3.6 Spying on Conversations

A late twentieth century device for spying on conversations employs much the same principles as Lissajous's mirror on a tuning fork. For about \$45,000, you can own one yourself. Of course, it is illegal to use it except if everyone inside knows you are recording their conversations. The role of the mirror is played by a window pane in the room where the conversation is taking place. Instead of a tuning fork causing the mirror to vibrate, the window vibrates ever so slightly in response to the sound waves from nearby conversations. An infrared (invisible) laser beam is directed at the window from the outside, perhaps from an adjacent building. Detection is by projection of the reflected beam onto a sensitive photodiode, which modulates electric current according to the amount of infrared light it is receiving. These modulations are picked up and transformed electronically into audible sound. The device was reputedly a favorite trick played on the U.S. Embassy in Moscow during the Cold War. It can be defeated by taping a small vibrator to the window, overwhelming the signature of a conversation with noise.

More modern versions can beam the laser to any object in a room, like a picture in a frame. The reflected beam is analyzed by *interferometry*, which uses a beam splitter to cause interference between the returning reflected laser beam and part of the outgoing beam. Tiny shifts in the phase of the returning beam, caused by vibrations of the object it reflected from, change the interference and in turn induce measurable light fluctuations at the photodiode, which again are amplified and turned back into sound.

3.7 Fourier Decomposition

In figure 3.11, the smaller fork is vibrating faster than the larger fork by factor of exactly 3/2. Using $f_1 = 200$ and $f_2 = 300$, the displacement of the smaller tines reads

$$y(t) = \cos(2\pi \cdot 200t) + \cos(2\pi \cdot 300t), \qquad (3.10)$$

where we have assumed equal amplitudes for both vibrations. This is the function that is plotted in figure 3.11. Notice that it is perfectly periodic. Its frequency is 100 Hz, lower than either the 200 Hz or the 300 Hz component. The two waves return to their starting phase at the same time after three periods of the 300 Hz wave and two periods of the 200 Hz wave; this time is 0.01 second, corresponding to a frequency of 100 Hz. Here is the first hint of a crucial issue, one that extends to theories of music and hearing. If any *single* pitch can be associated with this combination, it is certainly 100 Hz, but there is no 100 Hz frequency present at all. One can also easily "hear out," as it is called, the 200 and 300 Hz "atomic" components as individual sinusoidal pure tones embedded in the "molecular" whole. There is much more evidence to present regarding pitch perception, and we leave this subject for a full discussion in chapter 23.

By construction, the complex vibration represented in equation 3.10 comprises two known sinusoids. Suppose, however, that only the trace $y(t)$ of the net vibration had been presented, as a kind of puzzle, and you were asked to find the frequencies and amplitudes that together make up the signal $y(t)$—that is, to find the righthand side of equation 3.10 given only the lefthand side. This kind of puzzle solving can be done systematically and is called *Fourier analysis*, or *Fourier decomposition*; here the solution is simply equation 3.10. We could make up more difficult puzzles by adding together more sinusoids, but the solution would of course exist. All such explicit sums of sinusoids have solutions, however unrecognizable they might be to someone who had not seen their construction. The question is, can *any* function be represented this way?

Joseph Fourier (figure 3.16) answered this question in the affirmative around 1800. His motivation was heat flow, not sound or music, yet Fourier more or less unwittingly revolutionized the field of acoustics. Without his theorem, we would be at a loss to explain many disparate phenomena. It required other nineteenth-century scientific leaders like Georg Ohm, Hermann von Helmoltz, August Seebeck, and Rudolph Koenig to make clear the importance of Fourier's theorem, yet they divided into two opposed camps when it came to its implications for the mechanism of hearing. Partly because the biggest names took on the losing side of the issue, much of this confusion persists today. This delicate and subtle subject is approached in chapter 23.

Figure 3.16

Portrait of Joseph Fourier (1768–1830), the French mathematician who proved that any function can be written as a sum of sinusoidal waves.

Specifically, Fourier's theorem states that any periodic function $y(t)$ of period T can be expanded:

$$\begin{aligned} y(t) = {} & a_1 \cos(2\pi f t + \phi_1) + a_2 \cos(2\pi \cdot 2 f t + \phi_2) + \cdots \\ & + a_n \cos(2\pi \cdot n f t + \phi_n) + \cdots , \end{aligned} \tag{3.11}$$

where $f = 1/T$ and the ϕ's are phases, which could be any numbers between 0 and 2π. This sum is manifestly periodic, since advancing time by an amount $T = 1/f$ advances the phase of every term by multiples of 2π. However $y(t)$ could also be periodic with a shorter period—for example, if all the odd amplitudes a_1, a_3, and so on vanish, which would make it periodic with period $T/2$. The six sinusoids shown in color in figure 3.17 give the black wave when added together, approximating a sharp-cornered "sawtooth" wave. More sines give better and better approximations.

Nonperiodic functions can also be expanded in terms of sinusoidal functions. We have already mentioned that a function such as $y(t) = \cos(2\pi f t) + \cos(2\pi \cdot \sqrt{2} f t)$ is not periodic, having an irrational $\sqrt{2}$ ratio

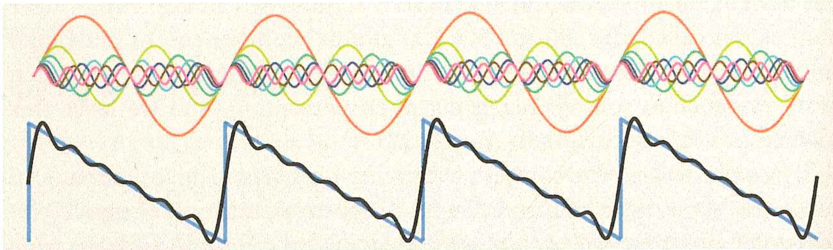

Figure 3.17

When added, the six sinusoidal terms in the Fourier series for a sawtooth wave when added together give the approximation shown in black.

of its component frequencies. Clearly, if this $y(t)$ were presented as a puzzle, the answer is just the equation in the preceding sentence, so this example hints at the fact that all nonperiodic functions also can also be written as a Fourier series. As Sir James Jeans, a famous mathematical physicist of the late nineteenth and first half of the twentieth century, and himself an author of an excellent popular book, *Science and Music*, put it, "Fourier's theorem tells us that every curve, no matter what its nature may be, or in what way it was originally obtained, can be exactly reproduced by superposing a sufficient number of simple harmonic curves [sinusoidal curves]—in brief, every curve can be built up by piling up waves."

We do not attempt to prove this here, but a little experimentation with the *Fourier* applet makes the theorem seem quite plausible. Any shape you draw with the mouse appears instantly reproduced as the sum of sinusoidal components. Fourier's theorem plays an enormous role in the theory of sound production, pitch perception, and indeed everything in this book. Fourier decomposition into many sinusoids was a matter of much debate in the past: is it purely a formal mathematical trick or something that real objects do when they generate complex signals? We take up this question again when we discuss vibrations of real objects and pitch perception.

The sinusoidal wave stands alone as the only shape that can be said to comprise a single frequency. This is essentially a tautology, yet it was by no means obvious to even very gifted natural philosophers working before Fourier's theorem became known. Any nonsinusoidal shape necessarily requires more than one term in its Fourier series.

The sinusoid is the fundamental and indivisible unit, the atom, of signal analysis and also of vibration, sound, and music. The sinusoids, each plain and colorless, together can describe any tone or indeed any sound. We shall see this principle in action many times in this book.

3.8 Power Spectra

The Fourier decomposition of a periodic signal (for example, a sound trace) into sinusoidal waves can be summarized in an important diagram called a *power spectrum*. Power is a measure of energy—in fact, the

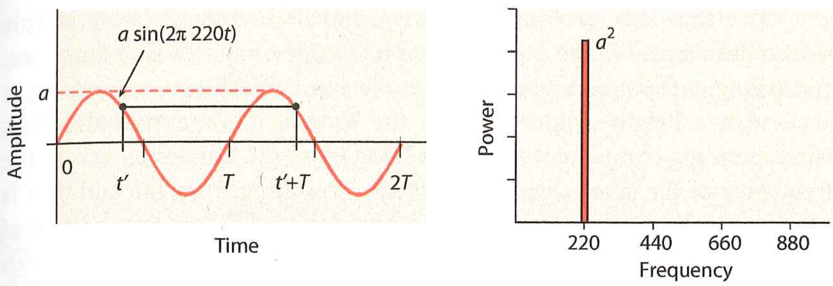

Figure 3.18
Power spectrum (right) of the sine trace (left) for the tuning fork; only one frequency at 220 Hz appears. Notice the black dots separated in time by the period T, illustrating $y(t') = y(t' + T)$, where here, $T = 1/220$ second. The power in the wave is proportional to the square of the amplitude.

amount of energy per unit time. Across many kinds of waves (light, sound, water, earthquakes), the energy in the wave increases as the square of the amplitude of the wave:

$$\text{energy} \propto \text{amplitude}^2.$$

For water, the amplitude is the height of the wave; for earthquakes, it is the pressure or the shear displacement, depending on the type of earthquake wave; for light, it is the strength of the electric field contained in the wave. The amplitude relevant to sound is the pressure variation δp above and below the background air pressure p.

The power spectrum of a simple sine wave is a graph with the frequency of the sine wave on the horizontal axis and the amplitude squared (power) in the wave on the vertical axis. With just one frequency present, at 220 Hz, the tuning fork power spectrum is very simple (figure. 3.18). Together two different sinusoids—for example,

$$\begin{aligned} y(t) &= a_{220} \sin(2\pi * 220\, t) + a_{440} \sin(2\pi * 440\, t) \\ &= a_1 \sin(2\pi * f_1\, t) + a_2 \sin(2\pi * (2 f_1)\, t), \end{aligned} \tag{3.12}$$

where $f_1 = 220$, and $a_1 = a_{220}$, $a_2 = a_{440}$ give the power spectrum shown in figure 3.19.

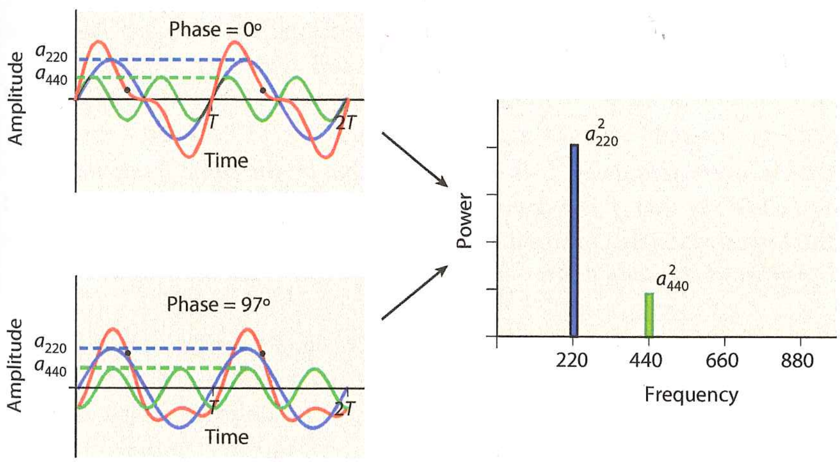

Figure 3.19
Power spectrum corresponding to the addition of two sinusoidal waves, one at 220 Hz (blue), and one at twice the frequency, 440 (green), corresponding to equation 3.12. The addition of the two waveforms is shown in red. The corresponding power spectrum is shown to the right. Each component alone would have a power spectrum corresponding to its own color. The relative phase of the two components affects the shape of the wave resulting from adding them (see the two cases on the left, for a relative phase of 0° and 97°), but the power spectrum (right) remains the same. That is, the power spectrum is insensitive to the relative phase of the partials. Notice the black dots separated in time by the period T, illustrating $y(t') = y(t' + T)$, where, here, $T = 1/220$ second.

Notice that the resulting wave $y(t)$, shown in red, is not a simple sinusoidal shape. We see explicitly that it is made up of two sine functions. Accordingly, the power spectrum for the red curve has two peaks in it, each with a height proportional to the square of the amplitude of its corresponding component sine wave, plotted on the horizontal axis at the frequency of the sine wave. The waveform resulting from the addition is now a complex signal, or musically a *complex tone*. It has a frequency of 220 Hz (it repeats every $T = 1/220$ seconds, but not sooner), even though 440 Hz is present. The 440 Hz sinusoid repeats twice as often as a 220 Hz sinusoid.

However, any periodic function of period T is also periodic at $2T$, $3T$, and so on. That is, if a signal repeats itself every second, it is certainly repeating itself every two seconds. A 440 Hz sinusoid is also periodic with the same period as a 220 Hz sinusoid. Adding any two signals that are periodic with the same period gives a result with that same period.

3.9 Periodic Functions

We need a rule or procedure for determining the period, if any, of a sum of individual partials. Discovering how to do this depends on the observation that advancing the phase in any sinusoid by an integer multiple of 2π is of no consequence: it has the same value it started with and is at the same point in its oscillation. The sum of many sinusoids, such as

$$f(t) = a_1 \sin 2\pi f_1 t + \phi_1 + a_2 \sin 2\pi f_2 t + \phi_2 + a_3 \sin 2\pi f_3 t + \phi_3 + \cdots , \tag{3.13}$$

will be periodic with period T if every sinusoid in the series has its phase advanced by an integer multiple of 2π when the time is advanced by T. Thus, we need $2\pi f_1 T = 2\pi K_1, 2\pi f_2 T = 2\pi K_2, \ldots$, where $K_1, K_2, \ldots$ are integers. But then we have

$$\frac{f_1}{f} = K_1, \ \frac{f_2}{f} = K_2, \cdots , \tag{3.14}$$

that is, *every frequency f_i is exactly divisible by the same frequency* $f = 1/T$. We say that f is a *common divisor* of $f_1, f_2, \ldots$. If we have picked the *largest* possible common divisor of f_1, f_2, and so on, we say that f is the *greatest common divisor*, or GCD. We have arrived at the rule:

> The addition of different periodic signals with frequencies $f_1, f_2, f_3, \ldots$ is periodic with the frequency given by the GCD f of all these frequencies, if it exists. If it does not exist, the resulting signal is not periodic.

The power spectrum is constructed by plotting the squares of the amplitudes a_n^2 at the corresponding frequencies f_n. Fourier's theorem states that there is enough flexibility in the amplitudes and frequencies to describe any function $y(t)$.

3.10 Aperiodic Signals and Vibrations

In the real world, there is no such thing as a perfectly periodic signal. A signal can last a long time, but ultimately it must have begun at some time, and it will come to an end in the future. In between, it can be very close to periodic. There can be other imperfections of perfect periodicity. Perhaps the signal doesn't quite repeat itself every period, but only comes close. We shall usually regard such signals as being periodic. In practice, this will often serve us quite well, with very little error in a practical sense. Other signals, such as the sound you get from a single clap of your hands, are not even close to being periodic. A telephone dial tone may be *steady*, but it is not periodic; the frequencies used are incommensurate.

A function that repeats itself only after a one-minute period may be periodic to a mathematician, but not to us in practice. Tones that have a period within our hearing range, 20 to 20,000 Hz, we shall call *audioperiodic*.

When we introduced the double tuning fork, we mentioned that the three possible frequencies could be in any relationship to each other, although the specific case we treated was the case of two frequencies with a simple integer ratio 3:2. We now return to the subject of what happens when the vibrational periods have no special relationship to each other.

The first notable aspect is that the power spectrum has no trouble with this. Given the principles of its construction, there is no reason why power could not exist at many arbitrary frequencies, not just those equally spaced and based on some fundamental frequency.

Suppose we have two frequencies, 77 Hz and 109 Hz (figure 3.20). Technically, both of these are different multiples of the same fundamental frequency—namely, 1 Hz. Other ways to say this are that the greatest

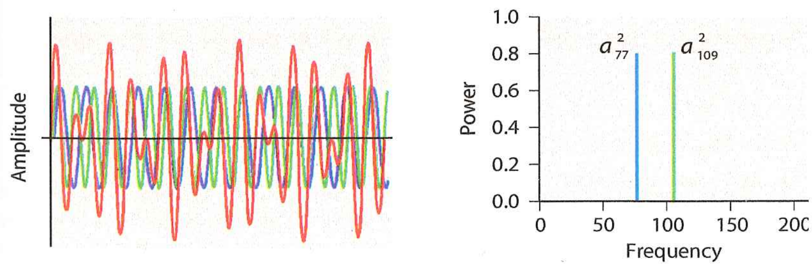

Figure 3.20
Superposition of two sinusoidal waves, one of 77 Hz frequency, and one of 109 Hz. The red trace is the sum of the green and blue traces. Notice that the result shows no sign of periodicity over this time interval, but the function is periodic, with period 1 s, or 109 oscillations of the higher-frequency component. The number 77/109 is rational. Substituting $\sqrt{2} \times 77 + 77 = 108.894\ldots + 77$ for $109 + 77$ gives a resultant that is almost indistinguishable on the time interval shown here, but is strictly aperiodic.

common divisor (GCD) of 77 and 109 is 1, or that the sum

$$s(t) = \cos(2\pi\ 77\ t) + \cos(2\pi\ 109\ t)$$

is indeed periodic, with period 1 second. On the other hand, the combination

$$\begin{aligned} s(t) &= \cos(2\pi\ 77\ t) + \cos(2\pi\ 108.894\ \ldots t) \\ &= \cos(2\pi\ 77\ t) + \cos(2\pi\ \sqrt{2}\ \times 77\ t) \end{aligned}$$

has no period at all or, in other words, the period is infinite. The two frequencies are related by the irrational number $\sqrt{2}$; the shape of this combination of only two sinusoids never exactly repeats itself! The GCD of 77 and $\sqrt{2} \times 77$ does not exist. Irrational combinations of frequencies are not available in *Fourier*, but MAX *Partials*[1] will give any combination quite accurately. The combination 77 Hz + 109 Hz is periodic, but not audio-periodic. To the ear, it is indistinguishable from 77 Hz + 108.89444 . . . Hz, which is not periodic at all. It is instructive to plot and listen to these sorts of combinations of sinusoids whose frequencies are not in any simple relation.

[1]MAX is a musical programming language that can create platform-independent applets that work with the freely downloadable MAX run-time software. Several applications were built by Jean-François Charles in MAX for this book.

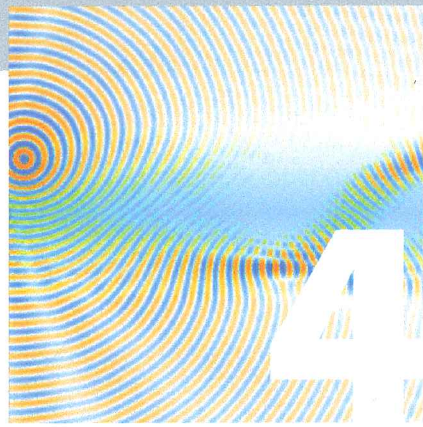

The Power of Autocorrelation

Autocorrelation is a powerful way to characterize time series data, such as the sound pressure at a microphone or the displacement of a surface. Autocorrelation is more directly related to the power spectrum than is the signal from which it is derived. Autocorrelation is a concept usually reserved for advanced treatises, but it is not difficult to grasp. The rewards for mastering it are manifold:

- *Autocorrelation provides a key summary of even very complex sounds.* Earlier, we called the sinusoid the *atom of sound*, and complex sounds, comprising many sinusoids, we called the *molecules of sound.* Molecules have collective or lumped properties such as mass and electrical charge, which are useful, even critical, aggregate effects. Autocorrelation for complex sounds plays an analogous role as a critically useful aggregate piece of information about the whole.
- *Autocorrelation leads directly to the power spectrum.* Autocorrelation is a function of time, power is a function of frequency; the two are each other's Fourier transform. This is the Wiener-Khinchin theorem, which we motivate but do not prove here.

These two properties are highlighted in this chapter. In future chapters, we shall reap an additional harvest from our understanding of autocorrelation. We preview these benefits now, without any hard evidence as yet:

- *A sudden impulse to a system supplies an autocorrelation, the* impulse response, *from which we can derive the power when driving a system sinusoidally.* By striking or hammering, velocity can be suddenly imparted to a small part of an object. The velocity of the stricken point as time progresses *is* an autocorrelation function called the *impulse response*, or IR. The power spectrum derived from Fourier analysis of the IR is simply the power that a sinusoidal drive would deliver to that point, as a function of the frequency of the drive. The IR is often quite easy to predict or understand, since it involves an initial impulse traveling away from the source, bouncing off boundaries and returning as an echo. This is a key point in chapter 9.

- *Autocorrelation plays a leading role in our hearing.* Autocorrelation of sound determines our sense of pitch, below a few thousand Hz. This is most likely a result of high-level data processing in the

Box 4.1

Autocorrelation Example: Temperature in Fairbanks

Nearly everybody has a sense of autocorrelation when it comes to the weather: the present temperature *tends* to be similar 24 hours later, less so 12 hours later; it *tends* to be similar 12 months later, less so 6 months later; it *tends* to be hot (cold) the next day if it was hot (cold) on a given day. The temperature change overnight is likely to be much less than the temperature change from summer to winter. If the autocorrelation of data is really useful, the events that "usually" happen ought to show up cleanly.

In figure 4.1, we show data that the author assembled from 5 years of temperatures taken every hour at Fairbanks, Alaska. Remember that the autocorrelation is not the temperature itself, but rather the *correlation* of the temperature at a given time with the temperature shifted to some number of hours later. Suppose the temperature data is a long column of numbers, one entry for each hour. We copy this column and place it next to the first one. The autocorrelation for a 6-hour shift is obtained by displacing the whole right column down by six entries, multiplying every entry on the left with its new righthand partner straight across, entering the result in a third column to the right of the first two, and finally adding all the numbers in the third column together, dividing by the number of entries. This gives just one number, the autocorrelation at 6 hours. What is plotted is the autocorrelation as a function of that time shift.

The following common trends are all reflected in the autocorrelation for 5 years of data, shown in figure 4.1: (1) It typically takes 1 to 3 hours for the temperature outdoors to change significantly. (Thus the autocorrelation does not change much faster than this.) (2) Daily temperature highs tend to become lows 12 hours later, and lows tend to become highs also in 12 hours, but those changes are less than the temperature swings from summer to winter. (The autocorrelation therefore dips relatively at 12 hours, 36 hours, and so on.) (3) 24, 48, 72, . . . hours after any given time is the same time on a different day, which tends to have a similar temperature (corresponding to the daily peaks in the autocorrelation). (4) Very warm or very cold weather comes in "spells" lasting typically a few days (corresponding to the 1 to 2 week decay in the autocorrelation seen in figure 4.1, lower left). (4) Summer is warmer than winter, so the envelope of the autocorrelation dips after 6 months, 18 months, and so on, but rises again at 1 year, 2 years, and so on. (5) Just because July 1, 2009, is a hotter than a normal day does not mean that July 1, 2010, will be (so that the peak in the autocorrelation reached at 1 year is not as high at it is at 0 days). The tendency to be *colder* every midnight still contributes to a *peak* in the autocorrelation at 24 hours. Similarly, cold winters tending to be cold again 12, 24, . . . months later contributes to the *peak* in the autocorrelation. The temperature in winter and 12 months later both lie at negative temperatures relative to the average; thus their product is positive.

Figure 4.1

Temperature data (top) and its autocorrelation (bottom) taken every hour of every day in Fairbanks, Alaska, for the years 2004 to 2009. The diurnal variation of the temperature can be seen in the upper left in a 20-day temperature record. On some days, expectations are brain, rather than built-in properties of the receptor systems in the ear. The sensation of pitch is the end result of our human autocorrelation algorithm. Even for periodic sounds, this can have surprising consequences, but the most important applications arise for nonperiodic sounds that still have a definite pitch. This point is taken up in detail in chapter 23.

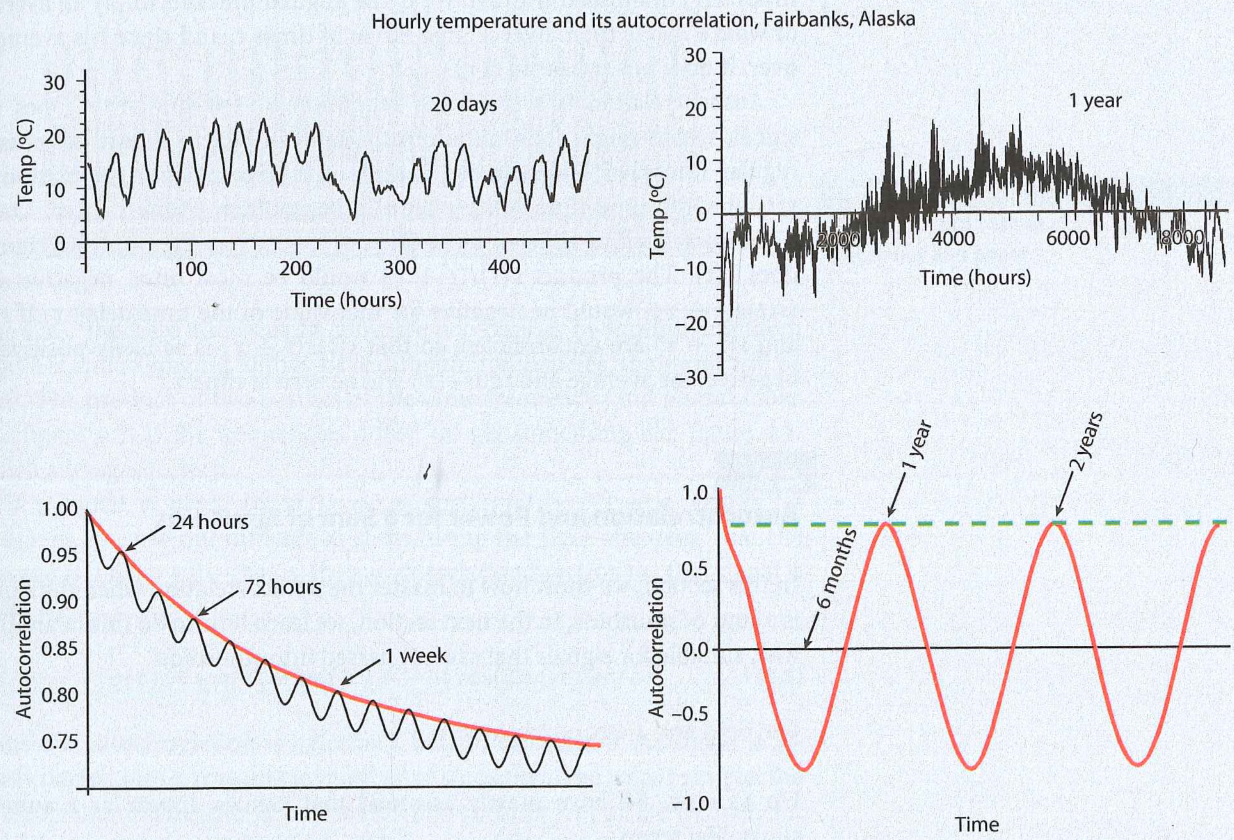

violated—for example, if a cold front arrives in the morning. The seasonal variation in temperature is seen in the one-year temperature record in the upper right. The black line in the lower left shows the computed hourly autocorrelation, averaged over the entire five-year string of data; the red line gives the envelope of the hourly autocorrelation data, which when plotted for three years (lower right) shows 6-, 18-, 30-,. . . month anticorrelations (negative correlations) corresponding to opposite seasons, and 12-, 24-, . . . month positive correlation corresponding to the same season of the year. The green dashed line shows that, while the temperature 12 months later tends to be positively correlated because it is again the same season, the correlation is not perfect, because, for example, a hot noon on July 1, 2005, does not imply a hot noon on July 1, 2006. Any given hour is always the same temperature as itself, so 0 days is always the strongest autocorrelation.

4.1 Obtaining Autocorrelation Functions

The autocorrelation $c(\tau)$ of a signal $s(t)$ is an average, defined over the history of the signal:

$$c(\tau) = \langle s(t)s(t+\tau)\rangle. \tag{4.1}$$

The quantity to be averaged is the product of the signal at two different times: $s(t)$ multiplied with $s(t+\tau)$. The angular brackets imply an average of what's inside them over a large range of times t, and since t is averaged over, it does not appear in $c(\tau)$.

Autocorrelation is a check for repetition or self-similarity. Does the signal, which might look quite unruly, tend in fact to repeat patterns at regular intervals? For example, if there is a large positive fluctuation in $s(t)$ at some time t, does there tend to be another a time τ later? Or, is it *anticorrelated* at time τ: most often $s(t+\tau)$ has the opposite sign as does $s(t)$. The product $s(t)s(t+\tau)$ would be most often negative and $\langle s(t)s(t+\tau)\rangle$ would be negative for that value of the time delay τ. If $s(t)$ and $s(t+\tau)$ are uncorrelated, so that $s(t)s(t+\tau)$ is as likely positive as negative, the average and thus $c(\tau)$ will be zero at time τ.

4.2 Autocorrelation and Power for a Sum of Sinusoids

In this section, we show how to master the autocorrelation when the signal is a sum of sinusoids. In the next section, we learn how to do this in another way, suitable for signals that are not parsed into sinusoids.

Getting the Autocorrelation

Up to now, we have mostly assumed that signals appear as a sum of sinusoidal terms:

$$\begin{aligned} s(t) &= a_1\cos(2\pi f_1 t+\phi_1)+a_2\cos(2\pi f_2 t+\phi_2)+a_3\cos(2\pi f_3 t+\phi_3)+\cdots \\ &= \sum_n a_n\cos(2\pi f_n t+\phi_n). \end{aligned} \tag{4.2}$$

For example, figure 4.2 is the sum of 10 such cosines. How can we find its autocorrelation? The key is what we call the *cosine averaging theorem*, or CAT, which states:

$$\langle\cos(2\pi f_1 t+\delta_1)\cos(2\pi f_2 t+\delta_2)\rangle = 0, \quad f_1 \neq f_2 \tag{4.3}$$

$$= \frac{1}{2}\cos(\delta_2-\delta_1), \quad f_1 = f_2. \tag{4.4}$$

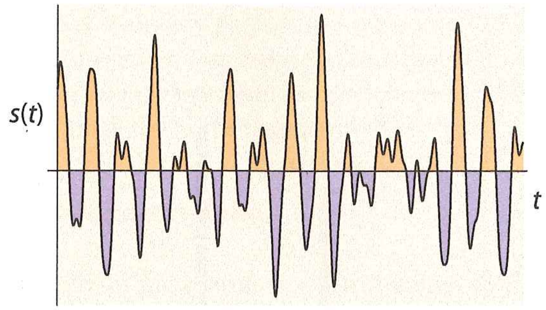

Figure 4.2
A sum of 10 sinusoids with different arbitrary frequencies and phases.

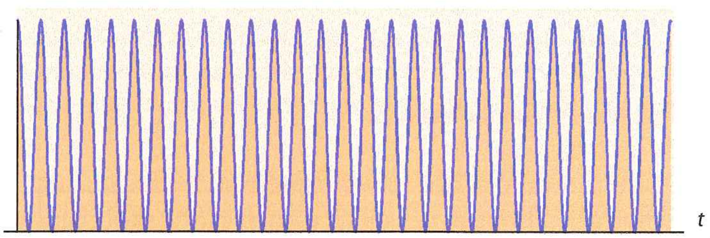

Figure 4.3
Product of two cosines of the same frequency and phase.

The CAT theorem allows us to compare two cosines by multiplying them together and averaging. If they have different frequencies, the result will be zero. The product of two cosines of the same frequency (and phase) looks like figure 4.3. If the frequencies differ, we get something like figure 4.4. which averages to zero.

If a signal is given by $s(t) = a_1 \cos(2\pi f_1 t + \delta_1) + a_2 \cos(2\pi f_2 t + \delta_2) + \ldots$, it is not difficult to show, using the CAT theorem, that the average $\langle s(t)s(t+\tau)\rangle$—that is, the autocorrelation function for that signal is given by

$$c(\tau) = \frac{1}{2}a_1^2 \cos(2\pi f_1 \tau) + \frac{1}{2}a_2^2 \cos(2\pi f_2 \tau) + \cdots . \tag{4.5}$$

Thus the autocorrelation is extremely easy to calculate if the signal is already parsed into a reasonable number of sinusoids. The rule is simple: the autocorrelation is a sum of cosine functions $\cos(2\pi f_i \tau)$ at the frequencies f_i of the partials, multiplied by half the squared amplitudes for each partial and added together.

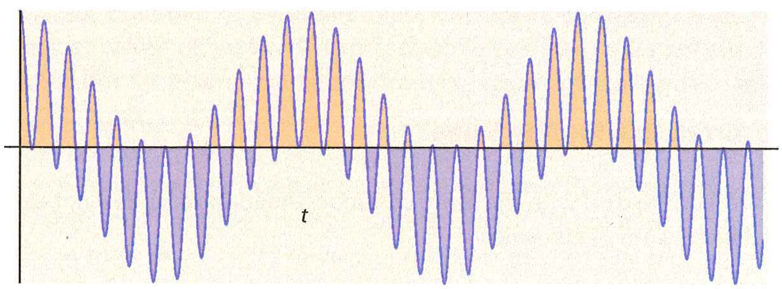

Figure 4.4
Product of two cosines differing in frequency.

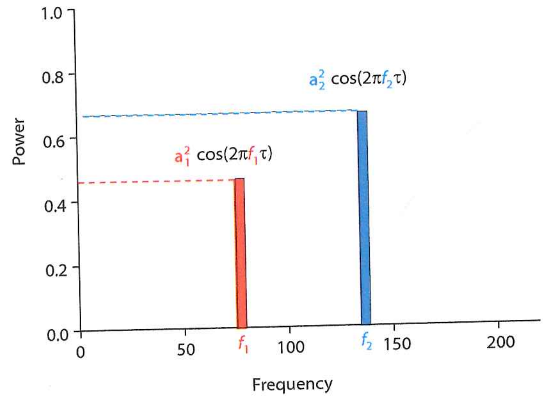

Figure 4.5
Each term in the autocorrelation, expressed as a sum of cosines, corresponds to the height and frequency of a corresponding line in the power spectrum.

Computing the Power Spectrum

We have just shown that the power spectrum corresponding to an autocorrelation given in terms of explicit cosines can be easily constructed: we simply read the amplitudes and frequencies of the cosines (figure 4.5). Vertical lines of height proportional to a_n^2 are drawn at frequencies f_n. Reversing the process, the frequencies and heights of the bars can be read from a power spectrum to yield the autocorrelation function. This completes the processing of a signal:

$$\text{signal} \rightarrow \text{autocorrelation} \leftrightarrow \text{power spectrum}. \tag{4.6}$$

One can go back and forth between the autocorrelation and the power spectrum, but one cannot get back unambiguously to the original signal because of the loss of phase information: for example, the amplitudes a_n of the original signal are squared, so their sign is lost. The original signal cannot be unambiguously constructed, but its power spectrum is unaffected by the sign of the amplitudes a_n or the phases ϕ of the sinusoids.

Autocorrelation functions are always sums of cosines with positive coefficients. They are maximal at time $t = 0$, regaining their maximum at one period, if the original signal is strictlly periodic. Most functions are not autocorrelation functions—for example, the function $b(t) = 2\cos(2\pi\, 100\tau) - \cos(2\pi\, 150\tau)$ is not an autocorrelation.

4.3 Autocorrelation for Any Signal

Now we learn to deal with the more common situation of signals not easily parsed into a sum of sinusoids.