PART I
The MESSAGE

Chapter 1
Audio

Sound

Sound is simply the vibration of air. When we speak, our vocal cords vibrate the air coming out; the sound travels through the air until it vibrates the ear drum in someone's ear, which eventually winds up in sending nerve signals to that person's brain.

In electronics, a microphone is used to convert the air vibrations into an electrical signal. The air vibrations move a thin metal or plastic plate (called a *diaphragm*) inside the microphone, which then converts the motion into an electrical signal. This signal is amplified and somehow sent from one place to another, and then converted back into air vibrations by a loudspeaker. Both the microphone and the loudspeaker (more often just called the speaker) are called *transducers*, a term which describes any device which converts energy from one form (such as mechanical vibration of the air) into another (such as an electrical signal.)

OK, now you see how detours work. Let's detour to talk about transducers.

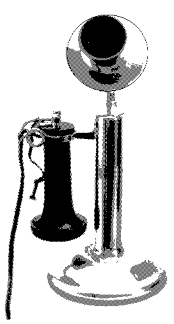

Microphones (also called mikes) come in many different types. Most common is probably the carbon mike, because it is used in every older telephone (the old-fashioned kind, not the modern little phones made in the far east. But not just the real old ones, as in the photo at the left; more modern ones too.)

Like every microphone, the carbon mike has a diaphragm, which vibrates when hit by a sound wave. The diaphragm in turn moves a thin metal sheet which presses on carbon granules in a small cup. The electrical connection is made to the metal sheet and to the bottom of the cup, as shown in Fig. 1-1.

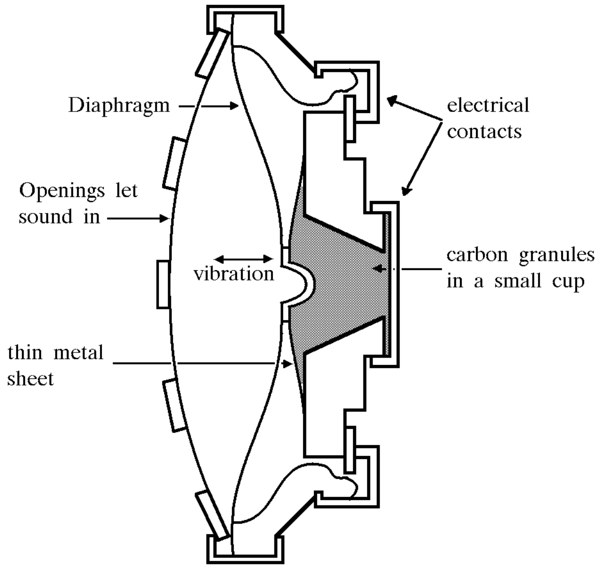

Fig. 1-1. Telephone-type carbon microphone

The carbon granules are small particles of carbon, which act as a resistance between the cup and the diaphragm. The value of this resistance depends on how closely the granules touch. As the diaphragm vibrates, it alternately squeezes the granules to increase the pressure between them, or releases the pressure. This changes the resistance of the mike in step with the vibrations of the air. The mike is in series with a battery (located in the telephone company's central office, not inside the phone), and thus varies the current in step with the sound. The current variations are converted to voltage variations when the current passes through a resistor or a transformer.

A carbon mike has very bad sound quality, but it has one advantage — the voltage variations in the series circuit can be quite large, which means that the electrical signal from the mike can travel long distances without any amplification. This was obviously necessary back in the early days of the telephone, before there were any amplifiers.

Carbon mikes were used in the early days of radio too, but they were soon replaced by better-sounding mikes.

One inexpensive mike still commonly available is a crystal or ceramic mike. This kind of mike uses a crystal (usually quartz) or ceramic material which is *piezo-electric*. This kind of material is a natural transducer — if you connect a set of terminals to a small block of the crystal or ceramic material, you get a voltage when you squeeze or twist the block; alternatively, the block twists or changes shape if you put an external voltage across it. So it naturally changes mechanical movement into an electrical signal, or vice versa.

Piezo-electric materials have many uses. For example, Fig. 1-2 shows the idea behind a crystal phonograph cartridge. The needle is attached to a small block of crystal. As the needle rides in the record groove, it twists the crystal, and the two wires attached to the other end generate a voltage proportional to the needle

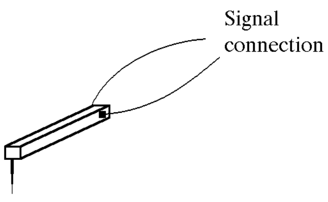

Fig. 1-2. Simplified crystal phono cartridge

movement. When sufficiently ruggedized, it also works backward. For example, some 50 years ago, Astatic made a recording head which used a crystal to make phonograph records. When fed with 50 or 100 volts of audio, the crystal would vibrate a sharp needle and cut the groove for a record.

Although crystal cartridges are no longer popular, piezo-electric materials are still often used. One modern application is for lighting a gas flame. When you push a button, a small weight hits a piezo-electric block, which in turn generates several thousand volts. This causes a spark which lights the flame.

In the crystal or ceramic mike, the diaphragm is coupled to the piezo-electric material so that the sound vibrations move the block; this in turn generates an electric voltage which is proportional to the sound signal. The mike can generate a volt or so of audio, though only at a small current.

Crystal and ceramic mikes are somewhat fragile, but back in the days before inexpensive IC amplifiers were available, they were popular because their output didn't need much amplification to be useful. But today's integrated circuits provide lots of cheap gain, so the most popular mike today is the dynamic microphone.

Dynamic mikes work on the same principle as electric generators in power plants or even in your car: when a coil of wire is placed in a changing magnetic field, the coil generates a voltage. This can be done in two ways — either keep the coil stationary and move a nearby magnet, or else keep the magnet stationary and move the coil.

Since the coil is usually lighter and smaller, it's more common to move the coil. In the dynamic mike, the diaphragm is attached to the coil, and the magnet is stationary. When the diaphragm vibrates, the coil moves in relation to the magnet, and thus produces a small voltage.

Dynamic mikes and dynamic earphones have much in common, since the earphone can work as a mike, and the mike can work as an earphone (except that it usually isn't rugged enough to produce much sound.) For instance, Fig. 1-3 shows the earphone from a typical telephone. In this case, an electric current through the coil (which is called the *voice coil*) produces motion. The voice coil is attached to a diaphragm which then vibrates and produces sound. Modern speakers have a

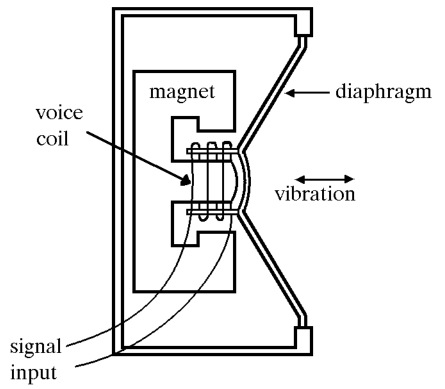

Fig. 1-3. A telephone-type earphone

similar construction, except that the diaphragm is much larger and is called the *cone*.

Dynamic mikes can produce very good sound quality, mainly because the mass of the diaphragm and attached coil is very small, and so they easily vibrate in step with the sound wave. Since the vibrations of sound occur very rapidly, a heavy diaphragm cannot move fast enough to accurately reproduce these sound waves. A carbon mike is worst, since it has to move a lot of carbon granules. A crystal or ceramic mike is better, but it still has to apply some force to the piezo-electric material and this increases the mass. A dynamic mike is still better because the diaphragm and coil can be very light. The best mike would be one where there is just a diaphragm, and nothing attached to it at all. There is a form of a dynamic mike, called a *ribbon microphone*, where the diaphragm is actually a thin strip of foil, which also acts as the coil. But these mikes are very fragile (a strong wind can ruin the foil) and also provide a tiny output, and so they are not overly common.

Instead, professional recording studios often use an excellent (though very expensive) mike called a *condenser microphone*. In the condenser mike, the diaphragm acts as one plate of a capacitor. As it moves, the capacitance changes, and an amplifier picks up that change and converts it into a voltage change.

Condenser mikes require a power supply, partially to charge up the capacitor, and partially to power an amplifier right inside the mike. The amplifier is needed because the condenser mike output voltage is tiny, and so it has to be amplified right inside the mike, before being sent out the cable. Professional recording studios usually have a 48-volt power supply inside the recording console to supply power to condenser mikes.

The cheap modern version of a condenser mike is the *electret microphone*. These mikes work on the same principle as older condenser mikes, but the diaphragm is made out of a permanently charged semiconductor ma-

Fig. 1-4. Carbon (left) and electret (right) microphones

terial which does not need a separate power supply. There is still an amplifier inside the mike, but the amplifier can be powered by a small battery. Electret mikes are not as good as professional condenser mikes (since their diaphragms are heavier), but they are cheap and small (see Fig. 1-4.) Radio Shack has some electret cartridges for $2; these cartridges are often found inside small cassette recorders. When you buy an actual electret mike, most of its cost is in the case and hardware, since the cartridge inside is often the same $2 cartridge (even cheaper in larger quantities.)

END OF DETOUR

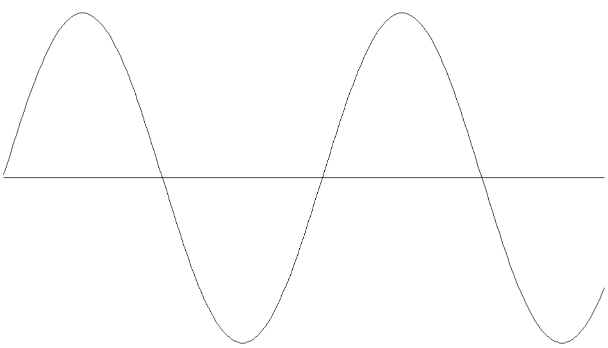

Fig. 1-5. Waveform of a whistle

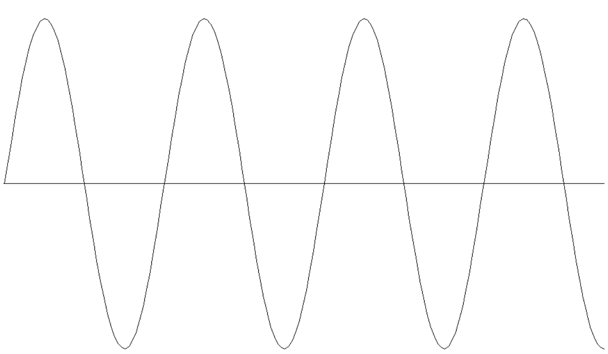

Fig. 1-6. Whistle an octave higher than Fig. 1-5

In simple sounds such as someone whistling, the resulting waveform may be a sine wave; in more complex sounds the waveform may be much more complex as well.

Let’s now look at the output of a microphone on an oscilloscope. What you see depends on the sound that the mike is picking up.

Suppose you stand in front of a mike and whistle a pure note into it (making sure to be far enough away so the mike isn’t picking up the sound of the air hitting its front.) You’d see the signal shown in Fig. 1-5. This kind of a wave is called a *sine wave* because of its relationship to the sine function from math.

On the other hand, suppose you whistle another note, but this time an octave higher. The word *octave* is a musical term, meaning eight white keys higher or lower on a piano keyboard. This time you’d see the waveform of Fig. 1-6.

Both of these waves have the same shape, but the second one goes up and down twice as often as the first. In electronic terms, its frequency is twice as high. More on this in a moment.

Now suppose you do the same thing, but this time look at the sound of some musical instrument, rather than whistling into the mike. You might see something like Fig. 1-7.

The sound in Fig. 1-7 has the same frequency as that in Fig. 1-6 (since it has the same number of cycles in a given time period,) but it looks very different. A musician might say that it has the same *pitch* (that is, it is the same musical note), but different *timbre* (a different sound quality.) Some repetitive sounds (like the pure tone of a flute) have a waveform almost like a sine wave; other repetitive sounds (like those from a violin or trumpet) have a waveform which has a basic frequency, but which looks much more distorted and "kinky" than a sine wave.

The frequency of a note determines the pitch — two different instruments playing the same note will have

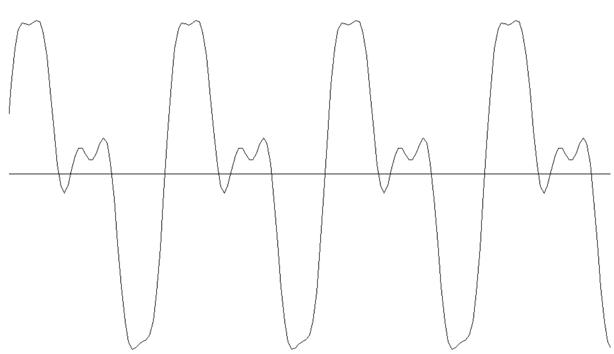

Fig. 1-7. More complex musical sound

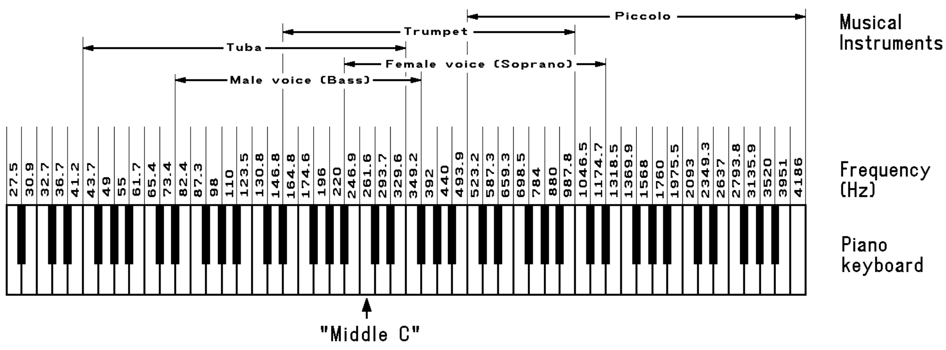

Fig. 1-8. Frequencies produced by musical instruments

the same frequency. But they will sound different because their waveshapes are different. Finally, note that the amplitude of the wave — its height — determines its volume. Quite often the amplitude of the waveform changes with time. For example, when you play a piano note, the amplitude builds up to a maximum fairly quickly when you hit the key, but then gradually decreases as the note dies away.

Note also that only repetitive sounds (like a whistle or the note of a guitar) have a definite frequency; other sounds (like the beat of a bass drum or the crack of a whip) do not.

Sound normally involves frequencies from about 20 to about 20,000 Hz, but many people cannot hear that entire range. Children often hear up to almost 20,000 Hz; as you get older, you hear fewer and fewer high frequencies. When you reach 60 or 70 years of age, you will be lucky if you can hear up to 10,000 Hz. On the other hand, many animals (such as bats or dogs) can hear much higher frequencies than humans can.

Fig. 1-8 shows the frequencies produced by each of the white keys of a piano. For example, if you look at the note labelled "Middle C" you will note that its frequency is 261.6 Hz. If you then go an octave higher — counting exactly eight white keys to the right — you get to the next C, which is at 523.2 Hz, exactly twice the frequency.

The piano has about eight octaves range, and its frequencies range from about 27 Hz on the left or bass end, up to almost 4200 Hz at the right or treble end of the keyboard. Fig. 1-8 also shows the frequencies produced by various other instruments. For example, the trumpet produces notes only in the range from about 160 Hz up to about 890 Hz.

That brings up an interesting question — if the frequencies of musical notes only range up to about 4200 Hz (and most musical instruments even have less of a range), why do hi-fi equipment manufacturers stress that their equipment goes up to 15,000 or even 20,000 Hz? The answer has to do with harmonics.

Harmonics or Overtones

So far, we've explained that

(1) The frequency of a sound determines its pitch or tone, and two instruments playing the same tone will have the same basic frequency. That basic frequency is called the *fundamental*.

(2) The amplitude of the sound determines its volume,

(3) The waveshape of the sound is what gives it its tone quality. For example, a trumpet and a violin can play the same note, but they will have totally different waveforms; that is what makes them sound different.

Consider, for example, a square wave like the one shown in Fig. 1-9. What is it that makes this wave look and sound different from a sine wave?

Suppose the above square wave has a frequency of 1000 Hz. An interesting thing happens when we send the above wave through some tuned band-pass filters; that is, filters which let through only one frequency. This is shown in Fig. 1-10.

When the 1000-Hz square wave is sent through a 1000-Hz filter, out comes a 1000 Hz sine wave! Nothing comes out of the 2000-Hz filter, while a small 3000-Hz sine wave comes out of the 3000-Hz filter. What's going on!?!

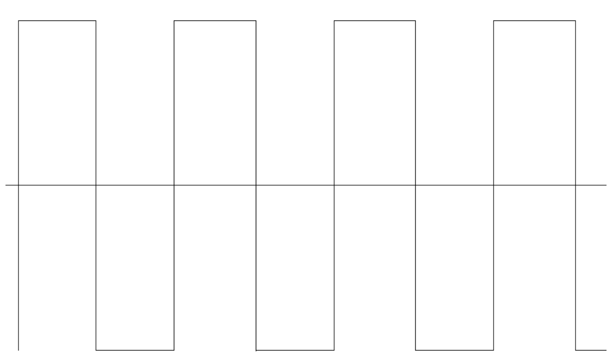

Fig. 1-9. A square wave

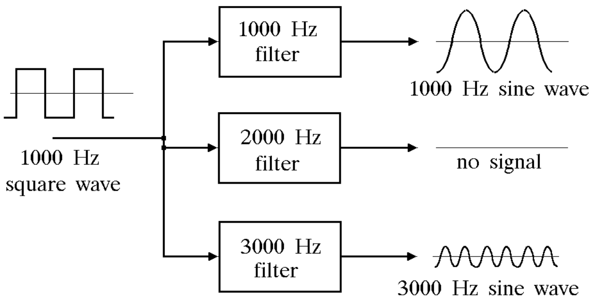

Fig. 1-10. Separating a square wave into components

Let's talk about filters for a moment.

Suppose you get several pieces of colored glass — one red, one green, one blue. When you look at a light bulb through the red glass, you see red light. Look through the green glass, and you see green light; look through the blue glass and you see blue light.

Don't think of the colored glass as a filter which is *changing* the white light into colored light. Instead, remember that white light consists of many colors, all combined together. The colored glass is simply a filter which lets one color through, while stopping all the other colors. We can demonstrate that easily be putting the green and red glasses together and trying to look through both of them. When the white light goes through the red glass, only red light comes out. There is no green light left, and so when that red light hits the green filter, it is all stopped and nothing comes out (assuming that the filters are good enough.)

If we have light from some unknown source, colored glass filters let us test that light to see what colors are in it. If we use a particular color glass, and nothing comes through it, then we know that that particular color was not present in that light. But we can also interpret this result in a different way: if what comes out of the glass filter looks different from what went in, then the original light must have had some colors in it which did not pass through the glass.

In the same way, electronic filters, like those in Fig. 1-10, let us test an electrical signal to see what its components are. When we put a 1000-Hz square wave into a 1000-Hz filter, but a sine wave comes out, this tells us that the original square wave must have some other frequencies in it which cannot get through the filter.

END OF DETOUR

Returning to Fig. 1-10, we have three filters all looking at the same original 1000-Hz square wave. We see three things:

- The 1000-Hz filter outputs a 1000-Hz sine wave, so there must have been a 1000-Hz component in the square wave.
- The 2000-Hz filter outputs nothing, so there wasn't any 2000-Hz component in the square wave.
- The 3000-Hz filter is also outputting a signal, and we see that it has three times the frequency of the square wave (there are three times as many cycles in the same amount of space) so it seems to be at 3000 Hz. Moreover, if we were to measure it carefully, we would see that it is exactly ⅓ the height of the 1000-Hz sine wave.

If we had more filters, we would see other frequencies as well. Any filter tuned to an *odd multiple* of 1000 Hz would show an output, while any other filter would show nothing coming out at all. For example, if we had a filter tuned to 5000 Hz, out would come a small 5000 Hz sine wave, and so on.

We can see from this that a 1000 Hz square wave consists of a large number of components. The 1000 Hz sine wave is called the *fundamental*, since it has the same frequency as the basic square wave. Each of the other components is an exact multiple of the fundamental; the frequency of the 3000 Hz signal is exactly 3 times the fundamental frequency, and is therefore called the *third harmonic*; the 5000 Hz component has a frequency of 5 times the fundamental, and is therefore called the 5th harmonic. In other words, a square wave consists of a fundamental sine wave signal (whose frequency is the same as that of the square wave), plus a large number (actually an infinite number) of extra frequencies, each of which is an exact odd multiple of the

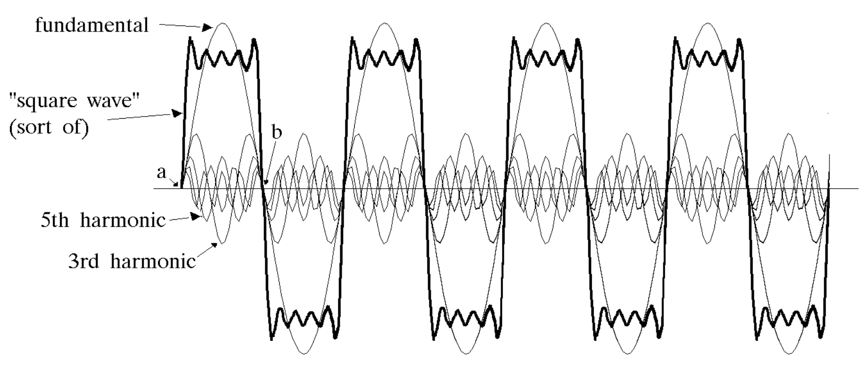

Fig. 1-11. Making a square wave out of sine waves

fundamental. We say that the square wave consists of a fundamental plus an infinite number of odd harmonics.

Fig. 1-11 shows a computer simulation of this situation. It shows a fundamental plus four harmonics (the 3rd and 5th are labelled; the 7th and 9th are not), each of which contributes a little to the output. When they are all added up, we get the squarish wave which sort of looks like a square wave, except that its sides don't go straight up and down, and the tops and bottoms are not quite completely flat. The reason it only approximates a square wave is that it only has a few odd harmonics, whereas a perfect square wave requires an infinite number of odd harmonics.

This is an important concept to understand — the square wave consists of an infinite number of components. Of course, the components must be just right — they have to have the right frequency, the right amplitude, and even the right phase. For the square wave, the rules are fairly simple:

- The harmonics must be exact odd multiples of the fundamental frequency. For example, the 93rd harmonic of a 1000-Hz fundamental would have to be *exactly* 93,000 Hz.
- Their amplitude must be just right. For example, the 3rd harmonic must be exactly ⅓ the size of the fundamental; the 5th harmonic must be exactly ⅕ the fundamental's size, and so on, all the way up. Thus the 93rd harmonic would have to be 1⁄93 the size of the fundamental. This points out that eventually the harmonics get so small that perhaps they can be omitted without making a noticeable difference in the square wave.
- Their phase has to be just right. To get the steep rise and fall of the square wave, all of the sine waves making it up have to go up together, and down together, as shown at points *a* and *b* in Fig. 1-11. If any one of them is out of step, the result will be some other wave, but not a square wave.

We can extend this concept in two important directions, both of which are critical to understanding many of the circuits in communications:

(1) Just like the square wave consists of a fundamental and harmonics, so *any repetitive waveform consists of a fundamental and harmonics,* although some of these might be missing in special cases. The sine wave is a special case since it has no harmonics at all; the square wave is another, in that only odd harmonics exist and there are no even harmonics. In a more general case, there might be both even and odd harmonics, some harmonics might be large and others small or missing, and they might have all sorts of strange phase relationships. You may have heard the term Fourier Analysis; it is simply a mathematical procedure for finding out what components make up any particular waveform.

(2) No matter how you generate that waveform, the harmonics are there even if you don't consciously put them in. For example, one way to generate a square wave is to set up a switch which rapidly switches between a positive and a negative voltage. This setup obviously doesn't put in any sine waves, yet if you look at the square wave with some tuned filters, the fundamental and harmonic sine waves are there.

Let's think about filters again. A filter lets you look for specific frequencies (or colors) in a signal (or in light). But using colored filters to look for specific colors in light is a hit-or-miss proposition if you don't know what to look for; you might need many specific color filters to identify the components in a particular light source. You can save a lot of time by looking at the light through a prism, which lets you see all the color components at the same time in the form of a rainbow or *spectrum.* The prism acts like a very large number of filters, all working together to check all the colors at the same time. You can then immediately spot whether a given color is in the light, or whether it is missing.

In communications there is an instrument which does the same job. Rather than looking at a signal through filters which look at just one frequency at a time, we can use a *spectrum analyzer* to break down a signal into its components and display them all on a scope screen as a *spectrum* (notice that we even use the same word as when we talk of colors.)

The spectrum analyzer measures the frequency components in a signal, and plots the voltage of each component against its frequency. For example, if you were to look at a pure 1000 Hz sine wave on the analyzer, you'd get a picture like Fig. 1-12.

If you imagine that 0 Hz is on the left of the screen, and each division to the right represents 1000 Hz, then the "blip" toward the left would be at 1000 Hz, and (in this case) have a height of 7 divisions. (Ideally, the blip would be just a thin line, but on the spectrum analyzer it is spread out so it looks like a very tall but thin bell.)

Leaving the analyzer at the same setting, Fig. 1-13 shows the spectrum of a 1000 Hz square wave. This time there is a big blip at 1 kHz (I added small numbers at the bottom of the figure to mark off kHz) that shows the fundamental, and progressively smaller blips at 3 kHz, 5 kHz, 7 kHz, and 9 kHz, showing some harmonics. If you examine Fig. 1-13 carefully, you will see that the blip at 3 kHz is ⅓ the height of the fundamental at 1 kHz, and so on.

The presence of harmonics has an important effect on communications. Whenever we talk about sending a

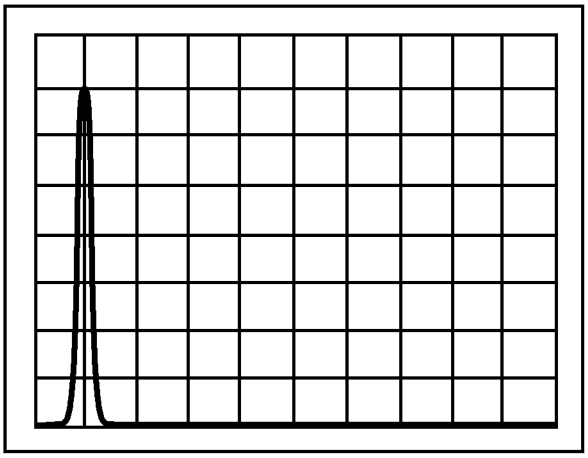

Fig. 1-12. Spectrum of a 1000-Hz sine wave

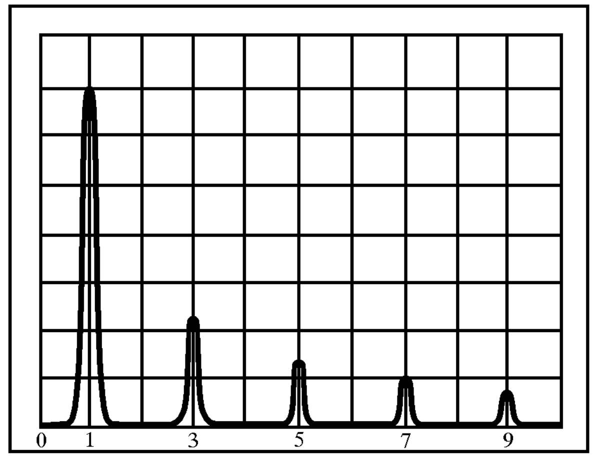

Fig. 1-13. Spectrum of a 1000-Hz square wave.

signal from one place to another, we have to make sure that all the components of that signal (or, at the very least, the *important* ones) get through as well. This brings us to the concept of bandwidth.

Bandwidth

Consider piano music. Obviously having a phonograph or tape recorder that covers the range from 27 Hz to about 4200 Hz will let through all the notes. letting us recognize the melody.

But a restricted range like that will not sound like a very good piano. To make it sound realistic, you must let through all the harmonics — or at least the ones you can hear. That is why modern hi-fi equipment typically reproduces up to 20,000 Hz (or, at least, that's what the manufacturers claim!) In fact, 15,000 or even 10,000 Hz would probably do for us older people whose hearing no longer extends to 20,000 Hz (pity!)

Let's look at the frequency ranges of sound that can be transmitted by various pieces of equipment:

- Two-way radio transceiver: about 300–2700 Hz.
- Telephone line: about 300–3400 Hz.
- AM broadcast station: about 100–10,000 Hz
- FM broadcast station: about 50–15,000 Hz
- Compact disc: about 5–20,000 Hz

Looking at these frequency ranges, we can clearly see which equipment will handle music best, and which is merely good enough for voice.

While we're on the subject of hi-fi equipment, let's discuss a few more terms.

It's not enough for a piece of hi-fi equipment to cover a wide range of frequencies; different frequencies in the range have to be treated equally. That is, an amplifier or tape deck that covers 20–20,000 Hz, but provides a lot less gain above, say, 1000 Hz than below, would sound very bassy. Ideally, hi-fi equipment should be able to handle signals of different frequencies equally well. Evenness of response is usually rated in *decibels* or db (see Appendix A for a lesson on decibels.). For example, a typical amplifier might have a rating of "20–20,000 Hz ±1 db," which means that the gain (how much it amplifies) does not vary more than plus or minus 1 decibel from some midscale value.

In addition to having a wide frequency range, the hi-fi device also should not distort the signal. That is, its output waveform should look like the input waveform (except for possibly being larger or smaller.) One way to rate distortion is as THD or *total harmonic distortion*. Remember that it's the harmonics which make one signal of a given frequency different from another signal of the same frequency. Hence if the output from an amplifier or recorder looks different from the input, its harmonics must somehow have been changed. The standard way of measuring this is to insert a pure sine wave test signal (which has no harmonics), and look to see whether there are any harmonics in the output. If so, then the signal got distorted. The THD number is a percentage that tells how much harmonic voltage got added to the pure signal. For instance, if the output from an amplifier (with a sine wave input which should have no harmonics) is 10 volts of fundamental and 2 volts of harmonics, then there would be 20% THD (a terribly high number, by the way. THD values of under one or two percent are more desirable.)

Actually, though, harmonic distortion is not nearly as bad as you think. Since music and speech normally have harmonics anyway, adding an extra percent or so of harmonics to them is not too noticeable. Amplifiers and other all-electronic hi-fi equipment tend to have low distortion, but tape recorders and mechanical components such as phonograph cartridges and speakers often have a high THD (sometimes as much as 5 to 10% for speakers.)

Much more dangerous is IM or *intermodulation distortion*, which introduces new frequencies which were not in the original at all. Even ½ or ¼% IM distortion is grating and unpleasant. Unfortunately, IM distortion is not very often listed in spec sheets for equipment; fortunately, IM distortion sort of goes hand in hand with THD, and a hi-fi device with low THD *probably* also has low IM distortion.

Finally, hi-fi equipment should have very little noise. Noise can appear in the form of a low-pitched hum (often caused by a bad power supply, bad grounding, or bad shielding of a wire) and a high-frequency hiss. Either one is bad. Hi-fi equipment specifications therefore often list the SNR or *signal-to-noise ratio*. This is the ratio between the loudest music it can handle, and the noise. For example, in a CD recording, the loudest music voltage is typically about 65,000 times higher than the noise voltage, while in a cassette recording it might only be 300 or 400 times stronger. In a telephone circuit, on the other hand, the ratio between the loudest voice signal and the noise might be as low as 10 to 1. In terms of db, the 65,000-to-one ratio is equivalent to about 95 db, the ratio of 300- or 400-to-one is about 50 db, and a ratio of 10-to-one is only about 20 db.

END OF DETOUR

So far, we've taken a short look at the nature of audio. We have seen that audio signals consist of frequencies in the range of about 20 to about 20,000 Hz, but that a narrower bandwidth (as well as some distortion and noise) will suffice if we're not too concerned with quality. For example, a typical telephone circuit can only handle the range of about 300 to about 3400 Hz.

A frequency range up to 3000 or 4000 Hz is good enough to understand speech and even to recognize the voice of the speaker, but it is certainly not hi-fi, and not good enough for enjoying music.

But there is more to the story — the bandwidth also affects how long it will take to get a certain message across.

Time

Let's say you have a friend somewhere far, far away, and you want to speak with him for a long time. But at the rate of perhaps $1 per minute for a phone call, you hesitate. So you come up with the following scheme:

Say you have 20 minutes of things to tell him. You decide to record it on tape at a low speed. Then you rewind the tape and play it back, but at double the speed so that it takes only 10 minutes to play. Can you thus send 20 minutes of speech, but pay for only a 10-minute phone call?

You can certainly do that, but your voice will sound like the Chipmunks and may not be too understandable. But suppose your friend records your voice on another tape recorder, but this time records at high speed and plays it back later at half-speed. This stretches the 10 minute tape back into 20 minutes. Will this work? (and if it does, can you speed up the tape by a factor of 10 and pay for only a 2 minute call?)

Yes... and no. What happens is that as you double the speed of your tape, every frequency on the tape is doubled too. A 1000-Hz component of your voice becomes 2000 Hz and so on. The problem is that the bandwidth of a typical telephone connection only goes from about 300 Hz to about 3400 Hz. Every component of your voice that is above 1700 Hz or so gets doubled to above 3400 Hz — and therefore doesn't make it through the phone line. In other words, your friend will only hear those components in your voice that are below 1700 Hz. (And if you tried to speed things up by a factor of 10, he would only hear those parts of your voice that are below 340 Hz.) As a result, you may have to speak a lot slower, and perhaps even repeat yourself a few times, in order to be understood.

In other words, the bandwidth of the telephone line (as well as the noise on it) limits how much information you can get across in 10 minutes. Even if you use a trick, such as speeding up the tape, it doesn't help — you wind up having to talk slower or repeat words several times.

If you used a higher-bandwidth line — such as the special lines that broadcast stations lease from the phone company for studio to transmitter links, which cover up to 10,000 or 15,000 Hz — you could easily speed up your tape by a factor of 3 or 5, and still get all your message across (though still only at normal telephone line quality.)

So there is a tradeoff between bandwidth and time. If you have a fixed amount of information to send, you can send it fast if you have a lot of bandwidth. but you have to send it slower if the bandwidth is small. That explains why, for example, a fax transmission can go through a regular telephone line, but a full-motion TV video image can't. The fax takes up to a minute to send one picture, whereas the TV has to send it in 1⁄30 of a second. In order to achieve that kind of speed, it needs more bandwidth.

Summary

Although we've rambled off and on about various aspects of audio, we've actually covered a lot of ground. We've discussed the characteristics that make up an audio wave — the frequency, waveshape, and amplitude of the signal. We discussed how harmonics affect the waveshape, and how the bandwidth of a system affects the sound quality that you can send through it.

Before we end, however, we should point out that real life is never as simple as theory, and real equipment is never as well behaved as the equipment described in textbooks. This chapter has shown several waveforms and spectra of typical sounds, but these pictures were all simplified for clarity. Fig. 1-14, on the other hand, shows (at the left) the waveform of a *real* trumpet playing a C at 523 Hz, one octave above middle C, and compares it with a theoretical sine wave (right) of the same frequency.

This is a waveform taken at just a tiny instant of time (it represents about 4 cycles out of the 523 that occur in one second, so it would last just about 4⁄253 of a second — a bit more than one hundredth of a second.)

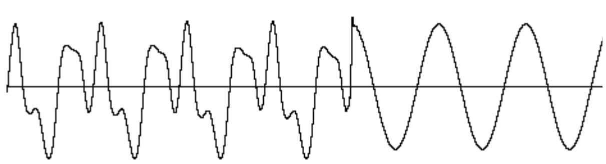

Fig. 1-14. A real trumpet waveform, and a sine wave

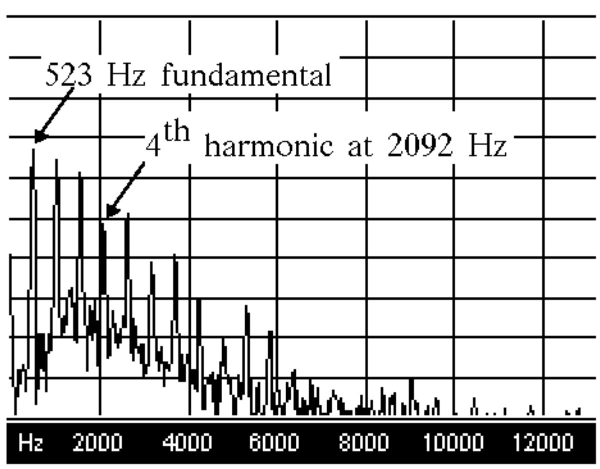

Fig. 1-15. Actual spectrum of a 523 Hz trumpet note

Yet even here, you can note that the four cycles are not quite identical. As a real instrument plays, the sound gradually changes with time, partly because a good musician will add a bit of *vibrato* to his sound — a slight wavering in the pitch of the note. As a result, the picture changes as well, and so does the spectrum.

Fig. 1-15 shows the actual spectrum of this real trumpet note. Note the big fundamental, and gradually decreasing harmonics. If you count them (we only had enough room to label the 4^{th} harmonic), you will note that some, like the 6^{th} or 9^{th}, are smaller than some of the higher ones. We also note that the picture is nowhere as neat as the theoretical ones we drew earlier in this chapter. That's just the way it is!

