

**A Stigmergic Model for Oscillator Synchronisation and its
Application in Music Systems**

5th September 2011

MSc Creative Systems

Candidate No: 59880

Registration Number: 20914003

Abstract

Non-linear and chaotic dynamics, predominantly used in engineering, have become a pervasive influence in contemporary culture. Artists, philosophers and commentators are increasingly drawing upon the richness of these systems in their work. This thesis explores one area of this territory: the synchronisation of a population of oscillators used for the generation of rhythm as applied in musical systems.

Synchronisation is taken as a basis for complex rhythmic dynamics. Through the self-organisation notion of stigmergy, where entities are indirectly influenced by each other, the Kuramoto model, a mathematical model of coupled oscillating processes, is discussed and modified to increase the model's stigmergic saliency. The notion of local field coupling is introduced as a qualitatively stigmergic alternative to Kuramoto and noise, distance, delay and influence are incorporated.

Stigmergy concerns the mechanics of environmentally mediated interaction and as such the importance of interactivity within the system is discussed. An interactive system of stigmergic synchronised oscillators is developed, which is open to be used across many fields. The user is allowed to become part of the stigmergy through influencing the environment. This system is applied to the field of music, generating rhythms and sounds by mapping the state of the system.

Introduction

This thesis explores a potential biological root of musical creativity: oscillator synchronisation. The thesis argues that through oscillator, interesting musical behaviour can be achieved.

In Chapter 1, the phenomena of oscillator synchronisation in the natural world is explored. The Kuramoto model is introduced as a powerful and elegant mathematical formula describing the phenomena. However, since synchronisation has its roots in self-organisation, the Kuramoto model encounters a problem and falls short of complete plausibility. Stigmergy, a notion where entities are indirectly influenced by each other, is used as the mode of exploration into self-organisation. An alternative model, *local field coupling*, derived from Kuramoto and other methods of oscillator synchronisation taken from biology and neuroscience, is described to solve this problem.

Chapter 2 concerns theories from the fields of chronobiology and biomusicology, which use oscillator synchronisation phenomena to explain many forms of behaviours in living systems. A clearly rhythmic, but not necessarily creatively musical behaviour is achievable through stigmergic synchronisation, termed *protomusical* behaviour. From here, this chapter explores the line where the protomusical may end and the creatively musical may begin. SymbioticA's *Silent Barrage (SB)* (SymbioticA, 2010) is used as an example of a stigmergic system which can also be considered as a creative system.

An interactive system developed for this thesis, *Crickets*, is detailed in Chapter 3. *Crickets* is an environment in which low-level creativity is achievable through biologically inspired protomusical behaviour. The protomusical behaviour generated by the system is able to be reused in many applications across disciplines.

Crickets sits in the same vein as systems such as *IanniX* (Coduys and Ferry, 2004), *StarLogo* (Resnick, 1997) and *SB*. *SB* is detailed in Chapter 2 and is the closest of the three to *Crickets*, and the theory behind *Crickets*.

IanniX is a visual sequencer comprised of simple objects such as time lines, play heads and triggers, the use of which can create complex non-linear and generative scores. *IanniX's* strength is its ability to be scripted and control itself through other time lines. *Crickets* takes a similar approach to *IanniX's* detachment of interface and sound engine. *IanniX* itself does not generate any sound but rather broadcasts its events to external programs. However, any behaviour added to *IanniX* that is beyond simple commands must be programmed into the system. As such it is a far more general tool in comparison to *Crickets*.

StarLogo is a type of experimentation environment for exploring self-organised phenomena. It allows agents to be scripted in such a way that naturally occurring phenomena such as a honeycomb structure or ant foraging can be allowed to emerge from simple rules. The environment is purely visual and internal, as it was originally intended for educational purposes. No information is sent to external programs and as such a self-organised system like *Crickets* could not be made in *StarLogo*.

The *Crickets* system is a hybrid of the above three systems in that it is an environment for exploring a specific self-organisation phenomenon whilst also acting as an interface for control of an external system, thus becoming a component in a larger system such as an art installation.

Chapter 1 - Oscillation and Self-organisation

Synchronisation

Oscillations can be observed practically everywhere. In nature, behaviour such as a honey bee's activity cycle, a fiddler crab's claw waving and an ant's alarm drumming are examples of living oscillations. On the microscopic level, every living system is controlled by internal biological oscillators, which affect the organism's physical, emotional and cognitive behaviour (Foster and Keirtzman, 2005), and even those systems are made up of matter, containing electrons oscillating about the nucleus of an atom. On a much grander scale, an organism's entire being is influenced by celestial oscillations. The Earth's orbit around the sun forms the day and night cycle and creates a circadian rhythm, defining the rhythm of life (Foster and Keirtzman, 2005). Furthermore, some systems that appear to be one oscillator, can in fact transpire to be constructed of a whole host of oscillators, collectively exhibiting one large synchronised oscillation (Knoester and McKinley, 2009).

The question which then arises in this case is; how do these systems perform their synchronisation? Are the oscillators connected in some manner and if so, how? The idea of oscillations in nature possessing identical natural frequencies is unlikely, especially when these oscillations are subject to the random forces of environmental noise. This lead Kuramoto (1984) to state that mutual synchronisation could be the only way the oscillation is both produced and maintained.

Kuramoto defines the term *synchronisation* as, “multiple periodic processes with different natural frequencies [that] come to acquire a common frequency as a result of their mutual or one-sided influence” (1984: 2). However, throughout this thesis, the term synchronisation is used in conjunction with the term *entrainment*. There is a key difference between the two terms regarding the way that the oscillators interact. Entrainment refers to a one-sided interaction: an owl's daily

activity cycle is *entrained* to the Earth's orbit around the sun; the owl has no influence on this orbit and thus entrains itself to it. In a synchronised oscillation however, there is a degree of mutual feedback where the two oscillators find a common frequency and phase.

For Knoester and McKinley (2009), learning more about synchronisation was an aid to gaining insight into making signals resilient to environmental noise. To study this they conducted an artificial evolution simulation, modelled on the firefly and its behaviour in a swarm. The firefly was chosen for its ability to synchronise the flashes of its abdomen to the swarm's common flash frequency and phase, even if the swarm spans a distance of many miles. The evolved fireflies were then subject to high levels of data loss when attempting to synchronise, thus simulating the environmental noise. A successful synchronisation was recorded when all flashes occur within a short space of time, rather than simultaneously.

The results of this experimentation showed that the evolved flashes were remarkably similar to a model developed by Ermentrout (see Knoester and McKinley, 2009) in which control signal flow was modified in the agent. This study provides evidence to suggest that a synchronisation is achieved via a similar method in reality. In the evolved artificial agents, control flow was adjusted via a recursive task taking a longer or shorter amount of time to complete. The length of this task corresponds directly to phase shifts in the agent which over time leads to a perceived change in frequency. Despite the significant loss rates of data in the simulation, a high proportion of evolved agents were able to synchronise successfully with this method, providing a noise robust solution.

In nature, it is unlikely that the exact phase shifting method evolved by Knoester and McKinley's fireflies is occurring. However, synchronisation through shifting phases is not a new discovery and has been explored in great detail in the Kuramoto model.

The Kuramoto Model

Kuramoto stresses the importance of synchronisation and entrainment, stating that they “cannot be

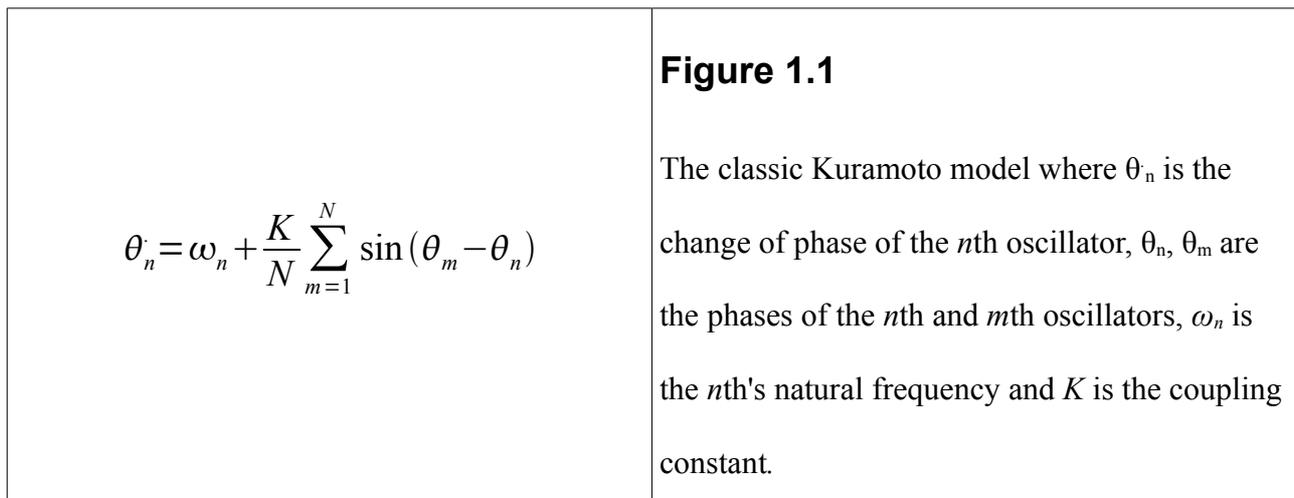
emphasized enough.” (1984: 2) This is due to the fact that the ubiquity of this phenomena even extends to non-living systems such as chemical reactions. Furthermore, “the collective rhythms in living systems that arise from the cooperation of the individual cellular oscillators seem to play important roles in the coordination of life processes.” (Kuramoto 1984: 60) This leads to the idea that the aforementioned examples of synchronised oscillation from nature may in fact be a collective oscillation on the microscopic, chemical level as well as the macroscopic, observable behaviour level.

Kuramoto formulated a mathematical model centred around oscillators that conserve energy and have a stable amplitude. This type of oscillation is commonly referred to as a limit cycle oscillation; a sine oscillator obeys a limit cycle of 1, for example. These oscillators were generalised into a basic functional unit, which when combined in some way can output various dynamical behaviours, forming what he termed a *many-body* theory of limit cycle oscillators, though a more common term is an oscillator *population*. In this way, the exact behaviour of each oscillator is abstracted away, allowing the same model to be used on several different oscillators, even allowing a population of many different oscillators to be synchronised (Kuramoto, 1984).

The oscillators are seen as being coupled via a diffusive field, leading Kuramoto to posit that reaction-diffusion equations, though literally meant to model chemical systems, can actually have a much wider scope, including the self-organisation phenomena of synchronisation. Figure 1.1 shows the resultant formula that has come to be known as the Kuramoto model for synchronised oscillators.

We can see from Figure 1.1 that the model specifies a global coupling where each oscillator is interconnected. They are connected via a sine interaction function, the output of which is reduced to zero where the phases are identical, or differ by π . The interaction is strongest where the coupled oscillators are in an anti-phase relationship, or differ by $\pi / 2$, as the output of the function is either at one or negative one at this point. This means that the oscillators coupled via this model are

attracted to a phase-locked synchronised state (Breakspear et al., 2010).



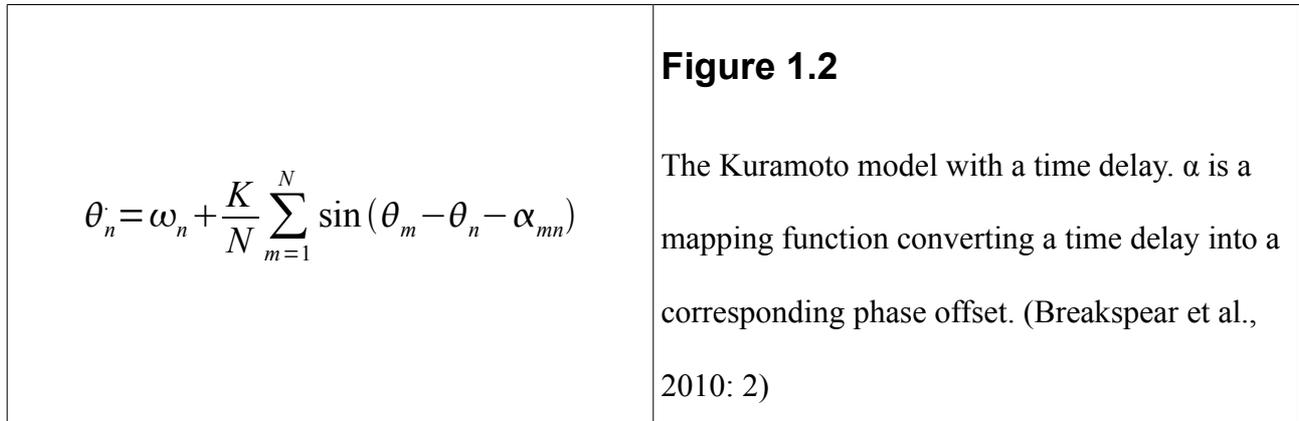
These three factors make the model a rather elegant solution to the synchronisation problem. The model only consists of a changing phase value, meaning that natural frequency is left untouched. However, since the phase is in flux, the oscillator gives the appearance of adjusting its frequency to the global average and thus the population becomes synchronised.

Similarities can be drawn between this model and the method evolved by Knoester and McKinley's fireflies. What Knoester and McKinley described as a change in control flow can be seen as a rediscovery of Kuramoto, whose model more explicitly describes the phase shifting.

Breakspear et al. (2010), working in the field of neuroscience, have criticised the Kuramoto model. According to them, this model has massive implications in their field, but lacks neurological plausibility. They identify the fact that neurons and their cortical oscillations are spatially embedded and therefore their coupling should be as well. Thus, Breakspear et al. added a time delay parameter to the model (Figure 1.2) so that oscillators situated further apart receive the information according to this delay.

Breakspear et al., observe that by adding the delay parameter into the oscillator population, complex synchronisation dynamics can be observed. The sinusoidal interaction would require a phase offset in order to synchronise, so even oscillators with equal frequencies would be in flux. In this circumstance, a stable synchrony becomes difficult, if not impossible, to achieve (Breakspear et

al., 2010). The end result of this delay parameter, may lead to the sort of complex rhythmic behaviour commonly observed in nature.



Emergence and Stigmergy

A self-organising system is a system that forms a pattern or order without a central control mechanism or external influence. The pattern is formed instead via interactions on a local scale, with each part of the system knowing nothing of the global effect of these interactions. Therefore, oscillator synchronisation is a self-organising phenomenon. Entrainment, on the other hand, cannot strictly be said to be self-organising, as there exists an external master oscillator and slave oscillators in the system. Self-organisation is interlinked with two other related terms, *emergence* and *stigmergy*, which seek to encapsulate self-organisation from differing viewpoints.

According to Beckermann (1992) an initial attempt to define emergent behaviour was made by Broad. Broad's notion formed part of a dual theory explaining how the phenomenon of life, or the characteristics a living entity exhibits in its behaviour, can be broken down into components that make up its physical presence. Broad named these two theories *mechanistic* and *emergent*. In the mechanistic view, an entity's behaviour can be reliably deduced through a close examination of these components. Emergence on the other hand occurs when the whole could not be predicted even with the most complete knowledge of the components (see Beckermann, 1992:16). Broad's terms were intended to refer specifically to living entities. However, examples of both mechanistic and

emergent behaviour can be seen in non-living systems. Newtonian physics, for example, can be seen to be a mechanistic model of the physical world, where an object's behaviour can be predicted using knowledge of all the forces acting upon it. Whereas a tornado can be seen as a chaotic emergent system.

The notion of *emergence* does not presuppose chaos. Whilst the resulting behaviour cannot be predicted upon inspection of the components, the behaviour itself that emerges could be predictable, at least in part, and once patterns are found can be reliably reproduced. In Conway's seminal cellular automata model, *The Game of Life* (see Gardner, 1970), many initial configurations of cells have been found that produce common behaviours.

Process is the key notion relating to a contemporary view of emergence.

The common non-specialist interpretation of the term *emergence* refers to revealing, appearing, or 'making visible' an event, object or outcome of a process. (McCormack and Dorin, 2001: 4)

The outcome or behaviour is said to be *emergent*, whereas the underlying rules produce *emergence* through the process. When considering emergence, Broad breaks down the behaviour into its constituent components for scrutiny. By doing so, a set of properties or rules is defined through which the outcome cannot be predicted. This is also a key point in understanding emergent behaviour: simple definable rules that result in a sophistication not present in the original design and is of an apparent higher level.

A famous computational example of exploring emergent behaviour is that of Reynolds' (2001) program, *Boids*. Boids is a model of bird flocking and fish schooling that is achieved by the agents obeying three simple steering rules: separation, steering to avoid overcrowding neighbouring agents; alignment, steering towards the average direction or the entire group; and cohesion, steering towards the average position of the neighbouring agents. Some implementations of Boids provide a

remarkably realistic and low-level simulation of the complex behaviour that is flocking.

Another property of emergence, and indeed its main criticism, is also present in Boids: that it is in the eye of the observer. It is only via external observation that emergent behaviour is defined, agents within the system, by their very nature, cannot *intend* to produce emergence as that will defeat the point. Furthermore, it is the observer that labels that outcome of the process a 'pattern' in the first place, let alone *emergent*. This leads to the area being difficult to study with great accuracy and therefore its importance is brought into question.

Stigmergy on the other hand circumvents this problem through its own definition. It is another term that has its roots in the natural sciences, being devised to explain the control of collective behaviour of social insect such as ants and bees (Theraulaz and Bonabeau, 1999). It is a notion common today in many agent based simulations, in that the agents remain independent entities. Their interactions with the environment affect the behaviour of the other agents, which in turn affects them. Stigmergy is therefore defined as pattern formation in a collective via an interaction with an environmental mediator. A common example is an ant following a pheromone trail to a food source. The ant is merely following a trail and has no knowledge of what lies at the end of it. The ant in turn leaves behind a pheromone trail of its own, thus strengthening the attraction to go down that path to other ants. The resultant collective behaviour is that of many ants taking the shortest possible route between the food source and the nest.

Stigmergy is an all encompassing term for mechanisms such as the ant pheromone example above, and so as Theraulaz and Bonabeau state,

[Stigmergy] has to be supplemented with an additional mechanism that makes use of these interactions to coordinate and regulate collective building in a particular way. At least two such mechanisms have been identified: quantitative stigmergy and qualitative stigmergy. With quantitative stigmergy, the stimulus-response sequence comprises stimuli that do not differ qualitatively [...] and only modify the probability of response of the individuals to

these stimuli. Qualitative stigmergy differs from quantitative stigmergy in that individuals interact through, and respond to, qualitative stimuli. (1999: 105)

An ant pheromone trail is an example of a quantitative stigmergic mechanism, whereas in qualitative stigmergy a specific environmental stimulus produces a specific response. An example of this is seen in the way wasps produce the honeycombed structure of their nests (see Theraulaz and Bonabeau, 1999: 107). Again, both these examples are from the animal kingdom, but stigmergy can exist in non-living systems, such as the path water takes down a surface. Oscillator synchronisation is quantitatively stigmergic.

The concepts of stigmergy and emergence are so closely related it may seem to be a matter of semantics. However, by focusing on environmentally facilitated interaction, stigmergy offers a more stable ground for the experimentation and exploration of self-organising phenomena compared to the notion of emergence, which is largely based on observation and trial and error in the formulation of behavioural rules. For this reason, this thesis focuses on self-organisation through a stigmergic perspective.

A Stigmergic Kuramoto

The notion of stigmergy is not present in the Kuramoto model. A single oscillator in the model has a map of the entire collective state, whereas stigmergy relies on indirect communication, facilitated by an interaction with the environment. Breakspear et al.'s addition of a delay parameter (Figure 1.2) goes some way to bringing in a concept of the environment. By spatially embedding the oscillators, the model now acknowledges the environment.

Further adjustments can be made to the model, to increase its stigmergic saliency. Firstly, the coupling constant K is made into a local value much like a connection weight in a neural network, this is a common adjustment to Kuramoto, although Breakspear et al. did not take this step. This allows the interaction between oscillator to be affected by many factors. For instance, K can become

linked to distance, allowing for the oscillators that are situated closer to have a greater influence on the synchronisation. Noise can also be added to the system to simulate both the signal degradation and random fluctuations present in the environment (Figure 1.3).

| | |
|---|---|
| $\dot{\theta}_n = \omega_n + p + \frac{1}{N} \sum_{m=1}^N K_{mn} \sin(\theta_m - \theta_n - \alpha_{mn})$ | <p>Figure 1.3</p> <p>An alternative Kuramoto model, K_{nm} is the coupling constant between oscillators n and m, p is a random perturbation.</p> |
|---|---|

However, even taking these steps we do not alter the model enough to consider it stigmergic. The interaction is still direct and in this sense it is more of an emergent formula than a stigmergic one. The inherent problem in the Kuramoto model in a self-organisational context is that stigmergy was not present in the initial design. To incorporate this, a radical rethink is required.

An Alternative Model

Further issues arise in the Kuramoto model when dealing with an oscillator that does not have a direct means of control for its phase value, or an entirely accurate frequency value. An example of one such oscillator is the Van der Pol oscillator (VDPO) (Figure 1.4). VDPOs were also inspired by biological systems in that it was used to model an extremely common biologically synchronised rhythm. Van der Pol used three oscillations of this form to model the human heartbeat. More recently, Camacho et al. (2004) have used VDPOs to study the circadian rhythm of melatonin in the human eye.

| | |
|---|--|
| $\ddot{x} - \epsilon(1 - x^2)\dot{x} + x = 0$ | <p>Figure 1.4</p> <p>The unforced VDPO equation, where ϵ represents a non-linearity co-efficient.</p> |
|---|--|

The VDPO obeys a limit cycle of 2 and is a relaxed oscillator, meaning that voltage or activation accrues over time and releases sharply. The non-linearity coefficient ϵ , controls the rate of this action and thus the frequency. Since the dynamics of the system obey a limit cycle, natural frequency can be measured by first measuring the period of the cycle.

These systems can be coupled in a multitude of ways (see Camacho et al., 2004), though the one believed to be the most useful here is what is referred to as the *bath* method. Camacho et al.'s coupling of two VDPOs via a bath is illustrated in Figure 1.5. The reason for the interest in this method is that it can be easily altered so that each oscillator has a local bath to which other oscillators contribute to a greater or lesser extent, thus making it a stigmergic coupling, since the local baths can serve as a perception of the environment.

| | |
|--|--|
| <p>a) $\ddot{x} - \epsilon(1 - x^2)\dot{x} + x = K(z - x)$</p> <p>b) $\ddot{y} - \epsilon(1 - y^2)\dot{y} + y = K(z - y)$</p> <p>c) $\dot{z} = K(x - z) + K(y - z)$</p> | <p>Figure 1.5</p> <p>Two VDPOs, x and y coupled via a bath, z. The parameter K represents a coupling co-efficient, much the same as the Kuramoto model.</p> <p>(Camacho et al., 2004: 2134)</p> |
|--|--|

What follows is an adaptation of Figure 1.4 into a more stigmergic form. Camacho et al.'s bath method is fixed in two dimensions, but since we require a population of VDPOs, this is firstly extended to become n -dimensional. In this thesis, an n -dimensional bath is termed a *field*.

$$\dot{z} = \sum_{n=1}^N K(\theta_n - z) \quad (1)$$

$$\ddot{\theta}_n - \epsilon_n(1 - \theta_n^2)\dot{\theta}_n + \theta_n = K(z - \theta_n) \quad (2)$$

Here the x variable is reassigned to the more common term for phase, θ .

Each field is now made local to that oscillator. In this thesis, this is termed *local field coupling*

(LFC)

$$\dot{z}_n = \sum_{m=1}^N K (\theta_m - z_n) \quad (3)$$

$$\ddot{\theta}_n - \epsilon_n (1 - \theta_n^2) \dot{\theta}_n + \theta_n = K (z_n - \theta_n) \quad (4)$$

Thus far, the coupling is still all-to-all as we have not yet localised the coupling co-efficient K , this is done next and parameter I is added as a scale or influence factor.

$$\dot{z}_n = \sum_{m=1}^N K_n I_{mn} (\theta_m - z_n) \quad (5)$$

where

$$\begin{aligned} 1 &\geq I \geq 0, \\ I &= f(D_{mn}) \end{aligned} \quad (6)$$

$$\ddot{\theta}_n - \epsilon_n (1 - \theta_n^2) \dot{\theta}_n + \theta_n = K_n (z_n - \theta_n) \quad (7)$$

and D_{mn} is the distance between the n th and the m th oscillators when placed in the field. This influence factor is a further step to that proposed by Breakspear et al. It arises from the notion in Reynold's Boids that a close neighbour will have greater influence on the agent than neighbours further away. It is a form of signal loss linked to distance.

To make this a viable real world model, the coupling must be resilient to noise. A noise parameter is added into (7), to account for this:

$$\ddot{\theta}_n - \epsilon_n (1 - \theta_n^2) \dot{\theta}_n + \theta_n = K_n (z_n - \theta_n) + p \quad (8)$$

where p is a random perturbation.

Furthermore, to satisfy Breakspear et al.'s neurological requirements, a time delay parameter is also added to (5):

$$\dot{z}_n = \sum_{m=1}^N K_n I_{mn} (\theta_m - \alpha_{mn} - z_n) \quad (9)$$

where

$$\alpha_{mn} = f(\tau_{mn}) \quad (10)$$

converting a time delay, τ_{mn} , into a corresponding phase offset.

Using (8) and (9) we now have a quantitatively stigmergic coupling model where the VDPOs are situated in a spatial field, feeding back into the local fields of the other VDPOs.

Testing Local Field Coupling

Experimentation with LFC has shown to be successful in synchronisation under certain conditions. To gather the data, a simulation was run where a population of VDPOs were attempted to be synchronised. The main purpose of the simulation was to ascertain the basic degree of synchronisation, for this reason, delay factoring was ignored and K was kept constant for all n . The simulation consisted of two rounds. In the first round (Figure 1.7a) the scaling factor I was ignored, remaining constant at 1, allowing for global coupling to take effect and making spatial positioning irrelevant. Therefore, this mode was closer to the classic Kuramoto model. In the second round (Figure 1.7b), I remained at 1 until a certain distance was reached, then linearly fell to 0 as distance increased. In this round, the oscillators were placed in a line so that each oscillator was able to influence two other oscillators.

Each round was allowed to run for 60 seconds with K set at 0, after which K was stepped up in integer increments every 60 seconds until it reached a value of 10. This was repeated for three levels of noise; 0, 1.5 and 3, which refer to the scaling factor for a random Gaussian distribution.

A measure of the standard deviation of the frequencies was taken, following the formula shown in Figure 1.6. Lower values denote higher levels of synchrony. The graphs in Figure 1.7 display this

value over time.

$$s_N = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2}$$

Figure 1.6

The standard deviation of the sample formula.

The graphs show that there are moments where the deviation spikes to a relatively high value. This is due to difficulty in obtaining a value for the frequency of each oscillator. The experiment relied on the fact that the VDPO has a limit cycle and therefore has a closed phase plane trajectory. Thus, its frequency f , is determined by its period, T :

$$f = 1/T$$

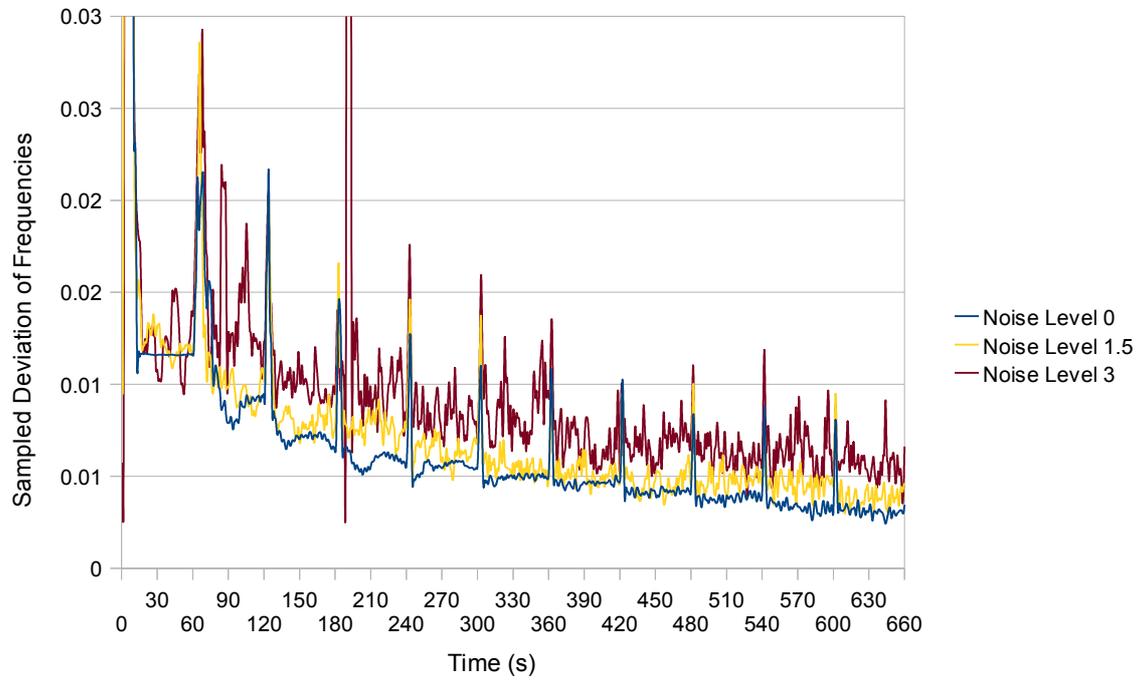
However, the VDPO also tends to have a low frequency, meaning that in the period of time where a VDPO's frequency is shifting, until the field finds a new synchronisation, we have fluctuating readings for T . In the worst scenario, this means that reliable values for the frequency deviation cannot be obtained for a number of seconds, while the mean frequency value adjusts. For a further discussion, and potential solution, see (Schuss et al., 1985); though this simulation attempted to overcome the discrepancy by allowing each K value to 'settle in' over a relatively long period of time.

The results show that the most stable results were obtained when global coupling was still in effect. However, with the correct settings for coupling level, LFC was shown to perform as well as global coupling and is a more viable solution in terms of stigmergy.

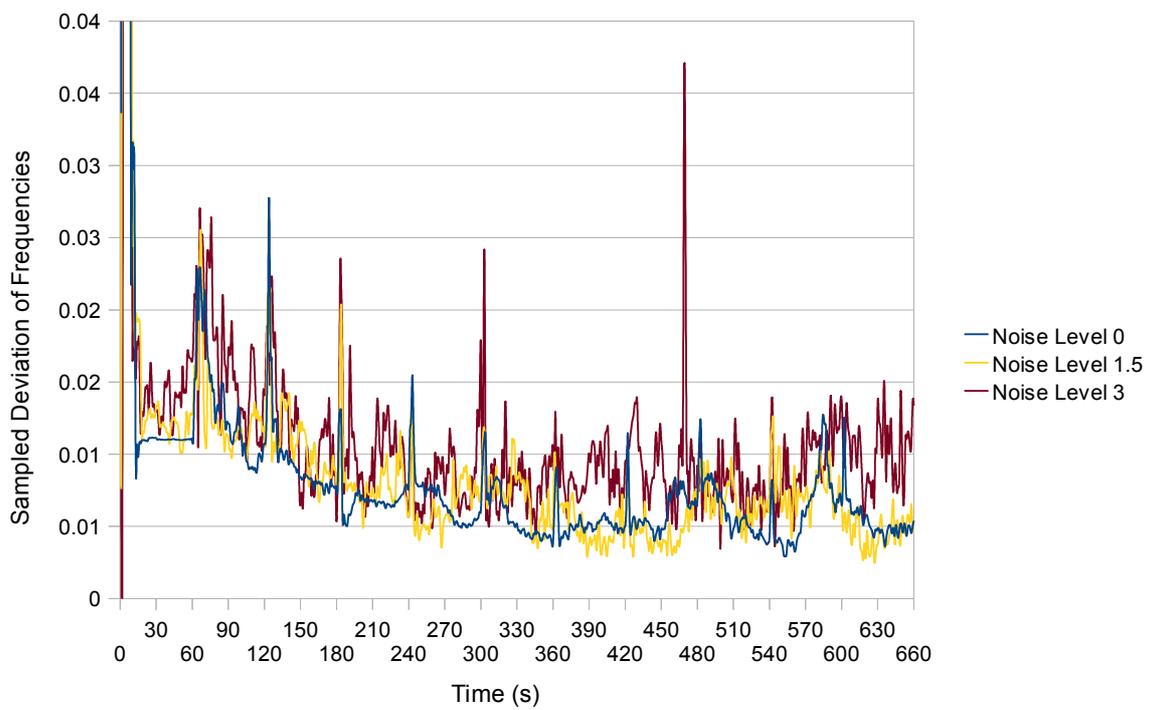
Figure 1.7

The results of tests using LFC coupling: a) shows global coupling and b) shows local coupling.

a)



b)



Discussion

There are many methods available to synchronise oscillators with differing natural frequencies. The Kuramoto model, shown to have strength in its simplicity, may not be viable, or indeed plausible, in all cases. In terms of stigmergy, it is limited by design in that its sinusoidal interaction function is both global and direct. In the case of the VDPO, coupling can be achieved in a self-organising, stigmergic manner, through the LFC method outlined above.

VDPOs coupled via LFC tend not to synchronise their phases, either remaining in their same phase relationship with the rest of the group, or often synchronising to the anti-phase. This is what Kuramoto refers to as repulsive coupling, its counter, attractive coupling, denotes the behaviour where phases lock. Kuramoto states that, “the population composed of identical oscillators coupled through an attractive interaction has been seen to behave like a perfectly self-entrained giant oscillator.” (1984: 78) With the interaction of LFC, it would be unlikely to see this macroscopic behaviour. If this is required, one might be able to turn to Camacho et al.'s (2004) work on the 'in-phase mode' of VDPOs synchronised via a bath and undertake the same localisation steps set out above. However, the aim of this thesis is to explore complex rhythmic behaviour, so with that aim in mind phase differences in oscillators with synchronised frequencies are welcomed.

LFC can be seen to have a 'memory' of past interactions. The local field variable z is in constant flux as its value changes over time. However, one of the strongest connections in each oscillator is its self-connection, which can lead to self-entraining taking place. When the local field oscillation and the self connection is combined, it is possible that the local field affects the behaviour of the oscillator long after the external influence has stopped.

LFC is a model in its infancy and remains to be fully realised, for this reason it still has many limitations and assumptions. By reducing a limit-cycle oscillator down to a basic functional unit, Kuramoto was able to generalise about a whole family of oscillators with one model. In contrast,

LFC was inspired, tested and explained through the use of VDPOs. There is potential scope to reduce the dependency on the VDPO: the field interaction in LFC is essentially an external input to the oscillator. However, the exploration and testing of these options is outside the remit of this thesis. This thesis uses the VDPO exclusively due to the fact that it was biologically inspired and has been used to model real biological oscillations. LFC has been shown to be an effective stigmergic coupling mechanism with VDPOs, therefore there is no concern for its use in this thesis.

It is important to note that the decision to exclude the time delay parameter was only taken for this preliminary experiment, to simply ascertain if LFC works. All further work will incorporate the Breakspear et al. inspired parameter in LFC.

Chapter 2 - Stigmergic creativity

Chronobiology and Biomusicology

Two scientific disciplines dominate the study of rhythms in biological entities: chronobiology and biomusicology.

Chronobiology, a sub-field of biology, concerns the study of periodic phenomena in biological systems including development cycles, behavioural properties and molecular activity etcetera.

Chronobiological studies often centre around the particular types of periodic cycles: circadian, infradian, and ultradian rhythms, which are terms for biological oscillations with periods of around one day, longer than one day and shorter than one day respectively.

A common feature of chronobiology is the notion of the essential nature of biological rhythms, and often the common acceptance that these rhythms are self-organised. For instance, Lloyd (2007) states that the species of yeast *Saccharomyces Cerevisiae* contains an ultradian rhythm that has a period of forty minutes. Lloyd states that the rhythm of this particular yeast species has the most detailed examination of an ultradian rhythm and that this rhythm has been found to be a ubiquitous phenomenon.

This phenomenon is not a mere curiosity, but an archetype. The strict rhythmic self-organisation evident here is essential in every living organism; its pervasive ubiquity underlines its importance. The time-base it provides is a necessity for all life processes. (Lloyd, 2007: 286)

The implications of this statement are substantial. Life and rhythmicity it seems are coupled together in that all living entities require biological oscillators. This makes sense in that many biological sub-systems are linked to the notion of time, such as respiration, reproduction, and growth. Furthermore, life is a temporal phenomenon as all living systems come to an end in death.

In this statement Lloyd also, quite casually it seems, discusses the ultradian rhythms of yeast in terms of self-organisation, it is this part that is stated to be essential. Attia (2004) goes even further and discusses stigmergic aspects of the behaviour of land snails. *Helix Pomatia* are common land snails which display a circannual (around a year) rhythm in their hibernation cycles and circadian rhythms in their activity cycles. Whilst prior observations have raised doubts in the importance of the environmental factors on biological rhythms beyond entrainment, Attia argues that this is due to a lack of knowledge on the interactions between these rhythms and the environment. Attia notes a definite endogenous component of biological rhythmicity, and observes that “in groups of animals [...] frequent contacts increase the individuals' activity levels in comparison with isolated animals” (Attia, 2004: 39). This means that the social interactions of each snail are having an affect on their internal biological clock. In addition Attia states that in the case of land snails, there is evidence of complex interactions existing with the environment. Although the term itself is not used in the paper, this is evidence of stigmergy.

Closely linked to chronobiology is the study of music and rhythmicity from a biological perspective, Biomusicology. Typical areas of biomusicological concern include the common musical features in biology, the origins of music, the biological processes behind music and their function in biological systems, how music evolved in animals and indeed the question over whether or not non-anthropomorphic music exists.

According to Marler (1999), animal calls may be a source of insight into the animal origins of music. There are many examples in nature where animals vocalise a clear and distinct intention. Some species of monkey, for instance, have distinct calls for raising an alarm, or communicating the discovery of some food. Several gibbons perform male-female duets. Many songbirds have a vast repertoire of songs that denote their social status and increase their reproductive prospects (Marler, 1999).

Chimpanzees engage in *pant-hooting*, which is a loud, structured, rhythmical hooting used often in

chorus to keep in touch with the rest of the troupe in the forest. Merker (1999) suggests that early synchronised choral behaviour such as pant-hooting may have arisen from a need to attract mobile females. A synchronised call produces a summed amplitude, and therefore a louder call, meaning passing females are more likely to be attracted to that group of males. This synchronised behaviour also occurred around the same time as bipedalism (walking on two feet), an act that requires complex rhythmic coordination in the brain (Merker, 1999).

The songs of certain songbirds and cetaceans are often passed from generation to generation through imitation, repetition and recombination (Marler, 1999). Marler states that “the ability to learn new vocalisations [...] greatly facilitated the emergence and exploration of phono coding” (1999: 39). The term *phono coding*, or *phocoding*, refers to the encoding or mapping of an abstract notion into a sound and according to Marler is a fundamental requirement for the evolution of speech, language and music.

Marler argues that the communicative behaviour of animals has more in common with music than with language. The type of notions that are communicated by animals in this fashion are abstract and emotional. They serve to fit an emotional purpose with an abstract meaning, whether that meaning is a warning or an attraction, rather than having a common phocoded vocabulary and syntax. However, Marler also notes that creativity is a fundamental requirement for the origins of music. As such the indications of creativity in certain musical behaviour in the animal kingdom are minimal. In the case of gibbons' duets, the fact that the behaviour appears to be innate and that there is usually one signature song per animal does not seem particularly creative. The winter wren, on the other hand is arguably a very creative animal due to their large repertoires of songs that develop by rearranging learned phrases (Marler, 1999). This clearly rhythmic, but not necessarily creatively musical behaviour in animals is termed *protomusical*, meaning that the behaviour exhibits common musical features without being defined as music.

Rhythmic Behaviour as Self-organised Synchronisation

Not only is a clear operational definition of 'good' music hard to come by, a definition of 'music' is often arbitrary at best. [...] Rhythm essentially refers to timing, both how long events last, and when they are scheduled to occur. [...] Music [...] may be beat-oriented or pulsed, but it need not be. [...] Even arrhythmic music is rhythmic. (Biles, 2007: 29-30)

Throughout this thesis, I have been referring to the notion of *rhythmic behaviour*, rather than discussing the notion of *music* in general. This is partly due to the difficulty of defining music and the relative ease of defining rhythm that Biles elucidates above. In the previous section, music as a term was discussed through a biomusicological perspective. Whilst this did not attempt to define *music*, the term *protomusical* was defined as a rhythmic behaviour, which is the basis for musical behaviour. Biles' second consideration above notes that the concept of rhythm is inextricable from merely events in time, therefore even random temporal events can be considered to be rhythmic. A further step is then needed to make rhythmic behaviour more than random noise: organisation.

Varèse coined a commonly received base definition of music: that it is 'organised sound' (see Goldman, 1961: 133). Merker (1999) notes that the pulse in a piece of human music is most often constant throughout the piece.

We hardly ever encounter music employing discrete, that is, stepwise (from one beat to the next) and frequent tempo changes as a structural device for generating variety. (Merker, 1999: 315)

Therefore in protomusical behaviour, as in measured music, the organising principle is pulse. Here, the same distinction is drawn between two loose categories of music as Merker defines; either music is measured, or it is not, which also echoes Biles' sentiment of arrhythmic music. This does not mean to say that in measured music the pulse is always fixed; Merker acknowledges that retardations or accents exist, but that these are deviations from a base pulse (Merker, 1999). Change

in pulse exists in juxtaposition to that base pulse, creating tension and leading to a resolution.

The pulse in music allows us to tap along and in doing so our tapping becomes entrained to that pulse. In many humans the behaviour of tapping along is subconscious, however it provides evidence for an interesting human trait: that humans have an extremely wide range of tempos to which we can entrain to. According to Merker, this is not so in other natural synchronisers such as insects, leading him to state that complex neural mechanisms must be present (Merker, 1999). It is not only the singular individual who has the ability to entrain to a pulse, but many individuals can mutually synchronise their pulses. Hence Merker states, “musical pulse is a cardinal device for coordinating the behaviour of those individuals in a joint, coherent, synchronised performance” (Merker, 1999: 316). It is a fundamental building block for musical group activity, even individual musical acts such as piano playing require a sufficient level of biologically internal synchronisation to occur.

According to Merker, there is good evidence to suggest that the synchronised chorusing of many animal species in the field is an epiphenomenon, arising out of a competitive strategy the males employ. Each male times their call in competition to be the first, thus causing a synchronising effect (Merker, 1999). Therefore this behaviour can be said to be self-organised synchronisation, as by definition self-organisation is also an epiphenomenon: it is a secondary occurrence, arising out of the primary behaviour.

Merker states that the human notion of music has evolved out of this protomusical synchronised rhythm, just as the human species and chorusing animals such as chimpanzees evolved from common ancestors.

Such an ancestral adaptation for entrainment to a repetitive beat would supply [...] an ancient biological foundation for the musical pulse no human culture has failed to feature among its musical means of expression. Indeed, if the present argument should turn out to have any merit, this adaptation for entrainment supplies an irreducible biological root of

human music (Merker, 1999: 319)

What Merker refers to as 'adaptation for entrainment' above can also be expressed as self-organised synchronisation. Here lies a crucial underlying question of this thesis: can self-organised synchronisation be considered an irreducible biological root of musical activity? To answer this question in great detail is out of the scope of this paper, however this important issue is explored.

Creativity and Computational Creativity in Music

Discussing *creativity* as a general notion has the same inherent limitations as discussing music. It is too vast a subject to cover here. It is important to note that whilst creativity is not limited to artistic activities such as music generation, matters related to creativity in general are discussed here solely in the context of music.

Leman offers a succinct definition of musical creativity which is linked to the idea of the novel.

A composer, performer, or improviser is considered to be creative when he or she does not merely repeat what has been learned or what others have done before but when a point of view is introduced that is unexpected and that adds new possibilities for further exploration (Leman, 1999: 285)

Boden (2004) divides this creative space further: a creative act may be novel in a personal context and not necessarily novel in a wider context. Boden refers to the former as personal creativity (P-creativity), and the latter as historical creativity (H-creativity). A P-creative act is just as creative as a H-creative act, however society may *value* a H-creative product more (Boden, 2004). This is not to suggest that musical creativity is focussed on the production of musical objects and their worth, as Leman states, “musical creativity is not a property of musical products but of persons that are involved in musical information processing” (1999: 285). The creativity lies in the subject undertaking a process, not the product.

Furthermore, the above definition does not consider Leman's own conception of musical creativity as a multifaceted concept, encompassing context, perspective, motoric and cognitive processes (Leman, 1999). To make a binary statement that A is creative and B is not is therefore too reductive. Pease et al. (2001) transform the question of creativity from the binary form above to the linear form: how creative is A? They also take the question to a multi-dimensional level: where does A lie in creativity space?

According to Leman, the romantic view of creativity, that people are “gifted” or touched by the divine, has been usurped by contemporary rational and pragmatic views. In the rational view, creativity is more closely linked to scientific creativity and as such formalised creative methods and automated processes have become part of the creative process (Leman, 1999).

Automation of music is by no means a new idea. Even in the romantic era, Mozart was composing music by throwing dice to generate the composition (see Cope, 1996). According to Leman, “the main effect of combining technology with musical creativity is that musical creativity was no longer believed to lay [sic] outside the human being” (1999: 288). Musical composition was not seen as a divine intervention, but coming from within the creator, or group of collaborative creators. Thus, the pragmatic view of creativity emerged, which embraces the fact that we know very little about the origin and nature of musical creativity. This spurred further research into computational creativity as a philanthropic and philosophic means to explore what it means to be creative.

When considering computer music systems, Leman states that often the musically creative idea originated in the human user or creator of a system. The system then 'unfolds' that idea computationally (Leman, 1999). Leman likens this to the contemporary practice in commercial music production where often the creative labour is divided between an 'artist' and the technical production staff. The technicians have just as much input into the final product, but it is still the artist who holds the intellectual property.

One major problem with creative systems such as AARON (Cohen, 1995), is that their output

quickly loses its novelty. The system itself is not able to adapt to the creativity space it is affecting. According to Elton (1995), this is an artefact of the fact that creativity cannot be modelled in mathematical terms. In mathematics, an established truth tends to stay true, whereas in art the notion of truth is in a state of flux. However, Colton (2009) recognises that we have more to deal with than clever hardware and software solutions to overcome this problem. He states that “building a truly computationally creative machine is as much a societal as a technical challenge” (2009: 1) and continues to cite the current limited skill, appreciation and imagination in a creative system as particularly large stumbling blocks.

According to Leman, “whether musical creativity can be fully automated remains a philosophical question” (2005: 118), due to the fact that no masterpieces have been created yet with fully automated methods. However, Colton takes a different approach to the subject, stating that “computers are not humans. This should be celebrated” (2009: 3). Introducing a machine as capable as a modern computer into the creative process requires a complete rethink of where the creativity lies with this extra dimension.

Case Study: Silent Barrage

Silent Barrage (SB) is a computational artwork installation piece which can be seen to be a stigmergic system.

SB can be broken down into three distinct parts. Firstly, there is a network of neurons, or 'brain', of the artwork. The brain is formed of real neurons, which have been grown and cultured under laboratory conditions away from the exhibition space. These neurons are interfaced with a multi-electrode array and send and receive data to and from the other parts of the installation via the internet. The second part of the work is the installation object itself, or 'body'. This consists of a grid of poles in the exhibition space, of varying number. Each of the poles corresponds to a region in the artwork's brain. Wrapped around each of these poles is a robot which responds to the signals sent

from the brain, by sliding up and down the pole and marking with a line around it the level of electrical activity. Over time, this provides a visual picture of the erratic neural activity in that particular region. Finally, there is perhaps the most important part of the artwork: the audience, or 'other', who participate by freely roaming the exhibition space. The space has been fitted with position mapping sensors and output of these sensors is fed back to the corresponding neural areas of the brain completing a closed feedback loop.

SB was created by the Australian artistic laboratory, SymbioticA, who operate as an interdisciplinary research group in biological arts. SymbioticA exists to offer “a new means of artistic inquiry, one in which artists actively use the tools and technologies of science – not just to comment about them – but also to explore their possibilities” (SymbioticA, 2010).

According to McRae, SymbioticA's work is mainly intended for artistic rather than scientific purposes, “[SymbioticA] is a group of artists who collaborate with scientists; therefore, what they are doing is making art” (2004: 7). Although there are areas of crossover into science, *SB* being a major example of such, where these crossovers do occur, it is conducted in the spirit of investigation rather than authority. SymbioticA place themselves in a position of questioning and challenging current thought on the cross over between biological science, technology and art; they promote debate rather than present answers.

SymbioticA state that *SB*, is “one of the very few real art and science works - in that it is both artistically meaningful and scientifically valid” (Gamblen et al, 2008). It is on the one hand an investigation into neural dysfunction, which may yield medical insights to conditions such as epilepsy, and on the other an expression of philosophical enquiry on the nature of thought and being, and the relation to humanity's own thought patterns.

SB markets itself as an Artificial Life (AL) artwork, and in early 2010 was awarded the VIDA prize for artistic creativity utilising new technologies and AL. In the world of AL, collaborations between art and science are relatively common. As McCormac and Dorin state, “less importance is attached

to distinctions between Art and Science than is attached to the philosophy used to approach the endeavour” (2001: 2); this certainly echoes the SymbioticA approach.

A major criticism of AL that *SB* overcomes is that of the 'ghost in the machine'.

Conventionally, [AL] is divided into three research fronts. *Wetware* is the attempt to create artificial biological life [...] *Hardware* is the construction of robots and other embodied life forms [...] *Software* is the creation of computer programs instantiating emergent or evolutionary processes [...] as long as AL programmes are considered to be simulations, any results from them may be artefacts of the simulation rather than properties of the natural systems (Hayles, 1999: 225).

What is atypical about SymbioticA's approach in *SB*, is that it incorporates all three fronts: the wetware is the brain, the hardware is the body and the software is the communication intermediary. Since *SB* is only in part a simulation, it raises the question of natural being, or at least natural neural activity, being just an artefact of the natural system of which it is a part. In other words, the posthuman notion of consciousness as epiphenomenon; consciousness as a self-organised phenomena.

Writing on emergence in AL, McCormac and Dorin warn of the “super-human ability to breathe real life into inanimate objects” (McCormac and Dorin, 2001: 3). This is what Hayles refers to as an AL narrative. AL narratives occur where “naming and interpretations function not so much as an overlay as an explication of an intention that was there in the beginning” (Hayles, 1999: 228). In *SB*, we can see this narrative come into play through SymbioticA's anthropomorphism of the exhibition into an entity: SymbioticA instruct the audience to 'calm down' the art work. This merges behavioural aspects with the machine aspects of the system and thus collides biological and informational domains. It also seals the position of the audience as a key component in the system. Without their interaction with the brain, an interaction facilitated by the exhibition space in which the body resides, the system would not be complete.

Through focusing on the audience's interaction with the system by actively involving the audience and positively encouraging their anthropomorphic view, *SB* creates a self-organised consciousness and can therefore be defined a stigmergic creative system.

Conversational Interactivity in Creative Systems

Hundreds of experimental [art] works are being created as artists struggle to understand the computer as a new medium for art. (Stern, 2001: 17)

As a theorist and game developer, Stern is interested in artwork that takes full advantage of the vast offerings a computer can deliver. Stern felt disappointed in what was on offer to him at the time which, whilst notably innovative, did not live up to his expectations (Stern, 2001).

Stern takes issue with the nature of the interaction between a user and work of art. Many art works on show are examples of what he terms *push-button art*. As a user interacting with this work, they are aware that if they push a certain button, the program will respond in a certain way. This reactionary model of interaction is based on more traditional ways of looking at the computer as a tool for achieving a set task, but Stern argues that it is not suited to the nature of interactive art. Instead, Stern advocates a focus on the *conversation* between the system and its user. Here Stern adopts the term conversation not to mean literal speech, although this forms part of the notion, but as a more general term for a meaningful communication between system and user; an exchange of information (Stern, 2001).

Conversational Interactivity (CI) is Stern's term for the aforementioned reactive aspects of interaction, combined with a sense of the system's own autonomy and adaptability. In a conversationally interactive artwork, when a user pushes a certain button, the system undergoes a process before taking any action and will perhaps respond differently, depending on a multitude of factors. Stern notes that the benefits of such a model include giving more power to the artist; as the experience of the artwork becomes much more personal and much less deterministic, the

communicative power of the work increases.

In terms of music and CI, we have to consider a conversation in more abstract terms than even Stern intended since the exchange of musical information cannot be said to be completely meaningful.

Nevertheless, links between language and music can be made: Marler (1999) states that human music was probably a precursor to human language.

A stigmergic creative system can be seen to be adopting the CI approach, even if there is no interaction with the user. Each agent is partaking in a conversation, facilitated by its interaction with the environment. If the user is allowed to become part of that interaction, by affecting the simulation environment, this too would be considered to be a CI system.

Stigmergic Creativity

According to Merker, the human notion of music has evolved out of protomusical synchronised rhythm processes. He also states that, “synchronous chorusing and dancing to a repetitive beat qualifies as music in the human sense” (Merker, 1999: 320), meaning that a transition from a synchronised rhythm to human-like music may not be that difficult to make. Etymologically, the Greek term for music originally included not only melody, but also dance and poetry, where the common feature is pulse based rhythmicity. Thus *music* had more in common with protomusic than a contemporary understanding of the term. Musicians, orators and even soldiers still share this same underlying protomusical synchrony inherent in the classical meaning.

According to Leman, in creative musical systems, “automated creativity [inevitably] becomes a high-level, and often even intelligent, tool for the organisation and manipulation of the creator's materials” (2005: 119). This thesis proposes to turn this notion on its head. Rather than focusing on high-level behaviour, the focus is on low-level creativity through biologically inspired protomusical behaviour. The intention is not to create an intelligent tool, but a tool for the stigmergic, self-organisation and manipulation of the creator's materials.

Chapter 3 – The Crickets System

System Aims

This thesis presents a system for simulating and generating protomusical behaviour. The system, named *Crickets*, provides an environment in which limit-cycle oscillators are subject to the stigmergic self-organised *local field coupling* (LFC) model set out in Chapter 1.

One of *Crickets'* aims is to follow a conversationally interactive (CI) paradigm. Each oscillator is abstractly partaking in a conversation with the other oscillators, through its interaction with the environment and its local field fluctuations. *Crickets* is also able to be interacted with by a user, who can control various LFC parameters, in addition to controlling the positioning of each oscillator. This provides an affect on the simulation without a direct control of the outputted data and thus avoids a 'push-button' approach.

Traditionally, the phenomenon of oscillator synchronisation has remained within an engineering discipline: chaotic dynamics. However, this science has now become a pervasive influence in contemporary culture, with artists, philosophers and commentators increasingly drawing upon the richness of these systems in their work. *Silent Barrage (SB)* is one such example of an artistic project that integrates many disciplines and has been widely praised in the art world for it.

According to Leman, artists are increasingly needing the support of scientists in order to explore these areas:

As tendencies toward integrated arts (theatre, music, visual art of all kinds) becomes more important, the domain of creative exploration gradually incorporates the whole audiovisual area. [...] This requires very advanced technology and scientific support. (2005: 113)

Colton (2009) agrees and hints that in the case where artists have explored these domains, a lack of skill has lead to the work not realising its full potential. This was certainly realised by SymbioticA,

the group behind *SB*, since the group was founded by a cell biologist, a neuroscientist and an artist, arguably giving their work more validity.

Leman goes on to state that,

Multimedia in art is no longer just a matter of bringing together different arts forms on the scene. Instead, the digital forum offers manipulation and integration of microlevels of information processing. (2005: 118)

This suggests that artists may benefit from modular scientific components in their art systems. The artist Lia, for example, has created many interactive generative software art pieces for different platforms. *PhiLia* and *Arcs 21* (Lia, 2011) are two examples of iPhone applications which incorporate flocking-like behaviour to create their visuals. This application, and many others inspired by flocking, could potentially have used Reynolds' (2001) *Boids* flocking model as a microlevel information processing module, to borrow Leman's term. In this case the artists would not need to concern themselves with the complete understanding of the model and its simulation in software, but can focus instead on how that model is utilised and realised in their work. This does not mean that *Boids* could be utilised in its current form for this purpose. An extra step would be required to turn *Boids* into a form suitable for use in this manner, to ensure it is flexible for use in software art projects.

This separation between scientific model and creative application was a key aim of *Crickets*. The final system is able to be utilised in a variety of different applications. As discussed above, protomusical behaviour has a variety of different applications in art such as poetry and dance, but it also has applications in science, in areas such as neural oscillation research (see Breakspear et al., 2010)

Crickets is a simulation tool, providing the microlevel of oscillator synchronisation modelling, and as such it is intended to be a component of a larger system. Here, *Crickets* is applied and tested in

music generation systems.

System Overview

Crickets encapsulates a qualitatively stigmergic oscillation synchronisation environment and interaction interface for that environment. A user interacts with the *Crickets* interface and the oscillation environment is affected. *Crickets* then broadcasts its data to other applications, known as patches. Each patch parses the data and produces some output as a result, be that audio or other. Thus the user is reacting to the output of the patch but interacting with the *Crickets* system, resulting in a closed feedback loop.

The *Crickets* application has been written in C++, using the OpenFrameworks toolkit.

OpenFrameworks is a cross-platform, open source collection of tools aimed at assisting 'creative' technology applications. Based on Processing, which is written in Java, it is used by many artists to create interactive visuals and installations which require the speed of a native C++ application. By choosing to develop *Crickets* with OpenFrameworks, it is hoped that other artists familiar with the framework would be able to understand the system with greater ease, allowing the potential for extensions to be build for it.

The majority of the computation in *Crickets* occurs in an implementation of the LFC model. In this implementation of LFC, the oscillators in the field are all Van der Pol oscillators for reasons stated in Chapter 1. The coupling co-efficient K has been implemented as a global parameter for every oscillator. This is mainly due to simplifying the interaction interface for the system. Allowing for specific coupling values between pairs of oscillators is possible, but creating an intuitive control proved to be a challenge. It seemed overly complex in terms of the usability of the system and so was dropped in favour of the global parameter.

Specific values for K between pairs of oscillators can still be achieved through the use of the influence variable, I . I is used to scale K depending on the distance between two oscillators.

Crickets implements one function for how I changes with distance. In this function, I remains at 1 until a certain distance is reached, then linearly falls to 0 as distance increases. This acts as a distance dependant weight between the two oscillators.

Appendix A shows a collection of screen shots which illustrate *Crickets* in action and the type of interactions that are available to a user. (Figure A-1) shows the visual interface. The darker area to the left represents the two dimensional space where the oscillators, referred to as crickets, are placed. Each cricket is represented visually as an expanding and contracting circle. One cricket in (Figure A-1) has a lighter gradient behind it which is displayed when the user hovers over the cricket or 'selects' it via a mouse click (Figure A-2). This also shows its identification number. On the right of the application is a control panel, which allows a user to add and remove crickets, place them randomly (Figure A-9) or in a grid formation (Figure A-8). All crickets can be selected by clicking the "Select All" button (Figure A-4). The control panel also contains a control to increase and decrease the simulation speed and a control to increase or decrease coupling. When coupling is a non-zero value, the crickets form a network of lines depending on the influences they have on each other (Figure A-5). When crickets are selected, more controls appear in the control panel that affect only the selected crickets. Noise represents the configurable intensity of a Gaussian random perturbation each oscillator's input is subject to; range controls how far the cricket can 'hear' (Figure A-6); Signal speed controls how large a delay there is on the signal, based on distance. In addition to these controls, the crickets can be dragged around in the environment through the use of the mouse, thus affecting their local fields and influences (Figure A-10).

Crickets uses the Open Sound Control (OSC) protocol to communicate its state with other programs. The OSC protocol was chosen for its flexibility in being able to communicate with many hardware and software systems. In computer music, it is often the protocol of choice and popular computer music software such as PureData, Max/MSP and SuperCollider are immediately compatible. In addition, hardware systems and circuit boards such as the Arduino units are

compatible with OSC. *Crickets* uses an XML document to allow the configuration of the OSC server and port.

There are two types of OSC messages broadcast: *update* and *field*. Update messages are sent for every cricket in the system on every simulation step, which currently runs at 60Hz. The message contains the cricket's identification number (ID), output level, selected state, and x and y coordinates. Field messages are broadcast once when the system starts up and subsequently when a cricket's local field has changed due to movement, change of range, or other interaction that affects the cricket's influence level. A field message contains the cricket's ID and information about the IDs and influence level of each cricket in its local field.

Three example patches have been created to demonstrate how these signals can be parsed and used for the creation of music. All of the examples are implemented using SuperCollider, a programming language for real-time sound synthesis. The first example, *Chaos Pad*, controls a group of continuous Saw oscillators through a band pass filter. The x coordinate of the oscillator controls the frequency of the Saw oscillation, the y coordinate controls the amplitude and the output controls the filter frequency. The second demonstration, *Drum Machine*, uses the output of each oscillator to trigger drum samples. As the output passes from below zero to above zero, a sound is triggered. The third and final example patch, *Bass Line*, is similar to *Drum Machine* in that the oscillators trigger musical events. However, *Bass Line* uses the Patterns library in SuperCollider to generate a simple melody and each trigger plays the next note in the generated sequence.

Results

Crickets' ability to be used for protomusical behaviour has been evaluated through user testing. An informal qualitative study was undertaken in which the system was explained to the users and they were given approximately one hour to interact with the system and the three patches outlined above. All the test users who took part were computer literate musicians with varying amounts of

experience using other music software. Some were also experts in fields such as performance art, teaching and design. The users were encouraged to talk aloud during their experiences and ask questions when they occurred to them. A printed questionnaire was completed by each user designed to assess their understanding, interest and the ease of use in interacting with the system. The users were slightly overwhelmed at first when they were shown *Crickets*. Since the number of oscillators at start up defaults to twenty, there was a lot of information presented to them with many control parameters at their disposal. This was before any sound was being produced at this stage. Many users felt that, after they had experimented with all of the patches, they understood the system more and wanted to go back to the earlier patches to try them again, this was the case even when the order that the example patches were presented was changed. It is clear here that an understanding of *Crickets* and what it is capable of comes through a variety of applications and users should be encouraged to try many patches.

All users found the *Chaos Pad* patch to be the most immediately intuitive. This is possibly due to some more direct mapping of x and y coordinates to the instantly audible results of frequency and amplitude. Patches where there were rhythmical events fired, such as the *Drum Machine*, caused some users to take great care over the positioning of each cricket. However, the users often did not take the time for the system to adjust to the new configuration before they changed some more parameters, causing a sense of frustration in both the user and the system, leading to noisy results. Whether immediately intuitive or not, all users found all patches engaging to the point where they felt one hour was not long enough to try all three completely.

A key theme in the users feedback was that of exploration. All users mentioned that by interacting with *Crickets* they felt they were exploring the different sounds and rhythms one could create. Often the results could not be predicted but a sense of order was present in the system. One user even commented on the emergent properties of the system, stating that complexity was achieved through simplicity.

All users said they felt like they understood the system after the demonstration was over, however when asked to describe the system each user had their interpretations. Some emphasised the idea of it being an interface for generative control, whereas others focused more on the interplay between the oscillator phase relationships. Some even described it as a potential musical instrument and here it was clear that the point where *Cricket*s ended and a patch began had become blurred.

Most users felt that their interaction with *Cricket*s was one of creative collaboration. *Cricket*s had the ability to innovate and output unexpected results, which the users interpreted as creativity in the system. Many users also felt the need to be creatively engaging back with their interaction to get the most from the system, showing evidence of CI.

All users immediately saw the potential for *Cricket*s to be applied in other fields and most expressed an interest to create their own patch for the system. Interfacing *Cricket*s with a video system was a common idea among the user group. Robotics, lighting, and musical effects were other areas of interest. All users also stated that they would be interested to see a more developed version for use in their own work, whether that be as part of an interactive installation, or sound exploration and generation for live music.

Chapter 4 – Conclusions

The phenomena of oscillation synchronisation is observable in countless places in the natural world. Living and non-living systems entrain and synchronise microscopic oscillators to form one large oscillator on the macroscopic level, which can in turn be synchronised. The way these oscillators achieve their synchronisation is largely accepted to be via self-organised processes and Kuramoto has provided a powerful and simple mathematical model for this. However, the neurological plausibility of Kuramoto's model has been questioned and the self-organisation aspect, in particular in terms of stigmergy, is not present. It has thus been reworked here into a stigmergic model for oscillator synchronisation: *local field coupling* (LFC).

The fields of chronobiology and biomusicology further elucidate oscillator synchronisation in living systems and the phenomena has been used to explain many forms of behaviours in those systems from their activity cycles to their development from birth to death. Even the animal origins of music have been suggested to arise out of synchronisation phenomena: a clearly rhythmic, but not necessarily creatively musical behaviour is achievable through synchronisation. This behaviour is termed *protomusical*: a behaviour that exhibits common musical features without being defined as music itself.

This thesis proposes that protomusical behaviour can be achieved through self-organised, stigmergic synchronisation, or in other terms: LFC. In addition, this could prove to be an irreducible biological root of musical *creativity*. A system, named *Crickets*, was developed in which low-level creativity is achievable through biologically inspired protomusical behaviour. This system is interactive and acts at Leman's (2005) microlevel of information processing, enabling the protomusical behaviour generated by the system to be reused in many applications across disciplines.

Many theorists have highlighted a problem in much contemporary art dealing with chaotic

dynamics or other highly specialised scientific discipline: a lack of scientific engagement, skill and experience (see Leman 2005, Colton 2009). As art exploring these notions becomes more pervasive, this problem will only widen. *Crickets* contributes a solution to this problem, as the user study has shown; it provides a tool that opens up the world of oscillator synchronisation interaction and helps artists to gain an understanding. It is also a useful scientific component able to be utilised in a multitude of art systems, a framework which other complex dynamical systems may follow.

Crickets does have its limitations and assumptions. Hayles (1999) has highlighted a stumbling block with many artificial life (AL) based simulations. Whilst they are simulated, and not embodied in hardware or other physical means, results and observations gleaned from them could be mere artefacts of that simulation and so discredited. *Silent Barrage (SB)* overcomes this problem through its combination of media (see Chapter 2), but *Crickets* remains a software system only. Many oscillators, Van der Pol oscillators (VDPOs) included, were originally implemented as hardware electrical circuits, however. A future consideration for *Crickets* would be to find a hardware implementation for *Crickets* using robotic, rather than simulated, agents. Perhaps a further step can be taken incorporating a wetware solution similar to *SB*.

According to Leman, “the ability to internally process and imagine sound is a feature of musical creativity” (1999: 286). This ability is not present in *Crickets*, since the embodiment of the sound is external to the protomusical environment. In addition, Marler's notion of *phocoding* is also implemented by the embodiment. According to Marler, *phocoding* is a fundamental requirement for the evolution of speech, language and music. To combat these issues, a further consideration would be to develop a way for the patch to communicate back to *Crickets* as a further feedback. In this way, the patch can process and imagine the sound and become a third party in the interaction, affecting the environment as the sounds are generated. The system could be fully autonomous if this were done, allowing a user to have less to be concerned with. Furthermore, if *Crickets* could accept incoming OSC messages, patches could be made solely for the intention of controlling the

environment. For instance, a physics engine could be created to move the crickets around, or an implementation of *Boids* could be made to simulate flocking taking oscillation into account. This has the potential to create some striking additional behaviours that dragging alone would not be able to achieve.

Whilst LFC has been shown to be resilient to noise (see Chapter 1), even when the noise level is reduced to zero in *Crickets* the environment is very noisy when there are more than two oscillators coupled. This results in the phases of the oscillators wandering between relationships, which can be beneficial in terms of generating new material, but not so much for a repetitive rhythm. This could be due to the relaxing action of the VDPO, with its slow build of output and its fast release. Since each VDPO in *Crickets* had a random damping parameter, which affects this build up, the output is accumulated at different rates and results in a chaotic value for the local field calculations. Camacho et. al. (2004) mention an 'in-phase' mode to VDPO bath coupling, which could be explored and added to LFC for greater accuracy in the local fields. With this in place, phase relationships would be explored through the delay function alone.

Crickets and its patches currently assumes that each oscillator is a VDPO and this is used in various parts of the system. A consideration for the future would be to experiment with different oscillator models to see if they perform as well as the VDPO in creating protomusical behaviour.

The interface for *Crickets* can be seen to be gesture based since the main interaction is dragging groups of crickets within the environment. According to Leman, a gestural interface “is an extremely important aspect of musical creation” (2005: 120) and with this in mind it may be worthwhile to improve the amount of gestural interaction in the system, perhaps through using a tactile user interface or a touch screen interface.

Some of the users evaluating felt that *Crickets* was a new type of musical instrument. Whilst this is an interesting perception by the user and shows a strong connection between interface and sound, it should be discouraged until patches for *Crickets* have fully considered that connection. This is

known as mapping. Hunt et al. (2002) state that in traditional acoustic instruments, the connection between the sound and the interface is due to the natural physics of that object. Thus, mapping has never been traditionally considered until recently when we have to consciously design it in systems such as *Crickets*. Hunt et al. argue that “the mapping layer affects the perception (and the playability) of an electronic instrument by its performer” (2002: 6) and is therefore an important consideration. This is not an impossible task with *Crickets*, since designers of patches for the system can consider it for each patch they create.

Bibliography

- Attia, J. (2004): 'Behavioural Rhythms of Land Snails in the Field', *Biological Rhythm Research*, 35 (1-2), pp 35-41.
- Beckermann, A. (1992) 'Introduction – Reductive and Nonreductive Physicalism', in Beckermann, A, Flohr, H and Kim, J. eds. *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism*. Berlin, Walter de Gruyter.
- Biles, J.A. (2007) 'Evolutionary Computation for Musical Tasks', in Miranda, E.R. and Biles, J.A. (eds.) *Evolutionary Computer Music*. London: Springer, pp 28-51.
- Boden, M. (2004) *The creative mind*. London, Routledge.
- Breakspear, M. Heitmann, S. and Daffertshofer, A. (2010) 'Generative models of cortical oscillations: neurobiological implications of the Kuramoto model', *Frontiers in Human Neuroscience*, 4, pp 1-14.
- Camacho, E., R and, R. and Howland, H. (2004) 'Dynamics of Two Van der Pol Oscillators Coupled Via a Bath', *International Journal of Solids and Structures*, 41, pp 2133-2143.
- Coduys, T., and Ferry, G. (2004) 'IanniX aesthetical/symbolic visualisations for hypermedia composition', *Proceedings International Conference Sound and Music Computing*. Paris.
- Cohen, H. (1995) 'The further exploits of AARON, Painter', *SEHR*, 4(2).
- Colton, S. (2009) *Seven Catchy Phrases for Computational Creativity Research*. London, Imperial College. Available from: http://www.doc.ic.ac.uk/~sgc/papers/colton_dagstuhl09.pdf (accessed 05th September 2011) .
- Cope, David. (1996) *Experiments in Musical Intelligence*. Madison, WI: A-R Editions, Inc.
- Elton, M. (1995) 'Artificial Creativity: Enculturing Computers', *Leonardo*, Vol. 28(3) pp. 207-213.
- Foster, R. and Keirtzman, L. (2005) *Rhythms of Life*. London: Profile.
- Gardner, M. (1970) 'Mathematical Games: The fantastic combinations of John Conway's new solitaire game "life"', *Scientific American*, 223, pp120-123.
- Gamblen, P., Ben-Ary, G., Gee, P., Scott, N. and Murray, B. (2008) *Silent Barrage*. (s.l.): (s.n.). Available from: <http://www.silentbarrage.com/project.html> (accessed 5th September 2011).
- Goldman, R. F. (1961) 'Varèse: Ionisation; Density 21.5; Intégrales; Octandre; Hyperprism; Poème Electronique. Instrumentalists, cond. Robert Craft. Columbia MS 6146 (stereo)', *Musical Quarterly* Vol. 47(1) pp 133–34.
- Hayles, K.H. (1999) *How we became posthuman: virtual bodies in cybernetics, literature, and informatics*. University of Chicago Press.
- Hunt, A., Wanderley, M. M., and Paradis, M. (2002) 'The importance of parameter mapping in electronic instrument design', *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02)*. Dublin.
- Knoester, D. B. and McKinley, P. K. (2009) 'Evolving Virtual Fireflies', *Proceedings of the 10th European Conference on Artificial Life, Budapest, Hungary, September 2009*. Berlin: Springer.
- Kuramoto, Y. (1985) *Chemical Oscillations, Waves, and Turbulence*. Berlin: Springer.
- Leman, M. (1999) 'Music', in Runco, M. A. and Pritzker, S. R. (eds.) *Encyclopedia of Creativity*,

Volume II. London: Academic Press.

Leman, M. (2005) 'Musical Creativity Research', in Kaufman, J.C. and Baer, J. (eds.) *Creativity Across Domains: Faces of the Muse*. London: Lawrence Erlbaum Associates.

Lia. (2011) *LIA – Software Art*. Vienna: (s.n.). Available from: <http://www.liaworks.com/> (accessed 14th August 2011).

Lloyd, D. (2007) 'Ultradian rhythms and clocks in plants and yeast', *Biological Rhythm Research*, 37 (4), pp 281-296.

Marler, P. (1999) 'Origins of Music and Speech: Insights from Animals', in Wallin, N. L., Merker, B. and Brown, S. (eds.) *The Origins of Music*, Massachusetts: MIT Press, pp 31-48.

McCormack, J and Dorin, A. (2001) 'Art, Emergence, and the Computational Sublime', *Proceedings of Second Iteration: a conference on generative systems in the electronic arts*. Melbourne, CEMA.

McRae, E. (2004) *A report on the practice of SymbioticA Research Group in their creation of MEART - the semi living artist*. (s.l.): (s.n.). Available from http://www.fishandchips.uwa.edu.au/project/Emma_text.pdf (accessed 5th September 2011).

Merker, B. (1999) 'Synchronous Chorus and Human Origins', in Wallin, N. L., Merker, B. and Brown, S. (eds.) *The Origins of Music*, Massachusetts: MIT Press, pp 31-48.

Pease, A, Winterstein, D, and Colton, S. (2001) *Evaluating Machine Creativity* Edinburgh, University of Edinburgh. Available from: <http://www.inf.ed.ac.uk/publications/online/0054.pdf> (accessed 5th September 2011).

Resnick, M. (1997) *Turtles, Termites and Traffic Jams: Explorations in Massively Parallel Microworld*. London, MIT Press.

Reynolds, C. (2001) *Boids: Background and Update*. (s.l.):(s.n.). Available from: <http://www.red3d.com/cwr/boids/> (accessed 5th September 2011).

Schuss, Z., Tier, C., Matkowsky, B. J., (1985) 'Frequency Fluctuations in Noisy Oscillators', *SIAM Journal on Applied Mathematics*, 45 (5), pp 843-854.

Stern, A. (2001) 'Deeper conversations with interactive art, or why artists must program', *Convergence: The Journal of Research into New Media Technologies*, 7 (1), pp 17-24.

SymbioticA. (2010) *About Us*. Perth: The University of Western Australia. Available from: <http://www.symbiotica.uwa.edu.au/home/about> (accessed 5th September 2011).

Theraulaz, G and Bonabeau, E. (1999) 'A Brief History of Stigmergy', *Artificial Life*, 5 (2), pp 97-116.

Appendix A – Screen shots

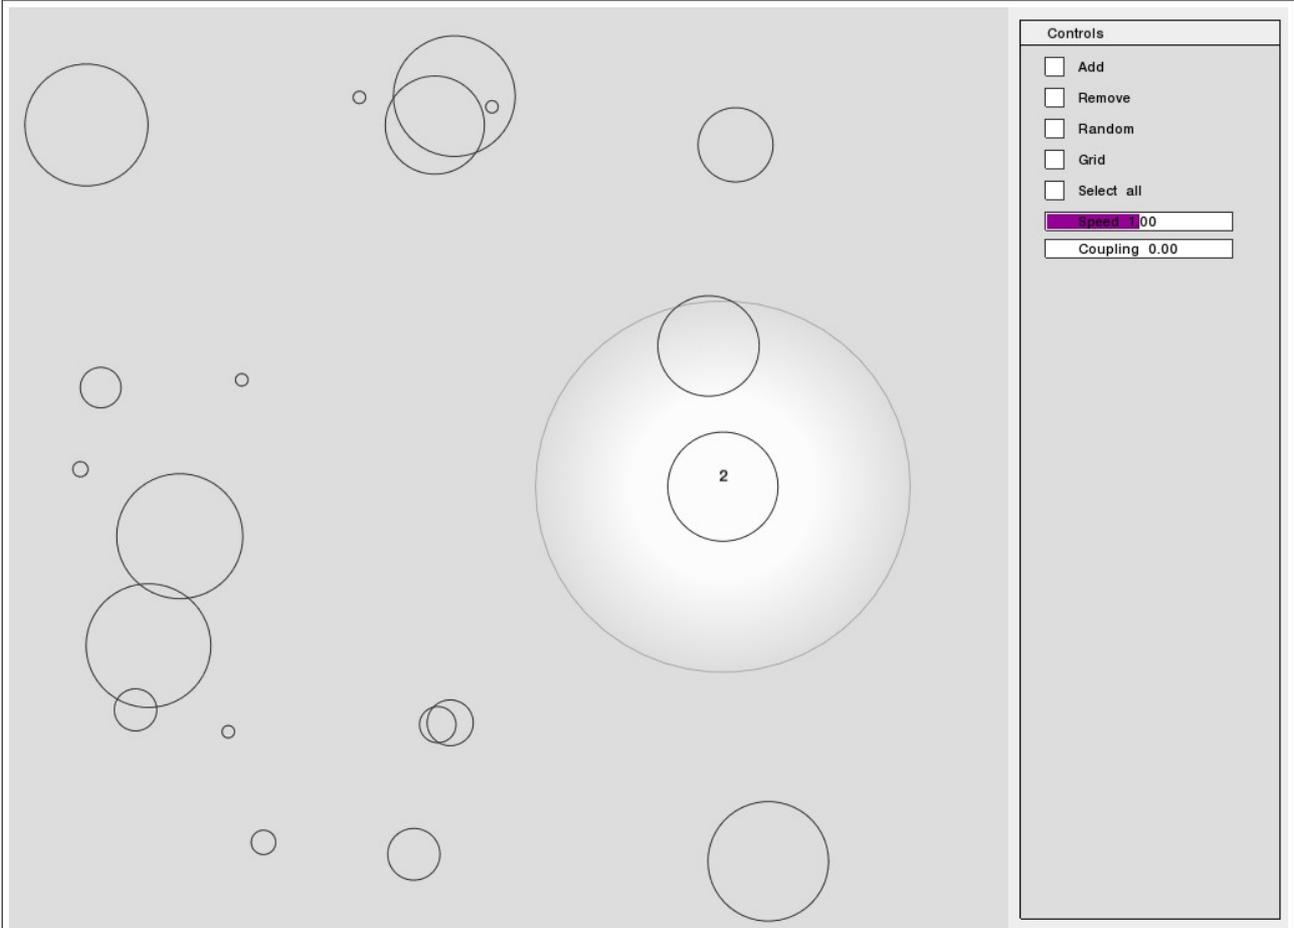


Figure A-1

The *Crickets* visual interface.

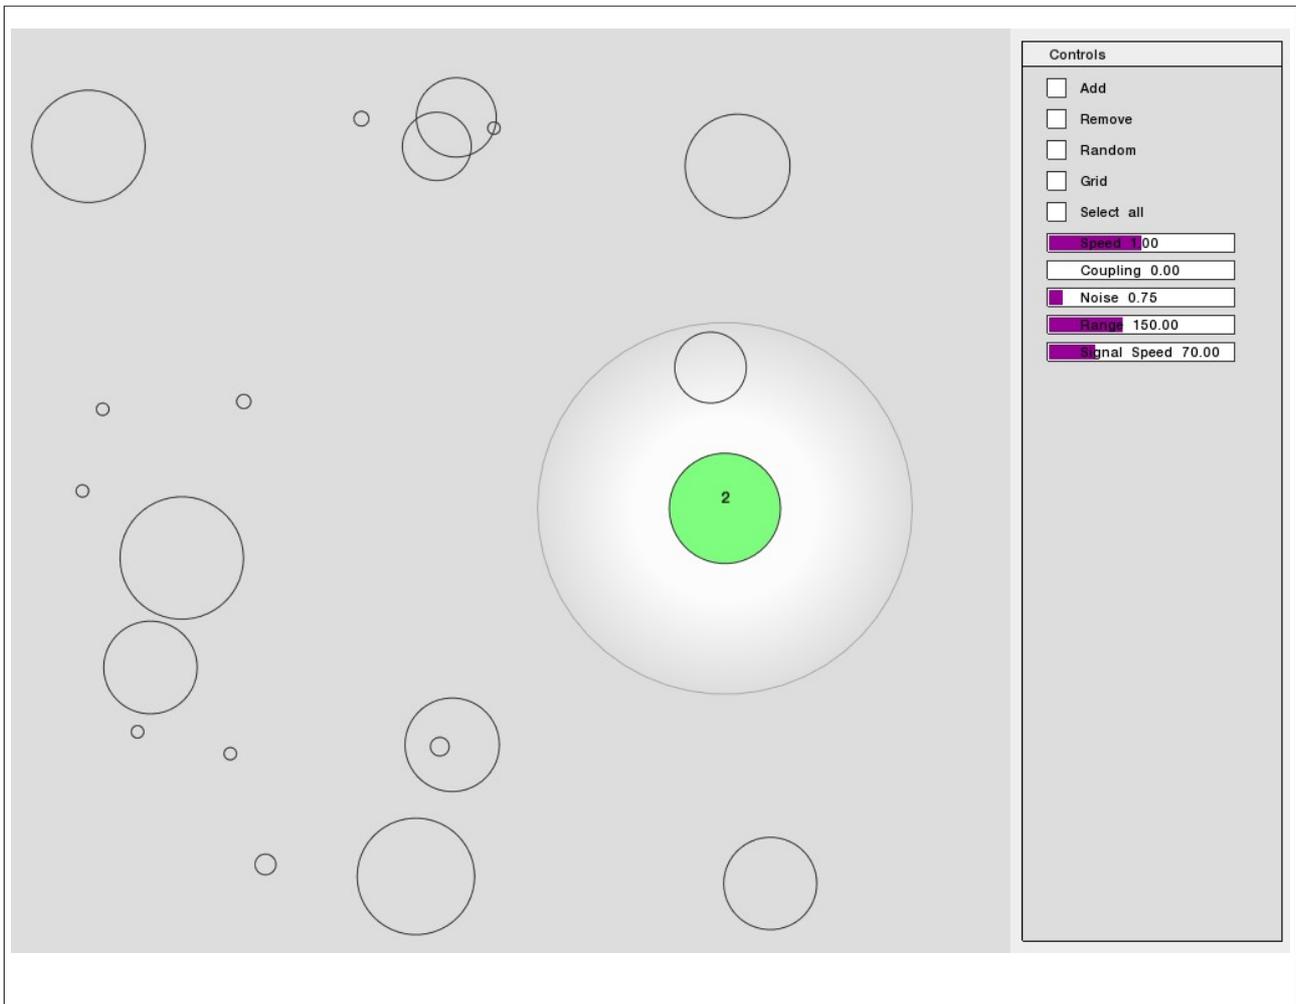


Figure A-2

Cricket 2 has been selected, bringing up the extra controls.

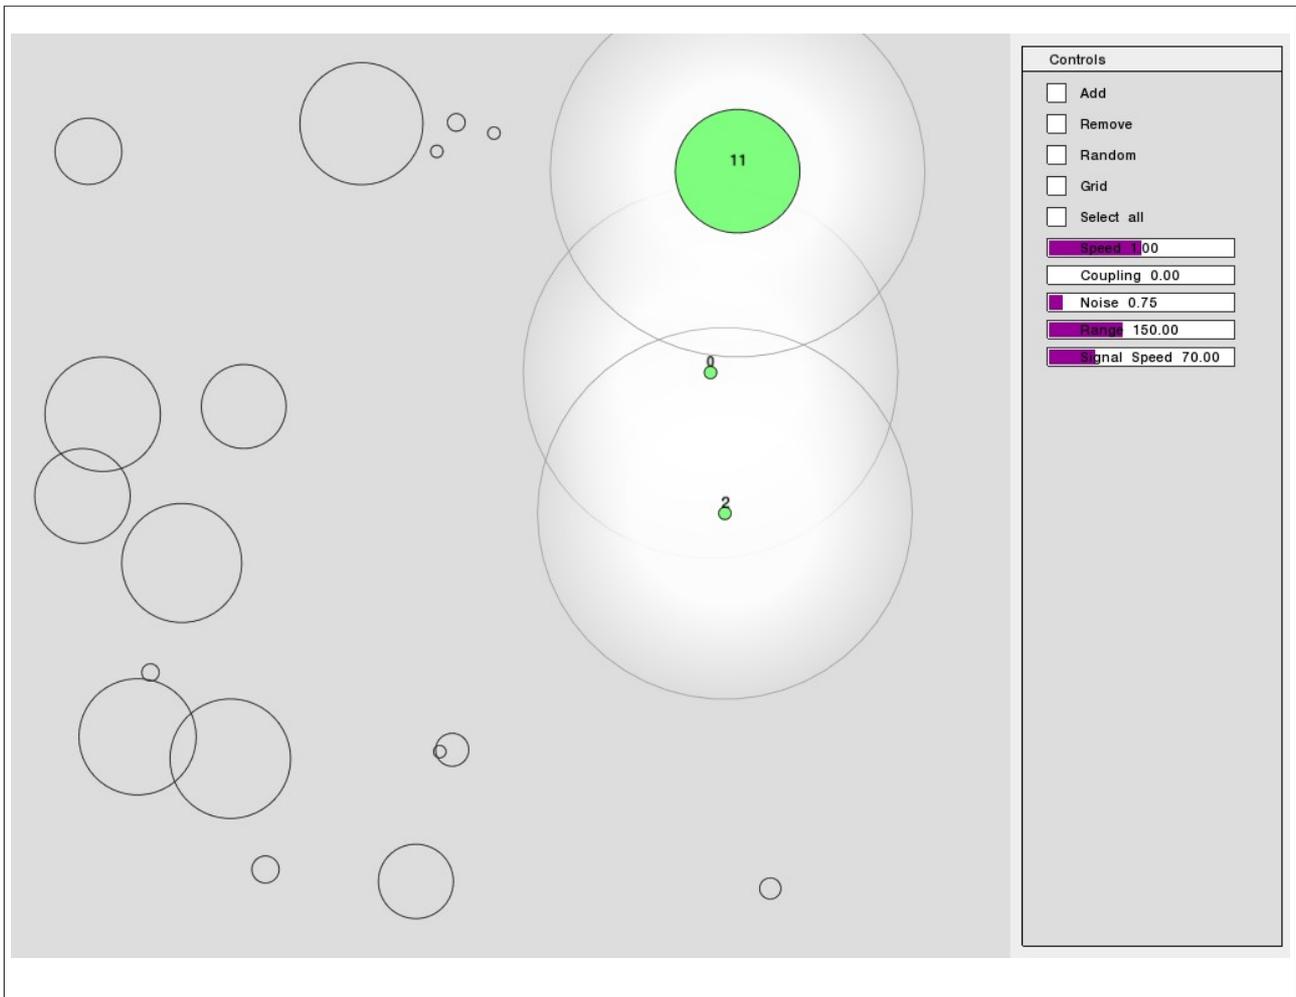


Figure A-3

Showing multiple crickets selected.

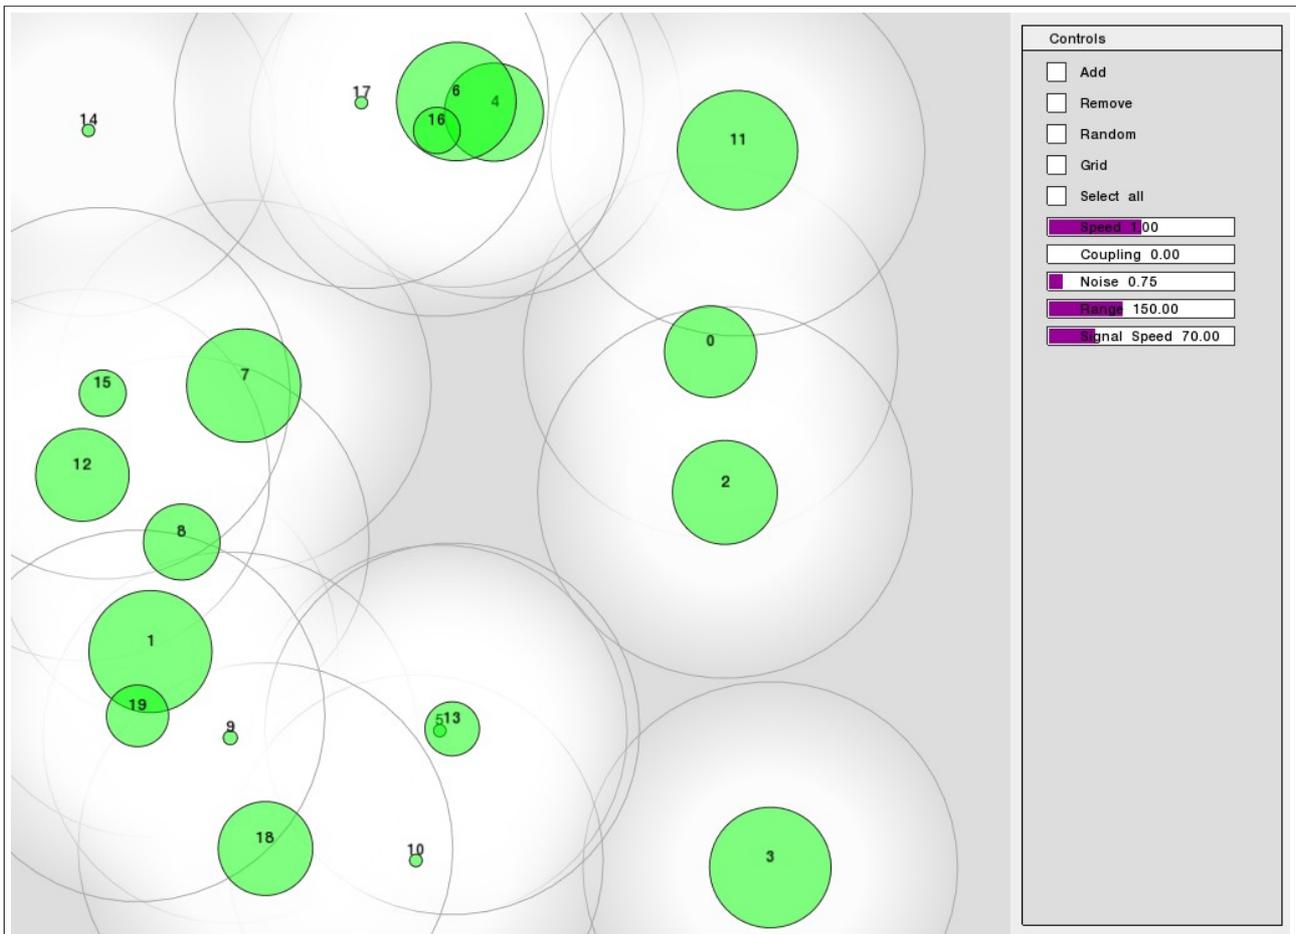


Figure A-4

All crickets selected.

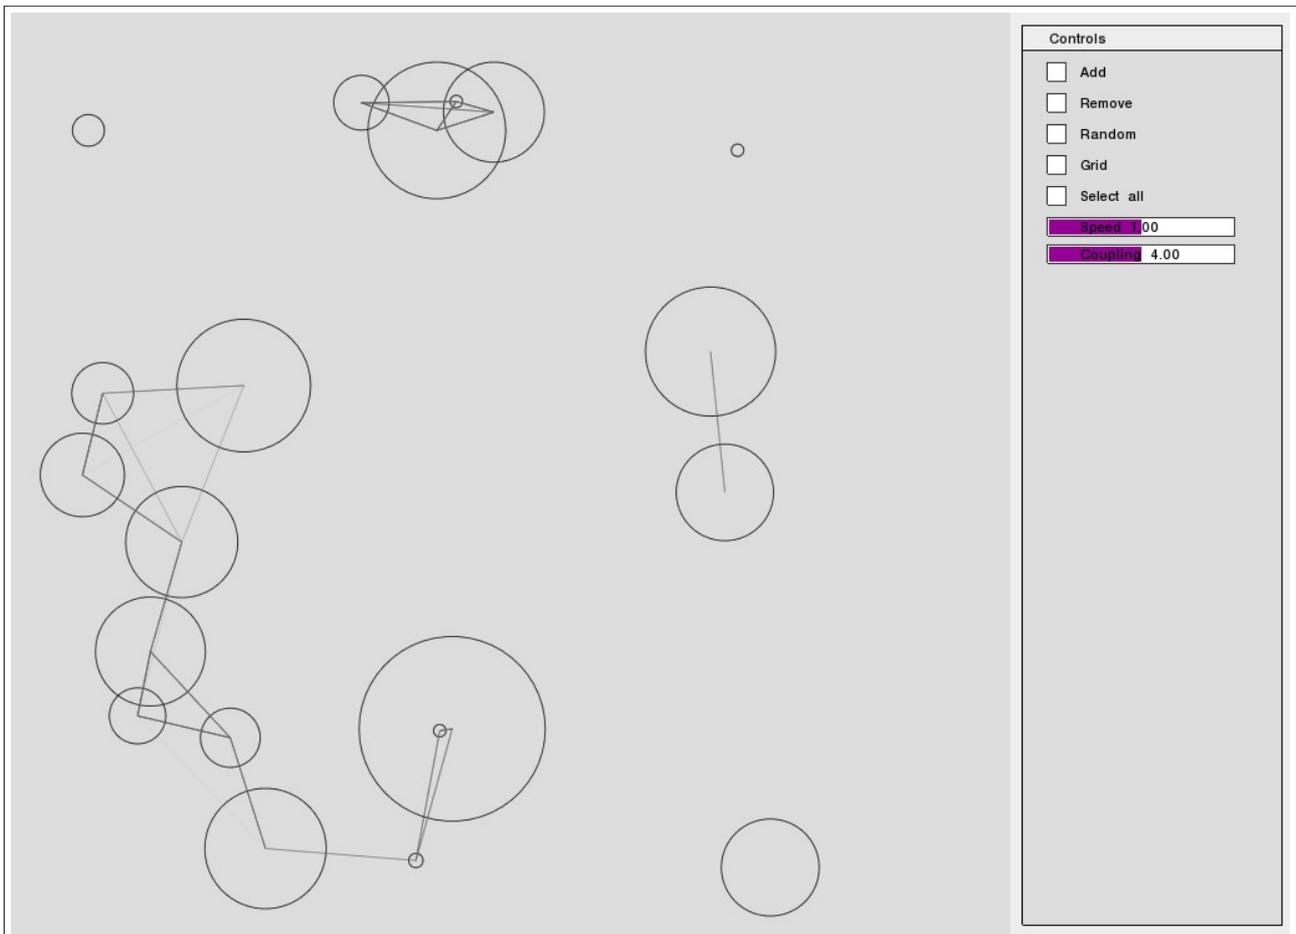


Figure A-5

Coupling enabled.

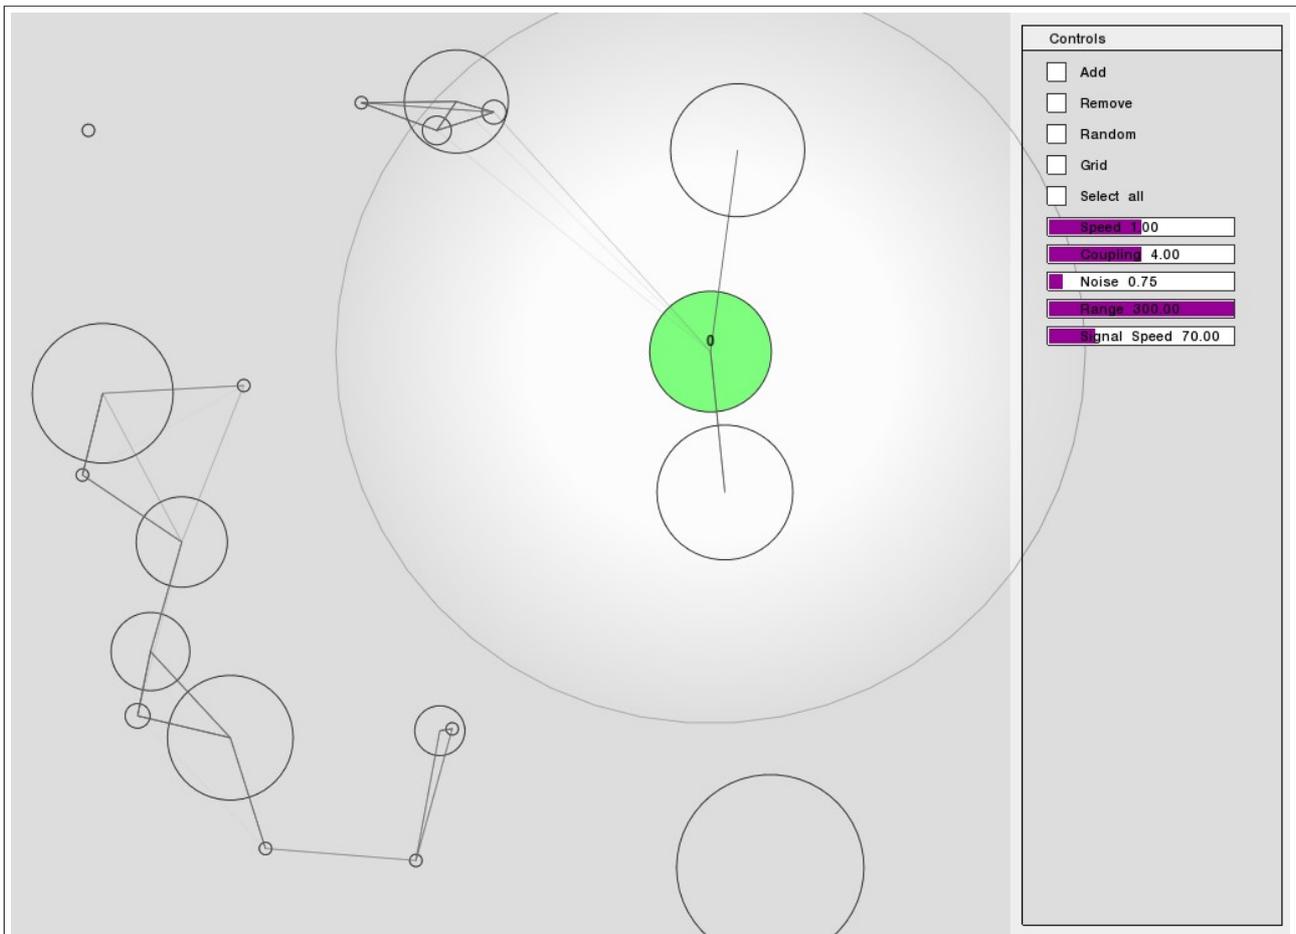


Figure A-6

One cricket's range increased.

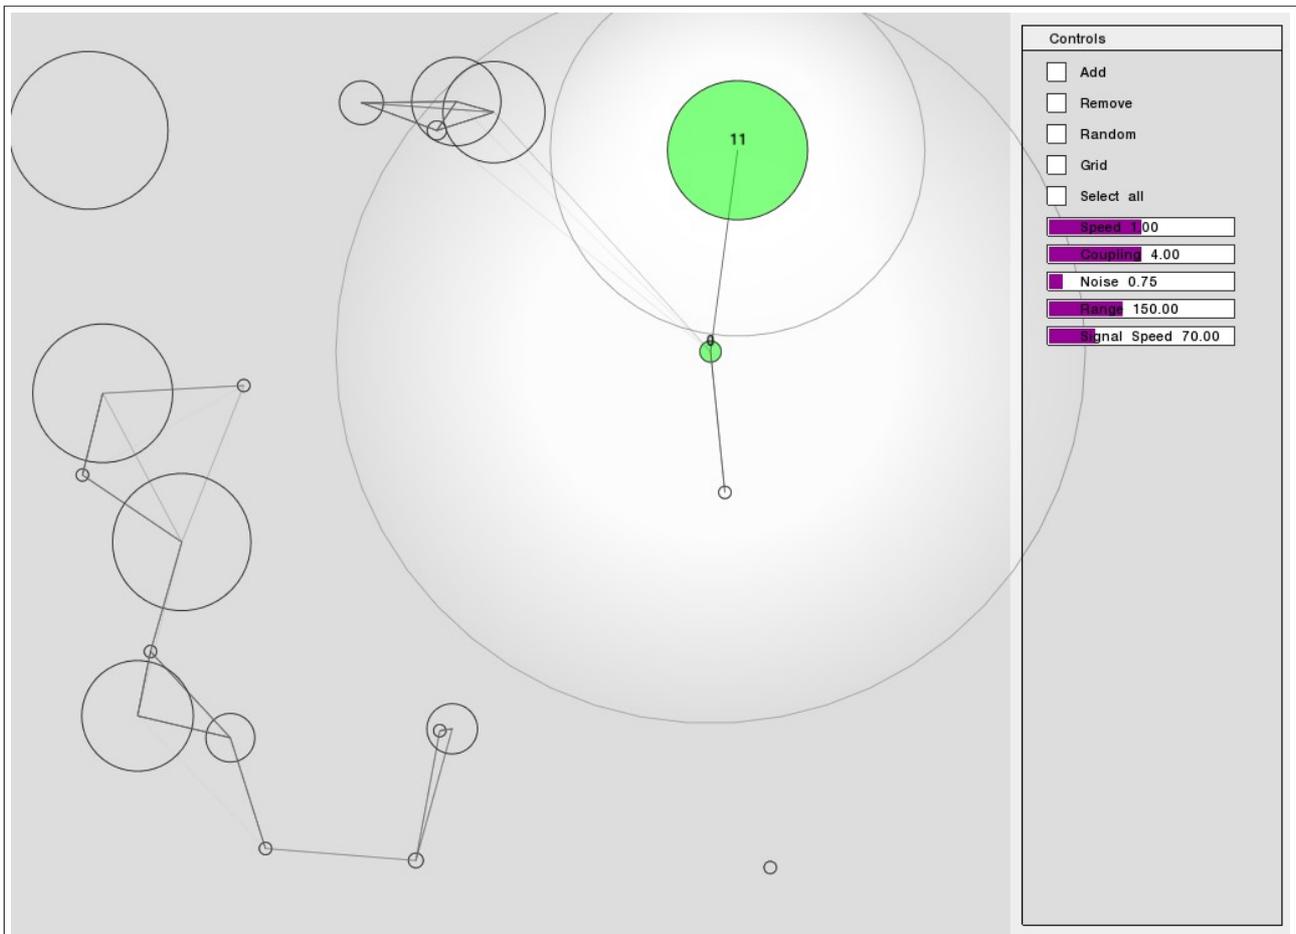


Figure A-7
Differing ranges

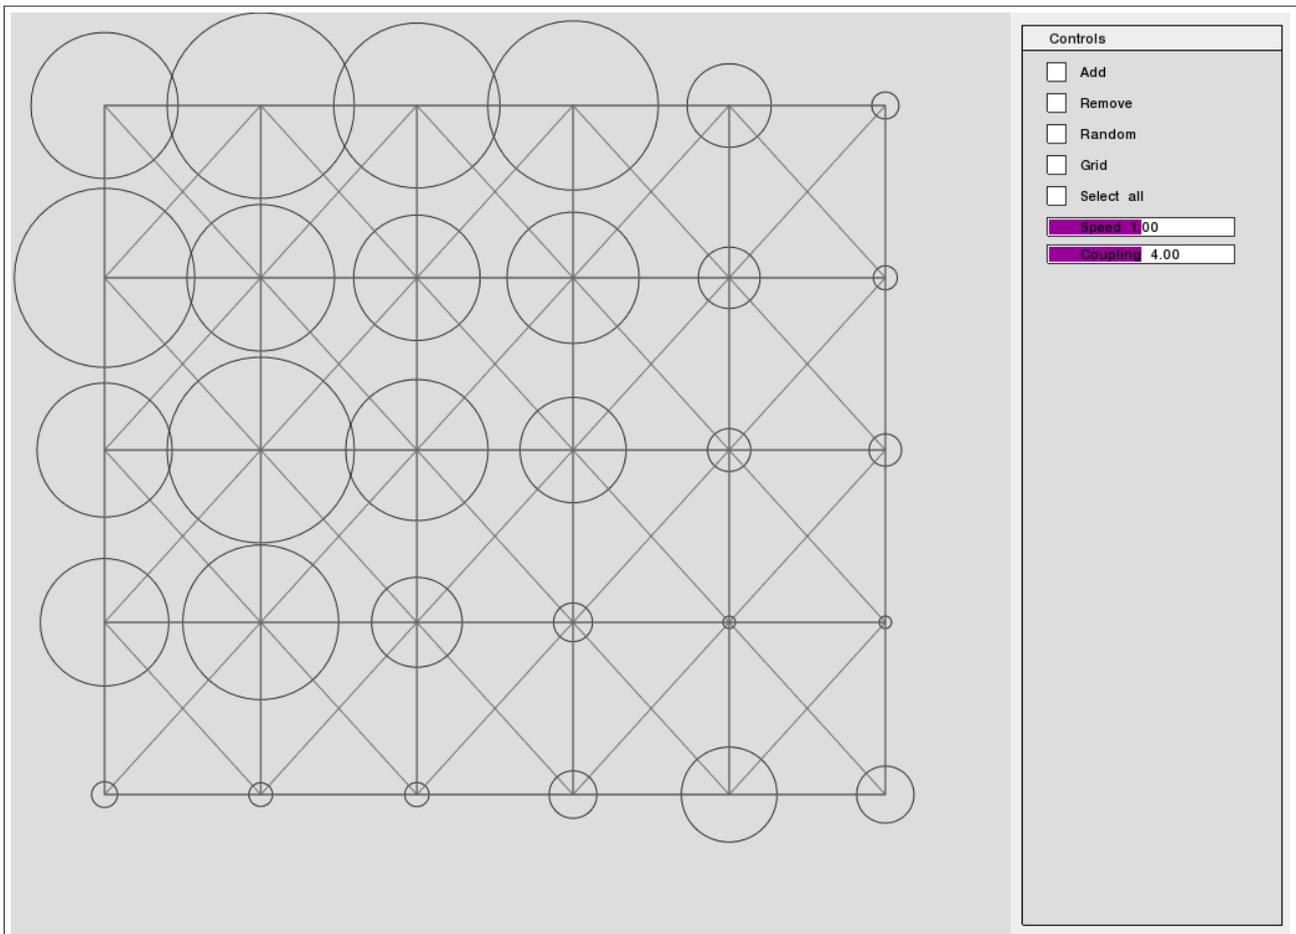


Figure A-8

Grid formation.

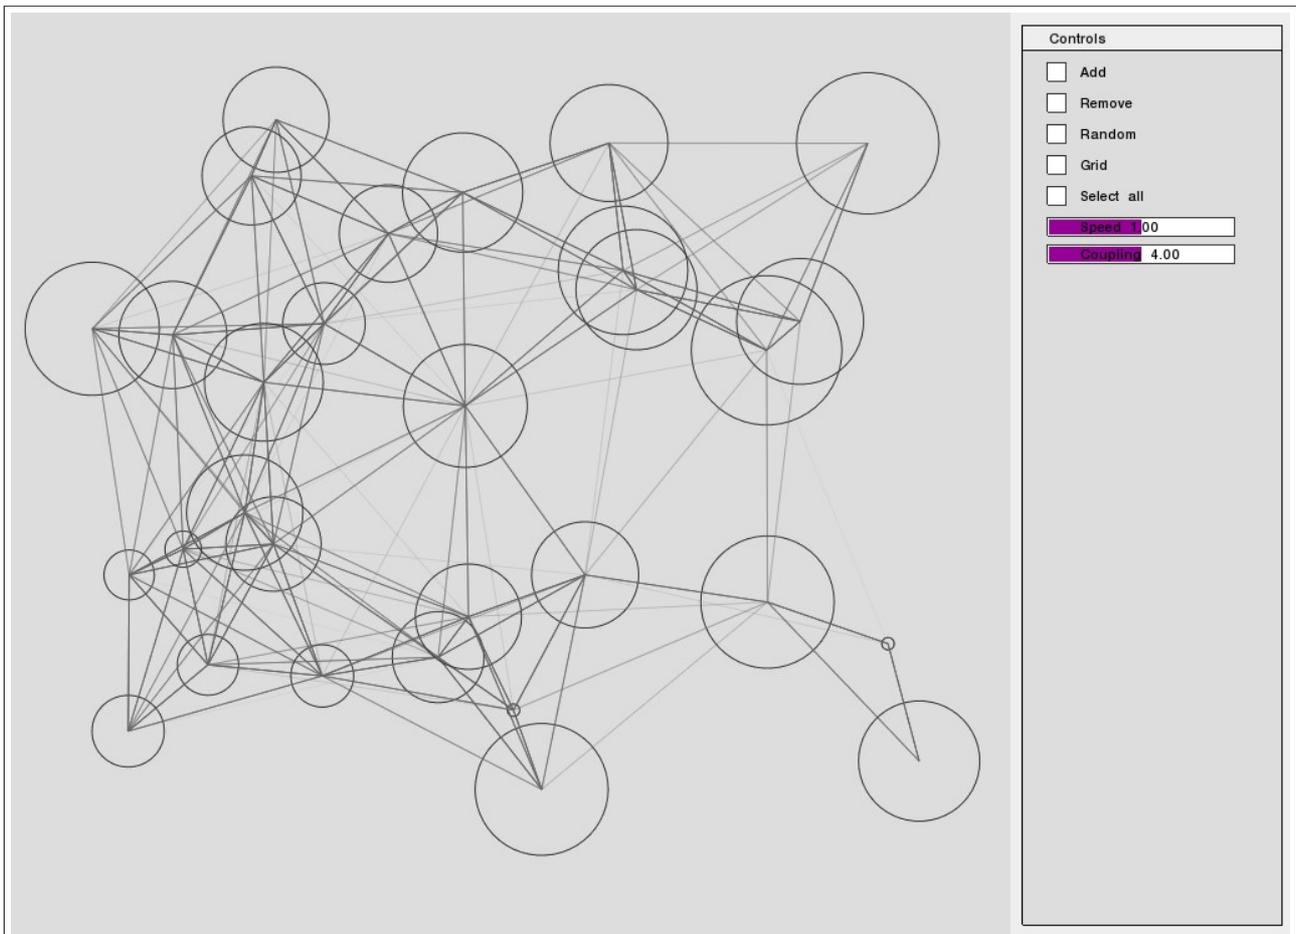


Figure A-9

Random formation.