Tolerance Mechanisms in Clock Distribution Networks

As described in the accompanying article, we are attempting to guard against a number of statistical tolerancing mechanisms, such as skew and jitter, that reduce the precision with which a clock signal can be delivered. Here we present an overview of these mechanisms.[1] For the purpose of considering system timing issues, it is useful to separate the system state architecture into a timing environment and a computation environment (see Fig. 1). The boundary between these two parts of the system is composed of the system state devices. Except for segment delay times and communications locality, we don't address the details of the computation environment here. The timing environment can be further broken down into three sections: the clock or phase generator, the clock distribution network, and the memory elements.

The clock generator supplies the signal whose edges eventually dictate when switching occurs throughout the system. The clock generator determines the period, pulse width, number of phases, and relative phase separation of the clock waveform. The primary attributes of the generator to be specified at design time are the waveform period and stability or jitter. For systems that use a processor chip, the period is usually specified by the manufacturer of the processor. Instability (jitter) in the waveform emerging from the generator detracts from either performance or reliability. Beyond these, there are frequently secondary issues and features that contribute to system testability—frequency and duty cycle adjustability, overtone suppression, modes (burst, single-step, fast, and slow), scan-path drive and timing, and others.

The state devices are flip-flops, latches, or memory devices of some type. New devices with enhanced testability features are appearing more frequently. The state devices play an important role in determining the low-level timing constraints in that their setup, hold, and minimum pulse width requirements must be satisfied at full clock speed.

The clock distribution network is a network of buffers and interconnects that conveys the clock signal to the clock consumers. It is responsible for fanout amplification and is generally tree-structured. In simpler systems, all of the fanout can occur in a single buffer. In larger systems, thousands of copies of the clock can be produced, requiring many levels of buffering (12 to 15 levels in some supercomputers). From a timing perspective, the ideal situation is for all of the copies of the clock waveform to emerge from the leaves of the clock distribution network at the same moment. However, the devices (both buffers and interconnects) that make

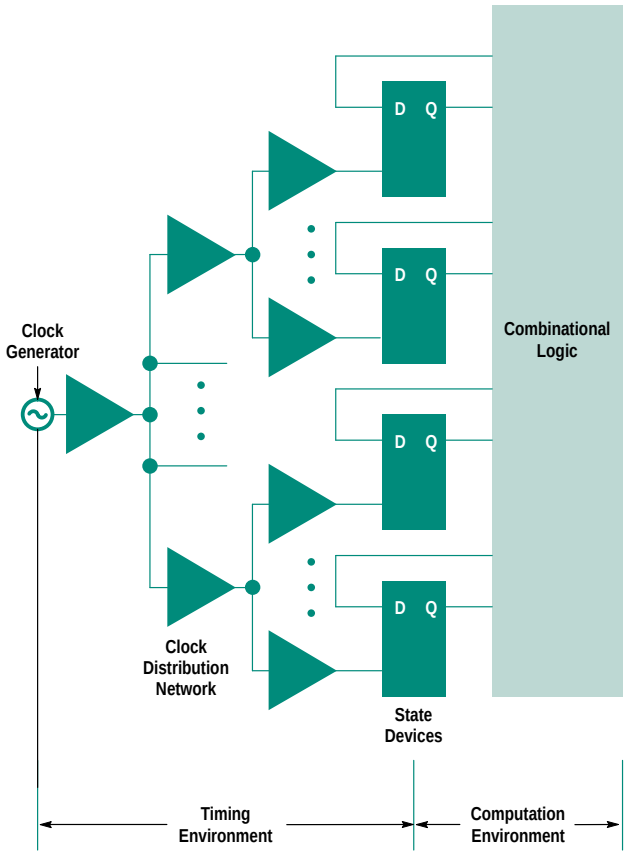

Fig. 1. State architecture model. Any synchronous digital system can be decomposed into a timing environment and a compute environment. The design issues specific to the timing environment are becoming critical in PC and workstation designs.

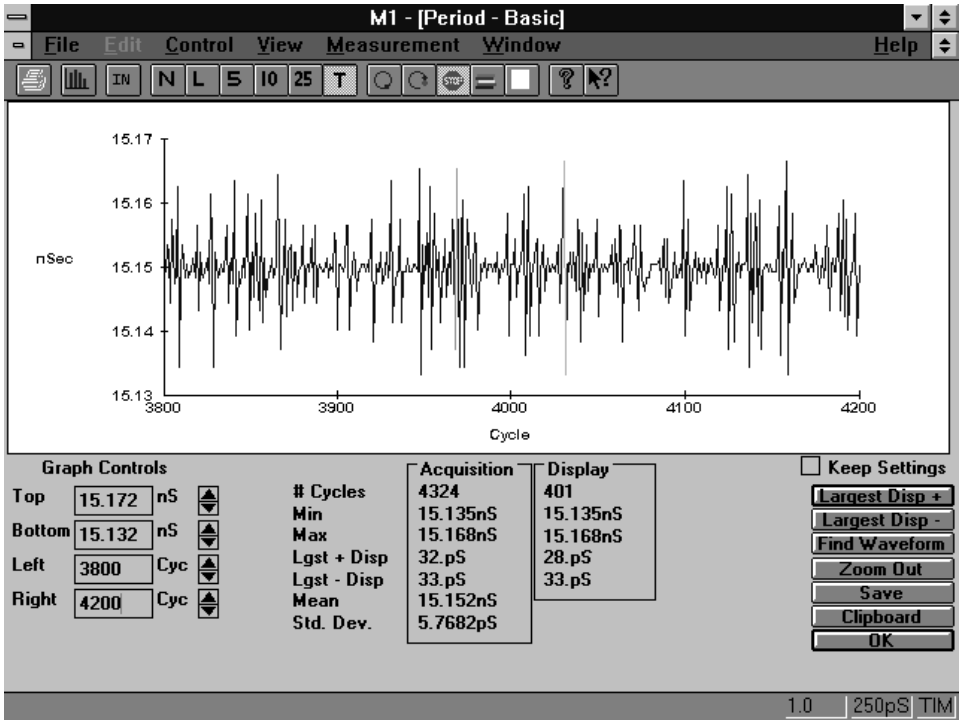

Fig. 2. There are a number of metrics of jitter. This measurement shows the cycle-to-cycle variation in the period of a 66-MHz clock. This was made using the Amherst Systems Associates M1 time interval measurement software, which analyzes digitized waveform data from an HP 54720D oscilloscope for jitter in a variety of ways.

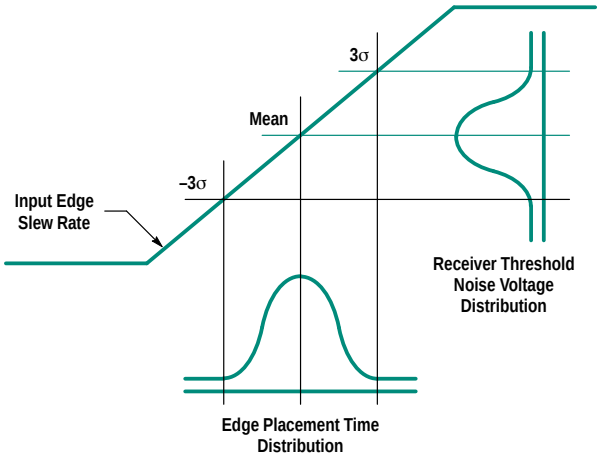

Fig. 3. Jitter, as it occurs in clock buffers, is generally the result of noise in the power environment (return currents, image currents, etc.) modulating the switching threshold of the buffer.

up the paths through the clock distribution network have a statistically distributed delay. These distributions can be time-invariant (static) or time-variant (dynamic). An example of a statically distributed tolerance is skew in clock buffers. This is the variation in delay either from pin to pin in a single package or from part to part. Interconnects can also exhibit tolerancing. This is most easily thought of as a variation in the propagation rate of several picoseconds per inch (10 to 40 ps/in). Interconnect tolerancing is frequently a source of unanticipated timing failures.

An example of a dynamically distributed tolerance is jitter. The placement in time of a waveform edge that has jitter varies from one cycle to another. It can be thought of as having a period that changes from one cycle to the next. Fig. 2 shows an example of this variation. Jitter can be added to the clock waveform in two places: at the generator or in the buffers. At the generator, jitter can occur through either internal noise or dynamic temperature or supply voltage instabilities. Jitter added in the clock buffers is caused primarily by noise in the power environment (return and image currents in power planes sweeping past power and ground pins, etc.) causing time-varying shifts in the device's switching threshold. This is illustrated in Fig. 3. Note that jitter (an expansion of the distribution of the edge placement) is increased when the noise voltage is increased or the edge rate of the signal arriving at the buffer is decreased. The management of jitter at consistent and acceptably low levels is perhaps the single greatest challenge for designers of systems that incorporate many of the new high-performance processors. A more in-depth discussion of jitter measurement and management can be found in references 1 and 2.

There are also statistical variations in how two identical parts are used. For example, one system may run a little warmer than another, another may have a little more noise in the power environment, and so on. Some of these tolerances are time-variant and some are not. As shown in Fig. 4, these device-level distributions can be statistically combined† to give a system-level distribution on the path delays in the clock distribution network. This system-level path delay distribution has a mean value that is sometimes called the nominal delay. By statistically combining the individual nominal delays along the path, one computes the nominal delay for that path.

When using the nominal delays, it is important to keep in mind that there is actually a delay distribution. This means that even if every path in the design is specified to be identical, when the product is manufactured there will be product-to-product variations in the propagation delay of any given path, there will be path-to-path variations within any given machine, and there will be cycle-to-cycle variations on a given path in a given machine. The result is that one must design the system in a manner that both suitably minimizes these tolerances and consciously considers the fact that the tolerances will always be nonzero. The design is said to be statistically stable when it has this characteristic.

When the tolerances in the system accumulate beyond the value anticipated by the designer, the design is said to be statistically unstable. In statistically unstable designs, some small fraction of the manufactured systems will experience timing failures despite the absence of any physical defects. In these systems, the clock can arrive at times other than the designer anticipated, and this can mean that one or more of the state device timing requirements (setup time, hold time, or minimum pulse width) will be violated.

Violations of any of the device-level timing requirements can result in statistically unreliable switching at the state devices. This can cause unpredictable deviations in normal system-level behavior. These faults can be extremely difficult and time-consuming to isolate. In fact, the failure modes exhibited by systems with internal timing problems are easily among the most difficult to diagnose using conventional troubleshooting methods. It is frequently necessary to employ an analytic approach to find these faults in any sort of efficient manner. These failure modes include:

- Intermittent or nonrepeating
- Low frequency of occurrence (minutes through weeks)
- Migration of the symptom location through the system
- Hibernation (failures occur as device parameters change slightly with age)
- Statistical.

References

1. M.K. Williams, "Distortion and Tolerance Mechanisms in High-Speed Clock Delivery," *Proceedings of the 1993 Hewlett-Packard High-Speed Digital Symposium,* pp. 4-1 to 4-41. Also available as Application Note ASA 93-1 from Amherst Systems Associates.

2. M.K. Williams, "Design Trade-offs in High-Speed Clock Distribution and Reception," *Proceedings of the 1993 Hewlett-Packard High-Speed Digital Symposium,* pp. 6-1 to 6-34. Also available as Application Note ASA 93-2 from Amherst Systems Associates.

† The combination of these subordinate distributions is more complicated than direct addition. It must also take into account correlations that occur in such tree-structured circuits, and other related mechanisms called tracking effects.

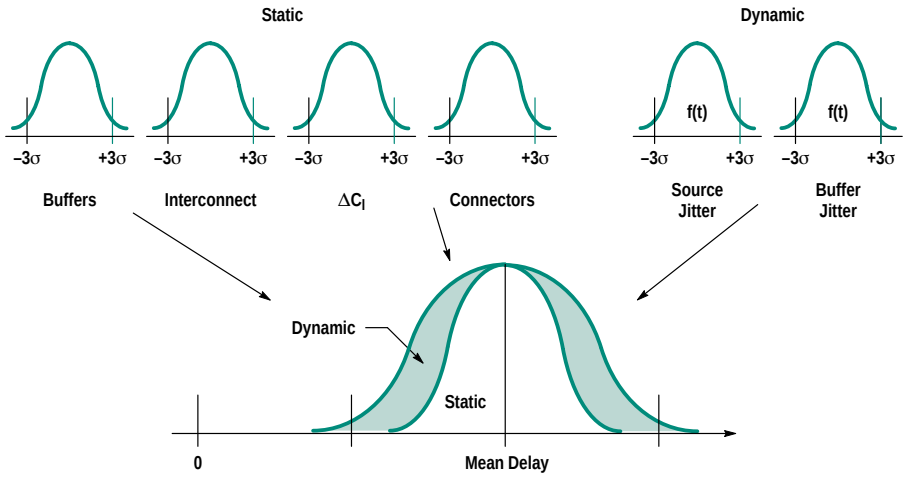

Fig. 4. A variety of tolerancing mechanisms contribute to the uncertainty in the arrival time of the clock edge at any clock load. Generally the only one that is available in catalogs or data sheets is the buffer tolerancing.