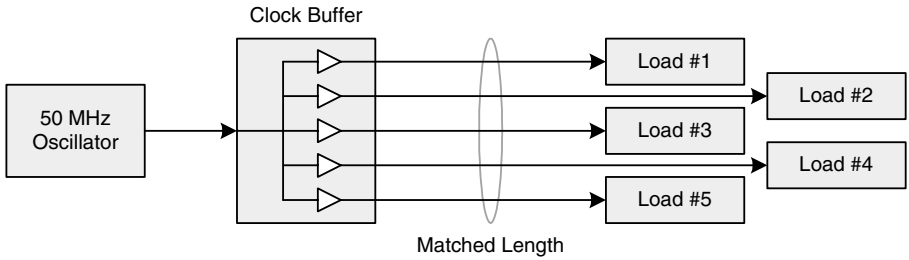

FIGURE 16.3 Single-buffer clock tree.

a single driver and a single load. It is assumed that each load is approximately the same as all other loads, which is usually the case when dealing with ICs. If one load is dramatically different from the others, it may cause its respective driver in the clock buffer to behave differently, and this could hinder the goal of creating a tree with balanced delays.

The absolute propagation delay from the oscillator to each load is immaterial in this case, because all that matters is that each load sees a clock with the same timing as the others. Therefore, the variables of interest are those that can potentially cause differential delays, or skew. The first such variable is the propagation delay variation from the buffer input to each output. Each output driver in an IC has slightly different physical properties that cause its propagation delay to vary within a narrow range. Low-skew clock buffers are specifically designed to minimize this variation. The acceptable amount of skew varies by application but generally declines with increasing frequency and with the use of faster devices, even at lower frequencies. Older low-skew buffers with pin-to-pin skews of better than 1 ns are still available for systems running at moderate frequencies. Newer buffers with skews of 200 to 50 ps are available for systems with very low skew budgets. Manufacturers of these buffers include Cypress Semiconductor, Integrated Circuit Systems, Integrated Device Technologies, and Texas Instruments. The second skew variable in this example is the variation in wire lengths connecting the buffer to each load. Wire delays are matched by matching wire lengths. Matching clock trace lengths when designing a PCB is one of the basics of proper clock distribution.

Signal integrity is tightly coupled with clock distribution concepts, because electrically clean clocks are a prerequisite for a digital system's reliable operation. Flops can be falsely triggered by nonmonotonic edges resulting from clocks that are improperly distributed in a way that causes their edges to become significantly distorted. Signal integrity is the topic of a later chapter, but a few basics are introduced here because of their relevance. The goal when distributing a clock is to prevent reflections on the transmission line that are formed by the wires that carry that clock signal. Reflections can result when a signal transitions rapidly relative to its propagation delay across the transmission line. It rises or falls so rapidly that one end of the wire is at a different voltage from the other end. Reflections distort a signal, because the instantaneous voltage on a wire is the sum of the driven signal and any reflected voltages that are present. A reflection occurs when a signal travels across an impedance discontinuity. Some energy passes through the discontinuity, and some is reflected back to its source. Terminating a transmission line with its characteristic impedance prevents reflections from occurring, because no impedance discontinuity exists.

One easy way of terminating a point-to-point wire that is always driven from one end is with *source termination*, also called *series termination*. This technique is illustrated in Fig. 16.4 as an enhancement to the single-buffer clock tree example. A resistor is placed at the driver and is sized such that its resistance plus that of the driver, Z_D, is approximately equal to the transmission line's characteristic impedance, Z_O. For this example, it is assumed that a transmission line with $Z_O = 50\ \Omega$ is implemented on a PCB and that the clock buffer's driver has $Z_D \approx 10\ \Omega$. Therefore, the series resistors

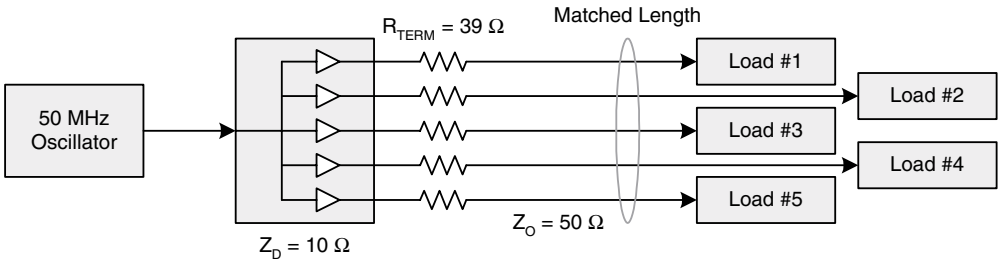

FIGURE 16.4 Series-terminated clock distribution.

are selected to be 39 Ω, the closest standard 5 percent value to $Z_O - Z_D$. In reality, it is difficult to precisely determine an output driver's impedance, because it is subject to variation with temperature and supply voltage. Trial and error may be necessary to select the best series termination value for a given clock buffer. Values between 39 and 47 Ω for 50 Ω for transmission lines are common.

Series termination does not prevent a reflection at the load, because the load is not terminated. The clock signal travels down a 50-Ω transmission line and then hits the high-impedance end load. The reflected energy travels back toward the source across the 50-Ω transmission line and then sees the 50-Ω load of the combined terminating resistor and the driver impedance. Therefore, the reflection is properly terminated and does not reflect back to the load and disrupt the clock signal. For series termination to work properly, the resistor should be placed as close as possible to the driver so that the resistances appear as a single lumped element rather than as two separate loads at high frequency.

Note that the wire between the oscillator and the buffer is not terminated. This is because the oscillator is usually placed very close to the buffer such that the unterminated transmission line is too short to cause reflections of significant amplitude. There are varying opinions on how long a transmission line can be before reflections become a problem. It is related to the signal's rise time. Many engineers use the rule that if a transmission line's delay is greater than one-fourth of the signal's rise time, it should be terminated. Electrical signals travel through a wire at approximately 6 in (0.15 m) per nanosecond. If a clock oscillator has a rise time of 1 ns, the wire between the oscillator and buffer should be no more than 1.5 in (3.8 cm) long.

Clock buffers have a number of outputs, and situations arise when one extra clock signal is needed. Rather than having to use a second clock buffer IC just for one more output, one pin can often drive two loads as shown in Fig. 16.5. A risk of this approach is that the more heavily loaded out-

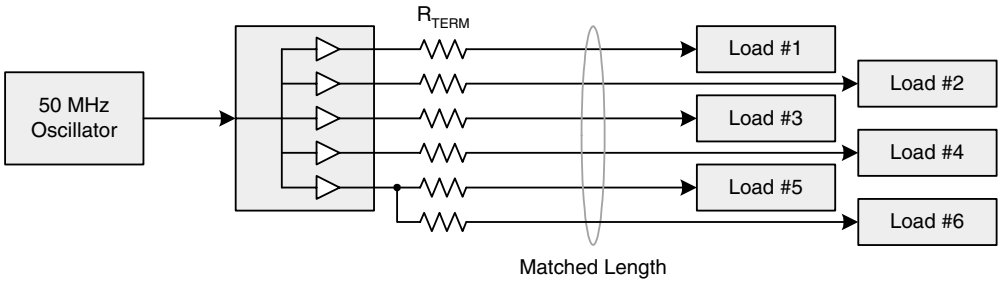

FIGURE 16.5 Single output driving two loads.