SCANNED SYNTHESIS: AN INTRODUCTION AND DEMONSTRATION OF A NEW SYNTHESIS AND SIGNAL PROCESSING TECHNIQUE

Richard Boulanger
rboulanger@berklee.edu
Berklee College of Music
Boston, MA 02115 USA

Paris Smaragdis
paris@media.mit.edu
MIT Media Lab
Cambridge, MA 02139 USA

John ffitch
jpff@maths.bath.ac.uk
University of Bath
Bath BA2 7AY, UK

Abstract

In this paper we introduce extensions and some unique features of the scanned synthesis algorithm. Scanned synthesis was introduced as a fairly simple principle, it did however contain some powerful ideas for sound generation. We will show extensions of scanned synthesis to higher dimensional shapes, using arbitrary scanning trajectories and how it relates to the idea of dynamic wavetables. We will also discuss its application to well known synthesis techniques that are traditionally based on static wavetables.

INTRODUCTION

Verplank, Mathews and Shaw first presented the idea of scanned synthesis in 1999. The system was based on a string model composed of masses and springs (Figure 1). The proper equations to emulate the physical behavior of a string were implemented (Verplank *et al.*). What was peculiar in this implementation though was the way that pitch was manipulated. Rather than change the number of the masses to simulate a string of different length, that number was kept constant. What changed was the rate at which the masses were scanned to produce the output waveform (by scanning we mean reading each mass displacement and translating it to an instantaneous amplitude value for the resulting output sound). This ended up being the basis of scanned synthesis. A physically inspired dynamical set of masses that was excited by external influences (so as to induce motion) and was then scanned at arbitrary rates (very much like a wavetable) to produce sound.

The nature of the scanned synthesis algorithm brought a special dynamic quality to the sounds that were produced. Without the use of external filters and other algebraic operations it was possible to create textures that otherwise required them[=].

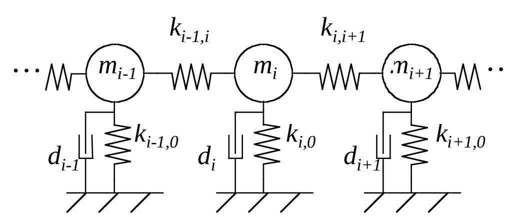

Figure 1. Three masses connected as a string. Each mass (m_i) is being connected with springs ($k_{i,j}$) to its adjacent masses, and to ground which simulates the effect of a centering force. These centering springs also have dampers (d_i).

[=] Although it is certainly true that the operations required to update the string space are indeed filters, they do not however apply on the resulting audio output. They are part of the waveform generating procedure.

EXTENSIONS TO THE INITIAL MODEL

After the initial implementation, variations of the original idea were introduced by the authors of this paper. The following sections will describe the most significant ones.

Mass Connections

The most significant extension was an arbitrary connection scheme for the masses. The initial model provided a system of masses which were connected in a string-like manner, thereby composing a structure similar to a one dimensional oscillator. By introducing additional springs that connected each mass to all of the rest (Figure 2), it became possible to produce arbitrary shapes. With this model it becomes possible for a user to disable unwanted filters, by setting their stiffness to zero, so as to construct any physical shape desired. If, for example, only the springs of adjacent masses had significant stiffness, we would then form a circular string (Figure 3)

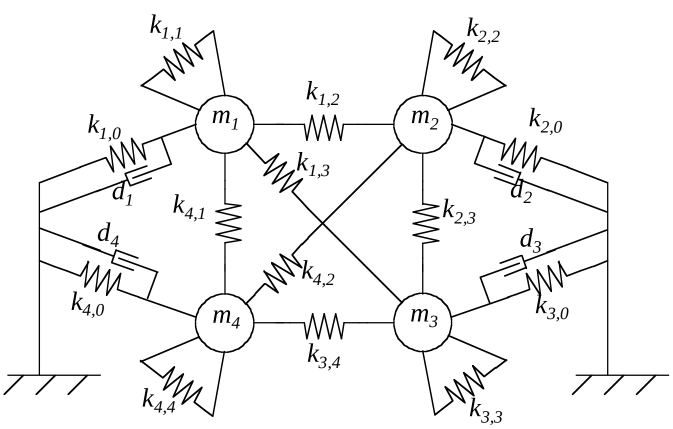

Figure 2. A fully interconnected system of four masses. In addition to all masses being connected to each other, they are also connected to the ground (which provides a centering force) and even to themselves.

For additional flexibility the springs were constructed so as to have an independent stiffness depending on the direction of the excitation. This makes possible the creation of strange constructs that propagate sound only towards specific directions.

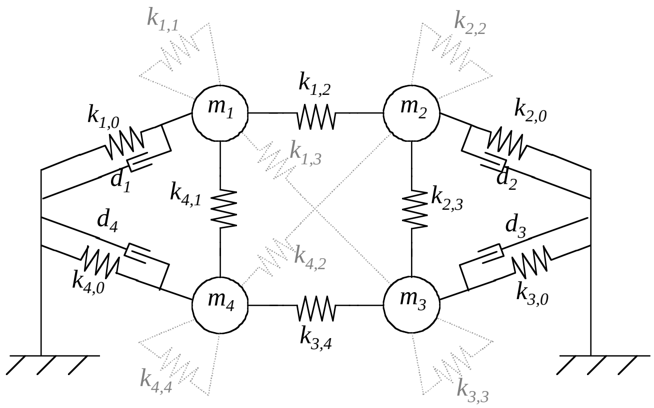

Figure 3. The system in figure 2 with the dimmed springs ($k_{1,3}$, $k_{4,2}$, and $k_{i,i}$) having insignificant stiffness, composes a circular string.

Scanning Trajectories

Given the power to design arbitrary shapes though, raises the issue of the scanning trajectory. In the initial scanned synthesis model, where we had a string structure, it was obvious that the way we scan the masses would be moving from each mass to its adjacent one. If we decide however to implement a higher dimensional structure, such as a grid or a torus, we immediately discover that there is no right way by which to scan the masses (Figure 4).

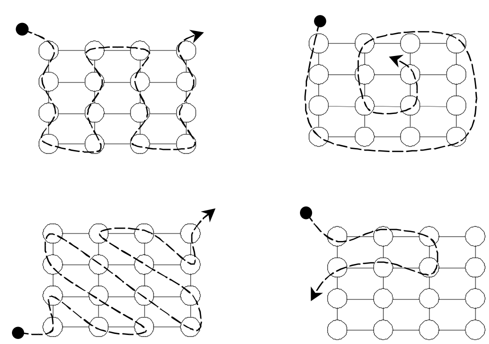

Figure 4. Samples of different ways by which to scan a 4×4 grid of masses.

The only rule we can supply is that the scanning trajectory should be moving between masses that are connected. This will insure that the scanning follows a continuous trajectory which will create a smooth sound (although, this rule is best if broken, so that the resulting discontinuities create additional periodicities). In general, the selection of a scanning trajectory is a non-intuitive procedure that depends on what effects the user desires to obtain from the system.

Unique Object Parameters

As mentioned above, one extension was to make all springs have a unique stiffness (on both directions). Unique parameters were also given in all aspects, so it is possible to have a different mass, different centering forces and damping for each point. Non-uniform parameter distributions have proved to be crucial in the production of interesting and dynamic timbres. Uniform and static settings produce a highly uninteresting ÒpluckÓ

Auditory Excitation

The original scan synthesis model was to be excited with exotic interfaces, initially with the Phantom, a popular haptic interface, later with the radio baton (Boulanger and Mathews). The obvious MIDI keyboard control came soon thereafter. However, most of these methods provided fairly simple ways to excite a string, which generally hit the string a specified point (with a specified hammer shape) and thus initiated movement.

A richer way to excite the string was to drive it with an audio signal. The approach taken was very straightforward. The excitatory audio samples were added to the displacement of the masses. The correspondence of time to mass was determined by the scanning trajectory. The first sample would displace the first mass in the scanning trajectory, the second sample the next mass, and so on. Once weÕre reached the last mass the next sample is added to the first mass and we start again. The reasoning behind this scheme was that due to the same adding and reading trajectory, we do get a processed version of the input in the same time order as it came in. If we were to overlay the input samples with disregard to the scanning trajectory we would introduce considerable discontinuities. The effect that this method produces, is a strange resonance of our structure to the frequencies of the excitatory input.

SCANNED SYNTHESIS AND DYNAMIC WAVETABLES

Scanning for Audio

It is probably apparent at this point that scanning the masses in order to produce sound, is equivalent to reading off a wavetable that is dynamically changing. For example in the case of the simple string we have a string similarity to a one-dimensional oscillator. A two dimensional grid, corresponds to a two-dimensional oscillator (such as those used in terrain synthesis). Likewise higher dimensionalities relate to increasingly more dimensional oscillators. The only feature that makes this approach distinct from wavetable lookup, is the fact that the wavetable shapes itself dynamically.

We feel that this is a significant representation for sound design in scanned synthesis. By using the wavetable abstraction we can scan the masses using popular sound synthesis methods. It is possible to adapt any algorithm that makes use of wavetables to use such a dynamic wavetable.

The obvious example here is frequency modulation synthesis (Roads). It is possible to use the dynamic wavetable as the source for either, or both the carrier and the modulating oscillator. Doing so will result in very rich sounds, which exhibit a strong dynamic character, especially so in the case where the dynamic wavetable is used to modulate.

Granular synthesis (Roads) is also a very good candidate for an implementation with dynamic wavetables. By using grains constructed from dynamic wavetables, it is easy to generate very rich textures, without having to introduce continuous parameter variations or inject more randomness.

Scanning for Control

In addition to the aforementioned ideas for producing sounds, it is also possible to use the mass displacement data for control purposes. Especially in the string model, the predictable wave propagation patterns can be used for controls requiring smooth wave-like motion.

Spectral shaping, is an example of control use. Upon obtaining the instantaneous amplitude spectrum of an audio signal we can scale it by the state of the mass displacements. Assuming a still string we would have no sound. As we pluck the string at the left side (which corresponds to the low frequencies), we allow these frequencies to sound. By proper setting up of the parameters we can obtain travelling waves that will implement multiple sweeping filters on our sound. If the input to be scaled is white noise or just an impulse, we would be manipulating a bank of oscillators, very much like an additive synthesizer.

DESIGNING SCANNED SYNTHESIS SOUNDS

Other than a couple of specialized programs, the only easily accessible version of a scanned synthesis implementation is in Csound, using the opcodes `scanu` and `scans` (Vercoe, Boulanger 2000).

Along our research we have come across many configurations yielding interesting sounds using scanned synthesis. The most important point to stress, is the use of non-uniform settings. A good example is the case where we have a sting configuration and the centering forces of the masses, are ramping from a low value towards a higher one. This will force the oscillations on the string to be Ôpushed Õ progressively towards one side of the string. This Ôpushing Õ of displacement towards the one end produces a sweeping filter effect as we scan the wavetable (since the wavetable transforms slowly from a triangle waveform to an impulse).

Similar examples are easy to generate by having such non-uniform settings for damping, mass, and even spring tensions. Interesting effects also occur when we construct one-way waves (using a configuration in which only one of the springs connecting two masses has stiffness, resulting on one-way propagation).

FUTURE EXTENSIONS AND CONCLUSIONS

For the sake of dynamic wavetables, it might be worthwhile to experiment with other ways to update the wavetable. Scanned synthesis is using a mass-spring model out of a specific series of circumstances that led to its discovery, but it could be replaced by any arbitrary updating function. It has been suggested that non-linear springs could be added, as well as additional physical clutter. However more abstract and mathematical models for updating can also be constructed, such as a chaotic system, a genetic algorithm, or cellular automata (Garcia 2000).

Perhaps the easiest avenue of research to pursue, is the effect of dynamic wavetables on wavetable-based synthesis methods. We have composed some simple experiments, but we still donÕt have a clear idea of how powerful this can be. We have plenty of heuristics on the effect of the physical parameters to the resulting sound for straightforward scanning, but how these will translate on FM, granular, or other methods is not clear yet. The creation of more wavetable updating methods will act as a constant supply of fuel for this particular work.

To our delight and frustration, scanned synthesis has opened many ideas that are hard to pursue all at once. We hope that in this paper we might have provided some excitement to get interested researchers working on some of these ideas.

Acknowledgements

The authors wish to acknowledge Bill Verplank, Max Mathews, Robert Shaw, Perry Cook and other consultants who participated in the scanned synthesis project at Interval research in the summer of 1999. Their help and discussions, as well as the financial support of Interval research, have been instrumental for the development of the ideas in this paper.

REFERENCES

Boulanger, R.C. 2000 Scanned synthesis tutorial available online at: http://cSounds.com/scanned/toot/

Boulanger, R.C. and Mathews, M.V. 1997 ÒThe 1997 Mathews Radio-Baton & Improvisation ModesÓ *proceedings of the International Computer Music Conference*, Thessaloniki, September 1997.

Garcia, R. 2000 Personal communication regarding the use of cellular automata as dynamic wavetables.

Nelson, J.C. 2000 ÒUnderstanding and Using Csound's GEN RoutinesÓ *The Csound Book*, Richard Boulanger ed.

Roads, C. 1996 ÒThe Computer Music TutorialÓ MIT Press, Cambridge MA

Verplank, W.L., Mathews, M.V., Shaw, R. 2000 "Scanned Synthesis", *Proceedings of the International Computer Music Conference*, Berlin, August 2000.

Vercoe, B. ÒThe public Csound reference manual: Canonical version 4.03Ó Media Lab MIT and contributors. Edited by John ffitch, Richard Boulanger, Jean Pich , and David Boothe.