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

Pat Byrne

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

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

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."

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.

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:

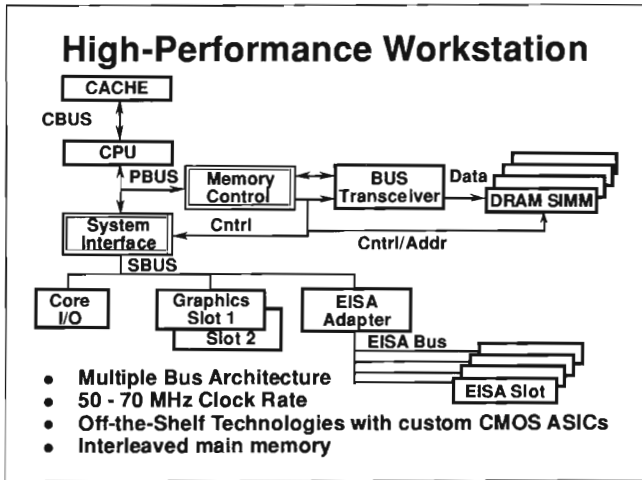
1. It is a Multiple Bus Architecture, where high-speed transfers are kept close to the CPU while low-speed functions are isolated by "smart" ASICs like those found in the bold boxes. These ASICs control the bus protocol, the interrupt sequence, arbitration, act as temporary storage through on-chip FIFOs, 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.


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.

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 repeatably 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

Bus Isolation

- Probe all busses along chain
- Isolate failure to one location and one transaction

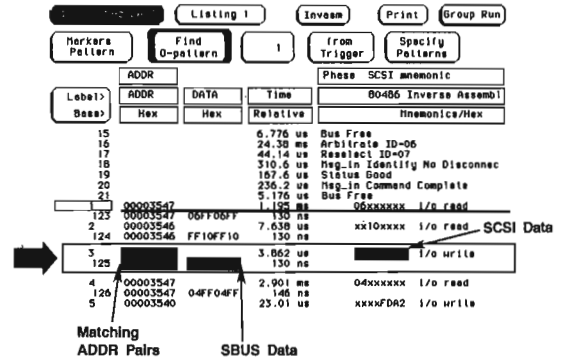


- **Result:** Failure isolated to PBUS read. Data good on PBUS write.

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

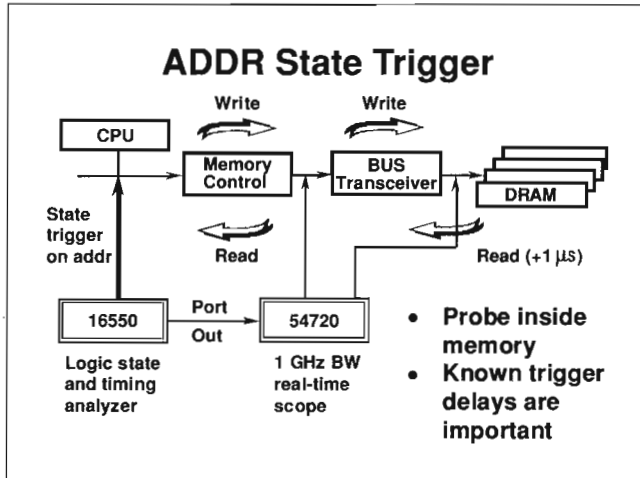
ADDR State Trigger



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 $1 \mu\text{s}$ 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.

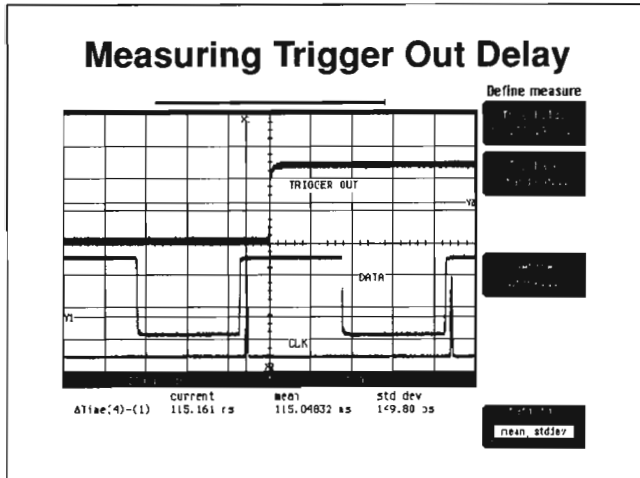
Slide #14



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.

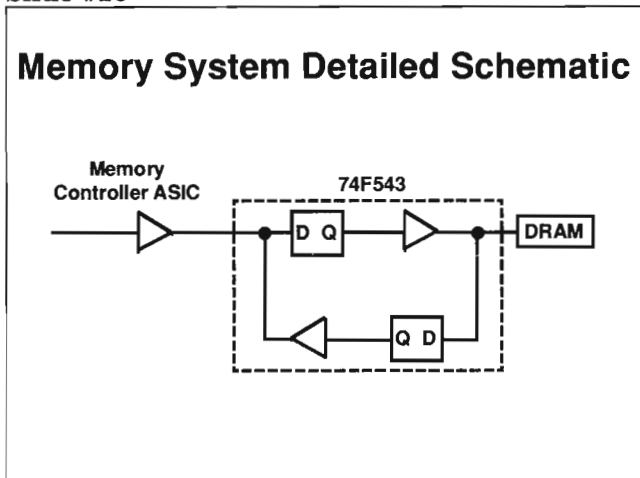
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

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

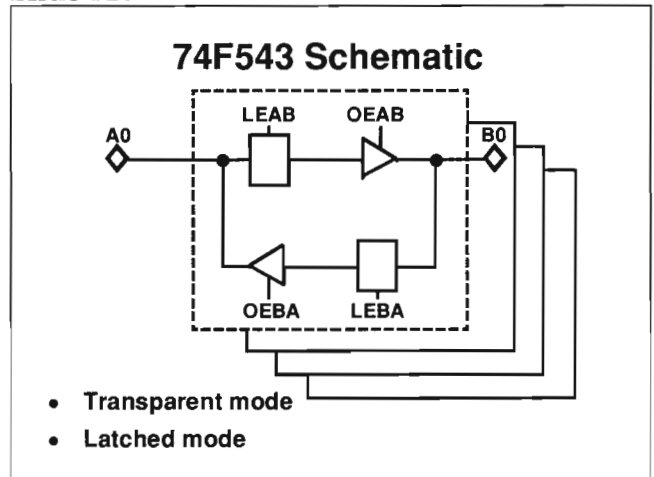


The root cause of the problem is found within the main memory system during the WRITE operation. The DATA passes through the Memory Controller ASIC without corruption but is found to be bad on

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

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.

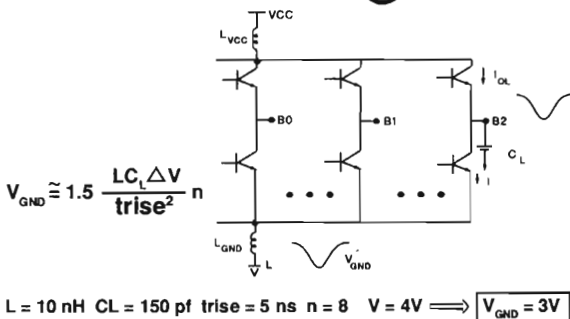
Note that the conditions to cause this failure are dependent on the state, timing, and loading conditions. All these combine to create the B output glitch conditions which make the tristate control block malfunction and cause bus contention on the A input. These are the most subtle and infrequent problems 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 (ΔV), and the transition time (T_{rise}). 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 (T_{rise}). 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.

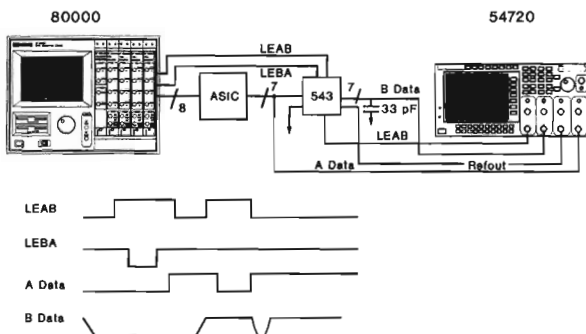
Generation System is a modular, high-performance Data Word Generator. It is capable of less than 200 ps edge speeds and 2 ps edge placement resolution. In the characterization process we are performing, we need to have better than 100 ps edge placement resolution so that we can characterize the ground 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

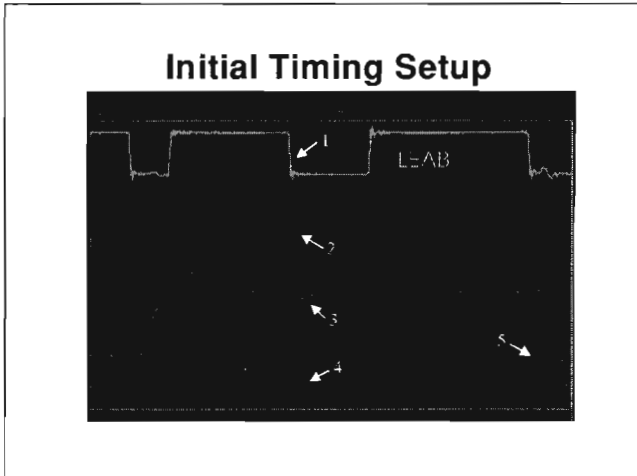
Debugging and Characterizing Ground Bounce Problems in High-Speed Memory System Hardware

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

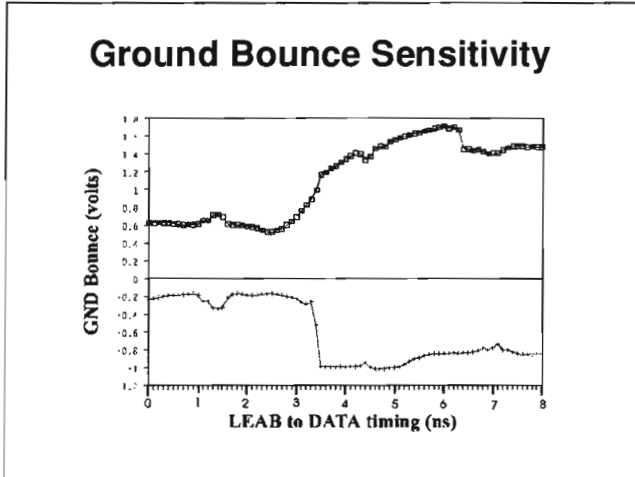
Ground Bounce Caused Instability



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

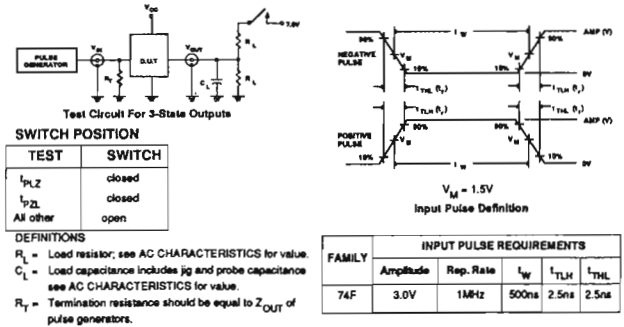
Principles of Chip Characterization for Ground Bounce Effects

- 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 Busses \Rightarrow more I overlap
Small Board Space \Rightarrow Higher L
Higher Integration \Rightarrow Fewer Ground Pins

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

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

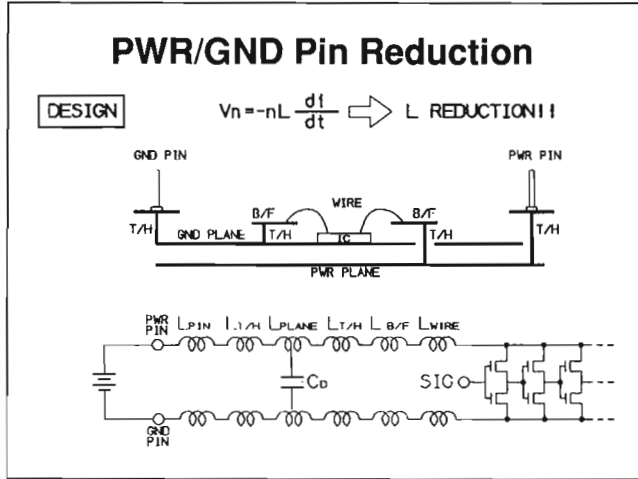
- Programmable output resistance for series termination
- Digital feedback to set drive based on turn-on conditions



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.