

## A Comprehensive Study of Analysis and Synthesis of Tones by Spectral Interpolation

Marie-Helene Serra, Dean Rubine, Roger Dannenberg  
June 1988  
CMU-CS-88-146

Computer Science Department  
Carnegie Mellon  
Pittsburgh, PA 15213

### Abstract

This paper presents a technique for the analysis and digital resynthesis of instrumental sounds. The technique is based on a model which uses interpolation of amplitude spectra to reproduce short-time spectral variations. The main focus of our work is the analysis algorithm—starting from a digital recording we are able to automatically compute the parameters of our model. The parameters themselves, harmonic amplitudes at selected times, are small in number and intuitively interpretable. The model leads to a synthesis technique more efficient than classical additive synthesis; moreover it allows dynamic spectral variations to be controlled with only a few high-level parameters in real time.

We have studied two analysis/synthesis methods based on spectral interpolation. The first uses only spectral interpolation. This method has allowed us to compress recordings of orchestral instruments to an average of 400 bytes per second without perceptible loss of realism, and to resynthesize these sounds with about 10 arithmetic operations per sample. The second method is a hybrid in which a sampled attack is spliced onto a sustain synthesized via spectral interpolation.

The spectral interpolation model has been applied successfully to different instruments belonging to the brass and woodwind family. We plan to extend the study to many more instruments.

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976 under contract F33615-87-C-1499 and monitored by the: Avionics Laboratory, Air Force Wright Aeronautical Laboratories, Aeronautical Systems Division (AFSC), Wright-Patterson AFB, Ohio, 45433-6543.

Further support for this work was provided by grants from Apple Computer, Inc., and the Yamaha Corporation, an IBM Fellowship Supplement, and a Ben Franklin Partnership Fund Matching Award, Agreement Number 402995-45176-308-001.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies, the U.S. Government, or Carnegie Mellon.

## Table of Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. Evaluation of the Existing Digital Synthesis Techniques</b>	<b>1</b>
2.1. The Criteria	1
2.2. Evaluation of Synthesis Techniques	2
2.3. Our Goals	4
<b>3. Synthesis by Waveform Interpolation</b>	<b>4</b>
3.1. Additive Synthesis Considered	5
3.2. An Attempt at Spectral Variation via Table Lookup	5
3.3. Table-Lookup Synthesis	5
3.4. The Waveform Interpolation Oscillator	7
3.5. Arbitrary Spectral Evolution via Waveforms Interpolation	8
3.6. Spectral Interpolation Analysis	10
<b>4. The Reconstruction Problem</b>	<b>11</b>
4.1. Reproducing Harmonic Sounds	11
4.2. Role of the Analysis	11
4.3. Non-Automatic Derivation of the Input Parameters in the Time Domain	12
<b>5. Analysis/Synthesis by Spectral Interpolation</b>	<b>13</b>
5.1. Digital Recording	13
5.2. Spectral Analysis	13
5.3. Spectral Paths and Spectral Ramps	15
5.4. Data-Reduction using the Time Linear Spectral Interpolation Representation	16
5.4.1. Spectral Ramp Interpolation Using Original Spectra	16
5.4.2. Spectral Ramp Interpolation Using Computed Spectra	17
5.4.3. First Results	19
5.5. Data-Reduction using the Nonlinear Spectral Interpolation Representation	20
5.6. Results	25
5.7. Combination of Sampling and of Spectral Interpolation Synthesis	28
<b>6. Related Work</b>	<b>31</b>
6.1. Mixing Waveforms to Generate Dynamic Sounds	31
6.2. Voice Point Interpolation	32
6.3. Timbre Interpolation	32
6.4. Combination of Sampled Sounds with Synthesized Sounds	32
<b>7. Directions for Future Work</b>	<b>32</b>
<b>8. Conclusions</b>	<b>33</b>
<b>9. Acknowledgments</b>	<b>33</b>

## List of Figures

<b>Figure 3-1:</b>	<b>Table-lookup oscillator</b>	<b>6</b>
<b>Figure 3-2:</b>	<b>The Waveform Interpolation Oscillator</b>	<b>7</b>
<b>Figure 3-3:</b>	<b>Interpolation of successive waveforms</b>	<b>9</b>
<b>Figure 5-1:</b>	<b>Three Spectral Ramps</b>	<b>16</b>
<b>Figure 5-2:</b>	<b>Synthesis by Time Linear Spectral Interpolation</b>	<b>21</b>
<b>Figure 5-3:</b>	<b>Synthesis by Nonlinear Spectral Interpolation <math>c=amp \cdot sramp</math> and <math>d=amp \cdot (1-sramp)</math></b>	<b>26</b>
<b>Figure 5-4:</b>	<b>Synthesis by Spectral Interpolation with Sampled Attack</b>	<b>30</b>

## 1. Introduction

This paper explores techniques for computer analysis and synthesis of musical sounds based on the interpolation of spectra. In section 2, after having examined the technical and musical performance of some existing synthesis algorithms, we list the goals of our work. In section 3 we introduce the method of waveform interpolation synthesis (of which spectral interpolation synthesis is a special case), then we describe the architecture and control of a waveform interpolation oscillator. Section 4 argues the need for an automatic and general analysis model based on interpolation. Here we discuss the conditions a sound must satisfy in order to be amenable to synthesis via waveform interpolation, and introduce the technique of analysis/synthesis by spectral interpolation. Section 5 describes each step of the analysis process using spectral interpolation: digital recording, spectral analysis, and data reduction. We propose two different data reduction algorithms based on "linear" spectral interpolation (meaning linear in time), and another algorithm based on "nonlinear" spectral interpolation. Then we analyze and compare the results obtained with these two algorithms on several types of instruments. This evaluation leads to an improved technique which combines sampling and spectral interpolation synthesis. This hybrid technique forces us to confront the problem of achieving an inaudible transition between sampled and synthesized sound. We discuss two approaches for solving this problem—one successful, the other not. Section 6 situates our research in relation to other relevant work. Finally we outline directions for future work.

## 2. Evaluation of the Existing Digital Synthesis Techniques

In this section we list some criteria by which we can assess synthesis algorithms; we then evaluate many current synthesis techniques according to these criteria. Given this groundwork, we argue that there are desirable points in this space of criteria which are currently achievable, but which until now have not been explored. This leads to the synthesis technique of waveform interpolation, which we describe and evaluate in subsequent sections of this paper. The reader should feel free to skip to section 3 if he wishes to learn about the technique before learning about its underlying motivation.

### 2.1. The Criteria

Each digital sound synthesis algorithm may be viewed as a model of sound production; the model takes as input a number of parameters which collectively determine the sound produced. The production model is usually represented by a mathematical formula whose variables are the input parameters and whose result is the sequence of output samples. Most of the sound generation systems let the user manipulate some of the input parameters to obtain different output audio signals.

Any synthesis algorithm may be judged according to how effectively it enables the musician to produce musically useful sounds. Effectiveness can be measured in terms of specific criteria such as the *generality* of the synthesis method, the computational *efficiency* of the corresponding algorithm, and the quality of the *control* over the sound available to the user. In what follows, we propose a definition of each of these criteria; we also explain why it is important that a digital sound system includes features that tend to maximize these criteria.

#### *generality*

Generality is the ability to generate many kinds of sounds. We can distinguish between two classes of digital sounds. The first is the class of the so-called "natural" or "natural-like" sounds which attempt to imitate the signals produced by acoustical systems such as orchestral instruments and the human voice. The other class of digital sounds, which we call "non-natural" are not attempts to sound like an acoustic system. The generality of a synthesis algorithm is determined by the variety of

sounds, both natural and not, which it is able to produce.

It is clearly desirable for a digital sound system to be general, so that it enables the composer to create a wide range of sonorities, resulting in music whose orchestration is rich and creative. The combination of electronic and physical worlds that a general system could provide may be musically fruitful [52]. The ability to digitally simulate many traditional instruments is good in its own right, since it allows a composer to have a full traditional orchestra always at his disposal. Furthermore, a general system that allows the modification and combination of digital representations of existing instruments to create new ones has interesting musical applications [5].

*efficiency*

Efficiency refers to the computational cost of the synthesis algorithm measured in computational operations per sample of synthesized sound. Greater efficiency implies some combination of reduced overall hardware cost, greater number of synthesized "instruments", and in non-real-time systems, shorter computation time. All of these are practical advantages.

*control*

By controlling the input parameters of the synthesis algorithm the user attempts to achieve a desired perceptual result. Depending on the nature of the input parameters, and on the nature of the algorithm itself, the synthesis system is more or less usable. We think that the following aspects of control are important: *timbral control*, *simplicity of control*, *intuitiveness of control*, and *flexibility of control*. Control over timbre is as desirable as pitch, loudness, and duration control. Indeed, as the musical properties of an instrument come from its characteristic variations of timbre, it is of primary interest to be able to reproduce such variations and to derive new sounds from them. Furthermore, timbre has become a main concern for computer musicians. Given the potential of digital instruments, it is now possible to write and produce dynamic timbral structures as part of a musical composition [53, 54]. It is desirable that control be simple, meaning that there are only a small number of parameters to drive. Control should also be intuitive, meaning that the user should be able to easily imagine the perceptual consequences of parameter modifications. Also, a flexible synthesis model is of interest for the musician because by modifying just a few parameters he can obtain a wide range of sonorities. Simplicity, intuitiveness, and flexibility are necessary to make the sound editing task attractive.

## 2.2. Evaluation of Synthesis Techniques

Currently no technique simultaneously achieves all the desirata mentioned above; existing techniques emphasize some of them to the detriment of the others.

- For example FM [12] is extremely efficient, flexible, simple to control and fairly general. By varying the modulation index and the frequency ratios between carrier and modulated waves, it is possible to get a rich ensemble of natural sounds, harmonic and inharmonic, and create new timbres as well. FM can also mimic the evolution of natural spectra [31, 43]. However, some problems remain in the design of a direct, automatic and general analysis algorithm which can derive FM synthesis parameters from instrumental sounds [34]. In addition, FM parameters are not intuitive.
- With additive synthesis it is possible to reproduce exactly any kind of natural sound (harmonic and inharmonic). The input parameters of the synthesis model (amplitude and frequency time functions of each partial) can be obtained by submitting the sound to a short-time Fourier based analysis: the heterodyne filter [31, 40] gives very good results for quasi-harmonic sounds, and phase vocoder analysis [32, 17] has been used extensively for characterizing and reproducing sounds with complex harmonic structures [45]. Because the input parameters give a very accurate definition of the time-varying spectrum, additive synthesis offers a powerful expressive control over the sound.<sup>1</sup> Furthermore, the parameters

---

<sup>1</sup>For example, interesting musical effects can be obtained by scaling the time-functions describing the evolution of the partials.

of the analysis and synthesis models lend themselves to a perceptual interpretation. Because of the properties of the representation (accuracy and intuitiveness) analysis-based additive synthesis has proved to be a very useful tool for carrying out physical and perceptual descriptions of acoustic instruments [21, 22] and fundamental studies on timbre characterization and timbre manipulation [19, 53]. Thus we see that additive synthesis is general and offers many musical possibilities. However, additive synthesis requires a large amount of data per voice (frequencies and amplitudes of every partial). This is the cause of several disadvantages: control is complex, signal production is expensive (requiring many operations per output sample), and sound representation is somewhat redundant, emphasizing local detail rather than global properties. To reduce cost, complexity, and redundancy, specific data reduction schemes can be applied (see section 3.1).

- Some early digital instruments generated sounds using fixed waveform synthesis [10] (see section 3.3). In contrast to additive synthesis, this technique uses much less data (one constant waveform, one frequency envelope and one amplitude envelope) making it simple and efficient. However it is neither flexible nor general. It can only generate quasi-periodic and harmonic sounds, and it does not provide any control over spectral modulations. For these reasons, fixed waveform synthesis is currently of minimal musical interest.
- Nonlinear techniques, such as waveshaping [39, 4], are economical, flexible, and simple to drive. They can generate and modify natural sounds. In waveshaping synthesis, the reproduction of instrumental sounds can be made by tuning the different modules (sinusoidal generator, nonlinear processors, filters, envelope generators) on a set of analyzed characteristics of the sounds [6]. Adjusting the modules is usually carried out by a trial-and-error phase that is laborious and does not always assure a perfect matching between the original and synthesized signal. For most nonlinear methods generality is limited and control is not intuitive.
- Subtractive synthesis models all consist of the output of a (usually harmonically rich) sound source passed through a parameterized filter. We distinguish between analysis-based and non-analysis-based subtractive synthesis. In analysis-based subtractive synthesis [25] a source function (usually white noise or train of periodic pulses) is fed into a time-varying digital filter. The coefficients of the filter may be obtained by linear predictive analysis of an input signal. In order to accurately reproduce instrument tones, high order filters whose coefficients change quite often are needed [31]. Evaluating analysis-based subtractive synthesis on our criteria, we see that its scores about the same as additive synthesis. The model is general, and intuitive and flexible to control. However, the large filters and high bandwidth needed make it inefficient and make control rather complex.
- Non-analysis-based subtractive synthesis has most often been employed in analog synthesizers. The filters employed in such systems are usually very simple, usually having at most three parameters. This simplicity limits the variety of sounds that can be produced. This results in an evaluation similar, though not identical, to that of FM. Non-analysis-based subtractive synthesis are simple, flexible, and intuitive to control, quite efficient, but not very general.
- In sampling machines [26, 29, 9] the output signals are obtained by playing back prerecorded digital sounds. With this technique it is possible in theory to reproduce virtually any sound, and as many as desired given enough memory. Sampling systems may include different control facilities among which are pitch transposition, filtering, amplitude envelopes, frequency envelopes, and looping of waveform segments. This combination of functions provides a large variety of means for modifying the library of recorded signals. However the design of the sounds, especially simulations of acoustic instruments, requires much effort from the user, the amount of data (the samples) is very large, and timbral control is

restricted.<sup>2</sup> Therefore, sampling is general and extremely efficient (there is little or no computation) but its control is not simple and lacks flexibility.

- Physical modeling synthesis directly simulates the propagation of acoustical waves inside an instrument. The instrument is usually described by a nonlinear source function exciting a linear resonator that may include acoustical losses. The waveguide filter approach maps the instrument into a set of digital filters or delay lines [48, 49], and the mechanical approach considers a set of equations expressing the movements of a system of masses and springs that models the instrument [1, 2]. Ultimately this kind of synthesis should lead to an accurate reproduction of any instrumental sound. However, the complexity of the analysis and the large amount of computation for the synthesis often make this kind of model very expensive. Furthermore, the control of physical models is not straightforward, requiring detailed knowledge of the driving forces applied by the musician to his instrument.

This succinct study about the properties of the principal digital synthesis techniques shows that there is a tradeoff between efficiency and generality (except for sampling) and that no technique scores highly on all criteria.

### 2.3. Our Goals

Our primary goal is to find a simple, efficient, flexible, and intuitive method of reproducing a wide range of traditional instruments. Among the techniques previously cited, only sampling, physical modeling, and additive synthesis lead to an accurate reproduction of natural sounds. We rule out sampling due to its inflexibility of control, and rule out physical modeling as it currently requires excessive computation and complex control. Additive synthesis is not entirely satisfactory as it is inefficient and requires a large amount of data. We desire a technique that keeps the advantages of additive synthesis, and simultaneously is more efficient and offers simpler control. We are willing to trade generality for efficiency and simplicity.

To achieve close similarity between natural sounds and their computed versions, we want a computer analysis model to drive the synthesis. In addition, analysis in the spectral domain has the potential advantage of timbral control and of intuitive control. In the next section we introduce an analysis/synthesis approach that has been designed to fit our goal. This approach has been inspired by both additive and table lookup synthesis.

## 3. Synthesis by Waveform Interpolation

This section introduces the idea of waveform interpolation as a low-cost synthesis method able to reproduce spectral variation. We first consider existing techniques for improving additive synthesis, and we see that while it is possible to reduce the bandwidth of the control data without sacrificing sound quality, synthesis is still expensive and control still complex. We thus abandon the idea of having an oscillator per partial, and take the fixed table-lookup oscillator as our starting point. We then proceed to derive an implementation of a waveform interpolating oscillator, and to demonstrate how spectral variation may be achieved using such an oscillator.

---

<sup>2</sup>Systems which enable velocity switching or velocity crossfading (velocity based selection or mixing of multiple samplers, respectively) allow velocity-driven spectral variation, but the resulting effects remain insufficient if the user wants to imitate the timbral inflections that occur with loudness in acoustic instruments. Also, timbral modifications are limited by the amount of memory.

### 3.1. Additive Synthesis Considered

Roughly speaking, the cost of a digital synthesis method is proportional to the number of independent oscillators and to the bandwidth of control information needed to reproduce a single voice. Many researchers have reduced the control information bandwidth using a number of different data reduction schemes. Grey [21] and Moorer [31], have proposed the use of piecewise-linear functions to approximate the amplitude and frequency functions for each partial. This method has been extensively used in digital synthesizers where the number of oscillators is large [50, 7]. Strawn [51] and Schindler [42] have extended the piecewise-linear approximation by organizing the envelopes breakpoints into hierarchical descriptions. Smith and Serra [49] use a tracking phase-vocoder analysis that retains only the essential features exhibited by the short-time spectrum. Charbonneau [11] and Sasaki [41] have sought data reduction methods that lead to a succinct and global representation of the spectral envelope.

Even when using one of the above data-reduction techniques, the cost of additive synthesis remains high. Additive synthesis requires one oscillator per partial, thus many oscillators are often necessary to synthesize a single voice. Similarly the number of parameters to control remains considerable. Therefore, to improve the efficiency of this technique, we have chosen to look further into ways of reducing the number of oscillators per voice.

### 3.2. An Attempt at Spectral Variation via Table Lookup

Fixed-waveform synthesis (or table-lookup synthesis) uses only one oscillator per voice. The oscillator scans a constant waveform table periodically, scaling the output by an amplitude envelope (see section 3.3). Thus, over a given time interval, only harmonic sounds with a constant spectrum can be generated. If the waveform currently addressed is instantaneously changed to a different one, the resulting spectrum will change, and by generalizing the same procedure to a set of waveforms, dynamic spectral variation can be produced. For example, smooth timbral transitions can be obtained by reading successively a sequence of waveforms that are "very close" to each other [15]. Unfortunately, avoiding perceptual discontinuities (clicks) at the switching points requires a large amount of input data [8, 44]. In other words, the changes between successive spectra need to be slight in order to be imperceptible, necessitating a large control bandwidth.

Thus, it is not very practical to use a single table-lookup oscillator to generate spectral variation, since it is costly to change smoothly between spectra. However, if we allow ourselves two table-lookup oscillators per voice we find that we can indeed change smoothly between wavetables. Using this insight, we can construct a wavetable interpolation oscillator based on two table lookups. It will be possible to drive such an oscillator by a low bandwidth control data computed by a specialized spectral analysis. Starting from the table-lookup synthesis algorithm, we now proceed to derive the wavetable interpolation synthesis algorithm.

### 3.3. Table-Lookup Synthesis

A periodic tone with purely harmonic partials and a slowly varying amplitude envelope can be efficiently generated with the table lookup synthesis technique [27]. The technique consists of reading repeatedly a sequence of numbers stored in a memory. The numbers represent the samples of one cycle of a digitized waveform and the memory is called a wavetable.

Table-lookup synthesis may be described by the following set of equations:



$$PhaseInc(n) = \frac{M \times Freq(n)}{srate} \quad (2)$$

$$Phase(n) = Phase(n-1) + PhaseInc(n) \text{ [mod } M] \quad (3)$$

$$y(n) = Amp(n) \times W[Phase(n)]$$

where

$n$  is the number (index) of the sample being computed,  
 $M$  is the (constant) number of samples in wavetable  $W$ ,  
 $srate$  is the (constant) sampling rate of the output signal,  
 $Freq(n)$  is the instantaneous frequency at sample  $n$ ,  
 $Amp(n)$  is the amplitude scale factor at sample  $n$ ,  
 $W[m]$  is the  $m^{\text{th}}$  sample of the fixed wavetable  $W$ ,

and

$PhaseInc(n)$  is the phase increment at sample  $n$ ,  
 $Phase(n)$  is the phase accumulator at sample  $n$ , and  
 $y(n)$  is the output signal at sample  $n$ .

Figure (Figure 3-1) shows a diagram of an oscillator based on these equations.




Figure 3-1: Table-lookup oscillator

It is quite straightforward to use table-lookup synthesis to create periodic tones. Setting  $Amp(n)$  and  $Freq(n)$  to constant functions insures periodicity (since  $PhaseInc(n)$  is constant,  $Phase(n)$ , and thus  $y(n)$  will be periodic). Furthermore pitch can be varied by varying  $Freq(n)$ , while an amplitude envelope may be imposed on the output by varying  $Amp(n)$ .

In practice,  $Amp(n)$  and  $Freq(n)$  are often specified as piecewise linear functions. In implementations containing ramp hardware,  $Amp(n)$  and  $Freq(n)$  may be updated on every sample (i.e.  $Amp(n) = Amp(n-1) + AmpPhaseInc(n)$ ). It is often the case that  $Amp(n)$  and  $Freq(n)$  are updated at a lower rate (i.e.  $Amp(n) = Amp(n-1)$  except at update times). This update rate imposes an upper limit on how fast the amplitude and frequency of the oscillator is able to vary.

It is worth noting in passing that if the control functions  $Amp(n)$  and  $Freq(n)$  are allowed to be functions of the output of another table-lookup oscillator, complex waveforms will be generated via amplitude and frequency modulation, respectively. Such modifications will not be considered here, and from now on we shall assume that  $Amp(n)$  and  $Freq(n)$  are slowly varying. Given these assumptions, since the table holds a single cycle of the waveform, table-lookup synthesis can generate harmonic signals with the ratios

between harmonic amplitudes always fixed.

### 3.4. The Waveform Interpolation Oscillator

Waveform interpolation is similar in implementation to table lookup (figure 3-1) but uses two phase-locked wavetables,  $W_L$  and  $W_R$ , each one loaded with a different waveform (figure 3-2).

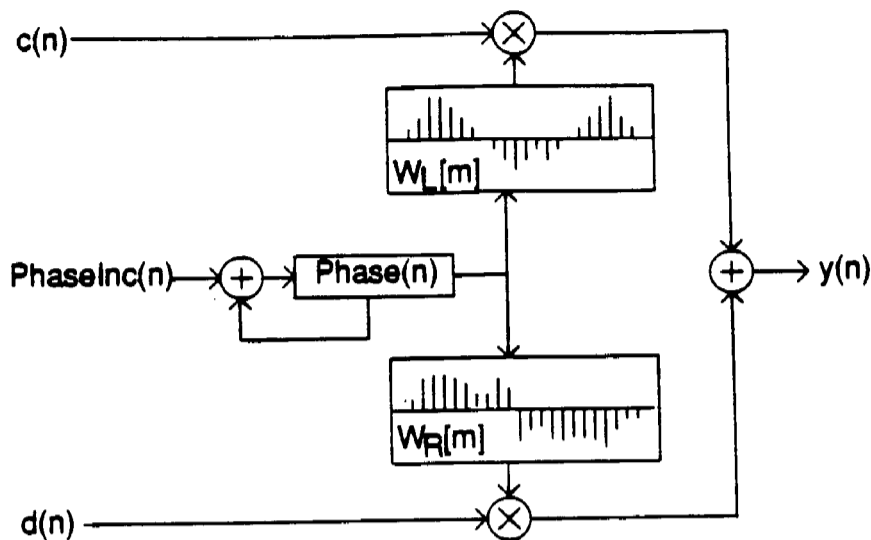


Figure 3-2: The Waveform Interpolation Oscillator

The two tables have the same length  $M$ , and are indexed with the same phase value  $Phase(n)$ . The interpolation signal can be expressed as

$$y(n) = c(n) W_L[Phase(n)] + d(n) W_R[Phase(n)] \quad (4)$$

where

$c(n)$  is the amplitude scale factor of the left wavetable at sample  $n$ ,

$d(n)$  is that of the right,

and

$n$ ,  $y(n)$ ,  $Phase(n)$ , and  $PhaseInc(n)$  are as before.

Thus, we see that a waveform interpolation oscillator is identical in effect to two table-lookup oscillators whose phases  $Phase(n)$  are constrained to be equal. An alternate view is to consider a waveform interpolation oscillator to be a table-lookup oscillator whose table is dynamically computed as the linear combination of two fixed wavetables.

Let us consider this latter view in a bit more detail. We rewrite equation (4) as

$$y(n) = W_{(n)}^E[Phase(n)]$$

where  $W_{(n)}^E$  is the "effective" wavetable in use at sample  $n$ :

$$W_{(n)}^E[m] = c(n) W_L[m] + d(n) W_R[m] \quad (5)$$

Are we justified in calling  $W_{(n)}^E$  the effective wavetable at sample  $n$ ? In other words, are the short-time spectral properties of the output at sample  $n$  close to those that would have resulted had we done simple table-lookup synthesis using fixed wavetable  $W_{(n)}^E$ ? To answer, we consider the short-time frequency domain behavior of the waveform interpolation oscillator at sample  $n$ . We expect that the output spectrum is a linear combination of the spectra of the two wavetables  $W_L$  and  $W_R$ , the combination coefficients being  $c(n)$  and  $d(n)$ . In fact, if one assumes that  $c(n)$  and  $d(n)$  are constant over the window duration, then it is easy to show that the short-time Fourier transform [37] at sample  $n$  of the sequence  $y(n)$  (call it  $Y_n(e^{j\omega})$ ) satisfies

$$Y_n(e^{j\omega}) = c(n)L_n(e^{j\omega}) + d(n)R_n(e^{j\omega}) \quad (6)$$

where  $L_n(e^{j\omega})$  and  $R_n(e^{j\omega})$  are the short-time Fourier transforms of  $W_L[\text{Phase}(n)]$  and  $W_R[\text{Phase}(n)]$ , respectively. Since  $c(n)$  and  $d(n)$  are not in fact constant, but are assumed to be slowly varying, equation (6) is only approximately true. To the extent that the spectrum  $Y_n(e^{j\omega})$  is the desired linear combination of the spectra of the left and right wavetables, we are justified in calling  $W_{(n)}^E$  the effective wavetable at sample  $n$ .

We see that the effect of interpolating waveforms is to interpolate the corresponding spectra. Thus, we propose to generate spectral variation through waveform interpolation.

### 3.5. Arbitrary Spectral Evolution via Waveforms Interpolation

We would like to use waveform interpolation (equation (4)) to generate a spectral evolution. By "spectral evolution" we mean a sound in which the ratios between the amplitudes of the harmonics change over time. There are two ways of achieving spectral evolution. First, as  $c(n)$  and  $d(n)$  may vary dynamically over time, the interpolated waveform, and thus the resultant spectrum, varies over time. As long as  $c(n)$  and  $d(n)$  are continuous the output of the oscillator will be free from clicks (being digital signals, we may interpret "continuous" as meaning "varying sufficiently slowly").

The second way of achieving spectral evolution is to switch the wavetables themselves over the course of a sound. For instance, by repeating the interpolation procedure (equation (4)) with different pairs of waveforms loaded in the right and left tables, one can get a succession of different dynamic spectral combinations. To avoid discontinuities at the point when a waveform is changed (the problem discussed in section 3.2), only one of the two waveforms is changed at any one time, and the change occurs when the scaling factor associated to the wavetable being changed is zero.

Figure 3-3 illustrates the interpolation of a succession of  $Q$  waveforms ( $Q=4$ ) taken in a time-ordered sequence  $\{(n_0, W_0), (n_1, W_1), \dots, (n_{Q-1}, W_{Q-1})\}$ , where  $n_i$  represents the sample at which the reading of waveform  $W_i$  starts. At any point in time, we are interpolating between one of the waveform pairs  $(W_0, W_1), (W_1, W_2), \dots, (W_{Q-2}, W_{Q-1})$ .

In the figure we have shown  $c(n)$  and  $d(n)$  as piecewise-linear functions each of which alternately has value zero at some  $n_i$ , rising linearly to its maximum at  $n_{i+1}$  and again reaching zero at  $n_{i+2}$ . It is not necessary that  $c(n)$  and  $d(n)$  have this form; all that is required is that they be continuous and have value zero at the points their respective wavetables change.

To describe the waveform switching mathematically, we define two functions  $L(n)$  and  $R(n)$ . Given the




Figure 3-3: Interpolation of successive waveforms

sequence of waveforms ( $W_0, \dots, W_{Q-1}$ ),  $L(n)$  and  $R(n)$  respectively determine the waveform in the left and right wavetable of the interpolating oscillator. In a hardware implementation, these functions are used to trigger the change of waveform in each table.

We have

$$y(n) = c(n) W_{L(n)}[Phase(n)] + d(n) W_{R(n)}[Phase(n)] \quad (7)$$

For convenience, we set  $n_{-1} = n_0$  and  $n_Q = n_{Q-1}$ . The left table will be filled with even numbered waveforms, and the right with odd ones as follows:

$$\begin{aligned} L(n) &= 2i & \text{where } n_{2i} \leq n < n_{2i+2} \\ R(n) &= 2i+1 & \text{where } n_{2i-1} \leq n < n_{2i+1} \end{aligned} \quad (8)$$

The necessary criteria to avoid clicks when switching waveforms are

$$\begin{aligned} c(n) &= 0 & \text{when } n = n_{2i} \\ d(n) &= 0 & \text{when } n = n_{2i+1} \end{aligned} \quad (9)$$

A particular and simple case of equation (7) is the *time-linear interpolation* of two waveforms. This is obtained by setting the two mixing coefficients  $c(n)$  and  $d(n)$  to two opposite linear ramps whose sum at each sample equals unity. We can rewrite this special case of equation (7) as

$$y(n) = (1-sp(n)) W_{L(n)}[Phase(n)] + sp(n) W_{R(n)}[Phase(n)] \quad (10)$$

$$sp(n) = \begin{cases} (n - n_{2j}) / (n_{2j+1} - n_{2j}) & n_{2j} \leq n < n_{2j+1} \\ (n_{2j} - n) / (n_{2j} - n_{2j-1}) & n_{2j-1} \leq n < n_{2j} \end{cases}$$

### 3.6. Spectral Interpolation Analysis

Problems may occur when interpolating between two waveforms whose phase distributions are different. Phase cancellation and phase shifting occur when interpolating between two waveforms whose corresponding harmonics are out of phase. Phase cancellation causes the amplitude of each harmonic in the interpolated signal to be less than expected [10]. Even when the amplitudes are the same, interpolating between two out of phase harmonics results in an unintentional but perceptible timbre change, often perceived as a frequency shift. We feel that if the phases are arbitrary, waveform interpolation often gives undesirable results.

To avoid these problems and to get an intuitive control over the interpolation process, we will only interpolate between spectra whose corresponding harmonics are all in phase. For natural sounds, it is in general not the case that the corresponding harmonics of each period are in phase. In order to conduct our experiments on waveform interpolation synthesis, we will rely on an assumption that has been extensively used in digital signal coding of speech and musical signals. The assumption is that the ear is not very sensitive to phase information, so this information can be thrown away [18]. In addition to ignoring phase shifts of given harmonic in time, we ignore the initial phase differences between different harmonics.<sup>3</sup> We use the term synthesis by *spectral interpolation* to mean that we have constrained the corresponding harmonics of each generator wavetable to be in phase.

Expressing the left and right wavetables of an interpolating oscillator as the sum of their respective harmonics, we have

$$\begin{aligned} W_L[m] &= \sum_{h=0}^{H-1} a_h^L \cos(2\pi hm/H + \theta_h^L) \\ W_R[m] &= \sum_{h=0}^{H-1} a_h^R \cos(2\pi hm/H + \theta_h^R) \end{aligned} \quad (11)$$

where  $a_h^L$  and  $\theta_h^L$  are the amplitude and phase of the  $h^{\text{th}}$  harmonic in the left wavetable, and  $a_h^R$  and  $\theta_h^R$  are defined analogously for the right wavetable.

Substituting (11) in (5) we obtain

$$W_{(n)}^E[m] = c(n) \sum_{h=0}^{H-1} a_h^L \cos(2\pi hm/H + \theta_h^L) + d(n) \sum_{h=0}^{H-1} a_h^R \cos(2\pi hm/H + \theta_h^R) \quad (12)$$

Assuming the corresponding harmonics in  $W_L$  and  $W_R$  are in phase (i.e.  $\theta_h^L = \theta_h^R = \theta_h$ ) we have

<sup>3</sup>Depending on the class of instrument, the frequency register, and others factors like pitch and loudness, natural tones show different degrees of initial phase shift [24]. As is often done in additive synthesis and the phase vocoder [32], we ignore the initial phases, at least for most of Section 5. We do, however, consider initial phase differences in Section 5.7.

$$W_{(n)}^E[m] = \sum_{h=0}^{H-1} (c(n)a_h^L + d(n)a_h^R) \cos(2\pi hm/H + \theta_h) \quad (13)$$

Thus, the effective amplitude of the  $h^{\text{th}}$  harmonic in the effective wavetable  $W_{(n)}^E$  is

$$a_h(n) = c(n)a_h^L + d(n)a_h^R. \quad (14)$$

Thus, given the constraints of spectral interpolation, the amplitude of a harmonic of the output at sample  $n$  is equal to the linear combination of the amplitude of the respective harmonic in the left and right wavetables, the combination coefficients being  $c(n)$  and  $d(n)$  as expected.

In this section we have described the process of generating an acoustic signal whose spectral composition can be dynamically modified through the interpolation of waveforms. In the next section we address the question of the analysis and reconstruction of an acoustic signal using such a model.

## 4. The Reconstruction Problem

Spectral interpolation can be used to synthesize a large class of instrumental sounds. We will discuss the types of sounds that are suitable for this type of synthesis and then address the problem of deriving control parameters by automatic analysis techniques.

### 4.1. Reproducing Harmonic Sounds

In section (3.3) we showed that a table lookup oscillator controlled with slowly-varying functions can generate only harmonic sounds. The same result applies to the waveform interpolation oscillator.<sup>4</sup> Thus, we will be able to reproduce only harmonic sounds.

What additional properties should a harmonic signal exhibit to make possible a quality resynthesis by waveform interpolation? Clearly, interpolation will best reproduce a signal with a slowly varying short-time spectrum. As a large class of instruments show gradual changes in their short-time spectra (especially in their sustain portions) we think that waveform interpolation is profitable in many situations.

This does not mean that it is not possible to reproduce signals with large, rapid spectral changes using waveform interpolation. Equation (7) may be flexible enough to adapt to such cases, as it allows the modification of the mixing coefficients of the wavetables, and switching of the waveforms themselves. A large control data bandwidth would be necessary to handle rapid spectral modulations.

### 4.2. Role of the Analysis

To reproduce a sound with the generation model described by the equation (7), we perform a sequence of interpolations between pairs of waveforms. The waveforms  $\{W_i\}$  used in equation (7) are called *generator wavetables*. Each generator wavetable corresponds to a single period of the signal. Because of the interpolation function, the number of generator wavetables will be less than the total

---

<sup>4</sup>This is not exactly true. It is possible to generate sounds with inharmonic partials by interpolation between waveforms containing out-of-phase harmonics [10], and this effect may be used, for example, to generate vibrato [38]. However, as we discuss in section (3.6), it is not likely that we can systematically use this effect to reproduce spectral variations of a given sound, and by excluding this effect we accrue advantages in both analysis and synthesis (see section 3.6).

number of periods in the signal. The time it takes for the interpolation to go from one generator wavetable to the next  $((n_{i+1}-n_i)/\text{srate})$  is called the *interpolation interval*. During the synthesis, the periods located over an interpolation interval are computed with the interpolation formula (7). The function of the analysis preceding the synthesis is to select the generator wavetables, interpolation intervals, and weighting functions  $(c(n)$  and  $d(n))$  such that the resynthesized signal sounds like the original.

#### 4.3. Non-Automatic Derivation of the Input Parameters in the Time Domain

One method of analysis consists in extracting the generator wavetables manually from periods of the original signal. In this method, the user studies a display of the signal over time, and selects the periods he believes to be the onset of significant spectral changes. These periods (each stretched to fill exactly  $M$  points) will be the generator wavetables, and the intervals between successive selected periods, the interpolation intervals. Time-linear interpolation may be used (in which  $c(n)$  and  $d(n)$  are determined from the interpolation intervals), or the user may specify (by, say, drawing) the mixing functions directly, subject to equation (9).

There are several inconveniences associated with manually selecting waveform interpolation parameters:

- The observation and interpretation of the time-domain evolution of the signal is laborious. This is due to the large amount of information displayed, and the difficulty of intuiting significant spectral changes from time-domain functions.
- Even in the simple case of time-linear interpolation, the process of extracting the generator wavetables manually from the original signal is inherently trial-and-error. Periods are selected, a signal is synthesized based on the selection, the result compared to the original, and the process is repeated until satisfactory results are obtained. In theory, this process has to be repeated for each new sound to be synthesized.
- If one desires to manually generate mixing functions as well, the problem becomes even more tedious, due to the interaction of mixing functions and generator wavetables on the synthesized signal.
- Further inconvenience occurs if the user wants to achieve as much data reduction as possible. In this case, when the resynthesized signal sounds good, the user is not sure that the number of generator wavetables can be decreased, and if the resynthesized signal is unsatisfactory, the user is not sure if he is using too few generator wavetables or just the wrong ones.
- When the waveforms are directly extracted from the original input signal, there may be phase differences among harmonics of the same order in successive generator waveforms. The resulting phase cancellation during resynthesis is hard to predict, making the selection process less intuitive. This problem can be somewhat alleviated by augmenting the manual selection process with some automatic phase analysis.

Non-automatic analysis is clearly difficult and labor-intensive. Thus, our effort has focused on building a general analysis model that computes the amplitudes and phases of the harmonics, then automatically selects (or computes) the generator wavetables, interpolation intervals, and weighting functions.

## 5. Analysis/Synthesis by Spectral Interpolation

In this section we present the analysis algorithm that precedes the spectral interpolation synthesis. The analysis algorithm takes as input the acoustic signal that we wish to regenerate, and outputs the control data for the synthesis.

In order to extract from the acoustic signal the relevant information for driving the synthesis, we follow several consecutive steps: digital recording of the sound, spectral analysis of the digitized sound, and data reduction. At the end of the analysis we arrive at a set of data describing (to some approximation) the original sound according to the spectral interpolation model. This set is then fed into the waveform interpolation synthesizer (or its software simulation) to verify the analysis.

### 5.1. Digital Recording

For the purpose of analysis we start to work with isolated tones played (as nearly as possible) at a constant pitch. These restrictions allow us to study separately the reproduction of the amplitude spectral variations by spectral interpolation. The analysis and synthesis described here will still work for tones whose pitch is not constant. However, since vibrato and connected pitches are facets of the tone we wish to control explicitly, we handle these at a higher level. This is briefly mentioned in Section 7 and will be described in a forthcoming paper.

The sound coming from the instrument is recorded using a microphone, and then sent to an analog/digital converter that transforms the analog signal into a stream of 16-bit samples. The sample rate and the anti-aliasing filter must be chosen according to the bandwidth of the signal. In addition we choose a sample rate that leads to an integer number of samples per period. For a given pitch, there are usually several possible sample rates in the system that can be suitable. We try to keep the sample rate close to 16 kHz, as we low-pass filter the signal at 6.4 kHz. If it is not practical to record at arbitrary sample rates, an interpolation and decimation procedure [16] or other resampling algorithm [38, 47] can change the sample rate efficiently. In order to fully experiment with our technique, we record several groups of notes for a given instrument. Each group contains notes played at constant pitch, and with different loudness and duration attributes. Thus we get different kinds of spectral evolutions to reproduce with the spectral interpolation model.

### 5.2. Spectral Analysis

There exist several methods of harmonic extraction based on the Short-Time Fourier Transform (phase vocoder, heterodyne filter, DFT). We use a pitch-synchronous DFT because it is simple and efficient. In addition, when the period is measured accurately, the DFT directly produces the amplitude and phase of each harmonic. Since we recorded the tone at a sample rate to insure an integral period, we need only to check that the period is correct. We do this using Moorer's optimum comb [30] and/or a simple peak detector.

Once the pitch modulations have been tracked and recorded, we can measure the evolution of the harmonics. For simplicity of presentation, we assume constant period  $P$  over the entire tone. The Discrete Fourier Transform of the  $i^{\text{th}}$  period of the discrete signal  $x(n)$  is defined by:

$$X_h^{(i)} = \sum_{n=0}^{P-1} x(iP+n) e^{-j2\pi hn/P} \quad 0 \leq h \leq P-1, \quad 0 \leq i < N_p \quad (15)$$



where  $P$  is the length of the period in number of samples, and  $N_p$  is the total number of periods in the tone. Equation (15) is equivalent to taking the Short-Time Fourier Transform of the signal  $x(n)$  by using a rectangular window of duration  $P$ , and sampling it every  $P$  samples.

The complex DFT at period  $i$ ,  $X_h^{(i)}$ , can be expressed in terms of its real and imaginary parts:

$$X_h^{(i)} = R_h^{(i)} + jI_h^{(i)}$$

The amplitudes of the DFT,  $a_h^{(i)}$ , are given by

$$a_h^{(i)} = |X_h^{(i)}| = \text{sqrt}((R_h^{(i)})^2 + (I_h^{(i)})^2)$$

The phase  $\theta_h^{(i)}$  is given by

$$\theta_h^{(i)} = \text{atan}\left(\frac{I_h^{(i)}}{R_h^{(i)}}\right)$$

As the signal  $x(n)$  is real, the DFT amplitudes are symmetric:  $|X_h^{(i)}| = |X_{P-h}^{(i)}|$ . Thus, we have at most  $\lfloor (P-1)/2 \rfloor$  harmonics (ignoring the DC component at  $h=0$ ). In actuality, we calculate  $H$  harmonics,  $H \leq \lfloor (P-1)/2 \rfloor$ , possibly ignoring some higher harmonics. We ignore the higher harmonics for a number of reasons: they often have insignificant amplitudes, the pitch-synchronous DFT computes them inaccurately, and we want to avoid aliasing when the tone is resynthesized at higher pitches. For simplicity, we evaluate the sum (equation 15) directly for each harmonic  $h$ ,  $1 \leq h \leq H$ . This is often more efficient than using the FFT to compute the entire DFT, since  $H$  may be significantly less than  $P/2$  (so there is no need to calculate all  $P/2$  harmonics), and  $P$  is in general not a power of two (making FFT methods awkward).

The filter interpretation of the short-time Fourier transform [37] shows that if the DFT is computed on a sequence whose length equals the period (or a multiple of the period), then the amplitudes  $a_h$  represent the exact amplitudes of the harmonics. If the input is not perfectly harmonic, or if the period is not an integral number of samples, there will be some error due to the fact that in the passband of the filter for harmonic  $h$  (centered at  $2\pi h/P$ ), more than one harmonic may be present, or the harmonic under analysis may not be in the center of the band. Nonetheless, we have obtained good results using the direct computation of equation (15). Also, we have been satisfied with the low time sampling rate of the short-time Fourier transform (one measure every  $P$  samples).

The DFTs result in a vector of amplitudes on each period of the tone. We call  $S^{(i)}$  the spectrum measured at period  $i$ .

$$S^{(i)} = (a_1^{(i)}, a_2^{(i)}, \dots, a_h^{(i)}, \dots, a_H^{(i)}) \quad (16)$$

The list of DFT spectra with their time indices is  $\{(n_0, S^{(1)}), (n_1, S^{(2)}), \dots, (n_{N_p-1}, S^{(N_p-1)})\}$ . We call this list the "spectral envelope" of the tone. To reproduce the tone using the spectral interpolation model, we want to transform the list of DFT spectra into suitable data for the waveform interpolation synthesizer. That is the subject of the next section.

### 5.3. Spectral Paths and Spectral Ramps

For the remainder of section 5 we describe several algorithms that process the time-ordered list of DFT spectra  $\{(n_0, S^{<0>}), (n_1, S^{<1>}), \dots, (n_{N_p-1}, S^{<N_p-1>})\}$ . The purpose of each algorithm is to obtain the control data needed to drive the spectral-interpolation oscillator (equations (16), (11), (7)): a list of  $Q$  spectra (with associated times) and the scale functions  $c(n)$  and  $d(n)$ . The number of spectra used for the synthesis  $Q$  should be much less than the number of spectra coming from the Fourier analysis  $N_p$ .

In order to explain how the data reduction works, we want to express the interpolation between two spectra  $S^{<0>}$  and  $S^{<1>}$  (successive in the synthesis) in terms of their individual harmonics.<sup>5</sup> Analogously to equation (14), we can express the instantaneous interpolated harmonics at sample  $n$ ,  $a_h^{<ij>}(n)$ , as<sup>6</sup>

$$a_h^{<ij>}(n) = c(n)a_h^{<i>} + d(n)a_h^{<j>} \quad n_i \leq n < n_j \quad 1 \leq h \leq H \quad (17)$$

The effective spectrum at sample  $n$ ,  $S^{<ij>}(n)$ , is

$$S^{<ij>}(n) = c(n)S^{<i>} + d(n)S^{<j>} \quad n_i \leq n < n_j \quad (18)$$

We call the sequence of spectra  $S^{<ij>}(n)$ ,  $n_i \leq n < n_j$ , the *spectral path* from  $S^{<i>}$  to  $S^{<j>}$ . A spectral path consists of the  $H$  loci which connect each harmonic of the initial spectrum  $S^{<i>}$  to the harmonic of same order in the final spectrum  $S^{<j>}$ . A spectral path is defined by a set of two amplitude spectra ( $H$  amplitude values for each spectrum), and a mixing function defined by the two coefficients  $c(n)$  and  $d(n)$ , for  $n_i \leq n < n_j$ .

The interpretation of a spectral path is straightforward when the interpolation is linear in time. In this case, the harmonic amplitudes (equation (10)) are given by:

$$\begin{aligned} S^{<ij>}(n) &= (1-sp(n))S^{<i>} + sp(n)S^{<j>} \quad n_i \leq n < n_j \\ a_h^{<ij>}(n) &= (1-sp(n))a_h^{<i>} + sp(n)a_h^{<j>} \quad 1 \leq h \leq H \quad n_i \leq n < n_j \\ sp(n) &= \frac{n-n_i}{n_j-n_i} \end{aligned} \quad (19)$$

Equation (19) is equivalent to:

$$a_h^{<ij>}(n) = a_h^{<i>} + \frac{n-n_i}{n_j-n_i}(a_h^{<j>} - a_h^{<i>}) \quad 1 \leq h \leq H \quad n_i \leq n < n_j \quad (20)$$

Here, the effect of interpolating between two spectra is that the amplitude of each harmonic ramps linearly from its value in the first spectrum to its value in the second. For this reason  $S^{<ij>}$  (consisting of  $H$  amplitude ramps) is called a *spectral ramp*. A *spectral ramp* is defined by the set of  $H$  initial and  $H$  final values of the harmonic amplitudes, together with the duration of the interpolation ( $n_j - n_i$ ).

Figure 5-1 shows three successive spectral ramps. Representing an amplitude spectrum evolution

<sup>5</sup>This explanation assumes that the spectra used in the synthesis have been selected from the DFT spectra computed in the analysis. While this is often the case, one of our algorithms (see section 5.4.2) computes the spectra used in the synthesis.

<sup>6</sup>For the remainder of section 5 we assume without loss of generality that the waveform for  $S^{<0>}$  is in the left wavetable and the waveform for  $S^{<1>}$  is in the right. If this is not the case, as happens every other pair of spectra, the roles of  $c(n)$  and  $d(n)$  must be interchanged. The equations as written generate sawtooth-like functions for  $c(n)$  and  $d(n)$  which after the interchange become the triangle-like functions as constrained by equation (9).




Figure 5-1: Three Spectral Ramps

with spectral ramps is similar to using a piecewise linear approximation for the individual harmonic amplitudes. However, in contrast to the usual representations used for additive synthesis [21], the breakpoints that define the piecewise linear function for each harmonic are simultaneous.

In order to ensure that the reproduced spectra are close to the original spectra, the data reduction algorithm is based on an error-minimization process. The various methods we use to determine the spectral paths in order to minimize our error criteria are described in the following sections.

#### 5.4. Data-Reduction using the Time Linear Spectral Interpolation Representation

We first consider using spectral ramps for resynthesis. The fitting of the spectral envelope with a small set of spectral ramps is based on the following process: starting with the spectrum of the first period of the tone, we compute the spectral ramp to the spectrum of each successive period in turn. For each successive period we calculate an error measure based on the deviation between the harmonic amplitudes of the original spectra and those of the computed spectral ramp. When the error exceeds a given threshold, the spectral ramp ending at the previous period is stored. The process is repeated using the end of this spectral ramp as the initial spectrum of the next spectral ramp. This loop is executed until the entire spectral envelope has been approximated.

We have tried two different ways for optimizing the spectral ramps. The first method is called *spectral ramp interpolation using original spectra*. It selects some of the DFT spectra in the spectral envelope of the original tone as endpoints of the spectral ramps. The second method is called *spectral ramp interpolation using computed spectra*. It uses a linear regression algorithm to compute the spectral ramps. We discuss each of these in turn.

##### 5.4.1. Spectral Ramp Interpolation Using Original Spectra

Consider equation (19) which defines an interpolated spectrum  $S^{(i,j)}(n)$  on each sample  $n$  within the ramp  $S^{(i,j)}$ . For the purpose of measuring the error on the spectral ramp  $S^{(i,j)}$ , we need to compare the successive computed spectra  $S^{(i,j)}(n)$ ,  $n_1 \leq n < n_2$ , to the corresponding sequence of DFT spectra  $S^{(l)}$  with  $i \leq l < j$  (since  $n_1 = iP$  and  $n_2 = jP$ ). Thus we need to sample the computed spectra at the same rate as the

DFT spectra, i.e. with a time interval equal to the period  $P$ . We consider that the interpolated spectrum on period  $l$  is equal to the spectrum computed at sample  $n_i = lP$ , i.e.  $S^{<l>} = S^{<l>}(n_i)$ . Given this notation equation (20) can be rewritten:

$$a_h^{<l>} = a_h^{<0>} + \frac{n_i - n_i}{n_j - n_i} (a_h^{<0>} - a_h^{<0>}) \quad i \leq l < j \quad (21)$$

Comparing the amplitudes of the harmonics of spectrum  $S^{<l>}$  with values given by equation (21) produces an error  $E_{<l>}^{<l>}$ .  $E_{<l>}^{<l>}$  is defined as the mean square error on the  $H$  harmonic amplitudes:<sup>7</sup>

$$E_{<l>}^{<l>} = \sum_{h=1}^H (a_h^{<l>} - a_h^{<0>})^2 \quad (22)$$

Using the following notation:

$$\Delta a_h^{ij} = a_h^{<j>} - a_h^{<i>} \quad \Delta a_h^{il} = a_h^{<l>} - a_h^{<0>} \quad \Delta n^{ij} = n_j - n_i \quad \Delta n^{il} = n_l - n_i \quad (23)$$

equation (21) becomes:

$$a_h^{<l>} = a_h^{<0>} + \frac{\Delta n^{il}}{\Delta n^{ij}} \Delta a_h^{ij} \quad (24)$$

From equations (24) and (22) we deduce:

$$\begin{aligned} E_{<l>}^{<l>} &= \sum_{h=1}^H \left( \frac{\Delta n^{il}}{\Delta n^{ij}} \Delta a_h^{ij} + a_h^{<0>} - a_h^{<0>} \right)^2 \\ &= \sum_{h=1}^H \left( \frac{\Delta n^{il}}{\Delta n^{ij}} \Delta a_h^{ij} - \Delta a_h^{il} \right)^2 \end{aligned} \quad (25)$$

The global error  $E^{<l>}$  within the spectral ramp  $S^{<l>}$  is defined as the sum of the errors on the individual spectra:<sup>8</sup>

$$E^{<l>} = \sum_{k=1}^{l-1} E_{<k>}^{<k>} \quad (26)$$

If the error  $E^{<l>}$  is less than the tolerated threshold  $E_{\max}$  we extend the spectral ramp to the next period  $<j+1>$  and compute the new error  $E^{<j+1>}$  using equation (26),  $E_{<j+1>}^{<j+1>}$  being computed with equation (25). Otherwise we store the data defining the previous the spectral ramp  $S^{<j-1>}$ , and we compute the next ramp starting at spectrum  $(n_{j-1}, S^{<j-1>})$ .

#### 5.4.2. Spectral Ramp Interpolation Using Computed Spectra

The *Linear Regression* algorithm is a way of fitting piecewise linear functions to a set of points [51]. We use a variant of linear regression called *anchored regression* [40]. Given a fixed point called the *anchor* and a set of data points, the anchored regression algorithm finds the slope of the line passing through the anchor which minimizes the sum of squared distances from the data points to the line. We use this algorithm for the computation of the  $H$  segments which define the spectral ramp. Instead of processing each harmonic separately (which would probably result in a set of segments of different

<sup>7</sup>The averaging factor  $1/H$  is omitted here as it is assumed constant for the tone under study.

<sup>8</sup>We may start the sum from  $i+1$  since as we use  $S^{<0>}$  as the initial spectrum in the spectral ramp, we have  $E_{<0>}^{<0>} = 0$ .

lengths) we perform  $H$  linear regressions on a fixed time interval, and we compute a global error on the resulting spectral ramp.

The anchors (or the endpoints of the spectral ramps) are notated  $S^{<i>}$ . The prime is used to differentiate the anchors from the DFT spectra. The anchored regression algorithm works as follows: We start with an anchor  $(n_i, S^{<i>})$ . For the first anchor we take  $(n_0, S^{<0>})$ <sup>9</sup>. For each successive consecutive DFT spectrum  $(n_j, S^{<j>})$ ,  $j > i$ , we compute  $H$  lines. The  $h^{\text{th}}$  line goes through its anchor  $(n_i, a'_{h<i>})$  and comes closest (in the least squares sense) to the set of points  $(n_{i+1}, a_{h<i+1>}), \dots, (n_j, a_{h<j>})$ . The error on the spectral ramp is computed as the sum of the errors of the  $H$  individual lines. If the error is below a threshold  $E_{\text{max}}$  the next DFT spectrum  $(n_{j+1}, S^{<j+1>})$  is examined. Otherwise the spectral ramp  $S^{<j-1>}$  is used. In this case the endpoint of the spectral ramp is  $(n_{j-1}, S^{<j-1>})$ . The  $H$  coordinates of  $S^{<j-1>}$ ,  $a'_{h<j-1>}$ , for  $1 \leq h \leq H$ , are computed using the slopes of the  $H$  best lines. The endpoint  $(n_{j-1}, S^{<j-1>})$  is then used as the new anchor and the process is repeated to get the next spectral ramp.

Using a simple least squares fit, it is straightforward to compute the best spectral ramp. Since each of the  $H$  lines forming the spectral ramp  $S^{<i>}$  must go through the coordinates  $(n_i, a'_{h<i>})$  each line is defined by an equation of the form

$$a'_{h<i>} < i > = m_h^{<i>} (n_i - n_i) + a'_{h<i>} \quad 1 \leq h \leq H \quad i \leq j \quad (27)$$

where  $m_h^{<i>}$  is the slope of the  $h^{\text{th}}$  line of the ramp.

Defining  $E_h^{<i>}$  to be the sum of the squared errors of each harmonic within the ramp  $S^{<i>}$ , we have

$$E_h^{<i>} = \sum_{l=i+1}^j (a_{h<l>} - (m_h^{<i>} (n_l - n_i) + a'_{h<i>}))^2 \quad (28)$$

The total error on the spectral ramp  $E^{<i>}$  is the sum of the errors on each harmonic:

$$E^{<i>} = \sum_{h=1}^H E_h^{<i>} \quad (29)$$

By using the notation previously defined (equation (23)) we can write:

$$E_h^{<i>} = \sum_{l=i+1}^j (\Delta a_h^{il} - m_h^{<i>} \Delta n^l)^2 \quad (30)$$

$$E^{<i>} = \sum_{h=1}^H \left( \sum_{l=i+1}^j (\Delta a_h^{il} - m_h^{<i>} \Delta n^l)^2 \right) \quad (31)$$

Differentiating the error with respect to  $m_h^{<i>}$ , we obtain

$$\frac{\partial E^{<i>}}{\partial m_h^{<i>}} = -2 \sum_{l=i+1}^j (\Delta a_h^{il} - m_h^{<i>} \Delta n^l) \Delta n^l \quad (32)$$

Setting each derivative  $\frac{\partial E^{<i>}}{\partial m_h^{<i>}} = 0$  gives the slope of each line:

<sup>9</sup>Here the prime is omitted because the anchor is not computed.

$$m_h^{<ij>} = \frac{\sum_{l=1}^j \Delta r^{il} \Delta a_h^{il}}{\sum_{l=1}^j (\Delta r^{il})^2} \quad (33)$$

As we mentioned before, if the total error  $E^{<ij>}$  is less than the threshold, we add the next spectrum  $S^{<j+1>}$  to our set, and calculate the error  $E^{<ij+1>}$ . Interestingly, we do not have to reevaluate from scratch the sums in equations (31) and (33) when we add the new spectra. Rewriting (30) gives

$$E_h^{<ij>} = \sum_{l=1}^j (\Delta a_h^{il})^2 - 2 m_h^{<ij>} \sum_{l=1}^j \Delta r^{il} \Delta a_h^{il} + (m_h^{<ij>})^2 \sum_{l=1}^j (\Delta r^{il})^2 \quad (34)$$

Now, each of the sums in (33) and (34) may be computed incrementally:

$$\begin{aligned} \sigma_{na}^{<ii>} &= \sigma_{aa}^{<ii>}(h) = \sigma_{na}^{<ii>}(h) = 0 \\ \sigma_{na}^{<ij>} &= \sigma_{na}^{<ij-1>} + (\Delta r^{i,j-1})^2 \\ \sigma_{aa}^{<ij>}(h) &= \sigma_{aa}^{<ij-1>}(h) + (\Delta a_h^{i,j-1})^2 \\ \sigma_{na}^{<ij>}(h) &= \sigma_{na}^{<ij-1>}(h) + \Delta a_h^{i,j-1} \Delta r^{i,j-1} \quad \text{for } i < j \text{ and } 1 \leq h \leq H \end{aligned} \quad (35)$$

Now,  $m_h^{<ij>}$  (equation (33)) and  $E_h^{<ij>}$  (equation (34)) can be computed as

$$\begin{aligned} m_h^{<ij>} &= \frac{\sigma_{na}^{<ij>}(h)}{\sigma_{rr}^{<ij>}(h)} \\ E_h^{<ij>} &= \sigma_{aa}^{<ij>}(h) - 2 m_h^{<ij>} \sigma_{na}^{<ij>}(h) + (m_h^{<ij>})^2 \sigma_{rr}^{<ij>} \end{aligned} \quad (36)$$

The incremental computation allows us to do a small amount of work (proportional to  $H$ ) to compute the new slopes and error when adding a spectrum to the regression. In other words, we can compute every error  $E^{<ij>}, E^{<ij+1>}, \dots, E^{<ij+H>}$  with the same effort as it takes to just compute  $E^{<ij>}$ .

#### 5.4.3. First Results

After computing the spectral ramps (using either method), we resynthesize the tone by evaluating equations (11) and (10) in software. The successive spectra defining the spectral ramps are sent to the waveform generator. The waveform generator computes one cycle of a waveform from a given amplitude spectrum. One period of the waveform is obtained by adding  $H$  sine waves, each scaled by the corresponding amplitude. The phases of the sine waves are alternately set to 0 or  $\pi$ , so that the waveform and its first derivative are close to zero at the beginning and at the end of the table. This is the technique used in the Bradford Musical Instrument Simulator [13]. After two waveforms defining a spectral ramp have been loaded in the wavetables, the wavetables are read, scaled by two opposite linear ramps (10) and added to form the output signal. The process is repeated until every spectral ramp has been synthesized.

The degree to which the synthesized signal approximates the input signal depends on the number of spectral ramps output by the data reduction algorithm. This number can be modified by varying the threshold  $E_{max}$ . For each tone, we run the algorithm several times using different thresholds. We then choose the synthesized signal with the smallest number of spectra (largest threshold) that is perceptually indistinguishable from the original tone. The chosen spectral ramps are an accurate (and usually succinct) representation of the tone.

We have obtained very good results on a number of instruments belonging to the woodwind and brass families (bassoon, clarinet, saxophone, trumpet, trombone). The average reduction rate on the number of spectra  $Q/N_p$  is on the order of 10%. The two algorithms, spectral ramp interpolation using original spectra and spectral ramp interpolation using computed spectra, were found to be almost identical in terms of the data reduction they achieve and in their computational cost. For an equivalent data reduction rate, the two algorithms lead to a similar perceptual output. Using computed spectra usually achieves a slightly greater data reduction, since by computing spectra a given amount of error can be spread over a longer spectral ramp than is possible using original spectra.

These promising results were obtained on sounds whose spectral evolution was rather regular. Figure 5-2 shows a sampled trombone tone<sup>10</sup>. Beneath the tone are the spectra which resulted from the DFT analysis of the tone. Under those are the spectra selected by the data reduction algorithm. The synthesized tone is shown under the selected spectra.

When applied to sounds whose amplitudes vary rapidly (as in a rapid tremolo), the data reduction was less effective, although always below 50%. In this kind of situation, a large number of spectral ramps were needed to track the changes in amplitude (and their corresponding spectral changes). As our algorithm proceeds sequentially without any global view of the signal, it does not take advantage of the oscillation between two close spectra, characteristic of the tremolo. To improve the data reduction (in these cases and in general), we have tried another method, based on *nonlinear interpolation*. By relaxing the constraint that the scaling factors  $c(n)$  and  $d(n)$  (equation (18)) be two opposite linear ramps summing to unity, we allow arbitrary combinations of the two waveforms. Thus instead of resorting to the computation of new waveforms, we look for more complex time-varying mixing functions, with the hope that these will result in greater data reduction. The nonlinear interpolation analysis is described in the next section.

### 5.5. Data-Reduction using the Nonlinear Spectral Interpolation Representation

In this section we consider that the paths connecting the harmonics (equation (17))  $c(n)$  and  $d(n)$  are arbitrary functions of time. To determine the optimal set of spectral paths we use a similar procedure as presented in the previous section. But, instead of optimizing the spectra defining the ends of the spectral path, we select two spectra in the DFT list  $S^{<u>}</u>$  and  $S^{<v>}</v>$  separated by an arbitrarily large time interval, and we minimize the square error within the spectral path by choosing the best coefficients  $c(n)$  and  $d(n)$ . If the error is larger than the threshold  $E_{max}$ , we try the same operation with a smaller interpolation duration, i.e. we try the spectral path that links together the spectra  $S^{<u>}</u>$  and  $S^{<v-1>}</v-1>$ , and so on.

Within the spectral path  $S^{<u>}</u>$ , the harmonics of the interpolated spectrum  $S^{<u>}</u>$  at period  $l$  are computed

<sup>10</sup>Actually, part of the sustain portion of the tone was removed so we could fit the signal on the page.

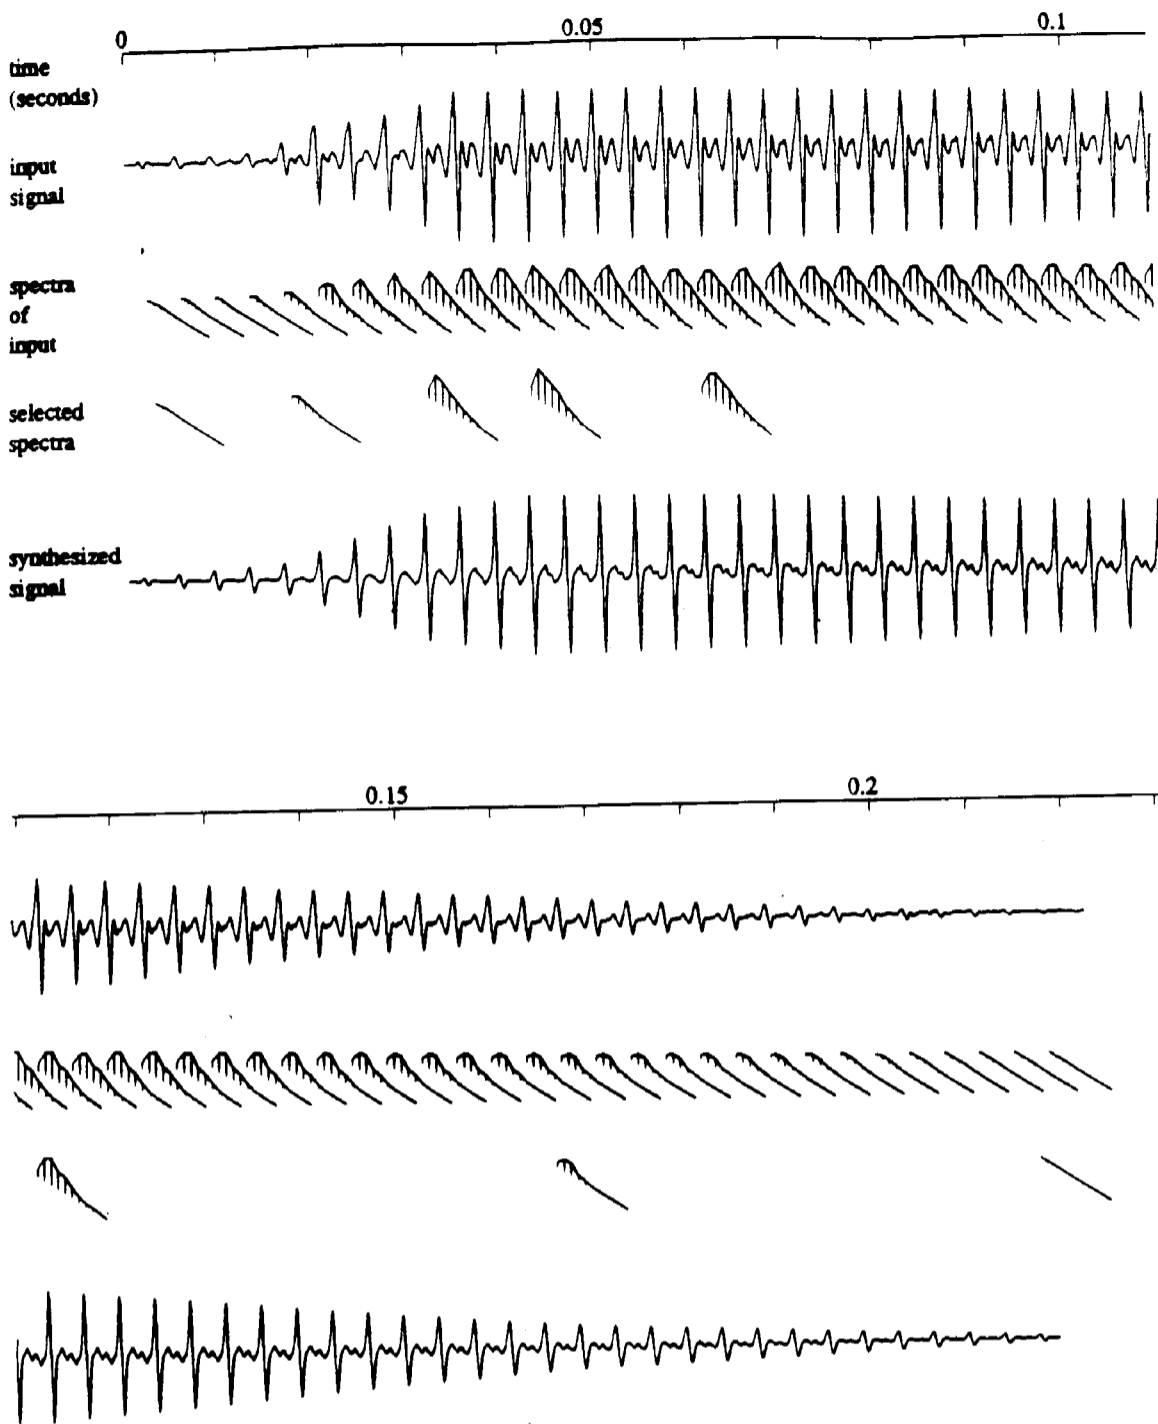


Figure 5-2: Synthesis by Time Linear Spectral Interpolation



as (17):

$$a_{\langle i \rangle}^{(j)} = c(n_i) a_{h \langle i \rangle} + d(n_i) a_{h \langle j \rangle} \quad \text{with } n_i = iP \quad i \leq j \quad (37)$$

Using the notation

$$c_i = c(n_i) \quad d_i = d(n_i)$$

equation (37) can be written as:

$$a_{\langle i \rangle}^{(j)} = c_i a_{h \langle i \rangle} + d_i a_{h \langle j \rangle} \quad (38)$$

or equivalently

$$S_{\langle i \rangle}^{(j)} = c_i S_{\langle i \rangle} + d_i S_{\langle j \rangle} \quad (39)$$

The error on spectrum  $S_{\langle i \rangle}^{(j)}$  within the ramp  $S_{\langle i \rangle}^{(j)}$  is defined as the mean square error on the harmonics<sup>11</sup>:

$$E_{\langle i \rangle}^{(j)} = \sum_{h=1}^H (a_{h \langle i \rangle} - (c_i a_{h \langle i \rangle} + d_i a_{h \langle j \rangle}))^2 \quad (40)$$

The error is minimized by setting the two partial derivatives  $\frac{\partial E_{\langle i \rangle}^{(j)}}{\partial c_i}$  and  $\frac{\partial E_{\langle i \rangle}^{(j)}}{\partial d_i}$  to zero:

$$\frac{\partial E_{\langle i \rangle}^{(j)}}{\partial c_i} = 2c_i \sum_{h=1}^H a_{h \langle i \rangle}^2 + 2d_i \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} - 2 \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle i \rangle} = 0$$

$$\frac{\partial E_{\langle i \rangle}^{(j)}}{\partial d_i} = 2d_i \sum_{h=1}^H a_{h \langle j \rangle}^2 + 2c_i \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} - 2 \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} = 0$$

This is a two equation linear system with unknowns  $c_i$  and  $d_i$ :

$$\begin{aligned} c_i \sum_{h=1}^H a_{h \langle i \rangle}^2 + d_i \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} &= \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle i \rangle} \\ c_i \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} + d_i \sum_{h=1}^H a_{h \langle j \rangle}^2 &= \sum_{h=1}^H a_{h \langle i \rangle} a_{h \langle j \rangle} \end{aligned} \quad (41)$$

The existence of a solution  $(c_i, d_i)$  to the system (41) depends upon the value of its determinant and on the values of the constant terms of the system. Consider the initial equation (equation (39)). If the two vectors  $S_{\langle i \rangle}$  and  $S_{\langle j \rangle}$  are linearly independent (in a vector space of dimension  $H$ ) it is possible to find a unique solution  $(c_i, d_i)$  which minimizes the error. Otherwise the problem remains underdetermined or impossible<sup>12</sup>.

The nonlinear interpolation model has been tried on a number of tones. In practice, there are two kinds

<sup>11</sup>Here again we omit the averaging factor  $1/H$ .

<sup>12</sup>There is an infinity of solutions when the constant terms are zero, otherwise there are no solutions.

of situations for which equations (41) have no unique solution:

- One of the vectors  $S_{<b>}^{<i>}$  or  $S_{<b>}^{<j>}$  is zero or close to zero; this happens usually when one of the spectrum is located in a low-energy segment of the signal (for example at the beginning or at the end of the note).
- The vectors are not zero but are linearly dependent or nearly dependent; this happens usually in the sustain part of the envelope, when the tone is played with a constant loudness or when the amplitudes of the harmonics are varying strictly linearly (which is rare on natural sounds)<sup>13</sup>.

We have investigated several methods for dealing with such cases.

In the first method, equation (39) is replaced by a one vector decomposition by, for example, assuming  $d_i=0$ .

$$S_{<b>}^{<i>} = c_i S_{<b>}^{<i>} \quad (42)$$

$c_i$  being a real number. Expressing equation (42) in terms of the harmonic amplitudes yields

$$a_{<b>}^{<i>} <b> = c_i a_h <i> \quad 1 \leq h \leq H \quad (43)$$

The error on spectrum  $S_{<b>}^{<i>}$  is now

$$E_{<b>}^{<i>} = \sum_{h=1}^H (a_h^{<i>} <b> - a_h <b>)^2 = \sum_{h=1}^H (c_i a_h <i> - a_h <b>)^2 \quad (44)$$

Setting the derivative of the error with respect to  $c_i$  to zero we obtain

$$c_i = \frac{\sum_{h=1}^H a_h <i> a_h <b>}{\sum_{h=1}^H a_h <i>^2} \quad (45)$$

The error on each spectrum within the spectral path can be computed by equation (44) and the error within the spectral ramp is taken as the sum of the errors on the individual spectra. If the global error exceeds the threshold a shorter spectral path will be tried.

This method is straightforward. Unfortunately it is not reliable. The main problem is an occasionally perceptible discontinuity between the two spectra  $S_{<b>}^{<i>} = c_{j-1} S_{<b>}^{<i>}$  and  $S_{<b>}^{<j>}$ . If it was the case that  $S_{<b>}^{<i>}$  and  $S_{<b>}^{<j>}$  were truly dependent, there would be no discontinuity. However, our test for dependency checks only that the determinant is very small, not that it is exactly zero, so it is usually the case that  $S_{<b>}^{<i>}$  and  $S_{<b>}^{<j>}$  are not strictly dependent. This difference can be occasionally perceived as a soft click in the signal. This is exactly the same problem as one gets when switching waveforms as discussed in section 3.2.

The second method uses a linear interpolation to avoid the discontinuity when we abruptly change spectra.

<sup>13</sup>When the two vectors  $S_{<b>}^{<i>}$  and  $S_{<b>}^{<j>}$  as well as  $S_{<b>}^{<i>}$  are close to zero but not strictly zero, we can find a solution.

$$S_{\langle \Delta \rangle}^{c_i} = c_i (S_{\langle \Delta \rangle} + \frac{n_l - n_i}{n_j - n_i} (S_{\langle \Delta \rangle} - S_{\langle \Delta \rangle})) \quad (46)$$

$$\text{so that } d_l = c_i \frac{n_l - n_i}{n_j - n_i}$$

The amplitudes of the harmonics in spectrum  $S_{\langle \Delta \rangle}^{c_i}$  are

$$a_h^{c_i} \langle \Delta \rangle = c_i (a_h \langle \Delta \rangle + \frac{n_l - n_i}{n_j - n_i} (a_h \langle \Delta \rangle - a_h \langle \Delta \rangle)) \quad 1 \leq h \leq H \quad (47)$$

The error  $E_{\langle \Delta \rangle}^{c_i}$  is expressed as

$$E_{\langle \Delta \rangle}^{c_i} = \sum_{h=1}^H (a_h \langle \Delta \rangle - c_i (a_h \langle \Delta \rangle + \frac{n_l - n_i}{n_j - n_i} (a_h \langle \Delta \rangle - a_h \langle \Delta \rangle)))^2 \quad (48)$$

By minimizing the error (48) with respect to  $c_i$  we obtain:

$$c_i = \frac{\sum_{h=1}^H a_h \langle \Delta \rangle \sum_{h=1}^H (a_h \langle \Delta \rangle + \frac{n_l - n_i}{n_j - n_i} (a_h \langle \Delta \rangle - a_h \langle \Delta \rangle))}{\sum_{h=1}^H (a_h \langle \Delta \rangle + \frac{n_l - n_i}{n_j - n_i} (a_h \langle \Delta \rangle - a_h \langle \Delta \rangle))^2} \quad (49)$$

The resulting error is obtained by replacing  $c_i$  by its value in equation (48).

An alternative formulation would be to take  $d_l = \frac{n_l - n_i}{n_j - n_i}$  so that

$$S_{\langle \Delta \rangle}^{c_i} = c_i S_{\langle \Delta \rangle} + \frac{n_l - n_i}{n_j - n_i} S_{\langle \Delta \rangle} \quad (50)$$

$$a_h^{c_i} \langle \Delta \rangle = c_i a_h \langle \Delta \rangle + \frac{n_l - n_i}{n_j - n_i} a_h \langle \Delta \rangle \quad (51)$$

This yields the optimal  $c_i$

$$c_i = \frac{\sum_{h=1}^H (a_h \langle \Delta \rangle a_h \langle \Delta \rangle - \frac{n_l - n_i}{n_j - n_i} a_h \langle \Delta \rangle a_h \langle \Delta \rangle)}{\sum_{h=1}^H a_h \langle \Delta \rangle^2} \quad (52)$$

These last two variants of nonlinear interpolation assures the continuity of the output signal, while allowing a high spectral reduction rate.

By applying the nonlinear interpolation analysis we have been able to greatly reduce the number of waveforms needed to accurately resynthesize the tone. The reduction is greatest during the sustain portion of the tone. However we now need many points (two numbers per period) to represent the scaling functions  $c_i$  and  $d_l$ . These scaling functions are usually well behaved, so we are able to approximate them with piecewise-linear functions containing a small number of breakpoints. This we do after we have computed the spectral paths representing the entire tone. We use the anchored regression algorithm of [40] to automatically perform the fitting.

We do not have an incremental algorithm for computing spectral paths as we had for computing spectral ramps (section 5.4.1). Thus, the effort expended in computing the optimal spectral path between  $S^{<P>}$  and  $S^{<P'>}$  does not help at all in computing the optimal path between  $S^{<P>}$  and  $S^{<P''>}$ , or between  $S^{<P'>}$  and  $S^{<P''>}$ . Our non-linear interpolation algorithm which iterates until the error threshold is crossed (as do the time-linear algorithms of section 5.4.1) is quite time consuming. We could reduce this time by doing binary search, or better yet by developing some kind of incremental algorithm. We have done neither of these, but have instead experimented with a fixed spectral path length, as we now describe.

To speed up the analysis process we have the option of using an arbitrary fixed interpolation interval, *i.e.* we approximate the tone with a list of spectral paths of identical length. On a given tone, we run this algorithm several times varying the spectral path length. We then listen to the resulting resynthesized tones, and choose the one with the longest spectral paths (fewest spectra) which sounds indistinguishable from the original. Figure 5-3 shows the same input tone as in figure 5-2 synthesized via non-linear interpolation method.


## 5.6. Results

The nonlinear interpolation analysis using fixed-length spectral paths (section 5.5) and the linear interpolation analysis (section 5.4) are both very efficient. The linear model is more efficient, partly due to the fact that the nonlinear model requires an additional smoothing of the synthesis input parameters. The nonlinear analysis generally results in an improved data reduction rate.

The table shows results of applying linear and nonlinear interpolation to several tones. All the tones shown in this table were digitized and resynthesized at a 16 kHz sample rate. The *duration* column gives the duration of the tone in seconds,  $P$  is the period of the tone in number of samples,  $N_p$  is the total number of periods in the tone (and thus the total number of DFT spectra),  $H$  is the number of harmonics in each spectrum, *method* is the algorithm (linear or non-linear) used for the analysis,  $Q$  is the number of spectra used in the synthesis (constant for a given tone),  $Q/s$  is the number of spectra per second sent to the synthesizer,  $R$  is the spectral reduction rate, and  $N_{bytes/s}$  is the number of 8-bit bytes per second sent to the synthesizer. The latter three columns are used to characterize the efficiency of the analysis/synthesis process:

- $Q/s$  is the number of spectra per second sent to the synthesizer,  $Q/duration$ .
- $R$  is the spectral reduction rate,  $Q/N_p$  in percent, and thus compares the number of periods in the synthesized tone to the number of periods in the original tone.
- $N_{bytes/s}$  is the number of 8-bit bytes per second sent to the synthesizer. It is equal to the number of spectra per second,  $Q/duration$ , multiplied by the number of harmonics,  $H$  (*i.e.* the amplitude of each harmonic is coded using one byte as in [13]). For nonlinear interpolation, the number of breakpoints per second of  $c(n)$  and  $d(n)$  are added assuming two bytes per breakpoint.

Each of the five tones shown in the table have been synthesized via both methods. For a given tone and method, the tone has been resynthesized using a number of different threshold values. Of these, the tone selected to be shown in the table is the tone synthesized using the smallest number of spectra that is judged to be indistinguishable from the original tone. Thus, the numbers in the table are subjective, but we believe they provide an accurate comparison between our two synthesis methods, and between our synthesis methods and other analysis/synthesis techniques.



**Figure 5-3: Synthesis by Nonlinear Spectral Interpolation**  
 $c = \text{amp} \cdot \text{ramp}$  and  $d = \text{amp} \cdot (1 - \text{ramp})$

<i>name</i>	<i>duration</i>	<i>P</i>	<i>N<sub>p</sub></i>	<i>H</i>	<i>method</i>	<i>Q</i>	<i>Q/s</i>	<i>R</i>	<i>Nbytes/s</i>
Bassoon1	1.44	204	111	40	Linear	16	11.1	14%	444
Bassoon1					Nonlinear	12	8.3	11%	388
Bassoon2	2.12	136	249	40	Linear	19	9.0	8%	359
Bassoon2					Nonlinear	14	6.6	6%	323
Bassoon3	1.06	81	211	30	Linear	18	17.0	9%	509
Bassoon3					Nonlinear	11	10.4	5%	400
Clarinet1	1.25	103	194	40	Linear	23	18.4	12%	736
Clarinet1					Nonlinear	20	18.4	10%	704
Clarinet2	3.77	91	663	25	Linear	21	5.6	3%	139
Clarinet2					Nonlinear	18	4.8	3%	175

In the table we did not mention which linear method was used to compute the spectral ramps (section 5.4.1 or section 5.4.2) as the two methods perform almost identically.

Since the result of spectral interpolation can be exactly computed using additive synthesis, it is interesting to compare the two methods in terms of efficiency. We will use the following parameters:  $r$ , the sample rate,  $N$ , the table size,  $h$ , the number of harmonics,  $w$ , the number of waveforms per second, and  $o$ , the number of operations per second.

For additive synthesis, we have one oscillator per harmonic ( $h$  oscillators) and each sample requires one multiplication for the amplitude, one addition to increment the phase, and one addition to sum the oscillator output into an accumulator. The total of operations per second is

$$o = 3rh.$$

This figure ignores the computation of amplitude ramp functions which could contribute as much as another addition and a comparison to each sample of each oscillator. In practice, amplitude ramps are often computed at low sample rates to reduce the computation requirement.

For spectral interpolation synthesis, a waveform computation requires a multiply and two adds for each harmonic and each waveform sample for a total of  $3hN$  operations. To increment the phase and interpolate between two tables requires three additions and two multiplies, or five operations per sample. Multiplying by the waveform rate ( $w$ ) and sample rate ( $r$ ), respectively, gives:

$$o = 3hNw + 5r.$$

The ratio of spectral interpolation to additive synthesis computation is

$$(5r + 3hNw) / 3hr = 5/3h + wN/r.$$

The  $5/3h$  term represents one interpolating oscillator versus  $h$  sine wave oscillators. The term  $wN/r$  represents the cost of summing harmonics into a wavetable versus computing harmonics at the sample rate.

Using actual data from our experiments, we find that the computation cost for spectral synthesis is

dominated by the cost of computing waveforms. Thus, the most critical factors are the size of the tables and the number of tables computed per second. The relative advantage of spectral interpolation over additive synthesis increases with the sample rate. For a table size of 512 samples and a sample rate of 16 kHz, the ratio varies from 0.30 to 0.63 in typical data. If the sample rate were increased to 48 kHz, the ratio would vary from 0.13 to 0.24. This represents a substantial reduction in the number of operations required by spectral interpolation synthesis with respect to summation synthesis, even if we ignore the cost of extra amplitude ramps and additional control information also required by additive synthesis.

The tones compared are those that gave the best results. These tones belong to the woodwind family. The resynthesis of brass instruments using similar data reduction rates produced acceptable but not identical tones when compared to the originals. In particular, the attacks of the trumpet and the trombone tones did not sound perfectly natural. In fact, increasing the number of spectra did not achieve the desired matching between the synthesized and original attack. We believe this is due to the fact that our analysis/synthesis method does not handle inharmonicity in general, and brass tones often show inharmonic properties in the attack. When we apply the pitch synchronous DFT on a non-harmonic attack we get erroneous spectral measures; furthermore, even if we resort to a more general spectral analysis model (such as the phase vocoder) it is still difficult to use spectral interpolation synthesis to reproduce inharmonicities (see section 3.3). In order to keep the advantages of the spectral interpolation model while achieving more natural attacks, we have investigated a technique in which a sampled attack is spliced onto a synthesized sustain, as we describe in the next section. An evaluation of this hybrid technique is given at the end of that section.

### 5.7. Combination of Sampling and of Spectral Interpolation Synthesis

We now consider the problem of reproducing a tone by connecting the attack portion of the tone to a synthesized sustain portion obtained by analyzing the tone using one of our algorithms previously described<sup>14</sup>. We first examine methods of deciding how much of the original signal to use. We then discuss the techniques we tried for connecting the sampled and synthesized portions.

For data reduction it is desirable to use as small a sampled attack portion as possible<sup>15</sup>. On the other hand, we need enough of the attack to create a convincing reproduction of the tone. As our synthesis method only reproduces harmonic sounds, the attack sample should include the significant nonperiodic or inharmonic portions of the tone. For our technique to be useful it is assumed that these inharmonicities occur at the start of the tone and are of short duration. As one might expect, the two methods we use for pitch detection (optimum comb and peak detection) show erratic behavior during the inharmonic portions of the attack. By observing where the pitch detectors settle down we attempt to get an idea of where the attack inharmonicities finish and the harmonic sustain begins. In actuality this method is not totally reliable (at least we have not worked at making it so) so we are occasionally forced to choose the attack end point manually by examining a plot of the signal.

---

<sup>14</sup>Here we are using the term 'sustain' to refer to the part of a tone after the attack. It includes what is normally thought of as the sustain portion as well as the release portion of the tone.

<sup>15</sup>There are other reasons to want a short attack sample. In particular, as the control of a sampling synthesizer is limited compared to that of the interpolation synthesizer (see section 2.2) we desire the duration of the sampled portion to be as small as possible.

Connecting a sampled attack to a synthesized sustain requires that for each harmonic the phase of the sampled harmonic matches the phase of the corresponding synthesized harmonic at the transition point. If some corresponding harmonic phases are not equal we may hear a click in the signal due to the instantaneous phase shifts. If the transition occurs at a zero crossing in the signal, the click is attenuated, but often still perceivable. We have tried two techniques for achieving smooth transitions, *phase interpolation* and *simple phase matching*.

In the phase interpolation technique we attempt to gradually shift the phases of the harmonics in the sampled attack to be zero or  $\pi$ <sup>16</sup>. To this end, we divided the attack portion into two segments. The first is used unchanged; the second is processed so that by the end of it the harmonic phases are zero or  $\pi$ <sup>17</sup>. We have to find a way of shifting the initial phases of the harmonics of the second segment to zero so that the frequency shift resulting from the phase shift is imperceptible. We tried to use a phase interpolation algorithm, based on [36], where the phases are progressively interpolated (using cubic polynomials) from their original value to zero or  $\pi$ . In order to reduce an eventual perceptual effect, we decrease the phase shift per sample by using a long interpolation time interval, from 5 to 10 periods or more, if necessary. Also, to reduce the global frequency shift, we tried to use opposite directions for each harmonic. In any case, the phase interpolation algorithm has been successful on those few cases where the phase shifts to be applied were very small. Unfortunately, for real signals such as trumpet and trombone tones, the shifts that we had to apply were quite large. In most cases, the phase interpolation was always perceived as a frequency shift, whose magnitude was not acceptable. Although this is a negative result, we found it surprising: *in general the relative phases of the harmonics of a signal cannot be arbitrarily rearranged in an imperceptible manner, even over many periods*. Consider adjusting the relative phase of the second harmonic. In the worst case the harmonic will be  $\pi$  radians out of phase, but this represents only  $\pi/2$  radians of the fundamental. If an adjustment is made over ten periods, we are shifting the second harmonic by at most  $\pi/20$  radians per period. Higher harmonics will have even less phase shift per period, and yet the shift is clearly perceptible.

Thus, we decided to compute the first waveform of the synthesis using the phases measured at the transition. This is the simple phase matching technique. As we can interpolate only between spectra having identical phases, all the waveforms used during the synthesis are computed with those phases. By doing this, the phase match at the transition is faultless. In addition, computing the waveforms with the phases found at the end of the attack did not alter the quality of the synthesized signal.

Using the simple phase matching technique we have obtained very high quality reproductions of brass instruments. Figure 5-4 shows the same input tone as the two previous figures synthesized using a sampled attack. The arrow in the figure points to the transition between sampled sound and synthesized sound.

---

<sup>16</sup>We had an ulterior motive for trying to do this. We wish to use the Bradford Musical Instrument Simulator [14] hardware for our synthesis. This hardware currently requires the phases of the harmonics in the wavetable to be zero or  $\pi$ . A number of benefits come as a consequence of this restriction: the transition between successive wavetable lookups is continuous, the update rate of the amplitude scaling factors in the oscillator can be low, only a small number of bits are needed to code the amplitude factors, and very cheap multipliers (gate arrays) can be used. We would have liked to maintain these advantages.

<sup>17</sup>Ideally, the point where the attack is broken should be one where the harmonic phases are as near as possible to their final values, zero or  $\pi$ . We wrote programs which searched for such points, but in practice it was never possible to find one where more than a few of the harmonics were close to their respective targets.



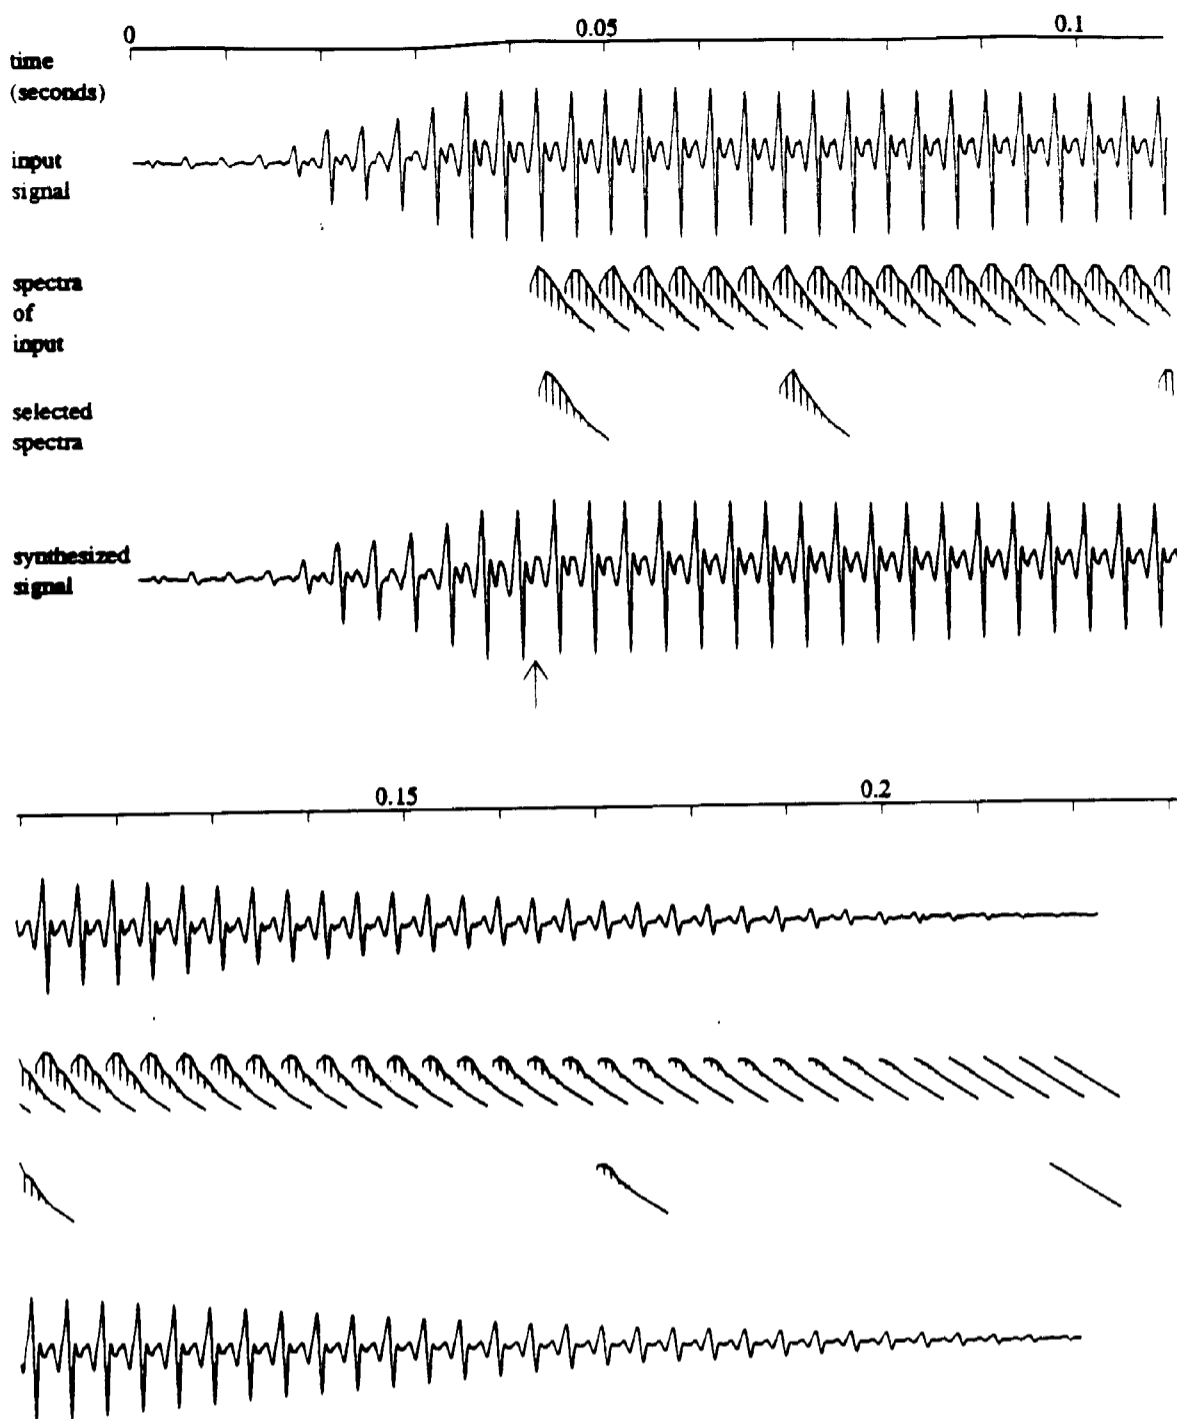


Figure 5-4: Synthesis by Spectral Interpolation with Sampled Attack

The following table depicts the data reduction rates obtained by combining sampling and spectral interpolation on several tones. The table uses the following parameters: *duration* is the total duration of the tone in seconds, *srate* is the sample rate in kHz, *P* is the period of the tone in number of samples,  $N_p$  is the total number of periods in the tone, *H* is the number of harmonics, *Q* is the number of spectra sent to the synthesizer, *attack* is the number of samples copied directly from the original tone, *Nbytes/s* is the number of 8-bit bytes per second sent to the synthesizer (excluding the attack samples), and *method* is the linear or nonlinear interpolation algorithm.

<i>name</i>	<i>duration</i>	<i>srate</i>	<i>P</i>	$N_p$	<i>H</i>	<i>method</i>	<i>Q</i>	<i>QIS</i>	<i>R</i>	<i>attack</i>	<i>Nbytes/s</i>
Trombone1	0.39	13513.5	58	81	25	Linear	5	14.2	6%	500	320
Trombone2	0.96	13513.5	58	210	25	Linear	5	5.4	2%	500	130
Trumpet1	0.55	13201	30	211	6	Linear	5	10.5	2%	1000	54
Trumpet2	0.99	13201	30	389	6	Linear	7	7.7	2%	1000	42
Saxophone1	4.82	16804	46	1666	20	Nonlinear	63	13.2	4%	500	300

The hybrid technique, sampling and spectral interpolation, can be applied without restriction to any kind of sound whose sustain is harmonic, which is the case for most orchestral instruments.

## 6. Related Work

We have presented and discussed the technique of waveform interpolation for the analysis and resynthesis acoustic instruments. We use a succession of wavetables that are dynamically mixed to reproduce analyzed spectral evolutions. Until now, waveform interpolation has been used in a number of other contexts by different people. Some of these past applications are similar to ours, some not. We now review this related work.

### 6.1. Mixing Waveforms to Generate Dynamic Sounds

The idea of mixing multiple sequences of samples to generate sounds whose spectra change over time is not new. Several commercial synthesizers have implemented this technique. The digital instrument built by Matsushita [33] reproduces acoustic instruments with a similar oscillator as in section (3.4). In this instrument, the waveforms are manually extracted from real tones. In fact, to avoid phase problems Matsushita [33] advocates using the Fourier Transform. The Matsushita digital instrument uses the same synthesis algorithm as we do. However, the work differs from ours as it does not use any automatic analysis.

The Prophet VS [35, 3] from Sequential Circuits uses "Digital Vector Synthesis". Each voice is the result of simultaneously mixing four wavetables. Unlike our oscillator, the wavetable contents of the VS are fixed over the course of a note. The four way interpolation is dynamically altered as a function of envelope generators, keyboard velocity, joystick position, as well as other factors.

The Keytex CTS-2000 is a *crossstable sampled synthesizer* [28]. It uses the technique of waveform interpolation. A voice is the sum of two oscillators, each with independent wavetables, frequencies, and amplitudes. Each oscillator can generate an interpolated succession of exactly three different complex wavetables. The synthesizer comes with a number of wavetables stored in ROM. Unfortunately, there is no way for the user to specify his own wavetables.

## 6.2. Voice Point Interpolation

Voice point interpolation is a technique whereby the user specifies signals or synthesis parameters at a number of points, and interpolation is used to compute the signal at points between the specified points. A point is a location in a space, the dimensions of the space most often being pitch, velocity, key pressure, and other controller input. Voice point interpolation is interesting as it attempts to efficiently model the timbre of an instrument as a function the control inputs.

Bradford Musical Instrument Simulator [13, 14] uses voice point interpolation to reproduce timbre variations with pitch and loudness for organ tones, among other instruments. For this two dimensional space they use two waveform interpolation oscillators. Spectral variation with amplitude is a result of the relative scaling of the two wavetables within each oscillator; spectral variation with frequency is achieved by varying the mix of the output of the two oscillators. It is easy to see how this scheme can be generalized to higher dimensional spaces, three dimensions requiring eight waveforms, four dimensions requiring sixteen waveforms, etc.

## 6.3. Timbre Interpolation

Research on timbre has also used interpolation of multiple samples or of spectra. Y. Lo [23] uses interpolation of portions of the signal (called breakframes) to reproduce timbre. The breakframes are normalized in length and recomputed by FFT for phase control. A similar model is applied for interpolating between different timbres. [21] to dynamically combine timbres. To interpolate between instruments, Grey [21] uses interpolation of envelope breakpoints defining the amplitude and frequency harmonics.

## 6.4. Combination of Sampled Sounds with Synthesized Sounds

X. Serra and J. Smith [46] independently combined sampled attacks with synthesized tones using the same phase-matching technique described in section 5.7. Their synthesis technique was based on summation of sinusoids rather than interpolation. The Roland D-50 [20] seems to be able to combine sampled attacks with synthesized sustains. At present we have no technical information on the method the D-50 uses to insure smooth transitions.

## 7. Directions for Future Work

We have simplified our analysis by restricting our data to relatively fixed-frequency examples. To relax this restriction, it should be possible to resample the input to obtain a fixed frequency, perform the analysis, and then use a time-varying frequency in the table lookup oscillator to reproduce the original frequency. To resample, the time-varying frequency must first be determined by some pitch extraction method. The signal can then be resampled [47] at a sampling rate that varies in proportion to the measured frequency. The effect is to "flatten" the frequency variations and obtain a fixed number of samples per period. The same frequency information can be subjected to linear regression for data-reduction before resynthesis.

We have concentrated on sounds with a highly stable period. It would be interesting to apply the same techniques to bow string and vocal sounds which might be expected to have some jitter in their periods. Our analysis/synthesis techniques will tend to remove the jitter, resulting in a perceptually altered sound. Can the jitter be introduced into synthesis for more realism? Can this extension to our basic techniques

be automated?

Finally, this study is the first stage of a two-stage project to meet the criteria of generality, efficiency and control. So far, we have not discussed the problem of control. Our plan is to develop a characterization of an instrument, be it real or artificial, as a multi-dimensional space of spectra. Typical dimensions would be amplitude and frequency, which are input parameters to this synthesis model. Other parameters, such as bow position, lip pressure, etc. can be introduced as additional dimensions. Variations in pitch, amplitude, and other parameters give rise to a trajectory in this space, and spectral interpolation can be used to reproduce the corresponding spectral evolution. Variations of this technique have been alluded to by Grey [21], Covitz and Ashcraft [15], Sasaki and Smith [41] and Bowler [8]. We have already produced synthesized crescendos using a one-dimensional space of spectra and we are now developing analysis software for multi-dimensional spaces.

## 8. Conclusions

We have described techniques for the automatic analysis and resynthesis of musical tones based on spectral interpolation. These techniques are interesting for several reasons. First, automatic analysis is important when it is desired to reproduce known sounds such as the sounds of traditional instruments. Second, we have achieved a high degree of data compression without perceptual degradation in quality or realism for a large and musically useful class of sounds. Third, the computation rate for synthesis is very low compared to other methods with equivalent generality. Finally, the data obtained from the analysis is in a form that can be modified and manipulated in various musically useful ways such as stretching, pitch changing, and interpolation between the spectra of different tones. We are currently studying an extension of this technique in which sequences of spectra are obtained not from a specific tone being reproduced, but by sampling an arbitrary trajectory in a precomputed spectral space. This extension promises a combination of simple and intuitive control, computational efficiency, and realistic production of traditional instrument tones.

## 9. Acknowledgments

We gratefully acknowledge the support of Apple Computer, Inc., and the Yamaha Corporation, whose grants enabled the undertaking of this work. We would also like to credit the Defense Advanced Research Projects Agency for additional support, particularly for providing first-rate research facilities. The Center for Art and Technology and the Computer Science Department both provided space and important organizational support.

Countless people have shared their viewpoints and insights as this work progressed. We would particularly like to thank Richard Stern, Julius Smith, and Andy Moorer for their input.

## References

- [1] J.M. Adrien and X. Rodet.  
Physical Models of Instruments: A modular Approach, Application to Strings.  
In *Proceedings of the 1985 Computer Music Conference*. Computer Music Association, 1985.
- [2] J.M. Adrien., R.Causse, X.Rodet.  
Sound Synthesis by Physical Models, Application To Strings.  
In *Proceedings of the 1987 International Computer Music Conference*. Computer Music Association, 1987.
- [3] J. Aikin.  
Prophet VS analog/digital synthesizer.  
Keyboard magazine.  
August, 1986
- [4] D. Arfib.  
Digital synthesis of complex spectra by means of multiplication of non-linear distorted sine waves.  
*Journal of the Audio Engineering Society* 27(10):757-768, 1979.
- [5] J.B. Barriere.  
Models of Continuity between Synthesis and Processing for the Elaboration and Control of Timbres Structures.  
In *Proceedings of the International Computer Music Conference*. Computer Music Association, August, 1985.
- [6] James Beauchamp.  
Brass-Tone Synthesis by Spectrum Evolution Matching with NonLinear Functions.  
*Foundations of Computer Music*.  
MIT Press, 1985.
- [7] L.Chouinard and L.N.Belanger.  
A Computer-Based Harmonic Analysis/Synthesis System.  
In *Proceedings of the 1985 International Computer Music Conference*. Computer Music Association, 1985.
- [8] I.Bowler.  
The Synthesis of Complex Audio Spectra by Cheating Quite a Lot.  
In *Proceedings of the 1985 International Computer Music Conference*. Computer Music Association, August, 1985.
- [9] D.Byrd and C.Yavelow.  
The Kurzweil 250 Digital Synthesizer.  
*Computer Music Journal* 10(1):64-86, Spring, 1986.
- [10] Hal Chamberlin.  
*Musical Applications of Microprocessors*.  
Hayden Books, 1985.
- [11] Gerard R.Charbonneau.  
Timbre and the Perceptual Effects of Three Types of Data Reduction.  
*Computer Music Journal* 5(2):10-19, 1981.
- [12] John M. Chowning.  
The Synthesis of Complex Audio Spectra by Means of Frequency Modulation.  
*Journal of the Audio Engineering Society* 21(7), 1973.
- [13] P.J. Comerford.  
Bradford musical instrument simulator.  
*IEE Proc.* 128(5), July, 1981.

- [14] P.J. Comerford.  
The Bradford Musical Instrument Simulator.  
In *Proceedings of the 1986 International Computer Music Conference*. October, 1986.
- [15] F.H. Covitz and A.C. Ashcraft.  
Analysis and Generation of Complex Sounds using Small Computers.  
In *Proceedings of Symposium on Small Computers in the Arts*. IEEE Computer Society,  
November, 1981.
- [16] R.E. Crochiere.  
Optimum FIR Digital Filter Implementation for Decimation, Interpolation, and Narrow Band  
Filtering.  
*IEEE Transactions on ASSP* ASSP-23(5), October, 1975.
- [17] Mark B. Dolson.  
The Phase Vocoder: A tutorial.  
*Computer Music Journal* 10(4), Winter, 1986.
- [18] J.L. Flanagan.  
Chapter 8: *Speech Analysis, Synthesis and Perception*.  
Springer-Verlag, New-York, 1972.
- [19] J.Gordon.  
Perceptual Effects of Spectral Modifications on Orchestral Instrument Tones.  
*Computer Music Journal* 2(1):24-31, 1978.
- [20] T. Greenwald.  
Roland D-50 Digital Synthesizer.  
Keyboard magazine.  
September, 1987
- [21] John M. Grey.  
*An Exploration of Musical Timbre*.  
PhD thesis, Stanford University, 1975.
- [22] John M. Grey.  
Perceptual evaluations of synthesized musical instrument tones.  
*J. Acoust. Soc. Am.* 62(2):24-31, August, 1977.
- [23] Yee-On Lo.  
Techniques of Timbral Interpolation.  
In *Proceedings of the 1986 International Computer Music Conference*. Computer Music  
Association, 1986.
- [24] D. Luce and M. Clark.  
Physical Correlates of Brass-Instrument Tones.  
*Journal of the Acoustical Society of America* 42(6), 1967.
- [25] J. Makhoul.  
Linear Prediction: a tutorial review.  
In *Proceedings of the IEEE*, 63(4), pages 561-580. 1975.
- [26] D. Massie.  
The Emulator II Computer Music Environment.  
In *Proceedings of the 1985 International Computer Music Conference*. Computer Music  
Association, August, 1985.
- [27] M.V. Mathews with J.E. Miller, F.R. Moore, J.R. Pierce and J.C. Risset.  
*The Technology of Computer Music*.  
MIT Press, 1969.

- [28] C. Meyer.  
Keytex CTS-2000 Crosstable Sampled Synthesizer.  
Music Technology magazine.  
1987
- [29] J. William Mauchly.  
*Advanced Sampling Guide*  
Ensoniq Corp., 263 Great Valley Parkway Malvern, PA 19355, 1985.
- [30] J. Moorer.  
The Optimum Comb Method of Pitch Period Analysis of Continuous Digitized Speech.  
*IEEE Trans. on ASSP* ASSP-22(5), October, 1974.
- [31] James A. Moorer.  
Signal Processing Aspects of Computer Music - A Survey.  
*Computer Music Journal* 1(1):4-37, February, 1977.
- [32] James A. Moorer.  
The Use of the Phase Vocoder in Computer Music Applications.  
*Journal of the Audio Engineering Society* 26(1/2):42-45, February, 1978.
- [33] M. Nikaido *et al*, Matsushita Elect. Inc.  
Wave Generating Method Apparatus Using Same.  
*United States Patent 4,597,318* (), Filed in January 1984, Issued in July 1986.
- [34] Russell G. Payne.  
A Microcomputer Based Analysis/Resynthesis Scheme For Processing Sampled Sounds Using FM.  
In *Proceedings of the 1987 International Computer Music Conference*, pages 282-289. Computer Music Association, 1987.
- [35] Curtis Roads.  
Prophet VS: User's Report.  
*Computer Music Journal* Vol.10, Number 2.  
1987
- [36] T.F. Quatieri and R.J. McAulnay.  
Speech Transformations Based on Sinusoidal Representation.  
*IEEE Trans on ASSP* ASSP-34, 1986.
- [37] Alan V. Oppenheim (editor).  
*Signal Processing: Digital Processing of Speech Signals*.  
Prentice-Hall, 1978.
- [38] T.A. Ramstad.  
Digital Methods for Conversion Between Arbitrary Sampling Frequencies.  
*IEEE Transactions on ASSP* ASSP-32(3), June, 1984.
- [39] Curtis Roads.  
A Tutorial on Nonlinear Distorsion or Waveshaping Synthesis.  
*Foundations of Computer Music*.  
MIT Press, 1985.
- [40] Dean Rubine.  
On the Analysis of Nearly Harmonic Tones.  
Unpublished 1985.
- [41] L.H. Sasaki and K.C. Smith.  
A Simple Data Reduction Scheme For Additive Synthesis.  
*Computer Music Journal* 4(1), 1980.

- [42] K.W. Schindler.  
Dynamic Timbre Control for Real Time Synthesis.  
*Computer Music Journal* 8(1):28-42, Spring, 1984.
- [43] Bill Schottstaedt.  
The Simulation of Natural Instrument Tones Using Frequency Modulation with Complex Modulating Wave.  
*Foundations of Computer Music*.  
MIT Press, 1985.
- [44] G.Schwartz.  
Sound Synthesis by Hierarchic Sampling.  
In *Proceedings of the 1985 International Computer Music Conference*, pages 62-73. Computer Music Association, August, 1985.
- [45] Xavier Serra.  
A Computer Model for Bar Percussion Instruments.  
In *Proceedings of the 1986 International Computer Music Conference*. Computer Music Association, October, 1986.
- [46] J.O.Smith and X.Serra.  
PARSHL: An Analysis/Synthesis Program for Non-Harmonic Sounds Based on a Sinusoidal Representation.  
In *Proceedings of the 1987 International Computer Music Conference*. Computer Music Association, August, 1987.
- [47] J.O. Smith.  
A Flexible Sample Rate Conversion Method.  
In *Proceedings of ICASSP*, pages 19.4.1-19.4.4. publisher, 1984.
- [48] J.O. Smith.  
Efficient Simulation of the Reed-Bore and Bow-String Mechanisms.  
In *Proceedings of the 1986 Computer Music Conference*. Computer Music Association, 1986.
- [49] J.O. Smith.  
Waveguide Filter Tutorial.  
In *Proceedings of the 1987 International Computer Music Conference*. Computer Music Association, 1987.
- [50] John Snell.  
Design of a Digital Oscillator That Will Generate up to 256 Low-Distorsion Sine Waves in Real Time.  
*Foundations of Computer Music*.  
MIT Press, 1985.
- [51] John Strawn.  
Approximation and Syntactic Analysis of Amplitude and Frequency Functions for Digital Sound Synthesis.  
*Computer Music Journal* 4(3):3-23, Fall, 1980.
- [52] Gabriel Weinreich.  
New Perspectives in the Physics-Music-Computer Triangle.  
In *1st Symposium on Computers and the Performing Arts*. Ann Arbor, Michigan, April, 1987.
- [53] David Wessel.  
Timbre Space as a Musical Control Structure.  
*Foundations of Computer Music*.  
MIT Press, 1985.



- [54] David Wessel.  
Control of Phrasing and Articulation in Synthesis.  
In *Proceedings of the 1987 International Computer Music Conference*. Computer Music  
Association, Urbana, Illinois, August, 1987.