Journal of Neuroscience Methods 206 (2012) 54–64

Contents lists available at SciVerse ScienceDirect

Journal of Neuroscience Methods

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

Basic Neuroscience

Detecting synfire chains in parallel spike data

George L. Gerstein[a,*], Elizabeth R. Williams[b], Markus Diesmann[c,d], Sonja Grün[d,e,f], Chris Trengove[d]

[a] Department of Neuroscience, University of Pennsylvania, Philadelphia, PA, USA
[b] Institute of Neuroscience, Newcastle University, Newcastle upon Tyne, UK
[c] Institute of Neuroscience and Medicine (INM-6), Research Center Jülich, Jülich, Germany
[d] RIKEN Computational Science Research Program, Wako City, Saitama, Japan
[e] Theoretical Systems Neurobiology, RWTH University, Aachen, Germany
[f] RIKEN Brain Science Institute, Wako City, Saitama, Japan

ARTICLE INFO

Article history:
Received 22 September 2011
Received in revised form 1 February 2012
Accepted 3 February 2012

Keywords:
Multiple spike trains
Synfire chains
Detection and analysis algorithms
Repeating spike patterns
Active identity intersection matrix

ABSTRACT

The synfire chain model of brain organization has received much theoretical attention since its introduction (Abeles, 1982, 1991). However there has been no convincing experimental demonstration of synfire chains due partly to limitations of recording technology but also due to lack of appropriate analytic methods for large scale recordings of parallel spike trains. We have previously published one such method based on intersection of the neural populations active at two different times (Schrader et al., 2008). In the present paper we extend this analysis to deal with higher firing rates and noise levels, and develop two additional tools based on properties of repeating firing patterns. All three measures show characteristic signatures if synfire chains underlie the recorded data. However we demonstrate that the detection of repeating firing patterns alone (as used in several papers) is not enough to infer the presence of synfire chains. Positive results from all three measures are needed.

1. Introduction

Ever since Hebb's (1949) book there is general agreement that neurons do not act alone and that they join dynamically into functional assemblies both for computation and representation. Despite vast amounts of information about anatomy and connectivity in the real nervous system, the formation of such assemblies, their function and their interaction remain obscure. One interesting theoretical suggestion for the organization of assemblies is the "synfire chain" as proposed by Abeles (1982, 1991). In its original and simplest form the synfire chain of length l consists of l sub-assemblies or links of w neurons each (the synfire chain width), with full unidirectional excitatory connectivity between successive links. Such a chain structure can propagate near-synchronous activity in the first or input link to successive links like a row of dominoes. With suitable conditions the synchrony of firing in a link increases as the activity passes down the chain.

Theoretical studies have defined detailed properties of various versions of such systems. Stability properties and temporal behavior have been examined as well as structural variations such as partial link to link connectivity or feedback when a given neuron occurs in more than one group (Aertsen et al., 1996; Diesmann et al., 1999; Mehring et al., 2003; Aviel et al., 2003; Goedeke and Diesmann, 2008; Gewaltig et al., 2001; Tetzlaff et al., 2005, 2009; Trengove et al., 2010b, 2012). There also have been studies of partially interconnected or overlapping synfire chains to define the conditions in which activity can be transmitted between chains (Hayon et al., 2004, 2005). Further theoretical studies have shown that coupled synfire chain structures can produce compositionality, the hierarchical representation of complex entities in terms of parts and their relations (Bienenstock, 1995; Hayon et al., 2004; Schrader et al., 2007, 2011).

As yet there has been little search for and no convincing demonstration of synfire chains in real nervous systems. Such searches require new analysis methods for multi-neuron spike train recordings and identification of the corresponding synfire chain signatures. One appropriate method is our previously described algorithm (Schrader et al., 2008) that directly detects and shows many details about synfire chains if they are present in the recorded data. The computation is based on intersecting the identities of neurons active within any one time window (like 3 ms) with those at all other time windows. The resulting matrix $M(\boldsymbol{i,j})$ of intersection values at times $\boldsymbol{i,j}$ shows the orderly progression of firing in a synfire chain run as a short diagonal stripe of high value pixels. Details, controls and some sensitivity studies are in Schrader et al. (2008). With low background activity, that work showed that if a synfire chain comprises about 5% of the large (40k) neuronal net, detection would require a random sample of about 100–200

* Corresponding author at: Department of Neuroscience, 215 Stemmler Hall, University of Pennsylvania, Philadelphia, PA 19104, USA. Tel.: +1 215 898 8752; fax: +1 215 573 0833.
E-mail address: george@mulab.physiol.upenn.edu (G.L. Gerstein).

0165-0270/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2012.02.003

neurons. One purpose of the present paper is to extend the intersection algorithm so as to achieve comparable sensitivity in the presence of higher background activity.

A second purpose of this paper is to examine the use of repeating spatio-temporal patterns as a signature of synfire chains. Existing pattern programs (Abeles and Gerstein, 1988; Abeles and Gat, 2001; Gerstein, 2004; Shmiel et al., 2006; Patnaik et al., 2008; Louis et al., 2010a) sort spike patterns whose duration is less than some criterion into groups according to pattern complexity and repetition number and identify those groups whose count significantly exceeds corresponding counts found in surrogate data (Gerstein, 2004; Louis et al., 2010b) or by using conditional probabilities (Sastry and Unnikrishnan, 2010). Such repeating patterns have been suggested as evidence for synfire chains (Abeles et al., 1993; Oram et al., 1999; Ikegaya et al., 2004; Shmiel et al., 2005; Luczak et al., 2007), and indeed synfire chains will generate many repeating patterns. However other mechanisms can be envisioned also as demonstrated toward the end of this paper. To uniquely identify the presence of synfire chains from repeating patterns requires further demonstration that the various repeating patterns have the interleaving temporal structures expected in synfire chain activity. Using data from simulated nets with embedded synfire chains we describe tools for this purpose when parallel recordings from some hundreds of neurons are available.

2. Test data generation and description

Much of the test data for the spike train analyses described in this paper were originally generated by a large scale simulation program which contained 40 thousand excitatory and 10 thousand inhibitory point neurons arranged into 50 synfire chains each 100 wide and 20 long. Each neuron participated in an average of 2.5 different chains (Schrader et al., 2008). For some of the present purposes we edited such spike train data so that the identities and connectivity of only one synfire chain remained while all other neurons were replaced by independent spike trains with gamma 4 interval distribution and either high (8 Hz) or low (1.7 Hz) firing rates as an unstructured background. Individual firing rates were modulated to mimic the average rate fluctuations of the original data file. Further editing in some cases involved random choice of smaller populations out of these single chain data sets. Since neuron identities had been randomly assigned, such smaller populations always contained a sample of the chain neurons.

Additional test data for parts of Figs. 6 and 7 came from two different simulations with 80 thousand excitatory and 20 thousand inhibitory point neurons. In the “111” data set there were 379 chains each 130 wide and (on average) 100 long. Each neuron participated in 62 chains. In addition each chain was connected to two “successor” chains in which it could probabilistically initiate activity, thus forming a large, complicated and probabilistic tree in which the activity could propagate. In effect this produced activity in synfire chains of varying length depending on when (probabilistically) a particular activation would die out (Trengove et al., 2010b). In the “558959” data set there were 914 chains each 148 wide and 32 long. Each neuron participated in 52 chains. All but one of these chains received 3 stimulations during the run. A single selected chain received 240 stimulations, in effect making it the only active chain in the data. Background rates were 10 Hz. This is a variant of a recent model demonstrating large-scale embedding of synfire chains (Trengove et al., 2010a, 2012).

Test data for Fig. 8 came from a simulation of a single chain with 20 links of 100 neurons. The data were edited so that the near-synchronous (within 3 ms) firing of each link was maintained, but so that the activity each group was independently dithered by ±100 ms relative to the orderly sequence of the original synfire activity. An additional 4000 background neurons fired with gamma 4 interval distributions and rates of 1.7 Hz.

3. Signal to noise and firing rate in the intersection algorithm

We have previously described a sensitive algorithm for detecting synfire activity if present in a multi-neuron spike train recording (Schrader et al., 2008). The basic computation detects repetitive activations of a synfire chain, by searching for repeated sequences of synchronous activities. In comparing any two runs of a particular synfire chain similar sets of neurons from link 1, then link 2, etc. would be active. We search for repeats of such activity by first binning the data on a scale that corresponds to the expected inter-link delay (usually 3 ms), and list the identities of all neurons whose spikes fall into a given time bin along the data; call the list $\boldsymbol{S}(\boldsymbol{i})$ at time bin $\boldsymbol{i}$. We then compare the activities at time bins $\boldsymbol{i}$ and $\boldsymbol{j}$ by calculating the size of the intersection of $\boldsymbol{S}(\boldsymbol{i})$ and $\boldsymbol{S}(\boldsymbol{j})$, i.e. the number of neuron identifications that appear in both sets. After appropriate normalization the result is shown as a matrix, called the intersection matrix, with shading proportional to the normalized sizes of the $\boldsymbol{i,j}$ list intersections. This is also essentially equivalent to treating each list as a (0,1) vector $\boldsymbol{V}(\boldsymbol{i})$ at time bin $\boldsymbol{i}$ and computing the angles within all possible vector pairs. The normalizations in both formulations are essential to deal with possible rate variations along the data.

The matrix entry at times $\boldsymbol{i,j}$ with appropriate normalizations can therefore be formally expressed in two ways:

$$M_S(i,j) = \frac{|\boldsymbol{S}(\boldsymbol{i}) \cap \boldsymbol{S}(\boldsymbol{j})|}{\min(|\boldsymbol{S}(\boldsymbol{i})|,\ |\boldsymbol{S}(\boldsymbol{j})|)} \tag{1}$$

$$M_V(i,j) = \frac{\boldsymbol{V}(\boldsymbol{i}) \cdot \boldsymbol{V}(\boldsymbol{j})}{\sqrt{|\boldsymbol{V}(\boldsymbol{i})|^2 |\boldsymbol{V}(\boldsymbol{j})|^2}} \tag{2}$$

Each pair of runs of any particular synfire chain involves a repetition of an ordered sequence of identity lists, and therefore leads to a short diagonal stripe of high values in the intersection matrix. Length of the stripe corresponds to the number of links in the chain. Note that this result is independent of different time delays along the synfire chain as long as the two chain runs have the same time and identity properties. The short diagonal stripe might have moire-like variations of value if the original data binning does not quite correspond to the average inter-link delay in the synfire chain. If the two runs of the chain differ in timing values, the corresponding short diagonal stripe may shift among several adjacent diagonals and appear somewhat crooked. For N runs of the synfire chain there will be $N(N-1)/2$ run-pairs and therefore a systematic array of that number of short diagonal stripes in the intersection matrix. Detailed explanations and examples are given in Schrader et al. (2008).

In that paper we also examined signal to noise performance, i.e. the extent to which such short diagonal stripes stand out over the general background values in the matrix which result from the random (non-synfire) intersections of activity. We showed that the signal to noise increases with the ratio of sample to total population while its variability decreases. With a single synfire chain of 2000 neurons out of a total population of 50,000 neurons (i.e. 4%) and low background firing rates (1–2 Hz) reliable detection of the synfire activity requires observation of about a 100–200 neuron random sample.

Now consider two 50,000 neuron data sets with identical runs of one synfire chain (as above), but with different firing rates in all the non-synfire neurons. For pixels on one of the short diagonal stripes the numerator of Eq. (1) will differ little since most of the contribution to these list intersections will come from the synfire activity.

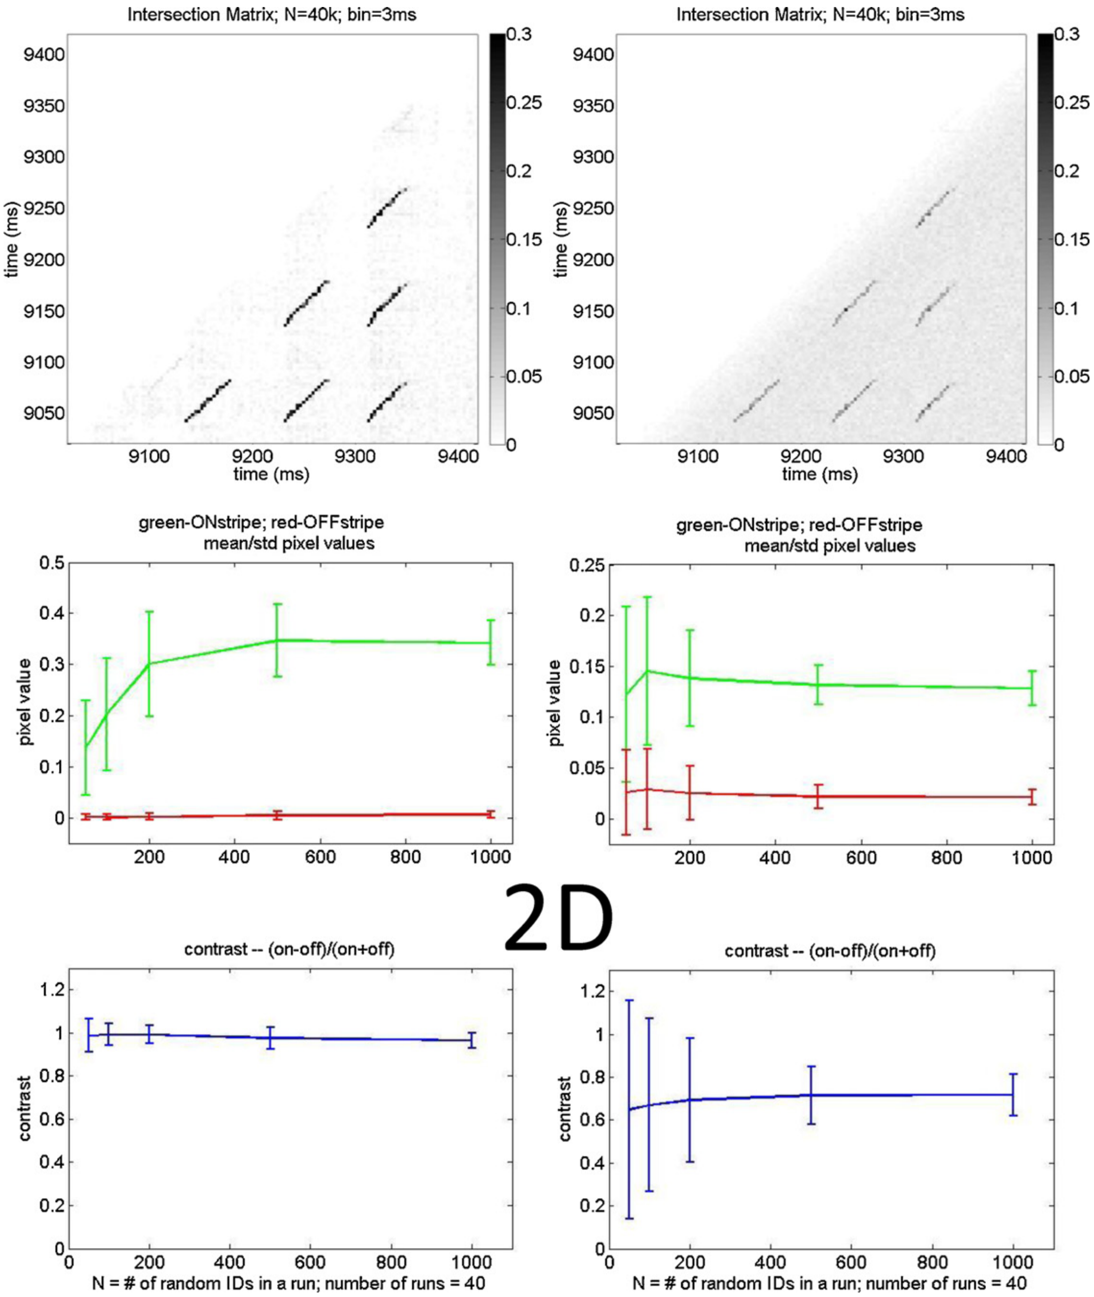

Fig. 1. Detection of synfire chain activity is harder with higher background firing rates. Two-dimensional analysis of simulation data that contains a single synfire chain and 38 thousand independent neurons with gamma 4 interval distributions. The left half data has independent neuron firing rate of 1.7 Hz, right half 8 Hz. Top panels are the intersection matrices for an arbitrarily chosen 0.4 s in the full data sets. The short diagonal stripes at the times of pairs of synfire runs are much fainter relative to background at the 8 Hz background rate. Middle panels are signal to noise analysis using total pixel value sum on or off a known short diagonal stripe. Each data point represents 40 random samples of the indicated size. On-stripe pixel values are lower and off-stripe values higher at the 8 Hz background rate. (Note that the y-axis calibration differs in the left and right middle panels.) Bottom panels are the corresponding contrast and show the reduction in contrast and increase in variability as the background rate increases.

But the length of lists in the denominator will depend directly on the background firing rates, and if this is high will reduce the pixel values on the short diagonal stripe. In addition, with higher background rates there will be more random intersections, and pixels *not* on the short diagonal stripes will have increased values. Both these effects of background rate are demonstrated in the matrices of the top two panels of Fig. 1, comparing background rates of 1.8 and 8 Hz for data with four closely spaced runs (and hence 6 run-pairs) of a synfire chain.

A more rigorous depiction of the signal/noise situation can be made by repeat calculations using varying numbers of randomly chosen neurons and comparing the average and standard deviation of intersection matrix values on and off a small diagonal stripe which had been detected with the full large data set. In order to have reasonable statistics we calculate over 40 disjoint random choices of each sample number of neurons and in each iteration calculate the average value over 15 pixels along 45 degrees both on and off the known short diagonal stripe location. Comparisons of results for 1.8 and 8 Hz background rates are in Fig. 1. The middle panels show mean pixel values on and off a (typical) short diagonal stripe (note the different ordinate scale at the different background rates), and the lower panels show the corresponding contrast. At the higher rate the mean pixel value on the short diagonal stripe is about 1/3 of that for the low rate, and the off-stripe values are about twice as high; the corresponding contrast drops from 1.0 to 0.6. Reliable detection of the synfire activity requires only 50 neurons with the low background rate, but 200 neurons at the high rate.

A simple way to improve the detection is to calculate intersections of activity lists at three instead of two times. The formal calculation is given by Eq. (3)

$$M_S(i,j,k) = \frac{\left|\boldsymbol{S}(\boldsymbol{i}) \cap \boldsymbol{S}(\boldsymbol{j}) \cap \boldsymbol{S}(\boldsymbol{k})\right|}{\min(|\boldsymbol{S}(\boldsymbol{i})|, |\boldsymbol{S}(\boldsymbol{j})|, |\boldsymbol{S}(\boldsymbol{k})|)} \tag{3}$$

Any three runs of a synfire chain will produce repeating sequences of identity lists, and hence a short diagonal stripe of high pixel values in the 3-dimensional matrix. For the same data set with 4 runs of one chain as used above, the top panels of Fig. 2 show the short diagonal stripes for the corresponding 4 run-triples in a 3-dimensional display. (Note that pixels in these displays do not have gray levels as in the 2-dimensional matrix display; all pixels above a threshold are shown at the same intensity.) The increased number of background pixels when data have higher background rate is clearly visible.

The quantitative analysis of signal to noise is in the lower panels of Fig. 2. As in the 2-dimensional analysis the pixel values on the short diagonal stripe are reduced at high background rates, but the off-stripe pixels remain low. As a result the contrast stays at 1.0 even at the high background rate, and the synfire chain detection level remains at samples of 50 neurons. Thus the 3-dimensional analysis is more sensitive than the 2-dimensional analysis in the face of higher background rates.

With the network and synfire chain parameters used above, approximately 5% of the neurons in each analyzed sample were members of the same chain. Thus inclusion of 3–5 chain neurons in a sample is needed for detection. Although not shown, if the data are diluted so that less than 5% of the overall network are chain neurons, detection requires a compensating increase in the sample size so that the necessary 5 or so chain neurons are included.

In the above study of sensitivity we compared pixel values on and off a short diagonal stripe whose location was determined from analyzing the full data set of 50,000 neurons (see the upper panels in Figs. 1 and 2). Obviously this is not applicable to experimental data limited by current technology to less than 200 neurons. However as shown previously (Schrader et al., 2008) the 2-dimensional short diagonal stripes become visually detectable (possibly with the aid of appropriate spatial filtering of the matrix) at around sample sizes of 100 neurons at low background firing rates. High background rates require larger samples of the population for synfire detection.

4. Repeating pattern spectrum richness with chains

We have previously described algorithms for finding all repeating temporal patterns among simultaneous recordings of multiple spike trains (Abeles and Gerstein, 1988). The basic idea is to represent the multiple observed spike trains as holes in an old time paper computer tape, different neuron identities across the tape, time along its length. Two copies of this tape are shifted in time (along the length) and regions where holes superimpose represent repeating spatial-temporal patterns. Our implementation extracts the longest repeating patterns that fall within a specified duration. It does not also count the repetitions of sub-patterns of the detected patterns as do the published methods based on bottom up data mining algorithms (Shmiel et al., 2005; Patnaik et al., 2008).

The results of our programs are presented as a pattern repetition spectrum as shown in Fig. 3. Here we have a bar plot of the logarithm of the number of detected patterns as a function of complexity and repetition number. Complexity is the number of spikes that are involved in a given pattern; this is not a measure of how many individual spike trains (neurons) are involved.

Examination of significance of counts in the pattern repetition spectrum can be based on multiple repetition of the spectrum computation with surrogate data. The generation of appropriate surrogates requires matching the original rate and inter-spike interval structures but destroying only patterning. Useful ways to do this depend on dithering of each individual spike, i.e. moving it some relatively small distance in time (Date et al., 1998; Hatsopoulos et al., 2003). However, additional details of the dither can produce better preservation of rate and interval. We have shown that a dither probability shape based on the square root of the joint interval histogram (JIH) of that spike train is very effective in preserving interval structure (Gerstein, 2004) although the rate structure is somewhat distorted. If there is a repeating underlying rate structure even better results can be obtained if the joint interval histogram is used in operational time (Louis et al., 2010b). Other types of significance test for patterns have been suggested using conditional probabilities (Sastry and Unnikrishnan, 2010). Some work using different surrogate generating methods (Baker and Lemon, 2000; Baker et al., 2001) concludes that repeating patterns in some experimental data occur at chance levels.

In Fig. 3A we show the pattern repetition spectrum for 500 independent spike trains, each with inter-spike interval distributions of gamma order 4 and with a rate of 1.7 Hz and slow rate modulation mimicking our original simulation with synfire chains. Note that all repeating patterns are clustered at low complexity and low repetitions. In Fig. 3B we show the pattern spectrum for 500 randomly chosen neurons from the original full data set with 50 synfire chains. There is now an extraordinary richness of patterns of high complexity and repetition. A similar sample in Fig. 3C from the data set that contains only one synfire chain of the 50 with the rest of the net independent, shows fewer but still very many patterns at high complexity and repetition. Finally in Fig. 3D we show the pattern spectrum from the same 50 chain data set as in Fig. 3B but dithered with a flat distribution in a ± 5 ms window. With these data the pattern richness has reverted to clustering at low complexity and repetition similar to the result from independent neurons.

The comparison of Fig. 3B and D also demonstrates that the high pattern counts at complexity >3 and repetition >2 in 3B are highly significant, since a dither surrogate almost completely eliminates them. This result might depend on data length or rates, so that it is always wise to make the simple dither test; if met it is generally not necessary to go through the procedure of many repeated surrogate calculations to assess significance.

It is interesting to examine what neurons participate in patterns of any particular complexity and repetition. From the data used in Fig. 3B we arbitrarily chose the five patterns at complexity 8 and repetition 19. These patterns are analyzed in the five panels of the left two columns of Fig. 4. Time structure of each pattern in milliseconds is indicated above its panel; neuron identities are different for each pattern but are omitted for space reasons. Recall that these data contained 50 synfire chains. In each panel we show how many neurons of the analyzed pattern (indicated in the title of the panels) belong to a particular chain. Since in these data each neuron participated in 2–3 different synfire chains we expect that the neurons in each pattern would be members of various chains (the low bars). However for every pattern *all* 8 neurons participating in that pattern are members of one particular chain, different for each pattern (the high bar). The panels in the right two columns of Fig. 4 show the same analysis for the data used in Fig. 3C with a single synfire chain. Here we arbitrarily chose the 7 patterns of complexity 5 and repetition 11. For each pattern *all* 5 neurons were members of chain 35 (the only one in these data). Although not shown, when the detected pattern includes intervals <2 ms some of the neurons are members of the same link in the chain.

Similar analyses (not shown) of the same data for the neuron identities in patterns represented at the edges and near the

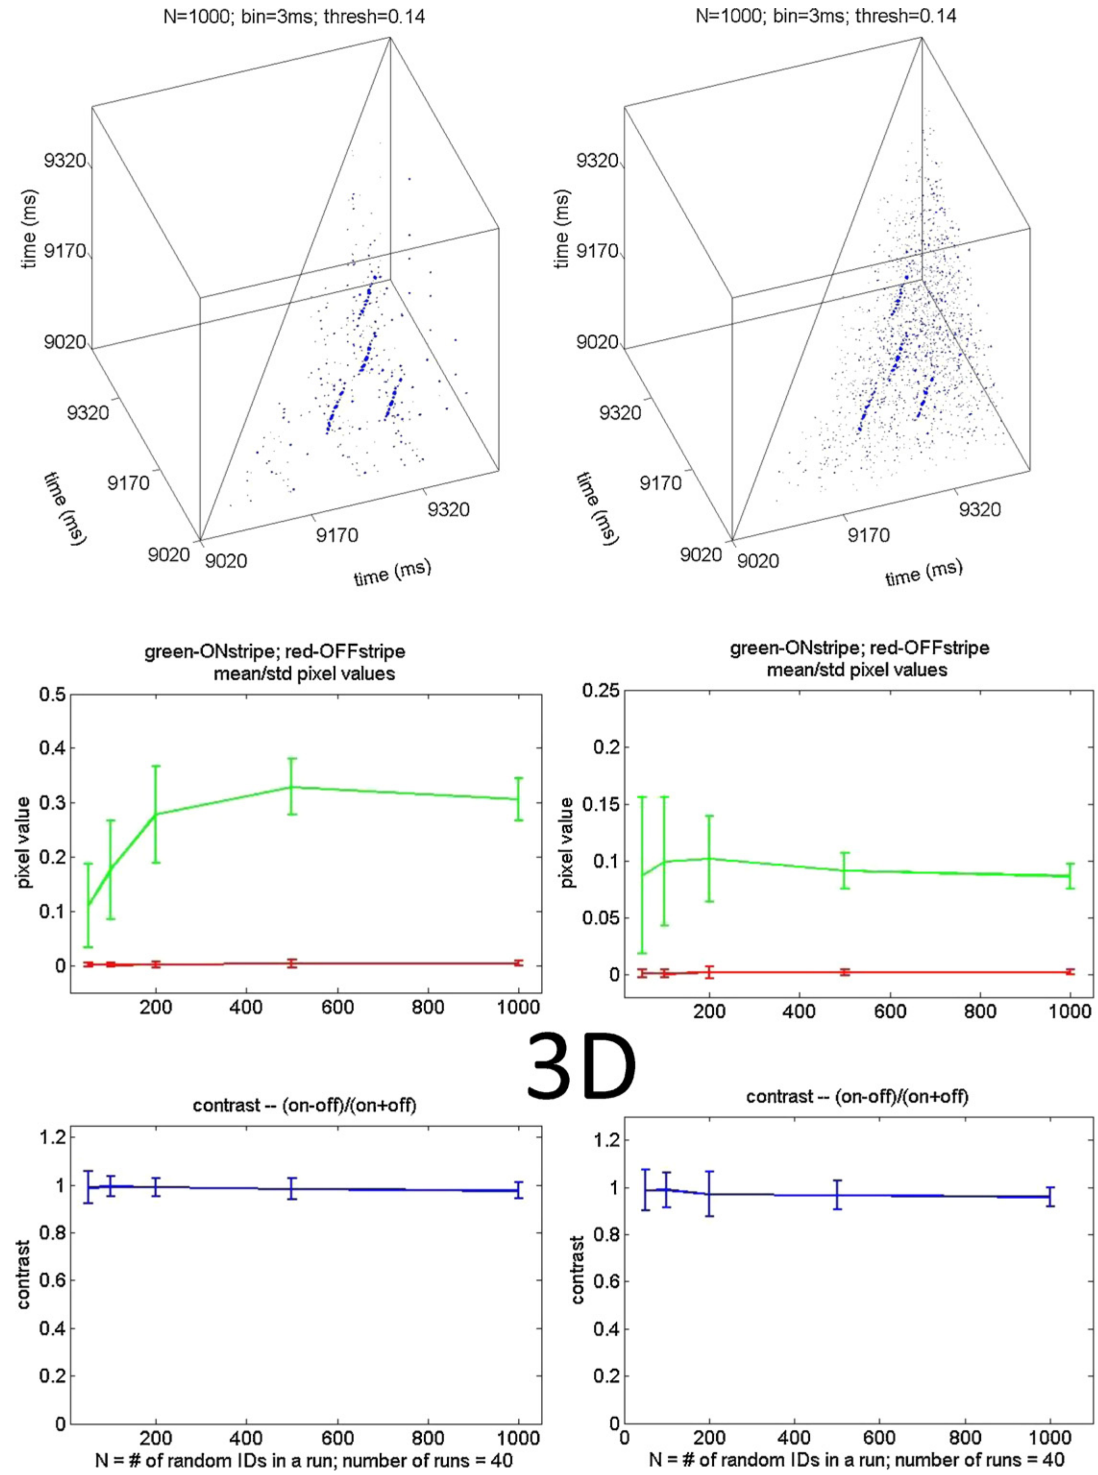

Fig. 2. Analysis of the same simulated data carried out with the 3-dimensional version shows that detection is less affected by background firing rate. Pixels in the 3-dimensional matrix indicate overlap of active neuron identities at three times. The short diagonal stripes remain strong at the higher background rate. The signal to noise analysis in the middle panels differs from the 2-dimensional results mainly in that the off-stripe pixel count remains low with the 8 Hz background rate. (Note that the *y*-axis calibration differs in the left and right middle panels.) The contrast remains at 1.0 even with the high background rate.

origin of the plane, i.e. with complexity 3 and/or repetition 2 (i.e. low complexities and repetition), would usually include or be entirely composed of non-synfire neurons. As in the examples above, the neuron identities in patterns towards the center of the plane (complexity >3, repetition >2) are entirely from the synfire chain.

These data had low spontaneous firing rates of the non-chain neurons. As will be shown below, increasing these firing rates results in patterns at complexity >3 and repetition >2 that also include non-chain neurons.

The total number of repeating patterns depends both on analysis parameters and on properties of the data. Relevant analysis parameters include maximum allowed pattern length, bin width or time resolution, and the random sample size drawn from the total population. Relevant data properties include data length, the percentage of synfire neurons in the

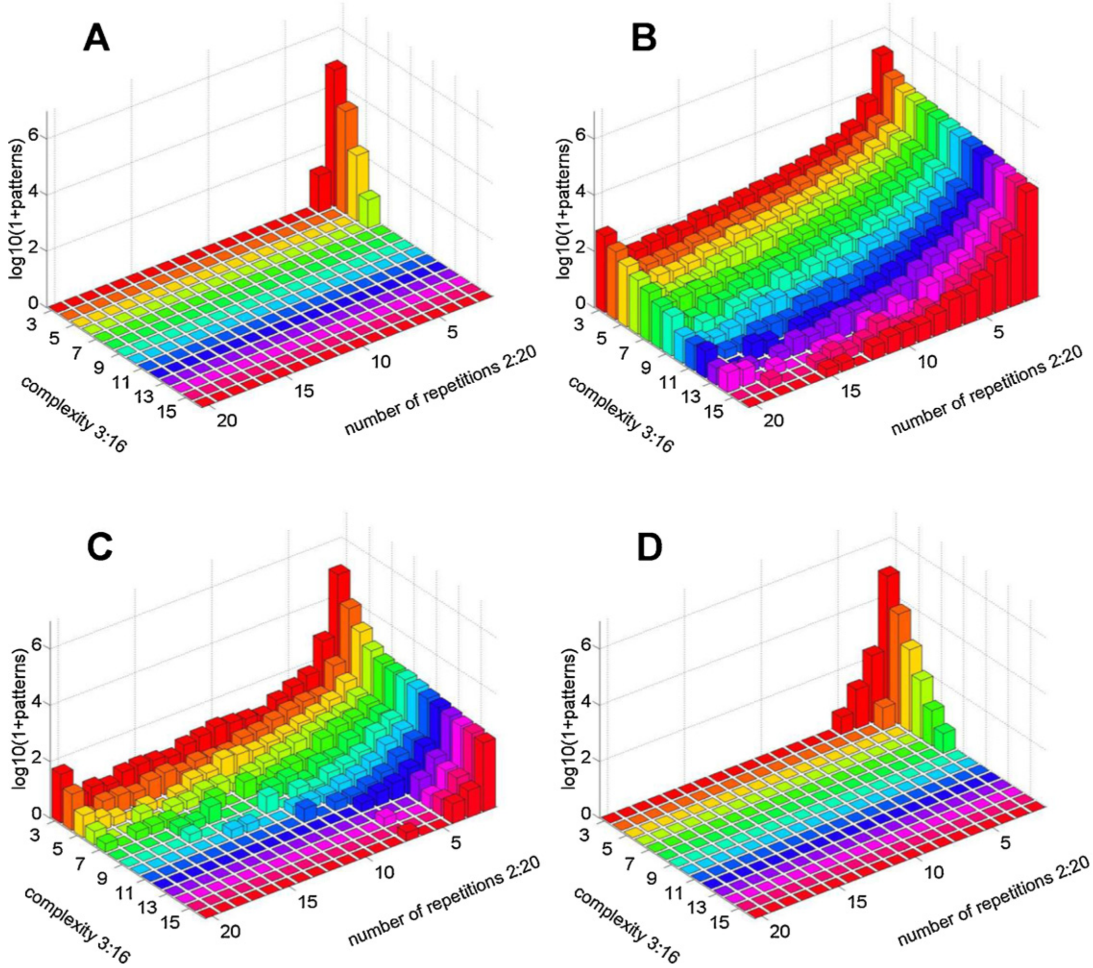

Fig. 3. Pattern repetition spectra for 500 randomly chosen neurons from each of four different simulations. Synfire chain data shows a rich set of patterns at high complexity and repetition. (A) Independent firing of 40 thousand neurons each with gamma 4 interval distribution and mean rate of 1.7 Hz. (B) 40 thousand neurons arranged as 50 synfire chains. Each neuron participated in an average of 3.5 chains. Chains were stimulated by independent Poisson sequences at a rate of 1 Hz. (C) One synfire chain and 38 thousand independently firing neurons as in (A). (D) The same data as in (B) but with each spike time dithered by ±5 ms.

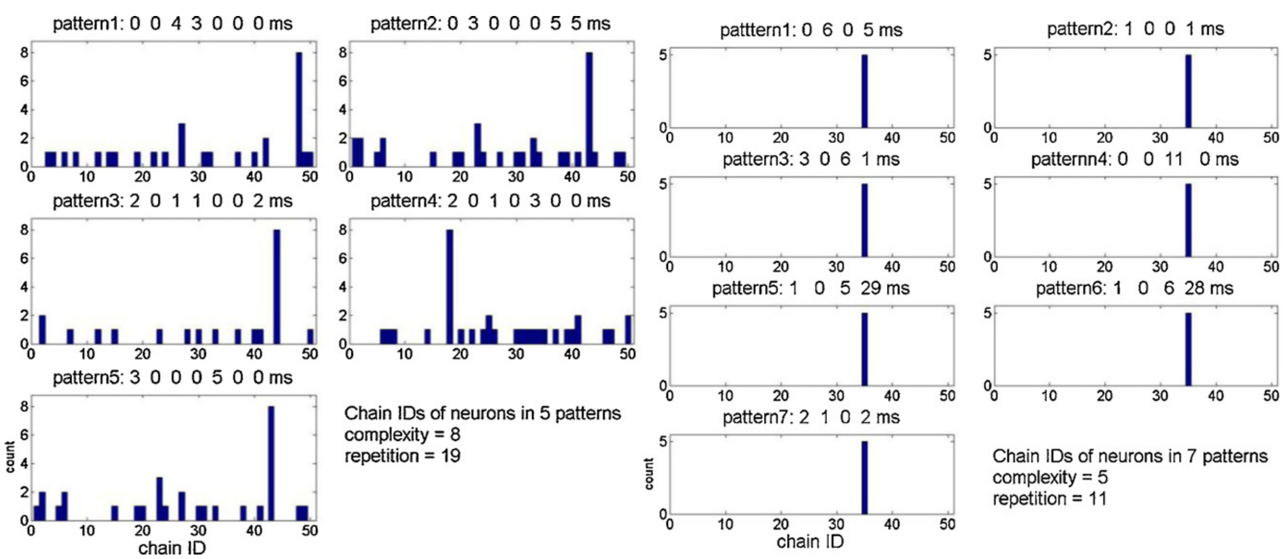

Fig. 4. Chain membership of neurons that participate in patterns arbitrarily chosen from the regions of complexity >3 and repetition >2 in Fig. 3. Each panel shows for one particular pattern how many neurons (*y*-axis) of the pattern are members of a particular synfire chain (*x*-axis: chain ID). Left panels: 5 different patterns of complexity 8 and repetition 19 from data with 50 synfire chains shown in Fig. 3B. The assignment of each neuron to several different chains produces the low bars; but in each panel all eight neurons come from some single chain (the high bar). Right panels: the same analysis for the 7 patterns of complexity 5 and repetition 11 from Fig. 3C. All five neurons in each different pattern come from the single synfire chain in the data.

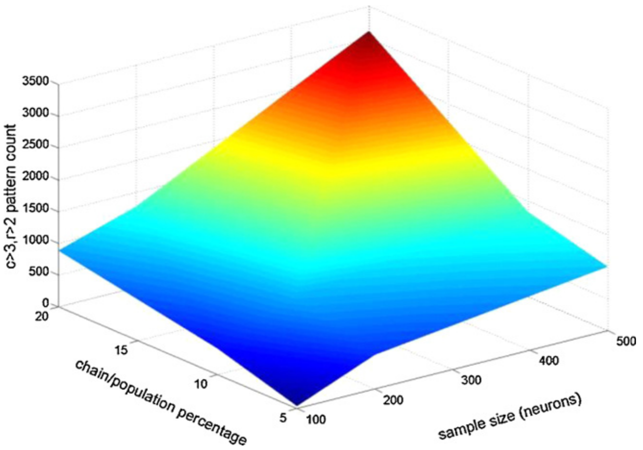

Fig. 5. Dependence of the total number of patterns for complexity >3 and repetition >2 (i.e. the center region of the complete pattern spectrum) on sample size (number of neurons) and the fraction of chain to non-chain neurons in the data.

total population (and in the random sample) and the properties of the non-synfire neurons (like rates, ISIs, temporal modulation, correlated firing, etc.). Using data sets similar to that used for Fig. 3C, Fig. 5 demonstrates the dependence of the total number of complexity >3, repetition >2 patterns on sample size and percentage of synfire neurons in the data set. Increase of either parameter produces a monotonic increase of the total number of patterns in the center part of the plane. Reciprocal changes of these parameters can keep the number of patterns constant.

Although not shown, we have examined pattern properties when the two parameters of rate and regularity are varied. Intuitively the number of patterns should increase with regularity, going to extremely high numbers for a clocklike set of neurons. For this investigation we used simulations with a single synfire chain as in Fig. 3C, and edited the properties of the 38,000 non-chain neurons. Regularity was specified by the order of the gamma function used to generate their inter-spike intervals. Higher gamma corresponds to a narrower ISI distribution and therefore higher regularity. For the non-chain neurons we varied gamma from 4 to 12, and rates from 1.7 Hz (same as in the single synfire chain) to 7 Hz.

Gamma order in this range produces smaller changes in total pattern counts or $c > 3$, $r > 2$ pattern counts than the variability produced by random selection of 500 neurons from the data. It is likely that much higher values of the gamma order would have a visible effect. By contrast increasing the rate of the non-chain neurons in the range examined increases the total pattern count by a factor of approximately 400, but cuts the $c > 3$, $r > 2$ counts in half. This last observation can be explained by recalling that at the original low rate of 1.7 Hz and gamma order 4 most of the $c > 3$, $r > 2$ patterns originate from the synfire chain as shown in Fig. 4. As the firing rate of the non-chain neurons is increased there is an increasing likelihood that additional (non-chain) spikes will appear in the existing synfire patterns and that this will occur repetitively. When this happens the assignment of the pattern moves on the plane to a higher complexity and a lower repetition (usually only $r = 2$). Thus, as also shown in Figs. 1 and 2, the effects of higher rates and/or very high regularity in the non-chain neurons tend to obscure the signature of synfire activity, and make its detection that much harder.

4.1. Pattern timing issues

In the above we have demonstrated that networks containing synfire chains produce a rich spectrum of patterns with high complexity and repetition. The converse, however, is not true. Some repeating patterns with complexity >3 and repetition >2 do occur randomly. Furthermore many such repeating patterns can be generated by other means than synfire chains. An example using approximately synchronous firing within each of several independent neural groups (cortical columns?) is given towards the end of the paper. Thus finding repeating patterns is by itself not enough to infer the presence of synfire chains. Additional measurement of the time structure of the repeating patterns is essential in order to make that inference. In order to develop these requirements we examine the detailed time structures of patterns observed in simulated nets with synfire chains as above.

For this investigation we have used four different large scale simulations (40,000 or 80,000 excitatory neurons) with one or more embedded synfire chains. (Parameter details for these simulations are in the section on data generation.) The pattern spectra of our four example simulations are shown in Fig. 6, with the left column for the simulations with many chains and the right column for the simulations with only a single chain. In each case a random sample of 100 neurons (repeated to check stability but repetitions not shown) was used to mimic the current limits of recording technology. All four spectra show many patterns at complexity >3 and repetition >2, i.e. the middle of the plane excluding the upper edge and right edge values. The effect of sample sizes on pattern counts is seen by comparing Fig. 3B with Fig. 6A and Fig. 3C with Fig. 6B and is explicitly shown in Fig. 5. The pattern richness is characteristic (but not unique) for synfire chains; simulations with only independent spike trains show spectra with patterns mainly at the edges and origin of the plane (see Fig. 3A).

For each such spectrum calculation we have examined temporal characteristics of all patterns of complexity >3 and repetition >2 by constructing (1) a raster of when each pattern (represented as a point) occurs in the data set; (2) each full pattern displayed as a raster and as a spike time histogram (like a PST histogram) in time lock (i.e. alignment) to the known synfire chain stimuli; (3) each full pattern displayed as a raster and as a spike time histogram both aligned on the first spike of the pattern.

The first of these measures, the overall raster, is useful to assess temporal heterogeneity, including for real data the possible modulation associated with behavior or the existence of recording problems. However this is of little use here since we are using simulations that are fairly steady state. The second measure, locking detected patterns to synfire chain initiating stimuli, indeed demonstrates that most patterns in these data are associated with runs of the synfire chain(s). However this measure depends on knowledge of the stimuli times which would generally not be available in an experimental situation. The third measure, aligning patterns on the first spike, turns out to be sensitive to the delay between links in the synfire chain. This should be applicable to experimental data. With simulated data this measure also clearly identifies some patterns which in part include firing of non-chain neurons.

For each of the four data sets in Fig. 6 the corresponding spike time histogram of complexity >3, repetition >2 aligned patterns is shown in Fig. 7. In each case there are prominent peaks which tend to fall in bins (1ms width) with spacing of about 3 ms. This is the time for propagation between successive links in the chain models used in these simulations so that patterns arising from synfire chain activity would have multiples of this basic timing in their interval structure. The background in the histograms below the peaks represents the patterns which involve the activity of non-chain neurons. Note also that the large peaks come early in the pattern except in case C. All these simulations have variability built into the link to link delays. These temporal variations build up during each run of a chain, so that exactly repeating patterns (to 1 ms time resolution) in two or more runs of the chain can originate only from a limited portion of the full chain length and thus have a duration less than that of a full chain run. The spaced peak structures in each panel of Fig. 7 are therefore an additional signature of synfire chain activity.

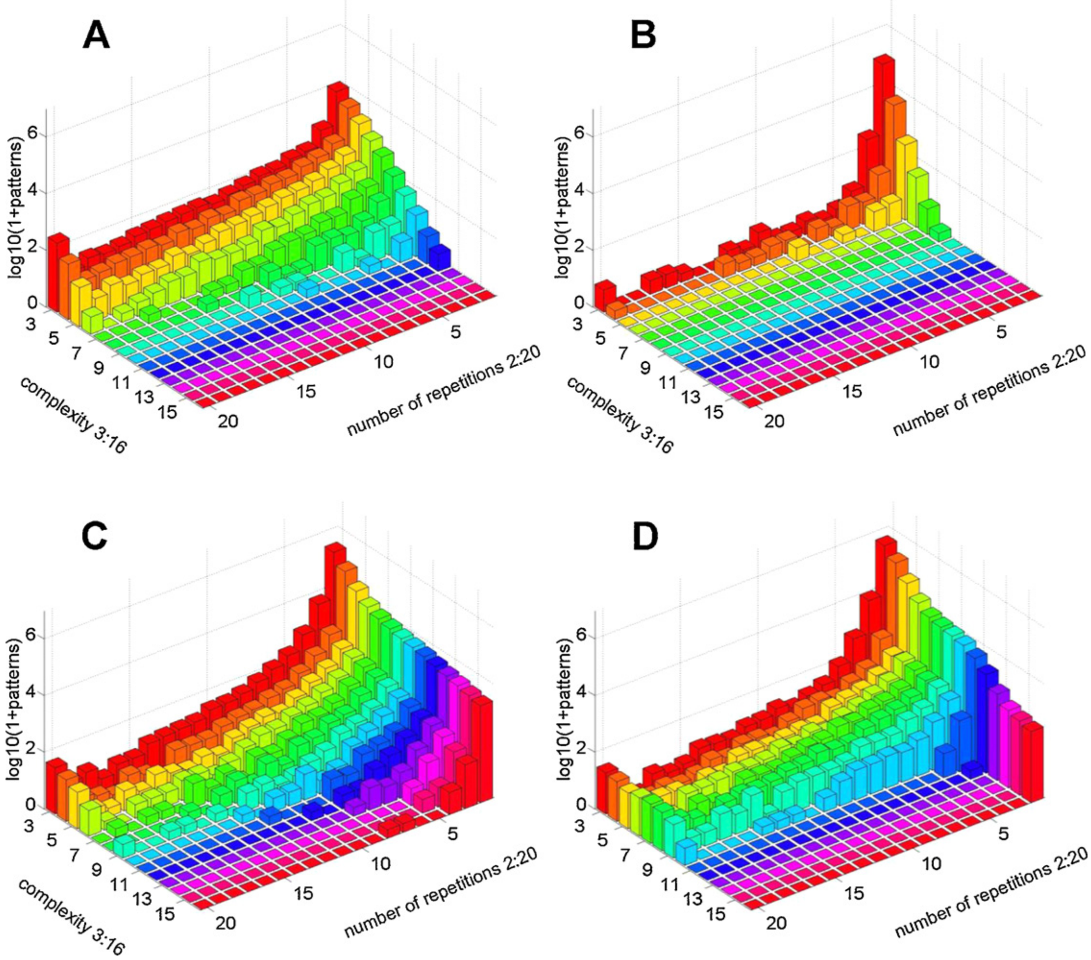

Fig. 6. Pattern repetition spectra for 100 randomly chosen neurons from four different simulations. (A) Same data as Fig. 3B. (B) Same data as Fig. 3C. These two panels demonstrate the effects of sample size on pattern numbers. (C) 100 thousand neurons in a multiply branching tree of synfire chains. (D) 100 thousand neurons arranged as many overlapping chains; in effect only one chain is stimulated. See section on data simulation for details of C and D. All four panels show the rich pattern repetition spectrum at complexity >3 and repetition >2 characteristic of synfire chains.

5. Group synchrony without chain structure

Several papers in the literature have claimed evidence for synfire chain activity based on repeating patterns or sequences (Euston et al., 2007; Ikegaya et al., 2004; Yosimura et al., 2005; Prut et al., 1998; Shmiel et al., 2005, 2006; Abeles et al., 1993). In some of these papers various surrogates were used to assess significance of the pattern counts. In this section we analyze simulation data that produces many complexity >3, repetition >2 patterns, but contains no synfire chains. Such data can be distinguished from synfire-containing spike trains by the two additional measurements set out in this paper: (1) the matrix of neuron identity intersections and (2) the time structure of the detected repeating patterns.

The data consist – as the above synfire chain simulations – of activity from 20 groups of 100 neurons (same parameters as used in many of the above synfire chain simulations) with the neurons in each group firing in near synchrony (within 3 ms) but without the groups being sequentially activated. The data were generated from the simulated data containing one synfire chain but here the synchronous spikes from the individual links were together dithered in time within a window of $\pm$100 ms. A raster display of a 0.7 s segment of the firing activity of the 20 groups is shown in Fig. 8A. There is no ordered sequential firing of the several groups as would be the case for links in a synfire chains. (Comparable rasters of simulated orderly synfire chain activity are in Schrader et al., 2008). The data also contained 4000 independent neurons firing at 2.7 Hz with gamma 4 interval distributions.

The neuron identity intersection matrix is for an arbitrary 700 ms segment of the data is shown in Fig. 8B. Unlike the matrices in the top panels of Figs. 1 and 2, there are no orderly short diagonal stripes of pixels, just an apparently random distribution of single pixels. This of course is expected, since these data have no repeating sequential ordering to the activity of the several groups.

The pattern spectrum of a random 100 neurons chosen from this activity is in Fig. 8C, and shows considerable repeating pattern counts for complexity >3, repetition >2. Since we have full knowledge of this simulation data, examination of the neuron identities contributing to these patterns (as in Fig. 4) shows that most patterns originate within some single group (not explicitly shown). There are few patterns with members spread among one or more groups and background activity as is the case in the ordered sequences of synfire activity. This prevalence of single group origin of patterns is made obvious by the histogram of spike times in aligned patterns shown in Fig. 8D. Unlike the synfire data in Fig. 7, many patterns here have durations of less than 3 ms, indicating their origin in a single group. Note that the other repeating patterns have a continuous distribution of inter-spike times. This contrasts with the spaced peaks in the histograms of spike times in aligned patterns shown in Fig. 7 which correspond to the delay structure of a synfire chain.

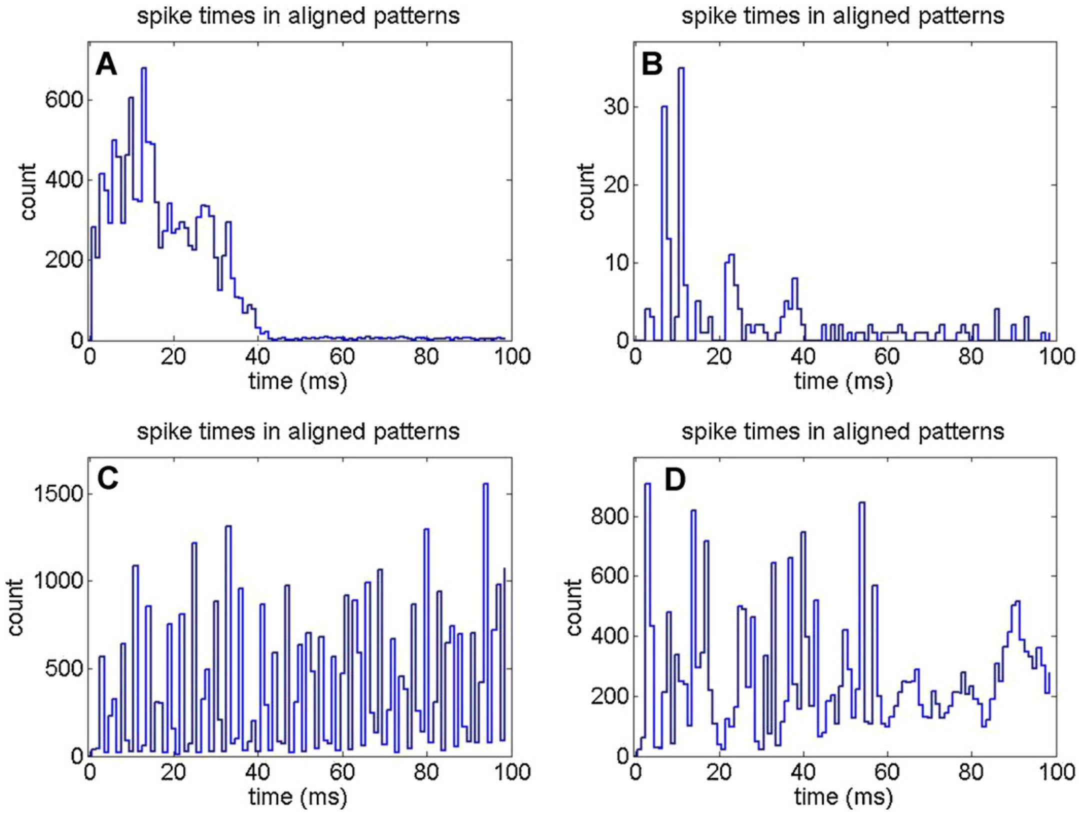

Fig. 7. Histograms of spike times in patterns aligned to first spikes. Data and arrangement are identical to Fig. 6, but only patterns of complexity >3 and repetition >2 are considered. Bin width of the histograms is 1 ms. Each panel shows various numbers of prominent peaks with 3 ms spacing. This corresponds to the interlink delays built into the several synfire simulations.

Thus with at least this one counterexample, repeating patterns alone are not enough to infer synfire activity. It is also essential to examine the matrix of active neuron identity intersections as well as the time structure of the detected repeating patterns.

6. Discussion

The synfire chain is an intriguing theoretical concept with interesting properties that should make it very applicable to understanding many aspects of brain function. Unfortunately there has to date been no convincing experimental demonstration of synfire chains due partially to the difficulties of multi-neuron recording but also to the lack of appropriate analytic tools. We have previously published one such analysis tool (Schrader et al., 2008). In the present paper, using simulated data, we have extended that work and in addition examined two new tools for detecting the presence of synfire chain activity in neuronal spike train recordings. Such analyses require data from at least one hundred simultaneously recorded spike trains, although larger samples would be preferred in order to improve reliability and sensitivity.

There are serious methodological difficulties in obtaining the necessary data. Of the currently available recording methods the multiple electrode approaches all suffer to some degree from the problem of clean separation of spike waveforms (Warren et al., 2001; Csicvari et al., 2003) from different neurons detected on the same electrodes (Harris et al., 2000). Such mixing of spike trains seriously degrades and makes ambiguous any inference from measurements of spike timing (Gerstein, 2000; Pazienti and Grün, 2006), and depending on the degree of contamination, would impact all three methods developed in the present paper. The older dye based optical methods have either spatial or temporal resolution difficulties as well as problems with bleaching and opto-kinetic damage to the neurons. There is hope however in the newer techniques with 2-photon calcium dye methods as well as with the rapidly developing genetically engineered fluorescent protein approaches (Knöpfel et al., 2003; Ohki et al., 2005, 2006; Miri et al., 2011; Miyawaki, 2005).

In a previous paper (Schrader et al., 2008) we described a simple analysis of the active neuron populations in the data that yields a unique signature of synfire activity if it exists. In the present paper we have further studied the signal to noise properties of this analysis and have extended it to the detection of triples rather than pairs of synfire chain runs. This has provided a considerable gain in sensitivity and particularly reduced problems with relatively high rate background neurons. Our signal to noise studies compared matrix pixel values on and off a short diagonal stripe whose location was determined by analyzing the large full set of simulated neurons. This obviously cannot be done in an experimental situation. However, results from the procedure should be applicable to the threshold for visual detection of short diagonal stripes (possibly aided by appropriate spatial filtering of the matrix) for experimental data. We have previously shown that visual threshold is about 100 neurons in the 2-dimensional pair analysis at low background rates. The 3-dimensional analysis of chain run triples with its better signal to noise properties should maintain this performance even at higher background firing rates. Nevertheless the sampling problem remains critical especially with high background rates.

Our second approach to detecting synfire chain activity is based on finding repeating patterns of firing. Methods for finding all such repeating patterns are well known (Abeles and Gerstein, 1988; Shmiel et al., 2005; Patnaik et al., 2008) and can be top down or bottom up. Data from simulations containing multiple synfire chains or even with a single synfire chain in a great sea of background neurons (Figs. 3 and 6) show a great richness of patterns of complexity

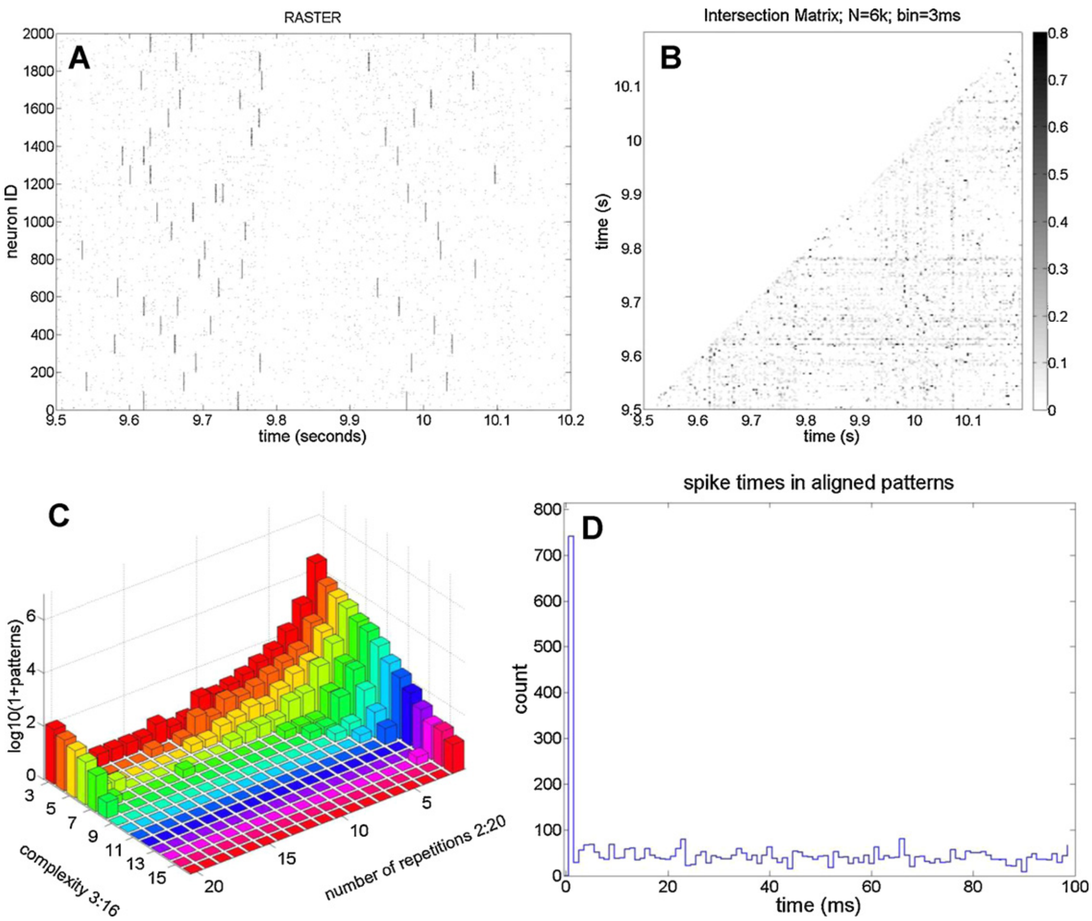

Fig. 8. Analysis of simulation data containing synchronous groups (3 ms precision), that were not active sequentially in time. The data were generated from simulated data containing one synfire chain. Same synfire chain data as in Fig. 3C but here the synchronous spike responses of the 20 individual groups of 100 neurons are dithered in time within a window of $\pm$100 ms instead of being sequentially ordered as in the original synfire chain. (A) Raster display showing the activity of the 20 groups over an arbitrary 0.7 s interval. Activity of the remaining 4000 gamma 4 background neurons is not shown. (B) Activity intersection matrix of an arbitrary 700 ms segment of the data. This is a similar time scale to that used in Fig. 1, but shows only randomly arranged individual pixels rather than short diagonal stripes. (C) Corresponding pattern repetition spectrum for a random sample of 100 neurons. (D) Corresponding time histogram of spikes in aligned complexity >3 repetition >2 patterns. The single peak at 1–2 ms indicates that members of a given repeating pattern were mostly in some single group. Repeating patterns with longer intervals do not show the discrete structure corresponding to synfire chain delays. Thus repeating patterns per se are not necessarily an indicator of synfire activity.

>3 and repetition >2. We have demonstrated in Fig. 4 for synfire data that the neurons involved in each such pattern are usually members of some single chain. Although directly unverifiable, one would hope this would also be true in patterns high complexity, high repetition data when found in experiments. Comparable multi-neuron spike trains from simulations of independent neurons do show repeating patterns but mostly at complexity 3 and/or repetition 2.

Our third approach to detecting synfire chain activity is based on examining the detailed time structure of any repeating patterns. The most useful such measure is to align the repeating patterns to their first spike. With synfire data a histogram of spike times of such aligned patterns tends to show peaks with spacing corresponding to the interlink delays of a synfire chain (Fig. 7).

The detection methods presented here also apply to a generalization of synfire chains known as braids (Bienenstock, 1995) or polychrony (Izhikevich, 2006) which are characterized by the rule that all direct and indirect paths between any two neurons have the same delay. In these feed-forward structures the neurons need not be grouped into discrete links, and they support reproducible, precisely timed spike sequences that need not be grouped into synchronous packets. The delays on the feed-forward connections preserve the key property of synfire chains that, during spike sequence re-activation, each neuron receives synchronously arriving spikes from upstream neurons. Repeated activations of braids, just as for chains, will give rise to diagonal stripes in the intersection matrix and will produce temporal patterns with high complexity and repetition. However, the comb-like structure seen in Fig. 7, being specific to chains, would not generally be found in braids.

Although synfire chains produce many high complexity, high repetition patterns so can other mechanisms. We have demonstrated one such situation with a simulation of near synchronous firing within neural groups but with each such group active at random times. Such data (Fig. 8) show rich patterning at complexity >3 and repetition >2, but as expected the intersection matrix has only randomly arranged pixels rather than short diagonal stripes. The histogram of spike times in aligned patterns reflects the prevalence of patterns limited to the near synchronous firing within each group. Patterns with longer intervals have continuously distributed intervals rather than preferred discrete values which would correspond to the chain delay structure.

Finally we suggest that experimental detection of synfire chains cannot rest exclusively on the demonstration of repeating firing patterns. It is necessary in addition to demonstrate an activity intersection matrix with short diagonal stripes, and a histogram of spike times in aligned patterns that shows that many patterns have a structure appropriate for expected inter-link delays in a synfire chain. We have so far had access to four multiple neuron recording data sets from four different laboratories and preparations. All of them show repeating patterns, but none of them have intersection matrices with short diagonal stripes or appropriate spaced peaks in the spike times of aligned patterns. The search continues.

Acknowledgments

Partially funded by EU Grant 269921 (BrainScaleS), DIP F1.2, the Helmholtz Alliance on Systems Biology (Germany), and the Next-Generation Supercomputer Project of MEXT (Japan). ERW was supported by the CARMEN project, funded by the Engineering and Physical Sciences Research Council, UK (www.carmen.org.uk).

References

Abeles M. Local cortical circuits: an electrophysiological study. Berlin: Springer-Verlag; 1982.
Abeles M. Corticonics neural circuits of the cerebral cortex. Cambridge, UK: Cambridge Univ. Press; 1991.
Abeles M, Bergman H, Margalit E, Vaadia E. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. J Neurophysiol 1993;70:1629–38.
Abeles M, Gat I. Detecting precise firing sequences in experimental data. J Neurosci Methods 2001;107:141–54.
Abeles M, Gerstein GL. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. J Neurophysiol 1988;60:909–24.
Aertsen A, Diesmann M, Gewaltig M. Propagation of synchronous spiking activity in feed-forward networks. J Physiol (Paris) 1996;90:243–7.
Aviel Y, Mehring C, Abeles M, Horn D. On embedding synfire chains in a balanced network. Neural Comput 2003;15:1321–40.
Baker S, Lemon RN. Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance level. J Neurophysiol 2000;84:1770–80.
Baker S, Spinks R, Jackson A, Lemon RN. Synchronization in monkey motor cortex during a precision grip task. I. Task-dependent modulation in single-unit synchrony. J Neurophysiol 2001;85:869–85.
Bienenstock E. A model of neocortex. Network Comput Neural Syst 1995;6:179–224.
Csicvari J, Henze DA, Jamieson B, Harris KD, Sirota A, Bartho P, et al. Massively parallel recording of unit and local field potentials with silicon-based electrodes I. Neurophysiol 2003;90:1314–23.
Date A, Bienenstock E, Geman S. On the temporal resolution of neural activity. Technical report, Division of Applied Mathematics, Brown University 1998 (Available at: http://www.dam.brown.edu/people/elie/papers/temp-res.ps).
Diesmann M, Gewaltig MO, Aertsen A. Stable propagation of synchronous spiking in cortical neural networks. Nature 1999;402:529–33.
Euston DR, Tatsuno M, McNaughton BL. Fast forward playback of recent memory sequences in prefrontal cortex during sleep. Science 2007;138:1147–50.
Gerstein GL. Cross-correlation measures of unresolved multi-neuron recordings. J Neurosci Methods 2000;100:41–51.
Gerstein GL. Searching for significance in spatio-temporal firing patterns. Acta Neurobiol Exp (Warsz) 2004;2:203–7.
Gewaltig M-O, Diesmann M, Aertsen A. Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. Neural Netw 2001;14:657–73.
Goedeke S, Diesmann M. The mechanism of synchronization in feed-forward neuronal networks. New J Phys 2008;10:015007, http://iopscience.iop.org/1367-2630/10/1/015007.
Harris KD, Henze DA, Csicsvari J, Hirase H, Buszaki G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J Neurophysiol 2000;84:401–14.
Hatsopoulos N, Geman S, Amarasingham A, Bienenstock E. At what time scale does the nervous system operate? Neurocomputing 2003, 25–29, 52–54.
Hayon G, Abeles M, Lehmann D. Modeling compositionality by dynamic binding of synfire chains. J Comput Neurosci 2004;17:179–201.
Hayon G, Abeles M, Lehmann D. A model for representing the dynamics of a system of synfire chains. J Comput Neurosci 2005;18:41–53.
Hebb D. The organization of behavior: a neuropsychological theory. New York: Wiley; 1949.
Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, et al. Synfire chains and cortical songs: temporal modules of cortical activity. Science 2004;304:559–64.
Izhikevich EM. Polychronization: computation with spikes. Neural Comput 2006;18:245–82.
Knöpfel T, Tomita K, Shimazaki R, Sakai R. Optical recordings of membrane potential using genetically targeted voltage-sensitive fluorescent proteins. Methods 2003;30(May (1)):42–8.
Louis S, Borgelt C, Grün S. Complexity distribution as a measure for assembly size and temporal precision. Neural Networks 2010a;23:705–12.
Louis S, Gerstein G, Grün S, Diesmann M. Surrogate spike train generation through dithering in operational time. Front Neurosci 2010b;4(September) (article 127).
Luczak A, Barthóf P, Marguet SL, Buzsáki G, Harris KD. Sequential structure of neocortical spontaneous activity in vivo. Proc Natl Acad Sci U S A 2007;104: 347–52.
Mehring C, Hehl U, Kubo M, Diesmann M, Aertsen A. Activity dynamics and propagation of synchronous spiking in locally connected random networks. Biol Cybern 2003;88:395–408.
Miri A, Daie K, Bardine R, Aksay E, Tank D. Regression based identification of behavior encoding neurons during large-scale optical imaging of neural activity at cellular resolution. J Neurophysiol 2011;105:964–80.
Miyawaki A. Innovations in the imaging of brain functions using fluorescent proteins. Neuron 2005;48:189–99.
Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC. Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature 2005;433:597–603.
Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC. Highly ordered arrangement of single neurons in orientation pinwheels. Nature 2006;442:925–8.
Oram MW, Wiener MC, Lestienne R, Richmond BJ. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. J Neurophysiol 1999;81:3021–33.
Patnaik D, Sastry PS, Unnikrishnan KP. Inferring neuronal network connectivity from spike data: a temporal data mining approach. Sci Program 2008;16:49–77.
Pazienti A, Grün S. Robustness of the significance of spike correlation with respect to sorting errors. J Comp Neurosci 2006;21:329–42.
Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, Abeles M. Spatiotemporal structure of cortical activity: properties and behavioral relevance. J Neurophysiol 1998;79:2857–74.
Sastry PS, Unnikrishnan KP. Conditional probability based significance tests for sequential patterns in multi-neuronal spike trains. Neural Comput 2010;22:1025–59.
Schrader S, Grün S, Diesmann M, Gerstein GL. Detecting synfire chain activity with massively parallel spike train recording. J Neurophysiol 2008;100: 2165–76.
Schrader S, Morrison A, Diesmann M. A composition machine for complex movements. In: Proceeding of the 31st Göttingen neurobiol conf TS18-1C; 2007, Available online at http://www.neuro.uni-goettingen.de/archiv/2007/pdf/Proceedings-Goettingen2007.pdf.
Schrader S, Diesmann M, Morrison A. A compositionality machine realized by a hierarchic architecture of synfire chains. Front Comput Neurosci 2011;4:154.
Shmiel T, Drori R, Shmiel O, Ben-Shaul Y, Nadasdy Z, Shemesh M, et al. Neurons of the cerebral cortex exhibit precise inter-spike timing in correspondence to behavior. Proc Natl Acad Sci U S A 2005;102:18655–7.
Shmiel T, Drori R, Shmiel O, Ben-Shaul Y, Nadasdy Z, Shemesh M, et al. Temporally precise cortical firing patterns are associated with distinct action segments. J Neurophysiol 2006;96:2645–52.
Tetzlaff T, Morrison A, Timme M, Diesmann M. Heterogeneity breaks global synchrony in large networks. In: Proceeding of the 30th Göttingen neurobiol conf 206B; 2005, Available online at http://www.neuroanatomie.uni-goettingen.de/neurobio_archiv/2005/pdf/Proceedings-Goettingen2005.pdf.
Tetzlaff T, Einevol G, Diesmann M. Synchronization and rate dynamics in embedded synfire chains: effect of network heterogeneity and feedback. BMC Neuroscience 2009;10(Suppl. 1):P258, doi:10.1186/1471-2202-10-S1-P258.
Trengove C, Leeuwen CV, Diesmann M. High storage capacity of synfire chains in large-scale cortical networks of conductance-based spiking neurons. Nineteenth Annual Computational Neuroscience Meeting: CNS*2010. BMC Neuroscience 2010a;11(Suppl. 1):F1, doi:10.1186/1471-2202-11-S1-F1.
Trengove C, Leeuwen CV, Diesmann M. Random compositional networks of synfire chains dynamically self-tune to the critical state for ongoing percolation of activity. In: Front Neurosci Conference Abstract Neuroinformatics; 2010b., doi:10.3389/conf.fnins.2010.13.00041.
Trengove C, van Leeuwen C, Diesmann M. High-capacity embedding of synfire chains in a cortical network model. J Comput Neurosci 2012, in press.
Warren DJ, Fernandez E, Normann RA. High-resolution two-dimensional spatial mapping of cat striate cortex using a 100-microelectrode array. Neuroscience 2001;105:19–31.
Yosimura Y, Dantzker JL, Callaway EM. Excitatory cortical neurons form fine-scale functional networks. Nature 2005;433:868–73.