CHAPTER 2

INTEGRATE-AND-FIRE OSCILLATORS

2.1 Introduction

The integrate-and-fire oscillator is possibly the simplest model of neuronal behavior. A variable associated with the membrane voltage of a neuron is allowed to increase from zero until a threshold is reached. Once the threshold is attained, the oscillator is said to "fire". This variable is instantaneously reset to zero and the process repeats. Although a gross approximation of neural activity, the integrate-and-fire model has been of heuristic value to neurobiology. Peskin [Peskin, 1975] created a network of globally connected integrate-and-fire oscillators to generate synchronous periodic activity; his purpose was to model the sinoatrial node - the group of synchronous neurons which generate the heart beat. Networks of integrate-and-fire neurons have been examined by physicists and mathematicians [Mirollo and Strogatz, 1990, Kuramoto, 1991, Bottani, 1995]. The neural network community has used these oscillators to model neural activity [Eckhorn et al., 1990, Chawanya et al., 1993]. Also, some types of integrate-and-fire oscillators are identical to the driven elements used in earthquake models, as noted by Corral et al. [Corral et al., 1995a]. Thus, these networks have direct links to studies of earthquake, avalanche, and forest fire models.

We examine only locally coupled networks. Local couplings are more neurobiologically realistic and it has been argued they allow for more powerful and flexible properties of information processing in Section 1.5. In simulations with locally coupled networks of integrate-and-fire oscillators, several authors have observed that they exhibit synchrony, i.e. all oscillators fire at the same time. [Mirollo and Strogatz, 1990, Corral et al., 1995a, Hopfield and Herz, 1995]. Mirollo and Strogatz [Mirollo and Strogatz, 1990] speculated that, a system of such oscillators would "end up firing in unison for almost all initial conditions, *no matter how the oscillators were interconnected*." Rigorous mathematical proof indicating that synchrony occurs in these systems remains undiscovered. The answers to many questions about synchronization in integrate-and-fire oscillators are unknown, such as "What is the rate of synchronization?", "How robust is synchronization in the presence of disorder?", "How does the rate of synchronization change in the presence of disorder?",

"How do the behaviors of the network change as the dimension of the system changes?", etc. In this Chapter, we examine the rate of synchronization in locally coupled networks. We provide numerical evidence and a heuristic explanation indicating that the time to synchrony increases with the logarithm of the system size for one- and two-dimensional networks of identical oscillators without noise.

Integrate-and-fire oscillators have been used in the neural network community to perform computational tasks [Hopfield and Herz, 1995], as have pulse-coupled neural networks [Johnson, 1994], a similar type of oscillator. Hopfield and Herz [Hopfield and Herz, 1995] indicate that in locally connected networks of integrate-and-fire oscillators, synchrony occurs on long time scales (more than a hundred periods in a 40×40 network). Based on this observation, they view synchrony as biologically irrelevant because computational decisions must be made rapidly by an animal. They then construct a network of integrate-and-fire oscillators which uses local synchrony to perform image segmentation. Local synchrony is a behavior in which clusters of oscillators fire in unison, but the entire network is not synchronous. In their network (as in Figure 3) each pixel of an image is associated with an oscillator. The potential of an oscillator is determined by the grey-level of its corresponding pixel. When one oscillator fires, it can induce its neighbor to fire if the potentials of the oscillators are similar. Different groups of oscillators corresponding to different homogeneous grey-level regions in the image fire at different times. One drawback of this network is that a region of smoothly varying grey-level intensities can be arbitrarily broken into smaller regions. Another drawback is that two different regions with the same intensity would fire at the same time and would thus be improperly grouped together. This occurs because their network does not have a mechanism which actively desynchronizes different oscillator groups.

In our examination of integrate-and-fire oscillators, we have duplicated the results of Hopfield and Herz [Hopfield and Herz, 1995] and further, we have found that synchrony can occur quickly (within a few periods) in these networks by appropriate adjustment of parameters. We use this property of fast synchrony to create a network for image segmentation (Section 2.9) that does not have the flaws of the network proposed by Hopfield and Herz [Hopfield and Herz, 1995]. Our network actively desynchronizes different oscillator groups so that only a single region fires at a time. Also, because our network achieves synchrony (as opposed to local synchrony), regions with smoothly varying grey-levels are not arbitrarily broken apart. The network architecture we use is directly based on that of Terman and Wang [Terman and Wang, 1995]. Because of the computational efficiency with which integrate-and-fire oscillators can be numerically integrated, we apply our network to real images and segmentation results for two such images are shown.

One of our purposes is to compare how different types of oscillators synchronize when locally connected. In Chapter 3 we study relaxation oscillators (see Appendix A for a quick review of relaxation oscillators). Relaxation oscillators and integrate-and-fire oscillators are frequently assumed to have similar properties. Both models have two time scales and both models are typically examined with an asymmetric, discontinuous coupling [Mirollo and Strogatz, 1990, Somers and Kopell, 1993, Terman and Wang, 1995, Ernst et al., 1995]. However, the two models have significant differences. For example, integrate-and-fire oscillators have an instantaneous interaction, while that between relaxation oscillators is of finite duration. Furthermore, integrate-and-fire oscillators contain

only one variable and are computationally simpler. On the other hand, they cannot exhibit amplitude variations as seen in relaxation oscillators. In Chapter 3 we examine relaxation oscillators and observe that in one-dimensional networks the time to synchrony is a power law in relation to the size of the network. In Section 3.7 we attempt to explain where these qualitative differences between relaxation oscillators and integrate-and-fire oscillators originate.

We first define our integrate-and-fire model in Section 2.2 and examine the behavior of a pair of coupled oscillators in Section 2.3. In Section 2.4 we present our data on the time to synchrony for locally coupled one-dimensional systems of integrate-and-fire oscillators. In Section 2.5 we present similar data for two-dimensional systems of locally coupled integrate-and-fire oscillators. We heuristically explain how the time to synchrony scales as the size of the system in Section 2.6. Our integrate-and-fire oscillator system contains only two parameters and in Section 2.7 we examine how these two parameters are related to the rate of synchronization. In Section 2.8 we examine how quickly synchrony is attained when errors of various sizes are introduced into a two-dimensional synchronized system. Section 2.9 demonstrates how to achieve desynchronization between different groups of integrate-and-fire oscillators while maintaining synchrony within each group. We use this ability of synchrony and desynchrony to perform a temporal labelling task related to image processing in Section 2.10. We discuss our results in Section 2.11.

2.2 Model definition

A network of integrate-and-fire oscillators is defined as follows,

$$\dot{x}_i = -x_i + I_0 + \sum_{j \in N(i)} J_{ij} P_j(t), \qquad i = 1, \ldots, n \tag{2.1}$$

where the sum is over the oscillators in a neighborhood, $N(i)$, about oscillator i. The variable x_i represents some voltage-like state and we refer to this as the potential of oscillator i. The parameter I_0 controls the period of an uncoupled oscillator. The threshold of an oscillator is 1. When $x_i = 1$ the oscillator is said to "fire". Its potential is instantly reset to 0 and it sends excitation to its neighbors. The interaction between oscillators, $P_j(t)$, is defined as follows,

$$P_j(t) = \sum_m \delta(t - t_j^m) \tag{2.2}$$

where t_j^m represents the m firing times of oscillator j and $\delta(t)$ is the Dirac Delta function. When oscillator j fires at time t, oscillator i receives an instantaneous pulse. This pulse increases x_i by J_{ij}. If x_i is increased above the threshold, then it will fire. Note that information is transmitted between oscillators instantaneously, thus the propagation speed is infinite. Figure 4 displays the temporal activity of a pair of integrate-and-fire oscillators. The oscillators initially have different potentials, but the interaction quickly adjusts their

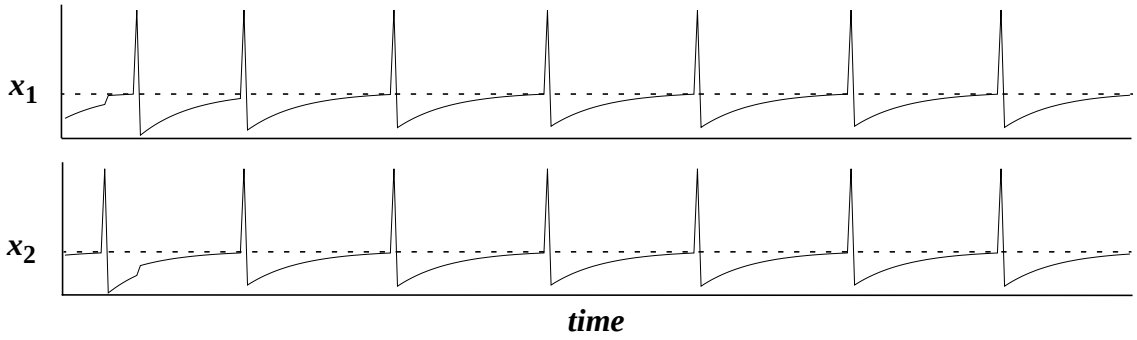

Figure 4. A diagram of a pair of integrate-and-fire oscillators with pulsatile coupling. The solid curves represent the potentials of the two coupled oscillators and the dashed lines represent the threshold. The initial potentials of the oscillators are chosen randomly. The oscillator labeled x_2 fires first and the potential of x_1 increases at that time. Similarly, when x_1 fires, the potential of x_2 increases. The phase shifts caused by the pulsatile interaction causes the oscillators to fire synchronously by the second cycle. The spikes shown when an oscillator fires are for illustration only.

trajectories so that they eventually fire in unison. When two or more oscillators fire at the same time we call them synchronous. The spikes shown in Figure 4 when an oscillator reaches the threshold are for illustrative purposes only.

The coupling is between nearest neighbors, i.e. an oscillator interacts with two neighbors in one-dimension and four neighbors in two-dimensions. The connection strength from oscillator j to oscillator i is normalized as follows

$$J_{ij} = \frac{\alpha_I}{Z_i} \tag{2.3}$$

where Z_i is the number of nearest neighbors that oscillator i has, e.g. $Z_i = 2$ for an oscillator i at the corner of a two-dimensional system. The constant α_I is the coupling strength and the subscript I is used to denote that this parameter is used for integrate-and-fire oscillators. The normalization ensures that all oscillators receive the same amount of stimulus, and therefore, have the same trajectory in phase space when synchronous [Wang, 1995]. Note that there are only two parameters in system (2.1), the coupling strength α_I and I_0. When oscillator i reaches its threshold, it will fire, and its value will be reset to zero. Oscillator i then sends an instantaneous impulse to neighboring oscillator j. If oscillator j is induced to fire, then its value is reset in the following manner,

$$x_j(t^+) = x_j(t^-) + J_{ji} - 1 \tag{2.4}$$

Since oscillator j fires, immediately oscillator i receives excitation and thus $x_i(t^+) = J_{ij}$. Because of this, the period of the synchronous system is shorter than the period of a single uncoupled oscillator. The synchronous period of the system is given by

$$\log\left(\frac{I_0-\varepsilon}{I_0-1}\right) \tag{2.5}$$

There are many reset rules to choose from and we choose (2.4) because it appears to have the greatest similarity to the system of relaxation oscillators studied in Chapter 3. This particular realization of a network of integrate-and-fire oscillators was called “Model A” by Hopfield and Herz [Hopfield and Herz, 1995].

2.3 A pair of integrate-and-fire oscillators

We now describe the behavior of a pair of integrate-and-fire oscillators. This section is a short summary of some of the results derived by Mirollo and Strogatz [Mirollo and Strogatz, 1990]. The trajectory of a single oscillator can be solved analytically. i.e. $x(\phi) = f(\phi) = I_0(1 - \exp(-\phi))$, where ϕ can be thought of as a phase, or a local time variable. Note that the function $f(\phi)$ increases monotonically, $f'(\phi) > 0$, and is concave down, $f''(\phi) < 0$. Using $f(\phi)$ and its inverse, $g(x)$, one can calculate the return map (Figure 5A) for a pair of pulse coupled integrate-and-fire oscillators. A line of slope 1 is also shown in Figure 5A for comparison. The horizontal axis represents the initial phase difference between the two oscillators and the vertical axis represents the phase difference between the two oscillators after they have both fired once. There are three different regions in the return map. The first region is in the range of initial conditions $\phi_1 \in [0, \phi_L]$, where $\phi_L = 1 - g(1 - \alpha_I)$. In this region, the oscillators are near enough so that when one oscillator fires, the second oscillator is induced to fire as well. We call this the *jumping region*, and it has a direct analog in a pair of relaxation oscillators. Once the two oscillators are in the jumping region, the oscillators always fire at the same time and it can be shown that their phase difference always decreases. The same is true for the analogous jumping region in a pair of relaxation oscillators [Terman and Wang, 1995]. The second region is in the range of initial conditions from $[\phi_L, \phi_U]$, where $\phi_U = 1 - g(f(\phi_L) - \alpha_I)$. For these initial conditions, when the first oscillator fires, the other oscillator receives excitation, but is not induced to fire at the same time (as in the first firing of the second oscillator in Figure 4). Similarly, when the second oscillator fires, the relative phase between the two oscillators again changes, but the two oscillators do not fire in unison. In this region there is an unstable fixed point for which the phase between the oscillators does not change. In the third region, the first oscillator fires and the second oscillator receives excitation, but does not fire immediately. When the second oscillator fires, the first oscillator receives excitation and is induced to fire a second time. We consider this third region as part of the jumping region. In summary, this return map contains a range of initial conditions for which the two oscillators fire together, and another set of initial conditions for which it may take several cycles before both oscillators begin firing together.

In Figure 5B we display the number of cycles needed before the two oscillators are in the jumping region. The horizontal axis in Figure 5B indicates the initial phase separation between the two integrate-and-fire oscillators and the vertical axis indicates the number of cycles needed until the two oscillators are in the jumping region. As expected, initial conditions near the unstable fixed point require more cycles before synchrony occurs.

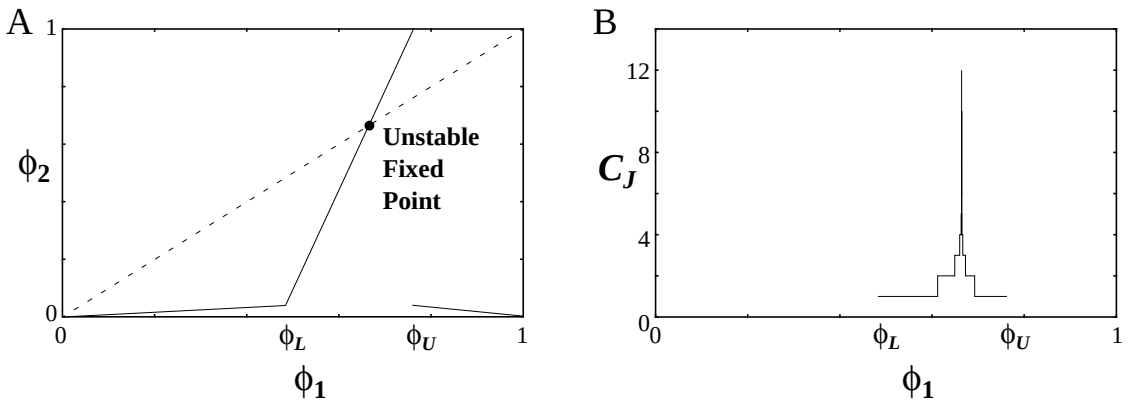

Figure 5. (A) The return map for two pulse coupled integrate-and-fire oscillators. The phase difference between the oscillators before they have jumped (ϕ_1, horizontal axis) and after they have jumped (ϕ_2, vertical axis). (B) A plot of the number of cycles needed, C_J, before the two oscillators are synchronous as a function of ϕ_1. Both plots use $I_0 = 1.11$

2.4 Synchrony in One-Dimensional Chains

We observe that synchrony occurs in all locally coupled networks of integrate-and-fire oscillators with positive coupling, with $f'(\phi) > 0$, and $f''(\phi) < 0$ for all initial conditions tested. This is consistent with the observations of several authors [Mirollo and Strogatz, 1990, Corral et al., 1995a, Hopfield and Herz, 1995]. They each observed that locally coupled networks of integrate-and-fire oscillators always synchronized.

The oscillators in (2.1) can be numerically integrated with an event driven algorithm. In all simulations we use the following procedure.

1. The potentials are chosen from the range [0,1].

2. Calculate the local time, or phase variable, for all oscillators.

3. Find the oscillator nearest to the threshold. The amount of time it needs to fire is calculated and all the oscillators are advanced using this amount of time.

4. The oscillator at the threshold fires. The potential of this oscillator is reset to zero and the potentials of its neighboring oscillators are increased using (2.3).

5. Check if any of the oscillators that have received excitation are above the threshold. If any oscillators are above the threshold, they are reset according to (2.4) and excitation is sent to their neighbors. Repeat this step until no oscillators are above the threshold.

6. Return to step 2.

All trials with locally coupled networks of integrate-and-fire oscillators have resulted in synchrony. Over 10^5 trials in which the initial conditions were chosen randomly and uniformly in the range $[0, 1]$ have been recorded. These networks were also tested with other, more correlated initial conditions. Networks in which the initial conditions were spin waves also achieved synchrony. The speed with which networks with spin wave type initial conditions attained synchrony was, on average, faster than the time to synchrony using random initial conditions. For long wavelength spin waves, the potentials of the oscillators are near to each other and one oscillator can cause many of its neighbors to fire. Several large groups, or blocks, of oscillators form and fire synchronously during the first cycle. For short wavelengths that are integer multiples of the lattice size, the oscillators also synchronize more quickly than with random initial conditions. Small blocks of synchronous oscillators form, and since these blocks are formed based on repeating patterns of initial conditions the blocks have a spatially repeating pattern. This process repeats until synchrony occurs. This implies that incommensurate wavelengths may take longer to synchronize because spatially repeating patterns of blocks do not form and their interactions with one another would not be uniform. This intuition does appear to be correct; incommensurate wavelengths tend to have longer synchronization times. However, we could not find any initial conditions whose resultant time to synchrony was an order of magnitude larger than the average time to synchrony with random initial conditions (over 10^4 incommensurate frequencies were tested). Similar tests in two-dimensional networks yield similar results. There are a few solutions which are not synchronous, e.g. initial conditions in which the phase difference between pairs of oscillators is at the unstable fixed point shown in Figure 5A. In numerical tests with these initial conditions, floating point errors eventually cause small perturbations away from this unstable solution and synchrony quickly results. Furthermore, in trials with periodic boundary conditions (a ring topology), solutions with travelling waves were never observed. Based on these observations, it is our belief that locally coupled networks of integrate-and-fire oscillators *always* synchronize. Although all of our data has been gathered using one or two specific integrate-and-fire oscillators, we claim that our results generalize to the class of integrate-and-fire oscillators with positive coupling, $f'(\phi) > 0$, and $f''(\phi) < 0$.

We display the temporal evolution of a one-dimensional network in Figure 6. In this graph we display the firing times of all the oscillators in a network of 400 oscillators. Time is shown along the vertical axis and the horizontal axis represents the index of the oscillators. Each dot represents the firing time of one oscillator and each line represents the firing time of a block of oscillators. Near the bottom of the graph, there are many single dots and small lines. These represent the fact that the oscillators have random initial conditions and initially have distinct firing times. But quickly, by the time $t = 5$, blocks of various sizes have formed. Just after time $t = 5$, at the lower left of Figure 6, oscillators 1-20 fire simultaneously. This block formed from three smaller blocks. Near $t = 30$ there is a single solid line shown, indicating that all the oscillators fired at the same time. Underneath this line, there are two separate blocks of oscillators. One might at first wonder why these two large blocks have merged in just one cycle. This represents the fact that the system has an instantaneous propagation speed. When the oscillator at the border of the left block receives excitation, it is induced to fire. When this oscillator fires it sends excitation to its left neighbor, which is also induced to fire, and this process repeats throughout the length

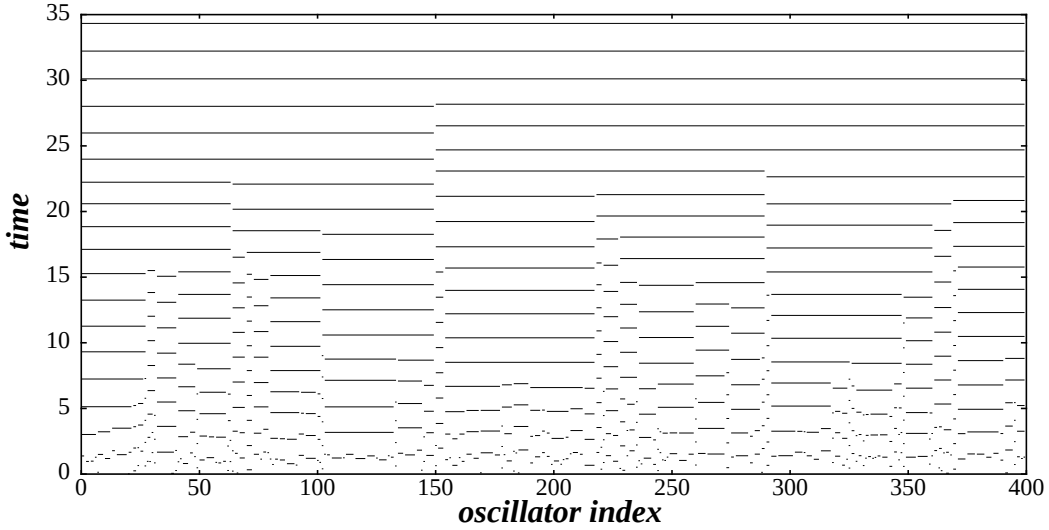

Figure 6. A diagram displaying the evolution of a one-dimensional network of integrate-and-fire oscillators. The vertical axis represents time and the horizontal axis represents the position of the oscillator in the network. Each dot (line) represents the firing time of a single (several) oscillator. The parameters are $\alpha_I = 0.2, I_0 = 1.11$.

of the left block. In the algorithm we use, the firing and reset of an oscillator are instantaneous, as are the excitatory pulses sent to neighboring oscillators. This results in an infinite propagation speed. Thus, no matter how large a block is, it can merge with a neighboring block in one cycle. The most striking feature of Figure 6 is that it is impossible to find a block that decreases in size. In Section 2.4.1 we prove that in a one-dimensional network, it is not possible for the interaction to break apart a block of oscillators.

In Figure 7 we display data indicating that the average time needed to synchronize a chain of size n increases in proportion to $\log_{10}(n)$. The averages are based on several hundred trials with random initial conditions. The averages appear to lie on a straight line for each of the three parameter pairs tested. Although only three data sets are displayed, our tests with other parameters only yield a change in the slope of the resulting line. The inset in this Figure is shown with error bars to indicate the standard deviation of the averages. The standard deviation for the other data sets are similar in that they remain nearly constant after the chain length becomes larger than 20. We tested various combinations of α_I and I_0 in the ranges $\alpha_I \in [0.0025, 0.96]$ and $I_0 \in [1.01, 20]$. In Section 2.7 we discuss how these two parameters relate to the slopes of the lines shown in Figure 7.

Only several hundred trials were needed to compute the averages because simulations indicated that the distribution of the synchronization times did not appear to have a long tail. In Figure 8 we display the distribution of the synchronization times for a chain of 500 oscillators for 20000 trials with random initial conditions. The data however, is not sufficient enough to determine whether the tail is an exponential decay with an exponent between 1 and 2, or a an algebraic decay with an exponent of 5 or greater.

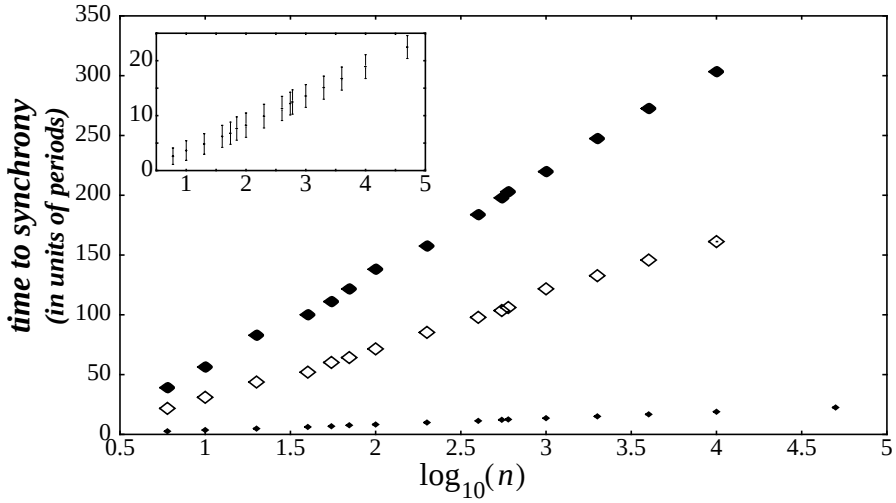

Figure 7. The average time needed for a chain of n oscillators to synchronize as a function of $\log_{10}(n)$. Three different symbols represent different parameters; black diamonds, $\alpha_I = 0.25, I_0 = 5.5$, open diamonds $\alpha_I = 0.025, I_0 = 1.1$, and small diamonds $\alpha_I = 0.2, I_0 = 1.11$. The data is based on approximately 300 trials with random initial conditions. The inset it displays the data for $\alpha_I = 0.2, I_0 = 1.11$ along with the standard deviation of the averages.

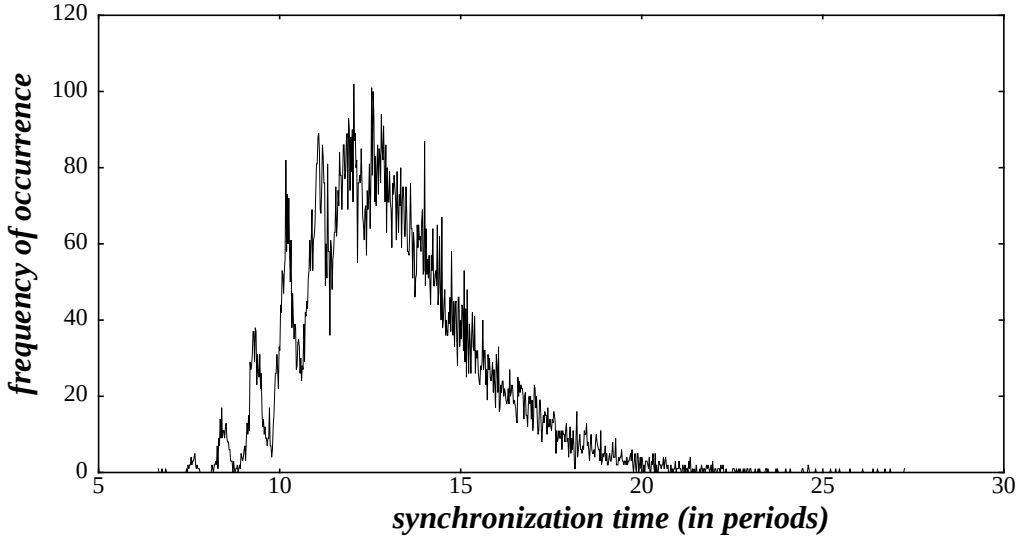

Figure 8. A histogram of the synchronization times for a one-dimensional chain of 500 oscillators. This data is based on 20000 trials with random initial conditions. The parameters used are $\alpha_I = 0.2, I_0 = 1.11$.

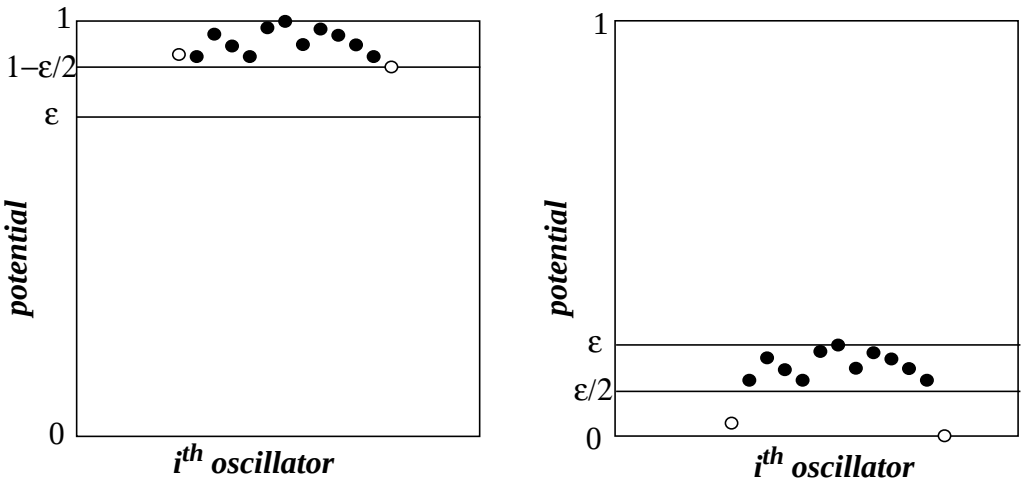

Figure 9. A block of oscillators immediately before (A) and after (B) they fire. The open circles indicate the two oscillators on the border of the block.

2.4.1 Block Size Only Increases

We now show that in one-dimensional networks, blocks of synchronous oscillators do not decrease in size. The first necessary step in arguing that blocks of oscillators do not break is to define a block of oscillators. A block of oscillators is a connected group of oscillators that has previously fired in unison. In Figure 9 we give an example of a block of oscillators. In Figure 9A one oscillator is at the threshold and when it fires, all of its neighboring oscillators are able to fire at the same time because they have an potential greater than 1–ε/2, or equivalently, the maximum phase difference between any two neighboring oscillators is less than

$$\varsigma = g(1) - g(1 - \varepsilon/2) \tag{2.6}$$

We are using the notation given in Section 2.3 where the potential of an oscillator is given by $x = f(\phi)$, where ϕ represents the phase of an oscillator, and the inverse is $\phi = g(x)$. After these oscillators have fired their values are reset using (2.4). Figure 9B shows the potentials of the oscillators immediately after the block has fired. Most of the oscillators in the block have two neighbors that also fire, and thus receive two pulses weighted by ε/2. The oscillators that receive two pulses cannot have an potential less than ε/2 after firing and these oscillators are the black filled circles in Figure 9B. The two oscillators at the ends of the block receive excitation from only one neighbor. These oscillator are distinguished by unfilled circles in Figure 9B and cannot have a value greater than ε/2. For all the oscillators in the bulk, the potential difference between them is maintained immediately before and after the jump. It is obvious that these oscillators will again jump together because the maximum difference between their potentials is less than ε/2. However, it is uncertain whether or not the two border oscillators will again fire with the block.

We examine how the two oscillators on the block boundary are altered when their neighbors which are not part of the block jump. Let us denote the oscillator on a border of a block as O_B, and let its neighbor which is not part of the block be O_C. We do not assume anything about O_C, but it must fire some time t^e in the range $t^e \in [0, g(1) - g(\varepsilon)]$ (we also assume that Figure 9B represents $t = 0$). The question that needs to be answered is, after O_C fires, can O_B still fire in unison with the block it originally fired with. There are only two cases that need to be examined. The first case occurs when O_B receives excitation from O_C and the resultant potential of O_B is greater than that of any oscillator in the block. We examine the worst case, when the greatest potential of any oscillator in the block has a value of just greater than $\varepsilon/2$ (the case of a block of size 2), and the potential of the other oscillator is just less than $\varepsilon/2$. We need to test the following relation,

$$g(f(t^e) + \varepsilon/2) - t^e < \varsigma \tag{2.7}$$

The LHS of this relation yields the phase difference between the bordering oscillator and oscillator with potential $\varepsilon/2$ after a time t^e has passed. The bordering oscillator receives excitation at time t^e. We first note that the derivative of the LHS is given by

$$g'(f(t^e) + \varepsilon/2)f'(t^e) - 1 \tag{2.8}$$

since $g'(f(t))f'(t) = 1$ we rewrite (2.8) as

$$\frac{g'(f(t^e) + \varepsilon/2)}{g'(f(t^e))} - 1 \tag{2.9}$$

Since $g'(f(t^e)) = 1/f'(t)$ we know $g'(f(t^e)) > 0$ and also $g''(f(t^e)) > 0$, which shows that (2.9) must be positive. This implies that the largest value of the LHS of (2.7) is given by $t^e = g(1 - \varepsilon/2)$ which yields ς. Since the actual value of t^e must be less than $g(1 - \varepsilon/2)$ (otherwise it implies that O_C had the same firing time as the block) the inequality of (2.7) is satisfied. Since the derivative is positive, the maximum value is given by this upper bound, which implies that the inequality is met for all t^e less than $g(1 - \varepsilon/2)$. The second case that must be examined is when the separation between oscillator in the bulk with the largest potential and O_B is at its maximum. In this case it might be possible for the block to fire without the bordering oscillator. We examine the following inequality

$$g(\varepsilon) + t^e - g(f(t^e) + \varepsilon/2) < \varsigma \tag{2.10}$$

The LHS represents the phase separation from the oscillator nearest to the threshold and the bordering oscillator. Again, we start with the derivative of the LHS, which is

$$1 - \frac{g'(f(t^e) + \varepsilon/2)}{g'(f(t^e))} < 0 \tag{2.11}$$

so the largest value of the LHS of (2.10) is given by $t^e = 0$, which yields $g(\varepsilon) - g(\varepsilon/2)$. This value is less than ς because $g(x)$ is an increasing function of x, therefore the bordering oscillator still fires with the block.

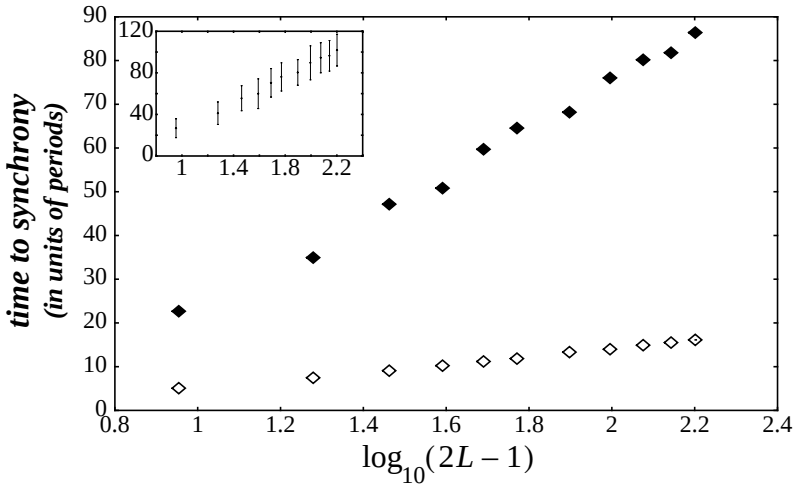

Figure 10. The average times for an $L \times L$ network of oscillators to synchronize are plotted as a function of $\log_{10}(2L-1)$. The solid diamonds are for the parameters $\alpha_I = 0.2, I_0 = 2.0$ and the open diamonds are for $\alpha_I = 0.2, I_0 = 1.11$. Each average is computed from approximately one hundred trials with random initial conditions. The inset indicates the standard deviation for one set of data.

We have shown that the bordering oscillators continue to fire with the block they originally fired with, regardless of the influences of its other neighbor. The potentials of the other oscillators in the block never change their respective ordering, although the relative distance between the oscillators within the block decreases because of the compression that occurs each period. Thus in one-dimension, blocks never break apart, and can only increase in size by merging with other blocks.

2.5 Two-dimensional systems of oscillators

We now present data indicating that the time to synchrony in two-dimensional networks of locally coupled integrate-and-fire oscillators increases as the logarithm of the system size. As mentioned previously, all trials with two-dimensional networks resulted in synchrony. We tested various size spin waves in the two directions with similar results to those in one-dimensional systems, namely synchrony was achieved regardless of the initial conditions. Travelling waves were never observed, even with periodic boundary conditions.

We display the average synchronization time for a two-dimensional system as a function of $\log_{10}(2L-1)$ in Figure 10. In this two-dimensional system, each oscillator is coupled to its four nearest neighbors and the longest distance between any two oscillators (in terms of lattice sites) is $2L$-1. The data indicate that the average time to synchronize scales logarithmically with the system size. We have tested more parameters than shown in Figure 10 and all tested parameters yield an identical scaling relation.

As we examine two-dimensional systems, a natural question is “How does the rate of synchronization vary as the dimension of the system changes?” We first define the rate of synchrony as follows. The data indicate that $\langle T_S \rangle \sim \frac{1}{r_s}\log(n)$, where $\frac{1}{r_s}$ corresponds to the slope of a line from Figure 7. We refer to r_s as the rate of synchronization. In tests where the value of α_I is held constant, but the dimension of the system changes from 1 to 2, we find that the rate of synchrony halves. When the individual coupling strengths between oscillators are maintained, i.e. α_I doubles, we find that the rate of synchronization remains approximately the same as the dimension of the system increases from 1 to 2. This indicates that the rate of synchrony is controlled by the individual connection weights between oscillators and not the total input to each oscillator.

2.5.1 Disorder in Integrate-and-Fire Oscillators

We have also tested this system with disorder in the form of different intrinsic frequencies, ω_i, of oscillators. The intrinsic frequencies are allowed to vary within a range $\omega_0 - \Delta < \omega_i < \omega_0 + \Delta$ (plus or minus 5 percent for example). With this type of disorder synchrony is still achieved in both one- and two-dimensional networks. Imagine that all the oscillators are initially given an potential of zero. The oscillator with the highest frequency reaches the threshold first and the rest of the oscillators have a smaller potential, but are still near the threshold. If the coupling strength is large enough, then the entire network can be induced to fire. So oscillators can repeatedly fire at the same time as long as their frequency difference is not too large. However, if one allows the intrinsic frequencies to have a Gaussian distribution, then the system behaves differently. If the width of the distribution is not too large then small one-dimensional systems can synchronize completely. As the system size becomes larger, then synchrony occurs less often. The system typically has several large synchronized clusters. Tests with two-dimensional systems yields a similar qualitative behavior. Unfortunately, we have not quantized any of the above observations and this remains a subject of study.

2.6 Heuristic argument for exponential block growth

Here we present an argument for the exponential increase in the block size as a function of time. We have not been able to directly relate the time to synchrony to the system size using equations (2.1)-(2.4). However, we are able to make the following heuristic arguments based on the two following assumptions; a block never decreases in size and its probability of merging with another block is independent of block size.

For mathematical simplicity we make the following arguments based on the breaking of blocks, which can be thought of as the time reversal of the synchronization process. We begin with a single block that is the size of the system, where *S* is the number of sites in the system. We use the following assumptions; a block never increases in size, a block can only break into two smaller blocks, when a block breaks into two pieces, there is no preference for the sizes of the resultant pieces. The probability of a block not breaking is given by $1 - \lambda$ and is independent of block size. Let $N_m(t)$ represent the number of blocks of size *m* at time *t*. With the above assumptions and notations, we write the following,

$$N_m(t+1) \;=\; (1-\lambda)\,N_m(t) + \sum_{j>m}^{S} P_{jm}N_m(t) \tag{2.12}$$

where P_{jm} is the probability that a block of size j breaks into block of size m and $j-m$. We write (2.12) as a continuous function of time

$$\frac{d}{dt}(N_m(t)) \;=\; -\lambda N_m(t) + \sum_{j>m}^{S} P_{jm}N_m(t) \tag{2.13}$$

Note that (2.12) and (2.13) are only valid for $m > 1$. For $m = 1$, the first term on the RHS of (2.13) disappears because blocks of size 1 cannot be broken. There is a conservation of the total number of sites, which can be written as

$$\frac{d}{dt}\left(\sum_{m=1}^{S} mN_m(t)\right) \;=\; \frac{d}{dt}(S) \;=\; 0 \tag{2.14}$$

Using (2.13) we rewrite (2.14) as

$$\frac{d}{dt}\left(\sum_{m=1}^{S} mN_m(t)\right) \;=\; -\lambda \sum_{m=2}^{S} mN_m(t) + \sum_{m=1}^{S} m \sum_{j>m}^{S} P_{jm}N_m(t) \;=\; 0 \tag{2.15}$$

In matrix form this becomes

$$\lambda \sum_{m=2}^{S} mN_m(t) \;=\; \begin{bmatrix} \cdots\; P_{21}N_2 & P_{31}N_3 & P_{41}N_4 & P_{51}N_5 & P_{61}N_6 & \cdots \\ & 2P_{32}N_3 & 2P_{42}N_4 & 2P_{52}N_5 & 2P_{62}N_6 & \\ & & 3P_{43}N_4 & 3P_{53}N_5 & 3P_{63}N_6 & \\ & & & 4P_{54}N_5 & 4P_{64}N_6 & \\ & & & & 5P_{65}N_6 & \\ & \cdots & & & & \end{bmatrix} \tag{2.16}$$

We set the individual terms N_k equal, yielding

$$\lambda k N_k(t) \;=\; \sum_{i=1}^{k-1} iP_{ki}N_k(t) \tag{2.17}$$

for $k > 1$. As mentioned previously, we assume that the probability of a block breaking into two pieces is independent of the size of the two smaller pieces (with the exception being the probability of block with an even numbered size breaking into two equal size

pieces). Because of this difference we compute the simplest case first, that of a block with an odd numbered size. In this case we assume $P_{km} = c_k$ where c_k is independent of m. With this assumption, (2.17) becomes

$$\lambda k = \sum_{i=1}^{k-1} iP_{ki} = c_k \sum_{i=1}^{k-1} i = c_k k(k-1)/2 \tag{2.18}$$

which leads to probabilities $c_k = 2\lambda/(k-1)$. For a block with an even numbered size, we assume that $P_{km} = d_k, m \neq k/2$ and $P_{k(k/2)} = d_k/2$. Substituting these assumptions into (2.17) yields $d_k = 2\lambda/(k-3/2)$. Now that the probabilities are known, we calculate other quantities. Of interest is how the average block size changes in time. Let μ be the average block size, or

$$\mu = \frac{\sum_{m=1}^{S} mN_m(t)}{\sum_{m=1}^{S} N_m(t)} = \frac{S}{\sum_{m=1}^{S} N_m(t)} \tag{2.19}$$

The derivative of μ with respect to time is

$$\frac{d\mu}{dt} = \frac{-S}{\left(\sum_{m=1}^{S} N_m(t)\right)^2}\left[\frac{d}{dt}\left(\sum_{m=1}^{S} N_m(t)\right)\right] = -\mu\left[\frac{\frac{d}{dt}\left(\sum_{m=1}^{S} N_m(t)\right)}{\sum_{m=1}^{S} N_m(t)}\right] \tag{2.20}$$

We solve for the derivative of the total number of blocks by using (2.13), which results in

$$\frac{d}{dt}\sum_{m=1}^{S} N_m(t) = -\lambda \sum_{m=2}^{S} N_m(t) + \sum_{m=1}^{S}\sum_{j>m}^{S} P_{jm}N_m(t) \tag{2.21}$$

Using the appropriate probability values derived above, and performing the sum of the second term on the RHS of (2.21) yields

$$\frac{d}{dt}\sum_{m=1}^{S} N_m(t) = -\lambda \sum_{m=2}^{S} N_m(t) + 2\lambda \sum_{m=2}^{S} N_m(t) = \lambda \sum_{m=2}^{S} N_m(t) \tag{2.22}$$

substituting this in (2.20) yields

$$\frac{d\mu}{dt} = -\lambda\mu\left[\frac{\sum_{m=2}^{S} N_m(t)}{\sum_{m=1}^{S} N_m(t)}\right] \tag{2.23}$$

There are three cases to consider for this equation. The first case is if there are no blocks of size one, or $N_1(t) = 0$, then the bracketed term in (2.23) is equal to one and the average block size decreases exponentially. The second case occurs if there are some blocks of size one, then the bracketed term in (2.23) is less than one, and the rate of change of the average blocks size is an exponential modified by a constant. The third case is when all the blocks are of size 1, then the bracketed term is zero and the average block size does not change. This last case makes sense because if all the blocks are of size one, then there can be no further changes in the average block size.

For the set of assumptions used, we obtain an exponential decrease in the block size until the smallest block size is attained. Since our equations describe the breaking of blocks, the inverse time evolution of the system is the merging of blocks and yields an exponential growth in the average block size. The linkage between this argument and integrate-and-fire systems is relatively weak. The above argument contains no topology and can be seen as a mean field model because there is no correlation between blocks.

In two-dimensional networks numerical simulations indicate that blocks can break apart. This seemingly qualitative change in behavior does not appear to alter the logarithmic relation between the time to synchrony and the system size. Attempts to incorporate this dynamic into (2.12) have not been fruitful. The addition of a positive nonlinear term to the RHS of (2.12) (implying that the block size can increase) yields equations that are not readily solved. However, the terms leading to exponential decay are still present and they might dominate the behavior of the average block size.

2.7 The Rate of Synchrony in Integrate-and-Fire Oscillators

We present data indicating that the rate of synchrony is approximately proportional to $\alpha_I/(I_0-1)$ for a specified range of parameter values. Remember that we have defined the measured rate of synchronization as the inverse of the slope of a line from Figure 7. We first consider the qualitatively consider the coupling strength, α_I, for a pair of oscillators. As the coupling strength goes to zero, the average time to synchrony must become infinite and r_s approaches zero. Also, with $\alpha_I = 0$ the return map is straight line with slope 1. A first guess is that the coupling strength is proportional to the rate of synchrony. Note that this cannot be true for large values of the coupling strength. As α_I nears 1, the period approaches 0; the oscillators fire frequently, but only slowly change their relative phase. As α_I nears 1 the time to synchrony must become infinite. At $\alpha_I = 1$ system (2.1) is meaningless because the oscillators are constantly firing and resetting. Now we examine I_0, which essentially controls the period and is related to the curvature of the potential curve, which is given by $I_0e^{-t}/(1+I_0^2e^{-2t})$. For values of $I_0 > 10$ the potential of the oscillator rises quickly to the threshold and the potential curve is nearly linear. For smaller val-

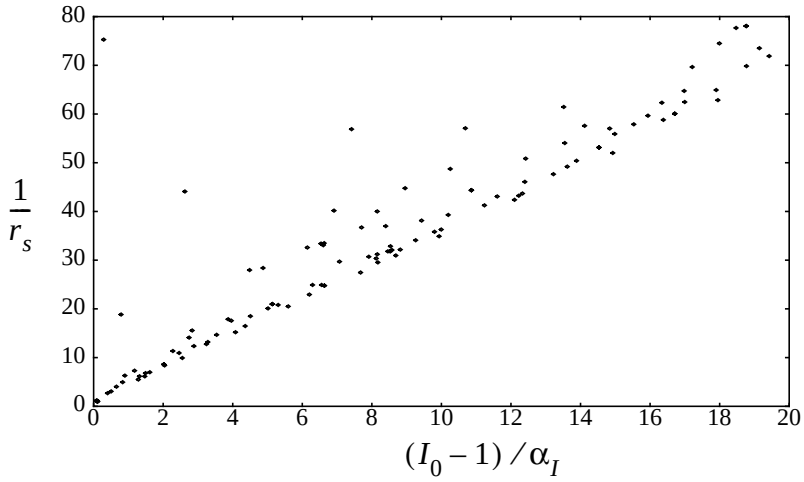

Figure 11. We display the inverse of the measured rate of synchrony as a function of $(I_0 - 1) / \alpha_I$. Most of the 140 different parameter pairs are near a straight line. An examination indicates that those points not near the line have values of $\alpha_I > 0.8$ and/or values of $I_0 < 1.05$. Each point represents the average time to synchrony calculated from 250 trials with random initial conditions.

ues of I_0, the potential curve become more nonlinear (see Figure 4 for example). A perfectly linear slope would yield a return map with a slope of 1 (except for the jumping regions) and a pair of oscillators would have a constant phase difference. As I_0 approaches 1, the curvature becomes highly nonlinear and yields a larger compression of the phase difference for each period. Our first guess is that the rate of synchrony is inversely proportional to $I_0 - 1$. We mention that Hopfield and Herz [Hopfield and Herz, 1995] used a value of $I_0 = 10$ in their simulations and this is the reason that they saw rather slow convergence times to synchrony. Both Peskin [Peskin, 1975] and Mirollo and Strogatz [Mirollo and Strogatz, 1990] note that the coupling strength and the degree of nonlinearity in the potential curve are important to the rate of synchronization.

In order to examine how the rate of synchrony varies with these two parameters, we performed trials with several different size chains, with randomly chosen values of $\alpha_I \in (0, 1)$ and $I_0 \in (1, 11)$. The size of the chains varied from 100 to 1000. In Figure 11 we display the inverse of the measured rate of synchrony as a function of $(I_0 - 1) / \alpha_I$. The data lie near a line, indicating that this simple first guess is not a bad approximation. An examination of the points not on the line indicates that they have values of $\alpha_I > 0.8$ and/or $I_0 < 1.05$.

A more direct evaluation of how the coupling strength alters the rate of synchrony is shown in Figure 12 where we vary only the coupling strength. The rate of synchrony is approximately proportional to the coupling strength for $\alpha_I < 0.8$. Similar tests in which I_0 alone is varied yield synchronization rates that vary approximately inversely with $I_0 - 1$. Figure 13 displays this relationship for two different values of the coupling strength.

We also attempted to determine how the values of the parameters relate to the rate of synchronization. We reasoned that the return map Figure 5 between two oscillators give one some idea of how fast two oscillators might synchronize. If the coupling strength is

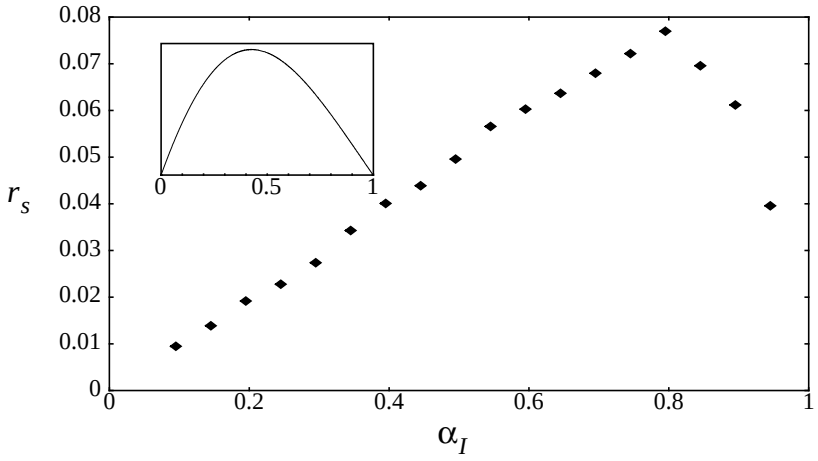

Figure 12. A plot of the measured rate of synchrony as a function of the coupling strength. The linear relationship between the coupling strength and the time to synchrony is maintained for $\alpha_I < 0.8$. A value of $I_0 = 5.0$ was used. The inset indicates the rate of synchrony predicted from the return map for two oscillators. We used a chain length of 500 oscillators. Each point represents the average over several hundred trials with random

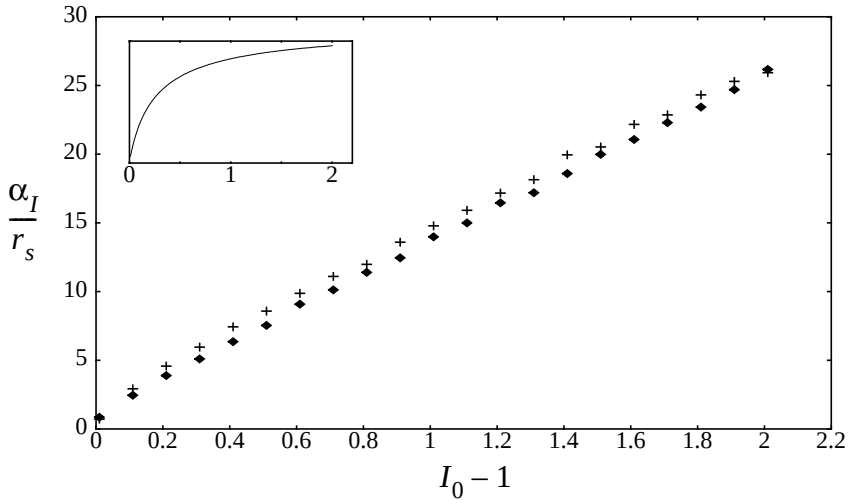

Figure 13. The inverse of the measured rate of synchrony (times α_I) is plotted as a function of I_0. The filled diamonds are for $\alpha_I = 0.25$ and the plus symbols are for $\alpha_I = 0.8$. The data is from a chain of 1000 oscillators using 100 trials with random initial conditions. The inset represents the rate of synchrony predicted from the return map for two oscillators.

zero, then the return map is a straight line of slope one and synchrony never occurs - also, the area under the return map, A_I, is equal to 0.5. As the coupling strength increases, A_I becomes less than 0.5 and the two oscillators synchronize. At a qualitative level, there is some link between A_I and the rate of synchrony, r_s. In the inset shown in Figure 12, we plot $1/A_I - 2.0$ as a function of the coupling strength. This theoretical curve does share several things in common with the numerically obtained data in that the endpoints are correct and there is only a single maximum. However, the location of this maximum is significantly different from the maximum given by the numerical data, indicating that this approach to deriving r_s is flawed. For completeness we also calculate A_I as a function of I_0 and display this graph in the inset of Figure 13. These graphs exhibit a much different shape and given further indication that this is not the correct route for analytically deriving the rate of synchrony. The area under the return map, A_I cannot be calculated analytically for the particular equations we use (see Section 2.3), although one can derive a power series expansion for the area, which results in

$$\begin{aligned} A_I \;=\; & -\phi_L \log\,(I_0-\alpha_I) \;+ S(1-\alpha_I,\;(I_0-1)\exp\,(\phi_L)) + S(1-\alpha_I,\;(I_0-1)) - \\ & (\phi_U-\phi_L)\log\,(I_0) - S(I_0\,(I_0-1)\;-\alpha_I^2,\;\alpha_I\,(I_0-1)\exp\,(\phi_U)) + \\ & S(I_0\,(I_0-1)\;-\alpha_I^2,\;\alpha_I\,(I_0-1)\exp\,(\phi_L)) + S(\alpha_I,\;\alpha_I\,(1-I_0)\exp\,(\phi_U)) - \\ & S(\alpha_I,\;\alpha_I\,(1-I_0)\exp\,(\phi_L)) - (P-\phi_U)\log\,(\,(I_0-\alpha_I)\;+\;\,) \\ & S(\,(1-\alpha_I)\,\alpha_I + I_0\,(1-I_0),\;(1-\alpha_I)\;(I_0-1)\exp\,(P)) - \\ & S(\,(1-\alpha_I)\,\alpha_I + I_0\,(1-I_0),\;(1-\alpha_I)\;(I_0-1)\exp\,(\phi_U)) - \\ & S(\alpha_I,\;(1-I_0)\exp\,(P)) + S(\alpha_I,\;(1-I_0)\exp\,(\phi_U)) \end{aligned} \tag{2.24}$$

where

$$\begin{aligned} S(a,x) \;=\; & \log\,(a)\log\,(x) \;+ \frac{x}{a} - \frac{x^2}{2^2a^2} + \frac{x^3}{3^2a^3} - \frac{x^4}{4^2a^4} \quad \ldots \qquad \left[x^2 < a^2\right] \\ \;=\; & \frac{(\log\,(x))^2}{2} - \frac{a}{x} + \frac{a^2}{2^2x^2} - \frac{a^3}{3^2x^3} + \frac{a^4}{4^2x^4} \quad \ldots \qquad \left[x^2 > a^2\right] \end{aligned} \tag{2.25}$$

2.8 Error Correction

We now examine how errors are corrected in noiseless, locally coupled two-dimensional networks of identical integrate-and-fire oscillators. We start with initial conditions for a two dimensional network such that a small central region of oscillators has random

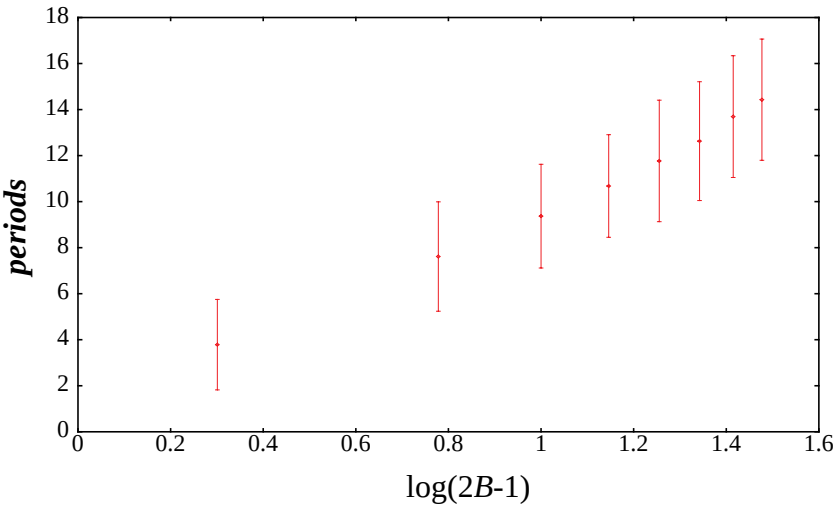

Figure 14. The time needed to correct a given size error in a two-dimensional network of integrate-and-fire oscillators. The error is a right isosceles triangle region that contains oscillators with random initial conditions. B represents the size of the base of the triangle.

potentials and the oscillators in the larger surrounding area are given the same potential. The small region within the larger synchronized block of oscillators is considered the "error". This error will take some amount of time to synchronize with the larger block and we call this the correction time. This particular form of error was found to take the longest time to correct. A rectangular block of oscillators in the center of the system with initial conditions such that its phase relationship with the rest of the network is near the unstable fixed point, shown in Figure 5A, results in synchronization on the same time scale as just two oscillators with the same phase relationship. This occurs because all the blocks in the error have the same potential and changes that occur on the boundary can propagate through the entire region in an instant. Other types of errors (spin waves) were also found to be quickly corrected, so a central block of random potentials was deemed the most difficult generic error to correct.

We now examine the average time needed to correct a central block of oscillators with random potentials. Numerical simulations indicate that the error is corrected at times that increase linearly with the logarithm of the perimeter of the smallest bounding rectangle. In Figure 14 we display the time needed to synchronize a triangular error (random potentials of oscillators surrounded by a synchronous grid of oscillators) as a function of the logarithm of the smallest square enclosing the error. The data indicate an error correction time that scales as the logarithm of the error size. This indicates that the system is robust with respect to errors.

Given the rigid lattice structure of this system, each oscillator is coupled to its four nearest neighbors, it is possible that differently shaped regions of error may exhibit different correction times. In tests with circular and triangular regions of error, we noted that these regions typically expanded in size until to the smallest rectangle enclosing the original error.

2.9 Desynchronization in Integrate-and-Fire Oscillator Networks

The integrate-and-fire networks examined so far have the property that synchrony is quickly achieved. But a system that only achieves synchrony is not very useful for information processing, since such a system is dissipative and almost all information is lost. In order to perform computations, some other mechanisms must exist that can store or represent information. In oscillatory correlation, the different phases of oscillators encode binding and segregation information.

We now describe a network of integrate-and-fire oscillators which can desynchronize different groups of synchronous oscillators. In order to create a network of integrate-and-fire oscillators for oscillatory correlation, we need a mechanism which desynchronizes different oscillator groups. This mechanism will be connected with every oscillators since the phases of different oscillator groups need to be desynchronous regardless of their position in the network. This unit is called the global inhibitor and the architecture of this network is identical to the locally excitatory globally inhibitory oscillator networks (LEGION) proposed by Terman and Wang [Terman and Wang, 1995]. Figure 15A displays a diagram of the LEGION architecture.

We now define a LEGION network which uses integrate-and-fire oscillators as its basic units. The activity of each oscillator in the network is described by

$$\dot{x}_i = -x_i + I_i + \sum_{j \in N(i)} J_{ij} P_j(t) - G(t), \qquad i = 1, \ldots, n \tag{2.26}$$

where n is the number of oscillators. $N(i)$ represents the four nearest neighbors of oscillators i. The parameter, I_i, is now dependent on the input image; we refer to this parameter as the stimulus given to an oscillator. In this section we segment a toy image that consists of binary pixels. The respective stimulus for each oscillator is either $I_i > 1$ or $I_i = 0$. If $I_i > 1$ we call oscillator i stimulated. If an oscillator does not receive stimulus, $I_i = 0$, its potential decays exponentially towards zero. As before the threshold for each oscillator is 1. The interaction term, $P_j(t)$, is the same as in (2.2). Only neighboring oscillators that both receive stimulus have a nonzero coupling strength. Thus, image information is contained in the connection weights between oscillators. Only oscillators that are part of the same connected image region can be coupled. This is also an implicit encoding of the Gestalt principle of connectedness [Rock and Palmer, 1990]. The connection strengths are normalized so that all stimulated oscillators receive the same sum of connection input and have the same frequency. However, we use a slightly modified version of (2.3) to reflect that the input to oscillator i is now normalized by the number of stimulated neighbors coupled with i.

The global inhibitor, $G(t)$, sends an instantaneous inhibitory pulse to the entire network when any oscillator in the network fires. It is defined as

$$G(t) = \Gamma \delta(t - t_j^m), \qquad \forall j, m \tag{2.27}$$

where t_j^m represent the m firing times of the j^{th} oscillator. The constant Γ is less than the smallest coupling strength. When an oscillator fires, the global inhibitor serves to lower the potential of all oscillators, but because this impulse is not as large as the excitatory signal between neighboring oscillators, it does not destroy the synchronizing effect of the local couplings (see [Terman and Wang, 1995]). In this fashion, a connected region of oscillators receiving input synchronizes as the system evolves in time. This region of oscillators has no direct excitatory connections with other spatially separate oscillators. It will however, interact with other groups of oscillators through the global inhibitor. This interaction inhibits other blocks of oscillators from firing at the same time and ensures that there is a finite amount of time between firings of synchronous groups of oscillators.

We now demonstrate the ability of this network to perform oscillatory correlation. In Figure 15B we display an input image and in Figure 15C we display the network response. The four graphs in Figure 15C display the combined potentials of every oscillator comprising each of the four objects. The oscillators have random initial conditions varying uniformly from 0 to 1. Initially, many oscillators fire and the effect of the global inhibitor can be seen in the jitter, or lack of smoothness, in the potentials of the oscillators during this time. As the system evolves, clusters of oscillators begin to form and the curves become smoother because the global inhibitor does not send inhibitory impulses as often. By the third cycle each group of oscillators comprising a distinct object is almost perfectly synchronous and the different oscillator groups have distinct phases. Oscillators that do not receive excitation (not shown) experience an exponential decay towards zero and are periodically perturbed by the small inhibitory signals from the global inhibitor.

In this network, there is an unlimited number of oscillator groups that can be segmented. In other words, the segmentation capacity is infinite. Imagine two groups of oscillators that have nearly the same phase. When the first group fires, the potential of the second group of oscillators decreases by Γ. Thus, the second group needs to traverse this distance Γ before it can fire. This implies that there is a finite amount of time between the firings of two consecutive groups. This also implies that as the number of groups increases, the period of the system increases. Simulations support the above statements and we have segmented more than 100 groups of oscillators.

2.10 Image Segmentation Using Integrate-and-Fire Oscillators

In the previous Section, we segmented four black objects on a white background in a small 20×20 image. Since numerical evidence suggests that there is a logarithmic scaling relation between the time to synchrony and the network size, we should be able to use this same network to quickly (in a few periods) perform image processing tasks with real images.

In order to segment gray-level images, we alter how the connection weights and values of I_i are chosen. The alterations made are variations of methods proposed in Wang and Terman [Wang and Terman, 1997]. Let the intensity of pixel i be denoted by p_i. If $|p_i - p_j|$ is less than a given threshold, then the two pixels are said to satisfy the *pixel difference test*. Two oscillators have a nonzero coupling strength only if they are neighbors (we now use the eight nearest neighbors of i) and if their corresponding pixel values satisfy the pixel difference test. The weights of the connection strengths are determined using

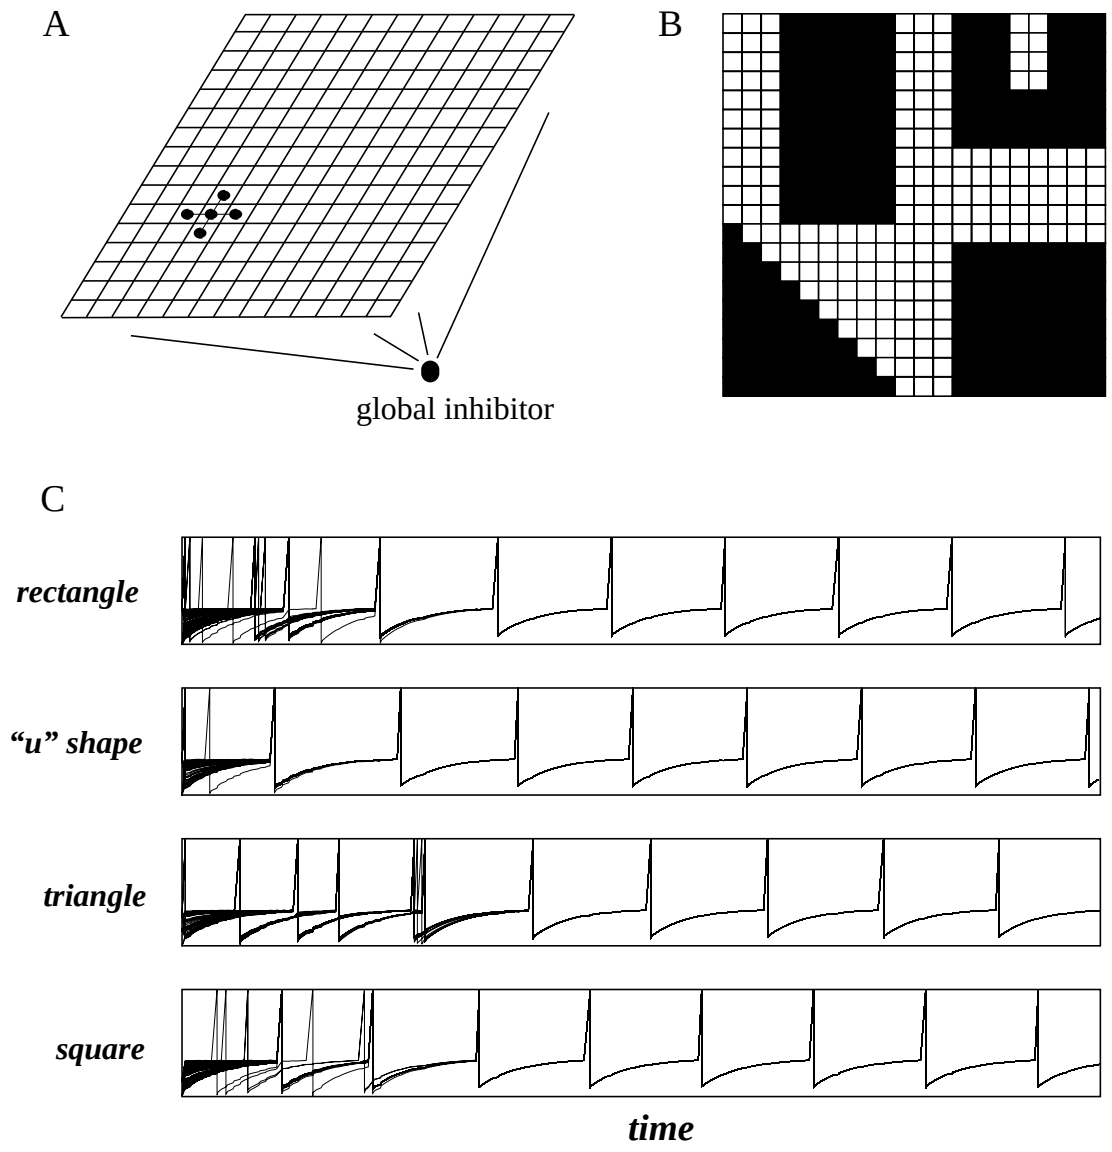

Figure 15. (A) A diagram of the network architecture. Each unit has local excitatory connections. The global inhibitor is coupled with every unit in the network and serves to desynchronize different groups of oscillators. (B) The input we use to demonstrate the behavior of our network. The black squares represent those units which receive stimulus and the units corresponding to the unfilled squares receive no stimulus. (C) We display the temporal activities of all units comprising each of the four objects in (B). The parameters used are $I_i = 1.05$ for those oscillators receiving stimulus, $\alpha_I = 0.2$, and $\Gamma = 0.01$.

(2.3), except that Z_i now represents the number of neighboring pixels of i that pass the pixel difference test. The stimulus I_i for each oscillator is chosen in the following manner. We examine a region $Q(i)$ centered on pixel i. $Q(i)$ is a neighborhood about oscillator i that contains more pixels than $N(i)$. If half of the pixels in $Q(i)$ satisfy the pixel difference test, then pixel i is possibly within a homogeneous region and we set the stimulus, I_i, to a value I_L which is greater than 1. Such an oscillator is called a *leader* [Wang and Terman, 1997] and is able to oscillate. If the region about pixel i contains no pixels that satisfy the pixel difference test, then we assume that this region is noisy, or contains high intensity variations. The corresponding oscillator i, receives no stimulus, and does not oscillate. If pixel i is not a leader, but its surrounding region contains one or more pixels that satisfy the pixel difference test, oscillator i is given a stimulus I_N, which is less than, but near 1, and is said to be a near threshold oscillator. A near threshold oscillator is able to fire only through interactions with other oscillators. In this fashion, only regions of sufficient size, and with smoothly varying intensities will contain oscillators which are leaders. These leaders will oscillate, and can induce neighboring oscillators that are near threshold to oscillate. The border of a region will consist of pixels whose neighboring pixels have sufficiently different grey-levels that they do not pass the pixel difference test and are not connected. Only sizeable regions of smoothly varying intensities will be connected and oscillatory. Regions with high intensity variations will not exhibit oscillatory activity and are referred to as the background.

The rules for the connection weights and oscillator stimuli described above have been implemented in an integrate-and-fire oscillator network and we display the segmentation results for two real images in Figure 16. Figure 16A displays an aerial photograph. In Figure 16B we display the segmentation results of our network. Each group of synchronous oscillators is represented by a single grey-level intensity. Inactive oscillators comprising the background are colored black. There are 29 regions shown, although it is not easy to discern every different grey-level. We also segment a CT (computerized tomography) image of a slice of a human head. The original grey-level image is shown in Figure 16C. The bright areas indicate bone structure. Our segmentation result is shown in Figure 16D and contains 25 different regions. The different bones are segmented, except for two of the smaller bones which do not contain many pixels. Regions of soft tissue are also segmented. These results are comparable with those in [Wang and Terman, 1997, Shareef et al., 1997], which used similar methods and images in their oscillator networks. Shareef et al. [Shareef et al., 1997] used a similar CT image and obtained almost three times as many different regions as we did; many of their regions were small however. Through appropriate parameter choices, we could also obtain more and smaller regions.

This segmentation process has conceptual relations to region growing algorithms in computer vision. In region growing, a region is grown outward from some chosen seed. Pixels on the border of a region are merged with that region if their grey-level intensities satisfy some conditions based on the existing region, i.e. the pixel intensity is within some threshold based on the average pixel intensity of the region. Some of these techniques (in particular [Adams and Bischof, 1994]) are similar to our network if one removes the dynamical aspects of oscillatory correlation and replaces it with an algorithmic labelling procedure. Many different segmentation algorithms exist. They all suffer from the same problems in that the segmentation results are parameter sensitive and judgement of the

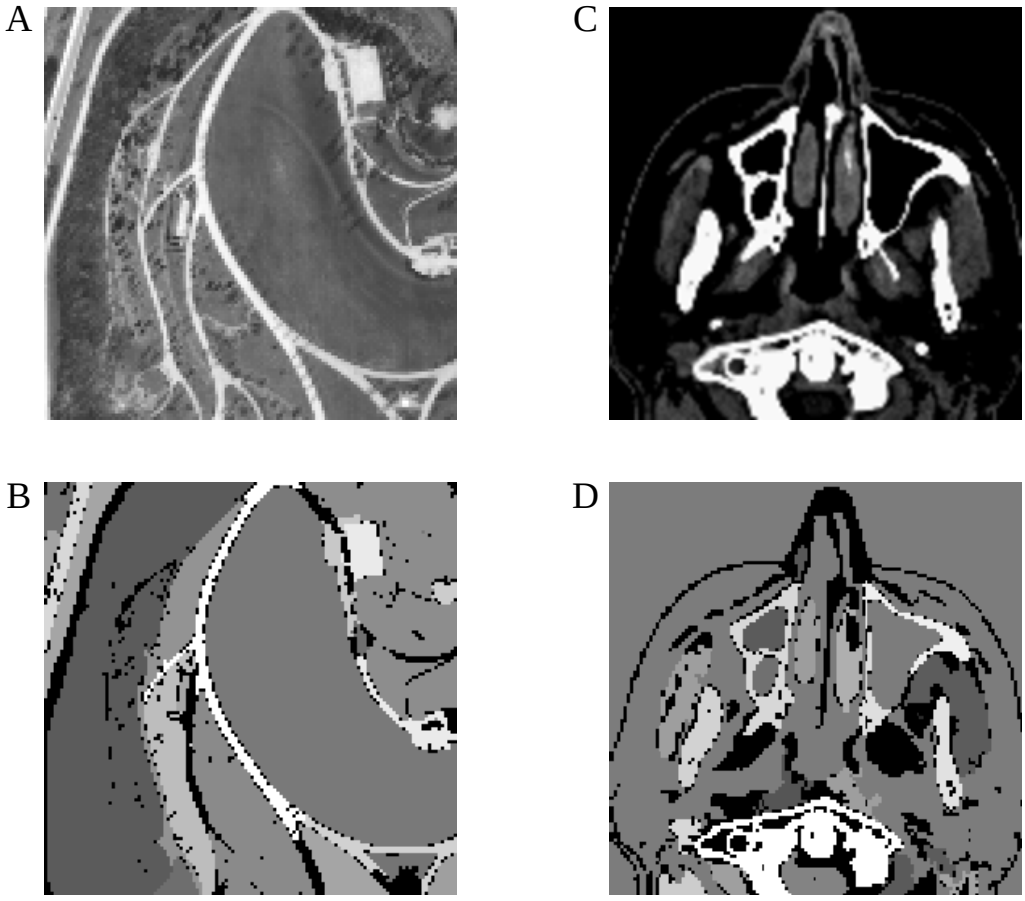

Figure 16. (A) An aerial image with 128×128 pixels and (B) the segmentation results for (A). The network produced 29 different synchronized groups. Each synchronized group is represented by a single grey-level. Black pixels represent those oscillators that do not exhibit periodic activity. The threshold for the pixel difference test is 19, $Q(i)$ is a region of size 7×7, with $I_i = 1.025$, $I_N = 0.99$, $\alpha_I = 0.2$, and $\Gamma = 0.01$. (C) A 128×128 CT image of a slice of a human head (the bright regions indicate bone). (D) The segmentation results for (C). The network produced 25 different groups of synchronized oscillators. Each synchronized group is represented by a single grey-level. Black pixels represent those oscillators that do not exhibit periodic activity. The threshold for the pixel difference test is 15 and $Q(i)$ is a region of size 9×9 and the other parameters are as listed above.

results is subjective. The advantage that our image segmentation network has over others is in its neurobiological basis and in the parallel and distributed nature of computation. The oscillators all operate in parallel and segmentation computations are made on a local basis (the global inhibitor does not make any decisions regarding which pixel belongs to which region, it merely assists in labelling the regions). This allows for comparisons with the distributed perceptual processes of the brain, and it also allows for VLSI implementation.

2.11 Discussion

Our investigation of locally coupled networks of identical integrate-and-fire oscillators has revealed a number of interesting properties which may have significant computational implications. Among the most interesting is that the time needed to achieve synchronization appears to scale logarithmically with the size of the system for one- and two-dimensional noiseless systems. Other properties of the system also appear to scale similarly. When an error is introduced into a synchronized system, in the form of a small region of oscillators with randomly chosen potentials, numerical results indicate that synchrony is attained at times proportional to the logarithm of the perimeter of the smallest square enclosing the error. This indicates a quick correction of faults and makes this oscillator appealing for use in networks where synchronization is a desirable feature.

If we introduce disorder in the form of different intrinsic frequencies, ω_i, of oscillators, and the distribution of the frequencies is $\omega_0 - \Delta < \omega_i < \omega_0 + \Delta$, then synchrony is still possible. This is quite different from a network of phase oscillators with diffusive coupling, in which synchrony is not possible using the same distribution of intrinsic frequencies.

It is interesting to note that slightly different versions of this network can yield drastically different behaviors. Corral et al. [Corral et al., 1995a] examine an almost identical oscillator network and derive parameter regimes in which three different behaviors arise, self organized criticality, periodic with a few clusters of oscillators, or globally synchronous. If we do not normalize the connection strengths, the oscillators along the borders have different periods from oscillators in the bulk. It appears that these different frequencies cause different dynamics to arise in these networks [Middleton and Tang, 1995, Corral et al., 1995a].

Hopfield and Herz [Hopfield and Herz, 1995] numerically examined several different types of integrate-and-fire oscillators and found that they attained synchrony on long time scales (100's of periods). Because of this, they discount synchrony as too slow to be useful in biological computations and instead use a network capable of local synchrony to perform image segmentation. Their network suffers from several problems. Because they rely on local synchrony, large homogeneous grey-level regions can be arbitrarily broken into several regions. Another problem is that they ignore the issue of desynchronization, and as a result different objects can have the same firing time and thus are incorrectly grouped.

In our investigation of integrate-and-fire oscillator networks, we duplicated the results of Hopfield and Herz [Hopfield and Herz, 1995] and found that by modifying the parameters, we can obtain synchrony on short time scales (a few periods). Because locally coupled networks of integrate-and-fire oscillators are able to achieve synchrony quickly, we

examined their potential as a feature binding network. We found that using the LEGION architecture [Terman and Wang, 1995] we are able to create an oscillator network for image segmentation. Our network is able to synchronize a large homogeneous region of smoothly varying grey-levels without breaking it into small regions. The LEGION architecture contains a global inhibitor, which serves the purpose of desynchronizing different groups of oscillators, while maintaining synchrony within each group of oscillators. This ensures that different oscillator groups have distinct firing times. This property allows for a temporal labelling of perceptually distinct objects as proposed by theoreticians [Milner, 1974, von der Malsburg, 1981] and we think that this is a promising framework for representing the results of feature binding.

There are differences between the network proposed here, and the network studied in [Terman and Wang, 1995]. A major difference is the segmentation capacity, or the number of different objects which can be desynchronized. Their network has a finite limit on the number of groups that can be desynchronized. In our integrate-and-fire network the period increases as the number of oscillator groups increases. Since there is no consequence of lengthening the period, there appears to be no limitation on the number of groups that can be desynchronized. This may be computationally more desirable, but is not psychophysically plausiable, since the number of objects that can be attended to is generally considered to be small, 7 plus or minus 2 [Miller, 1956].

Our results indicate that the form of the interaction between oscillators results in significantly different behaviors. We observe that with a pulsatile interaction integrate-and-fire oscillators synchronize at times proportional the logarithm of the network size (in one- and two-dimensions). If the coupling is a continuous function of the phase variables (diffusive for example) then the time to synchrony exhibits a completely different relation with the size of the network [Niebur et al., 1991b]. Also, the discontinuity in the interaction is partially responsible for the ability of the network to achieve synchrony in the presence of some forms of disorder. These results are in agreement with two other papers indicating that oscillator networks with discontinuous interactions have fundamentally different properties of synchronization than equivalent networks with continuous interactions. Daido [Daido, 1993b] achieves partial synchronization in a globally coupled network of phase oscillators with a normal distribution of frequencies using a step-like interaction. This same network with a continuous interaction achieves partial phase-locking (the oscillators have the same frequency but maintain a constant phase difference), but not synchronization. Somers and Kopell [Somers and Kopell, 1993] reported that a Heaviside type coupling in locally coupled networks of relaxation oscillators resulted in fast synchrony and they conjectured that the time to synchrony scales linearly with the size of the system. In Chapter 3 we further indicating that sinusoidal type oscillators also synchronize at times proportional to the network size with a Heaviside type interaction. This information not only indicates some generic classifications of the interactions in oscillator networks, but it may have possible practical applications as well. Several devices (Josephson junctions and resonance tunneling diodes [Young et al., 1988]) oscillate at megahertz frequencies. However, the output current of both of these devices is very small. There is interest in creating devices that can oscillate at these high frequencies and have a substan-

tial current output as well. At the moment, it is unknown how to synchronize the outputs of locally coupled arrays of these devices. Even a partial understanding of how to achieve quick synchrony in the presence of disorder and noise would be highly valuable.