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.

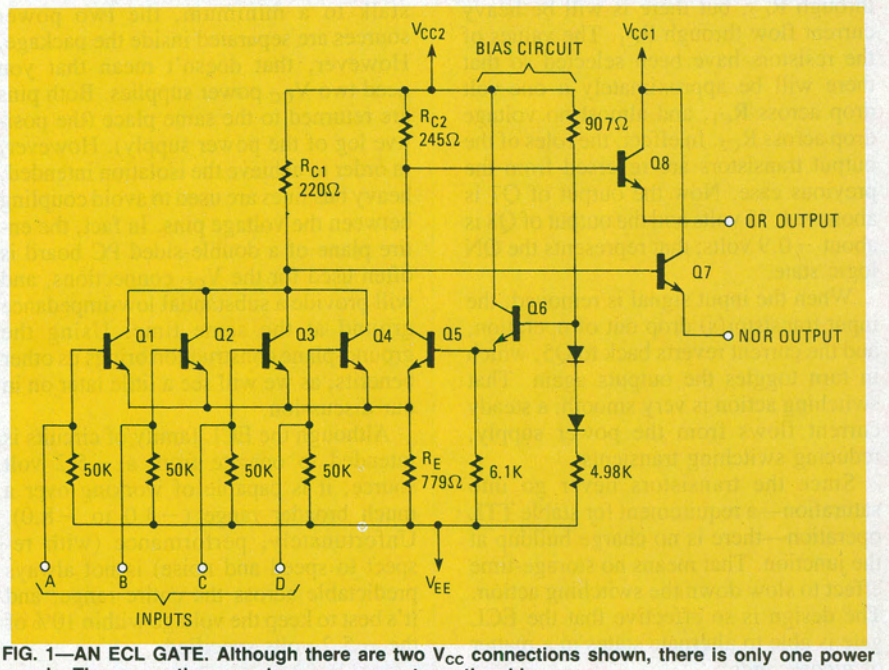

FIG. 1—AN ECL GATE. Although there are two V_{CC} connections shown, there is only one power supply. The connections are, however, separate on the chip.

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

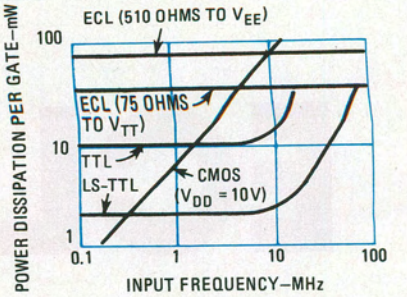

FIG. 2—POWER-DISSIPATION CURVES for various logic families shows that while ECL is often considered to be a power waster, at high frequencies it is quite miserly.

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-

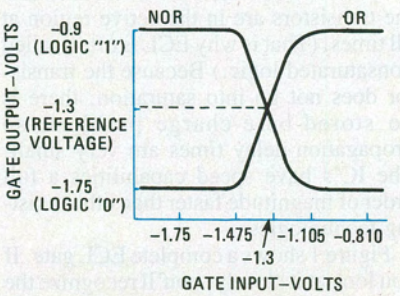

FIG. 3—TRANSFER CHARACTERISTICS of an ECL gate.

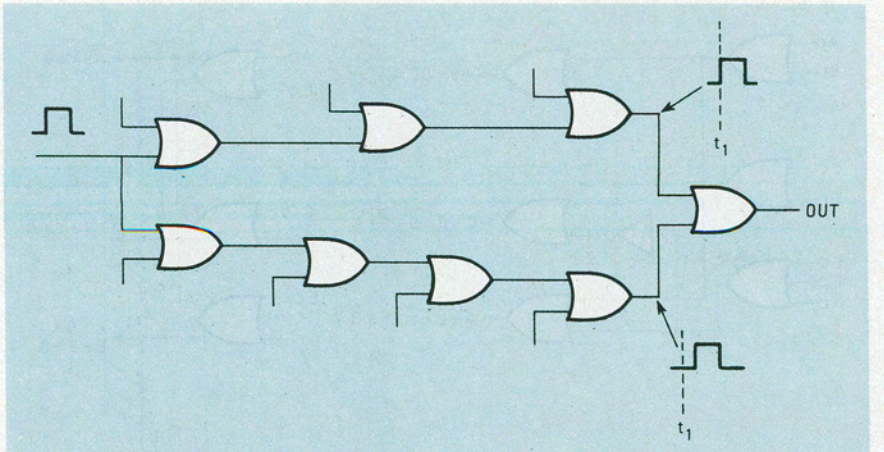

FIG. 4—THE EXTREMELY HIGH clocking speed used with ECL circuits demands that careful attention be paid to design.

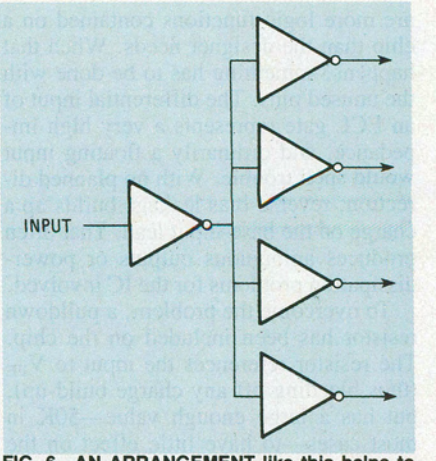

FIG. 6—AN ARRANGEMENT like this helps to minimize clock-skewing problems.

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!

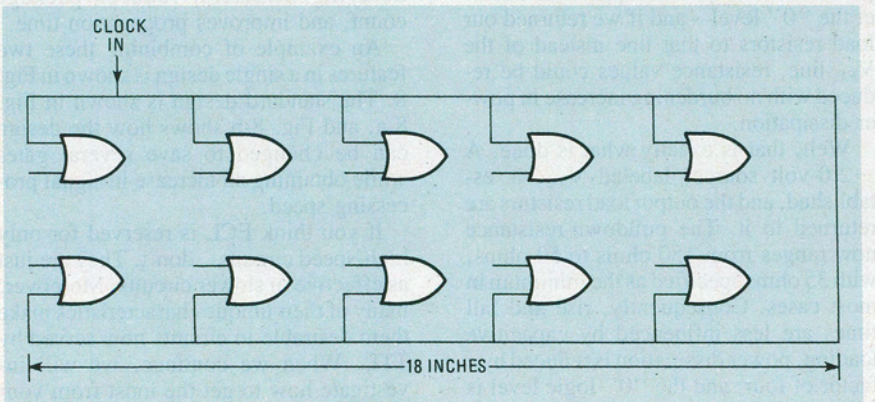

FIG. 5—LEAD LENGTHS become critical in ECL circuits.

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.

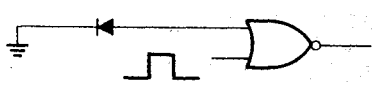

FIG. 7—A DIODE IS USED to imitate a "high" input for an ECL gate.

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.

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.

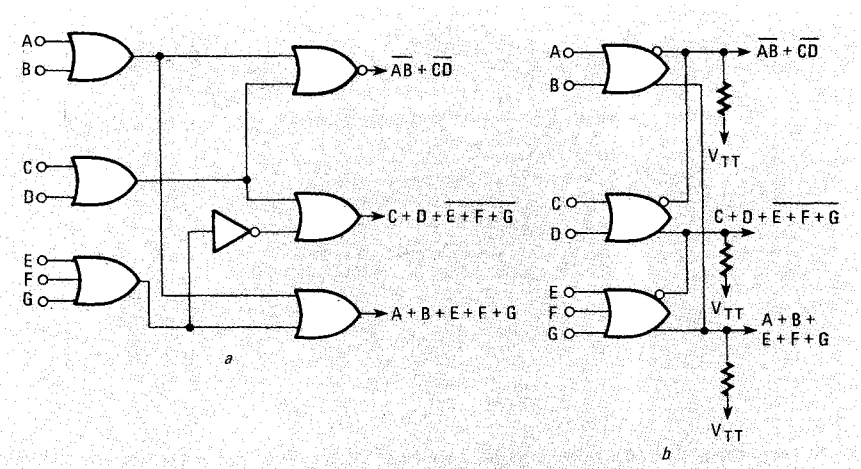

FIG. 8—THE NOR OUTPUT can be used to keep your designs more efficient, by reducing delay times and the number of gates required.

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**