ECL LOGIC CIRCUITS

Emitter-coupled logic, though not as popular as other logic families, is by far the fastest available commercially. Let's take a look at some ECL basics.

TJ BYERS

WHAT WOULD YOU SAY IF I TOLD YOU that there are IC's available that can process data at the rate of 500 million bits per second? Some new laboratory discovery, you ask? Hardly. I'm talking about ECL (*E*mitter *C*oupled *L*ogic) circuits. Although ECL IC's became commercially available in 1962, they are still not very popular with hobbyists. However, due to their ultra-high-speed switching properties, they are popular in the data-processing, test-equipment, and digital-communications industries.

ECL IC's can be purchased today for about the same price as you're paying for TTL IC's. Does all this sound too good to be true? Well, before you rush out to replace your "antiquated" designs with ECL, first let me tell you that the average ECL circuit is nothing at all like the standard TTL format you've grown to love and understand.

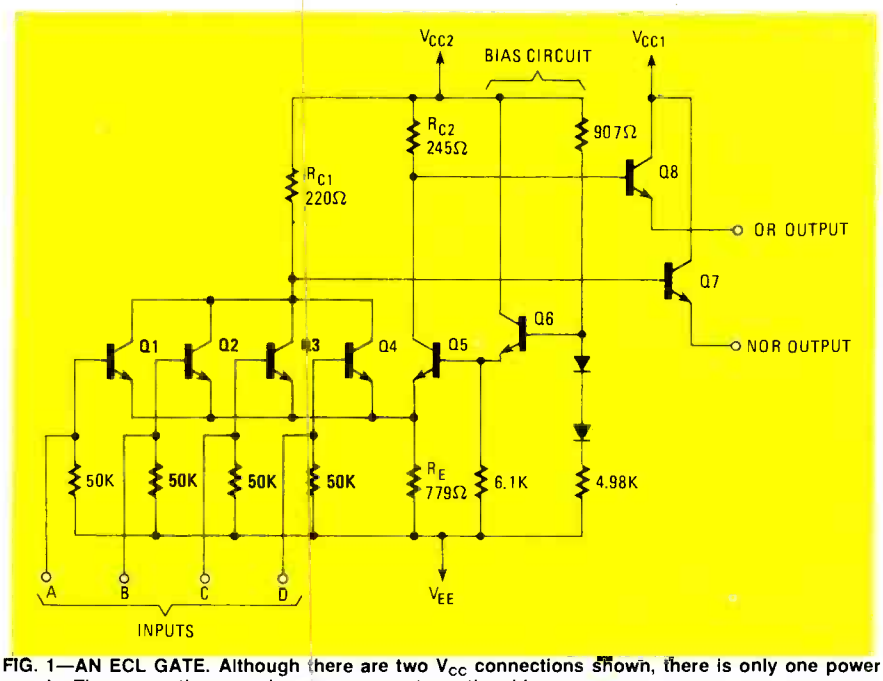

FIG. 1—AN ECL GATE. Although there are two V_{CC} connections shown, there is only one power supply. The connections are, however, separate on the chip.

ECL theory

Unlike the popular TTL-style logic, which exploits the two states of a transistor (either on or off), ECL is biased so that the transistors are in the active region at all times. (That is why ECL is often called nonsaturated logic.) Because the transistor does not go into saturation, there is no stored-base-charge problem; so propagation-delay times are very small. The IC's have speed capabilities a full order of magnitude faster than other existing technologies.

Figure 1 shows a complete ECL gate. If you look at it closely, you'll recognize the input section as being a slightly modified differential amplifier—similar in design to the input of an operational amplifier. However, also notice that one side of the differential "pair" includes four transistors, all in parallel. Each represents a separate gate input.

To the right of the differential amplifier is a biasing network that provides a stable, temperature-compensated voltage source. That voltage is input to differential-transistor Q5, establishing a reference level for differential operations. With no signal present on any of the inputs, Q5 conducts, resulting in a voltage drop across R_{C2}. Output transistors Q7 and Q8, whose base circuits are coupled directly to the collector resistors R_{C1} and R_{C2} respectively, monitor the status of the differential amplifier. Let's see what that means.

With no inputs, Q5 conducts, and little current flows in the other leg of the differential amplifier; there is no appreciable voltage drop across resistor R_{C1}. Thus the voltage at the base of Q7 is V_{CC}. Because

the output circuit is being operated as an emitter follower, the voltage developed at the emitter (output) of Q7 will be V_{CC} plus the voltage drop across the base-emitter junction—or approximately 0.9 volts less than V_{CC}.

Unlike R_{C1}, there *is* a voltage drop across R_{C2} when Q5 is conducting. In fact, with the resistor values shown, the voltage drop is very nearly 1 volt. As before, the voltage developed at the emitter of Q8 will be the base voltage plus the voltage drop across the base emitter junction. The emitter voltage of Q8 will be about 1.75 volts less than V_{CC}. With the outputs configured this way ($V_{E7} = -0.9$, and $V_{E8} = -1.75$). We define it as our OFF state.

If we now apply a voltage—which is in excess of the reference voltage established by the bias network—to the base of one of the input transistors, it will conduct. That increases the current flow through the emitter resistor R_E.

The increase in current raises the voltage drop across the R_E, causing a proportional decrease in the current flowing through Q5—as is the nature of a differential circuit. Very little current will flow through R_{C2}, but there is will be heavy current flow through R_{C1}. The values of the resistors have been selected so that there will be approximately a one-volt drop across R_{C1}, and almost no voltage drop across R_{C2}. In effect, the roles of the output transistors are reversed from the previous case. Now the output of Q7 is about -1.75 volts and the output of Q8 is about -0.9 volts; that represents the ON logic state.

When the input signal is removed, the input transistor(s) drop out of operation, and the current reverts back to Q5, which in turn toggles the outputs again. That switching action is very smooth; a steady current flows from the power supply, reducing switching transients.

Since the transistors never go into saturation—a requirement for stable TTL operation—there is no charge buildup at the junction. That means no storage-time effect to slow down the switching action. The design is so effective that the ECL gate is able to alternate states in a matter of *picoseconds*.

Power supply

In order to effectively use these fast switching times, which are equivalent to 1 GHz in some IC's, other parts of the circuit take on a different look. One of the changes you'll see is in the power supply.

Characteristic to ECL circuits, the V_{CC} voltage—which is our positive line—is grounded. That means that our V_{EE} line must therefore assume a negative potential (usually -5.2 volts). Although any voltage source of the ECL family can be designated as ground, there is a valid reason for selecting the V_{CC} source.

To better understand the reasoning behind this, again think of an ECL gate as a differential amplifier (which it is, basically). Differential amps have the capacity for very high common-mode rejection, allowing the input circuit to ignore many sources of noise and interference that are inherent in switching circuits. However, noise generated on the V_{CC} line is not canceled out by the differential circuit. However, by referring V_{CC} to ground, good noise suppression can be accomplished.

That changes our way of thinking somewhat, because a "1" level is now represented by ground potential. Consequently, the "0" logic assumes a negative value, which is consistent with our new thinking.

That brings up another point about the V_{CC} voltage. If you look again at the schematic in Fig. 1, you'll notice that there are *two* V_{CC} connections to the gate. That is quite intentional. The output transistors drive the load using emitter followers, and the currents (particularly surge currents) are at times very heavy. That is reflected back into the V_{CC} line as a voltage spike, or a glitch. At the speeds involved, crosstalk inside the integrated circuit is a real problem. To keep crosstalk to a minimum, the two power sources are separated inside the package. However, that doesn't mean that you need two V_{CC} power supplies. Both pins are returned to the same place (the positive leg of the power supply). However, in order to achieve the isolation intended, heavy bus lines are used to avoid coupling between the voltage pins. In fact, the entire plane of a double-sided PC board is often used for the V_{CC} connections, and will provide a substantial low-impedance ground at the same time. Using the ground-plane construction brings us other benefits, as we will see a little later on in our discussion.

Although the ECL family of circuits is intended to operate from a -5.2 volt source, it is capable of working over a much broader range (-3.0 to -8.0). Unfortunately, performance (with respect to speed and noise) is not always predictable across the entire range, and it's best to keep the voltage within 10% of the -5.2 volts specified.

Power dissipation

One of the major complaints that has been lodged against the ECL family is that of excessive power dissipation. Since the transistors operate in their active region, it is naturally assumed they will dissipate more power than other, saturation-type logic families. However, that is not always the case. If you look at the power-dissipation curve in Fig. 2, you will see that although the TTL-style logic does use less power initially, the condition soon changes as the input frequency increases. As a matter of fact, even the power-miserly CMOS gate (operating with a 10-volt source) soon dissipates more power than what is dissipated by an

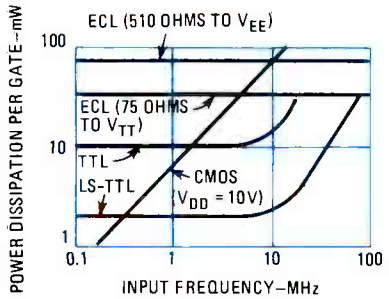

FIG. 2—POWER-DISSIPATION CURVES for various logic families shows that while ECL is often considered to be a power waster, at high frequencies it is quite miserly.

ECL gate—and before the operating frequency even reaches 10 MHz!

Even though TTL and LS-TTL seem to fare better at first, at high frequencies they, too, end up wasting as much power as an ECL gate. The ECL gate, on the other hand, keeps right on switching along, maintaining the same level of power regardless of the input frequency. In fact, in many situations the ECL design becomes a strong watt-for-watt competitor at frequencies as low as 20 MHz.

Logic levels

Since differential currents are used to represent the logic states, the logic voltages center around a threshold voltage that is established by the bias network. That can be seen from the transfer characteristics of the ECL gate, shown in Fig. 3. The "1" level is defined as -0.9 volts and the "0" level is -1.75 volts. Notice the use of negative values when expressing the logic levels, keeping consistent with our earlier observations.

The reference voltage for Q5 is set at approximately -1.3 volts. You'll see that the levels defined for our logic states pivot about the reference by just about 0.4 volts. By keeping the voltage levels confined to a narrow voltage band, rather than setting them at V_{CC} and V_{EE}, the *slew* effect is minimized. Slew is best defined as the amount of time it takes a pulse voltage to travel from one level to another. You probably know it better as rise time.

The greater the voltage transition, the more time it takes for the gate to switch. In other words, you can increase the op-

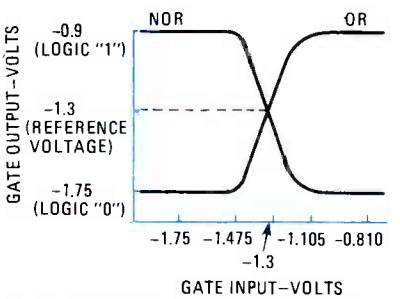

FIG. 3—TRANSFER CHARACTERISTICS of an ECL gate.

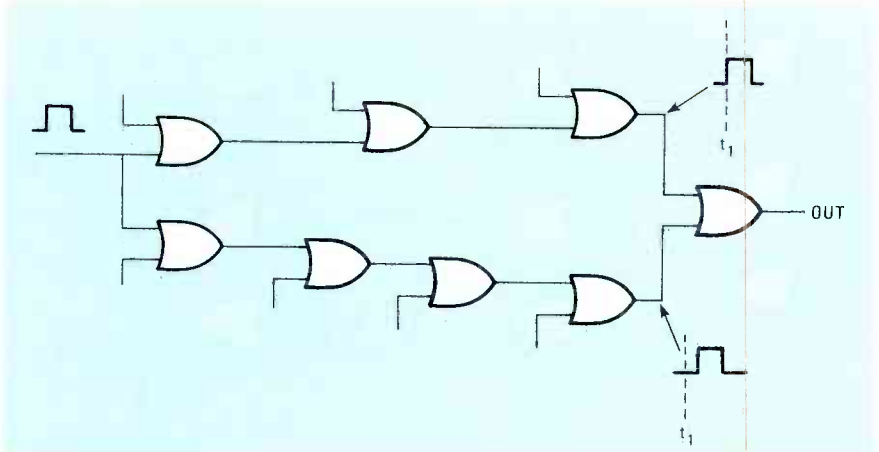

FIG. 4—THE EXTREMELY HIGH clocking speed used with ECL circuits demands that careful attention be paid to design.

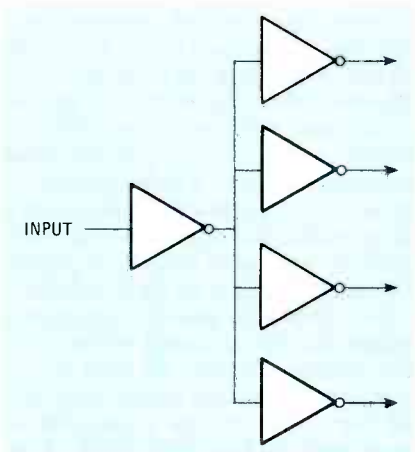

FIG. 6—AN ARRANGEMENT like this helps to minimize clock-skewing problems.

erating speed by limiting the voltage swing. That is not to say that the devices won't accept voltage levels beyond those defined as the logic states. They will. In fact, the situation often occurs at lower operating speeds. It is quite acceptable for the input pulse to run the entire gamut from V_{EE} to nearly V_{CC}.

High-speed logic

Basically, designing with ECL logic is pretty much like using standard logic IC's—if you remember that V_{CC} is ground. Because of the higher speeds involved, though, certain parameters require some important consideration. It's not unusual to find ECL gates operating comfortably at 500 MHz. At that speed, restrictions are frequently placed on pulse timing to a degree unimaginable with other forms of logic.

Assuming for the moment that the pulse width is 3 ns and that the delay per gate is 1.5 ns, the circuit in Fig. 4 will demonstrate our point. When a pulse is input to this configuration, it is split, with the signal traveling two different paths. (To simplify our example, only the actual signal paths are shown, but bear in mind that other inputs are present along the way, influencing the logical decisions.)

The pulse travels through the circuit and finally arrives at the last logic gate. However, the bottom signal arrives exactly 1.5 ns later than the top one. Now, one and one-half nanoseconds is not a very long time, and it would be ignored using standard logic practices. However, at the clocking speed we have set for ourselves in this circuit, it amounts to a full 25% of the clock cycle—and 50% of the pulse width!

At the beginning of the clock cycle, the final gate senses the inputs and makes its decision. However, 1.5 ns later the logic pulse from the lower path arrives at the input. Now, depending upon the logic states involved, the later signal could abruptly alter the status of the output. Furthermore, if there were a total of five gates involved in the lower path, the pulses would never coincide because they would be separated by 3 ns.

ECL clock skewing

Lead lengths must also be considered because of the high operating frequencies. A mere 12 inches of wire will delay a signal by about the same amount as one gate does. To emphasize that, let's take the circuit layout in Fig. 5, depicting, more or less, the physical positioning of several ECL gates on a PC board. The clock pulse is tied into the board at the top, and distributed to the IC's as shown. I'll bet you can see the problem already.

By the time the pulse reaches the fifth IC, the timing signal is skewed by a full 2.0 ns; on the back swing, the entire bottom row of chips is off by nearly 3 ns. As a matter of fact, the last chip is clocking more than 4 ns late; that is an intolerable situation!

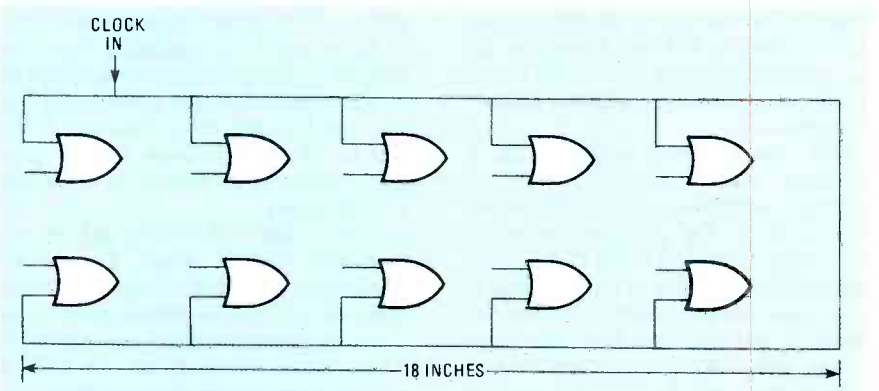

FIG. 5—LEAD LENGTHS become critical in ECL circuits.

To minimize such clock skewing problems, line delays must be matched to each other by better than 1 ns. A practical way to accomplish that is to use a logic-tree arrangement, such as is shown in Fig. 6. The timing signal is fed to one centrally located IC. From there, lines of matched length carry the timing pulses to the individual gates.

NOR logic should always be used when developing clock trees, and OR and NOR outputs should never be mixed within the chain. Due to lumped capacitance and other factors, loading of the outputs causes the gates to respond more slowly. So for high-speed performance, it is recommended that no more than four to six loads be applied to any one driver. Furthermore, when more than one gate is driven by a single output, it's better to run separate lines—of equal length—to each input rather than lump the loads at the end of a signal line. It is sometimes necessary (and usually desirable) to include an extra gate in a line to match delayed pulses from longer runs.

Parallel gates can be used to increase the bandwidth of an ECL driving gate when clock repetition rates are high, or when large capacitive loads are involved. Bandwidths can be extended by 40 or 50 MHz just by paralleling both halves of a dual gate.

As you may suspect, the timing problem becomes even more critical when more than one circuit card is involved in a system. Since the master clock is normally located on just one card, with output lines sent to the other cards in the system, it becomes imperative that all the feedlines be identical in composition (all coax, for example) and of *exactly the same length!*

Unused inputs

There always comes a time when there

are more logic functions contained on a chip than the designer needs. When that happens, something has to be done with the unused pins. The differential input of an ECL gate represents a very high impedance, and ordinarily a floating input would spell trouble. With no planned direction, reverse-bias leakage builds up a charge on the base-input lead. That often produces ambiguous outputs or power-dissipation problems for the IC involved.

To overcome the problem, a pulldown resistor has been included on the chip. The resistor references the input to V_{EE} (thus bleeding off any charge build-up), but has a large enough value—50K in most cases—to have little effect on the signal. As a result, all unused inputs can be left unconnected.

Be that as it may, there are exceptions. Several ECL devices don't have internal pulldown resistors; an example is a line receiver. In the case of a line receiver, one input must be tied to the V_{BB} pin (V_{BB} is a reference source used for Schmitt trigger applications) and the other returned to V_{EE} when the receiver is unused.

It frequently occurs during the analysis of a logic design that one of the inputs must assume a constant state. You see this configuration all the time in decision-making circuits that compare the input pulse to a fixed frame of reference. And, as has been the practice with standard logic, the inputs are hard-wired to either V_{CC} for a fixed HIGH, or V_{EE} for a LOW.

Although it is acceptable to tie the input to the V_{EE} line to simulate a low input with ECL gates, the reverse should not be attempted. Tying the input to V_{CC}—or ground, as is the case—to imitate a high input is not recommended. Due to their design, many of the ECL IC's won't operate properly when that method is used.

The proper way to force a logic "1" is to provide −0.9 volts at the input. You can use a resistor voltage-divider but most manufacturers recommend you use the forward voltage drop across a diode junction, as shown in Fig. 7, to provide the voltage required.

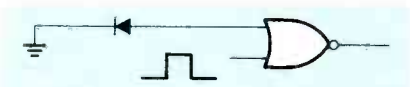

FIG. 7—A DIODE IS USED to imitate a "high" input for an ECL gate.

Outputs

Since the ECL output is an open emitter, an external load resistor is required to provide a current path for the output transistor. Normally, a resistance between 270 and 510 ohms is selected.

We should note that almost no dedicated driver circuits exist in the ECL family. That is because every gate is capable of becoming a driver. In fact, with the values so far specified, each output is capable of driving approximately 20 gate loads. The maximum fanout, however, is not limited by the gate's DC driving capacity as you may think.

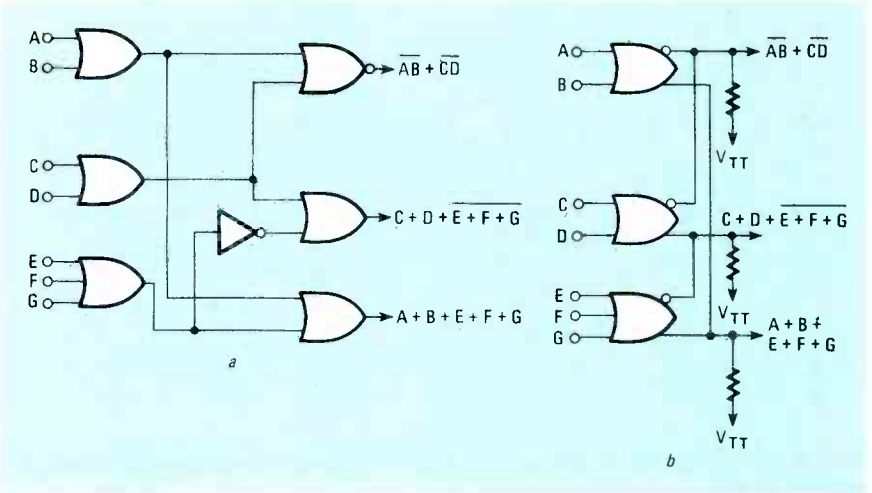

FIG. 8—THE NOR OUTPUT can be used to keep your designs more efficient, by reducing delay times and the number of gates required.

While DC loading, such as the pulldown resistors in the input circuits, can produce a shift in output voltage levels, it does little to alter processing times. AC loading, on the other hand, increases as the capacitance across the output circuit increases. The AC considerations are what determine the maximum fanout. If not for those considerations, the fanout drive capability would be greater than that which is required for most any practical application.

Driven through a typical output impedance of 7 ohms, the output transistor is able to force current through the capacitance, hastening the charge time and leaving the rise time less affected. Fall time, however, must depend on the discharge rate through the load resistor. Therefore, it is beneficial to keep the load resistance as small as practical. Unfortunately, as the load resistance decreases, power dissipation (in both the IC and resistor) increases. For that reason, 200 ohms to V_{EE} is suggested as the smallest value you should use. However, the gate input doesn't require that the logic voltage swing all the way to V_{EE}, remember? Any level in excess of −1.65 volts is all that's really necessary. Therefore, if we were to provide another power supply—one nearer the "0" level—and if we returned our load resistors to that line instead of the V_{EE} line, resistance values could be reduced with no burdening increase in power dissipation.

Well, that is exactly what is done. A −2.0-volt source, labeled V_{TT}, is established, and the output load resistors are returned to it. The pulldown-resistance now ranges from 150 ohms to 50 ohms, with 35 ohms specified as the minimum in most cases. Consequently, rise and fall times are less influenced by capacitive loading, power dissipation is reduced by a factor of four, and the "0" logic level is still within limits.

Of course, it must be decided during the course of design whether or not the cost and distribution of a −2.0-volt power supply is warranted. In many small systems, it isn't—in which case a resistor to V_{EE} is preferable.

Design shortcuts

The open-emitter drive makes it possible to engineer logic shortcuts into your designs. In particular is the wired-OR function. Wired-OR is the process of combining two (or more) outputs and OR-ing the results—without the use of a separate gate.

In general, it is recommended that the equivalent of only one pulldown resistor be used for the wired-OR design; although two resistors will improve fall times, it does so at the expense of added power dissipation. Due to LOW level current flow through the resistor, the number of gates paralleled in the wired-OR fashion should be limited to six in order to maintain a suitable level for LOW logic.

It is important that you keep the outputs involved physically close (from the same IC when possible). Anything else only aggravates the timing problems.

Another convenience of the ECL family is the inclusion of the complementary NOR output on most chips. Use of the NOR function eliminates the need for time-delaying inverters, reduces package count, and improves propagation time.

An example of combining these two features in a single design is shown in Fig. 8. The standard design is shown in Fig. 8-a, and Fig. 8-b shows how the design can be changed to save several gates while obtaining an increase in signal processing speed.

If you think ECL is reserved for only high-speed circuits...don't. They are just as effective in slower circuits. Moreover, many of their unique characteristics make them desirable in circuits now served by TTL. When we continue, we will investigate how to get the most from your ECL IC's. **R-E**

ECL LOGIC CIRCUITS

TJ BYERS

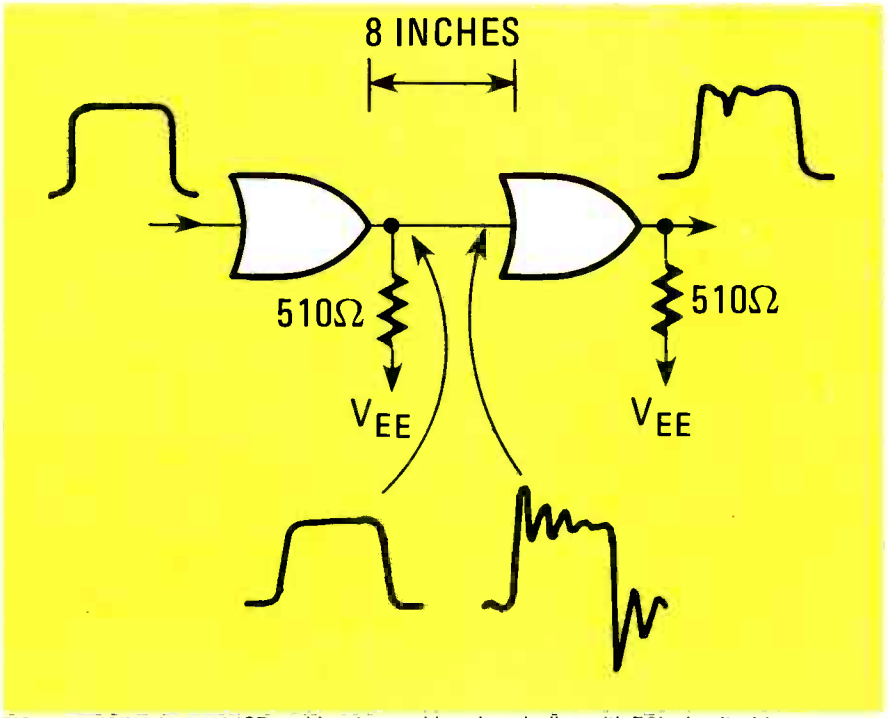

FIG. 1—RINGING IS A MAJOR problem to consider when dealing with ECL-circuit wiring.

Because of the high operating speed of emitter-coupled logic, standard wiring procedures cannot be used. Here we will look at the solution to the circuit-wiring problem.

Part 2 IF YOU READ THE FIRST installment in this series, then you have a general understanding of ECL (*E*mitter-*C*oupled *L*ogic) and its capabilities. However, to use ECL IC's, you have to understand more then the ECL gate. You have to pay close attention to the interconnections between devices. This month we'll investigate just that.

Wiring ECL gates

The application of ECL is identical to any other form of logic and, as with any other logic, the output of one gate must be connected to the input of the following gate. Normally, that is a routine kind of thing, and you simply run a wire from one to the other, just as we have done in Fig. 1. (Notice that, in compliance with the rules of ECL loading, a pulldown resistor is connected to the output pin of each gate.)

Before going any further, we should explain that at high frequencies, any wire connecting any two points (gate output to gate input, in our case) can be considered to be a transmission line. A transmission line has certain amounts of resistance, inductance, capacitance, and a time delay—all of which influence the signal traveling through it. We must also remember that if the transmission line is not terminated by its characteristic impedance, a portion of a signal flowing through the line will be reflected when it reaches the line's termination. Those reflections add to (or subtract from) the signal voltage. Reflections are present even at low frequencies, but in that case, they are usually masked by the relatively slow risetime of the pulse. However, when the delay time in the wire is longer than the risetime of the input pulse, the reflected power causes a *ringing* inside the line that affects the pulse. (If the ringing is limited to the risetime of the pulse it is not usually a problem, because the IC's are clocked after the steady-state levels have been reached.) For example, the lead length specified for our example in Fig. 1 would give a delay that is longer than the risetimes commonly encountered when using ECL gates. The result is shown—notice that a clean pulse enters the line from the gate output. But by the time it reaches the following input, it is distorted by ringing. The ringing is due to the reflected waves present in the transmission line.

ECL is forgiving to a certain extent, and some ringing is permissible. However, ringing on the input line does reduce the noise "safety" margin considerably and in some cases will even produce false triggering. Typically, an ECL gate will tolerate up to 35% overshoot and 15% undershoot. That's not a wide margin to work within!

Fortunately, there is a simple way to reduce ringing. By placing the load resistor at the end of the connecting lead—instead of at the gate output—the overshoot is attenuated. Instead of feeding a pulse down an open wire, the output circuit now sees a terminated low-impedance transmission line.

As shown in Fig. 2, that simple procedure gives us cleaner output waveforms. It now becomes apparent why the ECL-IC designers opted for an open-emitter driver and did not include a load resistor on the chip.

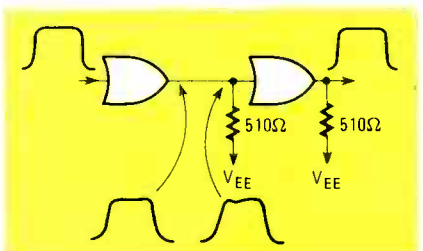

FIG. 2—CONNECTING THE OUTPUT RESISTOR at the end of the connecting lead can help to reduce ringing.

Ground planes

At higher frequencies, the noise picked up by an unshielded wire is prohibitive. One way to provide shielding, without using coaxial cables and the like, is to place the lead alongside a ground plane.

A ground plane is nothing more than a sheet of metal that is placed close to the interconnecting wire and is tied to the power supply return. (Remember, V_{CC} is our ground in ECL circuits.) Not only does it protect the signal from stray interference but it can help attenuate some of the unwanted reflections that occur in the line.

To emphasize the point, let's return to our circuit in Fig. 1. As you recall, Fig. 1 showed the tremendous amount of ringing present when the resistor was tied to the output pin of the gate. However, look at the same circuit and resulting waveform with only the addition of a ground plane (Fig. 3). You can see how much the overshoot is suppressed by the ground plane alone. Ground planes can be established in a number of ways. Probably the quickest—and easiest—is to use one side of a double-sided PC board.

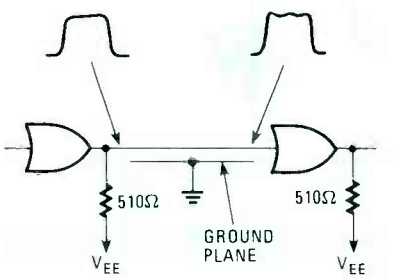

FIG. 3—USING A GROUND PLANE is another way to reduce overshoot and ringing.

Microstrip construction

A logical extension of the ground plane concept is microstrip construction. Microstrip design goes one step beyond the simple ground plane by allowing you to give a specific impedance to every line. In the ground-plane approach previously reviewed, no effort was taken to make sure that the impedance was constant. There are several advantages to being able to tailor the impedance of the transmission line. For one thing, it is much easier to match the load resistance to a line when you know its characteristic impedance. And, with a properly terminated line, a greater percentage of the input signal is seen at the other end. Thus, a wider margin for error is obtained. Moreover, it provides the highest possible noise rejection. A board using microstrip construction is shown in Fig. 4.

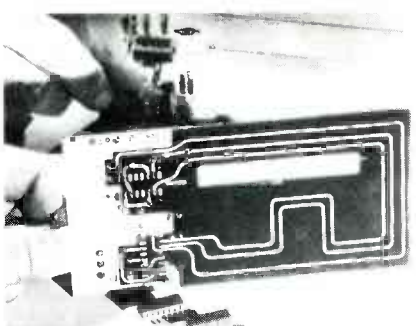

FIG. 4—THIS BOARD (an 83-MHz ring counter) uses 12-inch microstrip delay lines and a ground plane (seen in mirror).

Microstrip boards aren't hard to make—they're really no more difficult than a standard PC board except that the dimensions are more exacting. The microstrip transmission-line is characterized by a constant-width conductor on one side of a PC board, with a ground plane on the other side. The impedance is determined by the width and thickness of the conductor, the thickness of the board itself, and the dielectric constant, e_r, of the board material. That relationship between the impedance and those factors is summarized in the graph in Fig. 5.

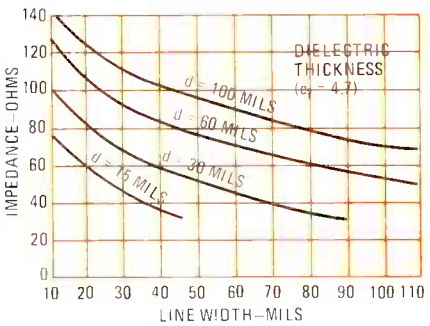

FIG. 5—THIS GRAPH SHOWS microstrip impedances for double-clad 1-ounce copper board; copper thickness = 0.0015 inches.

When laying out a microstrip board, certain precautions must be observed. First, there should be no squared corners in your leads—sharp bends should be avoided. For best performance, all bends should be given a radius no smaller than one-fourth the wavelength. Also, to minimize crosstalk, as much spacing as possible should be left between parallel lines. If you have no choice and have to separate two lines by less than 150 mils (0.15 inches), then a ground lead must be run between them.

For practical reasons, the characteristic impedance of the microstrip lines falls between 50 and 150 ohms. To achieve impedances greater than 150 ohms, the line width becomes prohibitively narrow; not that their construction isn't possible, but small imperfections in the etching process become more critical. That restriction, however, falls within the guidelines of good circuit design. As impedance increases, propagation time also increases. So, as far as speed is concerned, low-impedance lines are preferred. However, low-impedance lines require a low-value terminating resistance, which must—as we discussed last month—dissipate more power. An impedance of 68 ohms usually yields the best trade-off between power dissipation and speed—and happens to fall in the middle range of board construction. Of course, you are not restricted to using 68-ohm lines exclusively. You can use any impedance you deem necessary for the job. You can even mix the impedances on a board to tailor the performance for specific results, as we shall see shortly.

Line terminations

Anytime a transmission line is longer than the signal wavelength, termination of the line is a necessity. By using constant-impedance transmission lines, though, it becomes possible to terminate the line in more than one way and still achieve a good match with reduced overshoot.

We have already seen one—the use of a terminating resistor at the end of the line. That is called parallel terminations. It provides the highest speed while reducing the capacitance effect on the output of the gate. When one output drives several loads, however, there are a couple of variations to the parallel termination.

The first approach is to lump all the loads at the end of one transmission line, as seen in Fig. 6. Although that slows the risetimes and falltimes somewhat, because of the increased capacitance, it is desirable when all the inputs involved are located on a single IC. Notice that only one load resistor is used for all the inputs.

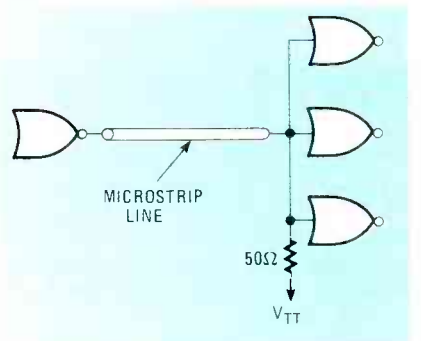

FIG. 6—PARALLEL TERMINATIONS. This approach uses only one load resistor.

An attractive benefit of a parallel-terminated line is the fact that the impedance is constant along its entire length. This makes it possible to tap the signal from any location along that length, as shown in Fig. 7. For proper distribution, though, the taps should be evenly spaced along the length of the line. You must keep in mind, however, that as the pulse progresses down the line, the delay increases. In other words, the first gate will receive its signal before the end gate.

A variation of the single line is the multiple-line mode. A representation of this method is shown in Fig. 8. Notice that the path to each input is through a separate transmission line. When the loads are scattered throughout the card, it is better to use that arrangement. You'll

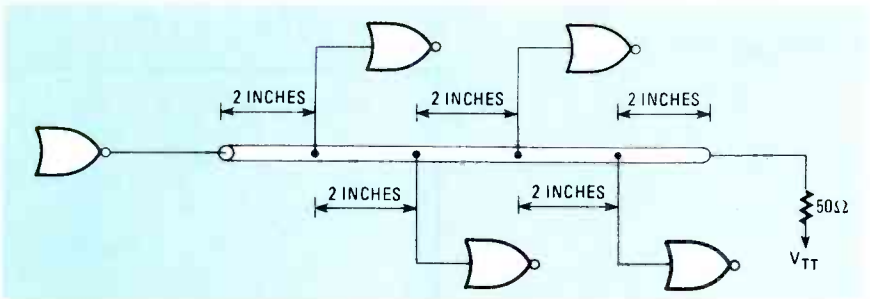

FIG. 7—BECAUSE THE IMPEDANCE is constant along its entire length, you can tap the signal at any point on the coaxial line.

also notice that each line is terminated by its characteristic resistance, which means that the power dissipation of the output gate increases as the number of lines increases. Therefore, it is best to use high-impedance lines so that the total lumped resistance doesn't exceed the DC limits of the output circuit. For instance, if we take the example in Fig. 8, the wise choice would be to run three 150-ohm lines to the inputs. In that way, the total load seen by the output will be 50 ohms—well within its operating parameters.

An obvious consequence of mixing impedances on a card, however, is that each impedance displays a different propagation time; delays increase as the impedance increases. Depending on the lengths involved, it's possible that pulse skewing could result even though you may have taken care to match wire *lengths*.

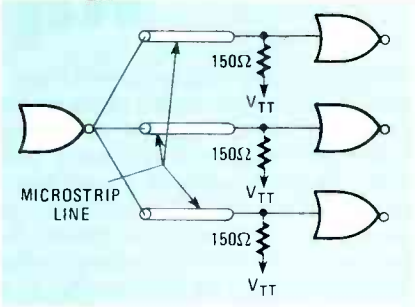

FIG. 8—THIS METHOD of PARALLEL termination uses multiple microstrip lines. The total impedance seen by the gate is 50 ohms.

Series terminations

The alternative to parallel terminations is series termination. Series termination is achieved by inserting a resistor in series with the transmission line, as shown in Fig. 9. The value of the series resistor is equal to the impedance of the line, less the output impedance of the gate. The typical output impedance of an ECL gate is 7 ohms; therefore, the proper series resistor for a 50-ohm line is 43 ohms. By placing the resistor in series with the line at the input, only half the voltage swing is transferred down the transmission line. When the signal reaches the end, however,

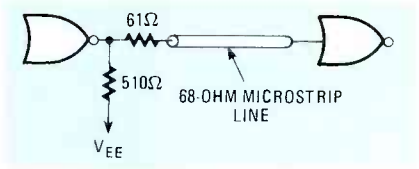

FIG. 9—SERIES TERMINATION is an alternative to parallel termination.

high-frequency reflections bouncing back and forth in the line combine to double the output voltage, thus re-establishing the original logic level.

To maintain clean wavefronts, though, the input impedance of the gate must be several times greater than the characteristic impedance of the transmission line. This requirement lends itself well to ECL circuits. Since the signal voltage is reinforced at the point of exit, it is possible to have more than one load on the output and still maintain proper voltage levels. However, the capacitance of the extra inputs has a greater effect on the rise and fall times than it does with parallel terminations. That is due in large part to the series resistor.

Some of the problem can be alleviated by decreasing the size of the series resistor, thus decreasing the R-C time constant. Unfortunately, less resistance means more ringing. Therefore, the series resistance must not go below the point where the ringing exceeds the limits imposed by the input. That approach is known as series damping, and a chart of the lowest acceptable resistor values can be found in Table 1.

A single load on a line doesn't present that problem. Therefore, it is better to run parallel lines to each input as shown in Fig. 10, instead of clustering them on one line. That is an excellent way to distribute a signal over a card without the increased power dissipation that's associated with multiple parallel-terminated lines. As before, the value of the series resistor for each line is equal to the impedance of the line.

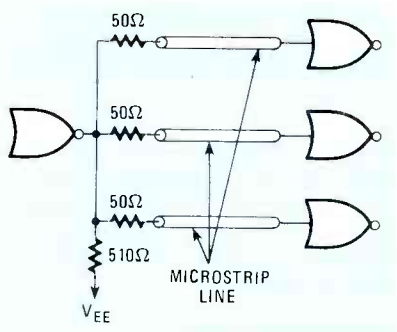

FIG. 10—SERIES TERMINATION using multiple microstrip lines helps to keep power dissipation down.

The size of the pulldown resistor, however, is affected by the number of lines the output must drive. If the value of the load resistor is too high, the output transistor will turn off during its transition from the high to the low state, creating a

TABLE 1—Minimum Series Resistance

Rise Time	Line impedance ohms	Series resistance ohms	Gate output impedance (ohms)
3.5 ns	50	9	15
"	68	18	"
"	75	21	"
"	82	25	"
"	90	29	"
"	100	34	"
"	120	43	"
"	140	53	"
"	160	63	"
"	180	72	"
1.1 ns	50	18	6
"	68	27	"
"	75	30	"
"	82	34	"
"	90	38	"
"	100	43	"
"	120	52	"
"	140	62	"
"	160	72	"
"	180	81	"

TABLE 2—MAXIMUM UNTERMINATED-LINE LENGTH

Rise Time	Line impedance (ohms)	FANOUT =	LENGTH (inches) 1	2	4	8
3.5 ns	50		8.3	7.5	6.7	5.7
"	68		7.0	6.2	5.0	4.0
"	75		6.9	5.9	4.6	3.6
"	82		6.6	5.7	4.2	3.3
"	90		6.5	5.4	3.9	3.0
"	100		6.3	5.1	3.6	2.6
2.0 ns	50		3.5	2.8	1.9	1.2
"	68		3.2	2.3	1.5	0.8
"	75		3.0	2.2	1.3	0.7
"	82		2.9	2.0	1.2	0.6
"	90		2.8	1.9	1.0	0.5
"	100		2.6	1.8	0.9	0.4
1.1ns	50		1.6	1.1	0.7	0.6
"	68		1.4	0.8	0.5	0.4
"	75		1.3	0.8	0.4	0.3
"	82		1.2	0.7	0.4	0.2
"	90		1.1	0.6	0.3	0.2
"	100		1.0	0.5	0.2	0.1

staircase effect on the fall time of the pulse. So as the number of lines increases, the load resistor must decrease. A 510-ohm resistor to V_{EE}, though, will easily drive up to four independent lines with no problem.

Unterminated lines

If the length of the transmission line—or wire, for that matter—is shorter than the wavelength of the input signal, the signal will pass through the conductor virtually unaffected by the reflections. Since many of the connections within a circuit are short and direct, they can be made with unterminated lines.

In a pulse circuit, the dominant frequency is determined not by the pulse repetition rate, but by the rise time of the pulse. The signal undershoot, which is the most critical of the two parameters, is held to about 15% if the travel time for a two-way trip through the conductor is less than the risetime. However, the propagation time through the line is determined by more than one factor. Involved are the length of the conductor, the dielectric constant of the board, the capacitance of the load, and impedance of the line. Those factors are often interrelated and variable, but Table 2 ties them together. With that table you can determine, at a glance, the longest unterminated line that you can use in a given situation.

As you can see, the shortest runs occur with those ECL IC's that have the fastest risetime. It is for that reason that a separate family of ECL IC's, the 10000 series, was developed. With deliberately slowed risetimes, they are able to take advantage of longer unterminated connections, thus easing circuit constraints. Unfortunately, their slower response time may not meet your system requirements in all cases.

System interconnections

In larger systems, more than one card is often involved. In that case, of course, connections between cards must be made. That presents a unique situation in that we must use all the transmission-line knowledge we have discussed so far. Furthermore, the parameters we discussed become more critical—and a new one comes to light.

This new parameter is attenuation. At the single-board level, attenuation is seldom a problem. But it must be taken into consideration when interconnections between modules and cabinets are made. Let's first take a look at the options open to us.

Although the mother-board arrangements can be used for tying cards together under special circumstances, it is better to use point-to-point wiring since few edge-connectors perform well at the frequencies involved. Single wires can be used if you respect their limitations. To begin with, they fall under the restrictions imposed by the rise-time versus lead-length rule. A practical example here would be a wire no more than 15-inches long, loaded with fewer than four gates. To prevent objectionable ringing, however, a ferrite bead must be placed at the end of the wire. To improve the signal somewhat, a 100- or 120-ohm resistor can be placed at the line ending and returned to the V_{TT} source. That resistance more or less matches the impedance of the line and thereby reduces some of the overshoot.

An open lead, unfortunately, is prone to pick up noise along the way, making it undesirable for many applications (particularly clocking pulses). A better approach is to make interconnections with coaxial cable. Not only does the

continued on page 90

Voice-Operated Switch for your Tape

WITH ONE HAND YOU DELICATELY BALance the tip of your scope probe on the tiny pin of an IC while your other hand fiddles with the scope's sync knob. Suddenly, the long-sought-after trace shows itself in all it's detailed glory! Now you can see the rise time, the overshoot, the pulse width, the DC level and all the other data needed to get the job done. But now what do you do? The notebook is across the room and you'll never remember all those precious numbers if you go get it.

You could dictate the readings to your secretary—if you had one. You could also lock the trace in your storage scope—if you had one. There has to be a better way. Let's see, your hands are full so you can't write—but you can talk! What's needed is cheap, hands-off recording gadget that would only record when spoken to, and would shut down during those long periods of utter silence when you are completely baffled by the peculiarities of your designs.

The answer is a voice-operated switch for a cassette tape recorder built from parts from your junkbox!

That is actually how the project came about. There was a real need to improve

ECL LOGIC CIRCUITS

continued from page 68

coax represent a well-defined impedance that is easily terminated, but it also provides good protection against cross-talk and noise.

There are several types of coax available for the job: RG58U, RG59U, etc. However, coax suffers from a noticeable attenuation of signal as the frequency increases. In other words, that type of connecting cable may not be suitable for all your interfacing, especially if the frequency is high and distances are long. The graph in Fig. 11 illustrates the point by showing you the maximum length of the coax as a function of input frequency

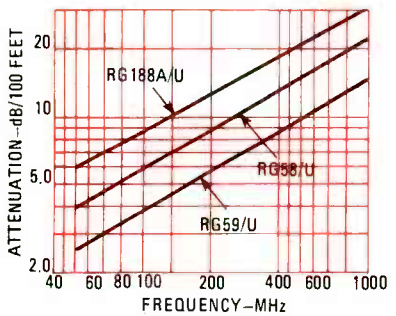

FIG. 11—ATTENUATION BECOMES A problem at high frequencies when using coaxial cable.

for three popular types.

Moreover, because of reactive loading, the fanout of a coaxial cable must be considered at high frequencies. For example, at 300 MHz it should be limited to no more than four. This is one of those situations where the logic tree comes in handy.

Unfortunately, both the open wire and the coaxial cable are afflicted by the shortcomings of a single-ended line. Things like ground loops, power-supply variations, and DC shifting from temperature differences must all be taken into account. Fortunately, there is another way to interface ECL IC's.

It will be easily understood if you first remember that an ECL gate is a differential amplifier. And because it is a differential amp, it has many of the desirable characteristics associated with differential design, including high common-mode rejection. As you recall, most ECL gates provide both OR and its complementary (NOR) output. Since the two outputs are always in mutual opposition, it presents the perfect opportunity to exploit the common-mode-rejection properties of an ECL IC. Making use of those properties allows us to connect two functions with nothing more than a twisted pair of wires.

The twisted pair is wired to both the OR and NOR outputs and connected to the input of an ECL line receiver. A line receiver is really nothing more than an ECL gate that has both inputs of the amplifier available to the user. Any noise that the twisted leads may pick up along the way will be induced in both wires equally; that is, the noise will have the same amplitude and polarity in both lines. This signal is then input into the line receiver and, as is the nature of differential inputs, the noise is cancelled out. That leaves us with only the digital information, which, of course, is what we desire.

Terminating twisted pairs

Thanks to differential design, twisted pairs provide the maximum noise immunity for any transmission line. As a result of this noise-free input, other parameters can be relaxed, including line terminations.

For reliable operation, the outputs of the driving gate must be terminated. The pull-down resistor is normally located right at the output pin, and more often returns to the −5.2-volt V_{EE} line, thus eliminating the additional V_{TT} supply, as we see in Fig. 12. You'll notice that both outputs are terminated similarly so that the driving source is balanced. Next, the twisted pair must be terminated at the receiving end. That is not a critical step, in contrast to the pains we took to assure proper termination of a single-ended transmission line.

The actual impedance of the line will vary depending on the wire gauge, insulation thickness and dielectric constant, and tightness of the twist. A 100-ohm resistor across the receiver inputs will usually be more than adequate. Any mismatch that may occur here is virtually ignored by the receiver.

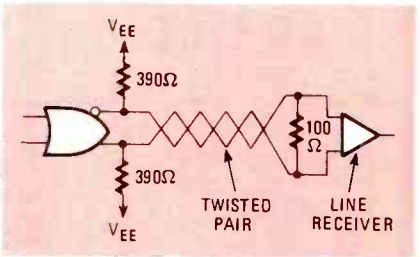

FIG. 12—A SIMPLE TWISTED PAIR connected to a line receiver can reduce common-mode noise.

And there you have it—a short course in ECL design. We must admit, though, that we have only touched on the subject. An interesting aspect of understanding microstrip theory and design, apart from its ECL applications, is that it is so applicable to many of the newer high-speed devices becoming available to the experimenter. Circuits like downlinks and ultraband communications rely almost exclusively on microstrip techniques, and are currently within the realm of practical experimentation. **R-E**

If the AP Products 8 page insert is missing circle 87 on free information card.