William A. Sethares
Department of Electrical and Computer Engineering
1415 Engineering Drive
University of Wisconsin
Madison, Wisconsin 53716, USA
sethares@ece.wisc.edu

Consonance-Based Spectral Mappings

This article presents a method of mapping the spectrum of a sound so as to make it tonally consonant with a given specified reference spectrum. One application is to transform inharmonic sounds into harmonic equivalents. Alternatively, this method can be used to create inharmonic instruments that retain much of the tonal quality of familiar (harmonic) instruments. Musical uses of such timbres are discussed, and forms of (inharmonic) modulation are presented. A series of sound examples demonstrate both the breadth and limitations of the method. [Sound examples will be included on the forthcoming compact disc for Volume 22 of *Computer Music Journal*.—Ed.]

Background

Wendy Carlos (Carlos 1987) gives several examples of the interrelationship between timbre and tuning, exclaiming, "Clearly the timbre of an instrument strongly affects what tuning and scale sound best on that instrument." The relationship between the spectrum of a sound and a scale or tuning in which the sound is most consonant has been formalized (Sethares 1993) using the idea of the dissonance curve: a plot of the calculated dissonance of a spectrum versus frequency interval. Such dissonance curves are based on a parameterization of tonal consonance data (Plomp and Levelt 1965), which is in turn closely related to Hermann von Helmholtz's work (Helmholtz 1863) on the beating of sine-wave partials of simultaneously sounding tones.

For example, Figure 1 shows the dissonance curve for a tone with seven harmonic partials, which has minima at intervals formed from simple integer ratios. Thus, harmonic tones are said to be *related to* (most tonally consonant when played in) just intonations—scales composed of intervals with simple-integer frequency ratios. Other kinds of sounds, such as those with inharmonic partials, are related to other (non-just) scales, and can be explored using computer sound synthesis.

Given a sound, it is straightforward (as shown below) to plot its dissonance curve and determine the scale in which it is predicted to sound most consonant. This scale and its "chords" provide a sensible starting point for the exploration of unfamiliar scales when played in conjunction with inharmonic timbres. Conversely, given a desired scale, it is possible to determine spectra (sets of partials) for timbres that will be most (tonally) consonant when played in that scale. Suppose a composer desires to play in some specified scale, say, eleven-tone equal temperament (11-TET). Many familiar harmonic sounds are very dissonant when played in 11-TET, because the minima of the dissonance curve occur far away from the scale steps, that is, many of the 11-TET intervals occur near peaks, rather than valleys of this harmonic dissonance curve. Observe that the top horizontal axis of Figure 1 is labeled with the scale steps of the familiar twelve-tone equal temperament (12-TET), while the bottom axis shows the 11-TET scale steps.

To increase the tonal consonance of a piece in 11-TET, it may be advantageous to create a new set of sounds, with spectra that cause minima of the dissonance curve to occur at the appropriate 11-TET scale steps. Figure 2, for example, defines a spectrum with a dissonance curve that has major dips at many of the locations of the 11-TET scale steps. These are locations where (tonal) consonance is maximized. Such sounds are typically specified as a desired set of partials, but a complete spectrum consisting of magnitude and phase must be chosen to draw the dissonance curve and to transform the sound into a time waveform for playback.

Computer Music Journal, 22:1, pp. 56–72, Spring 1998

Figure 1. Plots of the calculated dissonance of a spectrum versus frequency interval are called dissonance curves. This is the dissonance curve for a seven-partial harmonic spectrum. The minima of this curve occur at 1, 7/6, 6/5, 5/4, 4/3, 7/5, 3/2, 5/3, 7/4, and 2/1, which lie near many of the 12-TET scale steps (top axis) and relatively far from the 11-TET scale steps (bottom axis). Dissonance values (on the vertical axis) are normalized to unity.

Figure 2. Dissonance curve for the spectrum with partials at 1,a^{11},a^{17},a^{22},a^{26}, a^{28},a^{31},a^{33},a^{37}, and a^{38}, where $a = \sqrt[11]{2}$. In contrast to Figure 1, the minima of this dissonance curve occur at many of the 11-TET scale steps (bottom axis), and not at the 12-TET scale steps (top axis).

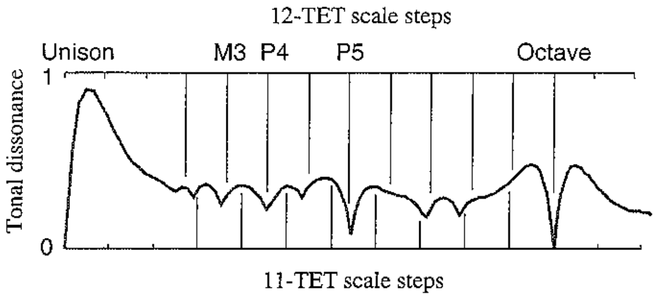

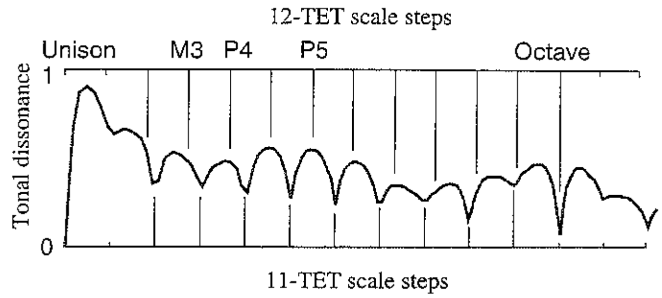

For the figure, all partials are assumed to have equal amplitudes, giving the sound a rich, organ-like quality.

The most straightforward approach to the problem of sound synthesis from a specified set of partials is additive synthesis, such as described by Jean-Claude Risset and David Wessel (Risset and Wessel 1982), in which a family of sine waves of desired amplitude and phase are summed. Though computationally expensive, additive synthesis is conceptually straightforward. A major problem is that it is often a monumental task to specify all of the parameters (frequencies, magnitudes, and phases) required for the synthesis procedure, and there is no obvious or intuitive path to follow when generating new sounds. When attempting to create sounds for new scales, such as the 11-TET timbre above, it is equally challenging to choose these parameters in a musical way. Making arbitrary choices often leads to organ- or bell-like sonorities, depending on the envelope and other aspects of the sound. While these can be quite striking, they can also be limiting from a compositional perspective. Is there any way to create a full range of tonal qualities that are all related to the specified scale? For instance, how can flute-like or guitar-like timbres be built that are consonant when played in this 11-TET tuning?

A common way to deal with the vast amount of information required by additive synthesis is to analyze a desired sound via a Fourier (or other) transform, and then use the parameters of the transform in the additive synthesis. In such analysis/synthesis schemes, the original sound is transformed into a family of sine waves, each with specified amplitude and phase. The parameters are stored in memory and are used to reconstruct the sound on demand. In principle, the methods of analysis/synthesis allow exact replication of any waveform. Of course, the sound to be resynthesized must already exist for this procedure to be feasible. Unfortunately, 11-TET flutes and guitars do not exist.

Once a sound is parameterized, it is possible to manipulate the parameters. For example, one technique interpolates the envelopes of harmonics to gradually transform one instrumental tone into another (Grey and Moorer 1977), while another exchanges the spectral and temporal envelopes among a number of wind-family instruments, and conducts tests to evaluate their relative significance (Strong and Clark 1967). Probably the first parameter-based analysis/synthesis methods were the vocoder (Dudley 1939) and its modern descendant, the phase vocoder (Flanagan and Golden 1966), which were designed for the efficient encoding of transmitted speech signals.

The consonance-based spectral mappings of this article are a kind of analysis/synthesis method in which the amplitudes and phases of the spectrum of the source sound are grafted onto the partials of a specified destination spectrum, which is chosen so as to maximize a measure of consonance (or more properly, to minimize a measure of dissonance). The goal is to relocate the partials of the original sound for compatibility with the destination spectrum, while leaving the tonal quality of the sound intact. Musically, the goal is to modify the spectrum of a sound while preserving its richness and character. This provides a way to simulate the sound of non-existent instruments such as the 11-TET flute and guitar. Figure 3 shows the spectral mapping scheme in block-diagram form. The input

Figure 3. Block diagram of a transform-based analysis-synthesis spectral mapping. If the mapping is chosen to be the identity, then the input and output signals are identical.

Figure 4. Two sine waves are sounded simultaneously. Typical perceptions include pleasant beating (at small frequency ratios), roughness (at middle ratios), and separation into two tones (at first with roughness, and later without) for larger ratios. The horizontal axis represents the frequency interval between the two sine waves, and the vertical axis is a normalized measure of tonal dissonance. The different plots show how the tonal consonance and dissonance varies, depending on the frequency of the lower tone.

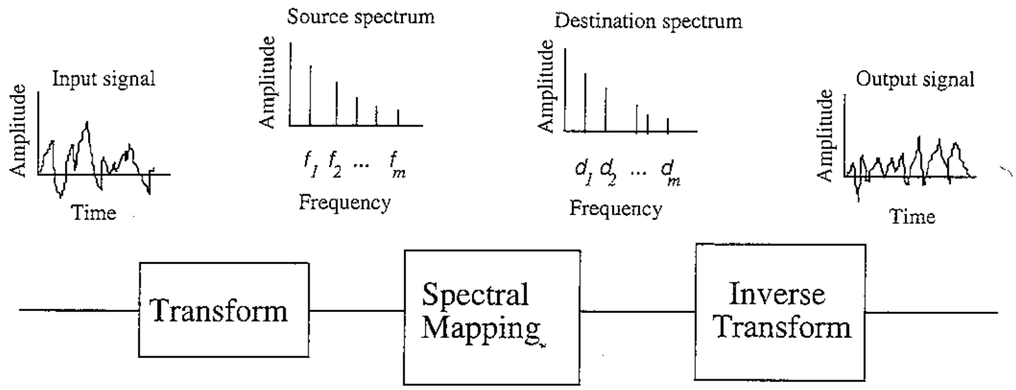

Figure 3

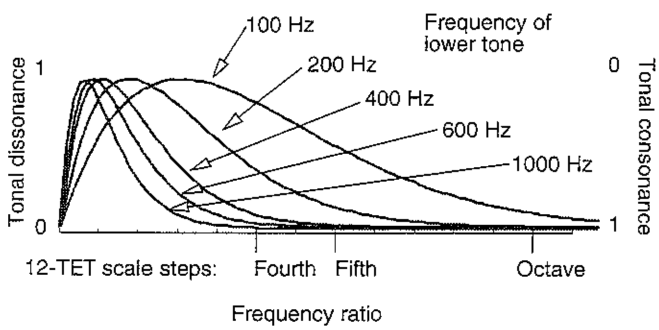

Figure 4

signal is transformed into its spectral parameters, the mapping block manipulates these parameters, and the inverse transform returns the signal to a time-based waveform for output to a digital-to-analog converter (DAC) and subsequent playback.

The next section of this article briefly reviews the construction of consonance and dissonance curves. The third section details the design of practical consonance-based spectral mappings, and the fourth section gives several examples that show both the strengths and limitations of the method. The fifth section discusses various aspects of timbral change in light of the sound examples, and suggests certain forms of inharmonic modulation.

Drawing Dissonance Curves

Existing psychoacoustic work provides a basis on which to build a measure of tonal consonance and dissonance that can be used to guide the choice of spectral mapping (Plomp and Levelt 1965). Reinier Plomp and W. Levelt asked volunteers to rate the perceived dissonance of pairs of pure sine waves, giving curves such as in Figure 4, in which the dissonance is minimum at unity, increases rapidly to its maximum somewhere near one-quarter of the critical bandwidth, and then decreases steadily back toward zero. When considering sounds with spectra that are more complex, dissonance can be calculated by summing the dissonances of all the partials, and weighting them according to their relative amplitudes. This leads to dissonance curves such as Figure 1 (for harmonic sounds) and Figure 2 (for the specified inharmonic sound).

To be concrete, the dissonance between a sinusoid of frequency f_1 with amplitude v_1 and a sinusoid of frequency f_2 with amplitude v_2 can be parameterized as

$$d(f_1, f_2, v_1, v_2) = v_1 v_2 \left[e^{-as|f_2 - f_1|} - e^{-bs|f_2 - f_1|} \right] \quad (1)$$

where

$$s = \frac{d^*}{s_1 \min(f_1, f_2) + s_2} \quad (2)$$

$a = 3.5$, $b = 5.75$, $d^* = .24$, $s_1 = .21$, and $s_2 = 19$ are determined by a least-squares fit. The amplitude term $v_1 v_2$ ensures that softer components contribute less to the total dissonance measure than those with larger amplitudes; d^* is the interval at which maximum dissonance occurs; and the s parameters in Equation 2 allow a single functional form to smoothly interpolate between the various curves of Figure 4 by sliding the dissonance curve along the frequency axis so that it begins at the smaller of f_1 and f_2, and by stretching (or compressing) it so that the maximum dissonance occurs at the appropriate frequency. Derivation, justification, and discussion of this model are available (Sethares 1993).

More generally, a spectrum F with base (or fundamental) frequency f is a collection of n sine waves (or partials) with frequencies $f, a_2 f, \ldots, a_n f$ and amplitudes $v_1, v_2, \ldots, v_n$, where the a_i are ordered and are all are greater than one. The intrinsic dissonance of the sound F is the sum of the dissonances of all pairs of partials:

$$D_F = \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} d(a_i f, a_j f, v_i, v_j), \qquad (3)$$

with the convention that $a_1 = 1$. When two notes of F are played simultaneously at an interval c, the resulting sound has a dissonance that is the same as that of a timbre with frequencies

$$f, a_2 f, \ldots, a_n f, cf, ca_2 f, \ldots, ca_n f.$$

Equation 3 can then be used to calculate the intrinsic dissonance of this concatenated spectrum. Equivalently, the transposed version of F can be defined as cF with partials at $cf, ca_2 f, \ldots, f$ and amplitudes $v_1, v_2, \ldots, v_n$. The dissonance of F at the interval c is

$$D_F(c) = \frac{1}{2}(D_F + D_{cF}) + \sum_{i=1}^{n}\sum_{j=1}^{n} d(a_i f, ca_j f, v_i, v_j), \qquad (4)$$

and the dissonance curve generated by the spectrum F is defined to be this function, $D_F(c)$, over an appropriate range of c.

There have been many uses of the words "consonance" and "dissonance" throughout the centuries. James Tenney has identified five distinct notions, all of which implicitly apply to sounds with harmonic spectra, when used in a diatonic setting (Tenney 1988). Dissonance curves such as these generalize Mr. Tenney's fifth notion of *tonal* consonance and dissonance to situations in which possibly inharmonic sounds are played in unfamiliar tunings. Except where explicitly noted, all use of the words consonance and dissonance are in this tonal sense.

Such a mechanistic approach to consonance is not without controversy, and its use has been attacked from at least two perspectives. First, the idea of tonal dissonance cannot hope to capture the functional ideas of musical dissonance as restlessness or desire to resolve, and the linked notion of consonance as the restful place to which resolution occurs (Cazden 1945). In contrast, tonal consonance is a static notion appropriately applied only to clusters of partials. It is the responsibility of the composer to impose motion from tonal dissonance to tonal consonance, if such a motion is desired. The second attack comes from certain experiments in psychoacoustics that address the relevance of beats and roughness to perceptions of intonation. Among these, Douglas Keislar (Keislar 1991) examines musicians' preferences to various Just and tempered thirds and fifths by manipulating the partials of the sounds in a patterned way. Mr. Keislar concludes that beating is not a significant factor in intonation, but other studies (Vos 1988), using different techniques, have found the opposite.

This article takes a pragmatic approach, whereby the implications of the tonal consonance theories are pursued to their logical musical conclusions. These conclusions are in the form of predictions of the ways that certain kinds of timbres and scales will relate. The sound examples are included so that readers may judge the validity of the predictions.

Mappings between Spectra

A spectral mapping is defined to be a transformation from a set of n partials $s_1, s_2, \ldots, s_n$ (called the source spectrum) to the partials $d_1, d_2, \ldots d_n$ of the destination spectrum, for which $T(s_i) = d_i$ for all i. Suppose that an N-point discrete Fourier transform (DFT) (or a fast Fourier transform, FFT) is used to compute the spectrum of the original sound, resulting in a complex valued vector X. The mapping T is applied to X (which presumably has partials at or near the s_i), and the result is a vector $T(X)$, which represents a spectrum with partials at or

Figure 5. Schematic representation of a spectral mapping. The first nine partials of a harmonic source spectrum are mapped into a non-harmonic destination spectrum with partials at f,2.1f,2.9f,3.8f,5.4f,5.8f,7f, 8.4f, and 8.9f. The spectrum of the original sound (from the G string of a guitar with fundamental at 194 Hz) is transformed by the spectral mapping for compatibility with the destination spectrum. The mapping changes the frequencies of the partials while preserving both amplitudes (shown) and phases (not shown).

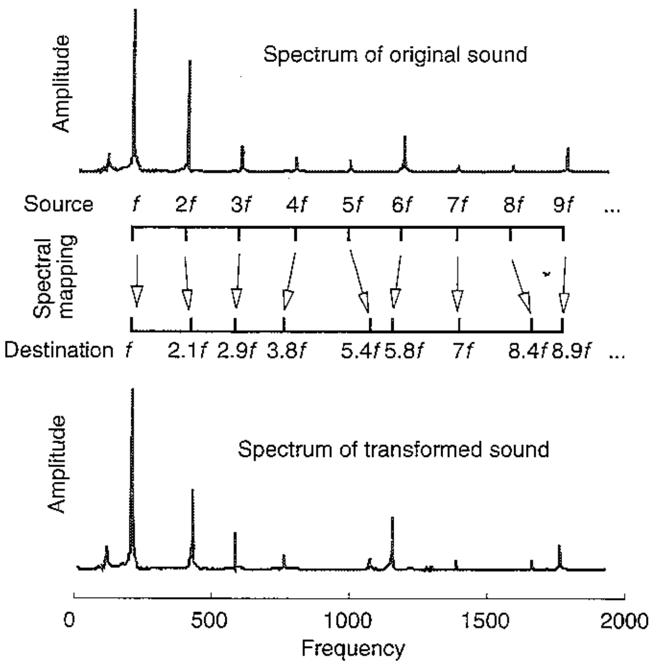

near the d_i. This is shown schematically in Figure 5 for an arbitrary destination spectrum.

The simplest T is a straight-line transformation:

$$T(s) = \left(\frac{d_{i+1} - d_i}{s_{i+1} - s_i}\right)s + \left(\frac{d_i s_{i+1} - d_{i+1} s_i}{s_{i+1} - s_i}\right),$$

where

$$s_i \leq s \leq s_{i+1}.$$

Smoother curves, such as parabolic or spline interpolations, can be readily used, but problems arise with such direct implementations, owing to the quantization of the frequency axis inherent in any digital representation of the spectrum. For instance, if the slope of T is significantly greater than unity, then certain elements of $T(X)$ will be empty. More seriously, if the slope of T is significantly less than unity, then more than one element of X will be mapped into the same element of $T(X)$, causing an irretrievable loss of information. In addition, it is not obvious how to sensibly combine the relevant terms.

A better way to think of the spectral mapping procedure is as a kind of resampling, in which the information contained between the frequencies s_i and s_{i+1} is resampled to occupy the frequencies d_i to d_{+1}. We use a standard polyphase implementation with an anti-aliasing low-pass finite-impulse-response (FIR) filter incorporating a Kaiser window. The examples filter ten terms on either side of x_i, and use $\beta = 5$ as the window-design parameter.

Figure 6. Resampling causes asymmetries in the transformed spectrum, which may cause audible anomalies.

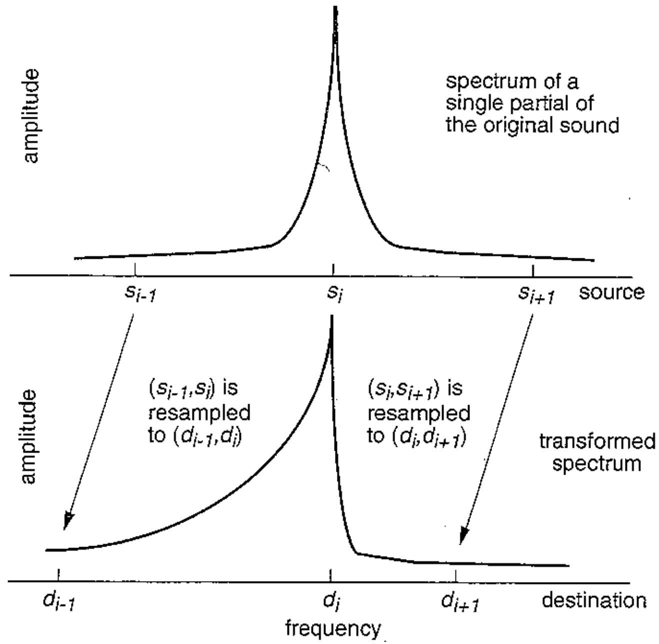

One presumption underlying spectral mappings is that the most important information (the partials of the sound) is located at or near the s_i, and is to be relocated near the d_i while being kept as intact as possible. Figure 6 shows an exaggerated view of what occurs to a single partial when performing a straightforward resampling with a non-unity spectral map, T. In essence, the left half of the spectrum becomes asymmetric from the right half, and the transformed spectrum no longer represents a single sinusoid. This is a kind of nonlinear distortion that can produce audible artifacts.

One way to reduce this distortion is to choose a window of width $2w$ about the s_i that is mapped identically to a window of the same width about d_i. The remaining regions, between $s_i + w$ and $s_{i+1} - w$, can then be resampled to fit between $d_i + w$ and $d_{i+1} - w$. This is shown (again in exaggerated form) in Figure 7. In this method of resampling

Figure 7. Resampling with identity windows reduces the asymmetry of the transformed spectrum.

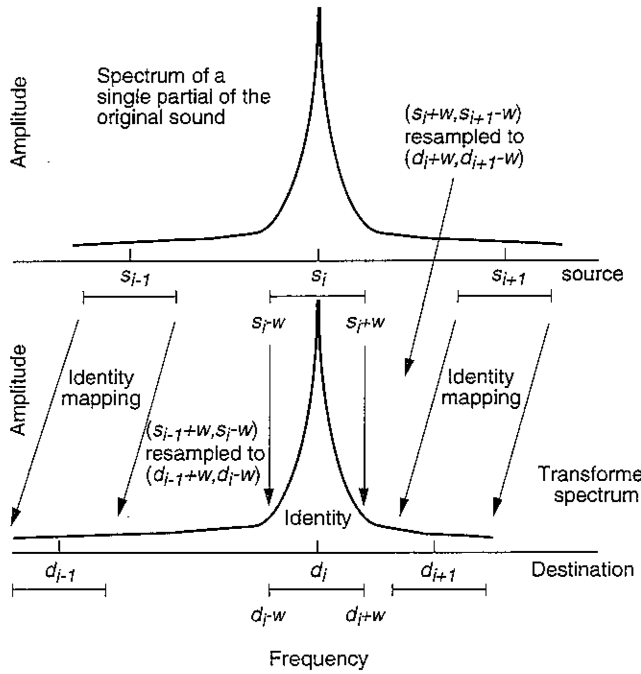

with identity window (or RIW), the bulk of the most significant information is transferred to the destination intact. Changes occur only in the less important (and relatively empty) regions between the partials. We have found window widths of about 1/3 to 1/5 of the minimum distance between partials to be the most effective in reducing the audibility of the distortion.

Spectral mappings are most easily implemented in software (or in hardware to emulate such software) in a program:

input spectrum = FFT(input signal),
mapped spectrum = T(input spectrum), and
output signal = IFFT(mapped spectrum),

where the function FFT() is the discrete Fourier transform or its fast equivalent, IFFT() is the inverse, and the RIW spectral mapping is represented by T. Other transforms, such as the wavelet or constant-Q transform (Brown 1991), might also be useful. Spectral mappings can be viewed as linear (but time-varying) transformations of the original signal. Let the signal be x, and let F be the matrix that transforms x into its DFT. Then the complete spectral mapping yields the output signal:

Figure 8. A filter-bank implementation of a spectral mapping. The input is band-pass filtered, and the signal is parameterized into n amplitude, phase, and frequency parameters. These are transformed by the spectral mapping, and the modified parameters drive n oscillators, which are summed to form the output.

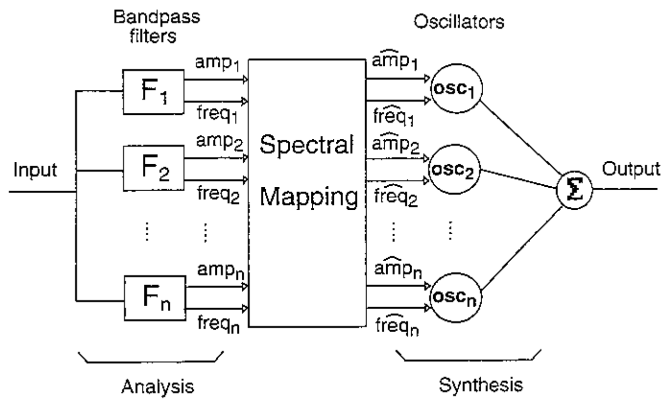

$$\hat{x} = F^{-1}TF(x)$$

where T is a matrix representation of the resampling procedure. This is clearly linear, and it is time varying because the frequencies of signals are not preserved. Often T fails to be invertible, and the original signal x cannot be reconstructed from its spectrally mapped version $\hat{x}$.

There are many possible variations of T. For instance, many instrumental sounds can be characterized using formants: fixed linear filters through which a variable excitation passes. If the original samples are of this kind, it is sensible to modify the amplitudes of the resulting spectra accordingly. Similarly, an energy envelope can be abstracted from the original sample, and in some situations it might be desirable to preserve this energy during the transformation. In addition, there are many kinds of resampling, and there are free parameters (and filters) within each kind. Trying to optimally choose these parameters is a daunting task.

It may be more computationally efficient to implement spectral mappings as a filter bank rather than as a transform, especially when processing a continuous audio signal. (A good modern approach to filter banks is available [Strang and Nguyen 1996].) This is diagrammed in Figure 8, which shows a bank of filters carrying out the analysis portion of the procedure, a spectral mapping to manipulate the spectrum parameters, and a bank of oscillators to carry out the synthesis portion. This does not change the motivation or goals of the mappings, but it does suggest an alternative hardware or software approach.

Figure 9. A square wave and its transformation into an 11-TET version. Maintaining the phase relationships among the partials helps the attack portion maintain its integrity.

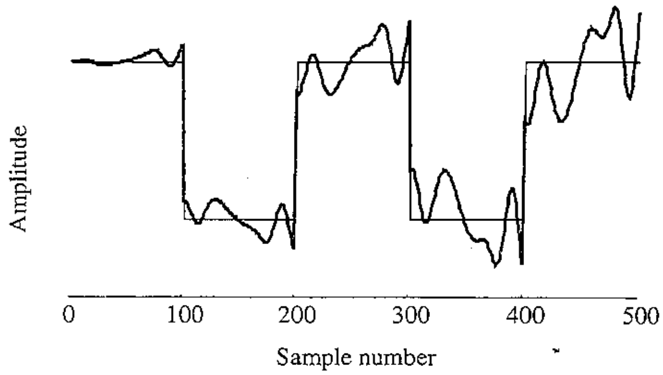

Figure 10. A 4,500-sample trumpet waveform, with the looped region indicated (a); the same waveform using a "fill with loop" rather than a "fill with zeroes" strategy to increase the length of the wave to 8 K samples (b).

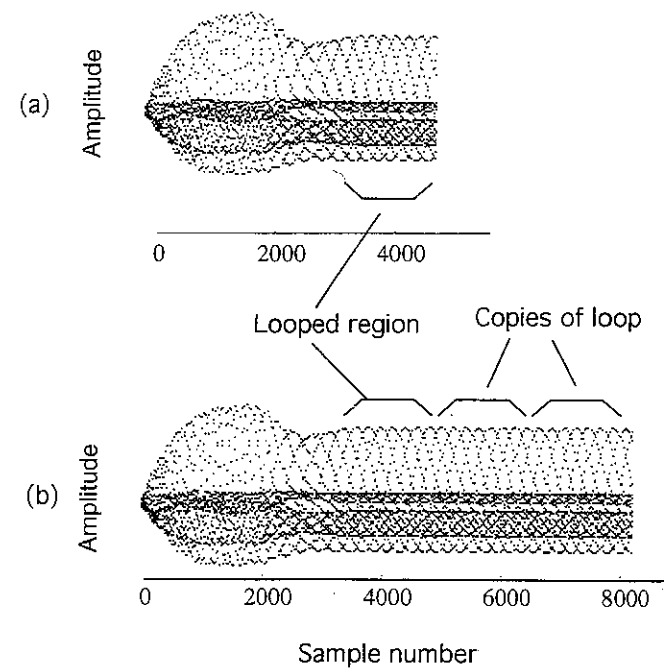

Maintaining Amplitudes and Phases

The tonal quality of a harmonic sound is determined largely by the amplitudes of its sinusoidal frequency components. In contrast, the phases of these sinusoids tend to play a small role, except in the transient (or attack) portion of the sound, where they contribute to the envelope. The transformation T is specified so as to keep each frequency component roughly matched with its original amplitude and phase. This tends to maintain the shape of the waveform in the attack portion. For example, Figure 9 shows a square wave and its transformation into the 11-TET timbre specified in Figure 2. The first few pulses are clearly discernible in the mapped waveform. Since the first few milliseconds of a sound are important in terms of the overall sound quality (Strong and Clark 1967), maintaining the initial shape of the waveform contributes to the goal of retaining the integrity of the sound.

Looping

A common practice in sample-based synthesizers is to loop sounds, which is to repeat certain portions of the waveform under user control. Periodic portions of the waveform are ideal candidates for looping. Strictly speaking, inharmonic sounds, such as those that result from transformations like the 11-TET spectral mappings, have aperiodic waveforms. Apparently, looping becomes impossible. On the other hand, the FFT induces a quantization of the frequency axis in which all frequency components are integer multiples of the frequency of the first FFT bin (for instance, about 1.3 Hz for a 32-K FFT at a 44.1-kHz sampling rate). Thus, true aperiodicity is impossible in a transform-based system. In practice, it is often possible to loop the sounds effectively using the standard assortment of looping strategies and cross-fades, though it is not uncommon for the loops to be somewhat longer in the modified waveform than in the original.

To be concrete, suppose that the original waveform contains a looped portion. A sensible strategy is to append the loop onto the end of the waveform several times, as shown in Figure 10. This tends to make a longer portion of the modified waveform suitable for looping. It is also a sensible way of filling or padding the signal until the length of the wave is an integer power of two (so that the more efficient FFT can be computed in place of the DFT). The familiar strategy of padding with zeroes is inappropriate in this application. Figure 11, for instance, shows the results of three different mappings of the 4,500-sample trumpet waveform of Figure 10. Calculating the DFT and applying the 11-TET spectral mapping of Figure 2 yields the waveform in Figure 11a. This version consists primarily of the attack portion of the waveform, and

Figure 11. Spectrally mapped versions of the trumpet waveform in Figure 10: using a DFT of the wave in Figure 10a (a); using an FFT of the wave in Figure 10a and the "fill with zeroes" strategy (b); and using an FFT of the wave in Figure 10b and the "fill with loop" strategy (c). Version (c) gives a longer, steadier waveform with more opportunity to achieve a successful loop.

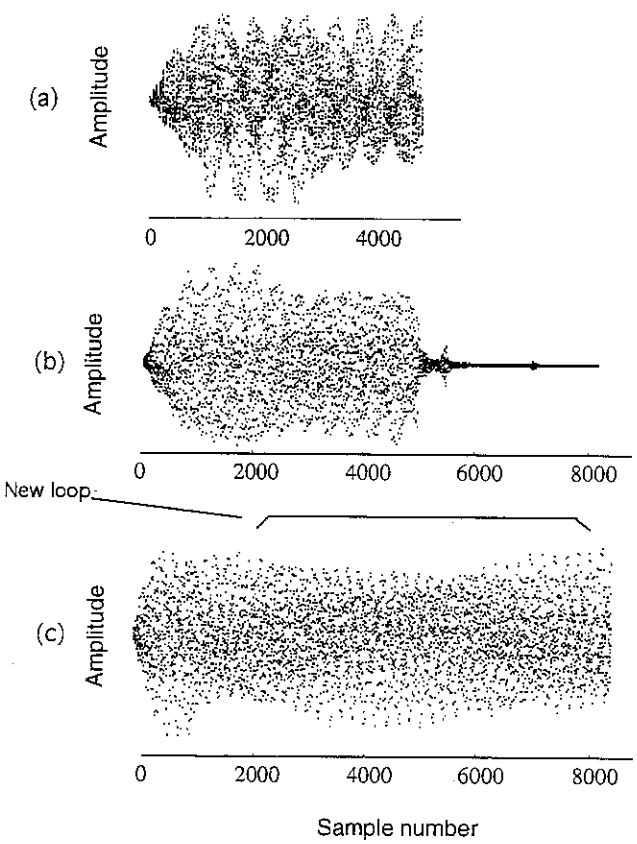

is virtually impossible to loop without noticeable artifacts. An alternative is to extend the waveform to 8K (i.e., 8192) samples by filling in with zeroes. This allows use of the FFT for faster computation, but the resulting stretched waveform of Figure 11b is no easier to loop than the signal in Figure 11a. A third alternative is to repeatedly concatenate the original looped portion until the waveform reaches the desired 8K length. The resulting stretched version, shown in Figure 11c, contains a longer sustain portion, and is correspondingly easier to loop.

Separating Attack from Loop

The attack portion of a sound is often quite different from the looped portion. The puff of air as the flute chiffs, the "blat" of the trumpet's attack, or the scrape of the violin's bow are quite different from the steady-state sounds of the same instruments. Indeed, it has been shown (Strong and Clark 1967) that it can often be difficult to recognize instrumental sounds when the attack has been removed.

Naive application of a spectral mapping would transform the complete sampled waveform simultaneously. Because the Fourier transform has poor time-localization properties, this can cause a "smearing" of the attack portion over the whole sample, with noticeable side effects. First, the smearing can sometimes be perceived directly as artifacts: a high tingly sound or noisy grating that repeats irregularly throughout the looped portion of the sound. Second, because the artifacts are nonuniform, they make it even more difficult to create a good loop of the mapped sound.

Thus, a good idea when spectrally mapping sampled sounds (for instance, those with predefined attack and loop segments) is to map the attack and the loop portions separately, as shown in Figure 12. The resulting pieces can then be pasted back together using a simple cross-fade. This tends to maintain the integrity of the attack portion (it is shorter, and less likely to suffer from phase and smearing problems), and reduce artifacts occurring in the steady state.

Often, a complete sampled instrument contains several different waveforms sampled in different pitch ranges and at different dynamic ranges. The creation of a spectrally mapped version should map each of these samples, and then assign them to the appropriate pitch or dynamic level. In addition, it is reasonable to impose the same envelopes and other performance parameters such as reverb, vibrato, etc., as were placed on the original samples, since these will often have a significant impact on the overall perception of the sound quality.

Examples

This section presents examples of spectral maps in which the integrity of the original sounds is maintained, and others in which the perceptual identity of sounds is lost. Examples include instruments mapped into a spectrum consonant with 11-TET, instruments mapped into the spectrum of a drum,

Figure 12. Transforming the attack and steady-state (looped) portions separately helps to maintain the tonal integrity of the sound.

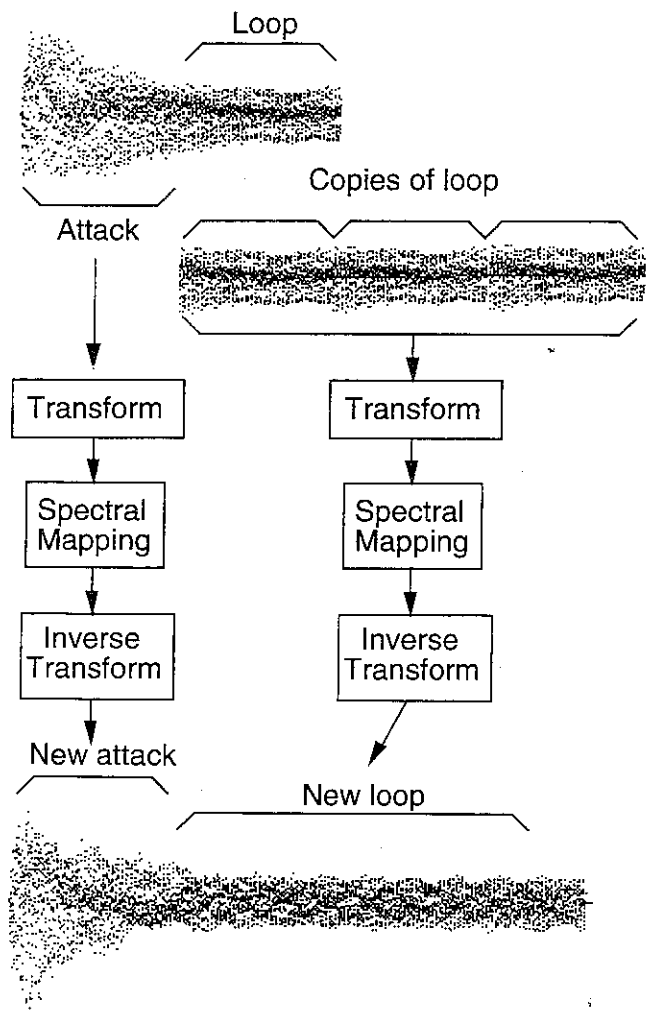

and a cymbal sound mapped so as to be consonant with 12-TET.

Timbres for Eleven-Tone Equal Temperament

Familiar harmonic sounds may be dissonant when played in 11-TET, because minima of the dissonance curve occur far from the desired scale steps, as in Figure 1. By using an appropriate spectral mapping, harmonic instrumental timbres can be transformed into 11-TET versions with minima at many of the 11-TET scale steps, as shown in Figure 2. These can be used to play consonantly in an 11-TET setting. The mapping used to generate the tones in the sound example maps a set of harmonic partials at

$$f, 2f, 3f, 4f, 5f, 6f, 7f, 8f, 9f, 10f, 11f$$

to

$$f, a^{11}f, a^{17}f, a^{22}f, a^{26}f, a^{28}f, a^{31}f, a^{33}f, a^{35}f, a^{37}f, a^{38}f$$

where $a = \sqrt[11]{2}$ and f is the fundamental of the harmonic tone. All frequencies between these values are mapped using the RIW method.

The waveforms were taken from commercially available sample CD-ROMs, and transferred to a computer running a MATLAB program that performed the spectral mappings. After looping (which was done manually, with the help of Infinity looping software), the modified waveforms were sent to an Ensoniq ASR-10 sampler. The performances were sequenced and recorded to digital audio tape. In all cases, the same performance parameters (filters, envelopes, velocity sensitivity, reverberation, etc.) were applied to the spectrally mapped sounds as they were used in the original samples.

Sound Example 1 contains several different instrumental sounds that alternate with their 11-TET versions:

(a) Harmonic trumpet compared to 11-TET trumpet
(b) Harmonic bass compared to 11-TET bass
(c) Harmonic guitar compared to 11-TET guitar
(d) Harmonic pan flute compared to 11-TET pan flute
(e) Harmonic oboe compared to 11-TET oboe
(f) Harmonic Moog synthesizer compared to 11-TET Moog synthesizer
(g) Harmonic phase synthesizer compared to 11-TET phase synthseizer

The instruments are clearly recognizable after mapping into their 11-TET counterparts. There is almost no pitch change caused by this spectral mapping, probably because some partials are mapped higher while others are mapped lower. Indeed, the third partial is mapped lower than its harmonic counterpart (2.92 versus 3), but the fifth is higher (5.14 versus 5). Similarly, the sixth is lower (5.84 versus 6), but the seventh is higher (7.05 versus 7).

Perhaps the clearest change is that some of the samples have acquired a soft, high-pitched inharmonicity. It is hard to put words to this, but we try. In (a), it may almost be called a whine; (b) has a slight lowering of the pitch, as well as a feeling that something else is attached; and (c) has acquired a high jangle in the transition. It is hard to pinpoint any changes in (d) and (f). In (e), as in some of the others, it becomes easier to hear out one of the partials in the mapped sound, giving it a kind of diminished feel; and the vibrato of (g) appears to have changed slightly, but the sound is otherwise intact.

Despite the fact that all sounds were subjected to the same mapping, the perceived changes differ somewhat from sample to sample. This is likely an inherent aspect of spectral mappings. For instance, the bass has a strong third partial and a weak fifth partial compared to the other sounds. Since the third partial is mapped down in frequency, it is reasonable to hypothesize that this causes the lowering in pitch. Because the fifth partial is relatively weak, it cannot compensate, as might occur in other sounds. Similarly, differing amplitudes of partials may cause the varying effects perceivable in (a) through (g).

Such perceptual changes may be owing to the way that inharmonicities are perceived. For instance, the question of how much detuning is needed before an inharmonic partial causes a sound to break into two sounds, rather than remain fused into a single percept, has been examined (Moore, Peters, and Glasberg 1985). Alternatively, the changes may be owing to artifacts created by the spectral mapping procedure itself. For instance, other choices of filters, windows widths, etc., may generate different kinds of artifacts. Certainly it is true that by doing the mapping foolishly, one can introduce strange effects. This was the major reason for separating the attack and looped portions of the sounds in the mapping procedure—separation reduces the smearing artifacts significantly.Isolated sounds do not paint a very good picture of their behavior in more complex settings. A short sequence of major chords are played:

(h) Harmonic oboe in 12-TET
(i) Spectrally mapped 11-TET oboe in 12-TET

As before, the individual sounds have only a small pitch shift. The striking difference between (h) and (i) may therefore be of interest to those who hold that consonance, dissonance, and the "out-of-tune" percept are caused primarily by pitch or interval relationships, and not by the structure of the partials of a sound. While (i) is not out of tune, it may be said to be "out of spectrum" or "out of timbre," in the sense that the partials of the sound interfere when played at certain intervals (in this case, the 12-TET major third and fifth).

The next segments contain 11-TET dyads formed from scale steps 0–6 and 0–7, and culminate in a chord composed of scale steps 0–4–6:

(j) harmonic oboe in 11-TET
(k) spectrally mapped 11-TET oboe in 11-TET

Examples (j) and (k) reverse the situation from (h) and (i). Because of the extreme unfamiliarity of the intervals (observe that 11-TET scale steps 4 and 6 do not lie close to any 12-TET intervals), the situation is perhaps less clear, but there is a readily perceivable beating of the 0–4–6 chord in (j) that is absent from (k). Thus, after acclimation to the intervals, (k) appears arguably less out-of-spectrum than (j).

Recall that it is possible to hear out the third partial of the sound used in (i) and (k). This suggests that the triads of (i) could be heard as dissonant six-note clusters (with the three extra notes arising from the detuned, unfused third partials), providing an alternative explanation of the dissonance in (i). The comparative smoothness of the three-note chords in (k), however, argues against such an interpretation; since these notes remain well fused.

Isolated chords do not show clearly what happens in genuine musical contexts. An excerpt of the piece *The Turquoise Dabo Girl* is played two ways:

(l) In 11-TET with all sounds spectrally mapped
(m) In 11-TET with the original harmonic sounds

The out-of-spectrum effect of (m) is far more dramatic than the equivalent isolated chord effect of (j), illustrating that the more musical the context, the more important (rather than the less important) a proper matching of the tuning with the spectrum of the sound becomes.

The excerpt from *The Turquoise Dabo Girl* may also suggest that many of the kinds of effects normally associated with (harmonic) tonal music can occur, even in strange settings such as 11-TET (which is often considered among the hardest tuning systems in which to play tonal music). Observe that many of the subtle oddities in the mapped timbres, as noted in (a) through (g) of the sound example, seem to disappear when contextualized. Even with careful listening, it is difficult (impossible?) to hear the inharmonicities or artifacts that were so clear when presented independently; except for the percussion, all the timbres used in *The Turquoise Dabo Girl* were demonstrated in (a) through (g). This may be owing to a simple masking of the artifacts, or to a kind of capture effect, in which the artifact or inharmonicity of one note is captured (or streamed with) other notes, and thus becomes part of the music. If this is the case, then the inharmonicities become a "feature" of the 11-TET setting, rather than a problem to be avoided.

Spectrum of a Drum

The spectral mapping of the previous example changes the partials only moderately. In contrast, mapping from harmonic tones into the spectrum of a drum such as a tom-tom changes the partials dramatically. The extreme inharmonicity of the sample is shown in Figure 13, and the severe mapping is readily heard as drastic changes in the tonal quality and pitch of the transformed instruments. A harmonic spectrum at g, $2g$, $3g$, $4g$, $5g$ is mapped to d, $1.67d$, $2.46d$, $3.2d$, $3.8d$ (which is precisely 245, 410, 603, 786, 934 for $d = 245$) using the RIW spectral mapping. Of the guitar, bass, trumpet, and flute, only the flute is recognizable, and even this is not without drastic audible changes. One listener remarked that the transformed sounds were "glassy—like a fingernail scratching across a glass surface." This description makes a certain amount of physical sense, since glass surfaces and drum heads are both two-dimensional vibrating surfaces.

Sound Example 2 contains several different instruments and their transformations into the tom-tom spectrum shown in Figure 13:

(a) Harmonic flute compared to tom-tom flute
(b) Harmonic trumpet compared to tom-tom trumpet
(c) Harmonic bass compared to tom-tom bass
(d) Harmonic guitar compared to tom-tom guitar

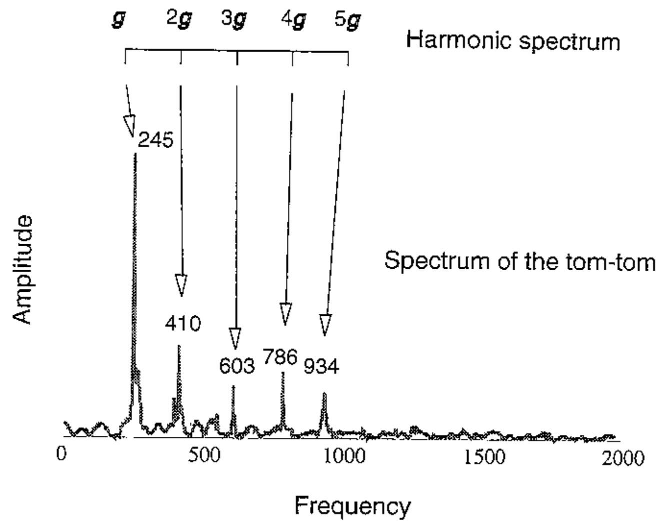

Figure 13. A harmonic spectrum with fundamental g is mapped into the tom-tom spectrum.

Clearly, this spectral mapping causes a large change in the character of the sounds. As before, it is unclear what aspects of the resulting changes are owing to the way inharmonic sounds are perceived, and what may be associated with the details of the spectral mapping procedure. For instance, each of the sounds undergoes a pitch change, but the pitch change is different for each sound. Presumably this is because the partials of the mapped sounds inherit the amplitudes of the original sounds. Likely, the ear picks out different "harmonic templates" (Moore, Glasberg, and Peters 1985) for each arrangement of amplitudes.

Again, it is hard to describe in words the kind of effects perceived. While (a) has a noticeable pitch change, it still sounds something like a flute. The trumpet undergoes a huge pitch change, and gains a kind of glassy texture. The single note of the bass becomes a "minorish" chord, and the guitar pluck also gains a chord-like sound along with jangly artifacts.

Though the transformed timbres do not sound like the instruments from which they were derived, they are not necessarily useless. The next audio track is an excerpt from *The Glass Lake*, which illustrates the transformed instruments (a) through (d) played in the related scale, with steps defined by

Figure 14. The dissonance curve for the tom-tom spectrum has the 11-note related scale, 1, 1.19, 1.3, 1.51, 1.55, 1.63, 1.98, 2.35, 2.46, 3.21, and 3.83, which covers a little less than 2 octaves.

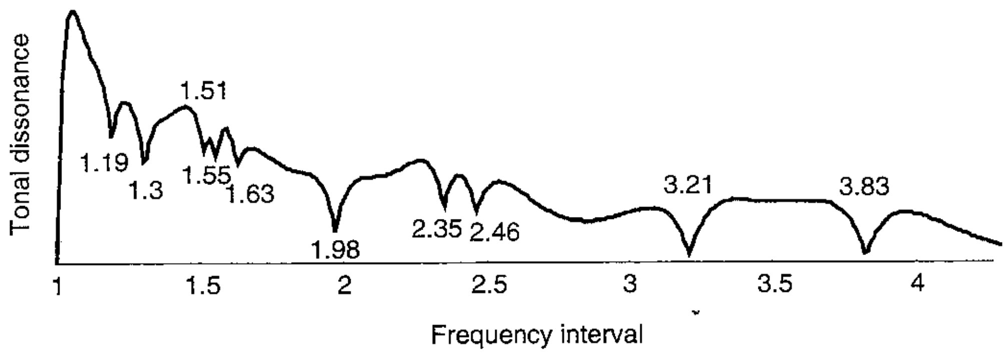

the dissonance curve of Figure 14. This scale supports perceptible chords, though they are not necessarily composed of familiar intervals. The piece is thoroughly *xenharmonic*, a word coined by Ivor Darreg (Darreg 1975), meaning music unlike anything possible in a traditional 12-TET setting.

A Harmonic Cymbal

The previous examples transformed familiar harmonic timbres into various unfamiliar timbres and scales. The third and final example uses spectral mappings to transform familiar inharmonic sounds into sounds maximally consonant with harmonic spectra. The spectrum of a cymbal contains many peaks spread irregularly through the whole audible range. For the chosen cymbal sample, the $N = 35$ largest peaks (labeled p_i, $i = 1,2,\ldots,N$) were fit to a nearby harmonic template by finding the fundamental f that minimizes

$$\sum_{i=1}^{N} (p_i - if)^2 .$$

The solution is $f = \frac{\sum i p_i}{\sum i^2}$, and the p_i and the if define the source and destination of the spectral mapping. The transformed sound retains some of the noisy character of the original cymbal strike, but has become noticeably more harmonic and has inherited the pitch associated with the fundamental f. The two brief segments of the third sound example are:

(a) The original sample contrasted with the spectrally mapped version

(b) A simple chord pattern played with the original sample, and then with the spectrally mapped version

The transformed instrument supports both chord progressions and melodies, even though the original cymbal strike does not.

Discussion

The discussion begins with a consideration of various aspects of timbral change, and then suggests additional perceptual tests that might further validate (or falsify) the use of spectral mappings in inharmonic musical applications. Several types of inharmonic musical modulation are discussed.

Robustness of Sounds under Spectral Maps

How far can partials be mapped before the sound loses cohesion or otherwise changes beyond recognition? It is clear from even a cursory listen that small perturbations in the locations of the partials (i.e., mappings that are not too distant from identities) have little effect on the overall tonal quality of the sound. Flutes and guitars in 11-TET timbres retain their identities as flutes and guitars. The consistency of such sounds through various spectral mappings argues that perceptions of tonal quality are not primarily dependent on the precise frequency ratios of the partials. Rather, there is a band in which the partials may lie without affecting the "fluteness" or "guitarness" of the sound. Equivalently, the partials of such a sound can undergo a wide variety of mappings without significantly affecting its inherent tonal gestalt. Sounds with stretched spectra have been used to investigate when the sounds fuse into a single entity and when they fission into separate partials (Cohen 1984).

Aside from the sounds demonstrated here, the author has spectrally mapped a variety of about 50 sounds into several different destination spectra, including stretched timbres with stretch factors from 1.5–3.0 (detailed discussions of stretched timbres are available [Slaymaker 1968; Mathews and Pierce 1980]), spectra designed to be consonant with n-TET for $n = 8, \ldots, 19$, and a variety of destination spectra derived from objects such as a tom-tom, a bell, a metal wind chime, and a rock. Overall, there is wide variation in the robustness of individual sounds. For instance, the sound of a tom-tom or cymbal survives translation through numerous mappings, some of them quite drastic. Only the flute still retains any part of its tonal identity when mapped into the tom-tom spectrum of Figure 13. Sounds such as the guitar and clarinet can be changed somewhat without losing their tonal qualities, surviving the transformation into the n-TET spectra but not into the more drastic tom-tom spectrum. Other sounds, like the violin, are quite fragile, and are unable to survive even modest transformations. Thus, not all mappings preserve the perceptual wholeness of the original instruments, and not all instruments are equally robust to spectral mappings.

Using the RIW spectral mapping technique of the previous sections, the attack portion is mapped separately from the looped portion, which tends to maintain the character of the attack. Since the envelope—and other performance parameters—are also maintained, changes in timbral quality are likely caused primarily by changes in the spectrum of the steady-state (looped) portion of the sound.

As a general rule, the change in timbral quality of instruments with complex spectra appears to be greater than that of instruments with relatively simple spectra. The flute and tom-tom have fairly simple spectra (only four or five spectral peaks), and are the most robust of the sounds examined, retaining their integrity even under extreme spectral maps. Sounds with an intermediate number of significant spectral peaks, such as the guitar, bass, and trumpet, survive transformation through modest spectral mappings. In contrast, sounds like the violin and oboe, which have very complex spectra, are the most fragile sounds encountered, since they were changed significantly by a large variety of spectral mappings.

Perhaps the most familiar spectral mapping is transposition, which modulates all partials up or down by a specified amount. As is well known, pitch transposition over a large interval leads to distortions in tonal quality. For instance, voices raised too far in pitch undergo "munchkinization." It should not be surprising that other spectral maps also have perceptual side effects.

Timbral Change

Is there a way to quantify the perceived change in a tone?

Even a pure sine wave can change in timbre. Low-frequency sine waves are "soft" or "round," while high-frequency sine waves are "shrill" or "piercing." Thus one aspect of timbral change is frequency dependent, which may explain timbral changes that are caused by transposition. A second element of timbral change is the familiar notion that tonal quality changes as the amplitudes of the harmonically related partials change. This idea is the likely explanation for the timbral differences between, say, a clarinet and a flute playing the same pitch. Spectral mappings suggest a third aspect of timbral change, namely, that modification of the internal structure of a sound (i.e., a change in the intervals between the partials) causes perceptual changes in the sound. Depending on the spectral mapping (and the partials of the sound that is mapped), this may involve the introduction—or removal—of inharmonicity.

Clearly, any measure of timbral change must account for all three mechanisms. It is reasonable to hypothesize that perceptions of change are:

Proportional to the amount of transposition
Proportional to the change in amplitude of the partials, and
Proportional to the change in the spectra under the mapping T

Table 1. Mappings for spectra consonant with *n*-tone equal-tempered scales.

Number of steps per octave	Partial number 0	1	2	3	4	5	6	7	8	9	10	11
5	1	5	8	10	12	13	14	15	16	17	17	18
6	1	6	10	12	14	16	17	18	19	20	21	22
7	1	7	11	14	16	18	20	21	22	23	24	25
8	1	8	13	16	19	21	22	24	25	27	28	29
9	1	9	14	18	21	23	25	27	29	30	31	32
10	1	10	16	20	23	26	28	30	32	33	35	36
11	1	11	17	22	26	28	31	33	35	37	38	39
12	1	12	19	24	28	31	34	36	38	40	42	43
13	1	13	21	26	30	34	36	39	41	43	45	47
14	1	14	22	28	33	36	39	42	44	47	48	50
15	1	15	24	30	35	39	42	45	48	50	52	54
16	1	16	25	32	37	41	45	48	51	53	55	57
17	1	17	27	34	39	44	48	51	54	56	59	61
18	1	18	29	36	42	47	51	54	57	60	62	65
19	1	19	30	38	44	49	53	57	60	63	66	68
20	1	20	32	40	46	52	56	60	63	66	69	72
21	1	21	33	42	49	54	59	63	67	70	73	75
22	1	22	35	44	51	57	62	66	70	73	76	79
23	1	23	36	46	53	59	65	69	73	76	80	82

Explanation: The table shows how the fundamental and the next 11 higher partials are mapped to scale steps of an *n*-tone equal temperament, for *n* from 5 to 23. For *n*-TET timbres, let $a = \sqrt[n]{2}$. Define the timbre by the spectral map that takes $f[1,2,3,4,5,6,7,8,9,10,11,12]$ to $f[1,a^{p_1},a^{p_2},a^{p_3},a^{p_4},a^{p_5},a^{p_6},a^{p_7},a^{p_8},a^{p_9},a^{p_{10}},a^{p_{11}}]$, where f is the fundamental of the harmonic sound, and p_i represents the partial numbers in the appropriate row.

Some general trends are suggested. Frequency shifts in a uniform direction (such as those of a stretched map, or in a transposition mapping) may not be as damaging to timbral integrity as those that shift some partials higher and others lower (such as the 11-TET mapping). Sounds with greater spectral complexity, like the violin, seem to undergo larger perceptual changes than simpler sounds, like the flute.

To minimize the amount of perceptual change, the mapping T should be defined so that all slopes are as close to unity as possible, that is, so that the mapping is as near to the identity as possible, while still consistent with the desire to minimize dissonance. For instance, when specifying timbres for *n*-tone, octave-based equal temperaments, it is reasonable to place the partials at frequencies that are multiples of $c = \sqrt[n]{2}$ to insure that local minima of the dissonance curve occur at the appropriate scale steps. A good rule of thumb is to define the mapping by transforming partials to the nearest power of c. Thus an 11-TET timbre may be specified by mapping the first harmonic to c^{11} $(= 2)$, the second harmonic to c^{17} (≈ 3), the third harmonic to c^{22} $(= 4)$, etc., as given in Figure 2. Analogous definitions of timbres for scales between 5–23 are given in Table 1.

Related Perceptual Tests

One way to investigate timbral change is to gather data from listener tests and apply a multidimensional scaling technique, as Reinier Plomp did (Plomp 1970). For instance, John Grey and John Gordon (Grey and Gordon 1978) swapped the temporal envelopes of the harmonics of instrumental tones, and tested listeners to determine how different the modified sounds were from the originals. Such a study could be conducted for sounds formed from various spectral mappings, giving a quantitative way to speak about the degree to which sounds retain their integrity under spectral mappings. The clustering technique used by Mr. Grey and Mr. Gordon found three dimensions to the sounds, which were interpreted as a spectral dimension, a dimension that represents the amount of change in the spectrum over the duration of the tone, and a dimension determined primarily by the explosiveness or abruptness of the attack. Sounds that undergo modest spectral mappings are likely to change in the first dimension and to remain more or less fixed in the latter two. Instrumental sounds that are mapped so as to be consonant with 11-TET (say) sound far more like the original instrumental samples than they sound like each other. An interesting question is whether the spectrally mapped sounds might cluster into a new dimension.

The sound examples of this article suggest caution in the interpretation of results such as the above, which rely on listening tests that lack musical context. Taken in isolation, 11-TET-mapped trumpet sounds are very similar to harmonic trumpet sounds, and thus should cluster nicely with harmonic trumpet timbres. But in a 12-TET musical context, the 11-TET trumpet will sound out of tune, for instance, when it is played in concert with harmonic instruments. Similarly, the harmonic trumpet will sound out of tune when played in 11-TET in an ensemble of 11-TET instruments. In this contextual sense, similarly mapped instruments should tend to cluster separately from harmonic instruments.

Increasing Consonance

Much of the current xenharmonic music is written in just intonations and other scales that are closely related to harmonic timbres. Many of the most popular equal temperaments (7, 17, 19, 21, and 31, for example) contain intervals that closely approximate the intervals of scales related to harmonic timbres. There is, of course, a body of work in tunings such as 11-TET that are unrelated to harmonic timbres. Some of these pieces revel in their dissonance, emphasizing just how strange xenharmonic music can be. Other composers have sought to minimize the dissonance. Bregman (1990) reports that the dissonance between a pair of sounds can be reduced by placing them in separate perceptual streams. This implies that musical parts that would normally be dissonant can sometimes be played without dissonance if the listener can be encouraged to hear the lines in separate perceptual streams. Skilled composers can coax sounds into streaming or fusing in several ways, including large contrasts in pitch, tone color, envelope, and modulation. These techniques have not gone unexploited in xenharmonic music, and they can be viewed as clever ways of finessing the problem of dissonance. They are solutions at the compositional level.

The spectral mappings of this article provide an alternative answer at the timbral level. It is possible to compose tonally consonant music in virtually any tuning by redesigning the spectra of the instruments so that their timbres are related to the desired scale. Of course, it is not always desirable to maximize consonance. Rather, the techniques suggested here are means for achieving increased contrast in the consonance and dissonance of inharmonic sounds when played in nonstandard tunings. Utilizing spectra that have dissonance curves with minima at the scale steps allows these intervals to be as consonant as possible, thus giving the composer greater control over the perceived consonance. (It is typically easy to increase the dissonance by playing more notes or more tightly clustered chordal structures; the hard part is to decrease the dissonance without removing notes or simplifying the spectra). That this is possible even for notorious scales such as 11-TET expands the range of possible moods or feelings in these scales. Similarly, consonance is only a part of the musical landscape. Even ignoring essential aspects such as rhythm, pattern, and modulation, it is certainly not desirable to simply maximize consonance. Indeed, silence is the most tonally consonant "sound."

Consonance-Based Modulations

Morphing from one set of related scales and timbres to another is a new kind of musical modulation. This might consist of a series of passages, each with a different tuning and timbre. For instance, a piece might begin with harmonic timbres in 12-TET, move successively through 2.01, 2.02, ... , 2.1 stretched octaves, and then return to harmonic sounds for the finale. Such consonance-based modulation can be extremely subtle, as in the modulation from 2.01–2.02 stretched. It can also be extremely dramatic, since it involves the complete timbre of the notes as well as the scale on which the notes are played. Alternatively, such modulations might move between various n-TET structures. By carefully choosing the timbres, the same instruments can play in different tunings, and the dissonance can be tightly controlled.

It is also possible to morph from one spectrum to another in the evolution of a single sustained sound. This can be done by partitioning the waveform into a series of overlapping segments, calculating a Fourier transform for each segment, applying a different spectral mapping to each segment, and then rejoining the segments. Such consonance-based morphing of individual tones can be used to smooth transitions from one tuning/timbre pair to another, or it can be used directly as way to control timbral evolution.

At a point when the mapping becomes too severe, individual notes can lose cohesion and fission into a cluster of individually perceptible partials. Albert Bregman has suggested several tone-manipulation methods that can be used to control the degree to which inharmonic tones fuse (Bregman 1990). Simultaneous onset times and common fluctuations in amplitude or frequency contribute to fusion, while independent fluctuations tend to promote fission. These can be readily used as compositional tools to achieve a desired amount of tonal coherence. For instance, a sound can be modulated from perceptual unity into a tonal cluster and then back again by judicious choice of such tools. In *Inharmonique* Jean-Claude Risset explored this type of modulation using an additive-synthesis approach (Risset 1987). Because spectral maps directly affect the amount of a tone's inharmonicity, a series of spectral maps can be used to approach or cross the boundaries of tonal fusion in a controlled manner.

Another form of modulation involves the boundary between melody and rhythm. For instance, when the cymbal of Sound Example 3 is played using the original sample, it is primarily useful as a rhythm instrument. When the same sound is transformed into a harmonic spectrum, it can support melodies and harmonies. Consider a series of spectral mappings that smoothly interpolate between these two. At some point the melodic character must disappear and the rhythmic character predominate. Careful choice of spectral mappings allows the composer to deliberately control whether the sound is perceived as primarily unpitched and rhythmic or as primarily pitched and harmonic, and to modulate smoothly between the two extremes.

Conclusions

Most of the sounds of the orchestra (minus certain members of the percussion family) and many of the common sounds of electronic synthesizers have harmonic spectra. Because the tonal quality of sounds is not destroyed under many kinds of spectral mappings, entire orchestras of sounds can be created from inharmonic spectra. These sounds can retain much of the character of the sounds from which they were derived, though they are not perceptually identical. For example, 11-TET sounds were created that clearly reflect their origins as guitar and flute samples. These are clearly perceived as instrumental in nature, and can be played consonantly in an 11-TET setting.

It is not necessary to abandon the familiar sound qualities of conventional musical instruments to play in unusual scales. The spectral mappings of this article provide a way to convert a large family of well-established, musically useful sounds into timbres that can be played consonantly in a variety of scales. Musical tastes change slowly, and it can be difficult for audiences to appreciate music in which everything is new. The creation of familiar sounds that can be played in unusual scales may help to ease the transition to xenharmonic musics.

Alternatively, extreme spectral mappings can be used to generate genuinely new sounds using familiar instrumental tones as raw material. When played in the related scales, these tend to retain familiar musical features such as consonance, even though the timbres and intervals of the scales are unfamiliar.

Spectral mappings can also be used to transform inharmonic sounds (such as certain cymbals) into harmonic equivalents. Using these sounds, it is possible to play familiar chord patterns and melodies using this new class of harmonic percussion instruments. Consonance-based spectral mappings make it possible to explore a full range of tonal possibilities for many different spectra.

Acknowledgments

The author would like to thank Tom Staley for extensive discussions about tuning, tonality, and temperament. Tom also helped write and perform the piece *The Glass Lake*.

References

Bregman, A. 1990. *Auditory Scene Analysis*. Cambridge, Massachusetts: MIT Press.

Brown, J. 1991. "Calculation of a Constant-Q Spectral Transform." *Journal of the Acoustical Society of America* 89:425–434.

Carlos, W. 1987. "Tuning: At the Crossroads." *Computer Music Journal* 11(1):29–43.

Cazden, N. 1945. "Musical Consonance and Dissonance: A Cultural Criterion." *Journal of Aesthetics* 4:3–11.

Cohen, E. A. 1984. "Some Effects of Inharmonic Partials on Interval Perception." *Music Perception* 1(3):323–349.

Darreg, I. 1975. *Xenharmonic Bulletin*, vols. 1–15. San Diego, California: Darrog (privately printed).

Dudley, H. 1939. "Remaking Speech." *Journal of the Acoustical Society of America* 11:169.

Flanagan, J. L., and R. M. Golden. 1966. "Phase Vocoder." *Bell System Technical Journal* (November): 1493–1509.

Grey, J. M., and J. W. Gordon. 1978. "Perceptual Effects of Spectral Modifications on Musical Timbres." *Journal of the Acoustical Society of America* 65(5):1493–1500.

Grey, J. M., and J. A. Moorer. 1977. "Perceptual Evaluation of Synthesized Musical Instrument Tones." *Journal of the Acoustical Society of America* 62:454–462.

Helmholtz, H. 1863. *Die Lehre vou der Tonempfindungen als physiologische Grundlage für die Theorie der Musik*, 1st ed. Braunschweig: vieweg. 4th ed. (1877) trans. by A. J. Ellis under the title *On the Sensations of Tone as a Physiological Basis for the Theory of Music* (London: Longmans, 1885). Trans reprinted (New York: Dover, 1954).

Keislar, D. F. 1991. "The Relevance of Beating Partials for Musical Intonation." *Proceedings of the 1991 International Computer Music Conference*. San Francisco: International Computer Music Association, pp. 202–205.

Mathews, M. V., and J. R. Pierce. 1980. "Harmony and Inharmonic Partials." *Journal of the Acoustical Society of America* 68:1252–1257.

Moore, B. C. J., B. R. Glasberg, and R. W. Peters. 1985. "Relative Dominance of Individual Partials in Determining the Pitch of Complex Tones." *Journal of the Acoustical Society of America* 77(5): 1853–1860.

Moore, B. C. J., R. W. Peters, and B. R. Glasberg. 1985. "Thresholds for the Detection of Inharmonicity in Complex Tones." *Journal of the Acoustical Society of America* 77:1861–1867.

Plomp, R. 1970. "Timbre as a Multidimensional Attribute of Complex Tones." In R. Plomp and G. F. Smoorenburg, eds. *Frequency Analysis and Periodicity Detection in Hearing*. Lieden: A. W. Sijthoff.

Plomp, R., and W. J. M. Levelt. 1965. "Tonal Consonance and Critical Bandwidth." *Journal of the Acoustical Society of America* 38:548–560.

Risset, J. C. 1987. *Sud, Dialogues, Inharmonique, Mutations*. INA/GRM compact disc INA C 1003.

Risset, J. C., and D. L. Wessel. 1982. "Exploration of Timbre by Analysis and Synthesis." In D. Deutsch, ed. *The Psychology of Music*. New York: Academic Press.

Sethares, W. A. 1993. "Local Consonance and the Relationship between Timbre and Scale." *Journal of the Acoustical Society of America* 94(3):1218–1228.

Slaymaker, F. H. 1968. "Chords from Tones having Stretched Partials." *Journal of the Acoustical Society of America* 47:1469–1571.

Strang, G., and T. Nguyen. 1996. *Wavelets and Filter Banks*. Wellesley, Massachusetts: Wellesley College Press.

Strong, W. J., and M. Clark. 1967. "Perturbations of Synthetic Orchestral Wind Instrument Tones." *Journal of the Acoustical Society of America* 41:277–285.

Tenney, J. 1988. *A History of "Consonance" and "Dissonance."* New York: Excelsior Music Publishing.

Vos, J. 1988. "Subjective Acceptability of Various Regular Twelve-Tone Tuning Systems in Two-Part Musical Fragments." *Journal of the Acoustical Society of America* 83(6):2383–2392.