

CHAPTER 1

INTRODUCTION

1.1 Coupled Oscillators

In 1665, Christiaan Huygens [Huygens, 1673] noted “*When we suspended two clocks so constructed from two hooks imbedded in the same wooden beam, the motions of each pendulum on opposite swings were so much in agreement that they never receded the least bit from each other and the sound of each was always heard simultaneously. Further, if this agreement was disturbed by some interference, it reestablished itself in a short time. For a long time I was amazed at this unexpected result, but after a careful examination finally found that the cause of this is due to the motion of the beam, even though this is hardly perceptible. The cause is that the oscillations of the pendula, in proportion to their weight, communicate some motion to the clocks. This motion, impressed onto the beam, necessarily has the effect of making the pendula come to a state of exactly contrary swings if it happened that they moved otherwise at first, and from this finally the motion of the beam completely ceases.*” The study of coupled oscillators has since become an active branch of mathematics, with applications in physics, biology, and chemistry. In physics, one encounters coupled oscillators in arrays of Josephson junctions [Chung et al., 1989, Blackburn et al., 1994], in modelling molecules [Sage, 1994], and in coupled lasers [Dente et al., 1990]. Coupled oscillators are also prevalent in biological systems. Most organisms appear to be coupled to various periodicities extant in our surroundings, such as the rotation of the earth about the sun, the alternation of night and day, or the tides. Not only do organisms exhibit periodicities due to their environment, but they also exhibit innate periodic behavior. Breathing, pumping blood, chewing, and galloping are examples of rhythmic patterns of motion.

One of our main concerns is understanding how synchrony is achieved in locally interacting populations of oscillators, where synchrony means that oscillators have the same frequency and phase. Some examples of biological synchrony occur between organisms, such as the synchronous flashing of fireflies in Southeast Asia [Smith, 1935], and the synchronous chirping of crickets [Walker, 1969]. Other types of biological synchronization occur between cells, such as the synchronous firing of the pacemaker cells of the heart [Peskin, 1975] and the synchronous neural activity seen in many different regions of the brain [Singer and Gray, 1995].

Our work is motivated specifically by synchronous neural activity observed in the brain. Synchrony in neural activity has been observed in several mammals [Singer and Gray, 1995], as well as in amphibians [Prechtl, 1994] and insects [MacLeod and

Figure 1. Synchronization in cortical and thalamic cells. The first three graphs are simultaneous extracellular recordings of neurons from cortical areas 5, 7, and the thalamus. The fourth plot is of the simultaneously measured EEG. The three plots at the bottom are cross-correlations of the firing times of pairs of the three different recordings of neurons.

Laurent, 1996]. Synchronous neural activity has been measured across 14 mm in the monkey motor cortex [Murthy and Fetz, 1992] and across different hemispheres of the brain in the visual cortices of cats [Engel et al., 1991]. The wide range of brain regions which exhibit synchronous activity, and the diversity of organisms in which synchrony is observed, suggest that neural synchronization may be fundamental to biological information processing.

To illustrate the phrase neural synchrony, we display data from Steriade [Steriade, 1994] in Figure 1, which shows several different recordings from three different brain regions of a cat. The top three graphs represent the firing times of neurons in cortical area 5, cortical area 7, and the thalamus. A simultaneously recorded electroencephalogram (EEG) is the fourth graph in the figure and the bottom row of graphs display the cross-correlations between the firing times of the neurons from the three different brain regions. The cross-correlation has a peak near zero indicating near perfect synchrony in the firing times of these cells. The peaks in this cross-correlation decrease to zero within a few periods, indicating that the period is not fixed. This figure demonstrates synchrony of the slow rhythm (approximately 0.6 Hz) which occurs during sleep.

The dynamical systems we study are based on simplified models of neuronal oscillators. While these oscillator networks have relations to neurophysiology, their links to physics are less straightforward. The concepts of energy and temperature in these networks resist definition and these systems are at best classified as non-equilibrium systems. Also, we study noiseless systems, and the mathematical techniques of equilibrium statistical mechanics are difficult to apply. However, many concepts useful in describing and understanding equilibrium systems, like spin systems, are useful in attempting to understand these non-equilibrium biological networks.

1.2 The Feature Binding Problem

Theoreticians have proposed that the temporal correlation of neural groups may be an efficient representation for feature binding [Milner, 1974, von der Malsburg, 1981]. In order to fully describe the feature binding problem, the following background is provided. First, neurons in the visual cortex appear to be functionally specialized. For example, some neurons respond best to stimuli of a certain orientation and speed, while other neurons respond best to certain color combinations [Zeki, 1993]. In later stages of the visual hierarchy, there appear to be entire cortical regions which are also functionally specialized, i.e. they are devoted to the perception of specific visual features such as color, motion, or depth [Zeki, 1993, Kandel et al., 1991]. Several clinical cases also offer convincing support for the functional specialization of the cortex. Reports exist of injuries causing loss of color vision [Zeki, 1993], while the perceptions of depth and motion remain. There is also one striking case of an injury resulting in motion blindness [Zihl et al., 1983]. The patient had difficulty perceiving motion. She had trouble crossing roads because a car that appeared far away would suddenly be nearby. The patient's perception of color, texture, depth, and her ability to recognize shapes were unimpaired. Within this framework, it becomes nontrivial to understand how these features, which are "processed" by different neurons in different cortical areas, are linked to form the perception of a coherent whole; this is the feature binding problem. Given that many objects exist in a typical natural scene, and that each object can have different features, such as texture, brightness, orientation, color, motion, depth, etc., it becomes obvious that there are a combinatorial number of linkages between features and objects.

There are two primary and competing approaches to bind the features of an object together. The first approach assumes that it is the average firing rate of a neuron which conveys information, which is the traditional assumption [Zeki, 1993, Singer, 1993]. Since the individual firing times of neurons do not encode any information, then the neurons in these separate regions of the visual cortex must be grouped and segregated by some other mechanism. It has been proposed that neurons become more selective as one proceeds from lower to higher visual brain areas [Barlow, 1972]. Instead of responding preferentially to a direction and velocity, a neuron may respond best to a single object. This is called the grandmother-cell representation. One would expect a large region of the brain to be devoted to this task given the large number of objects humans can recognize, but currently, evidence does not strongly favor this proposal. There are cortical areas that appear to be selective to classes of patterns [Fujita et al., 1992] (faces and hands for exam-

ple [Desimone et al., 1984, Perrett et al., 1987]), but these are not object-specific cells. Another problem with this proposal is that objects must be recognized in order to be deemed as objects, which makes the categorization of an object as “unfamiliar” problematic. Also, the features of an object can be bound together preattentively, suggesting that feature binding and segregation occur without recognition, and therefore without the aid of grandmother-cells. The second approach uses temporal correlation in the firing times of neurons to create linkages. Here the precise firing times of neurons are important. Both grouping and segregation can be encoded by the firing times of groups of neurons and other, more complex mechanisms need not be proposed. Using temporal correlation as a framework in which to represent feature binding appears elegant. Furthermore, it appears to be possible given the known physiology of the brain.

Several experiments have been proposed to test the hypothesis that correlations in the firing times of neural groups may have a role in information processing. Gray et al. [Gray et al., 1989] moved a single long bar of optimal orientation, velocity, and preferred direction through the receptive field of recorded neurons in the primary visual cortex of an anesthetized cat. A typical response to this stimulus is shown in Figure 2, (reproduced from Gray and Singer [Gray and Singer, 1989]). The local field potential is shown in the top graph, and the multiunit activity is shown in the second graph. The third and fourth plots are portions of the first two plots with an expanded time scale. The local field potential indicates the combined activity of many neurons (100's or 1000's). The multiunit activity indicates the firing times of a few neurons in the immediate vicinity of an electrode. The neurons generating these spikes fire at nearly the same time, although their activity is not rigidly periodic. This synchronous activity has an average frequency of approximately 40Hz and ranges from 35-50Hz. Gray and Singer found that the neural response to the stimulus exhibited synchronous neural activity over a distance of 7mm in the cortex and that synchrony occurred 25-50 ms after stimulus onset. When two separate and smaller bars moving in the same direction passed through the cat's visual field, there was a decrease in the synchrony between the neural responses to the separate bars. When two separate bars moved through the cat's visual field in opposite directions, there was no correlation in the neural responses between the two bars. This suggests that an object may be encoded by the synchronous firing times of neurons responsive to that object, and that the firing times of these neurons are uncorrelated with the firing times of the neurons responsive to different objects. The animals in these experiments were unconscious, which suggests that neural synchrony in the visual cortex may be a process that is not governed by recognition or attention. Also, in this area of the cat visual cortex the longest reported connection between neurons is approximately 4mm [Hirsch and Gilbert, 1991], thus implying that synchrony is achieved through local interactions.

Based on the above experiments, one can conjecture that a single object is represented by many neurons, each of which is firing in synchrony. Different objects may be represented by different groups of neurons which have different firing times. A simple implementation of this would be to represent each feature (or pixel) in an image with an oscillator. An example of this is shown Figure 3A. Here, each grid of the array consists of an oscillator coupled with its four nearest neighbors. The oscillators receiving stimulus (black squares) begin to oscillate, while the oscillators that do not receive stimulus (white squares) remain inactive. All active oscillators have the same frequency. The oscillators

Figure 2. The local field potential (first graph) and the multi-unit activity (second graph) in response to a moving bar. The third and fourth graphs are portions of the first two graphs with an expanded time scale.

comprising a single object begin to oscillate in synchrony, while the oscillators representing different objects have distinct phases of oscillation. This form of temporal correlation is called *oscillatory correlation* by Wang and Terman [Wang and Terman, 1995]. These authors have proposed a network that performs oscillatory correlation and the temporal activity of their network is shown in Figure 3B using Figure 3A as the input. The temporal correlation of the oscillators (in this case synchrony) allows the features (pixels) of a single object to be bound together, while allowing different objects to be segregated. The network can be said to represent the three objects simultaneously.

1.3 Relation to Computer Vision

In our work, the primary goal of performing synchronization and desynchronization is for visual processing (although feature binding through oscillatory correlation can be used in other sensory perceptions [Baird, 1986, Wang, 1996b]). The stage of visual information processing that these networks most clearly correspond to is that of segmentation. Segmentation is typically labelled as a “low-level” processing task [Pavlidis, 1973, Braddick and Atkinson, 1982]. In Figure 3 for example, the oscillator network segments the three objects from the background, and also from one another. In segmentation, the pixels of an image are grouped together using some basic visual features such as intensity, color, or

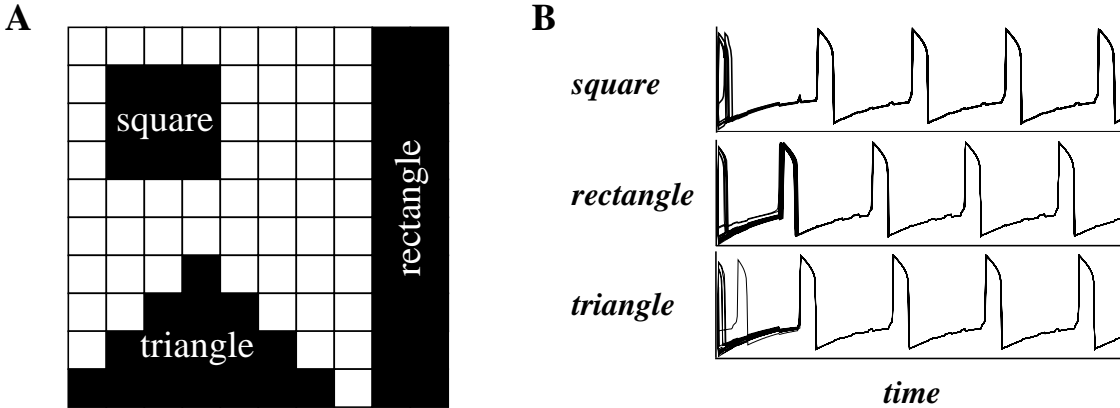


Figure 3. An example of temporal correlation to solve the feature binding problem. (A) each grid of the array is an oscillator. Oscillators receiving stimulus are black. (B) A plot of the temporal activity of the oscillators receiving stimulus. All the oscillators comprising a single oscillator become synchronous, and are desynchronized from the other

texture [Haralick and Shapiro, 1985, Sarkar and Boyer, 1993]. After segmentation, the image is then suitably prepared for “higher-level” tasks, such as pattern recognition. The reason for this is that it should be more efficient to perform recognition after the image has been broken into big chunks, or objects. Performing recognition at the pixel level would require much more computation. Scene segmentation is largely regarded as unsolved [Pal and Pal, 1993]. Despite decades of research (for reviews see [Zucker, 1976, Haralick and Shapiro, 1985, Pal and Pal, 1993]), it remains a computation roadblock for the computer vision community.

We present some real images which have been segmented using an oscillator network in Chapter 2. These results are obtained using procedures nearly identical to those developed in Wang and Terman [Wang and Terman, 1997]. Although our segmentation results are arguably no better or worse than other conventional computer vision approaches, our networks offer some advantages that traditional algorithms lack. The first advantage is that our networks are based on the brain, which offers the only working image processing paradigm. Our networks are related to the neurophysiology of the brain, and they make use of local computations. Many of the techniques from computer vision make use of global computations, such as averages [Pavlidis, 1977] or clustering [Jain and Dubes, 1988]. Experiments from visual psychophysics and neurophysiology indicate that global information is not always necessary to segment images [Nothdurft, 1994]. Our networks perform local computations, which is neurophysiologically supported, and further, more suitable for VLSI implementation.

1.4 Neural Network Approaches

The classification abilities of neural networks have been used to perform segmentation based upon pixel classification [Kohonen, 1995, Koh et al., 1995, Raghu et al., 1995]. However, these classification schemes require that the number of clusters be specified *a priori*. This requirement is a drawback since segmentation should reveal how many objects are in a given image. Also, all of these networks must be trained, but, as noted previously, scene segmentation does not require recognition. Further, these networks do not have the ability to represent multiple objects simultaneously - a serious drawback since most natural scenes contain many objects.

Temporal correlation offers an elegant means of representing multiple objects simultaneously. The discovery of synchronous oscillations in the visual cortex has created a flurry of activity in attempts to apply oscillator networks to the problems of scene segmentation. Oscillator neural network models simulating the experimental results have been proposed [Schuster and Wagner, 1990, König and Schillen, 1991, Grossberg and Somers, 1991, Sompolinsky et al., 1991, Chawanya et al., 1993, Grannan et al., 1993, Wang, 1993a], and have been used to explore the ability of oscillatory correlation to solve the problems of pattern segmentation and feature binding [Sporns et al., 1989, Baldi and Meir, 1990, Horn et al., 1991, Hummel and Biederman, 1992, von der Malsburg and Buhmann, 1992, Murata and Shimizu, 1993, Schillen and König, 1994, Terman and Wang, 1995, Hopfield and Herz, 1995]. To be used as a solution for the feature binding problem an oscillatory network must be able to perform synchronization and desynchronization. Of the models proposed, only Terman and Wang [Terman and Wang, 1995] prove that their network can achieve synchrony. The others demonstrate the behavior of their networks through computer simulations. Also, the authors (excepting [Terman and Wang, 1995]) do not examine how the time needed to achieve synchrony scales with the network size. This is an important issue because in the models proposed, no real images were segmented; only results for toy problems, as in Figure 3, were shown. How their networks can be utilized for image processing remains unknown unless the behavior of their networks at larger scales (say 256×256) are known. To the best of our knowledge, only Wang and Terman [Wang and Terman, 1997] segment real images with an oscillatory neural network. Also, most of these models do not address the task of desynchronization. Several authors consider this issue [von der Malsburg and Buhmann, 1992, König and Schillen, 1991], but they only demonstrate desynchronization between two different groups of oscillators. Terman and Wang [Terman and Wang, 1995] and Campbell and Wang [Campbell and Wang, 1996] provide two networks that are able to segment more than two objects.

1.5 Statement of Main Goal, and Related Issues

One of the main goals of this study is to understand how synchrony arises in finite, locally connected networks of neural oscillators. Because of the complexity of the problem, the networks we study are simplified as much as possible. The connections are only

between nearest neighbors, and all the oscillators are identical. The interactions between oscillators are also identical, although not always symmetrical. Noise is not analyzed, although it is used in simulations to test the robustness of these networks.

Local connections are examined not only because of their basis in neurophysiology [Kandel et al., 1991, Hirsch and Gilbert, 1991], but also because they may be required for scene segmentation [Sporns et al., 1991, Wang, 1993b]. Specifically, in a two dimensional array of oscillators, where each oscillator represents an element in a sensory field, all to all couplings indiscriminately connect multiple objects. All pertinent geometrical information about each object, and about the spatial relationships between objects is lost. Because this spatial information is necessary for object recognition and segmentation, it must be preserved as effectively and as simply as possible. The use of local couplings is one method of maintaining such spatial relationships.

We are not only interested in the property of synchrony, but in how fast it is attained. The neurophysiological data indicate that the neural response to visual stimuli synchronizes quickly, in 25-50 ms, or 1-2 periods of oscillation assuming a 40 Hz frequency [Gray et al., 1989]. Thus, in the brain some form of self-organization occurs within a relatively short time. As previously noted, the features of an object are bound together before recognition occurs. If feature binding is necessary for recognition, then recognition can only occur after feature binding. Psychophysics experiments [Biederman and Ju, 1988] indicate that recognition can take place within 50-100 ms after stimulus presentation. Though this does not support temporal correlation as a means of feature binding, this evidence does not refute it either. If synchrony underlies such basic cognitive processes as scene analysis and object recognition, then intuitively it must occur quickly to be useful in a rapidly changing environment.

Desynchronization is another topic we study. In order for different objects to be represented simultaneously, different groups of oscillators must be desynchronized while synchrony within each group of oscillators is maintained. In the networks studied here, in which there is only a single frequency, this corresponds to separating the phases of different objects. This is an important task. If two or more groups of oscillators happen to have the same phase, then they will be mistakenly grouped together. This probability of accidental synchrony [Hummel and Biederman, 1992] can be substantial when one considers that the effects of noise and the time needed to achieve synchrony require that a tolerance of phase differences be allowed. Also, as the number of groups increases, the possibility of accidental synchrony increases significantly. Some form of desynchronization is necessary to avoid this problem. In two of the networks we examine, mechanisms which perform desynchronization are demonstrated.

Another topic of interest is that of time delays in the interactions between oscillators. The speed of neural signal conduction is relatively slow, around 1 m/s in unmyelinated axons [Kandel et al., 1991, Traub et al., 1996]. For two neurons that are separated by 1 mm, the time delay is approximately 4% of the period of oscillation (assuming a 40 Hz frequency). How neural synchrony is achieved in the presence of significant time delays is an important question. We will examine time delays in a pair of neural oscillators in which an analytic solution can be found. Simulations are then used to help understand how time delays affect synchrony in larger populations of oscillators.

1.6 Background Knowledge

Achieving synchrony in locally coupled populations of oscillators has not been a widely studied topic. There are no general classifications describing the forms of interactions necessary for synchrony using a given type of oscillator. Below, we briefly review a few of the papers relevant to synchrony in networks of oscillators. This is not meant to be a comprehensive review, as more papers will be cited when relevant in later Chapters.

Much of the work done on populations of oscillators deals with the simplest description of limit cycle behavior, the phase oscillator,

$$\dot{\phi}_i = \omega_i + \sum_j J_{ij} G(\phi_i, \phi_j) \quad (1.1)$$

where the rate of change in the phase of oscillator i is given by its intrinsic frequency ω_i , and its interaction with its neighbors. The coupling strength and the extent of the coupling are given by J_{ij} . The interaction between two oscillators is defined by $G(\phi_i, \phi_j)$. Much work with phase oscillators has been done in globally connected oscillator networks in which each oscillator is coupled with every other oscillator (see [Kuramoto, 1984, Daido, 1993a, Strogatz and Mirollo, 1988] for example).

Niebur et al. [Niebur et al., 1991b] have derived results for locally coupled phase oscillators with a diffusive interaction, $G(\phi_i, \phi_j) = (\phi_i - \phi_j)$, in the presence of noise and with a distribution of intrinsic frequencies; both noise and disorder result in the absence of synchrony. If all the phase oscillators have identical frequencies and there is no noise, then diffusive coupling leads to perfect synchronization [Niebur et al., 1991b].

Another type of oscillator that has been frequently examined is the integrate-and-fire oscillator [Peskin, 1975], (see Section 2.2 for a description). The following authors report that synchrony occurs in locally coupled networks of integrate-and-fire oscillators [Mirollo and Strogatz, 1990, Corral et al., 1995b, Hopfield and Herz, 1995]. Hopfield and Herz [Hopfield and Herz, 1995] report that synchrony occurs, but that it occurs on slow time scales (several hundred periods). In Chapter 2 we reproduce their results and indicate how to adjust parameters to achieve fast synchrony. We also provide numerical evidence that synchrony occurs at times proportional to the logarithm of the network size for one- and two-dimensional systems.

Relaxation oscillators are two variable systems with limit cycle oscillations and they are also frequently studied (see Appendix A for an introduction). Relaxation oscillators are directly related to models of neurons [Fitzhugh, 1961, Morris and Lecar, 1981] and because of their biological relevance, they have a long history as models of biological oscillations (starting perhaps with [van der Pol and van der Mark, 1928]). Somers and Kopell [Somers and Kopell, 1993] and independently Wang [Wang, 1993a] indicated that networks of locally connected networks of relaxation oscillators, strongly coupled with a Heaviside type interaction, have fast properties of synchronization relative to sinusoidal type oscillators. Terman and Wang [Terman and Wang, 1995] further showed that a locally connected network of relaxation oscillators, with certain restrictions on the initial conditions, can synchronize at an exponential rate independent of the size of the network.

With more general initial conditions the rate of synchrony alters drastically. In Chapter 3 we provide numerical evidence that the time to synchrony increases as a power law of the size of the network for one-dimensional systems.

Work has also been done on desynchronization in locally coupled oscillator networks [Schillen and König, 1991, von der Malsburg and Buhmann, 1992, Schillen and König, 1994], but none of these authors demonstrate an ability to robustly desynchronize more than two groups of oscillators. Terman and Wang [Terman and Wang, 1995] are the only authors we know of to have performed a rigorous analysis detailing how desynchronization occurs. We present two novel methods for desynchronization of multiple groups of oscillators in Chapter 2 and Chapter 5.

1.7 Outline of Dissertation

Synchrony is investigated in three types of neurally derived oscillators. In Chapter 2 we examine the simplest model of neural behavior, the integrate-and-fire oscillator [Peskin, 1975]. One- and two-dimensional locally coupled networks without noise are examined and we observe that synchrony is always achieved. Furthermore, our data suggest that the average time to synchrony increases logarithmically with the size of the system. We present some heuristic arguments for this behavior. Numerical results are shown that indicate error-correction also occurs on a time scale proportional to the logarithm of the size of the error. Using integrate-and-fire oscillators we are able to desynchronize groups of oscillators without destroying synchrony within each group of oscillators. We then create an oscillator network which is able to segment real images using oscillatory correlation.

In Chapter 3, we perform a similar computational analysis for another, more complex model of neural dynamics, the relaxation oscillator [Fitzhugh, 1961, Morris and Lecar, 1981]. We first examine relaxation oscillators in the singular limit, and our numerical evidence suggest that the time to synchrony scales as a power law of the system size for one-dimensional networks without noise. There is a parameter regime for which the time to synchrony increases as the logarithm of the system size and we characterize the parameters required to observe this scaling relation. Relaxation oscillators are typically believed to share some properties with integrate-and-fire oscillators. The results of this chapter show that there are qualitative difference between these two models and we indicate where these differences originate. We also explore how the type of interaction affects the time to synchrony in networks of oscillators and find that a discontinuous interaction results in better properties of synchrony when compared to a smooth interaction. We believe that this statement might apply to many classes of oscillators and interactions.

In Chapter 4 time delays are introduced into the coupling between relaxation oscillators. In the specific relaxation oscillator we use, we are able to analyze the behavior of a pair of oscillators for a broad range of initial conditions and time delays. The time delays results in a phase relationship we call loose synchrony. We introduce a new measure of synchrony and use simulations to help understand the behavior of networks of relaxation oscillators with time delays in one- and two-dimensions. We also describe a range of initial conditions in which the degree of synchrony does not degrade as the network evolves.

We investigate synchrony and desynchrony in the Wilson-Cowan oscillator in Chapter 5. The Wilson-Cowan oscillator is derived from interacting populations of excitatory and inhibitory neurons [Wilson and Cowan, 1972]. We create a Lyapunov function for the piecewise linear approximation to this oscillator, and it is shown that synchrony is achieved if the coupling strength is large enough. Also, a mechanism is presented for achieving synchrony within one cycle for a finite number of oscillators. This mechanism uses strong coupling and is based on an interaction that dynamically creates and destroys fixed points in a fashion that assists synchronization. Also shown for this specific model is a means of achieving desynchronization for a small number of objects. Chapter 6 concludes the dissertation.

