PheroMusic: Navigating a Musical Space for Active Music Experiences

Kristian Nymoen, Arjun Chandra, Kyrre Glette, Jim Torresen
University of Oslo, Dept. of Informatics
{krisny, chandra, kyrrehg, jimtoer}@ifi.uio.no

Arve Voldsund
Kodebyraaet
arve@kodebyraaet.no

Alexander Refsum Jensenius
University of Oslo
Department of Musicology
a.r.jensenius@imv.uio.no

ABSTRACT

We consider the issue of how a flexible musical space can be manipulated by users of an active music system. The musical space is navigated within by selecting transitions between different sections of the space. We take inspiration from pheromone trails in ant colonies to propose and investigate an approach that allows an artificial agent to navigate such musical spaces in accordance with the preferences of the user, and a set of boundaries specified by the designer of the musical space.

1. INTRODUCTION

This paper presentes a musical system that invites *non-musicians*, i.e., people not knowing how to play a musical instrument nor knowing much about music theory, to participate actively in a musical experience. Such a concept of *active music listening* [1], is based on the idea that a user can engage with music somewhere along an imaginary axis between an instrumentalist and a passive listener. Thus an an active music experience enables a wider range of control possiblilities than the *play*, *pause*, *skip song* and *volume* controls of regular audio players, yet not requiring the instrumental skills of a musician.

In our setup, an *active music system* is built around a *musical space* that may be navigated within by a user. The musical space is pre-created by a composer/designer, and include various musical sections, that will be referred to as *soundscapes*. "Composing" a musical space for an active music system involves designing each individual soundscape, as well as a set of rules for transitions between these; including their durations and order.

During exploration of the system, the user can influence the order of soundscapes or sections, or influence some general parameters like the dynamics, mood, etc. We want to enable users to explore the space not only with immediate control actions, but with actions that influence how the piece of music evolves in the future.

Our active music system is realised by implementing an artificial *agent* utilising the *pheromone trail* mechanism from ant colony optimisation. Here the agent follows a pheromone trail between musical sections, which is initially specified by the designer of the space. The human user may then select a path for the agent, influencing the transitions taken by the agent in future repetitions. The following sections will present the design of the system, its implementation, and discuss its use.

2. OUR SYSTEM

Our active music system provides a limited set of high-level control actions to a user, mapped to a set of musical parameters. These parameters and their implementation are defined by the designer of the space, so that users are allowed to (1) navigate between different soundscapes, and (2) control the intensity of the music.

The designer of a musical space defines the *boundaries* of the space. This implies that the musical space must be flexible in order to give a certain amount of freedom to the user, to influence the music in accordance with their tastes or preferences. However, a space should not be so flexible that the interaction feels like composing or performing, i.e., the user should get the notion that he/she is *involved* with the music, not creating it from scratch.

As the music plays, an artificial agent navigates between the soundscapes according to the designer's specifications, or, new soundscapes are selected by the human user. We consider the selections made by the user to reflect the user's preferences within the boundaries laid out by the designer. To cater for these preferences, the system utilises a pheromone mechanism, allowing the agent to remember the path chosen by the user when making a transition between soundscapes in the future.

2.1 Boundaries

For the system presented in this paper, the boundaries are specified as:

- a set of soundscapes of the constituent parts of a musical space
- heuristics for the features of a set of musical elements in each soundscape, defining rhythmic, timbral, and harmonic properties
- various possible transitions between soundscapes, allowing users to explore the music based on such transitions

- a mapping of the control parameter *intensity* to a set of musical parameters, providing the user with one control dimension within each soundscape

In our system, each individual soundscape with its musical elements and heuristics has been implemented in Max. The pheromone-inspired approach for both automated and user-defined transitioning between soundscapes, including the boundaries for these transitions has been implemented in Processing. A video example showing the system in action is available online. [1] The various components of our system are described more in detail below.

2.2 Soundscapes

Soundscapes are the main parts of the musical space. In principle, a soundscape could consist of any type of electronic sound generation, not limited to any particular synthesis technique or genre. The soundscapes in our prototype make up one particular musical space, and is based on subtractive synthesis using white noise, envelopes, and resonant filters. A complete specification of each soundscape is beyond the scope of this paper, however, we will provide an overview of the sound synthesis patch.

The demo example shown in the included video consists of four main components: percussion, bass drone, melody, and DAFX, all driven by one **phasor**∼ object in Max. At various subdivisions of the phasor's output, sounds are triggered by multiplying an envelope with white noise passed through resonant filters. Each soundscape has a defined set of parameters, such as:

- filter gains, bandwidths, and centre frequencies
- envelope attack and decay slopes
- thresholds for triggering percussive/melodic tones
- a range of available tones to be played
- harmonic progressions
- tempo
- number of subdivisions in the triggering mechanism

Some parameters have fixed values in certain soundscapes, while a range of possible values, ultimately controlled by the user through the *intensity* parameter, is defined in other soundscapes. As non-musicians are the target group of our system, intensity is the only control parameter made available to the user. The mapping is one-to-many, or in other words, a single control dimension is used to manipulate a large range of synthesiser parameters in each soundscape, as defined by the designer of the musical space.

2.3 Transition Table

Table 1 depicts a *transition table*, representing the strength values of the link between any two soundscapes. The musical designer defines how likely it is for a transition to occur from one soundscape to another. The rows refer to soundscapes where the transition starts, and the columns to soundscapes towards which the transition happens. In other words, the columns represent the next soundscape in sequence from the soundscape in any row. The idea is to allow the designer to define a piece of music that is flexible even without user influence, having the artificial agent choose between different soundscapes. For instance, if the soundscapes refer to sections of a pop song, the designer may specify that transitions should be from a verse to prechorus and from prechorus to chorus, and that a transition directly from a verse to a chorus is less likely to occur.

Table 1. Example of a transition table for a musical space with five soundscapes

	1	**2**	**3**	**4**	**5**
1	3	1	1	1	3
2	1	2	1	2	0
3	1	1	1	1	0
4	1	0	0	5	0
5	1	1	0	1	1

Together with other specifics mentioned in Sections 2.1 and 2.2, the transition table with specified values completes our specification of the boundaries defined for the musical space. This table is used to initialise a *transition graph*, which we cover next.

2.4 Transition Graph

Figure 1 shows an example of a soundscape *transition graph*, in this case with seven soundscapes. It is a directed graph with loops, where nodes in the graph represent soundscapes. Each node in the graph connects with every other node and itself, via a link. The links in the graph represent transitions from one soundscape to the other. These links have an associated strength value, the thickness indicating this strength. Graphs are initialised using transition tables, as mentioned in Section 2.3. As such, navigating the composition amounts to navigating this graph in accordance with the designer's predefined boundaries.

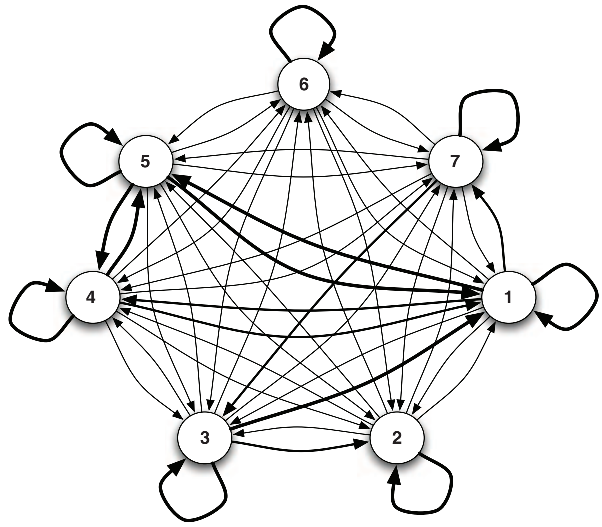

Figure 1. Nodes and links on a transition graph in a musical space with seven soundscapes

[1] http://vimeo.com/89418486

2.5 Navigating the Graph

The flexibility in the musical space that is laid out in this paper requires a mechanism for making transitions between the different soundscapes. While the user may be the one deciding the order of soundscapes according to his/her preferences, the active music piece should still continue when the user does not control the system, and transitions should be made according to the user's preferences. To achieve this, we take inspiration from the pheromone mechanism in ant colony optimisation, and let an artifical agent (ant) follow a pheromone trail between soundscapes.

To the best of our knowledge, ant colony inspired musical systems were first introduced in 2004 [2]. So far, the two main focuses in previous work have been the use of artificial ant colonies for automated music generation [3] and as enablers for automatic accompaniment for people with disabilities [4, 5]. The user interaction essentially involved assisting the ant colony to prevent the music from stagnating [4], while our approach in this paper invites the user to actively influence the coming transitions between musical sections. There has also been a lot of research in the past with regard to automated music generation, which makes use of tables similar to the one presented in Section 2.3, the case in point being Markov chains [6]. Although we use the notion of pheromone trails from ant colonies in our system, we focus on how the agent may cooperate with, and cater for the preferences of, the user within the pre-designed musical space, as opposed to automation of music generation or as enabling accompaniment.

The transition graph may be navigated by an agent based on the strength of the links between the nodes, in accordance with some *behavioural rule*. Given a soundscape and all the links that lead to the remaining soundscapes on the graph from this soundscape, we consider transforming the strength values of the links into probabilities of these links getting selected, thus the probability of selecting one of the soundscapes to be the next in the navigation sequence. The probability p_{ij} of choosing a soundscape j after i, is given by,

$$p_{ij} = \frac{\tau_{ij}}{\sum_{k=1}^{N} \tau_{ik}}$$

where τ_{ik} denotes the strength of the link between soundscapes i and k, and N is the total number of soundscapes the transition graph is composed of.

Thus, the agent navigates the graph based on this behavioural rule. As such, the stronger the link between two soundscapes, the more likely it is that this link will be chosen and traversed by the agent. The traversal of this link takes the agent to the next soundscape. The graph, the strength values for links, which are initialised using the transition table, and the traversal of the graph between soundscapes by the agent using its behavioural rule, provides a way for the agent to explore the active music space.

The agent should exercise the preferences of the user whilst navigating the graph. This would then warrant for the agent to memorise the choices that the user makes as it navigates. This memory is what we want the system to maintain. This is to have the system remember preferred navigation paths within the musical space, for the agent to not have to explore the space again from scratch, if it intends to revisit the preferred paths.

We want the system to allow for freedom of musical expression such that the preferences of the user are autonomously elicited and catered for, as compared with what would be available with a graph with fixed strength values laid out by a desinger. We provide for this freedom by making it possible for the system to have the strength values over the links on the graph be updated in real time.

We take inspiration from the notion behind pheromone trails in ant colonies to design the link update mechanism. In nature, an ant, when foraging for food, leaves behind some pheromones along the path, which acts as a form of memory that may be sensed by the same or other ants to evaluate the possibility of going along this path. The more a path is taken, the more pheromone get deposited along the path. The stronger the presence of pheromone along a path, the more likely it is that the said path gets taken. Over time, a fraction of the pheromone present on the path necessarily evaporates as well. We utilise this basic principle within our system for user preference elicitation.

Every time a user makes a selection in the navigation of a sequence of soundscapes, a pheromone trail gets updated in the form of strengthening of the links from one soundscape to the next. Moreover, each time a transition is made, a fraction of the pheromone on every link on the graph necessarily evaporates. This link strength update mechanism can be written as

$$\tau_{ij} = \begin{cases} (1-\rho)\tau_{ij} & \text{if link not selected by user} \\ (1-\rho)\tau_{ij} + \Delta & \text{if link selected by user} \end{cases}$$

Thus, the link that is traversed has a positive value $\Delta \in \mathbb{R}$ added to it. And, the links between the nodes necessarily have their strength decreased over time at a rate given by $\rho \in [0.0, 1.0]$, known as the *evaporation rate*.

With such a link strength update mechanism in place, if the user chooses a link, which initially was deemed a weak transition by the designer, this link becomes stronger by a value Δ. Over time, if this link is chosen over and over again, the strength of this link will increase. We assume such an increase in strength to reflect preference on the part of the user, for that particular transition. This allows for the sequence of soundscapes that gets navigated to have the possibility to easily be reproduced, if indeed preferred by the user. For transitions that might not get selected by the user much, the strength of the link associated with such transitions will decrease over time. The evaporation of link strengths thus allows for the system helping to forget unfavourable transitions.

We further consider using the transition table specified by the desinger (Section 2.3) as determining one of the boundaries for the system. We do so by having the link strengths in the graph thresholded to these transition table values. As such, the link strengths are not allowed to evaporate below the values in the transition table used for initialisation of the graphs.

2.6 User Interaction

Users of our system are encouraged to take part in controlling the musical output. Two main control dimensions have been defined: (1) soundscape selection and (2) intensity adjustment. A mobile device is used as controller, allowing soundscape selection by tapping on rectangular buttons on the screen (see Figure 2), and intensity adjustment through tiliting of the device. The interface has been implemented using Mobmuplat[2] and Puredata (vanilla), and passes control messages to the transition graph running in Processing on a laptop.

When selecting a soundscape, the user may tap a button to tell the system that upon the next scheduled soundscape transition, the transition should be made to that particular soundscape. Alternatively, the user may double tap on a button, to make the transition immediately.

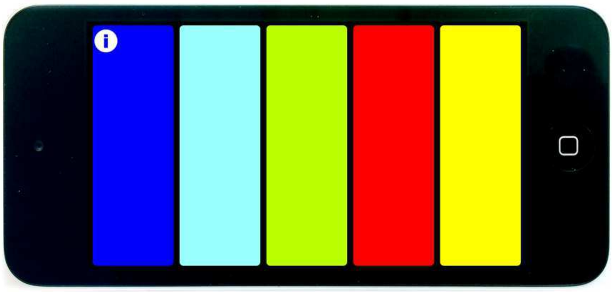

Figure 2. An Ipod with the control interface

3. DISCUSSION

While musical instruments and personal media players have been researched for centuries and decades, respectively, the concept of active music technology is still quite unexplored. Our approach is one out of many possible paths on the road towards flexible pieces of music to be actively modified by the listener. The pheromone approach taken in this paper is not only applicable to musical segments such as described here, but could be applied to transitioning between shorter snippets of music, for instance by replacing the soundscapes by sound samples, thus using the pheromone mechanism to navigate in a sample bank. Or, the approach could be taken on even longer sections of music, like songs in a playlist.

Our approach is not bound to any particular musical genre, nor to a particular sound synthesis technology. We have chosen to use low-level components such as white noise generators and resonant filters, but other sound technologies might work even better. For instance, remixing prerecorded music by navigating between different sections, and turning on and off tracks in a multitrack recording.

The main focus of this work has been on the development of a musical system that facilitates and adapts to user selections/preferences. Detailed analysis of the efficacy of the pheromone approach to such facilitation still remains to be conducted. Furthermore, friendliness in user interaction, namely the selection of new soundscapes by tapping on the screen of a moblie device, and adjusting intensity by tilting, also require further investigations.

[2] `http://www.mobmuplat.com`

4. CONCLUSIONS AND FUTURE WORK

We aim to establish a music system that invites non-musicians to actively experience music. We considered the issue of how a flexible space of active music may be navigated by an artificial agent, and how the agent may adapt to a user's preferences. Our pheromone inspired system enables an artificial musical memory for participants in active music. This memory may potentially make it feasible for non-musicians to effectively participate in the navigation of compositions. However, further investigations as to such effectiveness remain to be conducted.

We will continue this work in the future by embedding the comlete system in a mobile device. Our plan is to make a more generic framework for both users and composers, where users may select an active music space just as easily as selecting a song from a playlist, and composers/designers are provided with an interface for creating flexible musical spaces with the synthesis engine of their choice.

Acknowledgments

This research has received funding from EU FP7 under grant agreement no. 257906 (EPiCS).

5. REFERENCES

[1] M. Goto, "Active music listening interfaces based on signal processing," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 4. IEEE, 2007, pp. 1441–1444.

[2] C. Gueret, N. Monmarche, and M. Slimane, "Ants can play music," in *Fourth International Workshop on Ant Colony Optimization and Swarm Intelligence (ANTS 2004)*, ser. Lecture Notes in Computer Science, vol. 3172. Springer-Verlag, 2004, pp. 310–317.

[3] N. Monmarche and R. Clair, "Artificial ants for artificial art," in *Artificial Ants: from collective intelligence to real life optimization and beyond*, N. Monmarche, F. Guinand, and P. Sarry, Eds. ISTE — Wiley, 2010.

[4] R. Clair, N. Monmarche, and M. Slimane, "Interactions between an artificial colony of musical ants and an impaired human composer: towards accessible generative arts," in *XI Generative Art Conference*, Milano, Italy, December 16-18 2008, pp. 413–423.

[5] ——, "An interactive ant algorithm for real-time music accompaniment," in *Évolution Artificielle, 9th International Conference on Artificial Evolution*, Strasbourg, France, October 2009.

[6] K. Jones, "Compositional applications of stochastic processes," *Computer Music Journal*, vol. 5, no. 2, pp. 45–61, 1981.