PART III
TRADITIONAL METHODS

Chapter 7
Amplitude Modulation

In Chapter 6, we discussed wireless transmission. Specifically, we mentioned that the size of the antenna required to transmit or receive a radio signal depended on the frequency being transmitted, and so the frequency had to be increased in order to cut antenna size down to some reasonable size. This is done by modulating a high-frequency carrier wave with the desired signal (audio, video, data, or whatever). We now describe how this is done.

The process of putting our signal on a carrier is called *modulation.* In the transmitter it is done by a *modulator*, and in the receiver the desired signal is removed from the carrier by a *demodulator*, also called a *detector*.

The carrier itself is a high frequency sine wave. Although the carriers for the standard AM broadcast stations are in the range of 540 to 1600 kHz, carrier frequencies can be much lower as well as much higher. The US Navy operates some transmitters with carrier frequencies about 10 kHz; on the other hand, microwave transmitters often have carriers above 10 GHz — that is more than 10,000,000,000 Hz.

Modulating a carrier involves changing it in step with the signal (voice, music, picture, or whatever, but we will talk only about audio for now) that we want to send. Since the carrier sine wave has a frequency, an amplitude, and a phase, any of these three can be changed with modulation. When we change the amplitude, we produce *amplitude modulation* or AM; changing the frequency produces *frequency modulation* or FM, and changing the phase produces *phase modulation* or PM. We will begin by looking at AM, and leave FM and PM for a later chapter.

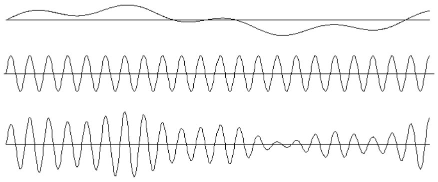

Fig. 7-1. Audio, carrier, and AM modulated wave

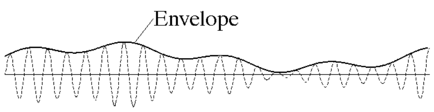

Fig. 7-2. The envelope

Amplitude Modulation (AM)

All of the radio stations on the standard AM broadcast band use amplitude modulation, a method which dates back to the very beginnings of radio.

Fig. 7-1 shows a sample of amplitude modulation. At the top, we see a typical audio signal; underneath it is the *unmodulated carrier*, a plain sine wave with a much higher frequency than the audio signal; at the bottom is the carrier with the audio modulated onto it. Note how the modulated carrier (bottom waveform) becomes bigger when the audio is positive, and becomes smaller when the audio is negative. If the audio wave is near zero volts (right in the center of the audio wave), then the modulated carrier is the same height as the unmodulated carrier.

If you take a pencil and carefully connect the tops of each cycle in the modulated carrier, you get a curve that looks just like the audio signal. This is the dark curve in Fig. 7-2, and it is called the “envelope”

The modulator in the transmitter therefore takes the audio and uses it to vary the amplitude of the carrier; the demodulator (detector) in the receiver then uses the envelope of the carrier to recover the audio, and throws the carrier away.

A simple AM receiver

Fig. 7-3 shows the diagram of a very simple AM receiver, called a *crystal radio*, that can be built in the laboratory with just a few parts. The unique thing about it is that it needs no battery or power supply, and so provides absolutely free radio reception.

But since “there is no such thing as free lunch,” there has to be a catch. The catch is that the radio really does need some power, and that power has to come from the radio station. This radio will only work if you

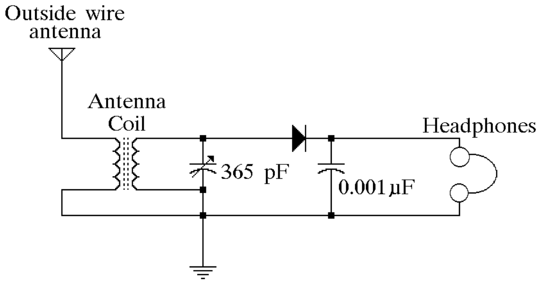

Fig. 7-3. A crystal radio receiver

are close enough to a radio station to receive a strong signal.

The antenna need not be one of the ones discussed in Chapter 6; a long length of wire, preferably strung somewhere outside the house (but away from electric power wires), will pick up very many different signals, including AM and FM radio stations, TV stations, taxi and police radios, and more. It may not be particularly efficient at any one frequency, but most AM signals are strong enough that such an antenna will pick up something. The first thing that must be done is to remove the stations we don't want, and keep only the one we do want. This is done by a tuned circuit consisting of the antenna coil and the 365 pF variable capacitor; the antenna coil acts both as a transformer as well as an inductor in the tuned circuit.

Years ago, these two components were very easy to get in almost every radio store; today they are hard to find. Suitable replacements can be made from more modern parts, but it is probably easier to buy a commercial crystal radio kit, such as the one from Radio Shack (which uses a slightly different circuit, but it works the same way.) When the right coil and capacitor combination is used, turning the capacitor from one end to the other will tune the radio through the 540 kHz to 1600 kHz range of the AM broadcast band.

Ideally, the output from the tuned circuit should contain only the modulated carrier from the one station we are tuned to; alas, that's not the way it usually works out. A single tuned circuit is usually not good enough to keep one station, but remove all the others; a normal radio needs several tuned circuits working together to accomplish that. What usually happens in the crystal radio is that we have the one desired station, plus the signals from a few adjacent stations. If the station we want is strong and the others are weak, then the radio will work well. But if the desired station is weak and the others are strong, then we might as well give up — we will not be able to hear the station we want.

The modulated carrier is now sent to the diode. Since a diode conducts in only one direction, only half of the signal gets through it. In this case, only the positive peaks of the carrier make it through the diode. When these peaks hit the 0.001 µF capacitor, they charge it up. (If you're familiar with power supplies, then you can think of the diode as the rectifier, and the capacitor as the filter which charges up to a DC value and removes the ripple.) The capacitor basically charges to the voltage of the envelope except that the voltage of the envelope keeps changing in step with the audio, and so the capacitor voltage also keeps changing. This voltage is then sent to the headphones.

Although the crystal radio circuit looks simple, actually all the components have to be just right, or it will not work well. For example, the antenna coil and capacitor must be the right values to tune to the band; the ratio of turns also has to be right to give the maximum signal. For best reception, the diode should be a germanium diode, not a silicon diode which needs more voltage to operate. The capacitor value is also somewhat important — it has to be large enough so it removes the carrier, but small enough that it doesn't remove the audio signal. (Actually, the radio will even work without this capacitor, but not quite as well.) Even the headphones must be carefully chosen — they must have a high resistance (1000 ohms or more) to prevent shorting the signal; that means that the kind of headphone usually supplied with the "walk-person"-type tape players will not work. (In their crystal radio kit, Radio Shack omits the 0.001 µF capacitor, and uses a crystal headphone which has a very high resistance.)

This kind of radio is called a crystal set because of the diode. Some 60 years ago, when crystal radios were very popular, tiny germanium or silicon diodes were not even invented yet. Instead, the crystal set used a small piece of galena crystal and a "cat's whisker". The cat's whisker was a thin, springy wire which pressed against the galena to form a diode junction. The cat's whisker was attached to a small handle with a knob, and you had to probe the crystal to find a "hot spot." Tuning and adjusting such a radio was almost an art (and a lot of fun!)

A simple AM transmitter

Now that we see how the AM receiver works, let's see how the transmitter works. Fig. 7-4 shows a simple circuit that can be wired in the laboratory for demonstration purposes. A commercial signal generator generates a carrier at some frequency in the AM broadcast band that is unused by a station.

The signal generator is fed to the base of the transistor. Since the transistor has no base bias, it only con-

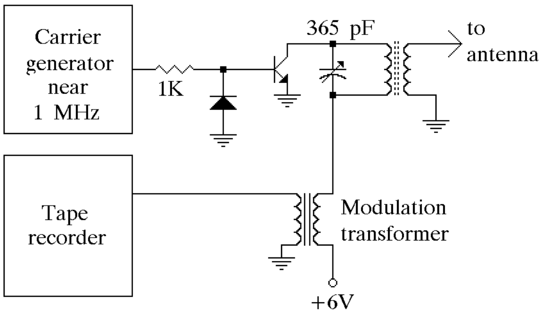

Fig. 7-4. A simple AM transmitter lab experiment

ducts when the generator's output is sufficiently positive (above about 0.7 volts) to bias it on. Thus the transistor only conducts some of the time. But when it does conduct, there is enough voltage and current in the base circuit to turn it on all the way.

The transistor therefore behaves like a switch, which is turned on and off roughly 1,000,000 times per second (depending on the frequency chosen for the signal). This applies a square wave on-off current to the tuned circuit at the carrier frequency. As we have seen in previous chapters, a square wave consists of a fundamental frequency plus harmonics. But the tuned circuit removes the harmonics, so the signal going out to the antenna should (hopefully) be just a sine wave at the carrier frequency.

The size (amplitude) of that sine wave depends on how much voltage gets switched by the transistor. Although the collector voltage supply is shown as +6 volts, the modulation transformer in series with the dc input changes that. When we play a tape on the tape recorder, the audio signal sent to the modulation transformer is alternately positive and negative, varying in step with the audio. The Transformer's secondary voltage is therefore also alternately positive and negative. When it's positive, it adds to the +6 volts dc to produce more voltage (and a bigger output carrier signal); when negative, it subtracts from the +6 volts to produce less voltage (and therefore a smaller output carrier.) In other words, the audio signal from the modulation transformer varies or modulates the amplitude of the carrier.

Modulation Percentage

There is a limit to how big that audio signal can get, because when the transformer output reaches –6 volts, it subtracts from the +6 volts to give zero; at this point, the signal disappears. If the transformer output were to get even more negative (such as –7 or –8 volts), the collector would go negative, and the carrier would be shut off completely for a while.

The point where the signal just barely disappears is called 100% modulation, and is the maximum that we can vary the amplitude. If the signal from the tape recorder is a sine wave at this point, the transformer output would have a maximum (peak) amplitude of ±6 volts, so that the collector voltage would vary from its normal +6 volts all the way down to 0 volts, and up to a maximum of +12 volts.

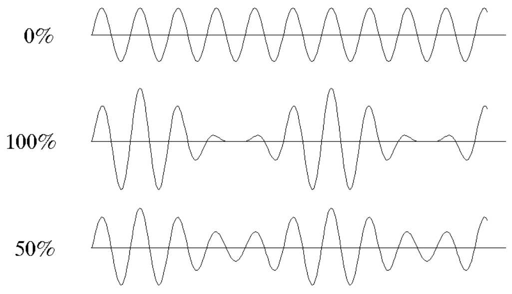

Fig. 7-5. Three different modulation percentages

100% modulation is a technical limit on how much you can modulate, but it is also a legal limit — if you tried to modulate a commercial transmitter more than 100%, your carrier would alternately go on and off; this would create a huge amount of interference to other stations, and the FCC would cite you for improper operation of your transmitter.

Fig. 7-5 shows a carrier with three different amounts of modulation — 0% (which is no modulation at all), 100% (which is the maximum permitted), and 50% (which is half-way between).

Actually, 100% modulation is just barely permitted, since it does cut off the carrier for an instant; most AM transmitters are set up so they will use a maximum of perhaps 95% or 98% modulation, just to avoid the possibility of accidentally going over 100%.

Let's take a closer look at Fig. 7-6, which shows some unknown percentage of modulation — how can we figure out the actual percentage?

It's actually quite simple. First, note that the maximum carrier voltage (called $V_{\max}$ in the figure) is 21.1 volts, while the minimum carrier voltage (called $V_{\min}$) is 3.7 volts. Assuming symmetric modulation (meaning that the audio signal goes up by the same amount as it goes down), this would place the average carrier voltage without modulation (shown as V_c at the left) halfway between the maximum and minimum. This is the average voltage, found from

$$V_c = \frac{V_{\max} + V_{\min}}{2} = \frac{21.1 + 3.7}{2} = \frac{24.8}{2} = 12.4 \text{ volts}$$

At the peak of the modulation, the voltage goes from 12.4 up to 21.1, which is an increase of (21.1 – 12.4) or 8.7 volts. At the valleys, the voltage goes from 12.4 down to 3.7, a decrease of (12.4 – 3.7) or also 8.7 volts.

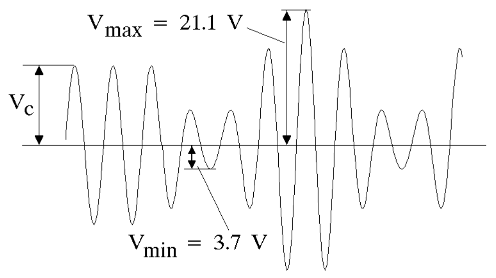

Fig. 7-6. Calculating the modulation percentage

To find the modulation percentage, we have to ask this question: a drop of 8.7 volts is how many percent of the maximum possible drop? Since the maximum possible drop (which would be 100% modulation) is 12.4 volts, 8.7 volts is what percentage of 12.4 volts? The formula is

$$\frac{8.7}{12.4} \times 100\% = 70\%$$

So Fig. 7-6 shows 70% modulation.

DETOUR

If you have, or can use, a PC-compatible computer, the following program lets you see a carrier with different percentages of modulation. It’s written in Basic, and can run with either IBM Basic, GWBasic, QBasic, or Quick Basic:

```
10 ’Program to display AM
20 PERCENT = 70 ’Enter percent here
30 SCREEN 2 : PRESET(0,100)
40 FOR X=0 TO 639
50   CARRIER = SIN(X/5)
60   AUDIO = SIN(X/50) * PERCENT/100
70   TOTAL = (1 + AUDIO) * CARRIER
80   Y = 100 + 40 * TOTAL
90  LINE -(X, Y)
100 NEXT X
110 IF INKEY$="" THEN 110
120 SCREEN 0
```

Enter the desired modulation percentage in line 20, and then run the program.

Line 30 of the program puts the display into graphics mode, and positions a black dot halfway down the left side of the screen. Lines 40 through 100 set up a loop to plot 640 dots across the width of the screen.

At each of the 640 positions, line 50 calculates a carrier voltage, line 60 calculates the audio (modulation) signal, and line 70 puts them together. By multiplying the two, it uses the value of the audio voltage to set the height of the carrier, and finally lines 80 and 90 calculate the height of the point on the screen and plot it.

END OF DETOUR

AM Sidebands

In Chapter 1, we learned that a square wave consists of a fundamental frequency and some harmonics. In fact, we made the sweeping statement that *any* repetitive waveform can be broken down into a fundamental and/or some harmonics. So that brings up the question — is this also true for the AM signal in Fig. 7-5? The answer is “Yes”, as long as we remember that some of these (the fundamental as well as the harmonics) might be missing.

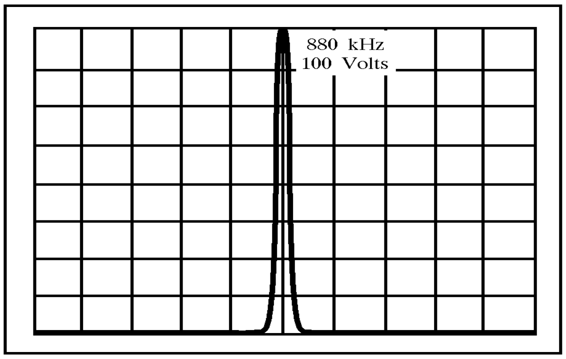

Fig. 7-7. An unmodulated carrier at 880 kHz

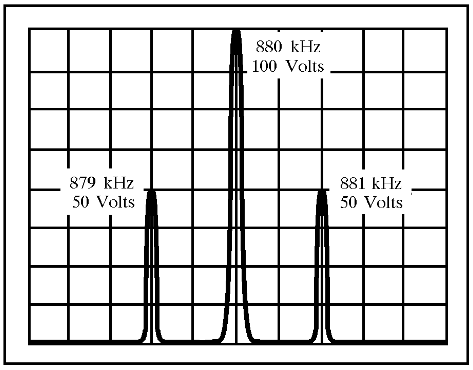

Fig. 7-8. 880 kHz carrier modulated 100% by 1 kHz

Let’s start off by putting down the general rule, and then we will explain it: *When a carrier with frequency C is amplitude modulated by an audio (or other) signal having frequency A, we get two sidebands, whose frequencies are C+A and C–A. Moreover, at 100% modulation, the two sidebands are exactly one-half the size of the carrier.*

Let’s consider an example. Suppose an AM radio station at 880 kHz transmits a carrier, whose amplitude is 100 volts when there is no modulation (that is, when there is no sound being transmitted.) If you looked at this signal with a spectrum analyzer, you’d see just a carrier, as in Fig. 7-7.

Now suppose the announcer steps up to the microphone and whistles a 1-kHz note into it, loud enough to produce exactly 100% modulation. If you looked at the transmitter output with an oscilloscope, you would see a 880kHz carrier, whose envelope goes up and down at a 1 kHz rate. But if you looked at that same signal with a spectrum analyzer, you’d see the picture of Fig. 7-8. You would still see the same 100-volt carrier at 880 kHz *(which never changed amplitude)*, plus a 50-volt signal (called an *upper sideband*) at 881 kHz (881 is 880 plus

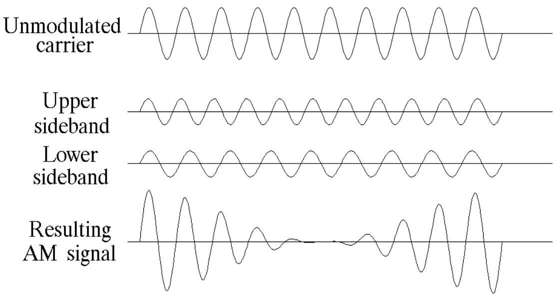

Fig. 7-9. The carrier and its sidebands

1 kHz), and another 50-volt signal (called a *lower sideband*) at 879 kHz (which is 880 kHz – 1 kHz).

So what is really happening is that the radio station is generating a 880 kHz carrier, and then AM modulating it (changing its amplitude) at a 1 kHz rate. But on the air, we actually have a 880 kHz carrier which has a constant amplitude, plus two extra signals called sidebands. These sidebands also have a constant amplitude. But when we see the combined signal on an oscilloscope, all three of these interact together to make it *look* as though the carrier is changing height.

Fig. 7-9 shows how this is possible. This is a computer-generated graph, which shows three sine waves at the top, and their sum on the bottom. The very top signal is the carrier, the second is the upper sideband, the third is the lower sideband. The bottom waveform is simply the point-by-point sum of the top three waves. At the far left and far right, the three top waves are pretty much in phase, and so they add up to a big result. In the middle, however, the two sidebands are out of phase with the unmodulated carrier, and so they cancel it out to produce a very small result.

DETOUR

Here's the Basic program which draws the same waveforms as in Fig. 7-9:

```
10 'Generate AM from sidebands
20 SCREEN 2
30 FOR X=0 TO 639
40    CARRIER=SIN(X/10.185616#)
50    UPPER=.5*SIN(X/10.185616#*1.1)
60    LOWER=.5*SIN(X/10.185616#*.9)
70    AM=CARRIER+UPPER+LOWER
80    PSET(X,20-20*CARRIER)
90    PSET(X,60-20*UPPER)
100    PSET(X,90-20*LOWER)
110    PSET(X,150-20*AM)
120 NEXT X
130 IF INKEY$="" THEN 130
140 SCREEN 0
```

END OF DETOUR

Look at Fig. 7-9 again. Suppose this graph covers exactly one second of time. Since the unmodulated carrier on top has exactly ten cycles, its frequency is exactly 10 cycles per second, or 10 Hz. It's a little harder to count the cycles of the AM signal on the bottom, but it too has exactly 10 cycles, and therefore is also exactly 10 Hz. Now look at the envelope of the AM signal; draw it in, if that will help. The envelope starts big, goes to small in the middle, and then becomes big again, so it has exactly one cycle during that second, and is therefore at 1 Hz. So we have a 10 Hz carrier, modulated by 1 Hz. (These frequencies are not exactly practical for real radio, but they are simple numbers, and so easy to visualize.) According to our previous discussion, the two sidebands should therefore have frequencies of 11 and 9 Hz (which is 10+1 and 10–1).

Sure enough, if you count the cycles of the two sidebands, you can see that the upper sideband has 11 cycles, while the lower sideband has only 9 cycles. Their frequencies are therefore 11 Hz and 9 Hz, respectively.

Note also that the AM signal is modulated to 100% — you can see that its amplitude goes all the way to zero. You can also see that the two sidebands are each exactly one-half the size of the unmodulated carrier.

Returning to the 880 kHz transmitter example, let's use the same notation that we used in Fig. 7-6. The carrier voltage V_c is 100 volts. When the two 50-volt sidebands are in phase with the carrier, all the voltages will add up to a $V_{\max}$ of 100 + 50 + 50, or 200 volts. On the other hand, when the two sidebands are out of phase with the carrier (look at the center portion of Fig. 7-9 to see how this happens), they all subtract to give us a $V_{\min}$ of 100 – 50 – 50, or 0 volts.

So we now have two ways of determining the percentage of modulation of an AM station — look at it on the oscilloscope, or look at it on the spectrum analyzer. For example, what is the modulation percentage for the signal in Fig. 7-10?

Let's see. The carrier has a height of about 6.4 divisions. Since we can't see the knobs on the analyzer, we don't know how many volts that is, but that doesn't matter — 6.4 divisions is good enough for us.

If that signal were 100% modulated, then the sidebands should be half of 6.4, or 3.2 divisions. But they are only about 1.7 divisions high. So the modulation is only 1.7/3.2 of its maximum. That works out to

$$\frac{1.7}{3.2} \times 100\% = 53\%$$

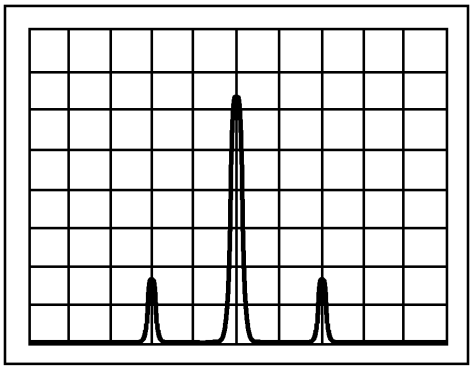

Fig. 7-10. What is the modulation percentage?

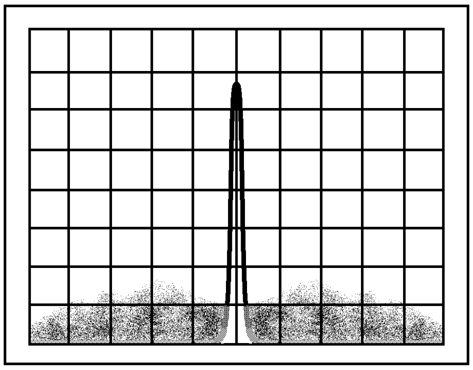

Fig. 7-11. An actual broadcast station

So far, so good. Now that we've gotten so good at this, let's look at a *real* AM broadcast station on the analyzer, and try to figure out what their percentage of modulation is. It's actually quite easy — we just have to throw some wire out the window and connect it to the analyzer's input, and we get Fig. 7-11.

We still see a nice carrier, but on each side we now see just plain fuzz, rather than a neat sideband. What's going on?

The difference is that the real station isn't broadcasting just plain tones — their audio consists of music and speech. This includes many different frequencies, all at different amplitudes but at the same time. Every single frequency in the audio generates its own pair of sidebands. So we have many sidebands, all occurring at the same time, and all constantly changing as the music or voice changes. The result? A fuzz that actually extends past the edges of the analyzer picture.

To avoid confusion, some people use the term *side frequency* when they describe the sideband from a single tone, as in Fig. 7-10. They would then say that all of these different side *frequencies* combine to make the two side*bands* in Fig. 7-11, a lower sideband to the left of the carrier, and an upper sideband to the right.

Bandwidth

Our simple example with the announcer whistling at 1000 Hz showed that the radio signal would consist not just of the carrier at 880 kHz, but also of sidebands (side frequencies) at 879 and 881 kHz, 1 kHz away from the carrier on each side. But AM broadcast stations normally transmit voice or speech with a frequency range of about 50 to about 10,000 Hz — not quite hi-fi, but still higher than just 1 kHz. Since the side frequencies lie at the carrier frequency plus and minus the audio frequency, we will now have side frequencies that lie anywhere from 50 Hz to 10,000 Hz away from the 880 kHz carrier. In other words, the sidebands will extend 10 kHz out from the carrier in both directions, down to 870 kHz on the left, and up to 890 kHz on the right.

The complete radio signal will therefore take up 20 kHz of space on the band (from 870 to 890 kHz), and so we say that the bandwidth of the AM signal is 20 kHz. We can summarize it this way: *The bandwidth of an AM signal is twice the highest audio frequency being transmitted.*

The bandwidth is related to how close stations can be placed to each other — the wider the bandwidth, the farther apart they have to be to avoid interference. To limit the bandwidth of commercial broadcast stations, the FCC limits their audio frequency range to a maximum of 10 kHz, which sets the maximum bandwidth at 20 kHz. Actually, though, AM broadcast stations are placed farther apart than their bandwidth would indicate, because the tuned circuits in most radios are not good enough to separate stations that are too close together. For example, in the New York City metropolitan area, stations are typically 50 or 60 kHz apart.

Sideband Power

When a transmitter sends out sidebands, it needs power to do that. Let's consider a commercial AM broadcast station which sends out a 5 kW carrier. If the cable leading from the transmitter to the antenna has an impedance of 50 ohms, then we can find the output voltage by solving the power equation, $P = V^2/R$, backward for the voltage:

$$V = \sqrt{P \times R} = \sqrt{5000 \text{ watts} \times 50 \text{ ohms}} = 500 \text{ volts}$$

When the station modulates that carrier at 100%, it still sends out a 5 kW (500-volt) carrier, but now it also sends out two 250-volt sidebands. Each of those also goes into the 50-ohm cable, so the power in each sideband is

$$P = \frac{V^2}{R} = \frac{250^2}{50} = 1250 \text{ watts}$$

The transmitter is now putting out 5000 watts into the carrier, plus another 2500 watts total into the two sidebands, for a total of 7500 watts.

Efficiency

The amount of sideband power depends, of course, on the modulation percentage. From the $P = V^2 / R$ equation, we see that the power is proportional to the square of the voltage. Thus cutting the modulation down to 50%, for example, drops the sideband voltage to ½ of the 100% value, and so drops the sideband power to ¼. Cutting the modulation percentage to 10% would drop the sideband voltage to $\frac{1}{10}$, and the sideband power to $\frac{1}{100}$. So in normal speech or classical music, where the volume is seldom at its maximum, the modulation percentage and sideband power tend to be much smaller than the carrier power. On the other hand, in modern popular music, where it seems like everybody is constantly yelling and screaming (to my jaded ears, anyway!), the average modulation percentage would be higher, and the sideband power also.

In any case, we now know that the carrier in an AM signal never changes; only the sidebands change, depending on the signal being sent. We also know that the carrier contains most of the power of the AM signal. But is it really needed?

If the carrier never changes, then it doesn't carry any information to the receiver. By the time it gets to the receiver, it is just a weak sine wave. In fact, if the transmitter somehow turned the carrier off (sending only the sidebands), and the receiver had a circuit which generated a substitute sine wave of the right frequency, amplitude, and phase, it would never know the difference. But generating a weak sine wave at the receiver is a lot cheaper than generating a high power carrier at the transmitter.

This is the basic idea behind several variations on AM; changes which make both the transmitter and receiver somewhat more complex, but which greatly increase the efficiency of the system by reducing the transmitter power. These major two methods are called DSB and SSB.

Double Sideband (DSB)

DSB is also sometimes called DSSC, which stands for Double Sideband Suppressed Carrier, and the name describes it well. In DSB, the transmitter eliminates the carrier and sends out only the two sidebands; the receiver then reinserts a substitute carrier.

Although the transmitter could produce regular AM and then filter out the carrier, this approach would waste power and be expensive. A better approach is to use a *balanced mixer* to produce the DSB directly. In the balanced mixer, two AM signals are produced at low power, and then combined in such a way that the carrier is canceled out, but the sidebands stay.

The major advantage of DSB is efficiency. Consider, for instance, the 5000-watt AM transmitter we discussed a moment ago. At 100% modulation, the transmitter must add 2500 more watts for the two sidebands, for a total power output of 7500 watts. On the average, considering typical voice or music, the sidebands might only contain perhaps a total of 1000 watts or so, but the transmitter must still output the carrier too. So its average output power might be closer to 6000 watts. If it could eliminate the carrier, then it would only need to output an average 1000 watts or so for the same sideband strength. The resulting signal would travel just as far.

Alternatively, if you didn't mind spending the money, you could put all 6000 watts into the sidebands, and get the same punch as if you had an AM transmitter of perhaps 20,000 watts or more!

And there would be still another advantage too. When two signals are close together in frequency, they interfere with each other. It turns out that most of the interference is between the carriers, not between the sidebands, so eliminating the carriers would eliminate much of the interference. In fact, during silent passages there is no interference at all since the sidebands are only there when there is something being transmitted.

By now you're asking, "If DSB is so wonderful, why doesn't everyone use it?" There are two answers to this: (1) It has some disadvantages, and (2) there's something even better.

One disadvantage is that it makes the receiver more complicated. In commercial broadcasting, the philosophy has always been to make the receivers as cheap as possible, so everyone can afford one — even if that makes it more expensive for the broadcasters. DSB doesn't fit into that pattern.

A second problem is that inserting a fake carrier in the receiver is not that difficult; but making sure that it is just the right frequency and amplitude is tricky. Putting it at the wrong frequency is the same as putting the sidebands in the wrong place — if the frequency difference between the carrier and the sidebands is wrong, the frequency of the audio will be wrong too. Even a slight difference — on the order of a tiny fraction of one percent — can make voices sound funny, and make music unlistenable.

So DSB would be OK for voice communications, but not really for music. (There are some places where it is used for music, as in stereo FM, but that's a special case, because some extra circuitry is used there to help the receiver put the carrier in the right place.)

But for strictly voice communications, there is something even better – SSB.

Single Sideband (SSB)

If you look at any of the spectrum analyzer pictures of AM — Figs. 7-8, or 7-11, for instance — you will note that the lower sideband is always the mirror image of the upper sideband. For every side frequency component in the lower sideband, there is also a matching side frequency component in the upper sideband. So why does the receiver need both sidebands?

It doesn't. As long as the receiver gets one sideband, it gets all the information there is. That is the idea behind SSB or Single Sideband. With SSB, the transmitter sends only one sideband — no carrier, no second sideband. There are two ways to do that. One way is to generate a DSB signal, but then filter out one of the sidebands; the other is to generate two DSB signals in such a way that adding them cancels out one of the sidebands because it is out of phase in the two signals.

Let's look at the advantages of SSB:

1. Much more efficient than plain AM, and even twice as efficient as DSB. No output power at all when there is no audio.

2. Even less interference to other stations than DSB, which also has no carrier.

3. Half the bandwidth of AM or DSB. Remember that the bandwidth of an AM signal is twice the highest frequency being sent. For example, a telephone-quality signal with audio from 300 to 3000 Hz would have an AM or DSB bandwidth of 6000 Hz. With SSB, on the other hand, the bandwidth is only 2700 Hz (since that is the frequency difference between the side frequency caused by a 300 Hz signal and the side frequency caused by a 3000 Hz signal.)

The bandwidth is important for several reasons. First, it means that twice as many stations can be crammed into the same bandwidth as AM or DSB (actually, more than twice, because there is so much less interference between them.) Equally important is the fact that the receiver can now have tighter filters, which can do a better job rejecting other noise as well. The noise power picked up by a receiver is proportional to the bandwidth; cutting the bandwidth in half cuts the noise power in half too.

What about the disadvantages? The major one is that it is even more difficult for the receiver to decide where to put the substitute carrier. With DSB, at least, the receiver can simply put the carrier exactly midway between the sidebands. With SSB, there is no midpoint. In fact, unless the user knows whether the transmitter is sending the upper sideband or the lower sideband, the receiver may not even know which side of the sideband to put the carrier on.

Tuning in an SSB signal is difficult. Even very expensive receivers require the user to adjust the tuning until "it sounds right." That may be close enough for voice, but useless for music. Hence SSB is used only for radio voice communications, never for music (although SSB is used for certain wire-line communications with better results, but there again there are special tricks used to give the receiver additional information to help it generate the right carrier.)

Vestigial Sideband

As we mentioned in Chapter 2, the bandwidth of the composite video signal in a TV is approximately 4 MHz. Since AM is used for transmitting the picture in commercial TV, modulating this composite video signal onto a carrier would normally result in a 8 MHz bandwidth (twice the highest frequency in the picture). But TV stations are only allowed 6 MHz bandwidth (and some of that has to be used for the sound carrier).

The solution used for TV is *vestigial sideband*. The word vestige means leftover or remainder. TV transmitters transmit the entire upper sideband, but only a part (the "vestige") of the lower sideband.

For example, TV's channel 2 occupies frequencies from 54 to 60 MHz, as shown in Fig. 7-12. The picture carrier is at 55.25 MHz; the upper sideband goes from 55.25 to 59.25 MHz, the full 4 MHz, while the lower sideband goes from 55.25 MHz down to 54 MHz, just 1.25 MHz. Even though most of the lower sideband is missing, the upper sideband contains all the information that the TV needs to properly receive the picture. We should mention that the top part of the 6 MHz channel is used for the sound signal, which has a carrier at 59.75 MHz. But this part is sent as FM or Frequency Modulation, so we will leave the discussion for later; for now, it's just important to mention that the TV picture and sound carriers are purposely different so that they cause

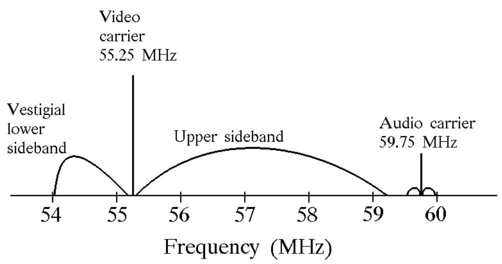

Fig. 7-12. TV Channel 2 spectrum

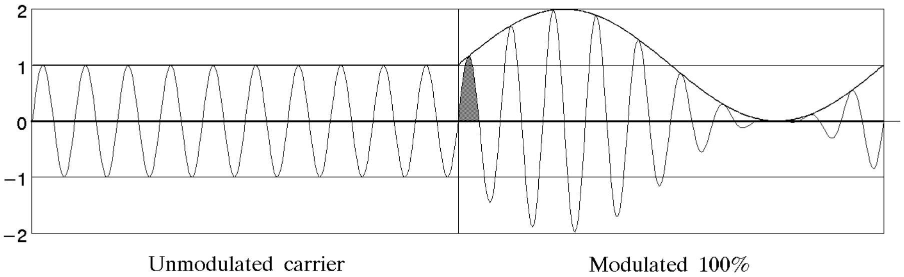

Fig. 7-13. Calculating the power of a 100% modulated carrier

as little interference with each other as possible. (The color subcarrier is also shown in Fig. 7-12.)

DETOUR

Here's an interesting calculation you can do yourself.

We've already said that, when an AM carrier is modulated 100%, each sideband is ¼ the power of the carrier. With two sidebands, that increases the power by ½. Here is how to check that.

Fig. 7-13 shows 10 cycles (20 half-cycles) of a 1-volt unmodulated carrier on the left, followed by 10 cycles (20 half-cycles) of a carrier modulated 100% on the right. Since the formula for power is $P = V^2/R$, the power is proportional to the square of the voltage. So let's square the voltage of each half-cycle of the unmodulated carrier on the left, and add up all the squares to get a some idea of the relative power (ignore the units.) Each half-cycle is 1 volt, whose square is also 1. For 20 identical half-cycles, the sum is 20.

Now repeat the same for each half-cycle of the modulated carrier. For example, the shaded half-cycle has a height of about 1.2 volts, so its square is about 1.44. Repeat that calculation for each of the 20 half-cycles (some are so tiny they are hard to see) and add up. You should get a sum of about 30, which is ½ more than 20.

This calculation shows that the total power of an AM signal increases by 50% when it is fully modulated. As we explained before, this extra power all goes into the two sidebands.

END OF DETOUR

Frequency Division Multiplexing (FDM)

If you put up a receive antenna, and then look at the signal coming down the antenna wire to the receiver, you will find that it containst the signals of many stations — radio stations, TV stations, telemetry, pagers, vell phones, and more. All of these signals travel down the same wire; the difference between them (and the factor that makes it possible for a receiver to pick out the one you want and reject all others) is the different frequency range of each signal. This is a perfect example of the general idea of *frequency division multiplexing* or *FDM*.

To *multiplex* means to combine many signals into one in such a way that they can later be *demultiplexed* — that, is, separated back into the individual component signals. With frequency division multiplexing, a fairly large band of frequencies is subdivided, and each of the desired signals is assigned a portion of that bandwidth.

FDM is used in other places besides wireless transmission. For example, an old FDM system used by telephone companies combines twelve 4 kHz-wide voice channels onto one cable by using the lower sideband of a single-sideband (SSB) signal like this:

Voice channel	Carrier Freq.	Lower sideband BW
1	64 kHz	60–64 kHz
2	68 kHz	64–68 kHz
3	72 kHz	68–72 kHz
:	:	:
12	108 kHz	104–108 kHz

The resulting signal has a bandwidth of 48 kHz.

In many cases, FDM has been replaced by time division multiplexing, which we will treat in a later chapter. Still, FDM still has an important place in communications.

Conclusion

Amplitude modulation is used in a number of important places. Commercial AM broadcasting uses it, of course, but so does TV (although with a vestigial sideband.) It's also used in the aircraft band, and by international broadcasters on the shortwave bands.

In amateur radio and other point-to-point communications, however, AM has been replaced by other modulation methods. SSB is important, but so is FM (Frequency Modulation, discussed in the next chapter) and various digital pulse methods (also to be discussed in later chapters.)