



Impact of Intermittent Faults on Nanocomputing Devices

Cristian Constantinescu

June 28th, 2007

Dependable Systems and Networks

Outline

- Fault classes
 - Permanent faults
 - Transient faults
 - Intermittent faults
- Field fault/error data collection
- Intermittent faults
 - Impact of scaling
- Mitigation techniques
 - HW vs. SW solutions
- Summary
- Q&A

Fault Classes

- **Permanent faults**, e.g. stuck-at, bridges, opens
 - Reflect irreversible physical changes
 - Occur at the same location, are always active
- **Transient faults**, e.g. particle induced SEU, noise, ESD
 - Induced by temporary environmental conditions
 - Occur at different locations, at random time instances
- **Intermittent faults**, e.g. manufacturing residues, oxide breakdown
 - Occur due to unstable, marginal hardware
 - Occur at the same location
 - May be activated and deactivated
 - Induce bursts of errors

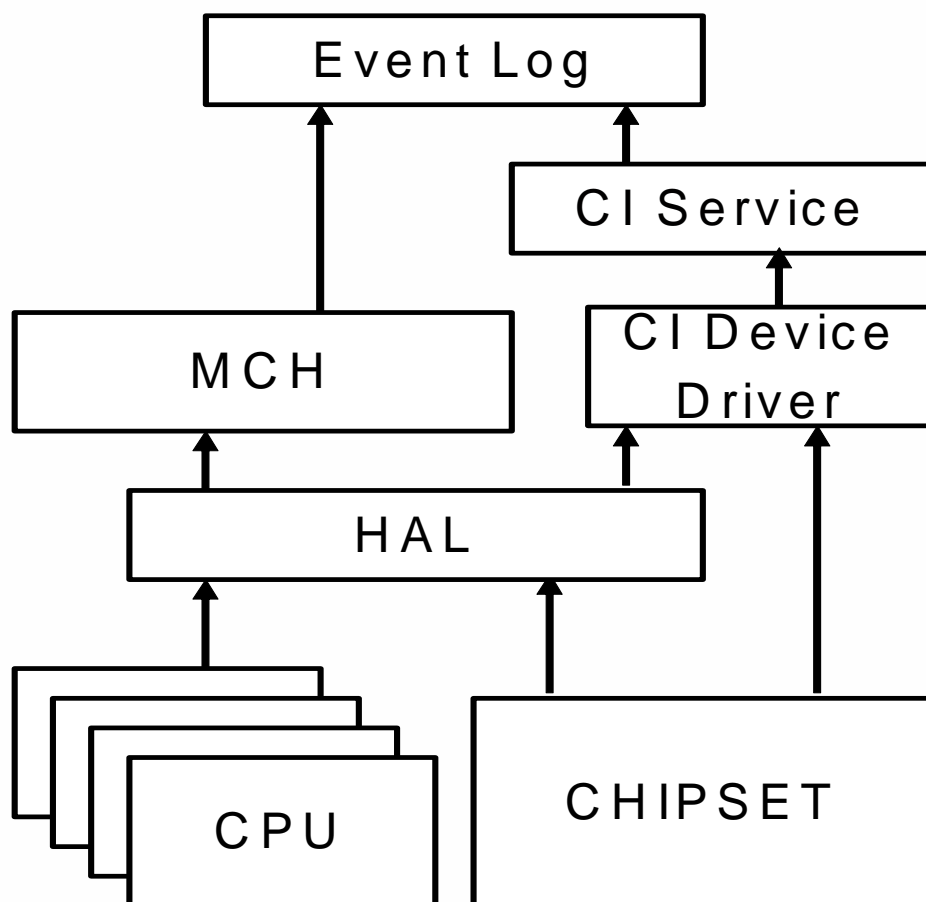
Fault/Error Data Collection

Fault/Error Data Collection Study

- Servers from two manufacturers were instrumented to collect errors
 - Manufacturer A: 193 servers, 16 months
 - Manufacturer B: 64 servers, 10 months
- Examples of reported errors
 - Memory
 - Front side bus
- Failure analysis performed when possible

Source: C. Constantinescu, SELSE 2006

Server Instrumentation



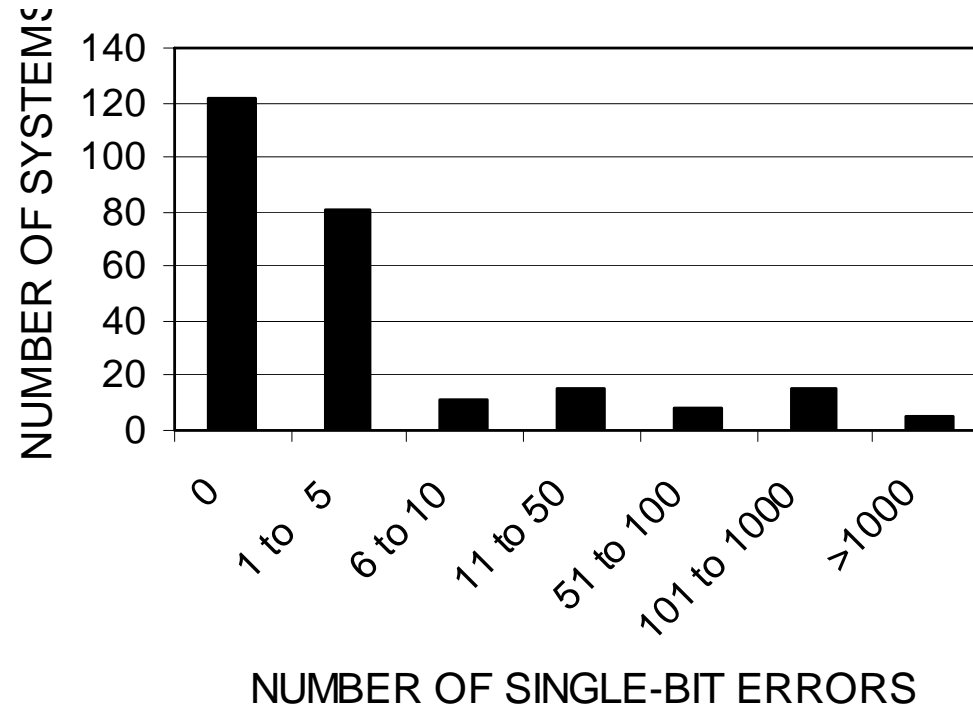
HAL – hardware abstraction layer

MCH – machine check handler

CI – component instrumentation

Instrumentation validated by fault injection


Corrected Memory Errors



- 310.7 server years
- Servers experiencing intermittent faults: 16 out of 257, i.e. **6.2 %**
- Corrected single-bit errors (SBE) induced by **intermittent faults**: 12990 out of 16069, i.e. **80.8 %**

Typical Signature of Memory Intermittent Faults

Daily number of corrected SBE



Failure analysis: SBE induced intermittently by poly residue, within memory chips



Source: Hynix Semiconductor

Processor Front Side Bus Errors

- Front side bus (FSB) errors
 - Bursts of single-bit errors (SBE) on data path
 - SBE detected and corrected (data path protected by ECC)


Server 1				Server 2			
P0	P1	P2	P3	P0	P1	P2	P3
3264	15	0	0	108	121	97	101
7104	20	0	0	-	-	-	-

- Servers experiencing FSB intermittent faults: 2 out of 64 (3%)
 - Burst duration examples: **7104 errors in 3 sec; 3264 errors in 18 sec**
- Failure analysis
 - **Intermittent contacts** at solder joints

More on Intermittent Faults

Timing Violations

BLM delamination



- Timing violations due to increased resistance; slow raise and fall times
 - **Intermittent** behavior occurs before the fault becomes permanent - specific for 90nm node and beyond
 - Permanent failures for previous technology nodes

Source: C. Constantinescu, SELSE 2006

Crosstalk Induced Errors

- Pulse induced by the affecting line into a victim line
- Timing violations due to crosstalk
 - Signal speedup or delay
 - Signal speedup – two adjacent lines switch in the same direction
 - Signal delay – two adjacent lines switch in opposite directions
- Process, voltage and temperature (PVT) variations amplify crosstalk induced skew
- Crosstalk increases with interconnect scaling and higher clock frequencies


Ultra-thin Oxide Faults

•Ultrathin oxide reliability

- Rate of defect generation decreases with supply voltage
- Tunnel current increases exponentially with decreasing gate oxide thickness

•Soft breakdown (SBD)

- **Intermittent** fluctuating current, high leakage
- SBD examples
 - Erratic erasure of flash memory cells
 - Erratic fluctuations of V_{min} in SRAM




SRAM V_{min}
90 nm technology

Source: M. Agostinelli et al,
IEDM 2005

Scaling Trend of the Vmin Sensitivity

Vmin sensitivity to gate leakage




Source: M. Agostinelli et al, IEDM 2005

Impact of Process Variations

- Increasingly difficult to accurately control device parameters
 - Channel length and width
 - Oxide thickness
 - Doping profile
- Intra-die variations, e.g., different transistor voltage threshold within the same SRAM cell
 - **Intermittent** failure of read/write operations
- Impact of process variations is increasing with scaling

Activation of Intermittent Faults



Voltage and frequency shmoo

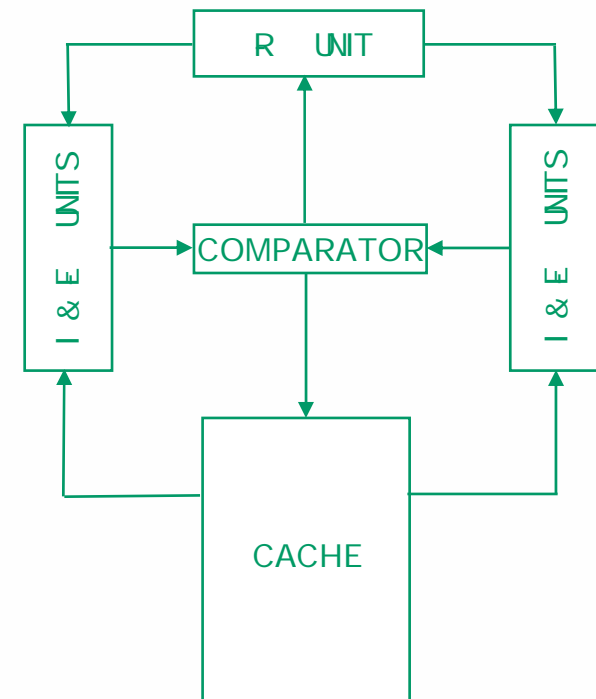
- Voltage
- Frequency
- Temperature
- Workload

Mitigation Techniques

HW Solutions: IBM G5/G6 CPU



- Mirrored Instruction and Execution units
- Comparator and register unit
- Compare outputs in n-1 instruction pipeline stage
 - No error: update checkpoint array (register content and instruction address into R-unit) in last pipeline stage and continue normal execution
 - Error detected: Reset CPU (except R-unit), purge cache and its directory, reload last correct state from checkpoint array, retry
- Transient faults are recovered from
- Error threshold can be used for intermittent faults
- Permanent faults require activation of a spare CPU under OS control



Source: L. Spainhower, T. A. Greg, IBM JR&D, 1999

HW Solutions: IBM G5/G6 CPU

- Pros

- Lower design complexity
- Shorter development and validation time
- No performance penalty (compare and detect cycles are overlapped)


- Cons

- Total circuit overhead about 40%
- It may not scale well with frequency

SW Solutions: AR-SMT

•Active-stream/Redundant-stream Simultaneous Multithreading (AR-SMT)

- Two copies of the same program run concurrently, using the SMT micro architecture
- Results of the two threads are compared
- A-STREAM errors are detected with a delay
- R-STREAM errors are detected before commit
- Recovery from transient faults (e.g. particle induced soft error) is possible
- Use committed state of R-STREAM



Source: E. Rotenberg, FTCS, 1999

SW Solutions: AR-SMT

- Pros

- AR-SMT relies on existing micro-architectural features, e.g. SMT
- No HW overhead

- Cons

- Increased execution time, 10% - 30%
- Increased performance penalty or even failure in the case of bursts of high frequency errors

Comparing Fault/Error Handling Techniques

- HW implementations are fast (e.g. ECC) - can handle bursts of errors induced by intermittent faults
- SW detection and recovery is slower
 - Performance penalty in the case of large bursts of errors
 - Near coincident fault scenario, in the case of high rate bursts of errors => SW fault/error handling may fail before recovery is completed
- SW solutions are better suited for failure prediction and resource reconfiguration

Summary

- Semiconductor technology is a two edge sword
 - Lower dimensions and voltages and higher frequencies have led to tremendous performance gains
 - Intermittent and transient faults have become a serious challenge to developers and manufacturers
- Designing for particle induced soft errors is too narrowly focused
- Software only techniques cannot effectively handle bursts of errors occurring at a high rate

FAULT TOLERANT CHIPS ARE THE FUTURE

Q & A

