DARPA HR0011-16-1-0006, Superior AI Phase 2
Quarterly Report - Q1, 9/12/2017 - 12/31/2017

Supplementary technical Report Task 3:
Graph/Network structures & processes supporting energy aware computing and memories

---

1. Design Principles of Network Architectures with the Capability of Serving as Massive Dynamical Memories

We study networks architectures with recurrent (feedback) connections and the propagation of activity patterns over these networks. We build on our results in Phase 1, where we identified lattice architectures with feedback connections between excitatory and inhibitory units producing spatio-temporal oscillations of varying complexity. We observe critical behaviors and phase transitions between fixed points when oscillations die out after some number of iterations; limit cycles when spatial patterns are repeated after a fixed number of time steps; and chaos-like dynamics when the spatial patterns keep on changing without repetitions within the observed time window, which may extend up to millions of time steps. We studied design principles of architectures that give rise to very long cycles and derived conditions to control the length of the oscillations exhibiting metastable spatial activity patterns. We demonstrate how input data can be stored and consecutively retrieved, which is the illustration of the capability to use these networks as associative memories.

We characterize the oscillations using phase portraits with time-delayed embedding. We evaluated phase portraits of locally smoothed activity patterns using smoothing window starting from small window limited to the direct neighbors of a pixel, and increase it until the window covers the whole lattice (mean-field). To visualize the complex dynamics in the case of very long cycles and when no cycles were observed cycles (chaos-like dynamics), we determined the Poincare section and Poincare maps. We also draw recurrence plots that exhibited high mixing properties with trajectories densely filling in the state space. The results showed behavior resembling strange attractors, possibly chaotic dynamics, although the presence of chaos cannot be shown rigorously under the given experimental conditions. These results show the potential of the dynamic regimes to maintain quasi-periodic oscillations that are key ingredients of the implementation of powerful dynamical memories.

2. Spiking Processes on a Neural Network with Excitatory and Inhibitory Nodes

2.1. **Introductory Remarks.** Previously we introduced a neural network $N(E, I)$. The excitatory layer of the network E has nodes and connections that are defined by the vertices and edges of the graph $G_{\mathbb{Z}^2_N, p_d}$ respectively. The inhibitory layer I has $\frac{N^2}{4}$ nodes that are connected in an all to all fashion and can be thought of as a fully connected graph $K_{\frac{N^2}{4}}$ whose vertices are the nodes. Moreover, each inhibitory

node is connected to four excitatory nodes at random in such a way that no two inhibitory nodes share any excitatory neighbors. Due to the convenient parallel with graph theory we occasionally borrow notation, and so by $V(E)$, $V(I)$ we mean the set of excitatory and inhibitory nodes, etc.

Each node in both E and I can take on one of two states: active or inactive. Let $\chi_v(t)$ define the potential function for node v in either layer at time t such that $\chi_v(t) = 1$ if v is active and $\chi_v(t) = 0$ if v is inactive at time t. The state of a node is completely determined at every time step by the state of its neighbors. To define this more formally, let $A^E(t)$ denote the set of active vertices in E at time t and similarly $A^I(t)$ denote the set of active inhibitors at time t. Furthermore, define $A^E(0)$ as a random subset of excitatory nodes that became active with probability p independently of all others and $A^I(0) = \emptyset$. Then for a vertex in E we say its state at time $t+1$ is

$$\chi_v(t+1) = \mathbb{1}\left(\sum_{u \in N(v) \cap V(E)} \chi_u(t) \geq k\right)$$

Similarly, for a vertex in I we have

$$\chi_v(t+1) = \mathbb{1}\left(\sum_{u \in N(v) \cap V(E)} \chi_u(t) \geq \ell\right)$$

In both cases $\mathbb{1}$ is the indicator function and $N(v)$ denotes the subset of nodes in the closed neighborhood of v (i.e the node v and its neighbors). Both k and ℓ are nonnegative integers that specify the number of active neighbors any given vertex needs to become active on the next time step in E and I respectively.

We conclude our review of $N(E, I)$ by presenting the inhibitory firing function we had introduced that causes all of the inhibitors to fire together once $m \in [0, \frac{N^2}{4}]$ inhibitory vertices are active during a time step. In other words an inhibitory node $v \in V(I)$ fires at time $t+1$ if

$$F_v(t+1) = \mathbb{1}\left(\sum_{u \in N(v);\ u,v \in V(I)} \chi_u(t) \geq m\right)$$

but v did not fire at time t. Notice that active inhibitory nodes fire simultaneously since they are in all to all connection with each other. At the time of firing, the inhibitory node sets the activity of all excitatory nodes connected to it and itself to 0. That is to say, in a firing step the following nodes become inactive: (i) all inhibitory nodes, and (ii) those excitatory nodes which were connected to an active inhibitory node that was firing at that step. After the inhibitory firing occurs both layers carry on by propagating activity (or the lack thereof) with whichever excitatory nodes were left in tact.

Before this network architecture, we had also shown that *assuming there is no inhibition*, for any $\lambda \geq 0$ with the $k = 2$ activation rule, E has an asymptotic critical probability $p_c = 0$. Therefore there always exists a system large enough such that for any choice of $p_{in} > 0$, regardless of how arbitrarily small, with high probability (whp) all of the vertices in E will eventually become active when the inhibitors are not allowed to fire. When we introduce the inhibitory firing rule, with a proper choice of parameters, periodic behaviour became a possibility. By

choosing parameters in such a way that the firing process extinguishes vertices while maintaining $\frac{|A^E|}{N^2} > p_c$ in the specified network, the activation in both layers likely begins growing once more until the next firing. For the sake of specificity

Definition 1. *Let Δt be a lapse in time such that there exists finite t where $|A^I(t)| = 0$, $|A^I(t+\Delta t)| = 0$, and $\nexists \Delta t' \in [0, \Delta t)$ such that $|A^I(t+\Delta t')| = 0$. We will call the sequence $\{A^E(t), A^E(t+1), \ldots, A^E(t+\Delta t)\}$ a **spike** of length Δt.*

As we saw in previous reports this the ability to produce spikes indicates oscillatory behaviour in this dynamical system for a certain parameters. We defined $p_c^{low}(\lambda, \ell, N)$ and $p_c^{high}(\lambda, \ell, N)$ as lower and upper critical probabilities respectively. It was proved in previous reports that these values exist and when the initial activation $p_{in} < p_c^{low}$ or $p_{in} > p_c^{high}$ then the system cannot exhibit oscillatory dynamics. These dynamics were also shown to impossible when we choose ℓ and m such that $\ell m > p_c^{high}(\lambda, \ell, N)$. Furthermore, we showed that $N(E, I)$ has an intriguing property we dubbed "memorylessness"; when the system is inside of the oscillatory region its periodic behavior is does not depend on initialization density.

3. Oscillatory Regimes and Chaos in $N(E, I)$

3.1. **Definitions & Recapitulation of Findings.** Our computer simulations revealed that, within the oscillatory region of parameter space, as m increased the length of the periods of $N(E, I)$ did too until $m \approx 1750$ at which point the system underwent a phase transition and began acting in a way that can be described as chaotic. Though strictly speaking chaotic regimes are impossible in any finite sized network, the behaviors displayed by parametrizations with $m \geq 1750$ can take so long ($> 10,000,000$ time steps) to repeat themselves that by all practical standards they are chaotic and behave chaotically. With this in mind

Definition 2. *Suppose there exists mutually unique spikes $\{S_1, S_2, \ldots, S_n\}$ of length $\{\Delta t_1, \Delta t_2, \ldots, \Delta t_n\}$ such that for some $t \in [0, \infty)$, $S_k = \{A^E(t + \sum_{j=1}^{k-1} \Delta t_j + y(\sum_{j=1}^{n} \Delta t_j)), \ldots, A^E(t + \sum_{j=1}^{k} \Delta t_j + y(\sum_{j=1}^{n} \Delta t_j))\}$ for all integers $y \geq 0$ and $k \in [1, n]$. We say our system is **n-periodic** if such a sequence of spikes exists and if there doesn't exist a sequence of less than n spikes that satisfy this property.*

We found that for a majority of parameters, n is finite and relatively small but, when m becomes larger than ≈ 1750, for practicality's sake we must define

Definition 3. *We say our system is **chaotic** if for some very large T, there does not exist a time $t \leq T$ for which our system is $n-$periodic.*

During the previous report we began a thorough analysis of this behavior when $\lambda = 0$. A method for counting the number of chaotic parametrizations was presented (given that the initialization is seeded). By steadily increasing the time for which we ran the system, we began seeing the rate at which the number of chaotic oscillations fell into a limit cycle. Through this we showed that, if the trend continues, we will have chaotic values of m until we run the system for over a hundred trillion time steps; this is strong evidence that these regimes are not just some transient. In this report we continue our analysis of these chaotic regimes.

In the discussion that follows, the reader should assume $N = 100$, $k = 2$, and

$\ell = 4$ unless it is explicitly stated otherwise. Also it should be assumed that any figures shown were created with an initialization dictated by a random seed.

3.2. **Continuation of Phase Space Analysis.** As we have witnessed in previous reports, the dimensionality of $N(E, I)$ grows incredibly fast with the size, N, of the network. This has made it difficult for us to create convincing visualizations of the system. For this reason we introduced the notion of a time-delay embedding for attractor reconstruction. The intuition behind this technique is that $N(E, I)$ is too high dimensional for us to both capture and visualize everything at once. On the other hand, we can create a time series $\{|A^E(t)| : t \in T\}$ of some telling summary statistic of the system's state during the time interval T. By creating the vector $[|A^E(t)|, |A^E(t+\tau)|, |A^E(t+2\tau)|, \ldots, |A^E(t+n\tau)|]$ we are mapping this time series into a higher dimensional space and effectively "unfolding" the series' trajectory by carefully choosing τ and n. Therefore, by taking this embedding we have summarized the complicated original system into a one dimensional time series, then we have taken the time series and lifted it into a higher dimension, and finally we hope this lifted version of the series is representative of the original system's movement through phase space while also easy to visualize. Due to the fact that $N(E, I)$ does not travel through phase space in a smooth nor continuous way, the guarantees normally associated with time-delay embeddings do not apply. Regardless this technique can still provide insights into $N(E, I)$'s dynamics with a careful choice of τ. We saw evidence of this in the previous report using the time series $\{|A^E(t)| : t \in T\}$ where T was the interval $[0, 10000000]$, but our visualizations were limited by the fact that it took meticulous observation to notice the difference between a long period and a chaotic regime.

Taking inspiration from the shortcomings in this application of time-delay embeddings, we first embedded $\{|A^E(t)| : t \in [950000, 10000000]\}$ into three dimensional space and treated the resultant vectors as if they were a form of step function in phase space. From this we were able to take what is known as a Poincaré section. We see this in Figure 1, where we plot the Poincaré section taken from embeddings of $m = 600, 1180, 1550,$ and1860 when $\tau = 5$. Similar plots has been built for $\tau =$ 10, and 20, but not shown here for compactness of the presentation. The intuition behind Poincaré sections comes from the notion that, as they move in time, dynamical systems with interesting regimes travel through lower dimensional manifolds in a ordered manner. Therefore by choosing a lower dimensional manifold in a tactful way and then keeping track of exactly where on the manifold this system passes *in one direction*, we have gathered a tremendous amount of information about the dynamic regime. In our case we have summarized $N(E, I)$ into 3 dimensions, placed a plane about halfway down the diagonal from $(1, 1, 1)$ to $(0, 0, 0)$, and observe how it passes through this plane in the backwards direction assuming it travels like a step function.

Notice that the 1-period appears as 1 point, a 2-periodic point appears as 2 points, and so on. This is because as the system travels in phase space a 1-period by definition only fires in one way, a 2-period in two ways, and so on for n-periods. Although still leaving much to be desired, this visualization already makes rather clear that the behavior of a chaotic-like regime differs tremendously from typical, short periodic oscillations. Part of what makes this picture hard to interpret has to do with something mentioned above during the description of embeddings, i.e. we should create the time series from a good summary statistic of the network's

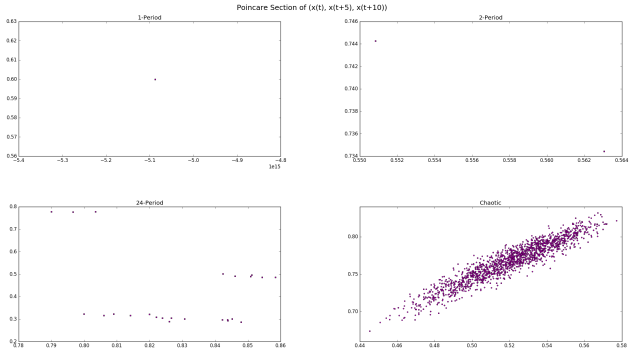

FIGURE 1. Poincaré Section for $m = 600, 1180, 1550, 1860$; $\tau = 5$

state. The time series we are using naively reports the number of active nodes at each time step. Though we are looking at these as a series, and so there is more implicit information, each individual datum in the series is equivalent to any other configurations of the network with the same number of active nodes. Therefore, we have in one sense captured too little by leaving out all of the topological information and we have simultaneously captured too much by obscuring away our knowledge of what portion of the network is active at any given point.

As a remedy to this we decided to toss out some information about the global activity in the system to instead capture more spatial information. To do this we focused on three kxk blocks of nodes that are each situated as far from one another as possible. We then captured the activity in all three blocks as a time series already in three dimensions and created Poincaré sections of this. To understand the importance of this spatial information we captured the activity in each block in one of two ways: (i) we simply summed the number of active nodes as before and (ii) we took a weighted sum of the activity in such a way that a node's weight exponentially decays with distance from the center-most node in the block. We see the Poincaré section for (i) in Figure 2 and for (ii) in Figure 3. Here the blocks are of size $k = 5, 10, 20, 25, 50$, $m = 1860$ and Top Left refers to the top left of the network when represented on a plane while the Farthest and Middle are relative to the top left.

By considering both of these side by side, we can see that in both the blocks of the same size capture the same general trend to varying degrees. This tells us some information about the trends in the local blocks which are spatially apart from one another. Even though we have gleaned some new insight into the dynamics of the system we have lost our ability to interpret this information globally in the network. Regardless, note the fact that just one block, though lacking most of the information, still captures something impossible to distinguish in the mean field by reporting local spatial activation. Furthermore we have blocks of gradually increasing size overlapping the same spatial location. By accounting for this spatial redundancy we are able to visually extract a huge amount of topological context while also being able to use the entire network's activity. We therefore took Poincaré sections of the time-delay embedding for the Top Left block of size $k = 5, 10, 20, 25, 50$ and compared it to the Poincaré section of the embedding of the entire system's activity. We did not use the exponential decay weighted sum

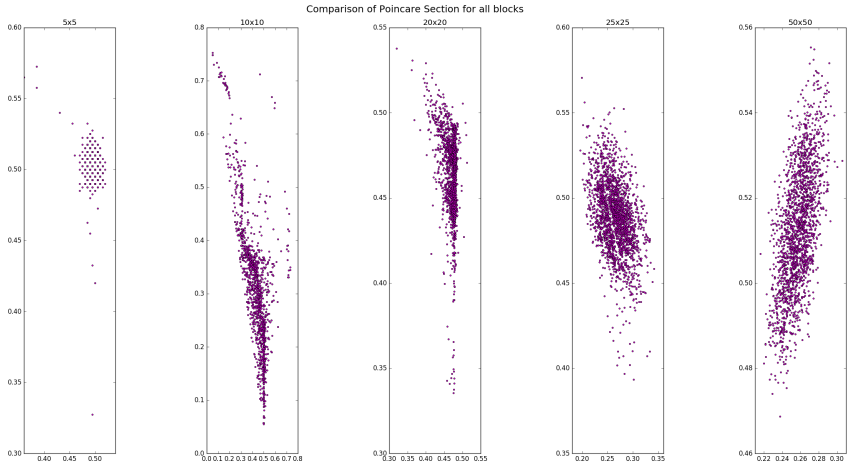

FIGURE 2. Comparison of Poincaré Section for blocks' of size $k = 5, 10, 20, 25, 50$ using simple sum.

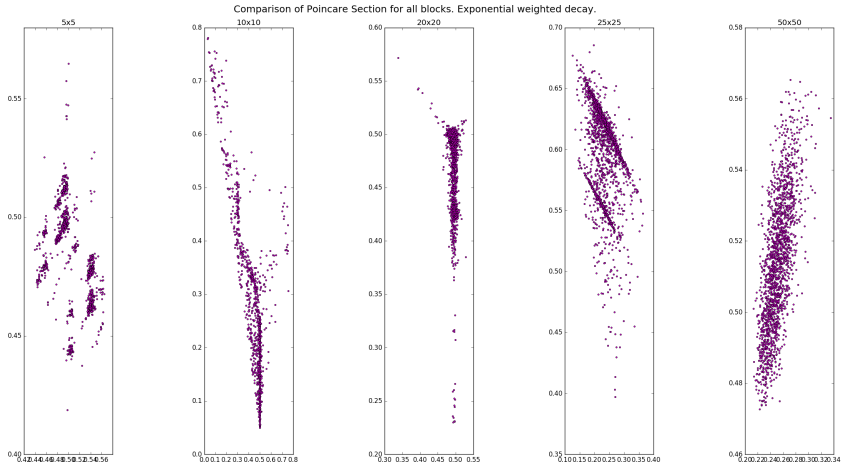

FIGURE 3. Comparison of Poincaré Section for blocks' of size $k = 5, 10, 20, 25, 50$ using exponential weighted sum.

to avoid having the blocks of different sizes report biased information due to where activity happens to be in the block[1] (e.g. if most of the activity is far from the center of big blocks it is hard to interpret between block sizes). Figure 4 shows us this progression with a time delay of $\tau = 5$, Figure 5 shows us this progression with a time delay of $\tau = 10$, and Figure 6 shows us this progression with a time delay of $\tau = 20$. All three figures were created with a time series from a chaotic trajectory with $m = 1860$.

[1]We intend to treat this more carefully in the future

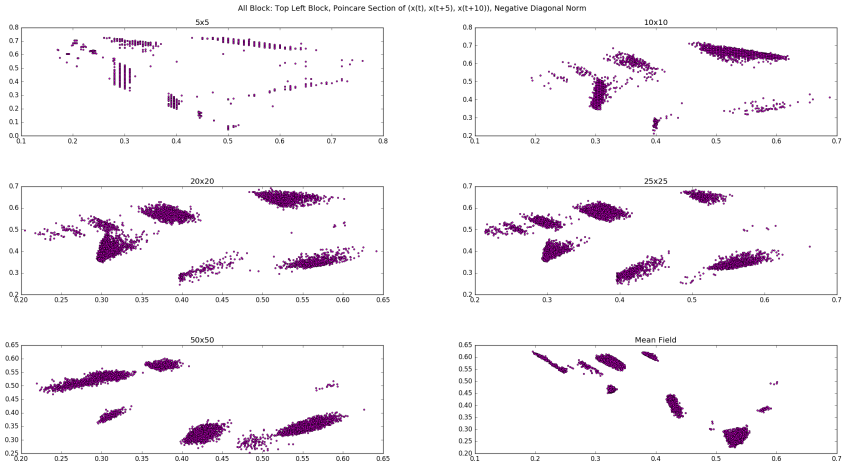

FIGURE 4. Comparison of Poincaré Section for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 5$

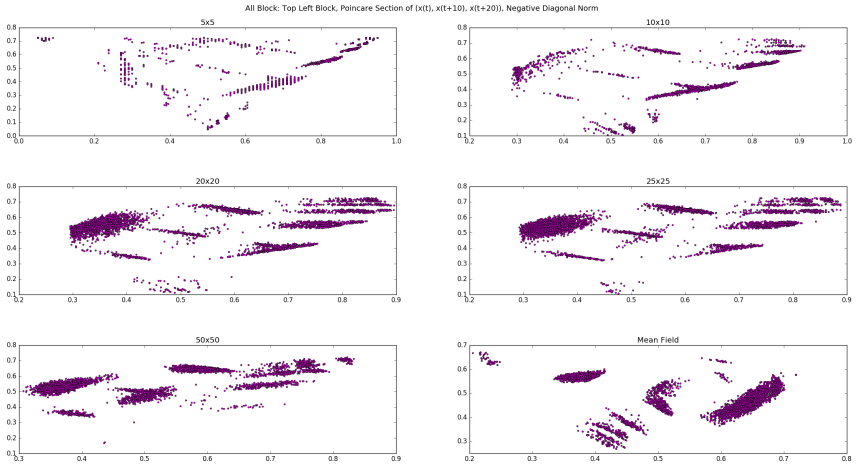

FIGURE 5. Comparison of Poincaré Section for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 10$

The difference in the shape being plotted and which (x, y) pairs, when considering these plots as a package, reaffirm what we know about average spike lengths and the standard deviation. However by seeing it in this context, we get a better notion of whether the deviations tend to be above or below the average; in some sense, this gives us an idea of where space gets distorted by the reset rule. This is demonstrated by the triangular shapes with a peak around the center of the x axis of the Poincaré plane. Unfortunately, because we are assuming the system is a step function and

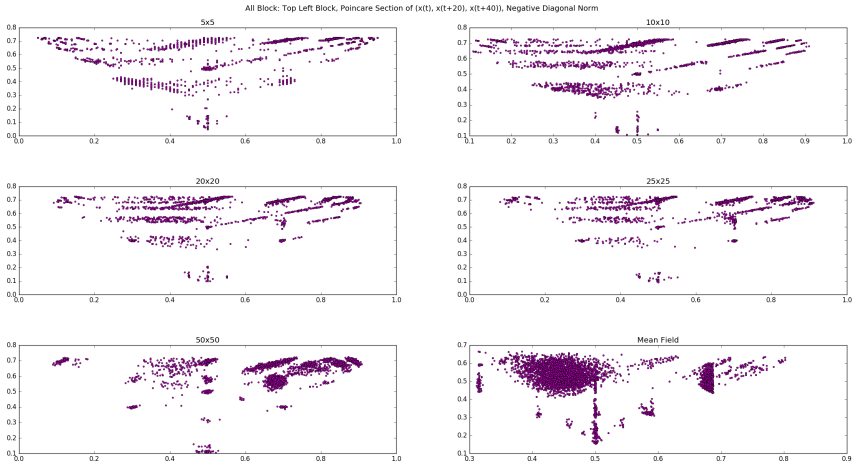

FIGURE 6. Comparison of Poincaré Section for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 20$

we are placing the plane somewhere in the middle of all the trajectories, the actual number being reported on each axis is hard to interpret and do not correspond to any specific time-delay axis. The actual shape shown in the plot however is typically accurate up to scaling and symmetries since we are looking at the plane "from the front". When we consider the Poincaré sections in order of increasing block size we begin to understand how the dynamics work within E itself; this process can almost be thought of as adjusting the focus on a camera. Notice that from the 5x5 block to the mean field the dynamics slowly gain more degrees of freedom as the space of possible densities grows. By considering where the points start to "appear" (i.e not overlap anymore) from one block size to the next we can see that certain parts of the network are where the excitatory layer differs between spikes. The different embeddings also reveal to us different information about this. For example notice that when $\tau = 5$ from the 5x5 block to the mean field there is a good difference but the structure of the shape stays the same. This implies that that general part of the lattice in this delay, though slightly more varied, had the same type of dynamics. On the other hand, when $\tau = 20$ from the 5x5 to the 25x25 block this is the case, but then from the 25x25 block to the mean field the dynamics underwent a rather fundamental change in their trajectory. This is likely because the larger value of τ has vectors that spans multiple spikes and reveals to us that where these spikes really begin to vary is somewhere between the 25x25 block to the mean field. In all of these cases, we see that the system has strong basins of attraction in the phase space.

These basins of attraction that are evident in the figures just discussed only become more interesting once understand how the system moves from one to the other. Unfortunately, without taking anything else into account, simply the points passing through the section do not have any insight into how this movement happens. However by treating those points as a new series $\{P(i)\}$ that captures the

order in which the trajectory crosses the plane and plotting $(P(i), P(i+1))^2$, we get a glimpse as to how the system moves from one basin into the next. This technique is referred to as a Poincaré map and is shown in Figure 7 for $\tau = 5$, in Figure 8 for $\tau = 10$, and in Figure 9 for $\tau = 20$.

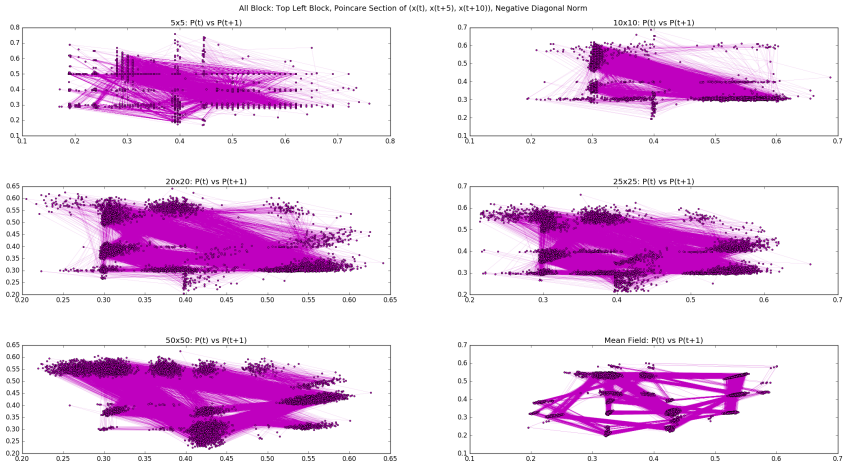

FIGURE 7. Comparison of Poincaré maps for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 5$.

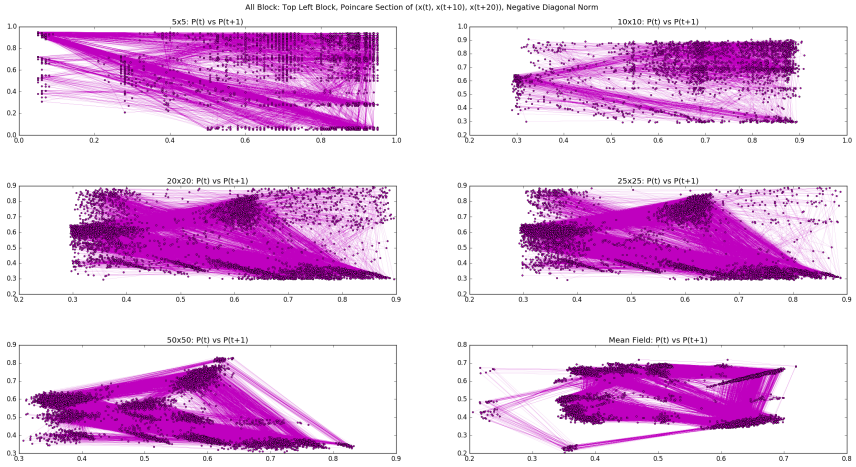

FIGURE 8. Comparison of Poincaré maps for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 10$

[2]This is in some sense another two dimensional embedding delay

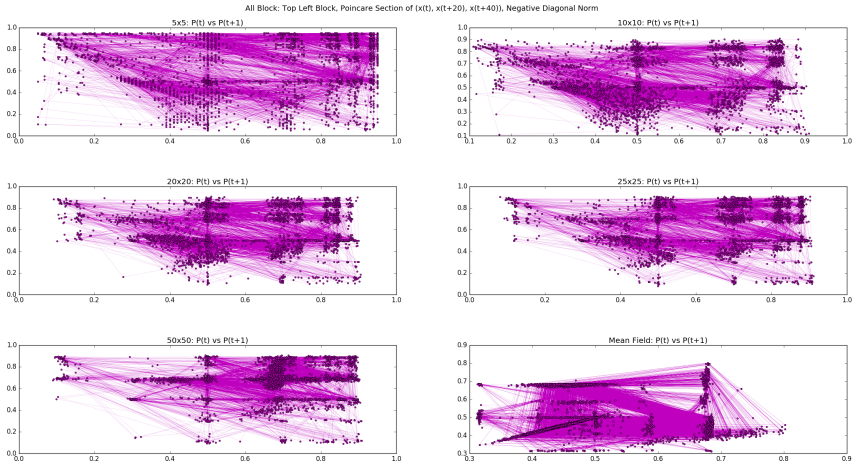

FIGURE 9. Comparison of Poincaré map for Top Left Blocks' of size $k = 5, 10, 20, 25, 50$ and the mean field (i.e $k = 100$). Here $\tau = 20$

It is immediately clear that regardless of block size and τ, the blocks are just in general less informative. The reason for this will literally be shown below but from these Poincaré maps we can already see that because the blocks do not have all of the information different basins overlap. Therefore the movement from one basin into the next looks almost random at times and so the huge number of crossing lines in the plots make things hard to see. The mean field on the other hand captures all of the basin's rather precisely and so they appear in the Poincaré map as concentrated clusters connected by differentiable bundles of lines. We can now be rather certain that the chaotic regimes of $N(E, I)$ behave in a very principled way. By comparing the Poincaré maps of different block sizes as we did above, we also get a glimpse of where these different basins live in the phase space and how they divide up this space in the actual configurations of the lattice. It is often noted that these Poincaré maps almost create a sense of three dimensions. This is because they are recreating the trajectory of the three dimensional system used to create them. Since we can see that $N(E, I)$ is travelling from one basin to the next in a well groomed manner we will look at the time-delay embedding in three dimensions without using the Poincaré sections. Furthermore, because looking at these figures in three dimensions makes a side-by-side comparison sensitive to the perspective by which we look at a trajectory anyway, we are no longer concerned with a biased comparison and will use data from the weighted block sum because it provides more spatial information than a regular sum. To help us establish a notion of orientation, which is not intuitive since we are treating time as a spatial dimension, we have colored the edges in the figure based on which time axis is firing from one moment to the next. The color corresponding to each direction is in the legend at the top of the figure. Though we have these figures for every sized box, we only include here a couple of telling examples. In Figure 10 we see the trajectories for the 25x25 block and in Figure 11 we see the mean field. Both are made with

$\tau = 5$. It should be mentioned that the mean field is also captured as a weighted sum based on the center-most node.

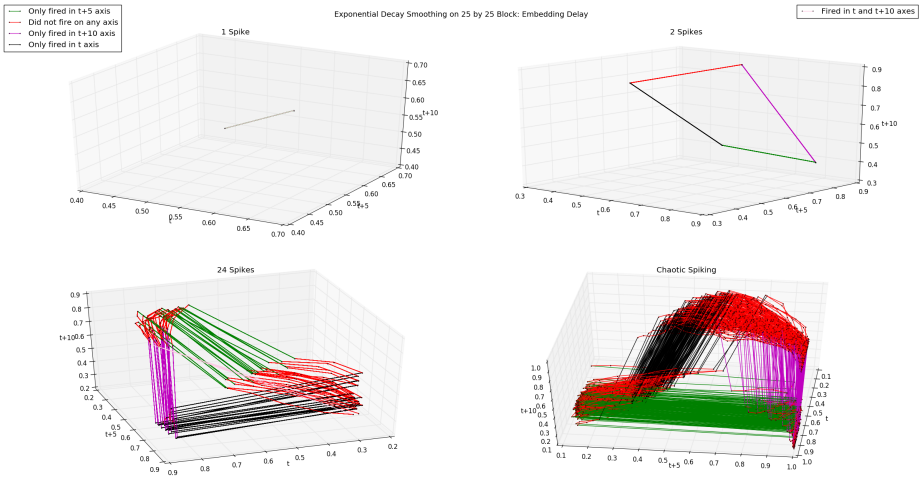

FIGURE 10. Three dimensional time-delay embedding with trajectory colored based on axes firing. For blocks of size 25x25 with exponential weighted decay sum on the nodes.

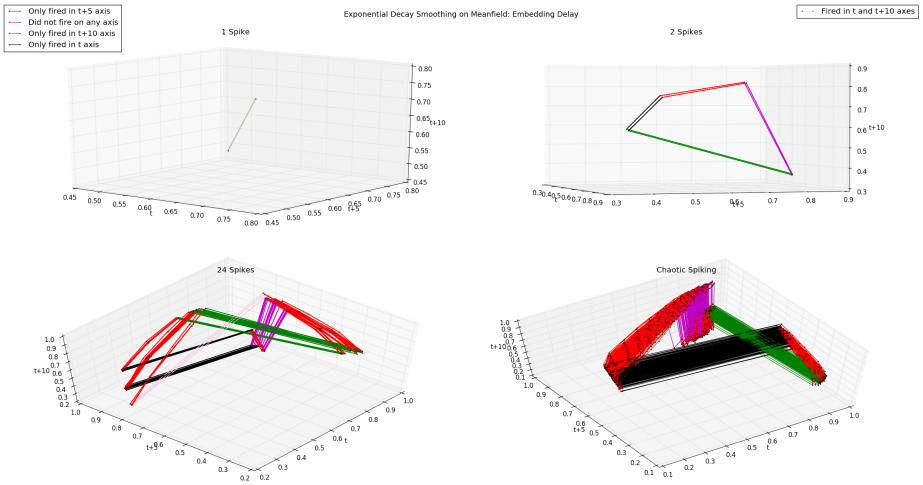

FIGURE 11. Three dimensional time-delay embedding with trajectory colored based on axes firing. For the entire lattice with exponential weighted decay sum on the nodes.

This confirms everything we'd seen in the Poincaré map and then some. In the Poincaré map we saw the basins overlap, but now due to the inclusion of a sense of depth we can see that in fact there is ordered movement in the blocks too. In the mean field a similar conclusion is reached. Previously we could see that in fact one

basin led to the next in a neat way, now we can confirm that this goes beyond that and is moving from one basin into the next in a morphologically repetitive way that requires a huge amount of detail to properly describe (i.e. it moves chaotically). Both of these points are further reinforced when we compare the trajectory for the different spike lengths.

Looking at these Poincaré section and delay embeddings we have caught glimpses into the spatial behavior of our network $N(E, I)$ in phase space. Considering all that we have seen in conjunction paints a rather vivid picture of how the system moves. Though we can understand the behavior and know what to expect it is by no means intuitive. As has been mentioned, this is because in a spatio-temporal context the trajectory takes sharp turns due to the discontinuous nature of the system coupled with its high dimensionality. Both of these difficulties arise in the spatial dimension and so by focusing on temporal representations, though we lose insight into the trajectory itself, we may get a clear view of the regime.

With that said, we continue to use the idea from the Poincaré section to capture some temporal information. All we do is have our plane run parallel to the time axis. This produces a figure such as Figure 12. Note that in this visualization it is as clear as can be whether the regime is periodic or not. Furthermore, because it still accounts for the density along one axis, we can also see the individual number of spikes. It does however present a new problem; chaos looks almost like noise. By completely removing the spatial axis we can avoid this.

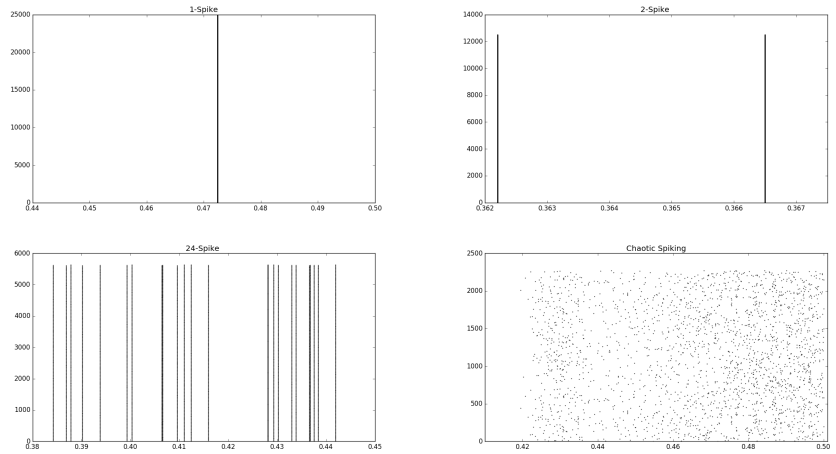

FIGURE 12. Temporal Poincaré section of meanfield time series

To accomplish this, we use what is known as a recurrence plot. Define our time series from the time-embedding delays as

$$x(t) = \left\{ \left[|A^E(t)|, |A^E(t+\tau)|, \ldots, |A^E(t+n\tau)| \right] : t \in T \right\}$$

This series, at the very least partially, recreates the network $N(E, I)$'s movement. Let t and t' be any two points in time. By plotting $(x(t), x(t'))$ whenever $x(t) \approx (t')$ we have a recurrence plot. As will be shown, this plot completely loses information about the individual spikes (including how many there are) but reports the length of time in a period. This is acceptable because our goal is to make absolutely

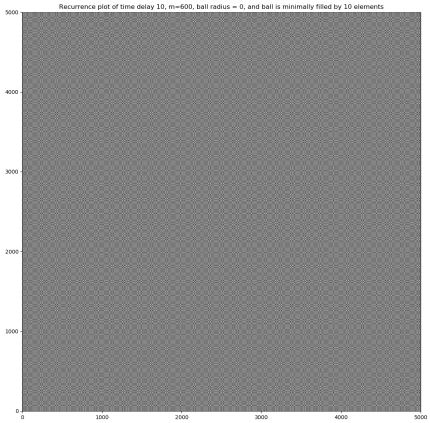

FIGURE 13. Recurrence Plot of 1-period from $m = 600$. Note 0 here represents the 9500000^{th} time step.

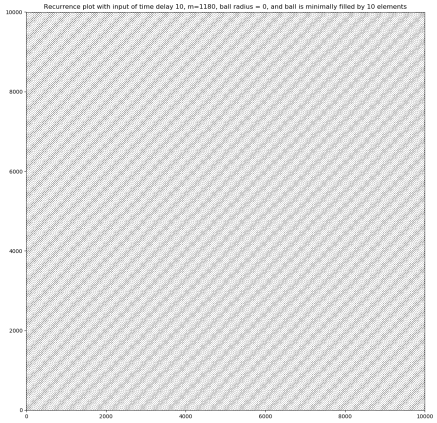

FIGURE 14. Recurrence Plot of 2-period from $m = 1180$. Note 0 here represents the 9500000^{th} time step.

clear the general type of dynamic regime, which it does, and we already have visualizations that can show us this information. In Figure 13 we see a 1-period gathered from $m = 600$, in Figure 14 we see a 2-period gathered from $m = 1180$, in Figure 15 we see a 24-period gathered from $m = 1550$, and lastly in Figure 16 we see chaotic spiking gathered from $m = 1860$. We created these plots using $t, t' \in [9500000, 10000000]$ and a ten dimensional embedding delay with $\tau = 1$. To capture some more information about the basins of attraction and also to make the plot less spare and easier to see, we plot $(x(t), x(t'))$ when both are within a ball

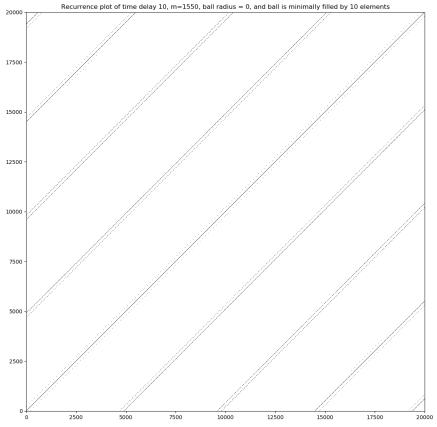

FIGURE 15. Recurrence Plot of 24-period from $m = 1550$. Note 0 here represents the 9500000^{th} time step.

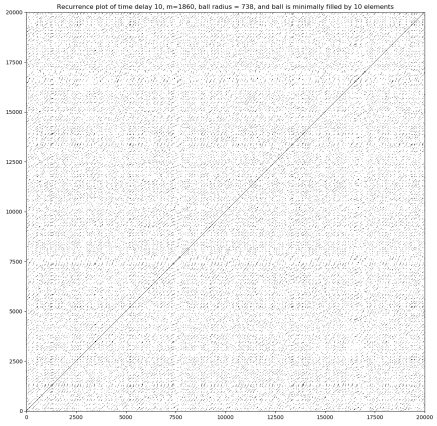

FIGURE 16. Recurrence Plot of chaotic spiking from $m = 1860$. Note 0 here represents the 9500000^{th} time step.

of radius that minimally contains 10 points. We note that every one of these plots loses details due to the extremely fine lines borne from having them made with ≥ 10000 time steps. Regardless, these plots still make clear the regime. For this reason we unfortunately cannot do a proper analysis on the entirety of the plot, but we would have noticed that periodic regimes form stripes while chaos is messy but structured. To see this we include Figure 17 as an example.

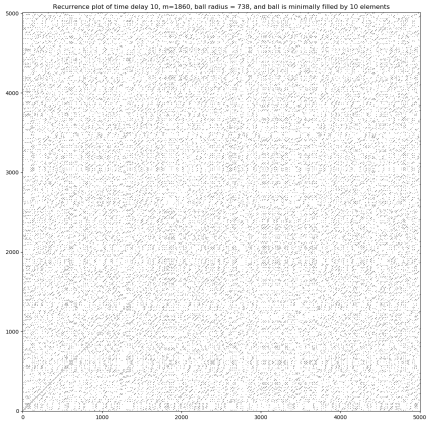

FIGURE 17. Recurrence Plot of chaotic spiking from $m = 1860$ zoomed in to only account for the first 5000 time steps

4. SUMMARY

We studied various network architectures with inherent oscillatory dynamics defined over these network structures. We implemented recurrent networks with feedback loops. Recurrent networks allow complex dynamics, such as oscillations in narrow and broad bands. We have observed very long oscillations, which have periods exceeding the monitoring window of 10,000 time steps. There is an intimate link between network structure and dynamics. Namely, we documented very complex (chaos-like) dynamics based on phase diagrams, Poincare sections, and recurrence plots. The produced spatio-temporal oscillations are self-sustained, thus they are very useful in the implementation of the the energy-efficient computing and learning tasks developed in other tasks of the project. These results, combined with learning rules developed in other tasks, have the potential of drastically increased memory capacity when these networks are used as associative memories to implement a machine learning task.