

Power Delivery Validation of Processor Front Side Bus

Mahadevan Suryakumar, Jiangqi He
Intel Corp
5000, W Chandler Blvd, Chandler AZ 85226
mahadevan.suryakumar@intel.com, 480-554-5193

Abstract

With lower I/O operating voltage and decreasing noise margin targets, it is imperative that the amount of noise generated on the voltage rails is kept within tight tolerances for high speed signaling. For this a good understanding of the behavioral model of the package and on-die passives is required in addition to the on-die transients generated on the power supply network when I/O's toggle. One of the major challenges is to accurately measure this dynamic power fluctuation (di/dt) for various I/O bit pattern excitation to understand the effectiveness of decoupling capacitors in the power distribution network. This paper discusses the power delivery validation methodology of the processor Front Side Bus (FSB) that allows the user to input a stream of data into the I/O's to measure the power supply noise and the resulting dynamic power fluctuation. A detailed three dimensional model of the package, motherboard was then created and simulated using Speed2K/PowerSI[1] to correlate with measured results.

Introduction

There are two methods currently used to validate the Front Side Bus (FSB) power delivery network. The first method involves using a network analyzer across the desired frequency range and map the AC impedance of the power delivery network. Although this method validates the power delivery network impedance, it doesn't give any insight into the dynamic power fluctuation or transients generated in the power delivery network when signals toggle. The second method involves using the MARS tool [2] to schmoos frequency domain victim-aggressor Simultaneous Switching Output (SSO) patterns that provide the worst case I/O margin. This method even though very powerful requires patterns to be tuned for optimal burst for a given processor and chipset configuration requiring the use of logic analyzer.

Proposed method: In this method a versatile DOS based tool was developed that can be tuned to output an optimal burst of data on the FSB data lines to measure the power supply noise through the test points on the package close to the processor die bumps. The tool was tested on prototype processor on an FCPGA (Flip Chip Pin Grid Array) package mounted on the validation board with Memory Controller Hub (MCH). First, the board and the processor were booted to DOS protected mode. The requested pattern is placed in a text file and the active I/O software was invoked through a series of commands. When the requested pattern is excited on the bus, an I/O switching from high to low or low to high is actually discharging or charging the capacitor that loads the I/O. The resulting di/dt is cumulative and gets higher if more number of signals toggle at the same instant. The processor drew a significant current (di/dt) through the power and ground

bumps when the signals toggle and this resulted in voltage droops. The scope was set to trigger when the pulse width of the signal match the requested pattern, and the voltage droops and recovery time were measured through the test points on the package. Several voltage droops were observed because local and remote power distribution responded differently in time. The measured voltage droops were then compared with the processor specification. A detailed electrical model of the processor FSB power delivery network was then created including the voltage regulator, motherboard, CPU package, MCH package, CPU and MCH die to correlate the simulation results with measured power supply noise.

Validation setup

The validation system consists of a Customer Reference Board (CRB) with a controlling PC. In the CRB, the prototype processor (CPU) and MCH communicate via the front side bus with 64bit data bus running at 400MT/s and 32 bit address bus running at 200MT/s. The processor was housed in a FCPGA package and was attached to the

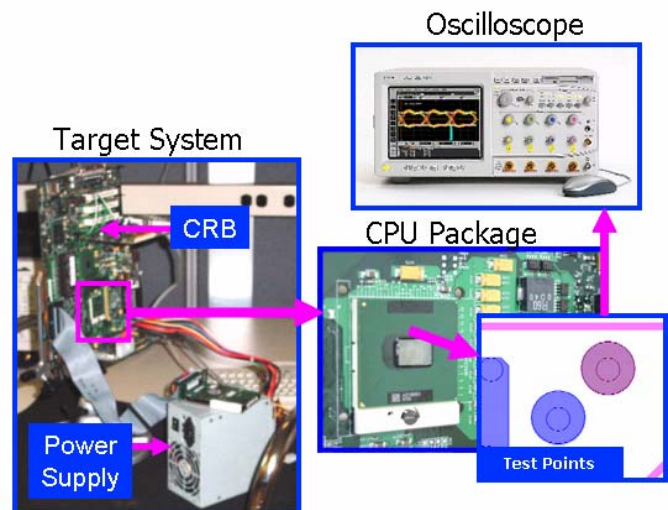


Figure 1: Validation setup

motherboard through an uPGA479 socket. The DOS based tool was invoked through a series of commands after booting the system in DOS protected mode. The program generates the same pattern endlessly and allows easy scope triggering. It uses an auxiliary utility to disable prefetching, sets cache mode and uses a pattern file to generate the requested data on the FSB endlessly by looping. The requested data is placed in a text file with each line containing a double-word and a pair of lines making a single 64bit pass on the bus. The corresponding bits of the data bus (D[0:63]) can be turned on or off and repeated many times over and over to stress the bus. The pattern file can be tuned to output any set pattern on

the data bus. The format of the pattern file is shown in Figure 2. Each hexadecimal digit is a nibble (lower 4 bits of a byte) and the number of lines in the pattern file corresponds to the length of data pattern. A batch file was set for read or write cache option for the appropriate memory destination by changing the cache utility parameters. In the read mode, few lines of data are read from main memory in write-back (WB) cacheable memory and using L2 cache structures (ways/sets), the data is repeated by reading the same lines in a loop. The following set of lines in the batch file describe the read operation.

```
mc21 -change 0x100000 0xffffffff UC
mc21 -change 0x100000 0xffffffff WB
```

where the second number is TOM -1 (TOM is Top of memory), in the above example this is 128MB and this can be checked by running the tom.exe utility. For write mode, write-back (WB) cacheable memory was used as the source and using WC as the destination, full lines of data are written to memory.

1	0	x	0	0	F	F	0	0	F	F
2	0	x	0	0	F	F	0	0	F	F
3	0	x	0	0	0	0	0	0	0	0
4	0	x	0	0	0	0	0	0	0	0
5	0	x	0	0	F	F	0	0	F	F
6	0	x	0	0	F	F	0	0	F	F
7	0	x	0	0	0	0	0	0	0	0
8	0	x	0	0	0	0	0	0	0	0
9	0	x	0	0	F	F	0	0	F	F
10	0	x	0	0	F	F	0	0	F	F
11	0	x	0	0	0	0	0	0	0	0
12	0	x	0	0	0	0	0	0	0	0
13	0	x	0	0	F	F	0	0	F	F
14	0	x	0	0	F	F	0	0	F	F
15	0	x	0	0	0	0	0	0	0	0
16	0	x	0	0	0	0	0	0	0	0

Figure 2: Pattern file example with four 1's and four 0's – an even mode pattern is excited on the lower order data bits of each data group.

In this way, various frequency domain patterns can be stimulated to study the effect of power delivery resonance as well as Inner Symbol Interference (ISI). For example, switching all FSB data lines except for one (aggressor line) creates the frequency domain stimulus. The left over line is considered as the victim line and can be switched in various relationships to the aggressor lines to look for worst-case I/O events. Adjusting the number of low and high cycles in a data pattern stimulates various FSB frequencies. Keeping close to a 50% duty cycle creates the largest amount of energy at one discrete frequency. Duty cycles varying from 50% will stimulate multiple frequency points with far less energy, but these could also be useful in platforms where multiple resonance are occurring instead of one dominant pole or zero. For this study, the bus inversion was enabled and “10101010” pattern was fired on 32 data lines at 200Mhz (400MT/s) with 8 signals switching in each of the four data groups. A high-speed oscilloscope (Tektronix TDS6604) was used to monitor the IO power supply voltage (VTT~ 1.05V) and the data

signal. The scope probes were attached to the power and ground locations closest to the processor which were the test points on the package while the probes for the data signal were attached between the signal and ground motherboard via on the secondary side of the board. The scope was then set to trigger when the pulse width of the monitored data on the signal line matches with the requested pattern.

Modeling and Simulation

Capacitor Characterization: To correlate the measured power supply noise would require accurate modeling of the power distribution network along with a correct representation of capacitor behavioral model under platform use conditions. As shown in Figure 3, the CPU package has a total of five 0805IDC land side capacitors with test points on the top layer of the package.

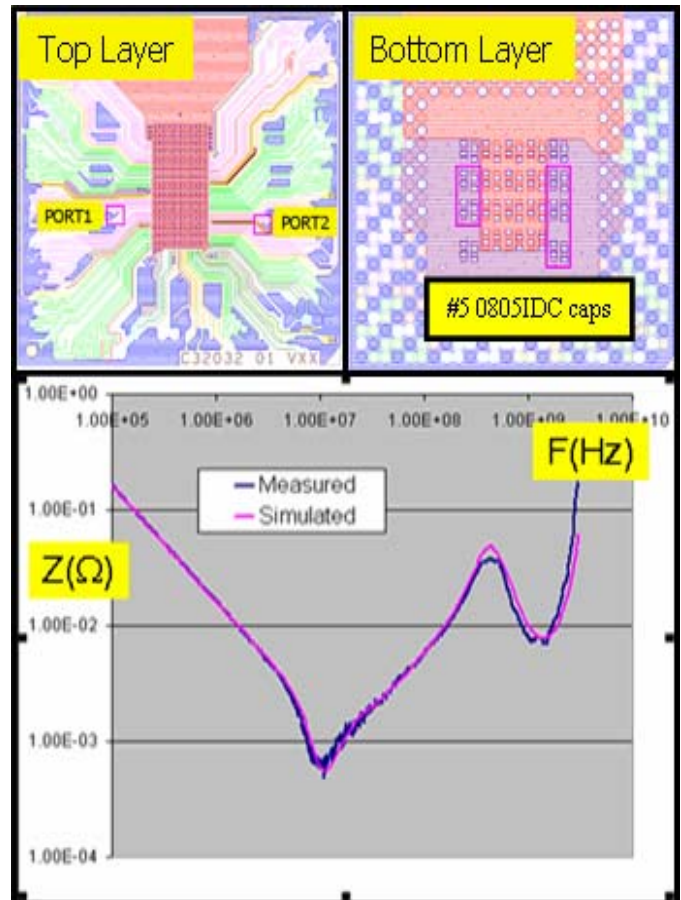


Figure 3: VNA measurements of package and die

Measurements were made using a VNA (Vector Network Analyzer) on the assembled CPU PGA package with package caps and silicon to map the AC impedance of the power delivery network. The package was then modeled using Sigrity-PowerSI tool with the ondie capacitance distributed on the power and ground bumps of the die matching the distribution in the silicon as closely as possible and by using higher order package capacitor models [3]. Since the IO Signal:Power:Ground (S:P:G) ratio on the CPU die is n:1:1, the ondie capacitance (Cdie) and effective series resistance (Rdie) of the on-die cap cell per bump pair were sized to match the Signal:Power:Ground ratio. [Example: Cdie per

bump pair = $n \cdot C_{die}$ per buffer and R_{die} per bump pair = $(R_{die}$ per buffer)/ n]. As seen from Figure3, a reasonable good match exists between measurement and simulation for C_{die} - R_{die} estimates. A resonance between the package inductance and ondie capacitance occurs around $\sim 450\text{Mhz}$ with an impedance peak of $\sim 48\text{m}\Omega$, while at 200Mhz an impedance of $14\text{m}\Omega$ is observed.

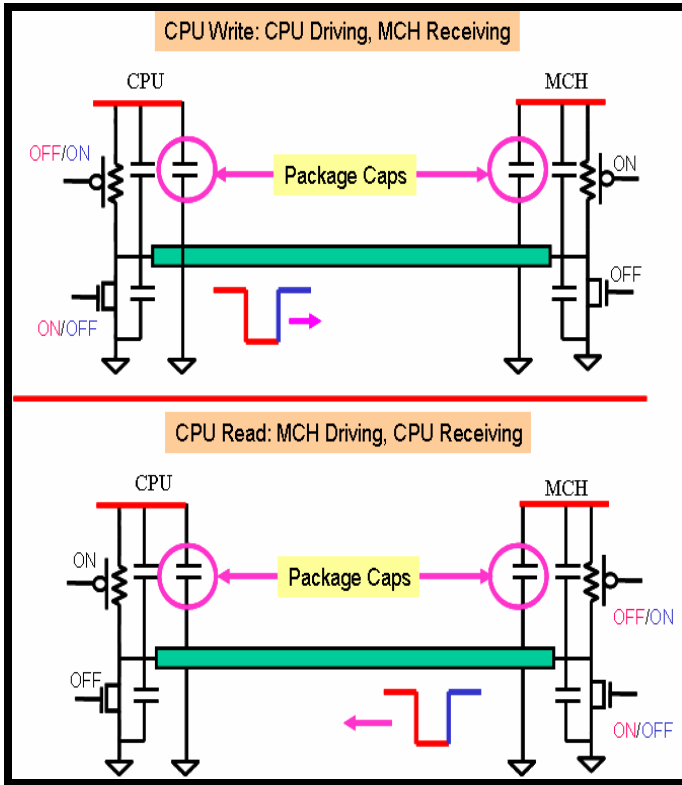


Figure 4: CPU Read and CPU Write topology – die, package parasitics omitted for readability.

The front side bus on the prototype processor uses CMOS (Complementary Metal Oxide Semiconductor) based AGTL+ (Assisted Gunning Transceiver Logic) signaling with dynamic termination on the driver [4]. This means the receiving agent will always have its pull-up device turned ON while the driver side, pull-down and pull-up devices will turn ON and OFF depending on whether high to low or low to high transitions are driven on the signal lines. In other words, to study the effectiveness of the CPU package IO power delivery connection, CPU read transaction must be used as the CPU package caps will be in the circuit only for the CPU read operation. See Figure4 for CPU read operation where MCH buffers toggle while the CPU is in receive mode. The package (R, L) parasitics are not shown in Figure4 for readability. In the CRB, the CPU silicon is connected to the MCH through the 6layer CPU PGA package, socket, and 8layer motherboard and 6layer MCH package. In addition to on-die capacitors both in the MCH and CPU silicon, capacitors are placed on the CPU, MCH packages and on the secondary side of the motherboard as shown in Figure5..

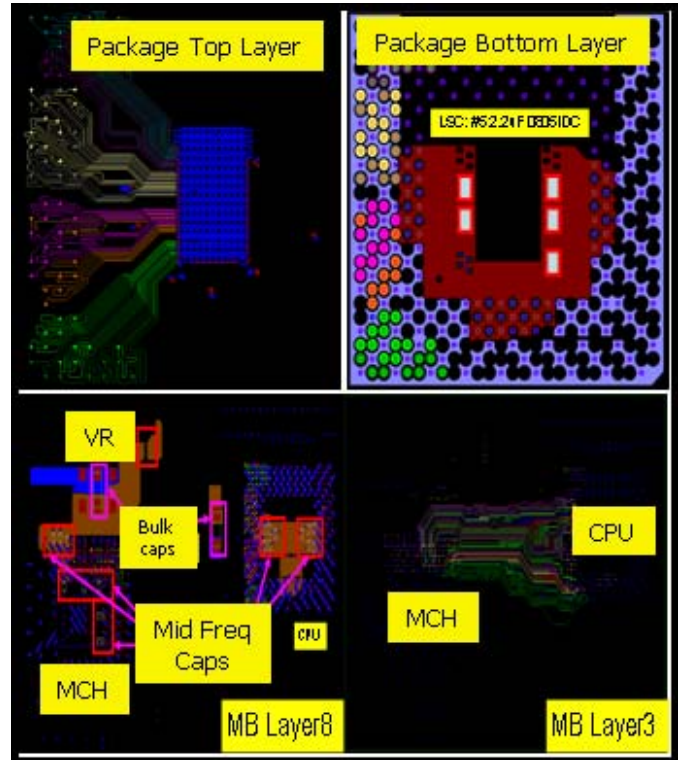
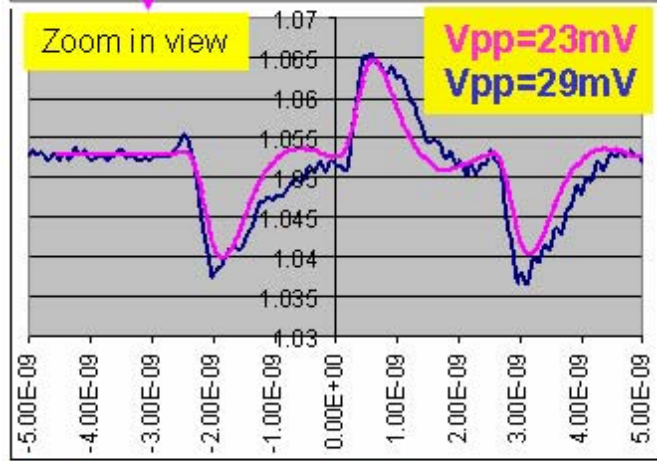
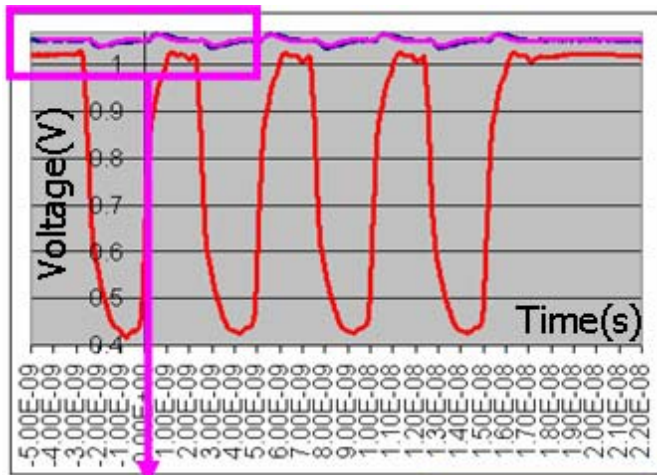


Figure 5: Speed2K Model of Package and Motherboard

Simulated vs. Measured Results

The package and board layout files are converted from cadence format (.brd, .mcm) to one that is recognized by Sigrity-Speed2K (.spd). The data signals along with the power and ground planes on the package and motherboard were extracted. The package and motherboard are then connected through vias of appropriate length and size to mimic the BOL (Beginning Of Life) socket resistance and inductance. The on-die capacitors and driver/receiver IBIS buffer models were included as sub circuits in Speed2000. Since the interest is only on the power delivery the crosstalk noise from trace to trace coupling was not included in the model. A data pattern of “10101010” was then excited on 32 data signals at 200Mhz (400MT/s) mimicking the validation setup and the voltage at the test point on the package was monitored. The VTT noise measured on the test points is averaged to reduce the noise from the measurement floor. Figure6 shows the measurement vs. simulation correlation. From Figure 6 one can observe that there is good correlation



- D1-measured
- VTT-Simulated
- VTT-Measured

Figure 6: Simulated vs. Measured VTT noise correlation at package test point

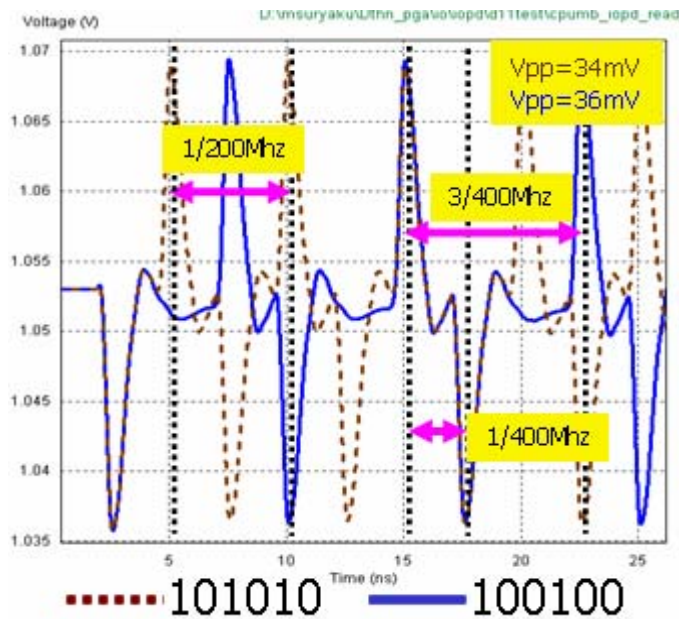
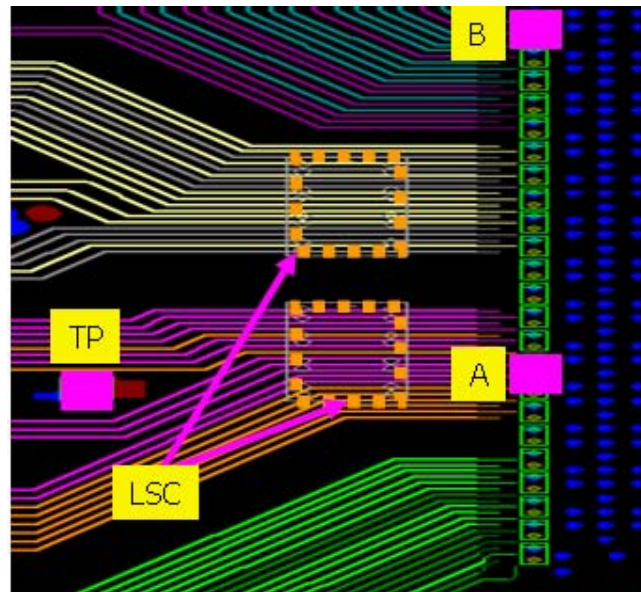
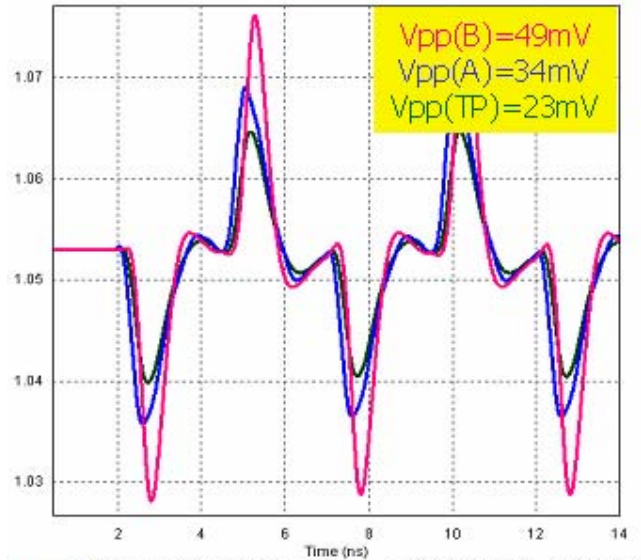


Figure 7: VTT noise at die bump for 101010 (200MHz) and 100100(400MHz) data patterns

on the shape and amplitude of the power supply noise measured on the package test point. As can be seen the system is critically damped when excited at 200Mhz. The high frequency voltage ripples on the measured waveform are resonances that exist between the ondie-on die capacitors at different locations on the die and also between package plane capacitance to on-die capacitance. A more detailed on-die capacitor model needs to be used to correlate the high freq resonances in the measured noise waveforms.



U:\msuryaku\lthn_pgav\topdtd1testcpump_ropd_read



- VTT noise (PointB)
- VTT noise (PointA)
- VTT noise (Test Point)

Figure 8: Spatial variation of VTT noise at different locations on the die for 101010 pattern

Now that we have the I/O power delivery model correlated with measurements, Sigrity-Speed2K simulations were done for different FSB data patterns. From figure7 one can observe that the system is critically damped and the peak to peak noise is comparable for both 101010 data pattern with peak energy

at 200Mhz and 100100 data pattern with peak energy at 400Mhz noise. In addition the noise at the bumps further away from the package capacitor is 2X higher than noise seen at the package test point. This is because the test point was close to the middle of the die and to the shadow of the landside package capacitor. Figure8 shows the spatial variation of VTT noise at different points on the die in relation to the package capacitors. One can observe that the noise on the test point is slightly higher than on the die bump (Point A) close to the package capacitor, however there is substantial difference in noise comparing to a location on the die (pointB) that is much further away from the package capacitor. This is due to lateral inductance contribution of the package plane layers connecting the land side package capacitor to the microvias that drop down from the die bumps at point B.

Conclusion

The active validation system design brings great benefit to the power delivery and signal integrity analysis as it is very easy to operate and is flexible in control of the board and the processor. These frequency domain and time domain measurements provide direct insight into the current profile of the processor as well as the performance of the designed power delivery network. This data together with the passive measurements on the package substrates provide a direct insight into manufacturing advancements for the upcoming technologies for high speed interconnects.

Acknowledgments

The authors would like to thank Lars Skoglund, Scott Niemeyer, for drilling small holes in the package for probing Ofer Rosner for active IO software, John Martinez and Michael Hill for helping with the measurement.

References

1. Speed2K and PowerSI user Guide --- Sigrity Corporation.
2. Jeff Fearing et.al "Mars methodology". Intel Internal Document.
3. J.H. Kim, Dan Jiao et .al "Characterization of discrete decoupling capacitors for high speed digital systems" ECTC 2003.
4. Michael Livshitz, et.al "Banias Processor System Bus power delivery validation". DTTC 2003 – Intel Internal Document.