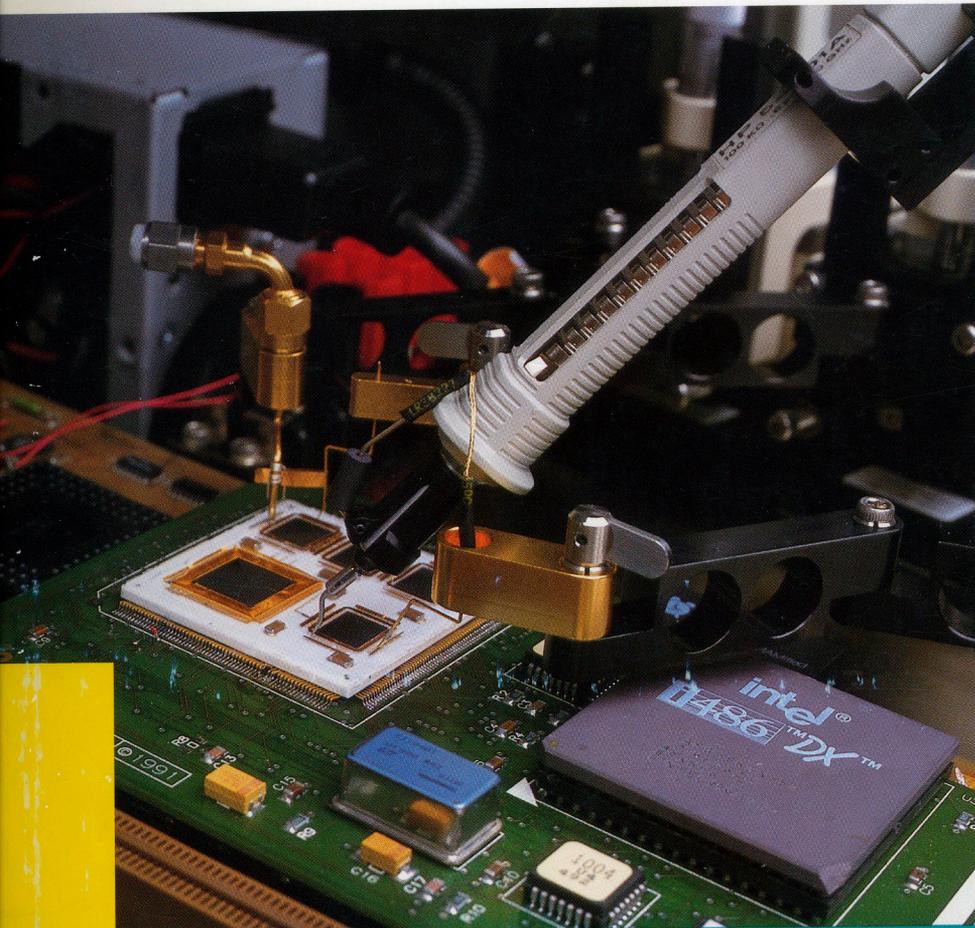


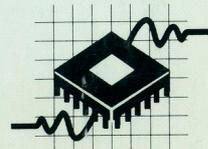
1993 High Speed Digital Design Symposium



Technical Papers

**High Speed System
Design Track**

**Digital Communications
Design Track**



Hewlett-Packard High Speed Digital Design Symposium Europe 1993

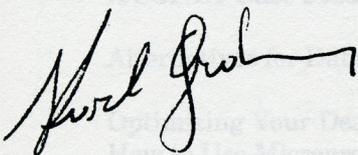
Welcome to this year's event!

As speeds and complexity of systems increase, so too do the design challenges facing today's digital design engineer. These challenges are coupled with an ever more competitive environment where product time to market is the key factor determining its success or failure.

Hewlett-Packard understands and is committed to support your complex digital design and test needs, thus helping you to bring your design to market faster. That commitment is demonstrated in the High Speed Digital Design Symposium.

We give you the opportunity to hear industry experts and consultants present an insight into the latest digital design, test and technology issues, design techniques and measurements solutions based on their own experiences. The 1993 event offers you a choice of two tracks, focussing on either high speed system design or digital communications design.

I hope you have an informative and enjoyable day at the symposium and that for you, today's design challenges become tomorrows opportunities.



Karl Grabner
Marketing Manager Europe
Digital Design Tools

Table of Contents

High-Speed System Design

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module	1-1
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery	2-1
Printed Circuit Design Techniques for the Control of Electromagnetic Interference	3-1
Advanced Methods for Noise Cancellation in System Packaging	4-1
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware	5-1
Glitches, Intermittents and Noise . . . Building in Reliability	6-1

Digital Communications Design

Understanding Evolving ATM Standards and ATM Design Verification	7-1
Physical Layer Design Issues for Serial Communications: A SONET Case Study	8-1
Alternatives for Data Transfer in High-Speed Systems	9-1
Optimizing Your Design Flow . . . How to Use Microprocessor and ASIC Emulators	10-1
Developing and Debugging an ISDN Terminal Adapter	11-1
A High-Performance Environment for Modelling and Simulation of Digital Systems	12-1

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Arthur Fraser

**TechKnowledge
5017 N. Amherst
Portland, OR 97203
Phone: (503) 289-2637
Fax: (503) 626-6023**

John Kaufman

**NCR Corporation
2850 El Paso Street
Colorado Springs, CO 80907
Phone: (719) 578-3415
Fax: (719) 473-0020**

Ken Smith

**Cascade Microtech
14255 SW Brigadoon Ct.
Beaverton, OR 97005
Phone: (503) 626-8245
Fax: (503) 626-6023**

**1993
High Speed Digital
Systems Design & Test
Symposium**

Doubling the Clock Speed
Feasibility Study on a
100 MHz Cache Module

Abstract

No one is paid to design something that goes slower. Rather, everyone wants to go faster, and digital engineers are often paid to efficiently iterate a current design to the next higher speed. The expectation is that if a system works at X MHz, than it should be almost trivially easy to make it work at 2X MHz.

Unfortunately, there are gotchas waiting for the unwary at all speeds. In general, it is about as difficult to move from 50 MHz to 100 MHz as it is to go from 500 to 1000 MHz. The unexpected effects or seemingly innocuous design constraints lurk, waiting for the engineer who simply puts in a new crystal, hoping everything will work OK.

This paper demonstrates several measurement and modeling techniques useful to working engineers designing their next higher speed system. The design example described herein is an NCR 50 MHz 486 cache module, where the objective is to obtain sufficient information from the current design to enable accurate signal integrity predictions of higher speed designs.

Authors

Arthur Fraser

Current Activities:

Co-founder of TechKnowledge, a company specializing in researching, packaging, and presenting technical information in human terms. Currently involved in medical application software marketing, guarded femto-amp current measurements, specialized travel for physically impaired, lecturing medical courses, and helping people start their own company for cheap.

Author Background:

Educated (MSEE) in philosophy, device physics, psychology, high speed & high frequency design, eastern religions, management, IC design, ionospheric wave propagation. Previously employed at TriQuint and Tektronix doing reliability physics, failure analysis, teaching high speed design, and helping custom foundry customers be successful. Written many papers and application notes for various companies.

Authors (cont'd)

John Kaufman

Current Activities:

Principal Design Engineer at NCR Microelectronics Products Division working in Advanced Development. He is involved in electrical model simulation, hardware design for various high speed MCM circuit designs. The MCM Advanced Development group can take an existing design and condense it into a tight package providing all the layout and simulation requirements resulting in a final MCM product. Thus the NCR MCM group relieves the customer of coordinating all the necessary suppliers and stream-lining the technology process.

Author Background:

John received his BSEE from Ohio State University in 1987. He started with NCR E&M Cambridge upon graduation as

a digital designer. In June 1991, John transferred to the MCM group. He provided design, debug and simulations support for the 50 MHz 486 MCM. John has experience in simulation using the Mentor Graphics CAE tools for board and MCM design and in debug using various high performance digital scopes and logic analyzers.

Ken Smith

Current Activities:

Vice President, module probing business unit at Cascade. Responsible for managing and developing probes and stations for characterizing IC interconnects, packages, testing and troubleshooting fine pitch digital boards and modules.

Author Background:

BS in General Engineering from Oregon State University in 1979 with a focus on electronic and thermal properties of materials. Ken has over 13 years experience in the design and manufacture of high performance hybrid microelectronics and systems. His work at Tektronix included development of: A multilayer thin film process for MCM substrates, high frequency packages and probes for GaAs ICs, and a miniature handheld oscilloscope probe.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

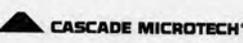
Slide #1

Doubling the Clock Speed

Feasibility Study on a 486-50 MHz Cache Module at 100 MHz

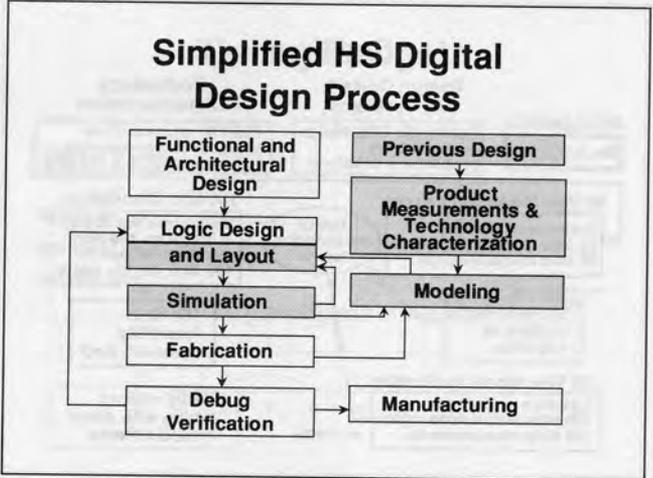







No one is paid to design something that goes slower. Rather, everyone wants to go faster, and digital engineers are often paid to efficiently iterate a current design to the next higher speed. The expectation is that if a system works at X MHz, than it should be almost trivially easy to make it work at 2X MHz. Unfortunately, there are "gotchas" waiting for the unwary at all speeds. In general, it is about as difficult to move from 50 MHz to 100 MHz as it is to go from 500 to 1000 MHz. The unexpected effects or seemingly innocuous design constraints lurk, waiting for the engineer who simply puts in a new crystal, hoping everything will work OK. This paper demonstrates several measurement and modeling techniques useful to working engineers designing their next higher speed system. The design example described herein is an NCR 50 MHz 486 cache module, where the objective is to obtain sufficient information from the current design to enable accurate signal integrity predictions of higher speed designs.

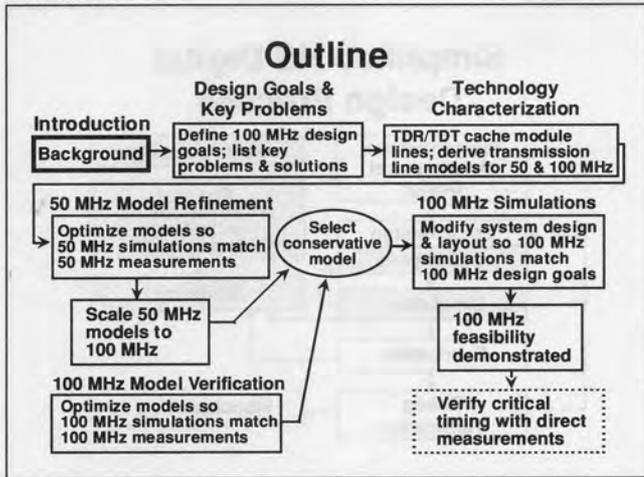
Slide #2



This paper focuses on measurement, characterization, and modeling techniques for engineers who have an existing working product and are designing a similar product with higher speeds or finer pitch layout rules. As the conductor pitch becomes smaller to accommodate higher densities, the measurement techniques change as hand-held probes are unable to contact increasingly smaller wires. Similarly, as speeds increase, measurement and modeling techniques change to accommodate additional physical effects, and to incorporate them into the simulation tools. Frequency-dependent interconnect losses, or *skin effect*, is an example of a physical effect which needs to be modeled as speeds increase. It is important to know accurately at what speed these effects need to be modeled, for the particular interconnect technology used.

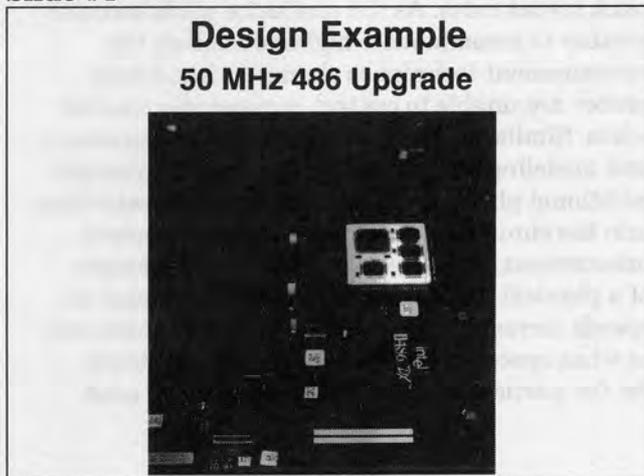
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #3



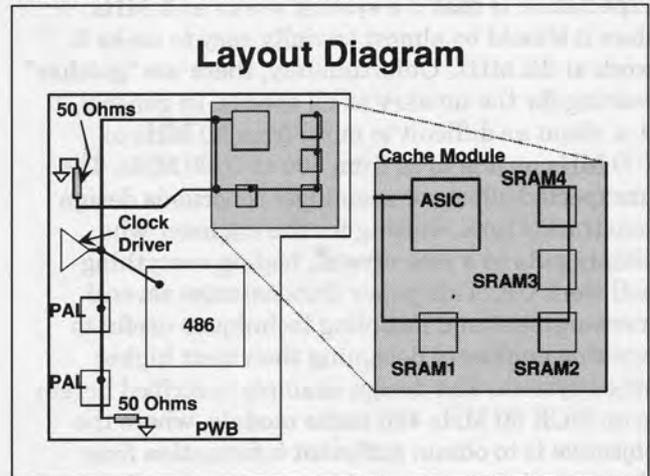
The product was first fabricated entirely on FR4 circuit board and electrically debugged. It measured 7 by 8 inches and was of course much too large to meet the final product mechanical requirements. In the current product, all the SSI/MSI were incorporated into PALs; the cache controller ASIC and the cache RAM were all placed on a single module built by NCR, fabricated in multilayer ceramic. These chips were TAB'ed onto the module, resulting in a very small assembly. This greatly helped to reduce the overall size and was intended to provide a fundamental high speed design which will be used for the next several generations.

Slide #4



NCR has graciously allowed the use of one of their product designs as the working example. This product is a 50 MHz 486 based system designed as an after market 486, 33-MHz upgrade. It is intended to plug into a 486 socket and enhance performance. This photo shows the 486, the clock generator/driver on the left, and the cache module on the upper right. Overall dimensions are 3.5 x 5.5 inches.

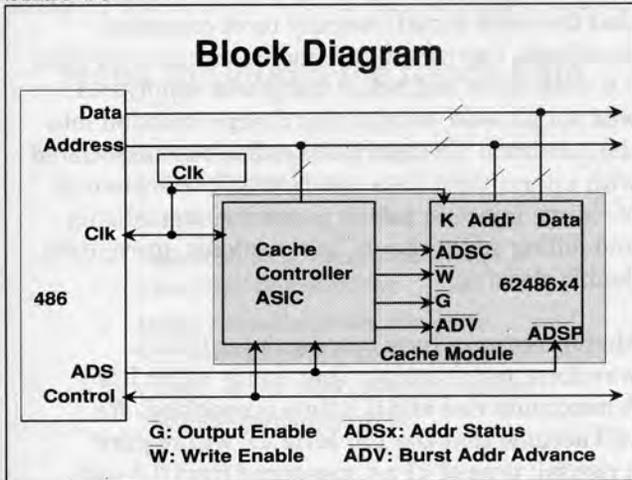
Slide #5



The clock distribution circuit is relatively straightforward, with the clock generator and buffer located on the printed wiring board (PWB), and the clock signal distributed to the 486, the PALs, and then to the cache module. On the cache module, the clock signal is distributed first to the ASIC and then to the 4 SRAMs. Note the two 50 Ohm termination resistors.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

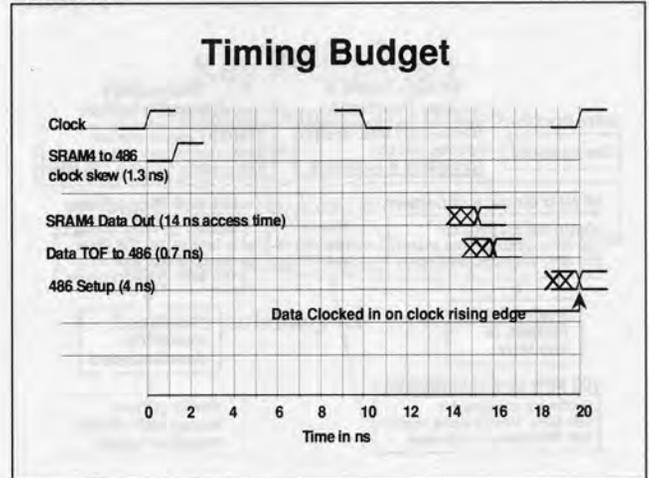
Slide #6



The basic functionality is similar to most 486 PC designs, with a 50 MHz processor and RAM cache communicating with a slower main memory and associated control chips. The RAM cache and its ASIC controller and associated clock distribution were the most difficult part of the original design work, so will be used as the design example in this paper.

When the 486 does not find the desired address in the first level 8K cache, an external read/write is generated. The cache controller ASIC determines if the requested addresses are in the second level cache, and if they are, completes the operation. If the next 3 addresses are also desired and are sequentially in the cache, the controller will initiate a burst mode read/write, executing the next three operations at 1 read/write per clock cycle. Each cache SRAM has a 2-bit address counter, which is incremented at each clock cycle rising edge when in the burst mode, facilitating the burst mode operation.

Slide #7

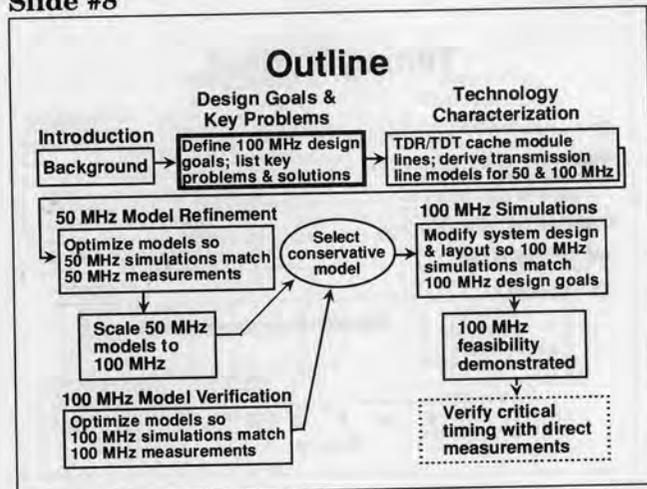


With a 50 MHz 486, and 14 ns SRAMs, the ASIC needs 1 clock cycle to decide if the data required is in the second level cache. If the data is cached, then the succeeding 4 reads or writes can occur in burst mode, at one operation per clock cycle. The timing diagram shown here is for a burst mode read operation from SRAM4. SRAM4 is electrically the furthest SRAM from the 486, so represents the worst case. At time 0, the clock rising edge starts another cycle. Because SRAM4 is located further from the clock driver than the 486 by 1.3 ns rising edge skew, SRAM 4 will start its operation 1.3 ns later than the 486 will. Fourteen ns later, valid data is available from SRAM4, and 0.7 ns later the data arrives at the 486. The 0.7 ns delay is due to interconnect time of flight (TOF) delay of 0.07 ns/cm for circuit board traces. The 486 requires 4 ns setup time before the clock latches the data into the 486 at the next clock rising edge.

If the clock rate were doubled to 100 MHz (10 ns clock period), and no other changes made, then 1 clock cycle operation would require 4 ns SRAMs: $(10 \text{ ns} - (1.3 \text{ ns skew}) - (0.7 \text{ ns data TOF}) - (4 \text{ ns processor setup time}) = 4 \text{ ns})$. A 100 MHz processor will probably have a 2 ns data setup time, and if you can cut the clock skew and data time of flight delay in half, then less expensive 7 ns SRAMs will work. If the layout is constrained to be the same, then 6 ns SRAMs are required. In this example, we are assuming that 7 ns RAMs will be used for the 100 MHz system.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #8



Implicit in these design goals was the requirement that the clock signal integrity meet company standards. Overshoot and undershoot greater than 0.6 volts above and below the power supply rails was not allowed, minimizing charge injection into the substrate. Charge injection has been associated with intermittent data loss related to the amount of charge injected, part type, and vendor. Rising and falling edges should be monotonic, preventing double clocking.

Microprocessors have more stringent clock waveform requirements than many other ICs. A maximum rise and fall time is specified. We will assume that the 100 MHz uP will require a rise/fall time of <1 ns, measured from 0.8 volts to 2.0 volts, rising and 2.0 volts to 0.8 volts falling. A minimum high and low time is also required, and we will assume the minimum high/low time is 3.5 ns each, with a high defined as > 2.0 volts and a low defined as < 0.8 volts. In this example, we are assuming the cache module ICs have the same waveform requirements as the 100 MHz uP. If the 100 MHz simulations do not meet these requirements, then modifications should be recommended that eliminate undesirable signal waveforms.

- Assumptions for the 100 MHz feasibility study:
- A 100 MHz microprocessor (100_MHz_uP) will be used similar in function to a 50 MHz 486
 - Basic system functionality will be the same
 - Only one clock driver will be used
 - The 100_MHz_uP data setup time is 2 ns max
 - 7 ns SRAMs similar to current product will be used
 - The cache module will be similar in design and layout
 - The cache controller ASIC will be upgraded to 100 MHz

While a 100 MHz 486 is not available, and may never be, it is useful for this feasibility study to pretend there is one. In this paper, the 100 MHz microprocessor will be referred to as the "100_MHz_uP."

Slide #9

100 MHz Design Goals

- Demonstrate 100 MHz 1 cycle burst mode
 - SRAM4 to 100_MHz_uP data TOF ≤ 0.35 ns
 - 100_MHz_uP to SRAM4 clock skew ≤ 0.65 ns
- Demonstrate acceptable clock signal integrity
 - monotonic clock edges
 - clock max/min not exceed 0.6 V outside power supply rails
 - rise/fall time ≤ 1 ns (0.8 to 2.0 volts)
 - min. high time & low time = 3.5 ns each
- Select correct cache clock line model for 50 and 100 MHz operation
- How fast will the cache module go?

Assuming that the clock edges, IC setup and hold times, and SRAM access times all scale linearly with clock speed, will this system work at 100 MHz with the current layout? This is the question NCR desired an answer to. If the answer was no, then what would be the minimum changes needed to reach 100 MHz? NCR also wanted to know "Just how fast can the cache module go?" The cache clock line is considered the critical path, so the question then became "What is the maximum speed of the cache clock line?"

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #10

Steps to Double the Clock Rate

- Characterize critical module transmission line
- Measure existing 50 MHz product
- Refine 50 MHz models to match measured performance
- Using refined models, simulate 100 MHz performance
- Compare 100 MHz simulations with 100 MHz measured data

The first step in reaching the design goals was characterizing the cache clock line with a TDR/TDT instrument. The TDR/TDT instrument quickly and easily measured transmission line characteristics, including impedance, electrical length, and maximum edge speed. With this information an accurate transmission line model was selected for simulation at 50 and 100 MHz.

The second step was measuring the existing 50 MHz product performance and comparing with simulations of the 50 MHz product. The simulations used the transmission line model selected in step 1. (If your measurements and simulations don't match well, then modify specific component values in the models obtaining a better match.) The clock driver model was modified during this step. This was expected as a good driver model was not available.

After the 50 MHz simulations matched the 50 MHz measurements, and we had confidence that the transmission line models were accurate at 100 MHz, then the third step was simulating the 100 MHz system. Ideally, in this situation, part or all of the system should be fabricated, so measurements could be made and compared with simulations to build confidence in the simulation models. In this example, a 100 MHz clock and driver were plugged into the 50 MHz system. While the complete system was not fully operating, sufficient information was obtained to justify confidence in the 100 MHz simulations.

Slide #11

Key Problems/1 Probing Small Structures

- Hand held probes too big
- Fixed pitch wafer probes won't work
- Requires variable pitch high speed probe

NCR encountered problems measuring signals on the cache module. Hand-held probes would not work. When the lead/interconnect pitch decreases below 25 mils (1mm), hand-held probing becomes impractical. It is simply difficult to consistently hold a probe on a selected lead, and as the pitch goes even smaller, it is difficult to even see and locate the interconnect. With TAB mounted chips, a slip of a probe can cause major damage to the TAB lead frame. Coplanar probes were tried, but did not work well as fixed pitch test pads had not been laid out. Wires were then attached to the coplanar probe ground contacts and connected to the module ground points. This not only damaged the modules under test, but the ground wire inductance distorted the measured signals.

In response to NCR's needs, Cascade Microtech and NCR worked together refining the design of several new probes for contacting fine pitch structures on circuit boards and modules. A wide range of probes are now available, including 1X, 10X, 20X, and 100X passive probes, as well as active probes. These all feature low inductance ground connections which are positioned independently from the signal probe, so special test pads are not required.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #12

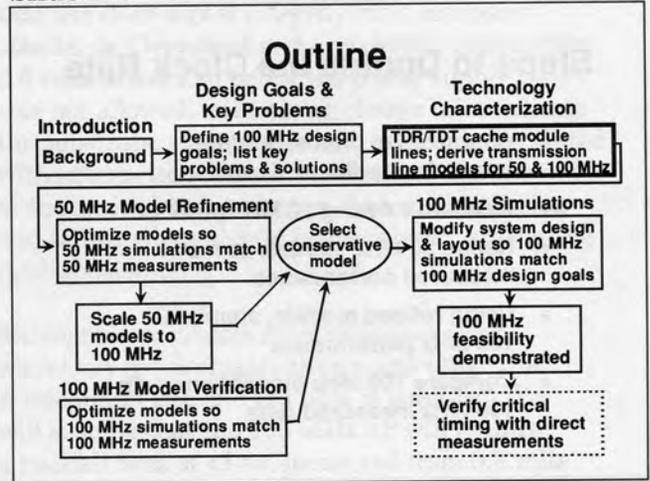
Key Problems/2 Which Transmission Line Model?

- Use simplest model until TDR/TDT or VNA measurements of worst case line indicate otherwise.
 - Ideal model (simplest)
 - Lossy model
 - Lossy model with skin effect

Much has been written on modeling circuit board and module interconnects. Beyond about 20 MHz, typical PWB interconnects must be modeled as transmission lines. However, when simulating transmission line behavior, should you use an ideal model, a lossy model, or skin effect model? The approach taken in this project was to use the simplest model until direct measurements on the worst case wire indicated a more complex model was required.

See Appendix A for additional information on transmission lines.

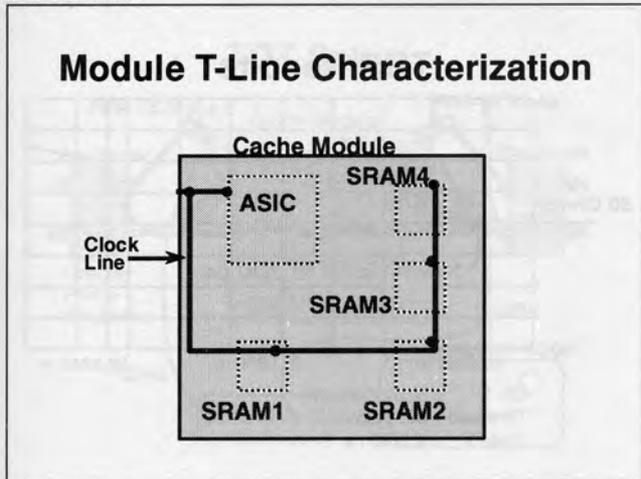
Slide #13



The technology characterization section involves using an HP 54120 time domain reflectometer (TDR) and a time domain transmission (TDT) instrument to measure the fundamental characteristics of the module clock line. From these measurements, you can select the simplest transmission line model that accurately models the measured behavior at the desired system speed (50 MHz and 100 MHz clock).

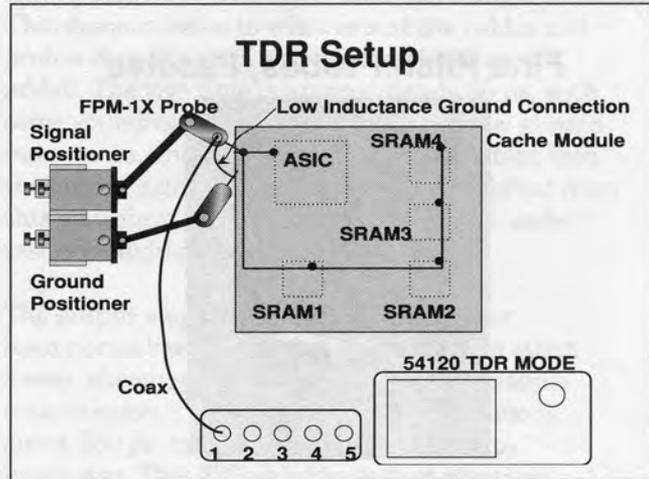
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #14



The cache module is laid out entirely in 50 Ohm microstrip for maximum signal integrity and minimum crosstalk. The clock line is the longest line on the module, so represents a worst-case line. The clock signal enters the module at the top left, connects to the cache controller ASIC, then proceeds to each of the 4 clocked 62486 SRAMs. The worst case access time is from SRAM4, as it receives the clock signal last.

Slide #15



This is a classic TDR setup, using an HP 54120 type instrument. Channel 1 outputs a 35 ps step, then looks for the reflected waveform which is displayed on the screen. Note that none of the ICs or other components are on the module during this measurement.

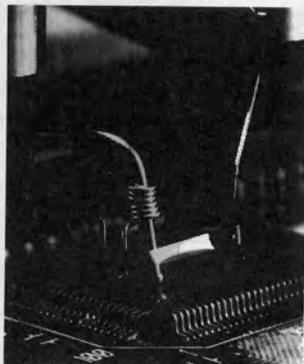
HP Equipment: HP 54120
Cascade Equipment: FPM-1X probe, FPD positioner, MTS-2000 base, Surrogate Chip Test Substrate.

The FPM-1X 50 Ohm probe is designed to work with the HP 54120 series instruments. It features a low-inductance, separately positionable ground. The surrogate chip includes the calibration elements (short, 50 Ohm resistors, and through connections) needed by the HP 54120 for calibration and normalization. The HP 54120 normalization function removes the cable and probe response from the measurement, providing higher accuracy measurements.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #16

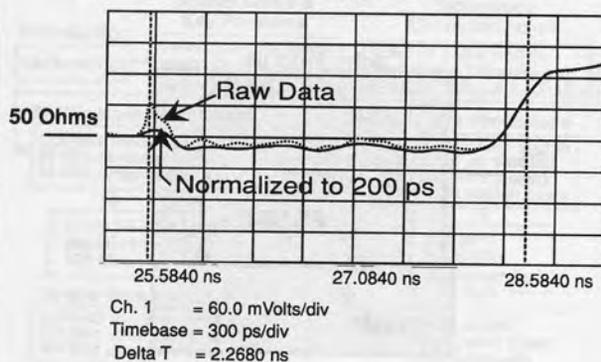
Fine Pitch Probes, Passive



This photo illustrates an FPM-1X 50 Ohm probe used in TDR/TDT measurements. The typical rise time is <math><60\text{ ps}</math> with 6 GHz bandwidth. The signal and ground contacts are independently positionable over a 2 inch range. Viewed and positioned through a stereo microscope, the contacts can be placed within a few microns. The 10X, 20X, and 100X FPM resistive divider probes have 20 ps rise times with 18 GHz bandwidth.

Slide #17

TDR Results

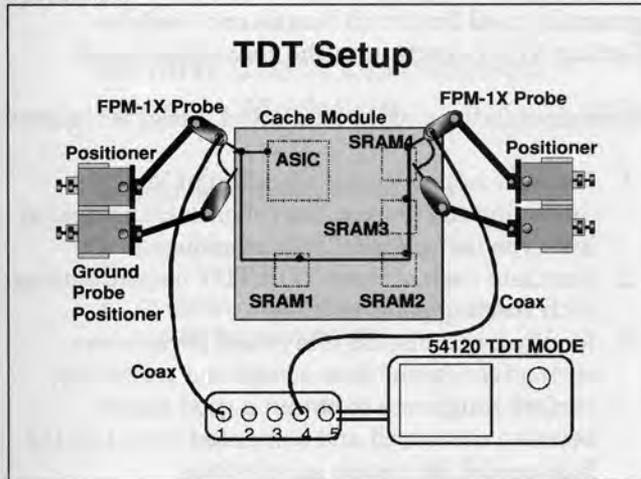


The raw data is what the HP 54120 displays as the TDR response with no calibration or normalization. Note, that even with a very low inductance probe, there is still a small inductive (positive-going) bump made visible by the fast 35 ps step. After the HP 54120 is calibrated (calibrated at minimum ground-to-signal spacing), the HP 54120 normalization algorithm will remove most of this effect. Cable skin effect and impedance inaccuracies are also removed.

The TDR provides very useful information. The electrical length is half the overall measured delay time (the step travels to the end and then back) and is 1.1 ns. Dividing by the physical length gives a propagation delay of 0.14 ns/cm. A perfect 50 Ohm line would have a flat horizontal response at the 50 Ohm point (see marker), so the measured line is a little less than 50 Ohms, or approximately 47 Ohm. Note the 47 Ohm response is fairly flat, indicating consistent transmission line behavior throughout the line, and little series loss. Series loss would appear as a slowly rising response, and would be modeled with a lossy transmission line model. Note this measurement was made without a termination resistor (open circuit) at the end, so the trace shows an open at the end of the transmission line.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #18



The setup for TDT (time-domain transmission) measurements is similar to TDR, except an additional probe and positioner are required to measure the output signal. The 35 ps step is introduced into the module by the left probe, travels through the module, out the probe on the right, then finally to the HP 54120 input. Note that none of the ICs or other components are mounted on the module during this measurement.
 HP Equipment: HP 54120
 Cascade Equipment: FPM-1X probe (2), FPD positioner (2), MTS-2000 base, Surrogate Chip

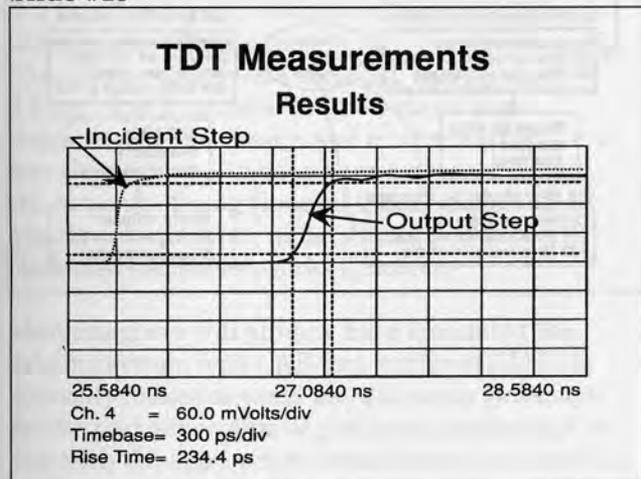
through connection on the Surrogate Chip*. This demonstrates to what extent the cables and probes slow the step, and what distortions are added. The rise time is approximately 50 ps, with some rounding of the corners due to probe ground inductance, and cable skin effect. The extent that the output signal from the module is different from this waveform is the degree to which line under test is changing the signal.

The output step (through the module) has been normalized, removing the cable skin effect losses, showing just the response of the module transmission line. The output step rise time is about 230 ps, much slower than the <50 ps input step. This 230 ps is the fastest rise time you can transmit through this line. For this paper, 300 ps will be used as the fastest edge the cache clock line will transmit, giving some margin for process variations.

What is causing the rise time degradation? Skin effect is probably playing a part. The classic skin effect degraded step is one where the initial 50-80% of the step occurs quickly, and the remaining few percent dribbles up slowly. With tungsten metallization (used in ceramic modules), the classic skin effect is modified by the surface roughness of the metal. As the high-frequency currents retreat to the outer "skin" of the conductor, the series resistance is higher than expected because the surface is rougher than at the center of the conductor. This further degrades high frequency performance. The cache clock line DC resistance is 2 Ohms, and this may play a small part in degrading rise time performance.

*The Surrogate Chip, available from Cascade Microtech, contains 50 Ohm, open, short, and through calibration elements used by the HP 54120 during calibration.

Slide #19



On the left is a raw (not normalized) display of the two probes connected together with a very short

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #20

TDR/TDT Conclusions

- Ideal transmission line model OK for 100 MHz simulations (risetimes >500 ps)
- Impedance: ~47 Ohms
- Propagation delay: 0.14 ns/cm
- Use lossy skin effect model for rise times <230 ps

From the TDR/TDT data, you can conclude that a simple ideal transmission line model will provide adequate simulation accuracy for transmission and reflection at 100 MHz (transition times >500 ps). We ignore cross talk in this example, which could be modeled as coupled lines or mutual C's & L's (see HP 1992 High Speed Digital Symposium). The impedance is approximately 47 Ohms, very close to the design value of 50 Ohms, and the propagation delay is 0.14 ns/cm. The TDR response of the transmission line is very close to flat (horizontal), indicating a consistent impedance throughout the length of the line and little series loss.

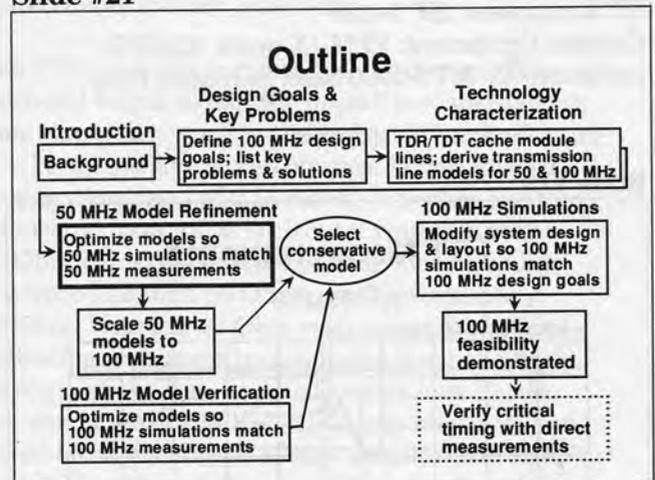
When simulating waveforms on this cache module with <300 ps transitions, a lossy, skin effect model is required. Otherwise, the simulator will predict faster performance than the transmission line will provide. Because time domain skin effect models are relatively new, we recommend that several line lengths be measured and the measurements be compared with modeled performance. With HSPICE, for example, you can input the physical values (resistivity, physical dimensions, dielectric constants) from which it computes the electrical performance. The model elements may need fine-tuning to precisely match measurements over the desired range of transmission line lengths and impedances. Recall, also, that buried lines will exhibit differing behavior than surface lines, so may require somewhat different model parameters. When using

HSPICE, the lossy U model, (PLEV=1, ELEV=1) is generally used for circuit boards and modules. Setting NLAY=2 turns on the skin effect model.

Recommendations when skin effect model is required:

1. Measure approximately 3 different lengths (including the longest line) of each transmission line type using TDR/TDT techniques.
2. Simulate each of these TDR/TDT measurements with starting point models.
3. Modify transmission line model parameters, such as conductor dimensions and resistivity, surface roughness to obtain a good match between measured and simulated data. Use the final model for system simulations.
4. Complex transmission line behavior may require a network analyzer to sort out what is really happening (call GigaTest Labs for example case studies at (408) 996-7500).

Slide #21



Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #22

50 MHz Clock Distribution Modeling

- Measure 50 MHz product
- Simulate 50 MHz product
- Compare and optimize clock driver model for best match

After the transmission line characterization was complete, and a model selected and calibrated, the entire 50 MHz clock distribution network was measured and compared with simulations. The sequence was:

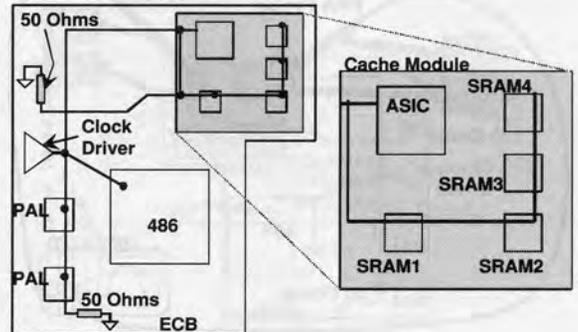
1. Measure clock performance of the current 50 MHz system
2. Compare with the 50 MHz simulations
3. Optimize the clock driver models obtaining a good match between measurements and simulations.

The critical areas to match were 1) clock transitions, 2) delays between signals, and 3) large anomalies due to reflections. In practice it takes several iterations to obtain a good match. For this example, 10 to 15 iterations were required, taking about 8 hours. It is important to understand basic transmission line theory, and have some feel how changing the various parameters will affect the response. Going through this exercise will greatly enhance your understanding of how clock distribution networks function.

Most designers will already have simulated the existing system on a CAE tool, such as QUAD. It generally makes sense to use this same simulator for the next generation of products, assuming it can deal with the appropriate transmission line models. In this example, HSPICE was used as the simulator because it is widely used, and the methods can apply to virtually any simulator.

Slide #23

50 MHz Clock Distribution

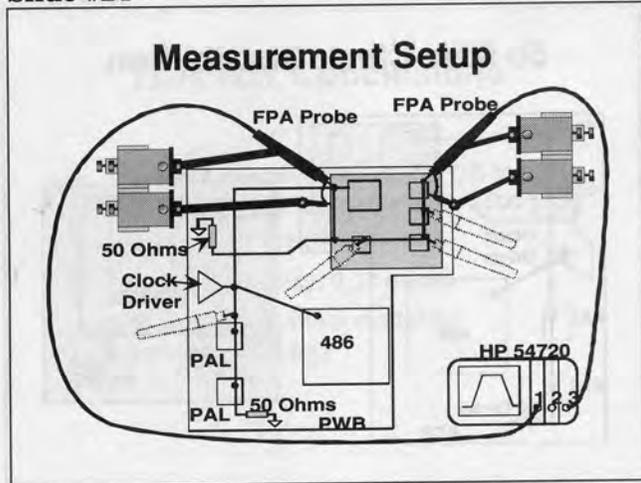


The clock distribution circuit is relatively straightforward, with the clock generator and buffer located on the printed wiring board (PWB), and the clock signal distributed to the 486, the PALs, and then to the cache module. On the cache module, the clock signal is distributed first to the ASIC and then to the 4 SRAMs. The clock distribution circuit is laid out as nominally 50 Ohm transmission lines both on the PWB and on the module. The circuit is terminated with 50 Ohm resistors in two locations, as shown. This presents a 25 Ohm DC load to the clock driver.

Because reflections can travel throughout all the clock tree, the entire clock distribution was simulated, rather than just connecting a ramp function to the cache module input. Ideal 50 Ohm transmission line elements were used to model the circuit board traces, with propagation delays of 0.07 ns/cm.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #24



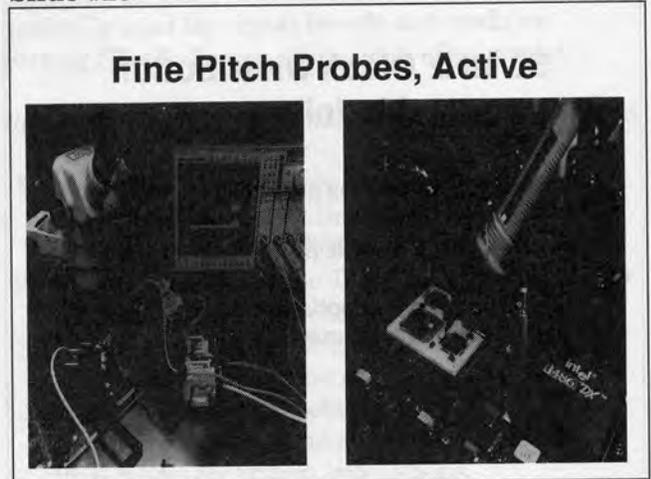
First, the measurements on a functioning 50 MHz system. An FPA active probe was positioned on the clock line at the ASIC, and the other probe moved to each of the SRAMs as well as another node at PAL1. The waveforms were saved as .PCX files and as .TXT files on a 3.5 inch floppy disk. The PCX files are basically a screen capture bit map. The .TXT files is an ASCII listing of the waveforms, and can be imported into a spreadsheet such as EXCEL.

HP Equipment: HP 54720 4-GSa/s Oscilloscope (1.1 GHz). Cascade Equipment: FPA fine pitch active probe (2), FPD positioner (2), MTS-1000 base unit

The HP 54720 was chosen for several reasons:

1. Sufficient bandwidth (1.1 GHz, rise time 320 ps) for accurately measuring a 100 MHz system (rise time approximately 1ns).
2. The above mentioned capability to save waveforms to disk is very useful when optimizing models and documenting results.
3. The capability of digitizing a single shot event, and triggering the scope with the logic function of several events. A typical example is looking at a burst mode read access time. This event occurs infrequently and is synchronized to the main clock, the read/write signal, and second level cache controller signals. These signals are logically combined in the HP 54720 to trigger the single shot data acquisition.

Slide #25

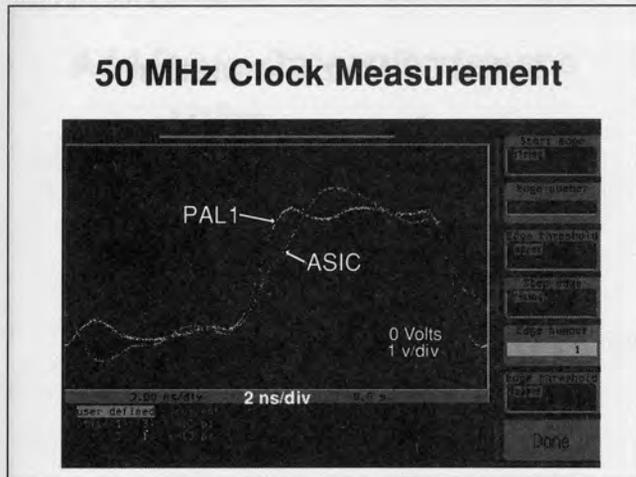


The left photo shows the measurement setup used including the active probes.

The probe on the right is based on the HP 54701 active probe which has been adapted to work with a positioner and a low inductance ground lead. Typical performance is 140 ps rise time, 100 kohm input resistance, and 0.6 pF input capacitance. The vacuum mount positioners attach to any solid base. The articulating arms provide a wide range of gross movements, allowing the user to roughly place the probes in the desired position. Fine positioning is then accomplished by adjusting the precision lead screw x, y, z positioners.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

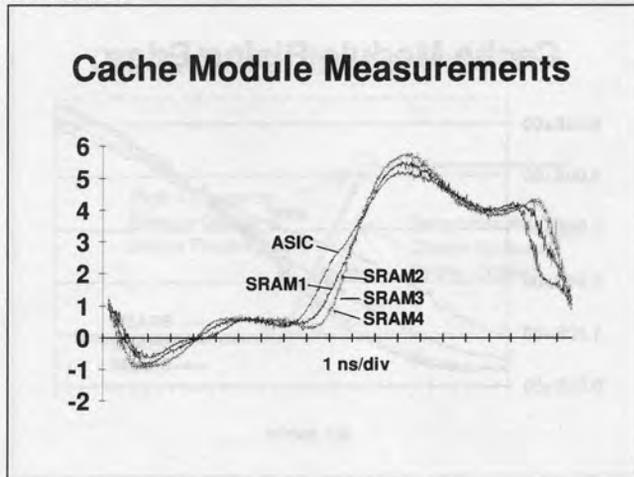
Slide #26



The clock at PAL1 is close to the clock driver, so the waveform looks good, with clean edges and little ringing. Note there is more ringing in the low level than at the high level, indicating the driver output impedance is not the same at both levels. By the time the clock signal reaches the module (ASIC), the effect of reflections are apparent, with more overshoot, and inflections beginning to appear in the transitions. This is an acceptable clock signal because the rising and falling edges are monotonic; the overshoot meets requirements (<0.6 volts above and below the rails), while the edge speeds, and the high & low times meet the 486 requirements.

Measurement system accuracy: Using the HP 54720 scope with a 4-GSa/s plug in, the measurement system rise time is 349 ps. This means that you can measure 1048 ps rise times with 5% accuracy and 760 ps edges with 10% accuracy. Conclusion: the above waveforms are being measured with less than 5% error.

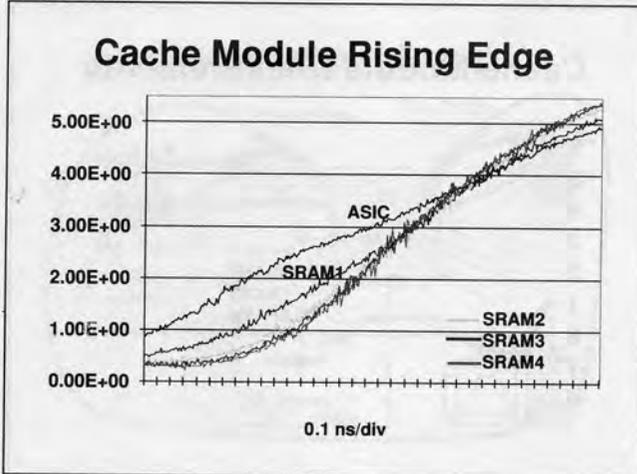
Slide #27



The HP 54720 waveforms were also saved as .TXT files, imported into EXCEL, charted, and saved as an EXCEL chart. One entire clock cycle is shown here, detailing the cache module waveforms at the ASIC, SRAM1, SRAM2, SRAM3, SRAM4.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #28

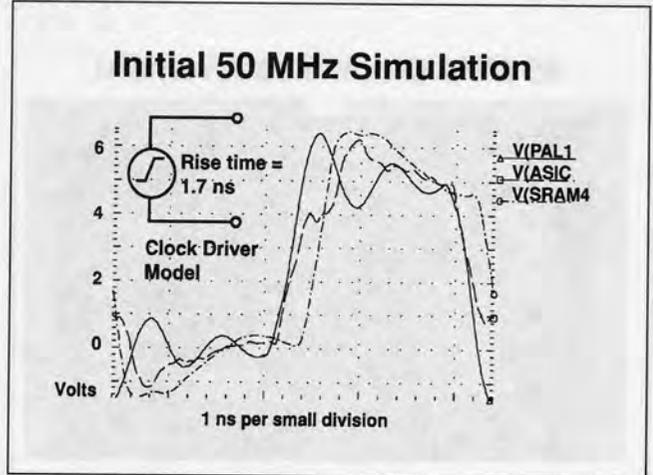


The positive going edge shown here is an expanded view of the same data shown in the previous slide "Cache Module Measurements." Within Excel, a smaller range was selected for each series, so that just the rising edge could be examined in detail. Note that because of the reflections, the waveform slopes are different, particularly V(ASIC), which has an inflection point where the slope changes. With faster clock edges, this inflection may degrade into a potential double clocking site.

The rising edge timing is defined generally at 2.0 volts for clock signals. Note that several of the signals have different slopes. This results in timing delay variations between signals if the voltage thresholds shift due to ground or power bounce.

When using the HP 54720 scope to measure time delay between two waveforms (at 2.0 volts for the rising edge, for example), the waveform noise will cause variation between successive measurements. The solution is to use the statistical functions within the HP 54720, giving a mean delay value and standard deviation.

Slide #29



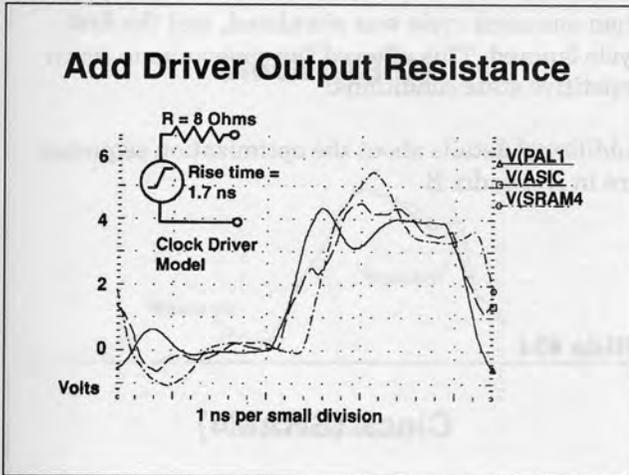
This is the initial 50 MHz simulation, using ideal transmission line models, capacitors & ESD clamp diodes for input gates. The clock driver is modeled as an ideal ramp, abruptly turning on at time equal to zero. Comparing these waveforms with the measured waveforms in slides "50 MHz Clock Measurement" and "Cache Module Measurements," there is a problem here. Note the excessive high frequency ringing, inflections, and other waveform anomalies not present in the measurements.

Each of the gate inputs is modeled as a capacitor with resistive ESD clamp diodes to V_{dd} and V_{ss}. These values were obtained from the manufacturer's specifications and from manufacturer's general technology handbooks. The PWB clock lines are modeled as ideal 50 Ohm transmission lines with propagation delay of 0.07 ns/cm. The cache module clock lines are modeled as 47 Ohm ideal transmission lines, with propagation delay of 0.14 ns/cm.

From the 50 MHz clock driver specification, the rise and fall times are 1.7 ns (max) into 50 pF. So the initial driver model is an ideal ramp, with 1.7 ns rise and fall times, and no output resistance.

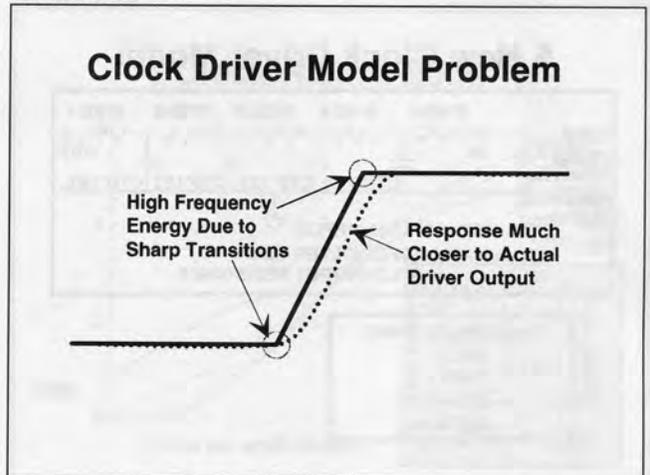
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #30



Real clock drivers have nonzero output resistance. From the 50 MHz clock driver specification, the rise and fall times are 1.7 ns (max) into 50 pF. Using $I = C (dv/dt)$, $I_{max} = 50 \text{ E-12} * 4\text{V}/1.7 \text{ ns} = 0.12 \text{ A}$, or assuming the driver ($V_{dd} = 5 \text{ V}$) meets specs at $V_{out} = 4.0 \text{ volts}$, then $R_{out} \sim 1\text{V}/0.12\text{A} = 8 \text{ Ohms}$. The above simulation is the same as the previous one, except the clock driver is now an ideal ramp in series with 8 Ohms. The PAL1 waveform has less ringing and somewhat better matches the measured waveform. V(ASIC), for example, still has non-monotonic edges. Something still is not right. (The measured waveforms, for visual comparison, are in slide #26 "50 MHz Clock Measurement" and in slide #27 "Cache Module Measurements.")

Slide #31

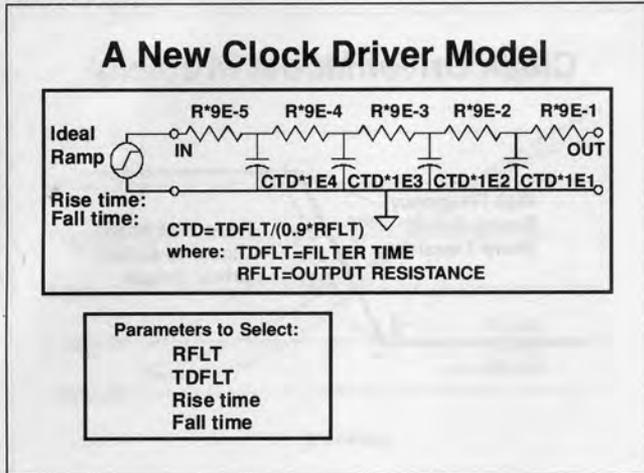


Using an idealized ramp signal to model a clock driver causes two errors: 1), is that an ideal ramp waveform contains high frequency energy due to the sharp transitions which is not present in the driver being modeled; and 2), when using ideal transmission line models, the additional high-frequency energy will be transmitted and reflected without loss. Real transmission lines will attenuate higher frequency energy more than lower frequency energy because of skin effect losses.

Ideally, you may have an accurate vendor supplied clock driver model. In many cases, they are not available, particularly when designing system speeds beyond those with currently available parts. In that case, you can construct a simple driver model (described in the next slide) provided in HSPICE (see chapter 3, *HSPICE Manual*, Meta-Software, 1300 White Oaks Road, Campbell, CA 95008).

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #32



Using a 4 stage RC output filter recommended by Meta-Software, an output waveform much closer to that actually generated by the clock driver will be demonstrated. This simple circuit is available as a macro in HSPICE.

Slide #33

Iterating a Clock Driver Model

- Select initial RFLT, TDFLT, t-rise, t-fall from data sheet
- Choose RFLT for best high level match at driver
- Choose t-rise & t-fall to match measured values
- Choose TDFLT to give best inflection & high frequency ringing match to measured

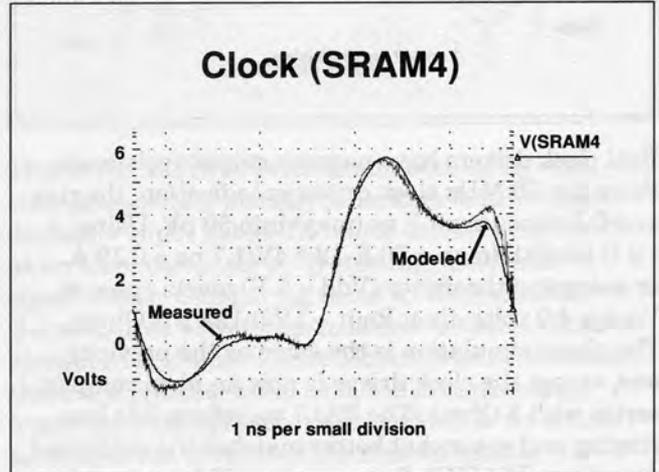
50 MHz Driver Model:
RFLT = 5 Ohms
TDFLT = 250 ps
Rise time = 1.9 ns
Fall time = 1.3 ns

Deriving a good clock driver model was an iterative optimization process, involving 10 to 15 iterations through HSPICE. This is less work than it first appears because the HSPICE simulations of the clock network including inputs, diode clamps, transmission lines, and clock driver ran in about

50 seconds on a 486, 33-MHz PC. Note that more than one clock cycle was simulated, and the first cycle ignored. This allowed the capacitors to reach repetitive state conditions.

Additional details about the optimization sequence are in Appendix B.

Slide #34



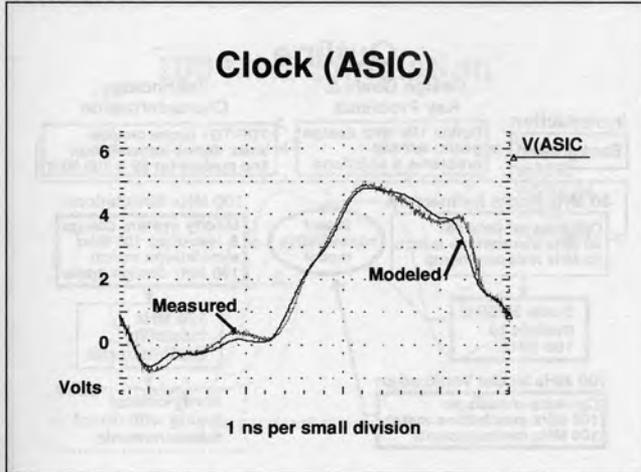
Using the clock driver model described, the simulated clock signal at SRAM4 matches the measured signal fairly well:

- RFLT = 5 Ohms
- TDFLT = 250 ps
- rise time = 1.9 ns
- fall time = 1.3 ns

Note also, the edges at 0.8 and 2.0 volts (clock low and high level, respectively) match very well and all measured anomalies are present in the simulation. The HSPICE plots were saved as HPGL files then imported into Persuasion. The HP 54720 measurements were saved as a .TXT file, converted in Excel to an Excel chart, and imported into Persuasion, and overlaid on the HSPICE waveform.

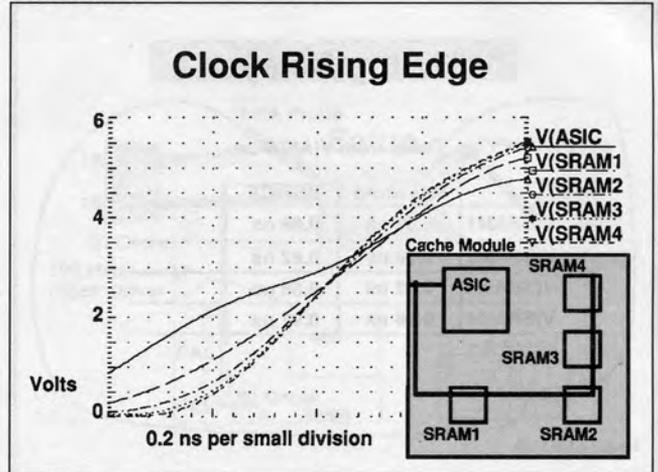
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #35



Again, using the clock driver model just described, the simulated clock signal at the cache controller ASIC matches the measured data fairly well. Note that all the major features are present in the HSPICE simulation.

Slide #36



Shown here is an expanded view of the simulated clock rising edge on the cache module. Note the differing slopes of the waveforms. At first glance, magic appears to be happening here. The electrical length from the ASIC to SRAM4 is 1.1 ns, yet the delay measured here is approximately 0.85 ns. What is happening here is the unterminated length of transmission line from SRAM2 to SRAM4 has an electrical delay of 0.58 ns, about 1/3 the clock driver transition time (1.5 ns typical). The reflections from this unterminated line result in differing transition slopes at SRAM1 through SRAM4. Used carefully, this technique can be useful. The problems are that anything which alters the voltage thresholds in each of these ICs will change the timing delay. Also, this technique will be sensitive to clock driver edge rates.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #37

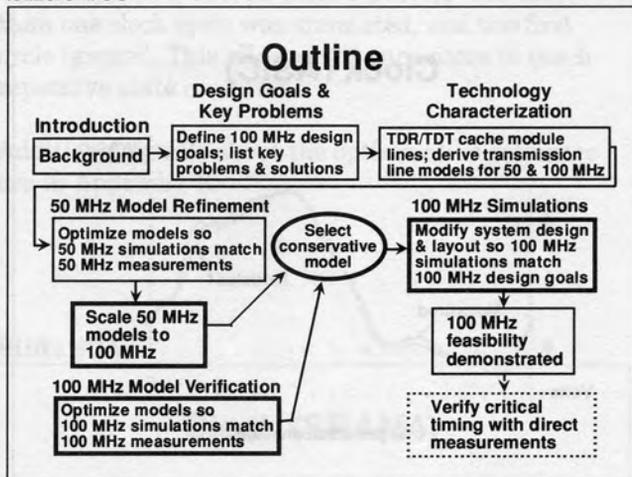
Timing Delays

Delay from V(ASIC to:

	Measured	HSPICE
V(SRAM1)	0.62 ns	0.69 ns
V(SRAM2)	0.78 ns	0.82 ns
V(SRAM3)	0.82 ns	0.86 ns
V(SRAM4)	0.86 ns	0.91 ns

Delay times were measured at 2.0 volts for rising edges, and 0.8 volts for falling edges. Only the rising delays are shown here, as the falling edges fell on top of each other.

Slide #38

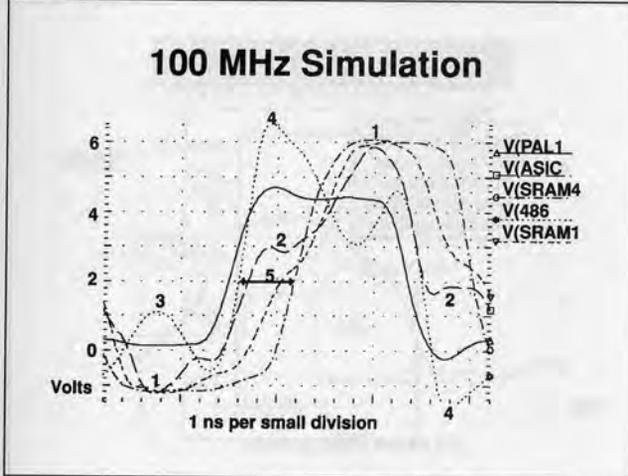


Now that you have confidence in the 50 MHz simulations, you can scale the clock driver model for 100 MHz operation, and simulate the clock distribution network at 100 MHz. Recall that our 100 MHz design goals were:

- 100 MHz Goals**
- SRAM4 to 486 clock skew <0.65 ns
 - SRAM4 data to 486 TOF <0.35 ns
 - Clock over/under shoot within 0.6 volts of the rails
 - No double clocking edges

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #39



Using the 50 MHz clock driver model, doubling the clock frequency, and making the following changes gives the simulation result.

- clock rise/fall times:
1/2 the 50 MHz values
- clock driver output resistance:
1/2 the 50 MHz value
- clock driver filter time:
1/2 the 50 MHz value

This gives a 100 MHz clock driver model:

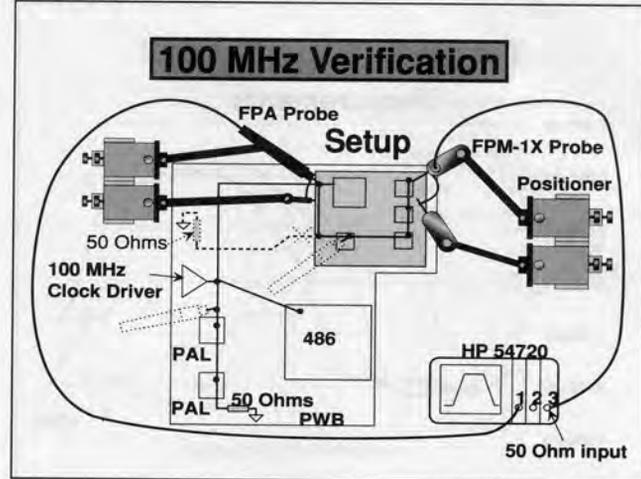
- RFLT = 2.5 Ohms
- TDFLT = 125 ps
- Rise time = 0.95 ns
- Fall time = 0.65 ns

The clock driver output resistance was halved, because if the transition times are cut in half, then the $C \cdot (dv/dt)$ current will double. We assumed that the capacitive loading at 100 MHz would stay the same as the 50 MHz version.

Note the cache module signals are now slamming into the ESD clamp diodes because the open stub causes larger reflections from the faster edges. The slight inflections in V(ASIC) transitions at 50 MHz are much larger and no longer monotonic. The clock waveform at the 486 has a potential problem site after the falling edge.

Note the rising edge (2.0 V) clock skew from the 486 to SRAM4 is 1.3 ns and the desired value is 0.65 ns for a 100 MHz system.

Slide #40

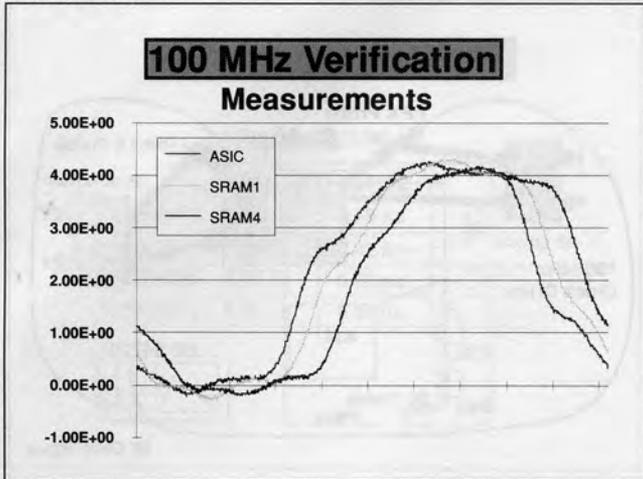


Before proceeding further with the 100 MHz feasibility study, it is useful to take time out and verify that the 100 MHz driver, transmission line, and gate input models are valid at 100 MHz. The 50 MHz xtal-driver was removed and replaced with a 100 MHz xtal-driver. While the system will not function at 100 MHz, the clock signals will still propagate throughout the clock network and can be measured and compared with simulations. One other change was made to the circuit in anticipation of a circuit improvement needed for 100 MHz operation. That was to disconnect the termination line at the spot marked "X", and to terminate the cache clock line at SRAM4. Unfortunately, no provisions were available on the module to connect a 50 Ohm chip resistor. Instead, an FPM-1X 50 Ohm probe provided termination. At first glance, this may not seem right. However, think of a rising edge propagating through the cache clock line, and as it reaches the FPM probe, propagates to the probe with minimal reflection. The reflection will be: $(50 - 47)/(50 + 47) = 0.03$ or 3%. The waveform will then travel through the probe and coax up to the 50 Ohm resistor in the HP 54720, and the resistor will absorb all the energy with no reflection, thus effectively terminating the cache clock line at SRAM4.

The signal from SRAM4 triggered the scope, thus assuring all the other waveforms measured have the same time reference.

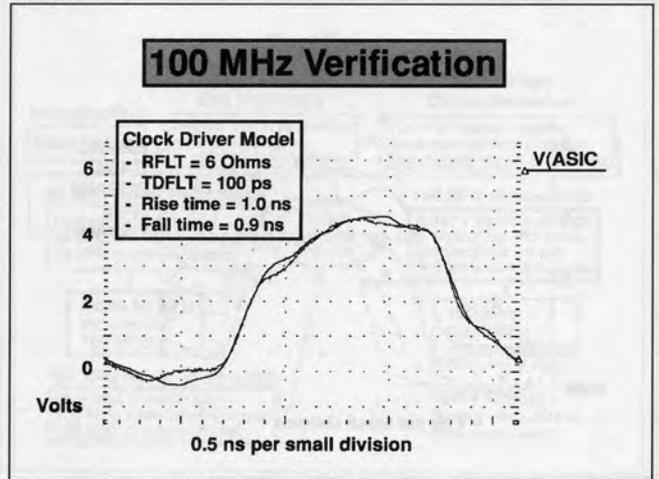
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #41



Using an FPA active probe, the cache module waveforms were measured, saved as .TXT files, placed into Excel, converted into a chart, copied from the work sheet window into the clipboard, and pasted into Persuasion. Note the reduced ringing because the 50 Ohm termination reduces reflections to a mere 3%.

Slide #42



Now that you have the 100 MHz measurements, can you get the simulator to match them? Using the same transmission line models and input gate models from the 50 and 100 MHz simulations and changing the netlist to tell the simulator the cache line is now terminated at SRAM4, instead of the previous arrangement, a good match can be obtained with the following clock driver values:

- RFLT = 6 Ohms
- TDFLT = 100 ps
- Rise time = 1.0 ns
- Fall time = 0.9 ns

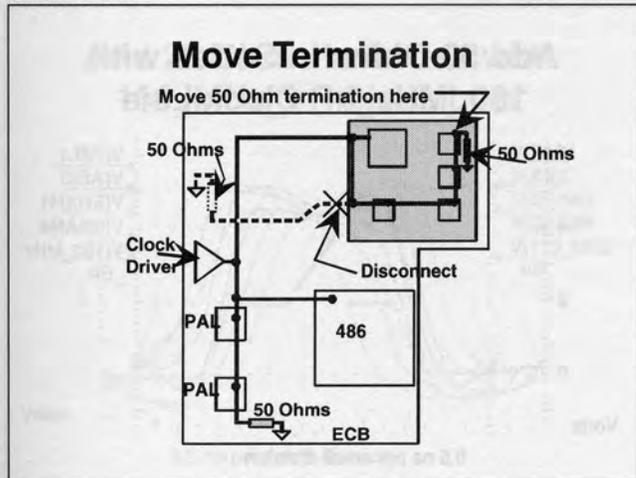
Shown above is simulated V(ASIC) overlaid by the measured V(ASIC). The question now is, for the 100 MHz simulations, do you use this driver model or the scaled 50 MHz driver model? Because the clock driver has not yet been chosen for 100 MHz design, you should pick the most conservative driver model, and that would be the model with the fastest transitions. So for the rest of the 100 MHz simulations the following clock driver model will be used:

- 100 MHz Clock Driver Model**

 - RFLT = 2.5 Ohms
 - TDFLT = 125 ps
 - Rise time = 0.95 ns
 - Fall time = 0.65 ns

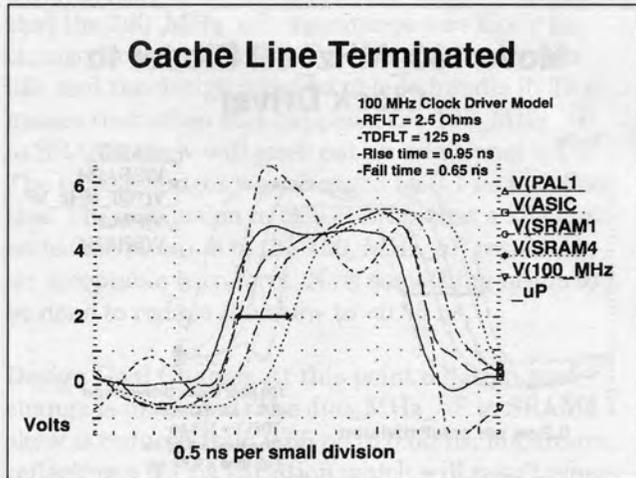
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #43



Back to the 100 MHz feasibility study. The first recommended change is to move the 50 Ohm termination from its present site to SRAM4, and to disconnect the transmission line connecting to the current termination resistor. This will eliminate the open stub on the cache module and minimize reflections caused by an unterminated transmission line whose electrical length is longer than the signal transition times.

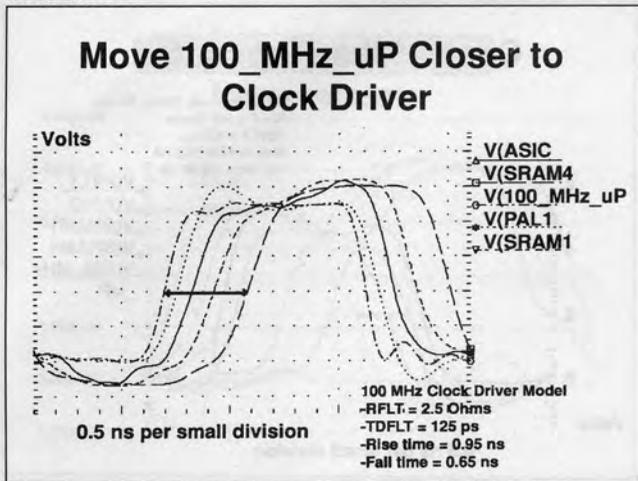
Slide #44



Eliminating the open stub cleans up the waveforms on the cache module as expected. The 100_MHz_uP waveform still does not meet the overshoot/undershoot requirements. One way to deal with it is to move the clock driver closer to the 100_MHz_uP, where the low impedance of the clock driver will control overshoot and ringing better. The problem with that is the 100_MHz_uP clock to SRAM4 clock skew goal is 0.65 ns. It now is 1.3 ns and will get larger if the clock to 100_MHz_uP distance is reduced (see layout diagram).

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

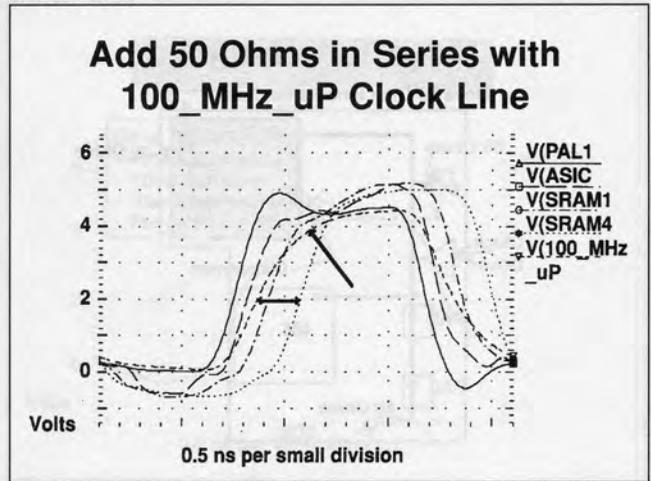
Slide #45



Reducing the clock driver to 100_MHz_uP distance from 3.5 cm to 0.5 cm reduces the overshoot and ringing at the 100_MHz_uP, but increases the 100_MHz_uP to SRAM4 skew to 1.9 ns. Recall that our design goal is 0.65 ns. There are a couple of options at this point. One is to retain the original 3.5 cm clock driver to 100_MHz_uP distance and use a series termination. Another option is to move the cache module much closer to the 100_MHz_uP and drive the module clock line in the center (or close to the electrical center) thus reducing the clock delay to SRAM4 and the ASIC.

Adding a 50 Ohm series resistor will reduce the ringing and add some delay: $R * C = (50 \text{ Ohms} * 15 \text{ pF}) = 0.75 \text{ ns}$.

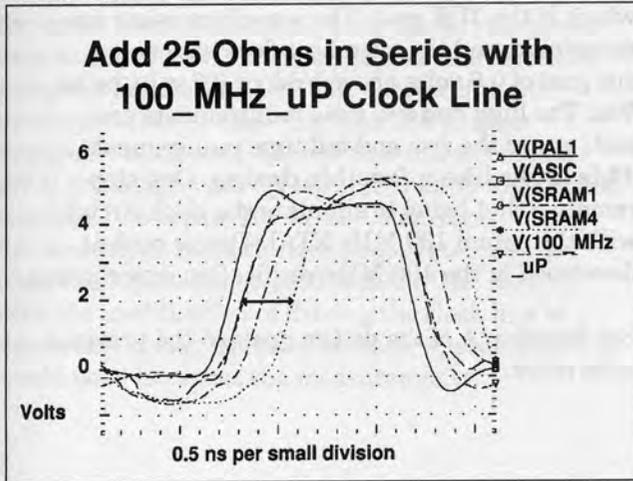
Slide #46



The 50 Ohm resistor in series with the 100_MHz_uP clock-in solves the overshoot/undershoot problem. The 100_MHz_uP to SRAM4 skew is now down to 1.1 ns, an improvement, but still not good enough. The series resistor also appears to be too large, giving a distinct RC exponential shape to the 100_MHz_uP clock signal. This, in itself is not a problem. However, input gate capacitance values are not well specified by IC manufacturers. If the design depends on a specific RC value, the circuit may malfunction if the input capacitance changes. So try a smaller resistor, say 25 Ohms.

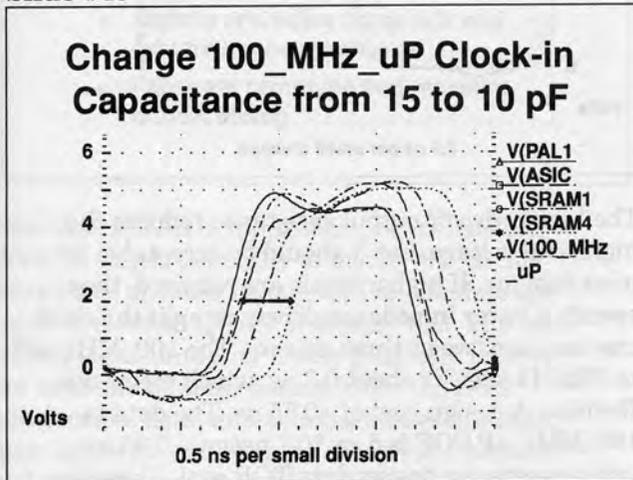
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #47



Very acceptable waveforms, 100_MHz_uP to SRAM4 skew is 1.4 ns. For all the simulations so far, the 100_MHz_uP clock-in capacitance value has been 15pF. A 100 MHz processor will probably have smaller capacitance, so let's try 10 pF and get a feel how sensitive a parameter it is.

Slide #48



Changing the 100_MHz_uP input capacitance from 15 pF to 10 pF changes the 100_MHz_uP to SRAM4 skew from 1.4 ns to 1.5 ns. A change of 0.1 ns, a very acceptable value but it is in the direction which

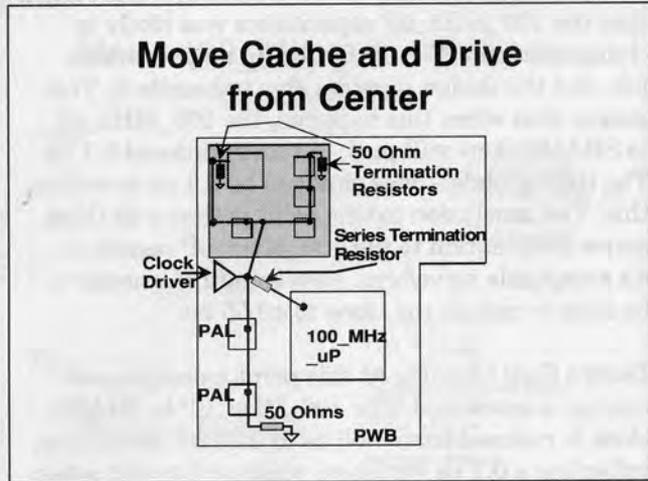
increases skew. In this feasibility study we assumed that the 100_MHz_uP capacitance was likely to change from 15 fF to 10 fF during its production life, and the design must be able to handle it. This means that when this happens, the 100_MHz_uP to SRAM4 skew will push out an additional 0.1 ns. The timing budget was changed by 0.1 ns to reflect this. The conclusion to this point is that a 25 Ohm series termination to the 100_MHz_uP results in an acceptable waveform. Now something needs to be done to reduce the skew to <0.55 ns.

Design Goal Change: At this point a design goal change is instituted. The 100_MHz_uP to SRAM4 skew is reduced from 0.65 ns to 0.55 ns, maximum, reflecting a 0.1 ns variation which will result when the clock-in capacitance is reduced from 15 fF to 10 fF by the manufacturer and we are not notified. Note that input capacitance is generally specified as a typical specification in any event.

- 100 MHz Goals (Rev. A):
- Demonstrate 100 MHz 1 cycle burst mode
 - SRAM4 to 100_MHz_uP data TOF <0.35 ns
 - 100_MHz_uP to SRAM4 clock skew <0.55 ns
- Demonstrate acceptable clock signal integrity
- monotonic clock edges
 - clock max/min not exceed 0.6 V
 - outside power supply rails
 - rise/fall time < 1 ns (0.8 to 2.0 volts)
 - min. high time & low time = 3.5 ns each

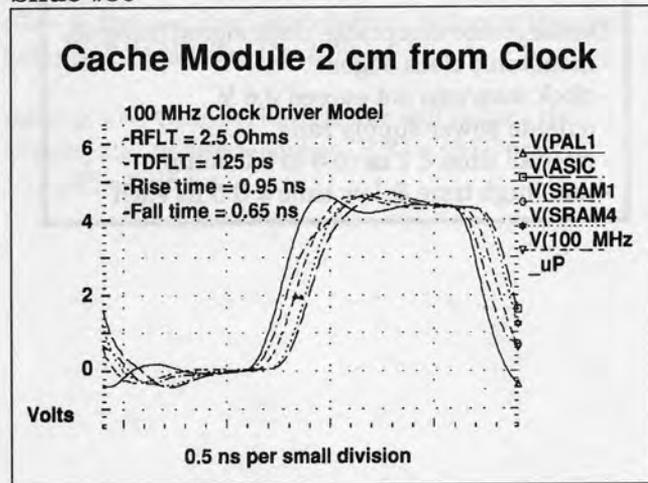
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #49



First, move the cache module physically closer to the clock driver and the 100_MHz_uP. Then connect the clock to the center of the cache module clock line, terminating each end with 50 Ohm resistors at the ASIC and at SRAM4.

Slide #50

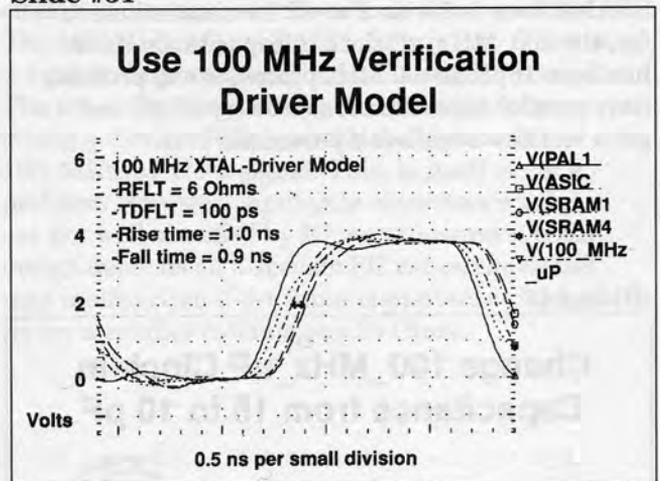


Moving the cache module next to the 100_MHz_uP (and driving the clock at the center of the cache clock line), reduces the 100_MHz_uP to SRAM4 skew to 0.3 ns which exceeds our goal of 0.55 ns. Because the cache module is now about 5 cm from the 100_MHz_uP, the SRAM4 data to the

100_MHz_uP is $< 5 \text{ cm} * 0.07 \text{ ns/cm} = 0.35 \text{ ns}$ which is the TOF goal. The waveform edges are monotonic, and the overshoot does not exceed our goal of 0.6 volts above Vdd, or 0.6 volts below Vss. The high and low time requirements are met, as are the rise and fall time requirements. **This looks like a feasible design.** One step remains: that being to simulate the clock circuit with the actual 100 MHz XTAL-Driver model developed in the 100 MHz verification experiments.

See Revision A of the design goals in the previous slide notes.

Slide #51



The higher driver output resistance reduces the high level voltage, but it should be acceptable for most designs. If higher levels are required, then specify a lower impedance driver, or split the clock tree and use two or three drivers. The 100_MHz_uP to SRAM4 skew is about 0.3 ns, which meets our Revision A design goal of $< 0.55 \text{ ns}$. The data to 100_MHz_uP TOF is $5 \text{ cm} * 0.7 \text{ ns/cm} = 0.35 \text{ ns}$, which meets the design data TOF goal. Therefore, this layout looks feasible for 100 MHz. Much work, however, remains to be done.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

An interesting aside is that now you are driving the cache module clock line at its center, the worst-case length is now 1/2 the length we characterized using the HP 54120 TDR. The shorter line will propagate a faster edge than the 230 ps edge previously measured. So how fast will the module go? If you conservatively assume that 50 MHz speed requires 1 ns edges, 100 MHz 0.5 ns edges, then the module, as characterized propagated 230 ps edge, then it should be good to 200 MHz. Now with the modification of driving the clock line in the center, it may function at faster speeds. One would need to repeat the measurements.

Slide #52

Next Steps

- 100 MHz operation feasible
- Modify layout and re-simulate complete system
- Simulate signal integrity on data and address busses
- Identify and solve cross talk and ground bounce problems
- Fabricate prototype and measure critical timing

Now that 100 MHz operation appears feasible, the next step is generally to modify the circuit netlist and physical layout, and re-simulate the entire system on QUAD or similar tool. The next critical signal integrity issues generally involve the data and address busses. Once the clock distribution, data, and address busses are solid, then issues related to crosstalk, ground bounce, and simultaneous switching noise can be identified and resolved.

After functioning prototypes have been built, then critical timing measurements can be performed. In this study, the timing of interest was burst mode data transfer between the microprocessor and the cache module. Observing and recording burst mode timing is ideally suited to the HP 54720. Because it can capture single shot events, you can look at individual timing, including how noise influences individual events, rather than having to look at an average of many events, which is what sampling scopes provide. An additional important feature of the HP 54720 is the ability to trigger from the logical combination of several signals. When looking at burst mode read timing, for example, it is desired to trigger the scope only when several events are true (positive clock edge, write enable high, and possibly other signals generated by the second level cache controller). The HP 54720 can do that, greatly simplifying these type of complex measurements.

Slide #53

Outline

- Introduction and Background
- Design goals and Key problems
- Technology characterization & 50 MHz Model refinement
- 100 MHz Simulation and Model verification
- Recommendations & Summary

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #54

Steps to Double the Clock Rate

- Characterize critical module transmission line
- Measure existing 50 MHz product
- Refine 50 MHz models to match measured performance
- Used refined models, simulate 100 MHz performance
- Compare 100 MHz simulations with 100 MHz measured data

Using an existing 50 MHz 486 (with a cache module) as a design example, the feasibility of using the cache module at 100 MHz was examined. After design goals were established:

1. The cache module transmission line characteristic were measured with an HP 54120 TDR instrument. The measurements established the maximum speed of the line and helped select the appropriate transmission line model for the desired edge speeds.
2. Clock line measurements were made with an HP 54720 and stored on disk. The measurement techniques were focused towards probing fine pitch (<20 mils) interconnects that are increasingly common on circuit boards and modules.
3. The 50 MHz simulator models were optimized to measurements of the 50 MHz performance.
4. The 50 MHz models were extrapolated to 100 MHz, and simulations run.
5. The 100 MHz simulations were partially verified by inserting a 100 MHz crystal-driver into the 50 MHz slot. Measurements were made and compared with simulations. Models were adjusted.
6. With additional confidence in the 100 MHz simulations, additional layout and circuit modifications were simulated, eventually meeting the 100 MHz design goals.

Slide #55

Conclusions

- 100 MHz feasibility demonstrated
- Cache module speed limit 230 ps (200 MHz)
- Clock driver model critical to accurate simulation. Edge speed critical to clock distribution performance
- Worst case T-line characterization using 54120 required to select right model

The example 486, 50-MHz system has been shown to be feasible at 100 MHz with a small number of layout and circuit changes. Much additional work remains, of course, for the final design, but overall clock timing and distribution looks OK.

The cache module speed limit is 230 ps edge speed, used as in the 50 MHz design. This corresponds roughly to 200 MHz. The propagation delay of 1.1 ns is a more limiting problem at 100 and 200 MHz (clock skew) than the edge speed degradation. If the module is modified by driving the clock line from the center (reducing the worst case prop delay to 0.55 ns) the maximum edge speed should increase. One should re-do the HP 54120 characterization measurements.

The clock driver model is critical to obtaining accurate clock distribution simulation. Ideal ramps generate high frequencies not present in the actual clock waveform. Small changes in the edge speeds can have a large impact on the performance of clock distribution networks with non-terminated lines, or large capacitive loads. Production products can fail if a clock driver with faster edges is used. Worst-case clock-driver models should include the fastest edges, as well as other variables.

Characterizing a transmission line with an HP 54120 provides the electrical length, propagation delay (knowing the physical length), impedance, maximum edge speed travelling through the line, series loss, and so helps determine which transmission line model should be used.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #56

Conclusions

- Previously impossible fine pitch measurements are now easy using Cascade FPA and FPM probes
- HP 54720 data acquisition and disk storage features very useful for model optimization
- With FPM-1X probes, one can easily inject signals or terminate fine pitch lines

Measurements which could not be made a year ago, are now easy using Cascade's new FPA (active probe) and FPM (passive probes). With their precise positioning ability (10 μm), users can now land these probes on fine pitch circuit board, module, and package structures.

The HP 54720, with its real-time digitizing capability and MS-DOS[®]-compatible waveform saving ability, is ideally suited for modeling and characterization activities. The ability to save a waveform and to then import it into a variety of MS-DOS and Windows applications and overlay simulations is very useful. Documentation of critical waveforms is now completely electronic.

A unique application of FPM-1X probes is to terminate 50 Ohm fine pitch transmission lines, using a 50 Ohm resistor on the probe, or the 50 Ohm load in an instrument. One can also use these probes to inject signals into fine pitch structures, performing in situ testing.

MS-DOS[®] is a U.S. registered trademark of Microsoft Corporation.

Slide #57

Recommended Resources

- HP:
 - 54120 series TDR
 - 54720 digitizing scope
- Cascade (503-626-8245):
 - FPM & FPA fine pitch probes
 - MTS-2000 fine pitch base units
 - Surrogate calibration chips
- Meta-Software (408-371-5100)
 - HSPICE simulator
- Conner Winfield:
 - Crystal-clock drivers

Slide #58

Recommended Resources Consultant Services, etc.

- HP SE help (phone #)
- Cascade App Engr: (503-626-8245)
- GigaTest Labs, Characterization Services (408-996-7500)
- NCR Microelectronics Products Division, Advanced Division (719-596-5795)
- Arthur Fraser (503-289-2637)

The authors gratefully acknowledge the invaluable help of Brad Frieden, of HP Colorado Springs, who successfully steered this paper between the shoals of disaster. Thanks for your comments, review, suggestions, and help with the equipment and measurements.

Thanks also to Art Porter of HP Colorado Springs, for helping us get on the air in one day with the HP 54720.

Thanks also to Wilton Hart, Tektronix, Beaverton, OR, (503-627-3035), for enlightening discussions on random SRAM data loss caused by charge injection due to clock waveform overshoot/undershoot.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #59

Appendix A Transmission Line Models

Appendix A. Transmission line models.

Recall that an ideal transmission line model assumes there is no loss, just time delay with no signal degradation. A lossy transmission line model adds frequency independent series and parallel resistive losses which are useful when simulating resistive conductors or situations with significant DC leakage between conductors. Frequency dependent losses are generally due to skin effect, so named because as frequency increases, the current flows in an increasingly thinner "skin" at the outer surface of the conductor. The degree of loss depends on the ratio of "skin depth" to the conductor thickness. Skin effect and other non-linear frequency dependent losses are difficult to model using time domain simulators. Impulse from HP and HSPICE from Meta-Software are two time domain simulators that feature lossy model which include skin effect.

In digital circuits, skin effect slows the signal transitions. In a typical skin effect degraded response, the first 50-80% of the edge occurs quickly, with the remaining signal slowly dribbling up to the final value. The term "dribble up" is often used describing this response. When selecting a worst case conductor for characterization, it is important to remember that skin effect loss is

roughly proportional to the square of the conductor length. This means that one should pick the longest line, all else being equal. Another factor is the conductor material. Conductors fabricated with high surface roughness may show increased skin effect losses compared with smoother conductors. As the current retreats to the rough conductor surface, less material is present to conduct the high frequency current, resulting in higher than expected losses. In this example, the cache module is fabricated with tungsten metallization, which has higher skin effect losses than copper conductors.

The general transmission line characterization guidelines are:

1. Select the longest high speed lines of each interconnect technology
2. Lines buried in a dielectric will have different characteristics than those on the surface
3. Conductor surface roughness and thickness affect skin effect losses
4. Dielectric losses may be a factor
5. Use TDR/TDT measurements to select and fine tune model

In general, nothing is better than direct measurements on actual transmission lines fabricated with the production processes. Complex models are useful in understanding the trade-offs between variables, but need to be calibrated with measurements for accurate simulations. There are too many variables to accurately model frequency dependent losses without actual measurements. A useful reference on these effects is "High Speed Digital Microprobing, Principles and Applications" from Cascade Microtech.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

Slide #60

Appendix B

Appendix B. Clock driver optimization sequence.

The clock driver model optimization sequence involved several iterations through HSPICE:

1. For this example, the starting point clock driver parameters were selected:
 - RFLT (driver output resistance) = 8 Ohms (calculated from spec sheet)
 - Rise time = Fall time = 1.7 ns (from spec sheet)
 - TDFLT (filter time constant) = 0.17 ns (start at 10% of rise time)
2. Using the clock distribution model described in the previous 2 slides and the initial filtered driver model, the clock tree was simulated using HSPICE.
3. RFLT (driver output resistance) was then chosen so the simulated clock output driver high level voltage matched the measured value. With the filter now in place, the high frequency ringing in the simulation was reduced, allowing a good estimate of the average high level output voltage. Minimum ringing was observed just before the transition to the low level voltage, and this value was used. A good match typically took 2 iterations through HSPICE. An output resistance of 5 Ohms was selected.
4. The rise and fall times were selected so that the simulated rise and fall times matched the measured values. Typically, waveforms electrically close to the driver output have cleaner edges, with fewer inflections and changing slopes. In this example, the cache module waveforms were of most interest, so V(ASIC) and V(SRAM4) was used. V(ASIC), however, changes slope during its transitions. What worked best, was to match to the initial, steeper slope. Note that when changing the rise and fall times in the simulation, the duty cycle must be compensated. An additional factor is that some clock drivers may not be operating at exactly 50% duty cycle. While this may not be a problem for the product, when overlaying waveforms (see below), a slight duty cycle difference is very apparent, which you may desire to compensate for by adjusting the simulated duty cycle value. In this example, the simulated duty cycle was adjusted to match the measured waveforms. 9 iterations through HSPICE were required to match the measured rise and fall times, and the duty cycle. A rise time of 1.9 ns and a fall time of 1.3 ns were chosen. When matching slopes and other waveform parameters, it was very useful to superimpose waveforms. The measured waveforms were saved as .TXT files, imported into Excel, converted into Excel charts, copied into the Windows clipboard (from the work sheet window), then pasted into Persuasion for this paper. The chart was then "ungrouped" and the cyan background deleted, then "re-grouped." Note that the .TXT files are large (>2000 elements), and the related charts may exceed the memory capabilities of many PCs when trying to Paste them into another program. Another option is to save the charts as Windows metafiles out of Excel, and Import them into other programs.

The HSPICE simulations were saved as HPGL printer outputs, renamed as *.plt files, and imported into Persuasion using the "Place" command.

Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

5. Finally, the TDFLT filter time was selected to match the transition inflections and degree of ringing. A value of 250 ps was chosen. 3 HSPICE iterations were required.
6. Because each of these parameters will partially interact with the others, an additional 2-3 iterations may be required to fine tune the over all fit. One should have the feel of an overall 5% fit, or better, at 2.0 volts on the rising edges, and 0.8 volts on the falling edges. All other waveform anomalies should be present in the simulation.

Note that some clock drivers may have significantly different high level and low level output resistances and may require a modified model.

Appendix B. Clock driver output resistance & capacitance (calculated from spec sheet)

For this example, the starting point clock driver parameters were as listed in Table 1. The clock driver model output resistance was involved in the HSPICE simulation.

1. For this example, the starting point clock driver parameters were as listed in Table 1. The clock driver model output resistance was involved in the HSPICE simulation.

2. The clock driver model output resistance was involved in the HSPICE simulation.

3. The clock driver model output resistance was involved in the HSPICE simulation.

4. The clock driver model output resistance was involved in the HSPICE simulation.

5. The clock driver model output resistance was involved in the HSPICE simulation.

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Clock Frequency Progression



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



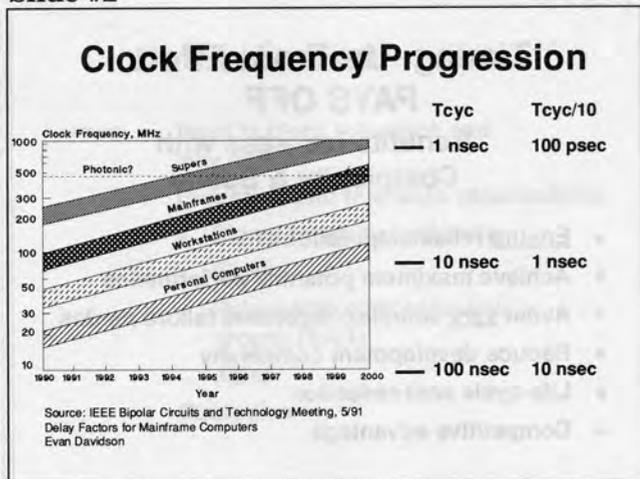
AMHERST SYSTEMS ASSOCIATES



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

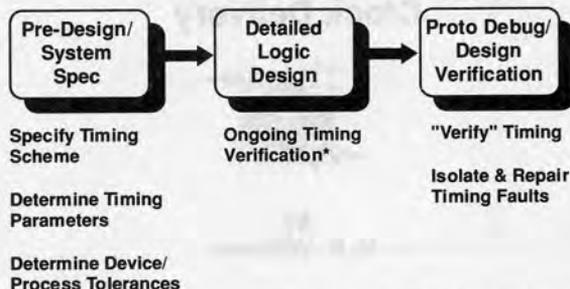
Benefits Increase with
Complexity & Speed

- **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



* TV Accuracy = (Timing Parameter Accuracy)

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

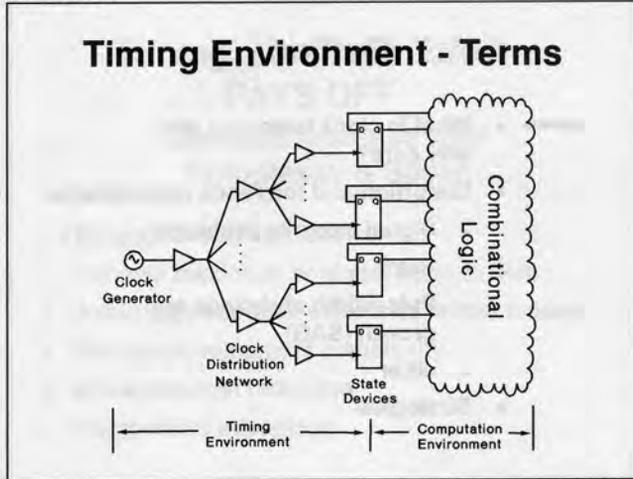
Outline

- ➔ • **What is clock tolerance and why care?**
- **Distortion and tolerance mechanisms**
 - **Signal integrity problems**
 - **Skew**
 - **Pulsewidth 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.

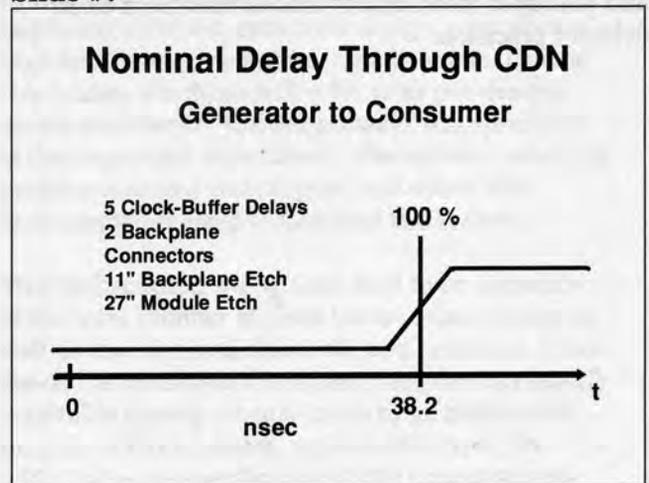
The clock generator supplies the signal whose edges dictate when things happen throughout the system. The generator determines the period, pulsewidth, number of phases and their relative edge placement. There are a large number of generator attributes to be specified in a typical design (for example, 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.

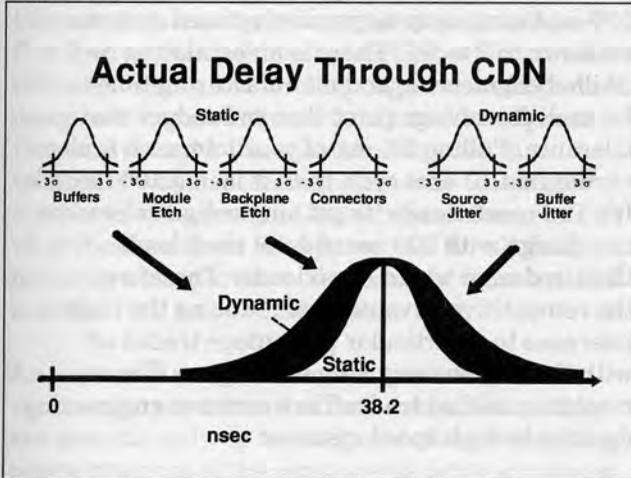
Slide #7



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

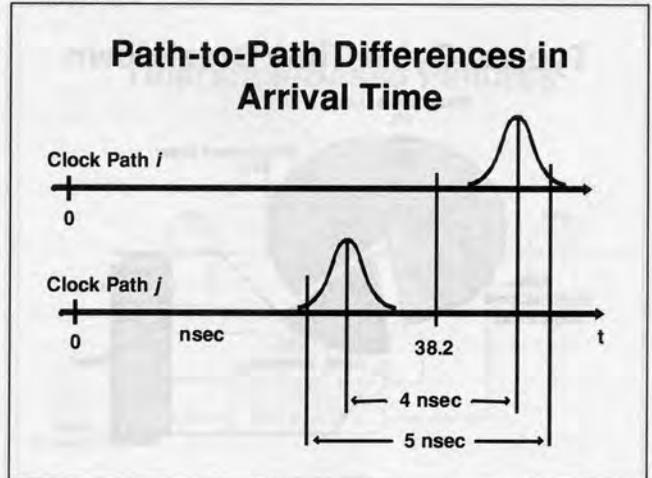
Slide #8



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.

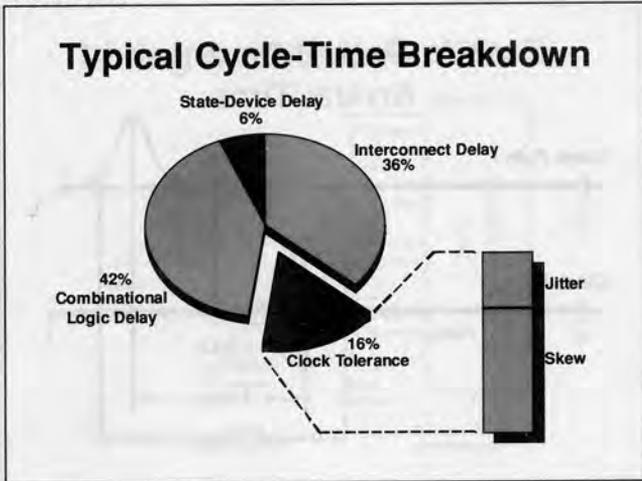
Slide #9



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:

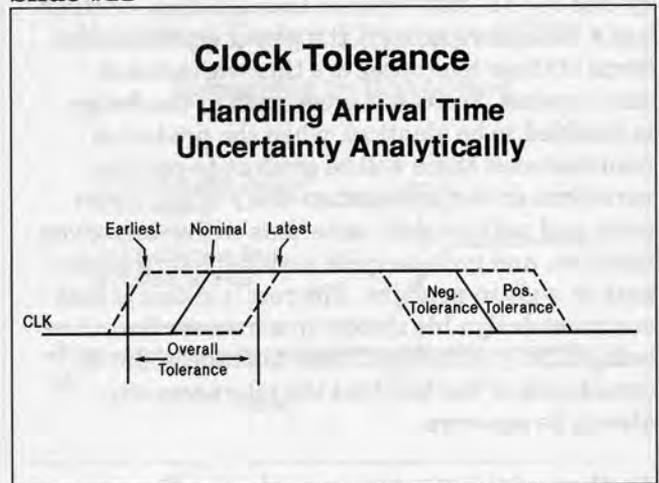
$$T_{cyc} > \text{max segment time} + \text{tolerance} + \text{setup time}$$

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.

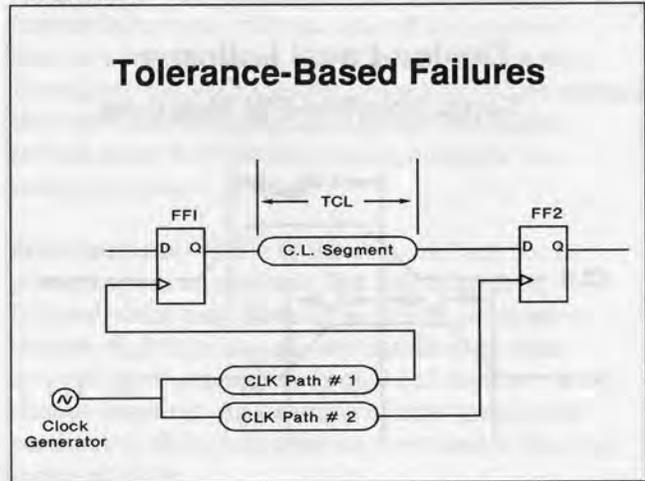
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

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.

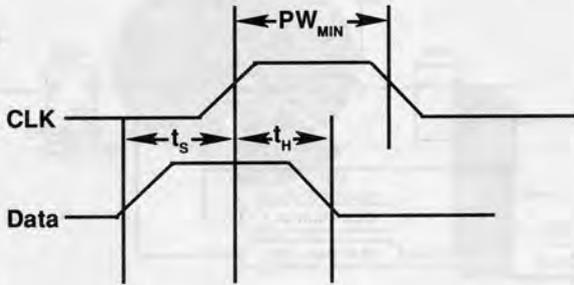
- If clock-path 1 is slower than nominal, or ...
 - If clock-path 2 is faster than nominal ...
- then FF2 will sample its input before the data wave emerges from the combinational logic segment → FAILURE.

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.

Slide #13

Device-Level Failures Setup/Hold/Min-PW Violations



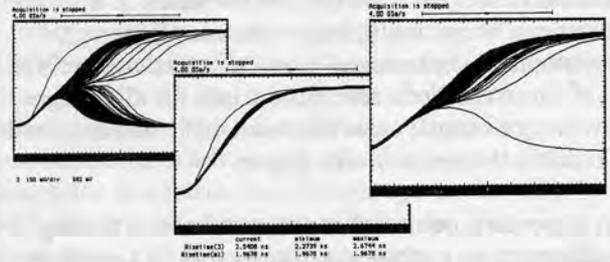
At the device level, a timing-related failure occurs anytime any state-device in a system fails to present, by the end of its propagation interval, a stable (that is, non-transitioning and non-metastable) copy of the data it was supposed to capture. This definition does not require that this erroneous 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

A Variety of Manifestations



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

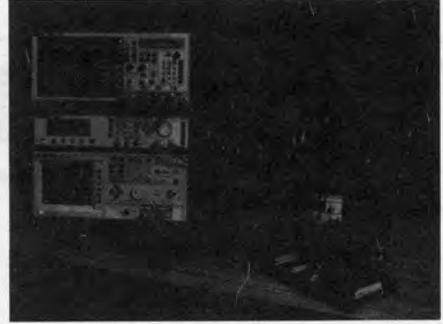
- What is clock tolerance and why care?
- ➔ • Distortion and tolerance mechanisms
 - Signal integrity problems
 - Skew
 - Pulswidth shrinkage and growth (SAG)
 - Jitter
- Strategies

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.

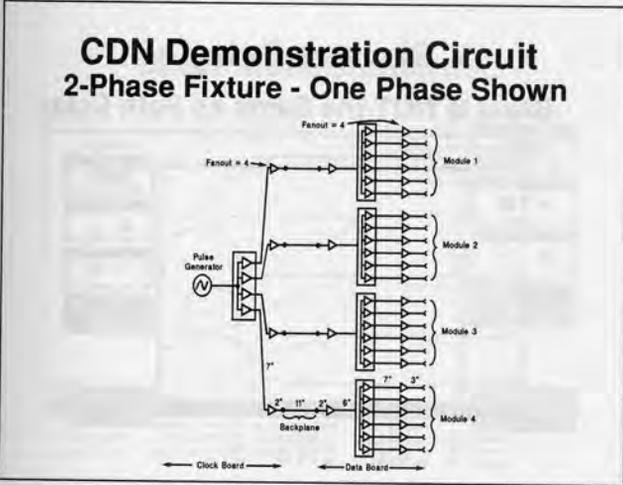
The HP 8110 Pulse Generators were selected to drive the two phases of the CDN. They have excellent edge-placement precision (10 ps), low rms-jitter (10 ps) and adjustable edge rates. Since one of the two phases is used to simulate various types of operational noise behavior, the high-degree of flexibility 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



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.

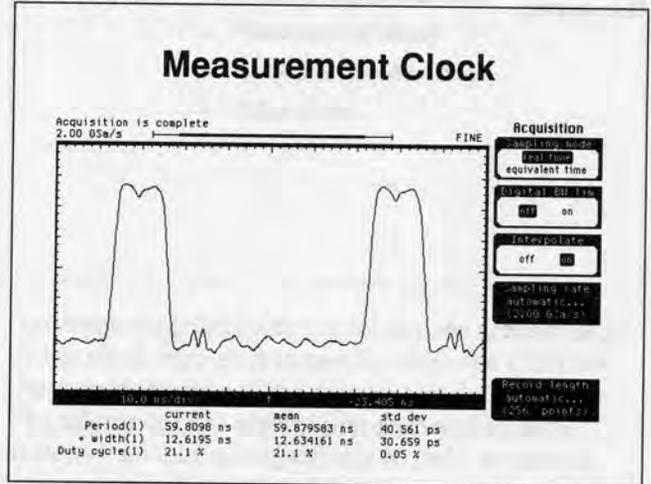
This test fixture for demonstrating CDN effects is an accurate representation of a clock distribution in a backplane-based, fast-CMOS cpu with two clock phases. Module sizes are specified to be identical to VXIbus C-size Eurocards, as is backplane spacing. The CDN has five buffering levels and total 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



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

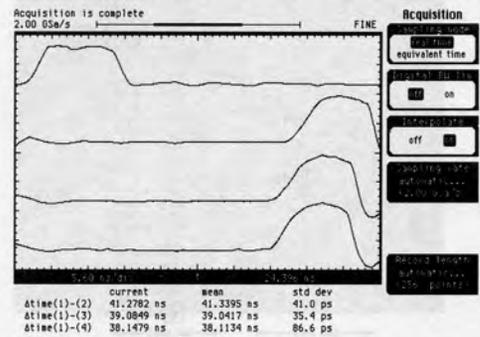
Distortion due to signal integrity problems can be a major source of trouble at elevated operating and edge speeds. The effects of these problems upon timing include multiple or unintended triggering (that is, extra clock edges), destabilization of the data during the setup-hold interval, or delay of 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

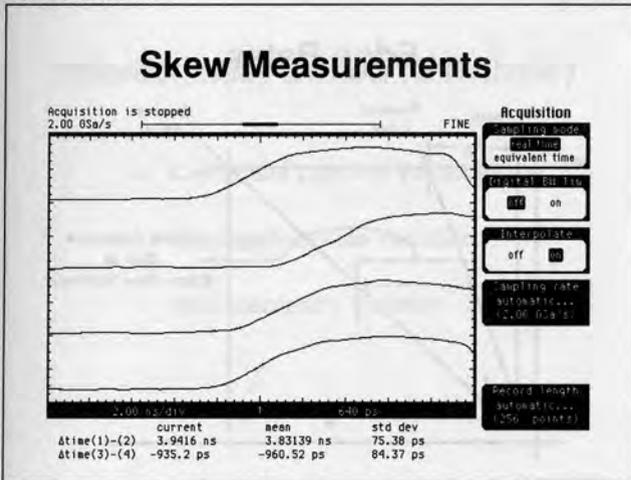
Skew vs Delay

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.

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

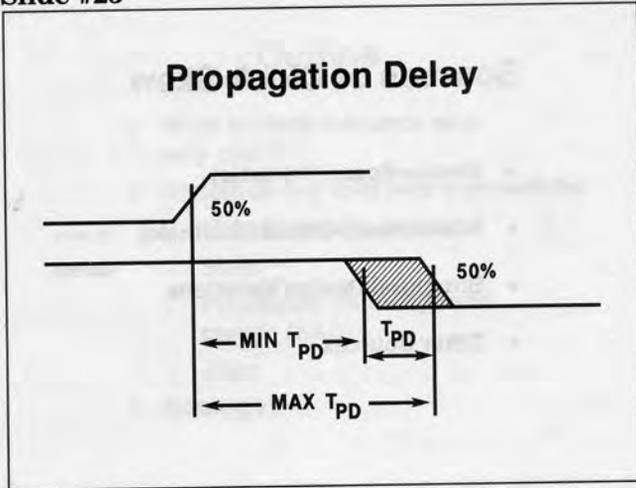
Device-Based (Intrinsic)

- 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.

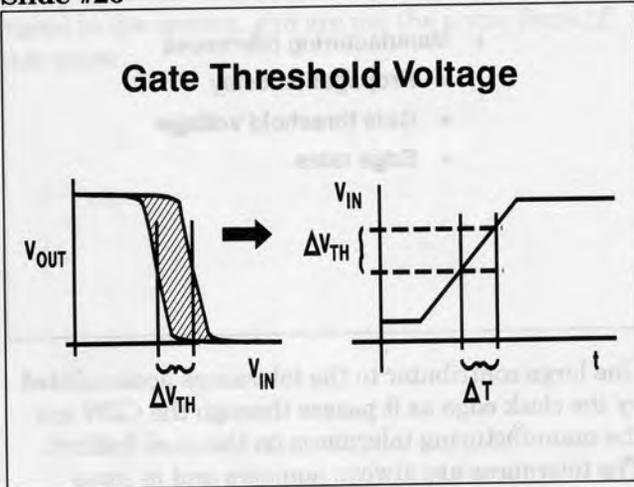
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #25



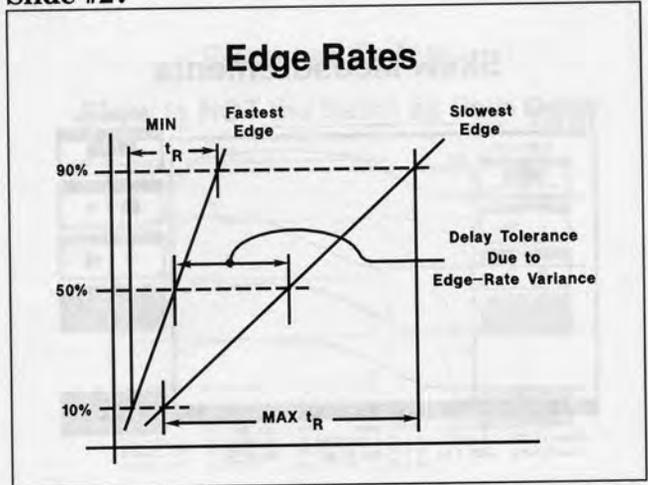
Buffer propagation-delay tolerances are a standard catalog rating for every buffer made. The figure illustrates the trailing-edge tolerance on the propagation delay of an inverting buffer. This tolerance, as well as most of the others in this section, can be specified separately for rising and falling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

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.

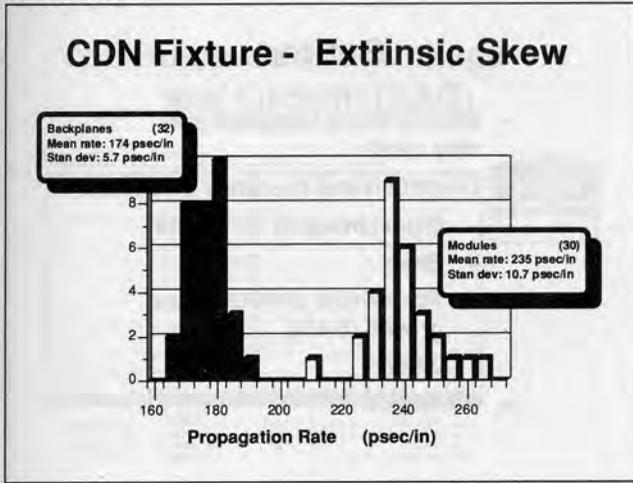
Slide #31

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.

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.

Slide #34

Other Sources of Skew

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

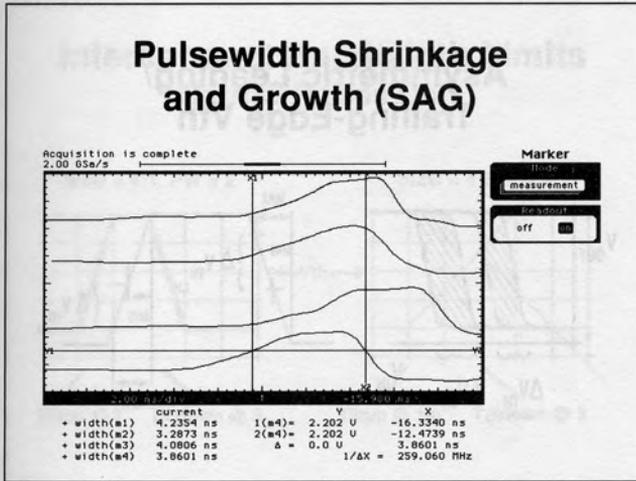
Given that device performance usually varies with temperature, and that thermal management is never perfectly consistent throughout the system, time-performance tolerances will exist even for manually sorted and matched buffers. This effect can be especially pronounced in mixed technology systems, where 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

Slide #36



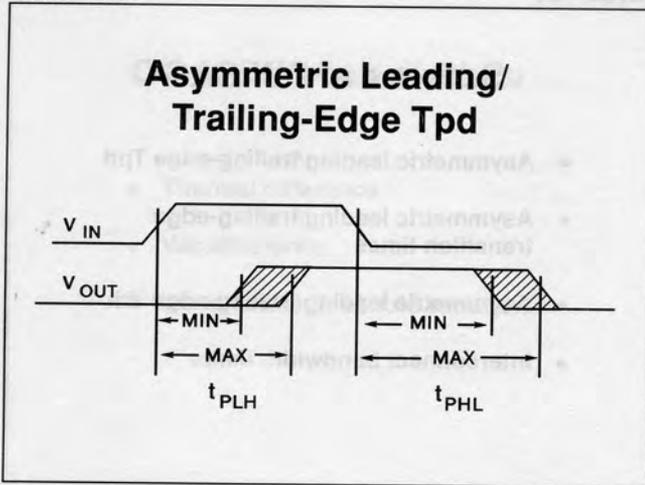
The measurement above shows several clock pulses emerging from the 5-level CDN fixture. The buffers in that system have a typical difference between their leading and trailing edge propagation delays of 500 ps. Note that for the four signals shown, all deviate considerably from the 4.00 ns pulsewidth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

Slide #37

Sources of PW SAG

- Asymmetric leading/trailing-edge Tpd
- Asymmetric leading/trailing-edge transition times
- Asymmetric leading/trailing-edge Vth
- Interconnect bandwidth limits

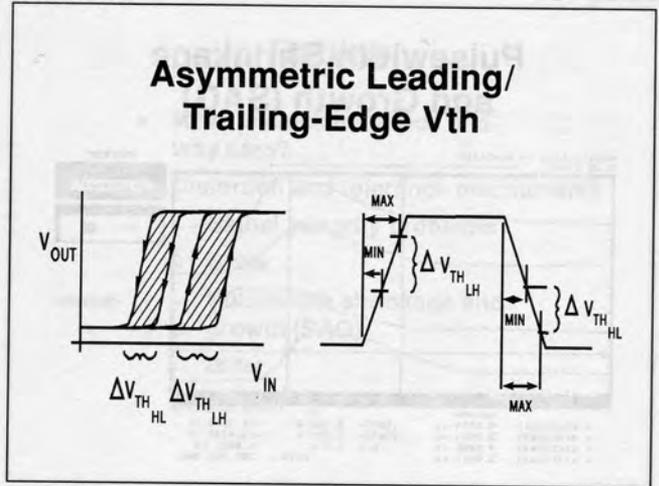
Slide #38



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.

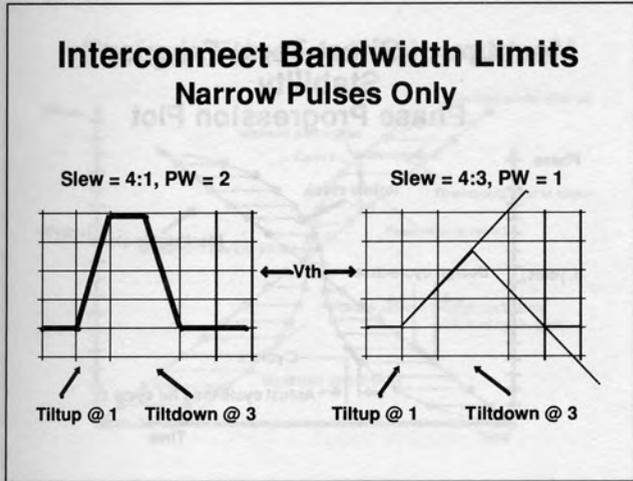
Asymmetric leading/trailing-edge transition times can also change the pulsewidth. The time to transition from a low to a high can be different than the time from low to high. This will cause unequal delays in driving the downstream logic to threshold, resulting in a change in the pulsewidth.

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.

Slide #40



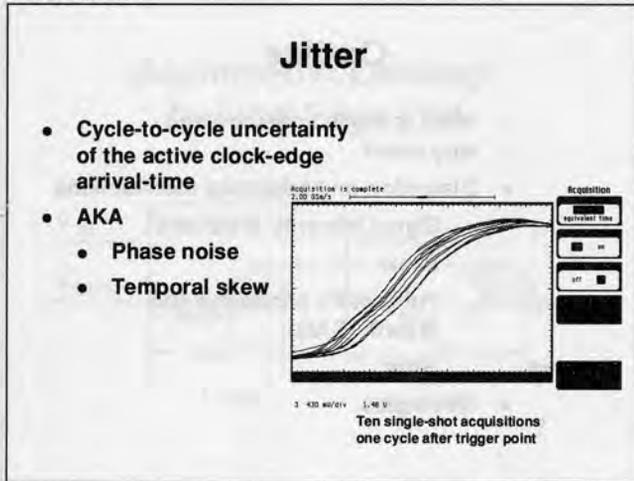
For sufficiently narrow pulses, a bandwidth limit in the transmission path can shrink the pulsewidth further. The bandwidth limit reduces the slew rates. A narrow enough pulse will not achieve full amplitude before it has to switch in the opposite direction. This will in turn reduce the time it spends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

Slide #41

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

Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #42

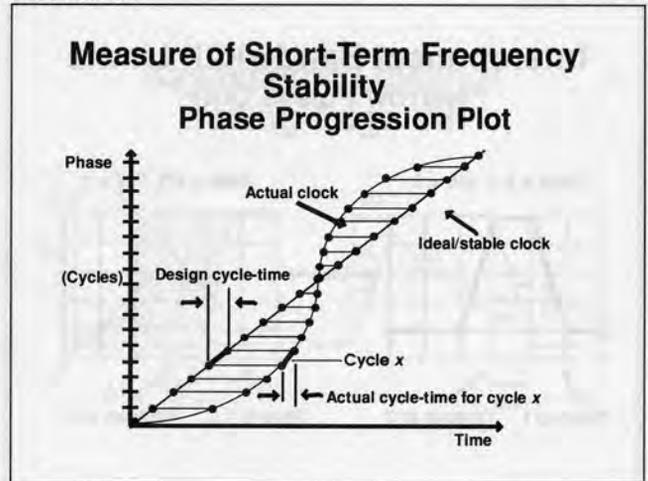


The two forms of distortion discussed thus far have involved time-invariant displacement of the active edge of a particular clock signal from its nominal arrival time. The statistical distributions of these displacements were not a function of time. There is a type of edge-displacement phenomena called 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

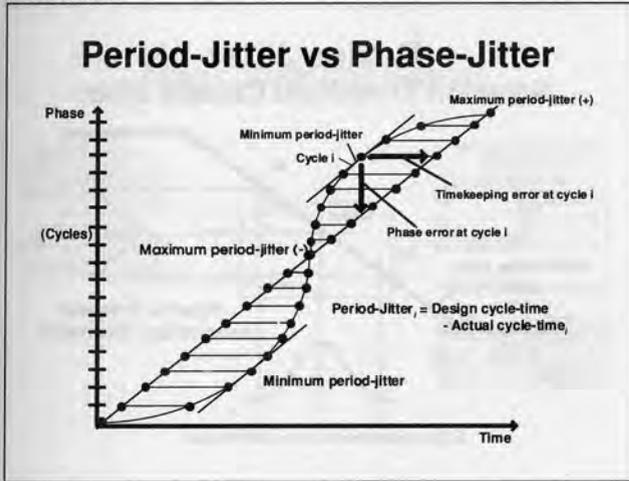


Long before jitter was a concern to digital system designers, it was a common problem for telecom and RF engineers. A number of analytical methods and models for considering the stability of a signal were developed and have some merit to the state-architecture synchronization case. One of these methods 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.

Slide #45

Sources of Jitter Fundamental Cause: Noise

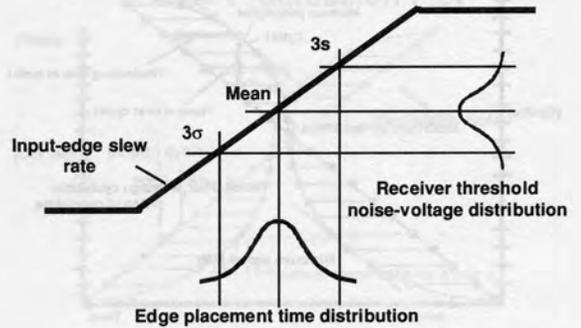
- Clock Source/Phase-Generator
- Buffer
- State-Device

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

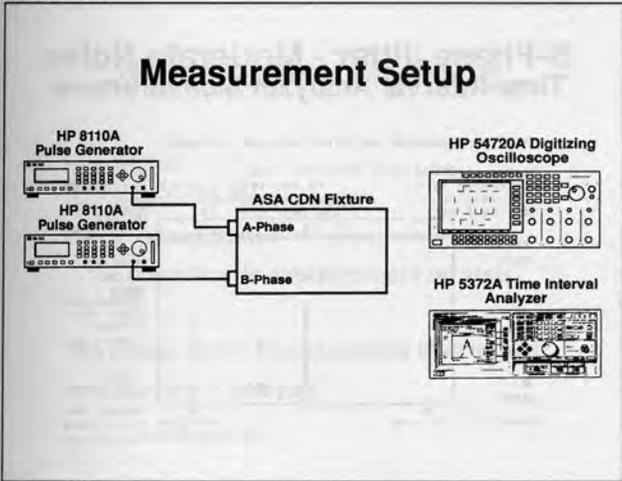
Device Jitter Noise on Threshold Causes Jitter



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



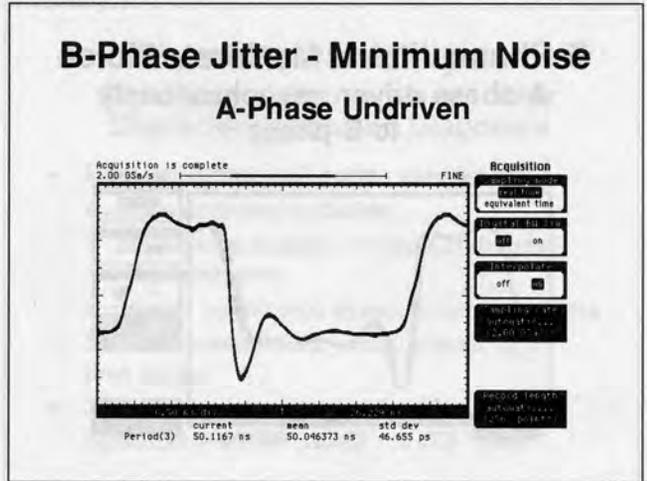
In this measurement, two separate pulse generators are used to drive the CDN fixture. Due to several minor signal-integrity faults in the fixture, when one phase is driven asynchronously with respect to the other at a different frequency, the power environment is corrupted by noise. This will cause the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

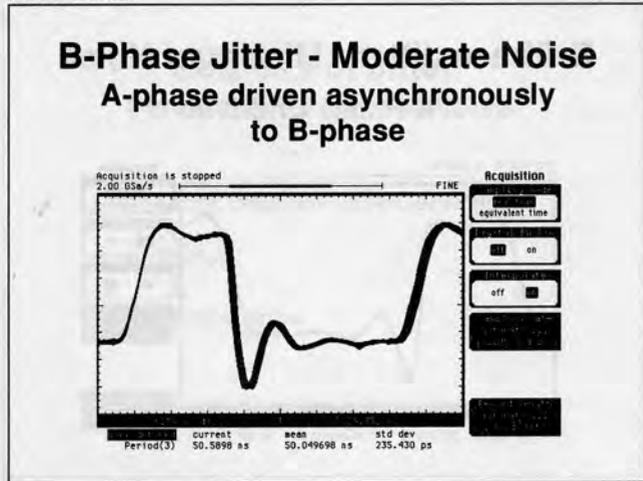


For this measurement of baseline jitter, the noise phase is undriven. The scope plot above shows one cycle of an output of the clock phase in infinite persistence. The trigger point is on the rising edge on the left. Jitter is best measured one cycle away from the trigger point (next rising edge).

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

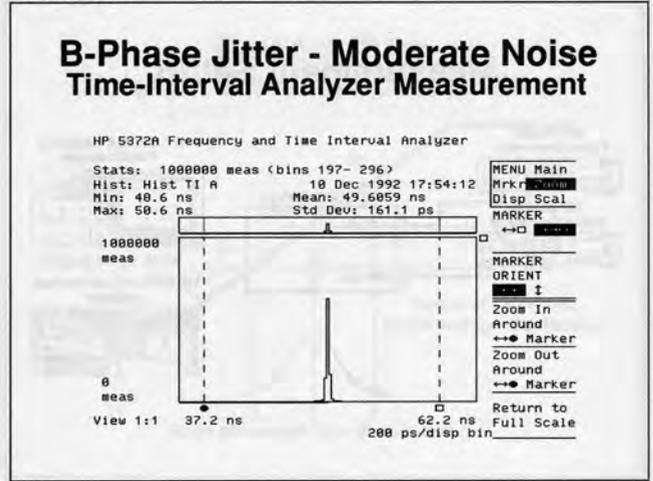
Slide #49



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.

Slide #51

Measurement Methods

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

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

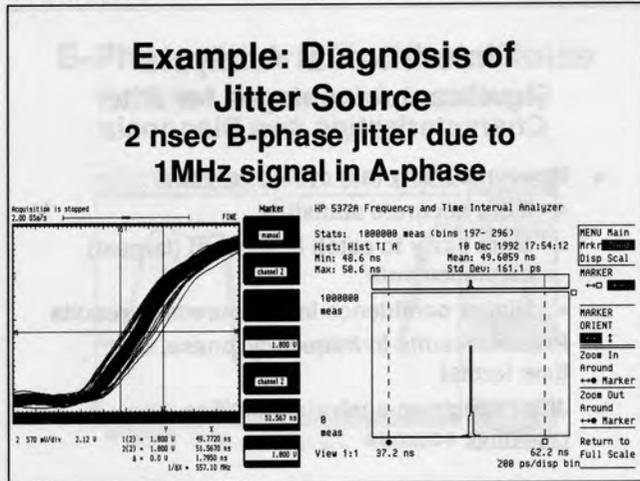
We have just seen a useful demonstration of jitter measurement using time-interval analysis. The fact is that several methods exist, and they all have their place. The author has found jitter and jitter measurement to be a very current topic with high-speed computer designers throughout the industry.

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.

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.

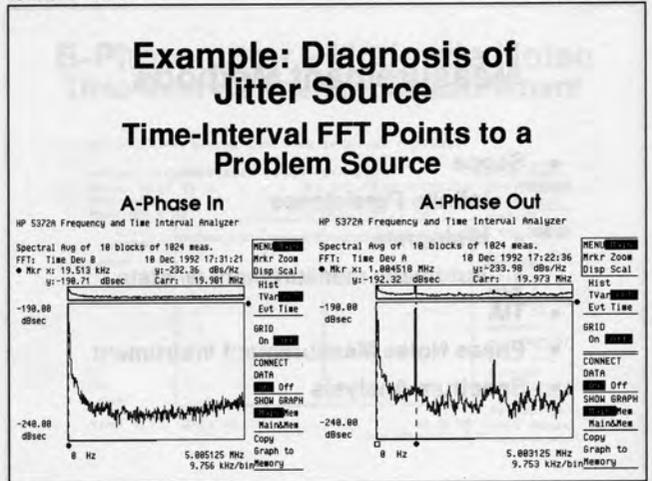
Distortion and Tolerance Mechanisms in High-Speed Clock Delivery

Slide #53



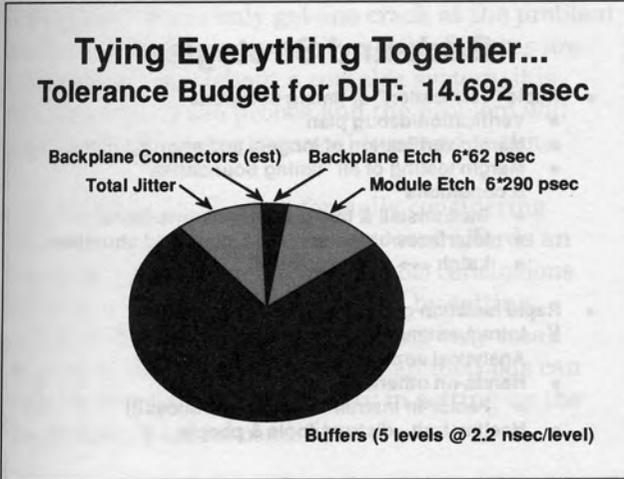
In this example, we're injecting a 1 MHz squarewave into the noise phase. The plots above show scope and TIA representations of the signal. The total jitter is about 1.4 ns. Either view shows us we have jitter, but we don't have any information about the source from these photos.

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.

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

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

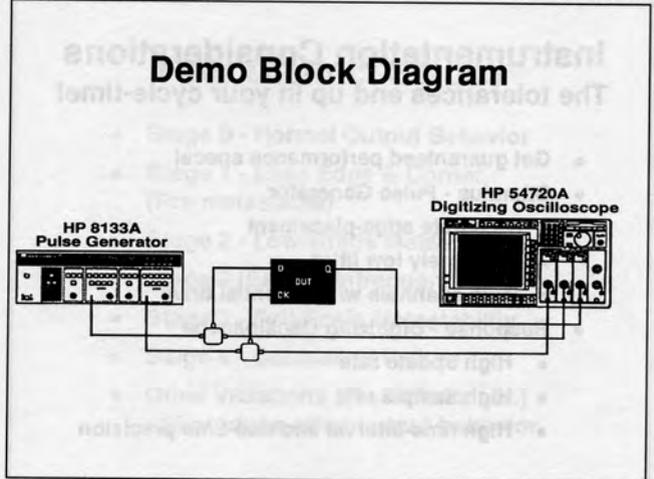
A design option for very aggressively timed systems is to employ self-characterization. It can be employed to factor out any excessive margins in the device manufacturer's guard-bands for rated specs, to determine device parameters that are not rated (for example, jitter susceptibility), or to simply 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.

Slide #64

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

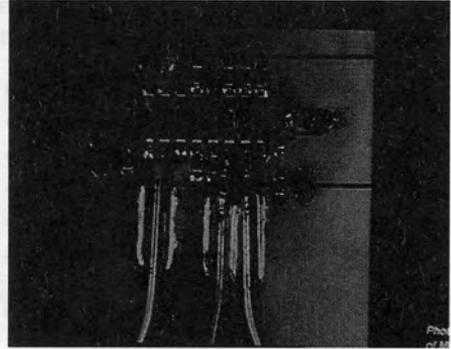
- Get guaranteed performance specs!
- Stimulus - Pulse Generator
 - Accurate edge-placement
 - Extremely low jitter
 - Two channels w/ differential drive
- Response - Digitizing Oscilloscope
 - High update rate
 - High sample rate
 - High time-interval and rise-time precision

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.

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.

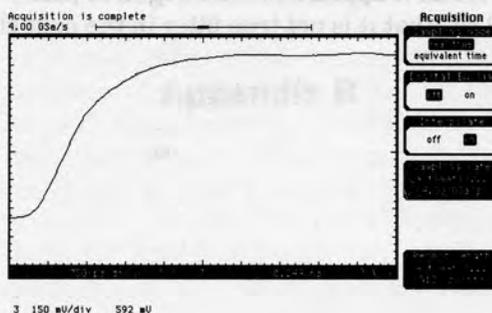
Slide #67

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

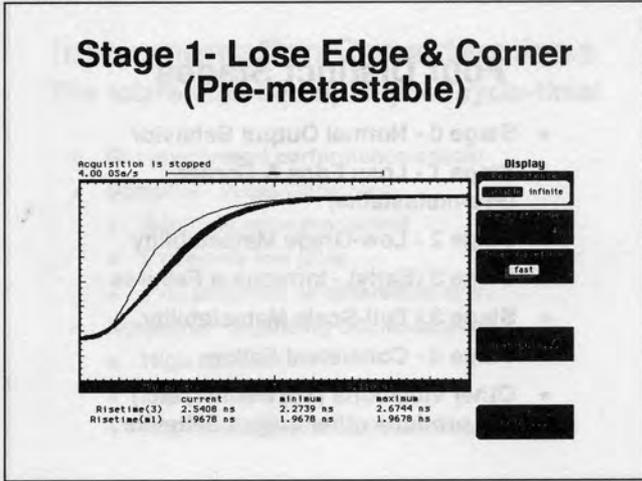
Slide #68

Stage 0 - Normal Output Behavior



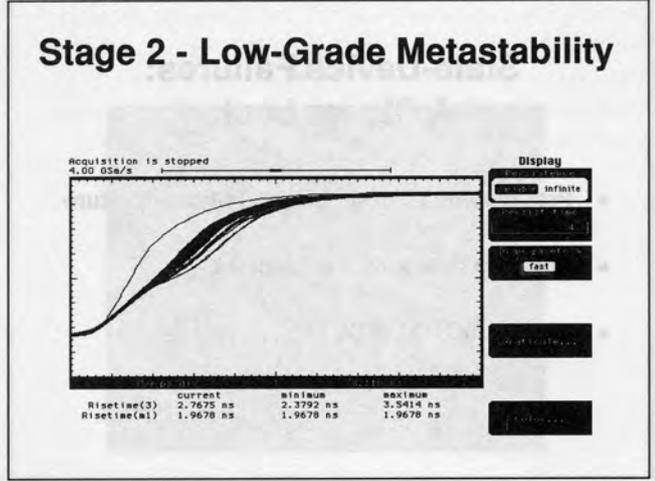
This photo shows the Q-output of the state-device when it is operating normally. It will be placed in the memory of the oscilloscope for use as a reference in subsequent measurements.

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

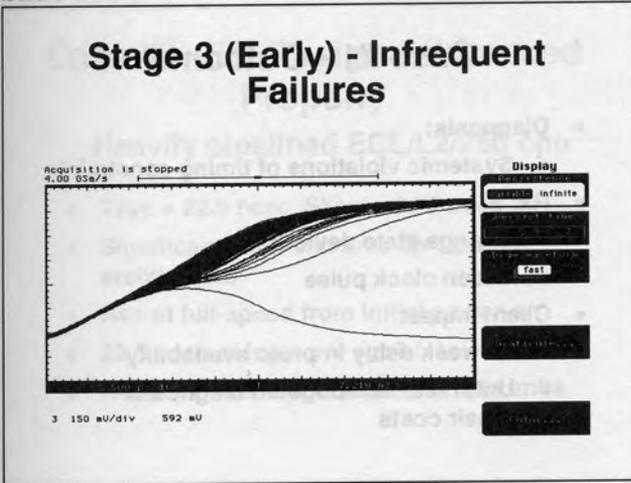


Decreasing the separation between the clock and the data edges a few more picoseconds, we see that the apparent jitter is actually low-grade metastable behavior. Jitter displaces the edge in time uniformly at all voltage levels. In this case, the lower half of the edge is stable. Note that none of the 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.



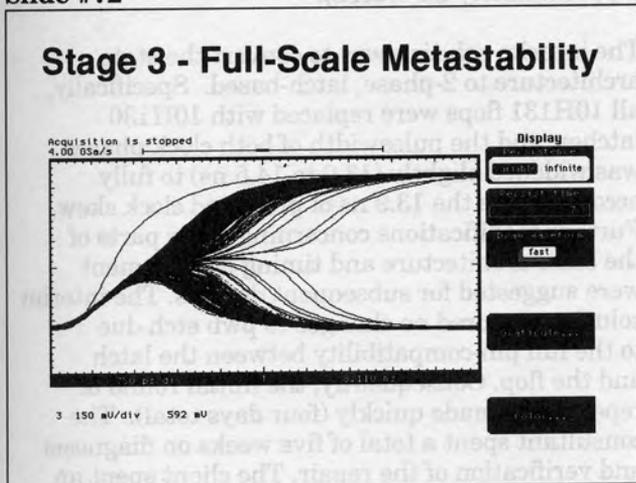
This plot shows the output of the state device when it is operating normally. It will be placed in the memory of the oscilloscope for use as a reference in subsequent measurements.

Slide #71



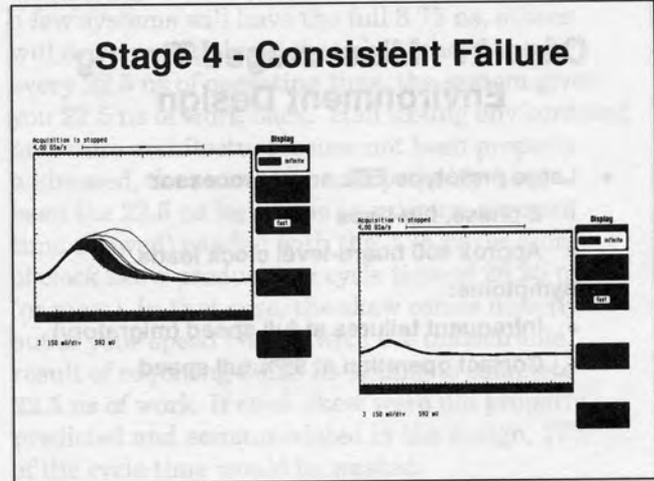
Advancing the edge one more picosecond, we now get a single failure (resolves to incorrect state) during the measurement period. This is one of the types of low-level failures that produces the difficult failure modes discussed earlier.

Slide #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.

Slide #73



Slide #74

Appendix B

olly populated with logic modules. The system state device is the J0H181 flip. All systems exhibited sporadic unrecoverable failures characterized by apparently incorrectly captured data when running at full speed (approximately 40 MHz/50ns). The system appeared to be up to approximately 80% 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 #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_{cy} = 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!

1988
High Speed Digital
Systems Design & Test
Symposium

Printed Circuit Design Techniques for the Control of Electromagnetic Interference

Michael L. Conn

Mikon Consulting
4248 Lake Santa Clara Drive
Santa Clara, CA 95054-1328
Tel: (408) 727-5697
Fax: (408) 727-5697

1993
High Speed Digital
Systems Design & Test
Symposium

Printed Circuit Design Techniques for the Control of Electromagnetic Interference

Abstract

This paper addresses the existence, origin, predictability, and control of generated electromagnetic interference (EMI) in high-speed digital designs. The performance bounds imposed by signal integrity issues are integrated with printed circuit board

(PCB) design techniques to control the radiation and reception of EMI. Detailed guidelines for PCB design are presented. Optional and newly emerging PCB design approaches are discussed with a focus on high frequency (40+MHz clock) digital systems.

Author

Current Activities:

Michael Conn (Mikon Consulting) offers problem diagnosis, mentorship services, and tutorials in analog design, analog-to-digital interfacing, signal integrity engineering, circuit modeling and analysis, all facets of electromagnetic compatibility engineering, and electronic systems engineering. Research, development, and implementation of challenging (and proprietary) concepts through the pre-production phase is a specialty.

Author Background:

Mike specializes in advanced power electronics designs for precision control systems, ordnance initiation, and dc power control and distribution. His experience with sub-microsecond switching of high power systems has necessarily led to the in-depth study of electromagnetic interference effects and their control through proper circuit and packaging design.

Mike has 34 years experience in research, development, test, and evaluation in aerospace, military, and high-technology commercial electronic systems. He has a BSEE and MSEE from Stanford University, plus 42 post-MS Honors Units from Stanford and the University of Santa Clara.

Slide #1

Printed Circuit Design Techniques for the Control of Electromagnetic Interference



Mikon Consulting



Slide #2

Presentation Overview

Melding EMC with Signal Integrity

- Inadvertent antenna creation
- Differential- and common-mode radiation
- Printed circuit board characteristics
- Field confinement and interception
- Current limiting and confinement
- PCB design guidelines
- Bench-top evaluation

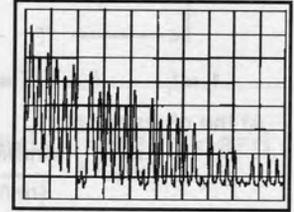
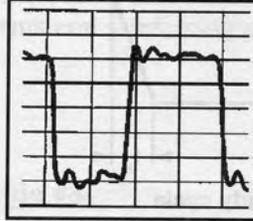
Allowable time constraints limit the scope of this paper to a focus on differential-mode and common-mode radiation as primary EMI design drivers for printed circuit boards.

Real-world design techniques will be correlated to their theoretical basis for preferred designs.



Slide #3

Time versus Frequency Domain

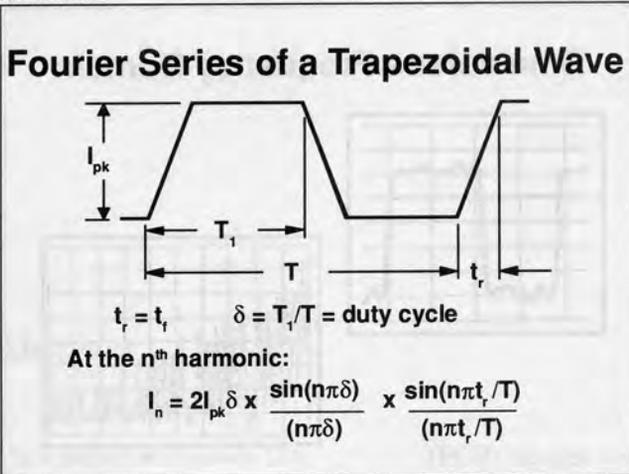


The commonly observed voltage traces in a digital circuit are related in a complex way with the radiated fields from those traces. The impedance levels, propagation times, trace and structure resonances, and coupled circuits all affect the magnitude and constituent frequencies of the radiated fields.

Time-domain reflectometry evaluation is normally done for signal integrity and noise margin evaluation, but spectral analysis of voltage and current waveforms is required for assessment of electromagnetic compatibility.



Slide #4



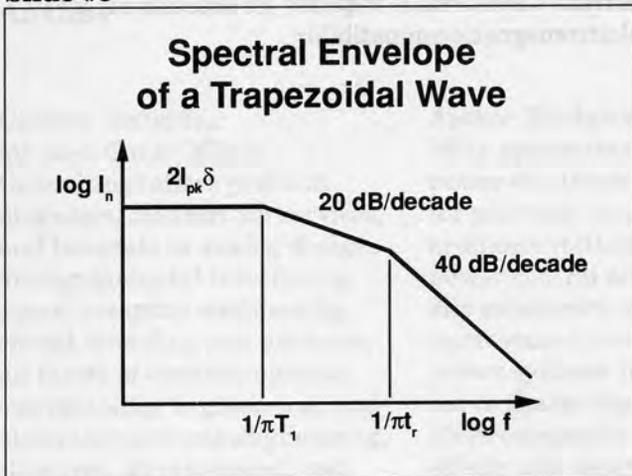
The prediction of radiation from a circuit requires identification of the frequency content and absolute magnitudes of coefficients through a Fourier transform of the time-domain waveform. This sample trapezoidal waveform assumes equal rise and fall times.

The curve rises as frequency decreases until frequency $1/\pi T_1$ (the pulse width) is reached, and the curve plateaus at the value $2I_{pk}\delta$ for all lower frequencies. The details on derivation of this spectral envelope can be found in the Fourth National IRE Symposium on RFI, June 1962.

Note that the edge rate (or rise time, t_r) is a key driver (or limiter) for the frequencies of concern. This fact alerts designers to use the slowest edge rate devices that are suitable for their application in order to minimize their potential EMI problems. As a point of reference, a 1 ns rise time yields a break frequency of 318 MHz.

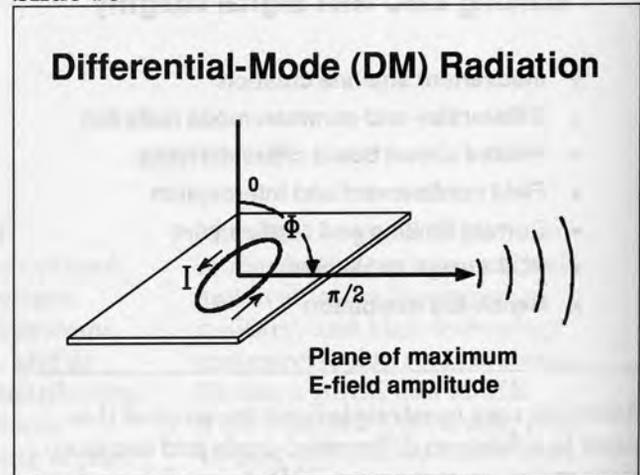
CAUTION: The rolloff of the spectral content illustrated in the slide can be offset by the efficiency of radiation exhibited by the radiating "antenna"—to be covered in later slides.

Slide #5



The magnitude of the Fourier coefficients are bounded by the curve shown. The $\sin X$ over X terms cause periodic attenuation/ripple in the actual magnitudes that you would measure. No energy is contained at frequencies less than the fundamental, and 99% of the energy is contained below $f = 1/\pi t_r$.

Slide #6



Take a simple loop and pass an alternating current through the loop. The current may consist of the harmonic frequency components of a digital signal on a printed circuit board (PCB).

The magnitude of radiation from the loop will vary in proportion to the current. The radiated electric field intensity at a distance d from the loop will be maximum when measured in the plane of the loop.



Slide #7

DM Radiation Prediction Far-Field Characterization

$$E = 1.316 \times 10^{-14} \times (A I_f^2 / d) \sin \Phi \text{ Volts/Meter}$$

- A = Loop area, square meters
- d = Distance from loop center, meters
- I_f = Current at frequency f, amperes
- f = Frequency (of harmonic), Hertz
- Φ = Angle from loop axis

There are three ill-defined distances referred to when studying EMI. They are near-field, intermediate-field, and far-field. In the near-field, some radiation terms roll off inversely with d cubed; the intermediate terms roll off inversely with d squared; and the far-field terms roll off inversely with d. Most certification tests at frequencies of concern are made in or near the far-field region. The boundary for the far-field region depends on the frequency and commences at $\lambda/2\pi$. At this distance, the radiated field approximates a plane (TEM) wave. For example, at a measuring distance of three meters (standard FCC test distance), the far-field definition applies to all frequencies above 15.92 MHz. At 10 meters, frequencies above 4.775 MHz qualify.

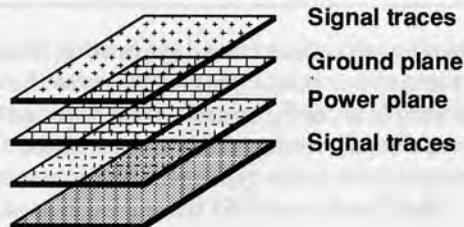
Note that the I_f in the equation is a particular current at a particular frequency (harmonic). Also, the magnetic field magnitude in the far-field region is simply $H = E/120\pi$, (the E-field divided by the radiation resistance of free space).

A sample calculation of the radiated field from the third harmonic (100 MHz) of a 254 mm (10 inches) long by 0.635 mm (0.025 inches) wide loop carrying a 33.3 MHz, 4 V_{p-p} signal into a 50 ohm load yields 41.6 dB μ V at the 3 meter test distance used for FCC Class B certification tests. Only 43.5 dB μ V is allowed at 100 MHz. The radiation will increase approximately in proportion to the square-root of the number of traces.

The equation given assumes the far-field. Therefore, assuming a worst-case reflection off the ground plane that *doubles* the signal is a recommended, conservative design approach that allows for a contribution of the intermediate (1/d²) radiation terms and reinforcing ground reflections.

Slide #8

Multilayer Printed Circuit Board I



Multilayer circuit boards radically change the character of circuit loops. The ground and power distribution conductors are typically embedded as planes. The return currents for signal traces now flow through a ground plane that is in close proximity to the trace itself. The smaller area of the current loop substantially reduces the magnitude of radiation from the loop. The plane of the current loop is now normal to the PCB; therefore, DM-radiation will be emitted directly off the face of the board instead of being emitted in the plane of the board.

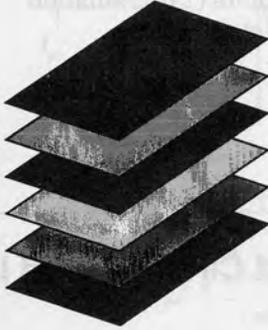
The use of ground and power planes provides the low-impedance power distribution necessary for good power supply decoupling.

The use of outer surface traces is commonly observed; however, placing ground and power planes on the outer surfaces affords superior EMI performance.



Slide #9

Multilayer Printed Circuit Board II



Enclosing signal traces between power and ground planes achieves a locally shielded enclosure that reduces radiation, radiated susceptibility, and ESD susceptibility.

The inverting of surface trace layers with their adjacent ground or power planes converts the traces to stripline configuration. The detailed configurations, attributes, and compromises between these two trace types will be discussed shortly. Note, however, that the signal traces are now *buried under a shield*, relative to the outside world. What better way to achieve 30 to 45 dB of attenuation and susceptibility isolation, and ESD protection?

The artwork for the individual layers of normal multilayer boards does not change if the design uses through-hole components, as opposed to surface-mounted parts. The latter require additional vias to be installed. This may seem an extreme penalty, but the preference for burying traces will become more obvious later.

Slide #10

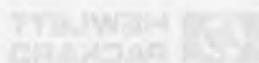
PCB Trace Configurations

Traces are Transmission Lines
($\lambda/4 < \text{trace length}$)

- Transmission Line Configurations
 - Microstrip
 - Stripline
- Trace Characteristics
 - Fields and impedances
 - Signal propagations
- What's Best and Why

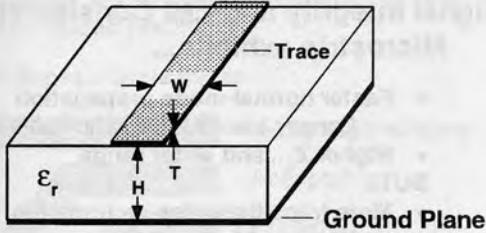
When designing high-speed circuits, the signal interconnections must properly be treated as transmission lines.

High-speed interaction effects become a signal integrity concern when the signal rise or fall time become less than about *twice* the propagation time on a given trace. At this point, the trace length approaches 1/4 wavelength for the highest frequencies contained in the signal. For the astute digital design engineer, this fact provides a guideline for partitioning of functional subcircuits to minimize signal integrity concerns.



Slide #11

Microstrip Structure



$$Z_0 = 87(\epsilon_r + 1.41)^{-0.5} \times \text{Ln}[5.98H/(0.8W + T)] \Omega$$

$$T_{PD} = 1.017(0.475\epsilon_r + 0.67)^{0.5} \text{ ns/foot}$$

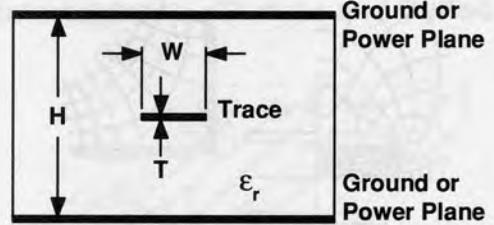
Approximations to the complex microstrip characteristic equations are acceptable for most applications. The relations presented in this slide are accurate within 3% for a W:H ratio of two or less. As an example, the propagation delay for FR-4 material is 1.73 ns/foot (or 144 ps/inch = 56.7 ps/cm) for a propagation velocity of approximately 7 inches/ns (or 18 cm/ns).

CAUTION: The relative dielectric constant of the chosen materials for the PCB should be controlled in critical applications as 5% tolerances are common and 20% tolerances are possible.

Signals with 1 ns rise time should be limited to 9 cm (3.5 inch) traces for microstrip constructed using FR-4 ($\epsilon_r = 4.65$ nominal) material, unless signal integrity techniques are carefully applied.

Slide #12

Stripline Structure



$$Z_0 = 60(\epsilon_r)^{-0.5} \times \text{Ln}\{4H/[0.67\pi(0.8W + T)]\} \Omega$$

$$T_{PD} = 1.017(\epsilon_r)^{0.5} \text{ ns/foot}$$

The added ground plane (or power plane) relative to the microstrip construction naturally adds capacitance between the signal trace and ground, resulting in a lower characteristic impedance. This same capacitive loading also slows the propagation down the trace (about 27% longer for FR-4).

The velocity is about 183 ps/inch for FR-4 material with $\epsilon_r = 4.65$.

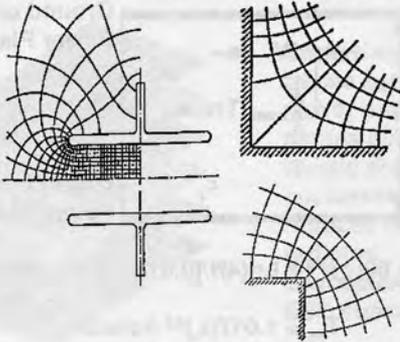
Be aware that multiple (gate) loads on a given trace effectively add a distributed capacitance load along the trace that will further increase the propagation delay by the factor $(1 + C_D/C_0)^{0.5}$, where C_D is the distributed capacitance loading per unit length and C_0 is the normal distributed capacitance per unit length of the trace. This same distributed load capacitance slightly lowers the line characteristic impedance.

Signals with 1 ns rise time should be limited to 7 cm (2.75 inch) traces for stripline constructed using FR-4 material, unless signal integrity techniques are carefully applied.



Slide #13

Visualizing Field Shapes



Orthogonal plotting of the EM field patterns around geometric shapes will give you insight into potential coupling problems before they arise. The symmetrical nature of the patterns simplifies the learning process.

Sketching of field patterns can lead you to solutions for difficult shielding, isolation, and impedance matching problems. Consider the differences in performance characteristics between multiple microstrips and multiple striplines as a practical example.

As a practical matter, the use of field solver software in the prediction of transmission line characteristic impedances, propagation velocities, and cross-coupling effects is desirable (and perhaps mandatory) in finalizing a sophisticated design. However, the creative design work that *leads* to the construction of probable design layouts will be accomplished faster with an understanding of field patterns.

Slide #14

Microstrip versus Stripline

Signal Integrity and EMI Considerations Microstrip exhibits...

- Faster normal-mode propagation
 - Longer traces for an allowable delay
 - Higher Z_0 ...and wider range
- BUT...**
- More transition edge degradation
 - Faster odd-mode propagation
 - More crosstalk
 - More radiation
 - Lower trace density

Visualization of the fields generated by the microstrip and stripline configurations makes these comparisons seem relatively intuitive. For example, the fields emanating from the microstrip surface are not guided to a controlled return conductor, but rather tend to terminate on adjacent traces (more crosstalk). Some of these fields escape the surface of the PCB totally and propagate or radiate outward.

To reduce the crosstalk associated with microstrip, the spacing between adjacent traces is necessarily widened, resulting in lowered interconnection density.

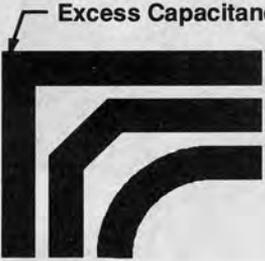


Slide #15

PCB Trace Corners

Consider signal integrity and thermal and flex-induced problems

Excess Capacitance



Avoid Right Angles
Acceptable
Preferred

The use of 90-degree corners causes excess capacitance to be introduced to the trace and represents a small, but unnecessary, impedance change in the characteristic impedance of the transmission line. The use of 45-degree turns with a minimum segment length of twice the trace width is better (and is offered by most auto-routing CAD programs). Continuously curved traces with an inside radius of at least the trace width is the best approach.

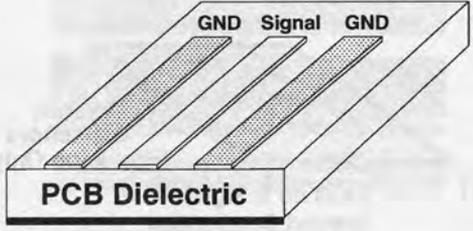
Spacing between adjacent active traces should not be less than the trace width to minimize crosstalk, but little additional benefit is gained for spacing of more than three times the trace width.

Although the electromagnetic effects of the 90-degree corners are secondary to most signal integrity effects, the use of such corners in high temperature or flex-circuit applications lead to degradation and reliability problems with cracks caused by stresses.

Slide #16

Co-Planar Ground Traces

Interceptors — Surface and Buried



GND Signal GND
PCB Dielectric

If a trace conducting high frequency currents is to be routed on the surface of a printed circuit board (PCB), Mikon recommends grounded traces be routed parallel to it to reduce both radiation and crosstalk. The ground traces should be connected to ground fill areas or ground planes at *varied* intervals not to exceed $\lambda/4$ at the highest frequency or harmonic expected. This recommendation applies to single- and double-sided PCBs, as well as microstrip lines. For example, use of vias-to-ground at spacings of five centimeters (two inches) or less with FR-4 ($\epsilon_r = 4.65$) would be satisfactory for harmonic frequencies approaching 685 MHz. However, waveform rise times should be greater than 1.5 ns to limit the magnitudes of such harmonics.

The adjacent grounded traces can be placed closer to the signal trace than other signal traces for more effective interception of emerging fields. The slight lowering of the line characteristic impedance can be compensated by thinning down the signal trace. The thinner, closer-spaced traces minimize the surface area required for this configuration while simultaneously adding more loss or damping to the signal line.

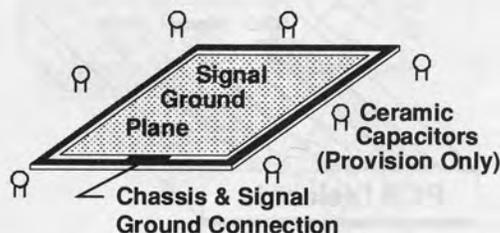
Asymmetrical striplines (buried microstrip) will also benefit from reduced crosstalk using this technique.



Slide #17

Co-Planar Chassis Ground Rings

Ring on each layer — vias < 5 cm apart

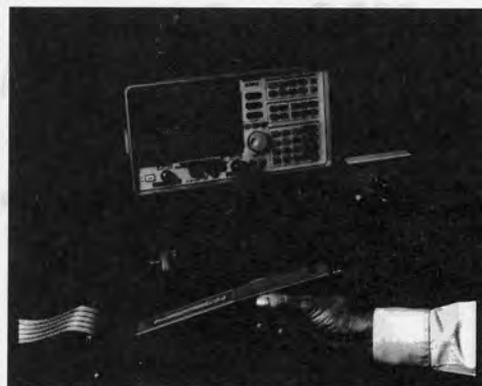


For field interception, Mikon recommends chassis ground rings, preferably wider than 2.54 mm (0.1 inch), be placed on the periphery of *each layer* of the circuit board and interconnected with vias at *varied* spacings up to 5 cm (2 inches) maximum.

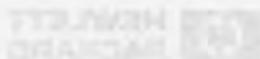
This construction presents a formidable shield (or field interceptor) to prevent radiation (or susceptibility) at the circuit boundaries. Experience (and recent publications of studies) has demonstrated up to 20 dB greater emissions from edge-located traces relative to traces well within the PCB borders. The chassis ground rings also act as a preferred path (interceptor) for electrostatic discharge, yielding a more robust design.

Since unwanted resonances may be created with this construction, *provisions* for "detuning" with ceramic rf capacitors or damping with low-value resistors should be made between this shield ring and the normal circuit signal ground plane. Capacitor values could range from 50 pF to 1000 pF depending on the frequencies of concern. Suitable damping resistor values range from 10 to 50 ohms, depending on the apparent RF impedance of the transmission line formed by the chassis ground rings and the signal ground plane.

Slide #18



Here is a typical test setup of a spectrum analyzer system. The HP 84100B system illustrated includes two close-field magnetic probes, a pre-amplifier, a spectrum analyzer, and an EMC "personality" card. The card customizes the general purpose spectrum analyzer to EMC measurements by including FCC, VDE, and CISPR test limits, and calibration compensation data for various antennas and probes. The appropriate test distances and receiver bandwidths required by the applicable test are also preprogrammed.



Slide #19

Radiation Test Results I



Log frequency (10 MHz to 1.5 GHz)

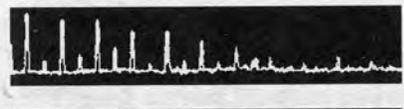
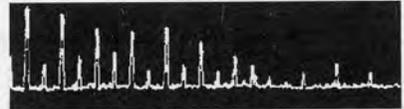
A PCB was designed with FR-4 materials to compare crosstalk, radiation, relative trace densities, and overall signal integrity of microstrip, guarded (coplanar ground traces) microstrip, stripline, and guarded stripline. After a successful verification of trace impedances by TDR evaluation, the respective radiation characteristics of the four transmission line implementations were tested with a benchtop spectrum analyzer system with the lines unterminated, except for surface pads with approximately 3 pF of capacitance.

The graphic above shows typical near-field test results for the standard microstrip, the guarded microstrip, and the standard stripline, all implemented with a duplicate, parallel trace to assess the effects of inter-trace coupling on 50 mil centers. A controlled rise and fall time generator was used to produce 1 ns rise and fall times with a 50% duty cycle operating at 66 MHz. The guarded stripline case is not shown as the performance, as determined separately by detailed TDR tests, was substantially better than the basic stripline configuration.

With minor deviations, the guarded microstrip improved (lowered) the radiation level relative to the standard microstrip by 6 to 8 dB to frequencies beyond 1 GHz. The stripline, as anticipated, made the radiation virtually undetectable. The noise floor was approximately 14 dB μ A/meter, and the measuring height was fixed (by a writing tablet) at 0.25 inch for all tests.

Slide #20

Radiation Test Results II



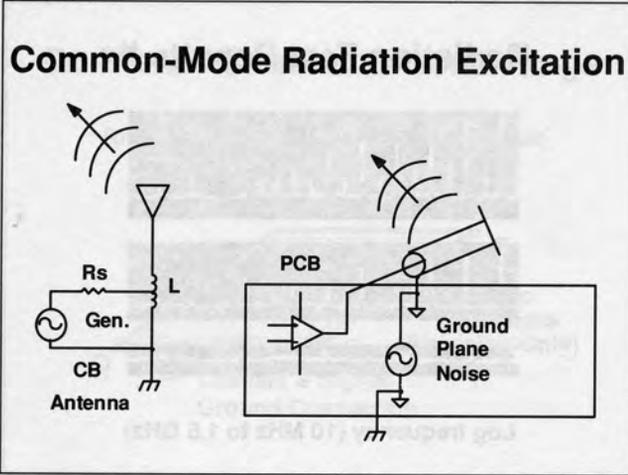
Log frequency (10 MHz to 1.5 GHz)

This second graphic compares the same transmission lines as in the previous slide, but with terminations added (47 ohm surface mount resistors). There are some minor differences in the amplitude distribution versus frequency, but the overall energy in the radiation fields detected were basically unchanged.

Note that a local ambient broadcast signal (near the center of the bottom trace in the previous slide) was not present. They had apparently reduced their signal strength for broadcast during the evening hours.



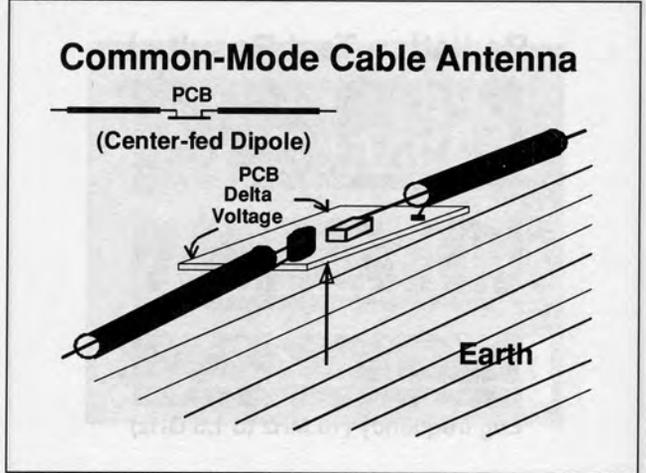
Slide #21



A portable citizen's-band (CB) transmitter typically uses an inductively loaded whip (or monopole) antenna that is considerably shorter than one-quarter wavelength. The inductor is tapped near the grounded, low-impedance end of the inductor for proper impedance matching to the generator source impedance.

Excitation of the ground plane on a circuit board can excite rf energy on cable leads and shields that then behave as monopole or dipole antennas.

Slide #22



This slide illustrates how ground plane noise can excite a pseudo-dipole antenna configuration. The magnitude of this phenomenon will depend on the containment and suppression techniques (if any) employed in the final design.

Note that the magnitude of the noise generated is limited to allowable design noise margins for proper logic functionality, but the ultimate success or failure in the radiated EMI tests is *not* directly correlateable to the noise margins.



Slide #23

Common-Mode (CM) Radiation Monopole Antenna Radiation

$$E = 4\pi \times 10^{-7}(fLI_c/d)\cos\Phi \text{ volts/meter}$$

f = Frequency, Hertz

L = Cable length, meters

d = Distance from cable, meters

I_c = CM current in cable at frequency f , amperes

Φ = Angle from normal to cable

Note the proportionality to frequency and the larger coefficient (by 100 million to one!) relative to the equation for DM radiation, which is repeated below for convenience.

$$E = 1.316 \times 10^{-14} \times (AI_c^2/d)\sin\Phi \text{ volts/meter}$$

for DM radiation

As an illustration, assuming the angle for maximum field strength (zero degrees), a 1 meter cable carrying only 10 μ A of common-mode current at 100 MHz will yield a field strength of 52.44 dB μ V/meter at the three-meter distance required for FCC testing. The FCC Class B limit at 100 MHz is 43.5 dB μ V/meter. For comparison, recall that a 33.3 MHz clock signal on a 50 ohm line that formed a 10-inch-by-0.025-inch loop would yield a DM signal of 41.6 dB μ V/meter at 100 MHz at 3 meters. The 100 MHz (third harmonic) DM current is calculated to be 16.98 mA, or approximately 1700 times higher than the 10 μ A CM current assumed.

Slide #24

Impulse Excitation of CM Radiation

- Displacement currents, power
 - Switching regulator/power supply
 - Non-synchronous, full-bridge motor driver
- Shoot-through current, power
 - Power MOSFET half-bridge switching
- Shoot-through current, medium power
 - Totem-pole MOSFET gate driver
- Synchronized gate transitions

Transient currents from any source that generates high di/dt conditions will generate ground and trace "bounce" voltages. The high di/dt generates a broad range of high-frequency currents that excite structures and cables to radiate in the common-mode manner.

For example, an on-board switching regulator producing a 500 mA dc output current at 5 Vdc can (depending on its mode of regulation) have a peak current when switching of 2 A or more. Switching the 2 A from one device into another (for example, a flywheeling diode) in 10 ns represents 4 A in 10 ns. The switching frequency may be anywhere from 20 kHz to 1 MHz. The obvious ground disturbance caused by this current redirection is further aggravated by the displacement current generated by the dV/dt of the switching device used to switch the primary inductance. The displacement current can create common-mode currents through a multitude of paths, depending on the insulators and shields used in the packaging of the devices. The design and termination of transformer shields, and the heatsinking of power transistors, have become specialty tasks with the proliferation of switching supplies and regulators.



Slide #25

Methods to Control CM Radiation Control Accomplished via Tradeoffs

- Cable regulations/specifications
- Low-Z ground and power planes
- Use of differential/balanced circuits
- Physical circuit current confinement
- Use of ferrites

If possible, control or specify the lengths, construction, and impedances of cables.

Use at least three skin-depths of copper at the frequencies of concern on the outer ground/power layers for maximum shielding and minimum CM impedance. Mikon recommends two-ounce copper (0.0726 mm or 2.86 mils thick) for most applications. Two-ounce copper provides three skin-depths down to 7.5 MHz, one-ounce copper down to 29.75 MHz, and half-ounce copper only down to 119 MHz.

Differential circuits are *not* always easily achieved, but they are designed to balance the flow of ground currents so as to cancel out, resulting in low CM voltage excitation.

Confinement of high-frequency currents to a local area of the PCB by moating limits the capability of those currents to excite an efficient radiator.

Ferrites are used to directly suppress CM energy by both reactive impedance and absorptive losses.

Slide #26

Moating — Current Confinement Interface Moats & Internal Moats

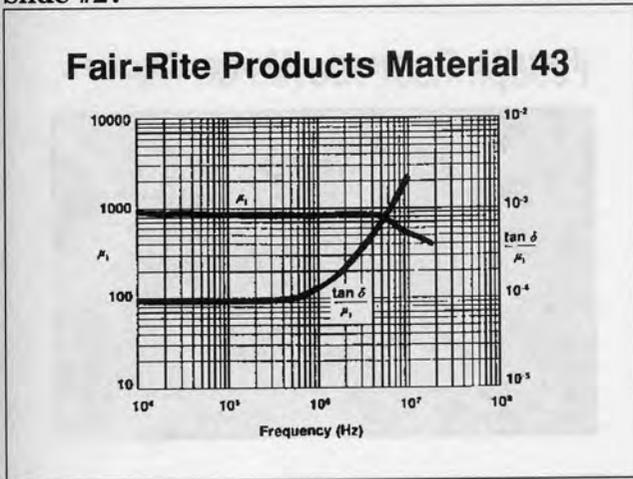
- Confine/corral current dispersion
- Guide I/O currents to a local area
- Bypass heavily for a "quiet" ground

The term "moat" as used here is the *removal* of a strip of copper on the power or ground planes that surround a particular circuit; for example, a switching regulator. The intent is to force the normal ground and power currents associated with the operation of the circuit into a specific area that can then be either heavily decoupled or chassis-grounded. The moat confines the high-frequency currents produced by the circuit so they cannot flow through adjacent circuits, potentially causing interference and radiation.

The moating technique is particularly effective at I/O connectors. The interface cable shield and the connector shell can be tied to the chassis ground directly, and the signal return/ground wires can be terminated at a focused "quiet" ground point that is heavily bypassed. The net result is a very low noise level at the cable interface, which minimizes the potential for radiation from the cable.



Slide #27

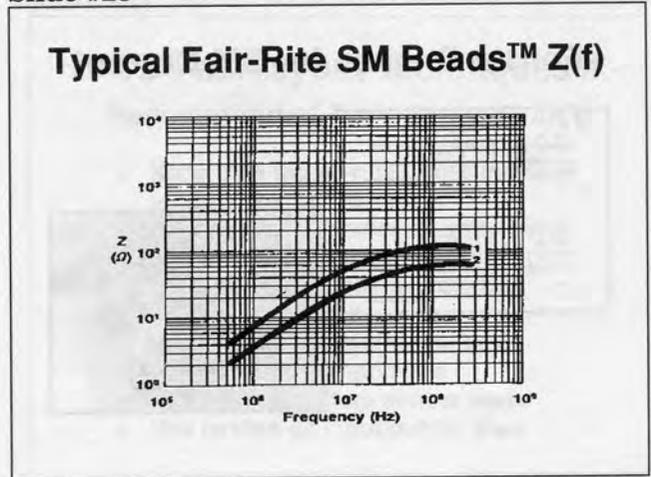


Most ferrite compositions are categorized as "soft" magnetic materials as opposed to "hard," or permanent magnet, materials. As such, they require very little energy to alter their magnetic flux density (B). The narrowness of the familiar magnetic hysteresis curve is an indicator of the amount of energy required to alter the flux. So, the smaller the area within the hysteresis loop, the lower the energy required to traverse the loop. This "soft" nature allows efficient use of ferrites as transformers and inductors at medium to high frequencies. However, as the frequency continues to increase, losses in the ferrite start to dominate its impedance and the inductor looks more like a resistor. This range of characteristics allows ferrites to fill a variety of roles in high frequency circuits.

Both the reactive and resistive (absorptive) characteristics of the ferrite serve to suppress the higher frequencies. These features can be used to effectively reduce the bandwidth and higher frequency energy content of digital signals, thereby suppressing unwanted radiation at high frequency.

The plots above illustrate the initial permeability (always measured at low flux density) and loss factor versus frequency for a popular nickel-zinc ferrite mix recommended for suppression over the 30 MHz to 200 MHz range.

Slide #28

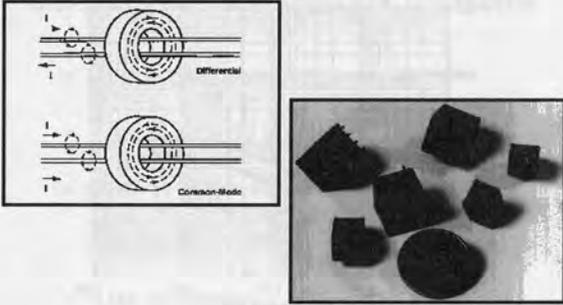


The two curves are representative of two material 43 beads and illustrate the effect of an increasing inductive reactance at the lower frequencies transitioning to a resistive absorption effect at the higher frequencies.



Slide #29

Ferrite Chokes for CM Current



Ferrites are very effective when used to reduce CM currents. By placing the ferrite around a signal *and* the return conductor carrying a differential signal, the fields developed in the ferrite core by the opposing currents cancel; hence, no effect is observed. However, the CM currents on the leads are in phase and their fields add. Therefore, any CM currents on the lines will experience an inductive reaction at low frequencies and a resistive loss at higher frequencies. The resistive (absorptive) transition occurs at frequencies above those where the selected ferrite material would be suitable for normal inductors or transformers.

The relatively low magnitudes of typical CM currents allow multiple turns through (or around) a ferrite core to be used before any threat of core saturation arises. The realized inductance increases as the square of the number of turns until self-resonance is reached and can achieve substantial impedance to the flow of CM current. The absorptive losses lower the Q of the resonance and extends the blocking impedance into higher frequencies, commonly achieving suppression over bandwidths exceeding a decade.

Slide #30

Ferrite Common-Mode Chokes



Ferrite absorbers come in all shapes, sizes, and material mixes. Solid (unbroken), split, hinged, and gapped versions are available for surface and through-hole PCB mounting. They are commonly used in cable clamp-on applications.

Multiple vendors supply ferrites especially designed for suppression at the interface connector pins. Both commercial and military-class connector add-ons are available.

Ferrite cores can also be effectively used in differential circuits for bandwidth limiting and high-frequency damping/absorption. Multiple-lead and multiple-turn ferrite cores are commonly used at digital data line interfaces for both transmit and receive lines. Use of these ferrites suppresses propagation from the transmitters, and incoming transients and CM noise is countered before penetrating the PCB circuits.

For higher current (power) circuits, gapping of the magnetic path through the ferrite is used to substantially extend the allowable dc current through the conductors before saturation reduces the suppression efficiency.



TELEVISION
DRAYTON



Slide #31

PCB Fab/Layout Techniques I Recommended Techniques

- Maximize use of stripline construction
 - Minimize surface conductors
- Border PCB with chassis ground strips
- Centrally locate clock circuits
 - Use remote distribution centers
 - Distribute symmetrically
 - Cancels fields
 - Shortens loops
 - Preserves signal integrity

Buried traces (stripline) confines fields and substantially reduces radiation from interconnections on the circuit board.

Chassis ground rings play a terminator/interceptor role to fields trying to leave the PCB. They also offer a protective intercept to ESD intrusion.

Clock circuits generate the highest toggle rates of all circuits and are the primary source of noise generation in most digital circuits. Clock timing and skew are critical factors affecting digital circuit performance and must be carefully controlled to achieve maximum design margins and robustness. These factors are best controlled by centrally locating the clock oscillator and distributing radially. Radiated fields from the outwardly flowing currents tend to cancel. Propagation delays are minimized and forced to be more synchronous throughout the board. Use of remote drivers for distribution minimizes the trace currents that are forced to travel long runs on large boards; therefore, their respective interference potential is reduced.

Slide #32

PCB Fab/Layout Techniques II Recommended Techniques Cont'd

- Keep high toggle-rate traces in-board
 - up to 20 dB reductions
- Locate line drivers/receivers at ports
 - Reduces CM noise
- Localize high frequency currents
 - Decouple locally
 - Use moating where practical
- Use shielded components
- Use bulk capacitors at/near ports
- Use ferrites on input/output lines

Recent studies by IBM have documented a logarithmic dependence of the radiation from a microstrip (6 mils wide) as it was moved from the center of the board towards the edge. An increase of up to 20 dB in radiation was reported. The same test series performed on stripline indicated no change in the far-field radiation as the traces were placed nearer the PCB edges. Therefore, stripline (versus microstrip) construction offers more flexibility in placement of functional circuits on the PCB.

Locating properly decoupled line drivers and receivers as close as is practical to their physical I/O interface reduces the coupling to other circuits on the board. This placement simultaneously reduces radiation from *and* susceptibility to the circuit board.

Localized decoupling with high self-resonance-frequency capacitors (for example, leadless ceramics) at individual integrated circuit packages confines that device's noise and maximizes noise margins. Where moating of a complete subcircuit is employed, sometimes filtering the subcircuit power via an inductor followed by a large (4.7 μ F to 8.2 μ F) tantalum capacitor can achieve additional decoupling. Of course the bulk capacitor must be accompanied by suitable RF bypass capacitors within the same circuit.



Where possible in high frequency circuits, use grounded, shielded component housings. The housing intercepts errant fields, shunting them to ground. Large plastic quad flat-packs (PQFP) are devices that typically will require installation in a shielded equipment enclosure for compliance, whereas the same devices in a grounded pin grid array (PGA) package may achieve compliance.

Bulk capacitors reduce system power fluctuation effects (typically <10 MHz).

Slide #33

PCB Fab/Layout Techniques III Recommended Techniques Cont'd

- Use narrow traces (4 to 8 mils)
 - Increases high frequency damping
 - Reduces capacitive coupling
- Minimize crosstalk
 - Use orthogonal crossovers for traces
 - Trace spacing-to-height ratio > 2
 - > 3 adds little additional benefit
 - Intersperse ground traces

Narrow traces offer many advantages. For example, these advantages can include higher density interconnections (and more dense packaging), higher Z_0 for lower currents/source loading, increased losses/damping at the higher frequency harmonics, and less coupling to other lines that pass at right angles (crossovers).

Caution: Be aware that for surface traces (microstrip), the magnitude of radiation increases with Z_0 .

Visualizing the field patterns associated with traces at a height H over a ground plane helps illustrate the recommended spacing-to-height ratio of 2 to 3. Closer than 2:1 will substantially increase crosstalk, and larger than 3:1 will impact the allowed density of interconnections.

Adding parallel ground traces better isolates critical signal traces for superior crosstalk and susceptibility performance.



Slide #34

PCB Fab/Layout Techniques IV

Advanced Techniques

- Use of low-dielectric materials
 - Faster propagation
 - Higher Z_0 values
- Multi-Wire™ techniques
 - Kollmorgen, Hitachi
 - High density, controlled Z_0
- Low impedance, Buried Capacitance™ ground/power plane sandwiches
 - Reduction in bypass capacitor count
 - Patented processing technology

Short propagation delays are synonymous with proper packaging for high-speed designs. Signal integrity concerns are minimized and radiation efficiencies are reduced allowing use of thinner traces of higher Z_0 value (lower source loading) with obviously higher trace densities. These desirable features require careful attention to component grouping/placing and must still address the mounting problem of crosstalk as traces are brought ever closer together. Stripline offers a substantial advantage over microstrip for narrow trace-to-trace spacings, especially for separations less than or equal to the dielectric thickness.

The use of lower dielectric materials, such as teflon and Duroid (2.2 - 2.5), Kapton (3.1 - 3.4), and other polyimide (3.8 typical) materials, offer distinct advantages over epoxy-fiberglass materials (4 - 6). These materials yield lower propagation delays, superior (lower) coefficients of thermal expansion, and higher operating temperature capability (>120°C). FR-4 (4.4 - 4.7 typical) requires strain-relieving leads on larger components to prevent long-term solder joint cracking problems caused by thermal cycling. These leads present inductance that compromises high-speed performance.

PCBs constructed with Multi-Wire™ techniques (Kollmorgen and Hitachi) typically combine epoxy-fiberglass layers with traces and polyimide-insulated wires (and sometimes coax) fused into the layered PCB construction. Low- and high-density interconnects and fast, controlled-impedance lines can be melded into the same PCB.

A recently patented innovation (Buried Capacitance™) that allows higher active component densities simultaneously with lowered production costs has been developed by Zycon Corp.



Slide #35

PCB Fab/Layout Techniques V

Advanced Techniques Cont'd

- Buried Capacitance™ features
 - Eliminates 90 - 100% of bypass capacitors (lower parts cost)
 - Higher active device density
 - Single-side SMT assemblies
 - One thermal/solder pass
 - Lower fab. time/cost
 - Increased reliability and MTBF
 - Most effective at 40 MHz and up

The use of localized, distributed capacitance in plate-form has been seen for years in "under-the-chip" parts from Rogers and other companies, but the extension of this concept to entire planes of circuit boards has been difficult because of quality control problems. The Zycon Corporation (See "Recommended Resources" at the end of this paper.) has refined and copatented a repeatable foil-dielectric-foil process that is now being licensed world-wide to implement built-in decoupling in PCBs.

The low-inductance, distributed capacitance nature of the product achieves superior radio frequency decoupling, allowing deletion of discrete decoupling capacitors that exhibit self-resonance, typically between 10 MHz and 30 MHz. By integrating the decoupling capability into the structure of the PCB itself, the majority of the capacitors can be eliminated from the assembly.

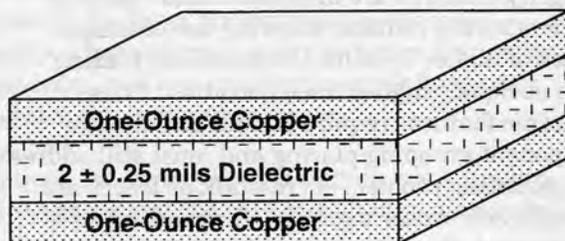
Because active parts can now occupy the space freed up by the capacitor deletions, the effective packaging density takes a quantum step upward. In many cases, assemblies with surface-mounted parts on both sides of a PCB can be transferred to a single side, requiring only one thermal/solder pass during manufacturing. This situation alone can nearly double the production rate of some manufacturing lines.

The parts count reduction can directly save costs, but the greatest gain may be in the improved reliability of the assembly achieved by the use of fewer parts. Further, the technical performance quality and repeatability of the decoupling function is improved by the lack of variations in discrete component self-resonances (and their interactions), and by the superior low impedance presented from a few tens of megahertz on up.

Slide #36

Buried Capacitance Construction

$C = A\epsilon/d$ and is approximately 0.5 nF/in²



On a typical PCB, the V_{CC} and ground planes are each replaced by a "sandwich" containing a V_{CC} plane and a ground plane. The sandwiches are comprised of 1-ounce, premium-grade copper foils separated by a 2-mil dielectric. The premium foil is used for quality control of surface irregularities. Each sandwich yields in excess of 0.5 nF per square inch, and the respective planes (V_{CC} and ground) are connected in parallel. Allowing for via and through-hole clearances, a typical net of 0.9 to 1.0 nF is realized for 2 sandwiches connected in parallel.

As noted earlier, the use of ground planes to reduce circuit loop sizes, provide predictable impedances for the control of signal integrity, and (for the proper construction) provide a shield or enclosure for



Printed Circuit Design Techniques for the Control of Electromagnetic Interference

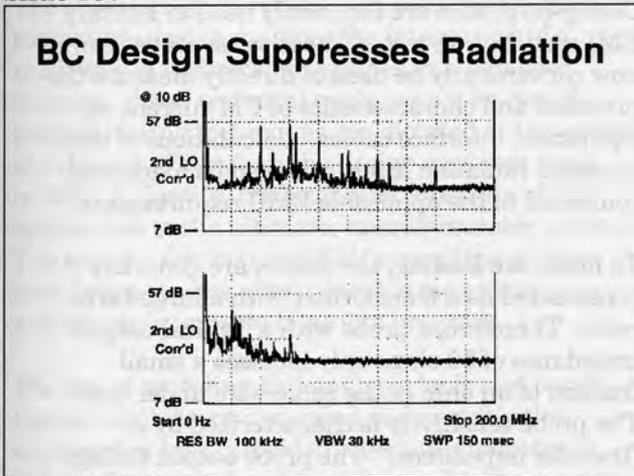
superior EMI performance is critical to the high-speed digital designer. The particular multi-layer PCB configuration shown above carries out those recommendations while adding 2 foil layers (assuming 2 sandwiches for the construction) whose cost is at least partially offset by the reduction in decoupling capacitors required in the final assembly.

For the Mikon-recommended construction, the sandwiches are located on the outside layers of the PCB and offer a *double shield* that should prove superior to a single layer shield of 2-ounce copper.

capacitors. The BC board used only 4 capacitors; namely, two 0.1 μF and two 0.01 μF . Even though low inductance connections were used for the 4 capacitors, their effect was minimal above 30 MHz. The response recorded was located approximately 4 inches from the capacitors; however, the results indicated were reported to be virtually identical to that found at all points on the PCB. Additional capacitors were reported to further reduce the emissions in the 10 MHz to 20 MHz range.

The BC manufacturing process used in the PCB of this example used a dielectric thickness of 2 ± 0.5 mils. Today the process is achieving thicknesses of 2.0 ± 0.25 mils.

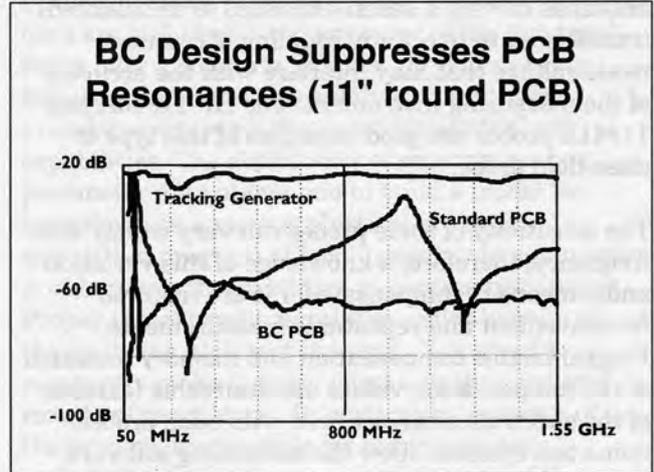
Slide #37



This "Before and After" example of a digital PCB design used in a line of computers is excerpted from the July 1991 issue of *Printed Circuit Design*. The boards used in the test were fabricated with dielectric and copper foils from the same lots, and tested in the same test setup. The standard 11-inch-by-14-inch PCB had 6 layers and the artwork for both boards was identical.

The board circuit contained multiple oscillators ranging from 14.745 MHz to 40 MHz, with the processor clocked at 20 MHz. The standard board was completely loaded and included 141 bypass

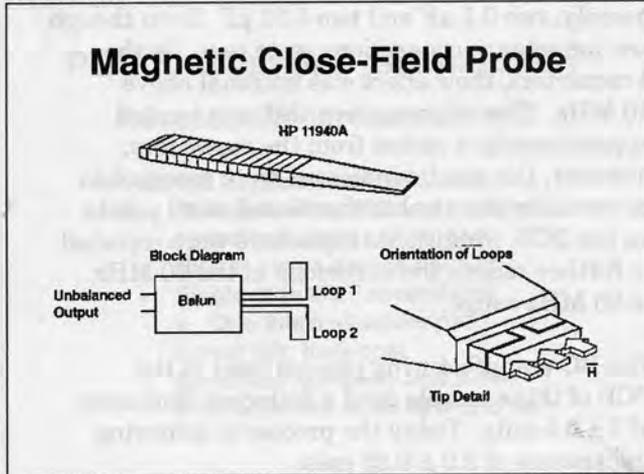
Slide #38



This slide provides a good example of resonance effects that occur on standard PCBs versus the same configuration of PCB using buried capacitance.



Slide #39

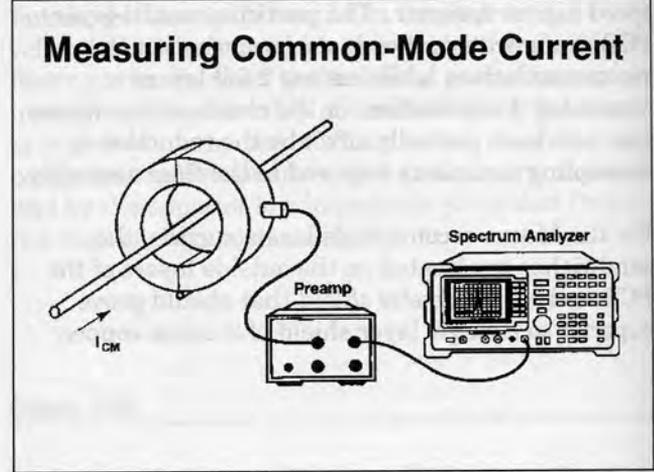


Close-field probes are valuable "sniffers" that allow detection of trouble spots in your circuit. Most probes are comprised of one or more magnetic loop antennas. The better probes use two loop antennas driving a balun (balanced to unbalanced) transformer to reject any coupling of common mode voltage that may interfere with the accuracy of the measuring instrument. The HP 11940A and 11941A probes are good examples of this type of close-field probe.

The sensitivity of some probes can vary widely with frequency; therefore, a knowledge of this variation and a means to compensate for it are required for convenient and repeatable measurements. Programmable compensation and memory retention of the compensation values are desirable features of the spectrum analyzer used with such probes. Some test systems allow the controlling software to compensate for these responses. The HP 859X series of spectrum analyzers, coupled with an EMC "personality" plug-in card, have this capability built in.

Other special probes for measuring surface currents are available. These probes also use balun techniques with magnetic loop antennas. They allow detailed maps of current flow in ground planes, conductive cases, and connector shells to be determined.

Slide #40



Clamp-on probes are commonly used in military EMI certification tests. These same probes can now conveniently be used to directly measure the presence and characteristics of CM current on equipment interface cables. Calculations of the expected radiation levels can then be made and compared to the applicable EMC requirements.

To minimize loading, the probes are generally constructed as a transformer with a large turns ratio. Therefore, a probe with a typical output impedance of 50 ohms only imposes a small fraction of an ohm in the cable/wire under test. The probe sensitivity is characterized by a "transfer impedance." The probe output voltage is the product of the transfer impedance and the common-mode current flowing through the cable/wire.

Slide #41

Conclusion & Summary Electromagnetic Compatibility Requires Knowledgeable Tradeoffs

- Any surface conductor will radiate
- Must meld signal integrity and EMC
 - Z_o and T_{pd} considerations
- Identify potential antennas/radiators
- Visualize, confine, and intercept fields
 - Think "ESD, shields, & skin depth"
- Model & analyze critical areas
- Evaluate with available tools early

Slide #42

Recommended Resources

Software Analysis Tools:

Field-solver for matrix parameters
Circuit simulators for analysis

Test Equipment & Accessories:

TDR system (HP 54120)
Spectrum analyzer system (HP 84100B)
Miscellaneous field probes

The successful generation of an electromagnetically compatible design requires the integration of a multitude of technical disciplines in a balanced manner. Some technical goals *cannot* be met if manufacturing techniques are dictated at the outset of a design. Therefore, the design engineer must be able to identify, trade off, and integrate viable approaches to the ultimate, manufacturable solution. The use of a few hours of EMI consulting at the first design review often prevents major cost and schedule impacts later in a program.

The use of modeling and analysis tools and bench-top test equipment can speed and simplify the engineering task, provided adequate competence and confidence is developed along with the models. "Sanity checks" on the predictions of analytical tools are a must — do not become only "terminal-literate."

A conscious effort should be made to predict the potential problem areas where EMC difficulties may arise. These areas should then be evaluated to a depth commensurate with their potential for disruption of the project. Always remember, "One person's signal is another person's interference!"

Multiple software packages ranging from a few hundred (US) dollars to tens of thousands of dollars are available. Many of the more sophisticated programs require reasonably constant use to achieve operator efficiency, but the less powerful ones are still useful and are more easily mastered. Some offer a separate program for determining matrix parameters and another program for evaluating crosstalk, but most integrate these capabilities. The availability of separate matrix parameter data allows you to build a model for insertion into a more sophisticated circuit simulation program. Circuit simulators again cover a wide range of capabilities at commensurate prices. Proper use of circuit simulators relies heavily on the modeling ability of the user. You should constantly make "sanity" checks on the predicted results of simulators. In particular, be aware of the prepackaged models for semiconductors included with the programs. Also, some of the better programs are *not* SPICE-based, but will import, run, and export SPICE models (like, Micro-Cap IV).

Time-domain reflectometry systems are "designers choice" as the speed, resolution, criticality, and cost factors needed for a given design are subjective. The more sophisticated (and more expensive) systems can sometimes pay for themselves by quickly yielding data difficult to obtain with alternate techniques. As an example, the HP 5412X series offers a range of options.



Printed Circuit Design Techniques for the Control of Electromagnetic Interference

For proper bench-top evaluation of most potential EMI problems, a spectrum analyzer system is highly recommended. A test range in excess of 1 GHz is normally required. Various sizes (and frequency ranges) of probes are available from multiple suppliers; for example, Hewlett-Packard, Electro-Mechanics Company, and Fischer Communications Company.

Mikon Consulting offers original design and troubleshooting services, as well as tutorials in electronic design techniques. Packaging for rugged environments, signal integrity and EMC engineering, mechanical and electrical/electronic system modeling, transient thermal modeling, and proprietary R&D for new electronic products are routine tasks. Complex design problems requiring the integration of multiple engineering disciplines are a specialty. Over 90% of Mikon's clients are repeat customers, testifying to the viability and cost-effectiveness of Mikon's services.

Slide #43

Recommended Resources Cont'd

Buried Capacitance™ PCBs:

Zycon Corporation (USA 408/241-9900)
ISOLA WERKE AG (Germany (0 24 21) 808-0)

Ferrite Components:

Fair-Rite Products (USA 914/895-2055)
European Contact (USA 914/895-1974)

Consulting, Tutorials, Design:

Mikon Consulting (USA 408/727-5697)

Zycon Corporation, located in Santa Clara, California, USA, fabricates printed circuit boards using their patented buried capacitance techniques to achieve higher component densities simultaneously with superior EMI suppression relative to standard manufacturing techniques. Their processes are licensed worldwide—call for the latest list of European suppliers.

Fair-Rite Products Corporation, located in Wallkill, New York, USA, produces a wide variety of ferrite products in several different formulations to cover all practical frequencies. Their products are marketed worldwide — FAX for identification of multiple European distributors.



Advanced Methods for Noise Cancellation in System Packaging

Henri Merkelo

Ultrahigh Speed Digital Electronics
University of Illinois
1406 W. Green Street
Urbana, IL 61801
Phone: (217) 333-2482
Fax: (217) 333-2736
e-mail: hm@lassi.ece.uiuc.edu

1993 High Speed Digital Symposium

Simplified HS Digital Design Process



Substantial improvements in signal quality both at component level and system level can be achieved by appropriately balancing the relative design of digital networks. Cancellation of noise created by components and technologies such as ICs and packaged devices is not detected in the early stages of the design. Delaying the cancellation of noise until later in the design process may result in a significant performance degradation. The author frequently finds that either the cancellation of noise is not considered or is considered too late in the design process. For cancellation to be effective, it must be considered early in the design process.

Henri Merkelo is a member of the Ultrahigh Speed Digital Electronics Research Laboratory at the University of Illinois where he is engaged in developing methods for predicting digital signal integrity in systems of high complexity. Both geometric complexity and network complexity are of issue. Substantial efforts are being expended toward developing simulation techniques capable of including numerous signal degradation effects in complex local networks. He is involved with the electronics industry in applying these methods to the analysis of advanced digital products.

Dr. Merkelo is on the faculty of the Department of Electrical and Computer Engineering of the University of Illinois and is Director of the Quantum Electronics Ultrahigh Speed Digital Electronics Research Laboratory which he established. His principal work has been on the engineering and physics of ultrahigh speed electronic and quantum electronic devices and on the study of propagation of short and ultrashort electronic and optical signals, especially as applied to modern packaging technologies. He has chaired numerous symposia and seminars on high speed technologies, holds patents on ultrahigh speed electronic and optoelectronic devices and has published numerous papers in his field. He has directed a number of short courses for the industry and abroad. He is a member of IEEE and in the electronics industry.

Abstract

Substantial improvements in signal quality both at component level and system level can be achieved by appropriately balancing the reactive design of digital networks. Cancellation of noise created by components, layout and technologies such as

vias, remote grounds and interposer contacts is demonstrated in networks operating from 50 to 200 MHz. This paper discusses the needed cancellation criteria, use of CAE tools and verification of design.

Author

Henri Merkelo

Current Activities:

Henri Merkelo is director of the Ultrahigh Speed Digital Electronics Research Laboratory at the University of Illinois where he is engaged in developing methods for predicting digital signal integrity in systems of high complexity. Both geometric complexity and network complexity are at issue. Substantial efforts are being expended toward developing simulation techniques capable of including numerous signal degradation effects in complex logical networks. He is involved with the electronics industry in applying these methods to the analysis of advanced and future products.

Author Background:

Dr. Merkelo is on the faculty of the department of Electrical and Computer Engineering of the University of Illinois and is director of the Quantum Electronics Ultrahigh Speed Digital Electronics Research Laboratory which he established. His principal work has been on the engineering and physics of ultrahigh speed electronic and quantum electronic devices and on the study of propagation of short and ultrashort electronic and optical signals, especially as applied to modern packaging technologies. He has chaired

technical sessions on the applications of high speed technologies, holds patents on ultrahigh speed electronic and optoelectronic devices and has published numerous papers in his field. He has directed a number of short courses for the industry on high speed digital electronics and has lectured extensively here and abroad. He is very active in IEEE and in the electronics industry.

Advanced Methods for Noise Cancellation in System Packaging

Slide #1

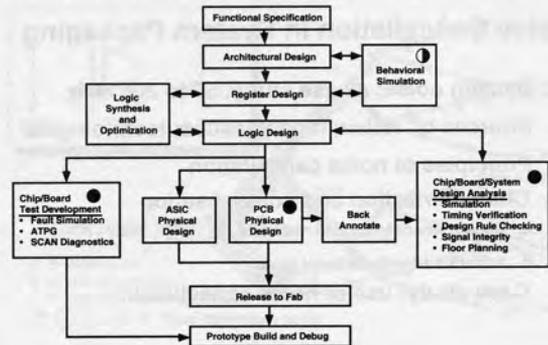
Advanced Methods for Noise Cancellation in System Packaging



Henri Merkelo
 Ultrahigh Speed Digital Electronics
 University of Illinois
 Urbana, IL

Slide #3

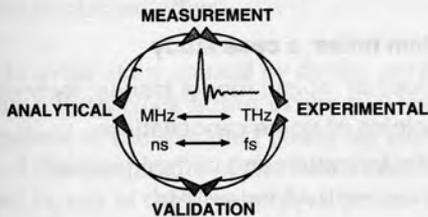
Simplified HS Digital Design Process



Many systems are designed in such a way that it is difficult to provide signal integrity estimation in the early stages of design because the design methods, the tools, and the technologies are not identified sufficiently early. Whereas it is common to have the IC products chosen at the beginning of the design process, the packaging and interconnection methods and products are often not selected in the early stages of the design. Delaying the decisions on the packaging technologies can result in either an inability of predicting system performance or in an inability of meeting performance expectations. The latter is frequently the case but either alternative has significant market implications. For these reasons and for reasons of being able to achieve the full performance potential of a system concept, reviews and assessments of the packaging technologies need to begin in the early stages of the design process. In particular, packaging design team selection and packaging and interconnection technology choices need to be in concert with IC selection and behavioral simulation.

Slide #2

High Speed Digital Electronics Program Scope



GOAL: Develop validated computer models for interconnections and devices suitable for computer simulation of high speed digital networks of high complexity. Emphasis on developing ultrahigh speed capabilities for digital electronics, photonics, and optoelectronics.

Advanced Methods for Noise Cancellation in System Packaging

Slide #4

Outline

Noise Cancellation in System Packaging

- I. System noise: a case study, 50 to 200 MHz
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

The position taken for this paper is that the best treatment of noise is its avoidance; if avoidance is not practical, then perhaps cancellation is possible or, at least, partial cancellation. In a self blaming attitude it can be said that digital system noise is designed-in. Therefore, it's only fair to ask whether it can be designed-out.

Fortunately, by studying the characteristics of noise, it can be concluded that noise, especially reflective noise (which causes the most harm) lends itself to cancellation when avoidance is not practical. This paper develops the needed criteria for noise cancellation. The implications of noise cancellation are discussed and demonstrated both on an individual component level as well as on a statistical basis for system level advantage.

Slide #5

Some Characteristics of Digital Noise

The good news

- Noise adds algebraically within the network of one stage of logic or within one clock distribution network.
- Destructive: $\uparrow + \downarrow = \downarrow$
- At each stage of logic noise is filtered out if it is below the noise margins
- Cancelling effects may leave timing unaffected

The bad news

- Constructive: $\uparrow + \uparrow = \uparrow$
- Noise signals that exceed noise margins may be amplified by successive stages
- Additive effects can propagate and magnify timing errors

Other than noise of thermal and radiative origins, digital system noise is caused by the arrangement and interconnection of system components and, therefore, has specific characteristics. In particular, within the interconnection network of any one stage of logic or the network of a clock distribution, digital signal noise is algebraic in nature in the sense that given an appropriate time relationship, noise adds or subtracts both to other noise and to digital signals. In this context, a rigorous statistical treatment of noise should follow the principles of the statistics of partial coherence of electromagnetic waves if carried out in the spectral domain. In real time, the treatment would be similar but carried out with waveform distributions, transfer functions and correlation functions expressed in the time domain.

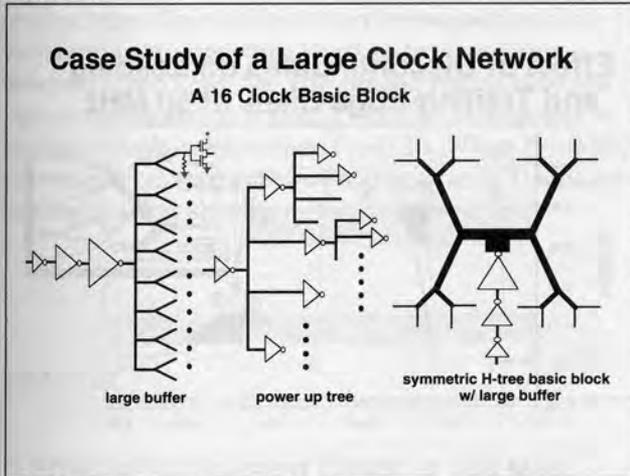
Slide #6

Noise Cancellation in System Packaging

- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

Within the trends of continually increasing computing speeds and continually increasing system complexities, the task of maintaining digital signal quality is also continually becoming more difficult. Even though the design and test tools for signal management have improved dramatically over the last few years, many deliverable systems nevertheless operate at speeds far below the potential speeds offered by the available component and device technologies and, therefore, do not fully reach their commercial advantage. Methods and tools for managing noise at a component level as well as on a system level are described and illustrated with actual case studies.

Slide #7

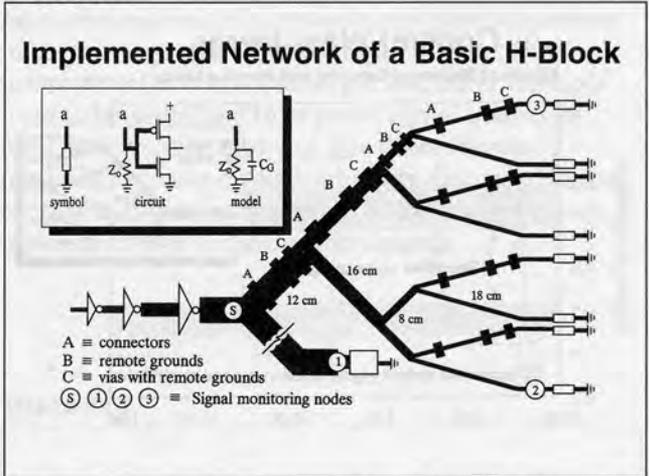


Several aspects of an actual design of a clock distribution network requiring a very large number of clocks, operating at 50, 100, and 200 MHz are used to illustrate and describe several effects. All illustrations are derived from a sub-block of 16 clocks. This number is sufficiently small for gaining some intuitive assessment of the effects of individual components but sufficiently large for observing significant statistical effects.

In order to avoid skew caused by device parameter variations, a large clock buffer driving 16 loads is chosen instead of the so-called power up tree option. Because of the complexity of the boards and multilayer layout of this assembly, a power-up tree is impractical anyway.

Each large buffer is then made to drive a so-called H-tree basic block which is designed to feed sixteen clocks or loads by a succession of transmission line divisions. Each time a transmission line forms a fan-out of two, the nominal impedance is increased by a factor of two. Ideally, in this arrangement, the signal flows toward the load without reflections. The diagram of the basic block is shown with progressively thinner lines suggesting an actual layout of microstrips or striplines.

Slide #8



The network is implemented with a combination of several microelectronic packaging technologies which are discussed later in the presentation. As in all systems of high complexity, the implementation of these technologies is imperfect. In this case, signal and ground paths separate for a short distance at the vias, the connectors are not uniformly controlled impedance, the signal and ground paths in the flex circuits (which are mated by separable interposer contacts) are unequal and the fine end-lines are relatively lossy.

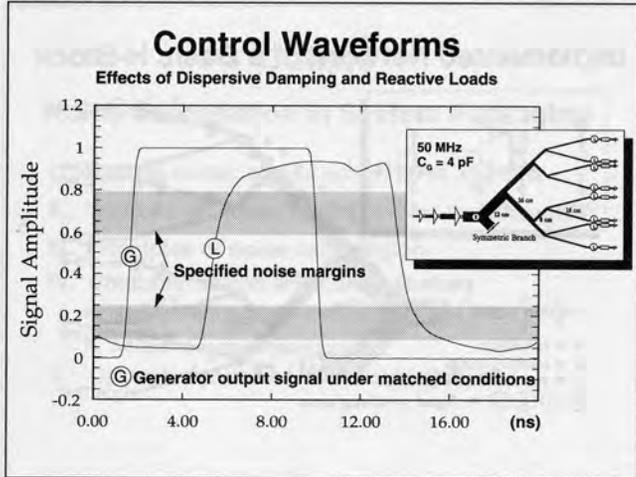
Slide #9

Summary of Network Features

- Half of each H-block is without discontinuities
- Seven arms with three discontinuities each
- Seven vias; seven remote grounds; seven connectors
- **Discontinuities:** $\Gamma_A \approx \Gamma_B \approx \Gamma_C \approx 0.15$; 1.6 to 3.5 cm in length
- **All lines dispersively lossy, particularly end lines**
- $C_G = 4$ pF @ 50 MHz; 2 pF @ 100 MHz; 1 pF @ 200 MHz
- $\tau_r = 0.5$ ns, 50 MHz; 0.25 ns, 100 MHz; 0.125 ns, 200 MHz

The important characteristics of the clock networks are summarized but the determination of these characteristics is saved for a later discussion.

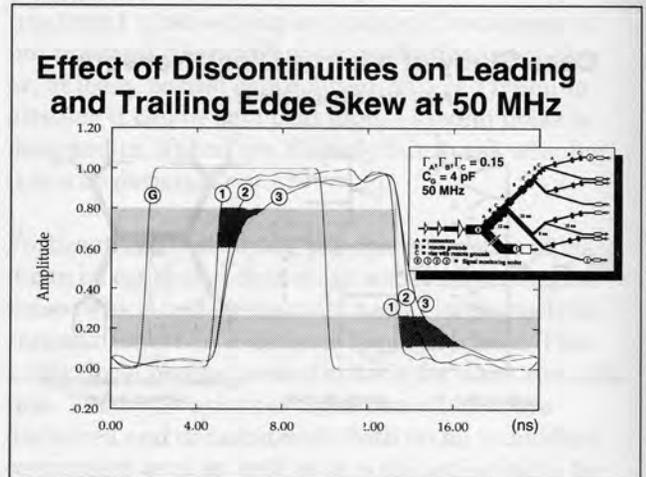
Slide #10



Propagation of signals is first simulated in order to observe the effects of frequency dependent loss, loss induced dispersion and the effects of reactive loads which are, for this purpose, capacitors placed in parallel with resistors $R = Z_0$.

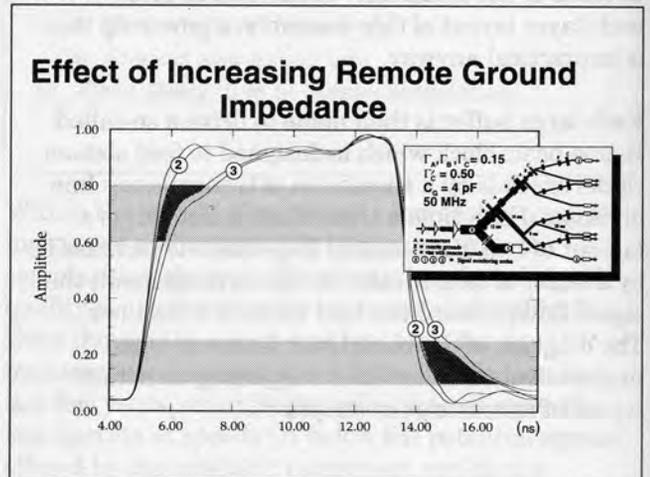
The dominant effects are low level reflections at the fan-out points with mild dispersion, dispersive damping, mild propagation dispersion and dispersive loading which all combine to give the propagation modified waveform shown for 50 MHz. Similar waveforms are obtained for higher frequency clocks when C_G is correspondingly reduced and dispersive damping maintained at the same level. These constitute the control waveforms. The details of carrying out the simulation are discussed in the section on computational tools.

Slide #11



When the twenty one discontinuities corresponding to the preliminary design are introduced into one half of the network (as per design and product specifications), clock signals develop a leading edge skew ranging from 400 ps to 2 ns and a trailing edge skew ranging from 0.5 to 2.3 ns for 50 MHz operation. The degree of skew depends on the level within the noise margins that the devices are actually switching. The locus of skew at each edge is shown shaded.

Slide #12

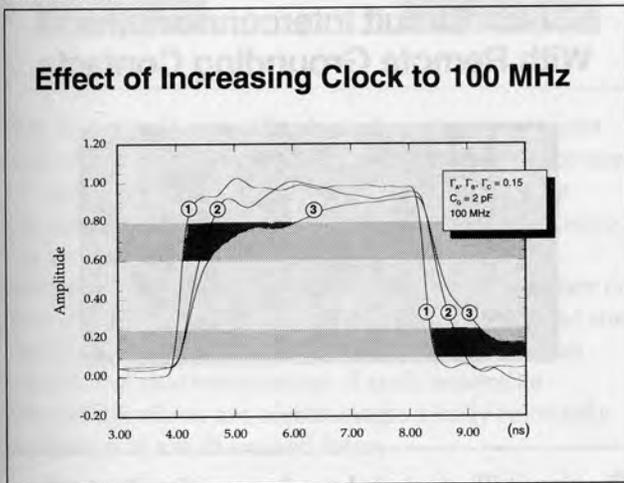


When only one of the connections \textcircled{C} in each branch is made to a more distant ground, increasing the reflection in the \textcircled{C} region to $\Gamma_C=0.50$, the observed skew between nodes $\textcircled{2}$ and $\textcircled{3}$ nearly doubles at the

Advanced Methods for Noise Cancellation in System Packaging

leading edge and also increases substantially at the trailing edge. This is illustrated by comparing the waveforms obtained with $\Gamma_C=0.15$ to waveforms obtained when $\Gamma_C=0.50$. The shaded region between signals monitored at nodes ② and ③ corresponds to the locus of edge skew when $\Gamma_C=0.15$. When $\Gamma_C=0.50$, the waveforms monitored at nodes ② and ③ separate substantially as shown.

Slide #13



The skew effects are accentuated progressively as the clock rate is increased even though the device sizes are adjusted such that the device charging rate does not control the operating rate. Signal ① in all cases serves as a relative reference since it is monitored in the portion of the network which has no discontinuities.

The fraction of the period occupied by edge skew is continually increasing and the most degraded signal ③ is becoming marginally acceptable. Moreover, since signal distortions are not equivalent at leading and trailing edges, pulse skew also develops. For applications that use both edges of the clock and for applications that specify controlled duty cycle requirements, the effects of the discontinuities severely degrade the useful clock rate range of this network.

Problems of this type, both in clock distribution networks and in logic networks continue to increase in severity as clock rates continue to rise. Since interconnections are seldom perfect, other methods need to be developed to improve signal quality in high performance systems. For these reasons methods for noise management are discussed in the context of the resources required for the implementation of known and new techniques.

Slide #14

Noise Cancellation in System Packaging

- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

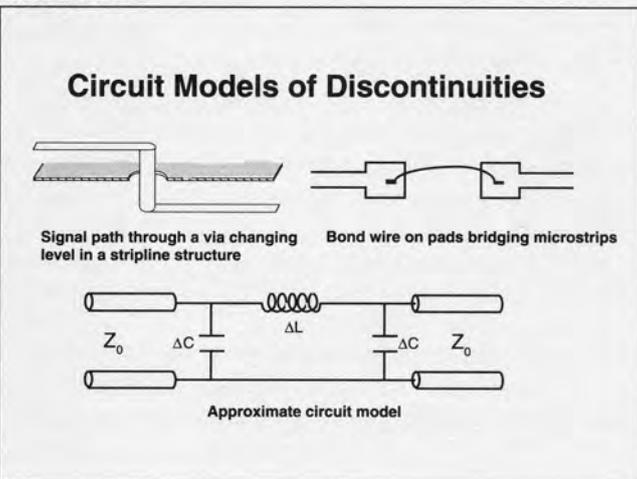
Slide #15

Sources of noise	Cures			
	Low loss lines	Load matching	Isolation	Reactive matching
Reflective @ load		●		◐
Reflective @ discontinuities				●
Crosstalk			●	◐
ΔI (w/ C_o)	◐			●
Ground shift (w/ C_o)	◐			◐

$C_o \equiv$ decoupling capacitors

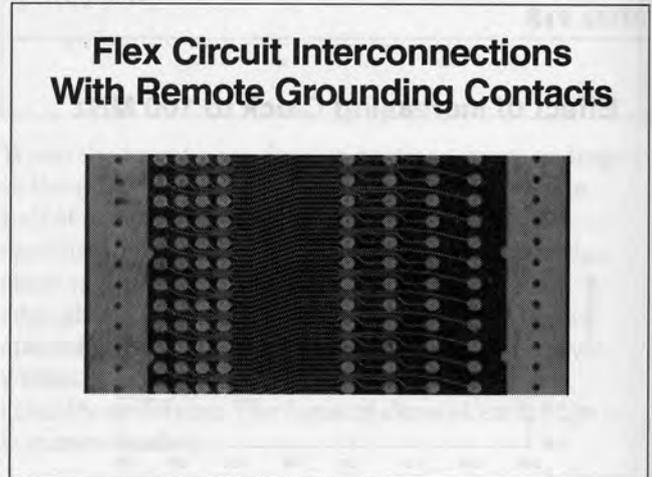
There are a number of causes for digital signal degradation. However, in practical systems, the many causes characterize more the physical description of the systems than the physical phenomena that cause signal noise. In systems being designed today all causes of noise are aggravated by impedance discontinuities and, therefore, can be cured or at least remedied by the reactive compensation techniques that form the subject of this paper. Particular attention is given to impedance mismatches created by such physical requirements as remote ground locations, ground loops, vias, bends, contacts, connectors, etc. which cause signal reflections at the respective locations, create the ΔI noise, accentuate the ground shift effect and accentuate crosstalk problems.

Slide #16



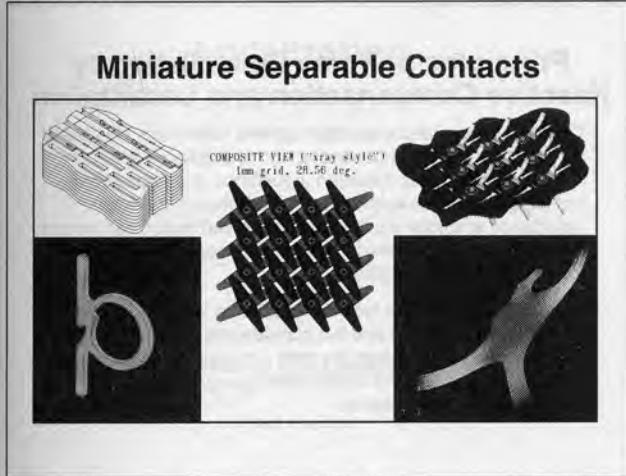
Various features and technologies used for implementing the clock distribution described earlier are shown as component examples. At first, the usual models derived from circuital notions are shown for these components. The circuit models tend to emphasize the notions of excess capacitance and excess inductance.

Slide #17



The circuit illustrated above features bending and flexing qualities and a relatively high packaging density. The connections at the pads and the grounds are made with miniature separable contacts described later. Note that different signal connections are made with different distances to grounds which are shown as wide strips. Such remote grounds add high impedance sections to the controlled impedance signal paths which are along the lithographically produced microstrips.

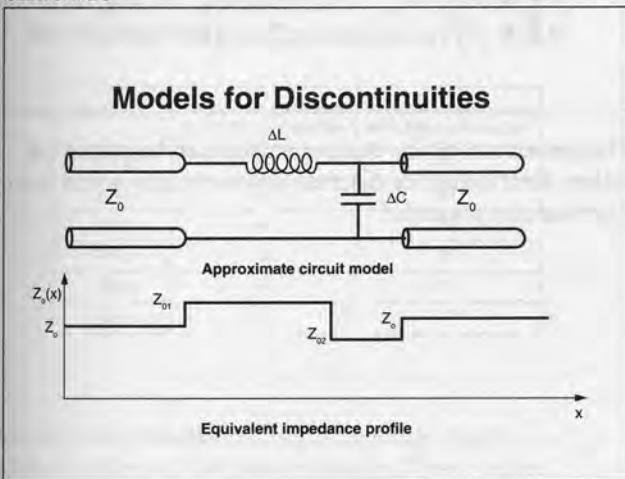
Slide #18



The flex circuit and other high density components such as multichip modules and large single chips are interconnected by unique, miniature, removable contacts used as interposers*. Interestingly, depending on the exact geometric configuration, these contacts can show high impedance, low impedance or impedance matched signal paths. In fact, some of the compensation principles described here were first used in the characterization of such separable assemblies which are electromagnetically relatively complex and are discussed later.

* MicroInterposer and Ampstar are the trademarks of AMP Incorporated for these contacts.

Slide #19



In a controlled impedance environment of a digital system, most signal transport occurs within some nominal characteristic impedance value Z_0 .

However, in various regions of a system network, the physical arrangements may require such measures as providing extra lead length as in bringing a ground from a remote location or providing extra metal surfaces for establishing contacts. Generally, extra lead length introduces excess inductance and extra metal surface area introduces excess capacitance. We have become accustomed to think of excess inductance and capacitance as evil and we call them parasitics. There are some historic reasons for this attitude. We have learned in analog signal analysis in general and in microwave applications in particular that the frequency dependence of the reactance of an inductor is ωL and that the frequency dependence of the reactance of a capacitor is $1/\omega C$. Therefore, any excess of one or excess of the other is incompatible with digital signal propagation since it is well known that digital signals contain many frequencies which are unequally affected by the dispersive nature of these reactances.

All of this, of course, is true but it tends to send the wrong message and misleads us. In particular, referring to inductances and capacitances as parasitics sets an adversary relationship between those who lay out networks and design connectors and contacts and those who specify such requirements as minimum capacitance and particularly minimum inductance. For some reason, inductance has a terrible reputation. It must be partially because of the notoriety of the ΔI noise phenomenon.

In sum, this disposition sends the message that all inductance is bad and that all capacitance is bad and that we should all have less of each. Nothing could be further from the truth. In fact, it will be seen that in high performance systems when excess reactance of one type or the other exists, the system can be improved by adding reactance as a complement. These arguments hold both on an individual component level as well as on a system level. The statistical implications at system level are particularly significant.

No rigorous proof is provided here for all the statements but even the heuristic arguments are more easily made and illustrated when any excess inductance is modeled as increased impedance and excess capacitance is modeled as lowered impedance as shown. The entire network can then be represented

Advanced Methods for Noise Cancellation in System Packaging

as a network of impedance profiles and their concomitant characteristics. Frequently, in electronic packaging, any excess reactance is somewhat distributed anyway and, therefore, showing it as a short section of transmission line is, in fact, more correct. The result of this modeling added to the requirements of incorporating other transmission line features reinforces the need for a simulator based fully on propagation principles as the one used to simulate the clock distribution network which will be discussed later.

This method of modeling and these considerations are helpful in formulating noise cancellation principles.

Slide #20

Noise Cancellation in System Packaging

- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

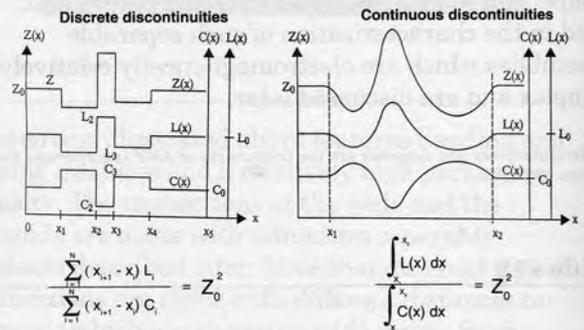
Slide #21

Principles of Noise Cancellation by Reactive Compensation and Localization

- **Compensation:** Since electromagnetic reflections are caused by discontinuities in impedances and impedances are measures of the ratio between inductance and capacitance, it is suggested that restoring the ratio between the total inductance and total capacitance in a given region can restore the matching and eliminate reflections if conditions for relative localization can be satisfied.
- **Localization:** Relative localization can be achieved (even when the mismatched and the compensating regions are not coincident in space) when the total propagation time through the mismatched and the compensating regions is much shorter than signal risetime.

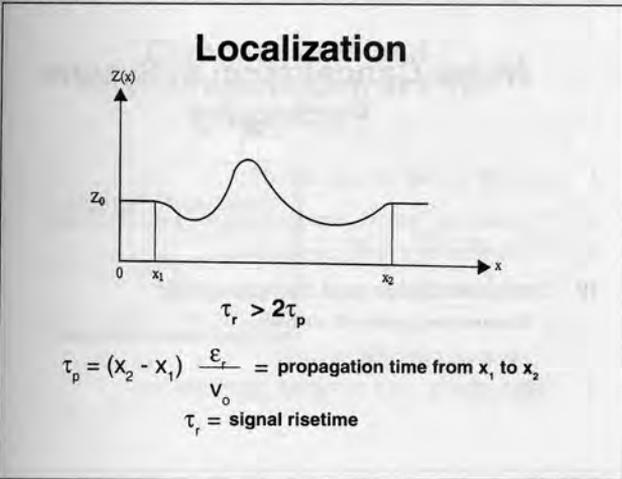
Slide #22

Compensation



The requirement for compensation can be stated for either distributed or discrete discontinuities and even discrete components.

Slide #23



When the condition for compensation is satisfied, noise cancellation is nearly complete when the signal risetime τ_r is greater than twice the propagation time τ_p through the discontinuity. It is important to note that noise cancellation continues to take place even when the localization condition is not satisfied well. However, the effectiveness of cancellation diminishes when $2\tau_p$ approaches or exceeds the value of τ_r .

Slide #24

Example of Localization Design Rules for Partial Noise Cancellation ($\epsilon_r = 4.0$)

Risetime, τ_r	Physical length in mm for ~ 90% cancellation	Physical length in mm for ~ 50% cancellation
1 ns	70	140
700 ps	50	100
500 ps	35	70
300 ps	20	40
100 ps	7	14
50 ps	3.5	7

Ideally, compensation should be done exactly at the location of the discontinuity. Then, when reactive compensation is complete, noise cancellation is complete. That is, wherever there is excess inductance, either the excess inductance should be removed or a corresponding amount of excess

capacitance should be introduced. That's generally the easiest way to provide compensation when these principles are applied at a sufficiently early stage of design. The additional motivation is that, in any given geometry, inductance and capacitance have reciprocal relationships. Generally, by providing additional capacitance, inductance in that region is automatically reduced and conversely.

Of course, such absolute localization is seldom possible, especially with geometrically complex components. Then compensation should be provided within the shortest distance possible to the mismatched region. For reflective noise, the amount of residual reflected energy is proportional to $2\tau_p/\tau_r$ where τ_p is the total signal propagation time through both the existing discontinuity and the compensation region and τ_r is signal risetime. Examples of physical lengths are given for $\epsilon_r = 4.00$ for ~90% and ~50% noise suppression.

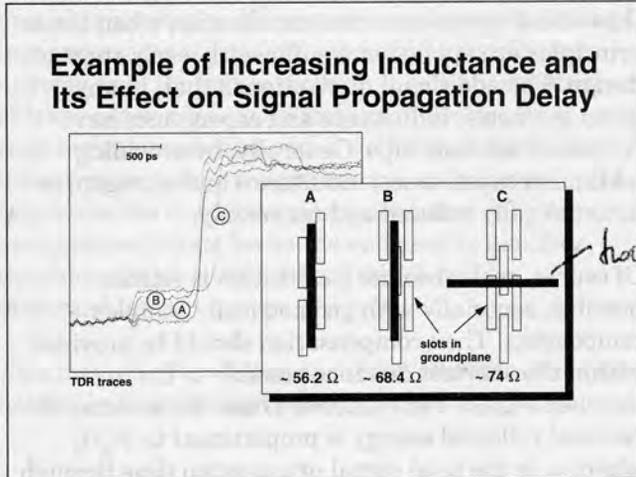
Slide #25

Some Features of Reactive Matching

- Reactive matching applies to both discrete as well as distributed regions.
- Correlation effects provide additional favorable conditions for system level statistics.
- Unlike in situations where resistive matching must be used, reactive compensation is without signal penalty other than in possibly modifying propagation delay.

Since reactive components can return to the system all the energy they store, cancellation of noise by reactive compensation is, in principle, without penalty. But, because there are many ways of increasing the inductance to capacitance ratio, the effective propagation distance may change and, therefore, the effective propagation delay may be affected without an actual change in permittivity or permeability.

Slide #26



As discussed before, reactive compensation techniques are without penalty except for possibly modifying propagation delay. This example serves to illustrate an untypical method for increasing the inductance to capacitance ratio of a microstrip or stripline without either changing the width of the strip or the strip to ground spacing. The connection is made on top of a ground plane perforated with elongated slots which create an anisotropic structure. In this case, the slots also serve as housings for the microinterposer devices described earlier.

Arrangements A, B, and C create progressively higher impedances as a result of modified ground currents and, for the same reason, produce different propagation delays. As seen on the TDR trace, structure B gives the shortest propagation delay and structure C gives a delay 45% longer than that of structure B. These are quite substantial effects, both on impedance and on propagation.

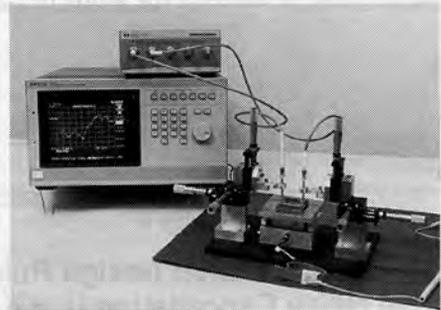
Slide #27

Noise Cancellation in System Packaging

- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

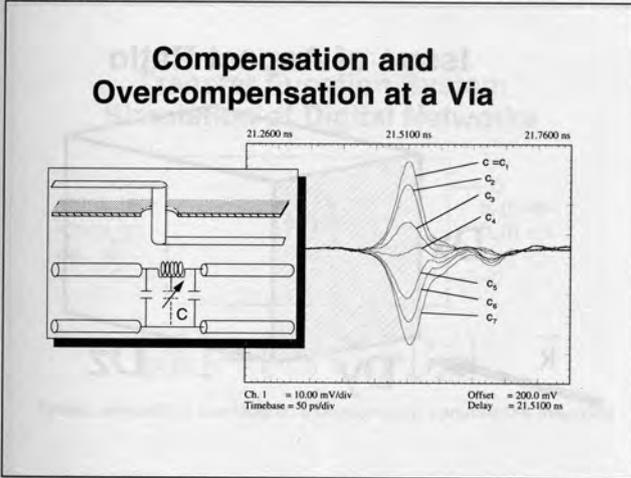
Slide #28

Frequency Domain Measurements on Network Analyzer



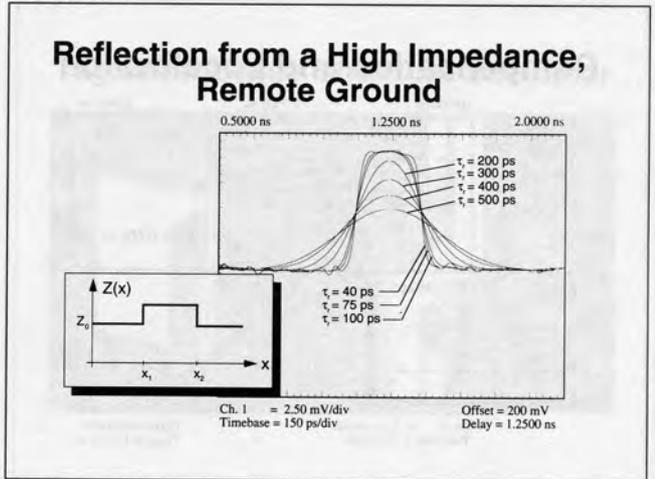
All time-domain measurements and verification of design are carried out on a time domain reflectometry system consisting of an HP 54121A Digitizing Oscilloscope, an HP 54121A Test Set for high speed pulse generation and sampling, an HP 9000, Model 310, computer controller which is networked with the HP-Apollo workstation ring by way of a LAN board and a Thin Ethernet Adaptor. This system is part of a substantial computational effort within which experimental time domain and spectral data can be merged with numerical analysis results for comparison, verification and validation. This facility is also networked directly to the National Center for Supercomputer Applications located near the laboratory.

Slide #29



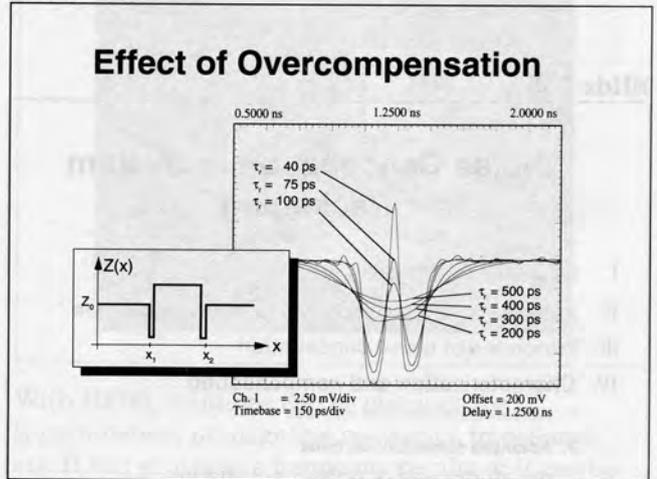
Compensation is very easily demonstrated on a prototype via in which the via region is left open such that additional grounding can be introduced, mimicking an increased capacitance such as produced by decreasing the via hole size. Because the via geometry is relatively small, localization is satisfied even for $\tau_r = 40$ ps of the TDR system. The illustration shows the reduction and near cancellation of the positive reflected signal as capacitance is introduced into the region. As more capacitance is introduced, the via becomes a low impedance structure and gives a negative reflection.

Slide #30



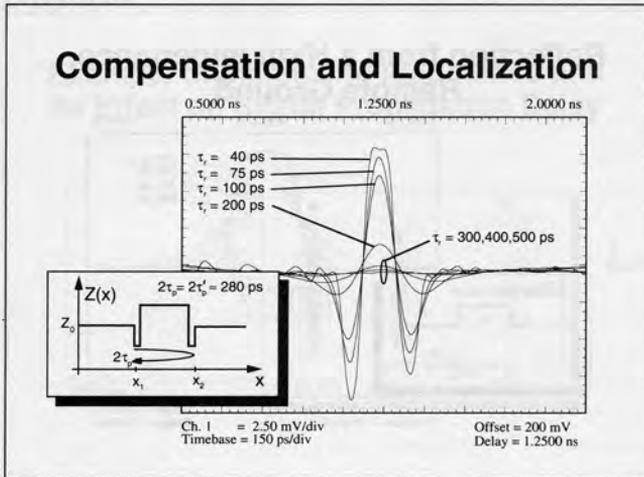
When a discontinuity is of high impedance, such as this remote ground connection of the flex circuit, the magnitudes of reflections from it do not change rapidly as the risetime of the signal is changed as shown on the traces of an HP TDR system. Thus, changing the risetime by more than a factor of ten spreads the reflection in time but reduces the peak magnitude only by a factor of less than two.

Slide #31



Similarly, when short sections of very low impedance are added at the ends, creating overcompensation, the reflection persists even for $\tau_r = 0.5$ ns.

Slide #32



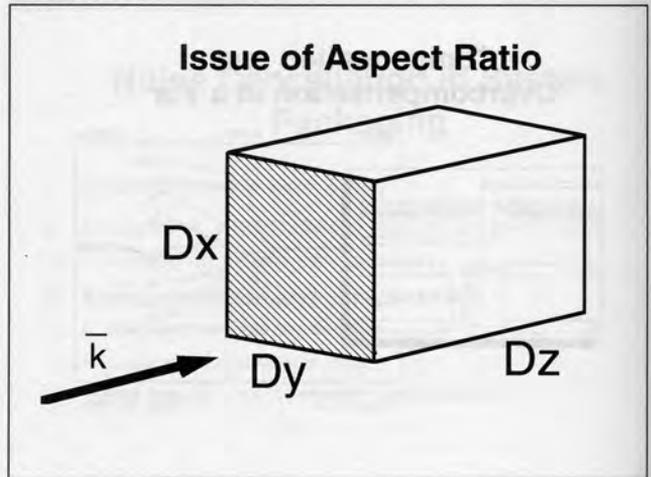
When compensation is of an appropriate amount, high speed signals still resolve the low impedance compensation and the high impedance remote ground. However, when the localization criterion is beginning to be satisfied such that the risetime τ_r starts exceeding $2\tau_p$, the reflection signal is reduced dramatically. Here, the overall length is approximately 3 cm in $\epsilon_r \approx 2.1$ which gives $2\tau_p \approx 280$ ps. Note that when τ_r is longer than ~ 300 ps, the reflection becomes very small and nearly vanishes when $\tau_r = 500$ ps.

Slide #33

Noise Cancellation in System Packaging

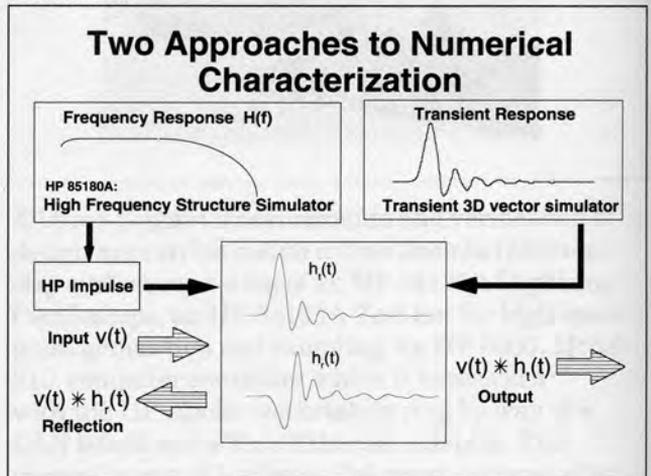
- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

Slide #34



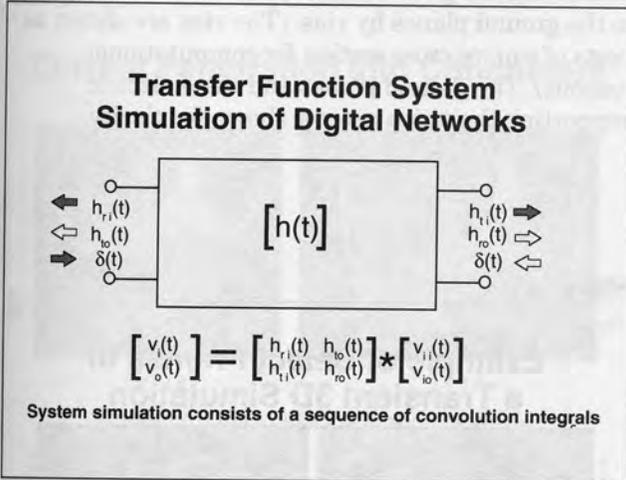
When the transverse dimensions Dx and Dy of an electronic interconnection are comparable to the longitudinal dimension Dz over which the object has a uniform cross section, neither static nor quasistatic principles are adequate to predict its performance for high speed applications. Many packaging features and components fall into that category, requiring a full wave, 3D-electromagnetic vector solver.

Slide #35



Whether the electromagnetic structure solver is based on harmonic analysis or transient analysis, the most convenient approach consists of obtaining an impulse transfer function from which the response to any waveform can be obtained by convolution.

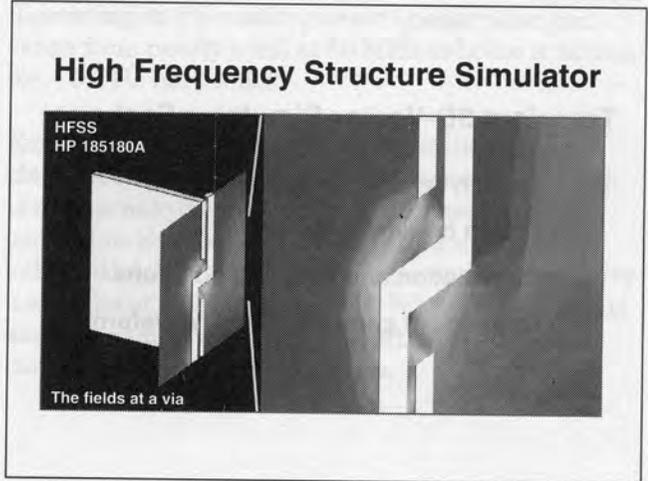
Slide #36



An interconnecting structure can be entirely characterized by obtaining the transmitted and reflected waveforms to impulse inputs. Impulse inputs $\delta_i(t)$ and $\delta_o(t)$ are numerically generated at both ends of the structure and the transmitted responses h_{ti} and h_{to} and the reflected responses h_{ri} and h_{ro} are determined. This method of characterization defines the reflective and transmitting signal response properties of a component and provides an analytical description suitable for the simulation of a network of components.

Thus, for numerical evaluation of the performance of different packaging technologies and methods, both an electromagnetic component characterization tool and a network simulation tool are required.

Slide #37



A broad menu of options is available with the HP High Frequency Structure Simulator (HFSS), HP 85180A, including transfer of files to HP Impulse for the determination of the impulse responses and for subsequent time domain analysis.

Slide #38



With HFSS, solutions can be obtained for S-parameters, propagation constants, impedance, etc. It can give single frequency results or it can be operated in a sweep mode. Clearly, the frequency data of HFSS are the natural complements to network analysis data that is obtained on an HP 8720/8510 Network Analyzer and the time-domain data of HP Impulse are the natural complements to TDR results.

Slide #39

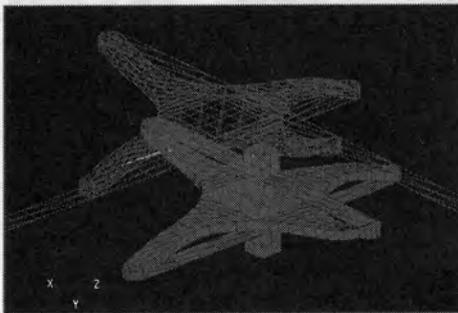
Transient 3D-Vector Simulator Features

- Launch impulse or any waveform
- Obtain transfer functions
- Convolution and post computations
- Graphics of currents, fields, waveforms

The transient simulator gives directly the impulse response and shows directly the distribution of fields when a time domain signal propagates through it. This is not a commercial product but is one of the tools used for characterizing microelectronic packaging products or design of new products in cooperative projects with industry.

Slide #40

Geometry of Interconnections Made with Interposer* Contacts



* Ampstar is a trademark of AMP Incorporated for this separable contact

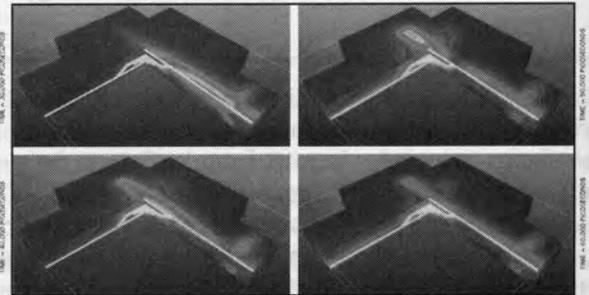
The dynamics of a full wave 3D transient solver is illustrated on a set of interposer connections.

The geometry of the interconnections consists of one star contact for connecting the strips of two microstrips and one star contact (foreground) for

connecting the grounding pads which are connected to the ground planes by vias. (The vias are shown as posts of square cross section for computational reasons). The ground planes and the dielectric supporting the strips are not shown for clarity.

Slide #41

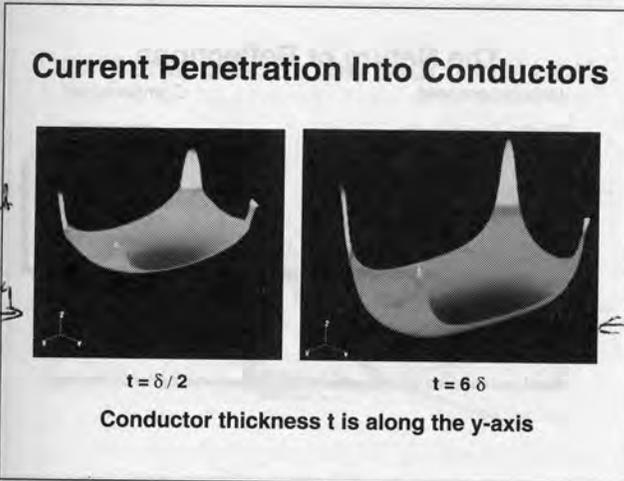
Example of Select Frames of a Transient 3D Simulation



The select frames show the fields developing at the contacts as the signal is approaching. The fourth frame shows a striking example of the field concentration at the tip of the contact as a positive reflection adds to the incoming wave.

Both HFSS and the transient simulator shown here take into account all aspects of the geometric complexity without approximations other than the discretization of space.

Slide #42



leads, current depths in conductors vary. For copper, depending on its quality, current penetration can range from nearly a mil at 50 MHz to below a micron for ~ 5 GHz harmonics.

Examples are shown of the manifestation of the classical skin effect when current penetration depth is comparable to cross sectional dimensions of the conductor. Even when current penetration δ is substantial, $\delta = 2t$, conductor edges and corners carry a significant amount of current. When penetration is small, $\delta = t/6$, a very large proportion of current is carried by the corners and edges.

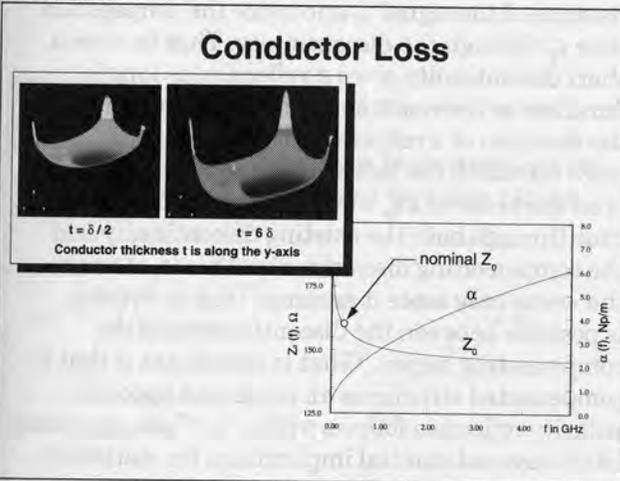
For particularly small conductors and in the presence of high frequencies, variations of the propagation parameters α , β and even the variation of the characteristic impedance Z_0 with frequency may be sufficiently significant to affect the resultant network signal. The particularly thin traces of the 18 cm, end-lines of the H-block network show the characteristics illustrated here.

Slide #44

Network Simulator Features

- Based entirely on propagation principles
- Takes into account dispersive and nondispersive:
 - propagation: $\beta = \beta(\omega) \leftrightarrow \alpha \neq 0$; $\beta = \beta(\omega) \leftrightarrow \epsilon_{\text{eff}} = \epsilon_{\text{eff}}(\omega)$
 - damping with dispersive loss: $\alpha = \alpha(\omega)$
 - reflections with $\Gamma = \Gamma(\omega)$
 - distributed or discrete discontinuities or components
 - linear, nonlinear loads
 - cross talk: dielectrically homogeneous, inhomogeneous
- Post-simulation analysis and graphics

Slide #43



When cross sectional dimensions of conductors are comparable to current penetration depth, detailed current distributions need to be obtained before $\alpha = \alpha(\omega)$, $\beta = \beta(\omega)$, and $Z_0 = Z(\omega)$ can be determined. Depending on the type and quality of metal traces or

After all the components of the intended technologies have been characterized, either analytically (HP HFSS, HP Impulse, Transient 3D, Current Profile) or with instruments (TDR, network analysis), the intended network is constructed for performance evaluation. Again, evaluation can be on a prototype network or with computer aided simulation.

In order to assess the effectiveness of a network design or the effectiveness of various improvements by analytical means, a network simulator capable of including all significant effects is required. For this case study, the clock distribution network is entered into a simulator based entirely on propagation principles and capable of taking into account all the propagation effects of a complex network, including propagation, damping, reflections, distributed as well

Advanced Methods for Noise Cancellation in System Packaging

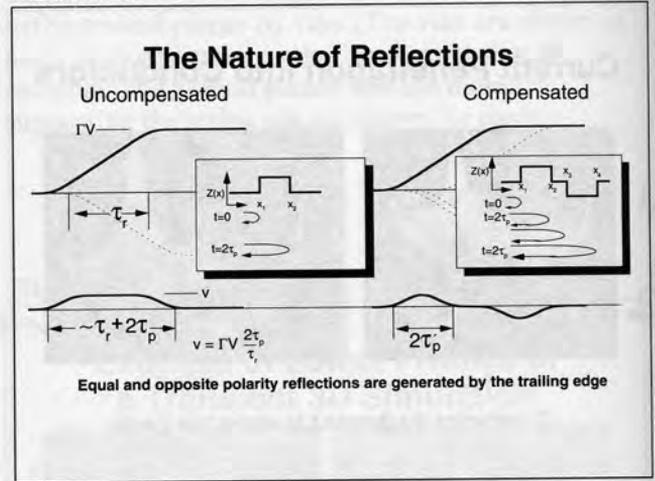
as discrete discontinuities and fan-in and fan-out. Moreover, all aspects of propagation, damping, reflections and loading can be dispersive such as in lossy lines or dielectrically inhomogeneous lines and in reactive loading. Provisions are also in place to take into account crosstalk in dielectrically homogeneous and inhomogeneous networks. This simulator is also not a commercial product but is again one of the analytical tools used in joint projects with industry.

Slide #45

Noise Cancellation in System Packaging

- I. System noise: a case study
- II. Sources of noise; model issues; technologies
- III. Principles of noise cancellation
- IV. Characterization and compensation
 - A. Measurement tools and methods
 - B. Advanced computational tools
- V. Case study: use of noise cancellation

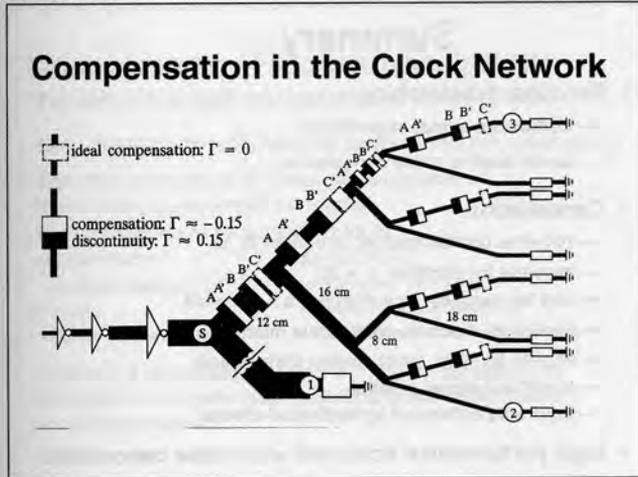
Slide #46



The effect of reactive compensation and noise cancellation was illustrated on individual components by making use of TDR instruments. The collective and statistical effects are best illustrated by implementing compensation on an entire network and evaluating its effectiveness with the network simulator.

The statistical implications of compensation are best illustrated by reviewing graphically the features of the reflections under compensated and uncompensated conditions. The duration of a reflection from a reactive discontinuity corresponds to the sum of the risetime of the signal τ_r and twice the propagation time τ_p through the discontinuity. That is, even a short discontinuity gives a reflection as long in duration as the risetime of the signal. In contrast, the duration of a reflection for a compensated network for which the localization criterion is met is on the order of $2\tau_p$ where τ_p is the propagation time through both the existing discontinuity and the compensating discontinuity. This is, of course, the worse case since it assumes that no overlap is possible between the discontinuity and the compensating region. What is significant is that in compensated structures an equal and opposite polarity reflection follows within τ_r . These characteristics have substantial implications for statistical analysis of system noise content.

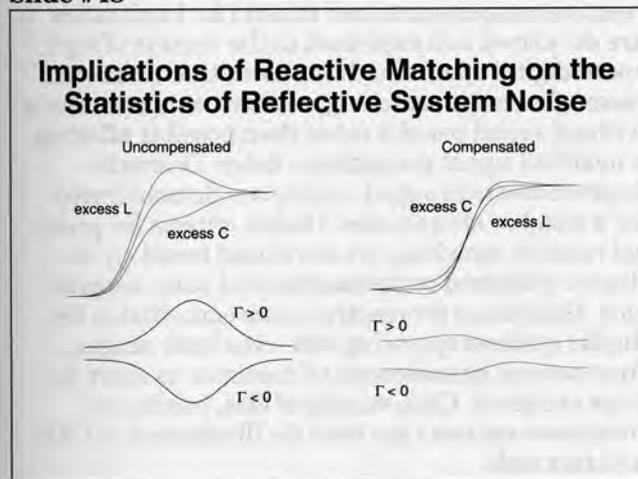
Slide #47



Returning to the clock distribution network, it was seen that the problems of skew were increasing with increasing clock rate. In order to illustrate the noise cancellation technique, the networks are modified as follows.

The remote grounds C are entirely eliminated and redesigned to conform to controlled impedance signal paths. The connectors B, however, and the vias A with remote grounds could not be eliminated or modified. For compensation, each high impedance of A and of B is followed by a low impedance A and B of similar length.

Slide #48



Even though the cancellation of reflective noise is incomplete when localization is imperfect, there is a secondary statistical advantage to the implementation of

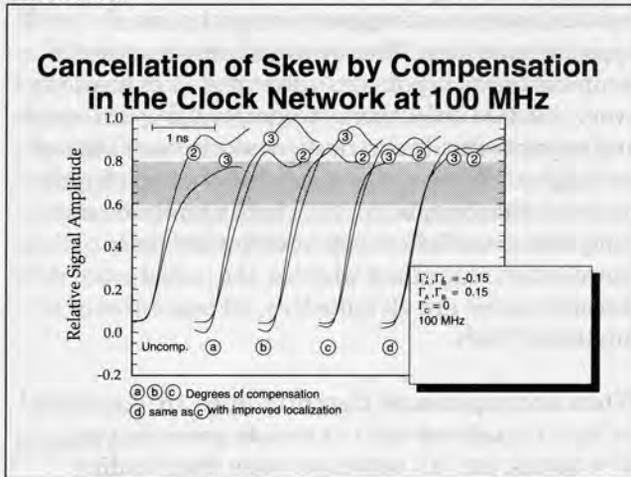
reactive matching since, as already illustrated, the residual components appear as equal pairs of opposite polarities. That is, when compensation is complete but τ_p is not zero (incomplete localization), every residual reflection is accompanied by an equal and opposite residual reflection within each edge of each signal. When a statistical distribution of such residual components exists in both amplitude and time, the overall effect is that of further noise cancellation regardless whether the initial discontinuities are all inductive, all capacitive or a mixture of both.

When uncompensated discontinuities are dominated by high impedance such as remote grounds, vias, wire bonds, etc., all reflective noise from leading edges is additive and is positive and all reflective noise from trailing edges is additive and is negative. The converse is true when uncompensated discontinuities are dominated by capacitive discontinuities.

Therefore, when no reactive compensation is provided, favorable statistical effects occur only when the reflective noise is statistically scattered over a time interval as large as the clock period T .

In contrast, when compensation is implemented, reflective noise needs to be statistically distributed only over an approximate time interval of $2\tau_r$ for favorable statistical effects to exist. The other circumstance which can show statistically favorable cancelling effects without providing compensation is when there is an equivalent mix and statistical distribution of high and low impedance discontinuities. Of course, that is simply a statistical statement of compensation.

Slide #49



The effectiveness of introducing distributed compensating discontinuities for noise cancellation is illustrated by providing different degrees of compensation. Illustrations are for leading edges only, for simplicity. Comparing the leading edges of (a) to the uncompensated edges, it is seen that the network is first overcompensated since waveforms (2) and (3) switch order. Waveforms (b) and (c) show improved degrees of compensation. Significantly, waveforms (a), (b) and (c) correspond to progressively improved compensation but the localization criterion is not met very well, particularly in the end lines. However, when compensation is placed in the end lines in such a way as to introduce compensation at the centers of the discontinuities (effectively decreasing all discontinuity lengths by a factor of two) waveform (d) is obtained which shows a nearly total elimination of skew in the leading edge. Signal quality is similarly improved at the trailing edges.

The statistical cancellation effects are dramatically illustrated by these waveforms since on a component by component basis a much greater attention to detail would have been required, especially regarding relative localization, for good cancellation to have taken place.

Slide #50

Summary

- **Reactive mismatching**
 - dominates signal degradation
 - lends itself to noise cancellation
- **Cancellation**
 - requires compensation: a budget of ΔZ
 - requires localization $\tau_r > 2\tau_p$
 - can be complete or achieved as a trade-off
 - applies to discrete, continuous mismatch
 - applies to logic, clock, power distribution
 - is without signal penalty
 - favorably enhanced by statistical effects
- **High performance achieved with noise cancellation**

Substantial improvements in signal quality both at component level and system level can be achieved by appropriately balancing the reactive design of digital networks. Proposed measures apply equally to logic, clock, and power distribution networks. In order to develop the needed criteria, methods for signal propagation analysis and testing in microelectronic digital networks are summarized and dominant issues relating to digital signal degradation are reviewed. Sources of noise are identified and characterized with particular attention being given to reflective noise caused by reactive mismatching such as remote grounds, vias, connectors, ground loops, etc. It is shown that with the exception of device loading, reactive mismatching is the dominant source of signal degradation in many digital networks that are being designed today. Principles for reactive compensation and criteria for localization are developed and explained in the context of high speed digital operation. It is shown that unlike in cases of resistive matching, reactive compensation is without signal penalty other than possibly effecting a modified signal propagation delay. Dramatic improvements in signal quality are demonstrated for a number of examples. Design criteria for practical reactive matching are developed based on the degree of desired compensation and noise suppression. Guidelines for reactive noise cancellation for digital systems operating with risetimes ranging from several nanoseconds to risetimes as short as 50ps are given. Case studies of vias, bends, and interposer contacts are used for illustration of CAE and test tools.

A system perspective is developed and the effects of reactive compensation on the statistics of system noise are discussed and illustrated.

Slide #51

Conclusions / Recommendations

For achieving high performance at system level

- Early attention to technologies (semiconductor, packaging)
- Selection of vendors and suppliers capable of implementing advanced concepts
- Rich environment of advanced tools

MEASUREMENT

- multiple channel TDR
- LCZ
- network analysis

ANALYTICAL

- 3D vector modeling
- impulse characterization
- simulation based on propagation

- Controlled impedance design with noise cancellation concepts on all networks including power distribution for ΔI suppression

Slide #52

Recommended Resources

- Equipment and accessories
 - HP 54121T TDR Oscilloscope
 - HP 8720/8510 Network Analyzer
 - Probe stations and fixtures
- Simulation Tools
 - HP HFSS
 - HP MDS/Impulse
- Consultant services
 - 3D characterization
 - Current density characterization: loss/dispersion
 - Advanced network simulation; supercomputations
 - System packaging, design, seminars

Pat Byrne

Hewlett-Packard
1900 Garden of the Gods Road
Colorado Springs, CO 80907
Tel: (719) 590-3601
Fax: (719) 590-2251

They have HFSS + one other field solver
with \downarrow HP they modeled the
interposer

1993
High Speed Digital
Systems Design & Test
Symposium

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Outline

1. Introduction
2. Characterizing critical components for debugging ground bounce
3. Debugging ground bounce in high-speed memory systems

Pat Byrne

Hewlett-Packard Company
1900 Garden of the Gods Road
Colorado Springs, CO 80907
Tel: (719) 590-3501
Fax: (719) 590-2251

Author

Current Activities
Patrick is an R&D Section Manager at the Hewlett-Packard Colorado Springs Division. He is responsible for logic analyzer and oscilloscope products that target high-performance applications.

Author Background
Patrick has nine years experience in bipolar ASIC design. He has designed ICs for HP workstations, peripherals and test instruments. He is the co-author of the "Best Paper" at ISDCC 1991 titled "A 4-GHz 8-bit Data Acquisition System".

1993 High Speed Digital Systems Design & Test Symposium

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory Systems Hardware

Abstract

Workstation and microcomputer memory systems have increasing bus bandwidths to support high performance graphics and compute-bound applications. When bus width, speed, and physical densities are improved to accomplish these goals, new hardware failure mechanisms begin to plague the design.

This paper describes a case study in debugging and characterizing a multiple-bus memory system. The physical mechanism of ground bounce is explained and is applied to design techniques that improve the performance and operating reliability of high-speed memory bus designs.

Author

Current Activities:

Patrick is an R&D Section Manager at the Hewlett-Packard Colorado Springs Division. He is responsible for logic analyzer and oscilloscope products that target high-performance applications.

Author Background:

Patrick has nine years experience in bipolar ASIC design. He has designed ICs for HP workstations, peripherals, and test instruments. He is the co-author of the "Best Paper" at ISSCC 1991 titled, "A 4 GHz 8-bit Data Acquisition System."

1993
High Speed Digital
Systems Design & Test
Symposium

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #1

Debugging and Characterizing Ground Bounce Problems In High-Speed Memory System Hardware



Patrick Byrne



This is a case study in debugging and characterizing software dependent hardware failures in a high-speed memory system design in a high-performance workstation. The principles derived from this case study are generally applicable for high-speed designs which are approaching the 50 MHz range, where ground bounce and other high-speed effects start to plague the operational reliability of a system.

Slide #3

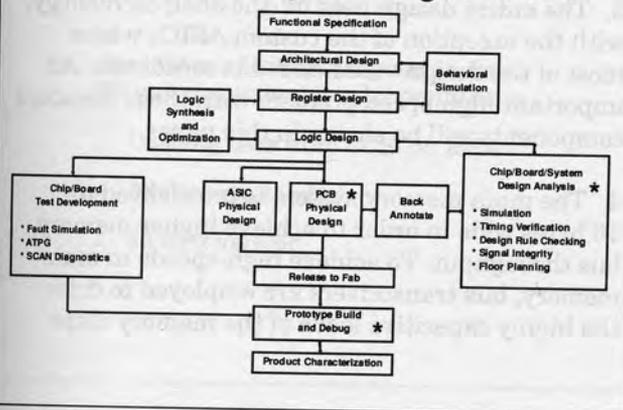
Outline

1. Debugging subtle problems in memory system HW
2. Characterizing critical components for ground bounce effects
3. Design approaches to reduce ground bounce in high-speed systems

The focus of this paper is to give the high-speed digital designer three sets of engineering principles: (1) The Debug Process from a high-level description at the OS crash level down to the root cause found in voltage versus time, (2) The Component Characterization Process where dependencies and sensitivities are quantified and, (3) The Design Process where ground bounce problems can be eliminated from new designs before the PC board and ASIC design is completed.

Slide #2

Simplified HS Digital Design Process



This is a simplified High-Speed Digital Design Process. This paper will focus on the three areas highlighted with an asterisk (*). These are the areas where most of the important ground bounce related issues are addressed for memory systems design.



Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

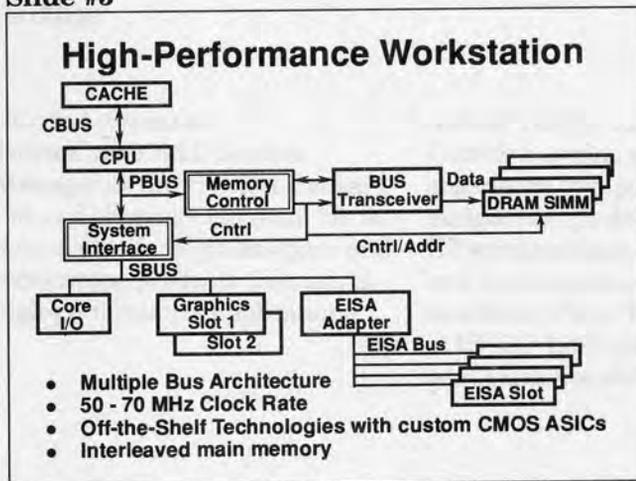
Slide #4

Outline

1. Debugging subtle problems in memory system HW
2. Characterizing critical components for ground bounce effects
3. Design approaches to reduce ground bounce in high-speed systems

The first process we will discuss is the Debug Process. This is where the engineer must repeatably isolate the design fault to the true root cause. The Debug Process found in this case study will highlight the use of a logic analyzer and high-speed digitizing oscilloscope to find the root cause of a SW dependent HW failure.

Slide #5



The Block Diagram of the target system in this case study is shown here. It is similar to most high-performance microprocessor-based systems.

The CPU has a cache bus (CBUS) for off-chip Code and Data cache, while the 32-bit Processor Bus (PBUS) port has the Memory Controller ASIC for reading and writing to DRAM and the System Interface ASIC which interfaces to the I/O subsystem (30 MHz SBUS). This same block diagram can be found in all high-end PCs, with the recent addition of a local bus like VL or PCI, for high-speed video and LAN support.

The key attributes of this block diagram are fourfold:

1. It is a Multiple Bus Architecture, where high-speed transfers are kept close to the CPU while low-speed functions are isolated by "smart" ASICs like those found in the bold boxes. These ASICs control the bus protocol, the interrupt sequence, arbitration, act as temporary storage through on-chip FIFOs, and can act as bus masters for block read/writes (for example, DMA transfers). In high-performance systems, these custom ASICs are often proprietary and are fine tuned to improve performance.

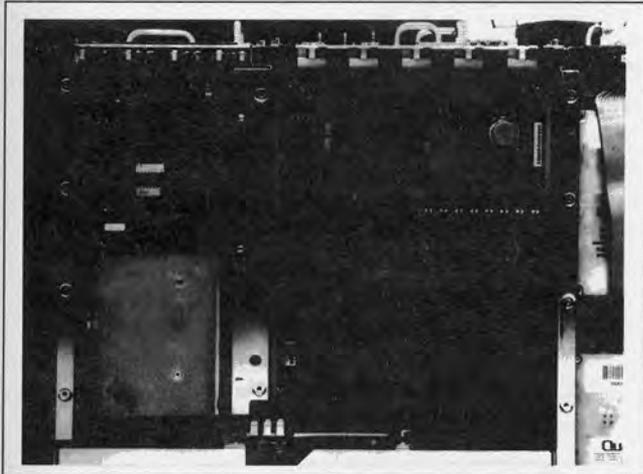
2. The CPU is operating at 50 to 70 MHz bus rates while the secondary busses are slower. At 50 to 70 MHz clock speeds, edge rates are approaching 1 to 2 ns. It is at these fast edge rates that digital designs start to exhibit effects like ground bounce, crosstalk, transmission line effects, etc.

3. The entire design uses off-the-shelf technology, with the exception of the custom ASICs where most of the design value-added is contained. An important high-speed problem with these standard components will be shown in this paper.

4. The main memory system is interleaved to 16 bytes wide in order to achieve higher memory bus throughput. To achieve high-speeds in main memory, bus transceivers are employed to drive the highly capacitive loads of the memory chips.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #6



This is a photograph of the target system. The Main PCB contains the CPU, Memory and System Bus Controller ASIC, and System Interface ASIC. Also shown is the 8 SIMM slots, with 8 Mbytes/slot. Next to the SIMM slots are the bus transceivers. Next to the CPU is the external cache memory. On the left of the picture is the EISA expansion slot, with an HP-IB card.

Slide #7

Problem Description

- **Specific application program causes OS halt due to multiple bit parity error**
 - Loading data from SCSI on EISA
 - 2X memory loaded
 - 50 MHz version

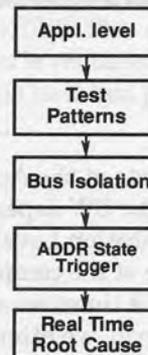
A problem was found in this design during the Prototype Build and Debug Phase prior to shipping. As is often the case, this is the time when an extensive characterization of application SW is combined with a variation of design options, for

example, different I/O configurations and different amounts of main memory. This is late in the design process, as shown in the Simplified High-Speed Digital Design Process. Problems found here can cause major costly redesigns.

In this case, the OS halts with a double bit error while attempting to execute an application from a high-speed, differential SCSI disk drive connected to an EISA expansion slot. It only occurred on the 50 MHz version (and not on the 70 MHz version) of the design and only with new, double-density DRAM SIMMs (Single In-line Memory Module) loaded. SCSI (Small Computer System Interface) is a system-level bus used to connect disk drives, tape drives, and other I/O devices to a computer system.

Slide 8

Debug Process Overview



This is the process used to debug this problem from that described in the Problem Description down to the root cause. This was a subtle problem, only found under special circumstances. These are the hardest problems because of their infrequent occurrence. The key issue in the Debug Process is making the problem occur repeatedly and in dependable sequence and then using real-time, or single-shot acquisition to capture the failure mode. Single-shot acquisition is important in order to capture the cause of the problem, as represented by the critical signals, in a single acquisition. At the high level, with the compiled application code running, the data sizes are large and the code is a

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

black box. However, we must start here because the SW dependencies are our only knowledge of the problem. As the debug process continues, we eliminate the SW dependencies as we discover more determinism in the HW dependencies. As the debug process proceeds, we also reduce the data sizes in the problem set.

Slide #9

Application Level

- Limited because code is compiled and long reboot times
- Isolate type of activity by monitoring CPU execution
- Isolated to DMA transfer / memory read

The Debug Process starts at the Application Level because this is where the SW dependence is created. Debugging at the Application Level is difficult due to the black box nature of the compiled application code and the long reboot times needed to recreate the problem. At this level, we attempt to identify the type of activity being performed by the application using a state analyzer to monitor CPU execution. On-chip cache can limit this capability because it hides CPU execution. This target system does not have on-chip cache so all execution is observable. It is determined that the error occurs on a memory read of data that was loaded using DMA (Direct Memory Access) transfer from the SCSI disk drive. During DMA transfers, the Memory and System Bus Controller ASIC operates as PBUS and SBUS master to transfer large blocks of code/data into main memory, without involving the CPU in the transfer.

Slide #10

Test Patterns

OS Running:

- Replace crashing application with test program.
- Binary search to reduce to Minimum Failing Sequence (MFS).

Modify Boot Code:

- Reduced to one SCSI block.
- Eliminate long reboot time. Add read to boot.
- Boot up to HW initialization.

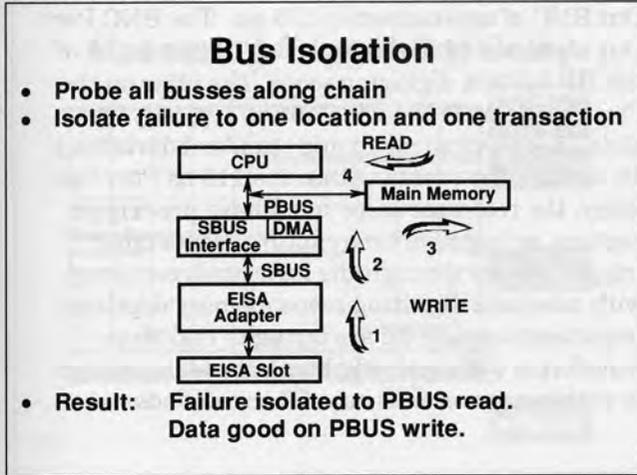
To further isolate the problem, we replace the crashing application program with a small test program that loads the same data identified by the CPU execution trace. This first step of the Test Pattern phase still has the OS loaded. To find the Minimum Failing Sequence (MFS), we reduce the problem using a binary search approach. In this approach, half of the data is eliminated on the first run, three-quarters by the second run, etc. This continues until the failing portion is inadvertently eliminated. The binary search is the fastest way to isolate the problem to data located in one SCSI disk block (1 K-byte). Sometimes a simple binary search routine can be inadequate because important sequential dependencies are eliminated. In this case, the search is complicated by the fact that ten percent of the time the error does not appear even when rerunning a previously failing sequence. It is later discovered that this anomaly is due to the asynchronous timing between when the MFS is executed and when the DRAM in main memory is refreshed. Each time the sequence is reduced, the test pattern is run and the assembly code is observed at the PBUS port.

Having reduced the MFS data sequence to one SCSI data block transfer, we now focus on eliminating the tedious and time consuming process of rebooting the OS every time. The boot code is modified to perform the SCSI block read once the critical HW registers have been initialized. We now have the ability to quickly and reliably recreate the problem, allowing us to begin the true debug process using real time acquisition tools.



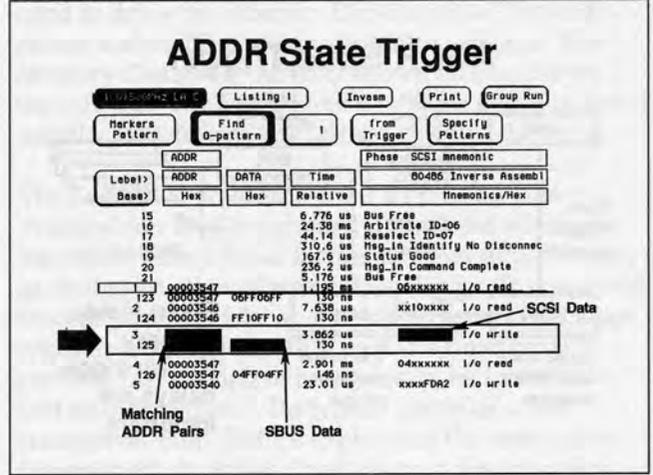
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #11



The Bus Isolation phase is intended to isolate the problem to a specific bus transfer. There are four possibilities in this case: (1) The WRITE from the EISA slot through the EISA adapter to the SBUS, (2) The WRITE from the SBUS to the PBUS, through the System Interface ASIC, (3) The WRITE into main memory through the Memory Controller ASIC, (4) The READ from main memory into the CPU. Since the ASICs in this design are custom designs, one big concern is the verification of the ASICs over operating SW. The ASIC verification tools (R&D IC testers) can rarely cover all the SW dependencies. In this case the problem was isolated to the WRITE/READ sequence labeled 3 and 4. The data and address were known to be good at the PBUS on the WRITE (#3) into main memory during the DMA sequence. The data was bad on the main memory READ (#4) after the Memory Controller ASIC had passed bus mastership to the CPU. Probing all sides of the bus transactions is key to confident isolation. The data and address pairs must be matched to verify a bus transfer. Sometimes this is difficult because of long latencies through the bus ASICs. Up to this point, only a state analyzer is required to isolate the problem.

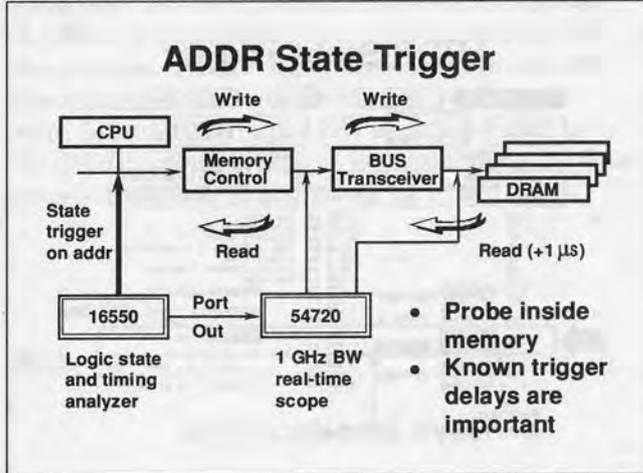
Slide #12



An example screen shot from the HP 16550A state analyzer is showing address/data pairs during bus transfers. Looking at the outlined box, you will see an I/O WRITE operation at the SCSI port. It is ADDR 00003544 and DATA XXXXXXXE0. The next line shows the ADDR/DATA pair on the SBUS port. Note the matching ADDR and DATA sets. The DATA is 000010E0. This capability to track DATA/ADDR pairs is required in the logic analyzer to complete the Bus Isolation phase.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #13



By this point in the Debug Process, we have isolated the problem down to: (1) a WRITE/READ sequence to and from main memory, observable at the CPU bus (PBUS) and, (2) we know that the DATA pattern is corrupted and, (3) we know the specific data bits which are incorrect. We need to know if the DATA is corrupted on the WRITE into main memory during the DMA transfer or on the READ out of main memory by the CPU. For this, we need to probe within the main memory system.

We probed the data transfers through the Memory Controller ASIC and the bus transceivers as shown in the slide. As this point, we started using an HP 54720A Real-Time Digitizing Scope, triggered from an HP 16550A Logic Analyzer, to look at the signal quality. We suspected signal quality problems because of the data dependencies and the highly capacitive loads within the DRAM SIMMs with the long transmission lines connecting DRAM ICs. The real-time scope and the logic analyzer must have repeatable trigger delays, even when the delays are long, on the order of 1 μs from ADDR trigger to real-time scope capture. These long delays with known, small jitter are required because of the long WRITE to READ latencies. In this case, the WRITE to READ latencies are on the order of PBUS arbitration times, hundreds of nanoseconds. Long delays will be encountered whenever bus mastership must be reassigned, in this case from the System Interface ASIC during DMA transfer to the CPU during READ operations.

The HP 16550A has a Trigger Out delay to the Port Out BNC of approximately 115 ns. The BNC Port Out is connected to the External Trigger input of the HP 54720A digitizing scope. The jitter on the Port Out delay is less than 150 ps typically (see slide 15 for typical mean and standard deviation). To capture the events within the 115 ns Port Out delay, the real-time scope must have pre-trigger capture, or negative time capture. Repeatable trigger latency through the logic analyzer along with adequate digitizing scope memory depth are important because we are trying to find scope waveforms with perfectly known time correlation to within nanoseconds out of microseconds.

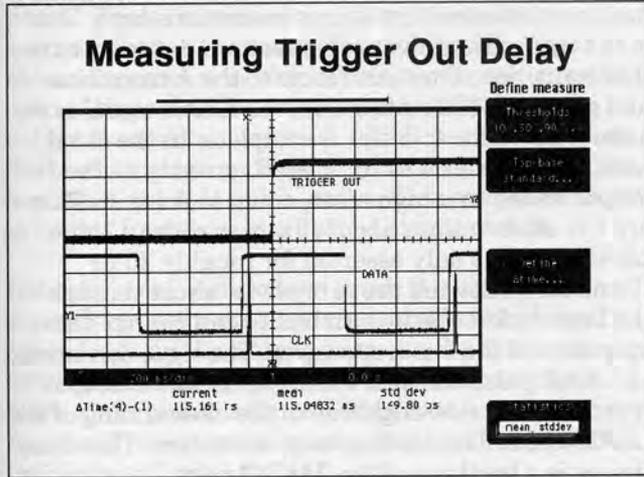
Slide #14



The measurement setup used for the debug phase consists of the HP 16500A Logic Analysis System (with one HP 16550A 100-MHz State/500 MHz Timing plug-in) and the HP 54720A Digital Oscilloscope (with two HP 54712A 1-GHz BW plug-ins and one HP 54721A 1-GHz BW plug-in with external trigger). The system under test is also shown.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

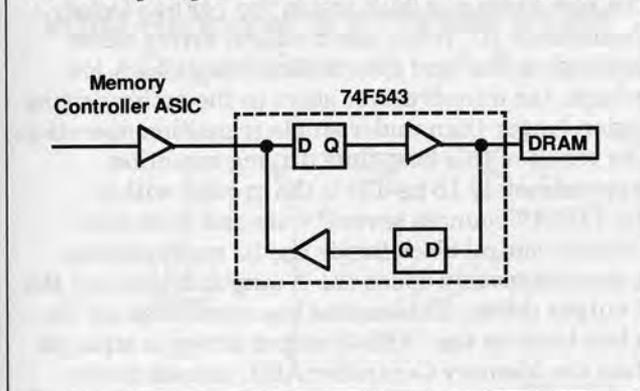
Slide #15



An HP 54720A scope screen shot capturing the HP 16550A Port Out delay is shown here. The mean and standard deviation of the delay is shown in the box. In this case, a common clock signal is sent to both the scope and the logic analyzer as a time reference and the delay is characterized. This Port Out delay is not part of the HP 16550A data sheet specifications but is important to know when performing cross-domain debug measurements.

Slide #16

Memory System Detailed Schematic

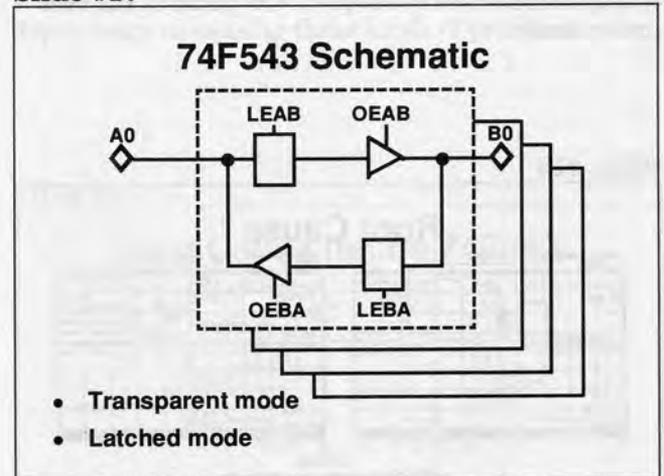


The root cause of the problem is found within the main memory system during the WRITE operation. The DATA passes through the Memory Controller ASIC without corruption but is found to be bad on

the output of the 74F543 Octal Bus Transceivers used to drive the Memory DRAM chips. The slide shows a simplified schematic of the problem. The Memory Controller ASIC is shown as a buffer on the left. The 74F543 is shown in some detail in the middle. It is driving the DRAM chips to the right.

The 74F543 is an off-the-shelf TTL Octal Bus Transceiver. It is a standard component which we have used many times before, in other applications, as well as in other places in this design. It is sourced from several vendors and is sold in several package types as a commodity TTL part. Its function is to perform fast driving of heavy loads, be tristatable, and be bidirectional. Its typical use is as a bus transceiver chip. Before explaining the root cause failure mode in detail, I will discuss the operation of the 74F543.

Slide #17



The 74F543 has 8 pairs of input/output bidirectional pins, labeled A0 to A7 and B0 to B7. Data is transferred from A-to-B or B-to-A through D-type level-sensitive latches. The latch can be enabled, in byte width, with LEAB (latch-enable-A-to-B) or LEBA (latch-enable-B-to-A). These signals are active low and should never be asserted simultaneously. The output drivers can be tristated using the OEAB (output-enable-A-to-B) and the OEBA (output-enable-B-to-A) signals. The B driver is capable of 70 mA DC load typically while the A driver is capable of only 25 mA DC load typically.

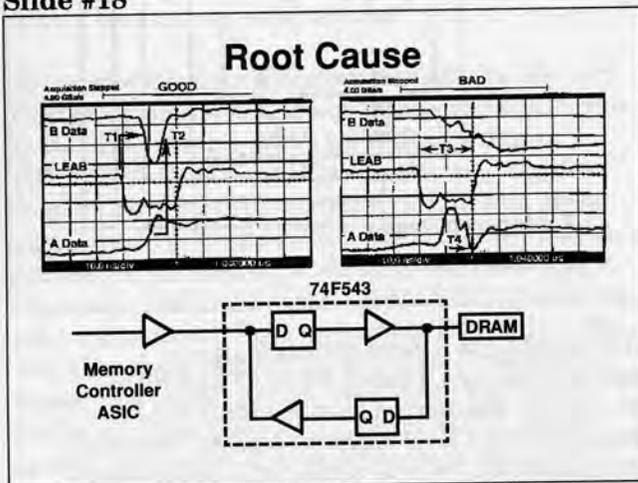
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

There are two normal operational modes, transparent and latched. In the transparent case, the latch is put in transfer mode and the corresponding output driver is enabled. A data transition on an input will cause a transition on the corresponding output, 5 to 10 ns delayed. In the latched mode, a data transition has already occurred on an input but does not appear on the corresponding output until both the Latch Enable and Output Enable signals are asserted. The typical Enable to Output transition time is 10 to 15 ns.

These two modes can legally be used in sequence. For example, data could be "waiting" on an A port in the latched mode and, after the Latch Enable signal is asserted to transfer the signal to the B port, the A signal can transition to the other polarity, utilizing the transparent mode.

It is this sequential case, under specific state, timing, and loading condition which caused the data corruption and the subsequent OS panic in the case under study.

Slide #18



The root cause is shown in the Good and Bad HP 54720 scope shots. In the Good case, the capacitive loading of the memory system was lowered while the Bad case is the one under study. Referring to the Good case, the 74F543 starts in the latched mode, with LEAB not asserted. The A and B data are in opposite states (A low, B high). When LEAB is asserted (going low), the B output starts to transition low (T1) to match the A input. The delay is, as

expected, about 10 ns. 10 ns after the LEAB signal has been asserted, the A data transitions high. This is the transparent mode of operation. Subsequent to this transition, B output reacts to the A transition and goes high (T2). After this, the LEAB signal is de-asserted and the transfer is complete. In the Good case, the transfer was completed correctly — the B output ended in a high state, equal to A input. There are two observations about this case so far, (1) the LEAB signal is only asserted for roughly 20 ns (10 ns/div), allowing the A input to "shoot through" the latch before the bus controller can change the operation of the bus transceiver. The legal minimum on LEAB pulse width is 4 to 5 ns, (2) The B output "glitches" high-low-high due to the close timing of the LEAB and A data leading-edge transition. This close timing is a legal use of the 74F543 part.

Referring now to the Bad case, where we have increased the memory loading to the double density case. In the Bad case, the B output starts to transition low, following the LEAB assertion. However, it stays low after LEAB is de-asserted. This is the data error! The A input, meanwhile, has suffered some kind of data corruption, leaving it in an undetermined state (neither high nor low). There seems to be bus contention on the A input, where one output driver is fighting another (one high the other low). We know that the Memory Controller ASIC is trying to drive the A input high, so it must be that the 74F543 is trying to drive the A input low and is partially winning the fight.

The root cause is a fault inside the 74F543 Octal Transceiver IC. When the B output driver slews fast high-to-low and then is forced to go back low-to-high, the output driver stays in the active driving region longer than under single-transition operation. The result of this long-time driving condition (approximately 15 ns-T3) is the ground within the 74F543 bounces several volts and then the Tristate control block inside the IC malfunctions. It simultaneously opens the A output driver and the B output driver. This causes bus contention on the A bus because the 74F543 output driver is stronger than the Memory Controller ASIC output driver. In fact, the Memory Controller ASIC output driver

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

has been put into Weak Drive Mode during this phase. Weak Drive Mode is where the output current sinking or sourcing capability of a chip is reduced occasionally to limit the power and current spikes. Now that the A input is corrupted, the B output is uncertain, as well. LEAB is de-asserted at the median point on the scope screen shot (T4). This is when A input is at its lowest. This causes B output to latch a low state, storing the incorrect data.

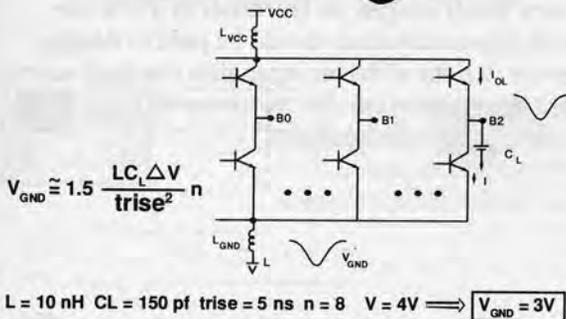
Note that the conditions to cause this failure are dependent on the state, timing, and loading conditions. All these combine to create the B output glitch conditions which make the tristate control block malfunction and cause bus contention on the A input. These are the most subtle and infrequent problems to find!

The problem is captured in real-time on an HP 54720A digital real-time scope running in 4 Gs/s (1 GHz real-time bandwidth) mode with HP 54721A and HP54712A plug-ins. A real-time high-bandwidth scope is required to capture the exact timing and voltage levels of the three signals. The timing and the voltage levels are critical to understanding the root cause. All three signals must be captured simultaneously to develop and verify the ground bounce/tristate control malfunction theory.

denoted B0 to B2. The L_{VCC} and L_{GND} are the package parasitic inductances in series with the power supply and ground leads, respectively. When an output goes through the high-low-high transition shown in the previous scope shots, the output current going through the pull-down transistor achieves a peak value determined by the capacitive load (C_L), the voltage swing (delta V), and the transition time (Trise). A factor of 50% high (1.5 multiplier) approximates the conduction overlap current found in TTL totem-pole output drivers when the upper transistor is still on while the lower transistor is driving low. The voltage drop across the ground inductance is controlled by the peak current spike, the inductance itself (L_{GND}), and the transition time. A factor of n multiplies the voltage swing. N represents the number of lines switching simultaneously. In the interleaved memory driving case, the peak value of the internal chip ground is -3V! This is not atypical in high-speed systems if the values are as shown. Note the square law on the transition time (Trise). This is the most important dependency in causing these kinds of problems.

Slide #19

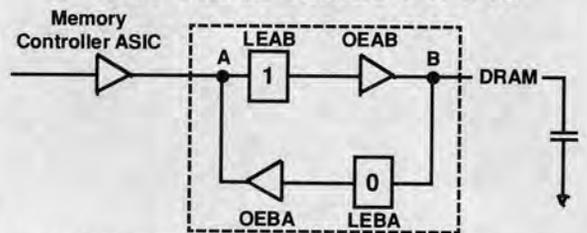
What Happens to a 74F543 When Output Goes ?



This shows the insides on any TTL totem pole output driver, in this case the 74F543, which has eight identical outputs, three of which are shown here,

Slide #20

Root Cause Inside 74F543



- DRAM loading and glitch switching causes 3V ground bounce
- Tristate control malfunctions
- Memory controller has Weak Drive Mode. 74F543 overpowers ASIC
- Only happens with critical LEAB to A input timing and data and control state conditions

Going back to the 74F543 block schematic. When the ground bounces to -3V, the tristate control block loses control of the LEBA output driver, causing it to temporarily turn on. Bus contention then exists between the Memory Controller ASIC and the 74F543 output driver. If the ASIC had a strong drive, it would have won the fight and there would

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

not have been a failure. However, the ASIC was in Weak Drive Mode during this phase of the transfer. Weak drive is used whenever the bus needs to respond quickly, while keeping current spikes to a minimum within the ASIC. The result of these loading, timing, state, and ASIC drive conditions is the error reported as an OS panic on the subsequent READ of this memory location.

The criticality of the LEAB-to-A input timing was discovered when we sought to understand why the 70 MHz version of this computer did not exhibit this failure mechanism. We discovered that the LEAB signal arrived 400 ps later relative to the A input due to different routing on the PCB (~ 2 inch longer route). This later timing allowed the critical ground bounce to stay within the acceptable range of the tristate control block of the 74F543. Even though this is a computer with 20 ns cycle time, parasitic effects within the devices are sensitive to timing 50X smaller!

Slide #21



This is a photograph of the new HP 16517A 4-GHz state and timing analyzer that works in the HP 16500A Modular Logic Analysis System. Using this analyzer in timing mode, entire octal transceivers can be probed simultaneously and glitches can be triggered on. The oversampled state mode can be used to find the intermediate transitions that can cause ground bounce. The exceptional time-interval accuracy, similar to a high-speed digitizing scope, can be used to characterize the critical simultaneous switching events, like LEAB-to-A data timing.

Slide #22

Principles of Memory System Debug

- Reduce OS code to MFS of data/addr pairs
- Analyze bus transfers with repeatable test patterns
- Trigger real-time scope from logic state analyzer. Trigger delays are critical
- Standard components can exhibit data dependent failure mechanisms. Beware of varying drive strengths on ASICs
- Ground bounce is present with highly capacitive loads and is especially sensitive to transition times
- Critical timing can be as low as 1% of the period

To summarize the Debug Process we have used, we need to start with the SW level at which the problem is reported, in this case the OS level. Reduce the failure to the MFS of DATA and ADDRESS pairs which you can track across bus control ASICs. Use a logic analyzer cross-triggering a high-bandwidth single-shot scope to capture the critical waveforms. The trigger delays must be well known to have accurate time correlation. A single-shot scope is required to capture all the important signals simultaneously on the failed condition.

Standard components like the TTL 74F543 Octal bus transceiver can exhibit data dependent failure mechanisms under certain critical conditions. Ground bounce is a particular problem because of fast transition times and highly capacitive loads. The timing conditions which cause the failure can be very small ranges, on the order of 1% of the period. Close attention should be paid to designs where a variety of driver strengths are used since weak drive modes can be overcome with glitching, but strong, tristated outputs.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #23

Outline

1. Debugging subtle problems in memory system HW
- ➔ 2. Characterizing critical components for ground bounce effects
3. Design approaches to reduce ground bounce in high-speed systems

We now move to the second process which the high-speed HW engineer must use to ensure high quality designs. This is the Component Characterization Process. We will isolate the 74F543 Octal bus Transceiver in a controlled stimulus-response characterization system to understand all the dependencies and sensitivities. The goal of this process is to understand all of the subtle timing and noise margin characteristics of the design so that the surrounding PCB and timing environment can be designed correctly.

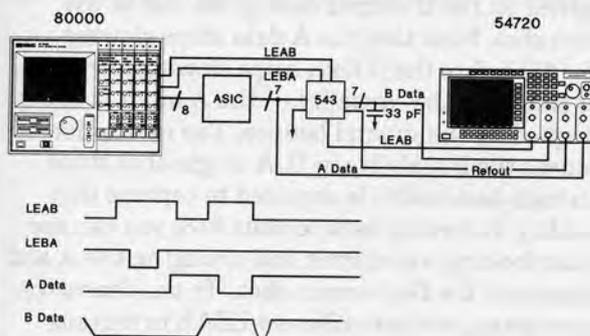
Generation System is a modular, high-performance Data Word Generator. It is capable of less than 200 ps edge speeds and 2 ps edge placement resolution. In the characterization process we are performing, we need to have better than 100 ps edge placement resolution so that we can characterize the ground bounce sensitivities to the critical LEAB-to-A input timing. An ASIC, with a Weak Drive Mode, drives the 74F543 to emulate the driving conditions of the case study. The HP 54720A digitizing scope is used because it was used in debug and it is helpful to capture repetitive single-shot events to show the changes in the response due to small changes in the controlled stimulus. This process will highlight the timing sensitivities to the ground bounce failure mechanism. We use the HP 54712A 1 GHz plug-in for the characterization so that we can see all the subtle wrinkles in the waveforms.

Seven of the DUT's outputs are loaded with capacitors to emulate the DRAM ICs. The characterization was completed with a range of capacitor values, from 15 pF to 150 pF. Note that all A inputs and B outputs switch simultaneously. Ground is monitored in this test setup by tying one output always low. This saturates the pull-down transistor and provides a handy ground monitor for ground bounce measurements.

The timing diagram shows the setup of the state conditions. The data transitions on the four inputs (LEAB, LEBA, A data, and B data) are needed to place the 74F543 DUT into the correct state conditions for the failure to occur. After these state conditions are achieved, the LEAB-to-A input timing is varied to change the amount of the B data glitch low and therefore the ground bounce.

Slide #24

Characterization Test Setup



The Characterization Test Setup uses the HP 80000 Data Generation System for accurate edge placement resolution. The HP 80000 Data

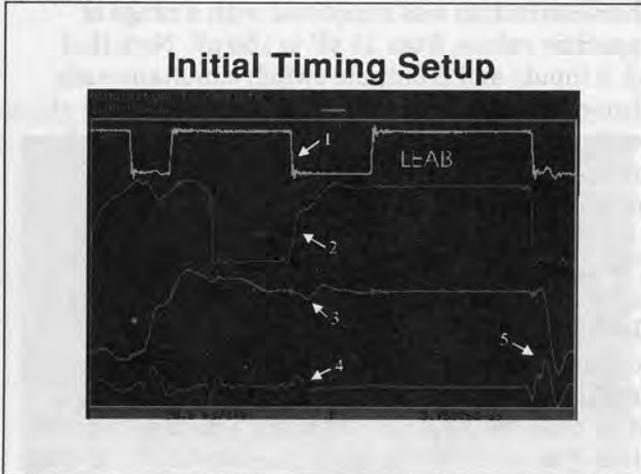
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #25



level ground bounce on the screen, in this case 2.61 volts ($V_{max}-V_{min}$). Although this is the worst-case ground bounce, it does not cause a failure because A and B ports are both in the same condition. Ground bounce exists in many places and times when it will not cause a problem. A failure requires the correct state conditions, as well.

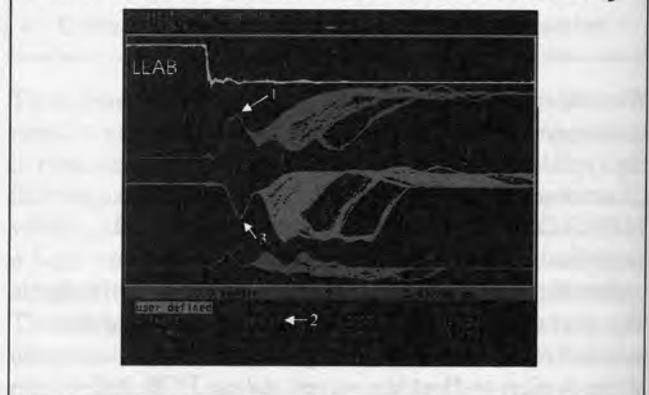
Slide #26



The Initial Timing Setup shows a screen shot from the HP 54720A. There are four waveforms shown. Prior to the #1 transition on LEAB, the 74F543 is in the latched mode and the A input and B outputs are in opposite states, like the case study. When the LEAB is asserted low near mid screen (#1), the B output starts to transition low (#3) but is returned high due to the A input transitioning high (#2). Note the ground bounce at this time (#4). Note also that the worst-case ground bounce happens later (#5) when all the A inputs and B outputs slew low simultaneously when the DUT is in the transparent mode. Voltage markers have been placed on the V_{GND} signal to record the worst-case high and low

Slide #27

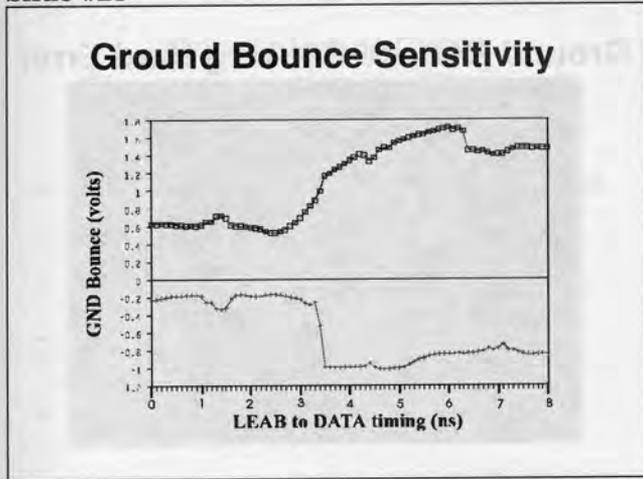
Ground Bounce Caused Instability



This screen shot from the scope is taken when LEAB falling edge is 4.2 ns before the A input rising edge. This is recorded in the delay time measurements at the bottom of the screen shot (#2). At 4.2 ns, the low going B data output is unstable due to ground bounce. The ground bounces enough to make the LEBA output driver open up, corrupting the A data input (#1). This corruption creates the unstable behavior on the B output during the rest of the screen shot. Note that the A data stops slewing high (#1) before the B data stops slewing low (#3). This indicates the causality — the A input is being corrupted by the ground bounce. The corruption of A causes the instability in B. A single-shot scope with high-bandwidth is required to capture this causality. Referring back to slide #18, you can see similar looking waveforms and timing on the A and B signals in the Bad screen shot. In the characterization setup, we have allowed LEAB to remain asserted throughout the instability so that we can record the duration and extent of the instability.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #28



This slide shows the sensitivity of positive and negative ground bounce to LEAB-to-A input timing (over an 8 ns range). Data points at the far left side of this graph (small LEAB-to-A input timing) correspond to the condition shown in Slide #26 where no failure occurs. Data points indicating much larger ground bounce with LEAB-to-A input timing greater than 3.5 ns correspond to the condition in Slide #27 where instability begins to occur. The 700 ps difference (4.2 ns versus 3.5 ns) is caused by cabling to the DUT (approx. 5 inches difference) and could have been calibrated out with the scope.

The critical information in this graph is the rapid increase in ground bounce over just 200 ps of LEAB-to-A data timing skew! This increase is caused by phasing of ground bounce with edge slewing to create the worst-case conditions. You must design your PCB routing and component characterization with this kind of accuracy to ensure good designs.

Slide #29



This slide shows two single-shot measurements to illustrate the tristate breakthrough in the design. The waveforms are recorded as sequential single shots. Examples of the two could be found within the many shown in Slide #27. Here two of them are isolated to highlight the ground bounce effects more clearly. Ground Waveform #1 corresponds to the lower A and B data waveforms, while Ground Waveform #2 corresponds to the upper A and B data waveforms.

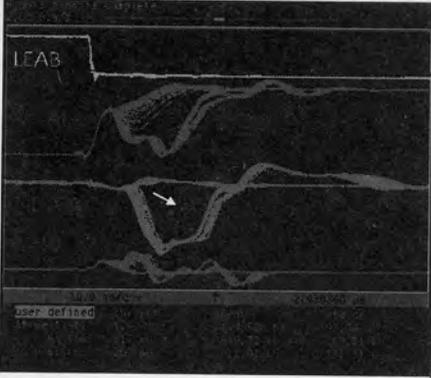
This screen shot illustrates an important feature of ground bounce. Ground always bounces in the direction to stop the desired transition by limiting the available drive current. When the B outputs are slewing low, the ground will slew high to cut off the output pull-down transistor, limiting its pull down current capacity. The subsequent negative slew on ground is in reaction to the positive slew on B data. You can see that the number of phase changes on ground (three times when there is a positive or negative peak on ground waveform #1) is equal to the number of phase changes on the B output slew edge.

As you can see, if the LEAB signal had been de-asserted after 20 ns of assertion time, the final states of A and B ports would be different for the two cases shown. Note that the separation points (labeled #3 on B data) of the A data and ground signals are approximately the same while the B data is 5 ns later. This verifies the ground bounce theory. A scope with good time interval accuracy between channels and good noise performance is important to see these waveforms in the correct order and shape.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #30

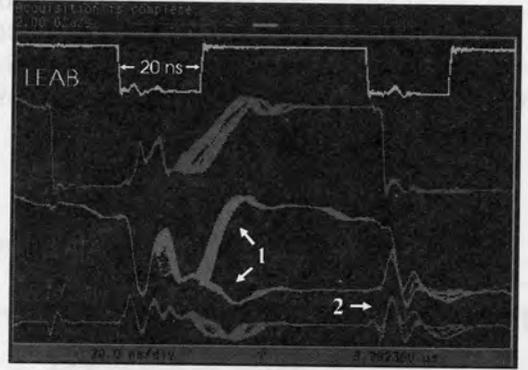
Exponential Temperature Sensitivity



One of the features of TTL parts is that the bipolar transistors driving the output have temperature sensitivities which tend to be exponential. The effect that this has in this case is dramatic. The DUT was cooled down to low temperature and then allowed to self-heat. Self-heating typically will force a temperature change of several degrees C over a few seconds. The scope is put in infinite persistence to capture the transition from one state to the other as the device self-heated. The repetition rate on the scope is approximately 100 ms. The part snaps from Bad (Cold) to Good (Hot) over only 2 single-shot captures, corresponding to only 200 ms of self-heating. This corresponds to only a few degrees C change in temperature. Ground bounce effects are very sensitive to temperature.

Slide #31

Ground Bounce Causing Hard Error



This screen shot shows the failure occurring with the LEAB asserted for only 20 ns, as in the case study (see Slide #18). Since the A input is in an indeterminate state, the B output can go either high or low, depending on how the D latch regenerates the B output (#1). The shapes of the waveforms are similar to those in the actual failure. Note the large ground bounce at time #2. As stated earlier, large ground bounces can be acceptable, depending upon the state conditions. At time #2, the A and B ports are equal so there is no failure, even though the LEBA tristate buffer is probably malfunctioning.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #32

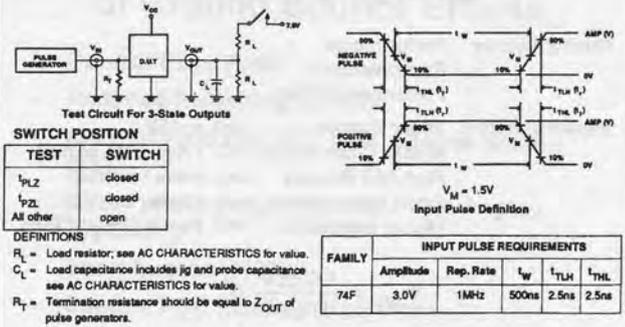
Principles of Chip Characterization for Ground Bounce Effects

- Need 100 ps resolution on controlled stimulus on ~ 10 ns parts
- Ground bounce is highly sensitive to temperature and power supply factors
- Glitches cause ground bounce - close timing, even on don't care transitions
- Tristate bus contention key focus
- Standard parts aren't. Vendors don't specify parts for these effects
- Real-time scope is helpful for controlled S/R

There are several principles which can be drawn from the Characterization Process used in this case study. These principles can be applied whenever you are trying to understand the effects of ground bounce on high-speed digital designs. As shown in Slide #28, there is a very high sensitivity to timing, on the order of 100 ps for 10 ns parts. As was stated earlier, the problem was not seen on the 70 MHz computer where the LEAB signal arrived 400 ps later due to different PCB routing. This confirms the 1% rule — during design of ASICs and PCBs, you should be concerned for edge placements down to 1% of the period, especially where there are glitches and heavy capacitive loads. Ground Bounce is very temperature and power supply sensitive so these environmental factors must be part of the characterization. With the effort in fast computers to fully utilize bus throughput and eliminate all bus deadtime, there is ample opportunity for bus contention and tristate malfunctions. Close attention should be paid to components which are involved in these close timings and potential contentions. The standard components which are used throughout digital designs are characterized with greatly simplified test setups. The 74F543 characterization schematic is shown in Slide #33 and does not take into account the conditions which created the ground bounce. Last, a precision Data Word Generator, like the HP 80000, and a high-bandwidth single-shot scope, like the HP 54720A, are essential to accurately characterize the parts and to understand all the causalities in the failure modes.

Slide #33

Test Circuit and Waveforms



The test circuit is taken from the TTL data book to illustrate that IC vendors don't specify these conditions. Note the simplified test schematic and the 1 MHz repetition rate. You must characterize your own parts for ground bounce conditions.

Slide #34

Outline

1. Debugging subtle problems in memory system HW
2. Characterizing critical components for ground bounce effects
- ⇒ 3. Design approaches to reduce ground bounce in high-speed systems

Lastly, we will discuss design approaches which can be used to eliminate these problems during the design phase.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #35

Ground Bounce is an Emerging Problem

- Driving Forces:** Performance
Performance/\$
Performance/Watt
- Implementation:** Wider Busses $\Rightarrow n = 32$
Higher Clock Rates $\Rightarrow t_{rise} = 7\%$ period
Fast TAT Buses \Rightarrow more overlap
Small Board Space \Rightarrow Higher L
Higher Integration \Rightarrow Fewer Ground Pins

$$V_{GND} = \frac{LC\Delta V}{trise^2} n$$

There are many effects which are driving the emergence of ground bounce as a problem. Wide busses, multiple busses (superscaler), higher clock rates which reduce the rise times, faster Turn-Around-Time (TAT) busses for optimal system performance partitioning, smaller board space, and higher integration are all making the problem worse. Particularly insidious is the higher levels of integration which make the problems due to ground bounce worse while making it much harder to expose the critical conditions because of the increased ASIC complexity and SW dependence.

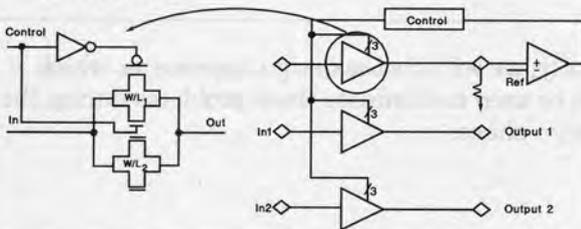
faster. A general rule of thumb is that the rise/fall times will be 7% of the period. This is a conservative number (faster than is often needed) and every effort should be made to keep rise/fall times no faster than required to meet system timing. The most important edge rates to control are those on wide busses, since their effect is multiplied by the bus width. The most important signals to have good fast edges on are the clocks since they control the system operational reliability more than any other signal. Several techniques have been developed to limit the rise/fall times in designs. The normal range of rise and fall times that arise from process, temperature, and voltage changes on chips is four to one. If this range arose on a design, it would change the ground bounce 16 to 1!

HP has patented a technique to control edge rates, shown here. There are several other techniques developed by other companies. In this approach, every output driver has several possible series output resistances, shown as pass gates with W/L1, W/L2, etc. On device power up the output series resistance makes a voltage divider with a precision off-chip resistor, R. A comparator forces the 3-bit control circuitry to force the on-chip resistance to a value set by the divide ratio and the off-chip precision reference voltage. With this effort, all the output drivers get an output resistance and therefore a rise and fall time, controlled within 25% of ideal, depending on the number of output drivers which are selectable. This approach can also improve EMI performance. This effort is worth it because of the square law on rise time.

Slide #36

Controlling Rise Times

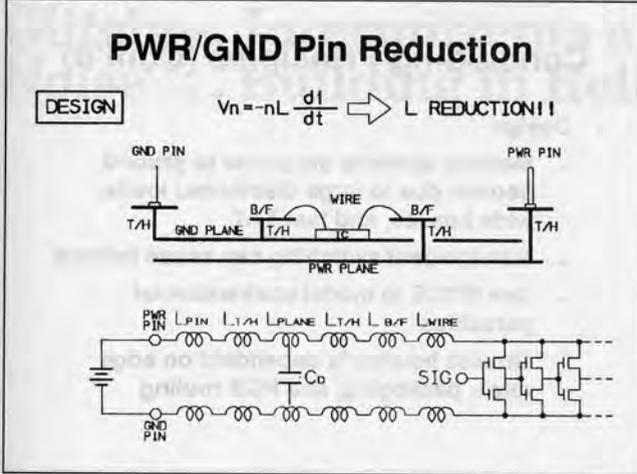
- Programmable output resistance for series termination
- Digital feedback to set drive based on turn-on conditions



Controlling edge rates is the most important design goal to manage ground bounce. There is a tendency toward faster edge rates as silicon processes get

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #37



Kyocera has developed a custom ceramic PGA process which integrated capacitors into the package. This reduces the loop inductance which the return current must go through. This on-package bypassing, C_D , reduces bouncing. The on-package bypassing tends to be very high quality capacitors, as well. Their self-resonance (the frequency where the inductance of the capacitor starts to dominate the impedance) is over 1 GHz. The inductance is controlled by the through-hole process (T/H in illustration). Some chip vendors have even put large on-chip bypass on large VLSI devices. DEC did this on the Alpha chip, published at ISSCC 92, to allow the several amps of clock current to have a return path without going off-chip.

Slide #38

Principles of Design for Reduction of Ground Bounce Effects

- Control Edge Rates
- Reduced Switching Activity
- More Ground and VCC pins; central ground pins
- Better packaging, including on-packaging bypassing
- Short busses (MCMs)
- Custom PCB Routing (3 pF/inch)

These are some techniques which can be used to reduce ground bounce on your high-speed digital designs. As mentioned earlier, the number one goal should be keeping edge rate under control, no faster than required to meet system timing. Although the second alternative is not always available because you are working with pre-defined bus standards, any effort to reduce the amount of simultaneous switching will improve the ground bounce. Adding ground and power supply pins will reduce the total ground and power supply inductance. My rule of thumb is to have half the pins on an ASIC dedicated to power supply and ground and to spread the ground and power supply pins around the chip since adjacent pins do not fully reduce the inductive effects. Mutual Inductance will result in two adjacent pins having only 30 to 50% less inductance than a single pin. Central ground pins on DIP packages are best since they have the shortest lead frame inductance. I have already reviewed the benefits of on-package and on-chip bypassing. MCMs (Multi-Chip-Modules) or other dense packaging alternatives can be used to reduce the total capacitive loading due to the package. MCMs are expensive and have worse coupling and are harder to debug, so their use should be justified by performance or system integration advantages. Last, every effort should be made to reduce the total routing capacitance on critical lines. In the case under study here, we found that the auto-place-and-route SW used to route the SIMM card resulted in twice the capacitive load on the PCB than what a custom PCB layout would accomplish. On 50 ohm FR4 PCBs, the capacitance is approximately 3 pF/inch.

Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

Slide #39

Concluding Principles

- Debug
 - Ground bounce is pattern dependent and therefore, SW dependent.
 - Real-time capture essential.
- Characterization
 - Few vendors understand ground bounce and how to specify parts
 - Failure mechanism is exponential and is highly sensitive to timing, temperature and power supply factors

I have discussed three processes related to the design of high-speed digital designs. These processes can be applied generally when clock rates approach 50 MHz and design techniques are employed to improve performance, such as multiple bus architectures, wide busses, and custom ASICs. The principles are outlined here for summary.

During the Debug Process, it is important to have real-time capture capability in logic analyzers and single-shot scopes because ground bounce is pattern dependent and so is, therefore, SW dependent. Accurate and repeatable time correlation must be retained throughout the debug process to find the subtle SW dependent HW failures.

For all critical parts, I recommend that you complete your own characterization using a controlled stimulus-response system like that shown in this paper for the 74F543 part. Critical parts are those whose timing is critical and where highly capacitive loads are being driven. Vendors of standard components do not characterize the parts for ground bounce effects. The characterization must be completed over all environmental and timing conditions to fully expose the failure mechanisms.

Slide #40

Concluding Principles (Cont'd)

- Design
 - Memory systems are prone to ground bounce due to large distributed loads, wide busses, and fast TAT
 - Non-incident switching can cause failures
 - Use SPICE to model environmental parasitics
 - Ground bounce is dependent on edge rates, packaging, and PCB routing

During the Design Phase, you should pay special attention to those conditions where heavy IC loads are combined with long, distributed transmission lines as well as wide, fast changing busses. All these conditions contribute to ground bounce related failures. Memory systems exhibit all these traits. A contributing factor related to long transmission lines is non-incident switching. This is where the driver is not strong enough to fully drive the long transmission line on the incident wave (one round trip from the driver to the load at 160 ps/inch on FR4). Since the load gate has not received the full voltage swing on the incident wave, the noise margin is less. This makes it susceptible to ground bounce effects.

Because of the parasitic effects in the package and IC which cause the failures we have discussed here, the only simulator I know that can fully reproduce these effects is SPICE. Only SPICE has the accurate second-order models on transistors, packages, and temperature which will show the real expected behavior. The long simulation times from SPICE are unfortunate but required to get a correct view.

Nothing replaces good modeling of edge rates, packaging, and handcrafted PCB routing to reduce and control ground bounce in high-speed designs. If these are done carefully, the turn-on phase will have fewer subtle problems, like the one we found during the Prototype Build phase.

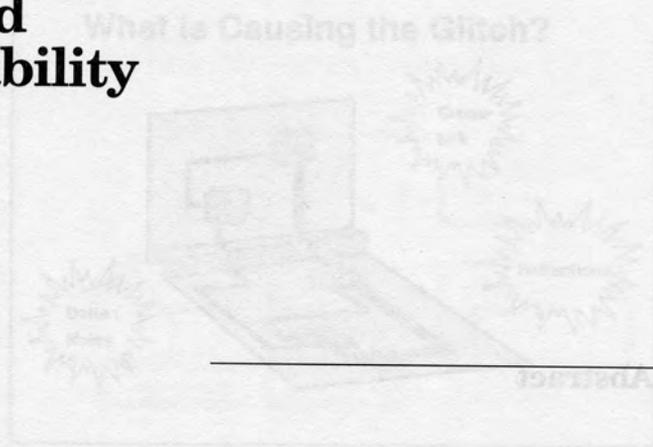
For this design, we modified the PCB routing to reduce the load capacitance and changed the timing to the 74F543. We also characterized several vendor's parts and chose the ones which exhibited the least sensitivity to ground bounce malfunction.

Glitches, Intermittents and Noise . . . Building in Reliability

Greg Doyle
Bernard J. Sheehan

Integrity Engineering, Inc.
1306 W. County Rd. F, Suite 100
St. Paul, MN 55112
Tel: (612) 636-6913
Fax: (612) 631-2241

1993 High Speed Digital Systems Design & Test Teleconference



What are the major causes of glitches and noise? How do they affect system reliability? What is the right time to build in reliability? What can be done to prevent glitches and noise? This paper discusses these issues and provides a methodology for building in reliability through design and testing. Case studies are used for illustration and examples.

Authors
Greg Doyle
Bernard Sheehan

Current Activities:
Greg Doyle is vice president of Integrity Engineering. He is currently active in the development of system level design and screening tools to ensure signal integrity and reliability. Greg has also given numerous workshops on high speed design.

Author Background:
Greg graduated from Michigan Technological University. He worked five years with Control Data in maintenance computer development before founding Integrity Engineering.



Abstract

What are the major causes of intermittent failure in digital designs? What is the origin of these causes? What can be done to minimize these failures? This paper discusses these issues and more including tips on building in reliability through noise budgeting. Case studies are used for illustration and examples.

Authors

Greg Doyle

Current Activities:

Greg Doyle is vice president of Integrity Engineering. He is currently active in the development of system level design and screening tools to ensure signal integrity and reliability. Greg has also given numerous workshops on high speed design.

Author Background:

Greg graduated from Michigan Technological University. He worked five years with Control Data in mainframe computer development before founding Integrity Engineering.

Bernard Sheehan

Current Activities:

Bernard Sheehan is the chief technical consultant at Integrity and has contributed proprietary methods for Boundary Element analysis, crosstalk simulation, and transmission line modeling including frequency dependent losses.

Author Background:

Bernard graduated from Carleton College, MN, and the University of MN. He also worked five years with Control Data in mainframe computer development before founding Integrity Engineering.



Slide #1

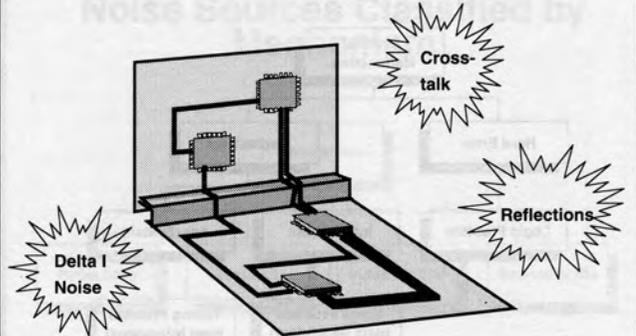
Glitches, Intermittents and Noise ...
Building in Reliability



Integrity Engineering, Inc.

Slide #4

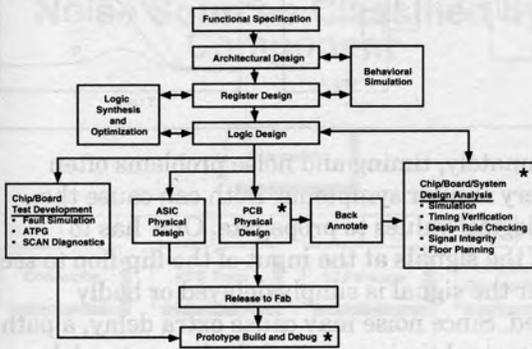
What is Causing the Glitch?



If your system has intermittents, how do you go about discovering the source of the problem? Where is the glitch coming from? How can you design to avoid glitches? Obviously, there is no simple, universal way to answer these questions. But knowing the characteristics of different noise sources and how to isolate and quantify them can assist you in "Zooming in" on the problem and eliminating its cause.

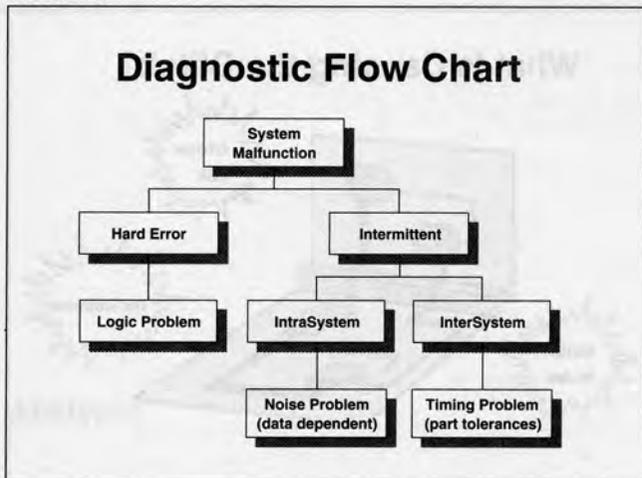
Slide #3

Simplified HS Digital Design Process



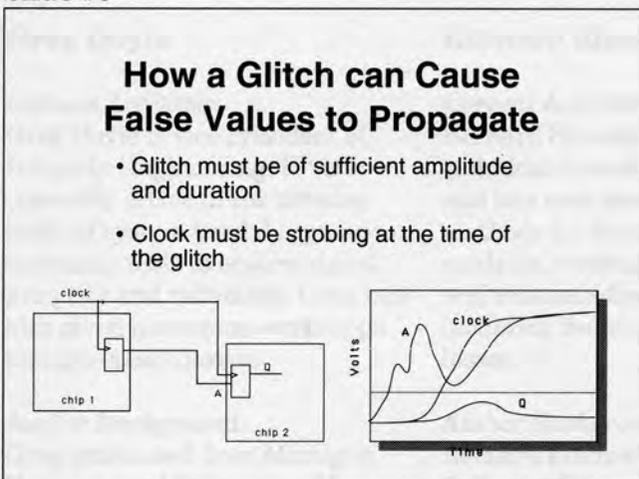
The areas of the Simplified Design Process this talk deals with are highlighted as shown here with asterisks.

Slide #5



Troubleshooting digital system problems is very much like detective work. Each case is unique and draws on the reasoning powers of the engineer solving the problem. This slide suggests a high level flow chart for isolating the cause of digital errors. A logical error is a hard error that will show up consistently in all systems. Timing errors may be present in some units and absent in others, depending on the variations of parts in each unit. Noise problems are intermittent within the same unit, depending on data patterns and activity in other areas of the circuit.

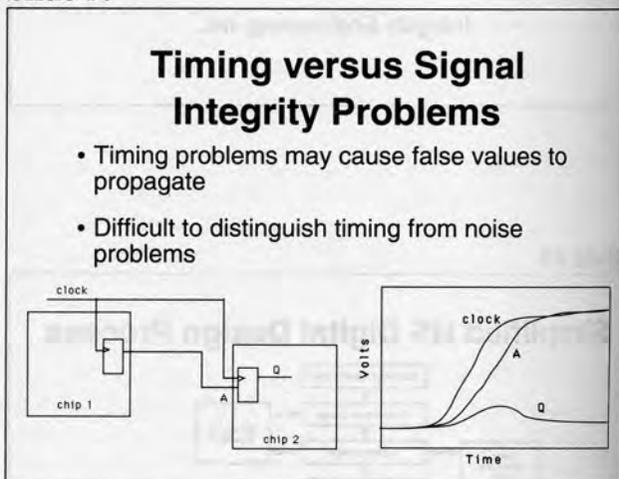
Slide #6



Normally, each stage of logic filters out the noise at its input, producing a clean output pulse for the next stage. This makes for very robust performance.

However, under certain conditions, the wrong value can be clocked into a flip flop. This will be the case, for example, if a large glitch occurs at about the same time as the clock strobe. This glitch must not only occur at the right time, but also be of sufficient amplitude and duration to cause the flip flop to interpret the input as a low, rather than of a high.

Slide #7



Unfortunately, timing and noise problems often have very similar symptoms. Both can cause the wrong logical values to propagate. One has to inspect the signals at the input of the flip-flop to see whether the signal is simply delayed or badly distorted. Since noise may cause extra delay, a path with marginal timing might fail when extra delay from signal noise is added. In this situation, the failure is due to a mixture of timing and noise problems.

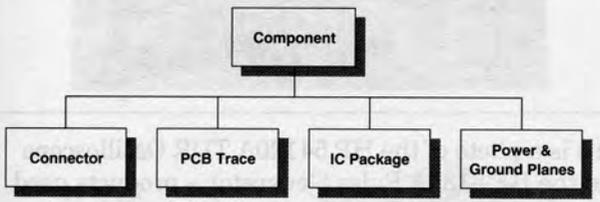
Slide #8



I. Principal Noise Sources
II. Test Board Measurements
III. Noise Budgeting
IV. Post Layout Screening

Slide #9

Noise Sources Classified by Component



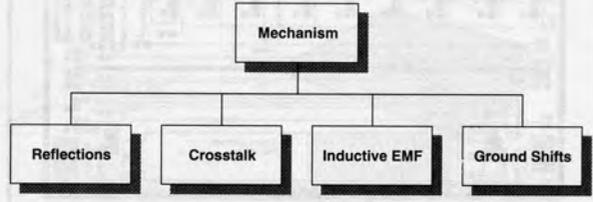
```

    graph TD
      Component[Component] --> Connector[Connector]
      Component --> PCBTrace[PCB Trace]
      Component --> ICPackage[IC Package]
      Component --> PowerGround[Power & Ground Planes]
    
```

One useful approach to track down the source of glitches and intermittents is to consider the noise generating abilities of each component in your system. The same component may be responsible for several forms of electrical noise. For example, an IC package may cause delta I noise, crosstalk, and reflections. Or PCB traces may cause crosstalk, IR drop, and stub reflections. Sometimes a glitch may be produced by a single component (like a connector with inadequate thru-grounds); other times it may be the sum of noise contributions from several components.

Slide #10

Noise Sources Classified by Mechanism



```

    graph TD
      Mechanism[Mechanism] --> Reflections[Reflections]
      Mechanism --> Crosstalk[Crosstalk]
      Mechanism --> InductiveEMF[Inductive EMF]
      Mechanism --> GroundShifts[Ground Shifts]
    
```

It is also useful to recognize the different mechanisms that generate noise. As already noted, the same mechanism may be operative in several components. Crosstalk, for example, may occur along the entire signal path—it may occur in the driver package, along the printed circuit traces, in any connectors along the path, and in the receiver package.

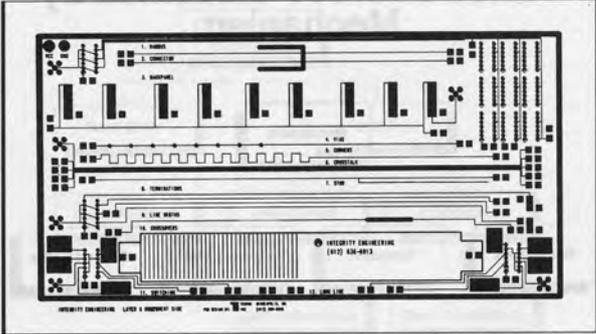
Slide #11



I. Principal Noise Sources
II. Test Board Measurements
III. Noise Budgeting
IV. Post Layout Screening

Slide #12

Signal Integrity Test Board



To illustrate various glitch producing mechanisms, a signal integrity test board with twelve "experiments" was designed. Measurements from this test board will be used to illustrate the noise effects being discussed in this paper.

The board is a 12" x 6" FR4 printed circuit board. The layout consists of a surface signal layer, a ground plane, a power plane, and a solder layer on the back. The default line width, which was designed using IET's CALIF software, was 11 mils; this resulted in an impedance very close to 50 ohms. Active components for driving and loading traces consist of 74F04 or 74HC04 hex inverters.

Slide #13

Measurement Setup

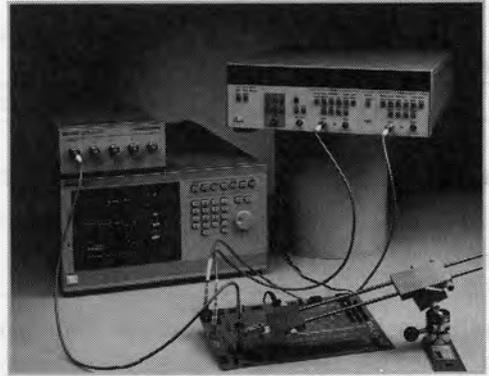
- HP 54120A Oscilloscope
- HP 54124A Test Set
- HP 8130A Pulse Generator
- HP 54720A (Single Shot Glitches, Intermittents)
- ICM TDR Probe Assembly

An HP 54120A oscilloscope was used for all of the measurements. Its TDR (Time Domain Reflectometry)

capability makes it ideal for investigating reflections and impedance discontinuity effects. For driving lines, an HP 8130A Pulse Generator was used. Its programmable period, pulse width, amplitude, and rise/fall time are valuable for looking at how crosstalk, for example, depends upon the pulse characteristics. HP 54006A Resistive Divider Probes were used to give wide bandwidth and low capacitive loading.

Slide #14

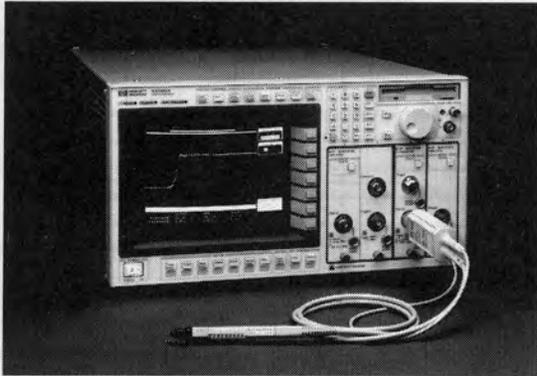
Measurement Setup (Cont.)



This is a photo of the HP 54120A TDR Oscilloscope and the HP 8130A Pulse Generator – products used in this talk.

Slide #15

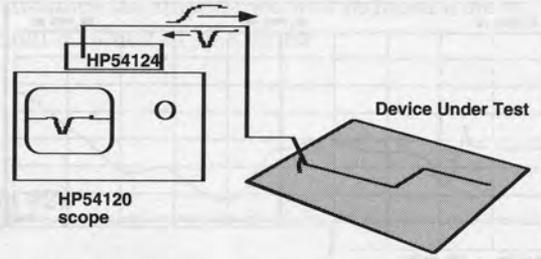
Measurement Setup (Cont'd)



This is a photo of the HP 54720A. Because of its fast digitizing rate (2 Gsa/s w/4 channels, 4 Gsa/s w/2 channels), it is an ideal tool for troubleshooting single shot glitches and intermittents. This particular aspect of troubleshooting will not be covered in this paper.

Slide #17

Time Domain Reflectometry



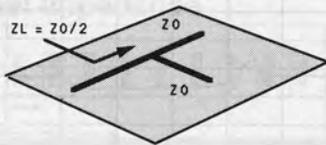
- TDR sends fast risetime pulse down line.
- It displays any reflections coming back.

Time Domain Reflectometry (TDR) is a powerful technique for characterizing the impedance control and reflection generating properties of components and interconnects. A TDR test injects a step voltage with a very short risetime down a cable into the device under test. It then observes reflections returning from the circuit. The temporal position of reflections can be related directly to the physical position of the glitch source.

Slide #16

Reflections

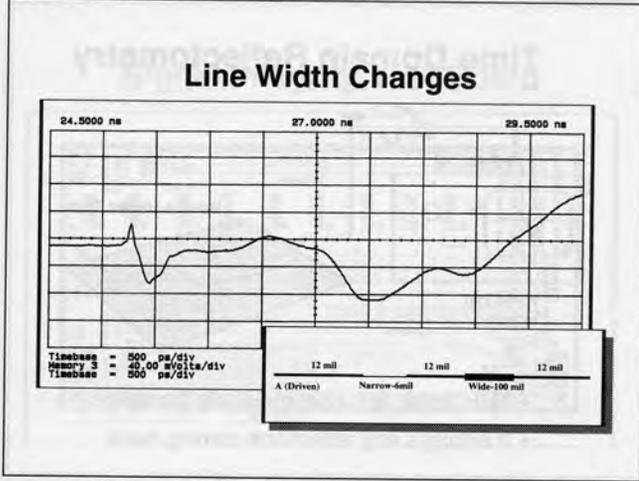
$$\rho = \frac{Z_L - Z_0}{Z_L + Z_0} = -\frac{1}{3}$$



- Routing stubs can produce significant reflections

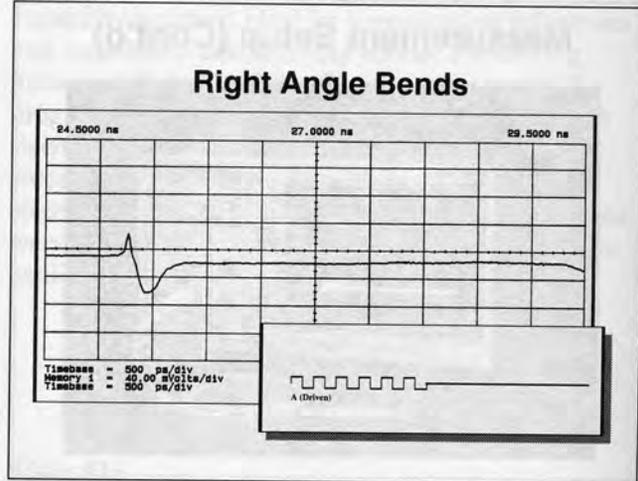
One key source of glitches are reflections. Routing stubs, for example, can cause negative reflections or glitches equal to 33% of the input signal swing. This noise source alone, can take up most of a logic family's noise budget. This is why high speed designs generally use daisy-chain routing or point-to-point routing, which avoids stubs.

Slide #18



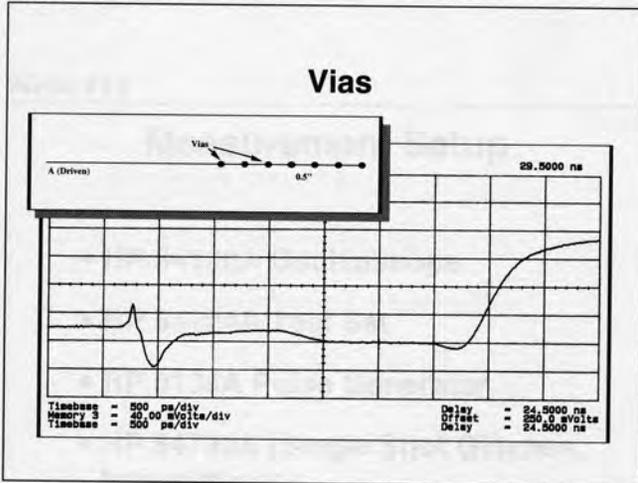
This is TDR measurement of a printed circuit board trace with a width of 12 mils except for a region where it narrows down to 6 mils and another region where it widens out to 100 mils. We see the positive reflection from the higher impedance section of narrower width (65.8 ohm) and the negative reflection from the low impedance section of a wider line (19.8 ohm). The narrow region we call "inductive" and the wide section "capacitive."

Slide #20



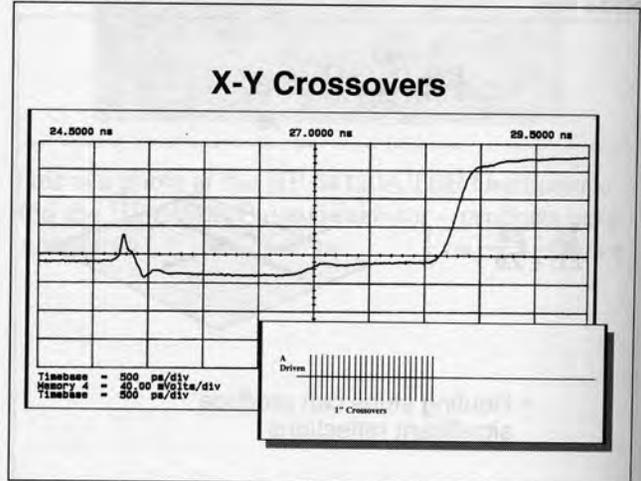
This slide shows the TDR plot of a line the first half of which has corners every 1/4" in its first half, and is simply straight in its second half. The corners have only a slight effect on impedance, dropping it from 50.6 ohms to 48.2 ohms.

Slide #19



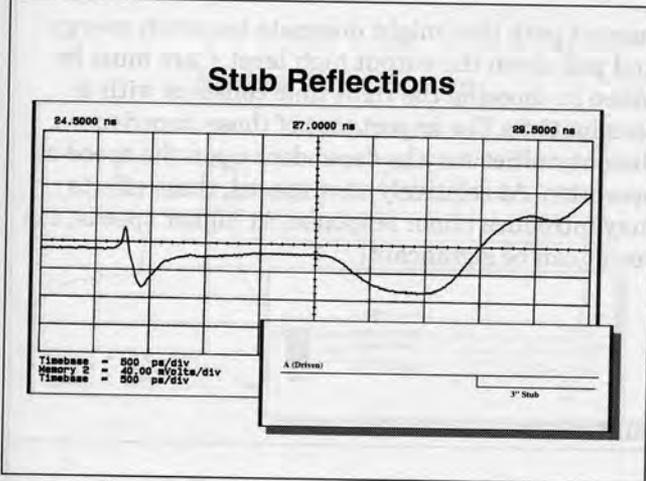
This slide shows the reflections from a line which has through-the-board vias every half inch in its second half. The presence of the vias dropped the line impedance from 50.6 ohms to 43.1 ohms.

Slide #21



This TDR measurement is of the first half of a microstrip line, which passes under a region of perpendicular crossover lines. The crossover traces are 1" long and are at a pitch of 100 mils apart. The presence of the crossovers reduced the impedance from 53.9 ohms to 43.8 ohms.

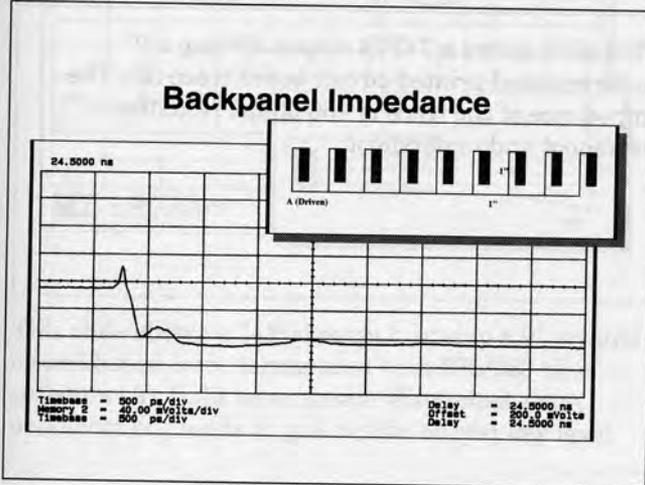
Slide #22



This slide shows the TDR measurement of a 9" long trace with a 3" stub branching from it. The stub causes the impedance to drop from 51.9 ohms to 27.2 ohms. This impedance drop, as already discussed in an earlier slide, is the result of the signal seeing two impedances in parallel.

The detrimental effect of stubs on signal integrity is one of the reasons daisy-chain routing is often used in high-speed designs.

Slide #23

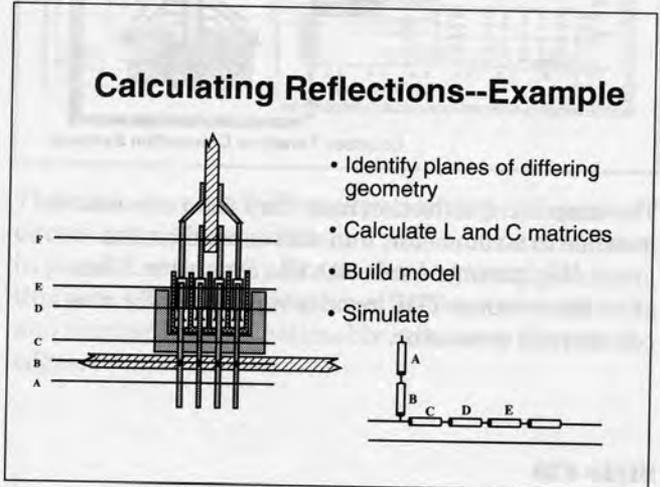


Electrically, a backpanel trace can be thought of as a transmission line with periodic stubs—the stubs representing the loading effect of the daughter cards.

This slide shows the TDR plot of a trace with periodic loads (our idealized backpanel). Just as a single stub

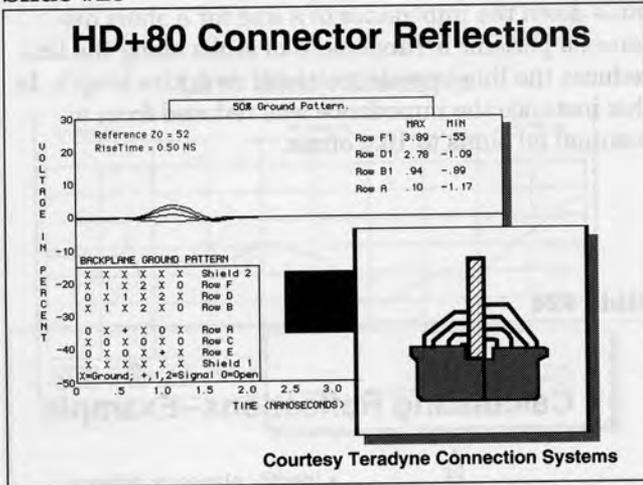
pulls down the impedance of a line for a short distance, a periodic arrangement of stubs along the line reduces the line impedance along its entire length. In this instance the impedance was reduced from a nominal 50 ohms to 19.4 ohms.

Slide #24



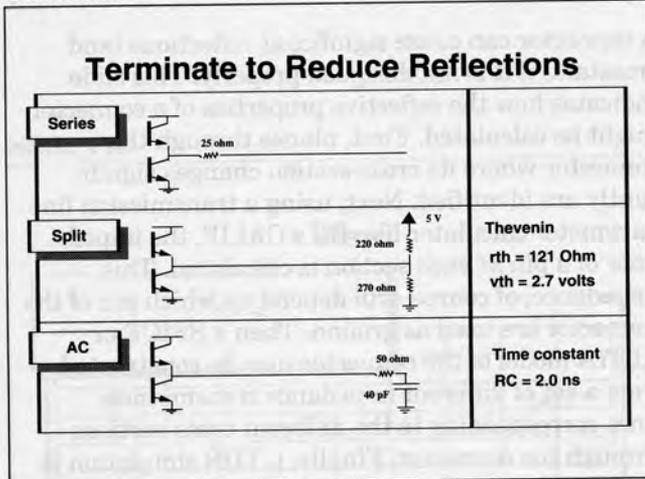
A connector can cause significant reflections (and crosstalk) if it is not designed properly. This slide indicates how the reflective properties of a connector might be calculated. First, planes through the connector where its cross-section changes significantly are identified. Next, using a transmission line parameter calculator like IET's CALIF, the impedance of a pin at each section is calculated. This impedance, of course, will depend on which pin of the connector are used as ground. Then a SPICE or ALTrA model of the connector may be constructed from a set of different impedance transmission lines corresponding to the different cross-sections through the connector. Finally, a TDR simulation is performed.

Slide #25



The amount of reflection may vary from one row to another in a connector, and will depend on the grounding pattern used with the connector. This slide shows some TDR results versus row for one commercial connector.

Slide #26

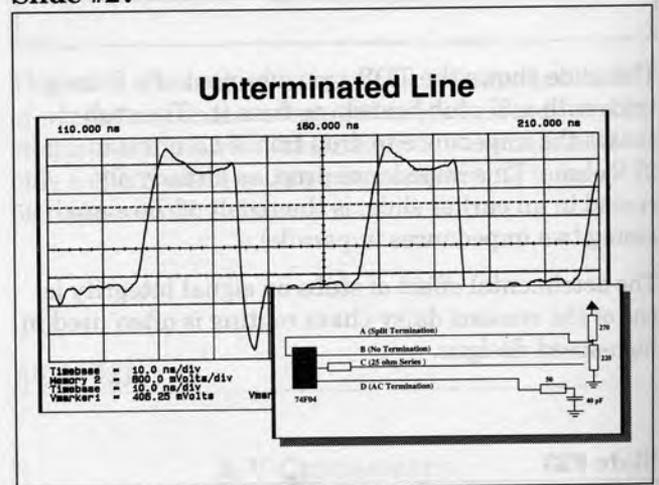


One way to reduce reflections and ringing on a net is by means of some termination scheme. Termination usually involves adding a resistor at one end of the line; there are various ways to do this. This slide illustrates some techniques used with TTL devices.

Series termination absorbs reflected energy but does not draw dc current. Split termination biases the line with a certain thevenin voltage and resistance. AC termination absorbs transients but avoids a dc

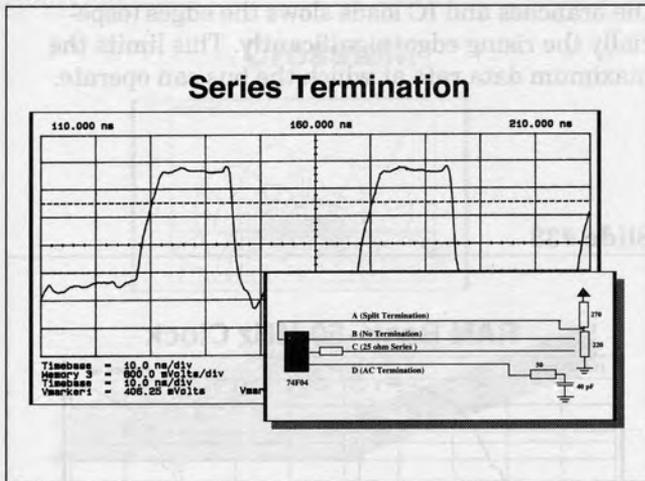
current path that might dissipate too much energy and pull down the output high level. Care must be taken in choosing the right time constant with ac terminations. The importance of these impedance discontinuities may be dependent upon the speed of operation. At relatively slow speeds, these effects may introduce minor response. At higher speeds, the result can be significant.

Slide #27



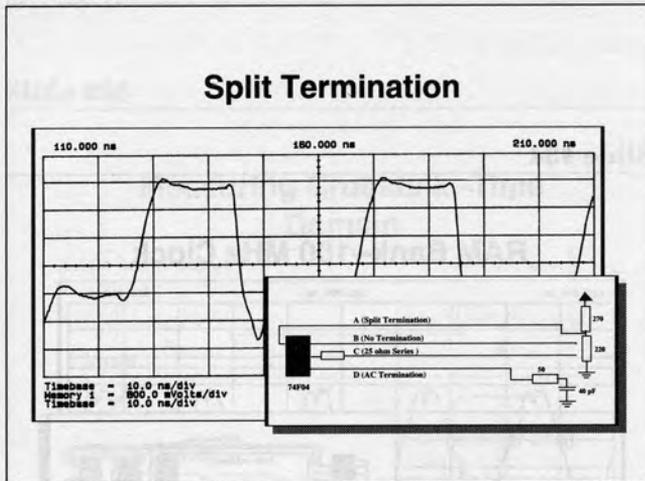
This slide shows a 74F04 output driving a 9" unterminated printed circuit board trace (B). The impedance of the trace is ~50 ohms. Note the overshoot and undershoot.

Slide #28



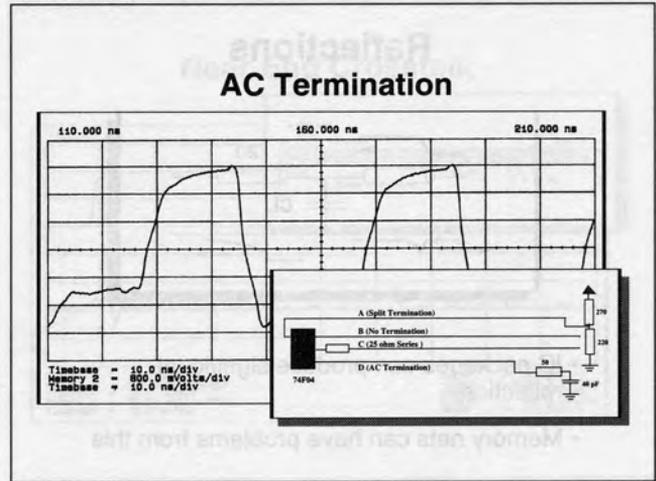
Adding a 25 ohm series resistor at the output (C) reduces both the overshoot and undershoot.

Slide #29



This slide shows a 74F04 output driving a 9" printed circuit board trace terminated with 270/220 ohm resistors (A). Split termination eliminates some overshoot but tends to pull up the output low level.

Slide #30

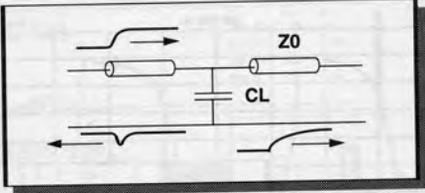


This slide shows a 74F04 output driving a 9" printed circuit board trace terminated with a 50 ohm resistor in parallel with a 40 pF capacitor to ground (D). In this case the AC termination eliminates overshoot and ringing but also noticeably influences the signal edges.

Glitches, Intermittents and Noise ... Building in Reliability

Slide #31

Reflections



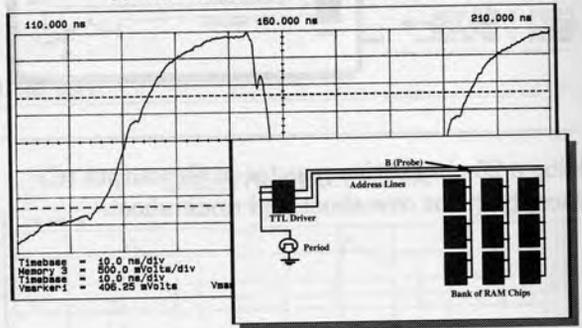
- IC packages can produce significant reflections
- Memory nets can have problems from this

Electrically, an IC package connected to a printed circuit board trace acts like a stub or capacitive load and generates a negative-going reflection. Package reflections become especially significant with large pin grid arrays, which can represent up to 12 pF of capacitance. This is the equivalent capacitance of 4" of 50 ohm transmission line. Designs with nets routing to several large IC packages are prime candidates for glitch problems.

the branches and IC loads slows the edges (especially the rising edge) significantly. This limits the maximum data rate at which the bus can operate.

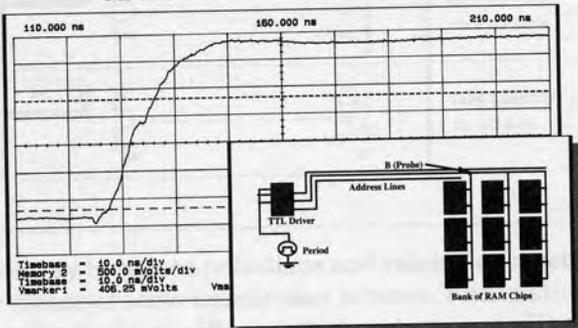
Slide #33

RAM Bank--50 MHz Clock



Slide #32

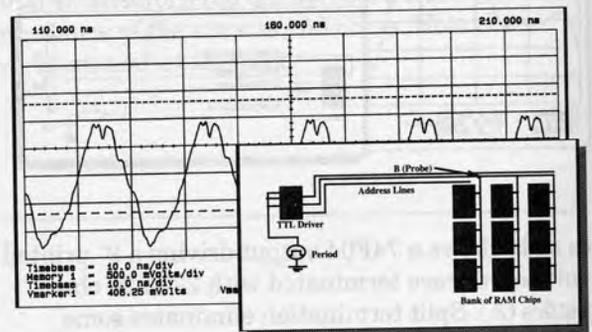
RAM Bank--10 MHz Clock



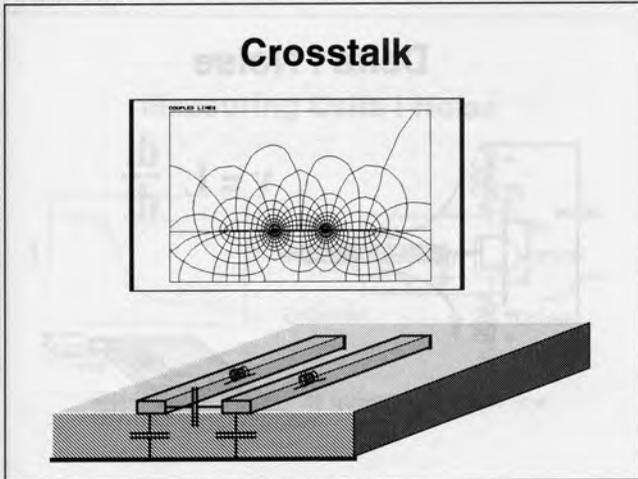
This example and the next two slides show how the signal looks when it is driving several banks of chips—a configuration that occurs in memory circuits. The signal is measured at point B on the net (at the foot of the first branch). The capacitance of all

Slide #34

RAM Bank--100 MHz Clock

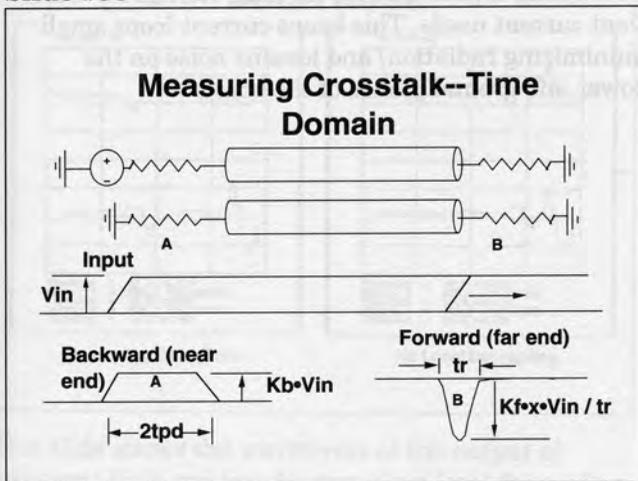


Slide #35



Crosstalk arises from the mutual capacitance and inductance between neighboring conductors. It can be a significant source of noise in densely routed printed circuit boards. Crosstalk may be particularly acute when busses consist of long sets of parallel lines. It can also be significant in connectors with inadequate thru-grounds.

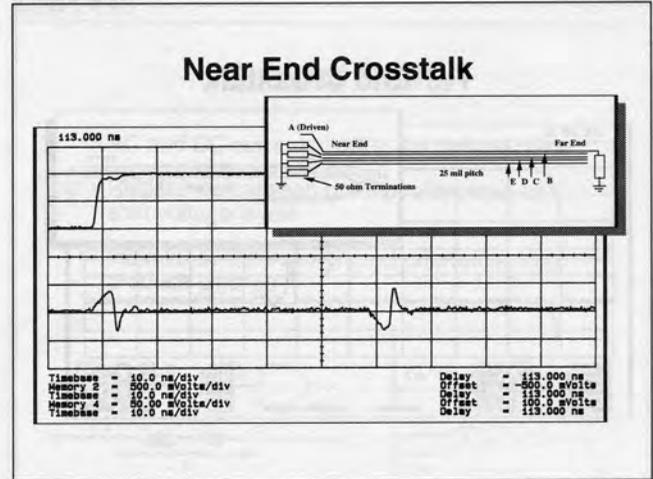
Slide #36



To measure crosstalk in the time domain, we drive one line with an HP 8130A Pulse Generator and observe the waveforms coupled into the near and far ends of neighboring lines with an HP 54120A Oscilloscope.

In addition to measurements, crosstalk can be accurately predicted with CAE tools like IEF's XTALK software.

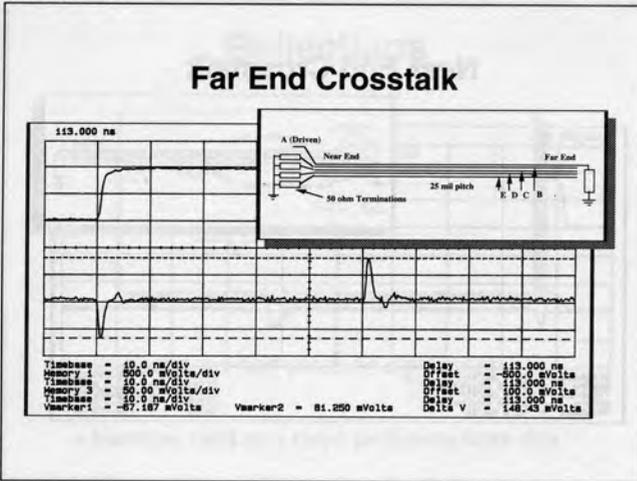
Slide #37



The test board experiment has five parallel 11 mil lines on a 25 mil pitch. The line lengths were about 10 inches. The upper line was driven at one end (point A) and terminated in 50 ohms at the other end. The other four lines are terminated in 50 ohms at their near ends and open at their far ends.

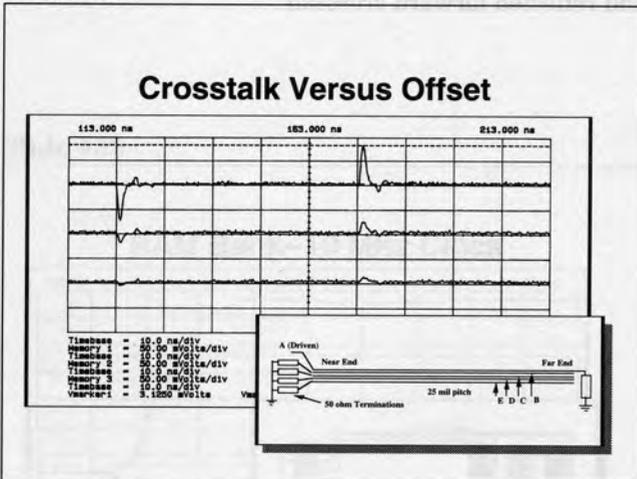
This scope plot shows the waveform at the near end of line B. The waveform includes backwards crosstalk and reflected forward crosstalk.

Slide #38



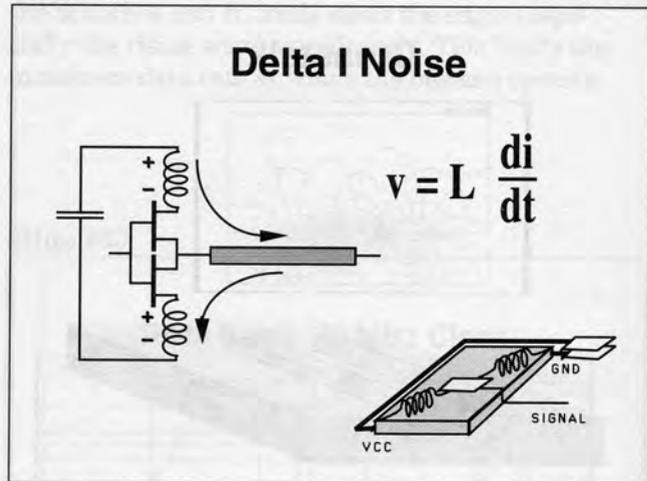
This scope plot shows the crosstalk at the far end of line B. Only forward crosstalk is present; the termination at the near end prevents backward crosstalk from being reflected.

Slide #39



This plot shows that the magnitude of forward crosstalk depends on the distance from the driven line. Starting from the upper waveform, the plot shows far end crosstalk into lines B, C, and D, respectively.

Slide #40



Delta I or simultaneous switching noise can be a significant noise source in digital systems. It arises from the parasitic inductance of the power and ground leads in an IC package. When multiple outputs switch, there is an abrupt change in the current passing through these power and ground leads. The inductive emf due to this sudden change in current causes a voltage spike on the chip's power and ground busses relative to the board power and ground.

Decoupling capacitors provide local storage for transient current needs. This keeps current loops small (minimizing radiation) and lessens noise on the power and ground planes of the board.

Slide #41

Measuring Delta I Noise

The diagram shows a circuit with a buffer and a current source I . A graph shows voltage vs. time with 'quiescent' and 'driven' states.

- Configure an output as quiescent high or low
- Drive all other outputs

This slide illustrates the technique of viewing the $L \cdot di/dt$ voltage spike on the IC's power rails by looking at the output of a 'quiescent' buffer.

Slide #42

Delta I Noise

The slide shows two oscilloscope waveforms comparing 'Decoupling Capacitors' and 'No Local Decoupling'. Both waveforms show voltage vs. time with a scale of 110.000 ns. The 'Decoupling Capacitors' waveform shows significantly reduced noise compared to the 'No Local Decoupling' waveform.

Decoupling Capacitors: Timebase = 10.0 ns/div, Memory 1 = 800.0 mVolts/div, Timebase = 10.0 ns/div, Memory 2 = 800.0 mVolts/div, Timebase = 10.0 ns/div

No Local Decoupling: Timebase = 10.0 ns/div, Memory 3 = 800.0 mVolts/div, Timebase = 10.0 ns/div, Memory 4 = 800.0 mVolts/div, Timebase = 10.0 ns/div

This slide shows the waveforms at the output of quiescent high and low drivers when local decoupling capacitors are present and absent. The device was a 74F04 part; four outputs were driven simultaneously. Of the remaining two outputs, one was held as a static high and the other held as a static low.

The decoupling capacitors assist in reducing switching noise, but clearly there is still much unbypassed inductance in the IC package.

Slide #43

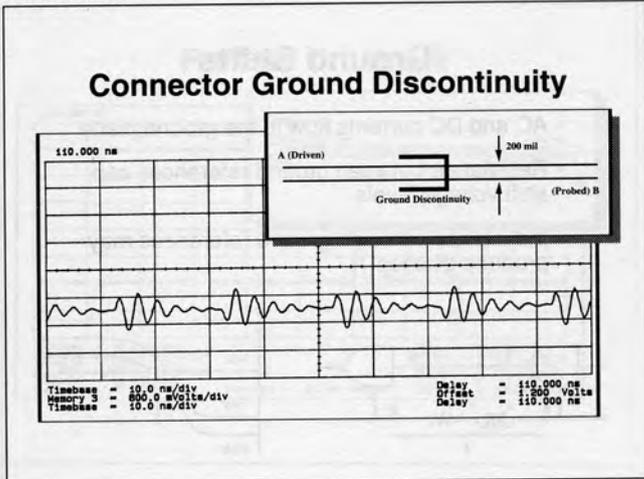
Ground Shifts

- AC and DC currents flow in the ground plane
- Resistance between ground references can shift voltage levels
- Inductance between ground references may produce glitches

The diagram shows a circuit with two inverters A and B, and a graph showing voltage levels V_{th} , $V_{aa'}$, and $V_{bb'}$.

The power and ground planes (or busses) in a printed circuit board carry both ac and dc return currents. DC currents flowing through the resistance of the ground (power) path shifts the voltage of the signal. Such shifts subtract directly from high or low noise margin. Similarly, AC currents flowing through the parasitic inductance of the ground (power) path will superimpose inductive voltage spikes on a signal. Breaks or interruptions in the ground (power) plane are likely sources of glitches from this mechanism.

Slide #44



This test board experiment shows how a signal crossing a break or cut in the ground plane sets up a potential difference across the grounds on the two sides of the cut. Moreover, this ground bounce (as it is called) couples into other signals across the same ground discontinuity.

The upper PCB line is driven with a 6V, 1 ns risetime pulse at point A. The scope plot shows the noise coupled into the lower PCB line. The two lines are spaced quite far apart (200 mils), so direct crosstalk between the lines should be negligible. The ringing seen at B is due to the disturbance at the ground discontinuity.

Slide #45



- I. Principal Noise Sources
- II. Test Board Measurements
- III. Noise Budgeting
- IV. Post Layout Screening

Slide #46

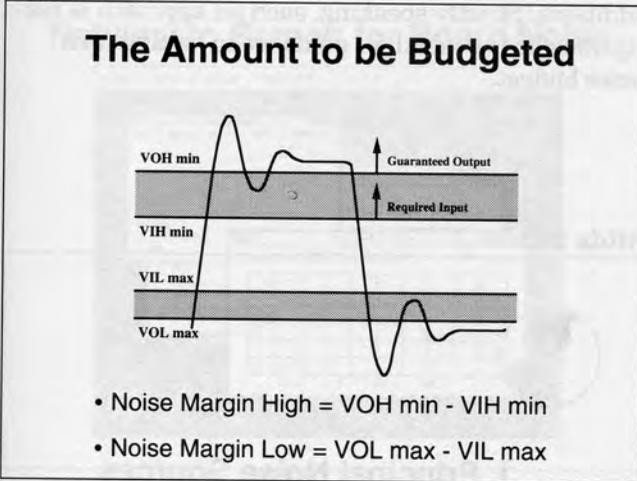
How to Ensure Reliability by Noise Budgeting

- A budget is the allocation of limited resources
- A circuit has limited noise margin
- A noise budget allocates the noise margin between potential noise sources



A noise budget is a disciplined way of designing for noise. When you draw up a noise budget, you allocate the available noise margin of your circuits among the various potential noise sources. Then you select/design each component to meet its noise allowance.

Slide #47



Whether electrical noise causes problems depends on two things: (1) circuit susceptibility to noise and (2) presence and intensity of noise sources in the system. It is common to focus attention on the second item, but the first is equally important. A digital circuit's susceptibility to noise is characterized by the circuit's noise margin. Strictly speaking, there are two noise margins, one for high and one for low logic levels.

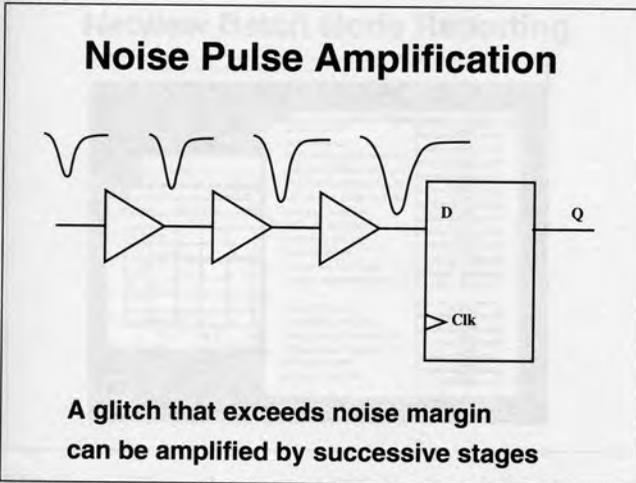
Slide #49

Noise Margin versus Logic Family

	Voltage Swing	Vth	Noise Margin	REL. Noise Margin
TTL	3.5	1.5	0.4	11%
CMOS	5.0	2.5	0.6	12%
ECL	0.8	-1.3	0.14	17%

Since the magnitude of noise produced is usually proportional to a logic family's signal amplitude (a 2V signal, for example, will produce twice as much crosstalk as a 1V signal), it is helpful to think in terms of relative noise margin. Relative noise margin equals the noise margin divided by the amplitude of the signal swing. For example, the noise margin of TTL is about 0.4V, compared to 0.14V for ECL. However, TTL signals are also much larger than ECL signals (3.5V swings compared to 0.8V swings). The relative noise margins of the two families are 11% and 17%, respectively.

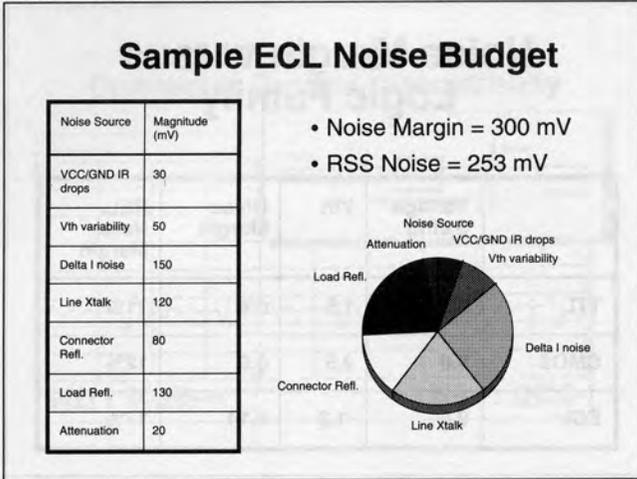
Slide #48



If a glitch due to reflections, crosstalk, delta I noise or the combination of these exceeds the noise margin, this pulse can be amplified from stage to stage, eventually reaching an amplitude sufficient to be clocked into a flip-flop or latch.

Glitches, Intermittents and Noise ... Building in Reliability

Slide #50



This slide shows a sample noise budget for an ECL design. If a design were being carried out under such a noise budget, each component would be selected with an eye to its noise allowance. For example, in choosing a connector, one would want to take care to choose a model that keeps reflections less than 80 mV or 10% on the input signal. This may dictate the grounding pattern needed for the connector.

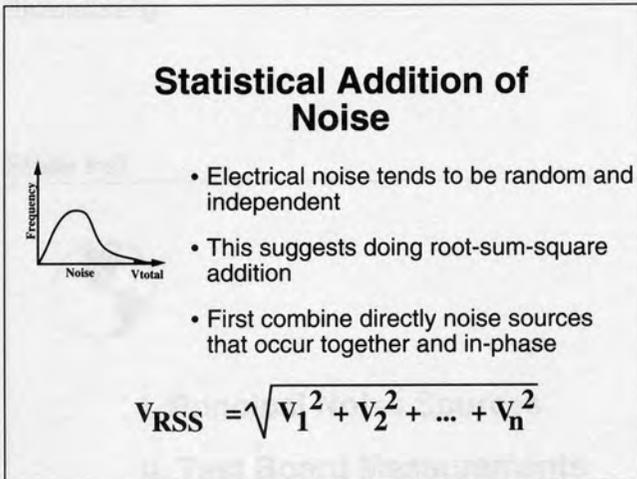
additions. Strictly speaking, such an approach is not rigorous but is useful for guidance in setting up a noise budget.

Slide #52

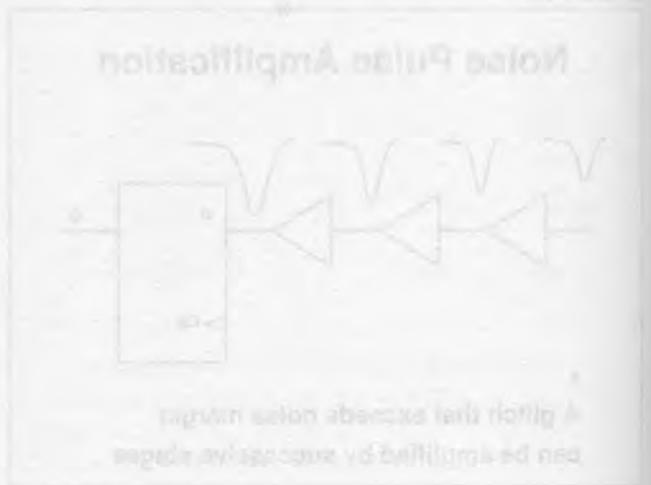


- I. Principal Noise Sources
- II. Test Board Measurements
- III. Noise Budgeting
- IV. Post Layout Screening

Slide #51



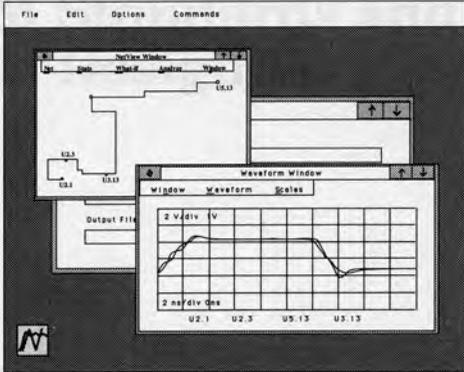
The total noise from many independent sources is not the arithmetic sum of the magnitudes of the individual sources—such a calculation is too pessimistic. The random nature of the noise sources is better taken into account by doing root-sum-square



Glitches, Intermittents and Noise . . . Building in Reliability

Slide #53

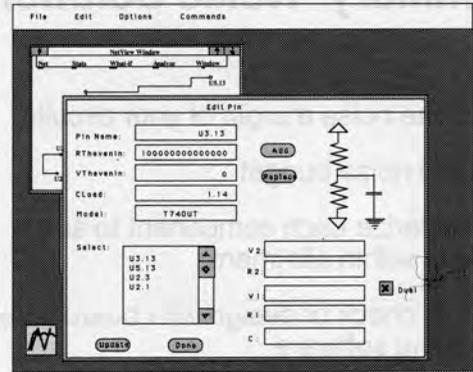
Netview to Screen for Board Noise



To detect potential glitches before committing to a board build, a tool like IEI's Netview in batch mode will simulate every net on a PCB and flag any potential problem nets. These can then be examined individually in Netview's interactive mode, where different possible fixes can be tried until simulation shows good signal integrity.

Slide #55

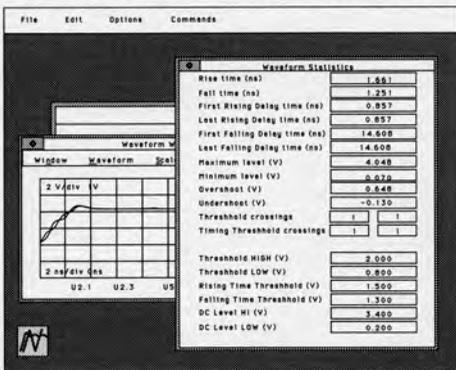
Netview What-if Capability



Netview also has interactive what-if capability, allowing you to try different termination schemes, for example, to fix nets with glitches.

Slide #54

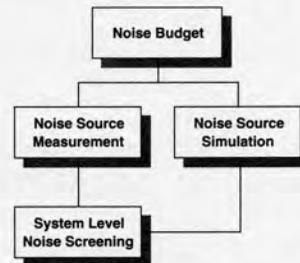
Netview Batch Mode Reporting



Netview will read a routed database and do a transmission line simulation of each net on the board. Waveforms are automatically scanned for excessive overshoot, undershoot, delay, or ringing and nets with potential problems are flagged.

Slide #56

Conclusion: Avoiding Glitches, Intermittents, and Noise



With so many noise generators lurking in your system—delta I noise, crosstalk, reflections, and ground shifts (plus many secondary sources not mentioned)—how can you ensure that your design will be reliable once it is finished? Intermittent noise problems are notoriously difficult to isolate and fix. Clearly a preventative approach is in order. Our opinion is to measure and characterize the glitch generative abilities of each component, and then draw up a noise budget to keep total noise amplitudes within reasonable limits.

Slide #57

Summary--Noise Solutions

- Know the noise margin of your circuits
- Prepare noise budget
- Characterize each component to see if noise is within allotment
- Do final check of design with board level screening software

The best approach to glitches and intermittents is prevention. Know and budget your noise sources.

The diagram illustrates a process for managing noise in system design. It starts with 'System Design', which leads to 'Noise Budget' and 'System Test'. Below this, there is an illustration of a person running with a briefcase, symbolizing the urgency of addressing noise issues. The text below the diagram discusses the importance of knowing noise margins and preparing a noise budget.

With so many generators lurking in your system, it's vital to establish a noise budget and identify the many secondary sources not mentioned—how can you ensure that your design will be reliable once it is finished? Intermittent noise problems are notoriously difficult to isolate and fix. Clearly a preventive approach is in order. Our proactive approach can help you identify the right components to use and ensure that your design is built to help you avoid noise problems that could cause your product to fail.

This screenshot shows a software interface for analyzing signal integrity. It features a graph with a grid, a legend, and various data points. The interface is designed to help users identify and analyze potential glitches and intermittents in their designs.

This screenshot shows a software interface titled 'Netview Batch Mode Reporting'. It displays a table of data and a graph. The interface is designed to help users analyze and report on the results of their simulations.

Understanding Evolving ATM Standards and ATM Design Verification

Rick Tinsley

TranSwitch Corporation
8 Progress Drive
Shelton, CT 06484
Phone: (203) 929-8810
Fax: (203) 926-9453

Dan Upp

TranSwitch Corporation
8 Progress Drive
Shelton, CT 06484
Phone: (203) 929-8810
Fax: (203) 926-9453

**1993
High Speed Digital
Systems Design & Test
Symposium**

Abstract

In the last two years, ATM has become the fastest evolving communication standard. Understanding where the standard is in its development can become a significant barrier to designing ATM-compatible equipment. This paper describes the latest advancements in the development of ATM standards and discusses tools that can provide ATM design verification.

Authors

Rick Tinsley

Current Activities:

Rick Tinsley is Director of Marketing responsible for TranSwitch's ATM product line.

Author Background:

Rick received his BSEE from Rensselaer Polytechnic Institute and MBA from the University of Dallas. Formerly with Texas Instruments, he held various marketing, sales, and business development positions. Prior to TI, he was an Analog Designer at General Electric.

Dan Upp

Current Activities:

Dan Upp is Vice President of Technology Development and a founder of TranSwitch Corporation.

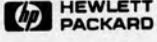
Author Background:

Dan received his BSEE and MSEE from Ohio State University and worked on satellite communications systems and antenna array systems at the OSU ElectroScience Labs. Subsequently, he was employed by North Electric Co. in the hardware design of the DSS-1 (later referred to as the ITT 1210) switching system.

Dan spent ten years at the ITT Advanced Technological Center where he was Director of Exploratory Systems, responsible for VLSI, hardware, and software development for packet switching, telephone switching, LAN and PABX products.

Slide #1

Understanding Evolving ATM Standards and ATM Design Verification



Slide #3

Current Technology

- **Circuit switched**
 - Fixed bandwidth
 - Low latency
- **Packet switched**
 - Simpler multiplexing
 - Delay variations and high latency

Current network technology generally falls into two classifications: circuit-switched technology in the traditional telephony arena and packet switching in the LAN/data networking community.

Circuit switching is characterized by fixed graduations of bandwidth, such as a 64kbit/s voice channel with prescribed subrate and superrate multiplexings. Circuit switching requires a continuous hold of the physical path for each connection, regardless of actual bandwidth used. This provides very low latency, and thus excellent performance for voice and other isochronous information.

Packet switching is more efficient with respect to bandwidth utilization since access is limited to the time required to transmit a given packet or frame. Switching and multiplexing are also simpler since routing is software oriented, but much slower as a result. Packet- or frame-oriented protocols include X.25, Frame Relay, Ethernet, Token Ring, and FDDI. All are characterized by higher latency and delay variations as compared to circuit switching.

Slide #2

Outline

- **Evolution of ATM Technology**
- **BISDN Protocol Reference Model**
 - Physical Layer
 - ATM Layer
 - ATM Adaptation Layer
- **Summary & References**

This paper explains the rationale for the development of ATM as a new, fast-packet networking technology. ATM standards are reviewed, including physical layers, the ATM layer, and the ATM adaptation layer. Reference material is listed on the last two pages.

Slide #4

What Would Be the Characteristics of an Ideal Network Technology?

- Routable at smallest possible level
- Low latency
- Common, scalable access
- Global interconnectivity

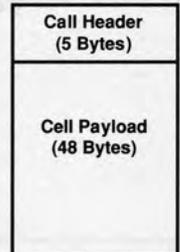
If you could define an ideal network technology, you would combine the best features of existing circuit and packet standards for optimum efficiency. To obtain maximum utilization of network bandwidth, you should route information at the smallest "molecular" level, or in other words, you should use very small, fixed-sized packets. This allows dynamic multiplexing or grooming of multiple signals on the same physical media. Small, fixed-size packets also enable the short, predictable delays required by constant bit-rate services.

Such a network technology would be well suited to all types of data, including voice, video, and bursty data. To be successful, governing standards would have to be internationally accepted and global in scope.

Slide #5

ATM: A New Network Technology

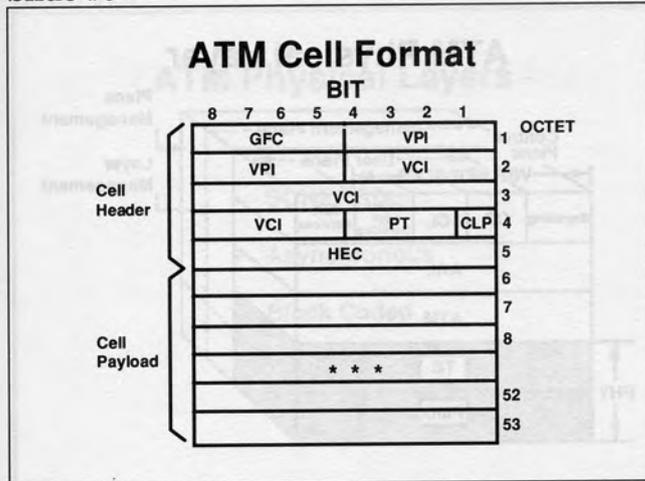
- 53-byte cell
- All types of data
- Transport & switching
- International standards



Asynchronous Transfer Mode (ATM) is a new networking technology based on international standards. It is intended and expected to be suitable for all types of information and provides the infrastructure for broadband networks beyond the year 2000.

With ATM, all information transfers are performed using standard 53-byte cells, each having prescribed structures and methods of formation. All switching and multiplexing is done one cell at a time with each cell being routed independently. All information required by the network to relay a cell from node to node is contained in the cell itself. Bandwidth and access may be dynamically allocated.

Slide #6



The ATM cell header is composed of the following fields:

- GFC Generic Flow Control (UNI only)
Note: The GFC is replaced with an additional 4 bits of VPI at NNI.
- VPI Virtual Path Identifier
- VCI Virtual Circuit Identifier
- PT Payload Type
- CLP Cell Loss Priority
- HEC Header Error Control (8-bit CRC).

The remaining 48 bytes form the cell payload.

Slide #7

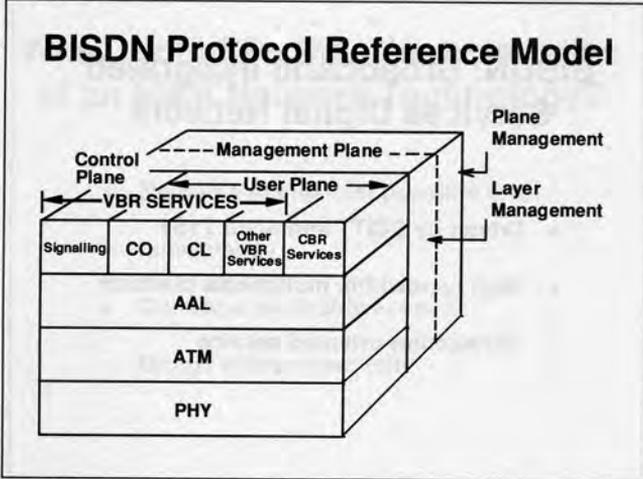
BISDN: Broadband Integrated Services Digital Network

- Driven by CCITT and ANSI T1S1
- High bandwidth, multimedia platform
- Connection-oriented service

ATM has been chosen by standards committees, including ANSI T1 and CCITT SG XVIII, as an underlying transport technology within many Broadband Integrated Services Digital Network (BISDN) protocol stacks. Transport technology relates to the switching and multiplexing techniques at the data link layer.

Understanding Evolving ATM Standards and ATM Design Verification

Slide #8

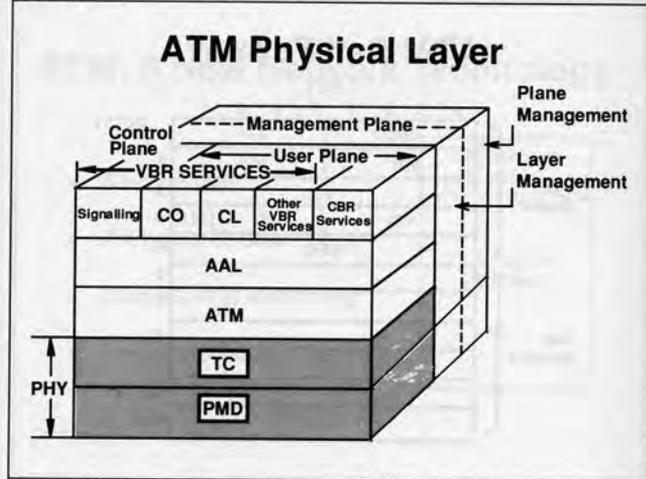


Implementing an ATM bearer service requires the specification of an ATM layer and a related physical layer. These two layers are service-independent and contain functions applicable to all upper layers.

The ATM Adaptation Layer (AAL) adapts the ATM bearer service to provide various networking services including Constant Bit Rate (CBR) and Variable Bit Rate (VBR) services.

The user plane provides the transfer of user-application information. The control plane deals with call establishment, call release, and other connection control functions. The management plane provides management functions and allows the interchange of information between the user plane and control plane.

Slide #9



The ATM Physical Layer consists of two sublayers: Transmission Convergence (TC) and Physical Media Dependent (PMD). ATM cell mappings correspond to existing physical layer standards such as SONET and SDH (synchronous optical networking hierarchies for North America and Europe, respectively), DS3 and E3 (level 3 asynchronous digital interface standards), Block Coding, and others.

The PMD sublayer deals with bit transmission over a physical link, such as fiber optic cable, coax, or copper twisted pair. Issues and specifications, such as line coding, electro-optic conversion, pulse masks, and clock recovery, fall within this sublayer.

Transmission Convergence generates and receives transmission frames and contains all the functions necessary to adapt the service offered by the physical layer to the service required by the ATM layer. In other words, the TC sublayer provides 53-byte cells to the ATM layer. All overhead functions associated with the transmission format are included.

Cell delineation is performed by the TC sublayer based on either explicit control signals or by identification of the HEC. For some physical layer standards such as SONET, the ATM cell payload must be scrambled prior to transmission and descrambled upon reception. Also included in the TC sublayer are HEC generation and verification. This is the 8-bit Cyclic Redundancy Check (CRC) that forms the fifth byte of the ATM cell header.

Understanding Evolving ATM Standards and ATM Design Verification

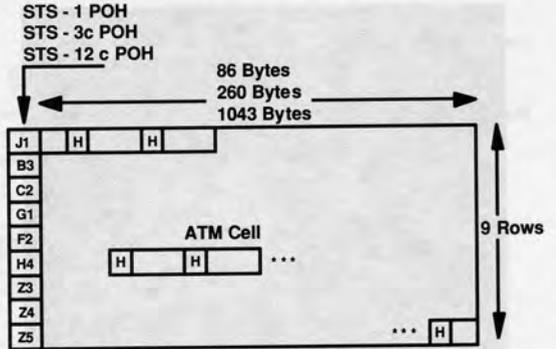
Slide #10

ATM Physical Layers

- SONET/SDH
- Asynchronous
- Block Coded

Slide #11

ATM Cell Mapping — SONET/SDH



The physical layers currently defined to transport ATM cells fall into three categories: synchronous, asynchronous, and block coded. The synchronous and asynchronous standards are borrowed from existing telecom transmission specifications. Block coding, such as the 100 Mbit/s, 4B/5B standard, is used in FDDI LAN applications. One of the strengths of ATM is that all switching and multiplexing is compatible via 53-byte cells, regardless of which physical layer is used to transport the cells. Physical layer compatibility is only necessary on a given physical link. For this reason, a variety of different physical layers, including some yet to be defined, may be deployed in production networks to address different price and performance requirements.

At present, most development work is directed at ATM cell mappings for various rates of the SONET/SDH hierarchy, various asynchronous standards such as DS3 and E3, and the 100 Mbit/s block coded protocol.

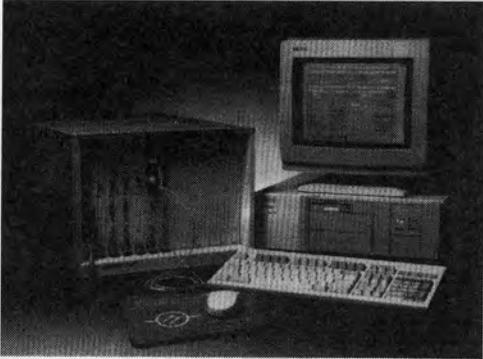
SONET (or SDH as it is referred to outside North America) is a set of international optical network interface standards enabling global network interconnection. It is expected that SONET/SDH interfaces will provide a means for attaining global interoperability in the long run for both public and private networks.

ATM cells are mapped into the Synchronous Payload Envelope (SPE) in a continuous fashion as shown. Upon termination of the SONET frame, cell boundaries are identified and delineated by observing the HEC sequence. The cell payload is scrambled to improve the efficiency of the HEC framing algorithm, as well as randomizing the data for more reliable transport.

Understanding Evolving ATM Standards and ATM Design Verification

Slide #12

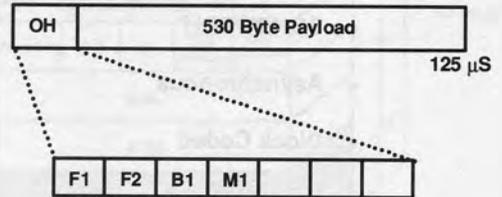
HP 75000 Series 90 ATM Analyzer



HP's ATM analyzer, shown above, can analyze ATM cell streams over a variety of physical interfaces including DS-3, SONET/SDH at 155 and 622 Mbit/s and Pure-ATM at 155 Mbit/s. Physical layer design verification can be performed to ensure that the physical transport system is capable of transmitting ATM cell streams. Physical layer tests include: exercising all of the SONET/SDH overhead functionality, all the PLCP overhead functionality within the DS-3 mapping, and physical OAM functionality within the Pure-ATM cell stream.

Slide #13

ATM Cell Mapping — E3



Although SONET/SDH is considered the preferred transport for ATM cells, it is not yet widely deployed. As a result, ATM cell mappings have been defined for the traditional Plesiochronous Digital Hierarchy (PDH) or asynchronous transmission standards. E3 is a 34.368 Mbit/s standard that is used in Europe and elsewhere. A new frame structure, shown above, is defined, whereby seven overhead bytes are followed by ten contiguous ATM cells. The cells are delineated by identifying the HEC within each cell. The cell payloads are scrambled to provide security against false cell delineation and cell payload replicating the frame alignment word.

Similar mapping has been proposed for other asynchronous and plesiochronous rates.

Understanding Evolving ATM Standards and ATM Design Verification

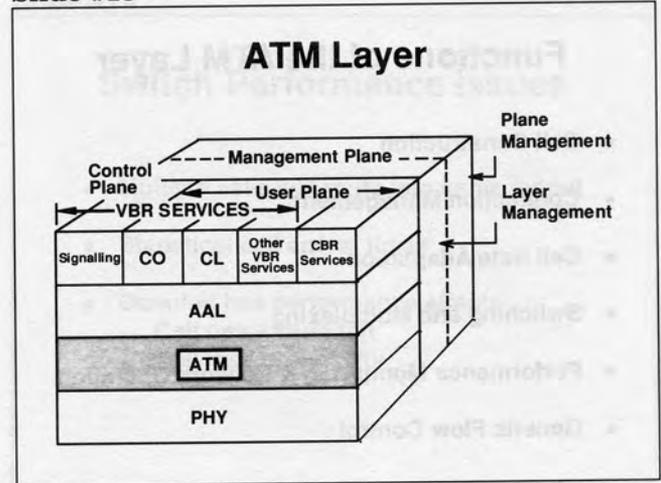
Slide #14

100 Mbit/s Private UNI Interface

- Less complex than SONET/SDH
- FDDI PMD Specification
- 4B/5B line coding — 125 Mbaud line rate
- Explicit asynchronous cell delineation

Since a private User-Network Interface (UNI) does not require the operations and maintenance complexity (nor the link distance provided by telecom standards, such as SONET/SDH), a LAN-like standard has been defined. The 100 Mbit/s standard is based on FDDI physical layer specifications and is intended to use multimode fiber and eventually copper. Unlike the various telecom standard cell mappings, the 100 Mbit/s interface specifies cell delineation based on explicit codes preceding each 53-byte cell. The cell transmission rate is fully asynchronous, and idle codes are sent continuously when no traffic exists. The 100 Mbit/s standard is an economical physical interface for lower performance LAN applications.

Slide #15



The ATM Layer provides transparent and sequential transfer of fixed-size data units between source and destinations with an agreed upon Quality of Service (QoS) and throughput. The ATM layer is service-independent meaning that all information transfers utilize the same cell formats and procedures.

Understanding Evolving ATM Standards and ATM Design Verification

Slide #16

Functions of the ATM Layer

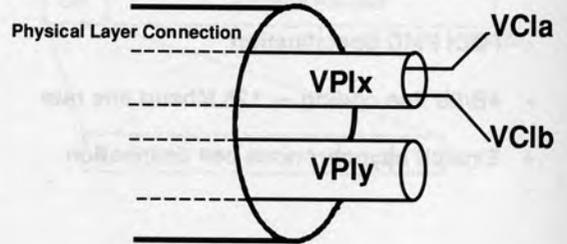
- Cell Construction
- Connection Management
- Cell Rate Adaptation
- Switching and Multiplexing
- Performance Monitoring & Network Operation
- Generic Flow Control

The functions of the ATM Layer are numerous and are categorized as follows:

1. Cell Construction
2. Connection Management
Connection Assignment/Removal
3. Cell Rate Adaptation
Unassigned Cell Generation/Extraction
4. Switching and Multiplexing
Cell Reception
Cell Header Validation
Cell Relaying
Cell Forwarding
Cell Multiplexing/Demultiplexing
Cell Copying
5. Performance Monitoring & Network Operation
Delay Handling
Cell Loss Priority Handling
Usage Parameter Control
Explicit Forward Congestion Notification
Cell Payload Type Discrimination
6. Generic Flow control

Slide #17

ATM Networking Connection Identifiers

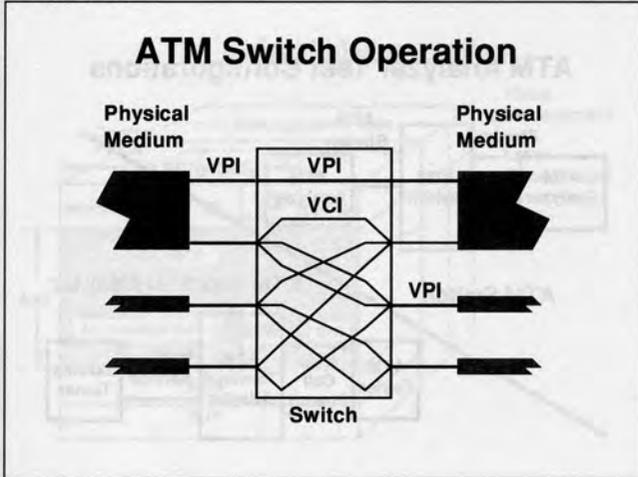


With ATM cell multiplexing, multiple information transfers may exist simultaneously on a given physical link. For this reason, it becomes necessary to distinguish between different transfers in a logical fashion.

A Virtual Channel (VC) is the basic unit of ATM switching and refers to an individual logical circuit. VCs are distinguished by a Virtual Channel Identifier (VCI), which is a routing field in the header of each cell. VCIs are defined unidirectionally on a link-by-link basis.

A Virtual Path (VP) is a logical association or bundle of VCs. The Virtual Path Identifier (VPI) field in the header of each cell is used to distinguish between different VPs.

Slide #18



In ATM networking, the switching function relates cells received on every port to the destination outlet port number and VPI or VCI number. The relations between the VPI/VCI assignments for a given inlet port and the VPI/VCI assignments for each outlet port are established as part of call setup.

Switching may be performed on the basis of Virtual Paths or Virtual Channels. With Virtual Path switching, Virtual Paths are not demultiplexed; cells are routed to outlet ports based only on VPI number. With Virtual Channel switching, Virtual Paths are demultiplexed and cells are routed based on combinations of VPI and VCI numbers.

Slide #19

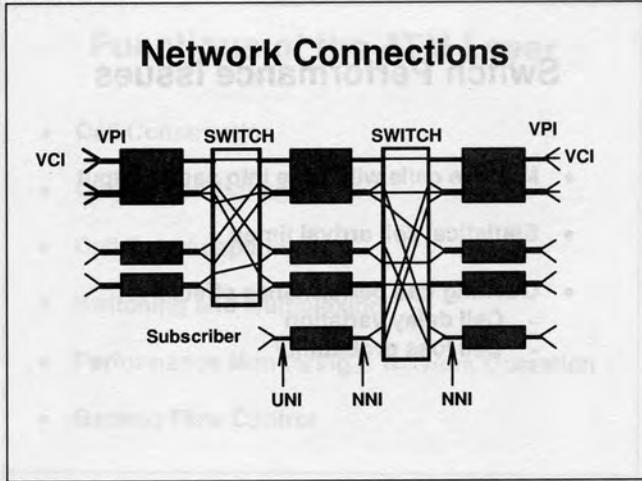
Switch Performance Issues

- Multiple cells will route into same output
- Statistical cell arrival times
- Queuing has performance effects
 - Cell delay variation
 - Cell loss probability

Since cells from multiple inlet ports may route to the same outlet port and the arrival time of all incoming cells is of a statistical nature, queuing is required to resolve the inevitable contention by multiple cells for the same outlet port. Cell queues are normally implemented in first-in first-out (FIFO) fashion and this produces two performance effects: cell delay variation and cell loss probability. Cell delay variation is a function of statistical FIFO length, while cell loss probability occurs due to FIFO overflow. In a simplistic sense, a tradeoff exists between cell delay variation and cell loss probability in an ATM switch.

Switch architects must balance traffic management parameters, such as average cell rates, peak cell rates, and burst duration per VPI/VCI, with required cell delay variation and cell-loss probability limits for proper network operation.

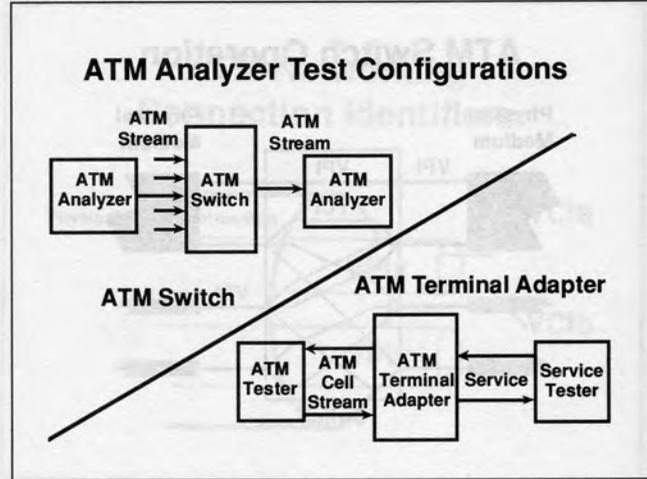
Slide #20



A Virtual Channel Connection (VCC) is an end-to-end connection formed by concatenating a series of VC links. Grade-of-Service (GOS), bandwidth, cell delay, and other traffic parameters are negotiated and allocated for each VCC.

A Virtual Path Connection (VPC) is an analogous concatenation of VP links. A VPC must sustain the GOS of the highest VCI which it contains.

Slide #21

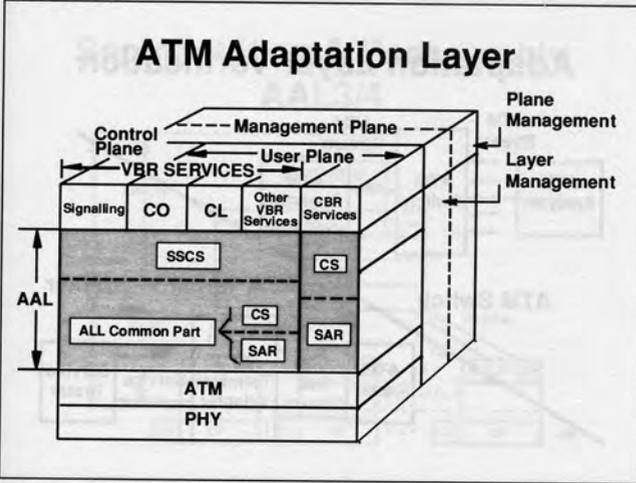


The ATM analyzer, a VXI-based measurement system, can be configured by its user to perform a variety of ATM design verification functions. It can be used to verify the ATM-layer functions performed by an ATM switch including: cell switching, cell performance analysis, and cell congestion evaluation. ATM Optical Load generators can be used to overload the ATM switch.

The ATM Analyzer's terminal adapter testing capabilities will be discussed later in this presentation.

Understanding Evolving ATM Standards and ATM Design Verification

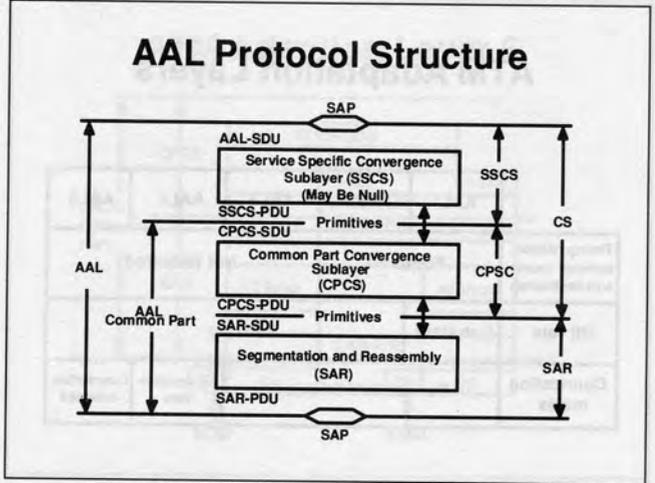
Slide #22



The ATM Adaptation Layer (AAL) defines the processes by which network terminal equipment segments user information into standard data units suitable for transport by the ATM Layer. The AAL matches diverse service requirements to the common format of an ATM cell payload. Networking efficiency is high since all types of information may utilize common resources and protocols for switching, multiplexing and physical layer mapping.

The sublayering of the AAL depends on the service; in particular, CBR and VBR services are handled separately.

Slide #23



The VBR service AAL protocol structure is shown. The Service Specific Convergence Sublayer (SSCS) may optionally provide services or may be null. Services of the SSCS include assured and nonassured data transport. The CPCS sublayer converts user information of an indetermined length into standard packets or into CS-PDUs to be segmented. The SAR sublayer in turn segments CS-PDU packets into 48-byte data units which form the payload of ATM cells.

Understanding Evolving ATM Standards and ATM Design Verification

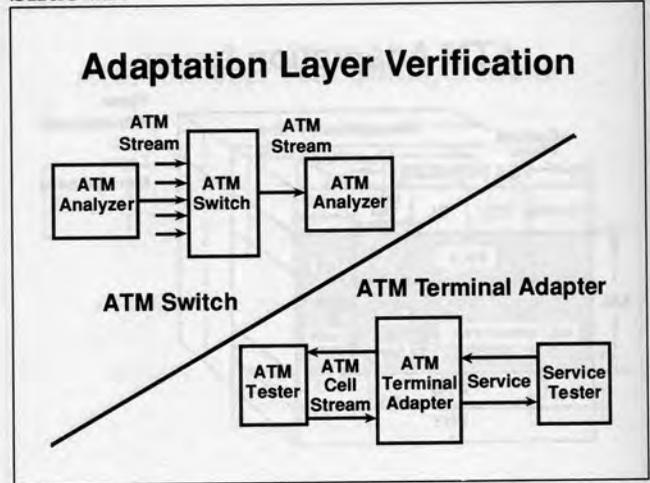
Slide #24

ATM Adaptation Layers

	AAL1	AAL2	AAL3	AAL4	AAL5
Timing relation between source and destination	Required		Not Required		
Bit rate	Constant	Variable			
Connection mode	Connection-oriented		Connection-less	Connection-oriented	

Multiple AALs have been defined to address various requirements as shown. AAL1 is for transfer of audio, continuous bit rate video, and other services having a constant rate related to network timing. AAL2 is used to transport variable bit rate information which has timing related to network timing. AAL3/4 is used for transport of data, MPEG compressed (bursty) video. AAL5 is an alternative data transfer methodology which has been promoted by commercial LAN interests.

Slide #25



Testing higher layer ATM protocols can be accomplished using either the HP ATM Analyzer or an HP Broadband Protocol Tester.

HP's ATM analyzer provides design verification of the AAL type 0, type 1, type 3, type 4 and type 5 adaptation layers. Included in the test suite is functional verification of the service segmentation and reassembly process and verification of the ATM adaptation protocols.

For testing above the ATM Layer, HP will introduce the HP Broadband Protocol Tester in April, 1993. The new test system is designed to help you develop Broadband switches, network equipment and networks. This tester is also the first DUAL-PORT, Broadband tester that provides fully BI-DIRECTIONAL and REAL-TIME measurements of:

- Higher-layer Broadband protocols
- Switch performance throughput
- Conformance testing

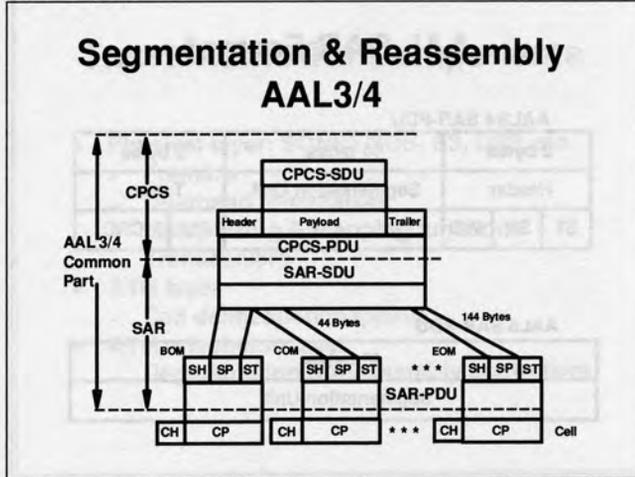
The tester also provides verification for all AAL types 3/4 and 5.

B-ISDN services testing includes:

- User network signalling (Q.93B, ATM Forum)
- Connectionless services (SMDS, CBDS, and I.cls/1.364)
- Limited test capabilities for the ATM layer and the physical layers

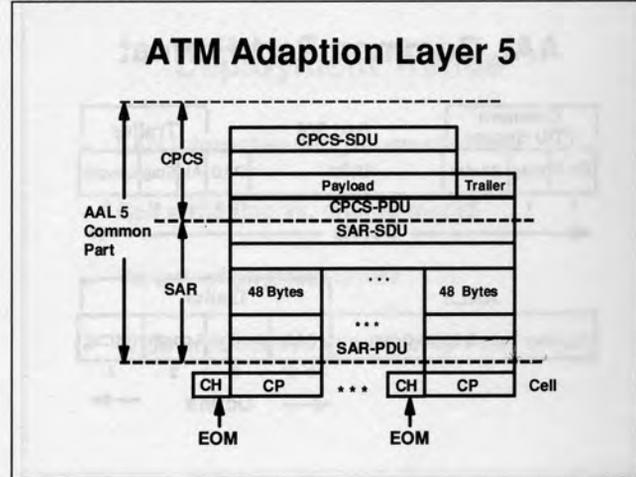
The HP B-ISDN protocol tester is packaged as either a VXI-based system or as a self-contained portable system.

Slide #26



The AAL3/4 segmentation and reassembly process is illustrated. The CPCS-PDU is segmented into contiguous groups of 44 bytes. A 2-byte header and a 2-byte trailer are appended to the 44 bytes forming a SAR-PDU or ATM cell payload. The SAR header and trailer contain information relating to reassembly and cell payload error checking.

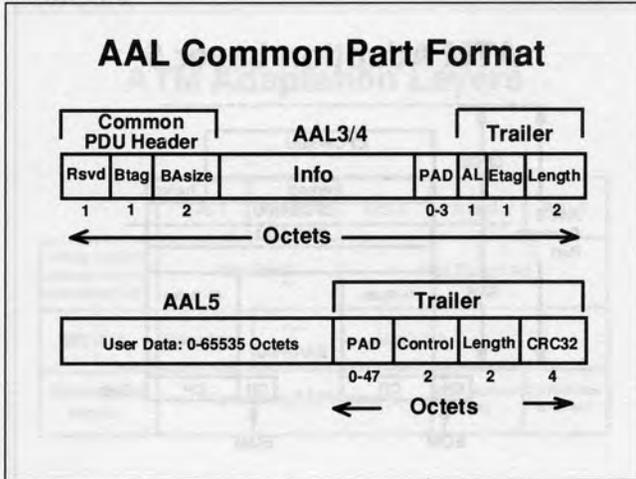
Slide #27



AAL5, which has also been referred to as Simple Efficient Adaptation Layer (SEAL), was developed by the computer and datacom community. The BISDN AAL3/4 protocol previously proposed for VBR traffic was perceived as being incomplete and inefficient for data communications. As a result, AAL5 was proposed particularly for local usage. AAL5 uses the full 48 bytes of cell payload and has no SAR-PDU header or trailer. Reassembly is based on VCI only and there is no MID field as in AAL3/4. One bit in the Payload Type field in the ATM cell header is used to distinguish between End-of-Message (EOM) cells and all others.

Understanding Evolving ATM Standards and ATM Design Verification

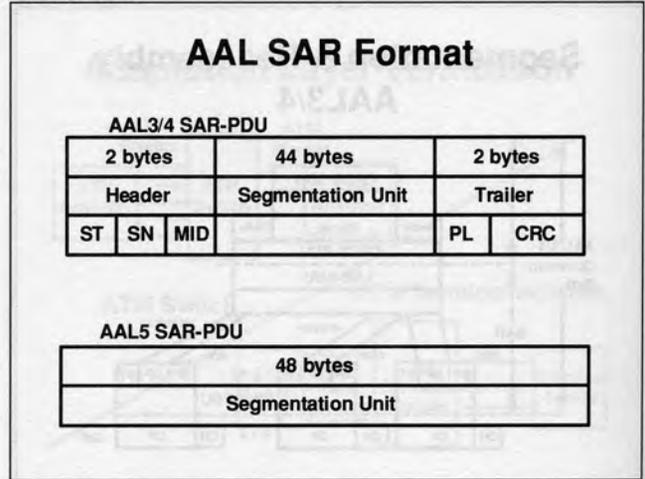
Slide #28



The CPCS-PDU or packet formats for AAL3/4 and AAL5 are shown above. The AAL3/4 information field is padded to a multiple of 4 bytes and bracketed by a header and trailer. Payload error checking is performed by a 10-bit CRC within each cell, not at the packet level. The BAsize field indicates the size of the packet such that upon reception of the first cell, the amount of information to follow is determined.

An AAL5 packet has no header, only a trailer. The length of the packet is not known until the final cell is sent, or is received as the case may be. Within the trailer is a 32-bit CRC covering the entire packet. There is no individual cell payload error checking or other information, which means that a full 48 bytes of payload may be used instead of only 44 in the case of AAL3/4. The packet is padded to a multiple of 48 bytes such that the AAL5 CPCS-PDU always fits exactly into an integer number of ATM cell payloads.

Slide #29



The SAR-PDU or cell payload for AAL3/4 contains 44 bytes of user data and the following header and trailer fields:

- Segment Type 2 bits
- Sequence Number 4 bits
- Message Identifier 10 bits
- Payload Length 6 bits
- Payload CRC 10 bits

The AAL5 SAR-PDU simply contains 48 bytes of user data.

Slide #30

TranSwitch ATM Components

- Physical layer: SONET/SDH, E3, DS3, etc.
 - Framers
 - Overhead terminators
 - Multiplexing & mapping functions
 - Transceivers
- ATM layer
 - Cell delineation/mapping
- ATM adaptation layer
 - Segmentation & reassembly controllers

TranSwitch designs and produces VLSI components for advanced telecom and datacom applications. A full family of physical layer devices are available for SONET/SDH, E3, and DS3 and all may be used in ATM applications. Cell delineation functions which implement transmission convergence and ATM layer functionality have been developed and a high performance AAL controller (SARA Chipset) which supports AAL3/4, AAL5, and CBR traffic at rates up to 155 Mbit/s is in production now. Evaluation boards, user documentation, and applications support are available for all devices.

TranSwitch's products may be configured in a variety of architectures to realize standards-based terminal, LAN, transmission, and switching applications. Such products can significantly reduce the development cost as well as the time-to-market for new system products, and allow system manufacturers to concentrate on additional value added functionality above the defined ATM standards.

Slide #31

Deployment Trends

- LAN internetworking — mostly data
- Cell switching vs. shared media
- Variety of physical layers
- Wide area access will develop more slowly

At present, the focus for near-term ATM deployment is on LAN internetworking or backbone applications. Initially, ATM is expected to be competitive with FDDI on a price and performance basis and be deployed in similar networking environments. The switching capabilities of ATM will be used to improve the performance, scalability, and management of private, local, and campus networks. In particular it is the switching characteristics which distinguish ATM from existing shared media LAN technologies. A variety of physical layers will be deployed to address varying price/performance requirements, while maintaining compatibility at the cell level for switching and multiplexing. Public wide-area network access will grow over time, although it will not be ubiquitous for many years.

Slide #32

Summary

- ATM is being driven by internetworking
- Deployment will enable new applications
- Strong standards support
- No real technology alternatives
- New design verification challenges

ATM is a new networking technology which is suitable for all types of information and is based on international standards. A strong commercial focus is accelerating standardization and deployment of initial systems. ATM has been selected as the underlying transport technology for B-ISDN and is expected to eventually be widely deployed in private and public networks around the world. Designers and developers can derive significant utility from tools which can recognize and verify ATM protocols.

Slide #33

References

- CCITT
- ATM Forum
- IEEE

CCITT

- I.113: Vocabulary of Terms for Broadband Aspects of ISDN
- I.121: Broadband Aspects of ISDN
- I.150: B-ISDN Asynchronous Transfer Mode Functional Characteristics
- I.211: B-ISDN Service Aspects
- I.311: B-ISDN General Network Aspects
- I.321: B-ISDN Protocol Reference Model and its Applications
- I.327: B-ISDN Functional Architecture
- I.361: B-ISDN ATM Layer Specification
- I.362: B-ISDN ATM Adaptation Layer (AAL) Functional Description
- I.363: B-ISDN ATM Adaptation Layer (AAL) Specification
- I.413: B-ISDN User-Network Interface
- I.432: B-ISDN User-Network Interface - Physical Layer Specification
- I.610: OAM Principles of the B-ISDN Access

ATM Forum

ATM User-Network Interface Specification

IEEE

- IEEE 802.6: Distributed Queue Dual Bus Subnetwork of a Metropolitan Area Network

Understanding Evolving ATM Standards and ATM Design Verification

Slide #34

References (Cont.)

- ANSI
- Bellcore

ANSI

T1S1.5/92-001 AAL SSCOP Baseline Document
T1.ATM-199X ATM Layer Functionality and Specification
T1.AL4-199X AAL 3/4 Common Part
T1.CBR-199X AAL for Constant Bit Rate Services Functionality and Services
T1S1.5/92-005 Connectionless Service Layer Functionality and Services
T1S1.5/92-010 AAL5 Common Part Functionality and Services
T1S1.5/92-111 Constant Bit Rate AAL Architecture

Bellcore

TR-TSY-000772: Generic Requirements in Support of Switched Multi-Megabit Data Service
TR-TSY-000773: Local Access Switching System Generic Requirements in Support of SMDS
FA-NWT-001109: Broadband ISDN Transport Network Elements Framework Generic Criteria
FA-NWT-001110: Broadband ISDN Switching System Framework Generic Criteria

FA-NWT-001111: Broadband ISDN Access Signalling Framework Generic Criteria for Class II equipment
TA-NWT-001112: Broadband-ISDN Used to Network Interface and Network Node Interface Physical Layer Generic Criteria
TA-NWT-001113: Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements
SR-NWT-001763: Preliminary Report on Broadband ISDN Transfer Protocols

Slide #35

Recommended Resources

- **Equipment and Accessories**
 - HP 75000 Series 90 ATM Analyzer
 - HP Eclipse Protocol Analyzer
- **Other Resources**
 - HP's BISDN Seminar
 - TranSwitch's ATM Technology and Applications Seminar

Physical Layer Design Issues for Serial Communications: A SONET Case Study

Josef Kreidl

PEP Modular Computers
Apfeltranger Str. 16
8950 Kaubeuren,
Germany

Gunter Rucker

PEP Modular Computers
Apfeltranger Str. 16
8950 Kaubeuren,
Germany
Phone: 0049 8341 43020

This paper will be distributed
at the symposium.

Robert Kraus

Motorola GmbH
Schatzbogen 7
8000 Munich, Germany
Phone: 0049 899 21030

Andreas Gunther

Motorola GmbH
Schatzbogen 7
8000 Munich, Germany
Phone: 0049 899 21030

1993
High Speed Digital
Systems Design & Test
Symposium

Alternatives for Data Transfer in High-Speed Systems

Josef Kreidl

PEP Modular Computers
Apfeltranger Str. 16
8950 Kaufbeuren,
Germany
Phone: 0049 8341 43020

Robert Kraus

Motorola GmbH
Schatzbogen 7
8000 Munich, Germany
Phone: 0049 899 21030

Gunter Rucker

PEP Modular Computers
Apfeltranger Str. 16
8950 Kaufbeuren,
Germany
Phone: 0049 8341 43020

Andreas Gunther

Motorola GmbH
Schatzbogen 7
8000 Munich, Germany
Phone: 0049 899 21030

**1993
High Speed Digital
Systems Design & Test
Symposium**

Alternative for Data Transfer
 in High-Speed Systems

Abstract

PEP Modular Computers has been designing VME systems since the early 1980s, specializing in the design of compact solutions for a wide range of market segments. These compact solutions use 3U VME (single height Eurocard) style boards and therefore suffer from relatively slow data transfer capability, fundamentally due to the limited size of the data and address busses. Recent enhancements to the

VMEbus specification (Revision D) have boosted the data transfer capability from about 20 Mbytes/s to 40 Mbytes/s. However as performance expectations of systems increase, even this significant improvement cannot avoid the fact that backplane connection still represents a bottleneck in many designs. The tremendous performance growth of CPUs and other system components, along with ever increasing

customer expectations, can be held primarily responsible for this view.

Design activities to enhance the data transfer capability over the parallel bus soon started to show the physical limitations of increasing operating frequencies. This paper looks at the way the problem was overcome and how it led to the birth of a new chipset which offers considerable potential for the future.

Authors

Josef Kreidl

Current Activities:

Josef Kreidl is president and major owner of PEP Modular Computers, the leading 3U VMEbus manufacturer for high quality industrial, medical, aerospace and telecommunication applications worldwide.

Author Background:

Josef was born in Innsbruck, Austria and studied electrical engineering at HTL Innsbruck. From 1970 - 1975, he was a design engineer at Olympia and Telefunken. In 1975, he founded PEP.

Gunter Rucker

Current Activities:

Gunter Rucker is technical director and part owner of PEP Modular Computers and is responsible for all VME designs and the AUTOBAHN transceiver chip.

Author Background:

Gunter studied electrical engineering at the Technical University of Augsburg. From 1968 - 1982, he was a design engineer at Olympia and a department manager at MBB. In 1982, he joined PEP Modular Computers.

Alternatives for Data Transfer
in High-Speed Systems



PEP Modular Computers

Authors (cont'd)

Robert Kraus

Current Activities:

Robert Kraus is an engineering manager responsible for Logic Integrated Circuits in Europe. A key area is the development and design of new products for communications and computer peripheral applications.

Author Background:

Robert was born in Munich, Germany and studied communication engineering at FH Muncih. He started with Motorola in 1985 as a design engineer, developing a precision AC test simulator for integrated circuits.

Andreas Gunther

Current Activities:

Andreas Gunther is a product engineer at Motorola and is responsible for the development and application of parallel bus interface IC and microprocessor support IC.

Author Background:

Andreas was born in Brehna, Germany. He studied physics and electronics and graduated from the Technical University Chemnitz in Germany.

Development of
Microprocessor Speed



With the performance of modern microprocessor-based systems increasing rapidly, designers are able to realize more and more sophisticated and complex. Real-time processing, high-resolution and powerful parallel processing are all achievable.

Systems have one thing in common: the need to move data around the system extremely quickly. In practice, it is this need that effectively limits the performance of the system. It is common to ask whether that is a single-board or multi-processor design employing a standard backplane. These inherent have their roots in the analog world. The complex interaction between conventional semiconductor technologies and the effects of the PC board they are mounted on.

As performance increases, we must the performance of the technologies used to achieve the increase. In particular, the methods and technologies used to transfer data within digital systems have experienced continual enhancement to achieve higher and higher data throughput rates. However, something of a crossroad has now been reached, where major electrical and physical limitations threaten this continual improvement. As the 25 to 30 MHz data rate boundary is crossed, new techniques become essential, but more importantly, new technologies become viable options for the performance systems of the future.



Alternatives for Data Transfer in High-Speed Systems

Slide #1

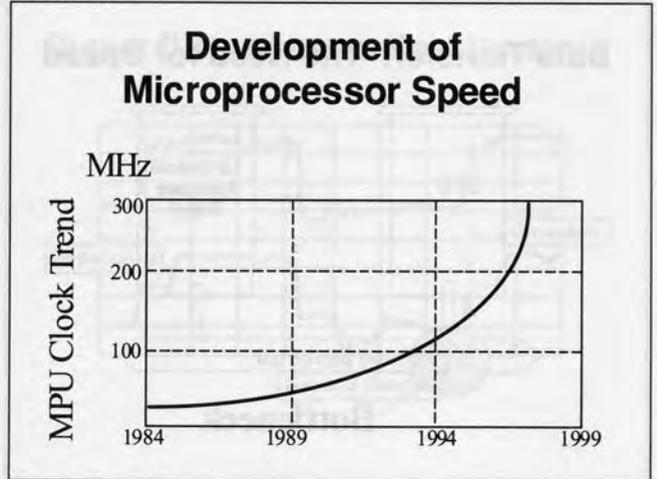
Alternatives for Data Transfer in High-Speed Systems

AUTOBAHN

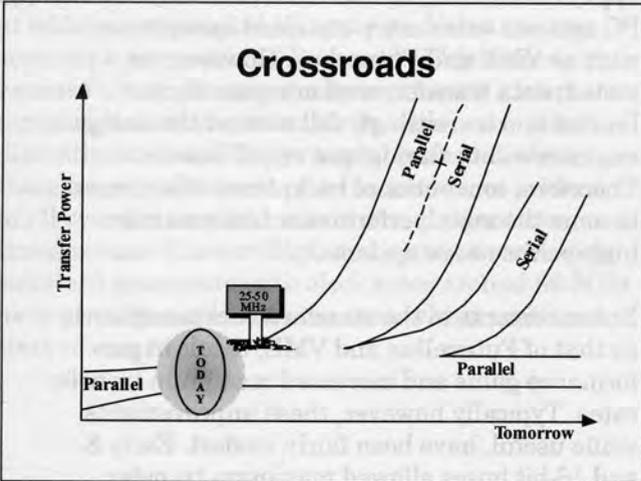
PEP Modular Computers




Slide #3



Slide #2



With the performance of modern microprocessor-based systems increasing rapidly, designers are able to realize more and more sophisticated and complex applications. Real-time processing, high-resolution imaging, and powerful parallel processing are all more easily achievable.

All these systems have one thing in common: the need to move data around the system extremely quickly. In practice, it is this need that effectively limits the performance of the system. It is common to all systems, whether that is a single-board computer or a multiprocessor design employing a proprietary or standard backplane. These inherent limitations have their roots in the analog world of complex interaction between conventional semiconductor technologies and the effects of the PC board they are mounted on.

As system performance expectations increase, so must the performance of the technologies used to achieve the increase. In particular, the methods and technologies used to transfer data within digital systems have experienced continual enhancement to achieve higher and higher data throughput rates. However, something of a crossroads has now been reached, where major electrical and physical limitations threaten this continual improvement. As the 25 to 30 MHz data rate boundary is crossed, new techniques become essential, but more importantly, new technologies become viable options for the performance systems of the future.

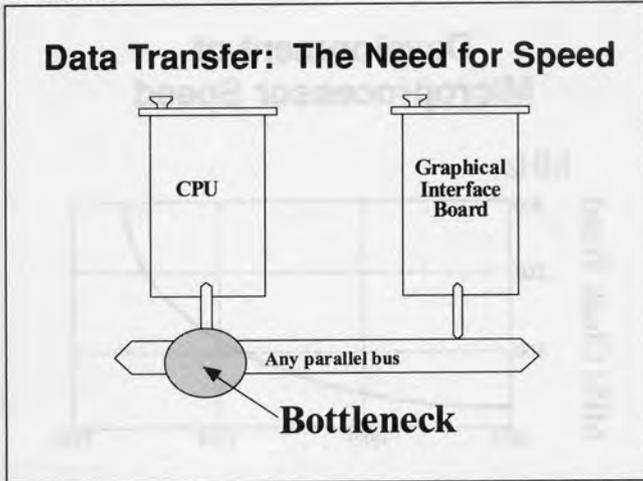


MOTOROLA



Alternatives for Data Transfer in High-Speed Systems

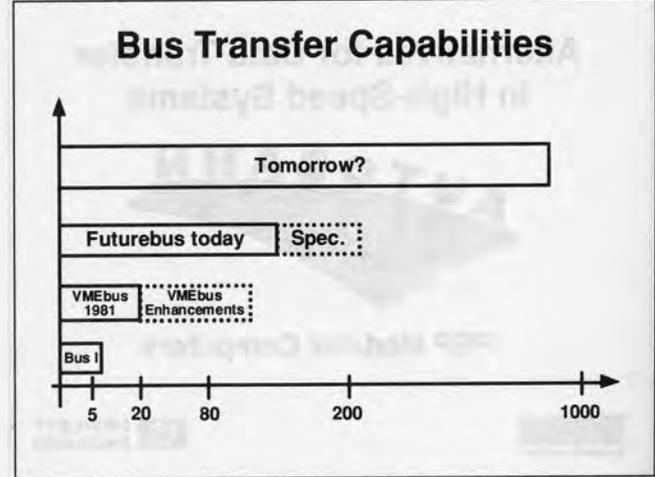
Slide #4



Consider a bus-oriented graphics application, where a central processor interfaces to a graphics processing board via a backplane connection. For normal operation, transfer rates in the region of 160 MByte/s between the CPU and the graphics card would be needed. For an HDTV application, this requirement could be much higher.

This level of data throughput clearly cannot be achieved because a typical 32-bit bus system, such as that employed in a VME or MULTIBUS II system, delivers a maximum of 80 MByte/s when running at 20 MHz.

Slide #5



The flexibility and versatility of parallel data transfer in multiple card systems made it the preferred choice for virtually all microprocessor applications, ranging from conventional PC systems to industry-standard backplanes, such as VME and Futurebus. However, as stated, data transfer rates are somewhat limited and increasingly fail to meet the design engineer's data throughput requirements. Therefore, conventional backplanes oftentimes become the main performance bottleneck for high-performance systems.

Enhancements to the standard backplane, such as that of FutureBus and VME, result in performance gains and increased maximum transfer rates. Typically however, these improvements, while useful, have been fairly modest. Early 8- and 16-bit buses allowed maximum transfer rates of only about 20 MByte/s. The latest VME enhancements allow transfer rates of up to 80 MByte/s. Today, FutureBus+ operation is specified at up to 160 MByte/s transfer rates, by using 64 data lines. However, this leads to a significant increase in the physical size of the system and the possibility of increased electrical problems due to the high toggle frequencies.

To achieve the high data rates that will be required for tomorrow's systems, two or possibly three times the performance of currently available systems will be a minimum requirement.



MOTOROLA



**HEWLETT
PACKARD**

Slide #6

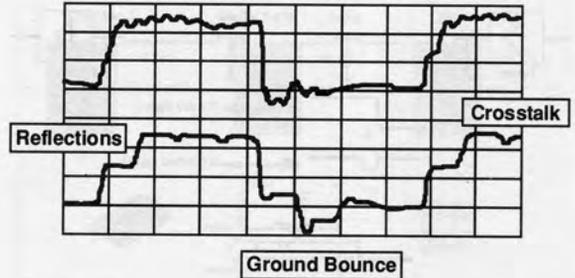
Backplanes vs. "PCB" Buses

- Backplanes run into problems at 20-25 MHz
- PCB boards show the same problems — just at higher frequencies

A typical parallel backplane design starts to experience signal integrity problems where clock rates break the 20 MHz barrier. This is approximately 10 to 20 MHz below the frequency at which severe signal integrity problems start to occur on a conventional PC board. This occurs because a backplane typically represents a more complex electrical environment, with many more discontinuities and longer signal paths. However, these same problems start to become more evident in all systems as clock rates rise and signal rise times reduce. Current high-end systems employing standard processors with clock rates around 66 MHz have already moved into the area where problems start to occur.

Slide #7

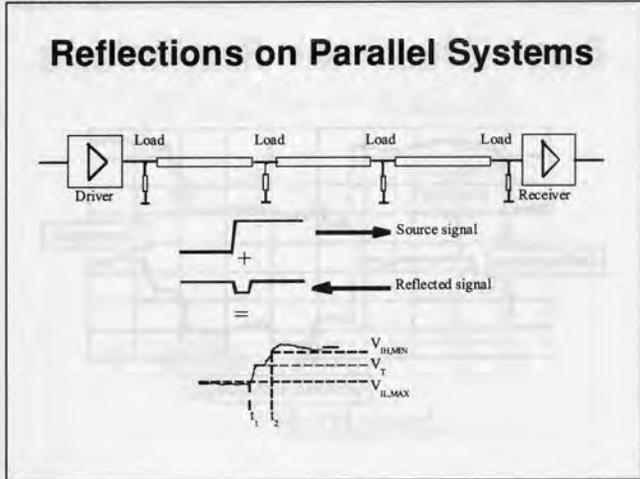
Signal Distortions in Real Systems



The root causes of poor signal integrity in high-speed systems are well known. Reflections, ground bounce, and signal crosstalk all contribute to the effect. The manifestation of these effects can be unreliable system behaviour or poor EMC/EMI performance. Analysing the mechanisms by which these problems arise reveal complex interaction between all system components, including semiconductor packages, connectors, cabling, and the PC trace itself. With increasing frequency and bus width, these intrinsic effects come to dominate the attention of the design engineer.

Alternatives for Data Transfer in High-Speed Systems

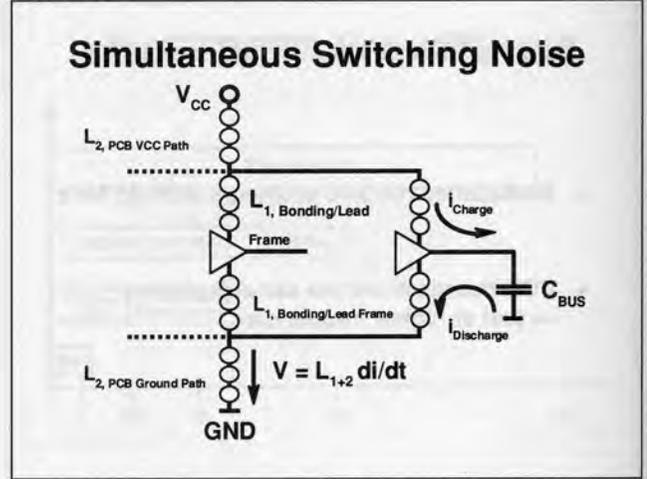
Slide #8



Reflections occur on any line that is not properly matched or terminated. In reality, it is never possible to achieve a perfect transmission line in a multipoint backplane environment. Every connector represents an additional electrical load and discontinuity. The larger the backplane, the more connections are present, and therefore, the more sources of reflection.

A typical example is shown above where the limited drive capability of standard TTL and CMOS devices can never achieve a stable -high+ condition at the initial transition. In order to establish a stable signal, a certain delay time must be allowed so that all reflections can settle. If this delay is not taken into account, the undefined state between t_1 and t_2 , exactly around the switching threshold for standard TTL or CMOS, may result in metastability or incorrect switching of subsequent stages. This need for a settling period is, of course, directly proportional to the speed of the system and to the environment of faster operation. The more loads added to a backplane system, the more settling time is stretched, thus limiting the performance. Potential solutions utilise powerful bus drivers with low output impedance. However, the required drive current is becoming beyond the capabilities of standard logic bus devices and the total power requirement of a wide 32- or 64-bit bus would be prohibitive.

Slide #9



Most EMI/EMC problems are directly attributable to simultaneous switching noise, or ground bounce. IC technologies have migrated towards faster and faster risetimes, driven by system design requirements toward higher speed to achieve more powerful microprocessor systems. However, the package evolution did not follow this technology trend. In most cases, the standard IC package is still the DIL (Dual In Line) package or surface mount SOIC (small outline) package. These packages, in particular, suffer from relatively high inductive load per pin, mainly due to the length of the bondwire inside the package. The consequences can be easily seen by referring to the following basic equation,

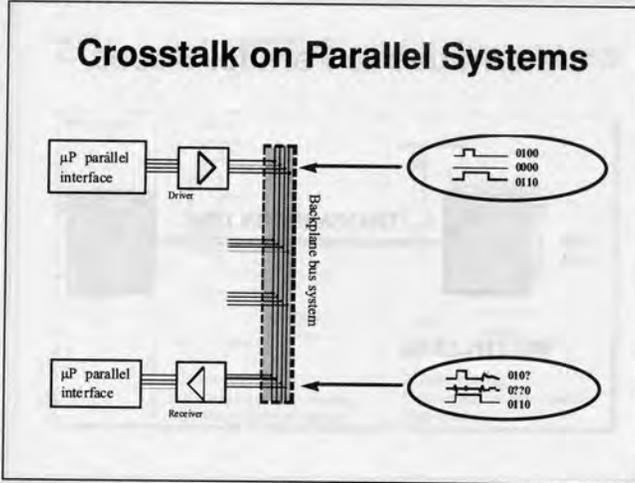
$$v = L di/dt ,$$

where L is the inductance and di/dt is the rate of change of current.

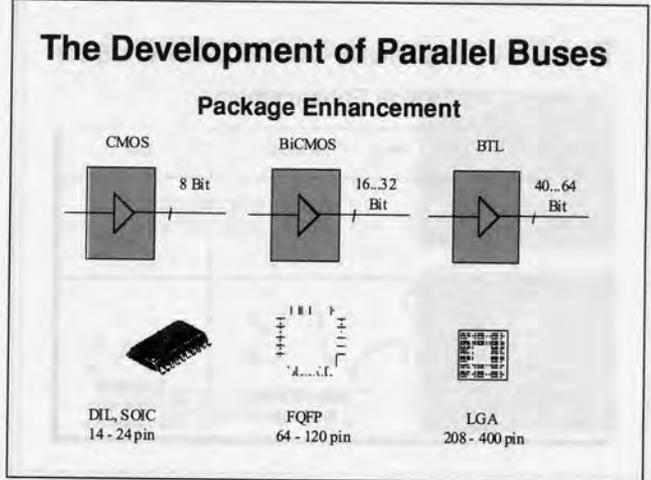
It can easily be seen that as system risetimes decrease (dt) even small values of bondwire inductance could potentially lead to very large voltage spikes. For example, the high i_oL current of buffered outputs simultaneously switching can lift the GND level to cause the output to change state. As predicted, this effect gets worse with signal risetime, and with technologies, such as CMOS, that employ full GND-to-VCC and VCC-to-GND output transitions. This becomes the single most critical signal integrity problem in the system.

Alternatives for Data Transfer in High-Speed Systems

Slide #10



Slide #11



Crosstalk is the electrical coupling of signals between one trace and another and occurs wherever signals are routed in close proximity to one another. Between two lines on a PC board there exists a coupling capacitance and a mutual inductance, through which electrical energy is exchanged. The size of these inductance and capacitance values depends on the geometry chosen for the board layout and the trace spacing employed. Large parallel buses are especially susceptible to crosstalk problems, with complex coupling impedances present. CMOS inputs are very sensitive to crosstalk effects because of their high input impedance. Even small amounts of coupled energy from an active parallel line can build up a critical parasitic signal that could cross the switching threshold and could cause the device to toggle. This kind of bit error can only be detected by comparing the original source signal to the received signal. The failure rate received within a given time interval — the Bit Error Rate (BER) — provides an invaluable guide to the degradation of performance in any bus system.

Because there is no complete alternative to parallel buses in most microcomputer systems, IC manufacturers are being challenged to develop enhanced products more robust against this type of disturbance. Although these signal integrity problems can never be fully eliminated on parallel buses, there are ways to at least reduce their impact and thus extend the maximum performance of the system.

For the IC manufacturer, one of the most important areas to improve is the device packaging itself. Recalling that simultaneous switching noise is the single most important signal integrity problem, a significant performance gain can be realized by reducing the VCC and GND pin inductance, formed by the bondwire. This can be done by increasing the number of supply pins or minimizing the interaction between different output buffers on the same chip. Using a square outline package can also help, as these have short, uniform length bondwires, which minimize the inductive effects and help to maintain balanced propagation delay times, even on wide output buffers.

A useful side effect of using a square package is that flat packages have much lower thermal resistance which helps to reduce the power dissipation problems that are also a feature of any high-speed bus application.

Alternatives for Data Transfer in High-Speed Systems

Slide #12

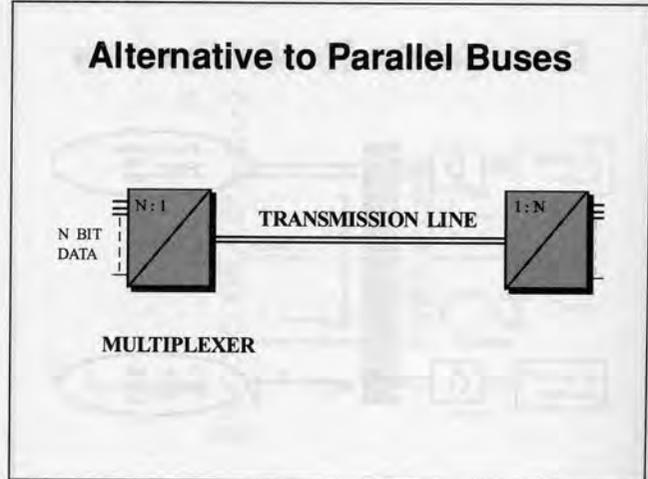
The Development of Parallel Buses

Design Enhancement

Bus Device Technology	CMOS	BiCMOS	BTL
Output Structure			
Output Signal Level			
	Full swing Sharp corners Steep edges	Reduced swing Moderate edges Steep edges	Low swing Moderate edges Low capacitance

Beyond packaging the IC, design and process technology can help to minimize parasitic effects. CMOS is most used for its economic power consumption, bipolar technology for its drive capability. Combining the two yields BiCMOS, where CMOS technology is used internally and output buffers are created using powerful bipolar transistors. Ideally, both advanced packaging and chip technology are used to achieve maximum performance. An example here would be the ALE_{EX}ISTM (Advanced Low Power EXpandable Interface Solution) product line from Motorola, that offers 16- to 20-bit-wide bus interfaces and transceivers in BiCMOS housed in a 64-pin Fine Pitch Quad Flatpack. A development of this will be BTL (Bus Transceiver Logic). As well as employing BiCMOS technology, this family also uses reduced voltage swing open collector outputs. These outputs also include a serial diode that significantly reduces capacitive load of disabled outputs, thus reducing excessive peaks of charge and discharge currents.

Slide #13



State-of-the-art process and package technologies may help to extend data rates to the 50 to 100 MByte/s range, which is a considerable improvement. However, there are physical limits to a backplane and any other parallel system that simply cannot be overcome.

Today's microprocessor systems use 16- or 32-bit-wide buses for data transfer, that result in a considerable board space requirement and the burden of the undesirable parasitic effects causing the signal integrity to deteriorate. These problems increase with speed and line length. Simply said, the more lines being used, the more problems to cope with. One alternative, and hardly a new one at that, is to employ serial data transmission, the premise being that it is easier to control the signal integrity of a single line than it is for a parallel bus. Obviously, it is an attractive solution, but one which brings with it a new dimension of ultra-high speed. Converting a 200 MByte/s signal (32 bit x 50 MBit/s NRZ parallel bus) results in a serial data stream of 1.6 GBit/s, equivalent to a frequency of 800 MHz on the serial line. This is completely out the range of any standard TTL, CMOS or BiCMOS technology, requiring a completely different approach.

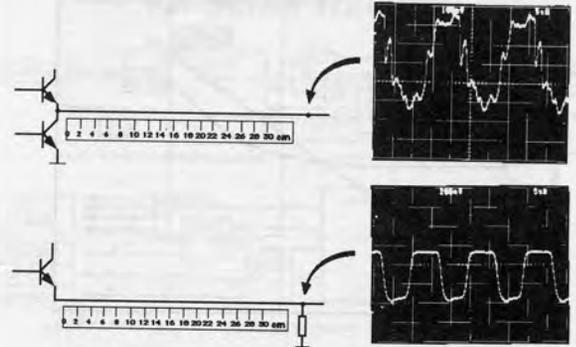
Slide #14

ECL vs. CMOS Output Structures

Technology	CMOS	
Output Structure		
Output Signal		
	Low drive capability Full voltage switching technology High charge/discharge currents	Differential design Transmission line drive capability High speed

Slide #15

An Alternative Solution: ECL

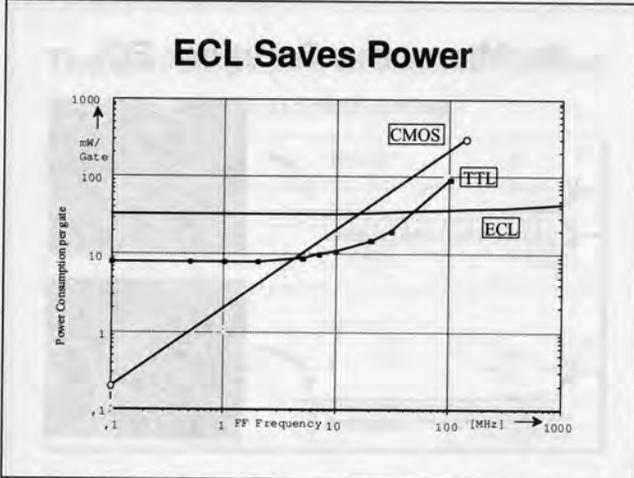


An enabling technology for these high frequencies is ECL. Here, a basically bipolar technology is used with a different circuit design, so that the output transistors are not switching type but are biased to operate in their linear range. This concept allows maximum switching speed of the transistor limited only by its cut-off frequency. In addition, ECL outputs can operate in differential mode, with output voltage swings of only 0.8 V, that guarantees excellent noise immunity to disturbances of the board environment and vice-versa. This is further supported by the fact that with such low impedance, ECL outputs can drive true terminated transmission lines with 50-ohm impedance, an ideal environment for high-frequency transmission.

The theoretical benefits of ECL are proven in real-world systems. With their fast on/off switching speed, conventional TTL/CMOS outputs contain very large high-frequency content, reaching well into the RF range. At the same time, uncontrolled impedances around the board cause reflections which heavily disturb the signal. As can be seen in the example above, a 25-MHz signal is routed around a board and after only 25 to 30 cm, a realistic distance for many PC boards and most backplanes, it is already badly corrupted and possibly unacceptable for the reliable operation of subsequent stages. Compare this to the superior performance of an ECL driven line. Even at 10 times the clock rate and 3 times the distance, the signal integrity is good. While the bit rate of this signal is very high, the reduced voltage swing technology means that it exhibits negligible RF emissions to the environment.

Alternatives for Data Transfer in High-Speed Systems

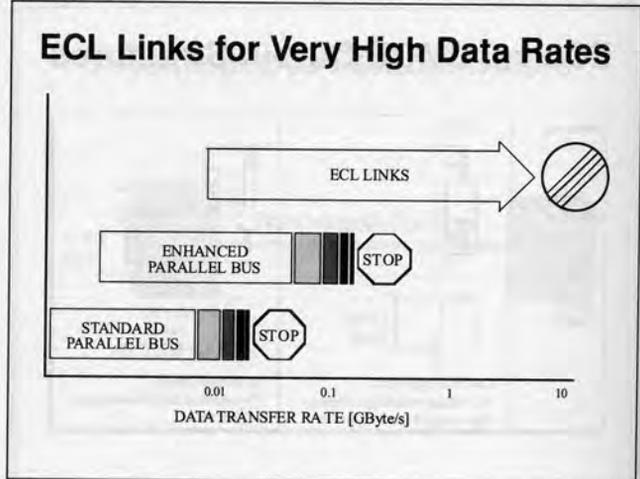
Slide #16



A historical concern of ECL has been its power consumption. Comparing ECL with CMOS and TTL structures, this is a valid concern at low frequencies. Under all conditions, the ECL transistors are biased, leading to relatively constant supply current requirement. CMOS devices however, with their complementary transistor structure, have a very low power consumption at low operating frequencies.

This picture changes drastically with frequency: CMOS transistors always switching full swing outputs have to completely charge and discharge all internal and external capacitors at every cycle, thus leading to a linear increase in the power consumption. Bipolar type TTL structures look somewhat better at certain frequencies; however, the switching output transistors also lead to excessive power drain beyond certain frequencies. This frequency is technology-dependent, but in the range around 30 to 70 MHz. This is not the case for ECL, where the output transistors are not actually switching rather just modulating around a bias point, the power consumption is very nearly constant over the operating range of the device. Depending on the individual application, the breakeven point in power drain occurs at around 30 to 40 MHz. Beyond that, ECL can dominate.

Slide #17



Looking at the best solutions for realizing data transmission, there is a scale of economy for both the parallel as well as the serial concept. Parallel buses at clock frequencies up to 60 to 80 MHz and further will be the first choice for bridging short on-board data exchange between microprocessor and peripherals. Serial, high-speed links will be used where maximum data throughput is required or longer distances are to be bridged such as board-to-board or board-to-external peripheral. There may be future supercomputers where the serial concept will be the most efficient solution, even on-board.

Latest ECL processes have cutoff frequencies of 25 GHz, such as Motorola's MOSAIC V technology (0.7 um, four layer bipolar; Motorola Oxide Self Aligned IC), which currently allows data rates of 10 GBit/s and beyond. Applications requiring such ultimate performance already exist today for example, real-time, high-resolution image processing, crosspoint switches for telecommunication PBX systems, and any type of large parallel processing system.



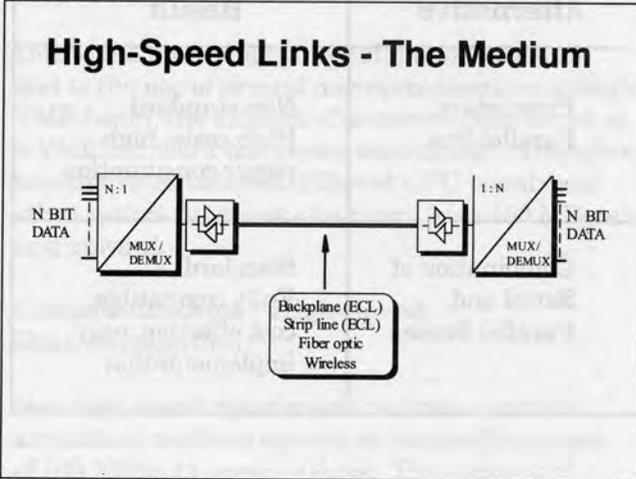
MOTOROLA



**HEWLETT
PACKARD**

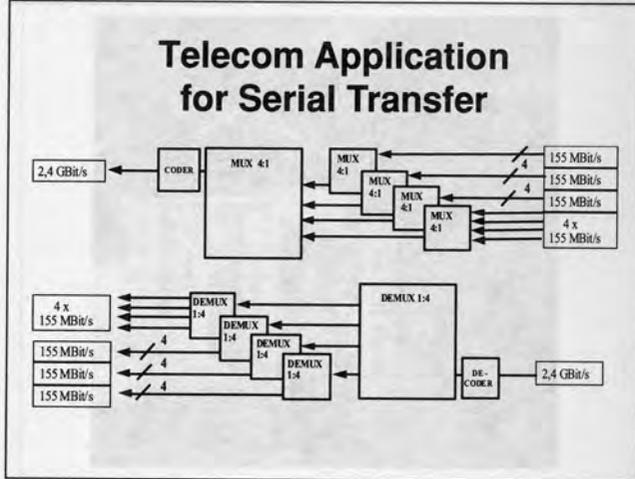
Alternatives for Data Transfer in High-Speed Systems

Slide #18



Once parallel data words are multiplexed into a high-speed serial data stream, some considerations are required on how to transport it, as its effective frequency may reach some 100 MHz to 2 GHz. With differential ECL, data transmission over distances of 0.5 to 1 m can be achieved on backplanes and PC boards as long good transmission lines are maintained. Even considering the additional components required for MUX/DEMUX functions, the serial link offers a good alternative for wide parallel multipoint or point-to-point connections through crowded, densely populated microcomputer boards. For longer distances, fibre optic links or wireless transmission is the recommended medium, where the multiplexed serial data could be fed directly into a laser driver or RF modulator.

Slide #19



Independent of the application—Telecom Switching Systems, HDTV, Video conferencing and digital mobile phone systems, such as GSM and DECT—international standards in Europe and USA have defined standard platforms to enable the different systems using common transmission channels to interface to each other. These standards named SDH (Synchronous Digital Hierarchy) in Europe and SONET (Synchronous Optical Network) in the USA are compatible on certain layers having fixed data rates:

Example:	STM-1	155 MBit/s	OC - 3
	STM-4	622 MBit/s	OC - 12
	STM-16	2.480 MBit/s	OC - 48
	STM-64	9.952 MBit/s	

These standards are based on multiplexing parallel data through several layers, utilizing high-density fibre optic links for transmission, allowing powerful and highly efficient data links to be realized using latest technologies in all stages of the system.



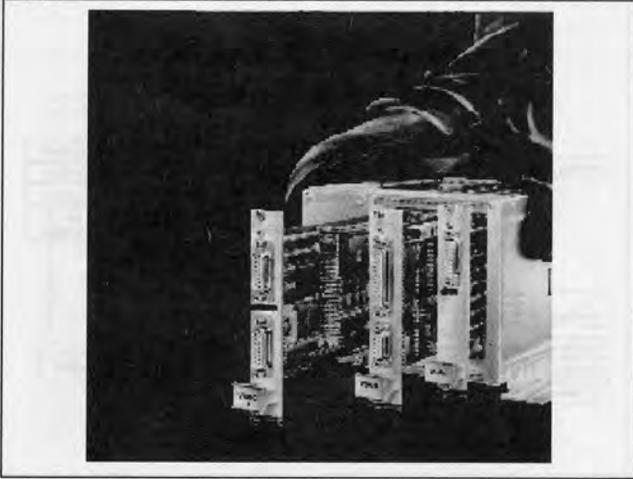
MOTOROLA



**HEWLETT
PACKARD**

Alternatives for Data Transfer in High-Speed Systems

Slide #20



PEP Single Height

Since 1975, PEP has produced modular microcomputer systems for a wide range of applications based on the small, single-height DIN form factor, using one single 96-pin connector. This allows a very compact and efficient design of complex high quality systems.

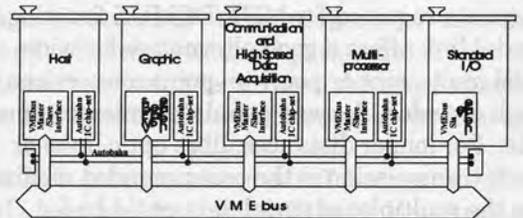
Through tremendous performance increases in the PEP product range throughout the 1980s, the parallel bus became the bottleneck. Many applications could not be solved this way. Very early on, we had to accept that the frequency increase of the parallel bus had physical limitations, for example, no solution.

Therefore, as many other companies, PEP was investigating resolving the problem with a specialized proprietary parallel bus. This would lead to a non-standard solution, with increased costs and power consumption. Many companies worldwide spent hundreds of man-years to develop such proprietary busses to meet their requirements. PEP was looking for a solution within the standard. The alternative solution was to combine the parallel bus, with an ultra-high-speed serial link, utilizing all advantages of either transfer method.

Alternative	Result
Proprietary Parallel Bus	Non-standard High costs, high power consumption
Combination of Serial and Parallel Busses	Standard Fully compatible, cost effective, easy implementation

Slide #21

Applications for High-Speed Data Transfer



This block diagram of a VMEbus system shows some typical applications where high-speed data transfers are needed between different components on the bus.

Graphics and Vision Systems

A single picture in color or printing quality format requires several MBytes of digitized data. In most graphics and vision system applications, numerous pictures/sec must be processed or displayed. This results in data transfer rates of 100 to 200 MByte/s between graphic and camera boards, as well as other system functions, like CPU or communications boards.



MOTOROLA



**HEWLETT
PACKARD**

Alternatives for Data Transfer in High-Speed Systems

Multiprocessing - Parallel Processing

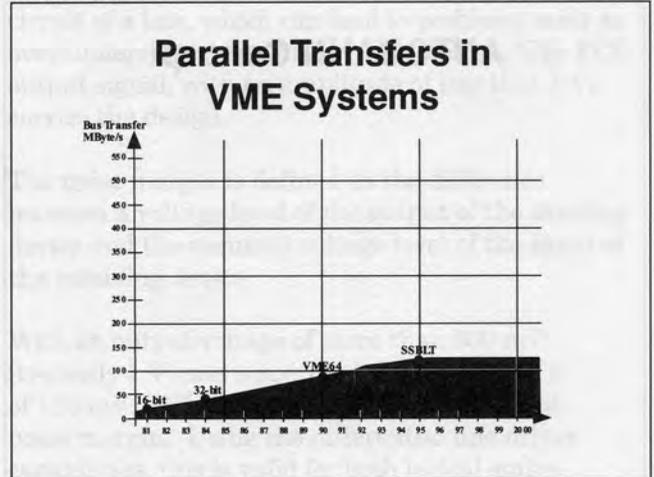
Due to the increasing power of microprocessors and to the use of several microprocessors on a single VME board, the amount of processed data found on a VMEbus board has grown enormously. Therefore, transfer rates between different CPU boards and other system functions also range from 100 MByte/s and upward.

Communications - High-Speed Data Acquisition

New high-speed communication links and data acquisition methods operate in bandwidth ranges of 100 Mbit/s to several Gbit/s. This amount of data must first be linked to a standard board on a VMEbus system, then transferred on the bus to another system function (CPU). This leads to bus transfer needs of hundreds of MByte/s. Typical high-speed applications are DAT, HDTV, FDDI, and so forth. In communications, logic analyzers and high-end research applications (such as super-colliders) require extremely high bandwidths for bus systems.

All of these application examples identify existing parallel bus structures as the throughput bottleneck of the system.

Slide #22



For a VME or any other standard backplane manufacturer or user, a key requirement of any design is its use of the standard interface. In the case of VME, enhancements to this interface have led to steady performance improvement. But none could really offer the ultimate performance requirement because of the constraints of standardization. Changes to the specification allow a system standard function to change but this process can be slow. Employing new technologies, such as the Motorola BTL (Bus Transceiver Logic) can offer significant improvement; but again, this path leads away from standardization and can potentially lead to significant cost issues.



MOTOROLA



HEWLETT
PACKARD

Alternatives for Data Transfer in High-Speed Systems

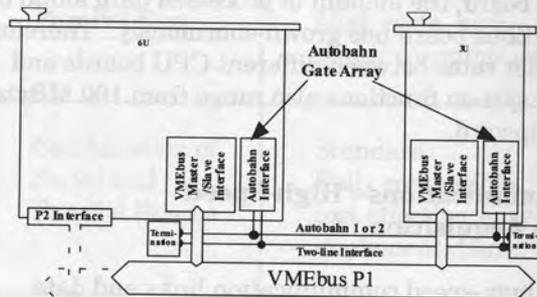
Slide #23

AUTOBAHN Concept

These constraints led to a different approach to the problem, from which the Autobahn concept emerged. The standardization issues of VME could be avoided by using parts of the system not defined in the VMEbus specification. The possibility of two spare backplane pins was of course known; thus, the idea of using them for serial data transfer between compatible boards became reality. The main benefit of this approach was that the defined parallel system remained unaltered, and hence the serial interface promised performance in addition to that already available.

Slide #24

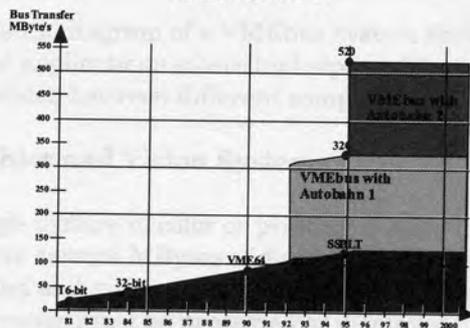
Implementation



In concept, the implementation was very straightforward with the inclusion of MUX/DEMUX capability on any board wishing to utilize the additional interface. However, the concept was complicated because to achieve the required data throughput of 32-bits at 50 MHz, a maximum frequency of 900 MHz on the backplane serial line would be required. However, the potential benefits of such a system justified the development investment.

Slide #25

Combination Serial/Parallel Transfers



The future solution is combining parallel transfers with ultra-high-speed serial data links.



MOTOROLA



**HEWLETT
PACKARD**

Alternatives for Data Transfer in High-Speed Systems

Combining both systems allows utilization of the best of each method.

The parallel bus is used twice; for parallel transfers (standard VME) and to arbitrate for and set up the high-speed serial link, including error handling.

As the serial bus does not need a protocol overhead for collision detection, the amount of net transferred data is nearly 100%.

Since parallel and serial transfer can run independently of each other, the performance of both transfers can be added. This leads to a very economical alternative for high-speed data transfers, eliminating the bottleneck in existing bus systems.

The higher the amplitude, the larger the reloading circuit of a line, which can lead to problems such as over/undershoot, crosstalk, and radiation. The ECL output signal, with an amplitude of less than 1 V, carries the design.

The noise margin is defined as the difference between a voltage level of the output of the sending device and the required voltage level of the input of the receiving device.

With an output voltage of more than 600 mV (typically 1 V) and a specified input hysteresis of 150 mV, ECL technology offers an excellent noise margin. Using the differential line driver capabilities, this is valid for both logical states.

Thanks to a very low specified input hysteresis of 150 mV, convenient waveforms, even sine waves, may be used to reduce jitter, ground bounce and crosstalk.

ECL devices, manufactured in MOSAIC V technology, offer an on-chip toggle frequency of 25 GHz. This gives a high transmission safety for AUTOBAHN data transfers and no need for additional circuitries, such as CRC or parity logic.

Slide #26

Selection of Technology

System Requirements	CMOS	ECL	GaAs
Transmission line drive capability			
Differential drive capability	---		
Low signal amplitude logic	---		
High noise margin	---		---
Moderate waveforms	---		---
High speed technology (1 GHz)	---		
Low cost			---
Low power at 1 GHz	---	?	?

The low output impedance and the high current drive capability of ECL makes it an ideal technology for driving transmission lines.

The differential amplifier does not switch on an off, but simply steers between two paths. This current stability greatly simplifies the design and avoids bounce effects. With common node noise rejection of 1 V or more, ECL line receivers are less susceptible to common node noise.



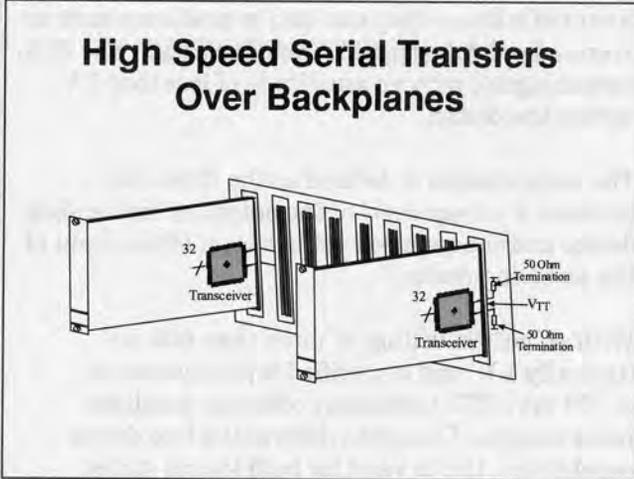
MOTOROLA



**HEWLETT
PACKARD**

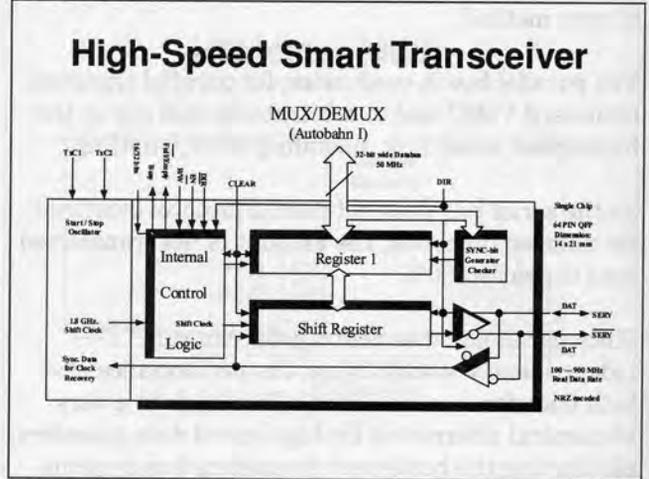
Alternatives for Data Transfer in High-Speed Systems

Slide #27



A standard 21 slot VME backplane is assembled with all the female DIN connectors fitted, thus allowing a partially loaded system to be expanded by adding extra cards. This complicates the design of the 50-ohm differential transmission lines because each of these connectors represents a signal stub, and hence a large potential reflection source when unterminated. For this reason it was decided to terminate each transmission line on the backplane, as close as possible to the DIN connector, thus preventing the problems associated with trace stubs. When a board is plugged in that uses the Autobahn serial interface, the signals are routed to the transceiver chip that is mounted within 25 mm of the DIN connector. When a card is plugged in that is not compatible with the Autobahn serial interface, it has no loading effect. As a result, it makes no electrical connection to the serial bus which is terminated in its characteristic impedance on the backplane. In this way, signal integrity is maintained at a high level throughout the system. Since the ECL signal is in this way terminated on the left and right end of the backplane, the signal transfer characteristic is independent of the slot transmitting the data; in other words, there are no slot-dependent transmit/receive requirements.

Slide #28



The high-speed transceiver allows a contiguous serialization of a 32-bit wide parallel data bus and vice versa. At a transfer rate of 200 MByte/s, the frequency on the parallel bus is on the order of 50 MHz, while on the serial bus, it is up to 900 MHz. Data is in NRZ format with one start bit added for each transmitted byte to allow received clock regeneration. This is achieved by an on-chip start/stop oscillator with a 3.2 GHz frequency, which is resynchronized by each start bit of the received data stream. This eliminates the need for an external VCO, Phase-Locked-Loop (PLL), and filters. In addition, it significantly reduces transfer time delays and jitter.

A critical part of the transceiver design was its power consumption. The Motorola PECL technology used offers extremely low power consumption coupled with the requirement for only a single +5 V supply. The 64-pin QFP chip consumes about 1 W when fully loaded, most of which is attributable to the parallel interface. This illustrates another reason to serialize high-speed data transfer.



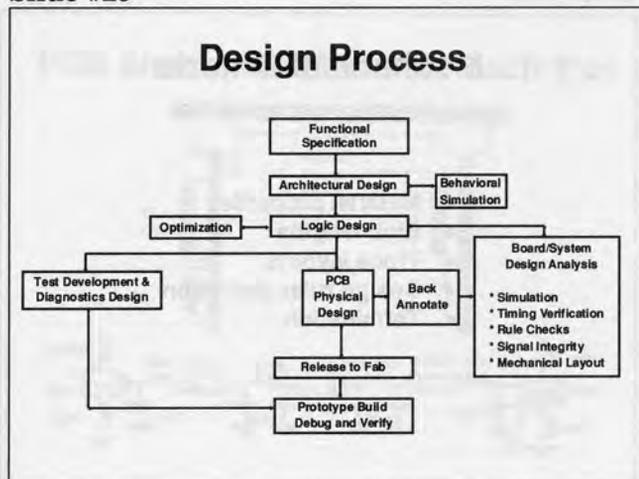
MOTOROLA



HEWLETT
PACKARD

Alternatives for Data Transfer in High-Speed Systems

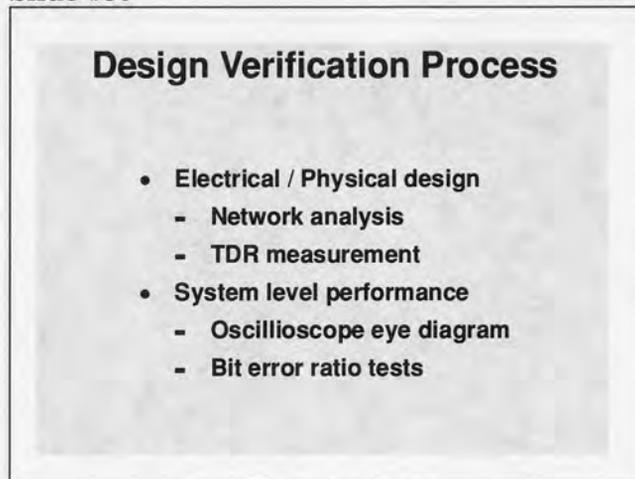
Slide #29



To ensure a reliable and robust design, the most important part of the design process was the PCB physical design, its simulation and modeling and the verification of timing and signal integrity performance. Mechanical layout was critical; although, in this case where a standard VME product was being enhanced, certain constraints were present. Design criteria can be separated into electrical and mechanical topics; although, as frequency increases, the line between the two becomes blurred.

Simulation and modeling of the combined effects is essential in the early stages using tools such as the Hewlett-Packard MDS or HDT Simulation products. The remainder of the paper concentrates on the measurement and verification of the post-simulation design.

Slide #30



The design verification of the components and complete system were completed in four logical steps. This first step was to verify the electrical properties of the backplane and confirm that the mechanical design chosen yielded a good electrical design. At higher frequencies, these two points are very closely linked. The main measurement techniques used here were Network Analysis and Time Domain Reflectometry (TDR). These measurements were used to measure attenuation at the predicted maximum data rate and to evaluate the impedance characteristics along the transmission line, which should represent a uniform, controlled impedance.

At a system level, transmission quality and bit error performance must be evaluated. This is achieved by looking at oscilloscope eye-diagrams to confirm signal integrity and noise margin performance, and by employing Bit Error Rate tests to display the received bit error rate of a known PRBS (Pseudo Random Bit Sequence) after transmission through the system, at a maximum data rate.

All of these tests represent those steps which must be undertaken by any designer creating a high-speed digital data transfer system. Only by comparing the predicted simulation results with real measurements and modifying the design accordingly can a completely functional and robust design be created.

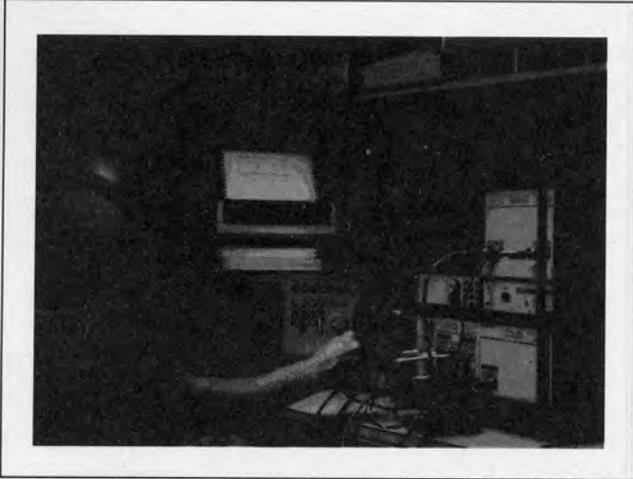


MOTOROLA



HEWLETT
PACKARD

Slide #31



Two instruments were chosen to perform much of the design validation processes. The HP 54120-series digitizing oscilloscope and TDR can perform Time Domain Reflectometry, Time Domain Transmission, and eye pattern analysis. In addition, it can also be used as a high-bandwidth scope, at up to 50 GHz. The HP 71600-series BER tester can perform a wide variety of BER measurements under all kinds of varying conditions, using industry standard algorithms.

Slide #32

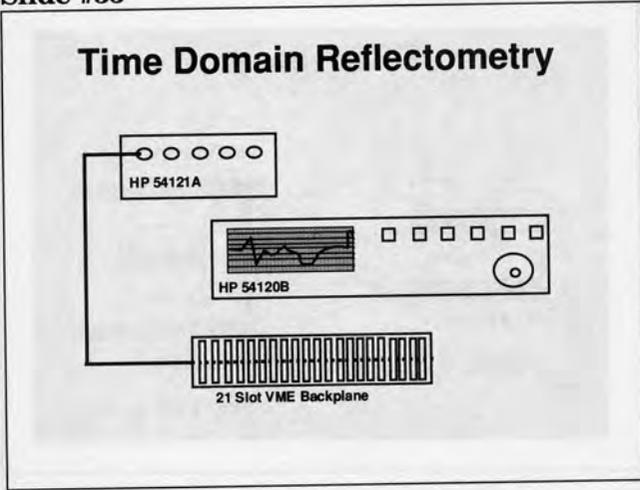
Mechanical Criteria

- Material properties
- Stub lengths
- Trace layouts
- Design rules definition
- Termination

With higher speeds come new concerns that have to be understood in order to make your model and simulation tools provide accurate representations of the design. Mechanical characteristics of the design are becoming increasingly significant aspects which these modelling tools must account for in order for the simulations to be accurate. These mechanical aspects of the design are starting to become part of modelling/simulation software packages, like HP's MDS and HFSS. Even common software tools, like HSPICE, are including some of the physical descriptions of the design in order to provide better accuracy, and consequently faster design cycles. Basically, the inclusion of these mechanical conditions help the software tools understand the electromagnetic properties of a design, which is a key aspect as speeds increase. When these considerations are included in the early design phase, fewer board cuts are required and designs are more likely to work.

Alternatives for Data Transfer in High-Speed Systems

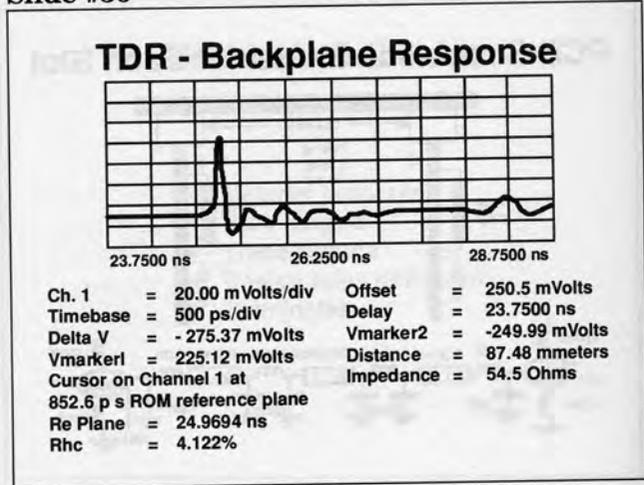
Slide #35



Following the initial simulation of the backplane, it was extremely important to verify the accuracy and validity of the models used. A TDR measurement was chosen for this purpose because of its accuracy and intuitive ease of interpretation. The measurement allowed the impedance along the backplane to be measured and the irregular effects of the connectors to be evaluated. This information could then be fed back into the simulation process if necessary as an aid to refining the simulator models.

The setup chosen is shown above, where the backplane was connected to the TDR channel of an HP 54121T Oscilloscope. A major benefit of the HP 54120-series is that it allows normalization of measured results to take account of real-life risetimes that are different from the risetime produced by the instrument's TDR step generator. In this case, the TDR step has a risetime of approximately 35 ps, whereas the Autobahn chipset has risetimes in the region of 150 ps. Therefore, normalization allows evaluation of the Autobahn backplane at realistic system risetimes.

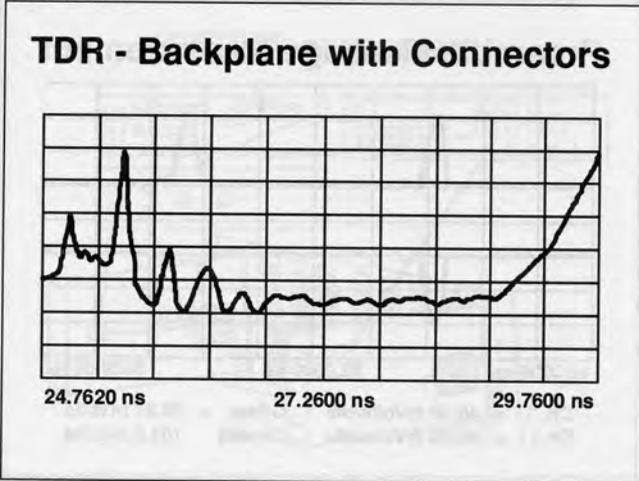
Slide #36



Shown in the diagram is the impedance of the backplane with no connectors fitted and with the end of the line terminated in its characteristic impedance. As can be seen from the diagram, the line exhibits a good uniform impedance, as simulated, even including the effects of signal stubs. Away from the launch point, these stubs cause an impedance variation of approximately 4 ohms. The first spike on the trace is the SMA connector used as the launch point. These connectors typically exhibit an inductive load as shown. The effect of the connector stubs is a small capacitance.

Alternatives for Data Transfer in High-Speed Systems

Slide #37



Slide #38

- ### Electrical Criteria
- Noise margins
 - Line terminations
 - Impedance characteristics
 - Signal / Waveform integrity
 - Reflection, crosstalk

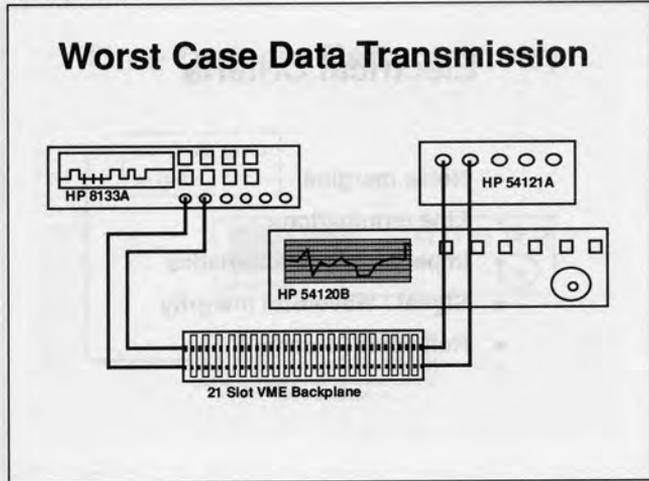
With connectors fitted, the impedance of the transmission line is very much more disturbed. In the diagram above, which is scaled at the same rate as the previous diagram, all connectors are fitted to an unterminated backplane and the TDR signal is launched through a standard VME male/female DIN connector. The overall effect is a general lowering of the impedance of the transmission line due to the capacitive loading of the connectors. This measurement compares the worst-case performance to a 35 ps edge which is considerably faster than the effective edge rate of the line. As the edge speed is reduced by normalization, the impedance discontinuities become less pronounced in their effect.

The TDR measurements allowed verification of the basic model used but indicated that the connector model needed some refinement. At the time of writing this paper, a new connector is under development which will improve the impedance profile of the backplane.

The topics discussed so far have primarily dealt with a standalone backplane with no active load connected. Connecting the driver/receiver boards with active transceivers completes the actual transmission line, and allows realistic system performance measurements to be made. The one aim of the design is to maintain good signal integrity under all operating or load conditions. Worst-case scenarios must be identified and performance verified under these conditions. In the case of a backplane-based system, this will almost certainly occur when all slots are occupied with operating transceivers. Noise margin calculations completed in the initial logic design phases must be validated under load conditions and the effects of impedance mismatch induced reflection or crosstalk reevaluated under real-life operating conditions.

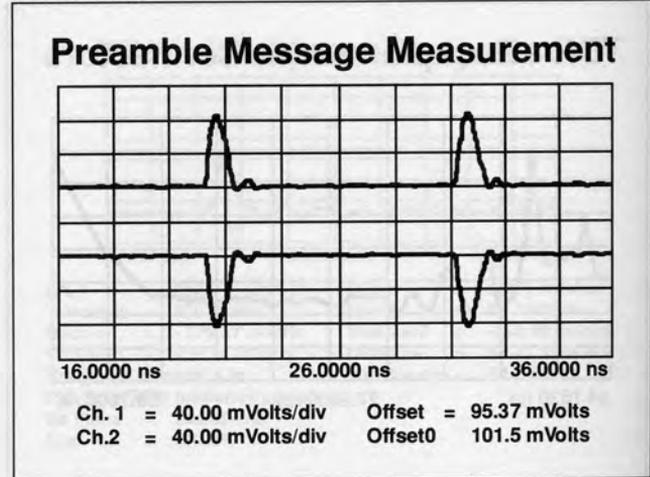
Alternatives for Data Transfer in High-Speed Systems

Slide #39



The next level of testing involved evaluating a loaded backplane's performance with real-life signals applied. The TDR measurements indicated a less than ideal transmission environment, even though the normalized measurements promised better performance. The following stimulus measurements allowed the backplane to be evaluated under realistic conditions. An HP 8133A was used to provide stimulus to an ECL driver/receiver set driving the differential lines on the backplane. This gave the flexibility of not only generating a bit-sequence at up to 3 GHz, but allowed pulse width variations to be simulated and their effects evaluated.

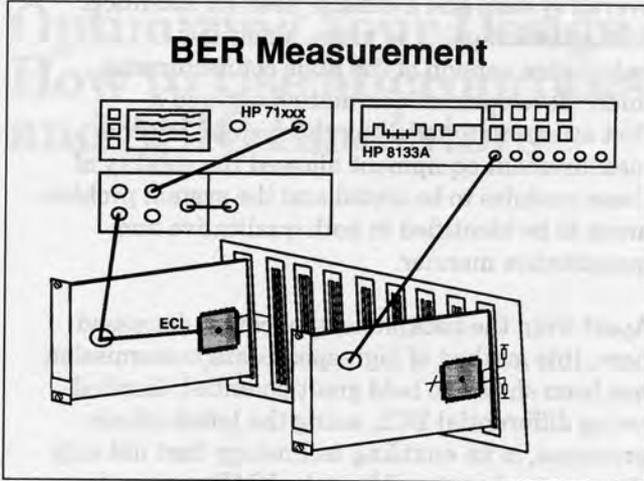
Slide #40



The diagram above shows the transmission of a preamble message across the backplane at 1.4 GBit/s which corresponds to a serial line frequency of 700 MHz or a parallel data transfer rate of approximately 155 MBytes/s. It can be seen that the signal transmission quality is good, even allowing for the connector mismatch problems of the backplane. The preamble message is a sequence of zeros separated by the Autobahn clock regeneration synchronization bit and represents one of the possible worst-case signals that can be present on the link.

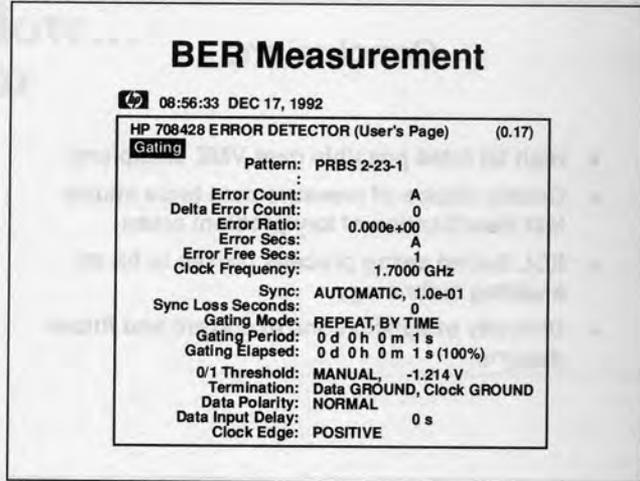
Alternatives for Data Transfer in High-Speed Systems

Slide #41



The measurements so far have proved qualitatively that the transmission of data at these rates is possible over the backplane. However, this must be verified quantitatively. This is done by use of a Bit Error Rate Test, where a PRBS (Pseudo Random Bit Sequence) is transmitted over the serial line from the HP 8133A through the ECL transceivers and the backplane and received by an HP 71600-series High-Speed BER Tester. This instrument detects a PRBS being transmitted by the pattern generator (it can also generate a full range of PRBS signals itself) and performs a comparison with the signal received and the predicted actual value, thus giving a quantitative measure of the quality of the data transmission.

Slide #42



The above shows a BERT listing made on the system discussed so far with the backplane driven by the ECL transceiver devices. It indicates that the current revision of the system can operate at 1.7 Gbits/s before it starts to experience bit errors. At 1.8 Gbits/s the received bit error rate was measured at 300 errors per million bits transmitted. With no connectors fitted, the error free figure was extended to 2.3 Gbits/s, which tends to confirm the findings of the measurements made with the oscilloscope and data generator and those of the TDR. Note that on the HP BERT system, the displayed bit frequency has a 1:1 relationship with the bit rate.

In practice, a true PRBS is not possible across the Autobahn link because there is always a sync bit added to each byte of the message for clock regeneration purposes. This clock regeneration method guarantees received data bits are always sampled in the center of the bit period, which improves bit error performance and helps minimize signal jitter.

Slide #43

Conclusions

- High bit rates possible over VME backplane
- Correct choice of measurement tools allows fast identification of key problem areas
- ECL limited swing process proves to be an enabling technology
- Diversity of applications in current and future designs

The study has shown that the transfer of data over the serial bus using the technology discussed in the paper is possible at the high data rates predicted. Initial measurements have identified

the din connector as a source of concern in the overall system and a definite area for attention. At the time of going to press, a new low, inductance version of the male connector was under development. Simulation allowed a first attempt at modelling the system, but the measurement equipment allowed the validity of these modules to be tested and the system problem areas to be identified in both qualitative and quantitative manner.

Apart from the backplane application discussed here, this method of high-speed data transmission has been shown to hold great potential. Limited-swing differential ECL, using the latest silicon processes, is an enabling technology that not only gives a new lease on life to the VME system, but promises much for high-speed PC board systems between microprocessors or high-speed subsystems. Here, the transmission line design will be much simpler because of the lack of complex connector structures used on the backplane.



**QUICKTURN
SYSTEMS**



**TESLA
CORPORATION**



**HEWLETT
PACKARD**

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Eric Amador

TESLA Corporation
1025 Buckland Avenue
San Carlos, CA 94070
Tel: (415) 637-9479
Fax: (415) 637-9361

Tim Chambers

Hewlett-Packard Co.
8245 N. Union Blvd.
P.O. Box 617
Colorado Springs, CO 80901
Tel: (719) 590-5570
Fax: (719) 590-7679

Naeem Zafar

Quickturn Systems, Inc.
325 East Middlefield Road
Mountain View, CA 94043
Tel: (415) 967-3300
Fax: (415) 967-3199

1993 High Speed Digital Symposium



Abstract

The dramatic increase in the complexity of microprocessor-based designs and their associated system software has caused companies to employ larger design teams, and consequently to incur higher hardware and software design costs. In addition, modern digital systems now invariably include custom digital circuits in the form of proprietary Application-Specific Integrated Circuits (ASICs). Managing the technical design issues associated with these projects, and meeting time-to-market constraints, requires the use of advanced software and hardware development tools.

This paper discusses two important tools useful to designers of complex digital systems. The first is Hewlett-Packard's HP 64700 Microprocessor Emulation System. The second tool is comprised of Quickturn's Enterprise and RPM ASIC Emulation Systems.

The HP 64700 provides in-circuit emulation of commercial microprocessors. In this role it serves as an advanced Hardware-Software development and integration platform. The various features and options of the HP 64700 are discussed, along with the emulator's role in assisting system design and integration.

Quickturn System's ASIC Emulation systems provide complete in-circuit emulation of proprietary ASICs. In addition to complete in-circuit functionality, these systems provide access to all internal circuit nodes of the ASIC being emulated. The architecture, options, capabilities and applications of these versatile instruments is discussed. Case study data demonstrates "real-life" applications.

Authors

Tim Chambers

Current Activities:

Tim is currently with Hewlett-Packard as a Sales Development Engineer working at HP's Colorado Springs Division. He is responsible for sales development of HP's Microprocessor Development Tools.

Author Background:

Tim has worked at HP for 11 years. Prior activities with HP included research and development of debugger simulators supporting the Motorola family of Microprocessors. Tim graduated from M.I.T. where he earned a degree in Computer Science and Engineering.

Eric Amador

Current Activities:

Eric is President of TESLA Corporation. He is involved in advanced system design consulting activities for TESLA Corporation's clients. In addition to managing a small team of highly skilled consultants, Eric's principal design responsibilities include development of high-performance VLSI ASIC's and the complex board-level assemblies in which these ASIC's operate. TESLA Corporation specializes in ASIC emulation, and advanced system simulation technologies.

Author Background:

Eric has been an independent system design consultant for the

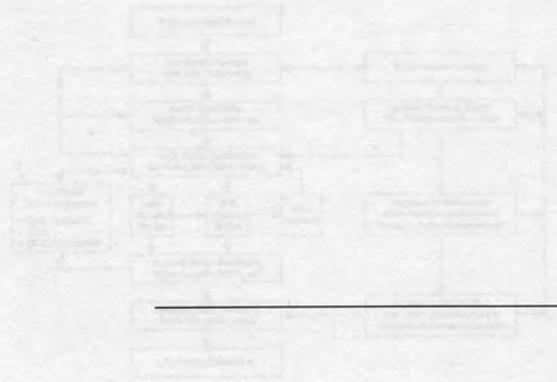
past 16 years. During this time he has consulted on numerous projects for major corporations. He is named inventor on several US and foreign patents. Before becoming a consultant, Eric was Engineering Manager at Arnold Magnetics Corporation where he was involved in the design and manufacture of high-density switching power supplies. Eric earned his Bachelor's degree in Biochemistry at Michigan State University.



Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Hewlett-Packard
Quickturn Systems
TESLA Corporation

Simplified HS Digital Design Process



Authors (Cont'd)

Naeem Zafar

Current Activities:

Naeem is Director of Product Marketing at Quickturn Systems, Inc. Naeem is responsible for all product marketing activities including new product definition. Naeem's responsibilities include the RPM Emulation system and the Enterprise Emulation system. Naeem is actively involved in promoting Quickturn's strategy in the emerging re-programmable hardware market.

Author Background:

Naeem has 12 years experience in the electronic design and CAE industries. Naeem designed VLSI chips and worked in the area of computer architecture at Honeywell, Inc. Naeem was founder of a company, XCAT, which developed hardware accelerators for logic and fault simulation. He joined Quickturn Systems in 1988 and has held positions in engineering, technical marketing and product marketing. Naeem has a Sc.B. degree in Electrical Engineering from Brown University and an MSEE from the University of Minnesota.

As shown here, in a complex system, the various engineering tasks in both hardware and software design have an iterative nature. The "top-level" aspects of a project are closely related to "bottom-level" activities. In order to save costs and meet project schedule constraints it is very important to minimize iterations through these "top-to-bottom" paths. Successful management of the complex interactions among these system design and production activities is greatly improved through the use of the advanced design methodology implemented in Microprocessor and ASIC emulators.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #1

**Optimizing Your Design Flow...
How to Use Microprocessor and
ASIC Emulators**



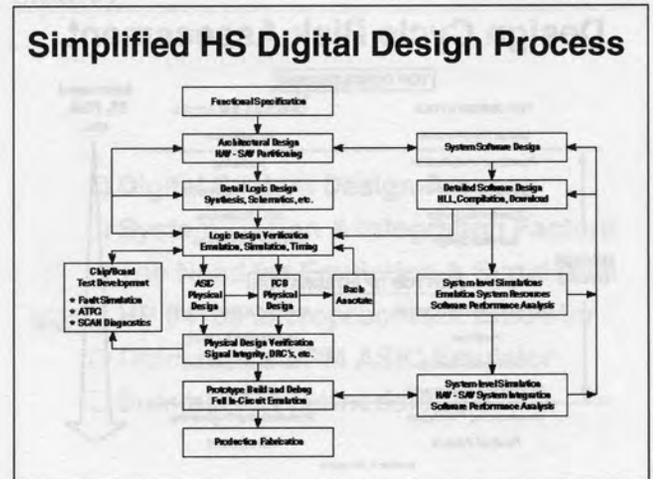
**Hewlett-Packard
Quickturn Systems
TESLA Corporation**

Slide #2

Agenda

- Digital System Design Process
- System Design & Integration Factors
- The Need for Emulation & Simulation
- HP 64700 Microprocessor Emulator
- Quickturn's RPM ASIC Emulator
- Summary of Methodologies

Slide #3



As shown here, in a complex system, the various engineering tasks in both hardware and software design have an iterative nature. The “top-level” aspects of a project are closely related to “bottom-level” activities. In order to save costs and meet project schedule constraints it is very important to minimize iterations through these “top-to-bottom” paths. Successful management of the complex interactions among these system design and integration activities is greatly improved through the use of the advanced design methodologies implemented in Microprocessor and ASIC emulators.



**QUICKTURN
SYSTEMS**



**TESLA
CORPORATION**



**HEWLETT
PACKARD**

Optimizing Your Design Flow...

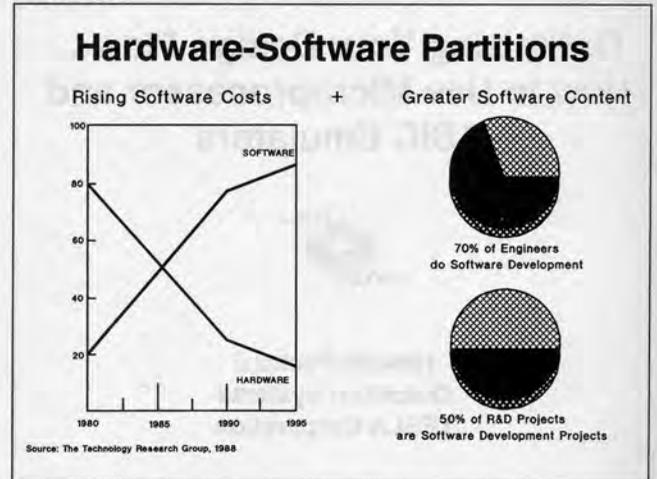
How to Use Microprocessor and ASIC Emulators

Slide #4



This diagram describes a "typical" design cycle and its associated check & balances. The "Increased Costs and Risk" along the right side of the diagram. Risk is illustrated by considering the effect on project costs which is caused by excessive iterations in the area defined as "Design Cycle". The role played by the Checks & Balances column in this description of the design cycle process is one which associates Design Tasks with Personnel Functions. That is to say, we are able to observe how personnel from various departments (Marketing, Software Engineering, etc.) must interact with each other as an automatic result of the performance of detailed design tasks (Logic Simulations, Diagnostics & Debug, etc.) This diagram serves to make clear the complex nature of present day system design. Excessive time spent in iterations of the "Design Cycle" dramatically increase the cost of a project and consequently place the success of the project at "Increased Risk".

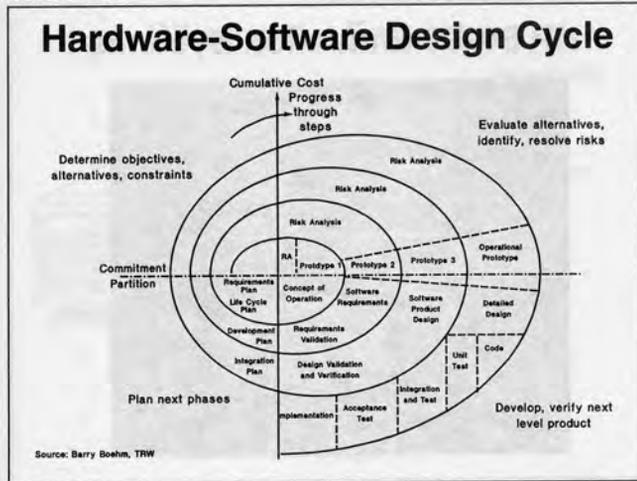
Slide #5



The cost allocations between Hardware & Software are shown here. During the past decade software costs rose dramatically while hardware costs declined proportionately. Of course, overall total project costs have increased significantly as well, giving added importance to the need for timely project completion. Two additional design relationships are described here in the form of pie-charts. The lower pie-chart shows the increased emphasis on software-only projects. The upper pie-chart shows the distribution of manpower between Software Development and Hardware Development engineers. As indicated, in 1988 70% of Engineers were involved in Software Development. Keeping software development costs and schedules under control is of paramount importance to the success of a complex digital system project. Of equal importance, hardware engineers can benefit significantly from the use of tools which assist them in their work with software engineers. In this regard, HP's 64700 Microprocessor Development Tools provide hardware and software development engineers with highly efficient and powerful hardware/software design and debug capabilities.

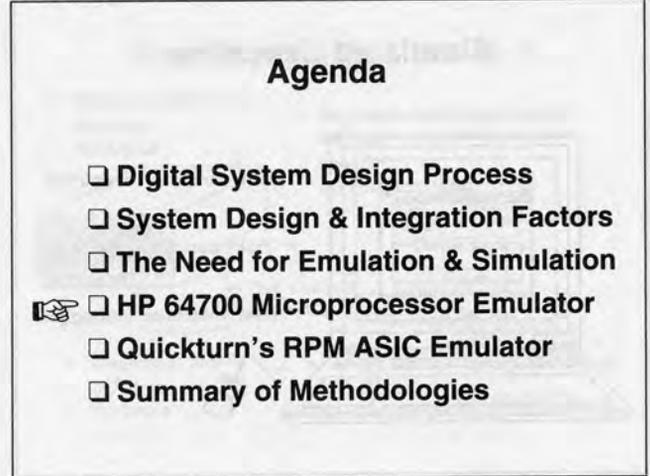
Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #6

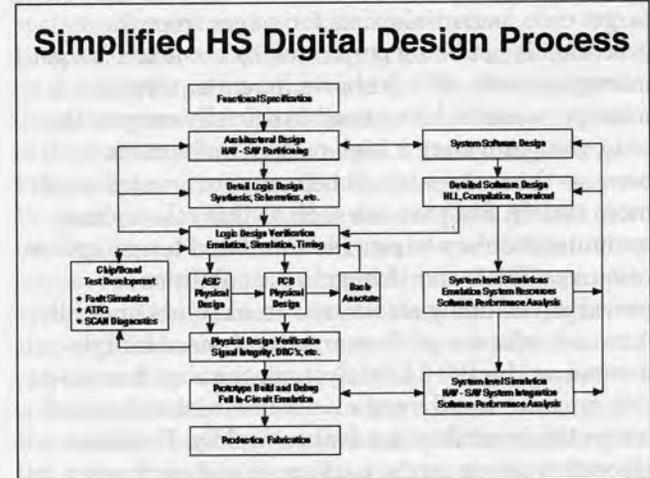


This diagram illustrates in a different fashion ideas similar to those described by the "Design Cycle Risk Assessment" slide. This representation accentuates the iterative nature of the design process. A typical system design begins at the center of the spiral with a Requirements Plan. The system progresses through Prototype 1 and Concept of Operation phases. This leads to a Life Cycle Plan. As shown here, at each iteration Risk Analysis is being performed. In our previous slide, the Risk Analysis was shown as "Checks & Balances"; here it is defined as an overall management function which takes place at various points in the lifecycle of a project. This chart emphasizes the planning phases of a major software project. As shown here, there are three Prototypes leading to the development of an Operational Prototype and a Detailed Design Specification. The subsequent phases following the Detailed Design are shown as ending in a finished Implementation of the Software System.

Slide #7



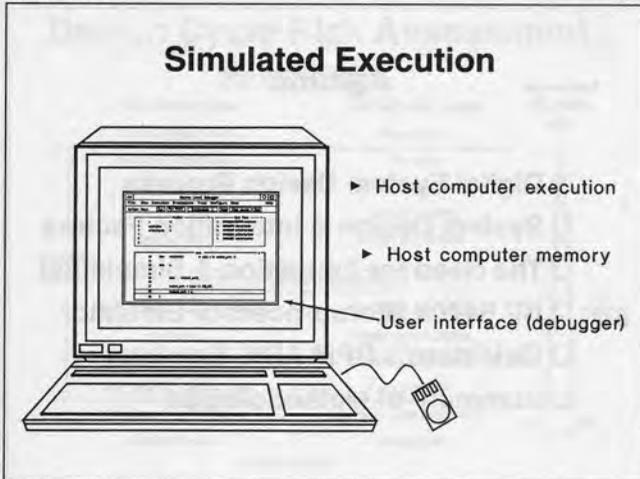
Slide #8



The HP 64700 Microprocessor Emulation system serves a valuable role in system design. As shown here, starting at the Hardware/Software partitioning function. The HP 64700 emulator plays a vital role in the design cycle all the way through the final stages of prototype verification. The emulator's real-time hardware/software analysis capabilities make it an indispensable tool initially during hardware debug and later during hardware/software integration.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

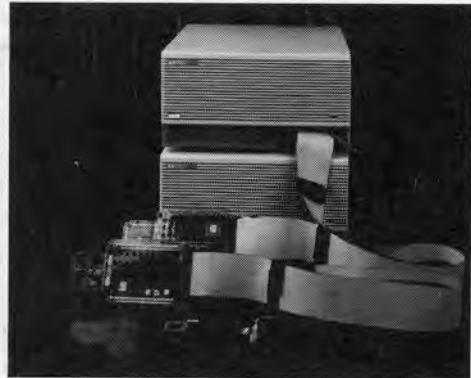
Slide #9



This slide shows an HP workstation being used to display the simulated execution of system software. This type of software simulation can also be performed using a remote host. In a case where a remote host is used to compile, run and debug the target code, potential exists for errors since the execution is not being performed by the actual target microprocessor. As it is shown here, the target microprocessor is being used to actually execute the code, thus providing a high measure of accuracy because the real system is being approximated much more closely. Simulations such as this rely on host computer memory to provide simulated target system resources. Extensive debugging capabilities are provided, including structured breakpoints and full-featured software performance analysis. As implemented on the HP 64000 System, simulated execution provides **hardware and** software development teams the capability to advance rapidly. Extensive algorithm debug can be performed without having to wait for target hardware prototypes. Once target hardware is available, the same environment can be incrementally mapped over during initial debug. When completed, execution of the target software via the Emulator permits exhaustive analysis and optimization of the target system.

Slide #10

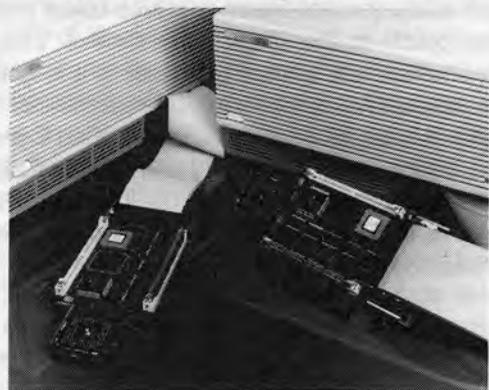
HP 64700 Emulation Systems



Two HP 64700 Emulator Systems are shown here. Their in-circuit pods are displayed, as is a small test-case printed circuit board. Comprehensive cross-triggering capabilities permit these instrument to be used in conjunction with each other, enabling hardware and software development of multi-processor designs.

Slide #11

HP 64700 Emulator Probes

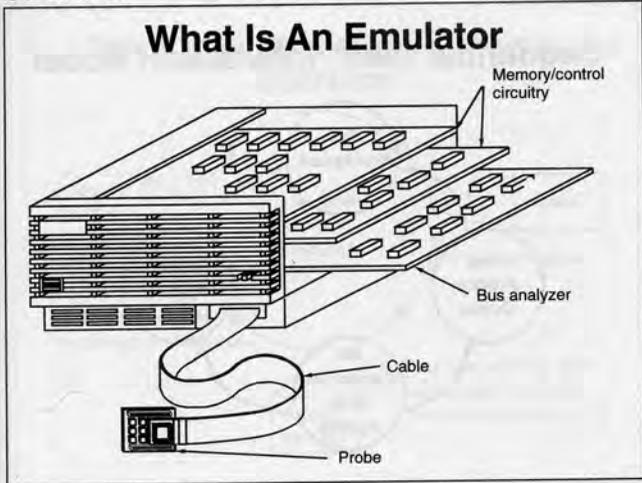


This photograph shows a close-up view of two HP 64700 in-circuit microprocessor pods. Real-time in-circuit performance is supported by the close proximity of the microprocessor device on the probe card. The ribbon cables connect the probes to their respective HP 64700 internal instrumentation cards. The "tip" of these probe cards provide a PGA plug which connects the emulator pod to the microprocessor socket(s) on the target hardware circuit boards.

Optimizing Your Design Flow...

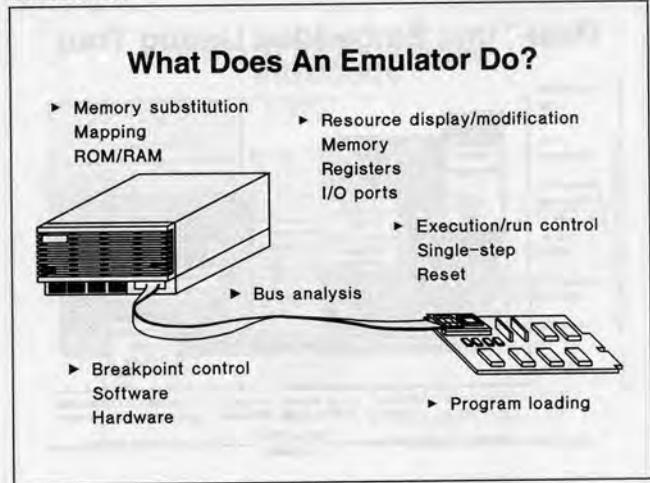
How to Use Microprocessor and ASIC Emulators

Slide #12



The architecture of HP's 64000 System is shown in this slide. As mentioned in the previous slide, the emulator permits the same code to be initially simulated and then later run in the target hardware. No changes to the user interface or tools are required. The probe assembly can be plugged into the target hardware system's microprocessor socket when hardware prototypes become available. Prior to a system prototype becoming available, the probe assembly can be used to execute code which is resident in memory provided by the HP 64700 emulator. The Bus Analyzer card permits the emulator to display cycle timing information, as well as other types of measurements useful during debug. Although not shown in this slide, the emulator chassis can hold additional options such as a State/Timing Analyzer which permits independent probing of target hardware. Another option, the Software Performance Analyzer (SPA) can perform detailed real-time system performance measurements, which are often valuable in the course of real-time system hardware design and debug.

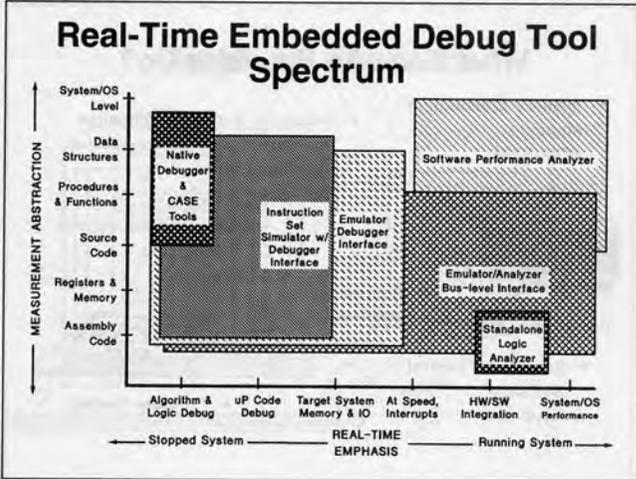
Slide #13



This slide depicts an overview of emulator functionality. The HP 64700 emulator provides memory resources which can be mapped so as to simulate target ROM or RAM. The mapping scheme is very flexible, thus allowing a hardware designer to rapidly change system memory configurations as required. Internal resources of the target microprocessor can be displayed and modified readily. This feature gives design engineers valuable insight into the operation of their selected microprocessor. Internal registers, I/O ports, memory, etc. can also be accessed. Program execution is under the complete control of the designer. Powerful breakpoint capabilities support parameters based on address, data, type of cycle, etc. The user can control execution in single-step and other modes. Of course, the microprocessor can be readily reset at any time. Loading of program code is extremely fast due to the real-time nature of the emulator itself. Complete analysis capabilities of the microprocessor's bus activity are provided directly to the hardware engineer via the probe's resources. Additional logic and timing analysis is supported via Logic Analyzer probes. In this fashion, both software and hardware environments can be fully observed and controlled during system development.

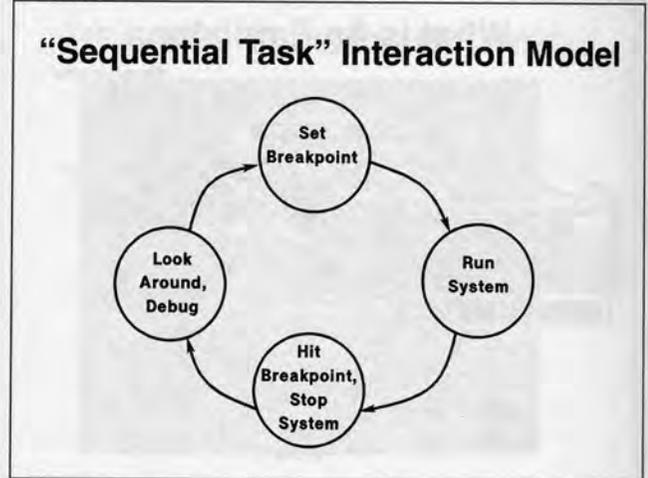
Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #14



Shown here are the various resources provided by HP's 64000 Microprocessor Emulation System. The vertical axis of this graph is used to represent Software Measurements, characterized by their "Level of Abstraction". The software level of abstraction is a measure of "closeness" to the execution of individual instructions. Thus, the Operating System itself (System/OS) is at the top and hence the most abstract element, while individual instruction are at the bottom and represent fundamental "primitive elements" of the target microprocessor. The horizontal axis represents the dependence of the target system on real-time issues. Overall performance is judged and optimized by adjusting the various elements which are placed along this axis. At the origin we show individual algorithms as the "primitive elements" of real-time algorithms as the "primitive elements" of real-time performance. At the extreme right of the horizontal axis we again show the System/OS performance, since it incorporates the elements of complex real-time system behavior.

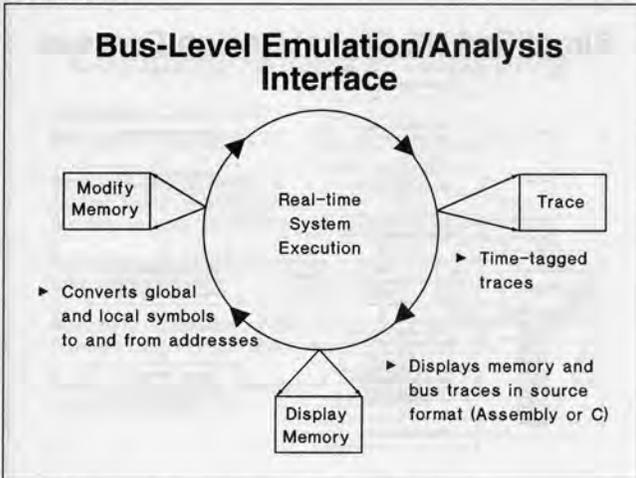
Slide #15



This slide graphically describes a "Sequential Task" debug strategy in which the user sets a breakpoint then starts execution of system code. Upon encountering the breakpoint, the system is halted and the user proceeds to examine the state of the system. The user can inspect registers, memory locations, etc. If the user uncovers a bug as a result of hitting the breakpoint, he proceeds to debug the system code or hardware, or both. The cycle is then repeated with the same or different breakpoints. This style of debug is sequential and, for the most part, not sufficient in the debug of complex systems. One of its disadvantages is the lack of real-time system monitoring. There are other disadvantages as well, including the lack of control and visibility over system hardware, etc. This strategy should be used only for simple debug scenarios where a single, static breakpoint will suffice.

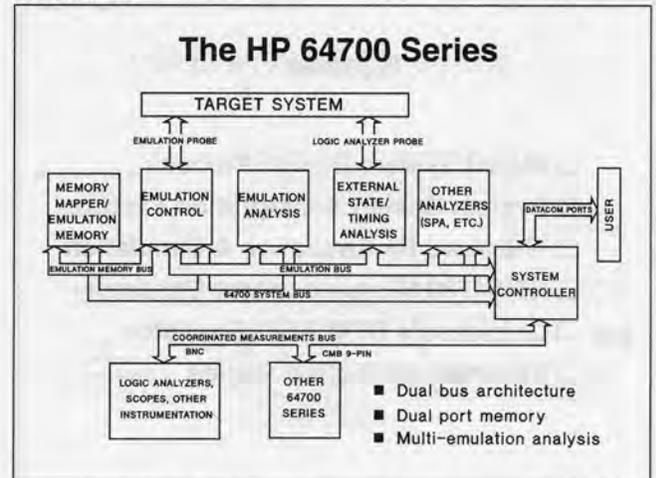
Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #16



This slide shows a "Parallel Task" debug strategy which makes full use of the Emulator's resources. At the center of this slide we show the target software as it executes code in real-time. Large amounts of data can be gathered and processed by the Software and Bus Analyzers. The user thus can derive useful conclusions regarding the operation of complex system software. In addition, this strategy provides time-tagged software execution traces permitting detailed timing analysis of software execution. Advanced high-level debugging facilities include global and local symbol mapping, as well as complete control over system memory. When combined with a Logic Analyzer, events in the hardware can be displayed correlated to microprocessor code execution. Triggering capabilities are synchronous across both Logic and Software Analyzers. Source level software debugging provides advanced trace capabilities which make source code debug possible in either Assembly or C. This slide summarizes state-of-the-art embedded controller debug techniques.

Slide #17



The HP 64700 Series architecture is described in this figure. The various functional assemblies are labelled as Memory Mapper/Emulation Memory, Emulation Controller, Emulation Analyzer, External State/Timing Analyzer, Software Performance Analyzer, etc. The HP 64700 System Controller, shown to the right, provides the user and external interfaces. Shown at the top is the Target System. Both Emulation and Logic Analyzer Probes are shown connecting to the Target System. Finally, along the bottom we have ancillary debug instruments which can communicate with the Microprocessor Emulation System. These instruments include external Logic Analyzers, Oscilloscopes, and other laboratory instruments. It is important to note that design and debug of systems comprised of multiple processor systems can be supported by using additional 64700 Systems. The HP 64700 contains an internal dual-bus architecture. The microprocessor emulator itself is controlled and analyzed using a high-speed Emulation Bus. This Emulation Bus is completely dedicated to providing real-time support of the Microprocessor Emulator and the various Analyzers. The second bus is the 64700 System Bus. This bus is used for slower non-real-time system operations such as downloading code or recovering trace information from the Analyzers for post-processing and display to the user. The combination of these two busses provides the high-bandwidth data processing capabilities required to sustain real-time system debug. A third bus, called the Coordinated Measurements Bus provides synchronous and state information useful when external instrumentation is connected to the HP 64700 System. Complex cross-triggering among the various internal Analyzers and/or external instruments is possible. Extensive timing and performance parameters permit extremely flexible instrumentation setups. The overall result is that users can readily perform complex debug tasks and accurately measure system performance.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

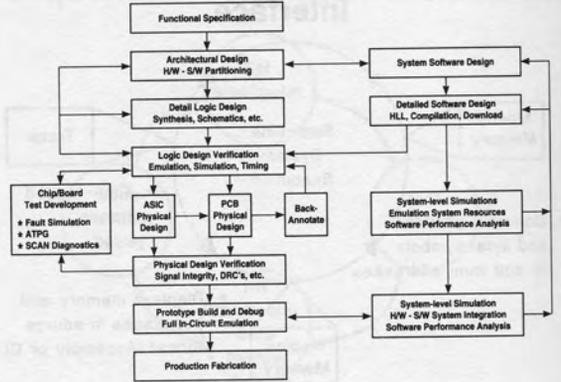
Slide #18

Agenda

- Digital System Design Process
- System Design & Integration Factors
- The Need for Emulation & Simulation
- HP 64700 Microprocessor Emulator
- Quickturn's RPM ASIC Emulator
- Summary of Methodologies

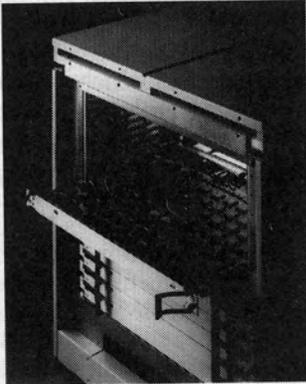
Slide #20

Simplified HS Digital Design Process



Slide #19

RPM Emulation System



QUICKTURN
SYSTEMS



TESLA
CORPORATION



HEWLETT
PACKARD

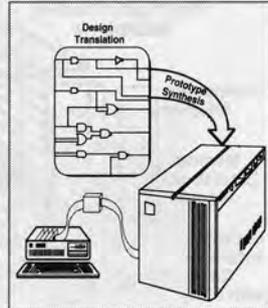
Optimizing Your Design Flow...

How to Use Microprocessor and ASIC Emulators

Slide #21

Computer-Aided Prototyping

- *Automatic Creation of Hardware Prototypes ...*
- Read any design Netlist & automatically partition, place & route on array of FPGAs
- Map memory & logic elements automatically
- Access to all internal nodes through built-in logic analyzer
- Ability to make incremental design changes



Computer-aided prototyping (CAP) is a methodology for achieving system design verification. It provides for early system integration and concurrent design verification. CAP includes CAE interface software, design-to-prototype translation software, timing analysis software, reprogrammable hardware (including logic and memory), debug instrumentation and in-circuit interface cabling.

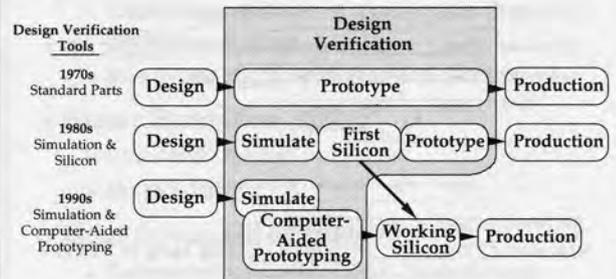
Interfaces to all popular design environments with netlists generated from schematics or synthesis when designs are captured in VHDL.

CAP provides design prototypes in hardware that can be verified in the system being designed where all software, hardware and interfaces can be thoroughly debugged before committing chips to silicon fabrication. It provides an environment for making hardware and software trade-offs, optimizing architectural approaches, and the development of more complete diagnostics software.

CAP prototypes are real hardware prototypes that use real gates and real wires. The implementation technology is reprogrammable logic using FPGAs. All interconnect between gates, interconnect to debug instrumentation, and interconnect to in-circuit interfaces is completed automatically with reprogrammable interconnect. In-circuit adapters plug directly into sockets where the ASIC chip will go when it has been fabricated in silicon.

Slide #22

Computer-Aided Prototyping



All systems are prototyped whether there is a plan to do so or not. The first article is always a prototype. Unfortunately most system prototypes are done with ASICs in silicon where debug and design changes are difficult and time consuming.

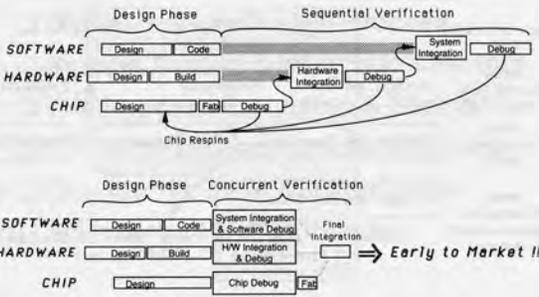
CAP allows the first article of an entire system to be prototyped before fabricating silicon so that design verification and debug can be completed in an environment where all internal nodes of ASICs can be observed, controlled, and changed quickly and conveniently.

CAP moves the first article of the system into the domain of design verification where it belongs and reduces the total development time, produces higher quality products, and reduces risk of schedule slips.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #23

Benefits of Computer-Aided Prototyping

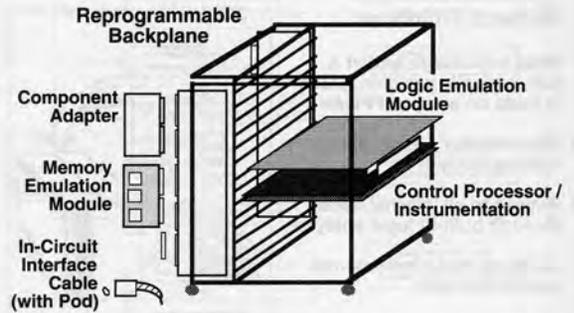


Various pieces of the product are designed by the user separately but they all come together at one point. Usually boards cannot be assembled until the chips are verified and the whole system verification has to wait for the hardware to be fully functional. This traditional approach is inflexible. Often problems are found very late in the development cycle and are not always fixable without changing hardware, compromising software or worse, re-spinning the chips.

CAP provides a way to do product integration much earlier in this cycle. All pieces of the system can be brought together and verified before ever fabricating chips. Silicon is made only after all pieces of software and hardware have been verified together. This results in much shorter system integration phase, higher quality due to fewer compromises and being earlier to market.

Slide #24

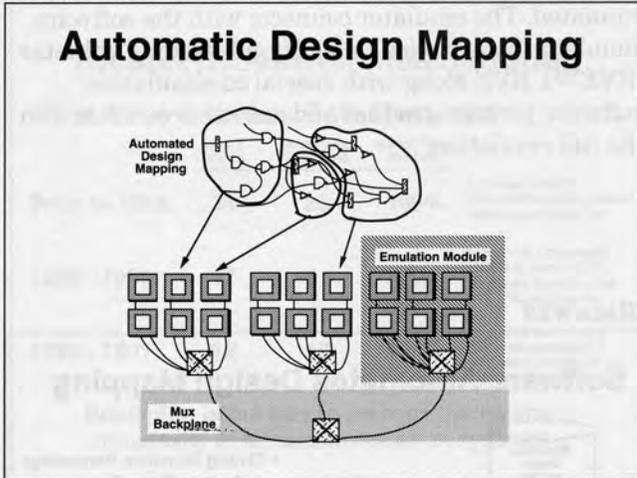
Anatomy Of An Emulator



The emulator accepts 11 cards, called Emulation Modules. Each emulation module has a usable gate capacity of 30K gates. There is a Control Processor card which is used to communicate to the host workstation and contains a 1024 channel logic analyzer. The interactive backplane provides electronic switching of signals among these Emulation Modules. Memory emulation cards are plugged into the backplane. These cards are automatically personalized by the mapping software and can emulate up to 2MByte of RAM per card, including multiport RAM (up to 32 port wide). User can plug up to 32 Memory Emulation Boards in a system. User has access to over 6000 I/O signals. Component Adapter cards allow users to place existing ICs and plug them into the emulator backplane. In-Circuit interfaces provide the connection to the user's target system.

Designs are automatically read in from just about any format including Verilog, EDIF and ASIC vendor formats. They are automatically partitioned and mapped on to the emulation boards. A graphical user interface gives user's complete access to the internals of their designs.

Slide #25



Designs are automatically read in from just about any format including Verilog, EDIF and ASIC vendor formats. They are automatically partitioned and mapped on to the emulation boards. A graphical user interface gives user's complete access to the internals of their designs.

A portion of the user's netlist is assigned to an FPGA and all FPGAs are connected via a custom interconnect chip (MIC) which acts as a cross-bar. Then all emulation modules can be connected electronically using another rank of MIC chips. This provides an expandable architecture where user can add more capacity modularly as his needs grow.

Slide #26

Enterprise Emulation System

- **Higher Capacity System**
 - » 330K Gates, 64MByte of Memory per Enterprise Emulation System
 - » 6M+ Gates per Multi-Enterprise Cluster
- **Faster Emulation Speed**
 - » 2 to 4 MHz typical
 - » Up to 8 MHz
- **Most Advanced Software**
 - » Connected with Major CAE Environment
 - » Most Efficient Mapping of Design

Enterprise emulation system is the next generation emulation system. Based on custom interconnect devices, latest FPGAs and a patented new Hierarchical Multiplexed Architecture this system delivers emulation for ASICs and custom ICs from 30K gate to 330K gates. Memory emulation is handled just as easily through special reprogrammable cards called Memory Emulation Modules. Enterprise emulation system connects to many CAE systems. It allows co-simulating mapped design from your software simulation environment before you may plug it in-circuit for system verification at hardware speeds.



QUICKTURN SYSTEMS



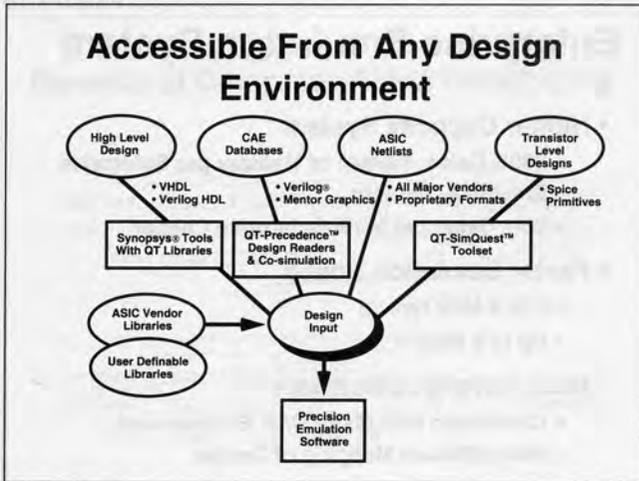
TESLA CORPORATION



HEWLETT PACKARD

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

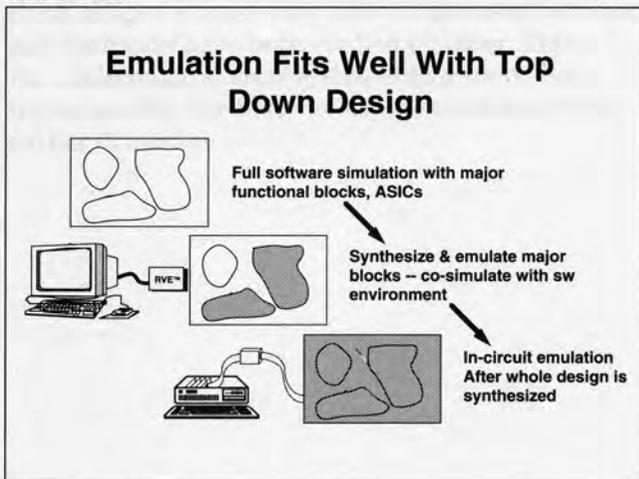
Slide #27



A rich set of design readers allow users to read-in a design from just about any design environment. Many translators allow users to read in design in SPICE, Verilog® or Mentor Graphics databases or just about any ASIC vendor format. Designs from VHDL can be synthesized using Synopsys™ emulation library for optimal synthesis for emulation.

Many ASIC vendor libraries are supported. User can develop their own libraries using the Library Development Kit for proprietary cells.

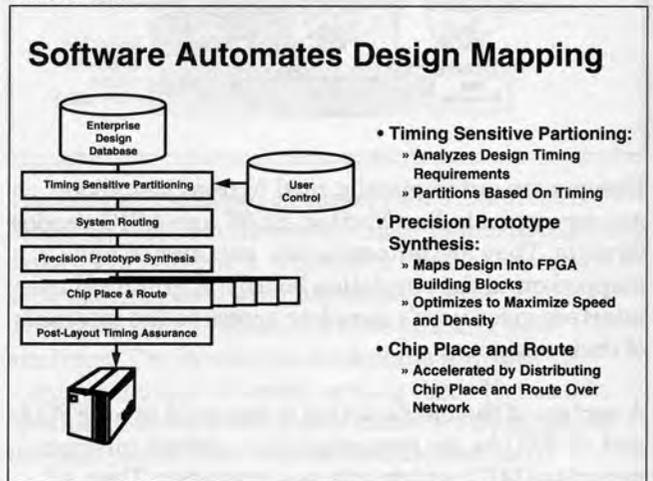
Slide #28



As part of the top-down design methodology users do a complete behavior simulation first. They synthesize block at a time and as a portion of the design

becomes available in gate level netlist, it can be emulated. The emulator connects with the software simulator via a device called Rapid Vector Evaluator (RVE™). RVE along with special co-simulation software permits gradual and smooth transition into the full emulation.

Slide #29



Precision Emulation Software™ automatically maps designs read-in by first partitioning into logic blocks which will fit into an FPGA. Once the design is partitioned, using timing driven techniques, logic cells are routed among FPGAs using system router through custom interconnect chip. The FPGA place and route can be spread among multiple workstation on the network. Once place and route is completed, a comprehensive post-layout static timing analysis is run to assure full functional equivalence. Software calculates emulation speed, identifies and fixes any potential hold violations and allows users to time certain asynchronous paths.



QUICKTURN
SYSTEMS



TESLA
CORPORATION



HEWLETT
PACKARD

Slide #30

Alcatel (Rockwell International)

First Time Success Rate Restored

	Total # ASICs	# Worked First Time	% First Time Success	
Prior to 1986:	36	36	100%	< 15K Gates/ASIC Well-Defined Specification Gate-Level Simulation
1986 - 1989:	17	4	23%	Increased Complexity (20K-50K gates/ASIC) Unproven Specification RTL + Gate-Level Sim.
1990 - 1991:	12	12	100%	Increased Complexity (30K-70K gates/ASIC) RTL, Gate Simulation + Quickturn Emulation

Emulation is the way to manage increasing complexity in telecommunication systems

Rockwell International— now Alcatel— were one of the first users of Computer-aided prototyping. They realized that as chips got above 20K gates the test vectors alone could not be relied up comprehensive testing of designs. The first silicon failure rate decreased significantly. New design methodology, including emulation, changed those statistics for a 100% success even after the designs grew over 50K gates.

Slide #31

5370 Superminicomputer Project Overview

New Generation 50 Series Dual Processor

- Same Architecture as the 6650 ECL Processor
- Uses CPU ASIC Components from 5340 CMOS Processor

Very Complex CISC Architecture

- Dual CPU with Crossbar Switches for Memory & I/O
- New Memory/Cache Controller

Required Development of Multiple New ASICs

Prime Computer of Natick, Mass. used CAP from Quickturn Systems to develop a new generation of superminicomputers, the 5370. The 5370 replaced a previous generation computer. ASICs were used to reduce the number of boards in the system and several new features were added including multi-processing.

The goal was to put all the logic that once required 8 boards onto a single board by integrating much of the logic into 50K CMOS gate arrays. The system achieved every objective and outperformed the previous generation computer by 25% to 41%.



**QUICKTURN
SYSTEMS**



**TESLA
CORPORATION**

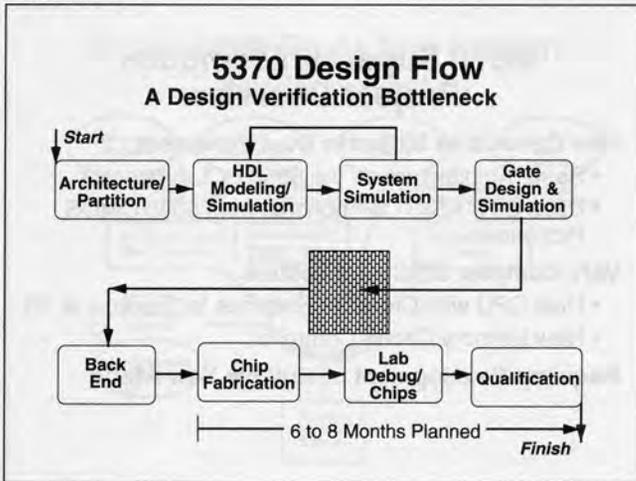


**HEWLETT
PACKARD**

Optimizing Your Design Flow...

How to Use Microprocessor and ASIC Emulators

Slide #32



Prime's design flow started with architectural partitioning and a description of the design in HDL. The design system was simulated with an HDL simulator and when the design was relatively stable the ASIC portion was recaptured with schematic capture at the gate level. Then gate-level simulation was performed using vectors captured during system simulation.

This design flow was an advance from previous flows by the addition of HDL system simulation because it reduced the time required to complete the total simulation task from the previous flow which used only gate-level simulation. Unfortunately Prime calculated that it would take thousands of years to simulate the equivalent of 30 minutes of diagnostics software running in real-time.

The original schedule for the time it would take to ship systems to customers after they sent ASIC data bases to fabrication was 6 to 8 months. This process included prototype chip fabrication, production chip ramp up, software check out, and final system Q/A and test.

Slide #33

Design Verification Bottleneck

Microdiagnostics Were Simulated

- Simulations Began to Take Multiple Days
- Clock Cycles Took 6 Seconds on Sun 4/330
- Diagnostics Would Have Taken 7600 Years

Simulation Alternatives

- More Compute Power
- Custom Simulation Language
- Hardware Acceleration
- Hardware Modeling
- FPGA Breadboards
- Computer-Aided Prototyping

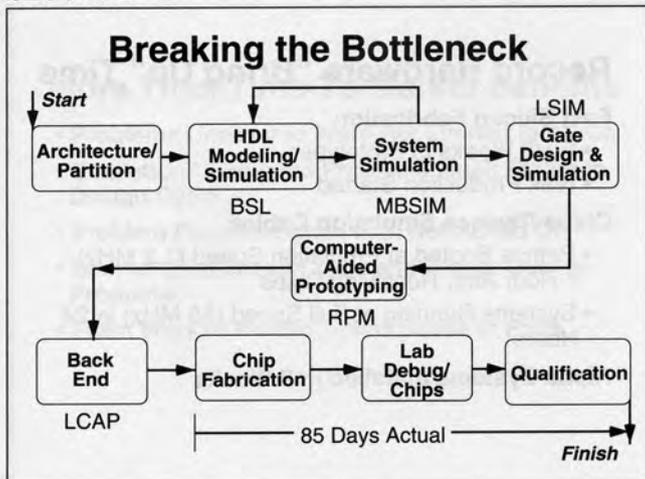
Prime's development strategy was to simulate portions of the diagnostic programs (called microdiagnostics) that were used with the previous generation of super-minicomputer and when the diagnostics executed correctly the system design including the ASICs would be known to be correct.

Simulations of the microdiagnostics began to take multiple days to complete. Prime knew it had to evaluate alternative approaches to design verification if it were to avoid major schedule slips.

Prime evaluated several alternative ways to accelerate the design verification process. They considered more compute power, a new simulation language, hardware accelerators for their simulators, the addition of hardware modeling to their system simulator, building FPGA breadboards manually, and CAP with Quickturn's RPM Emulation System.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #34



Prime selected CAP as the only viable way to achieve their product and schedule goals. They built a full system prototype using CAP for the ASICs. The design verification strategy was enhanced with CAP to execute the full diagnostics suite before fabricating silicon.

The design was simulated at the HDL-level, then the gate-level version of the design was recaptured and simulated at the gate-level to assure correlation with the HDL-level. Then a system prototype was built, the gate-level netlists were loaded into the RPM and plugged into the system in the sockets where the ASICs would eventually go when they were fabricated in silicon.

Slide #35

Breaking the Bottleneck

Timing Analysis

- Done in Parallel with Continuing Emulation Debug
- Used MDE LSIM and LCAP

Timing Problems Discovered Requiring Functional Changes

- All Functional Changes Reverified with Emulation
- ASIC Releases Held Several Times When Problems Discovered

Timing analysis was run in parallel with CAP because CAP is a functional verification methodology and is ASIC technology independent. Timing analysis uncovered several flaws requiring functional changes to the design. All functional changes to the design were reverified with CAP before designs were released for silicon fabrication.

The net result was a reduction of many months in the delivery time to customers. The delivery time was actually cut by 6 months compared to a previous project of lesser complexity.



QUICKTURN
SYSTEMS



TESLA
CORPORATION



HEWLETT
PACKARD

Slide #36

Breaking the Bottleneck

Run PXIO Assembly Language Diagnostics

- Emulation Took 12 Hours;
Real-Time Takes 1/2 Hour
- Several "Show Stopper" Bugs Discovered

Boot Primos Operating System

- Not Part of Emulation Debug Test Plan
- Uncovered Further Hardware "Show Stopper" Bugs
- Uncovered Many Microcode and System Software Bugs

Saved Multiple Respins

The full diagnostics suite was run on CAP and Prime was ready to send out for silicon prototypes. The diagnostics took about 12 hours with system clocks running at approximately 1.2 MHz. These diagnostics take about 1/2 hour in real-time. During the time the diagnostics were run several show stopper bugs were discovered that would have required silicon to be respun had they not been caught ... and simulation would not have caught them.

Prime had met its objective of executing full diagnostics before fabricating silicon and so they went ahead and sent data bases to LSI Logic for silicon fabrication.

It was not part of the plan but they invited the operating system people to run the Prime operating system on the test bed (system Prototype) while silicon prototypes were being fabricated. Much to their surprise they discovered additional show stopper bugs the diagnostics had not uncovered.

They stopped the silicon prototype fabrication process and continued to run the operating system software and found more bugs.

In the final analysis several silicon prototype respins were saved.

Slide #37

Record Hardware "Bring Up" Time

Fast Silicon Fabrication

- 2 1/2 Weeks to Prototypes
- Risk Production Started

Chips Replace Emulation Cables

- Primos Booted at Emulation Speed (1.2 MHz)
1 Hour After Receiving Chips
- Systems Running at Full Speed (55 MHz) in 24 Hours

Alpha Systems Installed in 2 Weeks

Prime achieved a record hardware "bring up" time after releasing the final designs for silicon fabrication. They received silicon prototypes 2 and 1/2 weeks after releasing design data bases. They unplugged the RPM Emulation Systems, plugged in the chips, and within 1 hour after receiving the prototype chip the system was running at 1.2 MHz.

Within 24 hours after the chips were received the system was running at the full design target design speed of 55 MHz and alpha sites were installed within 2 weeks.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #38

More Than Time-To-Market Benefits

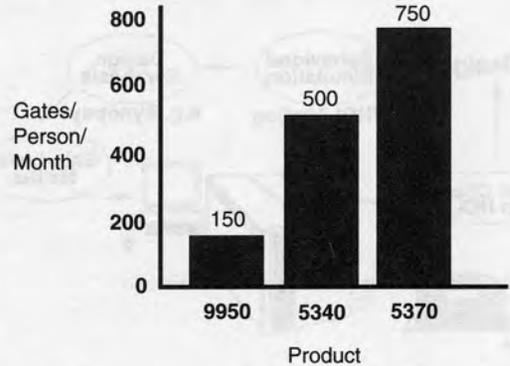
- Problems Uncovered Were Not Limited to ASICs
- Emulation Uncovered Problems Much Earlier In Design Cycle
- Problem Fixes Designed In ... Not Added On
- Shorter Qualification Time and Higher Quality Products!
- Team Worked Harder ... New Sense of Pride

Prime's experience demonstrates that CAP enhances the system design verification environment in ways beyond ASIC chip verification.

- Over 50 board-level bugs were discovered by debugging with CAP in the system that had nothing to do with the ASIC designs themselves.
- Bugs were discovered much earlier in the development cycle than would otherwise have been possible with simulation alone.
- Problems within the ASICs were discovered before chips were fabricated so the fixes were made by correcting the chip designs instead of making patches external to the chip or in software.
- The product qualification time was dramatically reduced because the complete product was so thoroughly verified before the qualification process was even started.
- The product quality was enhanced because the diagnostics programs that would be shipped with the product were refined for test coverage and changes to improve testability were made to the chip designs.
- Finally the team seemed to have a renewed sense of enthusiasm when they could get hands on the real system with CAP.

Slide #39

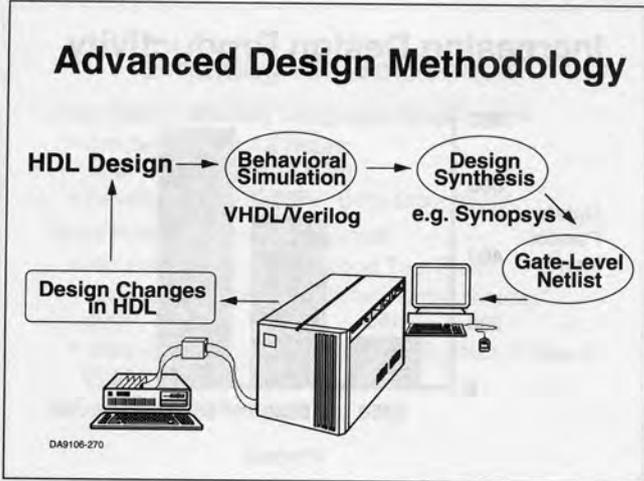
Increasing Design Productivity



Design productivity shot up dramatically on the 5370 project. They designed more gates in less time than the previous project such that the average number of gates per designer per month increased by 50%.

Optimizing Your Design Flow... How to Use Microprocessor and ASIC Emulators

Slide #40

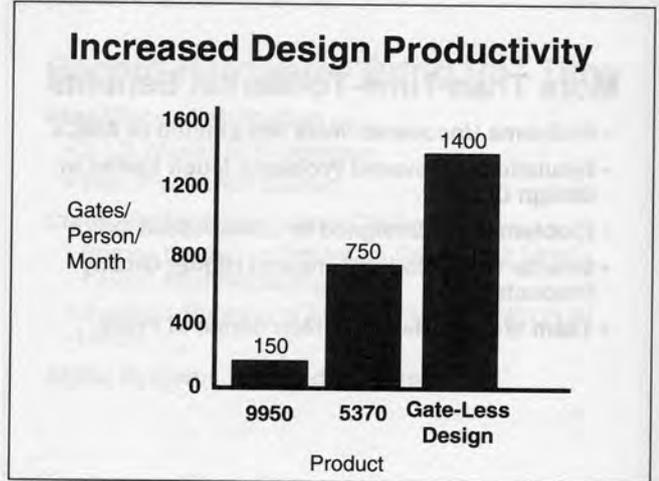


This advanced design methodology is getting to be very popular among the designers of high performance digital systems. A collection of advanced tools is resulting in increased productivity for the designer and allowing them to create and deliver more complex products in record time.

Using HDL to input design and doing behavioral simulation gives early affirmation of the design. Users can then synthesize the design and is able to see the design working in real hardware using the Computer-Aided Prototyping. Users debug the system with CAP and are able to modify designs as they debug, with a very short turn-around cycle.

Using CAP for comprehensive functional verification gives user more time to do timing analysis or timing simulation of their design using the conventional tools.

Slide #41



Prime expects to achieve another dramatic improvement in designer productivity by going to the gate-less design methodology. They completely eliminate gate-level design. This saves the time required to recapture the design at the gate-level and the time required to correlate it with the HDL representation.

Slide #42

Feature Summary

	HP 64700 MicroProc. Emulator	Quickturn RPM ASIC Emulator
Commercial MicroProcessors	Real-time S/W & H/W Support	
Commercial VLSI Devices		Limited Capabilities
Proprietary VLSI Devices		Near Real-time S/W & H/W Support
Internal Register Visibility	Complete Access	Complete Access
Software Performance Analyzer	Real-time S/W & H/W Support	
Logic Analysis Capability	Comprehensive Feature Set	Comprehensive Feature Set
VLSI Design Modifications		Incremental Netlist Changes
Software Design Modifications	Complete S/W Tools Available	
Scope of Support	Mainstream Microprocessors	Major ASIC Vendors Fully Supported

Hewlett-Packard, c/o: Dave Bellandi, 1900 Garden of the Gods Road, Colorado Springs, CO 80907



Developing and Debugging an ISDN Terminal Adapter

Jean Anne Booth

Advanced Micro Devices
5900 East Ben White Blvd., MS 561
Austin, Texas 78741
Tel: (512) 462-5879
Fax: (512) 462-5051

1993
High Speed Digital
Symposium

Integrated Services Digital Network (ISDN)

- All-digital network standards: voice, data, control signals
- Provides digital transmission capabilities
- Offers enhanced communication services:
 - Integrated services
 - Integrated networks
 - Integrated applications

Author

Jean Anne Booth

Background:
Prior to joining 30K Marketing, she managed the 30K Technical Support Center, providing hardware and software technical and application support to 30K Family customers. Before joining AMD, Jean Anne was a development engineer involved in the software implementation of real-time control systems. Jean Anne holds a BS in Electrical Engineering and an MS in Computer Engineering.

Current Activities:
Jean Anne Booth is a Senior Technical Marketing Engineer at Advanced Micro Devices. She has been with AMD for 6 years, and is currently responsible for technical marketing of current and future high performance 30K RISC microprocessors.

Abstract

The Integrated Services Digital Network digitizes voice signals at the telephone and sends both voice and control signals to the PABX or central office switch digitally. An ISDN terminal adapter interfaces between ISDN and non-ISDN equipment, typically at the subscriber loop level that interfaces between the customer's equipment and the

local telephone network office. This presentation introduces ISDN, describes an ISDN terminal adapter, and shows how to implement an ISDN terminal adapter in hardware and software. Using an evaluation board, PC, JTAG emulator, and HP 16500 logic analyzer, the presentation also shows how to develop and debug the ISDN terminal adapter presented here.

Author

Jean Anne Booth

Current Activities:

Jean Anne Booth is a Senior Technical Marketing Engineer at Advanced Micro Devices. She has been with AMD for 6 years, and is currently responsible for technical marketing of current and future high performance 29K™ RISC microprocessors.

Background:

Prior to joining 29K Marketing, she managed the 29K Technical Support Center, providing hardware and software technical and applications support to 29K Family customers. Before joining AMD, Jean Anne was a development engineer involved in the software implementation of real-time control systems. Jean Anne holds a BS in Electrical Engineering and an MS in Computer Engineering.

Slide #1

Developing and Debugging an ISDN Terminal Adapter



Advanced Micro Devices

Slide #3

The Integrated Services Digital Network (ISDN)

- All-digital network standard: voice, data, control signals
- Replaces subscriber loop with digital voice and data capability
- Global communications network benefits:
 - increased reliability
 - increased functionality
 - lower cost
 - worldwide standardization

Slide #2

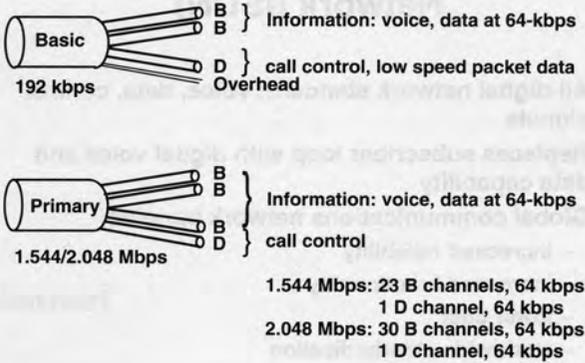
Developing and Debugging an ISDN Terminal Adapter

- Introduction to ISDN
- An ISDN Terminal Adapter
 - Am79C30A/32A Digital Subscriber Controller (DSC)
 - Am29200™ 32-bit RISC microcontroller
 - Am85C30 Serial Communications Controller (SCC)
- Developing an ISDN Terminal Adapter
 - Hardware
 - Software
- Debugging the ISDN Terminal Adapter
 - PC, JTAG emulator, HP 16500 logic analyzer, SA-29200 evaluation and expansion boards

The Integrated Services Digital Network, or ISDN, digitizes voice signals at the telephone and sends both voice and control signals to the PABX or central office switch digitally. Thus, ISDN replaces the last analog component of telephone circuitry, the subscriber loop, with a digital component capable of handling both voice and data information. ISDN brings the benefits of digital technology — increased reliability, new functionality, lower cost, and increased security and privacy — and worldwide standardization to the global communications network.

Slide #4

ISDN Classes of Service

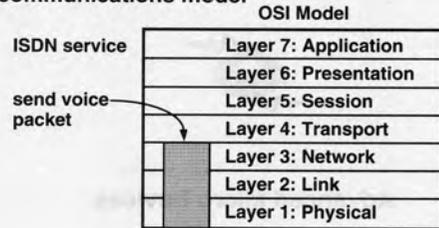


ISDN service is divided into two classes – *primary rate*, an expensive high-bandwidth connection, and *basic rate*, the type of subscriber connection most commonly used. The focus here is on the basic rate service, which provides three communications channels. The two B (or bearer) channels provide either voice or data service at 64 kbps. The D (or signaling/data) channel provides call control services and low speed packet data transmission (up to 9600 bps). Primary rate service is provided at either 1.544 Mbps (US, Canada, and Japan) or 2.048 Mbps (Europe). The channel structure for the 1.544-Mbps rate is typically 23 B channels and one D channel; the 2.048-Mbps rate is typically composed of 30 B channels and one D channel.

Slide #5

ISDN Standards

- CCITT responsible for ISDN standards
- Basic structure defined in CCITT I.411 recommendation
- Utilizes Open Systems Interconnection (OSI) communications model



The group overseeing the definition of ISDN is the CCITT, and the basic structure of the ISDN is specified in the CCITT I.411 recommendation. ISDN functions are further subdivided by the OSI (open systems interconnection) seven-layer communication model. The OSI model defines physical and logical services provided by each layer; a vertical “slice” of the model encompassing at least layers 1 through 3 provides one ISDN function, such as a user data transfer.

Slide #6

CCITT Recommendations for ISDN

I Series

- Complete set of recommendations for all standardization aspects of ISDN
- Cross references specifications from other series (Q Series for protocols, and V/X series for non-ISDN terminals)

Slide #8

CCITT Recommendations for ISDN

Q Series

- Q.920/Q.921 – Defines the Layer 2 protocol (LAPD) used by ISDN
 - Ensures error-free and correctly sequenced data transmission between Layer 3 entities
 - Specifies syntax of message format used within HDLC frames at Layer 2

Slide #7

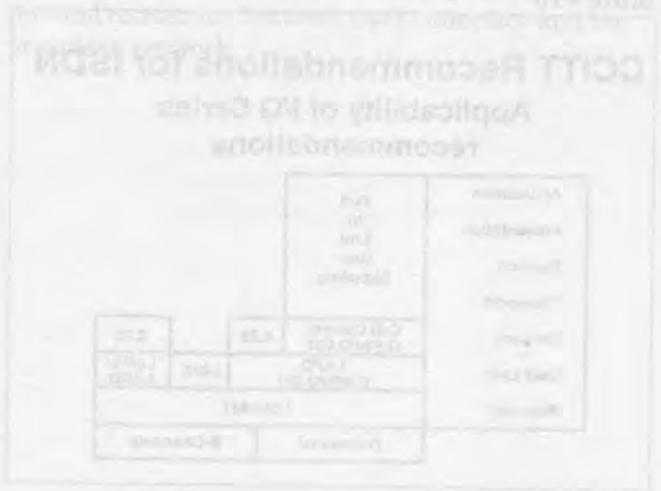
CCITT Recommendations for ISDN

I Series

- I.420/I.421 – Introduction to ISDN concepts and other I Series recommendations
- I.430 – Layer 1 interface specification ('S' and 'T' interface recommendations)
- I.431 – Layer 1 primary rate interface
- I.440/I.441 – Layer 2 protocol specification (LAPD) Cross references Q.920/Q.921
- I.450/I.451 – Layer 3 protocol specification. Cross references Q.930/Q.931

The Q.920/Q.921 recommendation defines the Layer 2 protocol used by ISDN, also known as LAPD. The LAPD protocol ensures error-free and correctly-sequenced data transmission between Layer 3 entities. The services provided by the LAPD protocol include both unacknowledged and acknowledged information transfer on the ISDN D-channel. The Q.920/Q.921 recommendation also specifies the syntax of the message format used within HDLC frames at Layer 2.

The CCITT I Series of recommendations is a complete set of recommendations for all aspects of ISDN. The I Series recommendations cross-reference specifications from other series, like the Q Series for protocols, and the V and X series for non-ISDN terminals. CCITT has defined many more I Series recommendations than the ones listed here.



Developing and Debugging an ISDN Terminal Adapter

Slide #9

CCITT Recommendations for ISDN

Q Series

- Q.930/Q.931** – Defines internationally agreed portion of Layer 3 (network layer) protocol for ISDN
- Provides packetizing and blocking of Layer 4 messages for Layer 2 conformance
 - Does not address supplementary service (defined by national committees)

The Q.930/Q.931 recommendation defines the internationally agreed-upon portion of the Layer 3 protocol used by ISDN. It details packetizing and blocking of Layer 4 messages for Layer 2 conformance, but does not address supplementary services, such as call waiting, call transfer, credit card calling, etc. Supplementary services are defined by national committees.

Slide #10

CCITT Recommendations for ISDN Applicability of I/Q Series recommendations

Application	End to End User Signaling		
Presentation	End to End User Signaling		
Session	End to End User Signaling		
Transport	End to End User Signaling		
Network	Call Control Q.930/Q.931	X.25	X.25
Data Link	LAPD Q.920/Q.921	I.46X	LAPB/ LAPD
Physical	I.430/431		
	D-Channel	B-Channels	

This is another look at the relationship between OSI and ISDN. As a network, ISDN is primarily unconcerned with layers 4 through 7, which you employ for exchanging information. Layer 1, defined in I.430 and I.431, specifies the physical interface for both basic and primary rate access. Because both the

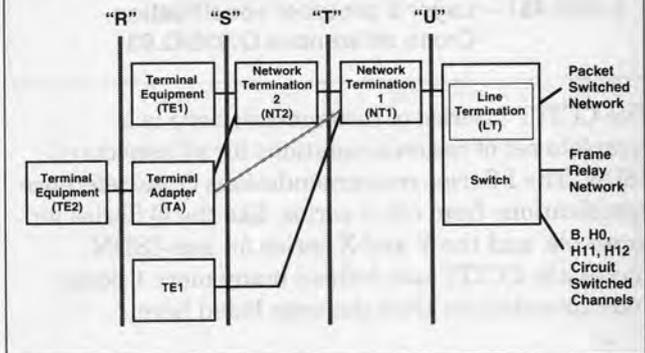
B and D channels share the physical interface, these standards apply to both types of channels. Above this layer, the protocol structure differs for the two types of channels.

For the D channel, the LAPD protocol defined in Q.920 and Q.921 is employed in the data link layer. For the B channel, the I.46X series of recommendations defines alternative protocols for interfacing existing equipment to the ISDN. Because the B channel can also be packet-switched, the LAPD protocol can also be used, in addition to using the more common LAPB protocol for data link layer transfers.

At the network layer, Q.930 and Q.931 define call control for the D channel. If the D channel is used to provide packet switching services, the X.25 level 3 protocol is used. For the B channel, the X.25 level 3 protocol provides network layer services for packet switching.

Slide #11

ISDN Network Topography



The network topography from the desktop to the switch is shown here, identifying certain classes of equipment that make up the network. "Reference points" are defined that represent various interfaces with CCITT standards for both hardware and software. In developing a terminal adapter, we will be dealing with the R and S/T interfaces.

Slide #12

ISDN Terminology

- **TE1:** ISDN compatible voice and/or data terminal
- **TE2:** Non-ISDN compatible terminal (such as V.24, X.21, X.25, SNA terminal)
- **TA:** Terminal adapter providing physical and/or protocol conversion between a TE2 and the ISDN
- **NT2:** Network termination providing switching and/or concentration (such as PBX) – not present in single line installations

In the ISDN network topography, a TE1 (terminal equipment type 1) refers to devices that support the standard ISDN interface, such as digital telephones, integrated voice/data terminals, or digital facsimile machines.

A TE2 (terminal equipment type 2) refers to any non-ISDN compatible terminal, typically existing equipment. Examples of TE2 are terminals with an RS-232 interface, host computers with an X.25 interface, and SNA terminals. TE2 devices require a TA (terminal adapter) to plug into an ISDN interface. TAs may provide physical conversions, protocol conversions, or both.

All terminal equipment, whether TE1 or TE2, provides protocol handling, maintenance functions, interface functions, and connection functions to other equipment.

An NT2 (network termination 2) is an intelligent device that provides switching and/or concentration functions. Examples of NT2s include digital PBX, terminal controllers, and LANs. An NT2 performs layer 2 and layer 3 protocol handling, layer 2 and layer 3 multiplexing, maintenance functions, and interface termination, in addition to switching and concentration.

Slide #13

ISDN Terminology

- **NT1:** Network termination providing physical and/or protocol conversion between the 'S'/'T' interface and the network-provided 'U' interface
- **LT:** Line termination performing physical and/or protocol conversion between 'U' interface and central office exchange internal highways

An NT1 (network termination 1) provides physical and electrical termination of the ISDN. It may also provide protocol conversion between the S/T interface and the network's U interface. The NT1 may be controlled by the ISDN provider, and forms a boundary to the network. The functions provided by an NT1 include line transmission termination, line maintenance and performance monitoring, timing, power transfer, layer 1 multiplexing, and interface termination, including multidrop termination employing layer 1 contention resolution.

An LT (line termination) provides physical and/or protocol conversion between the U interface and the provider's network.

The frame structure in the NT to TE direction is similar, except that some of the dc balancing bits are replaced by D-channel echo bits (E), which are a retransmission by the NT of the most recently received D bit from the TE. The echo bits are used to provide D-channel contention resolution in a network with multiple TEs in a passive-bus configuration. The activation bit (A) is used to activate a TE (bring it on-line). The N bit is set to one unless it being used in a multiframe structure. The M bit indicates a multiple frame. Multiframing is defined in I.430, and consists of 20 frames as defined here, where the auxiliary framing bit carries the Q multiframe data in the TE to NT direction, and S carries the Q multiframe data in the NT to TE direction.

Slide #16

Characteristics of the 'R' and Reference Point

The 'R' reference point allows non-ISDN terminals to be connected to the ISDN.

<u>Standard</u>	<u>Protocol</u>
V.120	LAPD packet protocol
DMI Mode 2	Bit stuffing
DMI Mode 3	LAPD/X.25
V.110/ECMA 102	Bit stuffing

The R reference point allows non-ISDN terminals to be connected to the ISDN. Some common non-ISDN interfaces and their protocols are listed here.

Slide #17

Terminal Adaption

- **Bit stuffing – ECMA102/V.110, DMI Mode 2**
 - Low cost
 - No error detection/re-transmission
- **Packetizing – V.120, DMI Mode 3**
 - Error detection/re-transmission
 - Statistical multiplexing
 - Higher throughput

An ISDN terminal adapter is the interface between ISDN and non-ISDN equipment, typically at the subscriber loop level that interfaces between the customer's equipment and the local telephone network office.

An ISDN data-only terminal adapter will interface either a bit-stuffing protocol or a packetizing protocol to the ISDN. A terminal adapter for bit-stuffing terminals has the advantage of low cost, but also has no error detection or retransmission. Terminal adapters for packetizing protocols are more expensive but also more reliable, with error detection and retransmission included in the protocol conversion. In addition, packetizing systems have a higher throughput.

Developing and Debugging an ISDN Terminal Adapter

Slide #18

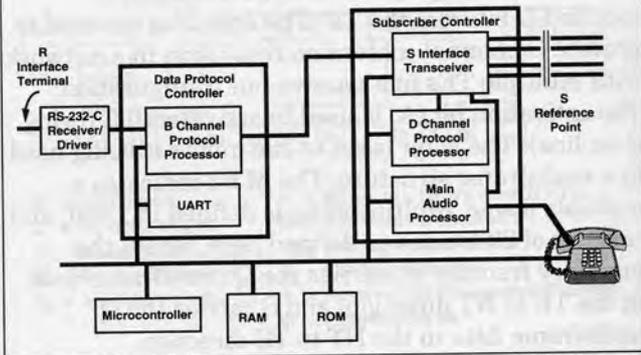
Terminal Characteristics

- High volume products
- Severe cost constraints
 - Low component cost
- Power constraints for voice products
 - Efficient CMOS devices
- High feature content
 - Complex software, powerful MPU
- Switch specific versions
 - Multiple software variants

An ISDN terminal adapter is a high-volume product because the installed base of non-ISDN compatible equipment is nearly the size of the entire installed terminal market. Like all high-volume products, they operate under severe cost constraints, requiring a low component count to maintain competitiveness with other vendors. Voice terminal adapters also have a strict power budget, so solutions with power-efficient CMOS devices are required. Because interfacing a non-ISDN terminal to the ISDN usually involves adding ISDN features that aren't a part of the existing analog solution, terminal adapters require powerful processors to implement these new features in software. To broaden the market potential for an adapter, manufacturers prefer to create specific versions of the same basic adapter that differ only in software to handle different protocols or implement required ISDN features not present in the non-ISDN terminal.

Slide #19

ISDN Voice and Data Terminal Adapter

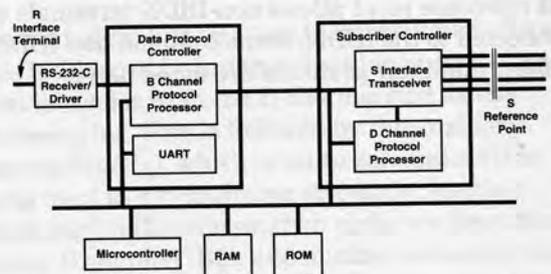


This shows a block diagram of an ISDN terminal adapter.

Our case study will build a terminal adapter using an Am79C30A Digital Subscriber Controller (DSC) for basic ISDN services and an Am29200™ RISC microcontroller for control of the DSC and ISDN protocol. An Am85C30 Serial Communications Controller (SCC) provides the serial interface to the R reference point. Using an Am79C30A allows ISDN-compatible transmission of both voice and data.

Slide #20

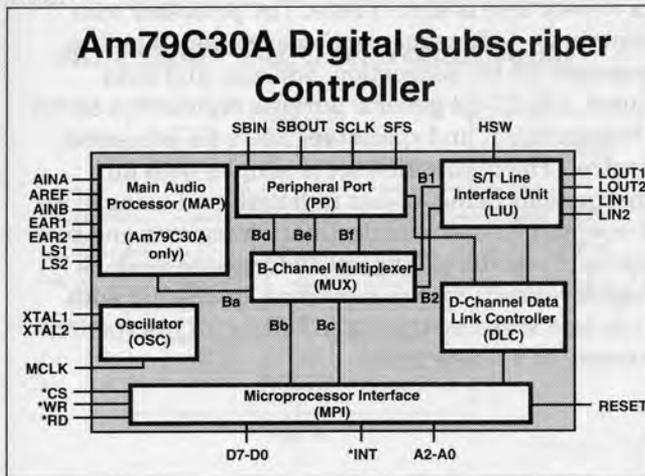
ISDN Data-only Terminal Adapter



An ISDN data-only terminal adapter could be built using an Am79C32A ISDN Data Controller (IDC) instead of the Am79C30A, as shown here.

Developing and Debugging an ISDN Terminal Adapter

Slide #21



The Am79C30A/32A provides a 192-kbps full duplex digital path over four wires between the TE (terminal equipment) located on the subscriber's premises and the NT (network terminal) or PABX linecard. All physical Layer functions and procedures are implemented, including framing, synchronization, maintenance, and multiple terminal contention. The Am79C30A/32A processes the ISDN basic rate bit stream. The B channels are routed to and from different portions of the DSC under software control. The D channel is partially processed and then passed to the microcontroller for further processing.

This is a block diagram of the Am79C30A DSC. The main audio processor, or MAP, is the only portion of the Am79C30A that is not present on the Am79C32A. The MAP uses DSP to implement a high performance codec/filter function. The MAP supports a loudspeaker, an earpiece, and two separate audio inputs. Gain, frequency response, and tone generation are programmable.

The S/T line interface unit, or LIU, provides the interface to an ISDN S or T reference point. It contains a hook-switch input and differential subscriber line inputs and outputs. The LIU monitors the S interface and hook switch during power down, allowing the microprocessor to be shut down to conserve power during idle periods.

The microprocessor interface unit (MPI) communicates with the processor controlling the terminal adapter. The address line inputs select source and

destination registers for read and write operations on the data bus. The data bus is used to exchange information with the controlling processor. An interrupt input informs the processor that the DSC needs service, and chip select and read/write signals are provided for system interfacing.

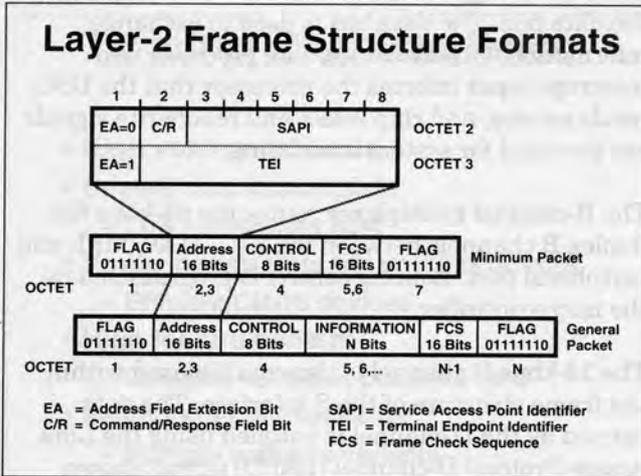
The B-channel multiplexer routes the 64-kbps full-duplex B channels between the LIU, MAP, MPI, and peripheral port. Routing control is programmed by the microcontroller.

The 16-kbps D channel is time multiplexed within the frame structure of the S interface. The data carried by the D channel is encoded using the Link Access Protocol D-channel (LAPD) format shown here. The LIU controls the multiplexing and demultiplexing of the D-channel data between the S interface and the D-channel data link controller (DLC).

The Am79C30A/32A will generate a maximum of one interrupt every 125µs. Once asserted, the interrupt will remain asserted until the microcontroller reads the DSC's interrupt register. Events that generate interrupts include DLC receive FIFO full, DLC transmit FIFO empty, LIU change of state (on hook/off hook), packet errors, and packet status (last byte, etc.).

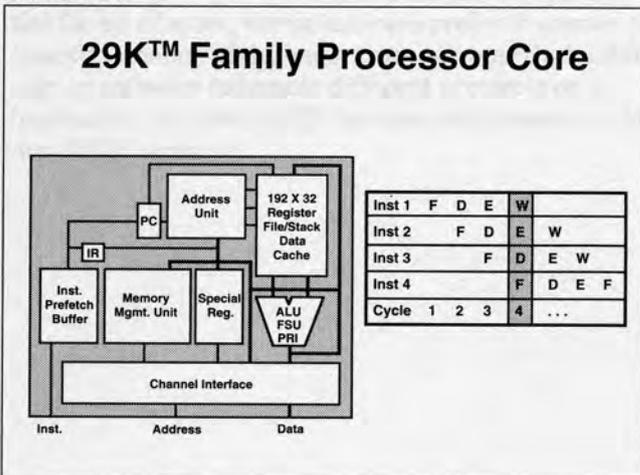
Developing and Debugging an ISDN Terminal Adapter

Slide #22



The DLC performs processing of Layer-1 and partial Layer-2 LAPD protocol, which includes flag detection and generation, zero deletion and insertion, and Frame Check Sequence (FCS) processing for error detection. Higher-level protocol processing is done by the external microcontroller. The DLC contains two 8-byte data FIFOs for receive and transmit data, and three 2-byte status FIFOs that make it possible to receive two back-to-back data packets.

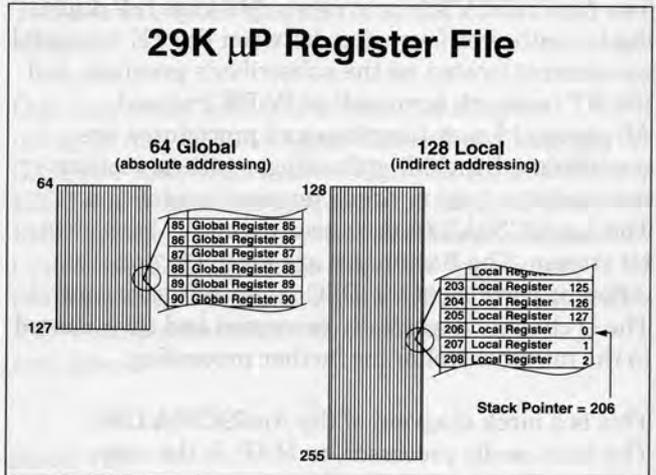
Slide #23



The Am29200 32-bit RISC microcontroller contains a 29K™ Family 32-bit RISC processor core, memory controllers, and integrated peripherals. The 29K

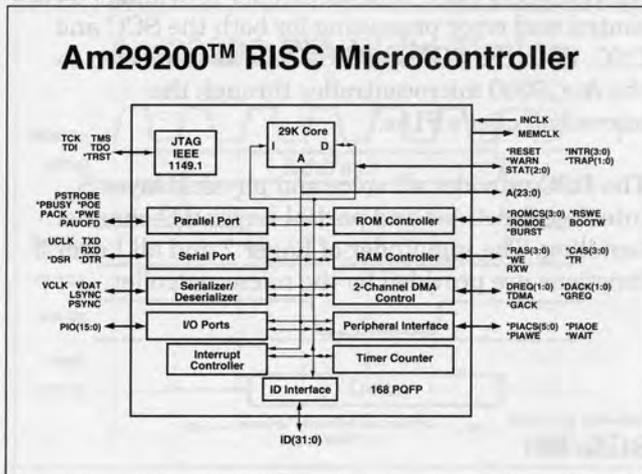
processor core is shown here. The processor core contains a 4-stage instruction/execution pipeline, separate 32-bit instruction, address, and data buses, 192 32-bit general purpose registers, a 56-bit timer/counter, and special registers for processor control. The instruction set is simple, with all instructions being 32-bits in length and using a 3-operand format (one destination register and two operand registers). The only addressing mode is register-direct; data is moved on- or off-chip with LOAD and STORE instructions. Nearly all instructions execute in a single cycle.

Slide #24



The 192 general-purpose registers are split into two groups, 64 global registers and 128 local registers. The local registers are organized internally as a circular queue. The register pointed to by global register 1 (**gr1**) is local register 0 (**lr0**); the register below **lr0** is **lr1**, and the register above is **lr127**. The local registers are used by compilers to implement dynamic register windowing, and are controlled by software convention.

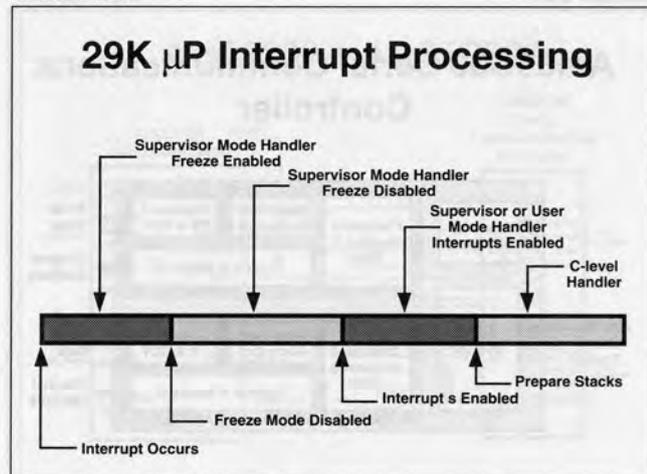
Slide #25



This is a block diagram of the Am29200 microcontroller. The ROM controller and DRAM controller implement a glueless interface to memory. The DMA controller provides two channels for transfer of data between the DRAM and internal or external peripherals. The peripheral interface adapter (PIA) implements a glueless interface to up to six generic peripherals, and will be used to connect to the DSC and SCC. The I/O port provides 16 programmable signals that can be outputs, inputs, or interrupt triggers. The 5-pin JTAG port is used for test and debug control of the processor.

Internal peripherals in the Am29200 microcontroller are addressed with LOAD and STORE instructions using a pre-defined memory mapped address. In addition, special registers are programmable for control and status for each peripheral.

Slide #26

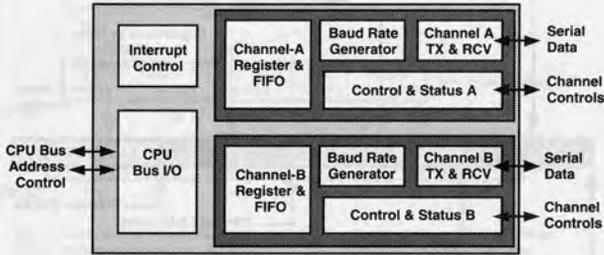


Interrupts in the Am29200 microcontroller are triggered by external interrupt pins, programmed pins in the I/O port, and internal sources, such as the DMA controller or software traps. When an interrupt occurs, the processor completes or cancels current bus activity, saves the Current Processor Status (**cps**) register into the Old Processor Status (**ops**) register, sets itself in Supervisor mode and freezes other processor status (so as not to corrupt the interrupted application), decodes the interrupt, and fetches the first instruction of the applicable interrupt handler. In keeping with the RISC philosophy of machine simplicity, no state or status other than the **cps** is saved. At the software engineer's option, the interrupt handler may execute in Freeze mode, utilizing only registers set aside for this purpose, or it may save more state and utilize more of the processor's resources.

Developing and Debugging an ISDN Terminal Adapter

Slide #27

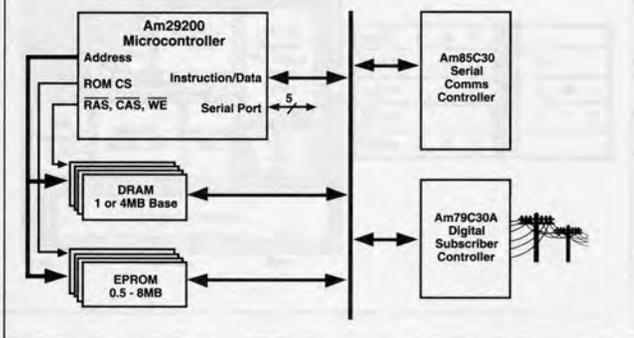
Am85C30 Serial Communications Controller



The Am85C30 Serial Communications Controller (SCC) is a dual channel multi-protocol data communications peripheral, handling both asynchronous and synchronous formats including SDLC/HDLC and BiSYNC. It contains two channels, and each channel has an independent oscillator, baud-rate generator, and digital-phase locked loop for clock recovery. The device is controlled by internal registers read and written by a microprocessor through the 8-bit data bus. The SCC interrupts its controlling processor for transmit complete, receive complete, and error conditions.

Slide #28

ISDN Terminal Adapter Functional Block Diagram

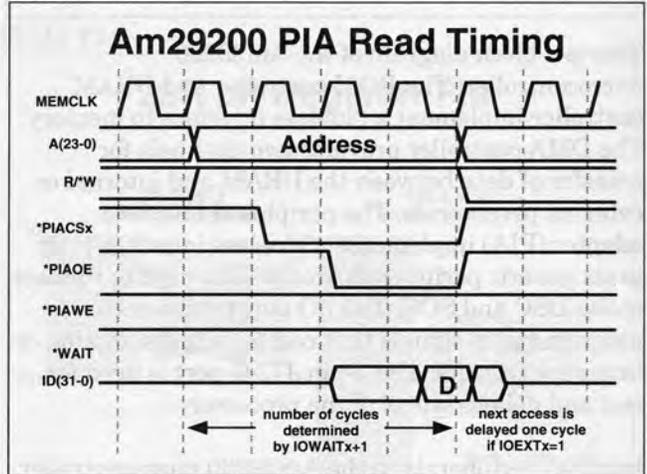


This block overview of the ISDN terminal adapter shows the SCC connecting to the ISDN R interface, the DSC connecting to the ISDN S/T interface, and

the Am29200 RISC microcontroller providing control and error processing for both the SCC and DSC. Both the SCC and DSC will be connected to the Am29200 microcontroller through the microcontroller's PIAs.

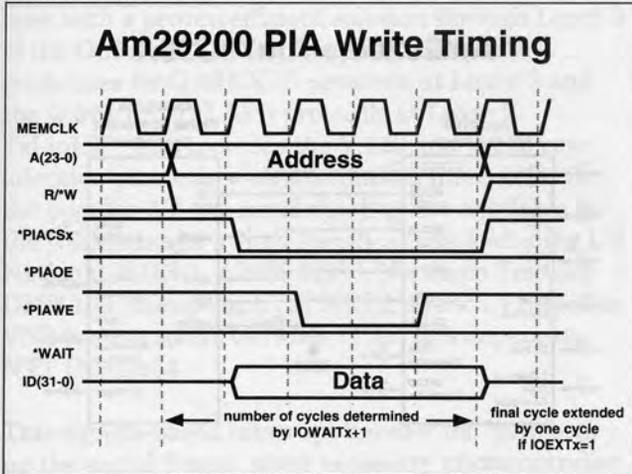
The DSC provides all voice and physical layer S interface functions and partial Layer 2 D-channel handling. The remainder of Layer 2 and all Layer 3 functions are provided by the microcontroller.

Slide #29



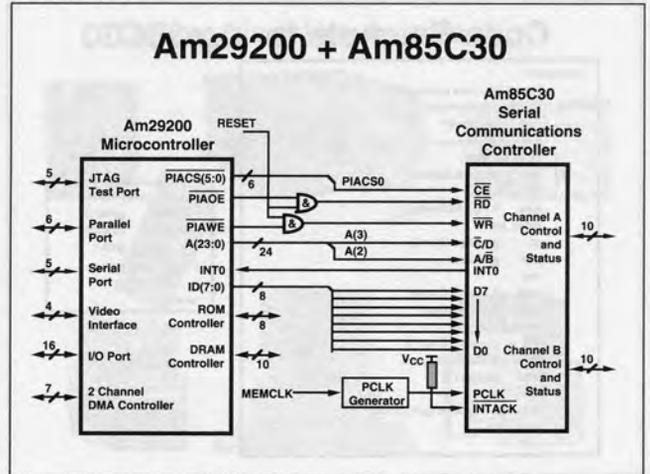
This is the timing of a PIA read cycle. The number of cycles until the PIA Chip Select (*PIACSX) and PIA Output Enable (*PIAOE) are deasserted is dependent upon the value of Input/Output Wait States (IOWAITX) field of the PIA control register in the microcontroller. The minimum access time for a PIA read is 3 cycles (2 wait states). If the Input/Output Extend (IOEXTX) field of the PIA control register is set, the next PIA access will be delayed one cycle for an additional cycle of output disable time.

Slide #30



This is the timing of a PIA write cycle. The number of cycles until the PIA Chip Select (*PIACsX) and PIA Write Enable (*PIAWE) are deasserted is dependent upon the value of Input/Output Wait States (IOWAITX) field of the PIA control register in the microcontroller. The minimum access time for a PIA write is 4 cycles (3 wait states). If the Input/Output Extend (IOEXTX) field of the PIA control register is set, the next PIA access will be delayed one cycle for an additional cycle of data hold time.

Slide #31



This is the interface between the Am29200 microcontroller and the Am85C30 SCC. Note that the SCC doesn't have an explicit RESET signal; to achieve a reset, the device expects *RD and *WR to be asserted at the same time. This interface shows an interrupt-driven communication mechanism between the microcontroller and the SCC. Assuming a 16-MHz microcontroller and an 8-MHz SCC, the SCC's PCLK is generated from the microcontroller's MEMCLK signal with a divide-by-2 circuit. A complicating factor is the access recovery time of the SCC. The minimum time from the leading edge of one command to the leading edge of the next command is defined as 3.5 times the PCLK cycle time. The access recovery time can be generated in software by inserting delay instructions, or in hardware by adding an external PAL with delay states or by using the microcontroller's WAIT signal. This design assumes that software assures the minimum access recovery time for the SCC.

Slide #32

Code Example for Am85C30

```

.SerialSend:
Save 6 Registers (gr96-gr101);
Control = 0x9000007
Data = 0x900000F
End = Start + Length -1;
Current = Start;
S3: store 0,1,gr98,Control ; TxRdy?
Delay_Macro
Delay_Macro
sll gr100,gr100,31-2
jmpf gr100,S3
nop
load 0,1,gr100,Current ; Send data
store 0,1,gr100,Data
Delay_Macro
add Current,Current,1
cpleu gr100,Current,End
jmpf gr100,S3 ; More?
nop
Restore 6 Registers (gr96-gr101)
jmpf lr0
nop

```

```

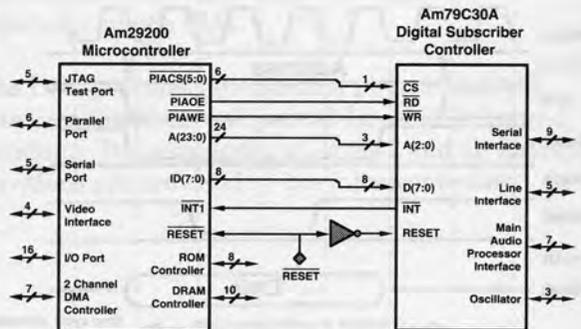
.macro Delay_Macro
mfsr gr98,lru
mfsr lru,gr98
const gr98,5
S1: jmpfdec gr98,S1
nop
.endm

```

When the SCC interrupts the microcontroller for service by asserting *INT0, the processor interrogates the SCC's status register RR2 to determine the source of the interrupt. This pseudo-code shows how serial information is sent using the SCC. Six microcontroller registers are saved before being used in the `_SerialSend` routine. The microcontroller sets the SCC control register for transmit mode, and then checks the SCC's response. When the SCC is ready to transmit, the microcontroller gets the address of the send buffer, and sends bytes until the message is complete. `Delay_Macro` consumes 6 cycles to meet the SCC's minimum access recovery time. Because the Am29200 microcontroller can overlap loads and stores in the pipeline for better performance, the delay macro includes a serializing instruction (`MFSR`, or move to special register), which won't allow loads and stores to overlap. The special register used in the serializing instruction doesn't matter; here special register `lru` (least recently used indicator) for the memory management unit is used.

Slide #33

Am29200 + Am79C30A



This is the interface between the Am29200 microcontroller and the Am79C30A DSC device. Note that the DSC `RESET` is an active-high signal, while the Am29200 microcontroller `*RESET` is an active-low signal.

Slide #34

Code Example Signal Interrupt Handler

```

.global _DSCInt
_DSCInt:
; Process signal
const i0,SIGUSR1 ; User defined signal number
sub msp,msp,4
store 0,0,i0,msp ; pass signal number
sub msp,msp,4
store 0,0,gr1,msp ; save rsp
sub msp,msp,4
store 0,0,rab,msp ; save rab
sub msp,msp,4
const i0,128*4 ; fix rab
sub rab,rab,i0 ; fix rab
pushsr PC0,PC1,PC2,CHA,CHD,CHC,ALLI,OPS ; save specials
;
mfsr i0,OPS
cfsr i0,i0 ; Ensure ints disabled
mfsr OPS,i0
;
sub msp,msp,4
store 0,0,tav,msp ; save tav
mfsr/chc,i0 ; ensure load/store does not restart
mfsr/pc1,SignalRegister
add i2,SignalRegister,4
mfsr/pc0,i2
int ; Goto algocode and C handler, ints disabled

```

Ensures that Handler is executed with interrupts disabled.

The easiest method to handle the DSC's interrupts is to utilize `C signals()` and `C interrupt handlers`. This minimizes the amount of code that must be written by hand in assembly language and allows the developer to take advantage of `TsLink3™` software from TeleSoft International, Inc. Written in

Developing and Debugging an ISDN Terminal Adapter

ANSI C, the TsLink3™ software provides the developer with a proven efficient solution through Layer 3 of the OSI model, and is compliant with ISDN guidelines for Q.391/X.25 protocols at Layer 3 and the Q.921 LAPD/LAPB protocols at Layer 2. TsLink3™ also includes the V.110 and V.120 rate adaption protocols with a command interpreter for the popular AT command set. Kits are available for the world's major switch specifications, including US National ISDN-1, AT&T 5ESS, Northern Telecom DMS-100, European ETSI NET3, French VN2 (with VN3 coming soon), German 1TR6, and Japanese NTT INSnet64.

This signals-based interrupt handler for *INT0 sets up the signal frame, saves necessary microcontroller special registers defining the current environment, and then passes control to the C-based signal interrupt handler. The interrupt handler executes with interrupts disabled.

Slide #35

```

void dscier()
{
    register ireg;
    register dsr2;
    unsigned char data;

    // loop while DSRI, DERI, LSRI, DRTHRESH, or DTTHRSR bits set in the DSC
    // Interrupt Register
    while((ireg = Rd_dsc(DSC_INT_REG)) & (DSRI | DERI | LSRI |
    DRTHRESH | DTTHRSR)) {
        /* Process Transmit FIFO Threshold Reached interrupt */
        /* (ireg & DTTHRSR) {
        Profile(DSC_INT_1);
        /* while Transmit Buffer Empty */
        while((data = Rd_dsc(DSC_DSR2)) & TBE && Ram->xmtcnt) {
            /* write a byte to FIFO */
            Wr_dsc(DSC_D_BUFF, *Ram->xmtnext++);
            Ram->xmtcnt--;
        }
        /* Process Receive FIFO Threshold Reached interrupt */
        /* (ireg & DRTHRSR) {
        Profile(DSC_INT_2);
        dsr2 = Rd_dsc(DSC_DSR2);
        /* while Receive Byte Available in FIFO */
        while(dsr2 & RBA) {
            /* read a byte from FIFO */
            *Ram->rcvfree++ = Rd_dsc(DSC_D_BUFF);
            if((dsr2 = Rd_dsc(DSC_DSR2)) & LBRP) {
                /* Set RCY_IP so won't read any bytes from d_buff in EORP code
                /* in dch_stat
                Ram->rcvstat |= RCY_IP;
                break;
            }
        }
    }
                
```

Code Example Signal Handler

Now in the C language signal handler, the actual interrupt is processed. First, the DSC interrupt register is read to determine what caused the interrupt, and then the condition (transmit FIFO threshold reached, receive FIFO threshold reached, etc.) is handled. Two C macros are used to communicate with the DSC: Rd_dsc() to read an interrupt register (thus clearing the interrupt), and Wr_dsc() to write a register in the DSC. This code is taken directly from the TsLink3 software.

Slide #36



The ISDN terminal adapter can be designed using an IBM-compatible PC (Am386(r) microprocessor-class or better), an SA29200 Evaluation Board and SA29200 Expansion Board, a Corelis Am29200 microcontroller JTAG-based emulator, an HP 16500 logic analyzer, and the TsLink3 software.

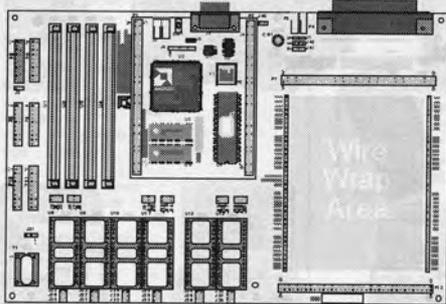
Slide #37



The SA29200 Evaluation Board is a small form-factor board containing an Am29200 microcontroller, DRAM, an EPROM with the MiniMON29K™ debug monitor, and a serial port connector.

Slide #38

SA29200 Extension Board



The SA29200 Evaluation Board plugs into an SA29200 Expansion Board which contains header sockets for extra memory, a parallel port connector, and a wire-wrap area in which the DSC and SCC can be wired.

Slide #39

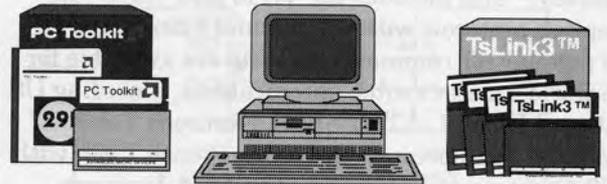
Corelis JTAG Emulator



The Corelis JTAG emulator is a low-cost in-circuit emulator that clips onto the Am29200 microcontroller and uses the 5-pin JTAG port to start, stop, and single-step the Am29200 microcontroller, and to request current status information from the microcontroller. It is a board that plugs into the expansion slot of a PC. A special cable connects the emulator board to the system under development.

Slide #40

PC Platform



The PC does triple-duty as a software development platform, the host side of the MiniMON29K debug monitor, and the host platform for the JTAG emulator.

Slide #41

HP 16500 Logic Analyzer



The HP 16500 is used to analyze system activity on a signal/bus level. It works with the JTAG emulator to provide comprehensive information about the system under development.

Developing and Debugging an ISDN Terminal Adapter

Slide #42

HP 16500 Analyzer Listing									
Label->	ADDR	AM29200 Disassembly	STAT	R/W	ROMOE/	ROMCS30			
Base->	Hex	memonics	RDP	Hex	Hex	Hex	Hex	Hex	Hex
5	000AAC	STORE 0,0x00,gr96,gr97	0	7	1	1			E
6	000AAC	ROM Inst. Read: 0x1E	0	5	1	1			E
7	000AAD	Byte #1	0	7	1	1			E
8	000AAD	ROM Inst. Read: 0x00	0	7	1	0			E
9	000AAE	Byte #2	0	7	1	1			E
10	000AAE	ROM Inst. Read: 0x60	0	7	1	0			E
11	000AAF	Byte #3	0	7	1	1			E
12	000AAF	ROM Inst. Read: 0x61	0	7	1	0			E
13	000AB0	CONST gr97,0x008C	0	7	1	1			E
21	000AB4	CONSTR gr97,0x8000	0	7	1	1			E
22	000AB4	ROM Inst. Read: 0x02	0	5	1	0			E
23	000AB5	Byte #1	0	7	1	1			E
24	000AB5	ROM Inst. Read: 0x80	0	7	1	0			E
25	000AB6	Byte #2	0	7	1	1			E
26	000AB6	ROM Inst. Read: 0x61	0	7	1	0			E
27	000AB7	Byte #3	0	7	1	1			E
28	000AB7	ROM Inst. Read: 0x00	0	7	1	0			E
29	000080	Idle Cycle		7	0	1			F
30	000080	Idle Cycle		5	0	1			F
31	000080	Internal data access: Serial Port Control Register		7	0	1			F
32	000AB8	LOAD 0,0x00,gr96,gr97	0	7	1	1			E
33	000AB8	ROM Inst. Read: 0x16	0	6	1	0			E

This HP 16500 display shows a part of the startup code for the Am29200 microcontroller. The startup code is located in a byte-wide EPROM, so the microcontroller does four byte fetches for each 32-bit instruction word. The lines labeled 5 through 12 show bus activity for the fetching of a single instruction (STORE 0,0x00,GR96,GR97) from byte-wide memory; each memory access takes two cycles. The RDP field indicates which DRAM or ROM bank was accessed; in this case, the accesses were all to ROM bank 0.

The microcontroller's three STAT (status) signals show what the microcontroller was doing during the previous cycle. A status of 7 indicates that the microcontroller was idle (data/instruction not valid). A status of 5 indicates that the microcontroller was executing during the previous cycle. A status of 6 indicates that the microcontroller executed an internal data access (to an internal peripheral), in this case to the serial port.

The R/*W signal indicates if a read access is taking place (1) or a write access is taking place (0). The *ROMOE signal is asserted (0) when the output enable to the EPROM is asserted; the analyzer listing shows that the microcontroller's ROM Controller was correctly programmed for two-cycle access memories. The romcs30 field shows the ROM chip selects that have been asserted; there is one chip select for each of the four allowed ROM banks. The startup code is all located in ROM bank 0, so this is the only chip select asserted during this listing. Note that during the internal access to the microcontroller's serial port, no ROM banks are enabled.

Slide #43

HP 16500 Analyzer Listing													
Label->	ADDR	AM29200 Disassembly	STAT	R/W	TR/OE	WE/	CAS30/	RAS30/					
Base->	Hex	memonics	RDP	Hex	Hex	Hex	Hex	Hex	Hex	Hex	Hex	Hex	Hex
16	000024	CONST lr3,0x40C4	3	4	1	1	1	F	7				
17	000024	DRAM Inst. Read from 0x4024	3	7	1	0	1	0	7				
18	000024	DRAM Inst. Read from 0x4024	3	7	1	0	1	0	7				
19	000028	CONSTR lr3,0x4000	3	7	1	0	1	0	7				
20	000028	DRAM Inst. Read from 0x4028	3	5	1	0	1	0	7				
21	000044	Idle cycle	3	7	0	1	1	F	F				
22	000044	DRAM Write: 0x4000C50 to 0x96E8	3	5	0	1	1	F	F				
23	0000E8	DRAM Write: 0x4000C50 to 0x96E8	3	7	0	1	0	0	7				
24	0000E8	DRAM Write: 0x4000C50 to 0x96E8	3	7	0	1	0	0	7				
25	000020	Idle cycle	3	7	1	1	1	F	F				
26	00002C	CONST lr2,0x0040	3	4	1	1	1	F	7				
27	00002C	DRAM Inst. Read from 0x402C	3	7	1	0	1	0	7				
28	00002C	DRAM Inst. Read from 0x402C	3	7	1	0	1	0	7				
29	000030	HALT	3	7	1	0	1	F	7				
30	000030	DRAM Inst. Read from 0x4020	3	5	1	0	1	0	7				
31	000034	ASNEQ 69,sp,sp	3	7	1	0	1	F	7				
32	000034	DRAM Inst. Read from 0x4024	3	5	1	0	1	0	7				
33	000038	CONST lr3,0x408C	3	7	1	0	1	F	7				
34	000038	DRAM Inst. Read from 0x4028	3	5	1	0	1	0	7				
35	000038	Idle cycle	3	7	1	1	1	F	F				
36	000038	Halted	5	1	1	1	1	F	F				

This HP 16500 display shows the execution of a sequence of code that includes a breakpoint set by the JTAG emulator. The code executes out of 3-cycle first access, 2-cycle page mode access DRAM located in DRAM bank 3.

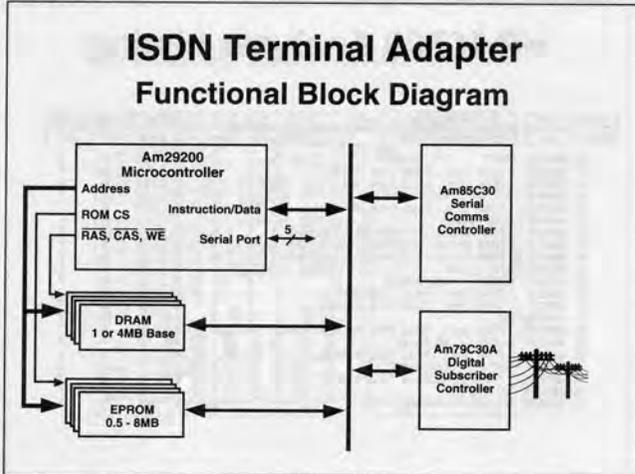
Again the STAT signals show what the microcontroller was doing in the previous cycle; a status of 4 indicates an external data access, in this case to DRAM bank 3. The *TR/*OE signal provides an output enable for read accesses. The *WE signal provides a write enable for write accesses.

The grouped signals *CAS3..0 and *RAS3..0 show the state of RAS and CAS for each of the four DRAM banks.

In this listing, a software breakpoint was set at address 0x030 in DRAM bank 3 using the JTAG emulator. The emulator replaced the instruction at address 0x030 with a HALT instruction and asserted the control signals necessary to have the microcontroller report to the emulator when the HALT instruction is encountered. Because the microcontroller is pipelined, the fetch of the HALT instruction occurs at lines 29-30 but isn't executed until line 36.

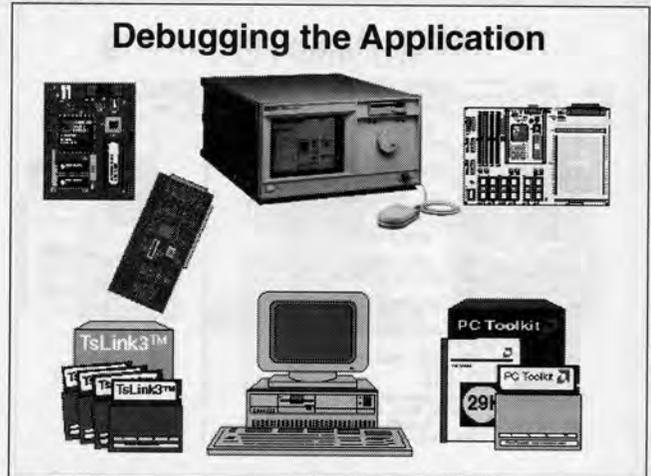
Developing and Debugging an ISDN Terminal Adapter

Slide #44



An ISDN terminal adapter has been developed using an Am29200 microcontroller, an Am85C30 serial communications controller, and an Am79C30A digital subscriber controller. This ISDN terminal adapter interfaces non-ISDN equipment to an ISDN network between the R and S/T ISDN reference points. With an ISDN terminal adapter, the user is able to send both voice and data information digitally, accruing the benefits of increased reliability, new functionality, and lower cost over current analog solutions.

Slide #45



This ISDN terminal adapter can easily be developed in the lab using an IBM-compatible PC (Am386(r) microprocessor-class or better), an SA29200 Evaluation Board and SA29200 Expansion Board, a Corelis Am29200 JTAG-based emulator, an HP 16500 logic analyzer, and the TsLink3 software.

Glossary of Terms and Conditions

Basic rate: The 192-kbps connection between the subscriber and the network. It contains two B-channels and one D-channel.

Bit Stuffing: A type of rate adaption that adds non-data dummy bits to bring the data rate up to 64 kbps. In addition, multiple channels can be multiplexed to bring up the data rate. However, bit stuffing does not support statistical multiplexing or error-checking and retransmission.

BiSYNC: A synchronous character-oriented transmission protocol. BiSYNC is used primarily by IBM. A special "sync" character precedes and ends the data being transmitted. The sync character is chosen such that its bit pattern is significantly different than the other characters being transmitted.

B-channel: A 64-kbps channel that can be used for either data or digitized voice communications.

CCITT: Consultative Committee on International Telegraphy and Telephony: The organization responsible for the ISDN standard, among others. The CCITT is part of the ITU (International Telecommunications Union), one of the oldest organizations in the United Nations.

Central office: The lowest level of switching in the public telephone network. A residential telephone or business PABX connects to the public network at a central office.

DLC: Data Link Controller: A functional block in the DSC that performs processing of Layer-1 and partial Layer-2 LAPD protocol for the D channel.

Developing and Debugging an ISDN Terminal Adapter

- DMI: Digital Multiplexed Interface:** DMI is a freely licensed specification from AT&T that contains four modes, three of which are commonly used in full data rate B channel transmission.
- DSC: The Am79C30A Digital Subscriber Controller.** An integrated chip that handles basic rate ISDN services for both voice and data.
- DSP: Digital Signal Processing:** A software method of digitally processing analog signals.
- D-channel:** A 16-kb/s channel provided by the ISDN basic rate interface. The D-channel is primarily used for call-control signaling functions. It can also be used for low priority, low-speed user data at rates up to 9600-bps.
- FCS: Frame Check Sequence:** A method of detecting errors in the LAPD protocol.
- HDLC: High-Level Data Link Control:** A bit-oriented synchronous communications protocol that uses a special bit pattern (the flag) to mark the beginning and end of data transmissions. If the data contains a flag, an extra 0 is inserted into the data stream; this is known as bit stuffing.
- IDC: The Am79C32A ISDN Data Controller:** An integrated chip that handles basic rate ISDN services for data only.
- ISDN: Integrated Digital Services Network:** An international standard for digital voice and data transmission over the switched telephone network.
- ISO: International Standards Organization:** An international organization that sets worldwide standards in telecommunications and other fields.
- JTAG: Joint Test Access Group:** The name of the group responsible for a 5-pin test and debug interface codified in IEEE 1149.1 specification. JTAG is also used generically as the name for any implementation that meets the IEEE 1149.1 specification.
- LAPB: Link Access Protocol Balanced:** A subset of the HDLC OSI Layer-2 communications protocol. LAPB is the accepted Layer-2 protocol of CCITT's X.25 packet switch specification, and establishes and maintains an error-controlled point-to-point link between a terminal and the packet network.
- LAPD: Link Access Protocol D-channel:** The OSI Layer-2 protocol defined by CCITT for use in ISDN's D-channel. LAPD can be used on the B channels as well (eg, V.120, DMI mode 3) and many times is preferred because only one software package needs to be supported for all channels.
- Layer-2 protocol:** Refers to Layer 2 (Data Link Layer) of the OSI communications model. Layer-2 converts an unreliable transmission channel into a reliable one; sends frames of data with a checksum; and uses error detection and acknowledgment. Standards that implement Layer-2 protocol include HDLC, SDLC, and BiSYNC.
- Layer-3 protocol:** Refers to Layer 3 (Network Layer) of the OSI communications model. Layer-3 transmits packets of data through a network. It is responsible for routing and congestion control. Standards that implement Layer-3 protocol include X.25.
- LIU: Line Interface Unit:** The functional block in the DSC that interfaces to an ISDN S or T reference point.
- LT: Line Termination:** Interfaces to the ISDN U reference point. The LT is located in the telephone company's switch, often at the central office. The LT performs Layer-1 functions for B and D channels, and Layer-2 and Layer-3 functions for D channels.
- MAP: Main Audio Processor:** The functional block in the DSC that provides a telephone audio interface.
- MPI: Microprocessor Interface:** The functional block in the DSC that interfaces to an external microprocessor or microcontroller.
- NT: Network Termination:** There are two types of NTs: NT1 and NT2. The NT1 acts as a repeater and performs two- to four-wire conversion (U to S interface). An NT1 deals only with Layer-1 of the OSI model. NT2s are intelligent and actively participate in the call routing/control process. PABXs and line concentrators are examples of NT2 devices. NT2 devices can be connected to multiple types of ISDN lines simultaneously. NT devices often form the boundary between equipment owned by the customer and equipment owned by the telephone company.

Developing and Debugging an ISDN Terminal Adapter

OSI: Open Systems Interconnection: A seven-layer model of communications services developed by the ISO. This layered model of communications divorces upper layers from changes in technology in lower layers.

PABX: Private Automatic Branch Exchange: A telephone exchange on the user's premises. It serves as a private central office and attaches to the public network at a central office on the network.

PIA: Peripheral Interface Adapter: The generic peripheral interface defined for the Am29200 microcontroller.

Primary rate: An expensive, high-bandwidth ISDN connection. The primary rate interface is composed of multiple B channels and one 64 kbps D channel.

R Reference point: The R reference point establishes the boundary between non-ISDN equipment and the ISDN network. Terminal adapters are used to convert the protocol used by the non-ISDN terminal to ISDN basic rate or primary rate protocol.

Reference points: CCITT-identified interfaces with established standards for both hardware and software.

S Reference Point: The CCITT designation for the connection between terminal equipment (TE) and the network terminator (NT2), or between TAs and TEs and NT1 if there is no NT2.

SDLC: Synchronous Data Link Control: An OSI Layer-2 bit-oriented synchronous communications protocol.

SNA: Systems Network Architecture: A structure of data protocols developed by IBM that predates the OSI model.

Subscriber loop: The connection between the central office or PABX and the user's equipment.

T Reference point: The CCITT designation for the connection between NT1 and NT2. If no NT2 is present, there is no T reference point.

TA: Terminal adapter: A device that connects a non-ISDN device between the R and S interface.

TE: Terminal Equipment: An ISDN-compatible device connection to the S/T reference points. A TE can be a computer, telephone, data terminal, etc.

U Reference point: The CCITT designation for the connection between the LT and NT1. Normally, a two-wire basic rate interface or a primary rate line is used, but the four-wire basic rate interface can also be used.

V.110: The CCITT recommendation for interfacing non-ISDN equipment to the ISDN using the bit-stuffing technique. ECMA 102 is the European Computer Manufacturer's Association's version of V.110.

V.120: The CCITT recommendation for interfacing non-ISDN equipment to the ISDN using packetizing techniques. V.120 uses LAPD, provides rate adaption via statistical multiplexing, and supports multiple logic connections.

X.25: The CCITT international packet-oriented protocol used primarily at Layer 3 of the OSI model. channel is primarily used for call-control signaling functions. It can also be used for low-speed user data.

A High-Performance Environment for Modelling and Simulation of Digital Systems

Silvio Forno

High Design Technology
Via Beaulard 64
10139 Torino ITALY
Phone: ++39.11.338434
Fax: ++39.11.3859967

Alberto Biondello

High Design Technology
Via Beaulard 64
10139 Torino ITALY
Phone: ++39.11.338434
Fax: ++39.11.3859967

Viscardo Costa

Italtel
Via R. Romoli
20019 Castelletto di Settimo Milanese (MI) ITALY
Phone: ++39.2.43888258
Fax: ++39.2.43888221

1993
High Speed Digital
Systems Design & Test
Symposium

Abstract

Today, digital system designers in the communication's industry need a variety of tools, sometimes manufactured by different vendors, to design their complex digital circuits. This paper describes a new design tool that integrates accurate time-domain measurement tools, with reliable modeling tools,

and fast simulation tools into one complete solution that can simulate entire digital systems which can include active components, transmission lines, connectors, etc.

Authors

Silvio Forno

Current Activities:

Silvio Forno is the software development manager of HDT. His particular areas of interest include device modelling and electronic CAD. He is currently involved in design and development of HDT products.

Author Background:

Silvio received the engineer degree (Dr. Ing.) from the Politecnico di Torino, in 1987 and Ph.D. in Computer Science in 1991. He joined HDT in 1991 after having spent the last three years in Computer Science Department of Politecnico di Torino working on telecommunication system simulation and real time systems.

A High-Performance Environment
for Modelling and Simulation of
Digital Systems



High Design Technology

Authors (cont'd)

Alberto Biondello

Current Activities:

Alberto Biondello is a member of the HDT customer support team. He works on passive and active device modelling and post-layout simulation of digital systems.

Author Background:

Alberto received the engineer degree (Dr. Ing.) in electronics from the Politecnico di Torino, in 1992, after 1 year stage at Olivetti, working on signal integrity issues in high-speed computers.

Viscardo Costa

Current Activities:

Viscardo Costa is responsible for EMC Physical Design in the Electronic Switching Division, R&D, of Italtel. His particular areas of interest include EMC design for EMI control in digital systems, and CAD and device modelling.

Author Background:

Viscardo received the engineering degree (Dr. Ing) from the Politecnico di Milano in 1986. He joined Italtel in 1986 where he has been involved in various aspects of EMC phenomena in digital circuits, semiconductor development and interconnections effects in high-speed ICs. He has published several papers on Computer Aided Design, Device Modelling and models to analyze EM emissions, susceptibility, and crosstalk for PCBs and cables.

Overview

- Emerging hardware requirements for TELECOM
- TDR/TDT measurements and behavioural time simulation
- Case study A1: High-speed PCB interconnect
- Case study A2: Pre-layout design of AT&T switch interconnects
- Complex system design & validation examples
- The new modelling and simulation environment
- Post-layout quality check

A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #1

A High-Performance Environment for Modelling and Simulation of Digital Systems



High Design Technology



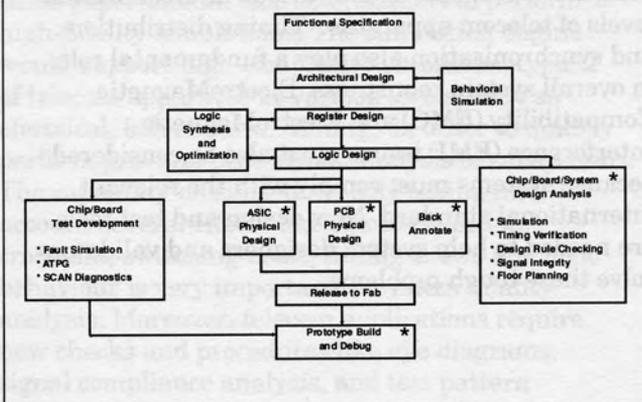
Slide #3

Overview

- Emerging hardware requirements for TELECOM
- TDR/TDT measurements and behavioural time modelling
- Case study #1: High-speed PCB interconnect
- Case study #2: Pre-layout design of ATM switch interconnects
- Complex system design & validation examples
- The new modelling & simulation environment
- Post-layout quality check

Slide #2

Simplified HS Digital Design Process



A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #4

Evolution of Telecom Systems

- Transmission standards
 - SONET
 - SDH (155 Mbit/s - 2.4 Gbit/s)
- Switching systems for BISDN
 - STM/ATM crossconnects
 - STM/ATM switches
- High-speed LAN MAN
- GSM
- HDTV
- High-density component technology
 - Fine pitch
 - MCM

The fast evolution of telecom systems, the related emerging transmission standards and the new switching system architectures push toward higher operation speeds.

The Synchronous Optical NETWORK (SONET) and the SDH standards require operating speed ranging from 155 Mbit/s (STM1) up to 2.4 Gbit/s (STM16). Broadband switches and crossconnects operating in STM (Synchronous Transfer Mode) or ATM (Asynchronous Transfer Mode) must be able to deal with these high-speed digital streams. On the other hand, the scaling down of integrated circuits and the high-density packaging and interconnection technologies like MCM (MultiChip Modules) give telecom manufacturers new opportunities to develop low to medium speed applications where miniaturisation, power dissipation, and reliability are key problems to solve. GSM (Groupe Speciale Mobile) terminals are a typical example of this kind of application.

All these new performance and quality requirements have a great impact on system design and test.

Slide #5

New Performance & Quality Issues

Signal integrity

Timing & synchronization

EMC/EMI

Hardware designers and validators have to face a new set of constraints. First of all, good signal integrity becomes a major goal, not only at transmission interfaces, but everywhere in the system. Issues like signal reflections, crosstalk, and switching noise must be controlled and kept below assigned thresholds at the various interconnection levels of telecom apparatus. Timing distribution and synchronisation also play a fundamental role in overall system robustness. ElectroMagnetic Compatibility (EMC) and ElectroMagnetic Interference (EMI) issues must also be considered because systems must comply with the relevant international standard. New design and test tools are needed to help system designers and validators solve these tough problems.

Slide #6

The Ideal Design & Test Toolset

- Wideband characterisation
- Fast & accurate modelling
- Fast & accurate simulation
- Telecom-oriented test procedures
- Full integration of previous items

The ideal tool set would require some important features. First of all, the capability of making fast and accurate characterisations of devices and subsystems with a time resolution compatible with the actual application bandwidth is important. Accurate models would be extracted quickly from these experimental characterisations to perform high-fidelity simulations. The simulation engine would support high-complexity situations, typical of telecom apparatus at various levels, such as electrical, behavioural, timing, in order to quickly perform pre- and post-layout analysis of subsystems. The capability of simultaneously taking into account several effects like signal reflection, crosstalk, switching noise, timing skews and logic behaviour is very important in system quality analysis. Moreover, telecom applications require new checks and procedures like eye diagrams, signal compliance analysis, and test pattern generation for performance verification. A strong integration of experimental characterisation, modelling techniques, simulation, and system testing is fundamental to fully achieve performance and quality goals of telecom apparatus.

Slide #7

Limitations of Many Conventional Simulation Tools for HSD Circuits

- Slow and difficult modelling procedures
- Slow & limited simulation engines for HSD designs
- Poor integration with experimental measurement

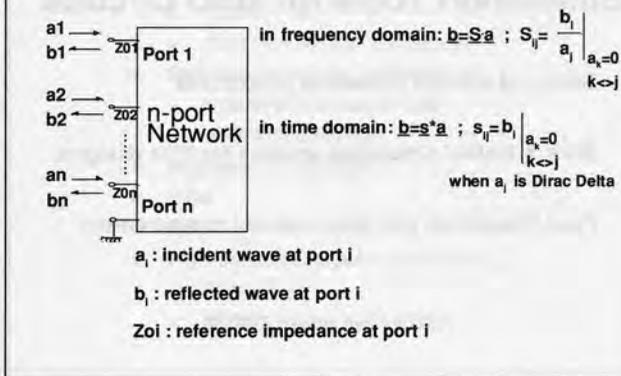
HSD = High Speed Digital

This paper will present a new set of experimental/simulative tools that overcome all the limitations of conventional tools. These new tools have strong integration with wideband test and measurement instruments, so that accurate electrical models are easily extracted. A telecom apparatus is simulated without any problem of convergence and speed, even in those complex situations (ones with tens of thousands of circuit elements are common nowadays) where, until now, it was not possible to perform even a simple analysis, especially when transmission line analysis is required.

A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #8

S-parameters Representation of Networks



The scattering, or S-parameter family is defined to relate incoming and outgoing waves at the ports of a network. In general, an n-port network has n^2 S-parameters associated, with n port reference impedances, Z_{oi} , $i = 1, 2, \dots, n$. A single reference impedance, Z_o , is usually chosen for all ports. For a linear n-port network, the defining equation at a given frequency is:

$$b_1 = S_{11}a_1 + S_{12}a_2$$

$$b_2 = S_{21}a_1 + S_{22}a_2$$

where a_i represents the incident wave and b_i the reflected wave at port i . The frequency domain S-parameters can be interpreted as reflection (S_{ii}) or transmission (S_{ij} , $i \neq j$) coefficients in matched conditions. They are, in general, complex numbers and their use is well known in microwave applications.

In the time domain, a convolution relationship applies:

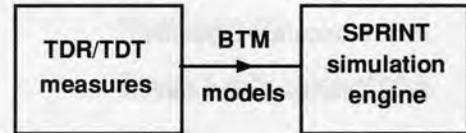
$$b_1 = s_{11} * a_1 + s_{12} * a_2$$

$$b_2 = s_{21} * a_1 + s_{22} * a_2$$

where s_{ij} is the generic reflected or transmitted wave in matched conditions when the incident wave is a Dirac Delta.

Slide #9

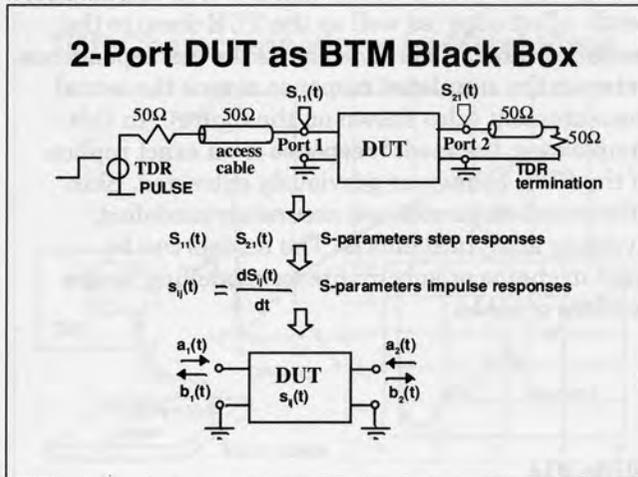
A New Approach to Modelling: BTM



A new modelling technique called BTM (Behavioural Time Modelling) is introduced. It is based on the fact that time-domain characterisation of components is the most straightforward and realistic way to get models in high-speed digital applications. Wideband TDR/TDT instruments, like those made by the HP 54120-series, are the optimal solution to perform this task, even if conventionally, they are utilised predominantly as verification tools. In fact TDR/TDT analysis of a fully terminated n-port device configuration corresponds to the measurement of the S-parameter step-response in time domain. The accuracy is good for digital applications because of the extreme precision of the TDR pulse, with amplitude aberration with respect to the ideal step limited to less than 1%. BTM directly utilises these S-parameter responses to extract models of the Device Under Test (DUT). Models obtained in this way are also suitable for most EMC/EMI applications. These wideband models can be validated by simulating the TDR/TDT set-up and comparing the simulated responses with the actual measurements.

The simulated responses of target systems, containing the modelled devices, can be also compared with the actual system measurements, performed with the same HP 54120 used as a wideband sampling oscilloscope, where the instrument's 50-GHz bandwidth helps to give extremely high confidence in the verified model. The overall procedure requires strong integration between the measurement and simulation environments.

Slide #10



Using BTM, the DUT is considered a "black box", accessible only through its ports. For instance, a 2-port linear device is fully characterised by its four S-parameter step responses: $S_{11}(t)$, $S_{12}(t)$, $S_{21}(t)$, $S_{22}(t)$. The device is reciprocal if, $S_{12}(t) = S_{21}(t)$. In the case of a symmetrical 2-port device, $S_{11}(t) = S_{22}(t)$. Symmetrical and reciprocal 2-port devices require only two S-parameter behaviour models. S-parameters impulse responses, $s_{ij}(t)$, are easily calculated as time-derivatives of step responses, $S_{ij}(t)$. When the DUT is connected to an external network, the following relationship applies between reflected waves, b, and incident waves, a, at its ports:

$$b_1(t) = s_{11}(t) * a_1(t) + s_{12}(t) * a_2(t)$$

$$b_2(t) = s_{21}(t) * a_1(t) + s_{22}(t) * a_2(t),$$

where the symbol "*" denotes the time convolution operator.

Slide #11

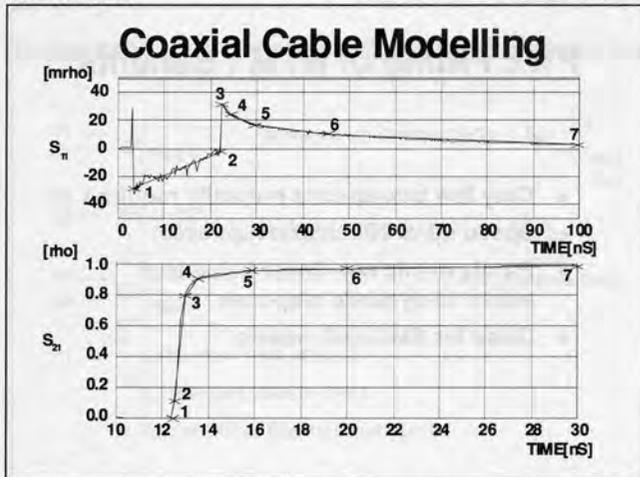
PWL Fitting of BTM : Benefits

- Only few breakpoints normally required
- Speed up of simulation runs
- Purely ohmic non-linear behaviour added to dynamic response
- Good for EMC/EMI models

Time domain simulations require time convolutions to calculate port signals when the BTM model is connected to an external network. A PWL (Piece Wise Linear) fitting of mentioned S-parameters can dramatically speed up this convolution process. As it will be shown in the following, for most situations only a few breakpoints are normally required to describe S-parameter behaviour, taking into account the accuracy constraints of digital applications (order of some %). This PWL fitting procedure is fully supported by the MCS (Model Capture System) of the graphical environment.

Another important consideration is that non-linear effects of I/O ports of digital devices (ICs) can be modelled with good approximation as purely static (ohmic) non-linearities, superimposed to a linear dynamic response. Simple examples will explain in further detail how to utilise this modelling technique in high speed digital design.

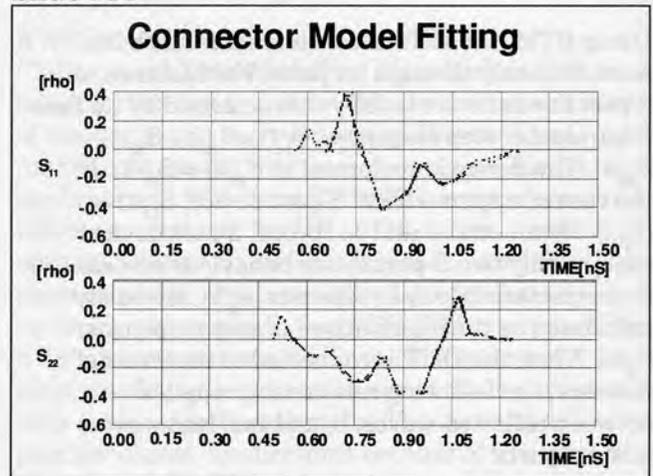
Slide #12



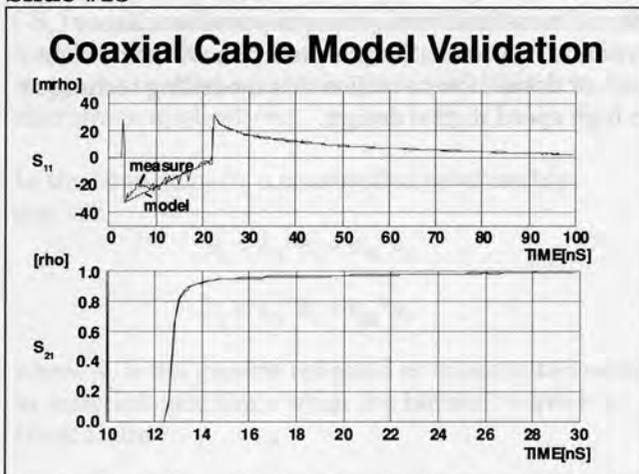
The slide shows the S_{11} and S_{21} characterisation of a microcoaxial cable 2 meters long and its related PWL fitting with 7 samples. Skin effects are clearly visible. In this case, the model is directly implemented by a 2-port block whose S-parameters are described by the PWL behaviour extracted from TDR/TDT measurements. To achieve a good model accuracy, it is important to activate the built-in normalisation algorithm provided by HP 54120 TDR before starting the characterisation process. This procedure will ensure the exact calibration of both time and amplitude (mrho) scales.

voltage generator (with 50-ohm internal impedance) sends a fast edge (as well as the TDR does) to the modelled cable. This slide shows the correspondence between the simulated response versus the actual measurement (also shown on the figure). In this simple case, the model response is an exact replica of the PWL behaviour previously extracted. Skin effects and dispersion are accurately modelled, avoiding analytical efforts. The models can be used in chains or subcircuits for modelling longer sections of cable.

Slide #14



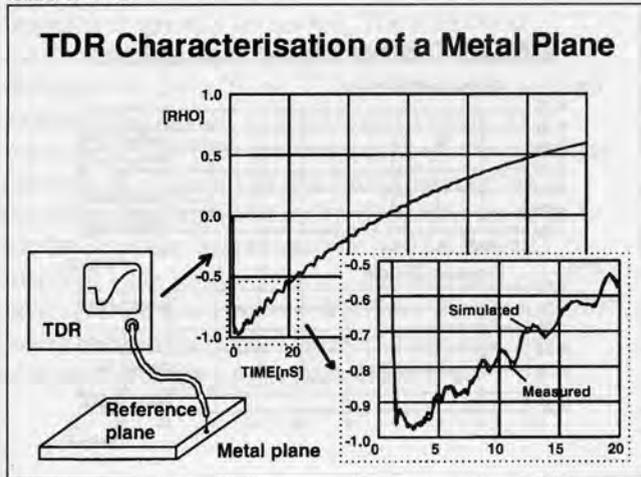
Slide #13



The behavioural model can be validated through a simulation of the measurement set-up, in which a

The methodology can also be applied to model asymmetrical devices (for example connectors), whose structures are very difficult to treat in terms of lumped parameters because of their discontinuities. A Behavioural model is more accurate and easier to build. In this case, 3 S-parameters (S_{11} , S_{21} , S_{22}) need to be known because the device is not symmetrical. The slide shows the reflectometer response for S_{11} and S_{22} and related PWL fitting of a PCB connector. The fitting starts after the first peak, which is the parasitic effect due to the launch cable, to the point where it is joined with the device under test. This portion of the response can be ignored.

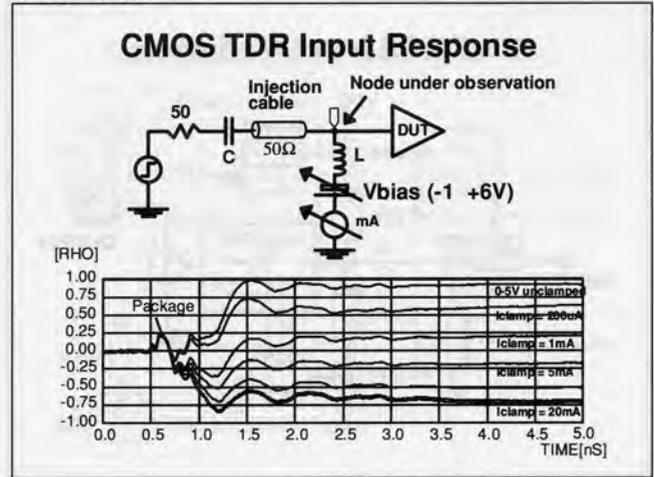
Slide #15



For fast edge operation, the power and ground distribution planes cannot be considered ideal. In fact, the current injected in a particular point of the plane (for example, by a switching driver or a termination) and its propagating phenomena cause noise that can affect other devices placed on the same substrate. Using a mesh of behavioural blocks, it is possible to build up two dimensional models of power and ground planes and take this effect into account. The TDR is an excellent vehicle of validation of this model.

The slide shows the comparison between the TDR measurement of a two-layer metal plane (the second plane acts as reference plane) and a simulation of the model in the same configuration. The global behaviour is roughly the typical reflectometer response of a capacitor. A detail of the first section of the graph shows the reflections of the TDR step due to the plane boundaries. It is interesting to point out the very good matching between measured and simulated results. Changing the signal injection point causes a strong modification of these reflections. Accurate ground and power plane modelling is fundamental to simulate the residual switching noise on multilayer PCBs or MCMs taking the effect of decoupling capacitors into account. Model parameters are optimized through a trial and error technique comparing the actual measured response with the simulated TDR response of the model. This process is fast due to short simulation time required.

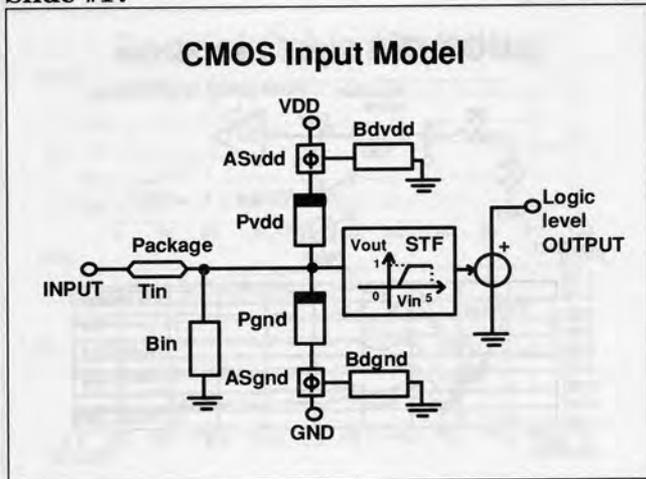
Slide #16



The BTM methodology described before for passive components applies as well for the modelling of the I/O interfaces of active parts. Using the TDR, it is possible to create an accurate model of the dynamic behaviour of an input or output in both normal operating conditions (within the logic swing) or non-normal conditions, such as, in clamping situations, including the package effect (lossy or lossless). Other modelling approaches (as simple Thevenin equivalents) are inaccurate or too slow (for example, SPICE MODEL cards) to be used for the simulations of large systems.

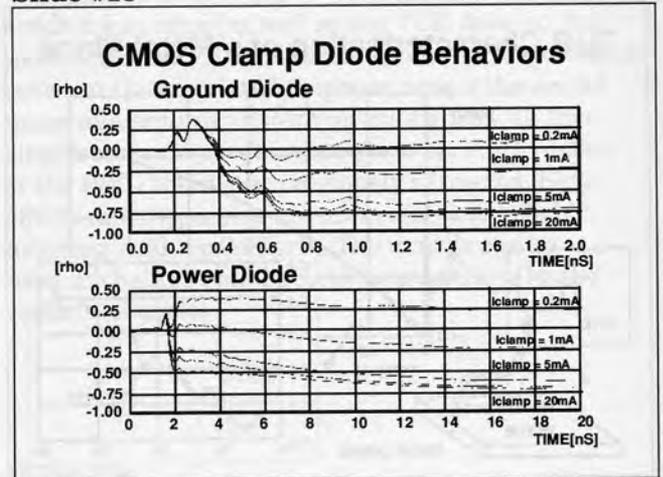
The measurement setup consists of a HP 54121T TDR and a bias generator. The 200mV step provided by the TDR pulse generator can be considered a small amplitude signal compared with the 5 V swing of CMOS output levels. A X10 matched attenuator, connected at the pulse generator output, will provide a "small" amplitude stimulus (20 mV) for low swing devices (such as, ECL). The inductor, L, has been inserted in order to present a high impedance path to the TDR pulse on the biasing stub. The capacitor, C, acts as a DC block in order to avoid injections of direct current into the TDR pulse generator and is practically a short circuit in the time window of the measurement. Obviously both L and C must have low parasitics in the frequency range of interest. Anyway, their non-ideal effects can be taken into account during simulation and model validation, effectively removing their effect by normalisation.

Slide #17



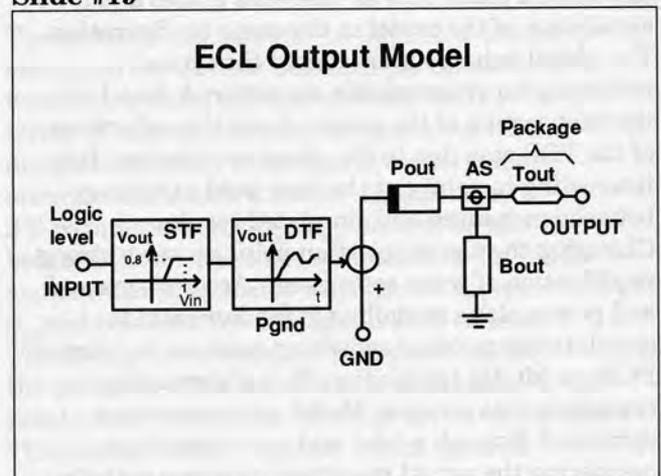
Few model architectures can represent all the major component families. For example, this slide shows a typical CMOS input model. The static characteristic is modelled for both the clamping diodes (Pvdd and Pgnd) and can be obtained by a V/I measurement, while the dynamic behaviour of the clamping diodes and the input in normal conditions is measured using the TDR. The Bin, Bdvdd, and Bdgnnd elements are described directly by measured samples or by their PWL fitting. The Bin block models the input of the device when both clamping diodes are off. The Series Adaptor blocks, ASvdd and ASgnd, are utilized to connect the behavioural blocks, Bdvdd and Bdgnnd, which take package effects into account, in series to power and ground nets, respectively. The Static Transfer Function (STF) can be utilised to convert the external analog levels to internal digital ("0", "1") levels. A set of library functions (AND, OR, etc.) can model internal logic and timing behaviour.

Slide #18



The slide shows two families of reflectometer responses related to power and ground clamping diodes versus clamping current of a CMOS EPROM input. It is interesting to point out the different behaviour of the two diodes. The ground diode shows a very fast response, so that low impedance levels are reached in about one nanosecond. On the contrary, the Vdd diode shows a slow transient (several nanoseconds long) before it reaches low impedance levels.

Slide #19

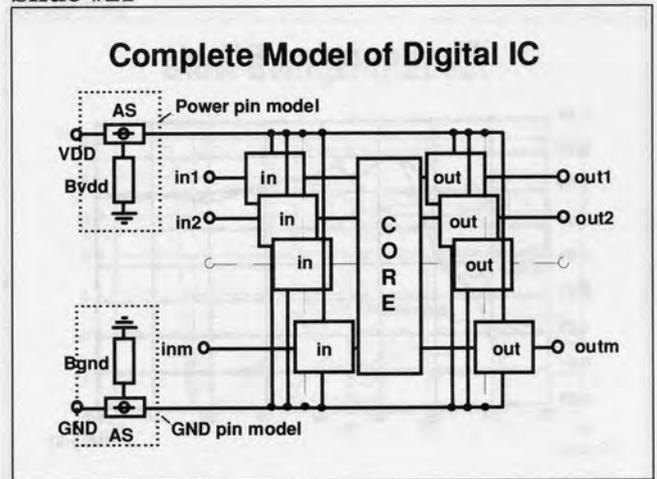


An ECL device presents a strong output resistance non-linearity at low current loads. During the falling transition, there are situations where the

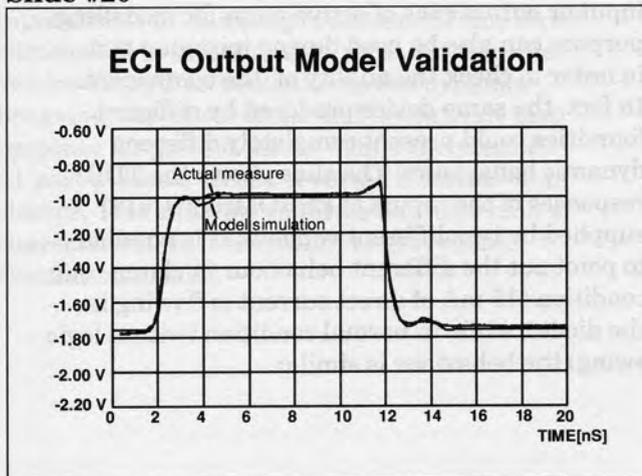
A High-Performance Environment for Modelling and Simulation of Digital Systems

output transistor goes near cutoff and its output impedance greatly increases. The reflection coefficient's dynamic behaviour (Bout) is usually the same for both "0" or "1" logic states, as well as the static output characteristic (Pout). The near inductive TDR response of the output emitter follower is modelled by the Bout block series connected through the series adaptor, AS. The Static Transfer Function (STF) translates the internal logic levels, "0" and "1", to output electrical levels, while the Dynamic Transfer Function (DTF), measured by a digital sampling oscilloscope, shapes the output waveform behaviour.

Slide #21



Slide #20

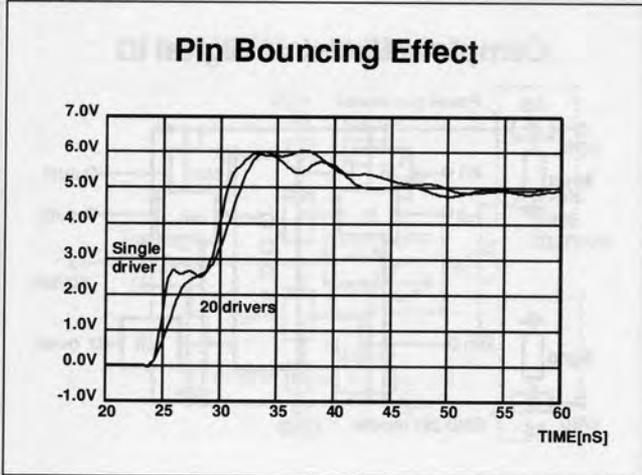


A model of an ECL output of a BiCMOS ASIC is validated by comparison with the actual measurement. The output load is 50 ohm to -2 V, so that all the elements of the model, shown in the previous slide, concur to determine the resulting waveshape.

All the previously mentioned models can be combined to build up a macromodel of an entire device composed of an input section, output section, a core section (at logical levels describing the logic/timing function of the device), and a behavioural model of power and ground pins created by means of TDR measurements. This model can also be used for an accurate simulation of the simultaneous switching noise, because all the input and output pins are coupled together by common power pins.

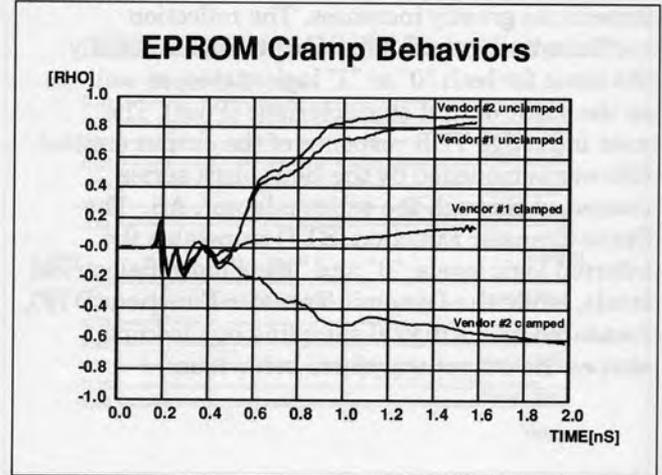
A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #22



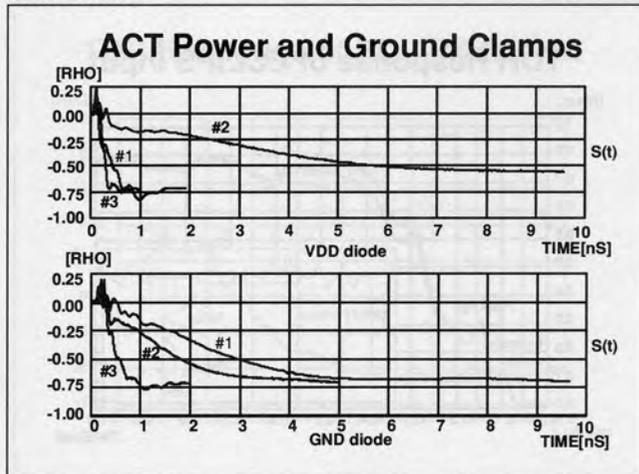
This slide shows the comparison between two simulations of an actual PCB interconnect with and without Pin Bounce effects. The driver's package contains 20 CMOS outputs driving 5 input loads each. The driver's model is of the type shown in the previous slide, where package coupling is modelled by behavioural blocks, Bvdd and Bgnd, obtained from TDR measurements on power and ground pins, respectively. In this case, the pin bounce determines a slow down of waveforms at receivers because the outputs switch simultaneously in the same direction. In a more general situation with random rise and fall transitions, a jitter effect would appear at these receivers. The amount of pin bounce strongly depends also on power and ground distribution nets, so that accurate models of them are needed.

Slide #23



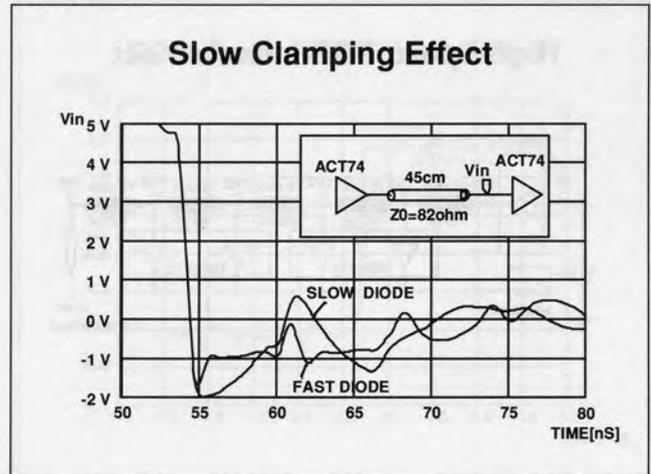
The same TDR measurements performed on input or output pins of active parts for modelling purpose can also be used during incoming test, in order to check the quality of the component. In fact, the same device produced by different foundries could present completely different dynamic behaviours. The slide shows the TDR responses of the inputs of EPROMS (27C512) supplied by two different vendors. It is possible to point out the different behaviour in clamp condition (15 mA of direct current is flowing in the diode), while in normal condition (within logic swing) the behaviour is similar.

Slide #24



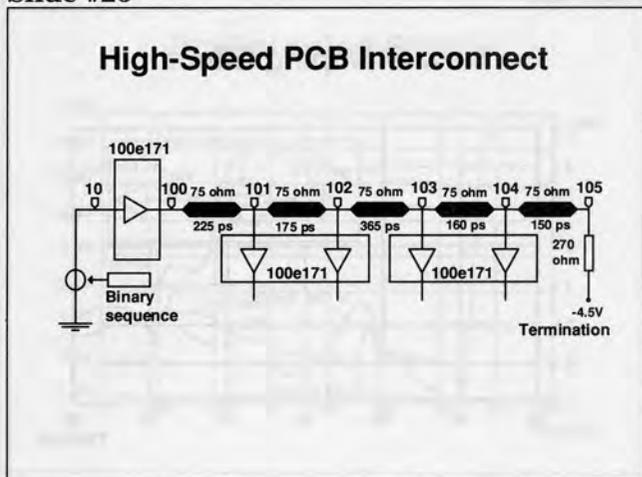
The slide shows the dynamic behaviours of ACT74 input diodes in clamping condition (20 mA direct current). It is interesting to point out that the vendor #2 part presents the slower dynamic behaviour for power diode and a relatively fast response for ground diode. While the vendor #1 part presents a fully reversed behavior of its clamps. This example confirms how this kind of characterisation is necessary to get reliable simulation results.

Slide #25



This slow behaviour has a great impact on the clamp action, so that if a voltage overshoot, due to reflection occurs, it will be effectively clamped only after the delay observed in the TDR characterisation. Due to the difficulties to forecast these effects, only this modelling approach, based on experimental measurements, can accurately take all the effects into account. The slide compares the simulation of a point-to-point interconnection 45 cm long using an ACT74 as receiver with the same static characteristic but different dynamic behaviour in clamping condition (vendor #1 and #2).

Slide #26

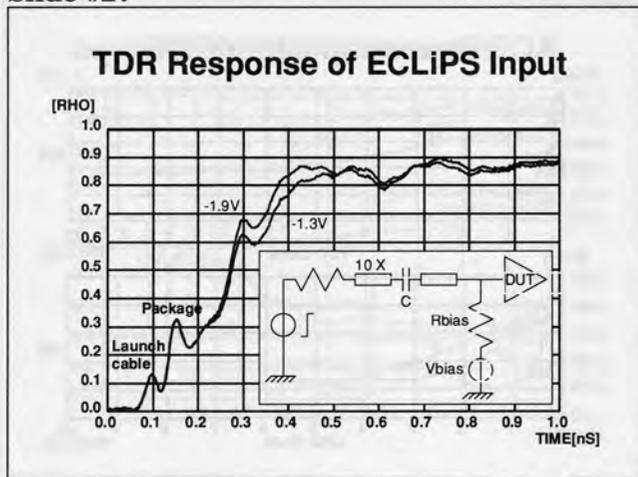


A PCB (Printed Circuit Board) interconnection among subnanosecond ECL (ECLiPS) devices has been utilised as simple example to show a typical high-speed design application.

The situation shown consists of an ECLiPS 100e171 multiplexer driving 4 receivers connected to a PCB bus. The receivers belong to two separate packages. All drivers are packaged in plastic chip carriers (20 pin PLCC) and are connected to the breadboard by sockets. The input stimulus can be generated internally by means of a simple loop oscillator, generating a 192 MHz clock signal, or externally by means of a high-speed pattern generator. The bus termination consists of a single pull-down resistor connected to VEE supply (-4.5 V), or a Thevenin termination. The single pull-down leads to an unterminated situation that has been analysed in order to check the effectiveness of the method, even in marginal operating conditions.

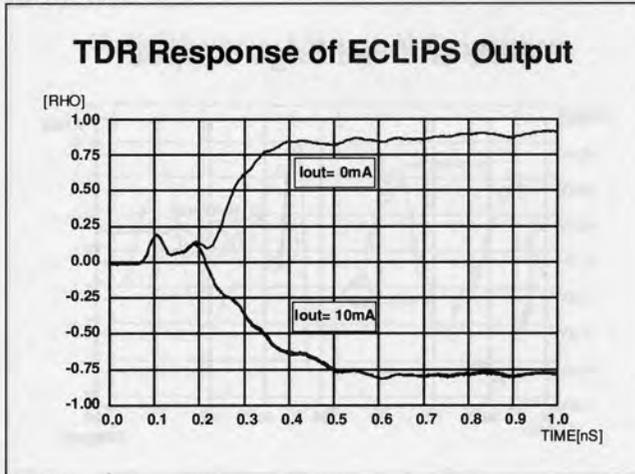
This example is used in the demonstration case study.

Slide #27



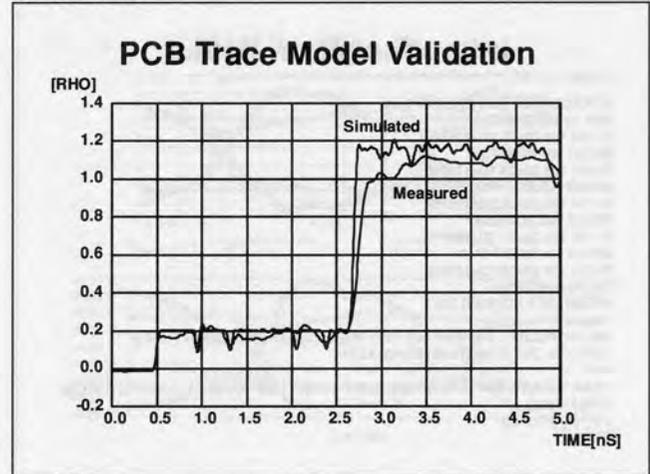
Shown is the behaviour of the measured reflection coefficient of a socketed 100e171 input (pin 6) at two biasing levels (-1.9 V corresponding to a logic "0" level and -1.3 V corresponding to VBB). These two responses are slightly different because at VBB the input draws current. The effects of launch cable and package assembly are clearly visible. It is easy to separate packaging effects from the active input contribution because the package response is independent of the bias levels. The resistive effect of biasing network (R_{bias}) is also visible because the asymptotic value of the input reflection coefficient is lower than 1. This effect can be cancelled by resimulating the setup connecting a negative resistance ($-R_{bias}$) in parallel to the behavioural block described directly by the acquired samples. From this new response a BTM model is extracted using a PWL fitting.

Slide #28



The output reflection coefficient of an unsocketed 100e171 is measured at several bias conditions. If the output current is constant (10 mA), the behaviour at both logic levels are practically the same. At 0-mA current, the emitter follower output shows a high-impedance to the TDR pulse. Launch cable connection and packaging effects are again clearly visible, because they do not depend on bias levels and these effects can be discarded when the model is built up. Output rise and fall edges, at constant current (10 mA), are also acquired because they practically represent the output model waveshape in unloaded conditions. Both input and output models are completed adding their non-linear static characteristics, modelled by a PWL resistor. The pairs of values (v, i) can be obtained by an automated power supply and precision digital multimeter setup.

Slide #29



A simple lossless model of interconnecting traces, including the short stubs toward receivers, is obtained extracting geometrical data from board layout. The microstrip delays and impedances are calculated starting from crosssection data through standard formulas. This slide shows the TDR model response of the trace left open without the device connected, compared with the actual response. There is a good match regarding delays of various trace pieces. Some impedance discontinuities due to geometrical changes in the crosssection are present, even if the average impedance level is in good agreement with theoretical value. Skin effect losses cause the typical slow down (about 80 ps) of TDR edge reflected by the open end of the interconnect. This lossless transmission line model can be utilised for preliminary interconnect simulations. A more realistic model can be obtained substituting lossless lines with lossy lines, modelled as 2-port S-parameter blocks with the BTM technique.

Slide #30

Interconnection Netlist

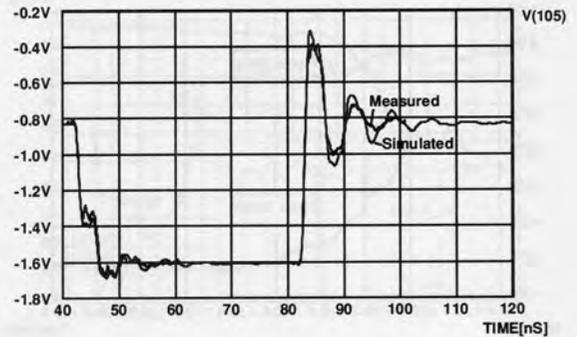
```

*****
INTERCONNECTION DESCRIPTION
XDR 10 100 DRE171
T1 100 101 Z0=75 TD=225PS
XREC1 101 RCE171P
T2 101 102 Z0=75 TD=175PS
XREC2 102 RCE171P
T3 102 103 Z0=75 TD=365PS
XREC3 103 RCE171P
T4 103 104 Z0=75 TD=160PS
XREC4 104 RCE171P
T5 104 105 Z0=75 TD=150PS
* pull-down resistor
VTERM 105 0 DC(-4.5V) 560
* edges
VIN 10 0 PULSE(1 0 40N 0 0 40N 80N) PSEQ (101000101111100001111000)
.OPTIONS DELAYMETH=INTERPOLATION
*****
.TRAN TSTART=40N TSTEP=10PS TSTOP=120NS LIMPTS=1000 V(100) V(101) V(102)
V(103) V(104)
+ V(105) I(T1,100)
.END
    
```

The simulation's netlist of the whole interconnection network is shown in the slide. Driver and receiver models are described as SPICE-like subcircuits. The driver model includes an input port that only has logical meaning when using 0 V and 1 V as input levels. This feature is very useful to define stimuli composed of sequence of "0" and "1", identified by PSEQ keyword in the PULSE statement extension of the simulator.

Slide #31

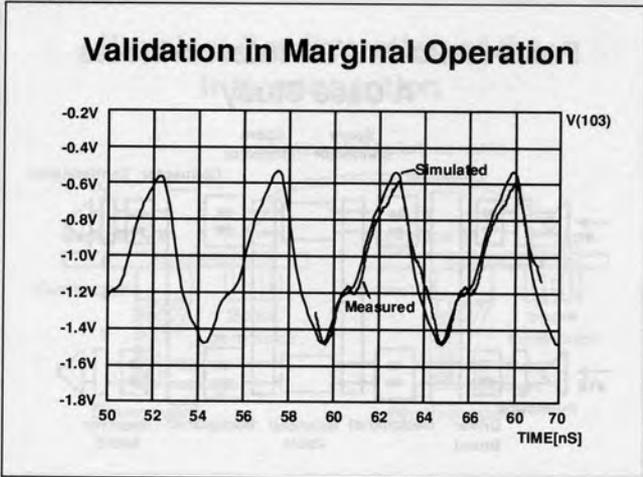
Rise and Fall Edge Shapes



The whole interconnection model has been simulated for various termination values. A typical simulation run requires about 1.5 s on a HP 750 workstation. As shown in the slide, a 560 ohm to -4.5 V termination causes a strong difference between falling and rising edges at the receivers. The falling pedestal edge is affected by a round-trip delay at -1.4 V, due to output transistor cut off, while the rising edge is followed by a great amount of ringing due to termination mismatch. Using a 270 ohm to -4.5 V termination, the operation becomes quasi-linear, so that the -1.4 V pedestal disappears, but both edges are followed by a consistent amount of ringing due to termination mismatch. To verify the accuracy of the models and the simulation, a comparison with measurements is performed on the actual breadboard using the HP 54120 digital oscilloscope. The results obtained are shown in this slide, where it is possible to point out the good match between measurement and simulation, even in these critical functioning situations.

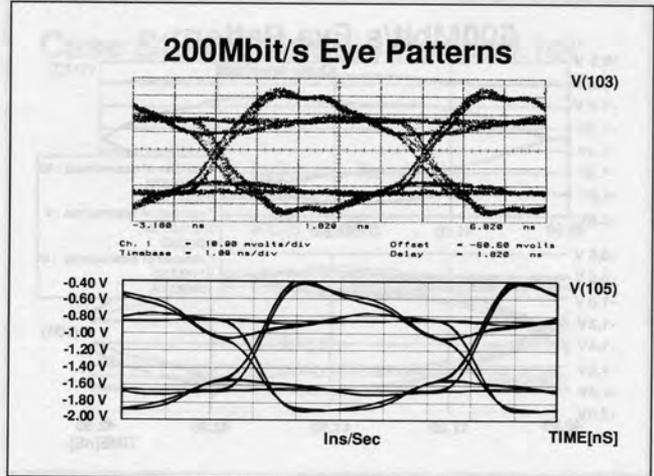
A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #32



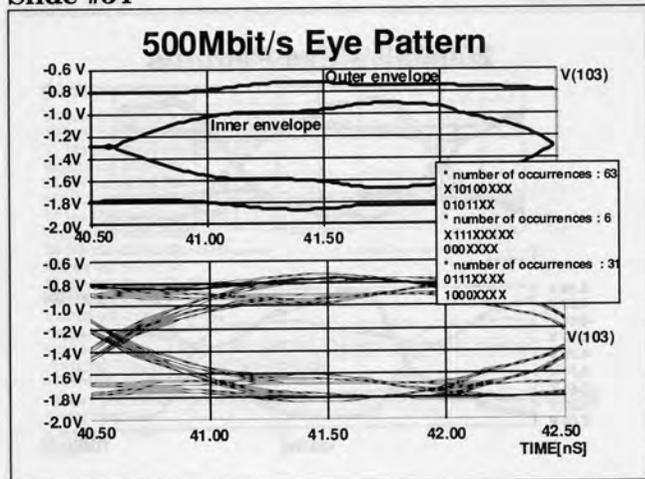
This slide shows a comparison between the measurement and the simulation when a 560 ohm termination and an input stimulus of an 192 MHz 0101 sequence, obtained from the internal oscillator, is used. In this non-linear situation, the edge difference and interaction cause a strong shift towards ground of the waveform swing, so that the waveshape is completely unacceptable. Even in this critical situation, the match with actual measurement still remains good.

Slide #33



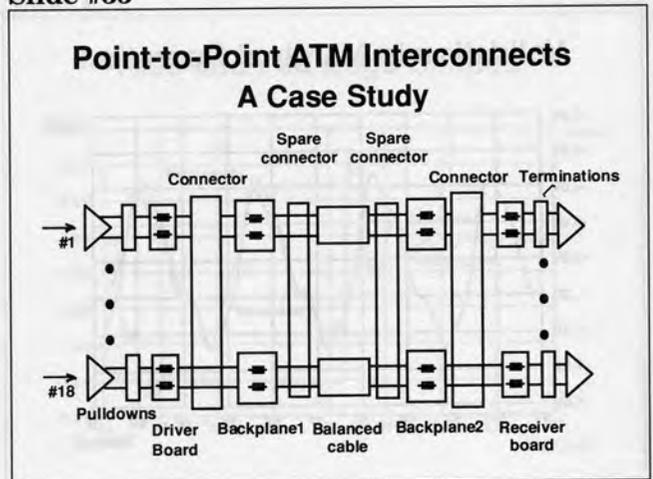
To obtain better operation, a lower termination value (270 ohm to -4.5 V) is evaluated. Measured and simulated eye diagrams at the input of third 100e171 input (V(103)) are compared in this slide. The input stimulus is an 63-bit NRZ Pseudo Random Binary Sequence (PRBS) at a 200 Mbit/s data rate. The dissymmetry affecting the eye's outer shape, due to the non perfectly linear driver operation, caused by the still high termination value, is still evident.

Slide #34



A higher speed (500 Mbit/s) operation is analysed with the interconnection terminated in a 75 ohm to -2.5 V Thevenin equivalent. In this case, the driver operation is linear, so that the eye pattern diagram can be used to determine the worst-case pattern that is causing the maximum eye closure at 500 Mbit/s. This pattern is then used as input stimulus for both interconnect model and actual breadboard. In this last case, the worst-case sequence is loaded on the high-speed pattern generator, such as the HP 80000A data generator system and injected at 500 Mbit/s at the input of driving 100e171. The eye diagram is compared with simulated eye diagram in this slide. The eye diagram mask facility is very useful to quickly evaluate the digital bandwidth of a system, starting from its simulated or measured step response.

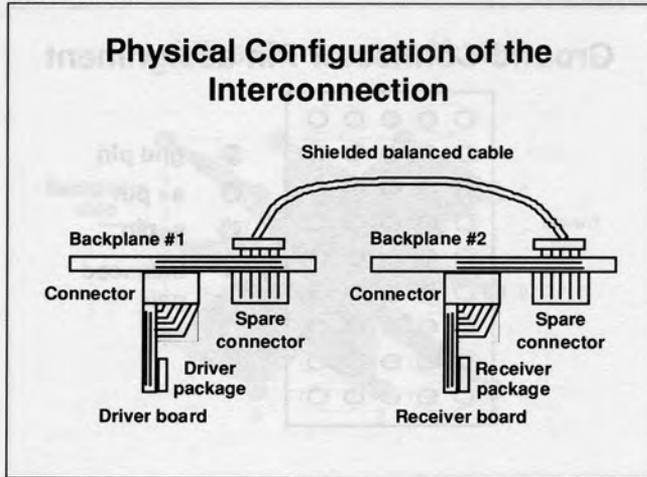
Slide #35



In telecom design, pre-layout analysis plays a very important role because of the reduced noise and timing margins (skew, settling time, etc.). Pre-layout analysis should be considered very early during the design phase: in this way errors in the physical system architecture can be avoided; consequently, improving quality and reducing cost.

This slide depicts a point-to-point transmission for an ATM (Asynchronous Transfer Mode) cross-connect system. Inter-stage connections of a broadband cross-connect require throughput on the order of several Gbit/s. Given the number of signals involved, a model related to a three-link module (18 wire) can be developed to verify the overall noise coupling effects.

Slide #36

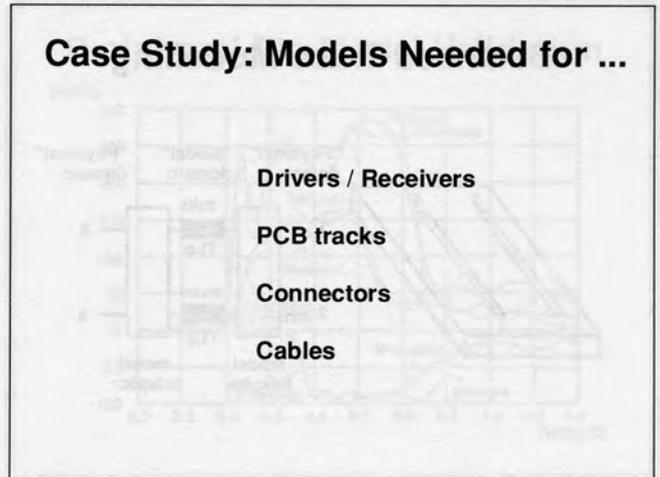


This slide shows the physical configuration of the interbackplane connection. Differential transmission was chosen to get the best result in terms of eye opening at the working bit-rate (155 Mbit/s). In fact, fully a balanced configuration ensures good common-mode rejection and minimises the effects of simultaneous switching noise.

A shielded balanced cable of between 3 m and 9 m in length is used to connect the backplanes. A spare connector is provided to replace the electrical link with an optical one when the connection length is greater than 9 m.

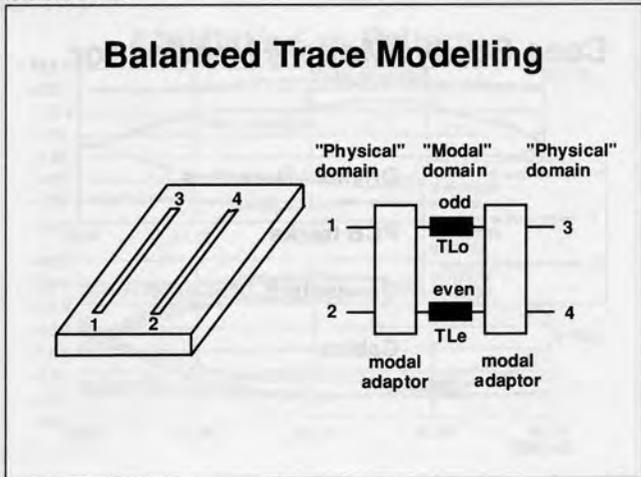


Slide #37



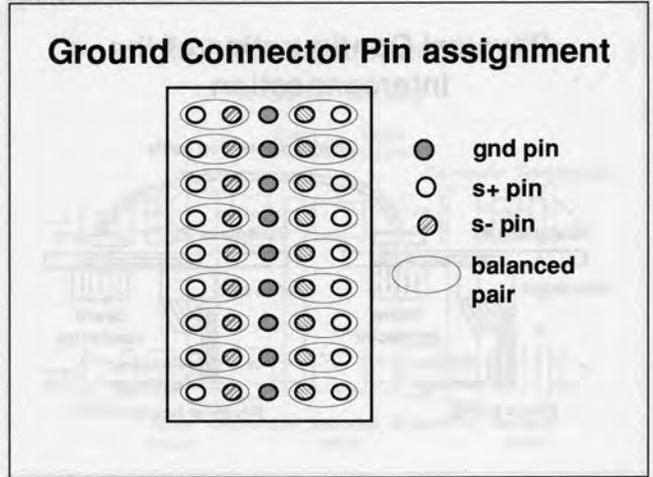
Several components have to be modelled in order to complete the simulation. Because active components (ECLiPS) have fast transition times (about 300ps), it is necessary to use very accurate time domain models of all component parts. Modelling techniques for ECLiPS components have been explained in the previous example.

Slide #38



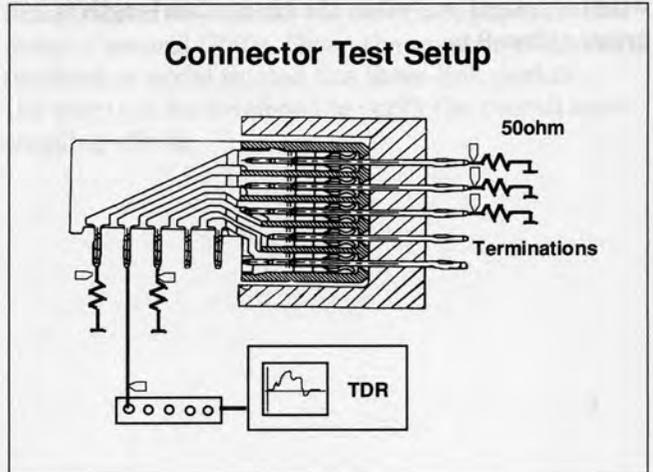
Each differential signal pair is carried by two coupled microstrips on the PCB. To model these traces, a transmission line model, based on modal analysis, is used. Two-mode decomposition is obtained using modal adaptor blocks. A modal adaptor converts physical waves of a multiconductor transmission line into modal waves according to the modal transformation, defined by its eigenvalues and eigenvectors. The simulator uses the modal adaptor as a primitive component. In the case of two coupled microstrips, only two modes (even and odd) are present. Transmission lines between modal adaptors represent even and odd mode propagation. Due to their short length (a few centimeters), you can assume these lines are lossless.

Slide #39



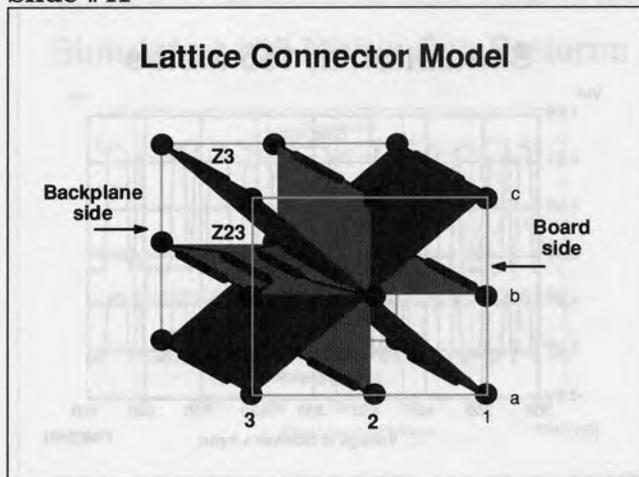
Considering the large number of signals and the use of balanced transmission, requiring two pins per link, a high-density PCB metric connector is usually chosen. It is necessary to position ground pins efficiently in order to preserve signal integrity with the minimum number of extra pins.

Slide #40



The first step is to build up an accurate connector model. A connector module of 9 rows and 5 columns is characterised by means of TDR/TDT techniques. Coupling voltages due to the TDR step are measured at adjacent pins. Direct coupling to non-adjacent pins is found to be negligible.

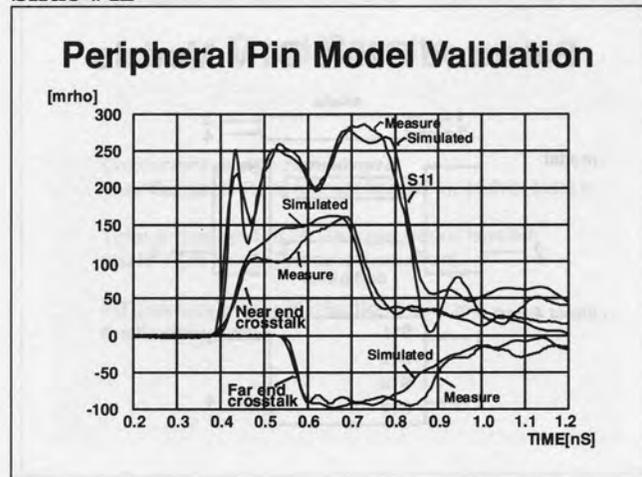
Slide #41



Starting from the previous statement, a lattice structure is selected to model the connector. Each pin has a direct coupling with adjacent pins. Due to fast operating edges and connector propagation delay (about 150 ps), a transmission line model (TLM) is the most suitable choice. This slide shows how this model couples each pin to its adjacent pins. The distributed intra-pin coupling is represented by a balanced transmission line while an unbalanced line models the coupling toward the reference ground plane.

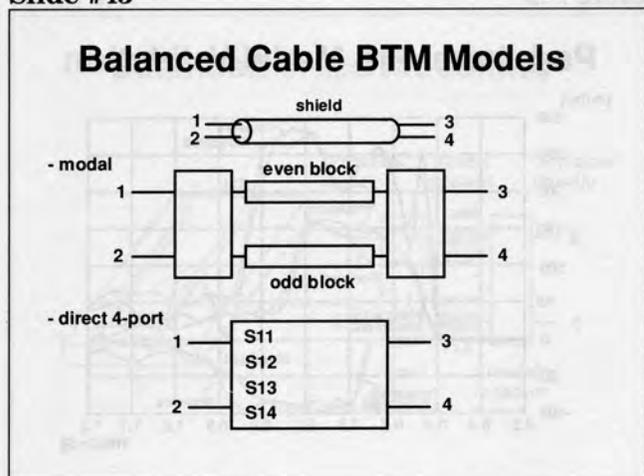
Using a trial and error fitting technique, the values of the model parameters are optimised through simulation of the experimental set-up. Despite the number of elements required to build up this 3-D model (about 700), the simulation runtime is very fast (about 2 seconds on an HP 750 workstation).

Slide #42



A comparison between measured and simulated TDR responses of two peripheral pins is shown in this slide. Transmission line effects are clearly visible, as is crosstalk behaviour. The lossless TLM models ensure good accuracy that cannot be achieved using conventional RLC models. A further accuracy enhancement can be obtained if lossless lines are replaced by lossy two-port S-parameters blocks.

Slide #43

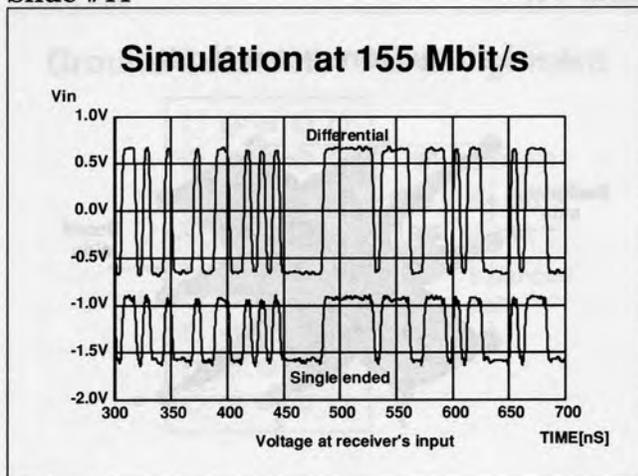


Because this application requires cable lengths of up to 9 m, cable models must include losses because they can have a significant impact on the fast signal edges (about 300 ps). The 4-port model of the balanced cable can be obtained in two ways. The first is through a modal decomposition, by means of two modal adaptors and two lossy lines representing even and odd propagation. Each line can be modelled using a PWL fitted 2-port S-parameter block.

Alternatively, a 4-port S-parameters block is used. Due to reciprocity and symmetry of the cable, only 4 different S-parameters are needed. This last model can be directly obtained from four unbalanced TDR/TDT measurements.

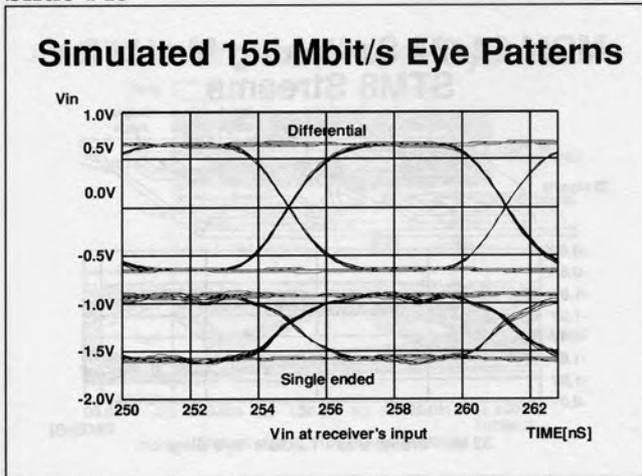
Both models are referred to the minimum interconnect length (3 m). Greater lengths (6-9 m) are obtained simply connecting two or three equal 3-m models in cascade configuration.

Slide #44



A complex model of the interconnection module (18 paths) is built up, in order to take all signal degradation effects into account. This model containing about 7,000 elements is stressed in simulated operating conditions; injecting 18 different 155 Mbit/s patterns into the parallel paths to pinpoint the effect of reflections, cable losses, and connector crosstalk. A typical result of these tests for a cable 9 m long, terminated at the receiving end, is shown in this slide.

Slide #45



Despite the fact that single ended signals at the receiver appear to be affected by a significant amount of noise, the differential responses are good thanks to the cancellation of common mode noise, as shown by the eye-diagrams in the slide.

Slide #46

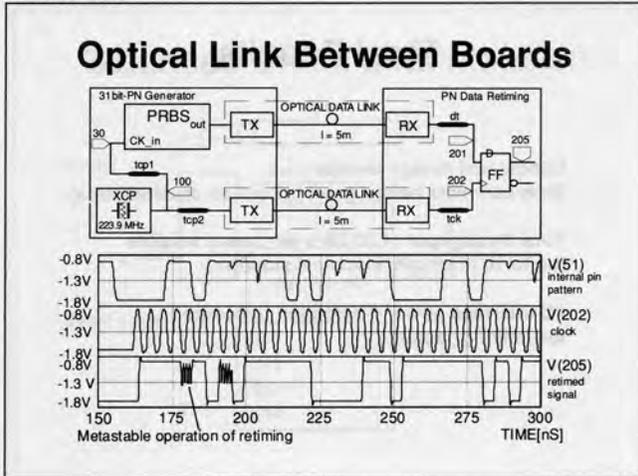
Final Results

- Concurrent design needed
(interconnects have a great impact on architectures)
- Total throughput of 30 Gb/s per board feasible
(up to 10 m length, electrical solution)
- I/O interface power dissipation becomes
the limiting factor

The results of the previous case study demonstrate the feasibility of an electrical backplane interconnect at a data rate of 155 Mbit/s per channel. The use of high-density connectors allows a global I/O throughput of about 30 Gbit/s per board, assuming that about 500 pins are used for high speed ATM streams (each balanced channel requires 2.5 pins). The major issue becomes power dissipation of I/O interfaces. The need to address interconnect problems as early as possible in the design phases (concurrent design) has again been emphasised in this case study.

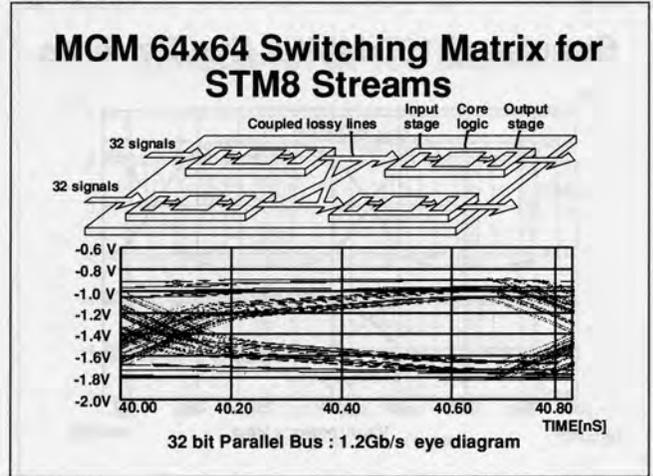
A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #47



Optical interconnects are playing a major role in telecom apparatus because they may overcome some of the bottlenecks of their electrical counterparts. This slide shows a digital subsystem including a 223.9 MHz quartz crystal oscillator, a Pseudo Random Binary Sequence (PRBS) generator board, and a retiming board connected at backplane level by two optical data-links carrying clock and data streams, respectively. The optical modules include an LED driver, a 1st-window LED (HP BR1414), a PIN-preamplifier module (HP BR2416), and an ECL post-amplifier. The optical cables are 50/125 μm , 5-meter long multimode fibers. The entire sub-system has been modelled and simulated at the electrical, timing, logical, and optical levels to check its overall performance. The faults due to timing problems, like metastable states, can be also easily pinpointed. About 30 seconds are required on an HP 750 workstation to carry out the whole simulation (1,300 elements x 6,000 timepoints).

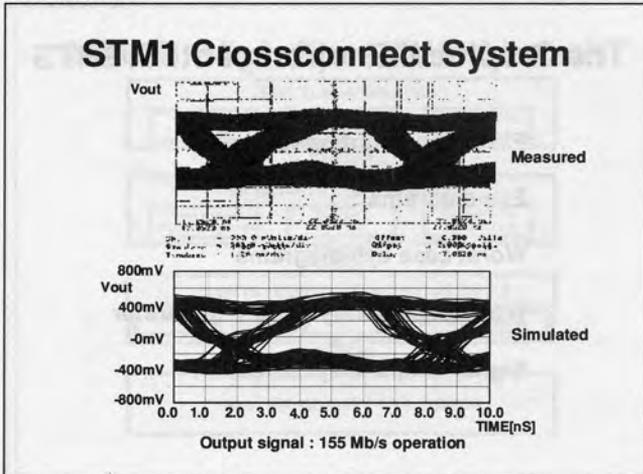
Slide #48



This slide shows a Telecom application regarding a 64x64 crosspoint switch operating at 1.2 Gbit/s implemented on MCM-D technology. All lines are modelled behaviourally as lossy interconnections taking into account skin effect and DC losses. Simultaneous switching and non-ideal power distribution are also modelled. Parallel buses are treated as lossy coupled lines, and crosstalk due to modal velocity differences is also taken into account. All models are obtained through TDR/TDT measurements of actual devices, using a deembedding procedure to discard the effects of test-fixtures and packages. Each bare-die holds 32 simultaneous switching drivers and 32 receivers. Core-chip functions are also implemented as VCVS generators. The overall model complexity is about 5,000 elements.

A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #49

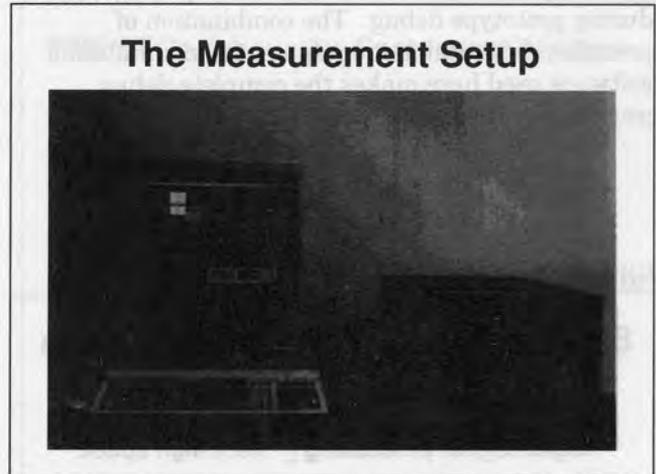


For telecom system validation, the BTM methodology allows the simulation of very complex systems with various input patterns. This slide shows the 155 Mbits simulated eye-diagram at the output of a digital crossconnect compared to the actual measurement.

The crossconnect system is composed of several switching boards, placed in different racks connected together and to STM1 peripheral ports.

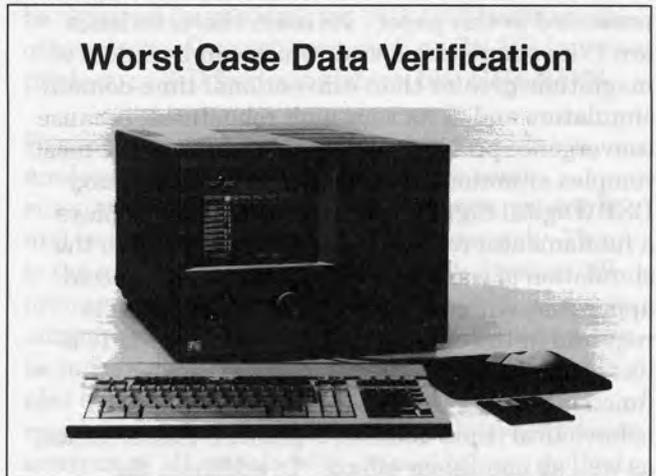
The complete model of the crossconnect counts more than 50,000 electrical/behavioural elements (switching elements, boards, connectors, cables). The 32 simultaneous input sequences are composed of 64 random bits each. All signal degradation effects, including pin bouncing and timing skews of each crossconnect, are taken into account. Simulation time of the entire system for 16,000 time points is about 1 hour on an HP 750 workstation.

Slide #50



Shown above is a picture of the measurement setup used in this paper. The setup contains an HP 54120-series high-bandwidth digital oscilloscope, an HP 750 workstation, and the modelling and simulation software by HDT, described on the next several pages.

Slide #51



SPRINT produces a worst case bit sequence from its simulation results that will cause maximum eye-closure of the simulated or measured signal. In the case study an HP80000 Data Generator System, as shown in the above photograph, was loaded with this bit sequence. It was then used in conjunction with an eye-diagram measurement made on an HP54120 oscilloscope to verify the performance of the ATM prototype.

A High-Performance Environment for Modelling and Simulation of Digital Systems

This stage is vital to verify real-world performance during prototype debug. The combination of precision measurement equipment and simulation software used here makes the complete debug process extremely effective.

Slide #52

SPRINT: The New Simulation Engine

Digital signal processing (very high speed

No convergence problems

Supports RLC, TLM, and BTM models

High complexity net simulations
(limited only by available workstation RAM)

SPRINT is the simulator used in all the examples presented in this paper. Its main characteristics are 1) its simulation speed, which can be orders of magnitude greater than conventional time-domain simulators and 2) its very high robustness, because convergence problems are avoided even in the most complex situations. Its speed, due to proprietary DSP (Digital Signal Processing) algorithms, plays a fundamental role in model validation and in the simulation of complex systems, typical of telecom apparatus, where a high level of interactivity is required in the optimisation process. The SPICE-like system description can include transfer functions taking into account logic, timing, and behavioural (time-domain, S-plane, Z-plane) issues, as well as non-linear effects. In addition, the simulation can be carried out simultaneously at all levels. Behavioural time models, as well as other modelling techniques, like TLM (Transmission Line Modelling), effectively take into account propagation effects, along with conventional RLC equivalent circuits, and are fully supported by the SPRINT simulation engine.

Slide #53

The Graphical Environment: SIGHTS

PWL fitting supported

Eye-diagrams

Worst case eye-diagrams

Worst case binary pattern generator

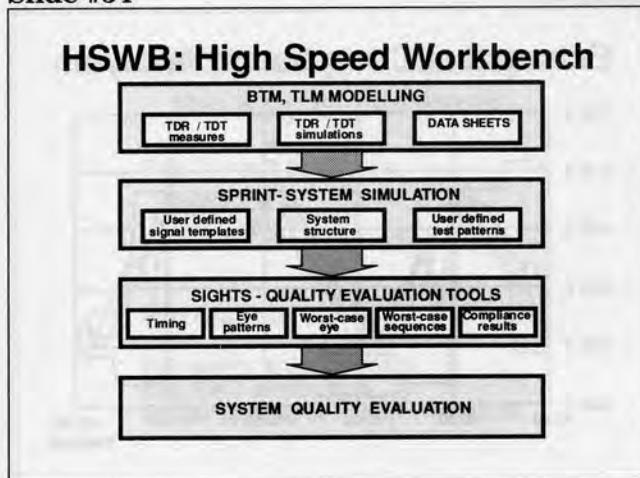
Signal template specification

SPRINT's output waveforms, as well as those coming directly from measurements, can be further processed within the graphic environment, SIGHTS, that also support a set of signal quality evaluation tools including:

- Eye-diagram display of signals for the evaluation of quality parameters including eye opening, noise margin, and time jitter.
- Calculation of worst case eye opening starting from a simulated or measured single transition. The digital bandwidth of components and systems can be quickly evaluated in this way.
- Worst case binary pattern generation. This pattern, causing the maximum eye closure at a given bit-rate, is very useful for generating test vectors to stimulate both simulated systems and actual prototypes. In this last case, the vectors can be loaded into a high-speed pulse pattern generators, such as the HP 80000A data generator system, for real-time verification.
- The signal template specification is a very effective technique used to carry out compliance testing at system's standard interfaces or to check signal quality within the system itself. Moreover, it is very useful to evaluate a system's global quality when a graphic display of a waveform is impractical, due to huge number of test points as usually happens in exhaustive post-layout simulation of PCBs or MCMs.

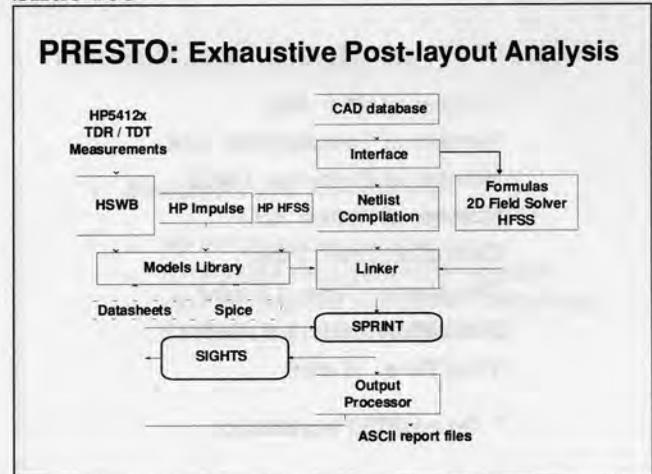
A High-Performance Environment for Modelling and Simulation of Digital Systems

Slide #54



The previously mentioned modelling, simulation, and quality checking tools are fully integrated in the HSWB (High Speed WorkBench) environment, that links wideband time-domain instruments to SPRINT & SIGHTS, as shown in the slide. The instruments are connected to the workstation via an HP-IB or RS232 serial interface and the acquired waveforms can be stored in SIGHTS format for further processing. Using HSWB, the user can quickly extract models from wideband TDR/TDT measurements, carried out on hardware components, including ASIC, packages, PCB traces, connectors, backplanes, cables, and even optical components. TDR/TDT component characterisations can also be performed in a simulative way using tools like SPICE or HP IMPULSE, if working at transistor level, or for ultimate accuracy from 3D field solvers like HP's HFSS, when technological and geometrical parameters are available.

Slide #55



PRESTO is the environment linking the hardware models obtained from TDR/TDT measurements to CAD databases to perform automatic pre- and post-layout analysis of entire PCB, MCM or Hybrid designs. Topological and geometrical information extracted from a CAD database is automatically converted into a SPRINT netlist, including electrical/behavioural models of interconnects and active devices. The experimental way is the most accurate method to extract models, but they can also be obtained from simulated TDR/TDT tests using other electrical simulators, such as the HP Impulse product, or 3-D field solvers, like HP's HFSS.

Exploiting SPRINT's features, it is possible to analyse the whole system within the same simulation run including transmission, crosstalk, and switching noise effects, simultaneously. This is the only way to get reliable results, because all drivers are loaded with their actual load. Input to output timing delays and even logic behaviour can be included in the models, so that simulation can also treat these issues with signal integrity in mind. The user can perform a complete compliance analysis of all signals with respect to user defined signal templates getting a global report that contains the list of nets with their violation errors. The effects of non-ideal power and ground planes, as well as EMI/EMC evaluation, will be available in the near future.

Slide #56

PREST0: Demo Board Test Results

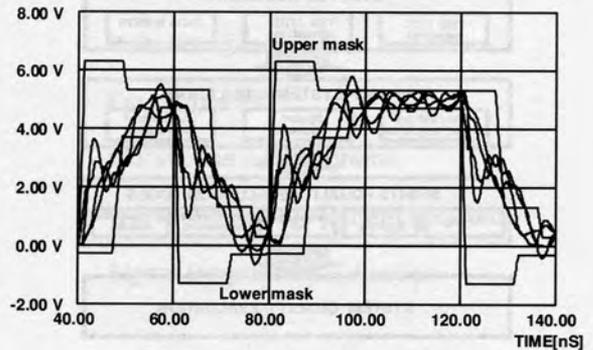
Number of Nets: 488
Number of Components: 228
Number of Elements: 13492
Number of Nodes: 6255
Extraction Time: 2 min*
Compilation Time: 1.5 min*
Simulation Time: 1.5 min*
Total Time: 6 min*

* On a HP750 workstation

Shown are the main results related to multilayer PCB post-layout analysis. This board contains about 200 components and 500 nets. The total analysis time including extraction, netlist compilation, and SPRINT simulation is about 6 minutes on an HP 750 workstation. Thanks to this speed, it is possible to dramatically reduce design and redesign cycle times and enhance product quality.

Slide #57

Exhaustive Compliance Analysis



Shown is an example of automatic compliance analysis of signals performed on the demo board previously described. The waveforms of a critical net and the user-defined signal templates are displayed to point out violation errors. An automatic violation error evaluation is carried out for each net and then stored in a report file. A global violation error is also detected to get a figure that represents the quality level of the whole board.

Slide #58

Performance & Quality: the Key of Success

Pre-layout analysis

Design rule setting
Performance optimization

Post-layout exhaustive checks

Reliability enhancement
EMI/EMC evaluation
Faster prototyping and validation

Focusing on performance and quality issues is the way to be competitive in the Telecom market. System manufacturers and service providers will be more and more deeply involved in these two major issues so that new design and test tools are strongly required to face them. This paper has shown some applications of an integrated measurement and simulation environment that offers unique help in design and validation of telecom apparatus. Tight integration among wideband time domain instruments, behavioural modelling procedures and a powerful simulation engine are the key features of this environment. A great number of actual applications, including design and validation of broad-band telecom switches, have demonstrated the effectiveness of this approach and its benefits.

Experience suggests the systematic use of this kind of tool to perform pre-layout analysis to set design rules to follow during the physical implementation. In this phase, it is also possible to optimise system performance to solve physical layer bottlenecks. Exhaustive post-layout checks, including signal compliance analysis and EMC/EMI evaluation, performed on routed boards and possibly extended to the entire apparatus, is the best way to verify and enhance the quality of hardware before its implementation and can greatly help both prototype debug and system validation.

Slide #59

Recommended Resources

- Equipment and accessories
 - HP 54120-series oscilloscope
 - HP 80000A data generator system
 - HDT high-speed workbench including: SPRINT, SIGHTS, and PRESTO

References:

- [1] "TDR Techniques for Differential Systems", HP AN 62-2
- [2] SPRINT&SIGHTS APPLICATION HANDBOOK, HDT
- [3] "Passive Components Modelling based on Reflectometer Measurements", HDT AN-01
- [4] "Modelling of Active Components", HDT AN-02
- [5] "Design and Validation of High Speed Optical Data-Links", HDT AN-13