CHAPTER 3

The Overtone Series and the Spectrum View

In the last chapter, we investigated a variety of sound properties using the waveform view of sound. That representation is useful in many instances, but it falls short with regard to timbre. The physical property related to timbre found in the waveform view is the waveform itself. However, the small collection of standard waveforms (sine, triangle, sawtooth, square, pulse) is of limited use when discussing real-world timbres. A better representation of sound for investigating timbre is the spectrum view. To understand the spectrum view it is useful first to consider the more familiar **overtone series**.

OVERTONE SERIES

Any note whose timbre is more complex than a sine wave contains the frequency that is heard as the pitch of the note—the **fundamental frequency**—*plus* some frequencies above that that are heard not as pitch but as the "color" or timbre of the sound. The overtone series represents the frequencies that are present in a single note, including the fundamental, and is usually shown using traditional music notation (see Figure 3.1; the fundamental is A).

This traditional notation is somewhat misleading, because it implies that every note is really a chord and that all the frequencies are distinct pitches. There is really just *one* pitch associated with an overtone series—the pitch related to the fundamental frequency. However, there are some useful features about the frequencies in the overtone series that can be seen with traditional notation. For example, for brass players, the fundamental frequencies of the pitches available at each valve fingering or slide position are found in the overtone series that is built on the lowest note at that fingering or position.

As implied by the name "overtone" series, the frequencies above the fundamental are often referred to as **overtones**. However, the terms **harmonics** and **partials** are also used. The distinction between these terms is subtle. The term overtone implies that those frequencies are over the fundamental, so the overtone series would consist of the

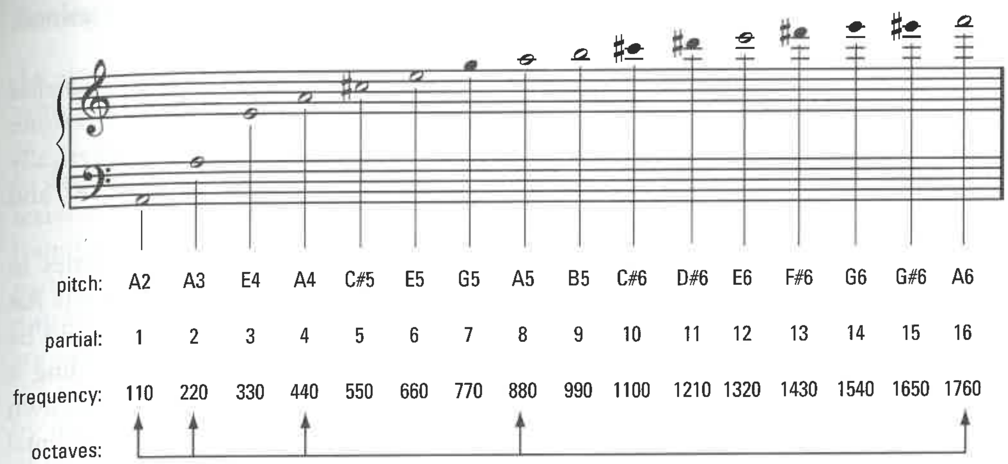

Figure 3.1 The first 16 partials of the overtone series built on A2. (From Holmes, 2008)

fundamental frequency plus the first overtone, the second overtone, etc. The term harmonic is somewhat ambiguous. It could refer to the frequencies above the fundamental: the fundamental, the first harmonic, the second harmonic, etc. However, it could include the fundamental as the first harmonic, so the series would be: first harmonic (fundamental), second harmonic, third harmonic, etc.

The term partial implies that all of the frequencies in a sound are all just parts of the sound: first partial (fundamental), second partial, third partial, etc. This term has a distinct advantage in that not every sound has frequencies that follow the overtone series, so the term partial can also be applied to frequencies of those sounds as well, whereas the terms overtone and harmonic only apply to sounds that follow the overtone series. Sounds whose frequencies follow the overtone series are referred to as "harmonic," and sounds whose frequencies do not follow the overtone series are referred to as "inharmonic"—not "enharmonic" as in G♯ and A♭, but "inharmonic" as in not harmonic. Most of the sounds in the world are inharmonic, but many of the sounds that we are concerned about in music, such as sounds made by many musical instruments, are harmonic. Inharmonic sounds will be discussed later in the chapter. This text will primarily use the term partial to describe frequencies in a spectrum and number them accordingly with the fundamental being the first partial (see Figure 3.1).

To see what the relationships are between the frequencies in the overtone series, we need to find some point of reference in this notational representation. Since this overtone series is based on A, it includes the familiar tuning A as the fourth partial, which has a frequency of 440 Hz. The only other fact we need to know is that octaves have a frequency relationship of 2 to 1. Armed with that information, we can see that the second partial, which is an octave below the fourth partial, would have a frequency of $\frac{1}{2} \times 440 = 220$. Similarly, the first partial is an octave below the second, giving it a frequency of 110, and the eighth partial is an octave above the fourth partial, giving it a frequency of 2 × 440 = 880. Table 3.1 shows the relationships derived so far (note that for the pitch-octave notation here, middle C is C4).

From Table 3.1, you can see that each partial's frequency is the *partial number multiplied by the fundamental frequency*. Applying this principle to the other partials in this overtone series, you get the frequencies given in Table 3.2. If the fundamental is more generically given as some frequency *f*, then the partial frequencies are 2*f*, 3*f*, 4*f*, 5*f*, 6*f*, 7*f*, 8*f*, and so on.

Table 3.2 shows an additional drawback to representing the overtones series in traditional notation. The "G5" (remember it's a frequency, not a note) in the table has a frequency of 770 Hz. If you were to play G5 on a piano, the frequency would be 783.99 Hz (see Figure 2.4 on page 21). This needn't be seen as too troubling a discrepancy, because the G5 played on the piano is a note in its own right with its own overtone series, and the "G5" that is part of the overtone series built on the fundamental A2 is part of the timbre of that sound. Nevertheless, the frequency relationships found in the overtone series have inspired many different approaches to tuning.

Table 3.1 **Frequency relationships for overtone series on A (incomplete)**

Partial #	*"Note"*	*Frequency*
1 (fundamental)	A2	110
2	A3	220
4	A4	440
8	A5	880

Table 3.2 **Frequency relationships for overtone series on A**

Partial #	*"Note"*	*Frequency*
1 (fundamental)	A2	110
2	A3	220
3	E4	330
4	A4	440
5	C#5	550
6	E5	660
7	G5	770
8	A5	880
9	B5	990
10	C#6	1100
11	D#6	1210
12	E6	1320
13	F#6	1430
14	G6	1540
15	D#6	1650
16	A6	1760
...	...	...

Tuning and Temperament

To derive the relationships between partial number and frequency above, we started with the fact that the ratio of two frequencies an octave apart is 2 to 1, often notated as 2:1. If we continue to look at the overtone series in traditional notation, it is possible to derive ideal ratios for other intervals as well. A perfect fifth is present in the overtone series as the relationship of partial 3 to partial 2, giving a 3:2 ratio. A perfect fourth is found between partial 4 and partial 3, giving a 4:3 ratio. A major third is found between partial 5 and partial 4, giving a 5:4 ratio, and a minor third is found between partial number 6 and partial number 5, giving a 6:5 ratio. It is important to note that these are *ideal* relationships. In practice they can present some difficulties.

One such difficulty is that it is possible to start at one note and get to another note by different intervals and end up with contradictory frequencies. As a classic example of this, if you start with one note and go up repeatedly by a perfect fifth (proceeding through the circle of fifths) until you reach the beginning note several octaves higher, and then do the same by octaves, you reach different frequencies. Starting on C1, the lowest C on the piano (with C4 being middle C), you can get back to the pitch C by going up twelve fifths, and to that same C by going up seven octaves. This is C8, the highest C on the piano.

Using the interval ratio derived from the overtone series for a fifth of 3:2, the frequency for each successive fifth is generated by multiplying the frequency of the previous note by 3/2. The results of this are shown in Table 3.3, where the frequency of C1 is given as f. Using the interval ratios for an octave of 2:1, multiplying by 2 generates the frequency for each successive octave. The results of this are shown in Table 3.4.

You might expect, because both series of intervals arrive on the same note, that $(3/2)^{12}f$ and 2^7f would be the same. However:

$$(3/2)^{12}f = 129.75f$$

and

$$2^7f = 128f$$

Table 3.3 **Going up from C1 to C8 by fifths**

C1	*G1*	*D2*	*A2*	*E3*	*B3*	*F#4*	*C#5*	*G#5*	*D#6*	*A#6*	*F7*	*C8*
f	$3/2f$	$(3/2)^2f$	$(3/2)^3f$	$(3/2)^4f$	$(3/2)^5f$	$(3/2)^6f$	$(3/2)^7f$	$(3/2)^8f$	$(3/2)^9f$	$(3/2)^{10}f$	$(3/2)^{11}f$	$(3/2)^{12}f$

Table 3.4 **Going up from C1 to C8 by octaves**

C1	*C2*	*C3*	*C4*	*C5*	*C6*	*C7*	*C8*
f	$2f$	2^2f	2^3f	2^4f	2^5f	2^6f	2^7f

The frequency of C1 is 32.7, so these two formulas give the frequency of C8 as:

129.75 × 32.70 = 4,243.24 Hz (going up by fifths)

and

128 × 32.70 = 4,186.01 Hz (going up by octaves)

If you look at Figure 2.4 in the previous chapter (page 21), you can see that going up by octaves gives you the frequency for C8 on the piano. Going up by fifths causes you to "overshoot" that frequency. The difference between going up by fifths and going up by octaves generates what's known as the **Pythagorean comma**.

There are a number of other discrepancies to be found by going from one note to another by different intervals, and overtone series built on different fundamentals can generate different frequencies for what is nominally the same note. What these and other discrepancies point to is that the ideal interval relationships derived from the overtone series by themselves don't form the basis for our familiar musical system. However, in isolation, intervals formed from the ideal ratios are said to sound more pure than any of the compromise tuning systems that have been developed. A variety of such compromise systems have been proposed and used over the centuries, including Pythagorean tuning, meantone intonation, just intonation, and **equal temperament**.

In equal temperament, the ratio between every semitone is exactly the same, so each interval, regardless of the starting note, is also exactly the same. Essentially, all intervals, with the exception of the octave, are slightly "wrong" in equal temperament, but it allows music using this tuning system to modulate to any key and still have intervals the same size as the original key. Johann Sebastian Bach's famous *The Well-Tempered Clavier*, which includes preludes and fugues in all 24 major and minor keys, shows the advantage of such a tuning system, though well-tempered and equal-tempered tunings are slightly different. Table 3.5 shows some ideal interval ratios and the approximate equal-tempered ratios.

Table 3.5 **Equal-tempered and ideal interval ratios**

Interval	*Ideal ratio*	*Equal-tempered ratio*	*Equal-tempered "error"*
Major second	9:8 = 1.125:1	1.122:1	Flat
Minor third	6:5 = 1.2:1	1.189:1	Flat
Major third	5:4 = 1.25:1	1.26:1	Sharp
Fourth	4:3 = 1.333:1	1.335:1	Slightly sharp
Fifth	3:2 = 1.5:1	1.498:1	Slightly flat
Minor sixth	8:5 = 1.6:1	1.587:1	Flat
Major sixth	5:3 = 1.667:1	1.681:1	Sharp
Major seventh	15:8 = 1.875:1	1.888:1	Sharp
Octave	2:1	2:1	None

The last column in Table 3.5 shows the qualitative error for each interval in equal-tempered tuning relative to the ideal ratios. In performance, many performers and conductors will adjust their tuning of chords to partially compensate for this error. For example, a performer holding the major third of a chord will often play it slightly flat relative to equal-tempered tuning to more closely approximate the pure intervals generated by the ideal ratios.

This brief discussion has really only scratched the surface of tuning issues, both historical and contemporary. There are many books and websites devoted to various tuning systems, particularly just intonation, and a number of contemporary composers have utilized such systems in their works. In addition, many hardware and software synthesizers contain resources for variable tuning.

THE SPECTRUM

The timbre of a note is determined in part by *which* frequencies are present in a sound and *how much* of them are present (their relative amplitudes). The notation representation of the overtone series has no way to show amplitude information, and shows frequency information inadequately given that the partial frequencies often don't match the equal-tempered pitches shown. To discuss timbre more generally, it is necessary to abandon traditional music notation altogether and use the **spectrum view** of sound.

The spectrum view represents sound as a graph of frequency vs. amplitude, as opposed to the waveform view that is a graph of time vs. amplitude. The spectrum view for the overtone series starting on A2 is given in Figure 3.2. The amplitudes of the frequency components are from a spectrum analysis of a trombone note.

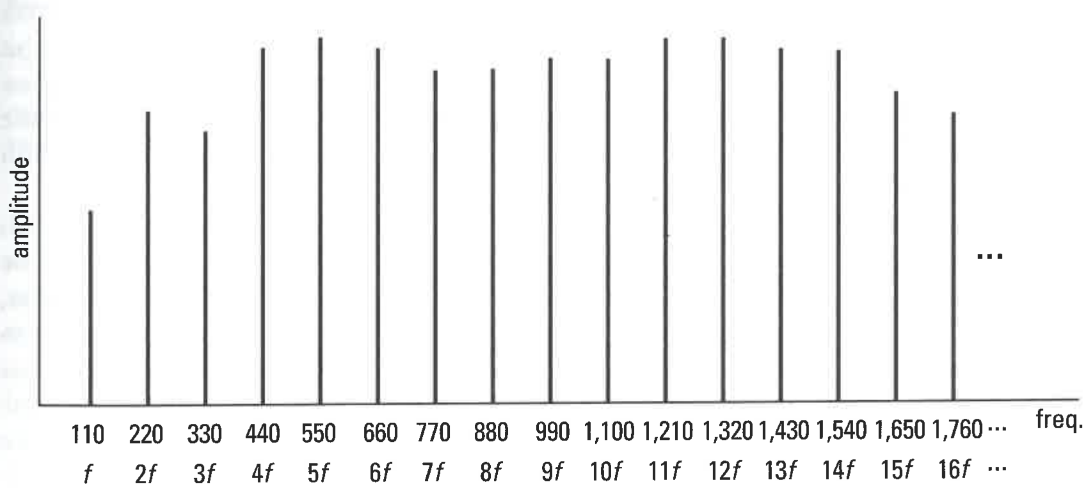

Figure 3.2 The spectrum view of sound with frequency on the *x*-axis and amplitude on the *y*-axis. The fundamental is A2 (110 Hz) and 16 total partials are shown, with the ellipses indicating that the partials continue. The relative amplitudes of the partials are based on a recording of a trombone. The frequency data is the same as in Figure 3.1 and Table 3.2.

There exist a number of standalone programs and plug-ins for DAW/sequencers that perform frequency analysis of live audio or of an audio file. The amplitude axis of this view differs from the amplitude axis of the waveform view in that the amplitudes in the spectrum view show the amplitude of each individual partial, whereas the amplitude in the waveform view is the overall amplitude of the sound. The overall amplitude in the spectrum view comes from a complex interaction between the amplitudes of each individual partial—it's not as simple as just adding them up.

It's also important to note that, in this spectrum view, there is no time axis as there is in the waveform view. To see how a spectrum changes over the course of a note or sound event—and it can change quite a bit—you would have to look at successive spectrum views that would provide a time-lapse view of the spectrum. There are variations on the spectrum view that allow three dimensions to be shown at once: frequency, amplitude, and time. One is the **spectrogram view**, which gives time vs. frequency and then shows the amplitude by the intensity of the color (see Figure 3.3).

Another of these spectrum view variations is the **waterfall spectrum**, which typically shows frequency vs. amplitude with time coming out of or into the screen in a simulated 3-D effect. There is a downward visual component to this time axis that gives it the name "waterfall."

Spectra of Basic Waveforms and Fourier's Theorem

In the previous chapter, you were introduced to a collection of basic waveforms that are largely derived from analog synthesis. To understand these basic waveforms in some more detail, we can look at their spectra. (See the book's website for audio of these examples.)

The **sine wave** is the simplest possible waveform, having only one partial: the fundamental (see Figure 3.4a). By itself, the sine wave has a pure, pale timbre that can be spooky in the right situation. The **triangle wave** contains only the odd partials (1, 3, 5, 7, etc.) but at very low amplitudes after the fundamental, so it is a little bit brighter than the sine wave and more suitable as the basis of a synthetic timbre (see Figure 3.4b). The amplitudes of the partials in a triangle wave are inversely proportional to the square of the partial number, so the third partial has a relative amplitude of 1/9, the fifth partial an amplitude of 1/25, and so on.

The **square wave** also contains only the odd partials, but in greater proportion than the triangle wave, so it sounds brighter than the triangle wave (see Figure 3.4c). The amplitudes of the partials in a square wave are inversely proportional to the partial number, so the third partial has an amplitude of 1/3, the fifth partial an amplitude of 1/5, and so on. In the octave below middle C it can sound quite a bit like a clarinet, which also has very little energy in the even partials in that register. A guitar can also produce a tone like this by plucking an open string at the twelfth fret. Timbres built from a square wave can be quite penetrating.

The **sawtooth wave** contains both even and odd partials (see Figure 3.4d). The amplitudes of the partials in a sawtooth wave are inversely proportional to the partial number, so the second partial has an amplitude of 1/2, the third partial an amplitude of

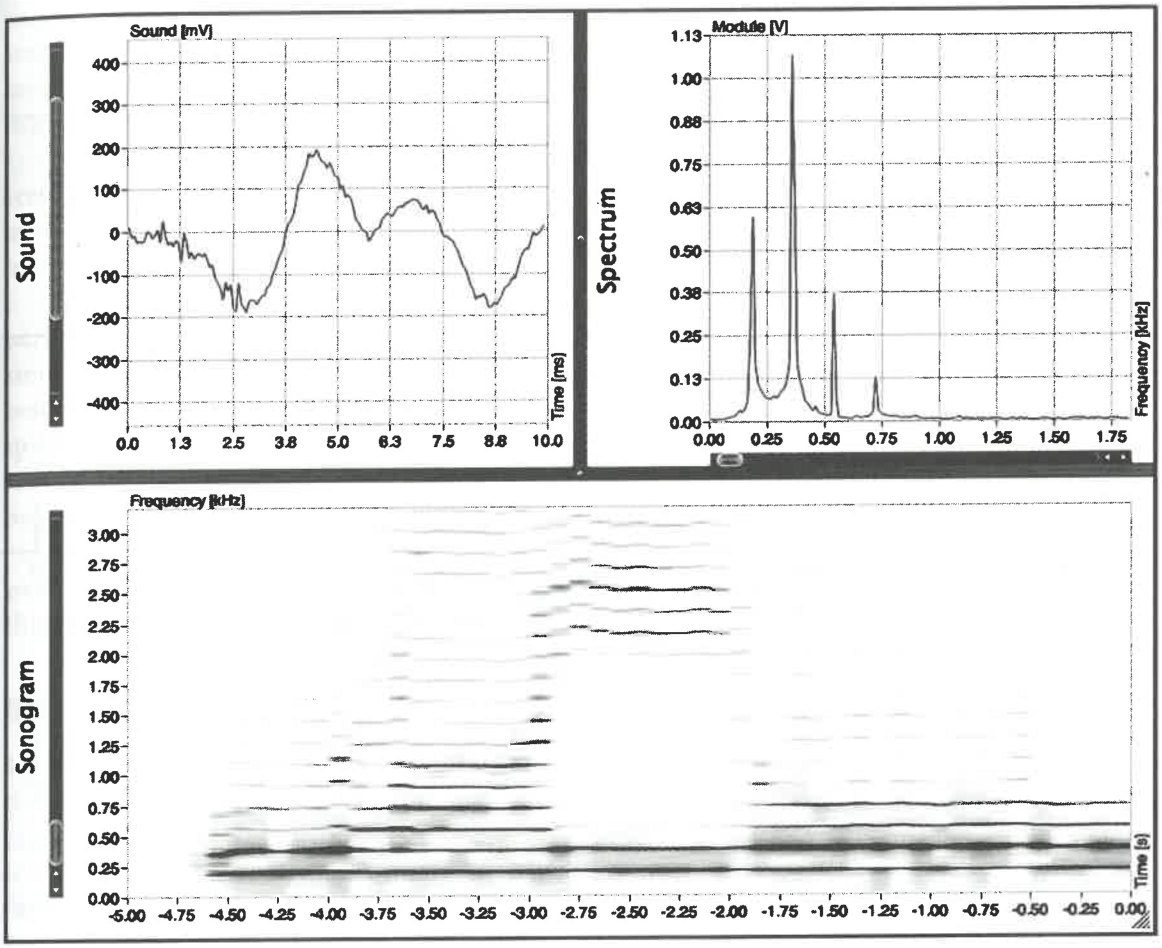

Figure 3.3 Three different views of a voice saying "oo-ah-ee-oh" on the same pitch. Upper left: waveform view frozen on "oh." Upper right: spectrum view frozen on "oh." Bottom: spectrogram view (also called the sonogram view) with time shown horizontally, frequency shown vertically, and amplitude shown by the intensity of the line. All four vowels are shown left to right. The different amplitudes of the partials for each vowel are shown clearly.

1/3, the fourth an amplitude of 1/4, and so on. As a result, the sawtooth wave is bright and nasal, which allows it to penetrate well when combined with other timbres. The sawtooth is one of the most common waveforms used for electronic timbres, particularly those meant to mimic analog synthesizer timbres.

Since the spectrum of a sine wave has only one partial, its fundamental, it is the most basic of the waveforms. In fact, each partial in the spectra of the other basic waveforms can be thought of as a separate sine wave. This implies that each of these spectra can be thought of as a sum of sine waves whose frequencies match the partial frequencies and whose amplitudes match the partial amplitudes. Figure 3.5 shows several sine waves adding together to form the beginnings of a sawtooth waveform. Many more partials would need to be added to create the sawtooth's characteristic shape.

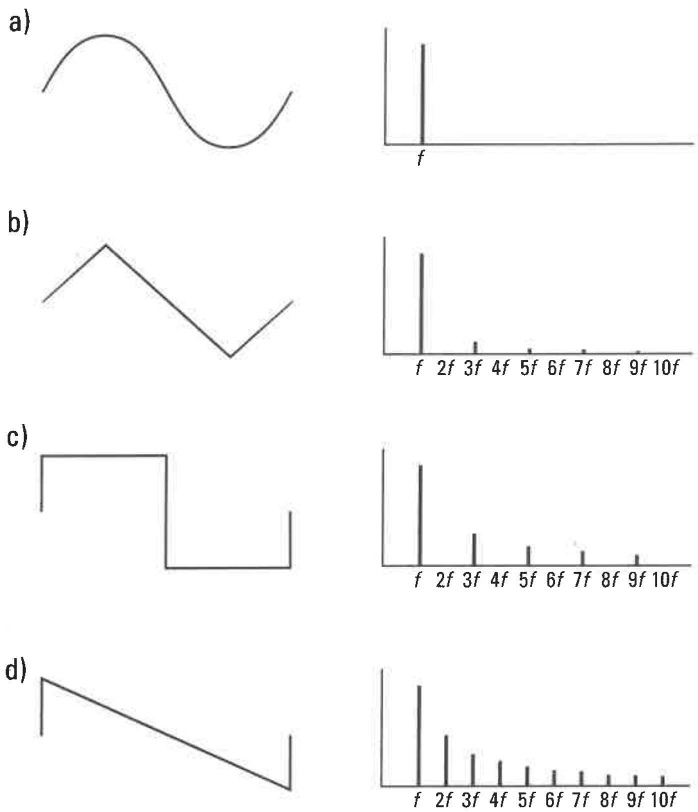

Figure 3.4 Basic waveforms and their spectra: (a) sine wave, (b) triangle wave, (c) square wave, and (d) sawtooth wave.

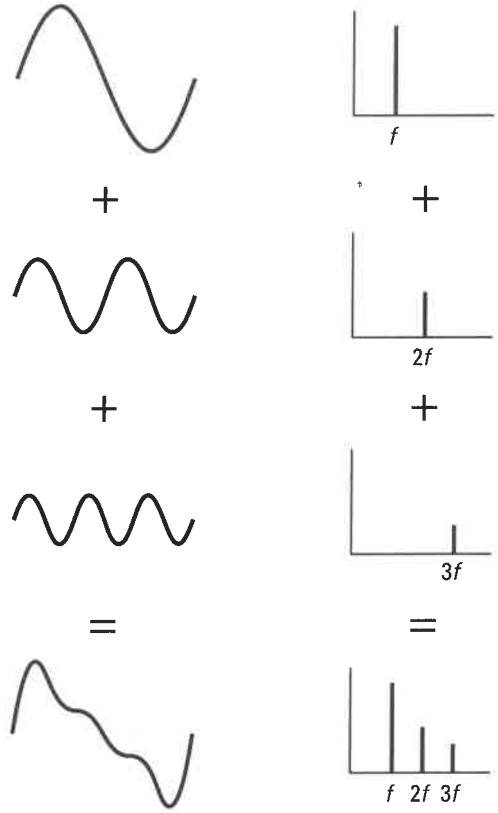

Figure 3.5 Waveform and spectrum view of three sine waves adding together to form a more complex waveform. If enough partials are added like this, a sawtooth wave will be formed.

The idea that a complex spectrum can be expressed as a sum of sine waves of various frequencies and amplitudes lies at the heart of **Fourier's theorem**. Fourier's theorem says that any periodic waveform can be expressed as a sum of sine waves. Periodic waveforms follow the overtone series, which means that most "musical" sounds made by winds, brass, strings, pianos, voices, and some percussion instruments all have spectra that can be thought of as sums of sine waves. Fourier's theorem is not only useful in the analysis of instrument spectra, but can also be used to synthesize new sounds through a technique known as additive synthesis, which is also referred to as Fourier synthesis. This method of synthesis along with a variety of others will be discussed later in the text in the chapter on synthesis methods.

Harmonic and Inharmonic Spectra

Thus far, we've been assuming that all sounds have partials that follow the overtone series. The spectra for those sounds are termed **harmonic spectra**. Despite the fact that sounds with harmonic spectra are only a small subset of all the possible sounds, it just so happens that many of the musical sounds we care about are part of this subset, including those made by brass, woodwinds, strings, pianos, voices, and certain percussion instruments.

The rest of the sounds in the world have partials that do *not* follow the overtone series and thus have **inharmonic spectra**. These sounds include everyday sounds such as ocean waves, car engines, and jackhammers, but there are also a number of musical instruments that have inharmonic spectra, such as bells and some kinds of percussion.

It is worth noting that even sounds whose spectra are essentially harmonic have partials that deviate from the precise ratios. The deviations in the piano spectrum that result in "stretched" octaves are perhaps the most famous example of this. In general real pipes, strings, and reeds have subtle physical characteristics that cause the resultant spectrum to deviate slightly from the pure overtone series. This deviation is sometimes termed **inharmonicity**. Nevertheless, they are still heard as being largely harmonic and belong in a different category from distinctly inharmonic sounds.

Figure 3.6 shows the spectrum of a bell sound. Notice that, while there are distinct partials, they do not form an overtone series of *f*, 2*f*, 3*f*, 4*f*, 5*f*, and so on. As a result, this spectrum is deemed inharmonic.

Noise Spectra

Noise does not have a harmonic spectrum, nor does it have distinct partials. Instead, the spectra of various kinds of noise are better conceived as a distribution of energy among bands of frequencies. **White noise**, for example, has a spectrum whose energy is distributed evenly among all the frequencies. This can be described as *equal energy in equal frequency bands*, so white noise will have the same amount of energy between 100 Hz and 200 Hz as between 200 Hz and 300 Hz, or 1,000 Hz and 1,100 Hz.

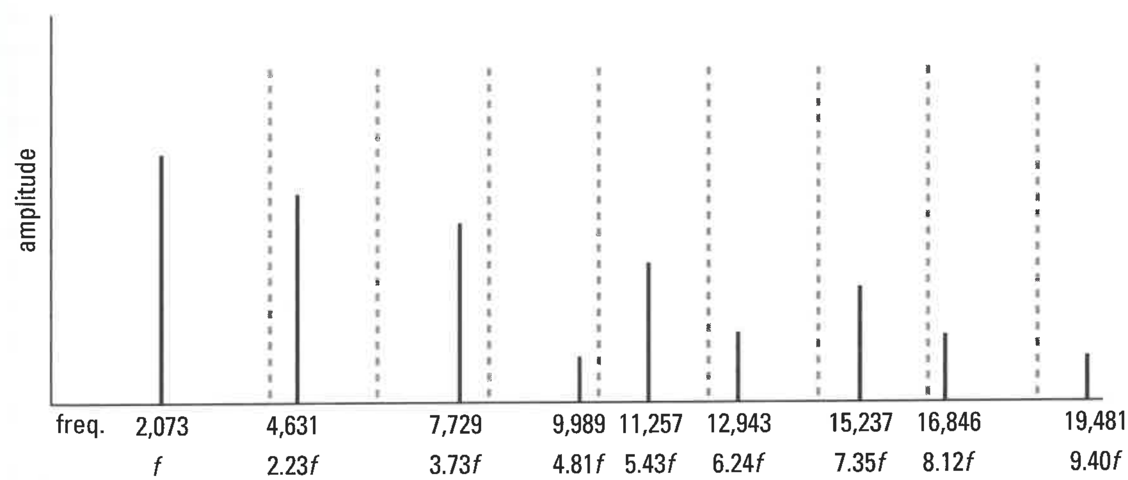

Figure 3.6 Inharmonic spectrum of a small bell. Note that the partials are not whole number multiples of the fundamental. Dashed lines indicate the positions of whole number multiples of the fundamental.

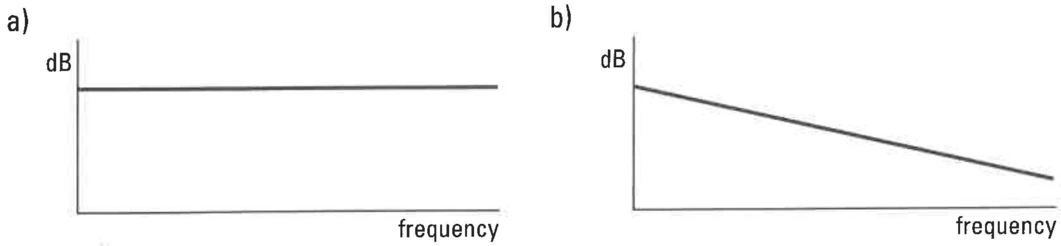

Figure 3.7 Two noise spectra shown as distributions of energy across the audible frequency range: (*a*) white noise, and (*b*) pink noise.

Pink noise, on the other hand, has *equal energy in each octave*, so there will be the same amount of energy between 100 Hz and 200 Hz as between 200 Hz and 400 Hz, or between 1,000 Hz and 2,000 Hz. Since we perceive these frequency bands as being of equal musical size (octaves), pink noise seems more evenly distributed and somewhat more pleasant to our ears. We perceive white noise as being louder at higher frequencies because the absolute size in hertz of musical intervals (thirds, fifths, octaves, etc.) gets larger as they go up in frequency. As a result, the white noise distribution contains more energy in higher octaves than in lower octaves. Figure 3.7 shows the frequency distributions of white noise (3.7a) and pink noise (3.7b).

MODIFYING TIMBRE

At first, the spectrum view may seem a bit esoteric, particularly because the waveform view of sound is so pervasive in audio recording programs. However, we actually have quite a bit of experience in manipulating timbre through the tone or **equalization** (**EQ**) controls of our home and car stereos.

Often stereos will have bass and treble controls, or bass, midrange, and treble controls. For each of these frequency bands, you can **cut** them (reduce the amplitude), leave them alone (flat), or **boost** them (increase the amplitude). Definitions of these ranges vary widely from device to device, but bass is roughly 20 to 200 Hz, midrange is roughly 200 to 5,000 Hz, and treble range is roughly 5,000 to 20,000 Hz. Many EQs have more than three bands and will often split the midrange up into two or three parts. Many manufacturers have their own definitions of these frequency bands.

Many stereos, other sound playback devices, and pieces of sound software have **graphic equalizers** that can adjust more than just two or three frequency bands. In consumer products graphic equalizers usually have presets that allow you to choose an appropriate setting for boosting and cutting the various frequency bands based on the type of music you're listening to. Figure 3.8 shows various settings for the graphic EQ in Apple's iTunes software. The "Hip-Hop" setting (3.8a) fittingly emphasizes the bass while the "Spoken Word" setting (3.8b) emphasizes the upper midrange, which improves the intelligibility of speech, particularly in a noisy setting.

More detailed discussions of EQ and timbre modification in general will be carried out later in the text in Chapters 4, 11, and 12.

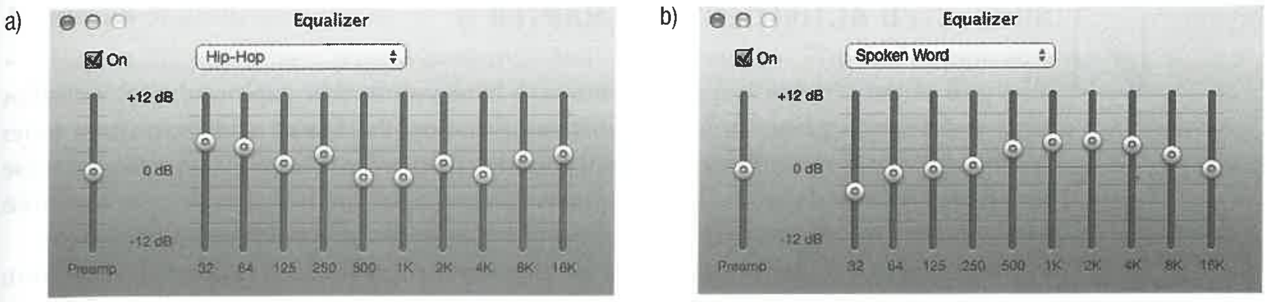

Figure 3.8 (a) The Hip-Hop setting of the graphic equalizer from Apple's iTunes; and (b) the Spoken Word setting of the graphic equalizer from Apple's iTunes. (Screenshots reprinted with permission from Apple Inc.)

REVISED SOUND PROPERTY SUMMARY

Table 3.6 is a slight revision of Table 2.6 from the previous chapter, and reflects the information presented in this chapter.

Table 3.6 **Revised perceptual and physical properties of sound**

Perceptual properties	*Physical properties*
Pitch	Fundamental frequency
Loudness	Amplitude
Timbre	Waveform *and* spectrum
Articulation	Amplitude envelope
Rhythm	Transient patterns

REVIEW OF KEY TERMS

- overtone series 36
- fundamental frequency 36
- overtones 36
- harmonics 36
- partials 36
- Pythagorean comma 40
- equal temperament 40
- spectrum view 41
- spectrogram view 42
- waterfall spectrum 42
- sine wave 42
- triangle wave 42
- square wave 42
- sawtooth wave 42
- Fourier's theorem 44
- harmonic spectra 44
- inharmonic spectra 45
- inharmonicity 45
- white noise 45
- pink noise 46
- equalization (EQ) 46
- cut 46
- boost 46
- graphic equalizers 46