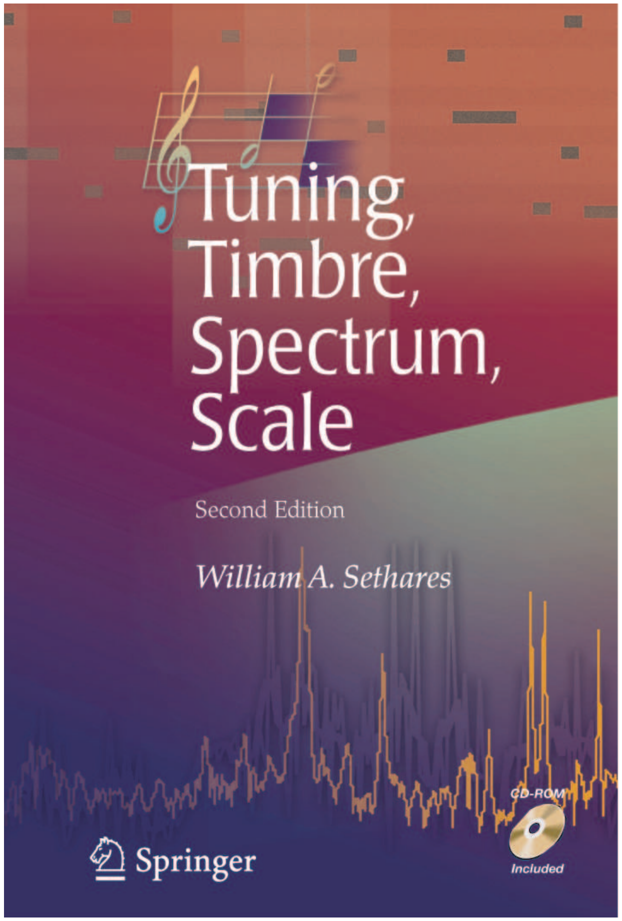

Tuning and Timbre: A Perceptual Synthesis

Bill Sethares

IDEA: Exploit psychoacoustic studies on the perception of consonance and dissonance. The talk begins by showing how to build a device that can measure the "sensory" consonance and/or dissonance of a sound in its musical context. Such a "dissonance meter" has implications in music theory, in synthesizer design, in the construction of musical scales and tunings, and in the design of musical instruments.

...the legacy of Helmholtz continues...

Some Observations. . .

Why do we tune our instruments the way we do?

Some tunings are easier to play in than others.

Some timbres work well in certain scales, but not in others.

What makes a sound easy in 19-tet but hard in 10-tet?

"The timbre of an instrument strongly affects what tuning and scale sound best on that instrument." – W. Carlos

What are Tuning and Timbre?

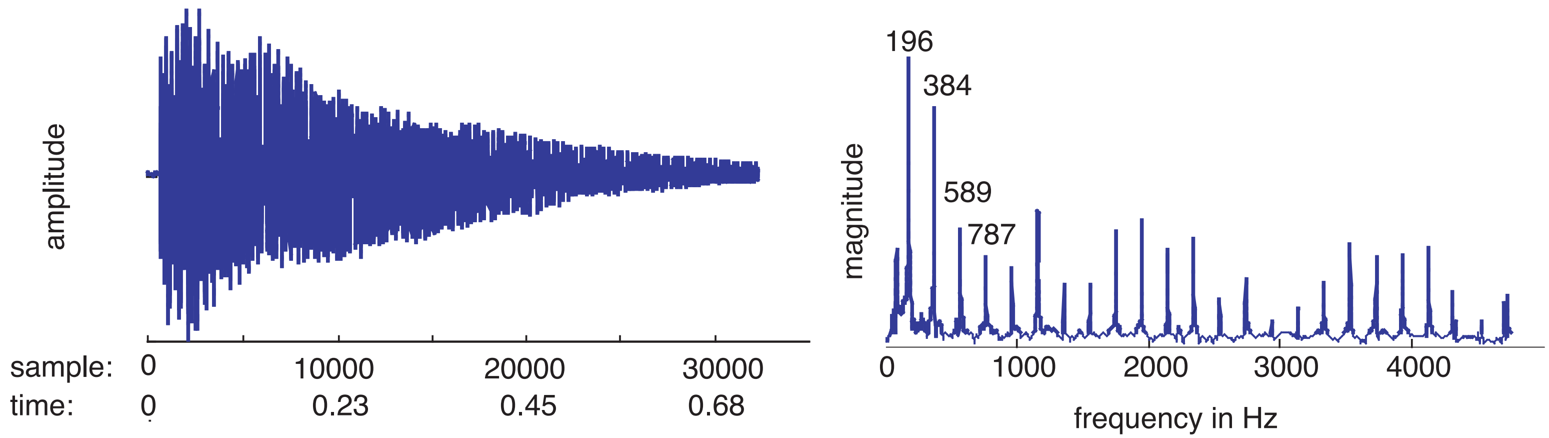

Tuning = pitch of the fundamental (in this case 196 Hz)

Timbre involves (a) pattern of overtones (Helmholtz)

(b) temporal features

Some intervals “harmonious” and others “discordant.” Why?

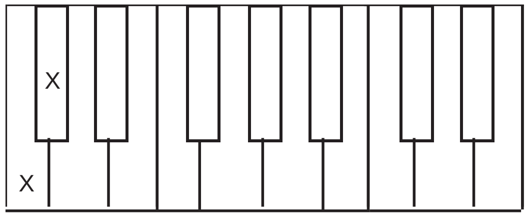

1.06:1

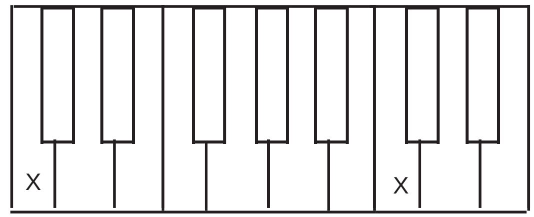

2:1

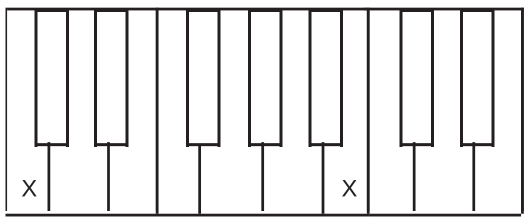

1.89:1

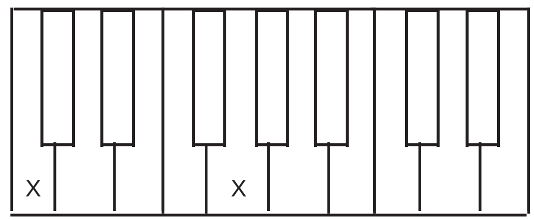

3:2

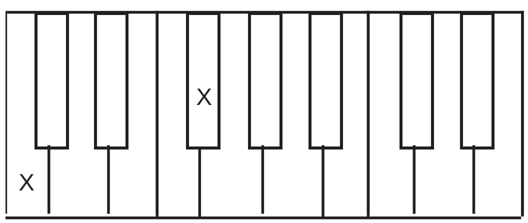

1.414:1

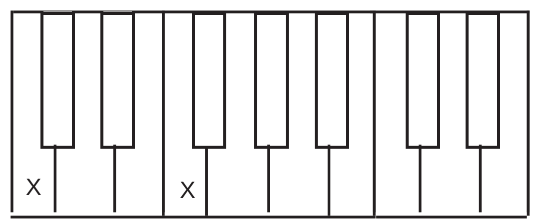

4:3

Theory #1: (Pythagoras) Humans naturally like the sound of intervals defined by small integer ratios.

small ratios imply short period of repetition

short = simple = sweet

Theory #2: (Helmholtz) Partials of a sound that are close in frequency cause beats that are perceived as “roughness” or dissonance.

absence of beats is called “consonance.”

A Short History of “Consonance” and “Dissonance” (after James Tenney)

- CDC#1: melodic consonance (e.g., of successive tones)
- CDC#2: polyphonic consonance (e.g., intervals between notes, “sounds good”)
- CDC#3: contrapuntal consonance (defined by role in counterpoint)
- CDC#4: functional consonance (relationship with “tonic” or “root”)
- CDC#5: psychoacoustic consonance (intrinsic to a sound)

What are Beats? (beats1-2-3.avi)

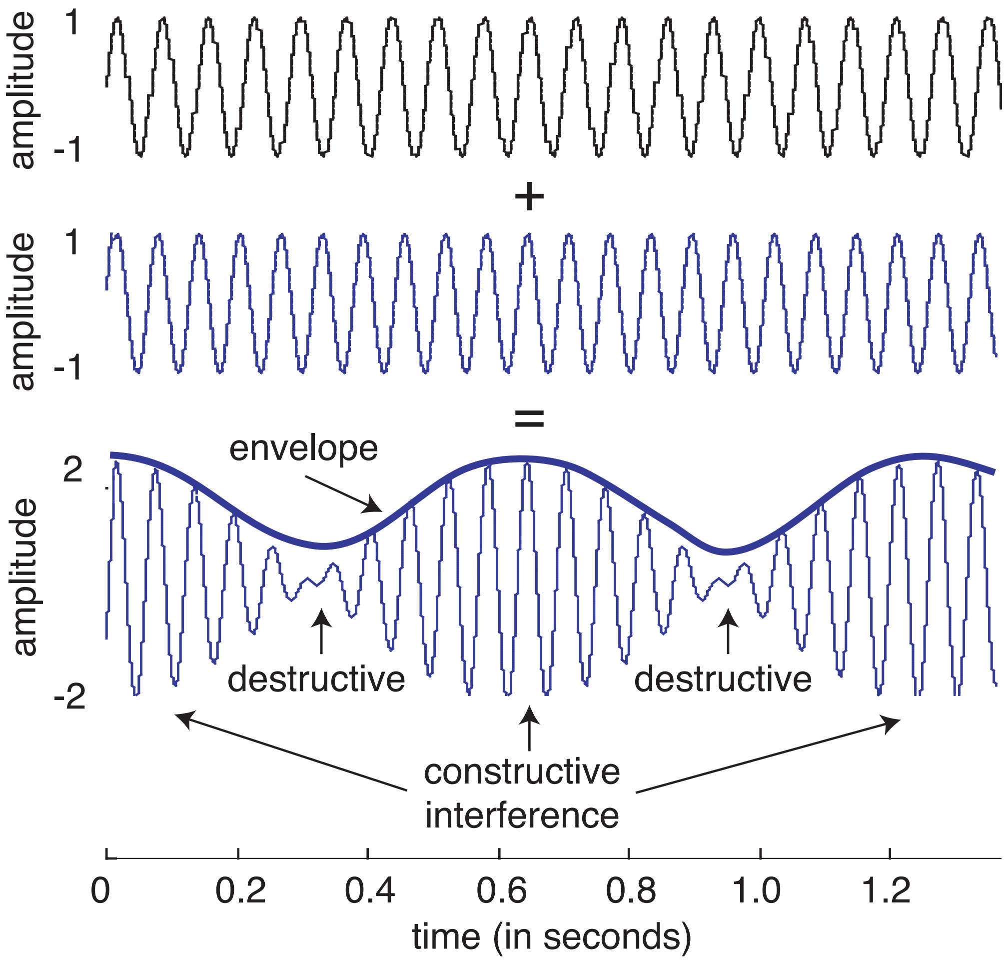

What should we expect to hear?

What happens when frequency of beats enters audio range?

100 Hz + 105 Hz = 5 Hz beats

100 Hz + 150 Hz = 50 Hz beats

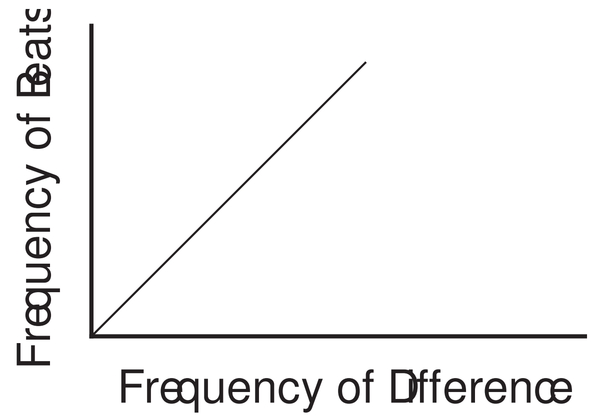

Experiment: (Plomp & Levelt) Fix w_1 and let w_2 scan through all frequencies. Ask listeners what they perceive.

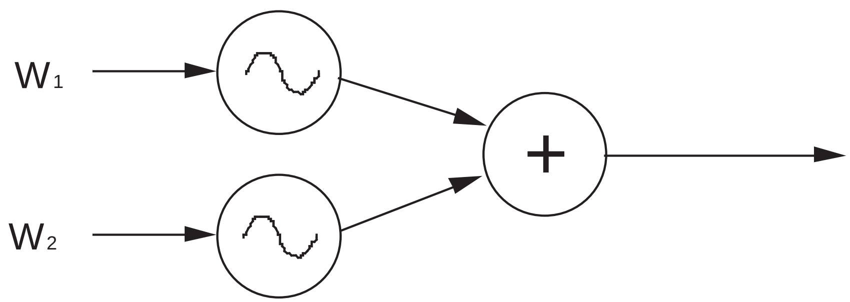

Two Sine Waves: One Fixed, One Sweeping (sinediss.avi)

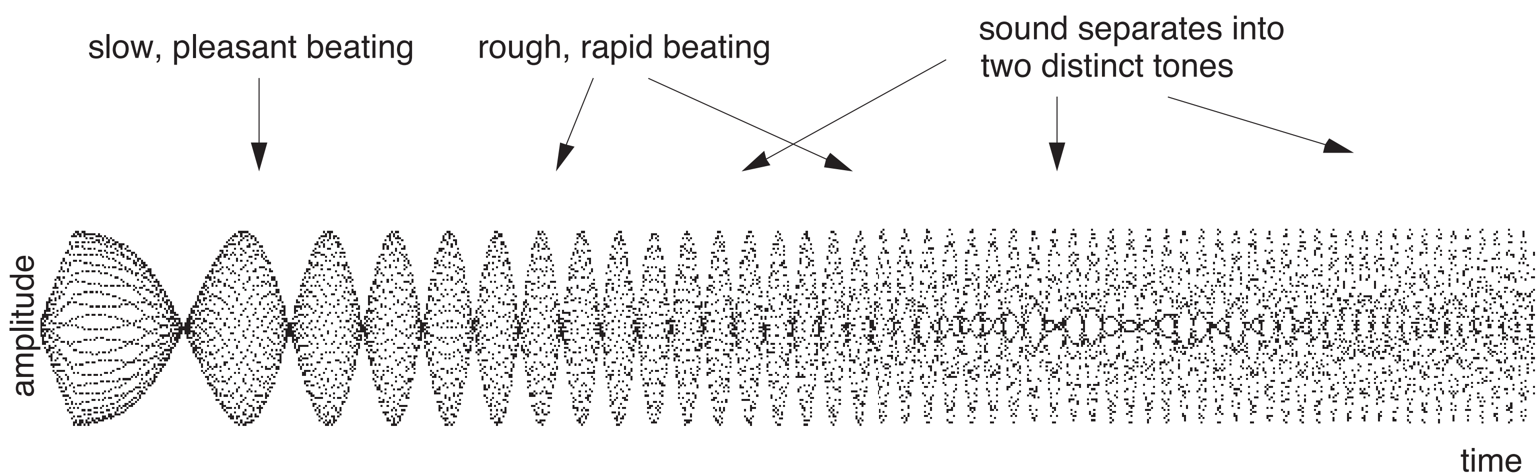

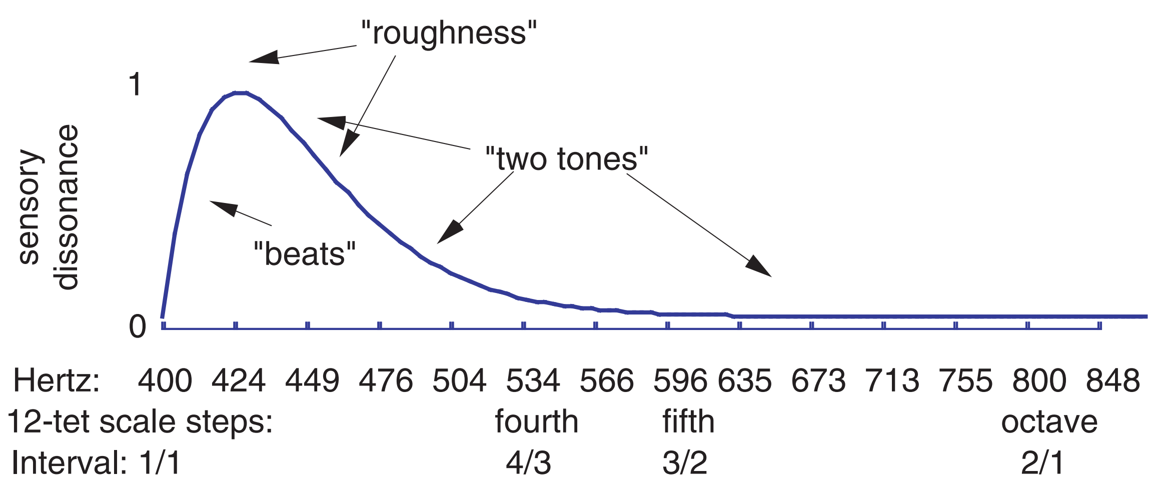

Parameterizing the Sensory Dissonance Curve

as the difference between a sum of exponentials

$$d(x) = e^{-ax} - e^{-bx}$$

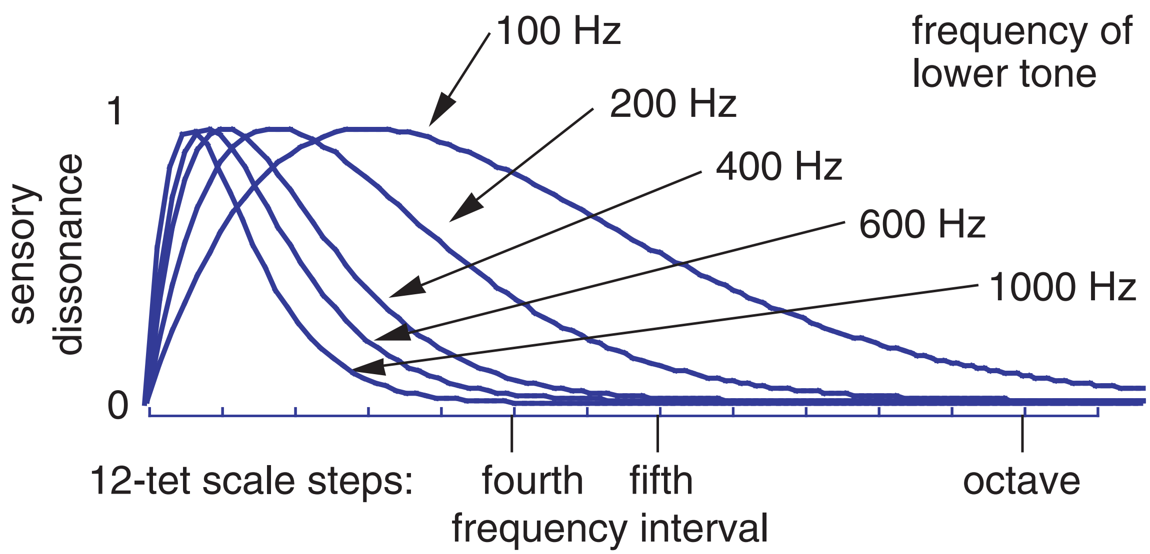

Some Implications...

Spectrum of a sound determines which intervals are most consonant

By using different kinds of spectra, can make almost any set of intervals sound consonant.

How to make a large variety of (organic, natural sounding) inharmonic timbres that work together? Using spectral mappings, can often maintain much of the original integrity of the sound. Examples:

consonant tritone — consonant pseudo-octaves

n-tone equal-tempered sounds

A Tritone Sound

Can use this parameterization to make predictions. Consider a sound composed of two partials:

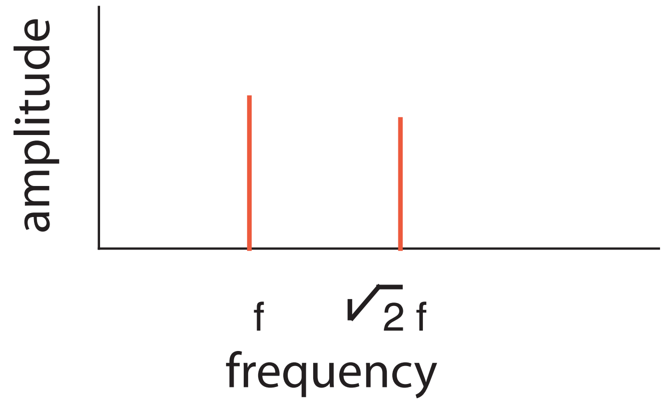

Summing all the dissonances between all pairs of partials gives the "dissonance curve" for this tone:

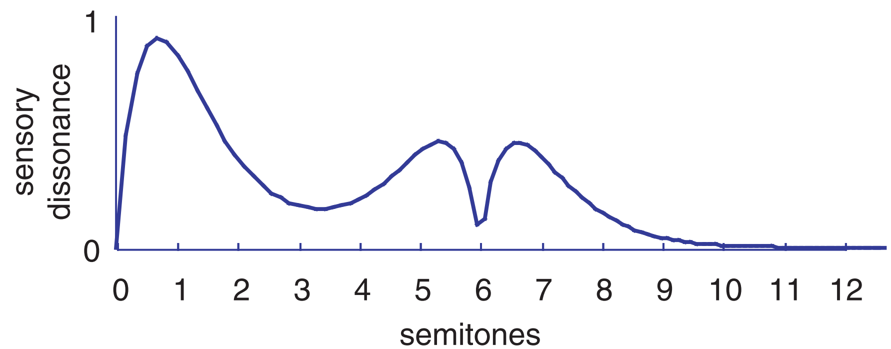

Tritone Sound II

Using the tritone sound, the dissonance curve predicts that the tritone interval $\sqrt{2}$ will be more consonant than the musical fifth or fourth.

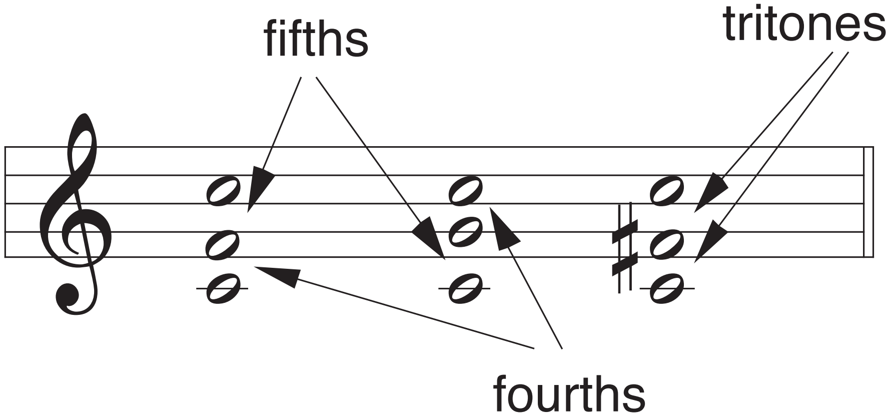

(trichime.avi)

Harmonic Dissonance Curve

Many musical sounds have harmonic partials, i.e., are (approximately) periodic with partials at

$$f, 2f, 3f, 4f, 5f, 6f, 7f, \ldots$$

Dissonance curves for harmonic tones have many minima that occur at simple integer ratios, which are close to the tuning of the keyboard.

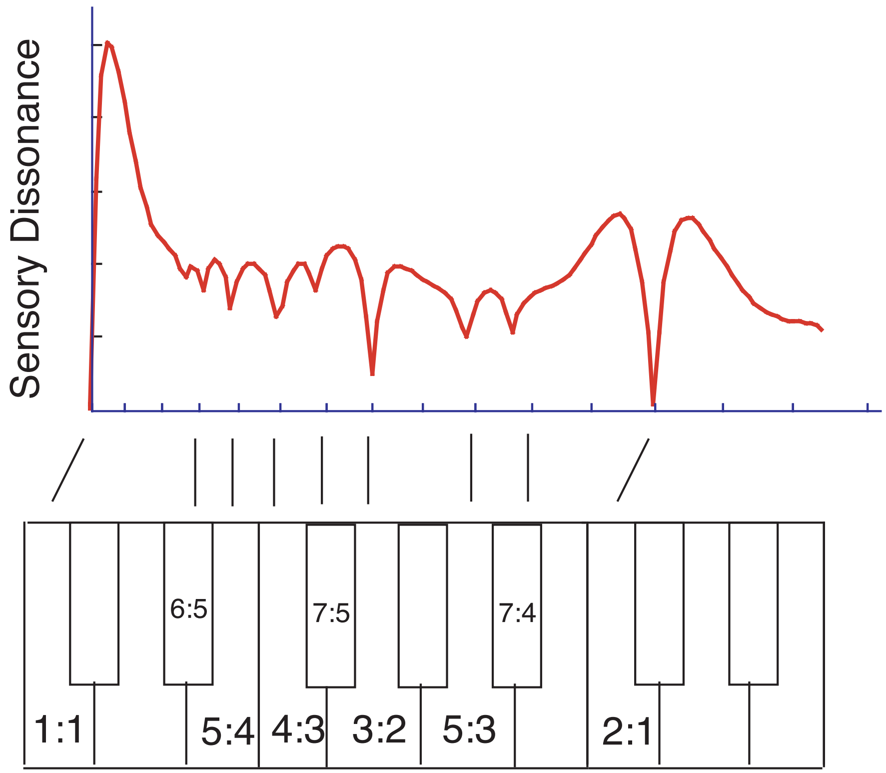

Helmholtz's Dissonance Curve

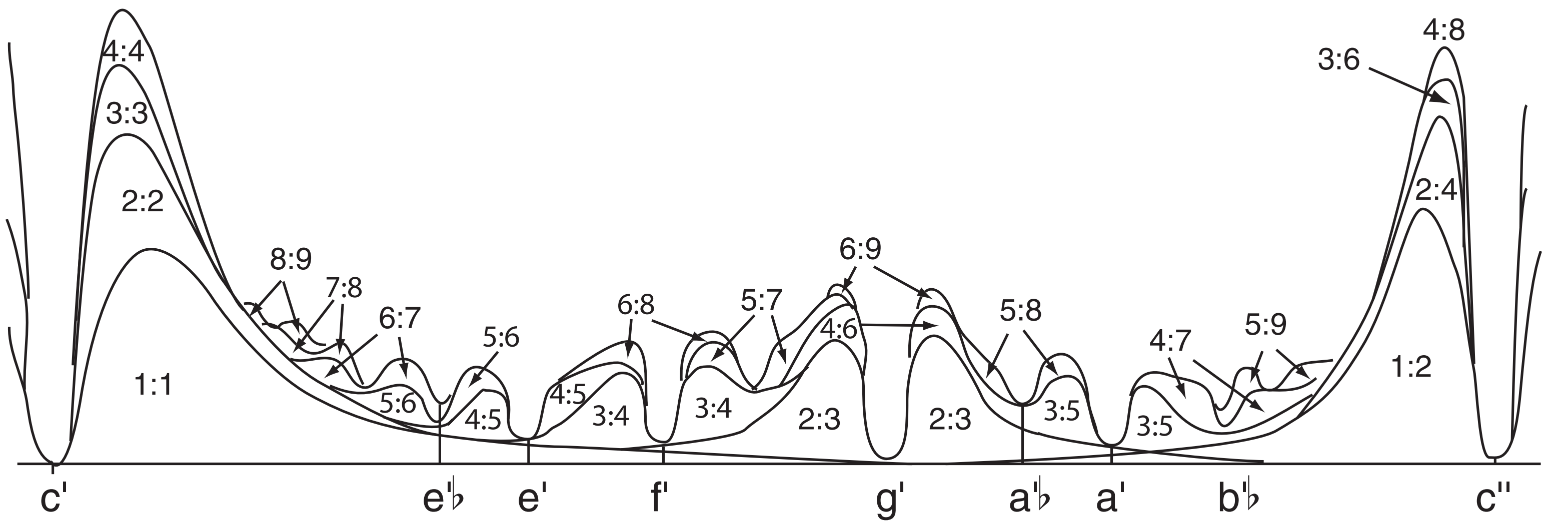

Two pitches are sounded simultaneously. The regions of roughness due to pairs of interacting partials are plotted over one another, leaving only a few narrow valleys of relative consonance. The figure is redrawn from *On the Sensation of Tone*.

Just Intonation

is a family of musical scales that contain many of the simple integer ratios (such as 3:2, 4:3, 5:3, etc.). A body of music exists in JI by composers such as Harry Partch (using his 43-tone per octave scale), Lou Harrison, D. Doty, Larry Polansky, and others. (* ⇒ just thirds, and <> ⇒ just fifths.) (Paradigms Lost)

A Just Intonation Scale in *C* and extension to a 12-note scale

	ratio	cents	
C	1/1	0	*<>
	16/15	112	*<>
D	9/8	204	*
	6/5	316	*
E	5/4	386	<>
F	4/3	498	*<>
	45/32	590	
G	3/2	702	*<>
	8/5	814	*<>
A	5/3	884	<>
	16/9	996	<>
B	15/8	1088	<>
C	2/1	1200	*<>

Relating Spectra and Scales

- Dissonance curves are drawn using the spectrum of a sound by summing the dissonances between all pairs of partials over a range of frequencies.

- Given a sound, the related scale is defined by the minima of the dissonance curve

- “Just Intonation” is the scale related to harmonic partials.

- What about other, inharmonic tones? Can dissonance curves be used to make predictions about what will sound “good”?

Stretched and Compressed Tones

have partials at $f_j = f A^{\log_2(j)}$

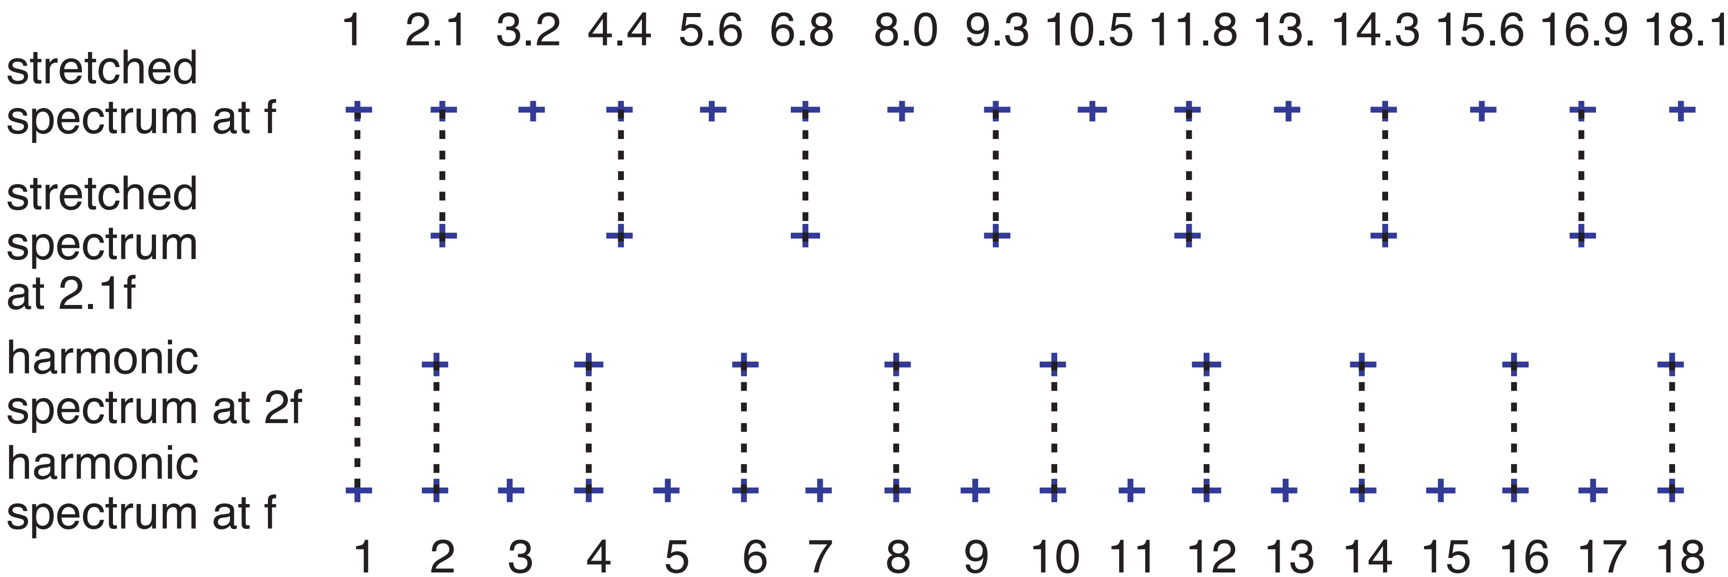

$A = 2$ gives harmonic tones, $A > 2$ are called "stretched," $A < 2$ are called "compressed"

Stretched and Compressed Tones II

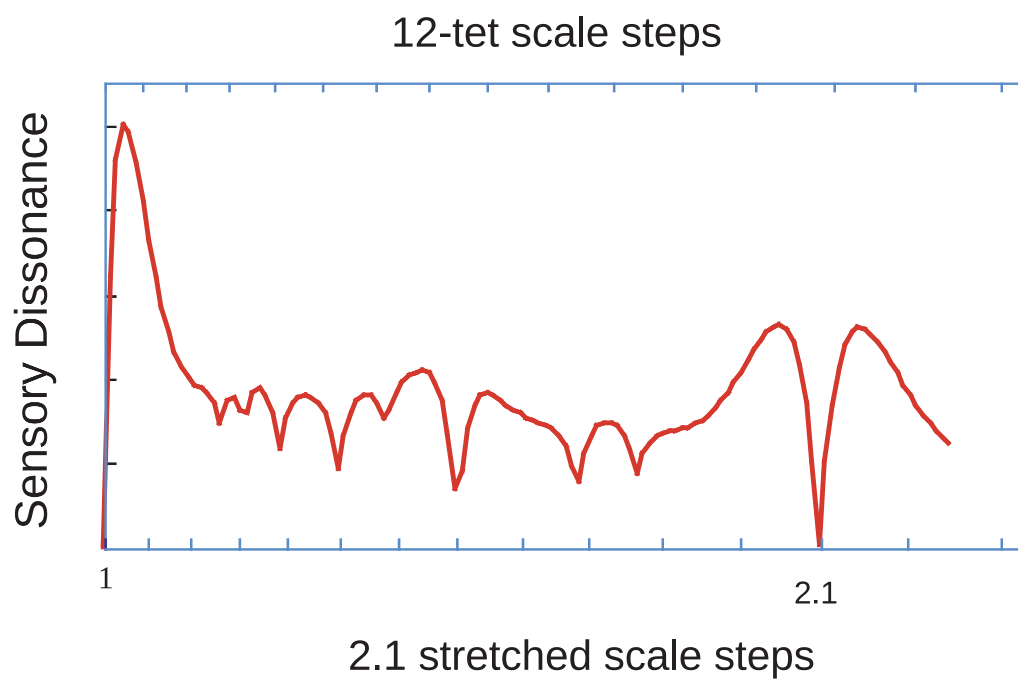

Predictions:

(simpletun1-2-3-4)

- harmonic tones in 12 tone = OK
- harmonic sounds in stretched scale = not OK
- stretched sounds in stretched scale = OK
- stretched sounds in 12 tone = not OK

Musical Implications

Olson: It is an established fact that the most pleasing combination of two tones is one in which the frequency ratio is expressible by two integers, neither of which is large.

Piston: Two sounds are said to be an octave apart when their frequency ratio is 2:1... The octave is the most consonant of intervals.

Such statements are found throughout the literature. Examples such as the stretched octaves show that either

(A) octave means a frequency ratio of 2:1, or

(B) octave means the most consonant interval (other than the unison).

(A) and (B) are the same only for sounds with harmonic spectra

Spectrum of the Tingshaw Bell

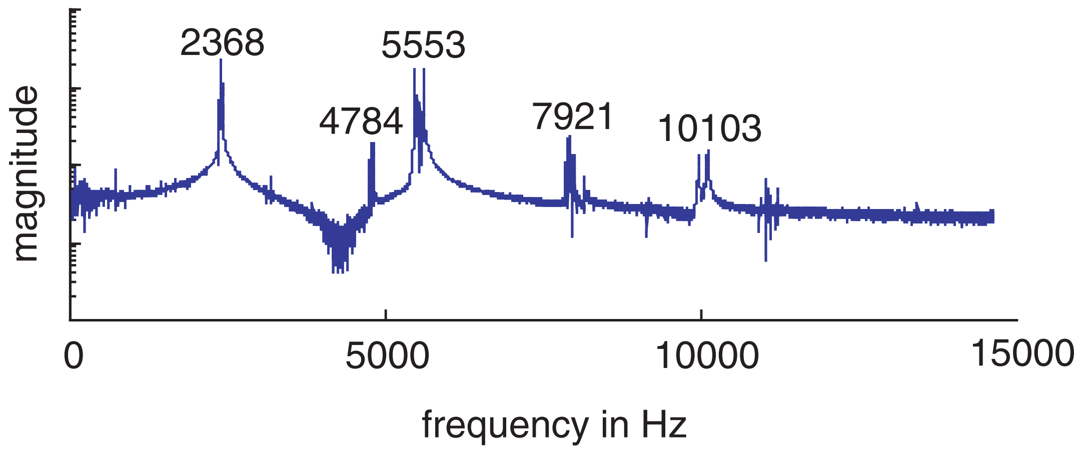

is simplified to a line spectrum

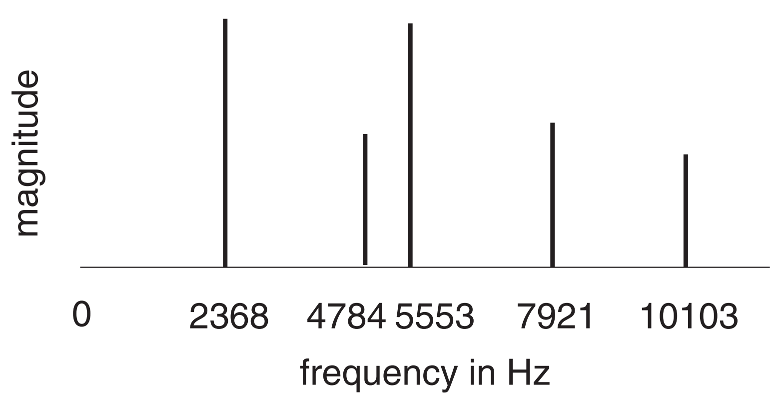

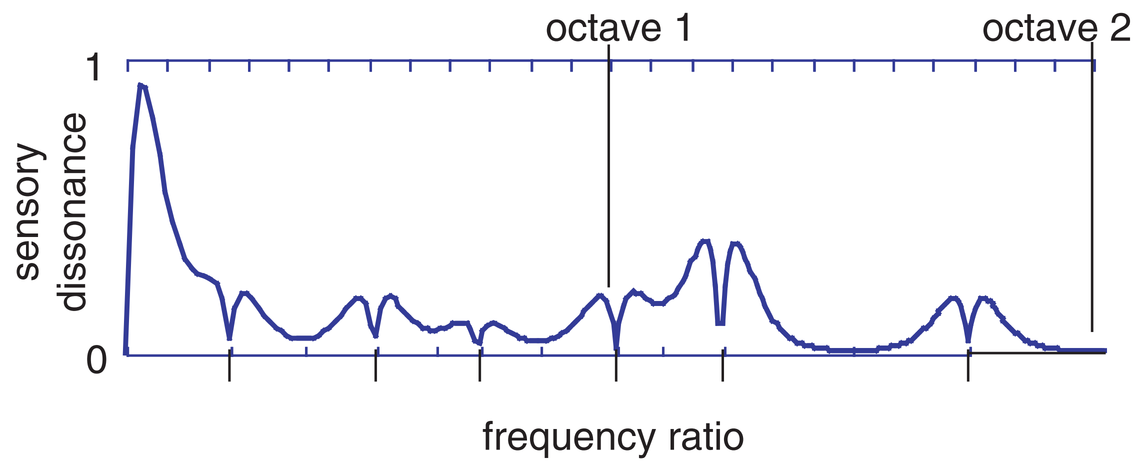

Tingshaw Scale

ratio	cents
1.0	0
1.16	257
1.29	441
1.43	619
1.56	770
1.66	877
1.81	1027
2.02	1200

Dissonance curve for the tingshaw bell has minima shown by tick marks. The mimimum at 2.02 serves as a pseudo-octave, because minima in the second pseudo-octave align with those in the first. (tingshaw)

Spectral Mappings

can be used to create new “instruments” that are consonant with a desired timbre yet retain much of the character of familiar instruments

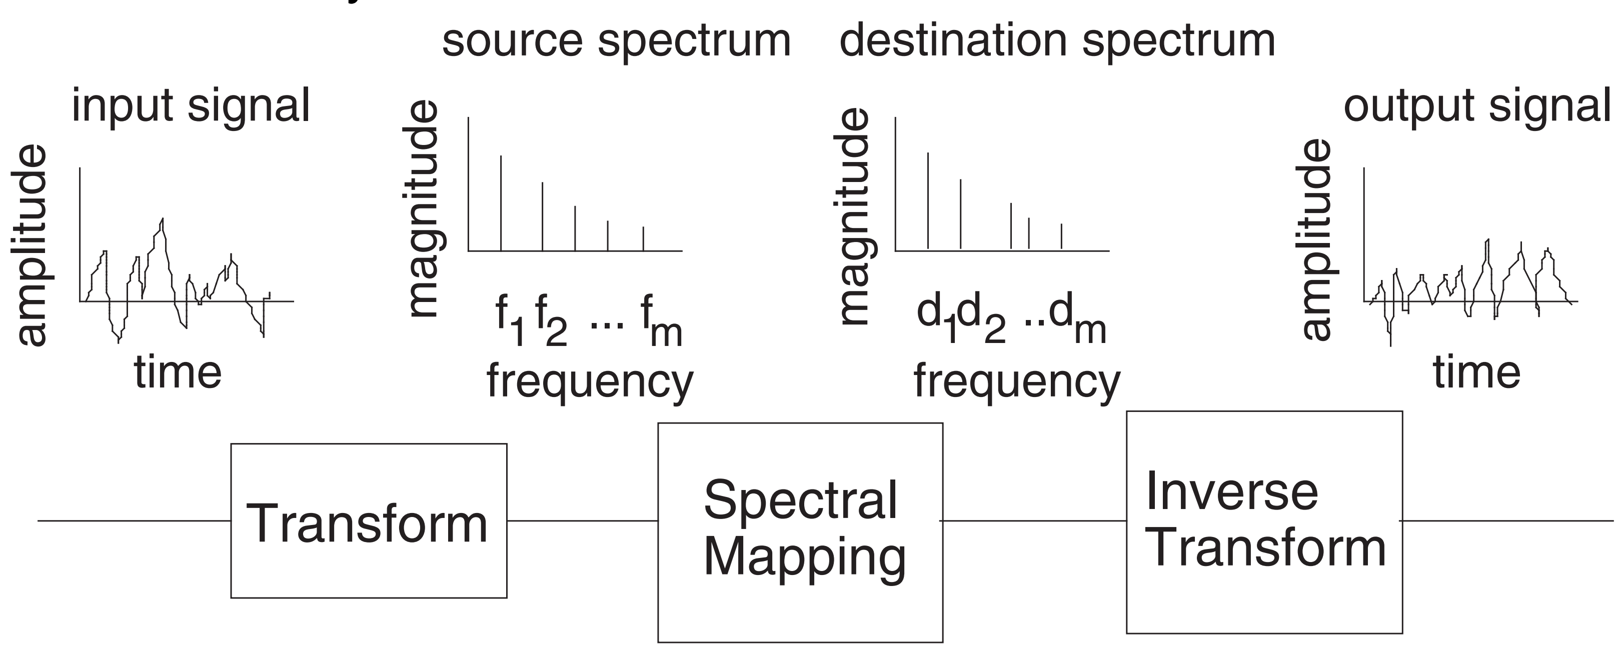

Spectral Mappings change frequencies, preserving magnitude and phase relationships, which helps to preserve the “character” of the original.

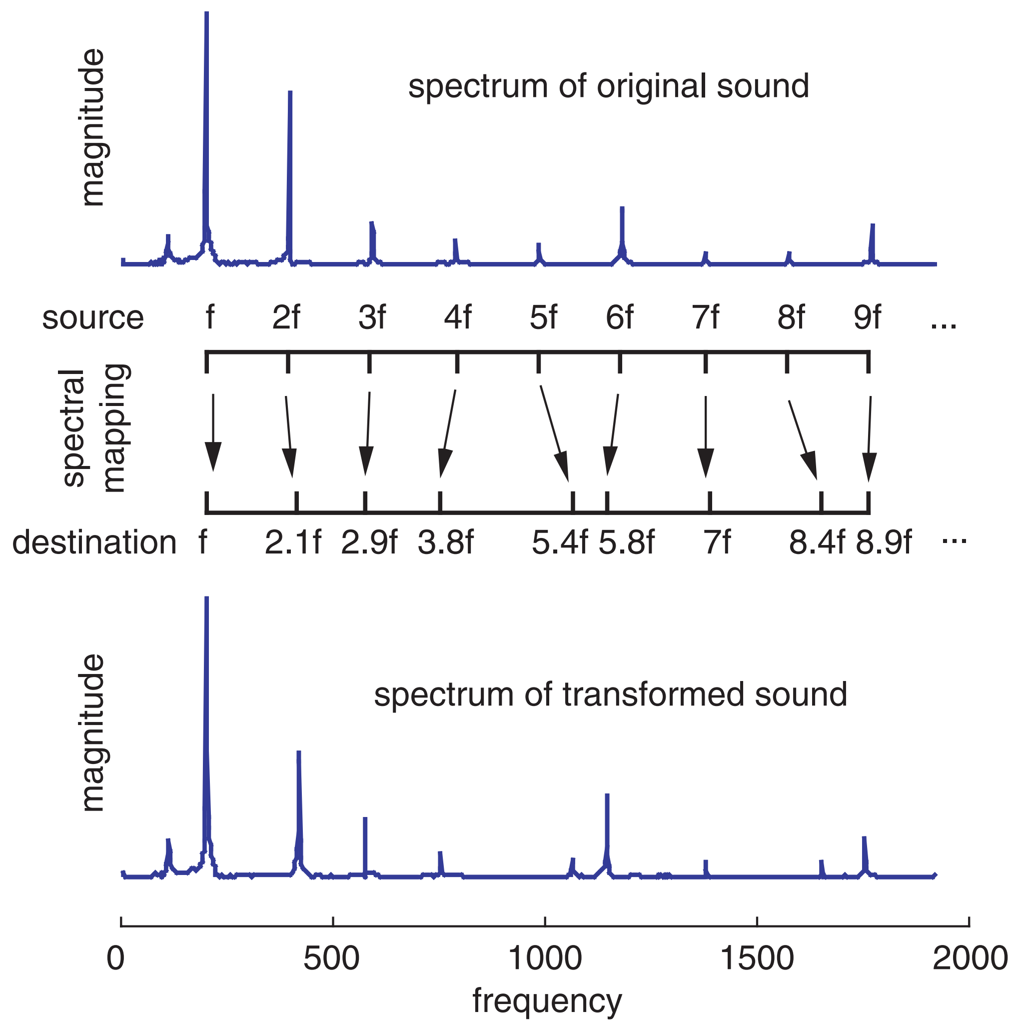

Sounds for 10-tone equal temperament

A timbre designed to be played in 10-tet has partials at $f, f\alpha^{10}, f\alpha^{16}, f\alpha^{20}, f\alpha^{26}, f\alpha^{29}, f\alpha^{30}, f\alpha^{33}, f\alpha^{36}, f\alpha^{39}, f\alpha^{40}$ where $\alpha = \sqrt[10]{2}$. The dissonance curve for this timbre is:

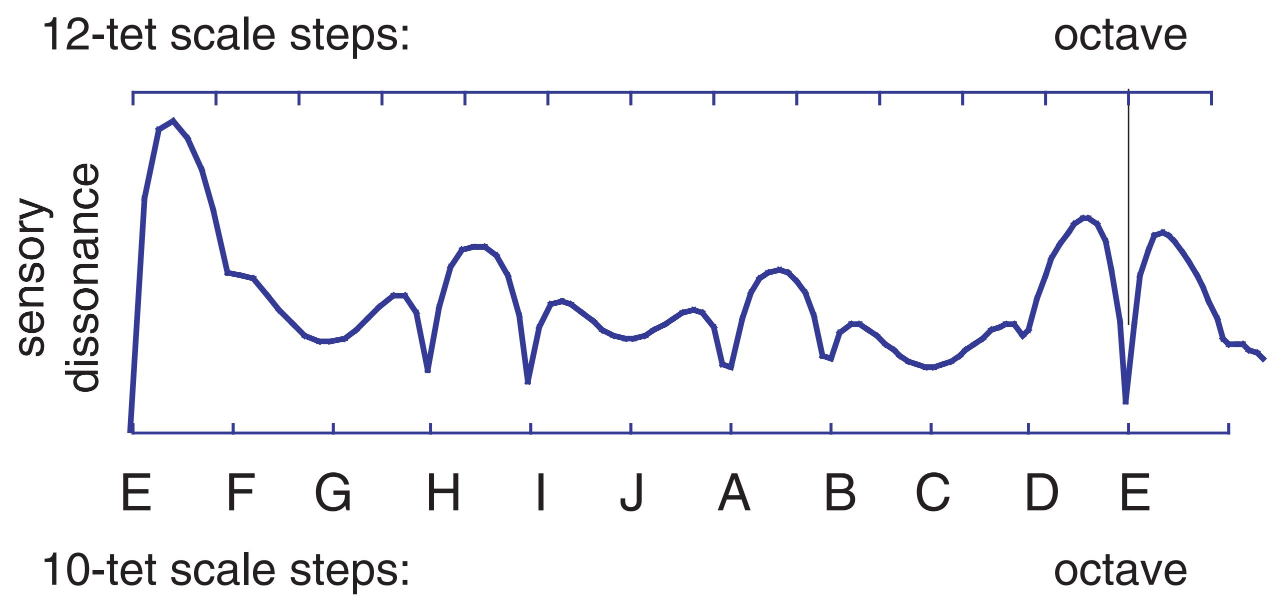

Minima of the curve coincide with steps of the 10-tet scale and not with steps of 12-tet. (Ten Fingers, MysteryX)

10-tet Circle of Thirds

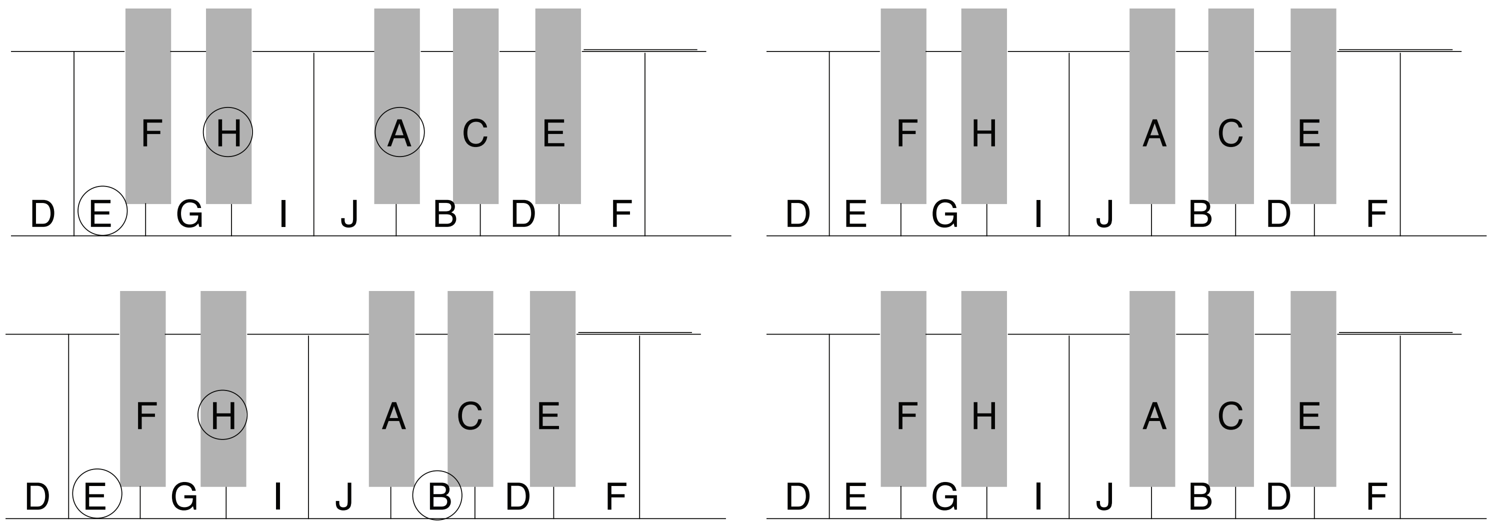

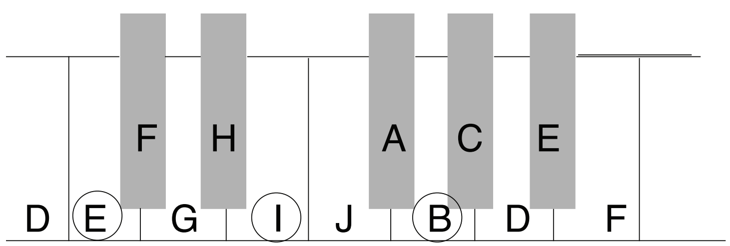

E->B->I->F->C->J->G
->D->A->H->E

There are many such music theories. You can invent new chords, new scales, new progressions – Bach didn't use them all up. (Circle of Thirds)

10-tet Chords

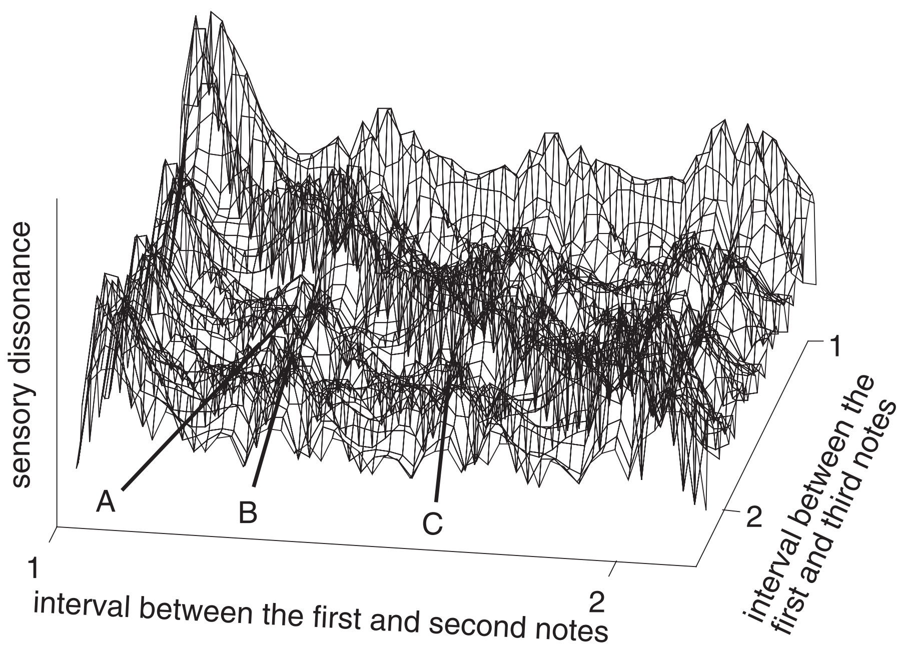

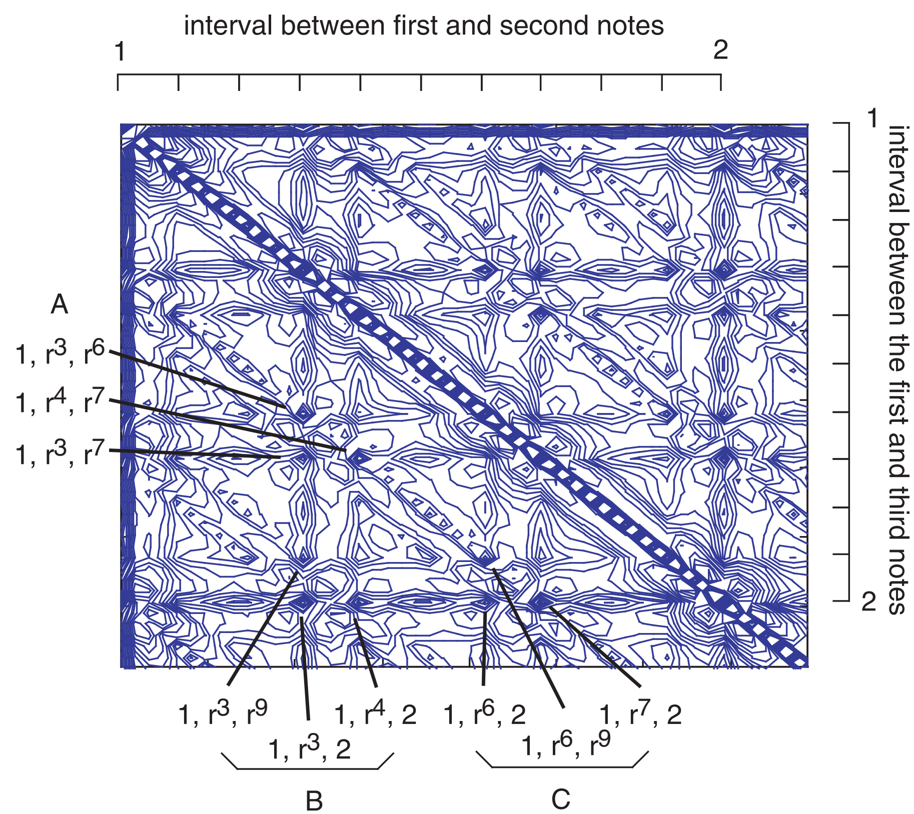

Sounds for 11-tet

- trumpet → 11-tet trumpet
- bass → 11-tet bass
- guitar → 11-tet guitar
- pan flute → 11-tet pan flute
- moog synth → 11-tet moog synth
- phase synth → 11-tet phase synth

Musical interludes played:

- In 11-tet scale with 11-tet timbres
- In 11-tet scale with regular (unmapped, harmonic) timbres

(tim11tet.avi, tim11vs12.avi, Turquoise Dabo Girl)

Create sounds consonant with the spectrum of a tom-tom: Harmonic spectra at $g,\ 2g,\ 3g,\ 4g,\ 5g, \ldots$ are mapped into the tom-tom spectrum and played using the related scale.

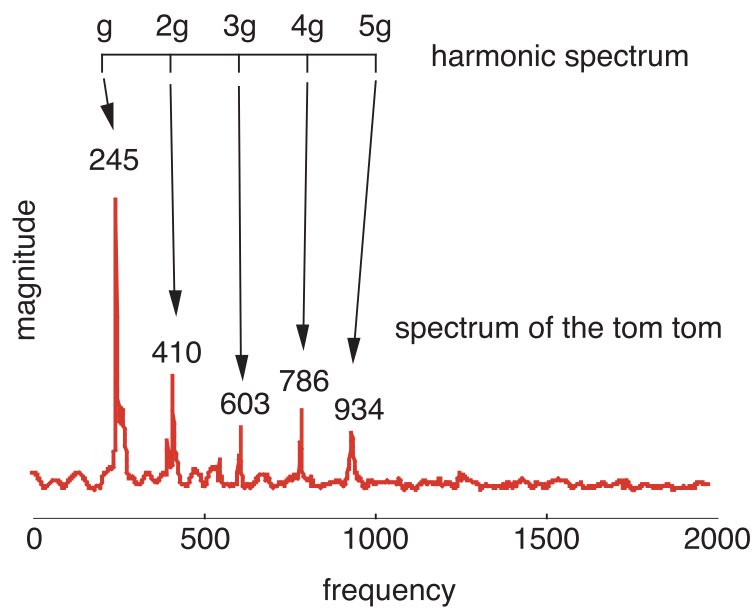

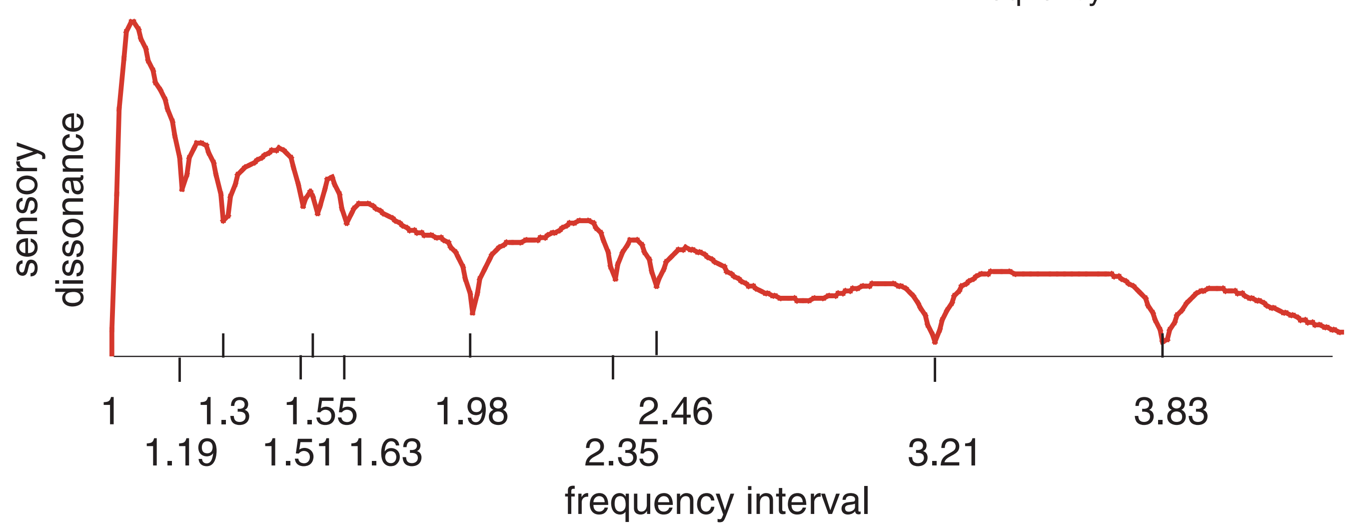

Sounds consonant with the spectrum of a tom-tom

Several instruments and their transformation into the spectrum of a tom tom. (tomspec.avi)

- guitar → tom tom guitar
- bass → tom tom bass
- trumpet → tom tom trumpet
- flute → tom tom flute

A musical passage illustrates the transformed instruments played in the related scale. (Glass Lake)

Mild transformations (like the 2.1 stretched and 10-tet) retain much of the character of the instrument from which they were derived.

Severe transformations (like the tom tom example) lose their tonal integrity. This does not mean that such sounds are musically useless!

The two sides...

Given a timbre, what is the related scale?

Just draw dissonance curve and find minima.

Given a desired scale, what are related timbres?

Solvable via iterative optimization methods

More Formally

Suppose the timbre T has n partials at $f_1, f_2, \ldots, f_n$. Let αT be the timbre with partials at $\alpha f_1, \alpha f_2, \ldots, \alpha f_n$. The dissonance curve generated by T is defined to be a plot of the sensory dissonance between T and αT over all intervals α of interest.

(1) Number of minima: dissonance curves have up to $2n(n-1)$ minima.

(2) Global minimum: the unison ($\alpha = 1$) is the global minimum.

(3) Asymptotic value: as the intervals grow large (as $\alpha \to \infty$), the dissonance approaches a value that is no more than the intrinsic dissonance of the timbre itself.

(4) Principle of coinciding partials: up to half of the local minima occur at intervals α for which $\alpha = \frac{f_i}{f_j}$ where f_i and f_j are arbitrary partials of T.

Timbre Selection as an Optimization Problem

Often one wishes to specify a desired scale. How can related timbres be found? Suppose there are m scale tones.

Try # 1: Choose a set of n partials and n amplitudes to minimize sum of dissonances over all m intervals.

Problems: (a) vanishing amplitudes (b) runaway frequencies

Fix by adding constraints: (a) chose set of amplitudes a priori (b) insure all partials lie in some predetermined range.

Revised problem: With the amplitudes fixed, select a set of n frequencies lying in range to minimize:

w_1 (sum of dissonances over all intervals) + w_1 (number of local minima)

Audio Signal Processing

These examples demonstrate that dissonance curves really do capture something crucial about our perceptions of desirable and undesirable sounds.

- predictions of good/bad sounding pieces are consistent with calculations
- minima of dissonance curve are good scale/chord indicators

Thus:

We can use the notion of consonance as a basis for audio signal processing devices.

Build into our machines knowledge of how our perceptions work.

What would such machines look like?

Idea: A keyboard "knows" what sounds you have played. Suppose you choose

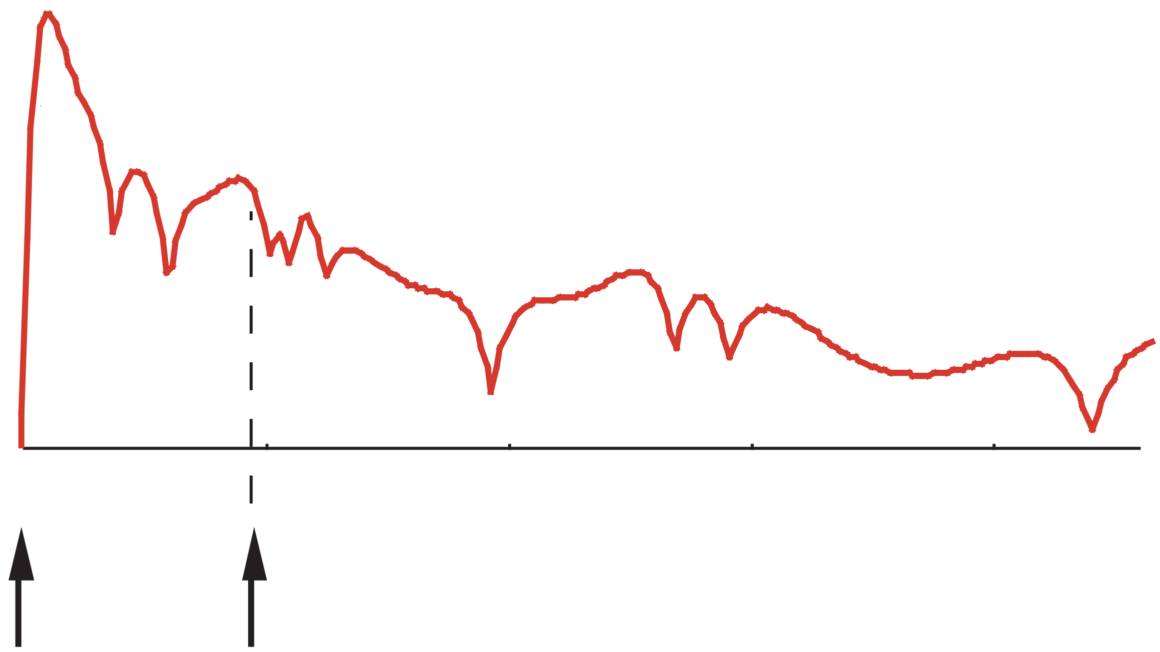

The keyboard could automatically adjust the tuning so as to minimize the dissonance by "sliding down" the dissonance curve.

Adaptive Tuning

A “Dynamic” or “Adaptive” tuning strategy can be implemented in a gradient style algorithm

$$\left\{ \begin{array}{c} new \\ frequency \\ values \end{array} \right\} = \left\{ \begin{array}{c} old \\ frequency \\ values \end{array} \right\} - \{stepsize\}\{gradient\}$$

that finds the nearest local minimum of dissonance curve to each commanded note.

Musical implications:

- a way to automatically play in JI when using harmonic timbres

- automatically play in related scale using nonharmonic timbres
- no knowledge of key or tonal center required (method operates on sound rather than on musical theory)
- many musicians (singers, violinists, horns) adjust their intonation in response to musical situation - provides a way for keyboardists to imitate this
- can act as a kind of intelligent portamento or “elastic tuning”

Listening to Adaptation

Three notes are played, each with timbre $f, 1.414f, 1.7f, 2f$

Initial ratios of fundamentals are: 1.0, 1.335, 1.587 (i.e., 12-tet notes C, F, G#). Final adapted ratios are 1.0, 1.414, 1.703

Example is played three times, with (a) extremely slow adaptation (b) slow adaptation, and (c) medium adaptation. (listenadapt)

- adaptation removes most prominant beats
- adaptation retunes all three notes
- remaining quicker beats are inherent to sound
- remaining slow beats ($\approx$ 1 per second) due to resolution of the equipment

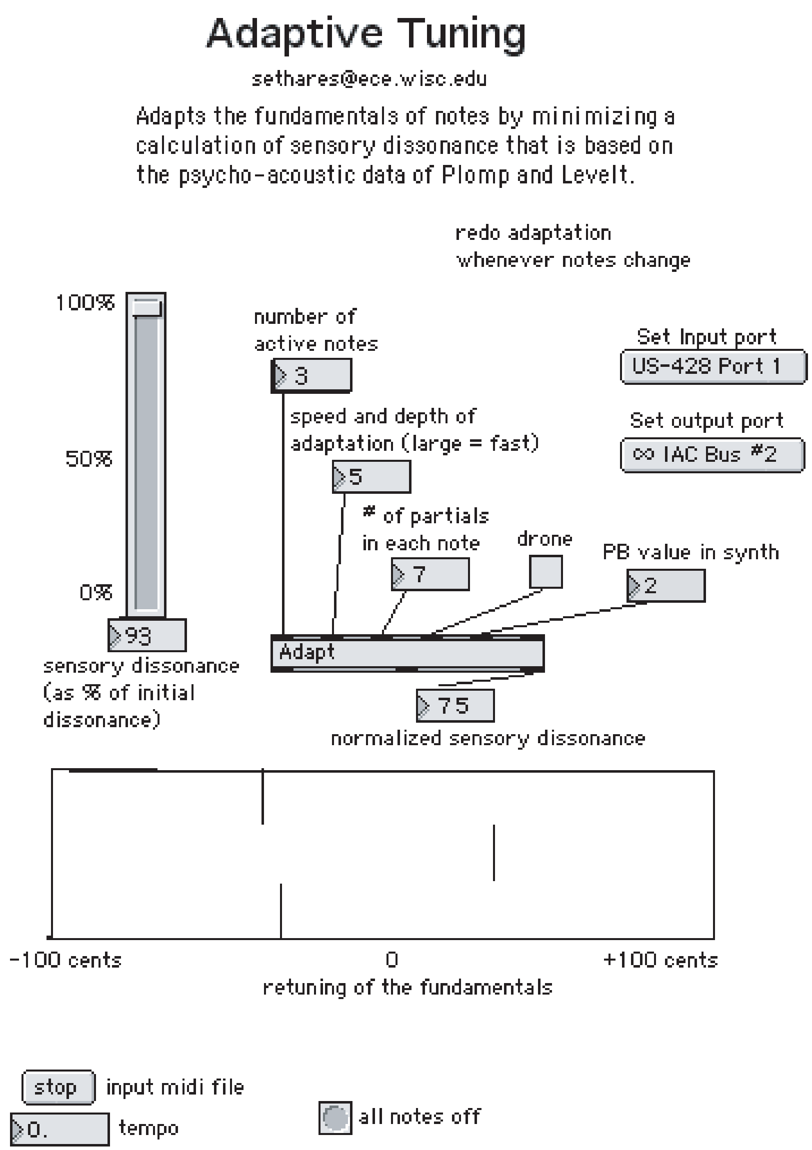

(local anomaly, aerophonius intent)

JI vs. 12-tet vs. Adaptive Tuning

An example of drift in Just Intonation: the fragment ends about 21 cents lower than it begins. 12-tet maintains the pitch by distorting the simple integer ratios. The adaptive tuning microtonally adjusts the pitches of the notes to maintain simple ratios and to avoid the wandering pitch. Frequency values are rounded to the nearest 0.5 Hz. (sytonJIdrift, synton12tet, syntonadapt)

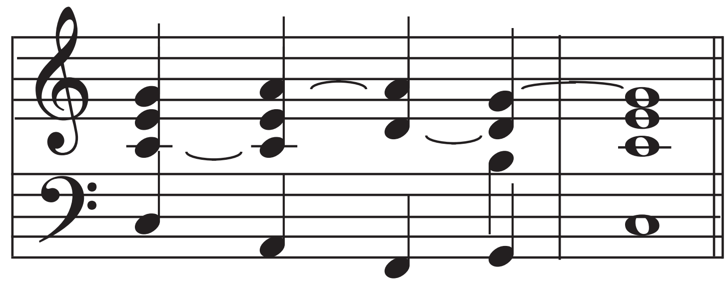

Frequencies when played in JI with held notes:	392.5	436---	436	387.5--	387.5
	327	327	290.5-	290.5	323
	261.5-	261.5	290.5	242	258.5
	131	109	87	96.5	129
Frequencies when played in 12-tet:	392	440	440	392	392
	329.5	329.5	293.5	293.5	329.5
	261.5	261.5	293.5	247	261.5
	131	110	87.5	98	131
Frequencies when played in adaptive tuning:	392.5	440	438.5	391	392.5
	327	330	292	294	327
	261.5	264	292	245	261.5
	131	110	87.5	98	131
Ratios when played in adaptive tuning and in JI:	6/5	4/3	3/2	4/3	6/5
	5/4	5/4	1/1	6/5	5/4
	2/1	6/5	5/3	5/4	2/1

Pitch/Intonation Correction

By making sounds more consonant, they become more “in-tune”

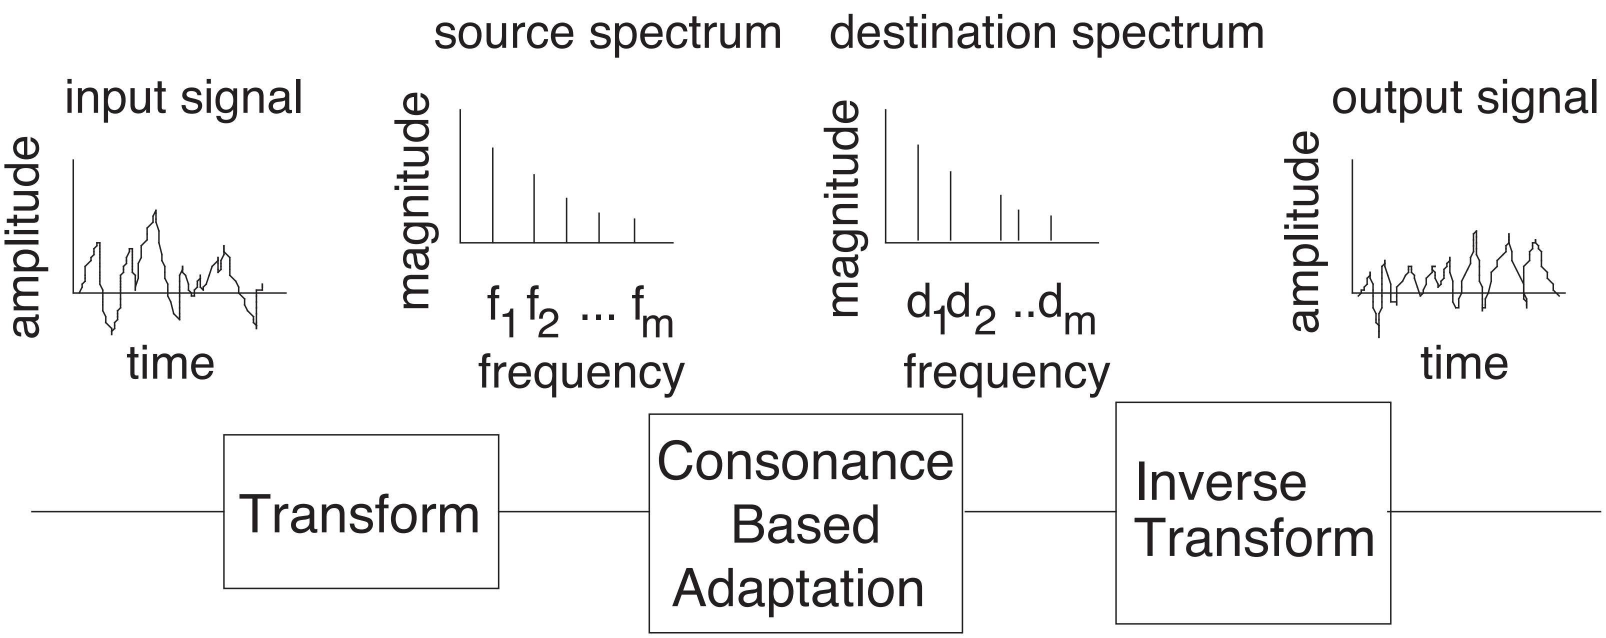

Consonance Based Sound Synthesis

Combines character of two (or more) waveforms in a musically intelligent (consonant) manner

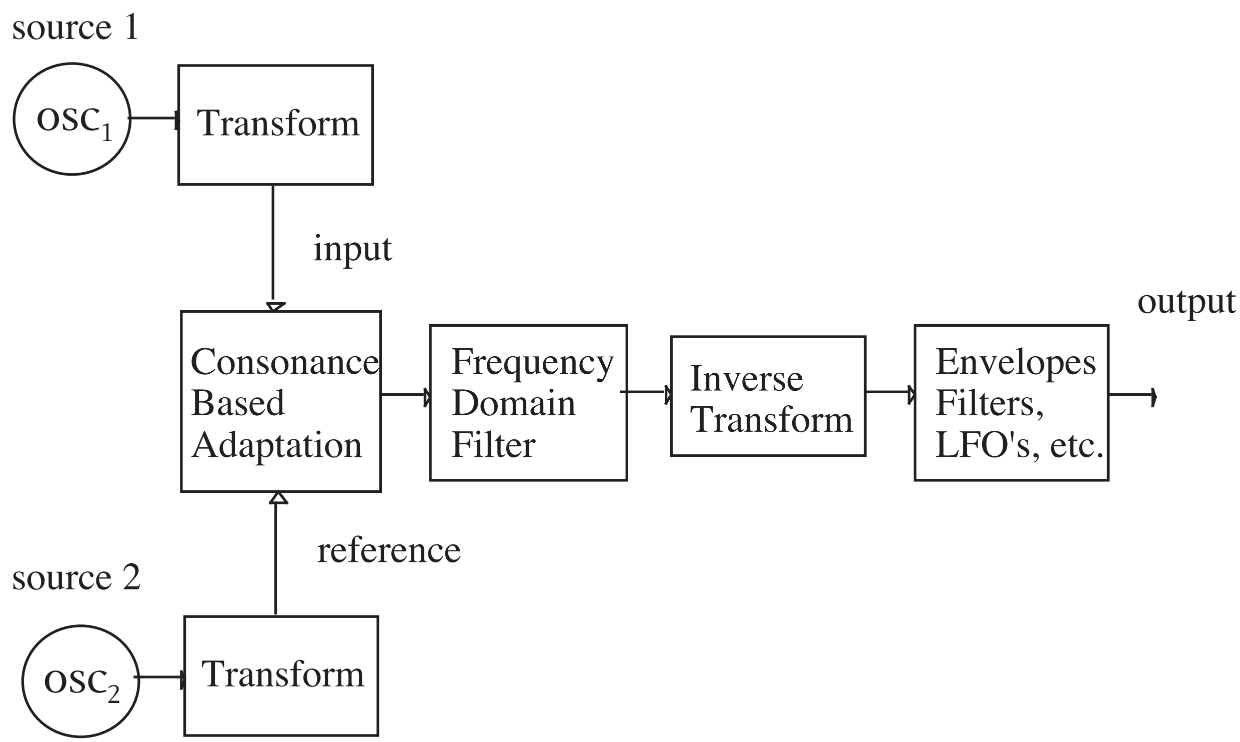

A method of sound synthesis that incorporates a model of the listener.

Special Effects Device

Algorithm adjusts partials of a sound to maximize consonance with reference.

- use of inharmonic reference
- voice with spectrum of xylophone
- snare drum made consonant with a flute
- use “backwards” to increase dissonance or to precisely control amount of consonance

(maxdiss)

Traditional Thai Music

is played on a collection of instruments (the gong circle) featuring the (xylophone-like) *renat*. (Sudsaboun)

Sorrell: Theoretically, the Thai scale has seven equidistant notes, which means that the intervals are “in the cracks” between our semitone and whole tone, and are equal...

Traditional Thai Music II

The spectrum of a typical key of the renat

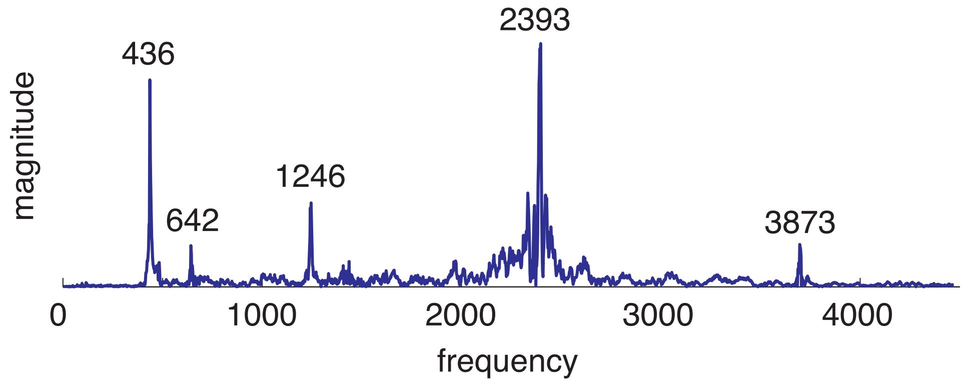

is very close to the spectrum of an ideal bar.

frequency Hz:	436	642	1246	2393	3873
ratio:	f	$1.47f$	$2.85f$	$5.48f$	$8.88f$
ideal bar:	f	$-$	$2.76f$	$5.4f$	$8.9f$

Traditional Thai Music III

But the dissonance curve for the ideal bar looks nothing like 7-tet. What's wrong? Observe that Thai music uses both inharmonic instruments (like the renat) and harmonic instruments (reeds, voice). Drawing the dissonance curve for both sounds simultaneously yields:

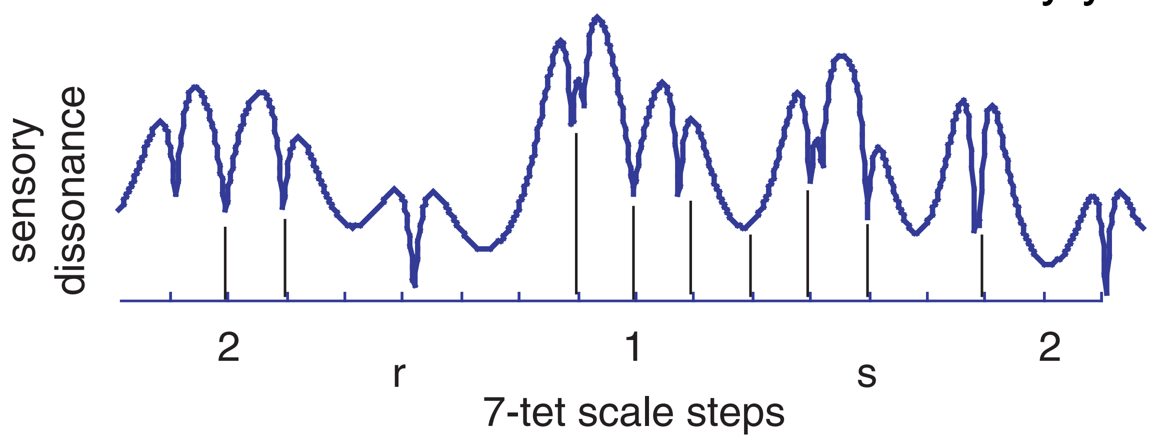

Traditional Indonesian Music

is played on a collection of metallophones including the *bonang*.
(Kebyar Duduk)

There are two kinds of scales: slendro is very close to 5-tet and pelog is a seven tone scale with unequal intervals (sometimes anotated S S L S S S L).

Traditional Indonesian Music II

Since the bonang has a unique bell-like shape, there is no ideal shape to which it can be compared. The spectrum of three different bonang kettles is shown, and a good generic bonang spectrum is

$$f,\ 1.52f,\ 3.46f,\ 3.92f.$$

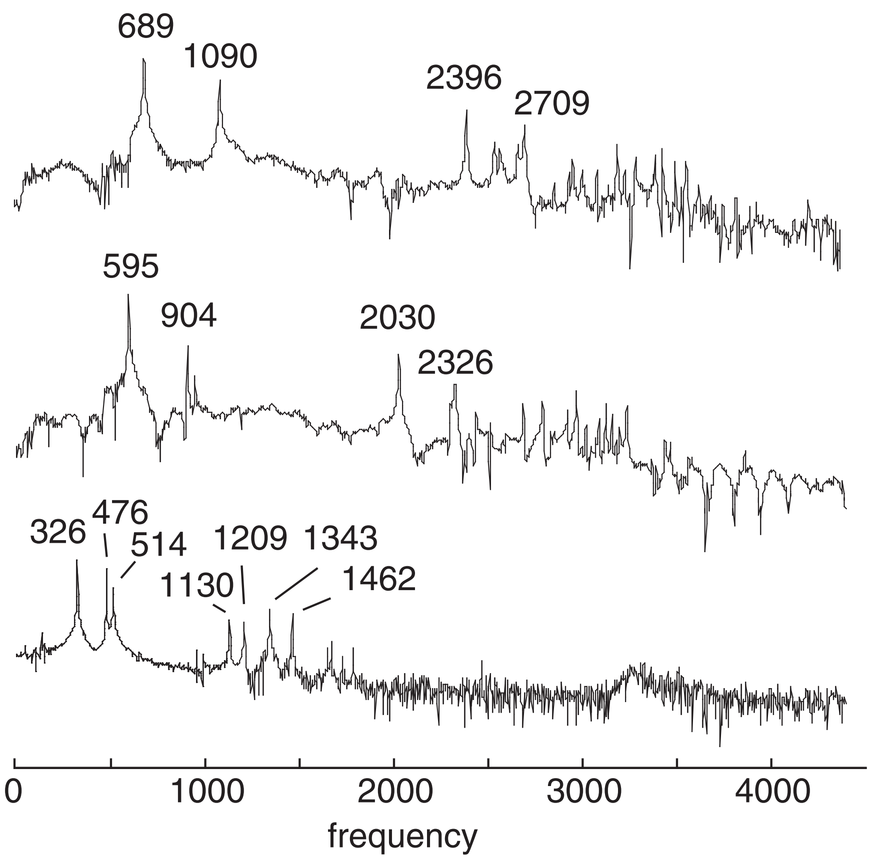

Traditional Indonesian Music III

As with the Thai instruments, the inharmonic bonang is often played together with harmonic instruments (flutes, voices). Drawing the dissonance curve for both sounds simultaneously yields

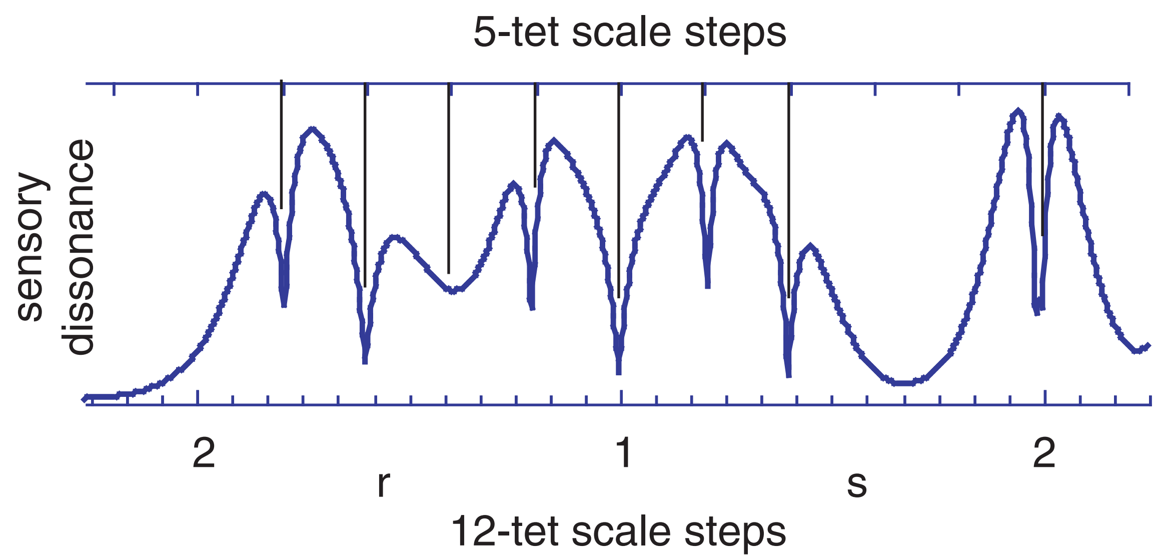

which has minima at or near all of the 5-tet scale steps. Similar analysis can be done for the pelog scale.

Current Work

- Last time we saw ways to adjust timbre/spectrum to fit a desired spectrum by minimizing sensory dissonance.
- Idea: take a melody (e.g., a clarinet) and dynamically re-align the timbre so that it "harmonizes" with a collection of sounds (e.g., multiphonics produced by the clarinet).
- a way to use arbitrary sounds in the role normally occupied by of a "chord pattern"

Example:

- the melody line (clarimel)
- the multiphonics (allmultiphonics)

Current Work II

How to align the multiphonics with the melody?

- one multiphonic all the way through (merge01)
- one multiphonic chosen randomly at each beat (clariphonicsrand2)
- divide into groups and choose from a small number (clariphonics-group2.mp3) (clariphonicsgroup2b.mp3)

Comments:

- clarinet-ness lost: it's not a clarinet or a multiphonic
- a bizzare kind of harmonization
- need to synchronize changes in "harmony" with sections of melody

Legend of Spectral Phollow

chooses at each instant a multiphonic to accompany the performed melody. The multiphonic is transposed, and this generates a kind of (inharmonic) harmonization.

What are the “rules” of multiphonic harmonization?

Use other sounds to derive other kinds of harmonies: gongs, bells, drums, anything with a rich/complex timbre will do.

Summary

Consonance and dissonance can be described in concrete terms and the predictions of the theory are readily verifiable.

Given a timbre, it is easy to find the related scale in which the sound can be played most consonantly.

Given a scale, it is possible to find related timbres that can be played consonantly in that scale.

Problem with *un*related scales and timbres is that there is little opportunity for contrast between consonant and dissonant passages - related scales and timbres allow composer/performer *control*.

Familiar “music theory” breaks down for inharmonic sounds. The consonance-based approach can help to give structure to inharmonic musical realms.

For harmonic sounds, the Western 12-tet tuning can be viewed as approximating related JI scales. The same ideas can be used to describe the 7-tet scales of traditional Thai music as well as Indonesian scales. The musical scales of these cultures can be derived from the timbre of pairs of instruments (rather than of a single instrument, as in the West).

Ideas lead to several consonance-based audio signal processing devices: adaptive keyboards, pitch/intonation correction, sound synthesis, timbral manipulation

Throughout, we have seen the enduring influence of the contributions of Helmholtz; the basic characterization of timbre in terms of spectrum, and generalizations of physical beating as a paradigm for sensory consonance.

Areas for Further Thought

Variable bandwidth transformations? wavelets or short FFTs?

Are there simple time domain operators that do interesting spectral maps?

How to (re)design acoustic instruments for play in other tunings by adjusting: mass and density of string — contour and geometry of tubes — shape and topology of resonator

Exploit other psychoacoustic phenomena in audio devices

- masking effects
- fusion/fissioning of sounds
- auditory illusions

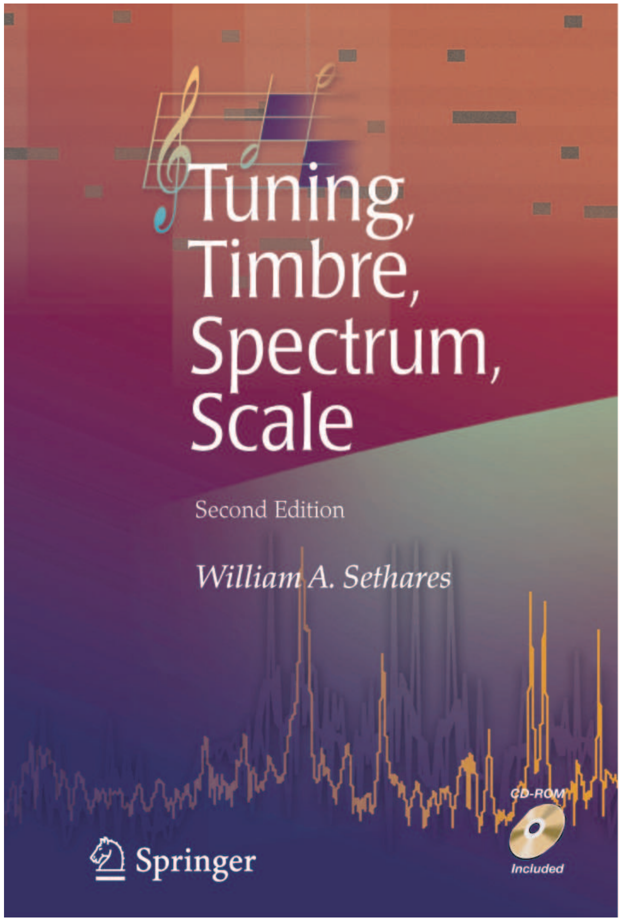

Tuning and Timbre: A Perceptual Synthesis

Bill Sethares

IDEA: Exploit psychoacoustic studies on the perception of consonance and dissonance. The talk begins by showing how to build a device that can measure the "sensory" consonance and/or dissonance of a sound in its musical context. Such a "dissonance meter" has implications in music theory, in synthesizer design, in the construction of musical scales and tunings, and in the design of musical instruments.

...the legacy of Helmholtz continues...