Journal of Theoretical Biology 328 (2013) 26–32

Contents lists available at SciVerse ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

Encoding with synchrony: Phase-delayed inhibition allows for reliable and specific stimulus detection

Badal Joshi [a], Mainak Patel [b,*]

[a] School of Mathematics, University of Minnesota, 127 Vincent Hall, 206 Church St. SE, Minneapolis, MN 55455, USA
[b] Mathematics Department, Duke University, Box 90320, Durham, NC 27708-0320, USA

AUTHOR-HIGHLIGHTS

- High threshold or phase-delayed inhibition can both decode synchronized oscillations.
- High threshold and phase-delayed inhibition detect absolute vs. relative synchrony.
- No noise: high threshold decoder and phase-delayed inhibition perform equally well.
- With noise: phase-delayed inhibition performs better than a high threshold.

ARTICLE INFO

Article history:
Received 28 January 2013
Accepted 12 March 2013
Available online 22 March 2013

Keywords:
Neuronal networks
Feed-forward inhibition
Synchronized oscillations
Synchrony detection

ABSTRACT

Synchronized oscillations are observed in a diverse array of neuronal systems, suggesting that synchrony represents a common mechanism used by the brain to encode and relay information. Coherent population activity can be deciphered by a decoder neuron with a high spike threshold or by a decoder using phase-delayed inhibition. These two mechanisms are fundamentally different – a high spike threshold detects a minimum number of synchronous input spikes (absolute synchrony), while phase-delayed inhibition requires a fixed fraction of incoming spikes to be synchronous (relative synchrony). We show that, in a system with noisy encoders where stimuli are encoded through synchrony, phase-delayed inhibition enables the creation of a decoder that can respond both reliably and specifically to a stimulus, while a high spike threshold does not.

1. Introduction

Coherent oscillations are an ubiquitous encoding feature of the brain, and have been observed in numerous neuronal systems within a broad assortment of animal species, e.g., see Eckhorn (1994); Friedrich et al. (2004); Gray (1994); Laurent and Davidowitz (1994); Murthy and Fetz (1992); Sridharan et al. (2011). The importance of synchrony in information processing tasks performed by the brain implies the existence of an appropriate neuronal decoding mechanism. If a set of encoders is representing information through synchronized oscillations, the simplest way to design a decoder (or read-out cell) is by employing a neuron with a spike threshold that is high relative to the strength of encoder inputs. With a high spike threshold, multiple encoders must fire simultaneously in order for sufficient summation to occur to drive the decoder's membrane potential above threshold. The read-out cell will therefore spike only if such a synchronous firing event transpires.

A second way to design a synchrony detector is through the use of phase-delayed inhibition (Fig. 1). In this scheme, the encoders provide direct excitation to the decoder, but en-route encoder axons collaterally innervate a group of inhibitory interneurons, which in turn provide powerful inhibition to the read-out cell. Due to delayed information flow through the inhibitory pathway, each excitatory input to the decoder is followed, with a characteristic temporal lag, by an inhibitory input. If the encoders spike haphazardly during an oscillation cycle, then the inhibitory interneurons follow accordingly, and the read-out cell remains shrouded in inhibition and unable to spike. If, on the other hand, the encoders fire synchronously during an oscillation cycle, then the decoder will receive concurrent excitation followed by pooled inhibition, allowing the read-out cell to respond within the window of excitation. Thus, the read-out neuron will fire only if encoder inputs are sufficiently synchronous.

Given that a decoder with a high spike threshold is architecturally easier to implement than one utilizing phase-delayed inhibition, one might expect to observe high threshold decoders in neuronal

* Corresponding author. Tel.: +1 919 660 2800; fax: +1 919 660 2821.
E-mail addresses: mainak@math.duke.edu, mainak_patel@hotmail.com (M. Patel).

0022-5193/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jtbi.2013.03.012

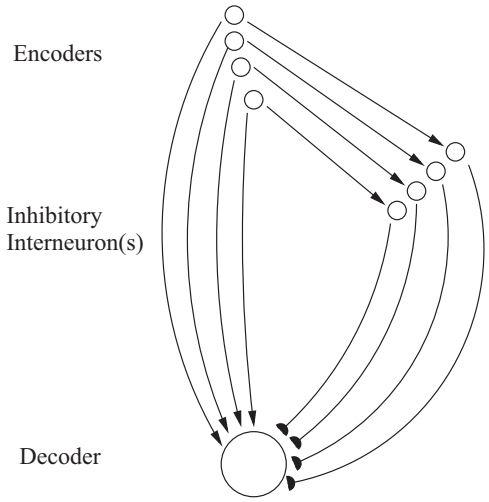

Fig. 1. Schematic of network architecture needed to implement phase-delayed inhibition. A set of excitatory encoder neurons innervate a set of inhibitory interneurons as well as sending convergent input to a decoder neuron. The decoder neuron also receives input from the inhibitory interneurons. Each excitatory encoder spike is followed, with a slight delay, by an inhibitory interneuron spike.

systems that require synchrony detection. However, in reality many systems employ phase-delayed inhibition, e.g., see Perez-Orive et al. (2004); Sridharan et al. (2011); Pouille and Scanziani (2001); Wehr and Zador (2003); Mittmann et al. (2005); Blitz and Regehr (2005). This suggests that phase-delayed inhibition may provide an inherent advantage over simply using a high spike threshold.

In upcoming work (Patel and Joshi, currently submitted), we showed that, while both phase-delayed inhibition and high spike threshold can create a read-out cell that behaves as a sharp synchrony filter (i.e., the response of the decoder jumps and rapidly saturates once the inputs exceed a critical level of synchrony), phase-delayed inhibition does so in a much more robust manner. With phase-delayed inhibition, the synchrony threshold (the critical level of synchrony that inputs must surpass in order to induce the decoder to spike), is impervious to noise – the value of the synchrony threshold remains relatively unchanged even when faced with considerable fluctuations in the number of input spikes. The synchrony threshold created using a high spike threshold decoder, on the other hand, exhibits no such robustness and is exquisitely sensitive to noise.

This difference in robustness arises because high spike threshold and phase-delayed inhibition detect synchrony in fundamentally different ways. During an oscillation cycle, a high spike threshold decoder requires a minimum number of concurrently arriving spikes in order to respond, and hence is sensitive to the total number of incoming spikes (absolute synchrony). A phase-delayed inhibition decoder exhibits little dependence on the total number of incoming spikes, but rather requires a minimum *fraction* of the incoming spikes during an oscillation cycle to be synchronous in order to respond (relative synchrony). This essential difference is portrayed in Fig. 2. In this work, we extend our previous results to show that if stimuli are encoded through synchronized oscillations by a population of noisy encoders, then phase-delayed inhibition can create a decoder that responds both reliably and specifically to a particular stimulus, while a high spike threshold cannot.

2. Results

Our approach is two-fold: we perform idealized analytical calculations to illustrate the fundamental distinction between phase-delayed inhibition and a high spike threshold, and we perform computational experiments using an integrate-and-fire decoder neuron in order to provide a more biophysically realistic setting. Our purpose is to show that, if stimuli are encoded with synchrony in a system of noisy encoders, a decoder that is both reliable and stimulus-specific can be constructed by phase-delayed inhibition but not by a high spike threshold. We therefore simulate two different stimuli by modulating the synchrony, rather than number, of encoder spikes, and we attempt to construct a decoder that responds in a reliable and specific manner to one of the two stimuli (i.e., if we denote the two stimuli by stimulus 1 and stimulus 2, then we attempt to design a decoder that responds during trials of stimulus 1 but does not respond during trials of stimulus 2).

2.1. Model construction

We fix a period of 50 ms within which to distribute encoder and inhibitory interneuron spikes, though similar results are obtained with all choices of period tested (so long as the time course of inhibition h is scaled appropriately with the period). We denote the total number of encoder spikes by N, where N is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation $\sigma_N = 25$. Stimulus 1 is represented by modulating the synchrony of a fraction of the total number of encoder spikes, with this fraction drawn from a Gaussian distribution with mean $\mu_1 = 0.55$ and standard deviation $\sigma_1 = 0.05$. The fraction corresponding to stimulus 2 is drawn from a Gaussian distribution with mean $\mu_2 = 0.275$ and standard deviation $\sigma_2 = 0.05$. The phase of stimulus-induced spikes is drawn from a Gaussian distribution with mean 0 and standard deviation $\sigma_{stim} = 3$ ms, while the phase of the remaining (noisy) spikes is drawn from a Gaussian distribution with mean 0 and standard deviation $\sigma_{noise} = 12$ ms. Unless otherwise indicated, these standard values are used in all figures.

For the analytical calculations, we idealize the distinction between high spike threshold and phase-delayed inhibition by appealing to the property that high spike threshold detects absolute synchrony, while phase-delayed inhibition detects relative synchrony. For the high threshold case, we set a threshold (in minimum number of synchronous encoder spikes) of f and we calculate the probability of decoder response as $\mathbf{P}\{N_w > f\}$, where N_w is the number of the N encoder spikes that fall within a $w = 3$ ms window centered at 0 ms. For the phase-delayed inhibition case, we set a synchrony threshold s ($0 \le s \le 1$) and we calculate the probability of decoder response as $\mathbf{P}\{N_w/N > s\}$ (see Methods for mathematical details).

For the computational studies, we simulate the decoder as an integrate-and-fire neuron with conductance-based excitatory and inhibitory synaptic inputs (see Methods for model details). If an encoder spike occurs at time τ, an excitatory input to the decoder is described as a square pulse beginning at time τ with amplitude A_e and duration c ms. The corresponding inhibitory interneuron spike is described as a square pulse beginning at time $\tau + d$ with amplitude A_i and duration h ms. We fix the synaptic delay at $d = 3$ ms; we fix $c = 3$ ms and $h = 5$ ms in accordance with the approximate time course of fast excitatory and inhibitory synapses within the brain, while the amplitudes A_e and A_i are varied during our simulations. The parameters are set such that the decoder spikes either 0 or 1 times during any particular 50 ms trial; the probability of decoder response is computed as T_s/T, where $T = 5000$ is the total number of trials and T_s is the number of trials during which the decoder spikes. The high threshold case is simulated by setting A_e to be relatively small with $A_i = 0$ (no inhibition with weak encoder inputs relative to the spike threshold), while the phase-delayed inhibition case is simulated by setting A_e and A_i at relatively large values.

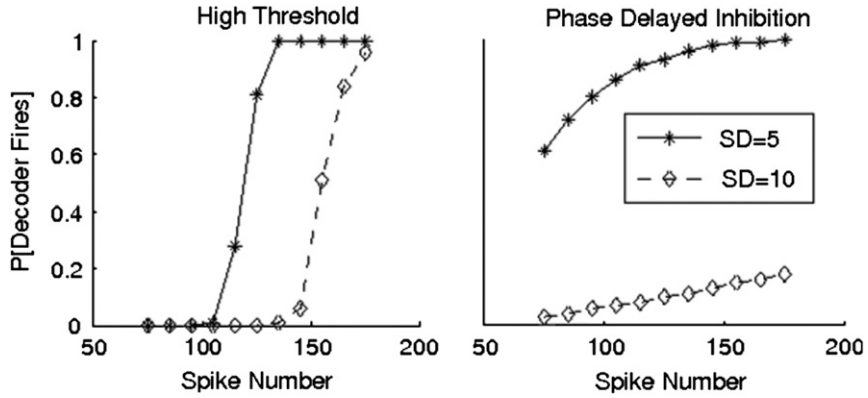

Fig. 2. Phase-delayed inhibition detects relative synchrony (i.e., the decoder responds if the fraction of incoming spikes that are synchronous exceeds a particular value), while a high spike threshold detects absolute synchrony (i.e., the decoder responds if a minimum number of the incoming spikes arrive synchronously). Both plots show the probability that an integrate-and-fire decoder spikes as a function of the total number of incoming encoder spikes (distributed throughout a 50 ms period); encoder spike phases have a standard deviation of 5 ms or 10 ms (mean spike phase is 0 ms, where the 50 ms window stretches from −25 ms to 25 ms). *Left*: in the high threshold case, there is no inhibition and the amplitude of excitatory inputs is set at a small value (0.001). The probability that the decoder responds jumps to high values once the total number of input spikes exceeds a certain threshold value. *Right*: in the phase-delayed inhibition case, the amplitude of excitatory encoder inputs is set at a large value (0.01), with each excitatory input followed, with a 3 ms delay, by a large amplitude (0.03) inhibitory input. The response of the decoder exhibits little dependence on the total number of input spikes; the decoder has a high response probability for a 5 ms spike phase standard deviation (highly synchronous inputs), and a low response probability for a 10 ms spike phase standard deviation (less synchronous inputs).

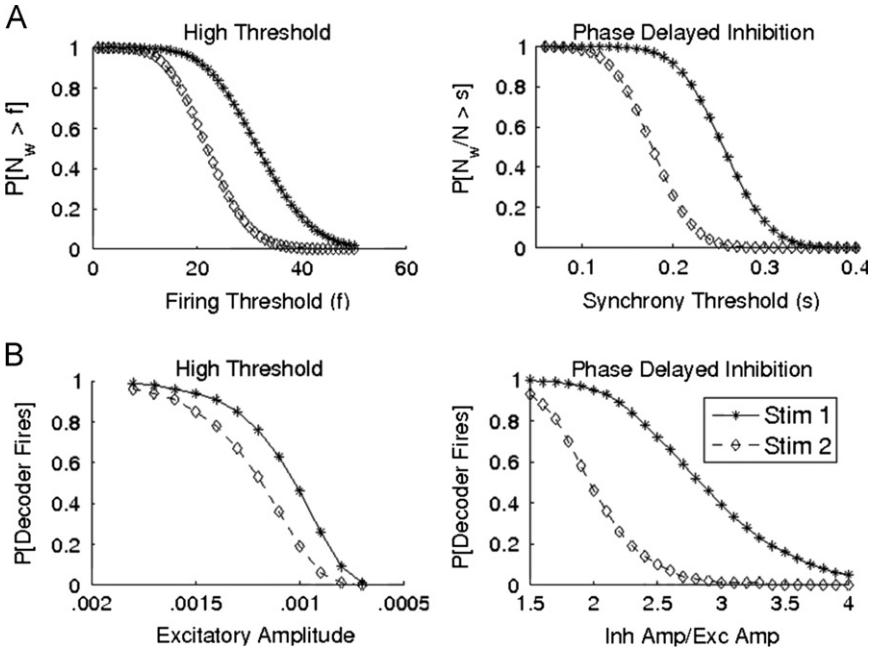

Fig. 3. The probability of decoder response for stimulus 1 and stimulus 2 is plotted in the high threshold case and the phase-delayed inhibition case, both for the idealized mathematical model and the biophysical integrate-and-fire decoder model (see text). (A) For the high threshold case, the parameter f represents the spike threshold of the decoder (i.e., the number encoder spikes that must arrive concurrently in order for the decoder to respond). For the phase-delayed inhibition case, the parameter s represents the synchrony threshold of the decoder (i.e., the fraction of encoder spikes that must arrive synchronously in order for the decoder to respond). (B) For the high threshold case, the amplitude of encoder inputs (A_e) is shown on the abscissa (since the integrate-and-fire decoder has a fixed spike threshold, lowering the strength of excitatory inputs is equivalent to raising the spike threshold). For the phase-delayed inhibition case, the amplitude of encoder inputs is fixed at $A_e = 0.01$, and the strength of inhibitory interneuron inputs A_i is varied, with the ratio A_e/A_i shown on the abscissa.

2.2. Phase-delayed inhibition vs. high spike threshold

In Fig. 3(A), we use the idealized analytical calculation to plot the probability of decoder response as a function of response threshold for both the high threshold and phase-delayed inhibition scenarios. In Fig. 3(B), similar plots are shown for the biophysical integrate-and-fire decoder. The response threshold is varied in the high threshold scenario by lowering the amplitude A_e of encoder inputs (since the spike threshold of the integrate-and-fire cell is fixed, lowering A_e effectively simulates a higher spike threshold relative to the strength of encoder inputs). In the phase-delayed inhibition case, A_e is fixed and the response threshold is varied by increasing the strength A_i of inhibitory inputs (increasing the amplitude of inhibition raises the level of relative synchrony required among the encoder inputs to induce the decoder to fire; this is explored in currently submitted work by Patel and Joshi). As expected, for both stimulus 1 and stimulus 2 the probability of decoder response drops from 1 to 0 as the response threshold is raised. However, since our interest is in creating a decoder that responds reliably to stimulus 1 without responding to stimulus 2, the important quantity is the difference in the probability of response to stimulus 1 and the probability of

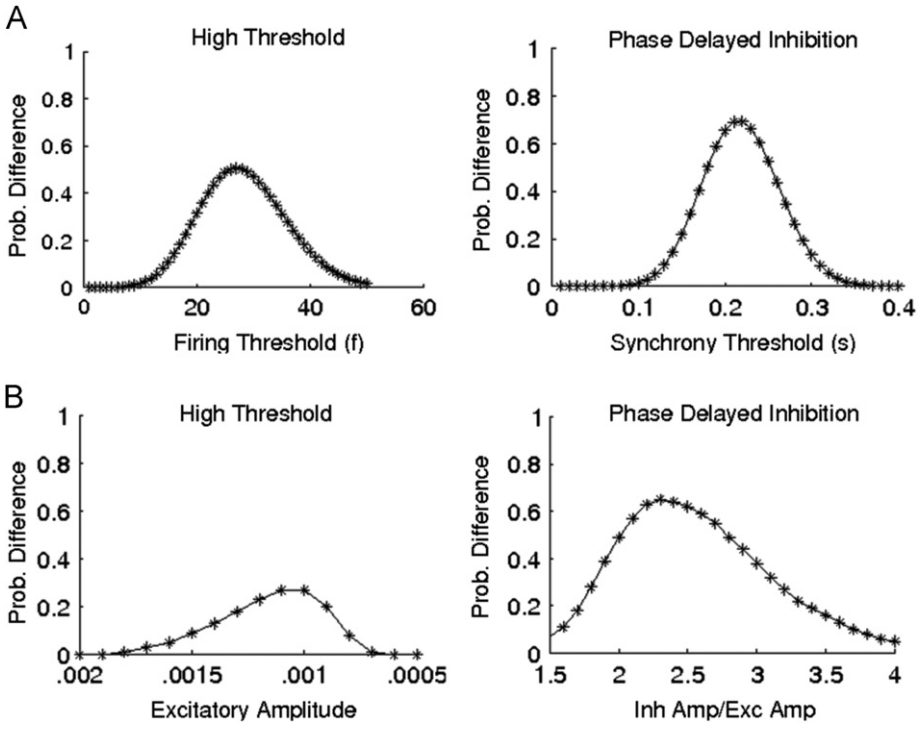

Fig. 4. Difference curves are shown for the plots in Fig. 3. The ordinate shows the probability that the decoder responds to stimulus 1 minus the probability that the decoder responds to stimulus 2, both for the idealized mathematical model and the biophysical integrate-and-fire decoder model (see text). (A) For the high threshold case, the parameter f represents the spike threshold of the decoder (i.e., the number encoder spikes that must arrive concurrently in order for the decoder to respond). For the phase-delayed inhibition case, the parameter s represents the synchrony threshold of the decoder (i.e., the fraction of encoder spikes that must arrive synchronously in order for the decoder to respond). (B) For the high threshold case, the amplitude of encoder inputs (A_e) is shown on the abscissa (since the integrate-and-fire decoder has a fixed spike threshold, lowering the strength of excitatory inputs is equivalent to raising the spike threshold). For the phase-delayed inhibition case, the amplitude of encoder inputs is fixed at $A_e = 0.01$, and the strength of inhibitory interneuron inputs A_i is varied, with the ratio A_e/A_i shown on the abscissa.

response to stimulus 2; the difference curves corresponding to the plots in Fig. 3 are shown in Fig. 4. For both the analytical calculation and particularly the computational results, a greater probability difference is observed in the phase-delayed inhibition scenario than in the high threshold case, suggesting that phase-delayed inhibition can create a decoder that responds both reliably and specifically to stimulus 1, while a high spike threshold can create a decoder that is *either* reliable or specific (but not both). Furthermore, Fig. 4(B) shows that in the high threshold case the probability difference is delicately sensitive to the amplitude of encoder inputs, implying that A_e needs to be finely tuned (within the range ~0.0012–0.0008) in order to achieve even a modest probability difference. In the phase-delayed inhibition scenario, a high probability difference can be generated in a robust manner – high probability differences are seen when the amplitude of inhibition A_i is ~2–3 times the strength of excitation A_e.

In Figs. 3 and 4, we employed the standard parameters described above, in which the total number of encoder spikes is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation $\sigma_N = 25$. The phase of stimulus-induced spikes has standard deviation $\sigma_{stim} = 3$ ms, while the phase of the remaining (noisy) spikes has standard deviation $\sigma_{noise} = 12$ ms. In Fig. 5, we keep these same parameters, except on the abscissa we vary σ_{noise} (i.e., we vary the synchrony difference between stimulus-induced spikes and noisy spikes), and on the ordinate we plot max{**P**[decoder responds to stimulus 1]–**P**[decoder responds to stimulus 2]} (i.e., the peaks of the curves in Fig. 4). Fig. 5 shows that as the difference in synchrony between stimulus-induced and noisy spikes increases, a high spike threshold and a phase-delayed inhibition decoder both become more reliable and specific; however, the phase-delayed inhibition decoder consistently performs significantly better than the high threshold decoder. Moreover, the phase-delayed inhibition decoder can both reliably and specifically respond to stimulus 1 even for relatively small synchrony differences between the noisy and stimulus-induced spikes (in the right panel of Fig. 5, the maximum probability difference for phase-delayed inhibition approaches ~0.5 for $\sigma_{noise} = 6$ ms).

The reason that phase-delayed inhibition performs strikingly better than a high spike threshold is because our system is noisy, in the sense that total spike number is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation $\sigma_N = 25$. If the total spike number were fixed, then it would presumably be easier to set the spike threshold of the decoder at a value that ensures both reliable and specific responses to stimulus 1, and hence the advantage of using phase-delayed inhibition would be lost. In order to assess the contribution of noise in spike number to the advantage of using phase-delayed inhibition over a high spike threshold, we keep the standard parameters described above, except we vary the noisiness of the system. Fig. 6 plots max{**P**[decoder responds to stimulus 1]–**P**[decoder responds to stimulus 2]} as a function of the noisiness of the system (σ_N). Fig. 6 shows that when there is no noise in the system ($\sigma_N = 0$), high spike threshold and phase-delayed perform comparably (both decoder schemes can create reliable and specific responses to stimulus 1), but as an increasing amount of noise is introduced into the system, the performance of the high spike threshold decoder rapidly declines while the phase-delayed inhibition decoder continues to perform at a high level.

3. Discussion

In this work, we have shown that in a system of noisy encoders where stimuli are encoded by synchrony, phase-delayed inhibition enables construction of a read-out neuron that can respond in

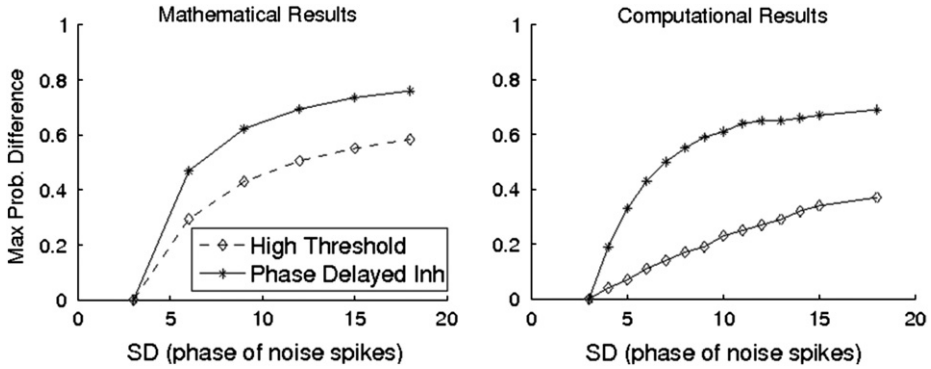

Fig. 5. max{**P**[decoder responds to stim 1]–**P**[decoder responds to stim 2]} (i.e., the peaks of the curves in Fig. 4) are shown on the ordinate. The total number of encoder spikes is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation $\sigma_N = 25$. The standard deviation in the phase of stimulus-induced spikes is fixed at $\sigma_{stim} = 3$ ms, while the standard deviation in the phase of the remaining (noisy) spikes (σ_{noise}) is varied on the abscissa. The fraction of total spikes that are stimulus-induced is drawn from a Gaussian with mean $\mu_1 = 0.55$ and standard deviation $\sigma_1 = 0.05$ for stimulus 1 or mean $\mu_2 = 0.55$ and standard deviation $\sigma_2 = 0.025$ for stimulus 2. Results are shown for the high threshold case and the phase-delayed inhibition case, both for the idealized mathematical model and the biophysical integrate-and-fire decoder model (see text).

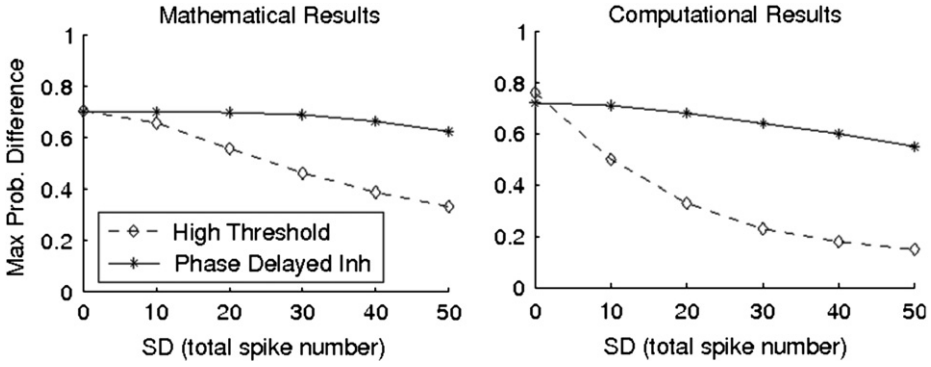

Fig. 6. max {**P**[decoder responds to stim 1]–**P**[decoder responds to stim 2]} (i.e., the peaks of the curves in Fig. 4) are shown on the ordinate. The total number of encoder spikes is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation σ_N, with σ_N varied on the abscissa. The standard deviation in the phase of stimulus-induced spikes is fixed at $\sigma_{stim} = 3$ ms, while the standard deviation in the phase of the remaining (noisy) spikes is fixed at $\sigma_{noise} = 12$ ms. The fraction of total spikes that are stimulus-induced is drawn from a Gaussian with mean $\mu_1 = 0.55$ and standard deviation $\sigma_1 = 0.05$ for stimulus 1 or mean $\mu_2 = 0.55$ and standard deviation $\sigma_2 = 0.025$ for stimulus 2. Results are shown for the high threshold case and the phase-delayed inhibition case, both for the idealized mathematical model and the biophysical integrate-and-fire decoder model (see text).

both a reliable and specific manner to a stimulus, while a high spike threshold decoder can respond in *either* a reliable *or* specific manner, but cannot exhibit both properties. The situation addressed in our model may arise when there is a system of n noisy encoders, with different stimuli modulating the synchrony of different (but potentially overlapping) subsets of the n encoders. For example, stimulus 1 may cause a subset of k encoders to fire in a synchronous manner, while stimulus 2 may cause coherent spiking of a different (but possibly overlapping) subset of k encoders. In such a system, in order to construct a decoder whose activity represents the presence or absence of stimulus 1, the decoder would read from a subset of encoders that includes the k encoders corresponding to stimulus 1. Thus, if stimulus 1 is presented, a large fraction of the encoders sending input to the stimulus 1 decoder will synchronize, while if stimulus 2 is presented, a smaller fraction of the encoders sending input to the stimulus 1 decoder will synchronize. Such a situation is observed in the coding of odor stimuli by the insect antennal lobe, in particular within the locust and moth (Christensen et al., 2000; Lei et al., 2002; Mazor and Laurent, 2005). In fact, within the locust, experimental and theoretical evidence has been provided suggesting that decoders of globally synchronized, 20 Hz antennal lobe activity employ phase-delayed inhibition (Assisi et al., 2007; Perez-Orive et al., 2002), though a recent investigation has disputed the existence of phase-delayed inhibition within this system (Gupta and Stopfer, 2012).

The phenomenon of phase-delayed inhibition has been observed in many systems. Gamma oscillations have been observed in the optic tectum of the barn owl, and outgoing tectal fibers to the nucleus rotundus of the thalamus have been shown to emanate collaterals to the intermediary GABAergic nuclei, which in turn provide phase-delayed inhibition to rotundal cells (Benowitz and Karten, 1976; Deng and Rogers, 1998; Sridharan et al., 2011). In several systems, phase-delayed inhibition has been observed as a means to decode synchrony in the absence of synchronized oscillations. One example is provided by Schaeffer collaterals, which have been shown to directly excite hippocampal pyramidal cells, as well as providing indirect phase-delayed inhibition to these targets with a ~2 ms time lag (Fricker and Miles, 2000; Pouille and Scanziani, 2001). Other examples include auditory cortical cells (Wehr and Zador, 2003), cerebellar Purkinje cells (Mittmann et al., 2005), and neurons within the lateral geniculate nucleus of the thalamus (Blitz and Regehr, 2005). The results of our investigation apply to these systems as well (to apply our analysis to these systems, we need only neglect residual dynamics from the previous oscillation cycle, and we find that doing so yields similar results). Our results may be able to furnish insight into these systems as well as others employing phase-delayed inhibition as a means to detect synchronized input.

4. Methods

We performed both idealized analytical calculations and more realistic computational experiments during the course of our investigations.

4.1. Analytical calculations

We assume that the total number of encoder spikes N is drawn from a Gaussian distribution with mean $\mu_N = 125$ and standard deviation $\sigma_N = 25$, in other words $N \sim \mathcal{N}(\mu_N, \sigma_N)$. The fraction F of stimulus-induced spikes is drawn from a Gaussian distribution with mean $\mu_1 = 0.55$ for stimulus 1 and $\mu_2 = 0.275$ for stimulus 2, while the standard deviations are given by $\sigma_1 = \sigma_2 = 0.05$. The phase of all spikes has a mean value of 0. In order to model the higher synchrony of the stimulus-induced spikes compared to the noisy spikes, we set the standard deviation of the phases of the stimulus-induced spikes to be less than the corresponding standard deviation of the noisy spikes.

We denote the number of stimulus-induced spikes by $N_s = FN$ and the number of noisy spikes by $N_n = (1-F)N$. Let the phases of the stimulus-induced spikes be represented by $\phi_1, \phi_2, \ldots, \phi_{N_s}$ and the phases of the noisy spikes by $\theta_1, \theta_2, \ldots, \theta_{N_n}$.

For the high threshold case we are interested in determining the number of spikes that occur in a window of size w centered around 0, while for the phase-delayed inhibition case we are interested in determining the fraction of spikes in the same window. We define N_w to be the number of spikes in a window of size w around 0. Let $I_{\phi_i \in (-w/2, w/2)}$ be the indicator of the event that the ith stimulus-induced spike has a phase less than $w/2$ in absolute value. In other words, the function $I_{\phi_i \in (-w/2, w/2)}$ takes the value 1 if $\phi_i \in (-w/2, w/2)$ and takes the value 0 otherwise. Similarly, let $I_{\theta_j \in (-w/2, w/2)}$ be the indicator of the event that the jth noisy spike falls within a window of size w around 0. Using the indicators, we may represent N_w as

$$N_w = \sum_{i=1}^{N_s} I_{\phi_i \in (-w/2, w/2)} + \sum_{j=1}^{N_n} I_{\theta_j \in (-w/2, w/2)}$$

In order to calculate the probability $\mathbf{P}\{N_w > f\}$, where f is an absolute threshold for the number of spikes required for the read-out neuron to fire, we condition on all possible values of N and F and then use the standard normal approximation. In the following, we abbreviate $I_{\phi_i \in (-w/2, w/2)}$ and $I_{\theta_j \in (-w/2, w/2)}$ by I_{ϕ_i} and I_{θ_j}, respectively, and let $k_n \in \{0, 1, 2, \ldots, n\}$. We then have

$$\mathbf{P}\{N_w > f\} = \sum_n \sum_{k_n} \mathbf{P}\left(\sum_{i=1}^{k_n} I_{\phi_i} + \sum_{j=1}^{n-k_n} I_{\theta_j} > f \right) \mathbf{P}(n) \mathbf{P}(k_n),$$

where $\mathbf{P}(n) := \mathbf{P}(N = n)$ and $\mathbf{P}(k_n) := \mathbf{P}(F = k_n/n)$.

Letting $p := \mathbf{P}\{I_{\phi_i} \in (-w/2, w/2)\}$, the binomial distribution $\sum_{i=1}^{k_n} I_{\phi_i} =: \mathcal{B}(k_n, p)$ may be approximated by the normal distribution $\mathcal{N}(k_n p, k_n p(1-p))$. We use a similar standard normal approximation for the noisy spikes as well.

For the phase-delayed inhibition case, the read-out neuron spikes when the fraction of the total number of spikes within the window crosses a synchrony threshold s. In other words, we need to determine the probability $\mathbf{P}\{N_w/N > s\}$ as a function of s. We calculate this quantity by using a similar conditioning on the values of N and F as follows:

$$\mathbf{P}\{N_w/N > s\} = \sum_n \sum_{k_n} \mathbf{P}\left(\sum_{i=1}^{k_n} I_{\phi_i} + \sum_{j=1}^{n-k_n} I_{\theta_j} > sn \right) \mathbf{P}(n) \mathbf{P}(k_n).$$

We used the usual normal approximation to the binomial, similar to the high threshold case, to perform computation of the above probabilities.

4.2. Computational procedures

Our decoder was modeled as an integrate-and-fire neuron. The membrane potential was governed by the following equation:

$$\frac{dV}{dt} = -g_L(V - E_0) - g_{exc}(t)(V - E_{exc}) - g_{inh}(t)(V - E_{inh}),$$

where $V(t)$ is the membrane potential, $g_L = 0.05$ is the leak conductance, $E_0 = 0$ is the resting potential, $E_{exc} = 4.67$ is the reversal potential for synaptic excitation, and $E_{inh} = -0.67$ is the reversal potential for synaptic inhibition. This is a reduced dimensional model, with a nondimensional membrane potential, time in units of ms, and conductance in units of ms^{-1}. A spike was recorded when $V(t)$ reached a threshold value $E_{thresh} = 1$, with $V(t)$ being instantaneously reset to rest following a spike. An absolute refractory period of 2 ms was simulated by holding the membrane potential at rest for 2 ms following a spike. Details of the reduced dimensional model are given in Tao et al. (2004). Simulations were carried out using the explicit Euler method with a time step of 0.01 ms.

A period of 50 ms was simulated, though similar results were obtained for all periods (20–100 ms) tested, so long as the time course of synaptic inhibition h was scaled appropriately with the period. Data were gathered from the second oscillation cycle during a trial in order to account for the effects of residual decoder dynamics from the previous period. Encoder spike number and spike phases were drawn from Gaussian distributions as described in the "Results" section. Each encoder spike was followed, with a 3 ms synaptic delay, by an inhibitory interneuron spike. Encoder and inhibitory interneuron inputs were described by step functions. If an encoder spike occurred at time τ, then $g_{exc}(t)$ was incremented by a value A_e at time τ and subsequently decremented by A_e at time $\tau + c$. To describe the corresponding interneuron spike, $g_{inh}(t)$ was incremented by a value A_i at time $\tau + 3$ and subsequently decremented by A_i at time $\tau + 3 + h$. The time course of excitation was fixed at $c = 3$ ms, while the time course of inhibition was fixed at $h = 5$ ms. The amplitude of excitation (A_e) and inhibition (A_i) were varied during the course of our simulations.

References

Assisi, C., Stopfer, M., Laurent, G., Bazhenov, M., 2007. Adaptive regulation of sparseness by feed-forward inhibition. Nat. Neurosci. 10, 1176–1184.

Benowitz, L., Karten, H., 1976. Organization of the tectofugal visual pathway in the pigeon: a retrograde transport study. J. Comp. Neurol. 167, 503–520.

Blitz, D., Regehr, W., 2005. Timing and specificity of feed-forward inhibition within the LGN. Neuron 45, 917–928.

Christensen, T., Pawlowski, V., Lei, H., Hildebrand, J., 2000. Multi-unit recordings reveal context-dependent modulation of synchrony in odor-specific neural ensembles. Nat. Neurosci. 3, 927–931.

Deng, C., Rogers, L., 1998. Organisation of the tectorotundal and SP/IPS-rotundal projections in the chick. J. Comp. Neurol. 394, 171–185.

Eckhorn, R., 1994. Oscillatory and non-oscillatory synchronizations in the visual cortex and their possible roles in associations of visual features. Prog. Brain Res. 102, 405–426.

Fricker, D., Miles, R., 2000. EPSP amplification and the precision of spike timing in hippocampal neurons. Neuron 28, 559–569.

Friedrich, R., Habermann, C., Laurent, G., 2004. Multiplexing using synchrony in the zebrafish olfactory bulb. Nat. Neurosci. 7, 862–871.

Gray, C., 1994. Synchronous oscillations in neuronal systems: mechanisms and functions. J. Comput. Neurosci. 1, 11–38.

Gupta, N., Stopfer, M., 2012. Functional analysis of a higher olfactory center, the lateral horn. J. Neurosci. 32, 8138–8148.

Laurent, G., Davidowitz, H., 1994. Encoding of olfactory information with oscillating neural assemblies. Science 265, 1872–1875.

Lei, H., Christensen, T., Hildebrand, J., 2002. Local inhibition modulates odor-evoked synchronization of glomerulus-specific output neurons. Nat. Neurosci. 5, 557–565.

Mazor, O., Laurent, G., 2005. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron 48, 661–673.

Mittmann, W., Koch, U., Hausser, M., 2005. Feed-forward inhibition shapes the spike output of cerebellar purkinje cells. J. Physiol. 563, 369–378.

Murthy, V., Fetz, E., 1992. Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. Proc. Natl. Acad. Sci. U.S.A. 89, 5670–5674.

Perez-Orive, J., Bazhenov, M., Laurent, G., 2004. Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. J. Neurosci. 24, 6037–6047.

Perez-Orive, J., Mazor, O., Turner, G., Cassenaer, S., Wilson, R., Laurent, G., 2002. Oscillations and sparsening of odor representations in the mushroom body. Science 297, 359–365.
Pouille, F., Scanziani, M., 2001. Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. Science 293, 1159–1163.
Sridharan, D., Boahen, K., Knudsen, E., 2011. Space coding by gamma oscillations in the barn owl optic tectum. J. Neurophysiol. 105, 2005–2017.
Tao, L., Shelley, M., McLaughlin, D., Shapley, R., 2004. An egalitarian network model for the emergence of simple and complex cells in visual cortex. Proc. Natl. Acad. Sci. U.S.A. 101, 366–371.
Wehr, M., Zador, A., 2003. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. Nature 426, 442–446.