Interactive Composing With Autonomous Agents

Michael Spicer

School of Digital Media and Infocomm Technology, Singapore Polytechnic, Singapore

mspicer@sp.edu.sg

Abstract

This paper describes the approach the author has adopted for creating real time systems whose musical output is created by the interactions of a human performer and an ensemble of software "improvisers". This approach has been used with notated compositions, free improvisation performances and for creating installations. The music produced typically consists of several distinct textural layers, where all the sounds produced are transformations of the sound made by the human performer. This type of system can be thought of as an "extended" instrument, where the performer effectively "plays" the ensemble. The design of the virtual performers is based on the concept of autonomous agents, a popular approach borrowed from Computer Science. One benefit of this approach is that it provides a unified way to encapsulate numerous algorithmic compositional techniques within one system. This system makes use of two broad categories of agent: performers and controllers. Performer agents transform the live sound in various ways, while controller agents' work at a higher structural level. They specify goal states and determine which agents are currently heard. Each performer agent has a way of transforming the audio input, and has its own internal strategies for determining what it does. The complexity of the performer agents note choice strategies ranges from simple harmony generators, to fairly complex algorithmic composition systems. This approach creates an environment where the human performer feels a sense of control over the entire ensemble while achieving a sense of collaboration with the system.

Keywords: Autonomous Agents, Interactive Composing, Extended Instrument

1. Introduction

This paper describes an approach to creating interactive performance/composing systems which has been used by the author for both improvisation and for fixed compositions. The music produced is a result of the interactions of a human performer and an ensemble of virtual "performers". The musical output typically consists of several distict musical layers, but all the sounds heard by the audience are derived from the sound made by the human performer. The approach can be thought of as a way creating an "extended" instrument, which is responsive enough that the performer can perceive how it reacts to performance nuances and note choices, and can effectively "play" the ensemble. The current systems are implemented in the programming environment PD, using a design based on the idea of autonomous software agents. The agent design was adpoted as it provides a neat way to encapulate various compositional techniques in a unified way.

2. Live Performance

The agent design described in this paper has evolved over a ten-year period, having successfully been used as an interactive control system for synthesizers in improvisatory performance contexts, and was interacted with the computer GUI (Spicer, Tan, Tan, 2003). In the few last years, this has changed dramatically. The agents in the new particular system listen to a live input signal, usually a flute or vocal performance via a microphone, and then make some sort of analysis that is used to guide the agents' behavior. The agent then uses the live signal as the raw material for creating its contribution to the musical texture. This creates a very tight feedback loop between the human performer and the agent ensemble. This aspect of all sonic material being derived from the live signal allows the player to leverage his/her instrumental technique to have significant impact on the ensemble sound. Modern extended performance techniques also effectively lend themselves for use with these systems.

3. Agents

Central to the system is the concept of autonomous agents. The idea of autonomous agents has been popular amongst the computer science community for many years (Russel & Norvig, 2003), and has been applied to solve many types of problems, including music. An agent can be considered to be as an entity that can perceive its operating environment in some way (in this case "listening"), and has some ability to change/interact in that environment ("performing"). It has some internal mechanism/process ("thinking") for deciding the exactly how it will interact with the environment. This "thinking" part of the agent, referred to as the "agent function", can be implemented in a wide variety of ways, and it is this flexibility that makes the agent model such a useful abstraction. It allows a neat way of encapsulating different compositional algorithms/ decision-making strategies in one, unifying design.

There are two broad categories of agent: performers and controllers. Performer agents transform the sound of the flute in various ways, while controller agents either control which agents are heard by the audience, or specify a set of goal states that determine the behavior of the performer agents.

4. Performer Agents

In this system, performer agents continully listen to the live audio signal produced by the human performer. The instantanious pitch and amplitude data, as well as cumulative statistical data, are used in the agents' decision making process which deterimines how it alters the live input signal.

There are two categories of performer agents. The simplest type are those which apply "traditional" sound effects, such as echo, chorus, flanging, spectral delay, ring modulation etc to the live signal. The agent alters the controls of these DSP effects in response to the input signal. These agents are mainly used to shape the overall ambience the ensemble output.

The bulk of the agents perform more of a textural/structural role. They take the live signal and transform it into distict musical lines, which may have distinctly different timbres, and perform distictly differenent musical functions.

This group of agents all shares the same basic operating principle, which is to:

- Determine the current pitch of the live signal.
- Calculate the next pitch the agent will play, using the strategy specific to that agent.
- Calculate the interval between the input pitch and the calculated pitch.
- Transform the live signal so that it has the required pitch using the signal processing technique built into the agent.
- Calculate the length of the new note.
- Play the note for the required duration.

There are a variety of algortihms employed in the various performer agents that determine exactly how the agent will behave. So far, the algorithms implimented are those that interest the author, and include various stochasitc algorithms, simple fractal generators, tracking and evasion and trajectory following. (Dodge & Jerse, 1997)

About twenty different types of performer agents have been implemented so far, and the signal processing techniques employed range from granular synthisis based pitch shifters to Karplus-Strong plucked string algorithm based synthesizers. The result of this is that the live signal is transformed into drones, parralel harmonies, counter melodies, ambience effects, and pointalistic clouds. As the audio is a transformation of the live signal, the human performer has a sense of control over the entire ensemble sound.

An important aspect of the performer agents is their ability to modify their internal states, so as to meet externally specified targets. The two targets used in this system, so far, are average pitch, and average duration. The target following capability is achieved by making use of a 96-element array for storing the outputs of the compositional algorithm. One array contains data representing the note duration, and another one contains data for the pitch. With the help of simple mapping functions, the agents use the contents of these two arrays to determine their musical output. In order to enable an agent to readjust its internal state so as to meet any particular supplied target, a simple form of gradient descent learning is used. Each agent periodically calculates the current average value of each array, and compares these to a target average pitch and a target average duration, producing an error measurement for pitch and duration. These two error measurements can then be used to bias the contents of each array slightly, so as to reduce the error. Periodically repeating this process eventually enables the agent to exhibit the target behaviour.

5. Controller Agents

In addition to the performer agents, there are two higher-level agents that affect the overall musical output. One agent provides target parameters that are used by the agent performers to individually adjust their internal state. The other agent acts as a "mixing engineer", and determines which performer agents are heard by the audience. Both of these agents are finite state machines that use analysis data of the live signal to shape the musical output. To set up these agents for each composition, the composer needs to:

- Determine 8 different target states for the system. These will be vectors consisting of the target volume levels for each agent or target average pitches and average durations for a performer agent. Setting these states is a very important compositional decision, as it determines the behavioural extremes of the piece.
- Assign each state to the 3D coordinates of the different corners of a cube, centred at the origin, with a side length of 2. Eg. (1,1,1)(1,1, -1) (1, -1,1)(1, -1, -1) (-1,1,1)(-1,1, -1) (-1, -1,1)(-1, -1, -1).

In operation, these agents:

- Periodically derive three values between -1 and 1 from the recently heard live signal using some mapping function applied to the pitch and duration data. The idea is to create a 3D coordinate from what the human performer has played. The exact way this is mapped is another major compositional decision that is made for each piece. Normally, the basic analysis of the input signal is done by fiddle~ (or sigmund~) and this data is accumulated for a particular time interval, and then some sort of statistical analysis is performed. For example, in one piece, the mean input pitch, the duration of the last phrase, and its range are each mapped onto this -1 to 1 range, to produce the 3D coordinate.
- Calculate the Euclidian distance of this 3D coordinate from each of the vertices of the cube.
- Set the system to the state that has the smallest Euclidian distance from the 3D coordinate.

The fact that these are deterministic processes means that there is an element of predictability in the behaviour of the system. This means that similar musical input gestures will tend to produce similar musical ensemble output. This enables the human performer to learn how to shape the response of the agent ensemble.

6. Applications

This system was developed specificly to realise a collection of notated compositions for flute and computer, called Pandan Musings. All the flute parts have been precomposed, in a consiously early 20th century style composers such as Poulenc and Debussy, and each piece uses its own optomised agent system. Because each piece is played with a fixed score, the deterministic high level agents enable the fine tuning of the various targets so as to achive specific musical results that are approximatly repeatable, or at least recognisable. This creates a consistancy in the ensemble part, whilst still leaving a lot of room for spontanious expression, and serindipitious moments to occur.

Another application of this system is an installation called "Free Voice", for "VOICES – A Festival of Song" in Singapore, December 2013. This consists of a two instances of the system, with two microphones, running in parralel. One system consists entirely of "sound FX" performer agents, whilst the other is entirly creating musical lines, using Karplus-Strong synthesizers that use short bursts of the live signal for exitation. Traces of intelligibility are heard in the resulting musical output.

In 2012, the author collaborated with the composer Robert Casteels, to create some electronic interludes in his piece "Venus Transit Symphony". These interludes made use of a variation of the agent based system to process the sounds produces by various celestial objects, as detected by radio telescocopes. In this case, the performer agents imposed various pitch contours on to the sounds from the Sun, Saturn and four quasars, while human performers adjusted the various targets in real time, via a MIDI controller.

While various versions of this system have been developed for specific compositions, the system is often used in improvisation contexts. The author usually performs with it using live input from instruments such as the flute, saxophone, Rhodes electric piano or analogue synthesizer. Minor adjustnments to the analysis paramaters are made to optimize the systems perfomance with each particular instrument.

7. Conclusion

The autonomous agent based system design has been a useful tool for the author to efficiently create effective sound transformation systems that have been used in real-world performance situations. The fact that this approach produces something that is a blend of extended instrument and interactive composing system is a very attractive aspect of these systems, as it means that the human performer can leverage his/her standard technique to react to and control the musical result, enabling each performance to be a unique experience.

References

Dodge, C. and Jerse, T.A. (1997). *Computer Music Synthesis, Composition, and Performance.* Schirmer Books, New York NY, U.S.A.

Russel, S, and Norvig, P(2003). *Artificial Intelligence* Prentice Hall, New Jersey.

Rowe, R. (1993)*Interactive Music Systems. Machine Listening and Composing.* The MIT Press, Cambridge Massachusetts.

Spicer, M.J., Tan, B.T.G. and Tan, C.L (2003)"A Learning Agent Based Interactive Performance System." In *Proceedings of the International Computer Music Conference*. San Francisco: International Computer Music Association.