

Department of Computer Science,  
University of Otago

UNIVERSITY  
of  
OTAGO



*Te Whare Wānanga o Ōtāgo*

---

Technical Report OUCS-2013-02

**Experiments on the effect of synaptic disruption on  
polychronous group formation: detailed methods  
and results**

Authors:

**Mira Guise, Alistair Knott, Lubica Benuskova**

Department of Computer Science, University of Otago, New Zealand

Status:

This is a document in support of a manuscript submitted to ICONIP 2013.



Department of Computer Science,  
University of Otago, PO Box 56, Dunedin, Otago, New Zealand

<http://www.cs.otago.ac.nz/research/techreports.php>

# Experiments on the effect of synaptic disruption on polychronous group formation: detailed methods and results

Mira Guise\*, Alistair Knott, Lubica Benuskova

*Dept of Computer Science, University of Otago, Dunedin*

---

## Abstract

Izhikevich (2006a) claims that groups of strongly connected neurons known as polychronous neural groups (PNGs) could provide the basis for memory in the brain. Polychronous groups exist as spatio-temporal patterns of connection lengths and weights; when activated by a matching input pattern, they are able to produce a causal chain of firing events that is not synchronous but is nevertheless precisely timed and consistently reproducible.

Izhikevich et al. (2004) have previously shown that synaptic disruption produces a dramatic decrease in PNG counts, leading them to the conclusion that PNG formation depends on activity-dependent changes in synaptic plasticity. However, despite the disruption some polychronous groups remain in the network, suggesting that polychronous groups can sometimes be formed by chance arrangements of synaptic weights in the network. In the following report we reproduce and extend the work of Izhikevich et al. (2004) by delving further into the nature of PNG formation: firstly, we examine the effects of synaptic disruption on the groups that remain; we then examine the considerable inter-network variation in PNG counts over the time-course of network maturation and categorize the resulting profiles into discrete classes of behavior.

The first experiment found a highly significant decrease in PNG count following synaptic disruption (paired  $t(19) = 9.0; p < 0.001$ ), reproducing the previous report (Izhikevich et al., 2004). However, we also found a highly significant decrease in both size (paired  $t(19) = 10.7; p < 0.001$ ) and temporal length (paired  $t(19) = 10.8; p < 0.001$ ) in the remaining polychronous groups.

A further experiment employing sampling at multiple time-points found two previously unreported phenomena: firstly, there is a consistent small peak in the PNG counts immediately following initialization; secondly, comparison of the temporal profiles of multiple networks over the course of maturation shows two broad classes of network behavior: either cyclic, or an initial burst of productivity followed by a slow decline. These results are discussed in the context of a model of PNG formation involving an activity-dependent interaction between supported and adapted groups which is at the heart of PNG formation.

*Keywords:* spiking network, polychronous neural group, synaptic plasticity

---

Although many leading theories of information exchange in the brain assume an encoding scheme based on mean firing rate (Shadlen and Newsome, 1994), many neuroscientists now believe that information transfer can also be mediated by firing patterns that include a significant temporal component (Theunissen and Miller, 1995; Gerstner et al., 1997; Maass and Bishop, 2001; Stein et al., 2005). Some of these temporally-oriented schemes are based

---

\*Corresponding author at: Department of Computer Science, University of Otago, PO Box 56, Dunedin 9054, New Zealand

*Email address:* mguise@cs.otago.ac.nz (Mira Guise)

on the firing patterns of single neurons (e.g. time-to-first-spike (VanRullen et al., 2005)), while others rely on precisely timed spatio-temporal patterns of firing within groups of neurons. In one such proposal, these spatio-temporal patterns are stored as distributions of synaptic weights and lengths within strongly connected groups of neurons known as polychronous neural groups (Izhikevich, 2006a).

Polychronous neural groups (or PNGs) exist in large numbers as subgraphs within the larger graph of the network, and are proposed to provide a basis for representation and memory (Izhikevich, 2006a). Izhikevich et al. (2004) claim that their formation requires activity-dependent changes in synaptic plasticity, citing evidence showing the sensitivity of these groups to synaptic disruption. However, the evidence for such a claim is weak and is worthy of further investigation. In this report the effect of synaptic disruption is examined in more detail with the intention of providing additional insight into the mechanisms of PNG formation.

Polychronous group formation is proposed to involve an interaction within and between groups of two different types. It is therefore worthwhile to provide the reader with additional background on the nature and variety of PNGs. In a spiking neural network with random connectivity and varying connection lengths, there exist many spatio-temporal subgraphs of connected neurons that have the potential to match an input pattern with the same spatio-temporal structure. Chance arrangements of the connection lengths of some of these subgraphs provide the potential for multiple convergent patterns of firing within the neural group. In order to realize this potential, neural groups must have strong connections between group members and sufficient congruency in the lengths and weights of these connections to allow the propagation of firing activity across the group.

Neural groups that meet this additional criterion for congruent weights are defined by Martinez and Paugam-Moisy (2009) as *adapted* polychronous groups, while those with incongruent weights are known as *supported* polychronous groups. Only adapted groups have the potential to propagate firing activity into deeper layers of the network: if activated with a matching pattern, the neurons in the group are said to *polychronize* in a causal chain of firing events that is not synchronous but nevertheless defines a consistently reproducible spatio-temporal pattern of response (Izhikevich et al., 2004; Izhikevich, 2006a).

Although supported groups can occur in any randomly connected network with variable connection lengths, the presence of adapted groups requires the network be *matured* by exposure to random input and subsequent reverberant network activity. Izhikevich et al. (2004) propose that self-organization of neurons into adapted polychronous groups does not occur by chance but is

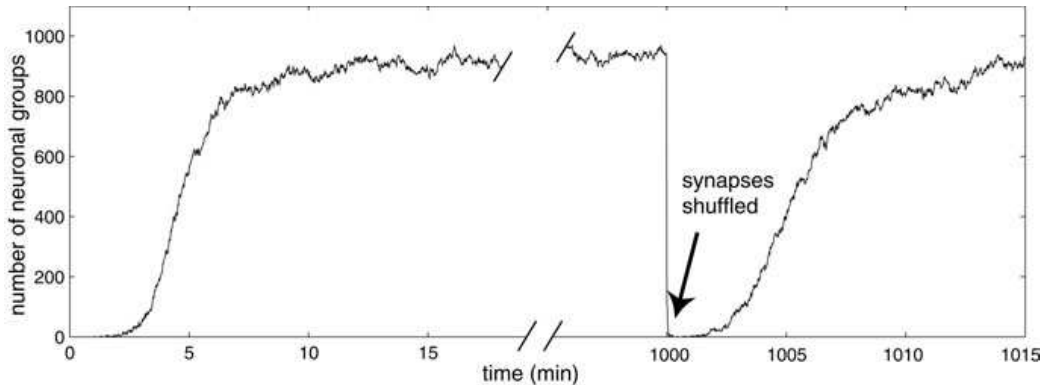


Figure 1: The number of polychronous neuronal groups that evolve in the network over the course of maturation, both before and after the synaptic weights are randomly shuffled. Maturation requires a sufficient period of exposure to network activity (produced in this example by 1 Hz random stimulation). *Reproduced from Izhikevich et al. (2004); permission granted.*

instead a consequence of activity-dependent changes in the synaptic strength of the connections between group neurons. Strengthening of these connections via spike-timing-dependent plasticity (STDP) creates congruent spatio-temporal structures in the network graph that are able to polychronize. Even in small networks these adapted polychronous groups exist in their thousands (Izhikevich, 2006a) and have the potential to selectively respond to a multitude of different input patterns.

Izhikevich et al. (2004) reasoned that if PNG formation is dependent on activity-dependent synaptic plasticity then adapted polychronous groups must be sensitive to the disruption of synaptic weights in the network. If synaptic disruption produces a significant drop in PNG numbers then the large number of groups found in mature networks are due to more than just chance arrangements of weights. Using an algorithm for discovering polychronous groups (Izhikevich, 2006b) they followed the PNG count over time as the network matured in the presence of a 1 Hz random input. This temporal profile of PNG productivity was generated both before and after synaptic disruption and is reproduced in Figure 1. With random input the network produces a rapid increase in PNG count over the first 15 minutes of maturation, followed by a prolonged period of stability. However, random reassignment of synaptic weights across the network’s connections produces a dramatic decrease in PNG counts, although the number of groups recovers rapidly with continuing maturation.

It is important to note that the algorithm used to determine the PNG counts in Figure 1 detects *adapted* neuronal groups; *supported* groups, in

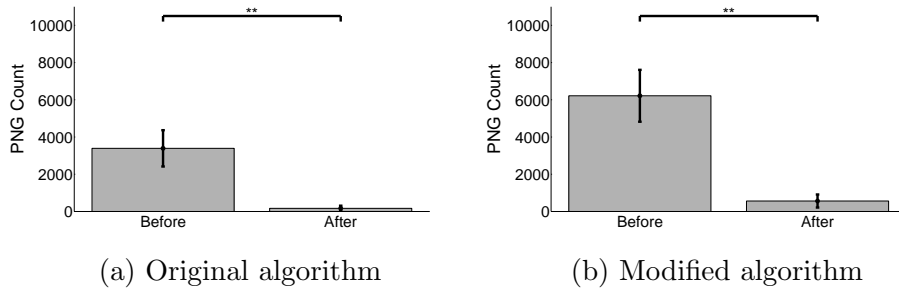
which the synaptic weights are not taken into account are of course immune to the effects of synaptic disruption. Izhikevich has designed two different PNG search algorithms which, for convenience, I have labeled the *Original* and *Modified* algorithms. Although both of these algorithms detect only adapted polychronous groups, the original algorithm can be easily adapted for discovering supported groups yielding many thousands of supported groups in a network with randomized connectivity and connection lengths (results not shown). In addition, preliminary experiments measuring the effect of disruption on adapted PNG counts suggest that a small number of adapted groups are still present in the network following synaptic disruption. With these observations in mind, a preliminary interpretation of Figure 1 is that adapted groups are generated by the effects of activity-dependent changes in synaptic plasticity on a very large reservoir of supported groups.

Given this account, a number of questions remain unanswered: What is the effect of synaptic shuffling on the adapted groups that remain? Does synaptic disruption reduce the number of neurons in the group? Does disruption affect the group's temporal length i.e. the time between the first and last firing of a group neuron? Preliminary evidence suggests that both the number of firing events in the group (the group size) and the temporal length are affected.

In addition, preliminary work suggests that there is considerable variation in the temporal profile between networks. While some networks follow the average trend shown in Figure 1, others produce a very different profile of PNG counts over time. We will therefore examine the temporal profile of PNG formation more closely and ask whether these varying profiles might be categorized into discrete classes of behavior. Each of these questions will be examined with the intention of revealing more on the nature of the interaction between supported and adapted groups which is at the heart of PNG formation.

#### *The single point shuffle experiment*

This first experiment is a simplified version of the original experiment (Izhikevich et al., 2004) that samples the counts from a mature network at a single time-point in the evolution of the network. Group counts are sampled both before and after the synaptic weights are disrupted by using a random shuffling procedure. The sampling procedure utilizes both of the PNG search algorithms provided by Izhikevich (2006b) (see the *Methods* section for details). In addition to examining the effects of shuffling on PNG count, this first experiment extends the results of Izhikevich et al. (2004) by looking for changes in the size and temporal length of the adapted groups that remain after shuffling.



Algorithm	Before	After
Original	3392	169
Current	6215	562

(c) Mean counts

Figure 2: The number of polychronous groups found in the network both before and after shuffling the synaptic weights. Counts for the original algorithm are plotted in Figure 2a and counts for the modified algorithm in Figure 2b. The corresponding mean values (averages over twenty mature networks) are shown in the accompanying table (Figure 2c). Error bars show the upper and lower confidence limits relative to the sample mean (99% confidence interval). Both the *Original* and *Modified* search algorithms show a highly significant decrease in the mean counts after shuffling: (*Original*: paired  $t(19) = 9.0, p < 0.001$ ) and (*Modified*: paired  $t(19) = 12.5, p < 0.001$ )

Performing this simple shuffling experiment on twenty independent networks produces the result shown in Figure 2. The PNG count averaged across all twenty networks shows a substantial drop after shuffling, relative to the pre-shuffle count. After shuffling, the mean PNG count for both algorithms falls to less than 10% of the pre-shuffle value.<sup>1</sup> This decrease is consistently seen across all twenty networks (individual results not shown) and for both PNG search algorithms, lending support to the idea that adapted polychronous groups are more than just random structures in the network but are instead produced by activity-dependent STDP.

Figure 3 shows the results of synaptic shuffling on the PNG size distributions. For both algorithms there is a shift towards smaller sizes following synaptic shuffling.<sup>2</sup> Comparing between algorithms, we can see that the groups found by the modified algorithm are generally smaller than those found by the original algorithm. This general pattern of results is also seen in the comparison of PNG temporal lengths before and after shuffling (see Figure 4). The polychronous groups discovered after synaptic shuffling show a shift towards shorter temporal lengths when compared to the groups found before shuffling. The modified algorithm also produces shorter temporal lengths relative to the original algorithm.

A summary of these results is shown in Figures 5 and 6. Following synaptic shuffling the mean PNG size is on average less than half of the pre-shuffling size (see Figure 5) and this size decrease is highly significant for both algorithms (paired  $t(19) = 10.7, p < 0.001$  and paired  $t(19) = 9.1, p < 0.001$ ; original and modified algorithms respectively). Mean temporal lengths in Figure 6 lose around a quarter of their value following synaptic shuffling and this decrease is also highly significant (paired  $t(19) = 10.8, p < 0.001$  and paired  $t(19) = 11.8, p < 0.001$ ; original and modified algorithms respectively).

The time-course of maturation depicted in Figure 1 suggests that the PNG count is very stable over time once the initial growth phase is complete. However, preliminary experiments have found that group counts can show considerable variation in group counts beyond this early period of growth. In addition, networks with reduced maturation time can sometimes produce higher group counts than the mature networks used in the current experiment; in one such productive network, the harvest of polychronous groups was substantially higher than any previously reported by Izhikevich (2006a)

---

<sup>1</sup>By algorithm, the post-shuffle values as a percentage of the original values are around 5% and 9% for the original and modified algorithms respectively.

<sup>2</sup>The apparent narrowing of the range of PNG sizes is likely to be an artifact of the reduced PNG numbers after shuffling.




Figure 3: Comparison of PNG size histograms before and after shuffling using both the original and modified algorithms for finding polychronous groups. The results for the original algorithm are shown in the top row, and those for the modified algorithm are shown in the bottom row. The plots on the left show the PNG size distributions before synaptic shuffling and those on the right show size distributions after shuffling. Each plot shows histogram results from the synaptic shuffling of twenty different networks, each of which had previously been matured for two hours network time. The PNG size data for each network was binned into fifty bins that were evenly distributed over the range 0 - 700, a range that includes the maximum PNG size for all networks. Note that the y-axis range differs between the before and after plots; the x-axis has the same range for all plots.



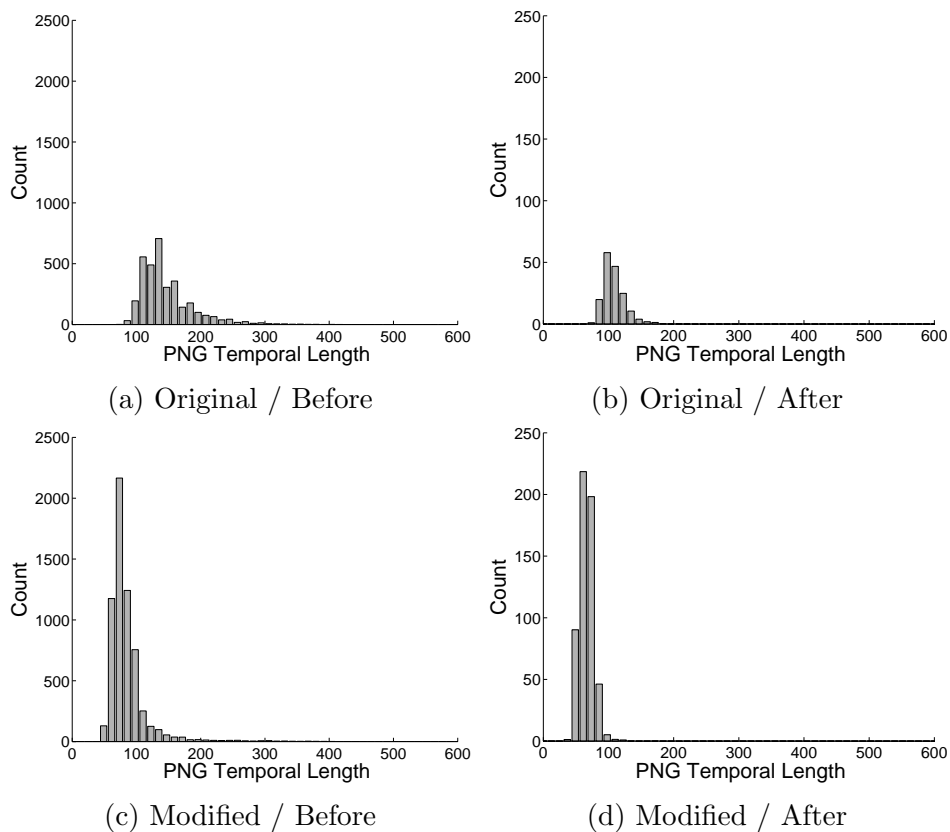


Figure 4: Comparison of temporal length histograms of discovered polychronous groups both before and after shuffling. The temporal length of a PNG is the period between the first PNG firing event and the last PNG firing event. The results for the original algorithm are shown in the top row, and those for the modified algorithm are shown in the bottom row. The plots on the left show the temporal length distributions before synaptic shuffling and those on the right show length distributions after shuffling. Histogram results from the synaptic shuffling of twenty different networks are shown. These are the same networks as shown in Figure 3. Length data for each network was binned into fifty bins that were evenly distributed over the range 0 - 600. Note that the y-axis range differs between the before and after plots; the x-axis has the same range for all plots.




Figure 5: Comparison of the mean PNG size before and after synaptic shuffling (averaged over twenty independent networks). Shuffling produces a consistent decrease in PNG size for both the *Original* and *Modified* search algorithms. Group sizes for the original algorithm are plotted in Figure 5a and sizes for the modified algorithm in Figure 5b. Error bars show the upper and lower confidence limits relative to the sample mean (99% confidence interval). The decrease in group size is highly significant: (*Original*: paired  $t(19) = 10.7, p < 0.001$  with mean values of before = 83 and after = 36) and (*Modified*: paired  $t(19) = 9.1, p < 0.001$  with mean values of before = 39 and after = 17)




Figure 6: Comparison of mean temporal length before and after synaptic shuffling. Shuffling produces a consistent decrease in temporal length for both algorithms. The temporal lengths for the original algorithm are plotted in Figure 6a and lengths for the modified algorithm in Figure 6b. Error bars show the upper and lower confidence limits relative to the sample mean (99% confidence interval). The decrease in temporal length is highly significant: (*Original*: paired  $t(19) = 10.8, p < 0.001$  with mean values of before = 143 and after = 104) and (*Modified*: paired  $t(19) = 11.8, p < 0.001$  with mean values of before = 85 and after = 65)

(results not shown). The next experiment therefore follows Izhikevich et al. (2004) in exploring a version of the experiment that looks at variation in group counts over time, both before and after synaptic shuffling.

#### *The multi-point shuffle experiment*

Group counts vary over the course of maturation, with a strong initial increase in early maturation followed by a plateau in the PNG numbers at each time-point (Izhikevich et al., 2004). Over this time-course, the distribution of excitatory synaptic weights becomes increasingly bimodal (see Fig. 6 of Izhikevich et al., 2004), with many weights either decreasing to zero or increasing to saturation. The bimodal weight distribution of mature networks reflects the influence of STDP, but may also be a direct result of the STDP-mediated formation of polychronous groups.

We can see in Figure 1 that the time-course in early maturation appears to be recapitulated after synaptic disruption. It appears that synaptic shuffling causes the network to regress to an immature state and the growth in group counts following shuffling is really just a re-growth. This might be as expected if the synaptic weights at the start of the pre-shuffle time-course and at the start of the post-shuffle time-course were equivalent. However, the pre-shuffle time-course begins with the synaptic weights being initialized to equal values. In contrast, the weights at the start of the post-shuffle time-course are produced by randomly swapping the weight values in a mature network. Because the shuffling procedure does not modify the weight values, the synaptic weights at the beginning of the post-shuffle time-course continue to maintain the strongly bimodal distribution of a mature network.

A more detailed examination of this re-growth in group numbers following disruption is afforded by the next experiment which samples PNG numbers at regular intervals over the course of maturation. Samples are taken at one minute intervals over a two hour period immediately before and immediately after the synaptic shuffling procedure. Each of the ten runs of the experiment is initialized with a unique starting network, allowing individual differences in network productivity to be exposed. For details of this experiment see the *Methods* section.

There are some clear differences between the results of this experiment (shown in Figure 7) and those previously described by Izhikevich et al. (2004, see Figure 1). Some previously unobserved features include an initial small peak in the PNG count, and the cyclic productivity of some of the networks. The initial peak occurs within three to six minutes following initialization, both before and after shuffling, and is consistent across all ten runs in Fig-




Figure 7: Emergence of polychronous groups as a function of time, both before and after synaptic shuffling. Ten runs of the experiment are shown, with a different starting network for each run. Individual runs plot 240 time points taken at intervals of one minute during maturation (1 - 120 minutes) and following synaptic shuffling (121-240 minutes). The vertical arrow marks the time at which synaptic shuffling occurred in the middle of each run. A line through the mean values over all ten runs is plotted in bold. Note the high variability between runs.

ure 7.<sup>3</sup> Also notable is that there are two classes of network behavior in Figure 7: the first class follows the average trend (bold line) with an initial burst of productivity followed by a gradual decline. However, some runs manifest a cyclic behavior in which productivity waxes and wanes with a regular period of thirty to thirty-five minutes. Of the ten runs, four in the maturing phase, and three in the shuffled phase show this behavior, although only two runs manifest this cycling behavior both before and after shuffling.

The average trend (bold line in Figure 7) also differs from that of Izhikevich et al. in that the initial strong surge in the PNG count is followed by a slow decline. In addition, the number of adapted groups was reduced to zero by synaptic shuffling in Figure 1, but remains non-zero (although greatly reduced) in the current experiment (PNG counts ranged from 17 to 233 over the ten runs).<sup>4</sup> However, although there are strong differences between the results of Izhikevich et al. and those shown in Figure 7, both sets of results do show similar behavior during the initial growth period. Izhikevich et al. observed a strong initial growth in PNG numbers, both within the first fifteen minutes and in the fifteen minutes immediately following synaptic shuffling. Similarly, the results in Figure 7 show a rapid initial growth in group counts in the first thirty minutes, both before and after synaptic shuffling.

The results in Figure 7 also highlight the high variability in productivity between networks, with some networks producing more than twice as many polychronous groups as others. However, synaptic disruption has the same effect on all networks, producing an immediate disappearance of most polychronous groups followed by a full recovery within thirty minutes.

## Discussion

Supported polychronous groups can be found in large numbers in a network with randomized connectivity and varying connection lengths. Izhikevich et al. (2004) maintain that *adapted* polychronous groups (i.e. those capable of polychronization) are created by the activity-dependent synaptic strengthening of these existing structures under the influence of STDP. Of course, these adapted polychronous groups could also occur in the network by chance. If they are purely chance arrangements of synaptic weights, then random shuffling of the synaptic weights should have no impact on the num-

---

<sup>3</sup>The PNG counts for this initial peak range from around 400 to 1200.

<sup>4</sup>One explanation for this difference is that the original version of the experiment may have been using an incorrect shuffling technique. The original C++ shuffling code (Izhikevich, 2006b) copies memory locations rather than swapping them. This was assumed to be unintentional and was fixed in the modified code library (see Methods).

ber of groups that are discoverable in the network. However, the results from both experiments in the current report echo the results of Izhikevich et al. (2004) in showing that group counts are substantially affected by synaptic shuffling. Therefore these strongly connected subgraphs are not just random features of the network.

The single point version of the shuffle experiment examined the effects of synaptic disruption at a single time-point, focusing in particular on the effect of shuffling on the remaining adapted groups. In contrast to the results of Izhikevich et al. (2004), not all groups disappeared following synaptic shuffling, although only a small proportion of the original groups remain. For these remaining groups, synaptic shuffling produced a significant reduction in the number of neurons that participate in each group, and a smaller drop in the mean temporal length of the remaining groups.

The fact that polychronous groups are still discoverable in the network after the shuffling procedure shows that these strongly connected subgraphs *can* occur in the network by chance. However, the effect of shuffling on the size, length, and particularly the overall numbers of polychronous groups in the network strongly argues that adapted PNG formation is more than just a chance arrangement of synaptic weights. Other evidence supports this conclusion, such as the substantial difference between the activation frequency distributions of polychronizing spike trains relative to inverted time versions of the same spike trains (Izhikevich et al., 2004). Izhikevich et al. make the stronger claim that these groups rely on activity-dependent synaptic plasticity for their formation. When viewed through this lens, the results shown here suggest that the effect of activity-dependent plasticity is to increase the size and overall numbers of adapted polychronous groups that are formed from a pre-existing pool of supported groups.

The results from the original multi-point experiment (Izhikevich et al., 2004) show a strong initial growth in the first fifteen minutes that is similar to that seen in the current experiment. However, the average trend in Figure 7 (bold line) more closely resembles the results of Chrol-Cannon et al. (2012, see Fig. 3a) who found an initial peak at around thirty minutes followed by a lower level plateau at around fifty minutes. Chrol-Cannon et al. found that increasing the stimulation intensity produced a flattening of the initial peak and an overall reduction in PNG count. They were also able to replicate the plateau in PNG counts observed in Izhikevich et al. (2004) by using a modified STDP rule (see Chrol-Cannon et al., 2012, Fig. 3b). Although the input protocol differed between the current experiment and that of Chrol-Cannon et al., the overall rate of stimulation between the lowest intensity experiment of Chrol-Cannon et al. and the current experiment was the same (1000 firing events per second) and both experiments used largely identical

network parameters following Izhikevich (2006a). In contrast, the original shuffle experiment of Izhikevich et al. (2004) was carried out on a considerably larger network (100,000 neurons) with both short-term and long-term plasticity mechanisms.

Chrol-Cannon et al. (2012) did not explore the variation in temporal profiles between networks and therefore did not observe two previously unreported features that are apparent in Figure 7. The first of these is an initial small peak in the group count that is consistently observable in the PNG counts following initialization of the maturation procedure, and also immediately after synaptic shuffling. The second new feature is the appearance of two broad classes of network behavior that can be categorized as follows: a *cyclic* class in which network productivity waxes and wanes with a regular period of thirty to thirty-five minutes, and a *default* class that more closely follows the average trend across networks i.e. an initial burst of productivity followed by a gradual decline.

Both of these network features may derive from the same underlying mechanism. One possible explanation is that these networks have a feedback mechanism that produces an oscillation in the balance of excitation versus inhibition, leading to cycles of group building followed by group pruning. According to this proposal, an increase in group number means that more polychronous groups are able to respond to a given stimulus and the level of excitation therefore increases. Building levels of excitation produce further increases in group building, leading to more responding groups and mounting excitation. However, many of the firing paths in a polychronous group terminate in an inhibitory neuron. The building levels of excitation therefore produce a mounting inhibitory drive that counters the excitation. As the level of inhibitory firing approaches a threshold, neurons in each polychronous group increasingly begin firing out of step due to the inhibition. Because many of these neurons participate in multiple groups, competition between groups produces pruning of the most inhibited groups leading to a slow downward trend in the PNG count. In keeping with this idea, Izhikevich et al. (2004) propose that PNG formation is a dynamic process: as input patterns come and go, the synaptic weights are in constant flux as they conform to one or the other spatio-temporal structures within the network.

The method used here (and previously, see Izhikevich et al., 2004; Izhikevich, 2006a; Chrol-Cannon et al., 2012) for counting polychronous groups utilizes the PNG search algorithms to identify structural groups in the network. These algorithms rely on the potential for structural groups to be activated in response to specific network inputs and entail an exhaustive search for group-activating inputs from within a selected subset of the possible network

inputs.<sup>5</sup> Group activation requires congruent synaptic weights that support the propagation of spatio-temporal firing patterns and is therefore sensitive to synaptic disruption.<sup>6</sup> The reduction in the PNG count returned by the PNG search algorithms following synaptic shuffling is therefore reflective of a reduction in the potential for group activation due to loss of congruency in the synaptic weights.

The overall results in these two experiments serve to demonstrate that polychronous groups are more than just random structures in the network. The dependency of these structures on a precise distribution of synaptic weights suggested to Izhikevich et al. (2004) that their formation is a result of activity-dependent changes in synaptic plasticity. The results presented here, while not providing definitive support, are nevertheless not inconsistent with this proposal. However, the current report additionally provides evidence in favor of a model of PNG formation that emphasizes the dynamic nature of the interaction between supported and adapted groups. In this model, the formation of adapted groups occurs through changes in synaptic plasticity governed by the selective reinforcement of supported groups that match the current inputs.

Izhikevich et al. (2004) has proposed that synaptic weights constantly fluctuate as individual neurons dynamically change the strength of their affiliations to the different neural groups in which they participate. However, this dynamic process also occurs at group level as a bidirectional flux between neural groups that exist as mere spatio-temporal structures and those capable of polychronization. The effect of synaptic shuffling is to disrupt the congruency between the lengths and weights of group connections that allows the neural groups to polychronize, and thus reduce the length of the paths capable of propagating firing activity.

For some networks, the presence of random input allows a stable equilibrium between supported and adapted groups, while for others the balance of excitation and inhibition undergoes repeated cycles, mediated by the number of responding groups. It is hoped that future work will employ a newly developed technique to more closely examine the dynamic nature of PNG formation.

---

<sup>5</sup>For performance reasons, the search domain is limited to input patterns composed of spatio-temporal triplets i.e. any combination of three neurons.

<sup>6</sup>The reliance on group activation in the modified algorithm means that some structural groups will not be counted if the algorithm is unable to find an input pattern that activates the group. An implicit assumption in this group counting method is that the proportion of unreported groups is independent of the effects of synaptic shuffling.



## Methods

Both experiments examined the effect of synaptic disruption on the PNGs produced from multiple independent networks using an STDP learning rule. Each network was composed of 1000 Izhikevich neurons with parameters as described in Izhikevich (2006a). Networks were matured by exposure to 1 Hz random input generated by a Poisson process. For the *Single Point Shuffle Experiment*, twenty networks were matured for two hours and PNG counts from each network were sampled at a single point in time. For the *Multi-point Shuffle Experiment*, ten networks were sampled at multiple intervals over the course of network maturation. The sampling procedure employs the PNG search algorithms of Izhikevich et al. (2004) to return a count of the number of polychronous groups that are discoverable in the network.

### *PNG definition*

The definition of a PNG in these experiments adheres to that employed by Izhikevich et al. (2004). This group defines structural PNGs in terms of both network structure and synaptic weights, what Martinez and Paugam-Moisy (2009) call an *Adapted Polychronous Group*. Martinez and Paugam-Moisy define an additional structural PNG type called a *Supported Polychronous Group*; however, this type is defined solely on the basis of network topology without reference to synaptic weights and is therefore not sensitive to synaptic disruption.

### *PNG sampling method*

Izhikevich has provided two versions of the PNG search algorithms which for the purposes of this report are labeled as the *Original* and *Modified* versions. The original fast algorithm was replaced with a slower but more precise version that appears to be the new de facto standard for finding PNGs. Both original and modified versions of these algorithms scan the structural groups in the network for those with synaptic weights that support the propagation of spatio-temporal firing patterns beyond a minimum length. A count of the number of polychronous groups that are discovered in the network can be made both before and after synaptic shuffling. For more details examine the code and documentation for either the original reference implementation (Izhikevich, 2006b), or for the Spinula software package (Guise, 2013).

### *Synaptic disruption*

The procedure for producing synaptic disruption in the network follows Izhikevich et al. (2004) in using a shuffling technique: for each connection between excitatory neurons, another connection between excitatory neurons

is randomly chosen and the synaptic weights are swapped (for more details see the code described in Guise, 2013).

### *Scripts*

The methods in this paper are explicitly described via a series of scripts that allow the reproduction of each experiment in detail. Execution of these scripts requires the installation of the Spinula software package (Guise, 2013), a small suite of dynamic link libraries that provide functions for network construction, execution and analysis (*SpikingNeuronLib*, *SpikingAnalyticsLib*, and *SpikingVisualisationLib*). Although some scripts reference the *SpikingNeuronLib* library directly, most scripts reference a wrapping library (*SpikingAnalyticsLib*) that provides a range of functions relevant to analysis of spiking data. A more detailed description can be found in the Spinula technical report (see Guise, 2013).

### *The simple shuffle experiment*

The script that conducts this experiment can be seen in Listing 0.1. Script parameters define the location of a saved state file for a mature network that is the subject of our experiment. An additional parameter specifies which of the two PNG search algorithms is utilized in a particular experimental run, providing an opportunity to compare results from these two algorithms. Note that the reported results show averages over twenty independent networks (each 1000-neuron Izhikevich network was matured for two hours prior to the experiment).

The script begins on lines 10 and 11 by generating a network specification for a 1000 neuron Izhikevich network with 100,000 synapses, and with a 4:1 ratio of excitatory versus inhibitory cells. Lines 14 and 15 produce a PNG count of the selected mature network prior to synaptic shuffling. Internally, this function uses the network specification to create a new network, loads the network state data from the selected state file, and executes the selected PNG search algorithm. For each discovered PNG a number of attributes are saved to a file (*before.txt*) in the source folder: these include the anchor neuron, the group size, and the temporal length. The same network state file is then reloaded in lines 18 and 19, and the synaptic weights in the loaded network are shuffled. The selected PNG search algorithm is once again unleashed in lines 22 and 23, and any PNGs found are saved to a file (*after.txt*).

### *The multi-point experiment*

This experiment reproduces the original shuffling experiment of Izhikevich et al. (2004) and proceeds with the following general outline: PNG counts are sampled at regular intervals over the course of network maturation; the

```

1 // Compare the PNG count both before and after shuffling
2 // the synaptic weights of a mature network
3 let ComparePNGCountBeforeAfterShuffling pathToStateFolder stateFileName algorithmVersion =
4
5     let numExcitatoryNeurons = 800
6     let numInhibitoryNeurons = 200
7     let numSynapsesPerNeuron = 100
8     let maxDelay = 20
9
10    let networkSpecifier = new IzhikevichNetworkSpecifier(
11        numExcitatoryNeurons, numInhibitoryNeurons, numSynapsesPerNeuron, maxDelay)
12
13    // get count of PNGs in the network state file before shuffling; save PNG details to a file
14    let beforeCount = Span.GetPNGCountsInStateFile(networkSpecifier,
15        pathToStateFolder, stateFileName, "before.txt", algorithmVersion)
16
17    // load the same network state file and shuffle the network
18    let shuffledNetwork = Span.ScrambleNetworkInStateFile(networkSpecifier,
19        pathToStateFolder, stateFileName)
20
21    // get count of PNGs in the shuffled network; save PNG details to a file
22    let afterCount = Span.GetPNGCountsInNetwork(
23        shuffledNetwork, pathToStateFolder, "after.txt", algorithmVersion)
24
25    printfn "%d %d" beforeCount afterCount

```

Listing 0.1

synaptic weights in the network are then disrupted by synaptic shuffling; finally, maturation is continued with regular sampling as before. Network maturation in the original experiment consisted of running the network for 24 hours network time in the presence of 1 Hz random stimulation. In the current experiment the maturation period was limited to just 120 minutes prior to synaptic disruption.<sup>7</sup>

Rather than performing the PNG counts directly, the experiment is split into two phases: in phase 1 the network state is saved at regular intervals; in phase 2, the PNG counts of each saved state are determined. Phase 1 generates a folder of network state files that sample the network state over time, both before and after synaptic shuffling. Phase 2 consecutively runs through these state files and counts the number of groups in each file. The result is a temporal profile of PNG counts before and after shuffling. To allow for variation in network productivity, the procedure is executed multiple times and the resulting profiles are overlaid on a single plot.

The multi-point version of the shuffle experiment utilizes two scripts, one for each phase: the first generates the folder of state files, and the second counts the number of PNGs in each file. The first script is shown in Listing 0.2. The script creates a 1000 neuron Izhikevich network and runs the network for two hours, saving state at one minute intervals. The network is then shuffled and network execution continues for a further two hours,

---

<sup>7</sup>Earlier work in which the maturation period was continued for 1000 minutes found that counts plateau beyond around 120 minutes.

with regular saving of the network state. The single argument to this script contains the path to the output folder where the state files, both pre- and post-shuffling, are saved.

The script begins by creating a network specifier (lines 13 and 14) and running the network to generate maturation files (lines 16 and 17). Note that the state files saved prior to shuffling have filenames of the form “maturingN.txt”, where N denotes the maturation age of the network in minutes (the corresponding form of the post-shuffle state files is “shuffledN.txt”). Lines 20 and 21 generate the name of the most mature (two hour) state file by substituting the value *120* in the filename. This file is then loaded and shuffled, and maturation continues (lines 24 - 26).

```

1 // Create a folder of network state files saved at regular intervals as a network matures
2 // both before and after shuffling the network state
3 let GenerateMaturationFiles pathToOutputFolder =
4
5     let numExcitatoryNeurons = 800
6     let numInhibitoryNeurons = 200
7     let numSynapsesPerNeuron = 100
8     let maxDelay = 20
9     let backgroundFrequency = 1
10    let runSeconds = 3600 * 2 // two hours or 120 minutes
11
12    // generate maturation state files before shuffling
13    let networkSpecifier = new IzhikevichNetworkSpecifier(
14        numExcitatoryNeurons, numInhibitoryNeurons, numSynapsesPerNeuron, maxDelay)
15
16    Span.GenerateMaturationFiles(networkSpecifier,
17        runSeconds, backgroundFrequency, pathToOutputFolder, "maturing")
18
19    // last state file in the "maturing" group
20    let runMinutes = runSeconds / 60
21    let maturedStateFileName = String.Format("maturing{0}.txt", runMinutes)
22
23    // shuffle the most mature state and continue generating maturation state files
24    Span.GenerateMaturationFilesFromScrambledState(networkSpecifier,
25        pathToOutputFolder, maturedStateFileName,
26        runSeconds, backgroundFrequency, pathToOutputFolder, "shuffled")

```

Listing 0.2

The second script (see Listing 0.3) counts the number of PNGs in each saved state file. It begins in lines 10 and 11 by defining a network specifier compatible with the saved state files resulting from execution of the previous script. Lines 14 and 15 consecutively load each state file in the *maturing* group and save the PNG counts to a file (*maturingCounts.txt*). Lines 18 and 19 perform the same function on the *shuffled* state files and save to a separate file (*shuffledCounts.txt*). Note that this script uses the original (faster) algorithm for finding PNGs.<sup>8</sup>

---

<sup>8</sup>Using the faster algorithm on a standard PC (Intel Core 2 Duo @ 2.66 GHz with 4GB RAM), the PNG count takes between 3 and 11 minutes (average 7 minutes) for each file, depending on the number of groups in the network. The total script execution time for each run (2 × 120 files) is therefore around 27 hours.

```

1 // Count the number of PNGs in each "maturing" and "shuffled" state file
2 let CountPolychronousGroupsInStateFiles pathToOutputFolder =
3
4     let numExcitatoryNeurons = 800
5     let numInhibitoryNeurons = 200
6     let numSynapsesPerNeuron = 100
7     let maxDelay = 20
8     let algorithmVersion = FindGroupsVersion.Original // original Izhikevich algorithm
9
10    let networkSpecifier = new IzhikevichNetworkSpecifier(
11        numExcitatoryNeurons, numInhibitoryNeurons, numSynapsesPerNeuron, maxDelay)
12
13    // count the "maturing" state files
14    Span.GetPNGCountsPerStateFile(networkSpecifier,
15        pathToOutputFolder, "maturingCounts.txt", "maturing", algorithmVersion)
16
17    // count the "shuffled" state files
18    Span.GetPNGCountsPerStateFile(networkSpecifier,
19        pathToOutputFolder, "shuffledCounts.txt", "shuffled", algorithmVersion)

```

Listing 0.3

## References

- Chrol-Cannon, J., Gruning, A., Jin, Y., Jun. 2012. The emergence of polychronous groups under varying input patterns, plasticity rules and network connectivities. In: The 2012 International Joint Conference on Neural Networks (IJCNN). IEEE, pp. 1–6.  
 URL [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=6252828](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=6252828)  
<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6252828>
- Gerstner, W., Kreiter, A. K., Markram, H., Herz, A. V., 1997. Neural codes: firing rates and beyond. *Proceedings of the National Academy of Sciences* 94 (24), 12740–12741.
- Guise, M., 2013. Spinula: software for simulation and analysis of spiking network models. Tech. rep., Dept of Computer Science, University of Otago, Dunedin.
- Izhikevich, E. M., Feb. 2006a. Polychronization: computation with spikes. *Neural computation* 18 (2), 245–82.  
 URL <http://www.ncbi.nlm.nih.gov/pubmed/16378515>
- Izhikevich, E. M., 2006b. Reference software implementation for the Izhikevich model: minimal spiking network that can polychronize.  
 URL <http://www.izhikevich.org/publications/spnet.htm>
- Izhikevich, E. M., Gally, J. a., Edelman, G. M., Aug. 2004. Spike-timing dynamics of neuronal groups. *Cerebral cortex* (New York, N.Y. : 1991) 14 (8), 933–44.  
 URL <http://www.ncbi.nlm.nih.gov/pubmed/15142958>

- Maass, W., Bishop, C. M., 2001. Pulsed neural networks. The MIT Press.
- Martinez, R., Paugam-Moisy, H., 2009. Algorithms for structural and dynamical polychronous groups detection. *Artificial Neural Networks ICANN 2009* 5769, 75–84.
- Shadlen, M. N., Newsome, W. T., 1994. Noise, neural codes and cortical organization. *Current opinion in neurobiology* 4 (4), 569–579.
- Stein, R. B., Gossen, E. R., Jones, K. E., 2005. Neuronal variability: noise or part of the signal? *Nature Reviews Neuroscience* 6 (5), 389–397.
- Theunissen, F., Miller, J. P., 1995. Temporal encoding in nervous systems: a rigorous definition. *Journal of computational neuroscience* 2 (2), 149–162.
- VanRullen, R., Guyonneau, R., Thorpe, S. J., 2005. Spike times make sense. *Trends in neurosciences* 28 (1), 1–4.