

HP Archive

This vintage Hewlett Packard document was preserved
and distributed by

[www. hparchive.com](http://www.hparchive.com)

Please visit us on the web !

Thanks to on-line curator: Istvan Novak

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Michael K. Williams

Amherst Systems Associates
P.O. Box 24
Amherst, Massachusetts
Tel: (413) 596-5354
Fax: (413) 596-5354

1993
High Speed Digital
Systems Design & Test
Symposium

Abstract

Examines the various tolerance issues which must be addressed in the design of clock distribution networks in high-speed digital systems. Clock distribution, skew, pulse distortion, marginal triggering/metastability, and other concerns are described in detail. It also discusses various design approaches, tools

available for characterizing device and system-level tolerancing, as well as methods and tools for isolating transient, anomalous behavior.

Author

Current Activities:

Mike is the Owner and Principal Consultant of Amherst Systems Associates — an engineering consulting firm specializing in timing-environment design/debug/analysis/instrumentation issues for high-performance digital systems. ASA also provides professional technical training in this and related areas to digital design groups, as well as technology transfer and applications-development services to other companies serving the high-speed digital design, and has served as a consultant to the computer, semi-conductor, and digital electronics design community since 1985.

Author Background:

Mike founded Amherst Systems Associates in 1985. From 1985 to 1989, he also served on the faculties of the University of Massachusetts at Amherst and National Technological University, where he performed research and taught courses in digital systems design. He was a designer for the VAX 8800 at Digital Equipment Corporation and prior to that started Williams Metal Testing — a metallurgical testing lab for jet-engine components. In another position with DEC, he designed test equipment and system controllers.

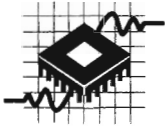
He holds a BSEE from Western New England College and an

MSECE from the University of Massachusetts, where he has also performed research toward a PhD. Research interests include timing-environment design in high-speed digital systems, the role of measurement and self-characterization in digital system timing verification, systems design methodology, and computer architecture.


Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #1


Distortion and Tolerance Mechanisms in High-Speed Clock Delivery



by
M. K. Williams



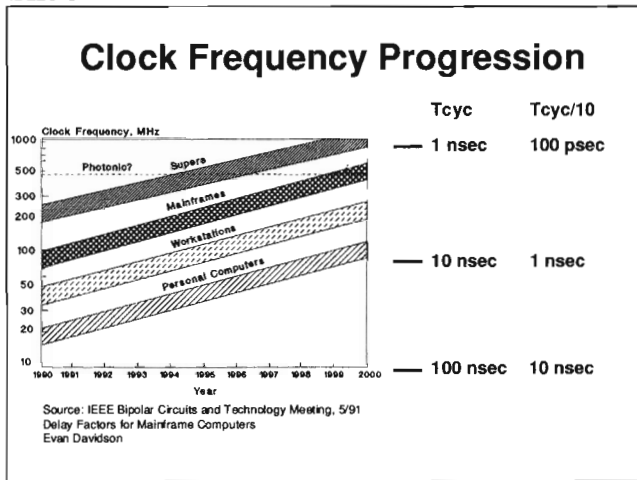
AMHERST SYSTEMS ASSOCIATES



HEWLETT
PACKARD

Rising system complexities and performance goals are constantly raising new concerns for digital system designers. The precision with which the edges of the system clock can be delivered has become an important factor in achieving competitive performance and reliability levels. The management of the various factors which impact this precision has emerged as a critical and challenging aspect of the design of high-performance digital systems. In this paper, we examine the underlying distortion and tolerance mechanisms which detract from that precision.

Slide #2



As the graph above shows, performance goals for systems of all types are rising exponentially. It is clear that most designers have or will come to a point where they must employ "high-speed design methods". One common definition for "high-speed" is from a signal-integrity perspective; it is the point where the analog effects within a circuit can no longer be ignored (Commonly accepted thresholds for this view of "high-speed" are clock speeds above 50 MHz and/or edge rates less than 2 ns).

An alternative view would be that you have crossed the high-speed threshold when system timing margins and device tolerances must be managed aggressively, and even represent an opportunity to gain a competitive advantage in your design. At this point, the careful and well-considered design of the system timing environment (TE) becomes an essential task in achieving the desired system performance and reliability levels. Our purpose in this paper is to describe the various mechanisms which reduce the precision with which the system clock can be delivered.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #3

Timing: Up-Front Effort PAYS OFF

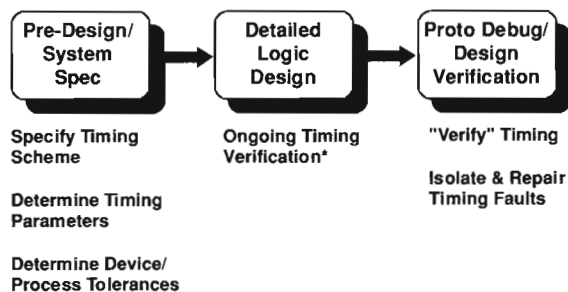
Benefits Increase with
Complexity & Speed

- **Ensure** reliable operation
- Achieve maximum potential performance
- Avoid **very** complex/expensive failure modes
- Reduce development complexity
- Life-cycle cost reduction
- Competitive advantage

Experience has repeatedly shown that knowledgeably addressing system timing issues up front (pre-design) yields a number of welcome benefits. Most importantly, one minimizes the risk of having to isolate and correct any of several very difficult failure modes described later in this paper. Beyond that, the timing environment can be a source of "free performance". That is, by employing methods which maximize the precision of the placement of the clock edges, you directly improve performance by reducing cycle time. And nowhere in the design cycle is the opportunity to usefully and intelligently apply the more sophisticated timing-environment design methods greater than in the system definition stage (that is, pre-design).

Slide #4

Timing Activities During the Design Process



During each phase of the system design cycle, there are important activities which pertain directly to some aspect of timing-environment design. Prior to beginning detailed, gate-level design, a number of high-level design decisions must be made. During this phase, which we will refer to as pre-design, senior members of the design team will specify all of the important structures in the system, establish performance and cost targets, and select the technology mix used to fabricate the system.

That information will in turn lead to an estimate of the total number of clock loads in the system, as well as the required clock rate and precision. From there, the timing-environment designer can specify a suitable timing scheme (cycle time, pulsewidth, number of clock phases, state-device type). In addition to the specification of the timing scheme, an important result of this stage is the determination of the timing tolerances of devices and processes used to fabricate the system. Those figures will be used by the timing verification software during the detailed design phase to check the timing of all of the data paths as they are designed. The importance of this step should not be overlooked as the accuracy of the timing verification is directly determined by the accuracy of the timing parameters used!

Once a prototype exists, the final timing activities can be carried out, as well as the other requisite proto-debug activities. Any gross timing errors, such as mismanaging skew may manifest as an inoperable prototype. Timing problems can be very

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

difficult to solve. The clock is distributed almost as widely as the power signals, and a systemic timing problem could produce numerous simultaneous failures which can be extremely difficult to diagnose.

As we will see later, subtle faults frequently have subtle symptoms (in other words, much less subtle than total failure). They have the potential for very low frequency of occurrence, migration of symptom locations, etc. and due to the statistical nature of device tolerances, may not even manifest in the prototype population at all. Given that, the prototype should not be regarded as a platform for formally verifying the correctness of system timing. It can tell you that you definitely have a problem, but it cannot tell you that you definitely do not.

In this paper, we will focus on the pre-design phase of the design process, and examine the distortion and tolerance mechanisms that detract from clock delivery precision.

Slide #5

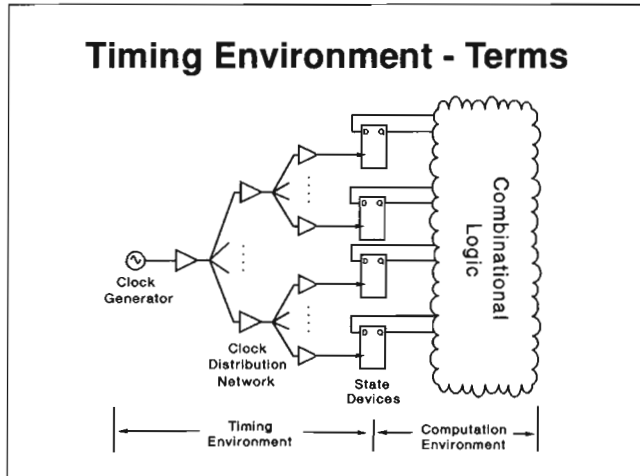
Outline

- • What is clock tolerance and why care?
- Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
- Strategies

It is important to first look at a typical timing environment, explain the basic terms and briefly explain what clock tolerance is and why it is important for system performance and operation. Next, each major source of clock tolerance and the mechanisms behind them are detailed. Finally, some strategies for controlling these effects are discussed.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #6



For the purpose of considering system timing issues, it is useful to separate the system state-architecture into a timing environment (TE) and a computation environment (CE). Note that the boundary between these two parts of the system is comprised of the system state-devices. Except for segment delay times and communications locally, we don't address the details of the CE in this paper.

The timing environment can be broken down into three parts: the clock/phase generator, the clock distribution network (CDN), and the memory elements.

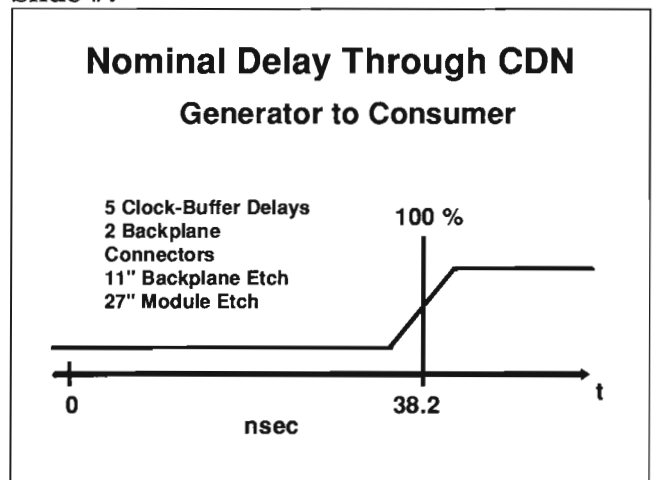
The clock generator supplies the signal whose edges dictate when things happen throughout the system. The generator determines the period, pulsewidth, number of phases and their relative edge placement. There are a large number of generator attributes to be specified in a typical design (for example, frequency stability or source jitter, frequency and duty-cycle adjustability, and overtone suppression) and frequently tradeoffs with system testability-enhancement features must be accommodated in the generator (burst/single-step/fast/slow modes, scan-path drive/timing, and so forth).

Contemporary state devices are either basic flops or latches, but new devices with enhanced testability features are appearing more frequently. The state-devices play an important role at the low-level in that their setup, hold, and minimum

pulsewidth requirements must be satisfied at full clock speed. For more complex timing schemes (such as, multiple phases) they also play very important roles, but those higher-level or structural timing considerations are outside the scope of this paper.

The CDN conveys the clock signal to the clock consumers. It is responsible for fanout amplification and is generally tree-structured for efficiency. The rest of this paper examines the impact of the CDN upon clock signal distortion and the impact of that distortion upon the state architecture. Be aware, however, that other timing effects must also be considered as systems get faster and larger.


Slide #7



All of the devices which comprise the paths of the CDN have a nominal (mean) delay. When you add the individual nominal delays along the path, you arrive at a mean delay for that path. For the circuit we will be using 38.2 ns.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery


Slide #8



Parts have statistical manufacturing tolerances. There are also statistical variations in how two nearly identical parts are used (in other words, one system runs a little warmer than another, another has a little more noise in the power environment). Some of these tolerances are time-variant and some are not. So 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, and path-to-path variations within any given machine, and cycle-to-cycle variations on a given path in a given machine. The result of this is that one must design his system in a manner that both suitably minimizes these tolerances and is considerate of the fact that the tolerances will always be non-zero.

Another subtle aspect conveyed by the diagram above is that statistics say that a small fraction of all the machines may have substantially more error in their path delays than the average path. That is, a path might be extremely faster or slower than nominal (or even 3-sigma). This general mechanism is called tolerance accumulation, and is the underlying mechanism by which many of the timing faults discussed in this paper (skew, SAG, and jitter) actually occur. When certain time tolerances of devices in the system accumulate beyond the value anticipated by the designer, the design is said to be statistically unstable. Despite the absence of any physical defects, some small fraction of the manufactured systems will experience timing failures.

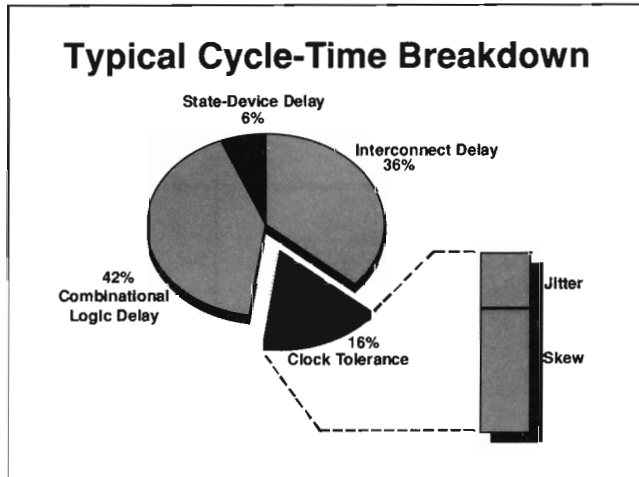
Slide #9



The diagram above shows two specific clock paths in a specific system that have ended up with different static delays (4 ns difference). Assuming that the dynamic tolerances of the two paths are uncorrelated with each other (safe assumption), we can expect a difference in the arrival time of the clock at these two points of 5 ns or more. Over the next couple of slides, we'll show two possible outcomes of such path to path differences—when you accurately predict the difference and when the difference is unexpectedly large.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #10



The smaller the cycle time is, the more results that a design computes per second. The lower limit on a design's cycle-time is determined primarily by the critical path delays. That includes combinational logic delay, delay through the upstream state-device, and all interconnect delays along the path (note that segment times are also distributed statistically). The cycle time, therefore, must be larger than the critical path delay (aka maximum segment time) plus a few other terms. For example, a simplified expression for the lower bound on a single-phase flip-flop based system is:

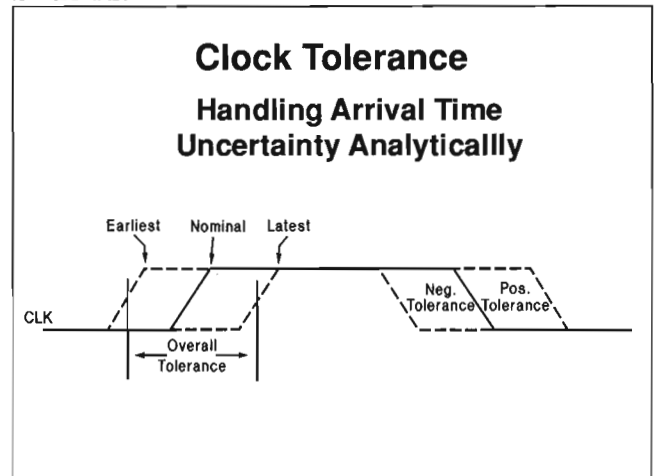
$$T_{cyc} > \text{max segment time} + \text{tolerance} + \text{setup time}$$

Therefore, to increase the performance of the design, you must minimize those terms which contribute to the cycle time. The setup time is usually small and there's nothing you can do about it one you've selected your system state devices. A major part of the detailed logic design phase is spent minimizing the critical path delays. Obviously by minimizing clock tolerance (both static and dynamic), you also increase performance. Clock tolerance represents unproductive delay and is as small as you design it to be.

The tolerance on a clock is generally less than about 15% of the cycle time, as shown above in the typical cycle time breakdown. Larger than 15% is usually uncompetitive. It is common to find designs that

have been carefully designed with less than 10% and some very aggressively-timed systems get down to 2 to 5%. There is an escalating cost (skilled engineering and manufacturing time) for each percentage point that you reduce that tolerance. Pulling 5% out of your tolerance is a lot harder in a 10 nsec cycle than it is in a 100 ns cycle. It's also much harder to get single-digit tolerances in a design with 300 board-level clock loads than in design which has 5 loads. Therefore, the competitive advantage of reducing the clock tolerance to a particular percentage trades off with the engineering costs of doing so. The resolution of that tradeoff is a common engineering decision in high speed systems.

Slide #11



From an analytical point of view, the earliest and latest possible clock edge arrivals are interesting events; the nominal arrival time is not. The computation of the working tolerances on the placement of the clock edges occurs at the pre-design stage. The principal use of these figures is during timing verification. Specifically, there is a separate tolerance computed for the leading and trailing edges, and they are usually different. For both multiclock and multiphase systems, there are separate working clock tolerances for each clock or phase.

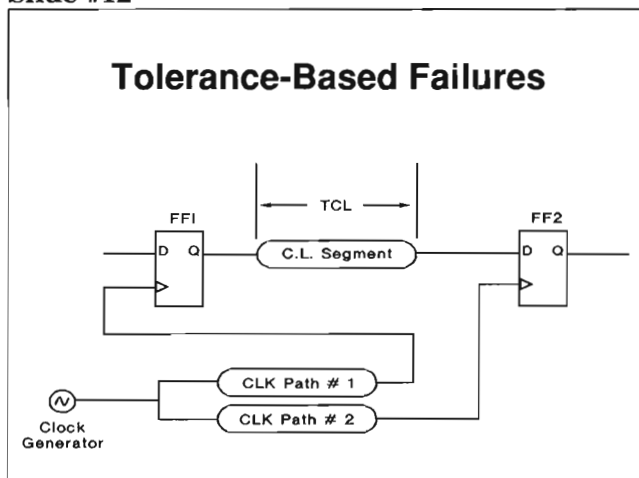
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

The distillation of the many complex statistical delay tolerances that comprise each clock path into just four figures per clock/phase (earliest and latest clock arrivals, leading and trailing edges) is an important design activity and should be allocated sufficient time and thought. This can be computed in a number of ways, including worst-case (catalog), statistical, measurement-based simulation, etc. Not all of these methods are appropriate for all designs, however, a comparative discussion of these methods is outside the scope of this paper.

An important point to keep in mind about timing verifiers, or any other simulation tool, is that they are accurate only up to the limits of the underlying representations of the components being simulated. The precision of these representations is the key to the precision of the computed result. A difference between the verifier's figures and real-world figures is not an uncommon source of prototype timing faults.

For any given clock signal in an actual system, the percentage displacement of the leading and trailing edges into their tolerance intervals will usually be similar, but not usually identical. You'll never see, for example, the leading edge appear at the beginning of its interval and the trailing edge at the end of its interval. The most typical case is that both edges appear with approximately the same displacements. Any significant difference from this is due to pulsewidth shrinkage/growth in the clock buffers, which is discussed later in the paper.

Slide #12



When tolerances are overestimated, we have seen that a performance penalty results. When they are underestimated, statistical failures are likely. Consider the circuit shown above as we illustrate such a failure.

Without changing the result, we can simplify the analysis by assuming ideal flip-flops ($T_s = T_h = T_{pd} = 0$) and wires ($T_{pd} = 0$). Then assume that the design clock period is just sufficient to permit the data arriving at FF2 to be stable and reliably captured if CLK arrives at both state devices precisely at its nominal time. FF1 and FF2 receive their clocks along different clock distribution paths. There are a number of components which populate these paths, and in any individual system manufactured, these paths would likely be built with different delays.

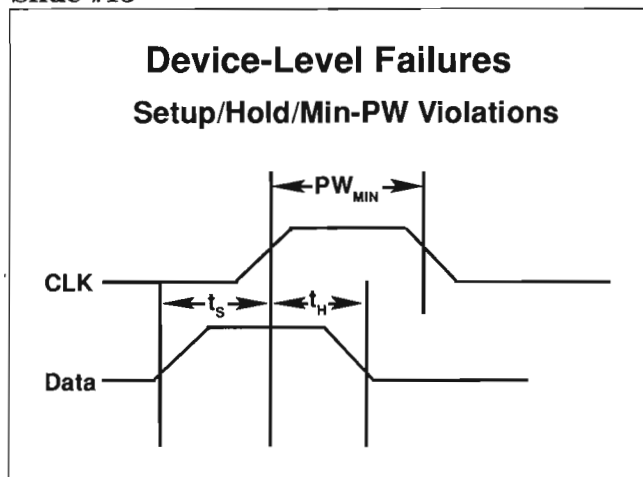
- If clock-path 1 is slower than nominal, or ...
 - If clock-path 2 is faster than nominal ...
- then FF2 will sample its input before the data wave emerges from the combinational logic segment → FAILURE.

If this example were run with non-ideal components, the result would be that the data is sampled either within the setup/hold aperture or after it, depending upon the severity of the tolerance. This would produce either intermittent or hard failures, respectively.

Keep in mind that the existence of dynamic tolerances (jitter) and data-dependent delays mean that the failure described above may not happen on every cycle. It could, in fact, happen extremely infrequently.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #13




At the device level, a timing-related failure occurs anytime any state-device in a system fails to present, by the end of its propagation interval, a stable (that is, non-transitioning and non-metastable) copy of the data it was supposed to capture. This definition does not require that this erroneous output manifest as incorrect behavior at the system level, such as during cycles when state-device outputs are ignored or not used in computations made during subsequent cycles. Instead, we will consider it to be a failure regardless of whether the condition is detected at the system-level or not. Examples of conditions which are defined as failures include:

Missed data—output in the opposite state

Unstable output—The output is still undergoing a transition at the end of the rated maximum propagation interval. This will be regarded as a failure, even if the device is transitioning to the proper state. This includes fully or partially metastable outputs. These failures can be brought about by missing the setup, hold, or minimum pulsewidth times, the presence of extra clock edges (signal-integrity problem), or a faulty state-device.

Slide #14



Indecisive switching, or metastability, is a type of aberrant behavior specific to state-devices. The measurements above show the Q-output of an ECL flop whose setup time has been intentionally violated. Note from the measurements that there are actually a variety of behaviors present. In general, one metastable trajectory will follow a different path through time-voltage space than another. The “hang-time” of a metastable trajectory is a random variable that decreases exponentially.

Metastability is the “marker” that is used to recognize state-device failure, either in a debug situation, or in the lab characterizing various state-device parameters. The differences from occurrence to occurrence make the behavior difficult to trigger upon. Time-qualified triggering can enable you to trigger on the runt-pulses that result when a trajectory starts up, and eventually resolves to a low, such as in the image on the right. In any case, since the behavior is not repetitive from cycle to cycle, it is necessary to use a real-time digitizing scope to see the behavior clearly. And the faster the scope update rate is, the higher the probability you will see any infrequently occurring behaviors. The appendix of this paper contains a much more detailed discussion of metastability and its detection/measurement.

Slide #15

System-Level Failure Modes

- Intermittence
- Low Frequency of Occurrence
- Migratory
- Hibernation
- Statistical

All of the device-level timing violations described earlier can manifest as deviations in normal system-level behavior. These 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 failures in any sort of efficient manner. These failure modes include:

Intermittent/non-repeating—Transient faults are difficult to diagnose since they are usually irreproducible. Systematically tracing aberrant behavior through a chain of devices is prevented when the failure can't be repeated. This characteristic can also make them difficult to capture on test equipment if you don't know the precise time, location and nature of the erroneous behavior.

Low frequency of occurrence—Timing problems have been known to occur at intervals far less than once per week.

Migratory—The location of the symptom of a timing failure can migrate around the system from one instantiation of the failure to the next. In conjunction with the two preceding failure modes, this can make a timing fault almost impossible to find using conventional debug methods in a practical amount of time.

Hibernation—Some faults may, in fact, not be present when a system is first manufactured. Instead, some may develop as device parameters change slightly with age, and manifest in some systems as an aggregate or chain of devices change together. An example of timing-oriented parametric change in a device over time is thermal incipient skew.

Statistical—Timing faults don't always manifest during proto debug. One reason is that prototype populations are typically much smaller than the production population. Given the statistical nature of many timing faults, they may occur infrequently enough to not manifest until well into the production cycle. One lesson to be learned from this is not to rely on proto debug to catch all timing faults. You have to do the analytical work too.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #16

Outline

- What is clock tolerance and why care?
- ➔ • Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
- Strategies

Distortion and Tolerance Mechanisms

As we stated earlier, the general mechanism by which statistical failures work their way into a system is tolerance accumulation along clock paths. There are four main distortion and tolerance mechanisms which affect the clock signal - signal integrity, skew, pulswidth shrinkage and growth, and jitter.

Slide #17

CDN Measurement Setup



The photo above shows two HP 8110 Pulse Generators, an HP 54720A Digitizing Oscilloscope, the HP 5372 Time-Interval Analyzer, and the Amherst Systems CDN Demonstration Fixture. This measurement setup will be used to make all of the measurements used in this paper.

The fixture provides a test bed for investigating and demonstrating the various timing-environment distortion and tolerance mechanisms. It is described further under the next slide.

The HP 8110 Pulse Generators were selected to drive the two phases of the CDN. They have excellent edge-placement precision (10 ps), low rms-jitter (10 ps) and adjustable edge rates. Since one of the two phases is used to simulate various types of operational noise behavior, the high-degree of flexibility in specifying the waveform (for example, pseudo-random bit-stream) was also desirable. Two are required, since we need two timebases for the jitter measurements.

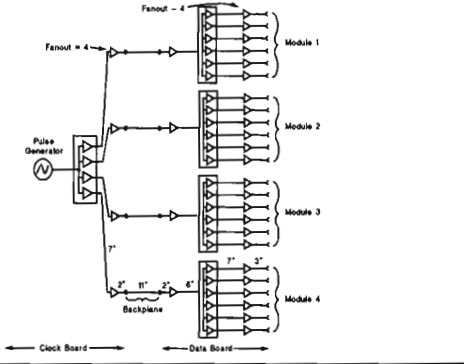
The HP 54720A scope configured for four channels was key in measuring skew and SAG.

The HP 5372A Time-Interval Analyzer permits jitter to be viewed in a manner not available with any other test instrument. It is an important tool when characterizing jitter (for example, determining the dynamic component of the working clock tolerance for the verifier), or when trying to locate the source of repetitive jitter.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #18

CDN Demonstration Circuit 2-Phase Fixture - One Phase Shown



Tolerance characterization is one of the most important timing-environment design activities of the pre-design stage. As part of the process of evaluating potential clock distribution schemes and verifying design decisions which impact clock distribution, it is common to build a technology board/system or test fixture. The fixture or system is used to measure CDN path delays and delay tolerances, and to spot any unanticipated signal integrity problems that can arise with a real CDN. In conjunction with statistical methods, these individual measurements can be used to project estimates of skew, pulsewidth shrinkage and growth, and jitter effects, such as failures per-thousand manufactured CDNs. The results of these measurements and their associated analysis are then used by the TE designer in making the final decisions about physically routing the clock through the system.

This test fixture for demonstrating CDN effects is an accurate representation of a clock distribution in a backplane-based, fast-CMOS cpu with two clock phases. Module sizes are specified to be identical to VXIbus C-size Eurocards, as is backplane spacing. The CDN has five buffering levels and total etch-length in all clock paths is controlled at 38". The clock buffers are MC74ACT241's, which is a common application for that part. They have catalog propagation delay ranges of 1.5 to 9.0 ns for both leading and trailing edges. The typical propagation delays are 6.5 nsec leading edge and 7.0 ns trailing edge. This difference in typical delays will be significant for the SAG demonstration.


Numerous measures were taken to ensure a minimum path-to-path variation in propagation delay. Radial distribution techniques are employed throughout to optimally balance all path delays. A separate copy of the clock signal is distributed to each module-slot on the backplane. The etch on the backplane is a controlled 11" from the central clock slot to all other module-slots. Loading at every level is identical across all paths.

All etch is run on the surface, has an AC-impedance of 50 ohms, employs lumped fanout, and is AC-terminated. Power and ground planes are employed, as are bulk and IC bypass capacitors. The power-supply voltage is 5.0 volts.

Several common signal-integrity faults (for example, insufficient number of ground pins on modules) were inadvertently built into the system during layout. These do not significantly affect the operation of the system when only a single clock phase is driven, but they do become significant when driven by two phases. The power environment noise resulting from these faults gives the system a higher susceptibility to jitter, which we show later in the paper.

Slide #19

Measurement Clock



Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #20

Outline


- What is clock tolerance and why care?
- Distortion and tolerance mechanisms
 - - Signal integrity problems
 - - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
- Strategies

Distortion due to signal integrity problems can be a major source of trouble at elevated operating and edge speeds. The effects of these problems upon timing include multiple or unintended triggering (that is, extra clock edges), destabilization of the data during the setup-hold interval, or delay of the clock arrival (distortion delay). The clock is the most important and widely distributed dynamic signal in the system and deserves the highest measure of consideration from a signal integrity perspective. Signal integrity measures for the clock are no different than they are for any other critical signal in the system, and are not the prime focus of this paper.

Slide #21

Skew vs Delay


Skew IS NOT the Same as Path Delay



This is a scope plot of the input clock waveform (upper trace), and three randomly selected clock signals at the output of the CDN. Real-time acquisition was selected so that dynamic path delay tolerances could be viewed simultaneously with static tolerances. The automatic measurements show that the mean delay through the CDN is different for all three paths. That is, the active edge of the clock passes through threshold at three distinctly different times. This is the fundamental definition skew. The standard deviation gives us a lower bound on the RMS jitter value. We will talk more about that later in the paper.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #22



This is a scope plot of four clocks emerging from the CDN. The two upper traces are the earliest and latest clock arrivals in the system. This is the global skew for this particular system. In this case the automatic measurements show this to be almost 4 ns. The lower two traces are clocks produced on the same module. This is referred to as the local skew. When there are several levels of packaging, or several levels of clock buffering on the current packaging level, there can be degrees of locality. This is sometimes referred to as the correlation of the skew.

Regarding oscilloscope selection for skew characterization work, the reader is strongly encouraged to use one with four channels. The author has gone through the process of evaluating the skew in large clock distribution networks a number of times using both two and four channel scopes. The four-channel approach has surprising productivity and accuracy advantages. Four channels lets you examine your local skew environment (for example, upstream and downstream flops of a critical path) in the context of the global skew environment (such as, the current earliest and latest arrivals if your sweeping across all of the clock nets in the system). With two channels, you end up keeping track of a lot of numbers on paper, slow and inaccurate.

Slide #23

Sources of Clock Skew

- Device-Based (Intrinsic)
- Interconnect-Based (Extrinsic)
- Structural/Design Variations
- Other Sources

The underlying causes of skew can be broadly broken down into three main types, as shown above.

Slide #24


Device-Based (Intrinsic)

- Manufacturing tolerances
- Propagation delay
- Gate threshold voltage
- Edge rates

One large contributor to the tolerances accumulated by the clock edge as it passes through the CDN are the manufacturing tolerances on the clock buffers. The tolerances are always non-zero and in some systems can add constructively to yield an arrival time much earlier or later than nominal (in other words, intoleranced). The device-based component of skew is occasionally referred to as intrinsic skew. Sources of intrinsic skew can be broken down into three principal types, as shown above.


Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #25




Buffer propagation-delay tolerances are a standard catalog rating for every buffer made. The figure illustrates the trailing-edge tolerance on the propagation delay of an inverting buffer. This tolerance, as well as most of the others in this section, can be specified separately for rising and falling output edges.

Slide #26



The gate threshold-voltage tolerance is a rating of how the input switching voltage can vary from one copy of a device to the next. The figure shows that this voltage tolerance also represents a time tolerance when input signals have non-zero (real) transition times.

Slide #27



The output pins of clock buffers also have a tolerance that affects timing. The diagram shows the fastest and slowest edges for a hypothetical clock buffer. That edge-rate tolerance also equates to a time tolerance, as shown in the figure.

There is some overlap among the three factors discussed above. Both edge-rate and threshold variations are contributors to propagation delay tolerances. However, there is another component of gate propagation delay variation unrelated to the other two effects. This other component is an internal gate propagation delay tolerance. Timing constraints for some aggressive systems may someday get so tight that it would be worthwhile to characterize these behaviors and match drivers and inputs that have an above average compatibility. For example, matching an output edge-rate to an input threshold voltage as a means of achieving balanced delays. "Handmade" CDN's have long existed in the form of clock-tuning and manually-characterized parts.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #28

Interconnect-Based (Extrinsic)

- Capacitive Loading Variation
- Propagation-Rate Variation
- Etch-Geometry Variation

The other large contributors to clock skew are the tolerances on the interconnect. This is sometimes called extrinsic skew. Sources of extrinsic skew can be broken down into three principal types shown above.

The interconnect component of the expression for critical-path delay in nearly any system is coming to dominate the expression. It pays to keep an eye on how much interconnect you have in your CDN during its development. A six-sigma tolerance of ± 25 ps/in on the interconnect hits you harder in a CDN with 40" paths than in one with 35" paths. The 40" CDN just has more opportunity to experience more parasitic encounters with nearby etch, vias, etch on adjacent layers, dielectric thickness variation.

Furthermore, as the dimensions of logic elements shrink, the ratio of interconnect delay to gate delay grows. The result is that the contribution of tolerances in the CDN attributable to extrinsic skew is a problem of increasing importance.

Slide #29

Capacitive Loading Variation

Path-to-Path Differences in the Capacitive Interactions Between the Clock Etch and:

- Adjacent Traces
- Nearby Vias
- Nearby IC Leads
- Signal/Power Planes

These are path-to-path differences in the capacitive interactions between the clock-etch and adjacent traces, nearby vias, nearby IC leads, and signal/power planes which can result in differences of signal risetimes. As we saw earlier, signals with different risetimes get to threshold at different times. Variations in the gate input capacitance, as well as path-to-path differences in the number of loads, are also included in this effect.

Slide #30

Propagation-Rate Variation

Path-to-Path/Board-to-Board Variation in:

- Dielectric Variation
- TxL-Geometry Variation
 - Micro-Strip: Approx 145 psec/in
 - Stripline: Approx 185 psec/in

These can occur due to manufacturing tolerances on various physical parameters that are determinants of signal propagation rate. Only two factors control the propagation rate of a signal on a conventional printed circuit board. One is the dielectric constant of the board material, ϵ_r , which is strictly a rating of the material, and the other is the geometry of the transmission line (for example, microstrip or stripline). Variations in the density or purity of the material which composes the board result in dielectric variations, which in turn result in propagation rate variations. These variations can be either board-to-board variations, where dielectric constants can vary 15% or so from lot to lot, or they can be variations in the dielectric constant across the surface of the same board, which will be much less.

Another aspect of this results when there is poor path-to-path control of the actual transmission line geometry. To move a signal anywhere around a board, it must change layers (geometry) to navigate around obstacles. The fundamental problem is that surface-etch/microstrip is faster (145 ps/in) than subsurface etch/stripline (185 ps/in). When a high-degree of extrinsic skew control is necessary, it is common to require that all clock-etch stay on the surface and require all other interconnect to navigate around it. This method can create a number of difficulties with regard to the control of radiated noise. When the tightest control of extrinsic skew is not required, a common method is to bury all of the clock etch in a dedicated layer of strip-line. A discussion of those trade-offs is beyond the scope of this paper. Another common approach is to precisely specify the number of inches each clock-path has on the surface and submerged. This requirement would normally be specified on a level-by-level basis.

Slide #31

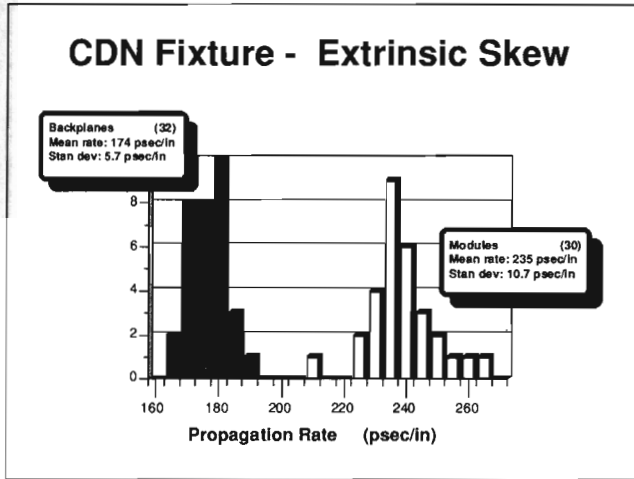
Etch-Geometry Variations

- Length/Position Tolerance & Variation
 - Time of Flight Differences
 - Frequency & Length-Dependent Attenuation -----> Clock Edge Degradation
- Thickness/Width Tolerance & Variation
 - Z0 variation -----> Clock Edge Degradation

In silicon, tolerances on wire thickness and width constitute the major source of skew. Variations in length, position, and thickness of board etch can also impact arrival time tolerances. Tolerances on interconnect length impact in two ways. One is that variations in path length means variations in time of flight. The other is that transmission lines of different lengths attenuate the high-end spectral content of the signal differently, since there is a frequency-dependent attenuation per unit length. Since the characteristic impedance of the transmission line is a function of its thickness, width, and dielectric thickness, any tolerances on those physical dimensions mean discontinuities in the characteristic impedance. That, in turn, means frequency-dependent internal reflections/dispersion, which means edge-rate degradation. Again, this delays the time the degraded edge takes to reach threshold. In some extremely fast systems, extrinsic skew on the backplane and backplane connectors is addressed by distributing the clock to individual modules in matched, high-quality coaxial cables. Of course, this represents a cost increase.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #32



The graph above shows the distributions gathered from a small number of measurements of the interconnect delays in the CDN fixture. All of the delay measurements were converted to propagation rate to permit easier comparison. The graph shows two distinct clusters — one for module etch (30 measurements) and another for backplane etch (32 measurements). The backplane etch has a mean propagation rate of 174 ps/in with a standard deviation of 5.7 ps/in. The module figures show a bigger spread at 235 ps/in and 10.7 ps/in, respectively. Six-sigma tolerances are 34.2 and 64.2 ps/in, respectively. The latter indicates poor manufacturing control for the module microstrip. The author regards anything above 40 ps/in as unacceptable for precise clock edge placement. A subsequent analysis of measured boards revealed poor control of path length (not all 7" paths were 7" long) as a primary contributor.

Despite careful measurement, the distribution reveals a fairly sloppy spread for this system. When these rates are projected along an entire clock path (27" module etch, 11" backplane etch), an unexpectedly high extrinsic skew results. The six-sigma limits for total module-etch path delay are 5.586 ns and 7.143 ns (mean of 6.537 with a 290 ps standard deviation). The six-sigma limits for the backplane are better, at 1.794 ns and 2.043 ns (mean of 1.908, SD of 62 ps). Full-path statistics can be easily computed from the separate module and backplane figures.

Slide #33

Structural/Design Variations

- Clock extraction from multiple levels of CDN
- Inconsistent use of inverting/non-inverting buffer outputs along all paths
- Inconsistent loading at each CDN buffer level
- Fanout schemes other than lumped fanout (point-to-point preferred)

In this case, we're talking about "designed-in" variations rather than manufacturing tolerances. By extracting the clock from only the leaves of the CDN, you have ensured that each clock edge has a similar set of "experiences". When copies of the clock are extracted arbitrarily from different levels, you create sets of clocks which are guaranteed to have experienced different delays. Of course, this technique can be used effectively to extract an "early" clock from the CDN to create a little more breathing room for an extra-long upstream segment. Clock systems where early and/or late copies are available in addition to the nominal are called multiclock (versus multiple-phase clocks).

In technologies such as ECL, which provide both true and complement outputs on buffers, the arbitrary use of both polarities to distribute the clock can be problematic. Specifically, when buffers have asymmetric leading and trailing edge propagation delays, care should be taken to ensure that all buffers at each level of the CDN use the same polarity output. In doing so, you guarantee the active clock edge experiences similar delays along every clock path.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #34

Other Sources of Skew

- Thermal difference
- Vcc difference
- State-device threshold variation

Given that device performance usually varies with temperature, and that thermal management is never perfectly consistent throughout the system, time-performance tolerances will exist even for manually sorted and matched buffers. This effect can be especially pronounced in mixed technology systems, where technology types are clustered. For example, in a system using both TTL and fast CMOS, any CMOS buffers placed in the vicinity of the TTL may operate differently than other CMOS buffers located in cooler parts of the system.


Slide #35

Outline

- What is clock tolerance and why care?
- Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - - Pulswidth shrinkage and growth (SAG)
 - Jitter
- Strategies

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #36



The measurement above shows several clock pulses emerging from the 5-level CDN fixture. The buffers in that system have a typical difference between their leading and trailing edge propagation delays of 500 ps. Note that for the four signals shown, all deviate considerably from the 4.00 ns pulsewidth generated by the clock generator. This effect is called pulsewidth shrinkage and growth (SAG).

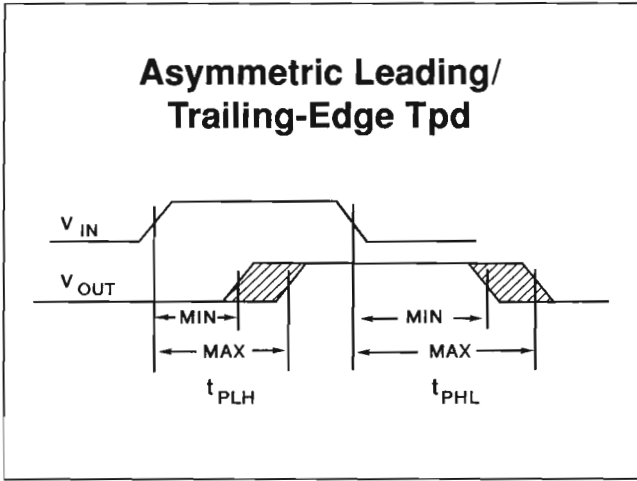
Slide #37

Sources of PW SAG

- Asymmetric leading/trailing-edge T_{pd}
- Asymmetric leading/trailing-edge transition times
- Asymmetric leading/trailing-edge V_{th}
- Interconnect bandwidth limits

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery


Slide #38



When there is a difference between the leading and trailing edge propagation delays as shown above, the active edge experiences delays differently than the inactive edge. For families with both true and complement outputs, it's a good idea to alternate output polarity from level to level in the CDN if this asymmetry exists.

Asymmetric leading/trailing-edge transition times can also change the pulsewidth. The time to transition from a low to a high can be different than the time from low to high. This will cause unequal delays in driving the downstream logic to threshold, resulting in a change in the pulsewidth.


Slide #39



Some device inputs are constructed with input hysteresis to enhance noise immunity (for example, Schmitt-trigger inputs). The figure shows there can be a range of thresholds (such as, a tolerance) for transitions in each direction. If the width of these bands are different, or if they operate at different voltage displacements from the beginning of their particular edge, the leading and trailing edge delays will again be different, and produce SAG.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #40



For sufficiently narrow pulses, a bandwidth limit in the transmission path can shrink the pulsewidth further. The bandwidth limit reduces the slew rates. A narrow enough pulse will not achieve full amplitude before it has to switch in the opposite direction. This will in turn reduce the time it spends above threshold.

Slide #41


- ### Outline
- What is clock tolerance and why care?
 - Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
 - Strategies

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #42

Jitter

- Cycle-to-cycle uncertainty of the active clock-edge arrival-time
- AKA
 - Phase noise
 - Temporal skew




The two forms of distortion discussed thus far have involved time-invariant displacement of the active edge of a particular clock signal from its nominal arrival time. The statistical distributions of these displacements were not a function of time. There is a type of edge-displacement phenomena called jitter which is a function of time, and we will discuss that in this section.

Conceptually, jitter is similar to skew in the way that it causes synchronization failures. The failure model for a logic segment we presented earlier is just as valid for jitter as it is for skew. Anything that causes a “late launch” into the segment, or an early sample at the end, can lead to a failure. The principal difference between jitter and skew is that the “computation” of whether the segment will fail must be made on a cycle by cycle basis for jitter.

A number of other terms are commonly used as synonyms for jitter, including phase noise and temporal skew. Phase-noise is a term from the RF world and does not, in the strictest sense, describe the part of the behavior we are most concerned about (state-device failures). This is also true for the casual/loose application of other terms which cross over from other applications. Furthermore, we shall see that there is more than one type of jitter. As you go through this section, keep that time scale in mind that we are discussing only those types of jitter that pertain directly to the synchronization of the state architecture.

Slide #43

Measure of Short-Term Frequency Stability Phase Progression Plot




Long before jitter was a concern to digital system designers, it was a common problem for telecom and RF engineers. A number of analytical methods and models for considering the stability of a signal were developed and have some merit to the state-architecture synchronization case. One of these methods is to look at the frequency stability of a clock. Jitter is a measure of short-term frequency stability, and is a direct threat to the synchronization of the state-architecture. Longer-term frequency instability is not a direct threat to that synchronization. It does have other effects which are outside the scope of this paper.

Consider the discrete-time phase progression plot above which shows the phase progress through time of an ideal clock (perfectly stable) and an actual clock. Each cycle (ticks on the vertical axis) for the ideal clock takes exactly the same amount of time (horizontal displacement). In this case, the trajectory of the actual clock varies around the ideal clock almost sinusoidally. The cycle-time of the actual clock varies with time (in other words, is dynamically distributed). In this diagram, each dot represents the arrival time of the active edge of the clock. We show a one-for-one correspondence between edges of the ideal and actual clocks.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #44



There are two types of jitter — phase-jitter and period-jitter. Understanding their difference can clarify the actual mechanisms by which synchronization failures occur. It can also give a very clear picture of what measurement method best suits your needs.

The diagram above shows the total phase and time deviation of the actual clock on cycle i . The phase deviation is the error in phase (units are cycles in our case) of the actual clock relative to its expected value at cycle i . The time deviation is the error in time, relative to its expected value, of the actual clock for that particular value of phase. These two values are a measure of the phase-noise (aka phase-jitter). They are useful measures of accumulated short-term errors, but they do not clearly describe the synchronization threat.

Period-jitter describes the cycle-by-cycle difference between the nominal and actual arrival-time of an edge, as shown in the expression for the i -th cycle jitter on the slide. An important analytical distinction between phase-jitter and period-jitter is that the former is defined with respect to an ideal clock. The latter is defined with respect to itself — the jitter of this cycle is defined with respect to the placement of the edge in the previous cycle.

Some other observations can be made from the phase progression plot. When the slope of the actual clock equals the slope of the ideal clock, both have the

same cycle time. Note that for sinusoidally varying phase noise, the points of minimum period-jitter correspond to the points of phase-jitter. Where the slope of the actual clock differs the maximally from the slope of the ideal clock, maximum period-jitter occurs, as does the maximum probability of synchronization failure. Also, note that when there is zero skew, the lines connecting the actual and ideal points are horizontal. Skew will give them a tilt.

This diagram shows a clock with period-jitter which has a repetitive component (sinusoid with a frequency of about one-twentieth the repetition rate of the clock). Jitter generally has a random component, and may or may not have a repetitive component. In the case of purely random jitter, the points along the actual clock have only two constraints: that they still fall along the horizontal lines and that they progress through phase and time monotonically.

Finally, note that if we were trying to measure the period-jitter of the actual clock shown above, our result would depend very much upon where we sampled it. If we only sampled the periods near the points where the actual clock's slope is equal to that of the ideal clock, we would presume we have no jitter. Furthermore, if the shape of the phase progression of the actual clock were such that it ran parallel to the actual clock for many cycles and then over the course of a few cycles, crossed over, it would be unlikely that any low sample-rate measurements would catch any of the largest jitter displacements. This brings out the point that to properly characterize jitter (for example, for use in your timing verifier), it is imperative that the measurement look at every or nearly every cycle. If there is a difference of several orders of magnitude between the repetition rate of the clock and the measurement rate of the instrument, the resulting measurement may substantially under-represent the actual jitter.

We will only concern ourselves with period-jitter for the rest of this paper.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #45

Sources of Jitter Fundamental Cause: Noise

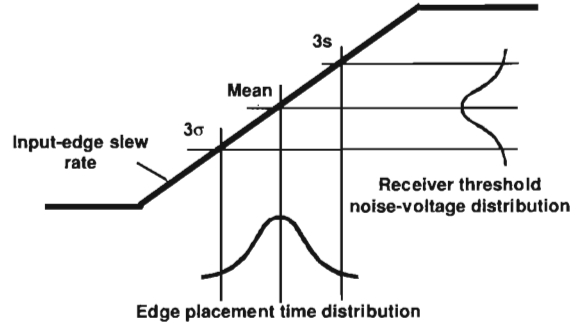
- Clock Source/Phase-Generator
- Buffer
- State-Device

There are two principal sources of jitter — active devices within the timing-environment, and the clock oscillator.

Jitter can occur in the clock generator through dynamic temperature and supply-voltage instabilities (external causality), as well as through internal noise processes (non-deterministic). The random noise includes shot noise, thermal noise, and other random internal types, as well as integrals of all of these. An excellent source of information on oscillators, and particularly the noise and frequency stability associated with them, is in numerous monographs available from NIST.

Slide #46


Device Jitter Noise on Threshold Causes Jitter



Device-based jitter comes about primarily when noise within the device (state-device or clock-buffer) causes time-varying shifts in the device's switching threshold. Noise on the device's power and ground lines can also cause these shifts (again, external causality).

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #47




In this measurement, two separate pulse generators are used to drive the CDN fixture. Due to several minor signal-integrity faults in the fixture, when one phase is driven asynchronously with respect to the other at a different frequency, the power environment is corrupted by noise. This will cause the phase we are treating as the clock to jitter in proportion to the amount of noise.

We are using two HP 8110 pulse generators — one to stimulate the “clock” phase and one to stimulate the “noise” phase. The stable output of the HP 8110 (10 ps rms-jitter) is an excellent choice for driving the clock phase. The flexible waveform specification of the HP 8110 also makes it a good choice to drive the noise phase. To inject frequency components which will result in repetitive jitter, a standard rectangular wave can be used. To cause non-repetitive jitter, the noise phase can be driven with a pseudo-random bitstream. For even more realistic simulation of actual noise processes, the ps-bitstream can be driven at three or four different voltage levels instead of the usual two.

The HP 54720A Digitizing Oscilloscope is used to observe jitter in the clock signal. Its time interval accuracy of 100 ps and 10 ps resolution are key for the measurement.

As we will see shortly, the HP 5372A time interval analyzer will give us a significantly better picture of jitter than even a high-performance scope. Specifically, we shall see that a time-interval analyzer can better help you quantify jitter and identify its sources.

Slide #48




For this measurement of baseline jitter, the noise phase is undriven. The scope plot above shows one cycle of an output of the clock phase in infinite persistence. The trigger point is on the rising edge on the left. Jitter is best measured one cycle away from the trigger point (next rising edge).

There is no spread in the trace at that point and the standard deviation of the automatic measurement of the period is about 46 ps. Since the rms-jitter of the scope is rated at about 6 ps, this is a valid measurement.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery


Slide #49



We now inject a moderate amount of noise. A 100 ns pulse with a 50% duty cycle is injected into the noise phase. The amplitude of that pulse is selected to be high enough to drive the noise phase intermittently. This was monitored on the power supply ammeter and a value was selected which drew an amount of current halfway between the current with the noise phase off and the current when the noise phase was reliably switching. This method requires the use of a source with precise and stable control of the amplitude of the waveform.

Note the “smearing” that shows up on the second rising edge. The automatic measurements now tell us we have about 240 ps of rms-jitter (1.4 ns total). That is a significant and realistic amount. The scope does a good job of alerting us to the fact that we have significant jitter present. However, despite the use of a high-performance scope, we may be seeing at most a hundred measurements per second which is several orders of magnitude below the repetition rate of the clock signal. Using the measurement we have thus far places us at risk of underestimating the largest jitter displacements present in the signal. This brings us to the role of the time-interval analyzer.

Slide #50



This plot is from an HP 5372 Time Interval Analyzer. It shows us a histogram of the time intervals of ten-million cycles (such as, a probability density function of the period of the clock)! Even for our high rep-rate clock, it measured every second or third cycle! It was obviously not subject to the under-sampling that can lead to jitter underestimation discussed earlier. Note that it captured some extremely infrequent high-amplitude displacements. The total jitter as measured by this method is about 3 ns, which is about 1.4 ns higher than we were able to measure before! You don't want to find that extra 1.4 ns in proto debug (or later). The total time to take the measurement was a few seconds.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #51

Measurement Methods

- Scope
 - Infinite Persistence
 - Histogram
 - Automatic measurement w/ stats
- TIA
- Phase Noise Measurement Instrument
- Spectrum Analysis

We have just seen a useful demonstration of jitter measurement using time-interval analysis. The fact is that several methods exist, and they all have their place. The author has found jitter and jitter measurement to be a very current topic with high-speed computer designers throughout the industry.

The most available test instrument to high-speed designers is the oscilloscope. A good digitizing scope can capture and present a wide range of useful information and is the best test option in a number of circumstances (for example, state-device characterization). Digitizing scopes also have a number of useful presentation modes in addition to just free-running. One is the use of infinite persistence which will show you the presence of an anomalous/infrequent event without having to continuously focus your attention on the display. In chasing down very-low-frequency of occurrence synchronization errors (for example, low-frequency metastability) over the years, the author has spent more than his share of time with the lights down, leaning over the scope, hands cupped around the display of an analog scope for hours trying to view an occasional metastable output. Infinite persistence works. Automatic measurements are also extremely useful features. The fact remains that for jitter characterization, anyone using an oscilloscope on a high-speed clock is probably not going to view the largest displacements. The scope, however, is usually the tool that alerts you to the fact that jitter is present, assuming that the scope jitter is well below the signal jitter.

Slide #52

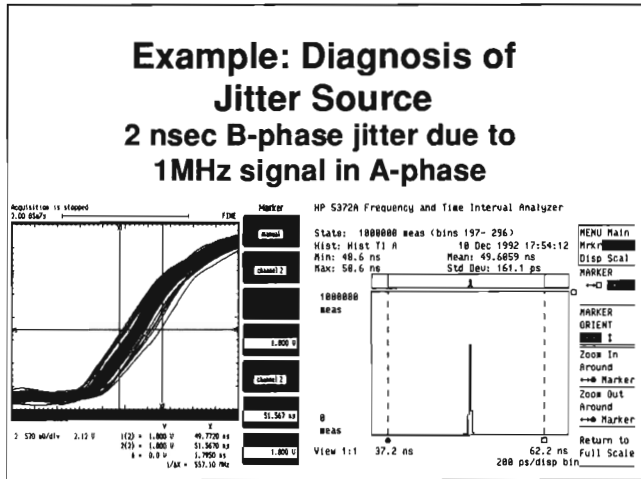
Time-Interval Analysis Significant Advantages for Jitter Characterization and Diagnosis

- Measures many (all) cycles very fast
 - More accurate statistics
 - More likely to catch infrequent (largest) displacements
 - Higher confidence in measurement results
- Present results in frequency, phase, or time format
- Jitter spectrum-analysis identifies repetitive sources

We have already seen the kind of high-confidence statistics we can generate very quickly using a TIA. The instrument can make a large number of other types of measurements and display the results in a variety of forms. Another useful timing application for a TIA is the characterization of PLL clock buffers. Before leaving the subject, however, we should also point out one other important application of the TIA — jitter spectrum analysis. Jitter spectrum analysis is more of a diagnostic tool than a characterization tool, as we shall see.

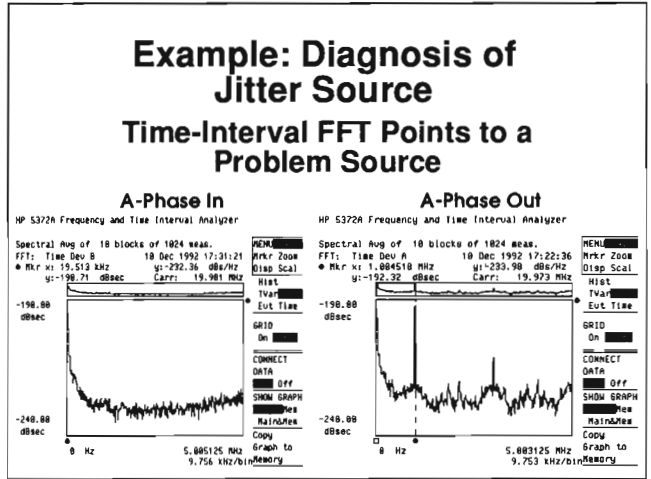
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #53



In this example, we're injecting a 1 MHz squarewave into the noise phase. The plots above show scope and TIA representations of the signal. The total jitter is about 1.4 ns. Either view shows us we have jitter, but we don't have any information about the source from these photos.


Slide #54



The plot on the right is a jitter spectrum analysis of the input to the clock phase of the CDN (HP 8110). It shows no repetitive jitter. The plot on the right shows a jitter spectrum analysis of an output of the clock phase of the CDN. There is a prominent spike at 1 MHz, and significant spurs at harmonics of 1 MHz. This measurement has essentially told us the source of our problem! The requirement for jitter spectrum analysis is that the jitter must have a repetitive component.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #55



The graph above shows us the tolerance breakdown for the CDN fixture. Clearly the biggest contributor is the clock buffer. This is the common result for a 5-level untuned CDN. The use of high-precision clock buffers would have reduced that component and the overall distribution considerably. Since there are some signal integrity faults in the system as constructed, the system has an excessive amount of jitter as well. Module etch tolerance should be about two-thirds of its current level.

Slide #56

- ### Outline
- What is clock tolerance and why care?
 - Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
 - ➔ • Strategies

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #57

Front-End TE Design Strategies Do the Work Up-Front

- Fully employ all appropriate ad hoc measures
 - High-precision clock buffers
 - Balanced/radial distribution
 - Minimize total # inches of clock path interconnect
 - Self-characterization of TE components
 - Path tuning, etc.
- Understand and apply higher-level timing schemes
 - High-speed CDN
 - Tolerance-insensitive schemes
 - Multiple phases/clocks
 - Regeneration/polychronic (semi-asynchronous), etc.

Controlling skew, jitter, SAG, and other clock signal distortions, as well as addressing all of the other functionality (scan, performance, and stalling) required of the clock, is NOT a simple process. Timing environment design is demanding work and for the fastest systems, requires specialized knowledge. However, proper tolerance management results in correct operation AND enables the computation environment to run at its maximum potential performance.

There are a large number of ad hoc methods that are commonly employed to enhance the precision of the placement of the clock edge. The list above is just a starting point. The reader should determine what methods are appropriate for the design and implement them.

One of the most effective methods for some designs/products is to build up a technology board or system for self-characterization, as discussed earlier. The possibility of a departure from worst-case design requires the designer first to carefully consider if that is suitable for his particular design. If so, he must then consider what design and maintenance methods to employ, and of which measurement methods to use. If you take this route, you will find it useful to characterize the behavior of the CDN components at both steady-state and through warm-up (this is a good idea for the whole system, not just the CDN).

An understanding of both the higher-level timing schemes and the alternative synchronization schemes will be useful in selecting a timing architecture that is best suited to the speed and logic complexity of your system.

Slide #58

Back-End Strategies

- Rapid "verification" of timing decisions
 - Verification/debug plan
 - Margin verification of longest and shortest paths
 - Margin testing of all "timing boundaries" & conditions
 - Stall/unstall & fast-clk/slow-clk interfaces
 - Interfaces between init'd & non-init'd structures
 - Latch <--> flop segments
- Rapid isolation of unanticipated timing faults
 - Infrequent/migratory/non-repeatable - Analytical approach
 - Hands-on otherwise
 - Factor-in instrumentation tolerances!!!
 - Healthy trait - distrust tools & people

Once a prototype exists, the final timing activities can be carried out, as well as the other requisite proto-debug activities. Some degree of timing problem at proto-debug is statistically likely, and should be anticipated. To avoid floundering, timing verification and characterization should be specifically addressed in the proto-debug plan. An integral part of that plan are contingencies for what tack to take if the proto is entirely nonfunctional. An important adjunct to that plan is to have "hooks" in place in the timing environment to facilitate timing debug (for example, the ability to run at slow-margin, fast-margin, and drive with an external clock). A detailed examination of how actual system timing compares to what was anticipated can usefully serve to "tune" your timing-environment design process for the next system.

If your system exhibits symptoms which indicate some type of timing fault (for example, infrequent and unrepeatable or migratory failures), how do you isolate the problem? At this point you must choose between an analytical and a hands-on approach. For repeatable symptoms which occur with enough frequency to probe efficiently, going straight to measurement makes sense. Otherwise, the author suggests that an initial analytical approach is the most efficient methodology. Then, as verification of analytically-developed conclusions is required, move to a hands-on/measurement phase. The reasoning behind this is that unrepeatable, infrequently occurring failures, in tandem with the migratory failure mode described earlier make efficient

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

diagnosis by traditional troubleshooting methods impossible. If you only get one crack at the problem per day (or week, or month), and the failures are "migrating" throughout a complex system, it is highly unlikely the probes and the problem will "get together" in any reasonable period of time.

On the other hand, after formally considering the symptoms present, measurement affords an excellent means of verifying (or not) conclusions about the source(s) of the problem by setting up controlled experiments involving only small regions of the circuits. Scan design methods can provide excellent controllability in setting up the experiment's initial conditions.

Slide #59

Resources

- HP 8110
 - 10 psec edge placement
 - 10 psec-rms jitter
 - Flexible waveform specification
- HP 54720A
 - 1.1 GHz BW
- Amherst Systems Associates
 - Timing Environment Design/Measurement/Training
 - M.K. Williams — Owner/Principal Consultant
P.O. Box 24, Amherst, MA 01004
(413) 596-5354

Slide #60

Appendix

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #61

State-Device Characterization and Measurement

- Can provide valuable insight into:
 - Actual device tolerances
 - Device-level failure modes encountered during proto debug
- Collect parametric timing data for timing verifier
 - Gain confidence in catalog figures
 - Some data not spec'd in catalog
 - Jitter susceptibility
 - Metastability resolution time-constant and aperture

As the receivers of the clock, the system state-devices play an obviously important role in timing. They establish the arrival-time constraints, which in turn dictate all other system timing decisions. Given their importance, it's a sound design practice to know as much about their behavior as possible. The process of characterizing the state devices you will be using in your design (or are evaluating for use) can give the designer valuable insight into how the device(s) will perform. I have found, for example, that a detailed knowledge of a device's actual parametric distribution (versus the catalog numbers), in conjunction with the device's functional and parametric behavior under a variety of marginal triggering conditions, to be invaluable during the debug of suspected timing faults. Of course, there are other reasons as well, including building your own parametric distributions for use in your timing verifier. Note that any characterization process should include an examination of their behavior in the same physical environment in which they will eventually operate.

In this section, we are addressing primarily designers of high-speed synchronous systems. Other engineers such as synchronizer designers (communication) and semiconductor test engineers (for example, producing specification sheets) need to perform state-device characterization as well, but those applications are beyond the scope of this paper.

Slide #62

Self-characterization - Build Your Own Distributions

"Heresy!!! What if the process changes?"

- Catalog tolerances typically have several elements
 - Process - you have no control over this
 - Supply voltage operating range - controllable
 - Thermal variation - controllable
 - Load variation - controllable
 - Instrumentation/measurement variation - controllable
- Supporting design & maintenance methods


A design option for very aggressively timed systems is to employ self-characterization. It can be employed to factor out any excessive margins in the device manufacturer's guard-bands for rated specs, to determine device parameters that are not rated (for example, jitter susceptibility), or to simply determine confidence levels for catalog figures to developing your own distributions for device timing parameters. In the limit, this becomes manual sorting and matching of parts (sometimes called "graded parts"). In choosing a self-characterization approach, one must specify maintenance procedures (device replacement policies, re-characterization, and restriction to use of pre-characterized part inventories) which recognize that changes in the manufacturer's process can eventually yield parts with a different distribution than originally characterized.

Parametric self-characterization is sometimes viewed suspiciously by designers the first time they hear of it, due to concerns over process variations over time. What must be kept in mind is that the rated minimum-maximum on most timing parameters are based upon a number of distributions and other factors. You don't have control over the process, so you generally can't "cheat" on that. But you do have control over thermal, loading, and power environment, and these you can characterize out. Some parts now even provide derating tables for this very process

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

(such as Motorola clock translators). Finally, there are also non-technical considerations occasionally built into ratings as well by lawyers and marketing engineers. The guardbands also contain factors to accommodate the manufacturer's instrumentation tolerances and measurement methods.

Slide #63



The figure above shows the HP 8133A Pulse Generator and the HP 54720A Digitizing Oscilloscope.

The HP 8133A was selected for its extremely precise edge placement (1 ps through the front panel, 300 ps over the HP-IB), and its very low rms-jitter (1.5 ps maximum, typically less than 1 ps). While not formally specified, the interchannel jitter is even lower and is an important consideration for making an accurate setup or hold time measurement.

The HP 54720A was selected to make this measurement for a variety of reasons. Its ability to make fast, accurate real-time measurements coupled with its very high update rate mean that you have the highest probability of capturing fast, infrequent events (such as, the trajectory of a metastable output) and then making very accurate measurements on that data.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #64

Instrumentation Considerations The tolerances end up in your cycle-time!

- Get guaranteed performance specs!
- Stimulus - Pulse Generator
 - Accurate edge-placement
 - Extremely low jitter
 - Two channels w/ differential drive
- Response - Digitizing Oscilloscope
 - High update rate
 - High sample rate
 - High time-interval and rise-time precision

A key to achieving an accurate characterization is to have a two-channel signal source which has the ability to precisely vary the delay from one channel to the other. It must also have jitter which is at least an order of magnitude below the minimum channel to channel delay you anticipate.

A second key to accurate characterizations is to use a digital scope. With analog scopes, measurement accuracy is dependent on the intensity setting, and the extremes of "hang-time" may not be sufficient to light the phosphor.

Slide #65

DUT



The device under test is a Motorola MC10KH131 ECL flip-flop. It is capable of switching rates in excess of 250 MHz, and its actual setup and hold times range from the high-tens to low-hundreds of picoseconds.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #66

State-Device Failures: A Closer Look

- Progression through the setup/hold aperture
- Advance data arrival w.r.t. clock
- DUT = 10KH131 (ECL FF)

The measurement is made as follows. Assuming a fixed data edge, the clock-edge (and its associated setup and hold interval) is walked in toward the data edge until anomalous output behavior is noted. That's your setup violation. Obviously, the smaller the steps are and the lower the stimulus jitter, the better your result will be.

To find the hold time, move the active edge of the clock back to a point before the setup point. Then reduce the pulsewidth of the data until anomalous behavior is noted. The hold time is the separation between the trailing edge of the data and the active edge of the clock. Depending upon the technology of the state device you're using, you need a delay resolution on the order of ones of picoseconds and a delay magnitude equal to the sum of the expected setup and hold times.

A marginally triggered state device can behave in a number of different ways. It depends upon the type and magnitude of timing violation, and the device being measured.


Slide #67

Four Distinct Stages

- Stage 0 - Normal Output Behavior
- Stage 1 - Loss Edge & Corner (Pre-metastable)
- Stage 2 - Low-Grade Metastability
- Stage 3 (Early) - Infrequent Failures
- Stage 3 - Full-Scale Metastability
- Stage 4 - Consistent Failure
- Other violations (Th, PWmin, etc.) will produce other output behavior

Slide #68


Stage 0 - Normal Output Behavior



This photo shows the Q-output of the state-device when it is operating normally. It will be placed in the memory of the oscilloscope for use as a reference in subsequent measurements.


Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #69



At this point, we have advanced the clock edge so that the data edge is now in or near the setup and hold interval. Note that the output has degraded (catalog risetime limit is 2 ns, this one measures over 2.2 ns). That is, relative to the previous measurement placed in memory, the edge rate has decreased (increasing the propagation delay through the part) and it appears to have begun to jitter. We will see that it is not true jitter in the next slide.


Slide #70




Decreasing the separation between the clock and the data edges a few more picoseconds, we see that the apparent jitter is actually low-grade metastable behavior. Jitter displaces the edge in time uniformly at all voltage levels. In this case, the lower half of the edge is stable. Note that none of the trajectories ever resolve into an incorrect state at this point, however, the delay through the part has increased (and is even time variant). If the segment downstream is long enough to be a critical path, this additional delay in launching could produce failures well down stream from this point.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #71




Slide #73




Advancing the edge one more picosecond, we now get a single failure (resolves to incorrect state) during the measurement period. This is one of the types of low-level failures that produces the difficult failure modes discussed earlier.

Slide #74



Slide #72



This shows about an equal number of trajectories resolving high and low. At this point we are well into the setup/hold interval. This behavior would be extremely easy to debug due to the high number of failures.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #75

Case #1 - Mismanaged Timing Environment Design

- Large prototype ECL array processor
 - 2-phase, flip-flops
 - Approx 400 board-level clock loads
- Symptoms:
 - Infrequent failures at full speed (migratory)
 - Correct operation at 95% full speed

As an example of the post-design headaches improperly considered timing can bring, consider a set of four essentially identical prototype ECL (10KH) array processor systems. Each system is physically-large, being comprised of 72" mil-racks fully populated with logic modules. The system state-device is the 10H131 flop. All systems exhibited sporadic, unrepeatable failures characterized by apparently incorrectly captured data when running at full speed (approximately 40 MHz/25ns). The systems appeared to operate correctly at up to approximately 95% of full design speed. The full-speed errors occurred on the order of two to four times per day per system. The mode was almost never repeated for any of the failures because the specific symptoms changed both their nature and location for each failure.

Slide #76

Case #1 - Outcome

- Diagnosis:
 - Systemic violations of timing constraints
- Rx:
 - Change state devices
 - Widen clock pulse
- Client impact
 - 11-week delay in proto availability
 - Unforeseen/unbudgeted diagnosis/repair costs

An analysis of the state architecture revealed inherent timing errors when running at speed due to marginal timing. Specifically, due to a misunderstanding of timing requirements multiphase flip-flop-based systems, a large number of system segment delays were slightly larger than would be allowed at full design speed. In this case, diagnosis was made analytically (approximately 1 week), followed by a period of taking measurements to verify the analysis (approximately 2.5 weeks).

The interim solution was to convert the state architecture to 2-phase, latch-based. Specifically, all 10H131 flops were replaced with 10H130 latches, and the pulsewidth of both clock phases was widened slightly (13.0 to 14.5 ns) to fully accommodate the 13.9 ns of predicted clock skew. Further modifications concerning other parts of the state architecture and timing environment were suggested for subsequent designs. The interim solution required no changes to pwb etch due to the full pin-compatibility between the latch and the flop. Consequently, the initial round of repairs was made quickly (four days total). The consultant spent a total of five weeks on diagnosis and verification of the repair. The client spent an additional five to six weeks prior to that on the problem. The impact of the 11-week loss had a very negative effect on the client's development plan and budget, and could easily have gone higher if any of the redesign involved modification to board etch. An intelligently-invested one or two man-weeks at the front-end would probably have obviated the several man-months of time and the additional expenses at the back-end.

Slide #77

Case #2 - TE Design Addressed Properly

Heavily pipelined ECL/L2/250 cpu

- T_{cyc} = 22.5 nsec, Skew = 3.75 nsec
- Significant front-end time on TE/state architecture
- Ran at full-speed from initial power-on
- 22.5 nsec of logic per cycle!
- No timing failures in protos or field units

Note that BEYOND correct initial operation, careful TE design produced a big win for performance. In every 22.5 ns cycle, the data is propagated through 22.5 ns of logic, despite the presence of up to

3.75 ns of global skew (of course, this is statistical—a few systems will have the full 3.75 ns, others will come in with less). Stated another way, for every 22.5 ns of operating time, the system gives you 22.5 ns of work back. Had timing environment and state architecture issues not been properly addressed, the outcome would probably have been the 22.5 ns logic time (maximum segment time allowed) padded with the 3.75 ns (or more) of clock skew producing a cycle time of 26.25 ns (or more). In that case, the skew comes directly out of your speed budget with the undesirable result of requiring 26.25 ns of time to complete 22.5 ns of work. If clock skew were not properly predicted and accommodated in the design, 17% of the cycle time would be wasted.

Stated another way, without the right approach, 62 days of the year you have a \$700 K, 9.2 KW, 1600 pound PAPERWEIGHT!