Communications Simplified, Part 1

Peter A. Stark K2OAW
PO Box 209
Mt. Kisco NY 10549

If you're into ham radio, then you are already into electronic communications as well. But if you are a beginner, or if you've limited yourself to contesting or chatting on your local repeater, you've probably missed out on a lot of the more interesting technical parts of communications. This series of articles will try to bridge that gap, but in a simple (and hopefully interesting) way.

These articles are the result of a community college course I have been teaching for some years. Since I haven't been able to find just the right textbook for my course, I've decided to write my own course notes. Eventually, these notes will become a textbook, but for now you're looking at just the notes.

But don't be scared by the idea that this comes from a college course—it isn't as complex as you think. First of all, I teach a survey course. That means that it's an introduction to communications for students who haven't taken any other communications courses before. So it starts fairly simply, and doesn't assume a lot of previous knowledge. In those areas where it assumes some knowledge from a prior course, or where it seemed like a good idea to discuss some related topic, I've inserted what I call "detours"—short discussions that temporarily break away from the main subject.

Second, I teach a survey course that tries to cover a big area in just one semester. That means that there isn't time to go into tremendous depth on any one topic. Hence, we have to stay fairly simple at all times. The course has a lot of descriptions and pictures, and almost no math.

Third, this is a community college course in the second year of a technical program. It's intended for students who will become technicians, not engineers or mathematicians.

Finally, Wayne has made me promise not to show off with fancy formulas and theorems, and to add a little sugar to make the material go down easier. (I guess that means I have to crack a bad joke now and then, such as "What do you get from a Mafioso college professor? An offer you can't understand!")

But before I start off, let me add a little fine print for the benefit of the many readers who probably know a lot more about this subject than I do. Please keep in mind that these articles are intended for beginners. As such, I will often provide explanations that may seem a little (maybe *very*) simple-minded to you. While my explanations will not be *wrong*, they may be very *incomplete*. Please don't barrage me with long treatises and reprints from other texts and journals if I've omitted or simplified too much (but do let me know if I say something that isn't true).

So let's get started.

Communications

Communications is simply the moving of information from one place to another.

An English teacher thinks of communications as the process of writing or speaking words, which tell your reader (or listener) something you want him or her to know. There is even something called mass communications, which is sort of like journalism, but for TV, radio, and film. In this course, on the other hand, we think of communications more as the electronic process and equipment needed to send information over long distances.

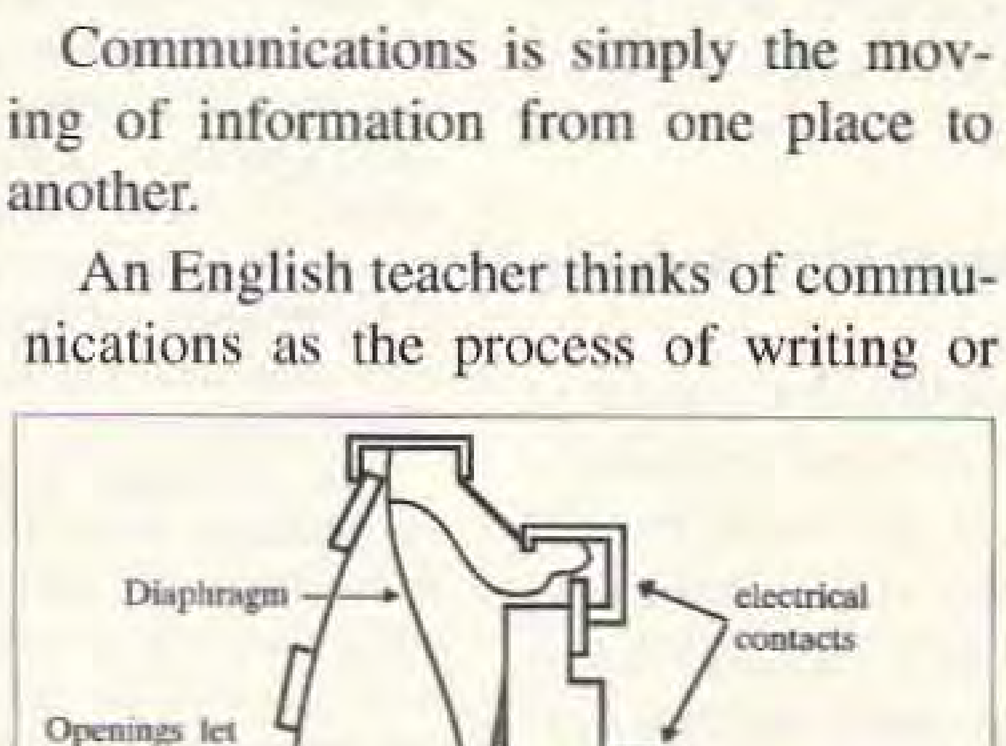

Fig. 1. Telephone-type carbon microphone.

For our purposes, we are interested in three aspects of communications:

• What is being sent? In other words, what kind of information are we sending; is it sound, pictures, or perhaps computer data?

• Through what is it being sent? That is, what is the *medium* through which it goes—wire, radio waves, sound, light, fiberoptics?

• How is it being sent? All by itself as in a telephone wire, or combined with other signals? Analog or digital?

It would be nice to be able to look at each one of these questions separately. Unfortunately, they all interact, and so our discussion will have to flit back and forth occasionally. Still, let's try to start with an orderly approach.

Let's start off with what is being sent. The information you send can be sound, video (pictures), or digital information. Let's look at each in turn.

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; a motion-to-electricity converter 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 describing any device that converts energy from one form (such as mechanical vibration of the air) into another (such as an electrical signal). (See Detour 1.)

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. 4. 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 lower on a piano keyboard. This time you'd see the waveform of Fig. 5.

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 signal created by the sound of some instrument such as a trumpet, rather than a whistle. Fig. 6 shows the resulting picture.

The sound in Fig. 6 has the same frequency as that in Fig. 5 (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 possessing 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 have the same frequency. But they 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 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. 7 shows the frequencies produced by each of the white keys of a piano. For example, if you look at the note labeled "Middle C," you will note that its frequency is 261.6 Hz. If you then go an octave to the higher—counting exactly eight notes to the right—you get to the next C, which is at 523.2 Hz, exactly twice the frequency. The piano has about an eight-octave range, and its frequencies range from about 27 Hz on the left or bass end, up to almost 4,200 Hz at the right or treble end. Fig. 7 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 range up to about only 4,200 Hz (and most musical instruments have even 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 or 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 Fig. 8. What is it that makes this wave different from a sine wave?

Suppose this square wave has a frequency of 1,000 Hz. An interesting thing happens when we send it through some tuned bandpass filters—that is, filters that let through only one frequency. This is shown in Fig. 9.

When the 1,000-Hz square wave is sent through a 1,000-Hz filter, out comes a 1,000-Hz sine wave! Nothing comes

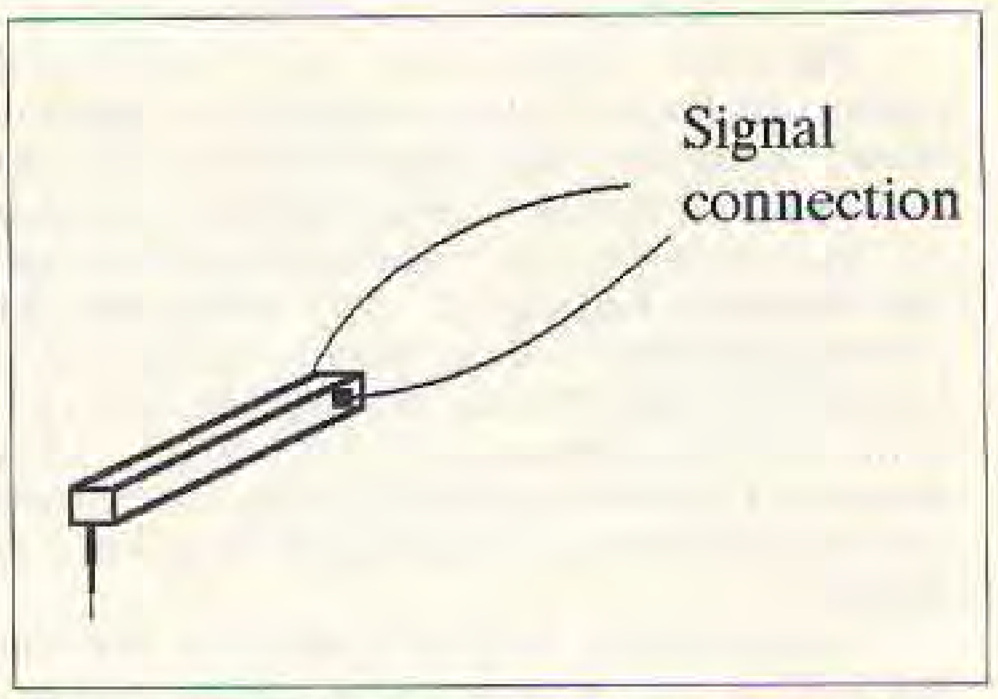

Fig. 2. Simplified crystal phono cartridge.

out of the 2,000-Hz filter, while a small 3,000-Hz sine wave comes out of the 3,000-Hz filter. What's going on!?! (See Detour 2.)

Returning to Fig. 9, we have three filters all looking at the same original 1,000-Hz square wave. We see three things:

• The 1,000-Hz filter outputs a 1,000-Hz sine wave, so there must have been a 1,000-Hz component in the square wave.

• The 2,000-Hz filter outputs nothing; thus, there was no 2,000-Hz component in the square wave.

• The 3,000-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 3,000 Hz. Moreover, if we were to measure it carefully, we would see that it is exactly one-third the height of the 1,000-Hz sine wave.

If we had more filters, we would see other frequencies as well. Any filter tuned to an odd multiple of 1,000 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 5,000 Hz, out would come a small 5,000-Hz sine wave, and so on.

We can see from this that a 1,000-Hz square wave consists of a large number of components. The 1,000-Hz sine wave is called the *fundamental,* since it has the same frequency as the original square wave. Each of the other components is an exact multiple of the fundamental; the frequency of the 3,000-Hz signal is exactly three times the fundamental frequency, and is therefore called the *third harmonic;* the 5,000-Hz component has a frequency of five times the fundamental, and is therefore called the fifth harmonic. In other words, a square wave consists of a fundamental sine wave signal (whose frequency is the

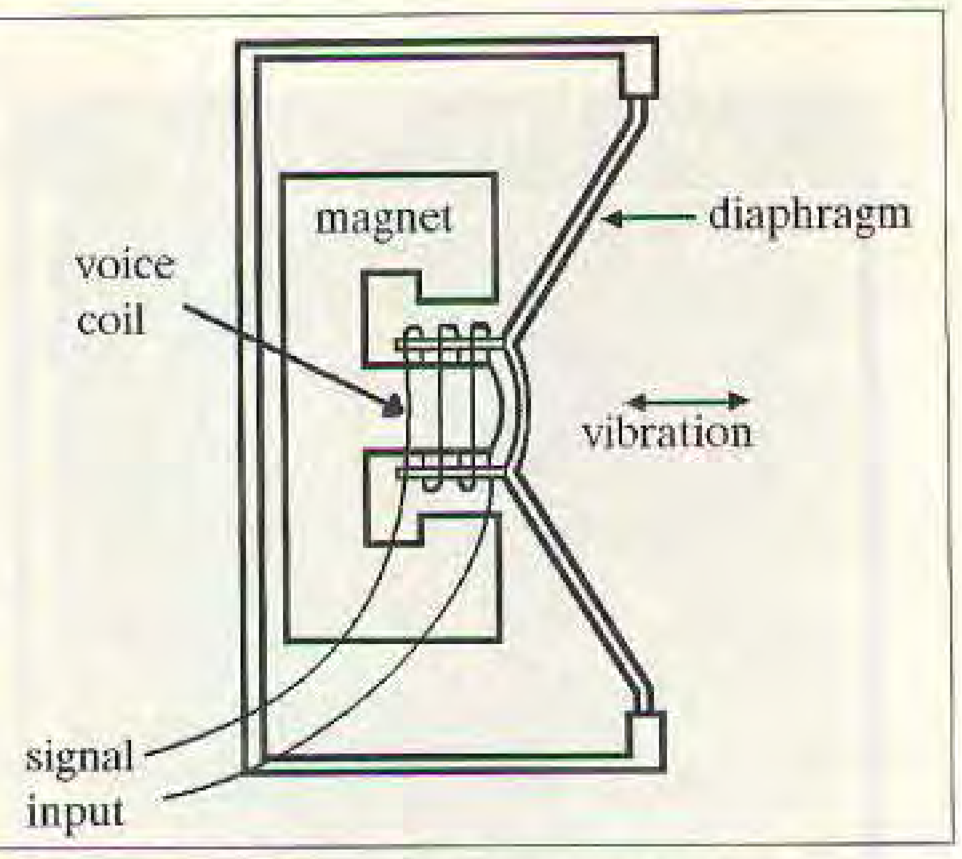

Fig. 3. A telephone-type earphone.

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

Fig. 10 shows a computer simulation of this situation. It shows a fundamental plus four harmonics (the third and fifth are labeled; the seventh and ninth are not), each of which contributes a little to the output. When they are all added up, we get the wave that looks square shaped, 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 has only a limited number of odd harmonics, whereas a perfect square wave requires an infinite number of them.

This is an important concept to understand: The square wave consists of an infinite number of components. Of course, the components have to be just right—they must 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 1,000-Hz fundamental would have to be *exactly* 93,000 Hz.

• Their amplitude must be just right. For example, the third harmonic must be exactly one-third the size of the fundamental; the fifth harmonic must be exactly one-fifth 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.

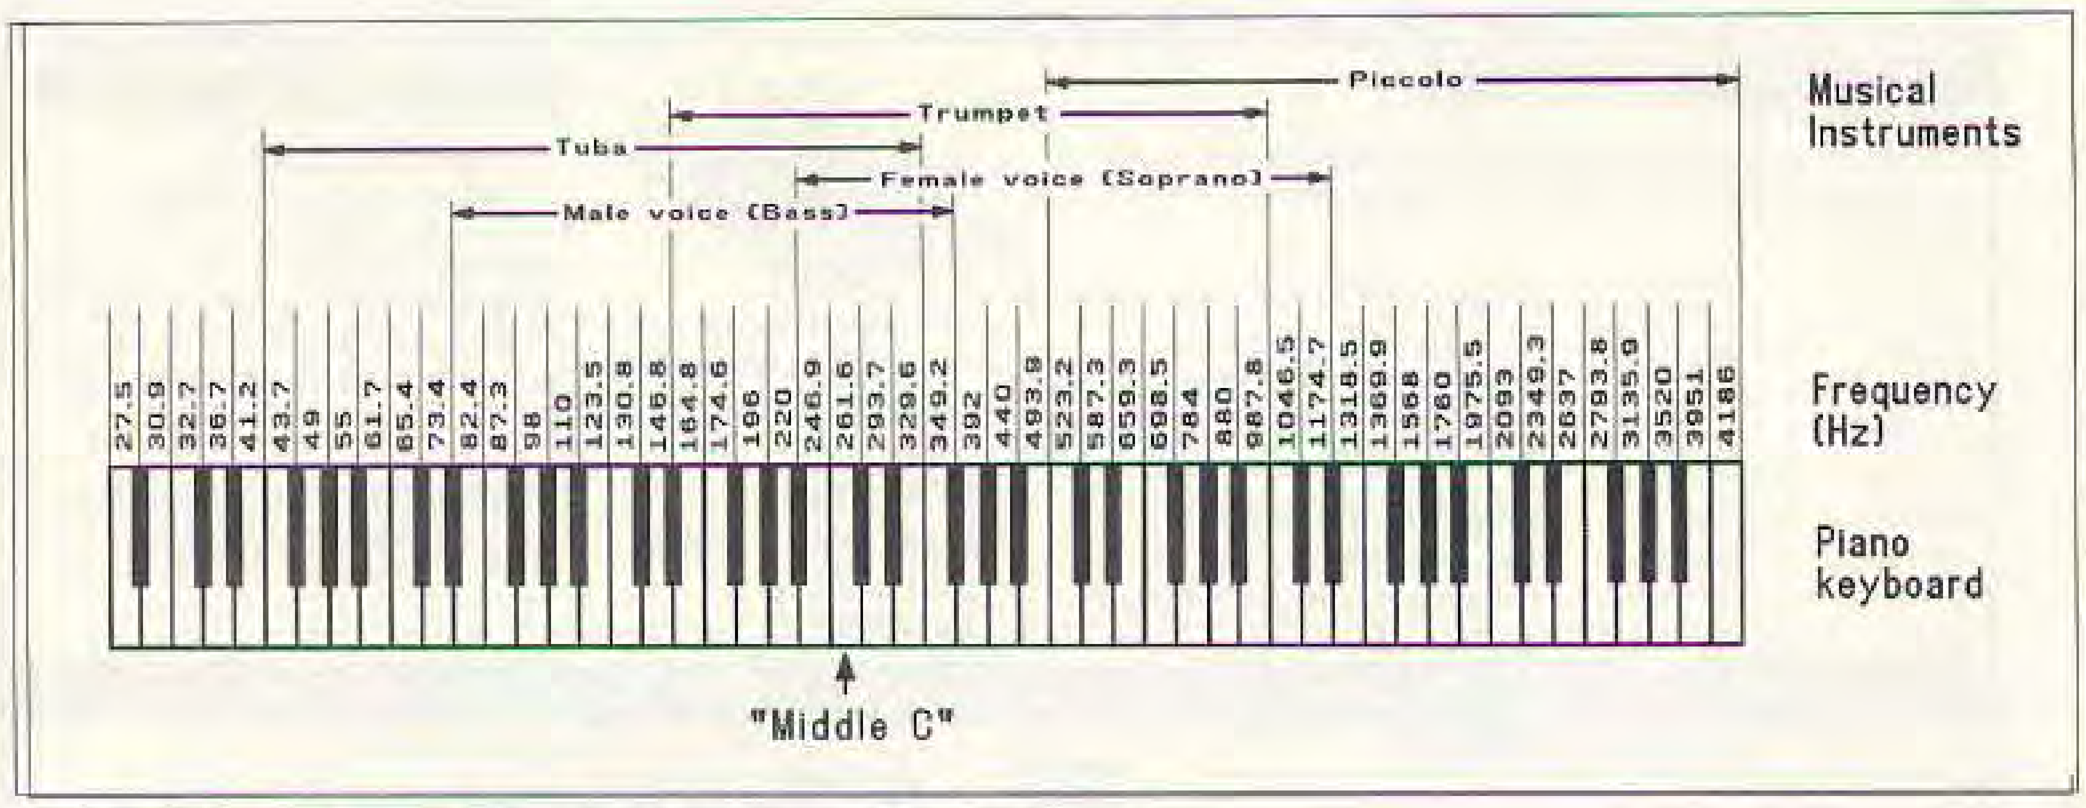

Fig. 7. Frequencies produced by musical instruments.

(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 that 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

DETOUR

DETOUR 1

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 newfangled little phones made in the Far East. But not just the real old ones, as in Photo A; 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 that 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.

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 (back in the telephone company's central office, not inside the phone), and thus varies the current in time with the sound. The current variations are converted to voltage variations when the current passes through a resistor or a transformer.

It turns out that 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 that 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. 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 movement. When sufficiently ruggedized, it also works backward. For example, some 40 years ago, Astatic created a recording head that used a crystal to make records. When fed with 50 or 100 volts of audio, the crystal would move a sharp needle and cut 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 that 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 that 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 contemporary mike 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 the earphone usually isn't rugged enough to produce much sound). For instance, Fig. 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 that then vibrates and produces sound. Modern speakers have a similar construction, except that the diaphragm is much larger and is called the *cone*.

As I mentioned above, dynamic mikes and earphones work in both directions. For example, some years ago, after placing a call from a pay phone, I discovered that someone had apparently stolen the mike from the handset. By yelling into the earpiece, I was able to complete my call anyway. But since dynamic mikes (and earphones!) produce much less output than a carbon mike, even with yelling my signal was very hard to hear.

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 worse, 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 an even greater improvement because the diaphragm and coil can be very light. The best mike would be one in which there is just a diaphragm, and nothing attached to it at all. There is a form of a dynamic mike, called a *ribbon microphone*, in which the diaphragm is actually a thin strip of foil acting 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 very 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 material that does not need a separate power supply. There is still an amplifier inside the mike, but the amplifier needs only a volt or two to run and so can be powered by a single battery. Electret mikes are not as good as professional condenser mikes (since their diaphragms are heavier), but they are cheap and small. 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).

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 that does the same job. Rather than looking at a signal through filters that 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 the similarity to the use of the word in referring to a spectrum of color).

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 1,000-Hz sine wave on the analyzer, you'd get a picture like Fig. 11.

If you imagine that 0 Hz is on the left of the screen, and each division to the right represents 1,000 Hz, then the "blip" toward the left would be at 1,000 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. 12 shows the spectrum of a 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) indicating the fundamental, and progressively smaller blips at 3 kHz, 5 kHz, 7 kHz, and 9 kHz, showing some harmonics. If you examine Fig. 12 carefully,

Continued on page 77

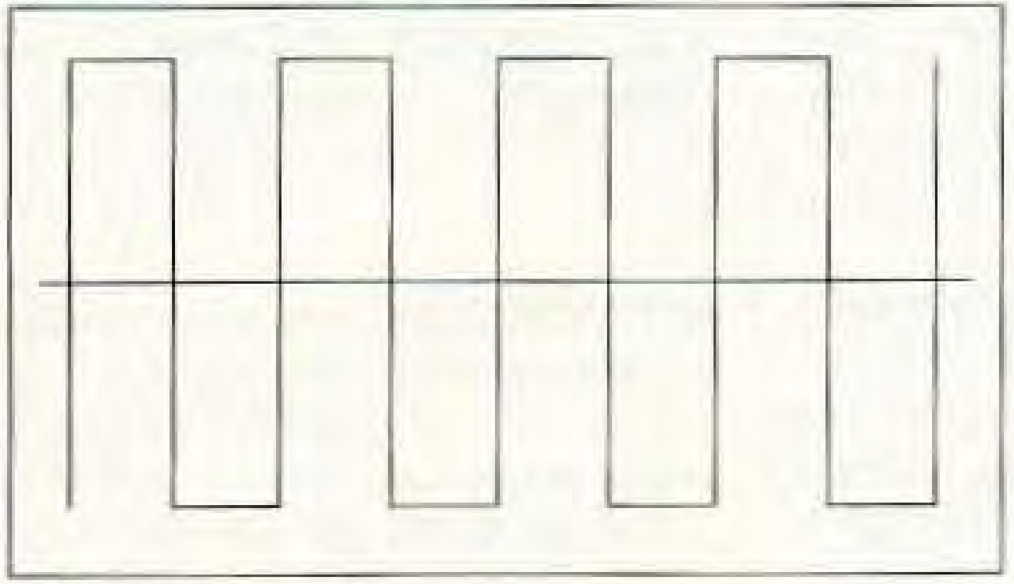

Fig. 8. A square wave.

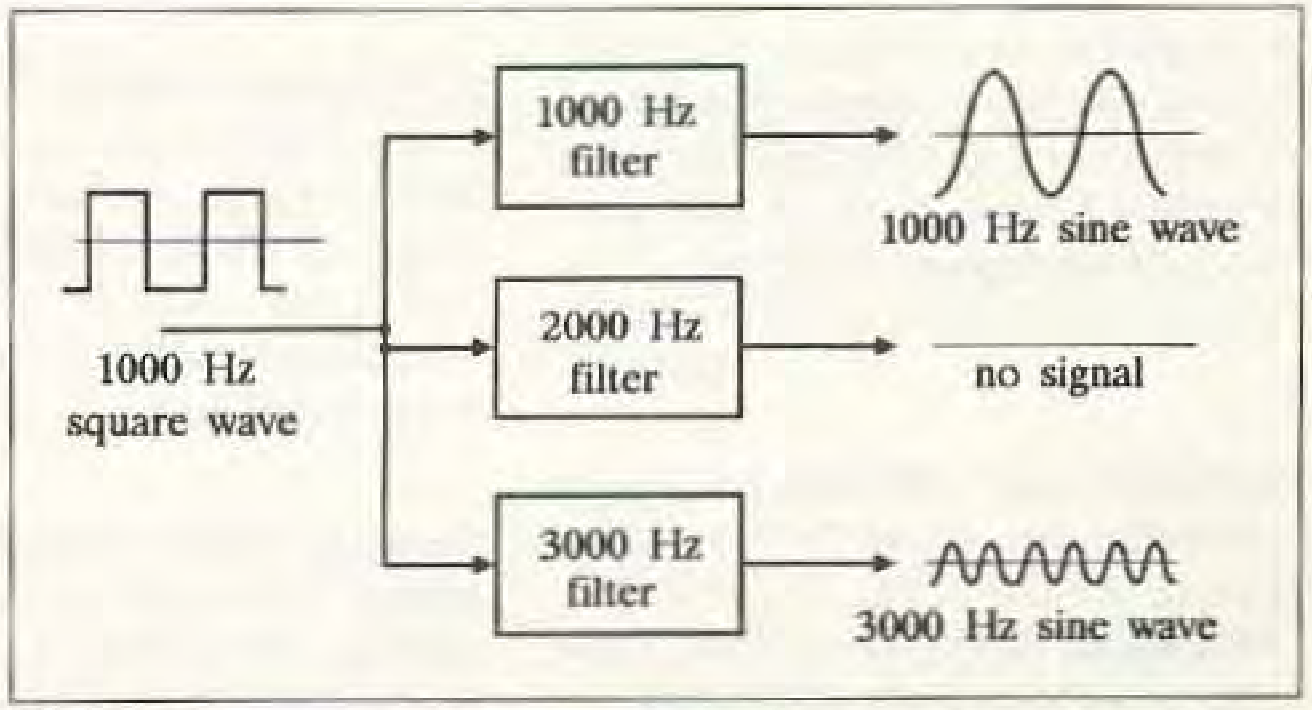

Fig. 9. Separating a square wave into components.

Continued from page 71

you will see that the blip at 3 kHz is one-third 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 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 that covers the range from 27 Hz to about 4,200 Hz will let through all the notes, allowing us to recognize the melody.

But a restricted range like that does 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

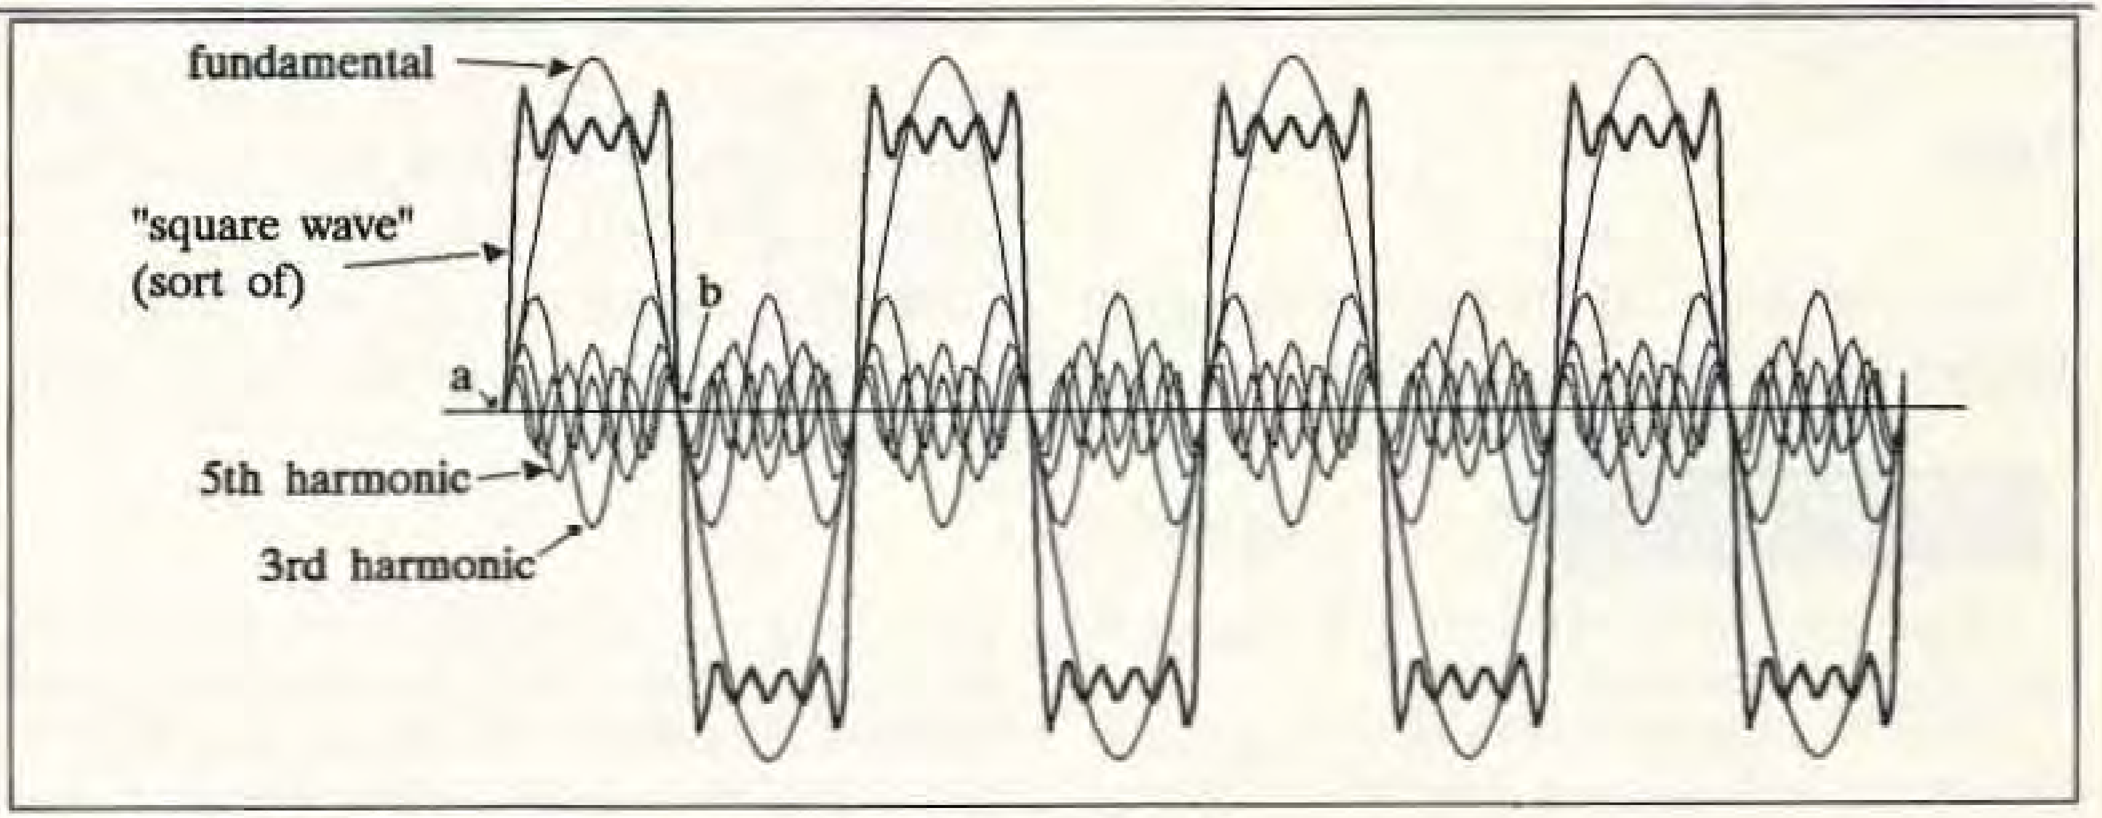

Fig. 10. Making a square wave out of a sine wave.

DETOUR

Detour 2

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 that 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 letting 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 generated by the unknown light source. But we can also interpret this result in a different way: if what comes out of the glass filter is different from what went in, then the original light must have had some colors in it that did not pass through the glass.

In the same way, electronic filters, like those in Figure 9, let us test an electrical signal to see what its components are. When we put a 1,000-Hz square wave into a 1,000-Hz filter, but a sine wave comes out, this tells us that the original square wave must have some other frequencies in it that cannot get through the filter.

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 covered by various pieces of equipment:

- SSB transceiver—about 300–2,700 Hz.
- Telephone line—about 300–3,500 Hz.
- AM broadcast station—about 100–10,000 Hz.
- FM broadcast station—about 50–15,000 Hz.
- Compact disk—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. (See Detour 3.)

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 Hz to about 20,000 Hz, but that a narrower bandwidth suffices if we're not too concerned with quality. For example, a typical telephone circuit can handle only the range of about 300 Hz to about 3,500 Hz.

A frequency range up to 3,000 Hz or 4,000 Hz is good enough to understand speech and even to recognize the voice of the speaker, but it is certainly not hi-fi. Let us now look where time comes into this.

Time

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 (that's how they do their voices!) 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 (try this with a Chipmunks record!) 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 gets doubled too. A 1,000-Hz component of your voice becomes 2,000 Hz, and so on. The problem is that every component of your voice that is above 1,750 Hz or so gets doubled to above 3,500 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 1,750 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 350 Hz.)

In other words, it's the bandwidth of the telephone line that limits how much information you can get across in 10 minutes. 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 transmit all of your message (though still only at normal telephone-line quality).

So the idea is that 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 more slowly 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.

Summary

Although the discussion has rambled off and on about various aspects of audio, we've actually covered a lot of ground. We have seen the characteristics that make up an audio wave—the frequency, wave shape, and amplitude of the signal. We have looked at how harmonics affect the wave shape, and how the bandwidth of a system affects the sound quality that you can send through it. Next time we will tackle transmission of video.

DETOUR

Detour 3

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, 1,000 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. 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. (We'll have a detour later to explain decibels.)

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 that 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 (that 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 that 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 not in the original at all. Even 1/2% or 1/4% IM distortion is grating and unpleasant. Unfortunately, IM distortion is not very often listed in spec sheets for equipment; fortunately, however, IM distortion sort of goes hand in hand with THD, and a hi-fi device with low THD *probably* also has low IM.

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 specs 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, the ratio between the loudest voice signal and the noise might be as low as 10 to 1. Since we're holding off on our discussion of decibels, let me just say at this point that the 65,000 ratio is equivalent to about 95 dB, the ratio of 300 to 400 is about 50 dB, and a ratio of 10 is only about 20 dB.