The design and application of electroacoustic instruments in composing for interactive performance works.

Synthèse: DMUS Composition

Zack Settel, été 1999
SETZ22075706

Introduction

The use of realtime technology in music composition and performance has long captured the imagination of composers. The application of analog electronics, and now, realtime computing technologies to music composition, performing, and distribution has occasionally provided significant new opportunities for musical expression, around which, an accompanying aesthetic has emerged. This text will focus on the author's approach to music composition and underlying uses of "interactive" technologies, with attention to technical, practical and aesthetic issues.

Over the years I have drawn on interactive realtime technologies as source for new approaches to creating and performing music. Often, this has resulted in alternative ways to deal with conventional problems in creating music (eg. form or variation in composition). My principal interests and application of these technologies have centered on electroacoustic instrument design [Lippe,1991], live performance mediation [Rowe, 1993], and compositional aids. [Xenakis, 1971].

1.0 Electroacoustic Instrument Design

It is important to distinguish my particular concept of "instrument" against the many others that have emerged with advances in technology and their application music, particularly in popular music (eg. Techno, Rave and the turntable). The conception of the "instrument" discussed in this text is largely based on the traditional instrument: you have to play it in order for it to produce sound, good playing should sound better than bad playing, and the playing of the instrument should be intuitive. Thus, the instrument should be causal, allow for extremely fine and direct control of the sound it produces, and employ coherent and appropriate control interfaces and control/audio mappings (this latter consideration determines the instrument's overall "transparency", from the performer's point of view). This way of regarding the instrument can be found in the work of Morton Subotnick and Todd Machover,

The process of defining and developing instruments for composition, as in the case of Cage's works for prepared pianos, or Parch's works in general, is of great importance to my work. Instrument choice and quality is essential to the compositional process and resulting piece. Using Electroacoustic techniques developed over the years, I have realized several pieces which include the use of what I refer to as "electronically prepared instruments" (as compared to mechanically prepared pianos for example). These typically involve traditional (referred to in this text as "basis instruments"), instruments whose sounds are modified in realtime via electronic means, in order to provide instrumentalists with access to additional timbral ranges of expression, or performance possibilities (eg. self accompaniment)—nonetheless based on their instrumental technique.

Of particular interest is the fact that electroacoustic instrument design (the definition and implementation of such instruments) involves a constant dialog between the musical material of the composition and the range of musical expression made possible through the design of the instrument; it is an approach analogous to composing a piece for muted horn while designing the mute at the same time. This dual approach to producing a composition, where the "instrument builder" and the composer are one in the same, fosters serendipity during composition; the choice of musical material can be altered by the possibilities of the instrument, or conversely, the design of the instrument can be altered to meet the demands of the musical material. This process has proven to be quite effective as a compositional tool or "discipline".

The following parts of this section describe my approach to the design and construction of electroacoustic instruments, and their application in music composition.

1.1 Electroacoustic Instruments and their structure: What's under the hood?

Based on compositional needs (including improvisation), I have concentrated on the use of these acoustic-based instruments to extend both the timbral range, and the range of potential musical material the instrument can play. To illustrate this idea, a simple example of an instrument incorporating both these extensions can be built around a monophonic instrument of limited spectral diversity. Using a flute with electronics, the spectrum of the instrument can be enriched with inharmonic spectra, while its tesitura can be extended; additional "voices" of the instrument can be generated to provide polyphony and/or the possibility to sustain notes beyond the duration of a breath. Compositionally, the resulting instrument is quite exciting to write for: it is still approached as a flute (as is Cage's prepared piano approached as a piano), however, the resulting instrument can play pitches and colors outside its acoustic range, and is capable of self-accompaniment. The "solo" form can be greatly expanded. There three main components involved in the definition of an electro-acoustic instrument are discussed at length below.

1.2 Controlling the instrument

Acoustic instruments are rarely considered in the sense of the sound generator and controller paradigm--probably since both sound generator and "controller" are essentially the same thing (consider the function of a trumpet player's lips). Nonetheless, it can be useful to view them in this sense in order to appreciate the role of control in the production of timbre. One good example of this is when a breath controller is used to modulate the amplitude of the tone generator it is connected to: a player with a nimble tongue can actually modify the spectral content via amplitude modulation by changing the wind pressure (and resulting amplitude) at rates greater than 20 times a second: the timbre is directly colored by the nature of the control signal. The better the control signal, the greater the degree (extent and precision) to which the resulting timbre can be controlled. In general, a good control signal needs to provide enough time and dynamic resolution to capture fast and subtle changes in the control source (a physical controller, or audio signal). The control source of greatest interest to me is the audio signal an acoustic instrument produces. It is a high-resolution signal, rich in performance information, in which many musical features (and physical measures) can be recognized and used for control purposes. The process of obtaining control signals (discrete or continuous) from an audio signal in realtime is described below.

1.2.1 Extracting control signals from audio signals.

Various musical events or features (such as those notated in scores: eg. dynamics, pitch, duration etc.) can be directly measured or extracted from the audio signal of an acoustic instrument. Continuous signal features such as amplitude, brightness, harmonicity, noise levels, and fundamental frequency, can be translated directly into control signals. Discrete events (segmentation of the audio stream) can be obtained by performing operations on these control signals; for example, a "note-on/off" event, such as the kind that typical MIDI pitch followers output, is produced making use of the frequency and amplitude measures of the input signal. Described below are the signal analysis techniques used by the author to obtain control signals from the audio input (acoustic sound) of the electroacoustic instrument.

Envelope follower:
Measure the amplitude of the input signal, outputting its current value at the control rate, or as an audio signal (when taken before the "snapshot" stage shown below). The envelope signal reports the energy level of the audio input, thus providing a measure of how loud the acoustic instrument is being played. The slope (derivative) of this signal indicates how its intensity is changing over time. Sudden changes in loudness, such as non-slurred note changes, or subito crescendi are easily recognized and can be used to segment the continuous audio input into desecrate events. Silences (rests) are also easily detected and converted into events which can be used for input by higher-level event detectors for, say, note duration.

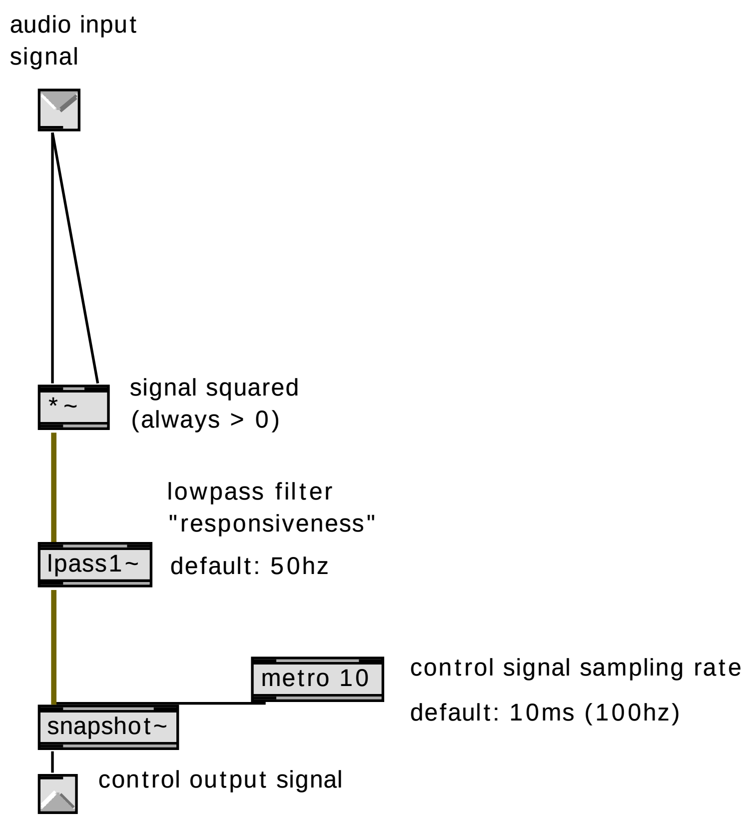

Band-limited regions of the input signal's spectral energy can provide useful signals, which in turn, can be used by higher-level objects that report on the center of spectral mass (energy). Using this approach for example, it was possible to track the position of a kettle drum's pedal.

Pitch following:

The pitch follower provides two types of information about the input signal. Fundamental frequency and note duration. Provided that the input signal is pitched, the follower provides a continuos signal representing its guess about the input signal's fundamental frequency. The note duration information produced by the follower is more complicated, taking into account the input signal's amplitude in order to determine when a note of a given pitch begins and ends. Since the segmentation of the audio input into note events can be done better (with greater flexibility) by higher-level objects that use input from envelope, pitch and other followers, only the instantaneous pitch information is used. As with the envelope follower, the slope of the instantaneous pitch signal can provide a measure of the change to the input signal's pitch, allowing for the detection of glissandi, and note boundaries.

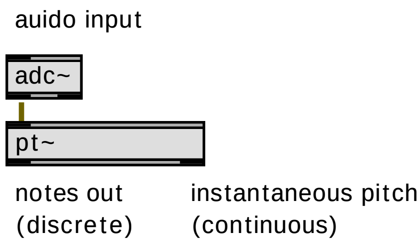

Noise detection (zero crossing density)

The noise detector counts the number of times that the input signal crosses the value of zero as it changes, providing information about the quality of the input signal; when analyzing the voice for example, sibilant sounds can be detected, due to the high density of zero crossings.

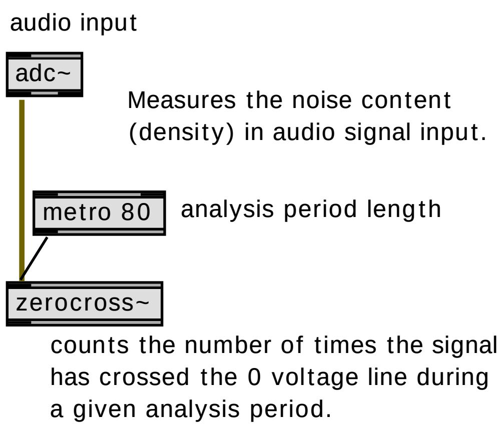

Using the above analysis techniques separately or jointly, signals and events can be generated from the audio input of an electroacoustic instrument; these are, in turn, used to control the signal processing/generation portions of the electroacoustic instrument. The following signals and events, obtained through analysis of the input signal, are most commonly utilized:

intensity (band-limited and overall)
intensity change (spike detection)
rests
pitch
pitch change (glissando detection)
noise density
center of spectral mass (spectral centroid)

1.3 DSP: signal processing and generation

An important part of the electro-acoustic instrument is the digital signal generation and processing module, which provides the sound that constitutes the electronic compliment to the Electro-acoustic instrument. Signal processing and signal generation techniques are used to provide timbral extensions to the basis acoustic instrument. Since the control parameters of these processors and generators are often driven by the control signals and events, produced by analysis of the acoustic instrument's signal, a tight coupling between the changes to the timbre of the acoustic instrument and the resulting electronic accompaniment is made possible.

Below are several DSP techniques, which are often incorporated (usually in combination).

1.3.1 Conventional time-domain techniques:

signal processing

Reverb, Harmonizers, Tapped Delays, Spatialization, Ring Modulators and Frequency Shifters.

signal generation

Sampling, Frequency Modulation, Resonant Filters

1.3.2 Frequency-domain techniques:

While the time-domain techniques are standard and well known, several of the of the frequency-domain techniques listed below are novel. An in-depth description of several of these frequency-domain techniques, implemented by the author [Settel,Lippe 1998], is provided.

signal processing

Hi-resolution Filtering, Spatialization, Dynamics Processing.

signal generation

Additive Synthesis, Cross Synthesis, Phase Vocoding

High-resolution filtering

Highly detailed time varying spectral envelopes can be produced and controlled by relatively simple means. A look-up table can be used to describe a spectral envelope in the implementation of a graphic EQ of up to 512 bands. The spectrum of the input signal is convolved, point by point, with the data in the look-up table, producing a filtered signal. Because we are able to alter the spectral envelope in real time at the control rate (up to 1kHz), we may modify our spectral envelope graphically or algorithmicaly, hearing the results immediately.

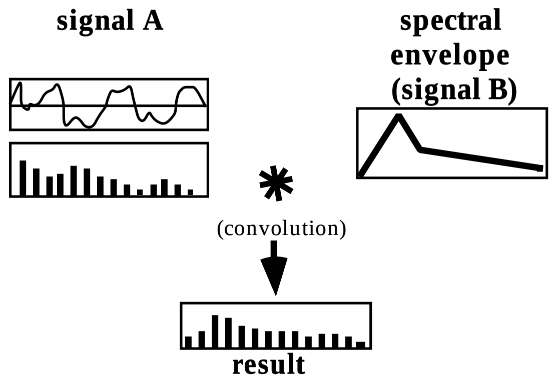

filtering with a user-specified spectral envelope

Using a noise source as the input signal, it is also possible to do *subtractive synthesis* efficiently.

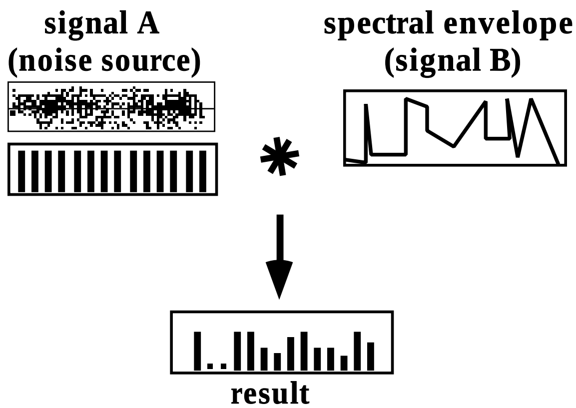

subtractive synthesis

Low dimensional control of complex spectral envelopes

The spectral envelope used in the above filtering application can also be obtained through signal analysis, in which case a second input signal, signal B, is needed. Signal B is analyzed for its spectral envelope, or amplitude spectrum, that describes how signal A will be filtered. Obtaining a spectral envelope from an audio signal offers several interesting possibilities: spectral envelopes can be "collected" instead of being specified, and can change at a very fast rate (audio rate), offering a powerful method of dynamic filtering. Also, audio signals produced by standard signal processing modules such as a frequency modulation (FM) pair (one oscillator modulating the frequency of another) are of particular interest because they can produce rich, easily modified, smoothly and nonlinearly varying spectra [Chowning 1973] which can yield complex time varying spectral envelopes. These spectral envelopes can be modified using only 3 (FM) parameters: carrier frequency, modulator frequency, and modulation index. Likewise, other standard signal processing modules such as an amplitude modulation (AM) signal generator, an additive synthesis instrument, or a band-pass filter bank offer rich varying spectral information using relatively simple means with few control parameters. One of the advantages of using standard modules is that electronic musicians are familiar with them, and have a certain degree of control and understanding of their spectra.

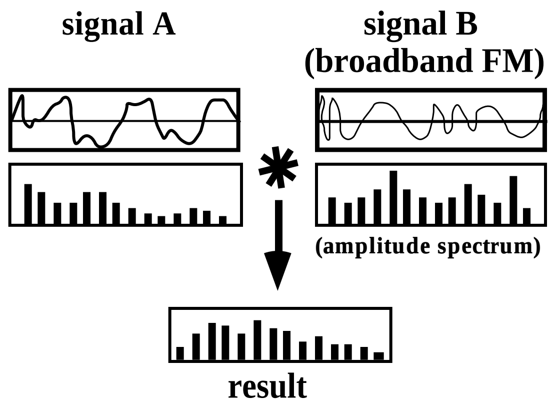

dynamic filtering of signal A using the spectral envelope of signal B

Cross synthesis

In this application two input signals are required: signal A's spectrum is convolved with the amplitude spectrum of signal B. Thus, the pitch/phase information of signal A and the time varying spectral envelope of signal B are combined to form the output signal. Favorable results are produced when Signal A has a relatively constant energy level and broadband spectrum, and when signal B has a well defined time varying spectral envelope. For example, when wishing to transform spoken or sung text, I assign the text material to signal B while specifying a pulse train, noise source or some other constant-energy broadband signal to signal A. Since the frequency information (pitch, harmonicity, noise content, etc.) of signal A is retained in the output, unusual effects can be produced when frequency related changes occur in signal A. In the following example of a *vocoder*, text can be decoupled from the speaker or singer's "voice quality", allowing one to modify attributes of the voice such as noise content, inharmonicity, and inflection, independently of the text material.

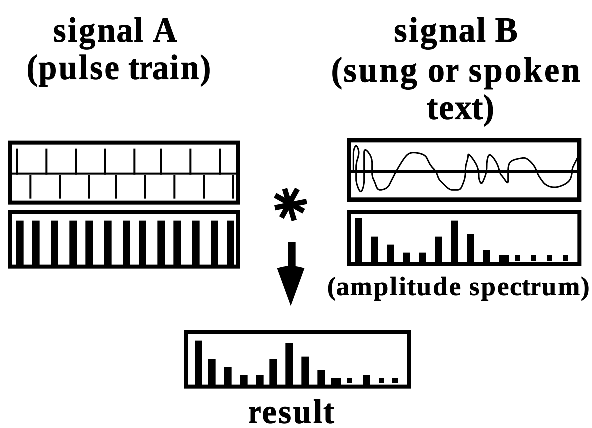

cross synthesis

Mapping qualities of one signal to another

A simple FM pair may be used to provide an easily controlled, constant-energy broadband spectrum for use in cross synthesis as signal A. Musically, I have found that in some cases, the relationship between signal A and signal B can become much more unified if certain parameters of signal B are used to control signal A. In other words, real-time continuous control parameters can be derived from signal B and used to control signal A. For example, the pitch of signal B can be tracked and applied to signal A

(FM) to control the two oscillators' frequencies. Envelope following of signal B can yield expressive information which can be used to control the intensity of frequency modulation (FM index) of signal A . In experiments incorporating the above, a mezzo soprano's voice was assigned to signal A, while her pitch and intensity were mapped onto signal B (FM), producing striking results akin to harmonization and frequency shifting.

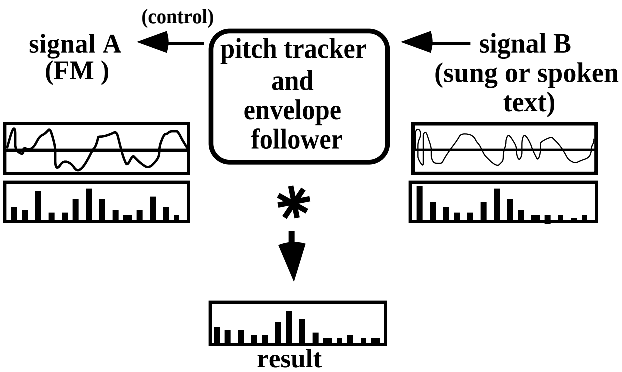

cross synthesis using signals with linked parameters

Finally, it should be noted that interesting transformations can be produced by simply convolving signal A's spectrum with signal B's spectrum. In this case, the phase (frequency) and spectral envelope information from each signal figures in the output signal. Transformations of broadband sounds, akin to, but more pronounced than flanging, can be produced when convolved with the signal of a high index, inharmonically tuned FM pair, whose frequency parameters are controlled by the pitch of the first signal.

Frequency-dependent spatialization

In the spectral domain, the phases of a given signal's frequency components can be independently rotated in order to change the component's energy distribution in the real and imaginary part of the output signal. Since the real and imaginary parts of the IFFT's output can be assigned to separate output channels, which are in turn connected to different loud-speakers, it is possible to control a given frequency's energy level in each loud-speaker using phase rotation. The user interface of this application permits users to graphically or algorithmically specify the "panning" (phase offset) for up to 512 frequency components.

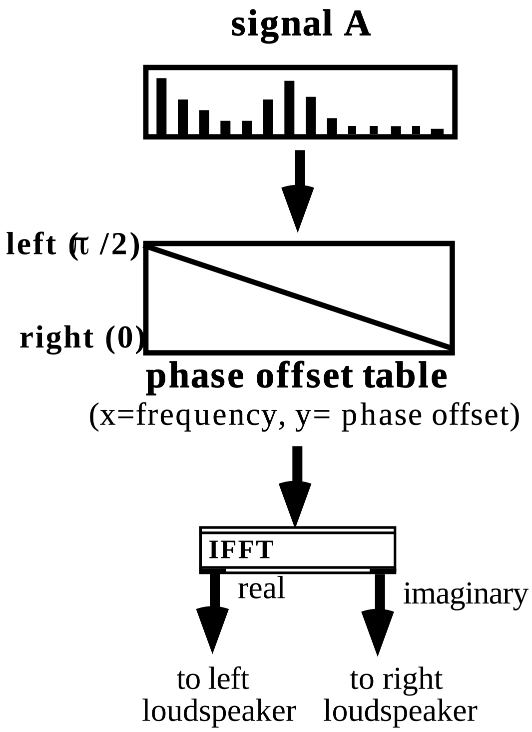

Band-limited spatialization

Band-limited energy dependent noise-gate

In the spectral domain, the energy of a given signal's frequency components can be independently modified. Our noise reduction algorithm is based on a technique [Moorer & Berger, 1984] that allows independent amplitude gating threshold levels to be specified for up to 512 frequency band-limited regions of a given signal. With a user-defined transfer function, the energy in a given frequency range can be altered based on its intensity and a user-specified threshold level. This technique, besides being potentially useful for noise reduction, can be exaggerated in order to create unusual spectral transformations of input signals, resembling extreme chorusing effects. Using non-linear transfer functions, it is possible to modify the relative intensities of the input's frequency components, allowing for example, *masked or less important components to be emphasized and brought to the aural foreground.

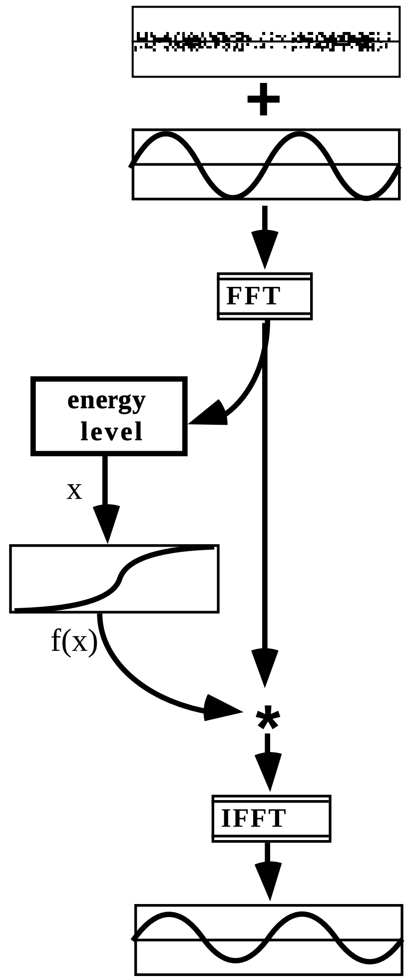

Gating low amplitude noise

Band-limited frequency dependent noise-gate

Similar to the noise-gate described above, this module functions independently of gain; its output depends on the stability of the frequency components of the input signal. Using a technique borrowed from phase-vocoding [Gordon, 1987], time-varying frequency differences of components in a given band-limited region of the spectrum are used to determine the stability of those components. Pitched components in the input signal tend to be stable and can thus be independently boosted or attenuated.

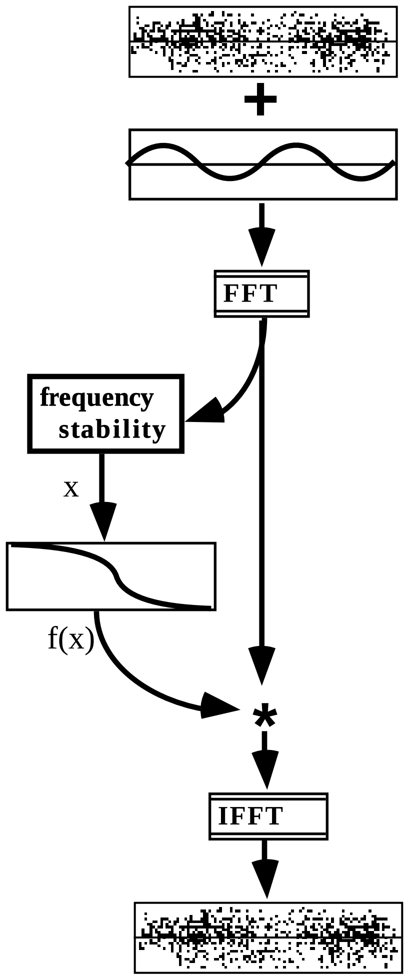

1.4 Mapping control parameters to DSP

Perhaps the single most important stage in defining an Electro-acoustic instrument, from the point of view of the performer, is the mapping stage. It is at this stage where the "feel" or "tightness" of the instrument (coherence between playing actions and resulting sounds) is specified.

1.4.1 Direct and complex mappings

Direct mappings: Experience has shown that "natural" mappings of control to sound processing or generation provide the performer with the greatest sense of coherence, or "oneness", in terms of the acoustic instrument and its electronic counterpart. For example, following direct mappings of playing features to DSP parameters have proven to be quite effective:

Performer's dynamic level ----> DSP gain
Performer's intonation ----> DSP tuning
Performer's brightness ----> DSP brightness
Performer's position ----> DSP panning.
Performer's register ----> DSP pitch
Performer's articulation ----> DSP event choice
Performer's attack ----> DSP event trigger
Performer's rest ----> DSP event termination

In order to create an instrument with a greater range of potential timbral expression, complex mappings of the above can be used, and usually are. For example, DSP brightness can be a function of both brightness and dynamic level, while panning is a function of intonation, and DSP pitch, a function of the sharpness of the player's attack.

1.4.2 Other mappings

The above mappings tend to work very well when timbral fusion between acoustic and electronics is desired. However, when the electronics serve more in the sense of an accompaniment, other mappings can be quite effective. For example, in **sound example 1.,** a monophonic instrument is capable of playing polyphonically. The number of harmonized copies a given note played by the acoustic instrument depends on the sharpness of the attack of that note. A sharp attack will produce up to four copies, while a modest attack will only produce one. Additionally, the degree of transposition of each harmonized copy of that note is changed each time a the performer plays a new note.

1.4.3 Control and audio rate mappings

One extremely valuable aspect of the MSP development environment is its flexibility for control. DSP modules can be controlled at the audio rate. As mentioned above, audio-rate control of the gain of a DSP module can contribute to the spectral content of the module in question, via amplitude modulation. When a performer's dynamic level is applied, as an audio signal, to control the gain of a DSP module, much of the player's fine and subtle performance information can be retained and "projected" on the resulting electronic sound. In **sound example 2.** , a sampler is controlled by the performer, whose sharp attacks are used to trigger note-ons. The playback samples are normalized, and have no predetermined envelope functions applied to them. Rather, the note-ons are gated by the intensity of the performer's playing. Thus, it is possible for the performer to control, to a very fine degree, the intensity and articulation of the sampler's currently playing note that his sharp attack triggered. Because the control is done at the audio rate, flutter tongue and other rapid performer articulations can be used to control the DSP.

1.4.4 Controlling Frequency-domain DSP: Audio-rate control of FFT

The high precision and new possibilities that frequency-domain DSP processing provide are quite attractive, however controlling these DSP modules can be a nightmare, due to the high number of parameters that these processes typically feature. For example, a

512-band vocoder has 512 gain parameters, one for each "channel", that must be updated at least, say, 50 times a second. As mentioned above, the Electro-acoustic instrument relies on the tight coupling of performer-derived control signals/events to the resulting electronics. Since the analysis of the performers input produces relatively few parameters, direct mapping strategies can not apply to the control of frequency-domain DSP. The author has developed a strategy that addresses this problem; it is described in detail below.

First a little bit about the DSP environment. The Max/MSP environment has two run-time schedulers: the Max "control" scheduler, which is timed on the basis of milliseconds, and the MSP "signal" scheduler, which is timed at the audio sampling rate [Puckette, 1991]. In FFT-based processing applications, where changes to the resulting spectrum are infrequent, MSP's control objects may be used to provide control parameters for the processing. This is both precise and economical, but limited in terms of spectral rate of change. For example, this implementation of FFT-based narrow-band "graphic" equalization uses a lookup table describing a filtering function, which is updated at the control rate; a given (frequency-domain) input signal is convolved with this function and spectral processing is accomplished. However, updating lookup tables at the control rate has bandwidth limitations. The rapidity with which a lookup table can be altered is limited by the bandwidth of the control system, giving the filtering certain static characteristics. Using 512 sliders to control individual FFT bins, drawing a filter shape for a lookup table with the mouse, or changing the lookup table data algorithmically provides only limited time-varying control of the filter shape. In addition, the amount of control data represented in a lookup table is large and cumbersome. Significant and continuous modification of a spectrum, as in the case of a sweeping band-pass filter, is not possible using MSP's control objects, since they can not keep up with the task of providing 1024 parameter changes at the FFT frame rate of 43 times a second (at the audio sampling rate of 44,100 samples per second).

Keeping in mind that the input signal to be modified is FFT data, a more dynamic approach to filtering is to update lookup tables containing a filter function at the signal rate (the audio sampling rate). The term "Spectral Processing Function" (SPF) will be used frequently in this text and refers to a lookup table-based function (actually a signal), whose length is that of the FFT. For each window of input signal (FFT data) we receive in real-time, we generate a corresponding SPF with which the input may be convolved. Dynamically, the SPF can describe a particular sequence of forms (or spectral envelopes), which determine the time-varying intensity of spectral processing by frequency component via convolution; a "form" describes the action of the spectral processor. Thus, this approach to dynamic processing spectra focuses on efficiently generating forms with a potentially high degree of detail or complexity, whose descriptions are simple (generated using few parameters, for low dimensional control) and intuitive [Settel,Lippe 1998].

I have explored two methods for generating SPFs: the first involves the generation of forms using the spectral envelope of a signal (via FFT), while the second makes use of table-based waveform generation techniques. In each case, low-dimensional parametric control of complex forms is achieved. In both cases, a time-varying SPF is generated for, and convolved with, each window of input signal to perform filtering or band-limited panning, as shown below:

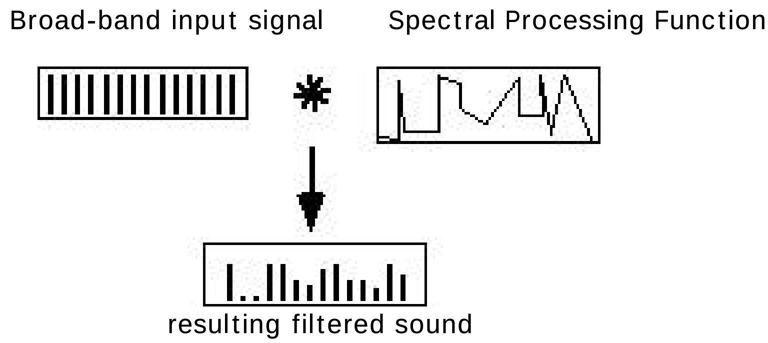

a signal is filtered (via convolution) using a Spectral Processing Function (SPF)

<u>Generating complex forms via Spectral Generation</u>

When hunting around for a suitable method for generating complex "forms" (or spectral envelopes) using simple techniques, the classic FM pair algorithm (generating complex spectra using frequency modulation), with its simple structure and inherent potential for generating highly detailed and complex spectra, leaps to mind [Settel,Lippe 1994,1995]. The SPF is simply the spectral envelope of the signal generated via FM , it provides a rich source of possible forms, whose shapes and complexity are determined using the FM algorithm's few parameters (carrier FQ, carrier:modulator ratio, modulation index). The intuitive mapping of the FM parameter values to the resulting spectral form make this method extremely easy to use. In this implementation, a particular set of static or time-varying parameters is specified for the FM algorithm; the parameters determine the shape of the resulting SPF. With each window of input signal to be processed, the algorithm's output is transformed via an FFT into a SPF (spectral envelope), which is then applied to the corresponding window of input signal via convolution. Thus, the maximum rate of change for the SPF (signal processing parameters) is given by the frame rate (function of window length) of the FFTs; a 1024 point FFT, with a sampling rate of 44.1 khz, would translate to a spectral update rate of approximately 43hz.

Needless to say, other sources for generating complex spectra, such as amplitude modulation, additive synthesis, or waveshaping, may be used in the implementation above. However, the clear advantage of using FM lies in the simple control of the highly complex spectra it offers. For example:

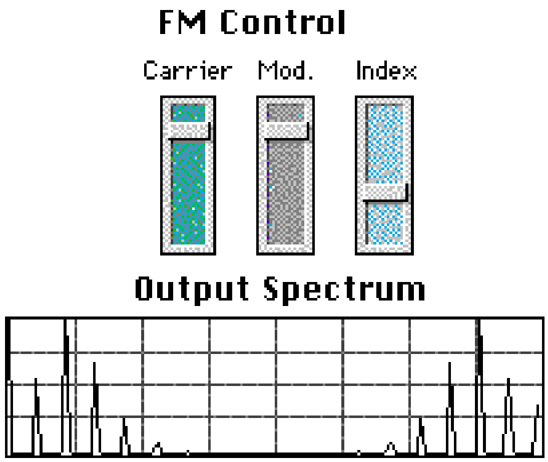

a complex spectral form depending on only three parameters

Generating complex forms using waveform generators

The use of wave tables and basic table lookup operations provides a general and flexible approach to the well-known waveform generation and synthesis techniques I use to generate the Spectral Processing Functions (SPF). Techniques such as FM, AM, waveshaping, phase modulation, pulse-width modulation, all have the potential to provide complex, evolving waveforms, which can be used as SPFs to provide a high level of flexibility and detail for spectral processing techniques such as filtering or spectral panning. The use of table lookup operations such as inversion, scaling, offsetting, wrapping, and nonlinear distortion (waveshaping), provides powerful means for dynamically modifying these SPFs. For example, nonlinear indexing of lookup tables can be employed to provide dynamic control for a constant-Q bandpass filter, where control on a nonlinear frequency scale is required. Most important, the parameters of these waveform-based techniques and are few, familiar and easy to understand.

Let's now focus on the lookup table operations underlying these techniques for the generation of a Spectral Processing Function. As discussed above, the FFT used in the implementation provides us with a phasor that is used as an index for table lookup. Note that the index is mirrored around the sampling rate/2, following from the symmetrical (real) spectrum output by the FFT.

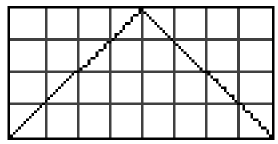

Phasor from the FFT's output

The following sequence of operations is performed on the phasor; note how its shape (form) is changed by each operation.

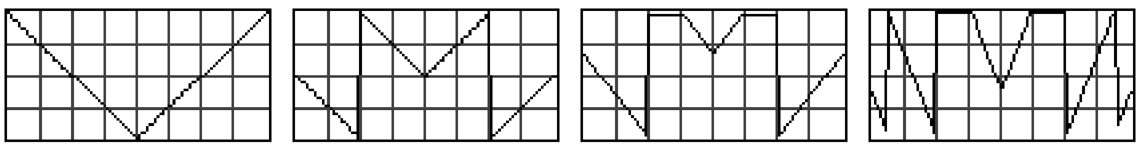

1. Inversion -1 | 2. Phase Rotation 50% | 3. Scale Y (w/ clip) 1.36 | 4. Period Multiply 2

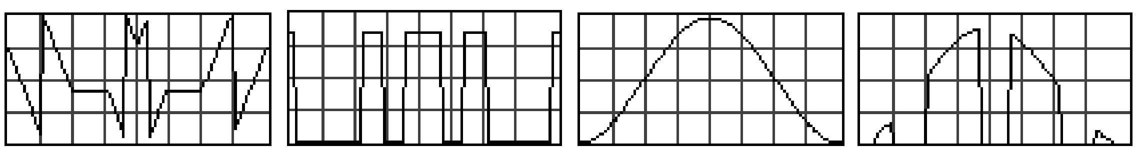

5. Phase Rotation 40% | 6. Lookup Pulse wave duty cycle = 50% | 7. Amplitude Mod. | Resulting SPF

lookup table operations

Note that steps 3, 6, and 7 entail loss of information; their placement in the sequence of operations can not be arbitrary.

While the above operations and parameter settings will produce a static SPF, it is important to note that the "phase rotation" operation also accepts a frequency parameter (LFO); a non-zero value will cause the phase rotation value to change constantly. Thus, the resulting SPF will vary periodically due to the constant change of its phase rotation parameter. This "LFO-based" technique provides control for periodic spectral modulation, and is useful in implementations such as swept band-pass filters (or spectral panners), comb filters, or "phase shifters".

To generate a more complex SPF, I start the same phasor above, then apply the following operations:

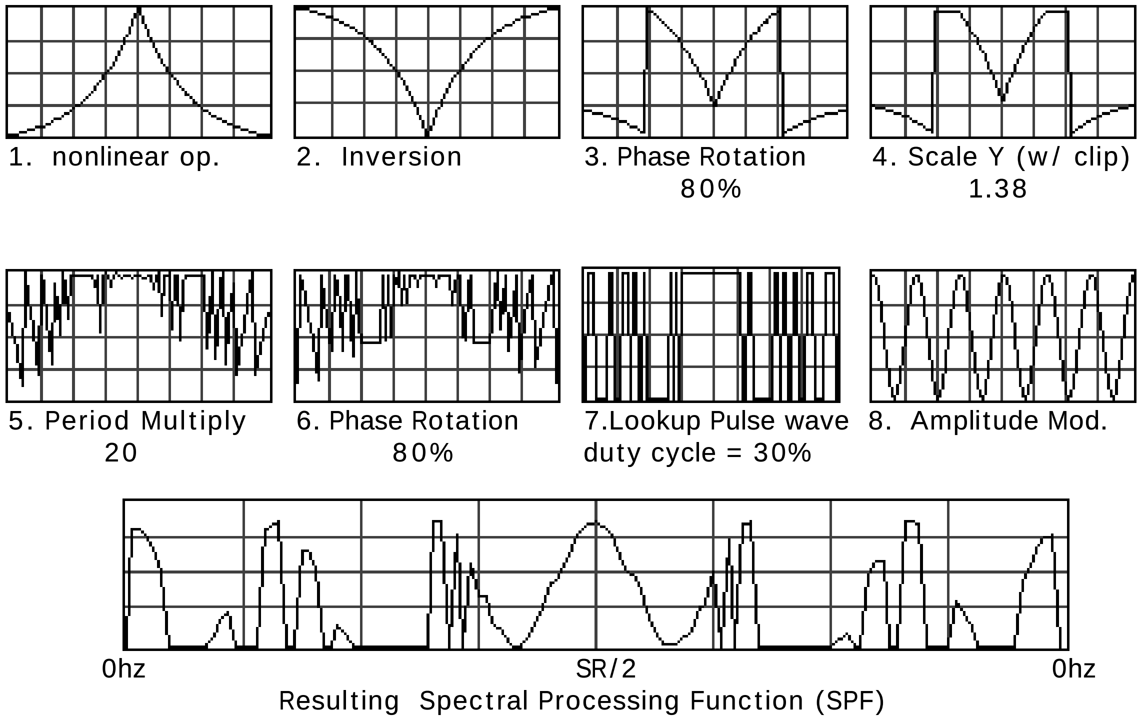

generating a complex SPF

Applying the techniques

The above techniques for dynamic spectral processing lend particularly well to applications such as high-resolution dynamic filtering, and spectral panning. The author has implemented such applications, and has experimented with spectral resolution, making use of 1024 to 4096 point FFTs, twice or four times overlapped. Given the implicit tradeoff between time and frequency resolution, and considerations of computational cost, the use of two overlapped 1024 point FFTs is preferable for real-time performance situations when audio signal latency (delay from input to output) is an issue. However, the author found that a longer FFT of 4096 points is generally preferable, despite the relatively low spectral update rate of 10hz (10 FFT frames per second), since a long window allows for a (correspondingly long) highly detailed spectral processing function (SPF). For example, the "period multiply" operation shown above requires a long SPF when a greater number of periods is specified.

In the case of high resolution filtering applications, the SPF specifies the degree of attenuation to be applied to each component (band-limited region or bin) of the input spectrum (*see sound example 8).* The input signal's spectrum is convolved with the SPF and filtering is accomplished. The nature of the filter is determined by the form of the SPF. However, in the case of an application for spectral panning (band-limited panning), the SPF is used to specify the degree of phase rotation (via convolution with a complex sinusoid) to be applied to the components of the input spectrum. The operation results in a change to a given component's energy distribution in its real and imaginary parts. The real and imaginary outputs of the IFFT are mapped to a corresponding stereo output. Dynamic spectral panning occurs when the phase of the input signal's spectral components is rotated by a changing amount. *see sound example 9*

The SPF is applied in each application as shown below:

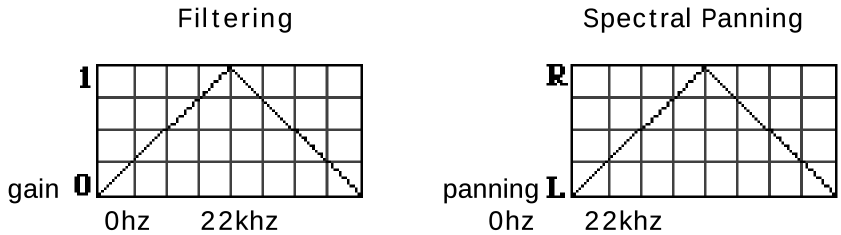

applying SPFs for filtering and spectral panning

The filter above defines a lowpass. The spectral panning distributes (pans) the energy of the input signal's components continuously from left to right, based on the component's frequency; for example: components near 0hz are panned Left, components near 12khz are panned L and R , and components near 22khz are panned right.

2.0 Live performance mediation

Finally, the Electro-acoustic instrument requires some degree of higher-level control. Up until now, we have been mainly looking at the Electro-acoustic instrument from an instrumental point of view, and have looked closely and its inner workings in this sense. In addition to timbre production, the Electro-acoustic instrument must also provide some degree of meta-control for the player, who may wish to invoke major changes to the Electro-acoustic instrument configuration. The technique of scorefollowing is often used to address this problem, and can be applied either loosely for improvisation, or rigorously for through-composed works.

During the late1980's and early-mid 90's, this technique was developed and greatly relied upon by IRCAM for many of the pieces produced there. While the results were sometimes quite remarkable, a general tendency of composers to "cater to the follower" emerged. This typically meant writing material that the follower could more easily recognize, such as unique pitches, repeated notes, or higher note values. Despite this precaution, in every IRCAM concert performance, a human was always required to follow the score follower, to correct it when it got lost.

In my own work, I have experimented with score following, and have found the follower's input constraints for robustness to be unacceptable--particularly in improvisation, when unintentional "input matching" can occur, causing important accidental changes to the instrument configuration. Thus, all meta-level control of the Electro-acoustic instrument (for changing among preset configurations, for example) is done via external controllers, using pedals and buttons (the fewer, the better).

Conclusion

The use of electroacoustic instruments in interactive performance opens up many possibilities for the composer in terms of compositional forms, writing for traditional instruments, and modes of expression in both notated and improvised music. Currently, work in signal analysis, signal processing and generation is starting to converge thanks to advances in real-time frequency domain processing. With further advances in computation and DSP/analysis technology, one can expect to see an increasingly rich and expressive evolution of electroacoustic instruments, and a corresponding degree of innovation and imagination in the compositions they will serve.

Sound Examples

Example 1, "Dygromitz", Z. Settel, 1999, for piano and live electronics. Duration, 12'39"

Example 2, "Throttle", Z. Settel, 1998, for saxophone duo and live electronics. Duration, 11'17"

References

Chowning, J. 1973. "The Synthesis of Complex Audio Spectra by means of Frequency Modulation" *Journal of the Acoustical Society of America* 21(7), 526-534.

Gordon, J. and Strawn J. 1987. "An introduction to the phase vocoder", *Proceedings, CCRMA*, Department of Music, Stanford University, February 1987.

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

Moorer and Berger, 1984. "Linear-Phase Bandsplitting: Theory and Applications", Audio Engineering Society (preprint #2132), New York: 76th AES convention 1984.

Puckette, M., 1991. "FTS: A Real-time Monitor for Multiprocessor Music Synthesis." Music Conference. San Francisco: Computer Music Association, pp. 420-429.

Rowe, Robert, 1993, "Interactive music systems : machine listening and composing" Cambridge, Mass.: MIT Press c1993.ISBN:0262181495

Settel, Z., and Lippe, C., 1994. "Real-time Timbral Transformation: FFT-based Resynthesis", Contemporary Music Review Vol. 10, United Kingdom, Harwood Academic Publishers, pp. 171-179.

Settel, Z., and Lippe, C., 1995. "Real-time Musical Applications using Frequency-domain Signal Processing" IEEE ASSP Workshop Proceedings, Mohonk New York.

Settel, Z., and Lippe, C., 1998. "Real-time Frequency-Domain Digital Signal Processing on the Desktop" International Computer Music Conference Proceedings, Ann Arbor, MI.

Xenakis I., 1971. "Formalized Music", Indiana University Press ISBN:0-253-32378-9.

Diskography

Cage, John, 1946-8, CD: "Sonatas and Interludes for Prepared Piano", Wergo WER60150-50

Hunt, Jerry, 1979, CD: "Lattice" , "Lattice, Transform Stream", CRI CD713

Hunt, Jerry, 1979, CD: "Ground" , "Chimanzzi, Bitom", OD OD disk #9

Lippe, Cort, 1993, CD: "Computer Music Series vol. 21, "Music for Sextet and ISPW", Centaur CRC 2255

Lippe, Cort, 1992, CD: "Computer Music Series vol. 24, "Music for Clarinet and ISPW", Centaur CRC 2310

Martirano, Salvatore, 1992, CD: "Computer Music Series vol. 22, "SATBehind Demo", Centaur CRC 2266

Rai, Takayuki, 1992, CD: "Computer Music Series vol. 21, "Three Inventions", Centaur CRC 2255

Settel, Zack, 1990, CD: "Electro Clips", "Skeewit Chupp", Diffusion Imedia IMED9004

Settel, Zack, 1993, CD: "Computer Music Series vol. 23, "Hok Pwah", Centaur CRC 2302

Subotnick, Morton, 1983, "Computer Music Series vol. 14, "Trembling", Centaur CRC 2170