

HP Archive

This vintage Hewlett Packard document was preserved
and distributed by

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

Please visit us on the web !

Thanks to on-line curator: Istvan Novak

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**

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.

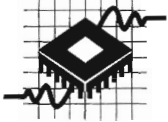
Author Background:

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

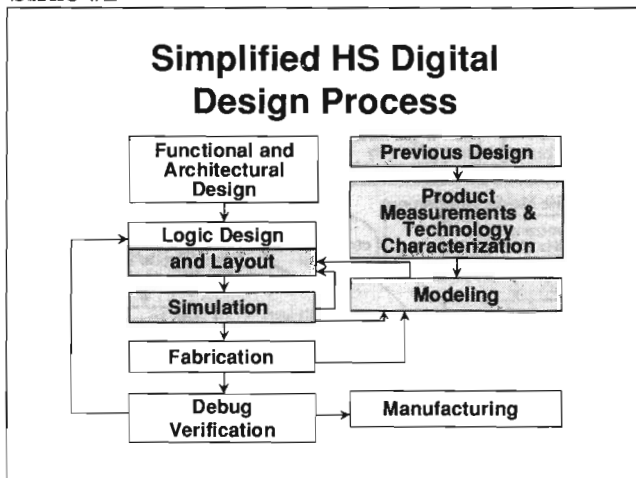
Doubling the Clock Speed
Feasibility Study on a 486-50 MHz
Cache Module at 100 MHz



TECHKNOWLEDGE **NCR** CASCADE MICROTECH[®] **hp** HEWLETT
Humanizing Technical Information PACKARD

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.

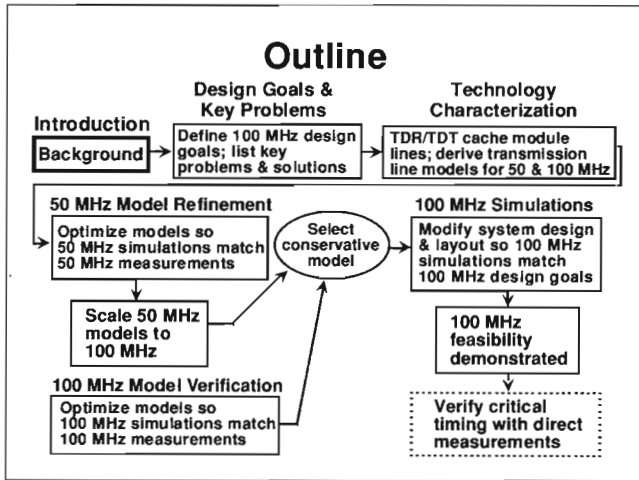
Slide #2



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



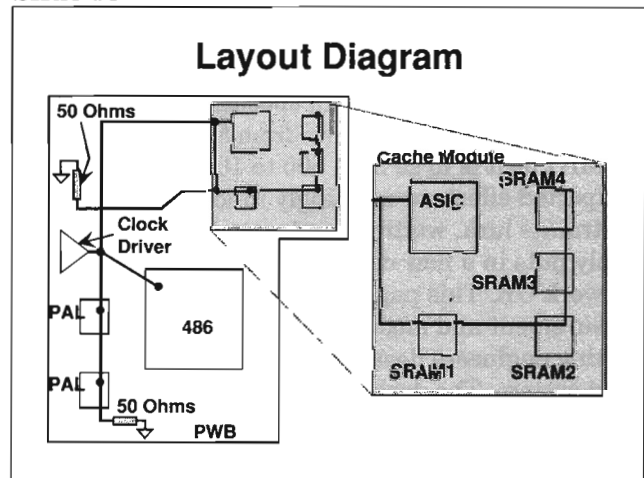
The product was first fabricated entirely on FR4 circuit board and electrically debugged. It measured 7 by 8 inches and was of course much too large to meet the final product mechanical requirements. In the current product, all the SSI/MSI were incorporated into PALs; the cache controller ASIC and the 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.

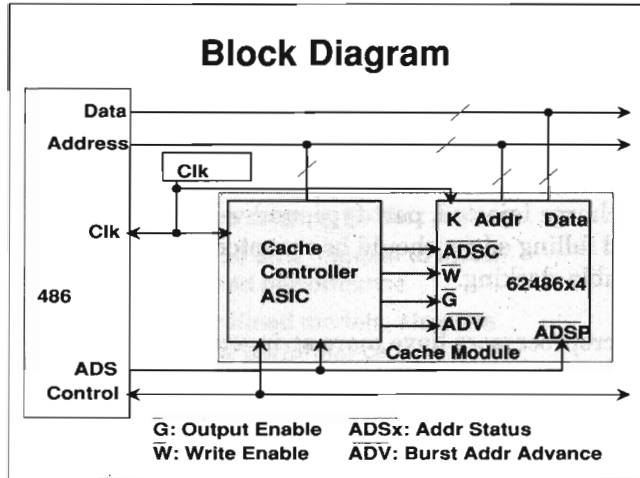
Slide #5



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

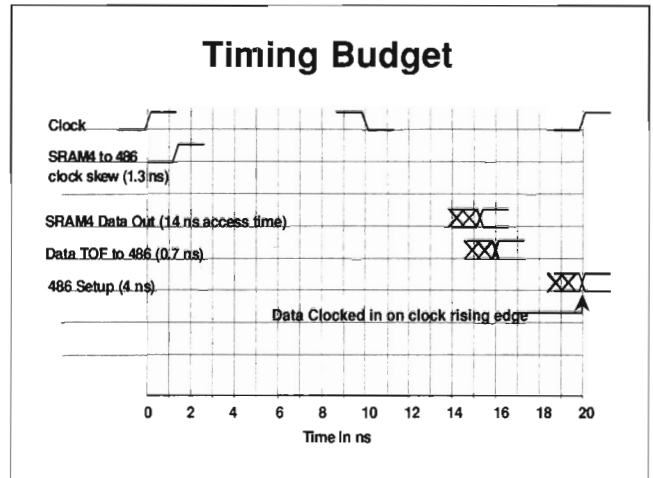
Slide #6



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

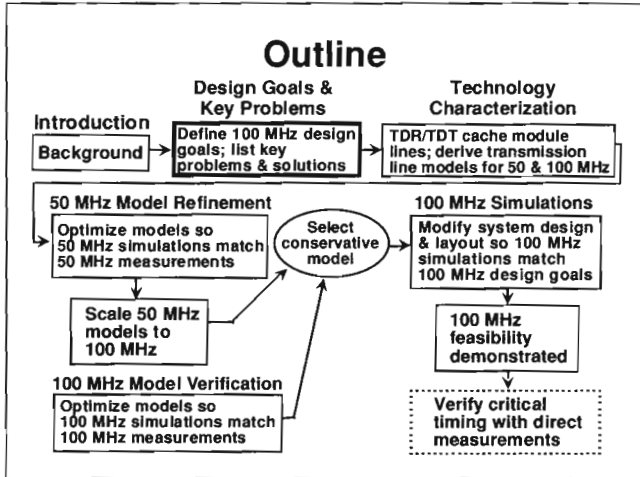


With a 50 MHz 486, and 14 ns SRAMs, the ASIC needs 1 clock cycle to decide if the data required is in the second level cache. If the data is cached, then the succeeding 4 reads or writes can occur in burst mode, at one operation per clock cycle. The timing diagram shown here is for a burst mode read 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



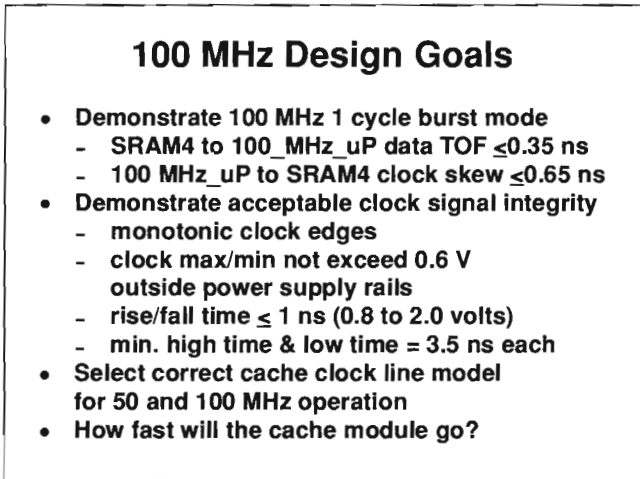
Implicit in these design goals was the requirement that the clock signal integrity meet company standards. Overshoot and undershoot greater than 0.6 volts above and below the power supply rails was not allowed, minimizing charge injection into the substrate. Charge injection has been associated with 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



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

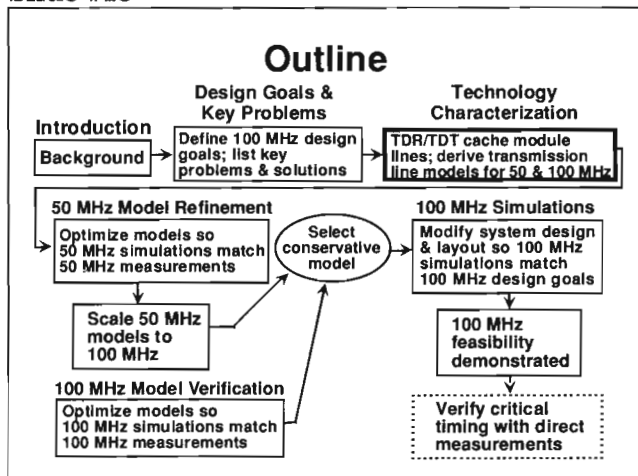
Which Transmission Line Model?

- 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.

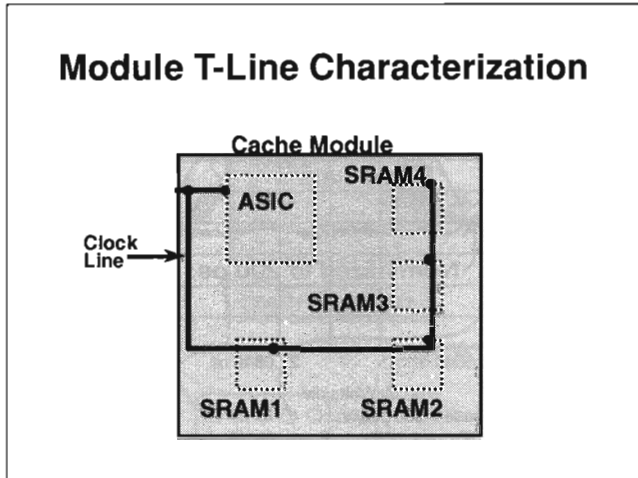
Slide #13



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).

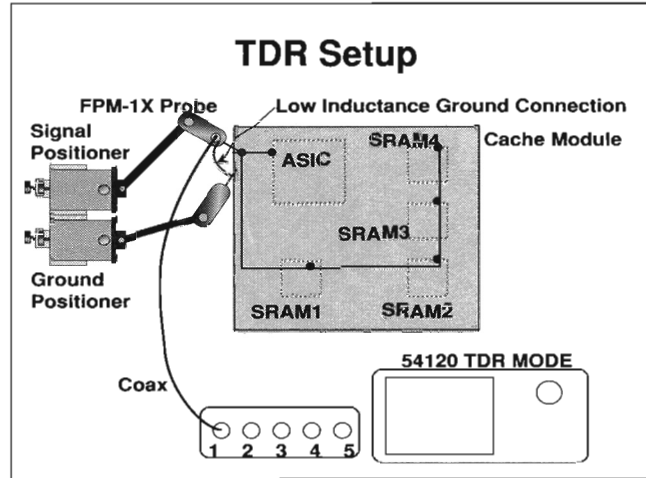
Doubling the Clock Speed Feasibility Study on a 100 MHz Cache Module

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

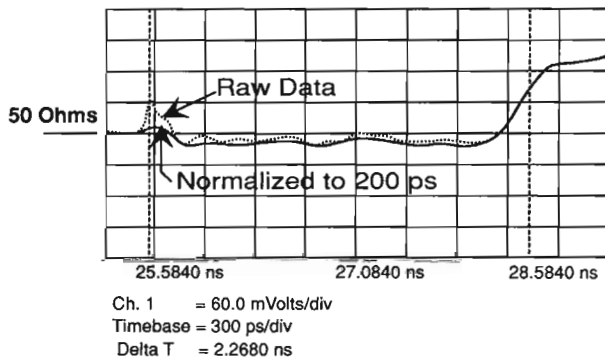
Fine Pitch Probes, Passive



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

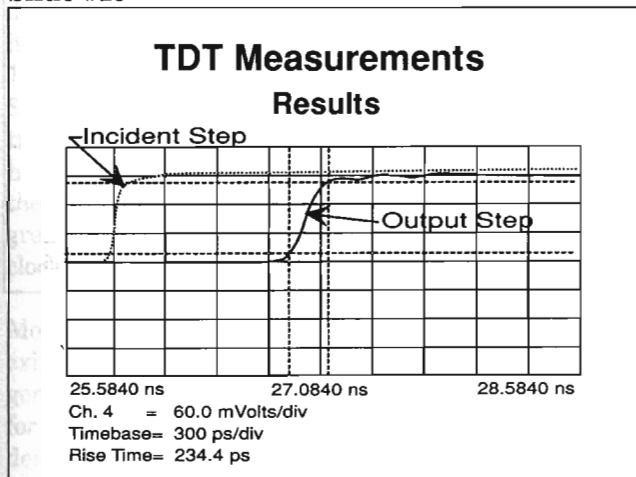
through connection on the Surrogate Chip*. This demonstrates to what extent the cables and probes slow the step, and what distortions are added. The rise time is approximately 50 ps, with some rounding of the corners due to probe ground inductance, and cable skin effect. The extent that the output signal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

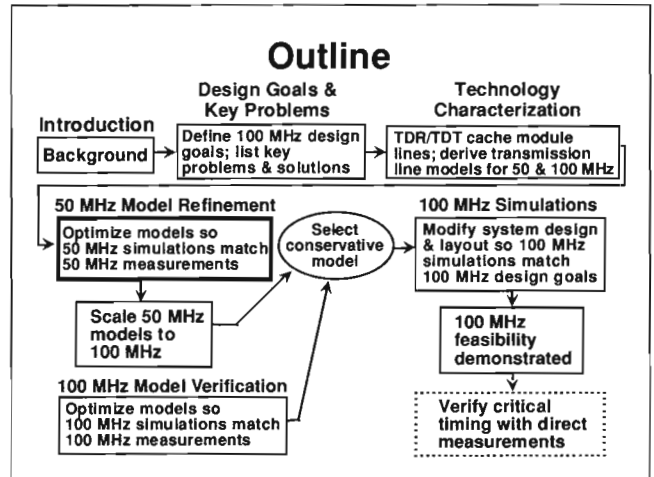
When simulating waveforms on this cache module with <300 ps transitions, a lossy, skin effect model is required. Otherwise, the simulator will predict faster performance than the transmission line will provide. Because time domain skin effect models are relatively new, we recommend that several line 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:

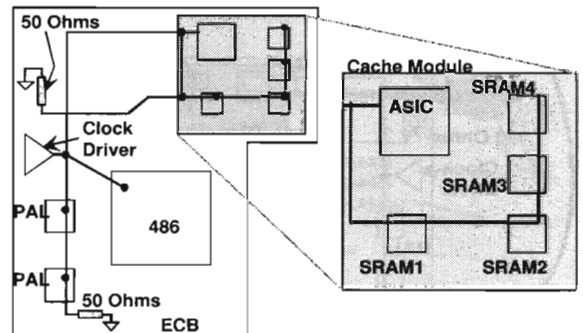
1. Measure clock performance of the current 50 MHz system
2. Compare with the 50 MHz simulations
3. Optimize the clock driver models obtaining a good match between measurements and simulations.

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



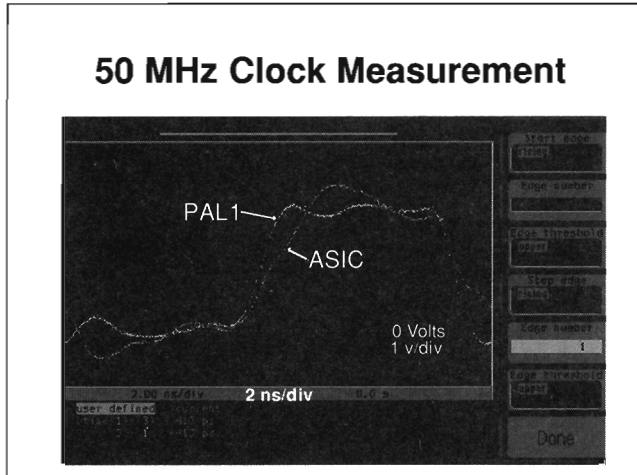
The left photo shows the measurement setup used including the active probes.

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 Vdd and Vss. 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

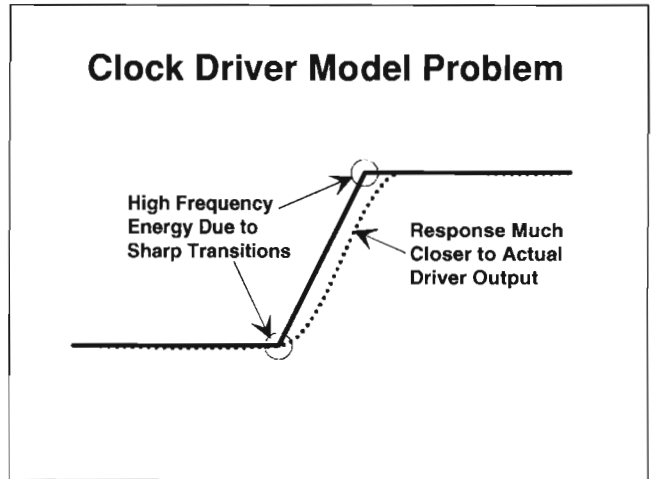
Feasibility Study on a 100 MHz Cache Module

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

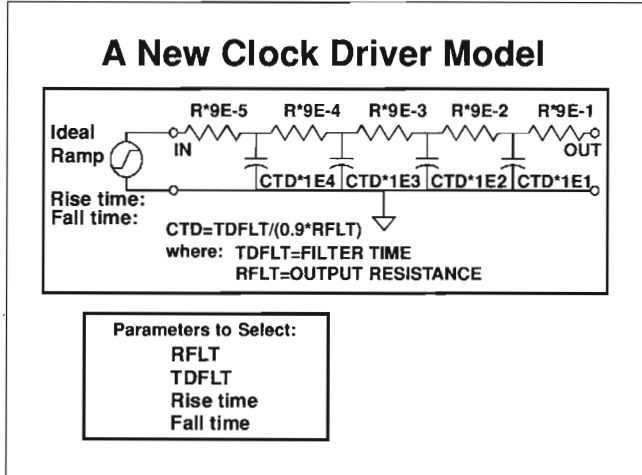


Using an idealized ramp signal to model a clock driver causes two errors: 1), is that an ideal ramp waveform contains high frequency energy due to the sharp transitions which is not present in the driver being modeled; and 2), when using ideal transmission line models, the additional high-frequency energy 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

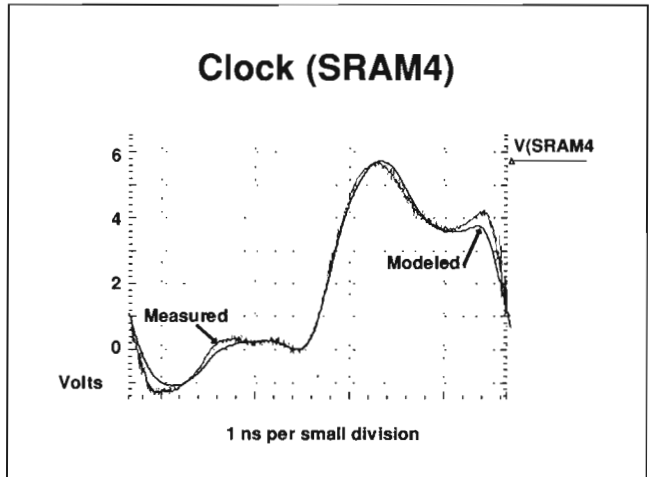
50 MHz Driver Model:
RFLT = 5 Ohms
TDFLT = 250 ps
Rise time = 1.9 ns
Fall time = 1.3 ns

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.

cr
st
ar
w
af
N
48
0.6

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

Slide #37

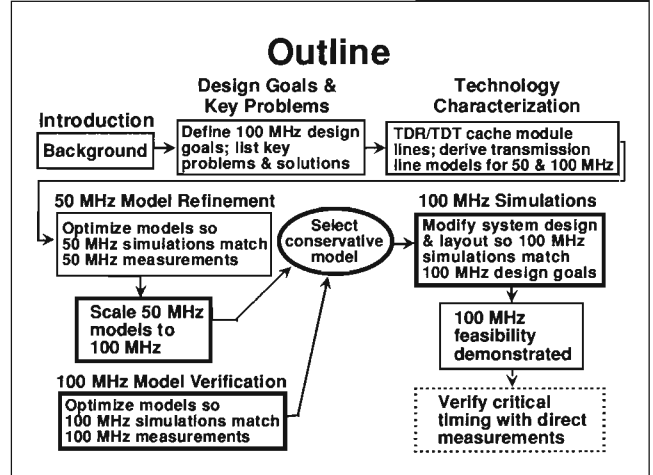
Timing Delays

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

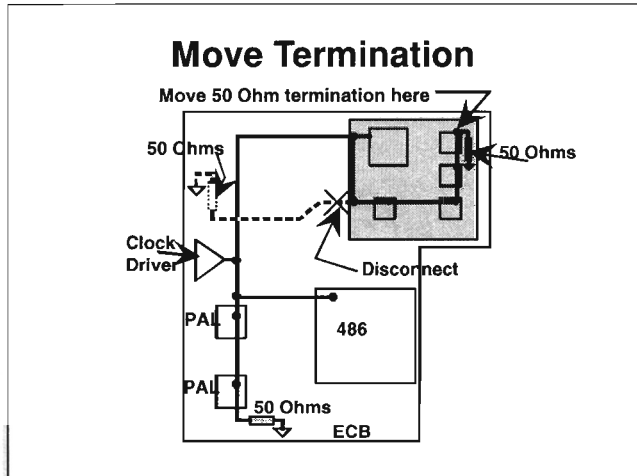
Shown above is simulated V(ASIC) overlaid by the measured V(ASIC). The question now is, for the 100 MHz simulations, do you use this driver model or the scaled 50 MHz driver model? Because the clock driver has not yet been chosen for 100 MHz design, you should pick the most conservative driver model, 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.

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.

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

- 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.

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 #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

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

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.

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 μ m), 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

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.

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

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 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.