Swarm Music: Improvised Music with Multi-Swarms

T. M. Blackwell

Department of Computer Science
University College London
Gower Street
London, UK
tim.blackwell@ieee.org

Abstract

This paper outlines Swarm Music, an improvisational musical system based on the dynamics of insect swarms. It is proposed that the interactions amongst performers in a freely improvising ensemble can be modelled using the self-organisational abilities of swarms. In a multi-swarm, each swarm is a musical entity which interacts with a target group. Each swarm is comprised of particles and each particle corresponds to a musical event. The targets are derived from other swarms, or from collaborating humans. The indirect interaction between swarms and humans is an example of stigmergy and can be expected to lead to temporal organisation. It is also proposed that the self-organisation of particles within a swarm leads to pleasing note patterns, so that each swarm is itself an interesting musically source.

1 Introduction

Real-time interactive music making with a machine is surprisingly rare. For example, a recent review finds only four systems that provide real-time music production (Miranda, 2001). Even rarer are systems that play improvised music – George Lewis' Voyager system is the most prominent (Lewis, 1999). Much effort has been spent on systems that only interact weakly with humans (for example, accompaniment programs that adjust pre-prepared material according to external parameters), or music composition programs, that embody some encoding of musical knowledge (as a grammar, for example) but do not produce music in real time. But the chief characteristic of improvised music is that it happens in real time and without revision. In fact interaction and improvisation have a symbiotic relationship, and this is to be heard most clearly in 'free improvisation'.

Freely improvised music (which differs from form-based improvisation as exemplified by jazz), is so named because it is free from many conventional musical constraints. In particular, no prior agreements are made on form or structure. Any temporal structure can only arise from the local interactions of the participants. In other words, it must be *self-organising*.

Examples of freely improvised music are to be found in the free jazz works of John Coltrane, Sun Ra and Cecil Taylor and, within the classical tradition, by composers such as John Cage, Cornelius Cardew and Karlheinz Stockhausen (this list is far from complete). A wide ranging review is given by Dean (1992).

Since improvised music operates without any prior agreement on tonality, pulse or dynamics, the question arises as to *what* the musicians actually respond to. One answer is that the musicians are attentive to the *expressive* qualities of their musical environment. Expressive qualities are high level descriptions of the music, including attributes such as event density, average loudness and pitch. The participants interact by either trying to match the expressive quality of the musical environment, as they perceive it, or by attempting to change it in some way.

One novel approach for the design of an artificial improviser is to abandon schemes with in-built musical rules, and to look towards nature for systems that are self-organising. Then, a mapping must be made from the extra-musical system onto actual musical events. This mapping would not encode musical syntax; musicality should be an emergent property (Blackwell, 2001, Blackwell and Bentley, 2002b). The collective behaviour of social animals offers a tempting analogy.

Insect swarms and avian flocks are examples of natural systems that exhibit self-organisation (Bonabeau, 1999) and seem to be analogous to improvised music. For example, a flock has a persistent shape and it may suddenly change its direction of flight, with every member almost spontaneously responding. Flocks and swarms may split, or they may join. Swarm members are 'attracted' towards the swarm as a whole, yet avoid colliding with their neighbours. This remarkable behaviour appears to be choreographed or led. In fact, merely local interactions between adjacent flock members are sufficient explanation (Reynolds, 1987).

The same de-centralised patterns are also evident in improvised music. Despite the lack of any pre-arranged form or material, improvisations can produce remarkable structures. For example, sudden changes in musical direction can arise whereby every member of the improvising ensembles departs on a new trajectory, as if orchestrated and conducted. And at a different level, the pattern of notes arising from the contributions of an individual may be swarm-like. The succession of notes may move around a key centre with a characteristic 'shape' – the pattern of rising and falling tones. Melodic variations can be visualised as gradual changes in shape. Notes may avoid 'collisions'. Similarly, the placement of notes in time gives patterns. There may be attraction to a beat, resulting in a simple organised pattern. Often, though, rhythms are a pattern of accents and placements about a beat, implying a degree of 'collision avoidance'. Chords are groups of notes with an avoidance of small intervals and an attraction to consonant intervals. Good harmony follows the rules of voice leading – notes make the minimum movement necessary for the harmonic progression. Chord progressions are 'attracted' to a key centre, but they may actually avoid the tonic for some time, delaying resolution.

The idea underlying Swarm Music is that the musical improvisations of an individual can be generated by a swarm of musical events. Individuals then interact with one another by expressive couplings. This interaction is realised by groups of attracting targets. Each target in a group is a captured event from another swarm, or from an external source. The swarm moves towards the targets, outputting musical events. These events are captured by other swarms, and become targets in their own target group. Here the term multi-swarm will refer to a colony of swarms with indirect target mediated interactions.

Previously published work on Swarm Music (Blackwell, 2001, Blackwell and Bentley, 2002b) outlines a single swarm system with simple animation. This paper extends that work by considering multi-swarms. The animation algorithm has been extended and the relationship of the parameters to the musical output is now better understood. Other improvements noted here are in event parsing and target placement.

Although Swarm Music may operate autonomously, it may also be directed by a process analogous to the conducted improvisations practised by the musicians Butch Morris and John Zorn and known as 'conduction'. In this approach, a vocabulary of gestures influences but does not specify the improvisations of a group of musicians. Recent work on conducted swarms is also described here.

This paper continues with an account of some of the theoretical principles underlying Swarm Music. The next section fills in some details and describes the three constituent modules in each swarming system of the multi-swarm. Then, the generation of a number of improvisations, is discussed in some detail.

The aim of Swarm Music is both to show that the dynamics within an improvising group can be modelled by interacting swarms, and to develop a practical application. In fact the two aims coincide, since a convincing performance by any such application would support the hypothesis. With this in mind, an attempt to evaluate the system from the perspective of the end user is given in section V. The paper closes with a discussion of the system and its place within the wider context of computer creativity.

2 Background

2.1 Self-organisation and stigmergy

The self-organisation of a system is the appearance of spatiotemporal structure due to low level interactions amongst its constituents (Bonabeau, 1999). There are believed to be four contributing mechanisms: positive feedback, negative feedback, amplifications of fluctuations and multiple interactions. For example, the pheromone trails of foraging ants are self-organised. The laying of a trail by an ant and the attraction of other ants to this trail is a positive feedback, but the evaporation of established trails is a stabilising negative feedback. Ants also wander randomly and may discover a new food source. This can then be amplified by positive feedback. Multiple interactions exist amongst the ants; they can make use of the results of their own activities and of the colony as a whole (i.e. they can follow their own trails, or trails left by other ants).

The interactions between individuals in a self-organising system may be direct or indirect. Direct interactions include actual contact between individuals, or contact mediated by sound, smell or vision. Indirect interactions occur when an individual modifies the environment, which other individuals respond to at a later time; this form of communication is known as stigmergy (Bonabeau, 1999).

2.2 Particle swarms

The animation of each swarm in Swarm Music is inspired by the classic flock animations of Reynolds (1987). In these animations, the individuals or boids are small objects with flapping wings, an orientation, a position and a velocity. Each boid experiences a number of accelerations. An attractive acceleration draws boids together, and a velocity matching acceleration encourages parallel motion. A third acceleration prevents collisions and enables the flock to fly around obstacles.

In a particle swarm, the finite-sized boids are replaced by point particles, and the velocity matching term is removed. In particle swarm optimisation, the attraction of a particle towards a local neighbourhood is replaced by attractions towards good solutions found by the particle and by the swarm as a whole (Kennedy and Eberhart, 1995). The particle swarm in Swarm Music, however, retains particle-swarm attraction, but introduces the concept of 'charge' (Blackwell and Bentley, 2002c). Charged particles experience Coulomb repulsion; this is collision avoidance and serves to enhance the diversity of the swarm, a useful property for dynamic optimisation.

2.3 Swarm Intelligence

Self organisation enables a colony to solve many problems (e.g. discovery of food sources, building nests). Swarm intelligent systems are directly inspired by this problem solving ability. Two recent publications demonstrate the use of ant colony and particle swarm algorithms in optimisation problems (Bonabeau, 1999, Kennedy and Eberhart, 2001).

Swarm Music uses self-organising particle swarms. Particles within the multi-swarm interact directly with fellow members from their own swarm, and indirectly with particles form other swarms via targets in their own target group. The size of the target group is finite and limited to the corresponding swarm size. This means that targets will be replaced as new events are captured – they have a finite life time. Moreover, the target groups will represent captured events from any swarm or external source, depending on the order of capture. The result is a stigmergetic relationship between the swarms. Positive feedback is present because targets are attractors for all particles, and because all particles in a swarm feel a mutual attraction towards the swarm centre. However negative feedback is also included in particle-particle and target-particle collision-avoidance accelerations. This negative feedback operates between particles in the same swarm and has the effect of introducing disturbances and maintaining the spatial distribution of the swarm (Blackwell and Bentley, 2002a). The disturbances are then amplified by stigmergy.

3 System Overview

Swarm Music has a number of independent systems, each responsible for one swarm in the multi-swarm. Each system has three functions: capture, animation and interpretation. The animator updates the positions of each particle. The swarm is then interpreted by mapping positions onto musical events in music space. Each event is then played on the internal synthesizer and visually displayed as a 3D animation. Concurrently with animation/interpretation, external events and events originating from other swarms are captured, parsed and placed as targets in the physical space of the swarm. There are a number of adjustable parameters for each function. These may be changed in real-time (conduction) or just left fixed (autonomous).

3.1Animation

Each swarm is made up of particles, and each particle has a position and a velocity in a Euclidean or 'physical' space. Each particle is subject to one or more of four accelerations. These accelerations are summed and added to the velocity. This updated velocity is then added to the particle position in order to advance the particle by one time step. The accelerations are indexed by *k* and described in Table 1.

Table 1: Particle Accelerations

k	***Description***
1	Attractive acceleration towards the swarm centre of mass.
2	Attractive acceleration towards the target group centre of mass.
3	Attractive acceleration towards a particular target.
4	Repulsive acceleration away from a neighbouring particle or from a particular target.

Each acceleration is a power-law of the form $\boldsymbol{a}_k = C_k |\boldsymbol{r}|^{n_k - 1} \boldsymbol{r}$ where $\boldsymbol{r}$ is the separation vector to a particle or target and n_k is the associated index. Negative values of C_k correspond to attraction, positive values to repulsion. For the case of negative index, the acceleration is set to a constant for separations smaller than a small value (core radius), in order to remove the singularity at $\boldsymbol{r} = 0$. Additionally, a limit of perception is set so that so that particles do not experience repulsions from other particles or targets at separations beyond this limit.

The first version of Swarm Music (Blackwell, 2001) used spring-like accelerations ($n_{1,2}$ = 1) towards the swarm and target swarm centres of mass and an inverse square inter-particle repulsion, n_4 = -2.0. Acceleration parameter C_3 was set to zero so that there was no attraction to *individual* targets.

Subsequently, different acceleration indices were tried, but the musical results did not vary much (whereas the indices would be more crucial for optimisation purposes). The important factors, from a musical perspective, are how much attraction and repulsion is present (determined by the acceleration constants) and which particles serve as attractors. Swarm Music currently implements all four accelerations, and allows the user to specify combinations and strengths.

The accelerations balance individual and group behaviour. For example, particles are attracted towards the swarm centre, which preserves the identity of the swarm as a whole, yet may also feel attractions to targets that lie outside the swarm, tending to fragment the swarm. Also, the inter-particle repulsion (which is only appreciable at small distances) balances the attraction towards the swarm centre, preventing collapse. Musically, this prevents the swarm from clumping together, thereby playing identical notes.

Table 2 specifies the main animation loop in detail.

Table 2: Main Animation Loop

k	***Description***
1	Attractive acceleration towards the swarm centre of mass.
2	Attractive acceleration towards the target group centre of mass.
3	Attractive acceleration towards a particular target.
4	Repulsive acceleration away from a neighbouring particle or from a particular target.

Here, each particle i, i = 1, 2…M has a 3-dimensional position $\boldsymbol{x}_i$ and velocity $\boldsymbol{v}_i$. The targets have a position $\boldsymbol{y}_i$ and are distributed in the target cube T of extent $\boldsymbol{x}_{max}$ in each dimension. A single swarm update or iteration now involves an update of each particle in turn. Each spatial dimension a = 1, 2, 3 has an associated clamping velocity $v_{max,\ a}$. No velocity component $\boldsymbol{v}_i.\boldsymbol{e}_a$ ($\boldsymbol{e}_a$ is the associated unit vector) is allowed to exceed this speed limit. Particles that do stray outside T are reflected back.

The **update parameters** are those specified in the main animation loop. The acceleration parameters C_k, limit of perception and core radius can be adjusted for each swarm. In addition, each of the four accelerations can be switched on and off. The components of $\boldsymbol{v}_{max}$ are adjustable between zero and an arbitrary maximum. If $|\boldsymbol{v}_{max}|$ is small in each dimension, the swarm will only vary slightly in shape from update to update. The result is a slowly evolving riff. Alternatively $\boldsymbol{v}_{max,\ pulse}$ may be small and $\boldsymbol{v}_{max,\ pitch}$ and $\boldsymbol{v}_{max,\ loudness}$ large. The swarm will then respond quickly to new events in the pitch and loudness dimension, but slowly to new pulses. New melodies will be generated, but each melody will share a rhythmic pattern. Finally, $\boldsymbol{v}_{max}$ can be automatically scaled according to particle-target separation so that the particles slow down as they approach a target.

3.2 Interpretation

Swarm interpretation is directly analogous to the interpretation of a score. A simple interpretation must be found that is musical and which does not presuppose too much. The motivation and details of interpretation have been described by Blackwell and Bentley (2001, 2002b). A brief summary will be given here.

Figure 1 shows a view of the swarm in physical space.

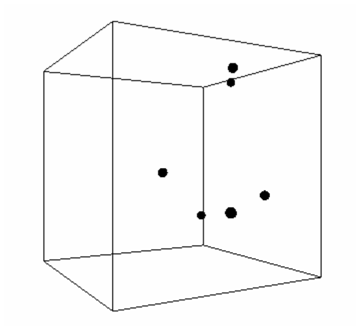

Figure 1: A view of the Swarm

Interpretation is a mapping from this space to the space of musical events – Music Space. Music space is populated by musical events, each one of which corresponds to a note played at a certain time and with a definite loudness. The three axes of music space are therefore loudness, pulse and pitch, which equates to the expressive parameters mentioned above (Blackwell, 2001).

Figure 2 shows a single event in music space.

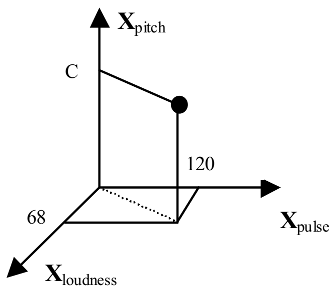

Figure 2: An event in Music Space

The event corresponds to middle C (C3), played at loudness MIDI 68, and sounding at a time interval corresponding to a beat of 120BPM after the preceding event. The mapping itself depends on the settings of the style parameters (see below) and is linear in loudness and pitch. The pulse mapping is more complicated and involves interpreting the physical separation along a given axis as a time interval Δt between notes. This is inverted to give a rate (i.e. beats per minute) and stretched so that very high and low BPM's are suppressed.

The **style parameters** govern this interpretation and how each note is played by the synthesizer. The swarm can be played by ascending order of pitch, and each note can be sustained. The rhythm of the notes can be stabilised by taking the pulse from the swarm centre of mass, rather than from each particle. The size of T may be adjusted to restrict the interpretation of a swarm to, for example, a specified instrumental range. Finally, the choice of MIDI instrument can be specified and the interpretation can be locked into a particular key and scale.

3.3 Capture

Each system captures both internal MIDI events generated by other swarms and external MIDI and audio events, parses these events for pitch, pulse and loudness, and sets targets in the physical space of the swarm. (The audio stream is split into events by a pulse height analyser – pitch is determined with a real-time fast Fourier transform.) At this stage there is much scope for directing the improvisations since the placement of the targets provides the stimulus for particle motion. Although the process can be completely transparent, a number of **scripts** can also be implemented to give the user some measure of control.

For example, one script will randomise the target group if the swarm converges on the target centre of mass. This will happen if there is no stimulus from outside the swarm, or if the stimulus is constant. The effect is to give the swarm a new musical direction in response to either an absence of ideas from its collaborators, or from too much similarity. In terms of the biological analogy, this is equivalent to the full exploitation of a food source; the swarm then scatters to enhance diversity until a new source is found.

Another script parses captured notes for key and scale (which can be tonal or synthetic). Interpretation is then made within this key and scale. Key and scale changes are invoked in response to harmonic changes in the input. The script also tolerates 'mistakes', only instigating key or scale changes if the number of incorrect notes surpasses a certain number. This is important since, within the context of improvised music where no form has been previously agreed, the only way to know if a note is an error or a deliberate dissonance (in which case no action should be taken) or an implied change (in which case some type of action *could* be taken) is to see if the event is repeated. The effect of this simple algorithm is to produce fascinating chases from key to key, from moments of bi-tonality (with two swarms) to periods of consensus.

Other new scripts include improved target placement. Targets can be scaled so that, for example, a swarm playing in a low bass range will adjust incoming events by subtracting octaves so that targets are placed in the appropriate pitch range. The responsiveness of a swarm to new target positions can now also be controlled by altering the target update time. Usually this is an almost continuous process, but the target update time can now be adjusted to longer intervals.

In another 'chordal' script, a number of particles can be allowed to sound at one time – part of the swarm is interpreted as a chord. This is governed by an adjustable parameter, Δt_{chord}, which is chosen by the operator; if n events are captured within Δt_{chord}, then n particles are sustained. The chord finishes probabilistically, where the 'notes-off' probability per iteration is also adjustable.

4 Recordings

A written analysis of musical improvisation is lengthy and difficult. However, two cds (Blackwell and Higgins 2001, Blackwell 2002) of Swarm Music improvisations with humans and of solo multi-swarm performances have been recorded and excerpts are available at http://www.timblackwell.com. Three of these excerpts are discussed below.

4.1 Follow Me (conducted improvisation)

This is an example of a conducted improvisation. Two three particle swarms were used. One of these swarms (A) plays sustained chords. At each update, three more notes are added to the growing chord. When the chord becomes too dense, the conductor kills the chord. Swarm B is playing a melody, mainly on a vibraphone, although the instrumentation does change during the piece. At certain places we can hear the swarm rapidly bouncing up and down in pitch. This was achieved by squeezing the pulse axis so that particles are interpreted within a narrow (and fast) range of beats per minute. At these times, and in pursuit of swarm A, swarm B began to oscillate with large amplitude in the pitch dimension. The 3D animation showed a gripping chase in music space. The behaviour was quite a surprise to the conductor!

4.2 Autumn Missed (autonomous)

Autumn Missed is a piano improvisation, performed by two five particle swarms - swarm A is playing left hand and swarm B is playing right hand. This was achieved by placing all targets below or above middle C for swarm A and B respectively. It is evident that both swarms are playing both melodies and chords. This was achieved by adjusting two sets of parameters. Firstly, MIDI "all notes off" messages are sent only at the start of each interpretation of the swarm (rather than after each particle has been interpreted) with probability one half. In other words, it is as if the sustain pedal is pressed down but momentarily released every five notes with probability one half – except the piano has two sustain pedals, one for each half of the keyboard! The second parameter that is set to achieve the melody/chord effect is $\Delta t_{chord.}$ This parameter is a time window for identifying chords. Since all MIDI events occur separately, the event parser needs to recognise chords as events arriving within a time interval. Normally this is set to a small value (5ms) so that chords can be clearly distinguished from very fast melodies. However, in this case, Δt_{chord} was set to 0.5 s. The event parser will force the interpreter to play interpreted notes as chords, where the number of notes in the chord corresponds to the number of notes in the captured chord. This setting will persist until two melody notes are heard, in which case all particles in the swarm will be interpreted as melody notes.

Despite the simplicity of these algorithms, when coupled with the self-organising properties of the two swarms, the effect is very musical, sounding very much like a piano improvisation of human origin.

4.3 Wind Up (autonomous ensemble with human)

This improvisation is an attempt to re-create a typical free jazz quartet with percussion, piano, bass and saxophone. In this recording a 3-swarm was used for the rhythm section (percussion, piano, bass) and a human supplied the saxophone improvisation. The bass can be heard to play a 'walking' line of steady tempo, only slowing towards the end of this excerpt. (This sense of pulse is typical of most of free jazz, which is not strictly freely improvised in the sense alluded to in the introduction. For example, much of the work of Ornette Coleman who is considered to be a founder of free jazz, retains a pulse and a song form, but is 'free' harmonically (Jost, 1994).) The percussion sounds were achieved by restricting the pitch dimension of one of the swarms to MIDI 35 – MIDI 52 and routing the output to MIDI channel 10, so that the each note corresponds to a percussion instrument on the standard GM drum kit. The bass and percussive swarms are well matched for pulse, but the piano-voiced swarm plays at a slower pulse with short melodies punctuated by chords. This was achieved with the chordal script described in section 3.3. The effect is certainly not unlike the 'comp' style of jazz accompaniment.

The correlation of activity within the multi-swarm is evident from the coherence of the roles of piano, bass and percussion throughout this excerpt. Particularly interesting from the perspective of self-organisation is a piano chord at 2:10 (min:s) near the end of the excerpt. At this point the pulse slows considerably form somewhere near 230 BPM to about 80 BPM. All three swarms take part in this temp change, possibly caused by the amplification of a fluctuation (as described in section 2.1). The saxophonist responds to the new direction.

The subtle response of the multi-swarm to the external input is harder to determine and requires careful listening. For example, two high notes from the saxophone at 1:20 prompt the bass line to rise at 1:25. Small crescendos and diminuendos involving the swarms and the saxophone are evident through the excerpt. In another pleasing episode, two multiphonic notes from the saxophone at 0:52 provoke another rising bass melody at 0:58 (targets being set from the higher harmonics of the captured multiphonic), and a flurry of saxophone notes is matched by a rapid beating from the percussion.

Overall, the excerpt does resemble free jazz, with evidence of interaction between the ensemble members, and a change in musical direction very characteristic of swarms.

5 Evaluation

5.1 Aims

The aim of Swarm Music is to produce a human-comparable improviser within the context of freely improvised music. This objective, if achieved, would indicate that freely improvised music is self-organising. The context is a ensemble improvisation, conducted improvisation or a solo improvisation. In conducted improvisation, the individual autonomy, as compared to an ensemble improvisation, is reduced since each member is expected to react to gestures emanating from the conductor. A solo improvisation which lacks interaction would have to be distinguished from a composition (also lacking in interaction) by the uncertainty of the musical output and the absence of any prepared musical structure.

5.2 Methodology

A modified Turing test is an obvious way of evaluating this aim. The test though is problematical since the listening experience of the subject is clearly very important. To the uninitiated (or even the most experienced) listener, free improvisation may always sound incoherent, in which case an artificial system would produce human comparable results just with a random output. On the other hand, a very experienced critic would begin to rate the performers, with the result that the artificial system would have to out-perform at least one of the humans if it were to pass. But it may be the case that the system is indeed human comparable but just not as good as the rest of the ensemble.

In the lack of such a test, a recent concert involving Swarm Music provides some indication as to how close Swarm Music is to fulfilling its objectives. Feedback from the musicians involved provides anecdotal evidence, and is prone to many problems of bias, but it can form the basis of a preliminary assessment.

5.3 Swarm Music in Performance

A concert, which was titled "Sonic Explosion", took place at Brunel University, UK, on February 26th 2003. Swarm Music was involved in three performances. The first was a duet with saxophonist Tim Whitehead who is a very experienced improviser. This was followed by a duet with Kathleen Willison, a professional jazz singer. The final piece was an ensemble improvisation with Whitehead, Willison, the composer Colin Riley (keyboards) and percussionist Robert Millet. The three pieces amounted to some 35 minutes of improvisation.

Both duettists had rehearsed with Swarm Music for 30 minutes prior to the concert, and interestingly chose different ways of working. Whitehead arranged three cues with the operator (me). Working with a single swarm, the style parameters were set so that the key and mode could not be changed by the swarm, but were only changed by the operator when cued to do so. This enabled Whitehead to force a change of key and mode. On the other hand, Willison wished to work with the set up of section 4.2 (Autumn Missed) - namely 2 autonomous swarms which can chose keys and modes from captured events. This results in periods of bitonality and introduces a degree of uncertainty since the pitch recognition algorithm can sometimes 'mishear', and identify a higher harmonic (Blackwell, 2001). However this can (and was) used to artistic effect by Willison.

It is important to comment on the involvement of the operator in the process, since the style parameters and scripts can be altered in real time, and whilst this does not amount to control of the multi-swarm, it clearly does have some influence. Again the improvisers had different preferences. Whitehead thought that any conduction could complicate his understanding of the swarm, so no adjustments were made to style parameters or to script. Willison was happy for there to be a degree of conduction, and the operator introduced the chordal script into swarm B from time to time, with occasional instrument changes.

The third ensemble improvisation was set up so that Swarm Music only received input from Willison. The ensemble could then react to the duet between Willison and the 2-swarm (and of course to each other). New musical directions initiated by the ensemble could therefore be indirectly passed to the multi-swarm via Willison. This configuration, which was chosen by Riley, is notable because he wanted to experiment with indirect as well as direct interactions between improvisers. This can be interpreted, from the point of view put forward in section 2, as stigmergy (I had not previously discussed these idea with him). In addition, the operator conducted the swarm, once more by switching the chordal script on and off and by making adjustments to the limits of the pulse dimension of either swarm.

5.4 Results

The audience seemed happy enough with the performance, and directly afterwards I interviewed the two duettists.

Whitehead thought that Swarm Music was "more assertive and much easier to converse with" in the ensemble improvisation. He also commented that "…the lyricism was so high that I felt compelled to follow it, mimic it, reply to it…" However he found the standards sampled sounds in the sound module were limiting "…needs some really good sounds!"

Willison was surprised to find in the first improvisation that Swarm Music seemed to be imitating her "(the swarm) hit the same note at the same time – the harmonies worked". However, there was some tension; "At times I would have liked it to slow down… it has a mind of its own…give it some space". Her solution to the "forward motion" of the swarms was to "wait and allow the music to catch up". In general she thought that "performers must be prepared to be open minded (and) sure of your notes or (the music) could fall apart… (but you) don't need to be the best improviser, it will react to you." Overall she did seem happy with the performance – "(it) felt like we were making music!"

5.5 Assessment

The positive reaction of the audience and the performers to the three improvisations, and indeed the willingness for professional musicians to participate, demonstrates that Swarm Music is, at some level, successful. However this does not by itself show that the aims have been achieved.

The remarks are again encouraging but need to be assessed critically. For example, it may be argued that the observation by Whitehead concerning "assertiveness" in the third piece, suggests that the increased activity of the operator is important. However, alteration of the style parameters and scripts does not determine in detail the melodic, rhythmic and harmonic content of the output, or the self-organisation of the swarms around the targets. Rather, the operator sets boundary conditions, and parameter values which *influence* the swarms. In other words, if the swarm produced random unorganised musical material, alteration of style parameters and scripts would not ameliorate the output. Operator meddling could of course diminish the musical output, and this might have been a factor in Whitehead's request for no conduction during his duet! The quality of his own response whilst playing, and his later comments on lyricism, suggest that he found Swarm Music to be a human-comparable partner.

Willison also found the swarms to be musically intriguing. Her frustration at the occasional disobedience of the swarms is also typical of human improvising partners, and is often a source of creative tension. To be disobedient also implies the ability to be obedient, and her remarks show that she did feel there were times when she and the swarms were collaborating, eventually giving her the impression that they were "making music".

Ultimately, the judgement will be aesthetic. For example, it is not enough just to demonstrate that the multi-swarms respond: they must be shown to do so in a musically fresh way. Also, the music must have appeal beyond the experience of the performers; with this end in mind, it is hoped that a recording of Sonic Explosion will soon be available.

6 Discussion

The recordings and concert demonstrate an ability to participate in improvisations with humans or with other swarms. An exciting new development, made possible with multi-swarms, is conducted improvisation. The recordings demonstrate that the multi-swarm is self-organising. Indeed, the results are surprisingly musical; little intervention is needed to sustain an interesting improvisation. In fact the result is almost too intricate. Whitehead remarked during the rehearsal for Sonic Explosion that there are "too many ideas!".

The crucial update parameters are now better understood. The introduction of velocity clamping in each dimension is a very useful device (evident on some of the recordings from the cd "Autumn Missed" (Blackwell, 2002) but not described here). The attractive accelerations are spring-like, and this might be expected to give rise to oscillations in each dimension. However the interactions amongst particles and between swarms distort this regular motion. Even so, a remnant of the oscillatory motion does persist and it might be one reason for the musicality of the melodies (Blackwell, 2001 and Blackwell and Bentley 2002b).

All the style parameters are selected by the operator (although scripts can override them - see below) and can be changed at any time. The system can be extended by adding dimensions to the physical space so that swarming would be possible along directions corresponding to other musical attributes. The duration of each note Δt_{event} is an obvious candidate.

Sonic explosion demonstrated that Swarm Music may perform autonomously, or may be 'conducted'. It would be interested to determine some heuristics that apply when an operator feels compelled to make adjustments to the style parameters and scripts. For example, if a random target script is switched on, then what prompted this decision? The music may have appeared boring to the conductor, but is boredom an absolute quality? The adjustable parameters will allow a performer to program in a 'boredom threshold', and maybe even a dynamic boredom threshold adjuster - but what determines this? These are complex issues and Swarm Music does not operate at this level.

However it would be very interesting if swarming in these higher levels could take place. To begin, the higher levels would have to be quantified so that attraction could be defined. For example, the boredom threshold (convergence radius around the targets) could itself become a target. This target would be placed according to analysis of captured events. If captured events clump together in Music Space for a certain time before separating and then re-grouping, an estimate of the boredom threshold of the source can be made, and a target set.

Another possibility for higher level swarming would be in frequency of repetition. Since a swarm can be made to riff by setting $\boldsymbol{v}_{max}$ zero, a riffing frequency of captured material can be worked out, and targets set along this dimension. Similarly, the rate of phrase development is determined by moving $\boldsymbol{v}_{max}$ away from zero, which would be another target, if this quantity was estimable from the captured events.

As it stands, Swarm Music can only structure its output by the self-organisation of low level events. It may, by chance, discover a sonata or a sixteen bar form, but this is very unlikely. Swarming in the space of possible forms is a fascinating idea, but it is hard to see how parsing the incoming low level events can lead to target setting in this space.

Although Swarm Music has been mainly used in an improvisational context where the music is performed without revision, it could also be used as a compositional tool. The context for this is set out in a publication on Creative Evolutionary Systems (CES) (Bentley and Corne, 2002). Within CES, systems evolution is open-ended rather than goal orientated, and this can lead to creative problem-solving. These systems have been used by designers, artists and jazz musicians to provide novel and creative solutions.

The multi-swarm is also not intent on a specific goal (such as the optimisation of an external objective function). On the other hand it is not aimless; attracting targets deliver a purpose, analogous to the search for food in an actual swarm. A composer can experiment with scripts which govern this search, retaining any interesting patterns that the swarms produce.

This paper has outlined some of the recent development in the Swarm Music project. Swarm Music will continue to perform with improvisers at various events. Undoubtedly the experience gained from live performance, where unprepared music must be generated by a machine, and in interaction with humans, will lead to further modifications and insights.

References

Bentley P.J. and Corne D.W. Creative Evolutionary Systems. *San Diego: Academic Press*, 2002

Blackwell T.M. Making Music with Swarms. *MSc thesis, University College London,* 2001

Blackwell T.M. Autumn Missed. *ownlabel cd 002, (available from author)*, 2002

Blackwell T.M. and Bentley P.J. Don't Push Me! Collision-Avoiding Swarms. *Proceedings of the Congress on Evolutionary Computation Piscataway* NJ:1691-1696, 2002a

Blackwell T.M. and Bentley P. Improvised Music with Swarms. *Proceedings of the Congress on Evolutionary Computation Piscataway NJ*:1462-1468, 2002b

Blackwell T.M. and Bentley P. Dynamic search with Charged Swarms. *Genetic and Evolutionary Computation Conference, Morgan Kaufman, San Francisco*:19-26, 2002c

Blackwell T.M. and Higgins R. Swarm Music. *ownlabel cd 001, (available from author)*, 2002

Bonabeau E., Dorigo M. and Theraulax, G. Swarm Intelligence: From Natural to Artificial Systems. *Oxford University Press*, 1999

Dean R. New Structures in Jazz and Improvised Music since 1960. *Open University Press*, 1992

Jost E. Free Jazz. *Da Capo Press, New York*, 1994

Kennedy J. and Eberhart, R.C. Particle Swarm Optimisation. *Proceedings of the IEEE International Conference on Neural Networks, IV*:1942-1948, 1995

Kennedy J. and Eberhart R.C. Swarm Intelligence. *Morgan Kaufman*, 2001

Lewis G. Interacting with Latter-day Musical Automata. *Contemporary Music Review* **18**(3):99-112, 1999

Miranda E.R. Composing Music with Computers. *Oxford: Focal Press*, 2001

Reynolds C.W. Flocks, herds and schools: a distributed behavioural model. *Computer Graphics,* **21**:25-34, 1987

Winkler, T. Composing Interactive Music. *Cambridge, Mass: MIT Press*, 2001