Three Meditations

For Prepared Piano and Computer

Thomas Ciufo
Brown University

April 2002

Contents

List of Figures

List of Listening Examples

Introduction

Three Meditations

For Prepared Piano and Computer

Three Meditations is a live performance piece that utilizes physical treatments inside an acoustic piano, as well as electronic treatments provided by computer-based digital signal processing. The piece is constructed in three movements, each with its own unique sonic signature. *Three Meditations* is a contemplation of the textural potential of the piano and an attempt to sidestep conventional notions about the instrument.

This paper will include a brief historical comparison/contextualization, as well as a general examination of the aesthetic concerns faced in *Three Meditations*. Technical concerns will focus on issues of software design and implementation and musician-software interaction. Simple block diagrams will show the design and structure of the software patches. The paper will conclude with a description of future developments and research directions.

Background and Context

There is no shortage of examples of composers using the acoustic piano in the context of unconventional/experimental music or electronic music. George Antheil's *Ballet Mécanique* is an interesting example that brings to mind the piano’s long history of mechanization and automation.

> Antheil wrote several versions of the piece. The very first, written in 1924 calls for 16 player pianos playing four separate parts, for four bass drums, three xylophones, a tam-tam, seven electric bells, a siren, and three different-sized airplane propellers (high wood, low wood, and metal), as well as two human-played pianos (Lehrman 2002).

The sonic alterations and extensions to the acoustic piano—often referred to as preparations—introduced by John Cage in the late 1930's were an important development and could be viewed as a precursor to many of the electroacoustic compositions that would follow. This connection is most evident in the use of electronics, both live and tape based, to extend the timbral range of the instrument. I have been influenced and inspired by Cage's works for prepared piano as they open the instrument up to the larger and more interesting world of complex timbres.

> The prepared piano is a regular grand piano into which various objects of wood, metal, rubber, and felt are placed. These objects, or preparations as they are called, are put between various strings of the piano and individually alter the sound of each note, often eliminating or obscuring pitch, coloring timbre, and emphasizing the percussive qualities of the sound. For the next decade, Cage wrote primarily for this instrument, his work culminating, between 1946 and 1948, in the piece many people still consider his most traditionally beautiful: the twenty *Sonatas and Interludes* (Duckworth 1995).

An early example of the piano as a sound source can be found in the musique concrete studies Pierre Schaeffer made in 1948. Both *Etude aux Tourniquets* and *Etude Noire* use piano sounds recorded for Schaeffer by Pierre Boulez (Chadabe 1997). Two early examples of electroacoustic pieces involving the piano are Austin's *Accidents* (1967) and Davidovsky's *Synchronisms #6* (1970). These pieces are additionally noteworthy as they represent the two main technical approaches to extended piano writing, live electronic augmentation (now more commonly accomplished with a computer) and piano with synchronized tape augmentation. Davidovsky's *Syncronisms #6*, as well as several more recent pieces for piano and tape can be heard on Aleck Karis's CD *Secret Geometry* (Karis 1996).

An early attempt at real-time digital signal processing of the acoustic piano was ambitiously undertaken by Rainer Boesch, with the help of Daniel Weiss and Nicolas Sordet. Together they developed a real-time sound processing board at the Swiss Center for Computer Music. This system was utilized in *Clavirissima* for piano and real-time processing (1987) written and performed by Rainer Boesch.

> The sound of the piano goes through a microphone and AD-converters into the sound-processor's memory, where it is more or less modified and then mixed to the acoustical piano-sound using DA-converters and loudspeakers.
> The program of the sound-processor analyzes the sound with a zero-cross technique, sets the beginning-pointer on such a zero-crossing and looks then for a sound-end, which allows it to produce all sorts of modified piano-sounds. The piece fixes precisely the reactions of the machine. The pianist triggers the reactions of the machine and reacts upon it (Kessler 2002)

Of particular interest, especially from a technical standpoint, is Cort Lippe's *Music for Piano and Computer* (1996) as it uses real-time digital signal processing instead of tape, and because is uses analysis of the audio input to control many aspects of the sound processing algorithms. Lippe's introduction to *Music for Piano and Computer* is as follows.

> *Music for Piano and Computer* (MPC) was composed in 1996 using signal processing Max (Puckette, 1988, 1991) running on the IRCAM Signal Processing Workstation (ISPW) (Lindemann, et al, 1990). The piece is divided into six sections, has a duration of approximately 18 minutes, and is part of a continuing series of interactive pieces by the author in which performers regulate many aspects of algorithms for both control and digital signal processing (DSP) in real-time performance situations (Lippe & Puckette, 1991). Analysis/resynthesis and convolution, both which require spectral domain analysis, are used extensively in the composition (Lippe 1997).

The use of real-time processing instead of tape is important to my work, as it makes an improvisatory approach much more practical. The use of analysis to drive the real-time dsp is also important as it goes much further towards creating an extended instrument or hyperinstrument (Machover 1992) as opposed to an acoustic piano with electronic accompaniment.

Other artists/composers that should be mentioned in the context of sonically interesting and unique piano works are Stephen Scott, Harold Budd, and Morton Feldman. Scott is important because of his development of unique acoustic piano performance techniques that greatly extend the sound range of the instrument.

> In 1977 he developed a 'bowed piano' technique: as many as ten players excite the strings of an open piano with monofilament bows and sticks coated with resin. The sounds produced resemble that of a mass of string instruments or a giant accordion, or occasionally electronic effects. His works in this medium range from short studies to the concert length *Vikings of the Sunrise*. His Bowed Piano Ensemble has toured extensively in the USA, Europe and Australia (Grove 2002).

Harold Budd also extended the sonic possibilities of the piano, primarily through his collaborations with noted producer Brian Eno (Tamm 1989). Together they crafted atmospheric ambient pieces that relied heavily on Eno's electronic treatments and exploited the possibilities of the recording studio. An especially colorful example of this work is *The Pearl* (1984) by Harold Budd and Brian Eno, with Daniel Lanois.

Morton Feldman should also be mentioned because of his highly original compositions for piano, such as *Intermission 5* (1952) *Piano* (1977) and *Palais de Mari* (1986). These pieces are especially interesting because of their spacious and contemplative quality, their beautiful lack of dramatic development, and the graceful way they move between sound and silence. Barbara Monk Feldman describes this phenomenon eloquently.

> The interplay of the varying decay of the sounds with silence renders a magnificent transparent quality: one could make analogies to light in Cézanne which no longer appears to be shining on the canvas but gives the illusion of emanating from it instead (Feldman 1989).

While none of the works mentioned above were a direct influence on *Three Meditations*, they are all part of my general musical awareness and have stimulated my interest in contributing to the diverse body of non-traditional music written for the piano.

Three Meditations: Conceptual and Aesthetic Concerns

Before starting to construct *Three Meditations*, I had a fairly clear and developed concept of how I wanted the piece to sound. This was especially important and beneficial because before I could work on the piece itself, I would have to build the instrument (or instruments) to play it on. I knew that timbre and texture were to be more important than melody, harmony, or rhythm, and for that reason I decided on the combination of physical preparations and digital signal manipulations. Because of this interest in timbre, I knew that playing inside the piano, in addition to playing on the keyboard, would be potentially interesting, and this technique turned out to be the main performance approach for the first movement.

I envisioned this piece as an ongoing exploration of sound possibilities rather than a fixed or finished musical entity. For this reason, I knew that structured improvisation would be a more flexible approach than using fixed notation. By structured improvisation, I am referring to a process by which certain overall characteristics such as pacing, mood, key or mode, and some melodic gestures stay the same each time the piece is played, but small to medium decisions are made freely in the moment. Improvisation was especially practical because of my commitment to perform all the pieces that I write, bypassing the need to translate my compositional ideas into any other language, such as notation. The improvisational approach was also beneficial because it allowed me to work on the hardware and software preparations while simultaneously developing and experimenting with the material for the piece. Using this approach, I hoped to discover how the instrument should be played and what type of musical gestures were most fitting. I have also learned after multiple performances of this piece that the hyperpiano—the combination of the acoustic piano, physical preparations and the way they interact with the signal processing and acoustics of the space—has unique sound characteristics each time it is set up, and it is important to be flexible and responsive to the instrument and setting.

I chose to construct the piece in separate movements for two primary reasons. First, I wanted to be able to explore separate and unique sonic and emotional spaces. I was hearing the piece as several self-contained vignettes that would hopefully fit together into a larger piece. The other motivation for separate movements was a long-range desire to add to the piece over the years, thus creating a large collection of movements that could be presented in unique and interesting combinations. At this point, I have written three movements, but I am sure there will be more.

My aesthetic approach to performance was to reject virtuosic playing in favor of a simple, contemplative style using relatively slow pacing and allowing time and space for each sound to be heard. This was a stylistic goal that also happens to fit well with my limited piano technique. My hope was that by constructing simple yet interesting sonic environments, I would be able to listen carefully and remain focussed, thereby performing this piece with sensitivity and care. This turned out to be a very difficult goal that I continue working towards each time I play this piece. To what degree this goal is achieved is up to the audience to decide.

Technical Concerns: Preparations

Experimenting with the physical preparations for the piano has been an ongoing process. I began with fairly conventional preparations, such as screws and bolts between the strings as well as other small metal or plastic objects. These can be fairly effective, especially if placed carefully. I often tune the placement of these items to create interesting multiple tones from a single key. I have purposely avoided damping treatments, as I am interested in letting the instrument resonate as freely as possible. To that end I always prop the sustain pedal to lift all the dampers from the strings.

Next I built several custom objects to be placed freely on the string so they can move inside the piano as they vibrate. These objects include rectangular pieces of thin metal, thin metal rods, and light gauge wire threaded through small metal nuts. At this point I progressed to playing the strings directly with various objects, including metal rods, and scraping the strings with small metal bells. I also developed a technique for lightly scraping a range of strings with an up side down Tibetan singing bowl. With this technique, it is possible to create resonant clusters of sound. One other unique performance approach is the use of repetitive bowing of individual strings with metal chains of varying design. This creates a raspy sustained sound that can be quite expressive.

I chose not to notate or otherwise make the preparations especially repeatable. This is in contrast to the various levels of detail used to notate piano preparations (Pritchett 1993). Because I am the only one who plays this piece, I prefer to experiment with and explore variations in the preparations each time I perform it. This allows me to customize the treatments to the specifics of each piano and acoustic space. One potential disadvantage of this variable approach is that the instrument can sound quite different each time it is prepared. For this reason it is ideal if I can arrange to have several hours to learn and explore the instrument before a performance. This is often impractical, in which case I try to embrace the spontaneity of the situation.

Technical Concerns: Amplification

Unlike the immersive or surround sound approaches used in some of my other works, for this piece I was interested in a very localized amplification method. The goal was to create a unified overall sound by combining the acoustic sound of the prepared piano with the amplified sounds of the computer based processing. After a fair amount of experimentation, I ended up placing the speakers directly under the piano and attempted to balance the perceived volume of the computer sounds with the acoustic piano sound. The best overall results were obtained by using at least four medium sized speakers under the piano arranged in a broad arc to maximize dispersion. I also found it helpful to send a small amount of the unprocessed piano signal directly to the speaker system to help further the illusion that the acoustic and electronic sounds were coming from the same place.

Although I have not been able to test this theory, I believe that better results, due to increased diffusion, would be achieved by using a more omni-directional speaker system such as the spherical speakers (multi-channel, outward-radiating geodesic speaker arrays) developed by Trueman, Bahn, and Cook (Trueman 2000). One potential challenge of the attempted localization and balance involves the dynamic range of the two sound sources. It was difficult (and perhaps not that important) to keep the acoustic and amplified sounds in balance over their complete dynamic range. This resulted in different perceived balance relationships at different dynamic levels. My approach to this was dependent on the specifics of a given movement. At times, in the first movement for example, I took advantage of this imbalance, and used the greater dynamic and timbral range of the computer treatments to highlight and exaggerate small gestures played inside the piano.

Technical Concerns: Sound Pickup

The other main problem of the localized amplification approach was sound isolation and gain before feedback. I would normally use open-air microphones placed above the strings of the piano to pickup the sound for computer processing. This approach worked fine during the early stages of experimentation when headphones were used to monitor the computer output. As soon as I started working with the localized speaker placement, feedback became a problem. Of the potential solutions available, I chose to explore two of the more practical approaches. First, I began to experiment with contact microphones, such as the hot spot pickups by K and K Sound Systems (K and K 2002). These are inexpensive surface mount contact transducers and as such have a very high gain before feedback ratio.

The configuration I have ended up with for a grand piano is three contact microphones placed on the under side of the sound board with one on the treble range, one on the bass range, and one near the middle of the instrument, further from the hammers. These locations are a matter of trial and error and are different for each instrument. The sound of these contact microphones is very different than the sound of open-air microphones, and in many applications they would sound too unnatural. In this situation, the sound was heavily manipulated in the computer, so the sound quality of the contact mics proved adequate. The sound can be optimized by careful placement and selective equalization of each microphone individually. Although I didn't use this technique, it is also possible to combine contact microphones with conventional microphones and balance them for best results. The other technique I used was to carefully gate the audio input in the software, thereby only selectively connecting the input to the output. This was only necessary for the second movement, which used several feedback prone synthesis methods.

Technical Concerns: Analysis and Real-Time Control

One important goal for this piece was to create the sound processing environment in such a way as to make it self contained and primarily driven by what is played on the piano (Lippe 1994). To this end, an analysis strategy was one of my first dsp concerns. Working in the Max/MSP environment (Zicarelli 2002) there are a number of signal analysis options available. One of the most useful objects for real-time signal analysis is the FFT-based fiddle object by Miller Puckette (Puckette 1998). This object runs an FFT analysis on an incoming audio stream and outputs a floating point pitch, a report for signal onset or attack, from 1 to 3 lists (pitch and loudness of a pitch track), signal power in dB, and a list of analyzed peaks with peak index, frequency, and amplitude for each peak (Puckette 1998).

This object was used in various ways in each movement and will be discussed in more detail in the individual movement analyses. It is a flexible and powerful analysis tool, but certain concessions must be made. First, the pitch accuracy is not infallible and depends greatly on the audio source. In this case, the prepared piano sound did not lend itself to sinusoidal decomposition because of the high quantity of inharmonic partials. This did not prove too problematic since I was using pitch tracking as a general purpose input derived control source, not for something specific, such as score following (Puckete 1992). Tracking errors were not a problem because the data was still loosely related to the input sound. A slightly more problematic aspect of the fiddle object was encountered when attempting to use the analyzed peaks for additive resynthesis. Again, because I was not interested in accurate resynthesis, I was able to use the multi-peak frequency analysis in a very stylized or effected way, resulting in a musically expressive form of resynthesis. In general, a combination of signal analysis (primarily pitch, overall amplitude, attack, and multiple-peak output) and simple automated control algorithms were used to drive the real-time signal processing and control various aspects of the sound transformation.

Movement One: <Listening example #1>

The first movement of this piece is the most complex from a timbral perspective because of the playing technique—inside the piano—and the software design. This movement uses two fiddle objects for analysis; one on the incoming prepared piano sound, and one on the output of the computer processed sound. In this way, the software can know about what I am doing inside the piano and also have some awareness of what the processing algorithms are producing. The most widely used analysis is a general amplitude follower, which is sent to various parts of the patch. This is simple and effective for two primary reasons. It is very fast, because it doesn't require buffering or windowing, and it is a good overall reflection of the level of performance activity. In addition, pitch analysis is used to control the transposition of a granulator. Attack is used to randomly reset delay times, which are kept short to make the sound reasonably responsive.

The audio processing for this movement consists of two variable delays with feedback, two transposing delays with feedback, two stereo granulators, a modulated filter, two ring modulators and a reverberator. Two of the delays have modulated filters (resonant bandpass and comb) in their feedback loop so the timbre is modified each time it is passed through the loop. The send level to the delays is driven positive by the amplitude of the incoming piano sound. The delay time is randomly reset each time the fiddle object detects an attack. The output of a modulated filter is also fed into these delays. The output of the delays is fed to the patch mixer as well as to the granulators, driven by inverted amplitude. There are also two transposing delays with feedback that have their send level driven by inverted amplitude. Because of this inversion, these delays are being fed more signal when I am playing softly, whereas the other delays are fed more signal as I crescendo. The amount of transposition of the delays is also driven by amplitude. The two granulators are fed from the first two delay lines (based on inverted amplitude) and have their grain duration and transposition modulated by amplitude and pitch, respectively. There are also two simple ring modulators that multiply the sum of two of the delays by one of the remaining delays. The amount of the ring modulators that get mixed into the output is driven by a combination of the input amplitude and a measurement of the overall output level of the combined processing. There is also a resonant filter that is fed a mix of all the above-mentioned processors. The filters output is sent to the main mix as well as being fed back to the first two delays. The center frequency and overdrive of this filter is modulated by amplitude and pitch, respectively. A main reverberator is fed a full mix, except for the filter, and is then mixed to the main outputs.

The following schematic illustrates a simplified version of the audio routing and processing. All control logic is omitted.

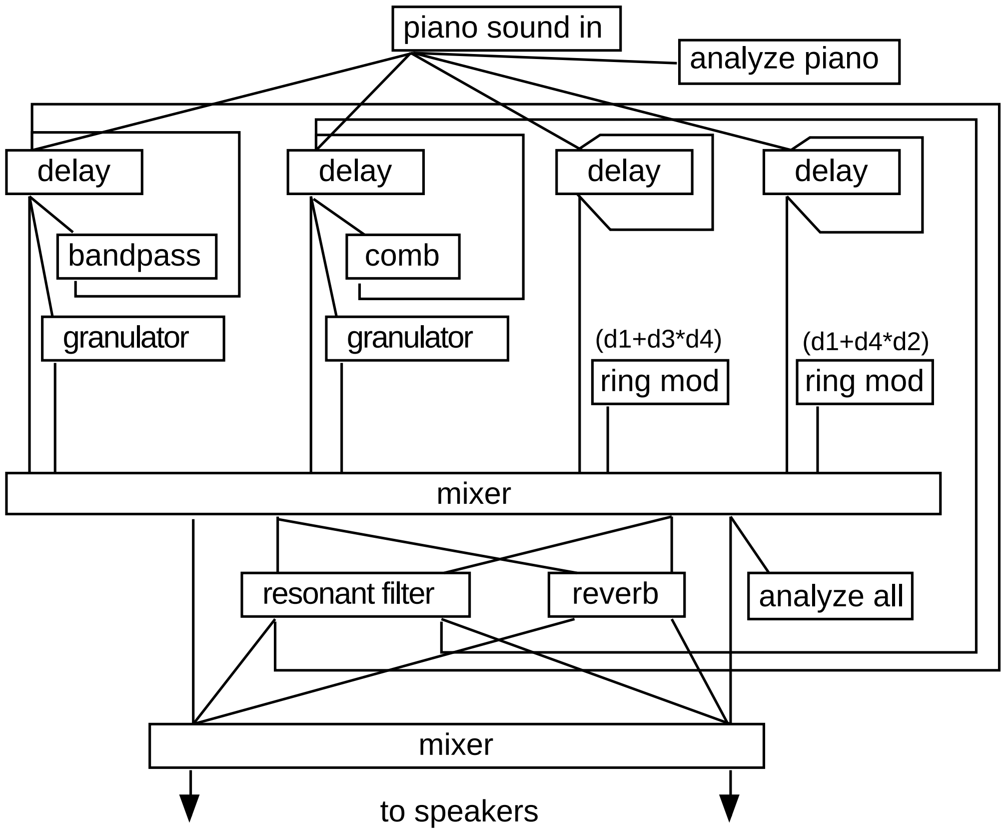

Figure 1: Simplified schematic. Movement #1

This dsp design, with multiple interconnected modules and multiple feed back paths, creates a complex and somewhat chaotic sound processing environment. The complexity of this movement is further compounded by the wide range of interior piano sounds fed into the computer. Additionally, the variety of analysis based control structures and their multiple destinations make for a very dynamic hyperinstrument.

Movement Two: <Listening example #2>

By comparison, the second movement is less chaotic and explores a smaller, more subtle sonic space. The performance approach is also different, as I am playing primarily single note gestures on the keyboard and am not playing inside the piano. The primary dsp approach is crude additive resynthesis (Dodge 1997) so again, analysis and the fiddle object play a central roll in this movement. Two fiddle objects are used: one is used just to analyze attacks and control a gate function. Whenever an attack is registered, a timed gate is opened allowing the audio input to be analyzed for peaks by a second fiddle object. This gated input is also sent to a bank of resonant filters that have their center frequencies controlled by the peak output of the second fiddle object. The additive resynthesis and resonant filtering are both feedback prone, so the gate was used to selectively break the feedback path. Amplitude analysis is used to control the attack and decay rates of the additive resynthesis.

The approach to resynthesis is fairly basic and the design intent is aesthetic rather than realistic. Two additive synths are used, one forty oscillator synth using sine waves, and one ten oscillator synth using a more complex waveform that is also transposed down one octave. Both synths derive their oscillator frequencies from the peak output of the fiddle object, which is set to output the eight loudest peaks. The amplitude of each peak is only used to tell a gate whether or not to pass on the associated frequency. This way, very low-level peaks are not resynthesised. The frequencies that pass the gate are sent to the multi-oscillator synth modules. The level control for the oscillators are driven by simple attack/decay envelopes with the attack and decay time modulated by the amplitude of the overall piano signal. The way this is mapped, louder playing will cause the resynthesized sound to attack and decay more quickly than softer playing. This simplified approach to resynthesis, applied to the difficult to analyze prepared piano sounds, results in a unique sound; it is a kind of ethereal halo around the sound that is clearly related to the source sound, but certainly not replicative.

The other main sound manipulation in this movement is a set of six resonant filters. These filters are fed the piano signal and their center frequencies are controlled by the same peak analyzing fiddle object that is used for the resynthesis. The filter Q is controlled by the piano amplitude with the Q getting narrower as the amplitude increases. The tracking errors in the peak analysis are made less noticeable by smoothing the frequencies before they are sent to the filter.

The output of the filter is also sent back to its input, creating a feedback loop that is always on the edge—and sometimes over the edge—of self-oscillation. A bit of ring modulation roughs up the sound of the resynthesis. The whole mix is passed through a reverberator before being sent to the speakers.

The following schematic illustrates a simplified version of the audio routing and processing and some of the analysis / control routing.

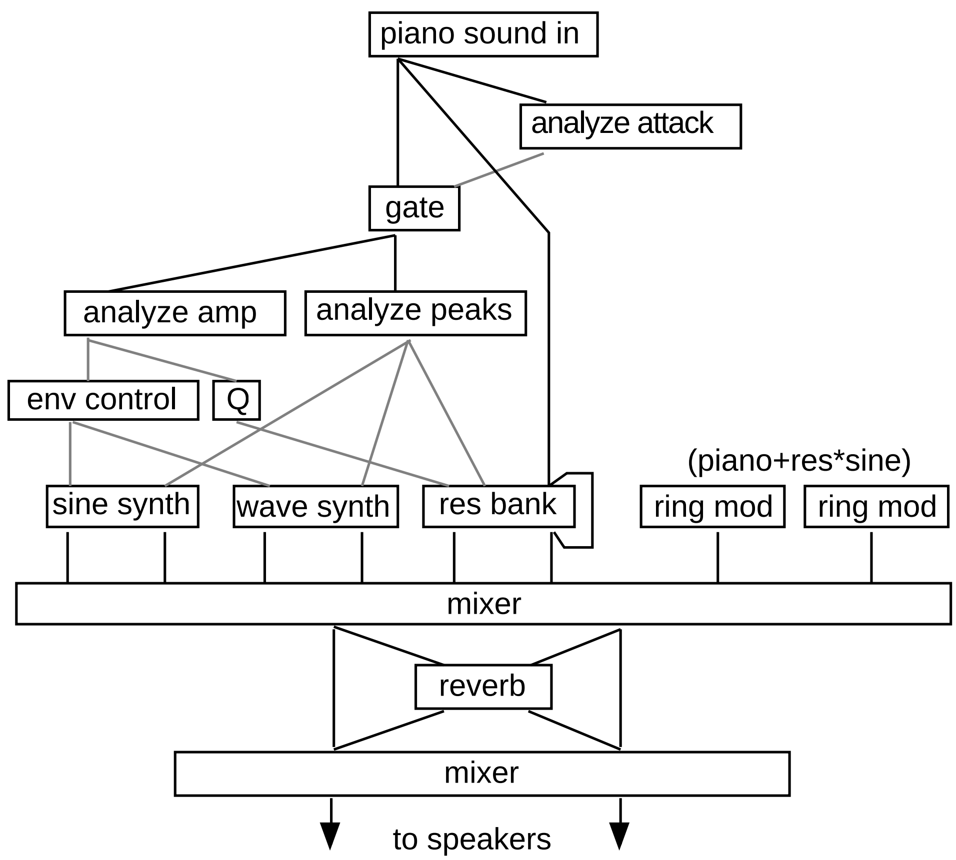

Figure 2: Simplified schematic. Movement #2. Control data is shown with gray lines.

In this movement, the combination of resynthesis, input driven resonant filtering, and sensitive dynamic control create an expressive and delicate sound. The success of this movement depends on a focussed and careful performance and requires a close rapport with the hyperinstrument.

Movement Three: <Listening example #3>

The third movement is in some ways the most direct of the three, as it relies almost exclusively on only two types of interrelated processing: modulated time compression/expansion and amplitude driven infinite reverberation. My concept for this movement centered on the ability to put sound gestures into the software and have them reappear at different times and, more importantly, at widely different playback rates, while still maintaining original pitch. I also wanted this processing to be dynamically controlled by performance gestures so that I could reintroduce material by continuing to play and freeze or capture a given sound when I stopped playing.

There are several different ways to accomplish time and/or pitch compression/expansion, including granular methods, real-time harmonizers, phase vocoding, wavelets, and linear predictive coding (Roads 1996). These techniques can be computationally intensive and/or difficult to implement in real-time performance systems. The technique implemented for this piece is a time compression/expansion algorithm using windowed granular sample playback. This implementation is very efficient and can be done in real-time by continuously writing into RAM buffers that are then read from at varying rates. This technique produces sonic artifacts, but in this application the processing byproducts are reasonably interesting and fit within the aesthetic of the piece. The basic controls implemented in this algorithm are window size, time factor, and buffer location. An elegant implementation of this technique that also includes pitch transposition was done by Dobrian (Dobrian 1999). In my implementation, both window size and time factor (i.e., playback rate where a factor of 1 equals real-time) are modulated based on the analysis of pitch and amplitude of the piano signal. To create a complex texture with multiple time domains, four time manipulators are run in parallel, reading from four separate buffers. The looping buffers are writing and overwriting the piano signal as it is played. The output levels of the four time processors are driven by input amplitude so that when I stop playing, the stretched signals fade out. The next piano attack will ramp the buffer players back up, but they may be playing from widely different points in their buffers due to variations in playback rates.

In addition to time stretching, an infinite hold reverberator is implemented and was fed a mix of the time stretched signals. The reverb decay time and output level is driven by the piano input signal amplitude.

This mapping results in progressively shorter decay and quieter reverb as I play louder, with a long decay when I stop playing. Additionally, a resonant filter, with cutoff driven by input amplitude, is inserted between the buffer playback and the reverb and output mixer.

The following schematic illustrates a simplified version of the audio routing and processing and some of the analysis / control routing.

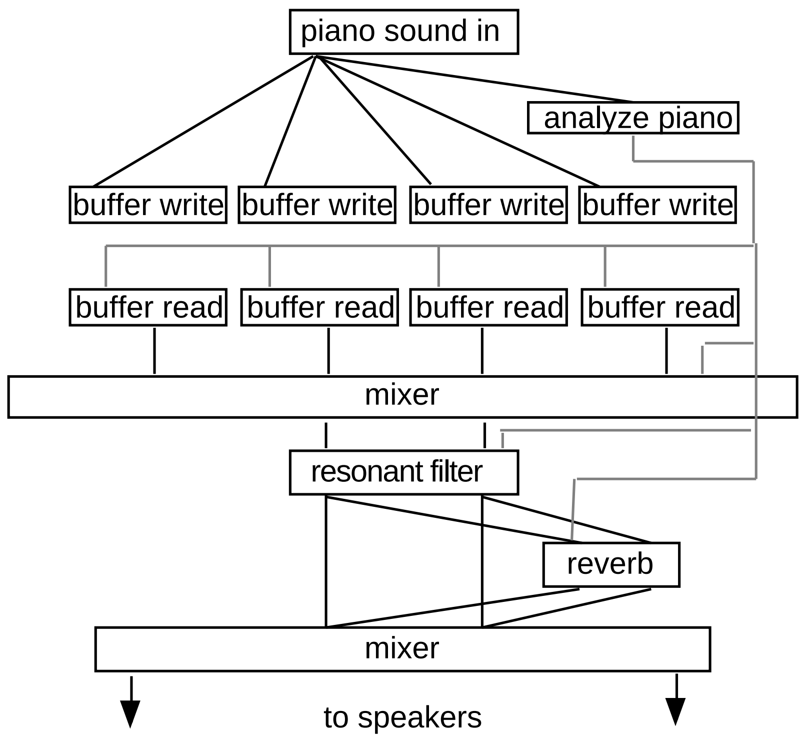

Figure 3: Simplified schematic. Movement #3. Control data is shown with gray lines.

This reverb decay combined with the time stretching creates an ongoing disjointed temporal history by combining various time lines and capturing or freezing certain moments. As is the case with the instruments for the first and second movements, this instrument requires sensitivity and careful selection of sound gestures.

Future Directions

Working on this piece, and the hyperinstruments to play it on, has set a foundation for further research and development. This research can be broadly described by the categories of complex analysis and response systems, musical and improvisational extensibility, and expanded personal listening/sounding abilities. These first two areas are intertwined and have to do with an instrument or sound environment's ability to listen and respond to unexpected inputs. This is sometimes referred to as extensibility or expressive range (Dobrian 2001). The inability to deal with unexpected input is often referred to as brittleness (Miranda 1999). Even though *Three Meditations* is based on structured improvisation, the range of input gestures is not that wide. For this piece, I am comfortable with the range of sounds available and the relationship between input gesture and output sound. The audio analysis is relatively simple but fairly responsive and expressively mapped within the context of several short movements.

In my current projects, I am trying to create sound spaces that can deal with a wide range of unknown and widely varying input gestures, and respond with expressive and interesting sonic contributions. I am exploring two main methodologies to meet these goals. First, I am working towards more complex and informative analysis techniques, creating software that goes beyond pitch, amplitude and attack detection to include analyzing timbre, noisiness, long term dynamic envelope, and event segmentation (Lippe 1994). I am exploring options for parsing, interpreting, and mapping this additional information. My hope is that the effective interpretation of this additional information, combined with musically intelligent mappings and expressive response algorithms, will result in a dynamic, pliable sound environment capable of dealing with a wide range of extended sonic improvisations.

In addition to working on better software, I am continually working on my own abilities to listening to and make sounds. One of the main challenges with performing *Three Meditations* is my ability to listen, adapt, and respond to the hyperpiano in real-time. The instrument is different every time it is set up and depending upon the performance schedule, I may not have the desired time to learn the instrument. My ability (or lack thereof) to quickly find the inner voice of the instrument, and to be focused and present, is key to the performance of this or any improvisational piece. In this spirit, I continue to listen and sound with as much attention as possible and to attempt the "Deep Listening" described and taught by Pauline Oliveros (Oliveros 1999, 2000).

References

Chadabe, J. 1997. *Electric sound: The Past and Promise of Electronic Music*. Upper Saddle River, NJ: Prentice-Hall.

Dobrian, C. 2001. Aesthetic Considerations in the Use of “Virtual” Music Instruments. Submitted for the CHI Workshop on New Interfaces for Musical Expression. http://music.arts.uci.edu/dobrian/articles.htm

Dobrian, C. 1999. Programming New Real-Time DSP Possibilities with MSP. In *Proceedings of the 2nd COST G-6 Workshop on Digital Audio Effects*.

Dodge, C. 1997. *Computer Music: Synthesis, Composition, and Performance*. New York, NY: Schirmer Books.

Duckworth, W. 1995. *Talking Music: Conversations with John Cage, Philip Glass, Laurie Anderson, and Five Generations of American Experimental Composers*. New York, NY: Schirmer Books.

Feldman, B. 1989. Liner Notes from, *Piano.* Therwil, Switzerland: Hat Hut Records.

Grove Music Online. 2002. Stephen Scott. http://www.grovemusic.com/data/articles/music/4/426/42693.xml?section=music.42693.1

Jehan, T. and B. Schoner. 2001. Audio-Driven, Spectral Analysis-Based, Perceptual Synthesis Engine. In *Proceedings of the Audio Engineering society 110th Convention 2001*. Amsterdam, Netherlands.

K and K Sound Systems. 2002. Home Page. http://www.kksound.com/welcome.html

Karis, A. 1996. *Secret Geometry: Music for Piano and Tape*. New York, NY: CRI.

Kessker, T. 2002. New Possibilities of Live-Electronic Music for the Composer Today. http://www.music.upm.edu.my/MusTech/Kessler.htm

Lehrman, P. 2002. About the Ballet Mécanique. http://www.antheil.org/balletmec.html

Lindemann, E., Starkier, M., and Dechelle, F. 1990. The IRCAM Musical Workstation: Hardware Overview and Signal Processing Features. In *Proceedings of the 1990 International Computer Music Conference*. San Francisco: International Computer Music Association.

Lippe, C. 1997. Music for Piano and Computer: A Description. In *Proceedings of the U.S.A/Japan computer music conference 1997*. Tokyo: Kunitachi Conservatory.

Lippe, C. and M. Puckett. 1994. Getting the Acoustic Parameters from a Live Performance. In *Proceedings, 3rd International Conference for Music Perception and Cognition.* Liege.

Lippe, C. and Puckette, M. 1991. Musical Performance Using the IRCAM Workstation. In *Proceedings of the 1991 International Computer Music Conference*. San Francisco: International Computer Music Association.

Machover, T. 1992-2002. Hyperinstruments. http://brainop.media.mit.edu/Archive/Hyperinstruments/

Miranda, E., ed. 1999. *Readings in Music and Artificial Intelligence*. The Netherlands: Harwood Academic Publishers.

Oliveros, P. 2000. Quantum Listening: From Practice to Theory (To Practice Practice). Plenum Address, International Congress on Culture and Humanity in the New Millennium: The Future of Human Values Chinese University, Hong Kong. http://www.deeplistening.org/pauline/writings/quantum_listening.html

Oliveros, P. 1999. Quantum Improvisation: The Cybernetic Presence. Keynote address Improvisation Across Borders Conference. http://www.deeplistening.org/pauline/writings/quantum.html

Pritchett, J. 1993. *The Music of John Cage*. New York, NY: Cambridge University Press.

Puckette, M. et al. 1998. Real-time audio analysis tools for Pd and MSP. In *Proceedings of the 1998 International Computer Music Conference*. San Francisco: International Computer Music Association.

Puckette, M. and Cort Lippe. 1992. Score Following in Practice. In *Proceedings of the 1992 International Computer Music Conference*. San Francisco, CA: Computer Music Association.

Puckette, M. 1991. Combining Event and Signal Processing in the Max Graphical Programming Environment. *Computer Music Journal* 15(3).

Puckette, M. 1988. The Patcher. *In Proceedings of the 1988 International Computer Music Conference*. San Francisco: International Computer Music Association.

Roads, C. et al. 1996. *The Computer Music Tutorial*. Cambridge, MA: MIT Press.

Tamm, E. 1989. Tamm on Eno. http://www.geocities.com/pitkin_family/tammeno.html

Trueman, D. et al. 2000. Alternative Voices for Electronic Sound: Spherical Speakers and Sensor-Speaker Arrays (SenSAs). . In *Proceedings of the 2000 International Computer Music Conference*. San Francisco: International Computer Music Association.

Zicarelli, D. 2002. Max/Msp. http://www.cycling74.com/index.html