the expansion board, more pins or wires must be used to make the required connections. The increased connectivity increases cost yet does not improve the design's flexibility and tolerance of future changes.

A clock's purpose in a system is to provide a regular timing reference. Clock signals do not communicate information, because they are repetitive, and no single clock cycle has unique properties to set it apart from others. Therefore, a clock can be arbitrarily delayed without losing or gaining information. As with sine waves, periodic signals complete one full cycle over 360°. If a clock is delayed by an integer multiple of 360°, its delay cannot be detected, because each cycle is identical to all others, and its transition edges still occur at the same relative points in time. This important principle enables solutions to the clock distribution problems just posed and others like them. While it is physically impossible to construct a clock distribution circuit with zero delay, the end result of zero delay can be replicated by purposely adding more delay until an integer multiple of a 360° delay has been achieved.

Without a special trick, it would be extremely difficult to adjust a circuit for a perfect 360xN-degree delay. Propagation delays through semiconductors and wires have too much variation for this to be practical. The secret to achieving near perfect delays is with a closed-loop control system known as a *phase-locked loop*, or PLL. A PLL continually monitors a feedback clock signal and compares it against a reference clock to determine phase, or delay, errors. As soon as a phase shift is detected, the PLL can respond and compensate for that delay. This continuous compensation makes a PLL exceptionally able to deal with varying delays due to time, temperature, and voltage. PLLs are analog circuits and were originally invented for use in radio modulators and receivers. FM radios, televisions, and cellular telephones would not be possible without PLLs.

Figure 16.9 shows the structure of a generic PLL. A PLL is built around a *voltage controlled oscillator*, or VCO. The oscillating frequency of a VCO is proportional to its analog control voltage input. A phase detector drives the VCO's control input through a lowpass filter. When the reference and feedback clock edges occur at the same time, the input and output clocks are aligned by 360xN degrees, and the phase detector emits a neutral error signal, enabling the VCO to maintain its frequency. If the VCO begins to wander, as all oscillators do to a certain extent, the feedback clock edges begin to shift relative to the reference. As soon as this happens, the phase detector generates an error signal, which causes the VCO to slightly increase or decrease its frequency to move the feedback clock edge back into phase with the reference. One of the beauties of a PLL is that it can use an inexpensive VCO, because the feedback loop keeps the VCO at the correct frequency. A phase detector responds very quickly to phase errors, and the potential exists for an unstable control system resulting from rapid overcompensation. The lowpass filter stabilizes the control and feedback loop to average the phase detector output so that the VCO sees a more gradual control slew as compared to rapid swings.

PLLs are analog circuits and their design is a non-trivial task. Fortunately, numerous companies manufacture special-purpose digital PLLs that are designed for clock distribution applications. All

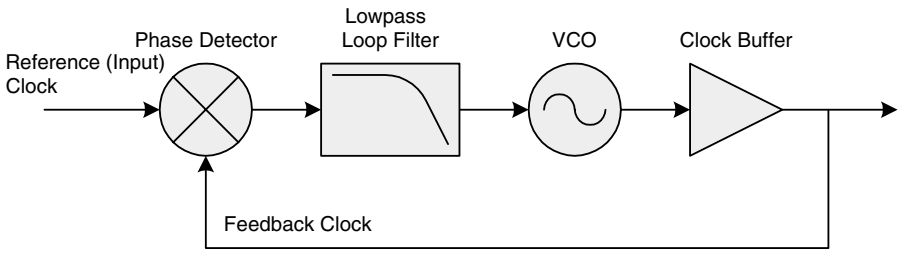

FIGURE 16.9 Generic phase-locked loop.

of the loop stability problems have been taken care of on a single IC, allowing engineers to use an integrated PLL as a tool and focus on the task of zero-delay clock distribution.

A zero-delay clock distribution circuit contains a PLL whose feedback path is chosen such that there is a nearly perfect phase alignment between the PLL reference clock and the clock loads. Figure 16.10 shows a general zero-delay clock distribution scenario. The feedback path has a low-skew relationship with the other clock loads. It is driven by the same low-skew buffer and is purposely routed on a wire that is length matched with the other clock loads. The result is that the feedback clock arrives back at the PLL at the same time that the other clocks reach their loads. Of course, a low-skew buffer has finite skew, and this skew limits the PLL's ability to align the output clocks to the reference clock. PLL feedback clocks are typically carried on wires that have been artificially lengthened by serpentine routing, because the PLL and buffer are usually close together while the clock loads are further away. It is actually the norm for a digital PLL to be integrated onto the same IC as a low-skew buffer, because the two functions are inextricably linked in digital clock distribution applications. Companies that manufacture low-skew buffers commonly offer integrated PLLs as well, often under the term *zero-delay buffer*.

Turning back to the microprocessor clock distribution example in Fig. 16.7, it is now apparent that a PLL with an integrated low-skew buffer would solve the problem. The microprocessor clock would drive the PLL reference, and matched-length wires emanating from the buffer would go to each clock load as well as the PLL feedback input. The resultant skew between the microprocessor's bus clock and the clock seen at each load is the sum of the propagation delay through the wire that connects the PLL to the microprocessor, the propagation delay through any small wire length mismatches, and the skew inherent in the PLL and buffer circuits. At an approximate signal conduction velocity of 6 in (0.15 m) per nanosecond, the skew due to wiring propagation delay is minimal if the PLL is located near the microprocessor and the output wires are matched reasonably well. As with any IC, the skew specifications for a zero-delay buffer should be compared against the system's requirements to select a suitable device.

The expansion card clock distribution problem in Fig. 16.8 can also be addressed using a zero-delay buffer as shown in Fig. 16.11. A normal low-skew buffer can be used on the base board, because the master clock is an oscillator with no skew requirement between that oscillator and the distributed clocks. All outputs from this buffer are length matched—including the wire going to the expansion board. As soon as the clock passes across the expansion connector, it drives a zero-delay buffer that has matched length outputs going to the expansion ICs as well as a feedback path back to the PLL. This ensures that the expansion ICs observe the same clock that was delivered across the connector. Attempts can be made to reduce skew introduced by the connector and routing to the zero-delay buffer by shortening the wire on the base board that drives the expansion connector. The wire should be shortened by the approximate distance represented by the connector and the short wire length on

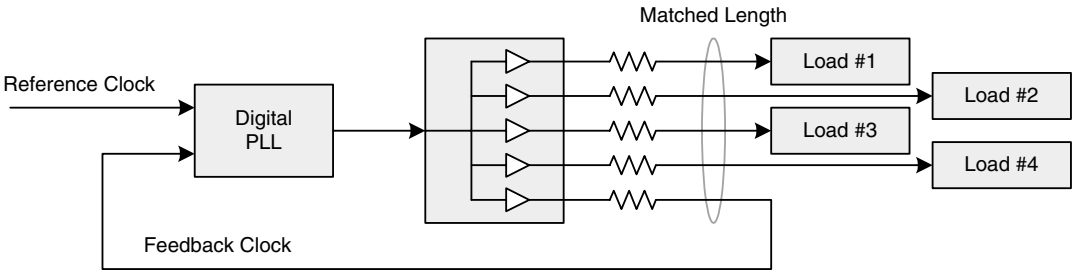

FIGURE 16.10 Zero-delay clock distribution.