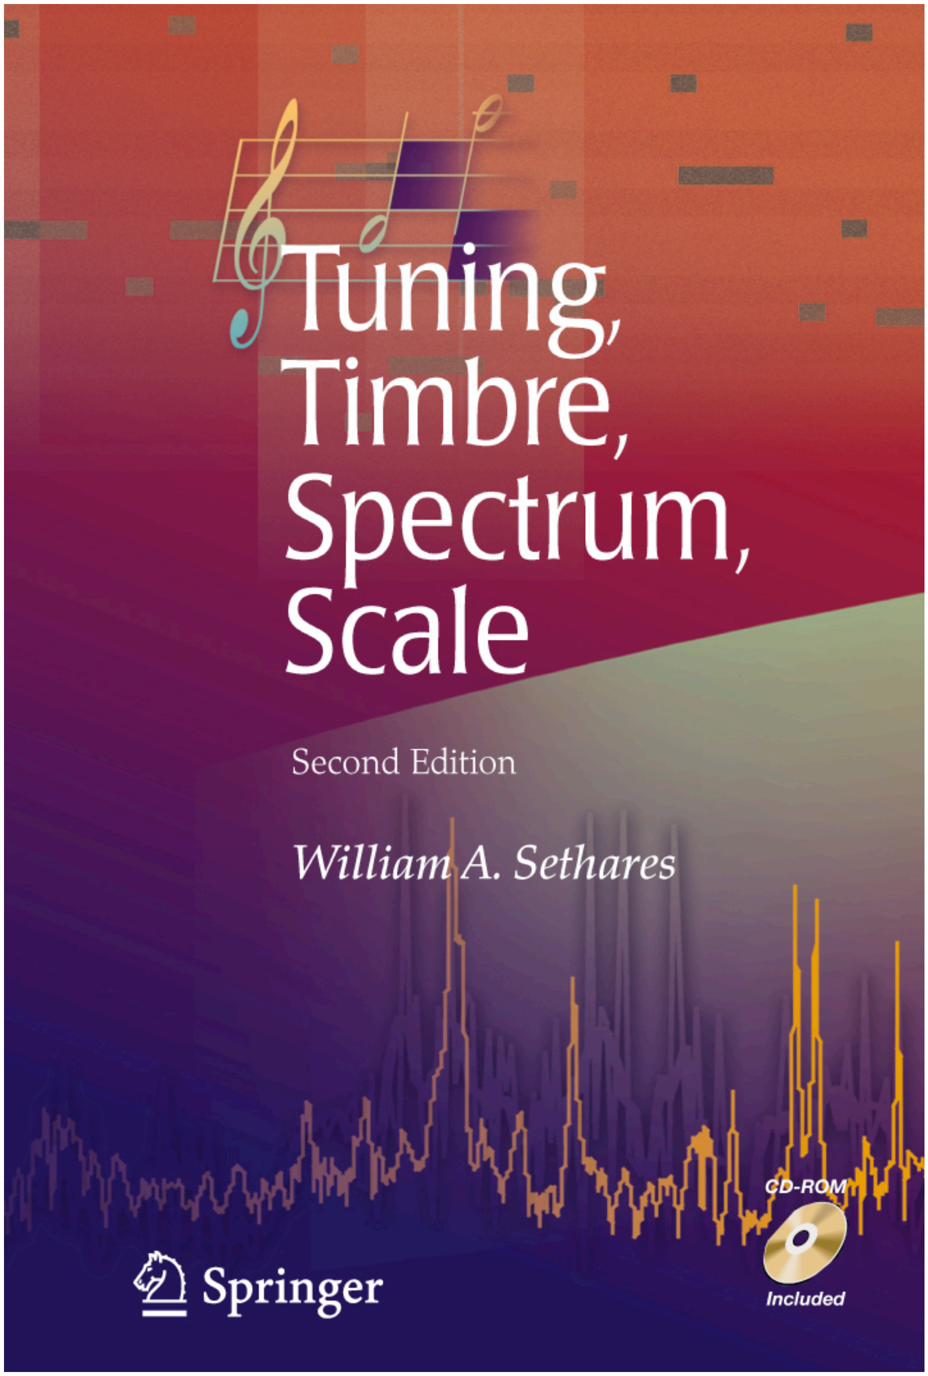
Tuning,
Timbre,
Spectrum,
Scale
Second Edition
William A. Sethares
Springer
CD-ROM
Included

Prelude

The chords sounded smooth and nondissonant but strange and somewhat eerie. The effect was so different from the tempered scale that there was no tendency to judge in-tuneness or out-of-tuneness. It seemed like a peek into a new and unfamiliar musical world, in which none of the old rules applied, and the new ones, if any, were undiscovered. F. H. Slaymaker [B: 176]

To seek out new tonalities, new timbres...
To boldly listen to what no one has heard before.

Several years ago I purchased a musical synthesizer with an intriguing feature—each note of the keyboard could be assigned to any desired pitch. This freedom to arbitrarily specify the tuning removed a constraint from my music that I had never noticed or questioned—playing in 12-tone equal temperament.[1] Suddenly, new musical worlds opened, and I eagerly explored some of the possibilities: unequal divisions of the octave, n equal divisions, and even some tunings not based on the octave at all.

Curiously, it was much easier to play in some tunings than in others. For instance, 19-tone equal temperament (*19-tet*) with its 19 equal divisions of the octave is easy. Almost any kind of sampled or synthesized instrument plays well: piano sounds, horn samples, and synthesized flutes all mesh and flow. 16-tet is harder, but still feasible. I had to audition hundreds of sounds, but finally found a few good sounds for my 16-tet chords. In 10-tet, though, none of the tones in the synthesizers seemed right on sustained harmonic passages. It was hard to find pairs of notes that sounded reasonable together, and triads were nearly impossible. Everything appeared somewhat out-of-tune, even though the tuning was precisely ten tones per octave. Somehow the timbre, or tone quality of the sounds, seemed to be interfering.

The more I experimented with alternative tunings, the more it appeared that certain kinds of scales sound good with some timbres and not with others. Certain kinds of timbres sound good in some scales and not in others. This raised a host of questions: What is the relationship between the timbre of a sound and the intervals, scale, or tuning in which the sound appears "in tune?" Can this relationship be expressed in precise terms? Is there an underlying pattern?

[1] This is the way modern pianos are tuned. The seven white keys form the major scale, and the five black keys fill in the missing tones so that the perceived distance between adjacent notes is (roughly) equal.

This book answers these questions by drawing on recent results in psychoacoustics, which allow the relationship between timbre and tuning to be explored in a clear and unambiguous way. Think of these answers as a model of musical perception that makes predictions about what you hear: about what kinds of timbres are appropriate in a given musical context, and what kind of musical context is suitable for a given timbre.

Tuning, Timbre, Spectrum, Scale begins by explaining the relevant terms from the psychoacoustic literature. For instance, the perception of "timbre" is closely related to (but also distinct from) the physical notion of the *spectrum* of a sound. Similarly, the perception of "in-tuneness" parallels the measurable idea of *sensory consonance*. The key idea is that consonance and dissonance are not inherent qualities of intervals, but they are dependent on the spectrum, timbre, or tonal quality of the sound. To demonstrate this, the first sound example on the accompanying CD plays a short phrase where the octave has been made dissonant by devious choice of timbre, even though other, nonoctave intervals remain consonant. In fact, almost any interval can be made dissonant or consonant by proper sculpting of the timbre.

Dissonance curves provide a straightforward way to predict the most consonant intervals for a given sound, and the set of most-consonant intervals defines a scale *related* to the specified spectrum. These allow musicians and composers to design sounds according to the needs of their music, rather than having to create music around the sounds of a few common instruments. The spectrum/scale relationship provides a map for the exploration of inharmonic musical worlds.

To the extent that the spectrum/scale connection is based on properties of the human auditory system, it is relevant to other musical cultures. Two important independent musical traditions are the gamelan ensembles of Indonesia (known for their metallophones and unusual five and seven-note scales) and the percussion orchestras of classical Thai music (known for their xylophone-like idiophones and seven-tone equal-tempered scale). In the same way that instrumental sounds with harmonic partials (for instance, those caused by vibrating strings and air columns) are closely related to the scales of the West, so the scales of the gamelans are related to the spectrum, or tonal quality, of the instruments used in the gamelan. Similarly, the unusual scales of Thai classical music are related to the spectrum of the xylophone-like *renat*.

But there's more. The ability to measure sensory consonance in a reliable and perceptually relevant manner has several implications for the design of audio signal processing devices and for musical theory and analysis. Perhaps the most exciting of these is a new method of *adaptive tuning* that can automatically adjust the tuning of a piece based on the timbral character of the music so as to minimize dissonance. Of course, one might cunningly seek to maximize dissonance; the point is that the composer or performer can now directly control this perceptually relevant parameter.

The first several chapters present the key ideas in a nonmathematical way. The later chapters deal with the nitty-gritty issues of sound generation and manipulation, and the text becomes denser. For readers without the background to read these sections, I would counsel the pragmatic approach of skipping the details and focusing on the text and illustrations.

Fortunately, given current synthesizer technology, it is not necessary to rely only on exposition and mathematical analysis. You can actually listen to the sounds and the tunings, and verify for yourself that the predictions of the model correspond to what you hear. This is the purpose of the accompanying CD. Some tracks are designed to fulfill the predictions of the model, and some are designed to violate them; it is not hard to tell the difference. The effects are not subtle.

Madison, Wisconsin, USA *William A. Sethares*
August 2004

Contents

1

The Octave Is Dead . . . Long Live the Octave

1.1 A Challenge

The octave is the most consonant interval after the unison. A low C on the piano sounds “the same” as a high C. Scales “repeat” at octave intervals. These common-sense notions are found wherever music is discussed:

> The most basic musical interval is the octave, which occurs when the frequency of any tone is doubled or halved. Two tones an octave apart create a feeling of identity, or the duplication of a single pitch in a higher or lower register.[1]

Harry Olson[2] uses “pleasant” rather than “consonant”:

> An interval between two sounds is their spacing in pitch or frequency... It has been found that the octave produces a pleasant sensation... 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.

W. A. Mathieu[3] discusses the octave far more poetically:

> The two sounds are the same and different. Same name, same “note” (whatever that is), but higher pitch. When a man sings nursery rhymes with a child, he is singing precisely the same song, but lower than the child. They are singing together, but singing apart. There is something easy in the harmony of two tones an octave apart - played either separately or together - but an octave transcends *easy*. There is a way in which the tones are identical.

Arthur Benade[4] observes that the similarity between notes an octave apart has been enshrined in many of the world’s languages:

---

1 From [B: 66].
2 [B: 123].
3 [B: 104].
4 [B: 12].

> Musicians of all periods and all places have tended to agree that when they hear a tone having a repetition frequency double that of another one, the two are very nearly interchangeable. This similarity of a tone with its octave is so striking that in most languages both tones are given the same name.

Anthony Storr[5] is even more emphatic:

> The octave is an acoustic fact, expressible mathematically, which is not created by man. The composition of music requires that the octave be taken as the most basic relationship.

Given all this, the reader may be surprised (and perhaps a bit incredulous) to hear a tone that is distinctly dissonant when played in the interval of an octave, yet sounds nicely consonant when played at some other, nonoctave interval. This is exactly the demonstration provided in the first sound example[6] [S: 1] and repeated in the first video example[7] [V: 1]. The demonstration consists of only a handful of notes, as shown in Fig. 1.1.

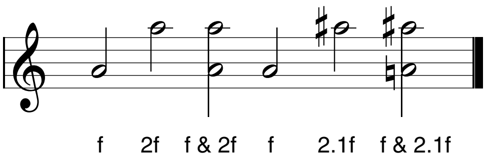

Fig. 1.1. In sound example [S: 1] and video example [V: 1], the timbre of the sound is constructed so that the octave between f and $2f$ is dissonant while the nonoctave f to $2.1f$ is consonant. Go listen to this example now.

A note is played (with a fundamental frequency $f = 450$ Hz[8]) followed by its octave (with fundamental at $2f = 900$ Hz). Individually, they sound normal enough, although perhaps somewhat "electronic" or bell-like in nature. But when played simultaneously, they clash in a startling dissonance. In the second phrase, the same note is played, followed by a note with fundamental at $2.1f = 945$ Hz (which falls just below the highly dissonant interval usually called the augmented octave or minor 9th). Amazingly, this second, nonoctave (and even microtonal) interval appears smooth and restful, even consonant; it has many of the characteristics usually associated with the octave. Such an interval is called a *pseudo-octave*.

Precise details of the construction of the sound used in this example are given later. For now, it is enough to recognize that the tonal makeup of the sound was carefully chosen *in conjunction with* the intervals used. Thus, the "trick" is to choose the spectrum or timbre of the sound (the tone quality) to match the tuning (the intervals desired).

---

5 [B: 184].

6 Beginning on p. 381 is a listing of all sound examples (references to sound examples are prefaced with [S:]) along with instructions for accessing them with a computer.

7 Beginning on p. 393 is a listing of all video examples (references to video examples are prefaced with [V:]) along with instructions for accessing them with a computer.

8 *Hz* stands for *Hertz*, the unit of frequency. One Hertz equals one cycle per second.

As will become apparent, there is a relationship between the kinds of sounds made by Western instruments (i.e., harmonic[9] sounds) and the kinds of intervals (and hence scales) used in conventional Western tonal music. In particular, the 2:1 octave is important precisely because the first two partials of a harmonic sound have 2:1 ratios. Other kinds of sounds are most naturally played using other intervals, for example, the 2.1 pseudo-octave. Stranger still, there are inharmonic sounds that suggest no natural or obvious interval of repetition. Octave-based music is only one of a multitude of possible musics. As future chapters show, it is possible to make almost any interval reasonably consonant, or to make it wildly dissonant, by properly sculpting the spectrum of the sound.

Sound examples [S: 2] to [S: 5] are basically an extended version of this example, where you can better hear the clash of the dissonances and the odd timbral character associated with the inharmonic stretched sounds. The "same" simple piece is played four ways:

[S: 2] Harmonic sounds in 12-tet
[S: 3] Harmonic sounds in the 2.1 stretched scale
[S: 4] 2.1 stretched timbres in the 2.1 stretched scale
[S: 5] 2.1 stretched timbres in 12-tet

where *12-tet* is an abbreviation for the familiar 12-tone per octave equal tempered scale, and where the *stretched scale*, based on the 2.1 pseudo-octave, is designed specially for use with the stretched timbres. When the timbres and the scales are matched (as in [S: 2] and [S: 4]), there is contrast between consonance and dissonance as the chords change, and the piece has a sensible musical flow (although the timbral qualities in [S: 4] are decidedly unusual). When the timbres and scales do not match (as in [S: 3] and [S: 5]), the piece is uniformly dissonant. The difference between these two situations is not subtle, and it calls into question the meaning of basic terms like timbre, consonance, and dissonance. It calls into question the octave as the most consonant interval, and the kinds of harmony and musical theories based on that view. In order to make sense of these examples, *Tuning, Timbre, Spectrum, Scale* uses the notions of *sensory consonance* and *sensory dissonance*. These terms are carefully defined in Chap. 3 and are contrasted with other notions of consonance and dissonance in Chap. 5.

1.2 A Dissonance Meter

Such shaping of spectra and scales requires that there be a convenient way to measure the dissonance of a given sound or interval. One of the key ideas underlying the sonic manipulations in *Tuning, Timbre, Spectrum, Scale* is the construction of a "dissonance meter." Don't worry—no soldering is required. The dissonance meter is a computer program that inputs a sound in digital form and outputs a number proportional to the (sensory) dissonance or consonance of the sound. For longer musical

[9] Here *harmonic* is used in the technical sense of a sound with overtones composed exclusively of integer multiples of some audible fundamental.

passages with many notes, the meter can be used to measure the dissonance within each specified time interval, for instance, within each measure or each beat. As the *challenging the octave* example shows, the dissonance meter must be sensitive to both the tuning (or pitch) of the sounds and to the spectrum (or timbre) of the tones.

Although such a device may seem frivolous at first glance, it has many real uses:

As an audio signal processing device: The dissonance meter is at the heart of a device that can automatically reduce the dissonance of a sound, while leaving its character more or less unchanged. This can also be reversed to create a sound that is more dissonant than the input. Combined, this provides a way to directly control the perceived dissonance of a sound.

Adaptive tuning of musical synthesizers: While monitoring the dissonance of the notes commanded by a performer, the meter can be used to adjust the tuning of the notes (microtonally) to minimize the dissonance of the passage. This is a concrete way of designing an adaptive or dynamic tuning.

Exploration of inharmonic sounds: The dissonance meter shows which intervals are most consonant (and which most dissonant) as a function of the spectrum of the instrument. As the *challenging the octave* example shows, unusual sounds can be profitably played in unusual intervals. The dissonance meter can concretely specify related intervals and spectra to find tunings most appropriate for a given timbre. This is a kind of map for the exploration of inharmonic musical spaces.

Exploration of "arbitrary" musical scales: Each timbre or spectrum has a set of intervals in which it sounds most consonant. Similarly, each set of intervals (each musical scale) has timbres with spectra that sound most consonant in that scale. The dissonance meter can help find timbres most appropriate for a given tuning.

Analysis of tonal music and performance: In tonal systems with harmonic instruments, the consonance and dissonance of a musical passage can often be read from the score because intervals within a given historical period have a known and relatively fixed degree of consonance and/or dissonance. But performances may vary. A dissonance meter can be used to measure the actual dissonance of different performances of the same piece.

Analysis of nontonal and nonwestern music and performance: Sounds played in intervals radically different from those found in 12-tet have no standard or accepted dissonance value in standard music theory. As the dissonance meter can be applied to any sound at any interval, it can be used to help make musical sense of passages to which standard theories are inapplicable. For instance, it can be used to investigate nonwestern music such as the gamelan, and modern atonal music.

Historical musicology: Many historical composers wrote in musical scales (such as meantone, Pythagorean, Just, etc.) that are different from 12-tet, but they did not document their usage. By analyzing the choice of intervals, the dissonance meter can make an educated guess at likely scales using only the extant music. Chapter 11, on "Musicological Analysis," investigates possible scales used by Domenico Scarlatti.

As an intonation monitor: Two notes in unison are very consonant. When slightly out of tune, dissonances occur. The dissonance meter can be used to monitor the intonation of a singer or instrumentalist, and it may be useful as a training device.

The ability to measure dissonance is a crucial component in several kinds of audio devices and in certain methods of musical analysis. The idea that dissonance is a function of the timbre of the sound as well as the musical intervals also has important implications for the understanding of nonwestern musics, modern atonal and experimental compositions, and the design of electronic musical instruments.

1.3 New Perspectives

The dissonance curve plots how much sensory dissonance occurs at each interval, given the spectrum (or timbre) of a sound. Many common Western orchestral (and popular) instruments are primarily harmonic, that is, they have a spectrum that consists of a fundamental frequency along with partials (or overtones) at integer multiples of the fundamental. This spectrum can be used to draw a dissonance curve, and the minima of this curve occur at or near many of the steps of the Western scales. This suggests a relationship between the spectrum of the instruments and the scales in which they are played.

Nonwestern Musics

Many different scale systems have been and still are used throughout the world. In Indonesia, for instance, gamelans are tuned to five and seven-note scales that are very different from 12-tet. The timbral quality of the (primarily metallophone) instruments is also very different from the harmonic instruments of the West. The dissonance curve for these metallophones have minima that occur at or near the scale steps used by the gamelans.[10] Similarly, in Thailand, there is a classical music tradition that uses wooden xylophone-like instruments called *renats* that play in (approximately) 7-tet. The dissonance curve for renat-like timbres have minima that occur near many of the steps of the traditional 7-tet Thai scale, as shown in Chap. 15. Thus, the musical scales of these nonwestern traditions are related to the inharmonic spectra of the instruments, and the idea of related spectra and scales is applicable cross culturally.

[10] See Chap. 10 "The Gamelan" for details and caveats.

New Scales

Even in the West, the present 12-tet system is a fairly recent innovation, and many different scales have been used throughout history. Some systems, such as those used in the Indonesian gamelan, do not even repeat at octave intervals. Can *any* possible set of intervals or frequencies form a viable musical scale, assuming that the listener is willing to acclimate to the scale?

Some composers have viewed this as a musical challenge. Easley Blackwood's *Microtonal Etudes* might jokingly be called the "Ill-Tempered Synthesizer" because it explores all equal temperaments between 13 and 24. Thus, instead of 12 equal divisions of the octave, these pieces divide the octave into 13, 14, 15, and more equal parts. Ivor Darreg composed in many equal temperaments,[11] exclaiming

> the striking and characteristic moods of many tuning-systems will become the most powerful and compelling reason for exploring beyond 12-tone equal temperament. It is necessary to have more than one non-twelve-tone system before these moods can be heard and their significance appreciated.[12]

Others have explored nonequal divisions of the octave[13] and even various subdivisions of nonoctaves.[14] It is clearly possible to make music in a large variety of tunings. Such music is called *xenharmonic*,[15] strange "harmonies" unlike anything possible in 12-tet.

The intervals that are most consonant for harmonic sounds are made from small integer ratios such as the octave (2:1), the fifth (3:2), and the fourth (4:3). These simple integer ratio intervals are called *just* intervals, and they collectively form scales known as *just intonation* scales. Many of the just intervals occur close to (but not exactly at[16]) steps of the 12-tet scale, which can be viewed as an acceptable approximation to these just intervals. Steps of the 19-tet scale also approximate many of the just intervals, but the 10-tet scale steps do not. This suggests why, for instance, it is easy to play in 19-tet and hard to play in 10-tet using harmonic tones—there are many consonant intervals in 19-tet but few in 10-tet.

New Sounds

The *challenging the octave* demonstration shows that certain unusual intervals can be consonant when played with certain kinds of unusual sounds. Is it possible to make *any* interval consonant by properly manipulating the sound quality? For instance, is it possible to choose the spectral character so that many of the 10-tet intervals became consonant? Would it then be "easy" to play in 10-tet? The answer is "yes,"

[11] [D: 10].
[12] From [B: 36], No. 5.
[13] For instance, Vallotti, Kirchenberg, and Partch.
[14] For instance, Carlos [B: 23], Mathews and Pierce [B: 102], and McLaren [B: 108].
[15] Coined by Darreg [B: 36], from the Greek *xenos* for strange or foreign.
[16] Table 6.1 on p. 97 shows how close these approximations are.

and part of this book is dedicated to exploring ways of manipulating the spectrum in an appropriate manner.

Although Western music relies heavily on harmonic sounds, these are only one of a multitude of kinds of sound. Modern synthesizers can easily generate inharmonic sounds and transport us into unexplored musical realms. The spectrum/scale connection provides a guideline for exploration by specifying the intervals in which the sounds can be played most consonantly or by specifying the sounds in which the intervals can be played most consonantly. Thus, the methods allow the composer to systematically specify the amount of consonance or dissonance. The composer has a new and powerful method of control over the music.

Consider a fixed scale in which all intervals are just. No such scale can be modulated through all the keys. No such scale can play all the consonant chords even in a single key. (These are arithmetic impossibilities, and a concrete example is provided on p. 153.) But using the ideas of sensory consonance, it is possible to adapt the pitches of the notes dynamically. For harmonic tones, this is equivalent to playing in simple integer (just) ratios, but allows modulation to any key, thus bypassing this ancient problem. Although previous theorists had proposed that such dynamic tunings might be possible,[17] this is the first concrete method that can be applied to any chord in any musical setting. *It is possible to have your just intonation and to modulate, too!* Moreover, the adaptive tuning method is not restricted to harmonic tones, and so it provides a way to "automatically" play in the related scale (the scale consisting of the most consonant intervals, given the spectral character of the sound).

New "Music Theories"

When working in an unfamiliar system, the composer cannot rely on musical intuition developed through years of practice. In 10-tet, for instance, there are no intervals near the familiar fifths or thirds, and it is not obvious what intervals and chords make musical sense. The ideas of sensory consonance can be used to find the most consonant chords, as well as the most consonant intervals (as always, sensory consonance is a function of the intervals and of the spectrum/timbre of the sound), and so it can provide a kind of sensory map for the exploration of new tunings and new timbres. Chapter 14 develops a new music theory for 10-tet. The "neutral third" chord is introduced along with the "circle of thirds" (which is somewhat analogous to the familiar circle of fifths in 12-tet). This can be viewed as a prototype of the kinds of theoretical constructs that are possible using the sensory consonance approach, and pieces are included on the CD to demonstrate that the predictions of the model are valid in realistic musical situations.

Unlike most theories of music, this one does not seek (primarily) to explain a body of existing musical practice. Rather, like a good scientific theory, it makes concrete predictions that can be readily verified or falsified. These predictions involve how (inharmonic) sounds combine, how spectra and scales interact, and how dissonance varies as a function of both interval and spectrum. The enclosed CD provides

[17] See Polansky [B: 142] and Waage [B: 202].

examples so that you can verify for yourself that the predictions correspond to perceptual reality.

Tuning and spectrum theories are independent of musical style; they are no more "for" classical music than they are "for" jazz or pop. It would be naive to suggest that complex musical properties such as style can be measured in terms of a simple sensory criterion. Even in the realm of harmony (and ignoring musically essential aspects such as melody and rhythm), sensory consonance is only part of the story. A harmonic progression that was uniformly consonant would be tedious; harmonic interest arises from a complex interplay of restlessness and restfulness,[18] of tension and resolution. It is easy to increase the sensory dissonance, and hence the restlessness, by playing more notes (try slamming your arm on the keyboard). But it is not always as easy to increase the sensory consonance and hence the restfulness. By playing sounds in their related scales, it is possible to obtain the greatest contrast between consonance and dissonance for a given sound palette.

1.4 Overview

While introducing the appropriate psychoacoustic jargon, Chap. 2 (the "Science of Sound") draws attention to the important distinction between what we perceive and what is really (measurably) there. Any kind of "perceptually intelligent" musical device must exploit the measurable in order to extract information from the environment, and it must then shape the sound based on the perceptual requirements of the listener. Chapter 3 looks carefully at the case of two simultaneously sounding sine waves, which is the simplest situation in which sensory dissonances occur.

Chapter 4 reviews several of the common organizing principles behind the creation of musical scales, and it builds a library of historical and modern scales that will be used throughout the book as examples.

Chapter 5 gives an overview of the many diverse meanings that the words "consonance" and "dissonance" have enjoyed throughout the centuries. The relatively recent notion of sensory consonance is then adopted for use throughout the remainder of the book primarily because it can be readily measured and quantified.

Chapter 6 introduces the idea of a *dissonance curve* that displays (for a sound with a given spectrum) the sensory consonance and dissonance of all intervals. This leads to the definition of a *related* spectrum and scale, a sound for which the most consonant intervals occur at precisely the scale steps. Two complementary questions are posed. Given a spectrum, what is the related scale? Given a scale, what is a related spectrum? The second, more difficult question is addressed at length in Chap. 12, and Chap. 7 ("A Bell, A Rock, A Crystal") gives three detailed examples of how related spectra and scales can be exploited in musical contexts. This is primarily interesting from a compositional point of view.

Chapter 8 shows how the ideas of sensory consonance can be exploited to create a method of adaptive tuning, and it provides several examples of "what to expect"

[18] Alternative definitions of dissonance and consonance are discussed at length in Chap. 5.

from such an algorithm. Chapter 9 highlights three compositions in adaptive tuning and discusses compositional techniques and tradeoffs. Musical compositions and examples are provided on the accompanying CD.

The remaining chapters can be read in any order. Chapter 10 shows how the pelog and slendro scales of the Indonesian gamelan are correlated with the spectra of the metallophones on which they are played. Similarly, Chap. 15 shows how the scales of Thai classical music are related to the spectra of the Thai instruments.

Chapter 11 explores applications in musicology. The *dissonance score* can be used to compare different performances of the same piece, or to examine the use of consonances and dissonances in unscored and nonwestern music. An application to historical musicology shows how the tuning preferences of Domenico Scarlatti can be investigated using only his extant scores.

Chapter 14 explores one possible alternative musical universe, that of 10-tet. This should only be considered a preliminary foray into what promises to be a huge undertaking—codifying and systematizing music theories for non-12-tet. Although it is probably impossible to find a "new" chord in 12-tet, it is impossible to play in n-tet without creating new harmonies, new chordal structures, and new kinds of musical passages.

Chapters 12 and 13 are the most technically involved. They show how to specify spectra for a given tuning, and how to create rich and complex sounds with the specified spectral content.

The final chapter sums up the ideas in *Tuning, Timbre, Spectrum, Scale* as exploiting a single perceptual measure (that of sensory consonance) and applying it to musical theory, practice, and sound design. As we expand the palette of timbres we play, we will naturally begin to play in new intervals and new tunings.