Designing for Emergent Behavior: a John Cage realization

Christopher Burns

Center for Computer Research in Music and Acoustics (CCRMA), Stanford University
cburns@ccrma.stanford.edu

Abstract

A new software realization of John Cage's Electronic Music for Piano *is described. Inspired by dedicatee David Tudor's work with feedback circuits in live electronic music, the principal audio component of the software is an audio feedback network with a number of chance-determined, time-varying parameters. The behavior of this feedback system is unpredictable, so the software interface is organized around higher-level, statistically-oriented controls. The result is an instrument which displays emergent behavior, guided but not controlled by its audio inputs and its operator.*

1 Introduction

Electronic Music for Piano (1964), by John Cage, is one of the composer's most permissive scores. (Cage 1968) While the ensemble (piano or pianos and electronics) is more circumscribed than in works for indeterminate performers like Cage's *Variations* series, *Electronic Music for Piano* lacks the systems of discipline associated with that series (often considered among the most open of Cage's works). *Electronic Music for Piano* is dedicated to David Tudor; presumably the "loose" nature of the indeterminate and improvisatory characteristics of this work have much to do with Cage's trust in his friend and colleague, and the shared culture developed through extensive collaboration in both the realms of piano and live electronic performance.

In keeping with the increasingly conceptual and improvisatory nature of Cage's approach to music with live electronics during the 1960s, the handwritten prose score of *Electronic Music for Piano*, complete with crossouts and emendations, is suggestive rather than prescriptive. There are lists of potential technical means: "feedback, ~~and~~ changing sounds (microphones, amplifiers, loudspeakers – separate system for each piano)." There are also a number of provocative suggestions about the performers' potential activities and strategies, in many instances drawn from compositional and performance techniques used in Cage's other works. For instance, the phrase "observation of imperfections in the silence in which the music is played" refers to a compositional method used in *Music for Piano*, in which Cage placed notes on the page where he noticed imperfections in the paper he was using. (Pritchett 1993) References to other works include *Atlas Eclipticalis*, *Cartridge Music*, and the *Variations* series.

2 Score and Realization

There are an enormous number of potential realizations of *Electronic Music for Piano*. Indeed, it would be difficult to specify criteria for the validity of a performance. Nevertheless, the score provides a rich interpretive context for realization. A number of fragments from the score were crucial stimuli for our performing version of the work.

Cage's notations include the phrases "feedback" and "for David Tudor." During the 1960s and 70s, Tudor gradually reoriented his career from the performance of avant-garde works for piano to the creation of live electronic music, with electronic feedback systems as the defining component of his work. (Adams 1997, Chadable 1997). In appreciation of Cage and Tudor's work in the domain of feedback, we designed this realization around an acoustic feedback network, implemented in software.

One fragment from the score reads "*Music for Piano 4-84* / single tones (K, M, P) and / noises (I, O)." Cage refers to the notation of his *Music for Piano* (1952-56) series, in which "K" stands for keyboard, "M" for muted, "P" for pizzicato, "I" for inside the piano, and "O" for outside the piano. (Cage 1960) Another fragment, "(no ~~observation of~~ notation)," leaves ambiguous whether the pianist is to perform part or all of *Music for Piano 4-84*, or improvise in the style of that piece (single notes and sonic events, with occasional overlaps between the resonance of one event overlaps with the onset of the next). In performances of our realization, pianist Christopher Jones has elected to play *Music for Piano 69-84* from the score.

Cage writes "as though there were / ~~take~~ a drawing of the controls / (volume, tone) available and – / on a transparency – transcription / from astronomical atlas ~~suggesting~~ / which (^were it superimposed) ^gives / suggestions for use of controls (~~not explore~~)." This refers to the compositional process for *Atlas Eclipticalis*, in which Cage traced points from star charts and then superimposed musical staves. Initially the notion of a "music of points" seemed far removed from the prolonged events characteristic of acoustic feedback. Cage's suggestion eventually came to fruition in the realization as a series of generators producing randomly determined breakpoints for time-varying values. The automation of these parameters dramatically reduced the performance complexity of the instrument, and became a major component of the emergent behavior described in Section 5.

This notation also points to Cage and Tudor's interest in music beyond the limits of their own personal aesthetic taste. Beyond the score of *Electronic Music for Piano*, Tudor explained in a 1987 interview, "What I would like to find is an improvisation that is not descriptive of the performer, but is descriptive of what happens, and which is characterized by an absence of intention.... The performer... is not apt to make a discovery spontaneously. I want to find ways of discovering something you don't know at the time that you improvise.... The first way is to play an instrument over which you have no control, or less control than usual." (Kostelanetz 1987) In this realization, the electronics and particularly the feedback network are designed to behave unpredictably, allowing the software to speak for itself, and creating the sort of unstable environment which Cage and Tudor sought out in their work. The system becomes an equal partner in the improvisational process.

An additional decision – to downplay the performance aspect of the electronics part – came from outside Cage's score. Despite Cage's "~~piano~~" notation, an onstage pianist is necessary in many concert halls, and provides a desirable visual focus for the performance. The electronics operator is then free to perform offstage, in the center of the audience. This location optimizes the operator's ability to hear the spatialized audio, and minimizes any visual distractions caused by the use of simple and standard computer input devices (MIDI faders, QWERTY keyboard, and mouse).

3 Design Process

The *Electronic Music for Piano* realization was developed through an iterative process of coding, rehearsal, and discussion. Four separate revisions of the software were developed before and in between rehearsals. The rapid prototyping capabilities of Miller Puckette's Pd language meant that refinement of the control algorithms and other software development could also take place during rehearsal and conversation with the pianist. In general, the number of features and controls grew over time, with occasional weedings to prevent the number of performance parameters from becoming difficult to manage.

The feedback network was the first section of the instrument to be implemented and tested. It quickly became apparent in rehearsal that the instrument would need additional capabilities in order to mediate between the decaying, fixed-pitch sounds of the acoustic piano and the sustaining, swooping, utterly unpianistic character of the feedback network. The ring modulation and resonator sections were added as a means to impose glissando to the piano's sound, connecting it to the sonic world of the feedback network. (Having the piano sound from the loudspeakers was another simple but important point of connection with the electronics). A "stompbox distortion" section was tested alongside ring modulation and resonance, but its most satisfying use was as a simple amplifier – so it was stripped of features in order to form the amplification section. In addition to mediating between feedback and piano, the availability of a variety of electronically modified piano sounds pointed towards Cage's "~~piano~~" instruction.

Another characterizing feature of the instrument grew directly from improvisation in rehearsal. Early drafts of the software included mutes on the audio inputs and outputs, controllable from the QWERTY keyboard. In rehearsal, we found it appealing to switch these mutes in and out quickly, imposing a layer of rapid rhythmic activity on top of the relatively slow rate of change in the piano. This idea eventually grew into the gating sections of the instrument, which contribute substantially to the emergent behavior of the realization as a whole.

Finally, as is often the case, serendipity played a role in the design. A separate, simultaneous project to develop constant-power eight-channel panning software with randomly-generated trajectories produced good results. It proved easy and effective to sum the outputs of the amplification, ring modulation, and resonator sections and distribute them in space with the panning system.

4 The Software Instrument

The *Electronic Music for Piano* realization is organized in a number of modules. Stereo audio inputs from microphones in the piano are passed to the stereo inputs of the amplification section, the ring modulation section, and the resonator section in parallel. Each of these three modules has a stereo output which passes through a gating section. The three "left" gates are then mixed and passed to an eight-channel spatialization module, while the three "right" gates are combined and passed to a second such spatializer. The feedback section is largely independent of this signal path; it is fed by the microphone inputs and sends it output directly to the eight loudspeakers.

4.1 Amplification Section

The amplification section is the simplest module in the instrument. A single gain parameter controls the amount of dry piano signal sent to the gating and spatialization sections. Like all the parameters in the instrument which are exposed to the operator and offer a range of values, it is controllable by mouse or MIDI fader. Table 1 summarizes the performance controls of the instrument.

4.2 Ring Modulation Section

This section multiplies each of its stereo audio inputs by a separate sine wave oscillator. Both oscillator frequencies track the pitches played by the pianist via MIDI note-on messages. (In performances without a MIDI-equipped piano, Pd's `fiddle~` pitch-tracker has proven an adequate substitute. The pitch tracking, and the ring modulation more generally, are used for coloristic effect and the imposition of glissando, not for precise frequency control). The oscillator sweep time interval, as it transitions from pitch to pitch, is set independently for each oscillator. The electronics operator has some control over the sweep time interval. As with many of the controls in the instrument, the operator determines the maximum available value, and then the software generates random positive values for each oscillator beneath the maximum. Another parameter for this section allows the operator to adjust the output volume of the module.

Section	Function	Range
Input	Input mutes	on/off (x2)
Amplification	Gain	0 - 1.0
Ring mod	Gain	0 - 1.0
Resonator	Gain	0 - 1.0
Resonator	Feedback coefficient	0 - 0.99
Ring / Reson.	Max. sweep time interval	0 - ∞
Gating	Maximum gate interval	0 - ∞
Gating	Probability of open gates	0 - 100%
Feedback	Gain coefficient	0 - 0.99
Feedback	Delay length scaling factor	0 - ∞
Feedback	Sweep time scaling factor	0 - ∞
Feedback	Sustain time scaling factor	0 - ∞
Feedback	Output gain presets	8 presets
Output	Output mutes	on/off (x8)

Table 1. Operator-controlled parameters for the *Electronic Music for Piano* realization.

4.3 Resonator Section

The resonator section functions in much the same way as the ring modulation section, using piano-keyboard tracking to match the length of a pair of delay lines with feedback to the presumed frequency of the input audio. The output of this section produces the closest points of timbral contact between transformed piano sound and the feedback network output. As with the ring modulation section, each delay line has an independently calculated sweep time, and the resonator section shares the operator's sweep time interval parameter with the modulation section for the sake of simplicity. A second parameter provides gain control.

4.4 Gating Sections

Stereo gating sections are used to process the outputs of the amplification, ring modulation, and resonator modules. As described in Section 3, the gates grew out of a simple muting and unmuting of the instrument's outpu t audio. In their final form they are somewhat more sophisticated. Each side of the stereo gate switches in and out automatically and independently, while the electronics operator is given two global parameters with which to guide the process. In analogy with the ring modulation and resonator modules, the operator can determine the maximum time interval of the gates, with the software randomly selecting positive values below that maximum. Additionally, the operator can control the probability that a gate will be open in any given time interval. Manipulating these two parameters in concert can lead to a wide variety of rhythmic textures and layering, from spare to busy, and burbling to sustained. The possibilities increase further as the amplification, ring modulation, and resonator sections are brought in and out of play with their respective gain controls. The interactions between these five parameters exemplify the type of control being pursued for this realization. The quality of control is simple, intuitive, and effective. At the same time, especially when long maximum time intervals and low probabilities are requested, moment-to-moment results are unpredictable.

4.5 Spatialization Sections

The left and right sides of the three gating modules are then fed into a pair of independent constant-power panning modules. Each of these modules pans its input between randomly selected pairs of loudspeakers, chosen from the eight available. Pans tale place over a scale of randomly selected durations, from three to 192 seconds in increments of three seconds. No parameters for this process are exposed to the operator. Spatial interactions between the piano, the transformed piano sound, and the feedback network are left entirely to chance. In practice, serendipitous interactions occur with great reliability and no predictability.

4.6 Feedback Section

The feedback section of the instrument passes the two microphone inputs into a circular chain of delay structures. Each of the eight nodes arranged around the circle contains two time-varying delay lines. The delay lengths, the sweep duration between each length, and the sustain duration at each length, are randomly and independently generated for each delay. Delay lengths correspond to the frequency range available on the piano keyboard, an idea suggested by Cage's inclusion of blank staves and ledger lines indicating the complete range in the score. Sweep and sustain times vary between fifteen and fifty seconds. The operator has access to global scaling factors for all of these parameters; it would be needlessly complex to set scale factors individually for each of sixteen delays. The scaling factors can compress or expand the parameter ranges given above considerably. The operator also has a global control over the gain coefficient of the network; this is the "strongest" parameter in the feedback section of the instrument.

As new energy is continuously added to the feedback system through the microphone inputs while preexisting signals continue to circulate, gain control is essential. The chance-derived input gains prevent the kind of "coefficient balancing" typical of waveguide networks. (We have performed with a variety of different configurations within and between the delay nodes, but have avoided use of "conventional" upper and lower waveguide rails). Instead, gain control is implemented through nonlinear waveshaping functions. These "soft clipping" functions prevent amplitude peaks from exceeding the system limits, and color the resulting sound of the network. (Sullivan 1990, Burns 2003)

The feedback network uses a different approach to spatialization from the other sections of the instrument. Each node is associated with a particular loudspeaker; the delay outputs are routed directly to that loudspeaker, with an independent gain control. As sound propagates through the feedback network, complex panning and other spatial effects occur without any performer or software intervention. The performer does not have control over each individual gain stage. Instead, there are eight selectable presets, each of which modifies all eight of the gain stages associated with the different loudspeakers. The presets are triggered via a keystroke, and are ordered roughly from softest to loudest, with varying spatial distributions and weightings of different segments of the feedback network.

Both the operator, through the parameters mentioned above, and the pianist, via the microphone inputs, have influence over the feedback network. However, they do not have command of the process; the network speaks in idiosyncratic and unpredictable ways, sometimes imitating onsets and pitches played at the piano very precisely, sometimes remaining quiet during busy passages, sometimes bursting into noise in the middle of a long silence. The operator can reliably squelch the system output by turning the gain scaling factor down to zero; it can usually but not always be encouraged to sound by bringing that scaling factor near unity. Long sustain times, which produce stable delay lengths, and allow resonances to build, also encourage the system to perform. Typically when the network is "quiet," it is actually producing DC.

The sonic character of the feedback network is equally varied and idiosyncratic. Complex, swooping pitch contours with continuous micro-alterations of timbre are typical, while the continuously varying delay lengths produce shifting, inharmonic harmonic movement. Depending on the gain settings, punctuating noisy explosions may also be frequent. The network has a strong heritage in waveguide synthesis. While it does not attempt to model any particular acoustic instrument, it does display the same complex sounding behaviors that waveguide physical models do. (Essl 2003) Indeed, by disregarding some of the traditional tuning techniques for physical models, this network is more likely to enter marginal and turbulent states. The system's behavior is genuinely emergent; the output is musical, articulate, and often surprising.

5 Conclusions: In performance

In this realization, signal processing and control are carefully tailored to one another, in accordance with the inspirations and directions provided by the score of *Electronic Music for Piano*. In particular, the majority of the control data is produced through algorithmic processes, with the operator focused on higher-level "steering" of those algorithms in performance. This general strategy for control is applicable in almost any situation where the parameter space for signal processing, synthesis, or other generated output is large. In this instance, the decoupling of the operator from certain details of performance is at the musical core of the realization.

Through a number of performances, the instrument has offered a substantial vehicle for the operator's musical expression, engaging and layering the treatments of the pianist's performance, and guiding the feedback network to achieve a cohesive, unified, and interesting performance. At the same time, the less predictable aspects of the software (especially in the gating sections and the feedback network) repeatedly created musical challenges for both the pianist and operator to solve, illuminating David Tudor's idea that unstable musical situations can free performers from taste and memory and lead to new discoveries.

This is the unusual aspect of this realization and instrument; the electronics are designed to guide the operator's musical choices just as the operator guides the electronics. There is a symbiosis of piano, pianist, software, and operator; in performance the situation is one of improvising with the electronics, rather than using the electronics to improvise. The unpredictable and emergent aspects of the electronics' behavior help foster intense listening and communication between the pianist and electronics operator in performance.

6 Acknowledgements

Thanks to Tamara Smyth for waveguide theory and practice, and especially to Christopher Jones for his collaboration on the realization of *Electronic Music for Piano* in rehearsal and performance.

References

Adams, John D. S. (1997). "Giant oscillations: the birth of Toneburst." *Musicworks* 69, 14-17.

Burns, Christopher. (2003). "Emergent Behavior from Idiosyncratic Feedback Networks." In *Proceedings of the International Computer Music Conference*. San Francisco: International Computer Music Association.

Cage, John. (1968). *Electronic Music for Piano*. New York: C.F. Peters.

Cage, John. (1960). *Music for Piano 69-84*. New York: C.F. Peters.

Chadabe, Joel. (1997). *Electric Sound*. Saddle River, New Jersey: Prentice-Hall.

Essl, Georg. (2003). "The Displaced Bow and Aphisms: Abstract physically informed synthesis methods for composition and interactive performance." *Florida Electronic Music Festival*, University of Florida, Gainesville, Florida.

Kostelanetz, Richard, ed. (1987). *Conversing with Cage*. New York: Limelight.

Pritchett, James. (1993). *The Music of John Cage*. Cambridge, England: Cambridge University Press.

Puckette, Miller. Pd. http://www-crca.ucsd.edu/~msp/

Sullivan, Charles. (1990). "Extending the Karplus-Strong Algorithm to Synthesize Electric Guitar Timbres with Distortion and Feedback." *Computer Music Journal* 14(3), 26-37.