

Luke Harrald

Elder School of Music
Faculty of Humanities and Social Science
Adelaide University, 5005
luke.harrald@adelaide.edu.au

Fight or Flight: Towards the modelling of emergent ensemble dynamics

Abstract

This paper explores the 'ENSEMBLE' system, an agent based musical 'Prisoner's Dilemma' tournament designed by the author and implemented in cycling 74's MAXMSP programming environment. The aims for this system are wide ranging, with initial experiments focusing on an algorithmic composer's assistant that works in both real and non-real time and is able to manipulate both auditory and visual media. Extensions of the system involve the incorporation of real-time interaction with live performers allowing for collaborative music making.

The main premise behind this system is the modelling of the ensemble dynamics produced through the use of performance indeterminacy, pioneered by the 'New York School' (Earl Brown, John Cage, Morton Feldman, Christian Wolff) in the 1950s and equally having profound effect on experimental music through the 1960s (Morgan, 1991:359). Also often broadly referred to as non-idiomatic improvisation, the term performance indeterminacy is preferred by the author as to avoid confusion with Jazz based forms of improvisation which tend to dominate peoples perception of the nature of these situations.

Introduction

Since the late 1980s, with the popularisation of scientific endeavours such as Chaos, Artificial Life and Artificial Intelligence, many composers have looked to adapt various algorithmic approaches from these fields as a novel approach to the generation of music. Indeed, a number of impressive algorithmic systems have been produced (Miranda, 2001). The roots of these endeavours stretch back to the mid twentieth century, mirroring the history of the Artificial Sciences, and many concepts have been forced to wait for computer technology to catch up delaying further development; for a thorough prehistory see Johnson, 2003.

While many researchers have adopted biologically inspired models, particularly Cellular Automata (Bilotta et al. 2001, Harrald, 2003) and

Genetic Algorithms (Jacob, 1995, Todd and Werner, 1999, Wascha II, 2001), a number of researchers interested in generative music have become more focused on the simulation of more specific musical problems. Ambitious research in this direction has been conducted by Miranda through the use of *mimetic* models. Drawing inspiration from research in linguistics that aims to reveal the origins of language using computer modelling and simulation, Miranda's research has equally lofty goals, aiming to model the emergence of culture, and the origins of music through using adaptive games (Miranda, 2002).

The ENSEMBLE system presented here draws inspiration from this type of approach as it focuses on the simulation of the dynamical systems of music, particularly the social dynamics involved in music performance. Essentially modelling an ensemble of improvisers, in spirit, this work is closely related to other behavioural explorations of improvisation, notably Blackwell, 2002.

Background**Performance Indeterminacy**

While the term 'improvisation' holds many connotations for different people, particularly with regards to sentiments about deep personal expression, and equally cultural background, perhaps a discussion of 'Performance Indeterminacy' will remove some of this emotional baggage, and shed light on the types of musical behaviours explored through ENSEMBLE.

After composing 'Music of Changes' in 1951, an early work using 'chance operations', or randomness in computer science terms, American Avant-Garde composer John Cage faced a dilemma. Through using chance, and consequently removing meaning from his work, he had reduced the performer to 'what you might call an automaton' (Cage, in an interview, Corbett, 1994). In order to fix this imbalance, Cage redefined the role of the performer through using performance indeterminacy.

Generally, indeterminacy gives the performer freedoms either to make choices about how they will play given material, or choose the material they will play in a given framework. Through these freedoms, the performer makes choices spontaneously while performing the piece, and as such,

no two performances of the work are alike. Equally no individual performance of the work is definitive, just an example of how the work could sound. In describing this situation, Cage points out "the function of the performer is that of a photographer who on obtaining a camera uses it to take a picture. The composition permits an infinite number of these, and, not being mechanically constructed, it will not wear out, it can only suffer disuse, or loss" (Cage, 1961:8).

While all the members of the New York School (Brown, Cage, Feldman, and Wolff) incorporated performance indeterminacy into their works in different ways, it is the work of Christian Wolff that has the most relevance to this research. Wolff's exploration of indeterminacy involved setting up unpredictable chains of performance situations that could only arise through the act of performance. These works present the performer with a precise set of rules that are triggered by the actions of the other performers. This is most easily explained through an example:

"For an orchestra made up of at least 15 players, each of whom chooses one to three sounds, fairly quiet. Using one of these each time, you have to play as simultaneously as possible with the next sound of the player nearest to you; then with the next sound of the next nearest player; then with the next nearest after him, and so forth until you have played with all the other players present."
(excerpt from 'Burdocks', Wolff, 1971).

What is interesting about this example is that it meets many of the criteria that would qualify the musical result as being 'emergent', in many respects reminiscent of flocking behaviour (see Reynolds, 1987).

Firstly, the work is non-linear in that its construction is impossible to derive by analysing a recording of the work in a musicological sense. Secondly, the performers of the work act autonomously in that once Wolff has outlined the process that governs the work, the actual way in which it plays out is purely governed by the interactions between the performers. Finally, governed by simple local rules, the performers are

able to produce a complex global result, allowing an interesting musical structure to emerge. The work of Cornelius Cardew has also been discussed in this light (Johnson, 2003) and equally, more recent examples can be seen in the game pieces of John Zorn (2002).

The exploration of the emergent dynamics of these types of musical works is the main premise under which ENSEMBLE has been developed. The main goal has been the design of an agent model that is able to simulate the ensemble dynamics produced by the performers in this type of performance environment, in turn, leading to the bottom up generation of musical forms through the autonomous interactions of the agents.

The Prisoner's Dilemma Game

As ENSEMBLE aims to simulate what is essentially a social interaction problem, it was suspected that the Prisoner's Dilemma Game could offer a possible framework for further exploration. The Prisoner's Dilemma has been found useful in a wide range of interaction problems ranging from understanding the dynamics between foraging fish and divorcing people (Axelrod, 1997: 6). As this single model is able to capture the fundamental features of such disparate situations without the need for representing the details, it stands to reason that it could also offer insight into the social interactions of music.

A typical 'Prisoner's Dilemma' scenario is as follows: Two drug dealers are caught by police and placed in separate rooms to be interrogated. While the police have the prisoners on possession charges which carry a possible sentence of one year in jail, they suspect that they are dealers but have no proof of this. The public prosecutor then offers both prisoners a deal. Confess to being a drug dealer and give evidence against your partner and you can go free while your partner serves five years in jail, if you both confess you will receive three years in jail for pleading guilty. The dilemma for the prisoners here is that while if neither confesses to dealing, they will go to jail for a year, they can't guarantee that the other won't confess and take the plea-bargain, so the only rational thing to do is confess as a three year sentence is a better outcome than a possible five years in jail.

While in a one off situation like the one outlined above there is no incentive for the two participants to co-operate, the scenario changes somewhat if the process is iterated allowing for multiple encounters between the participants. Suddenly a whole myriad of strategies are possible as the participants are able to take into account pre-

vious interactions in order to shape their response in future interactions.

The question of what is the 'best' strategy in this situation was approached by Robert Axelrod through the organisation of two computer tournaments in the early 1980s. Experts in game theory were invited to submit strategies for the iterated prisoner's dilemma which were played off against each other in a round robin tournament (Axelrod, 1984: 27-54). The winning strategy in both tournaments was the 'TIT FOR TAT' rule by Anatol Rapoport, although it has been found that in the Prisoner's Dilemma, much like in real life, there is often no best strategy, but that each strategy's strength lies in the environment in which it functions (Campbell & Sowden, 1985).

The ENSEMBLE model

The model at the heart of ENSEMBLE is a deterministic Prisoner's Dilemma tournament that draws inspiration from Axelrod's work but has some clear distinctions enabling agents in the system to more closely model a group of interacting musicians. These differences are hardly surprising, as pointed out by Waschka in a discussion of musical applications of Genetic Algorithms, as music is a unique problem, and therefore the models used require unique treatments (Waschka II, 2001).

The current model consists of eight interacting agents, each of which is preset to a particular strategy prior to the first round of interactions. The strategies typically are randomly selected, but can also be individually set by the user. Currently seven rule sets are implemented:

TIT FOR TAT: co-operate in the first round, do whatever the opponent did in the previous round for all subsequent rounds.

RANDOM: a random or irrational selection.

VENGEFUL: co-operate until defected against, then defect for the next five rounds regardless of the opponents response.

COPYCAT: do whatever the player with the highest score did in the previous round.

PAVLOV: Win stay the same, lose change. (Currently the system implements two variations of this rule, one is aggressive, and counts a cooperate/ cooperate result as a loss, while the other is passive, and counts this situation as a win.)

DOWNING: do what the most players in the previous round did.

The agents interact through a number of rounds, the length of which are defined randomly by an external clock and the range of which can be set by the user, allowing some control over the density of the musical material produced.

In each round of interactions, the agents encounter each other in randomly selected pairs. There is an equal chance that an agent will encounter the same opponent in the subsequent round as a new opponent. One important difference between the model used here and Axelrod's tournament is that all agents interact in pairs simultaneously, and can 'see' all of the other interactions that are occurring, not only the interaction with their current opponent. This means that if an agent 'suckers' an opponent in the current round, it may receive retribution for that defection even if it encounters a new opponent.

Another difference between this model and Axelrod's is in the memory of the agents; the agents really are like goldfish in that their memory is limited to only the previous round of interactions. In this sense, the model in some ways is similar to the spatial prisoner's dilemma (Nowak et al. 1995) although space is unimportant; each agent has an equal chance of encountering any other agent in each round.

At the conclusion of each round of interactions points are awarded to each agent according to a points matrix (see figure 1) and tallied depending on the outcome of the interaction, and these results along with the responses from the other agents are stored in the agents memory to determine the outcome of the next round.

		The Prisoner's Dilemma payoff matrix	
		Note: the payoffs to the row chooser	
		Column Player	
		Co-operate	Defect
Row Player	Co-operate	R=3, R=3 Mutual Co-operation	S=0, T=5 Row: Sucker's payoff Column: Temptation
	Defect	T=5, S=0 Row: Temptation Column: Sucker's payoff	P=1, P=1 Mutual Defection

Figure 1. A typical prisoners dilemma points matrix outlining the payoffs for players according to the outcomes of their interactions. (Axelrod, 1984:8)

Musical Considerations

In keeping with the hypothesis of this research, modelling emergent ensemble dynamics, the musical engine of Ensemble is rather simple, and draws inspiration from the scenario outlined in the Wolff example. Much like the scenario in the

Prisoner's Dilemma itself, the agents only have two choices. Either play the same material as the previous round or introduce new material through a random selection. There are no means by which the agents have any control over the hierarchical structure of the music they produce in that they are not programmed with any expert knowledge that relates to musical structure at all; any sense of musical structure produced by the system is a result solely of the interactions between the agents.

In a round of interactions, each agent confronts its opponent and issues a response. The overall point scores of the pair are compared and the agent with the highest score becomes the dominant player in the interaction. A mutual co-operation will result in the dominant player's musical material from the previous round being continued in the current round by both players. If either player is 'suckered' they will take on the opponent's musical material while the defector makes a random selection of new material. Mutual defection results in a random selection for both players; see figure 2.

Ensemble: possible musical outcomes for a single interaction	
Agent 1	Agent 2
Current Score: 246	Current Score: 233
Previous sound= 23	Previous sound= 2
Response: co-operate	Response: co-operate
Result: Both agents play sound 23.	
Response: Defect	Co-operate
Result: Agent 1 randomly selects a new sound, Agent 2 plays sound 23.	
Response: Co-operate	Defect
Result: Agent 1 plays sound 2, Agent 2 randomly selects a new sound.	
Response: Defect	Defect
Result: Both agents randomly select a new sound.	

Figure 2. An outline of the possible musical outcomes of an interaction between two agents in the Ensemble environment.

The actual transformation of these musical choices into a musical surface has been handled in a number of ways, and this will be explored further through discussion of specific examples of musical works produced using 'Ensemble'. At this point the system has been used to manipulate sound at a macro level, manipulating pallets of samples in a DJ-like fashion, the micro level through controlling the combination of sine waves and at the mezzo level through generating streams of midi events to generate musical scores.

Discussion

The interpretation of the actions of the agents in musical terms has obvious implications for the musical surface produced, but also influences the implementation of the Prisoner's Dilemma model itself. For example, while the agents in the model have a short memory, they are able to 'see' all the other interactions that occur in each round and equally 'hear' what the other agents are playing. This global information allows the agents to survey the dynamic landscape as it unfolds, while each interaction allows them to focus their attention on particular agents, much like a musician listening to the whole group of performers, and within that moving their focus around between individual performers.

This has implications for the way in which the various strategies are realised, for example, a typical DOWNING strategy surveys the history of its previous interactions with an opponent to determine the probability of its opponent's response in the current round. From this information it in turn determines its response (Axelrod, 1984:34). In the 'Ensemble' model, DOWNING is tweaked somewhat as it can see all the responses of the other agents from the previous round, but its memory is restricted to only that round. Hence it looks at whether the majority of agents are co-operating or defecting, and uses this information to determine its response. In this sense it is able to assess whether it is experiencing a co-operative or hostile environment, and respond accordingly as opposed to assessing its direct opponent. If the majority of agents are behaving in a hostile fashion, it will presume that its opponent is likely to be hostile regardless of its behaviour in the previous round.

While the agents are able to consider global attributes in their decision-making, the selection and transfer of musical material between agents is solely governed by the outcome of the local interaction of each pair of agents in a particular round.

This creates the musical dilemma for the agent, in that in order for their material to be played by their opponent and indeed possibly be passed on to other agents in future rounds, co-operation is required. At the same time, if their musical material is to persist, they must either have a higher score than their opponent so that they are the dominant player in the encounter, or 'sucker' them which either way means they must risk defection at some point. Of course, because their action can be seen by the other agents, they are always at the risk of swift retribution in the next round, which will obviously cause their musical material to be ignored through the introduc-

tion of new musical material by their opponent. Equally, through co-operation, they also run the risk of being 'suckered' causing them to be forced to take on their opponent's musical material regardless of their score.

This behaviour has strong links to the dilemma faced by performers in an indeterminate setting, as outlined by Blackwell in his description of the way structure emerges in non-idiomatic improvisation;

Structure in improvised music derives from spontaneous (ie. unplanned) changes in musical direction... Each musician is faced with a constant dilemma: new expressive initiatives may be followed, or they may be ignored. The tension between expressive attraction and avoidance leads to the sudden, apparently orchestrated, changes that characterise this music. (Blackwell & Bentley, 2002:1462).

Through modelling this type of scenario, it is obvious that a 'best' strategy could be searched for, where by a particular agent is able to dominate the others and basically impose its musical will in a kind of musical 'Core Wars' (Coventry & Highfield, 1995:240). At this point this is not the aim of this research, not to say such research would not produce interesting results, just that more emphasis has been placed on exploring the dynamics that arise from the interactions between the agents rather than searching for a 'best' strategy. Indeed, the main question posed by the model is what combination of strategies can lead to the emergence of interesting global musical structures.

Equally, another fascinating question is raised with regards to the deeper more fundamental understanding of free-improvisation, where one could argue that groups of improvising performers are indeed caught up in a form of prisoner's dilemma. While the ENSEMBLE model suggests a rather strong analogy supporting this argument, certainly further research needs to be conducted to verify such a statement.

System Design and Implementation

ENSEMBLE is implemented in Cycling 74's MaxMSP programming environment (Cycling 74, 1990-2004). Essentially the system is modular in nature, enabling easy modification to suit different musical tasks in different compositions, and to accommodate future development. At this time, the system consists of three main com-

ponents; a time module that triggers each round of interactions, a processing module that incorporates the Prisoner's Dilemma model, and an output module that handles both the audio and visual output. In the future, an input module will also be added to allow live performers to interact with the system in real-time.

Time

Time within the system is discreet, and controlled by a module that can be set for either real-time performance situations, or generative duties. The module also controls the pairings of the agents in each round of interactions. In real-time mode, the module varies the length of each round randomly according to limits set by the user. These are generally set above four seconds, and are aimed at the future interactive type of situation; human performers generally are unable to deal with decision-making processes shorter than two seconds so a four second round allows a buffer for this threshold. In generative mode, the rounds are triggered in a metronomic fashion at 500 millisecond intervals and the length of each round is simply delivered to the agents as an integer, allowing fast score generation in non-real time.

Processing

The processing module consists of the ensemble of agents. Their environment is solely made up from their interactions and essentially the only way the agents can communicate with each other is through the sequence of their own behaviour. Each agent consists of a behavioural engine, a memory, 'eyes' and 'ears'.

The behavioural engine contains all the agents cognitive abilities, effectively incorporating the strategies for interaction, outcome assessment, comparative functions to decipher musical and visual output and generative capabilities allowing the agent to produce new musical and visual output as needed. The agents have the facilities to implement all of the behavioural strategies in the system, but are preset to a particular strategy prior to the first round of interactions. This has been a deliberate limitation of the system to allow for clarity in the agents interactions.

The musical and visual capabilities of the agents have been refined through a number of versions which have been used in the construction of different compositions. While the agents are only able to produce a single sound in each round of interactions, the level of control they have over this sound has changed considerably with subsequent versions.

In initial experiments agents simply picked a sound from a palette of 32 sampled loops and these were played for the entire length of each round. Visual output was handled in a similar fashion with agents choosing from a palette of images, the choices of the eight agents were subsequently combined into a composite image.

Current midi versions of the system allow the agents to assess the length of each round and control the entry point and duration of their sound, as well as pitch and velocity. At this point, each agent is allocated a midi channel, so that their individual musical choices are easily assessed, but future versions will also allow for agent control over timbre. Synthesis based versions of the system will also allow agent control of attack, and envelope.

The memory of each agent, while being restricted to the previous round of interactions, stores a substantial amount of information about its experience in the previous round and equally that of its opponents. The memory receives information from the behavioural engine, as well as the agent's 'eyes' and 'ears'. Effectively, the 'eyes' and 'ears' give the agent feedback on the musical and visual output of the entire system occurring in the performance space, while the behavioural engine monitors and assesses the interactions between agents. At the end of each round, the agent's memory is purged, and updated with the feedback from the current round to inform the agent's behaviour in the next round.

Output/ Case Studies

The output module of the system handles both performance duties in real-time and score generation in non-real time. Several modules have been developed to handle output in order to meet the requirements for different compositions.

CONflict, a multi-media real-time work, for example, utilises an output module that allows the agents to select from a range of justly tuned sine waves and visuals. This particular work utilises eight speakers in a cube arrangement, to spatialise the agent's sounds in the performance space. Simultaneously, the visual selections of the agents are combined into a composite image that is also projected for the audience.

Several other experiments have been conducted with a non-real time output module that generates musical scores. This midi module is based around the Max 'detonate' object allowing the generation of midi files that in turn are directly converted into musical scores to be performed by live performers; see figure 3. While the main goals of ENSEMBLE revolve around real-time generation of material, and in the fu-

ture, interaction, the generation of musical scores has allowed snapshots of the musical behaviours produced by ENSEMBLE to be explored, pinning down the often transitory nature of the system for closer analysis.



Figure 3. Raw output from the ENSEMBLE score generator. This output became the second movement of 'Surroundings' premiered on 17th April, 2005.

The score generator has successfully been used to generate the musical works *Surroundings* for Piano Trio, which was premiered with little revision, and the work *Irene's Myth* where ENSEMBLE took on an 'algorithmic composer's assistant' type of role. For this work ENSEMBLE output was organised into slightly more traditional hierarchical forms. These works have highlighted the potential of ENSEMBLE in organising sound into pleasing temporal structures that are autonomously governed by the agents interactions within the system.

Conclusions/ Future Work

Evaluation of generative creative systems like ENSEMBLE is generally difficult, particularly when aiming for innovation. A discussion of this dilemma and equally what makes 'good' music can be found in Todd and Werner, 1999. Equally a notable omission from this paper has been a discussion of aesthetics. The only comment to be made on this issue is that at this time the main aesthetic consideration has been to be as transparent to the process as possible, and attempt to let the model 'speak for itself'.

As it stands, as an autonomous music generator, ENSEMBLE has shown that a bottom up sense of musical form can be produced through the interactions of the agents within the system. Through defection, the agents are able to introduce novelty to the musical surface, while co-operation allows for the consolidation and persistence of previously generated musical material. This can work at the micro, mezzo and macro levels of music, and creates a sense of musical movement without the need for any kind of traditional hierarchical musical form.

Experiments with ENSEMBLE have shown that the Iterated Prisoner's Dilemma appears to be a viable model for the simulation of the types of 'emergent' ensemble dynamics apparent in indeterminate performance situations. It is hoped that through the implementation of real-time interaction between ENSEMBLE and live performers that truly collaborative music making will result, with the system able to be both responsive to the performers input, and suggestive as it has the ability to both co-operate with and defect against the performers musical initiatives. It is also hoped that through further exploration of this simulation that a deeper, more fundamental understanding of the rules which govern performers in more open 'free-improvisation' performance situations may be revealed in order to improve the implementation of interactive music environments in the future.

Future work will also revolve around further refinement of the Prisoner's Dilemma model used in the system, particularly moving beyond the deterministic nature of the current model, towards systems where agents are able to change strategies to try to improve their performance, and equally, evolve new strategies.

References

- Axelrod, R. 1984. *The Evolution of Co-operation*. Basic Books, New York.
- Axelrod, R. 1997. *The Complexity of Co-operation: Agent-Based Models of Competition and Collaboration*. Princeton University Press, New Jersey.
- Blackwell, T.M. & Bentley, P.J. 2002. "Improved music with Swarms", *Proceedings of the Congress on Evolutionary Computation*. Piscataway, NJ: 1462-1468.
- Bilotta, E. & Pantano, P. 2001. "Artificial Life Music Tells of Complexity", *Proceedings of the workshop on artificial life models for musical applications (ALMMA 2001)*. Editorial Bios, Cosenza, Italy: 17-28.
- Cage, J. 1961. *Silence: Lectures and Writings by John Cage*. Wesleyan University Press, Middletown, Connecticut.
- Campbell, R. & Sowden, L. 1985. *Paradoxes of Rationality and Cooperation*. The University of British Columbia Press, Vancouver. (1985).
- Corbett, J. 1994. *Extended Play: Sounding off from John Cage to Dr. Funkenstein*. Duke University Press, London: 181-191.
- Coveney, P. & Highfield, R. 1995. *Frontiers of Complexity: The Search for Order in a Chaotic World*. Fawcett Columbine, New York.
- Cycling 74. 1990-2004. *MaxMSP*. software. see www.cycling74.com, last viewed 29/5/2005.
- Harrald, L.A. 2003. "Artificial Life: a model for musical innovation?", *Proceedings of the Australian Conference on Artificial Life (ACAL 2003)*. UNSW, Canberra: 128-141.
- Jacob, B.L. 1995 "Composing with Genetic Algorithms", *Proceedings of the 1995 International Computer Music Conference*. ICMA, San Francisco: 452-455.
- Johnson, C.G. 2003. "Towards a prehistory of evolutionary and adaptive computation in music", *Lecture Notes in Computer Science*. Vol. 2611. Springer-Verlag, Berlin: 522-509.
- Miranda, M.R. 2001. *Composing Music with Computers*. Focal Press, Oxford.
- Miranda, M.R. 2002. "Emergent Sound Repertoires in Virtual Societies", in *Computer Music Journal*, 26: 2, Summer 2002. MIT Press: 77-90.
- Morgan, R.P. 1991. *Twentieth-Century Music: a history of musical style in modern Europe and America*. Norton & Company, New York.
- Nowak, M.A. May, R.M. & Sigmund, K. 1995. "The Arithmetics of Mutual Help", *Scientific American*. 272:50-55.
- Reynolds, C.W. 1987. "Flocks, Herds, and Schools: A Distributed Behavioural Model", *Proceedings of SIGGRAPH '87*. In : Computer Graphics, 21 (4): 25-34.
- Todd, P.M. & Werner, G.M. 1999. "Frankensteinian Methods for Music Composition", In: Griffith, N., and Todd, P.M. (eds.): *Musical Networks*. MIT/ Bradford Books, Cambridge: 313-339.
- Waschka II, R. 2001. "Theories of Evolutionary Algorithms and a 'New Simplicity' Opera: Making Sappho's Breath", *Proceedings of the workshop on artificial life models for musical applications (ALMMA 2001)*. Editorial Bios, Cosenza, Italy: 79-86.
- Wolff, C. 1971. *Burdocks for one or more orchestras*. Musical Score. Edition Peters, New York.
- Zorn, J. 2002. *COBRA: John Zorn's Game Pieces (1984) Volume 2*. Compact Disc. Tzadik, New York.