ECL LOGIC CIRCUITS

TJ BYERS

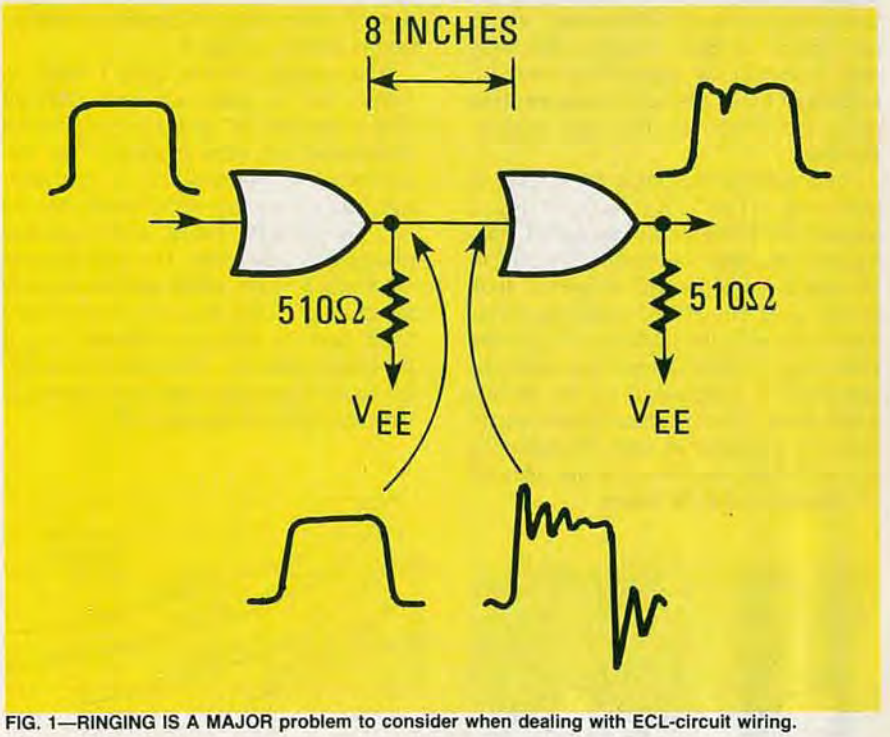

FIG. 1—RINGING IS A MAJOR problem to consider when dealing with ECL-circuit wiring.

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.

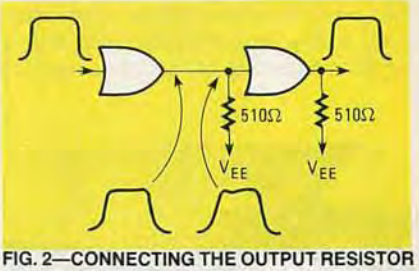

FIG. 2—CONNECTING THE OUTPUT RESISTOR at the end of the connecting lead can help to reduce ringing.

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.

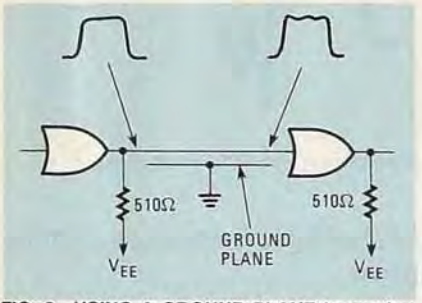

FIG. 3—USING A GROUND PLANE is another way to reduce overshoot and ringing.

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.

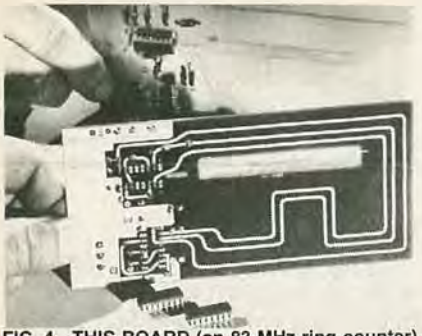

FIG. 4—THIS BOARD (an 83-MHz ring counter) uses 12-inch microstrip delay lines and a ground plane (seen in mirror).

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.

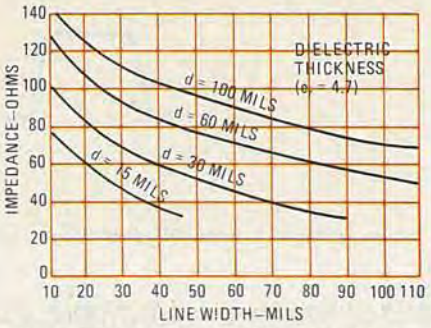

FIG. 5—THIS GRAPH SHOWS microstrip impedances for double-clad 1-ounce copper board; copper thickness = 0.0015 inches.

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.

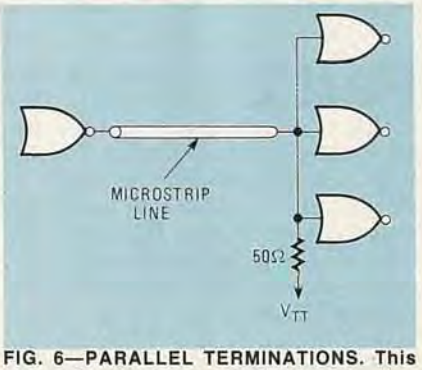

FIG. 6—PARALLEL TERMINATIONS. This approach uses only one load resistor.

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

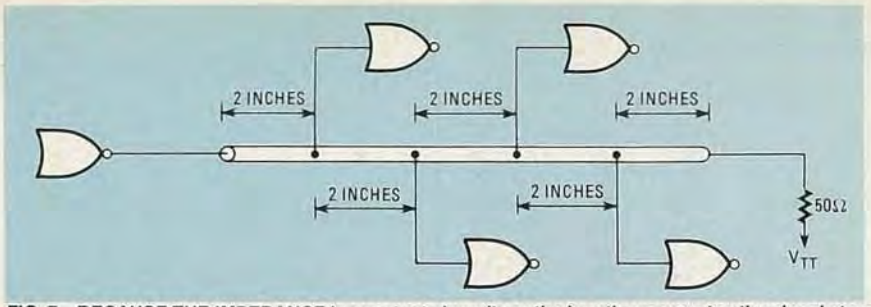

FIG. 7—BECAUSE THE IMPEDANCE is constant along its entire length, you can tap the signal at any point on the coaxial line.

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*.

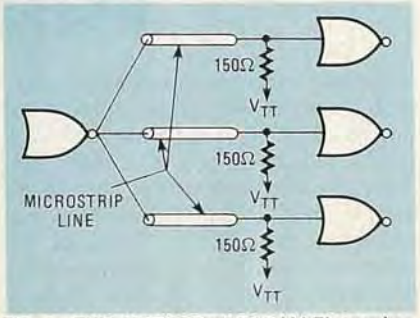

FIG. 8—THIS METHOD of PARALLEL termination uses multiple microstrip lines. The total impedance seen by the gate is 50 ohms.

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

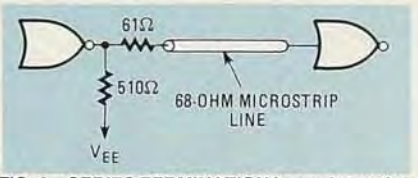

FIG. 9—SERIES TERMINATION is an alternative to parallel termination.

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.

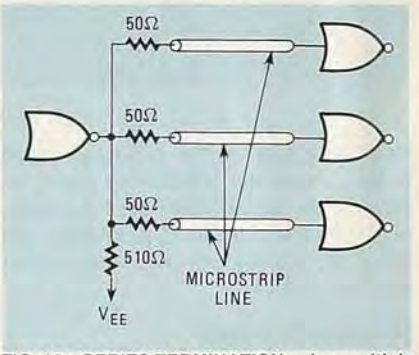

FIG. 10—SERIES TERMINATION using multiple microstrip lines helps to keep power dissipation down.

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

ECL LOGIC CIRCUITS

continued from page 68

coax represent a well-defined impedance that is easily terminated, but it also provides good protection against crosstalk 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

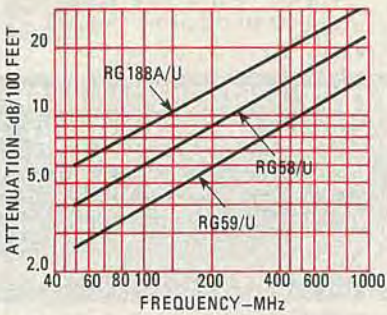

FIG. 11—ATTENUATION BECOMES A problem at high frequencies when using coaxial cable.

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.

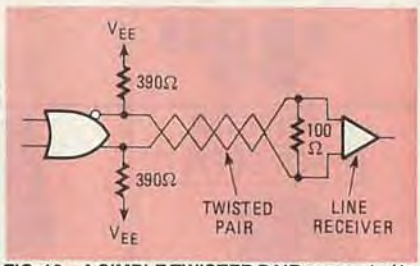

FIG. 12—A SIMPLE TWISTED PAIR connected to a line receiver can reduce common-mode noise.

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**