangan, A. V., & Young, L.-S. (2013a). Dynamics of spiking neurons: between homogeneity and synchrony. *Journal of Computational Neuroscience*, 34, 433–460.
- Rangan, A. V., & Young, L.-S. (2013). Emergent dynamics in a model of visual cortex. *Journal of Computational Neuroscience*, 35, 155–167.
- Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., et al. (2010). The asynchronous state in cortical circuits. *Science*, 327, 587–590.
- Rodriguez, E., George, N., Lachaux, J.-P., Martinerie, J., Renault, B., & Varela, F. J. (1999). Perception's shadow: long-distance synchronization of human brain activity. *Nature*, 391, 430–433.
- Roe, A. W., & Ts'o, D. Y. (1999). Specificity of color connectivity between primate v1 and v2. *Journal of Neurophysiology*, 82, 2719–2730.
- Sederberg, P., Kahana, M., Howard, M., Donner, E., & Madsen, J. (2003). Theta and gamma oscillations during encoding predict subsequent recall. *Journal of Neuroscience*, 23, 10809–10814.
- Sincich, L. C., & Horton, J. C. (2005). The circuitry of V1 and V2: integration of color, form, and motion. *Annual Review of Neuroscience*, 28, 303–326.
- Sohal, V. S., Zhang, F., Yizhar, O., & Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459, 698–702.
- Steinmetz, N. A., Koch, C., Harris, K. D., & Carandini, M. (2018). Challenges and opportunities for large-scale electrophysiology with neuropixels probes. *Current Opinion in Neurobiology*, 50, 92–100.
- ter Wal, M., & Tiesinga, P. H. (2017). Phase difference between model cortical areas determines level of information transfer. *Frontiers in Computational Neuroscience*, 11.
- Tiesinga, P. H., Fellous, J.-M., José, J. V., & Sejnowski, T. J. (2001). Computational model of carbachol-induced delta, theta, and gamma oscillations in the hippocampus. *Hippocampus*, 11, 251–274.
- Uhlhaas, P. J., & Singer, W. (2010). Abnormal neural oscillations and synchrony in schizophrenia. *Nature Reviews Neuroscience*, 11, 100–113.
- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10, 1321–1371.
- Vogels, T. P., & Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of Neuroscience*, 25, 10786–10795.
- Whittington, M. A., Traub, R. D., Kopell, N., Ermentrout, B., & Buhl, E. H. (2000). Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *International Journal of Psychophysiology*, 38, 315–336.
- Womelsdorf, T., Schoffelen, J.-M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., et al. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 358(1), 1609–1612.
- Xing, D., Shen, Y., Burns, S., Yeh, C. I., Shapley, R., & Li, W. (2012). Stochastic generation of gamma-band activity in primary visual cortex of awake and anesthetized monkeys. *Journal of Neuroscience*, 32, 13873–13880.
- Zandvakili, A., & Kohn, A. (2015). Coordinated neuronal activity enhances corticocortical communication. *Neuron*, 87, 827–839.