which is seen to be a sum of terms that are harmonics of the base frequency $f = c/2L$.

8.8 Traveling Waves on the String

Plucking one bead (with large enough N) creates pulses that travel in both directions along the string, reflecting off the ends with a change of sign but with little or no change in shape. Another trigonometry identity you thought you would never need quickly reveals the nature of the waves propagated by the string. The identity is

$$\sin A \cos B = \frac{1}{2}\left[\sin(A - B) + \sin(A + B)\right]. \tag{8.18}$$

Now write

$$y_n(t) = \sin k_n x \cos 2\pi f_n t = \frac{1}{2}\left\{\sin\left[k_n(x - ct)\right] + \sin\left[k_n(x + ct)\right]\right\}, \tag{8.19}$$

where we have used $2\pi f_n = ck_n$. We notice that this is a function of $x + ct$ plus a function of $x - ct$. Suppose we add two solutions—that is, two different modes—say, the combination $\alpha\, y_n(t) + \beta\, y_m(t)$. This too is a function of $x + ct$ plus a function of $x - ct$. In fact, we can add any number of the modes together and this will still be true. Therefore, since we can make any shape by adding up the sinusoidal modes, the most general wave we can have traveling on the string is

$$y(x, t) = f(x - ct) + g(x + ct), \tag{8.20}$$

that is, some fixed shape traveling to the right and some other fixed shape traveling to the left. The assembly of beads has now given birth to traveling waves!

Loaded String illustrates these points with great finesse. Pick a large number of beads, turn damping off, and shape the string any way you choose. Using Display Left and Right, the left-traveling and right-traveling waves are shown in color below the picture of the total wave—in other words, the sum of the left- and right-traveling waves. Because the applet assumes the initial bead velocities are zero, there are always equal amplitude left- and right-moving waves present if you shape the string by hand. It is very instructive also to select Display Modes and move the mouse pointer over the amplitude stalks to see the shape and contributing amplitude of each mode required to generate the initial shape.

Standing versus Traveling Waves

Equation 8.19 makes clear that a sinusoidal *standing wave*, described as a sinusoidal shape function $\sin k_n x$ multiplied by a sinusoidal time function $\sin 2\pi f_n t$, is in fact a sum of two counterpropagating sinusoidal traveling waves of equal amplitude.

Fourier Again

Fourier's theorem describes the decomposition of any function into a sum of sinusoids. Initially, what concerned us was a time signal arriving, for example, at a microphone. Now, we have seen that Fourier's theorem is also describing the initial shape of the string, that is, a function of position along the string. Fourier's theorem tells us that *any* shape of the string can be represented as a linear combination of some or, if needed, all of the modes. Each mode is a standing wave, sinusoidal in shape, oscillating sinusoidally in time.

Starting with one and then a few vibrating beads, and ending with a continuous string and its collective motion, has led us through much that is of wide application to sound and vibration. Given enough beads, we get collective, long-wavelength sinusoidal periodic vibrations that are themselves harmonic oscillators, with the mass and springiness distributed over the vibrating string.

In chapter 1, we motivated how sound waves are generated by the pushing and shoving of adjacent air cells. Now, we have seen that wave behavior arises in a very explicit model of a stretched string, its parts—namely, little beads connected by tensioned filament—playing the role of the air cells introduced in connection with sound propagation. In analogy with the air cells, they pull their neighbors up and down as they are themselves pulled. The results of this tug-of-war include the harmonic modes of a string tied down at both ends, together with right-and left-propagating pulses.

Ends and Boundaries

Up to now, we have skirted an important issue: the string ends have been fixed, as is normally the case in musical instruments. Nonetheless, we should consider the possibility of free ends, arranged by connecting the end of the string to a small massless ring allowed to frictionlessly ride up and down a vertical post. Pulses behave differently when they meet a free versus a fixed end, in analogy to pressure pulses in the air columns arriving at closed and open ends of tubes. It is best to think of the string *displacement* as analogous to pressure: positive displacement (positive y) of the string corresponds to compression above the background atmospheric pressure, and negative displacement corresponds to rarefaction. Using this analogy, a string tied down at the end corresponds to an *open* tube, where the pressure is being held *fixed* at the background atmospheric pressure. Therefore, we expect a pulse arriving at a fixed end of the string to return inverted, just like the pressure pulses did at the open end of the tube. You can easily check this in *Loaded String*, by launching a pluck in the middle using many beads. This will send a right-traveling wave and a left-traveling wave heading toward the fixed ends, which will bounce off with a change of sign. Physically, traveling pulses in a string can be created, for example, by plucking or shaking one end.

Box 8.3

Experiment with *Loaded String*

Once a sound is recorded, it can be analyzed for its amplitude and frequency content. This is done as follows: Using *Audacity*, *Amadeus*, or a similar program, record a pluck of the string in the *Loaded String* applet for a few seconds. If you use *Audacity*, set the sample rate for 8000 or 16,000 in Preferences. This has the effect of emphasizing the lower frequency regions in the spectrum to be calculated later. Create a system of beads, and then pluck the system after checking the Clear button and setting Mouse = Shape String. Notice the difference in tone when different beads are plucked, and notice too that the amplitude stalks are different in each case. Starting with your recording in *Audacity*, select Analyze . . . Plot Spectrum from the top menu bar. You should get a series of peaks, as in figure 8.9, each one corresponding to the frequency of one of the normal modes of the bead system that have all been simultaneously excited with your pluck. You can check the veracity of the whole arrangement by going back to the applet and using the amplitude stalks to create the initial conditions for the bead system.

The corresponding spectrum in *Audacity* should reflect the stalks you picked. This will build confidence in the capabilities of the sound analysis program. If you create waveforms using only the lower frequency amplitudes, leaving the higher frequency modes undisturbed, the wave will be approximately periodic. This is a direct result of the excited frequencies being nearly equally spaced. If, however, you use higher frequencies in making your waveforms, the resulting shape is not nearly periodic. The reason for this is the unequal spacing of the higher frequencies, which have been excited if you choose them among the amplitudes, or if you pluck only a single bead, which does the same thing.

Figure 8.9

Experiments with a 30-bead string in *Loaded String* and *Audacity*. Two plucks were created in *Loaded String* (bottom); the sound output was captured and fed into *Audacity* (top) for a waveform and spectrum analysis. (Upper left) Case A: Waveform of the sound recording for the case of a sharp pluck of the last bead on the string. It is highly nonrepetitive. (Lower left) A 30-bead system showing the initial pluck and the resulting amplitudes, all of which are quite low but nonzero. (Upper right) Case A: spectrum of the sound recorded for the sharp pluck. Notice the presence of many peaks at high frequency, and also that the peaks are roughly equally spaced at lower frequency, but pile up at high frequency. This plot is linear on the frequency scale, but logarithmic in the sound amplitudes, so that very small amplitudes are magnified compared to larger ones. At low frequency, a 60 Hz peak is due to ambient noise in the room, a quite common source of noise coming from alternating current power. (Upper left) Case B: Waveform for the case of only five of the lowest modes excited. Notice that the waveform is nearly periodic. (Lower right) Screenshot of the 30-bead system with these initial amplitudes and displacement. (Upper right) Case B: Spectrum showing five strong peaks corresponding to the five amplitudes initially selected. The populated peaks are nearly equally spaced. The peak near 60 Hz, closest to the origin, should be discounted as electrical hum. The sound file for both cases is available in 30beads.wav.

Periodic or Not?

Equation 8.15 carries another surprise: *any* shape we choose for the string initially will be repeated periodically at the frequency $f_1 = c/2L$. This is simply because all the cosines in equation 8.15 are periodic with the same period $\tau = 1/f_1$. This can be verified in *Loaded String* using a large number of beads, and being careful not to excite the highest frequency modes. The reason for this exclusion is that these modes are increasingly inharmonic, and only if the modes have frequencies that are integer multiples of a given fundamental frequency will the shape be perfectly periodic. If, instead, 20 or 30 beads are taken and plucked rather sharply, the

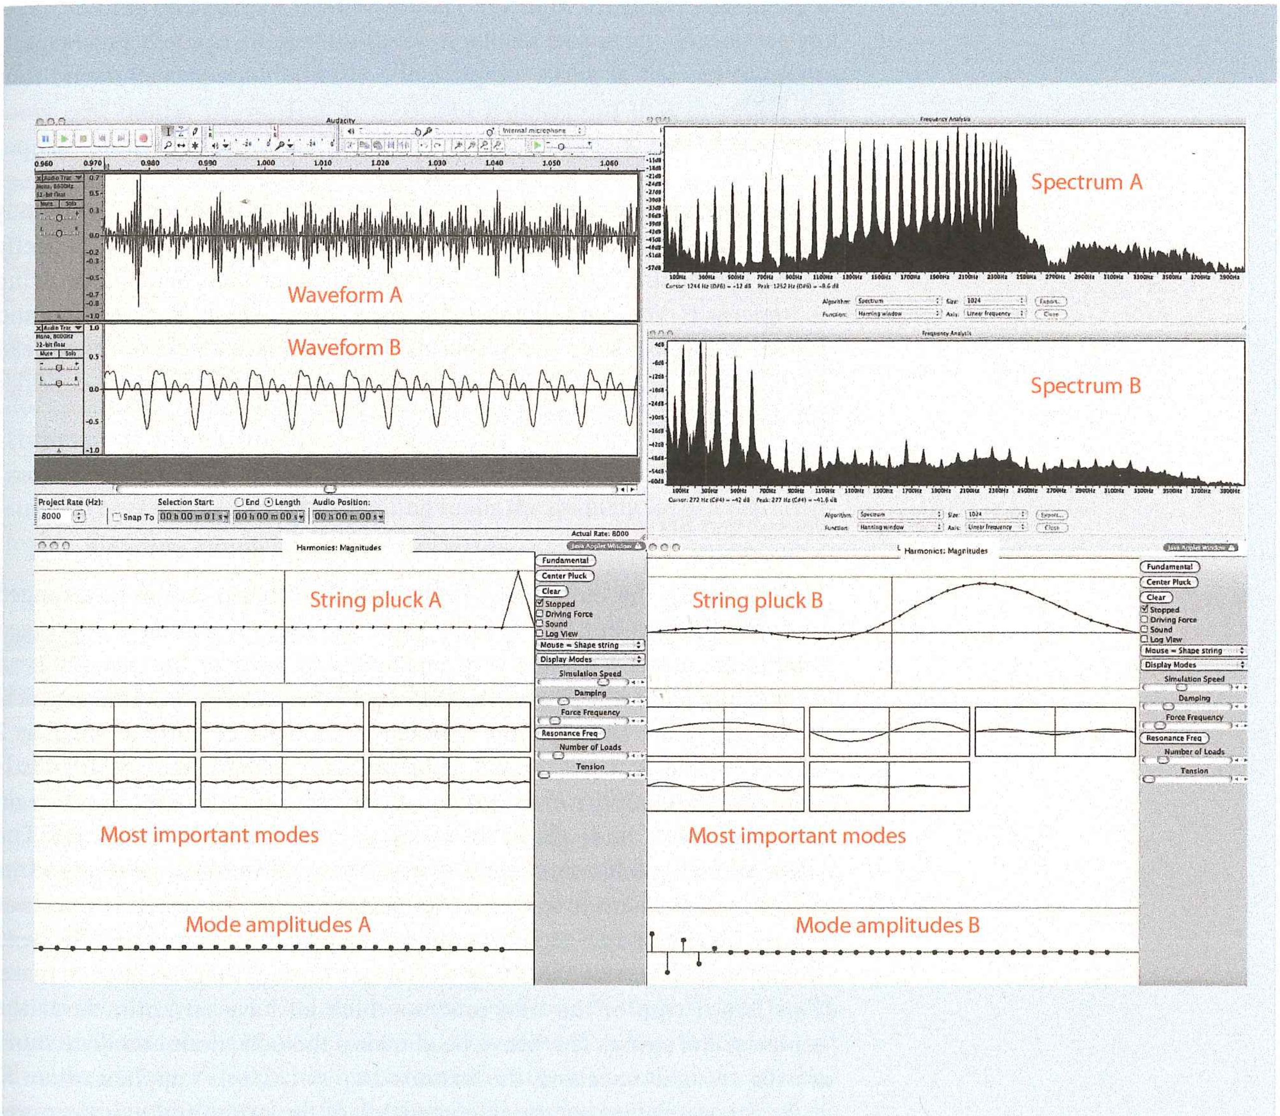

shape is only approximately regained after one period, and the subsequent attempts degrade after that.

> If the normal modes of a system have frequencies that are integer multiples of a given (fundamental) frequency, the motion will always be periodic. Sound produced by the vibrating body by any mechanism will also be periodic in time.

We have seen that ideal strings have a harmonic spectrum, and we shall see that (ideal) air columns in tubes do also. Equally spaced partials (called a *harmonic spectrum*) are the exception rather than the rule, and in fact never exactly true of real objects, although musical instruments approach this ideal fairly well.

8.9 The Imperfect String

Up to now, we have considered a uniform string of beads or a continuous string that has a uniform mass density ρ. Successive modes are perfectly evenly spaced in frequency, and their shape is sinusoidal. What if the string is imperfect? We can't exhaust this topic here, which is quite rich and subtle, but we consider two scenarios: First, we place an extra, heavy bead in the middle of an otherwise perfectly uniform string. Second, we consider the consequences of finite thickness of real strings and wires.

Weighted String

Unfortunately, the "heavy bead in the center" scenario cannot be arranged in *Loaded String*, but this is easily done in reality if you have a stringed instrument around, using a very small piece of putty or the material used for sticking pictures up on walls. The results are dramatic and, for a typical violin string, are already quite apparent for a piece of putty as small as 2 mm on a side. The exact center of a uniform string tied at both ends is a special place: all the even modes—the second, fourth, sixth, and so on partials—have a node there; the string does not move at that point. The shape of any specific mode and its frequency do not change if any of its nodes is tied down: it is already tying itself down, in effect. It does not matter if we place a weight, the extra heavy bead, at the node. The bead would not move, and the string would not "know" that the bead is there. This is not true of the odd modes, which all have an antinode at the center of the string. The heavy bead makes the odd modes oscillate more slowly—roughly speaking, the formula $2\pi f = (K/m)^{1/2}$ applies, where K is the force constant resisting deformation of the string, and m is the mass.

By adding a heavy bead, we are increasing the mass of those modes that move the bead, without increasing the resistance to bending. The frequency would be expected to drop.

The upshot of these arguments is that the even modes (2, 4, 6, . . .) are not changed by the presence of the heavy bead or putty, but the odd modes (1, 3, 5, . . .) are lowered in frequency. Perfect even spacing of the successive mode frequencies is ruined. Plucking the string results in a strange, sour note. If the extra load is heavy enough (still very small though), the string refuses to bow under the usual range of pressures on the bow; a screech is heard instead.

Ruination of the equal spacing of frequencies also spells ruination of periodicity of the shape of the string. Any pluck that activates both the even and the odd partials involves modes not all of which are oscillating at frequencies that are integer multiples of a fixed frequency. The string will not regain its shape periodically. Plucking the string in the center excites only odd modes ($n = 1, 3, 5, \ldots$), but it turns out that these are not equally shifted down by the heavy bead, so even that special pluck is not periodic.

Real Strings

A real string has thickness and resists bending of its own accord, independent of the tension. For low-frequency, long-wavelength vibrations, this does not make much difference, since the curvatures of the string are low. However, at higher frequencies and shorter wavelengths, the bending is more severe, since so many wavelengths are packed into the length of the string. This extra resistance to bending sends the frequencies of a real string higher than they would be for the ideal string.

8.10 Membranes as Stretched Bead-filament Systems

In complete analogy with the string, a two-dimensional thin stretched membrane may be constructed from an array of beads and filaments under tension. The beads are allowed to vibrate vertically, exerting forces on each other in exact analogy to the string. Once again, for N beads there are N modes, each with a different frequency. As the number of beads grows, the new membrane modes recapitulate all the old ones and add new ones, just as in the case of a string. Figure 8.10 shows a rectangular array, but the array can be made circular or indeed of any shape. None of these shapes has an equally spaced, or harmonic, spectrum. This corresponds with the fact that drumheads when struck don't give musical tones in the sense of a plucked string, although the degree of dissonance depends strongly on where the

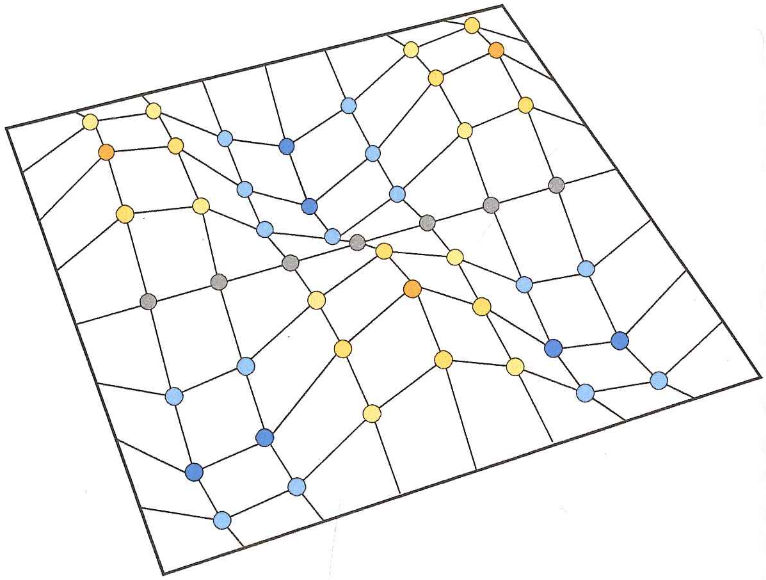

Figure 8.10

An array of beads on stretched filament vibrates in one of its modes. This mode approximates the smooth membrane that would result if a very large number of beads were used. A nodal line is seen, as indicated by the gray (uncolored) beads. There are two other nodal lines running in a direction perpendicular to this one, between rows of beads.

membrane is hit, which affects the amount each of the modes is excited. Section 15.8 has a discussion of these effects in the case of a kettle drum.

The tension T is assumed uniform in all directions along the membrane, and the density of the membrane is ρ, given in kg/m^2. Then the wave speed is $c = (T/\rho)^{1/2}$. Paul Falstad's *Oscillating Membrane* simulates a rectangular or square membrane, showing the shape and motion of various individual modes and the evolution of various kinds of plucks and hits. It also supplies the sound appropriate to each mode or combination of modes (corresponding to a strike somewhere on the membrane). Each mode has a shape that is actually a product of sinusoids, one in x and another in y. The motion of the membrane at each point is sinusoidal in time, at the frequency of the mode.

It is immediately apparent upon striking the membrane that the spectrum is not a harmonic series. The timbre (see chapter 24) depends on where the membrane is hit, and with what kind of stick or mallet. Unlike the *Loaded String* applet, *Oscillating Membrane* does not provide frequencies for a finite number of beads, but rather for the continuous membrane.

A typical spectrum is shown in figure 8.11. Hundreds of mode frequencies pile up as frequency increases. The rate of decay of the sound can be adjusted in *Oscillating Membrane* from slow to fast. Fast decay of the sound renders the spectral features broader, in agreement with the time-frequency uncertainty principle. This is seen by comparing the red and black spectra (obtained by capturing the sound from *Oscillating Membrane* and conducting a spectrum analysis for the same strike point on the membrane with different damping times selected). We consider membranes and shells (that is, thicker membranes) again in chapter 15.

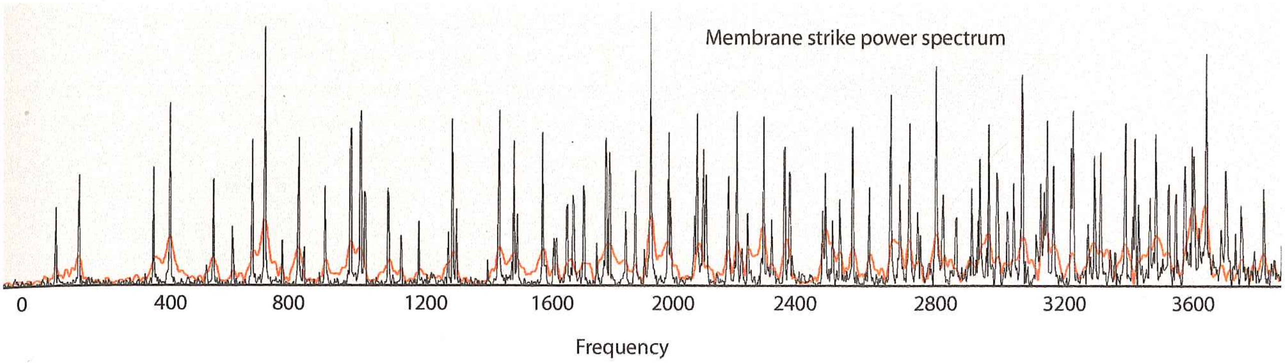

Figure 8.11

Power spectrum of the sound recorded from a hit of the membrane in *Oscillating Membrane*. Red: strong damping. Black: weak damping.

8.11 A Metal Chair

We compare the N-bead string and the membrane to hitting the seat of a metal folding chair with a knuckle. The waveform and the power spectrum of the sound are shown in figure 8.12. The spectrum crudely resembles that of a membrane. There are narrow peaks at many different frequencies—meaning that the clang of the chair (sound file chair.wav on whyyouhearwhatyouhear.com) is the sum of many pure sine tones. The chair, if sinusoidally driven at just one of these frequencies, would sound a pure sine tone even after the drive was turned off. (We will be looking into this sort of scenario in the next chapter.)

At the beginning of the discussion of beaded strings, we noted that it was worthwhile to spend time building them up because they show the way to more general systems consisting of many different modes and parts, oscillating at many different frequencies. The modes of the chair that show up as sharp peaks in the spectrum in figure 8.12 are each analogous to one of the string modes. The chair spectrum is far from harmonic, and certainly it is more complicated, but the principle of many modes at different frequencies, each its own harmonic oscillator and each excited in proportion to its amplitude at striking point, is the same.

8.12 Decomposing Complex Vibrations

One of the key concepts in this book is the notion that objects can vibrate in many ways, or *modes*, individually or in combination. Each mode has a different frequency, and each has a unique pattern of deformation that it repeats sinusoidally at the frequency of that mode. Modes can be combined; the sum of several modes is nothing more and nothing less than

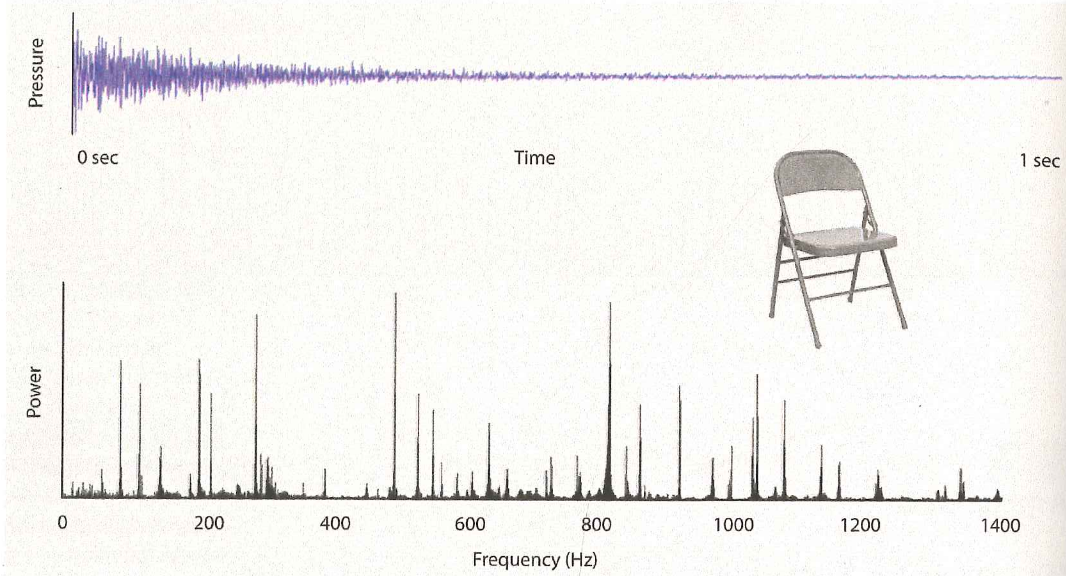

Figure 8.12
Sound pressure versus time and the power spectrum computed from its autocorrelation for a metal folding chair struck with a knuckle in the middle of the seat. The *Q* values calculated from $f/\Delta f$ vary from about 100 to over 1000. File: chair.wav on whyyouhearwhatyouhear.com.

its parts.[4] Our first hard encounter with this principle was with the string; a good review of the concepts that we encounter again here can be found by returning to the *Loaded String* applet and playing with the number of modes excited various ways and listening to the results. Each mode ultimately generates sound waves in air with greater or lesser efficiency, depending on the nature of the body that the string is attached to. The body does most of the pushing of air and is responsible for almost all the sound production. Linear superposition asserts that each mode is undisturbed by whether the others are activated or not.

Mersenne and Sauveur

The notion of simultaneous vibration of many modes of the same object, each with a shape and frequency of its own, was a difficult concept for some very smart people, partly because the notion of individual pure modes was not well understood. The motion can look complicated, making it difficult to accept that it is the sum of much simpler oscillations that are sinusoidal in time.

[4] At least, this defines so-called linear systems, which are usually a good approximation for small-amplitude vibrations.

The Minim Friar Marin Mersenne, educated by Jesuits, was an undeniable genius and strong supporter (and also a bit of a scientific competitor) of his contemporary, Galileo. (We first encountered Mersenne in chapter 2.) We have already mentioned that Mersenne made the first measurements of the speed of sound, using echoes. He was nonetheless confused about simultaneous vibrations of many different modes of the same object. He found them "impossible to imagine." (No professor, therefore, should ever get frustrated with a student facing the same difficulty!) Even though he could hear up to five partials in a single complex tone, he could not quite come to accept that a string was somehow vibrating in at least five ways at once.

This point was clearer to the younger mathematician and physicist Joseph Sauveur (1653–1716), who coined the terms *fundamental* and *harmonic* in 1701. A good namer of things, Sauveur also coined the terms *node* and *acoustics.* (Naming phenomena is a surprisingly important part of science, often earning the namer a lot of credit, whether or not they were the first to discover the phenomenon.) Sauveur explained that in pure upper partials a string vibrates in parts, which is quite true: the motion of the string at the second partial is in two halves, with a node in the middle. He was the first to associate the faint upper partials that could be heard when the string is plucked with specific modes of the string. He did not shy away from saying that the these modes could exist simultaneously on the string. This has to be one of the most fundamental discoveries in all of the history of musical sound, yet it is fairly difficult to find reference to it. Portraits of him seem to be scarce to nonexistent. He was mute until the age of five and quite hearing impaired, certainly an unexpected profile for someone who made such fundamental contributions to sound and its perception.

The nodal point at the center of the string belonging to the second mode might just as well be pinned down in that mode and indeed all the even modes, since it isn't moving anyway. Once pinned, *we see that the second mode is really the fundamental mode of a string of half the length.* The frequency is doubled compared to the fundamental of the whole string, since $f \propto 1/L$, where L is the length between the pinned ends. Sauveur explained correctly that the motion of the second partial with its two parts could be going on while the lower partial with twice the wavelength was also activated. It is possible to check this idea by watching the motion of the nodal point of the second partial with the first partial also excited: that point should execute pure single-frequency sinusoidal motion, since the second partial is not causing any displacement there.

Try this in *Loaded String.* When the program starts, only the lowest mode is excited. Drag the adjacent stalk for the second mode, and the string will distort away from its sinusoidal shape. Then, roll the mouse pointer over the first stalk, whereupon the first mode will appear in yellow. The white string position and the yellow first-mode position will intersect halfway along the string at all times, since the second mode adds nothing there to the overall motion, being a node of the second mode. The same will hold true if more even-numbered mode stalks are added—the 4th, 6th, and so on. If at any time you subtract the lower partial motion, you just recover the first partial, which is, so to speak, "riding on the back" of the first partial. To an excellent approximation, partials live independent lives, almost as if they were excited on two different strings, yet they share the same string simultaneously. This is *linear superposition*.