LOW-DIMENSIONAL AUDIO-RATE CONTROL OF FFT-BASED PROCESSING

Cort Lippe

University at Buffalo
Department of Music
Hiller Computer Music Studios
222 Baird Hall
Buffalo, NY, USA 14260
lippe@buffalo.edu

Zack Settel

McGill University
Music Faculty
555 rue Sherbrooke Ouest
Montreal, Quebec H3A 1E3
CANADA
zack@music.mcgill.ca

ABSTRACT

While the use of the Fast Fourier Transform (FFT) for signal processing in music applications has been widespread, applications in real-time systems for dynamic spectral transformation have been quite limited. The limitations have been largely due to the amount of computation required for the operations. With faster machines, and with suitable implementation for frequency-domain processing, real-time dynamic control of high-quality spectral processing can be accomplished with great efficiency and a simple approach. This paper will describe some recent work in dynamic real-time control of frequency-domain-based signal processing. Since the implementation of the FFT/IFFT is central to the approach and methods discussed below, the authors will provide a brief description of this implementation, as well as of the development environment used in our work.

1. INTRODUCTION

the bin number (index). The IFFT is the complement of the FFT and expects as input, real and imaginary values in the same format as the FFT output.

As seen in Figure 1, the index values provide a synchronization phasor, making it possible to identify bins within a frame, and recognize frame boundaries. The index values can be used to access bin-specific data for various operations, such as attenuation or spatialization, and to read lookup tables for windowing.

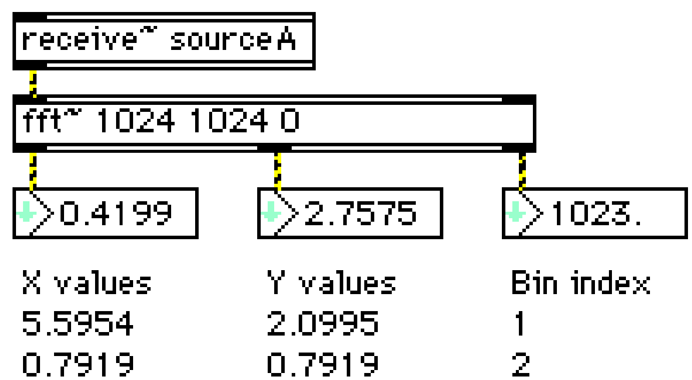

audio sampling rate of 44,100 samples per second). A more dynamic approach to filtering is to update lookup tables containing filter functions at the signal rate (the audio sampling rate.) The term “Spectral Processing Function” (SPF) will be used frequently in this text and refers to a lookup table-based function (actually a signal), whose length is that of the FFT. For each window of input signal (FFT data) which 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 individual frequency components via convolution; a “form” describes the action of the spectral processor. Thus, our approach to dynamic processing spectra focuses on efficiently generating forms with a potentially high degree of detail using simple and intuitive descriptions [5]. In this paper, we will explore 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 both cases, low-dimensional parametric control of complex forms is achieved, and a time-varying SPF is generated for, and convolved with, each window of input signal to perform filtering, band-limited panning, or “spectral mixing” of two or more signals.

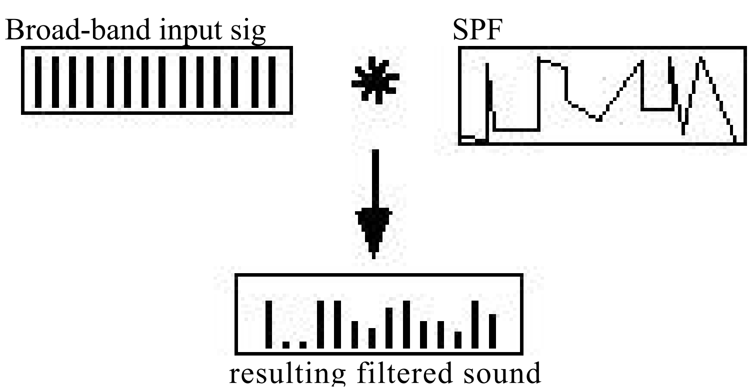

of the FM parameter values to the resulting spectral form make this method extremely easy to use. In our implementation, a particular set of static or time-varying parameters is specified for the FM algorithm; the parameters determine the shape of the resulting SPF.

Other sources for generating complex spectra, such as amplitude modulation, additive synthesis, or waveshaping, may also 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.

2.1 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 used to generate our Spectral Processing Functions (SPF). Techniques such as FM, AM, waveshaping, phase modulation, and 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, spectral panning or spectral mixing. Table lookup operations such as inversion, scaling, offsetting, wrapping, and nonlinear distortion (waveshaping), provide 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 are few, well-known, and easy to understand. As discussed above, the FFT used in our implementation provides us with a phasor that we use as an index for table lookup. Note that the index is mirrored around the frequency at half the sampling rate, following from the symmetrical (real)

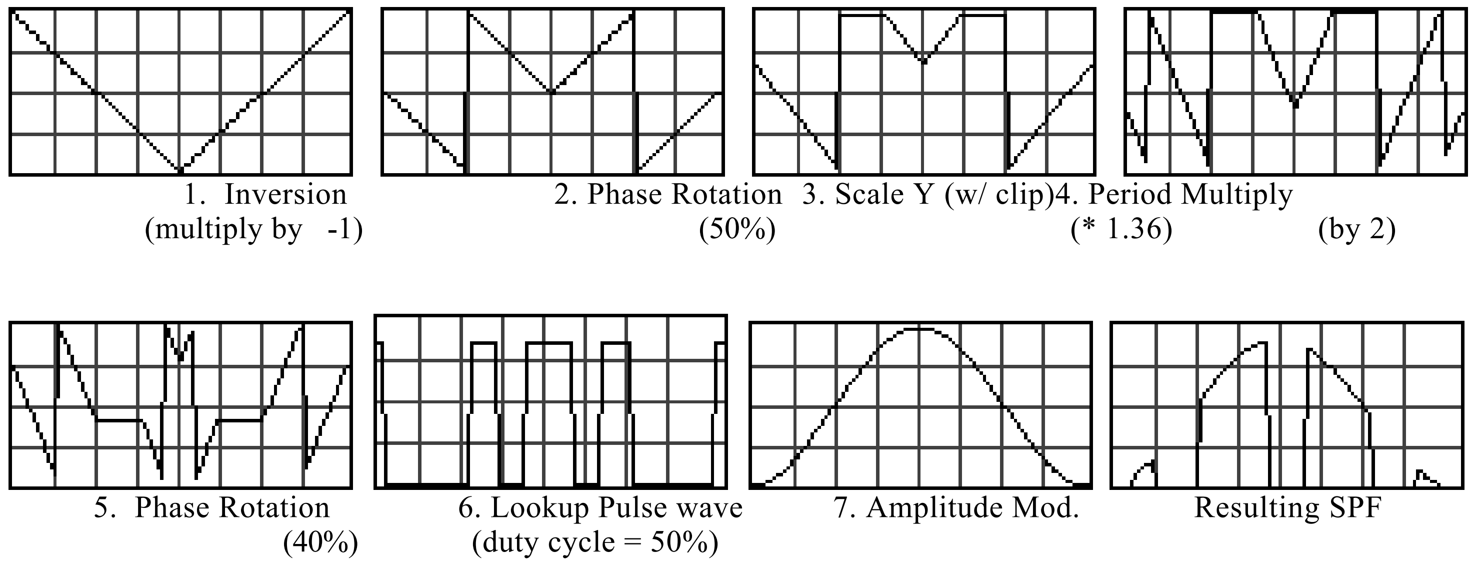

Figure 4. *lookup table operations*

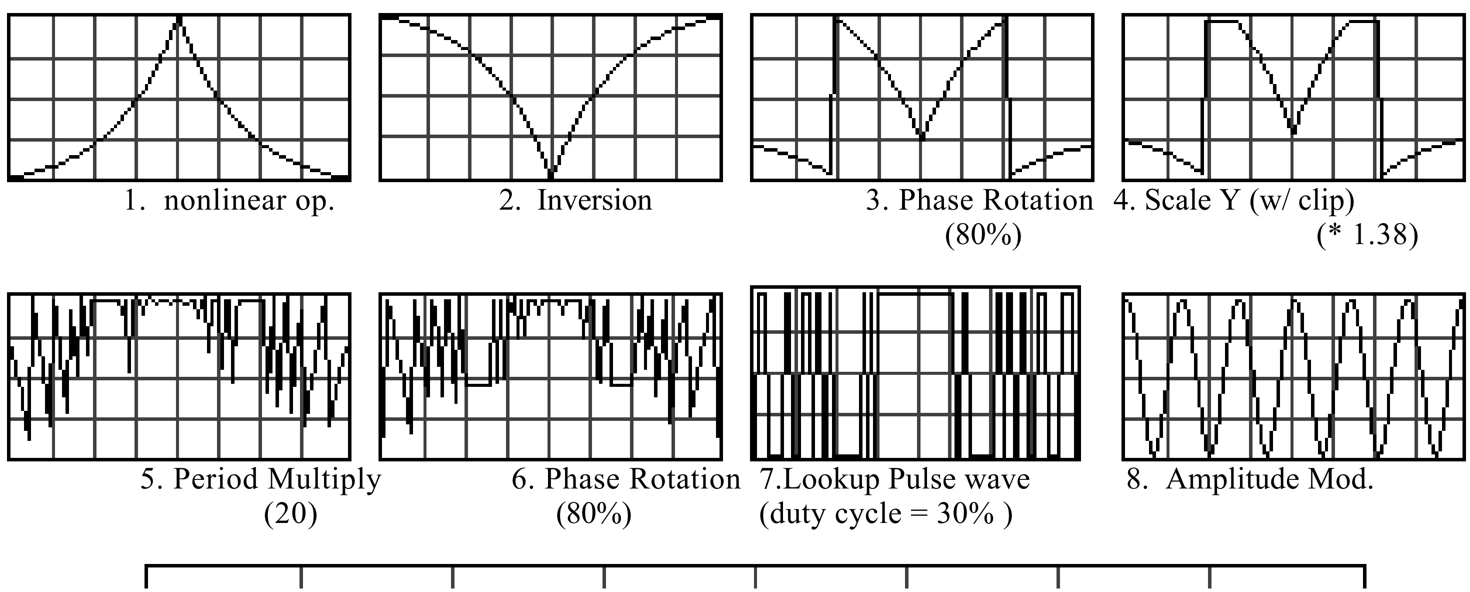

preferable, despite the relatively low spectral update rate of 10Hz (10 FFT frames per second), since a long window allows for a (correspondingly long) highly detailed spectral processing function (SPF.) For example, the "period multiply" operation shown above requires a long SPF when a greater number 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. 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. 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. And, in the case of spectral mixing of two or more signals, the SPF is used to specify the degree to which a given frequency bin will be present for a given signal.

Compressor/limiters are examples of signal processors that analyze the input signal to determine how it is in turn to be modified. A signal passing through such a processor can be thought of as a self-modifying signal. When an analysis stage (such as an envelope or pitch follower) is added to our spectral processing implementation, the resulting input analysis information may be used to dynamically control one or more parameters of the spectral processing. Since our interface for spectral processing requires only a small number of parameters, and since each parameter can significantly alter the shape of the SPF (offering a wide potential range of transformational possibilities), the mapping of few (even just one) input-derived control streams to processing parameters can provide a very high degree of signal self-modification.

There are many other possible mappings of input features to spectral processing parameters. In each case, a given aspect of the input signal will potentially cause the spectrum of that same signal to be modified in a significant way. The challenge is to recognize the practical/musical sense of certain choices (mappings), which in turn, can serve a musical purpose in performance and/or composition. The approach of playing an instrument while processing its sound (as described above) has proven to be a very effective way to evaluate the mappings and discover their particular musical tendencies. As with dynamics processors, the choices of input signal, analysis mapping and processing technique produce results that cover a wide range of timbral possibilities. Future work will be concerned with discovering and establishing musically effective combinations of these choices.

4. CONCLUSIONS

The use of spectral generation and waveform-based techniques for low-dimensional audio-rate control of FFT-based processing has great potential: the parameters are few, familiar and easy to control, and direct mappings of real-time audio input from musicians to the control of FFT-based DSP is made possible. While the applications in this paper has been limited to filtering, spectral panning, and spectral mixing; audio-rate control of FFT-based processing applies equally well to any FFT-based applications where a high degree of processing control is required at the frame rate. The authors have also implemented the techniques mentioned above in applications for denoising and dynamics processing. We believe that these techniques hold great promise for control-intensive FFT-based applications, especially for live-performance DSP, where a new level of interactivity for a relatively new set of frequency-domain DSP techniques, is now possible.