GA2014 – XVII Generative Art Conference

Philippe Kocher
Daniel Bisig

Installation: **Sheet Music**


Abstract:

Sheet Music is a sound installation whose physical setup consists of several piezoelectric film speakers distributed in space. The sonic content of the installation is generated by employing time-delayed recurrent networks as sound synthesis systems. The installation setup embodies some of the algorithmic principles that underlay the generative sound synthesis process. Each speaker corresponds to a specific point in the network and renders the network's activity at this point audible. In addition, the physical distances among the speakers are proportional to the time delay applied to the signals between those points. Hence, some of the otherwise hidden properties of the generative algorithm are made visible. Furthermore, each display of the installation is unique and site specific, as the distances between the speakers depend on the particular spatial properties of the venue. The installation represents an attempt to establish a clear correspondence between a generative system and its physical manifestation.

Topic: Music

Authors:
Philippe Kocher
Zurich University of the Arts, Switzerland – Institute for Computer Music and Sound Technology
http://www.icst.net

Daniel Bisig
Zurich University of the Arts, Switzerland – Institute for Computer Music and Sound Technology
http://www.icst.net

An early prototype of a piezoelectric speaker film (left) and one of the actual objects of the installation (right).

Contact:
philippe.kocher@zhdk.ch

Keywords: Sound installation, sound synthesis, recurrent time-delayed networks, piezoelectric speaker films

The Sound Installation "Sheet Music"

Philippe Kocher
Zurich University of the Arts
Institute for Computer Music and Sound Technology
Switzerland
http://www.icst.net
e-mail: philippe.kocher@zhdk.ch

Dr. Daniel Bisig
Zurich University of the Arts
Institute for Computer Music and Sound Technology
Switzerland
http://www.icst.net
e-mail: daniel.bisig@zhdk.ch

Abstract

Sheet Music is a sound installation whose physical setup consists of several piezo-electric film speakers distributed in space. The sonic content of the installation is generated by employing time-delayed recurrent networks as sound synthesis systems. The installation setup embodies some of the algorithmic principles that underlay the generative sound synthesis process. Each speaker corresponds to a specific point in the network and renders the network's activity at this point audible. In addition, the physical distances among the speakers are proportional to the time delay applied to the signals between those points. Hence, some of the otherwise hidden properties of the generative algorithm are made visible. Furthermore, each display of the installation is unique and site specific, as the distances between the speakers depend on the particular spatial properties of the venue. The installation represents an attempt to establish a clear correspondence between a generative system and its physical manifestation.

1. Introduction

The use of feedback and delay is commonplace in electroacoustic music. There is a large number of applications that use these techniques for audio signal processing and sound synthesis. Likewise, these mechanisms of feedback and delay also play an important role in artificial neural networks that process temporal data. Despite the shared interest, there exists very little research concerning the adoption of recurrent neural networks as mechanisms for signal processing and sound synthesis [1, 2]. We believe that the arrangement of feedback and delay units in network-like structures has a promising potential to be an attractive artistic tool for computer music. An ongoing research project at the *Institute for Computer Music and Sound Technology* of the *Zurich University of the Arts* investigates these audio feedback networks [3].

Audio feedback networks exhibit a complex and non-linear behaviour. From an artistic point of view, this is a most desirable property: it is the basis of all the unpredictability, variety and richness of the acoustic textures thus generated. On the other hand, due to their complexity the internal operations of such networks are fairly obscure. Hence, there is a concern that these networks become incomprehensible black boxes used as readymade systems but never fully understood [4]. We propose that this issue could be addressed by rendering some of the network's internal properties directly perceivable and tangible. Therefore, we try to establish an *ontology alignment* via a number of correspondences between the sound synthesis network and the physical objects of the sound installation. These correspondences include on the one hand the representation of each network node as an individual sounding object and on the other hand the matching of the distances between the objects and the duration of the time-delays of the connections between the nodes.

Sheet Music is an attempt to prototypically demonstrate the application of an audio feedback network in a sound installation context. It is also an attempt to explore the effectiveness of ontological correspondences. The structure of the sound synthesis algorithms as well as the specific, custom-made hardware are described in the following sections.

2. Hardware

2.1 Piezoelectric Film Speakers

The sound installation is formed by eight sound-emitting objects distributed in a room. Each object consists of a piezoelectric film speaker mounted on the foot part of a music stand (Fig. 1). The operation of the piezoelectric film speakers is based on the piezoelectric effect, which refers to the accumulation of electrical charge in response to an applied mechanical force. Materials that show this piezoelectric effect also exhibit the reverse piezoelectric effect, i.e. a mechanical deformation resulting from an applied electrical field. The latter allows piezoelectric films to be used as sound transducers. Piezoelectric film speakers are characterized by a highly nonlinear frequency response and a variable sound emission directionality, both of which depend on the shape and curvature of the film (Fig. 2). These peculiar acoustic traits render the speakers interesting as musical artefacts that strongly colour the sonic emission in relation to their shape.

Figure 1. Piezoelectric speaker films mounted on the foot part of a music stand form the sounding objects of the installation.

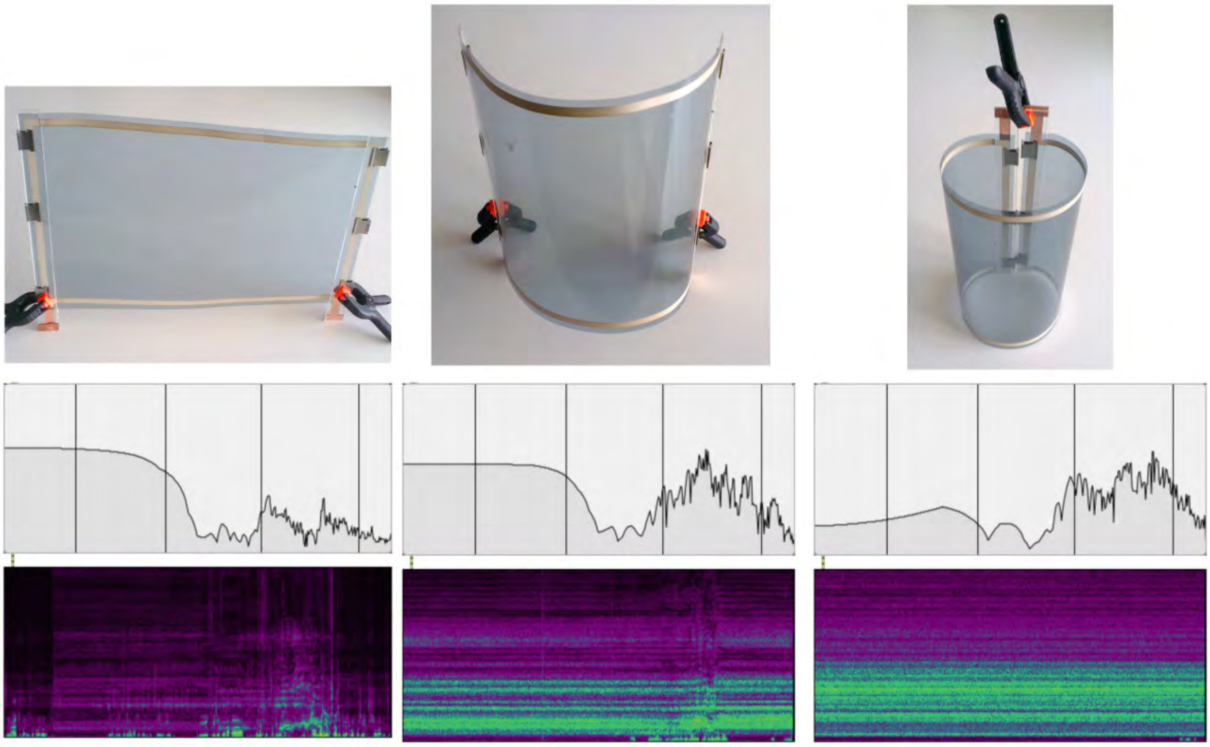

Figure 2. Curvature dependency of the frequency response of piezoelectric speaker films. The figure depicts the speaker films (top row) together with the corresponding amplitude spectra (middle row) and spectrograms (bottom row) of the speakers' audio emissions. The input audio signal is white noise.

Apart from its function as sound emitting device, the piezoelectric film speakers are also an important visual component of the sound installation. Being thin, transparent and weightless, the piezoelectric films possess several material properties unlike those we normally associate with loudspeakers. Therefore, they do not iterate the concept of loudspeakers and sound generation being two detached processes. Rather, the film speakers are readily accepted as sounding objects in their own right, which consequently reduces the gap between sound reproduction and the underlying generative algorithm.

2.2 Audio Amplifier Electronics

Conventional audio amplifiers are not suitable for our needs, as the piezoelectric films require a high voltage to function as a speaker. Therefore, we decided to use custom-made audio amplifiers. These special amplifiers have been developed and manufactured at the *Institute for Computer Music and Sound Technology* (Fig. 3). The amplifiers contain a toroidal transformer that is stepping up the drive voltage to the high level required. The electric power supply is provided either by a DC wall-plug adapter or a lithium-polymer battery, which can be charged from USB. To make the technology accessible to the community, the circuit board design is released under an open hardware license. The amplifiers' electronics are fitted in a plexiglass box, which turns the amplifiers into aesthetically appealing objects. Consequently, they can well be used as visual elements of the installation.

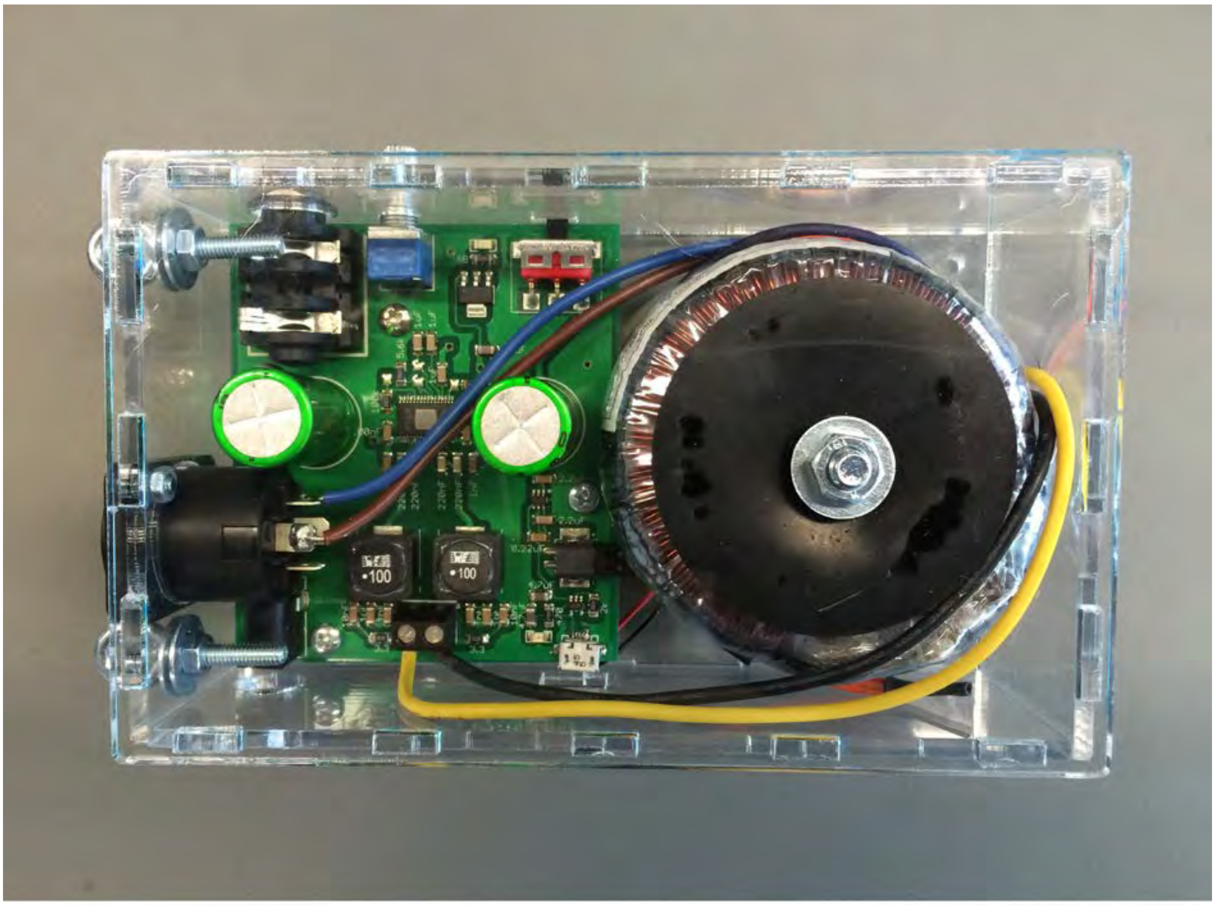

Figure 3. Audio amplifier electronics. Custom designed amplifier board to drive the piezoelectric speaker films.

3. Sound Synthesis

The sound synthesis method that drives the audio output in this setup consists of units organized in a network (Fig. 4). These units – nodes and connections – have certain functionalities. The nodes' function is mainly that of a control mechanism. Since the network contains many feedback loops, audio signals tend to accumulate which is very likely to lead to runaway conditions. This is avoided by switching to another connection whenever the volume rises over a certain threshold. In addition, a peak limiter comes into action when the signal remains above the threshold despite the rerouting. Finally, the summed input of every node is connected to a speaker (Fig. 5). The connections contain a delay line and a low-pass filter. Both the delay-times and the filters' cutoff frequencies have an important impact on the sonic characteristics of the output. In order to compensate for the loss of energy caused by the low-pass filter the signal is multiplied by a gain factor (Fig. 6).

Due to the feedback loops, this sound synthesis system is self-sustained: Once excited with a short audio signal it continuously keeps on sounding. As a result of the different delay times and the ever so often change of the signal path, a variety of different rhythmical gestures and textures emerge. This, however, also depends on the size of the network: It requires a certain quantity of nodes in order to provide a suffi-

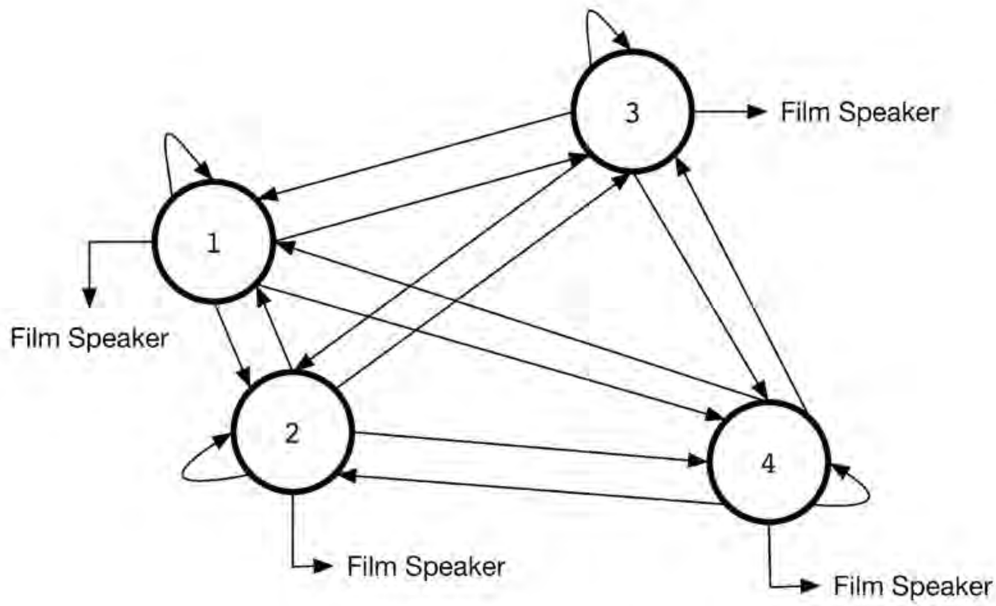

Figure 4. The network topology used for the sound synthesis. Each node is connected to every other node and to itself. (This diagram is simplified; the actual number of nodes in Sheet Music is eight.)

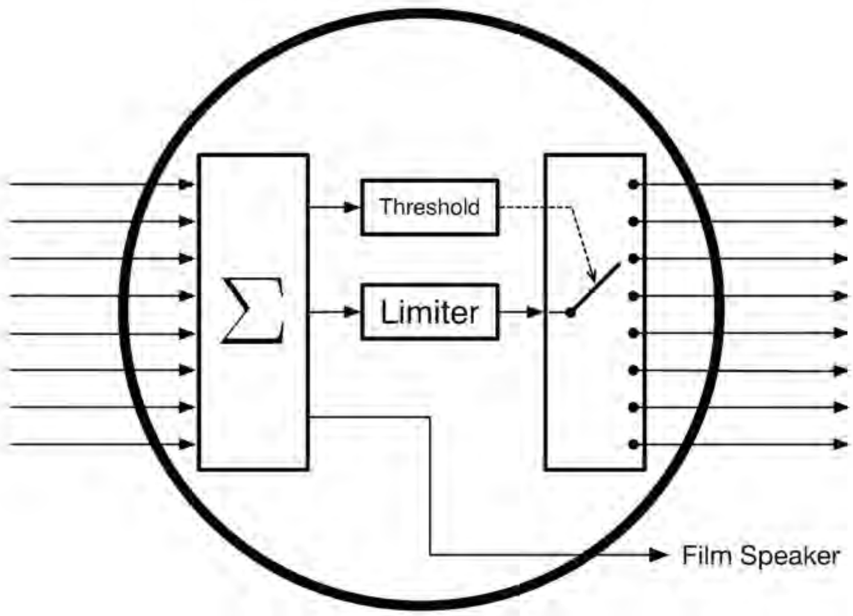

Figure 5. Signal flow inside a node. All input signals are summed and routed to one single output connection at a time, switching to the next connection whenever the level is above a given threshold.

ciently high number of different signal paths. This sound synthesis approach is neither intended to mimic natural sounds (e.g. musical instruments) nor to model any existing physical property. Rather, its artistically most compelling feature is the fact that this synthesis method permits to shape the timbral as well as temporal properties of the sound within the same formalism.

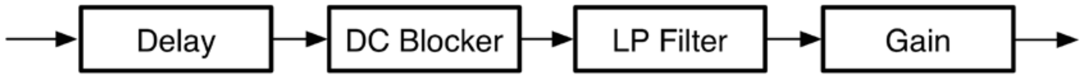

Figure 6. Signal flow inside a connection. The signal is delayed, the DC component of the signal is removed, the signal is low-pass filtered and a gain factor is applied.

4. Discussion

The sound installation *Sheet Music* is an attempt to realize a work that relates the sound generation to the spatial positions of the physical objects. The effect of this relation is twofold: it furthers the tangibility of the underlying generative sound synthesis algorithms and it increases the site specificness of the installation. The spatial layout of the objects makes the network explorable, i.e. the listener can move among the object and hear the sound of the network "from within". The correspondence of the delay times and the distances between the objects is particularly well perceivable when the sounds are short and of a percussive quality. This is achieved by the switching between the connections, which occasionally cuts the sound into short bits. The signal propagation inside the network is reflected in spatial effects, i.e. sounds moving about. As far as the site specificness is concerned, the correlation of the installation's layout and the sound synthesis is also meaningful. The actual sound of the installation depends on a given spatial layout. Due to the fact that this layout is always to a certain extent determined, or at least influenced, by the size and shape of the exhibition space, every rendition of this installation is unique.

In view of this relation between sound synthesis and spatial position of the objects, it must be asked how the sound synthesis algorithm has to be designed in order to obtain the results postulated above. In our case, there are two major aesthetic criteria: First, the algorithm must produce a sufficiently large number of different musical gestures and, second, it must maintain a certain rate of change yielding an output that is neither too static nor too repetitive. No matter what the spatial layout of the objects might be, the algorithm must work flawlessly with the parameters derived from that layout. Therefore, in order to meet the aesthetic criteria in any case the algorithm must be quite robust. At the same time, these parameters should have a sufficiently height impact, in order that the algorithm actually yields a perceivably different result in a different setup. The fact that this installation can take on many different forms is most important. The work is unfinished and only completed when realized in a certain environment. In this vein, the sound installation *Sheet Music* is what Umberto Eco calls a *work in movement* [5].

The relation between the generative algorithm and some of the physical properties of the installation's objects encourages the listener to engage with the work both on a perceptual and a conceptual level. Such an ontology alignment between the material and the immaterial aspects is particularly well realizable in an audio-visual, i.e. multimodal, environment. Yet, establishing a relation between the algorithm and the physical properties of an artwork serves not only the purpose of understanding the particular algorithm of this artwork. It provides also an example of how generative art can play a role in the discourse about today's information society where objects have become more and more de-physicalized and where there is an existential need to find ways to understand the principles of networked structures [6].

The ontological alignment would be even more compelling in an interactive situation, as this would take the tangibility one step further. We therefore intend to address interactivity in future works, for instance, by making the objects movable. This would allow the listener to alter the delay-times or even the topology of the feedback network and listen to the sonic changes in real-time.

5. References

[1] Kenichi Ohya, *Sound Variations by Recurrent Neural Network Synthesis,* in: Proceedings of the International Computer Music Conference, Ann Arbor 1998, pp. 280–283.

[2] Alice Eldridge, *Neural Oscillator Synthesis: Generating Adaptive Signals with a Continuous-Time Neural Model,* 2005.

[3] Daniel Bisig and Philippe Kocher, *Early Investigations into Musical Applications of Time-Delayed Recurrent Networks,* in: Proceedings of the Generative Art Conference, Milano 2013.

[4] Alan Dorin, Jonathan McCabe, Jon McCormack, Gordon Monro and Mitchell Whitelaw, *A framework for understanding generative art,* in: Digital Creativity, 23:3-4, 2012, pp. 239–259, here: p. 240.

[5] Umberto Eco, *The Open Work,* Harvard 1989, p. 12.

[6] Alice Eldridge, *The Poietic Power of Generative Sound Art for an Information Society,* in: Proceedings of the International Computer Music Conference, Ljubljana 2012, pp. 67–81.