ORACLE®

System Level Issues due to Technology Trends

60th Meeting of the IFIP 10.4 Working Group
on Dependable Computing and Fault Tolerance

Alan Wood
July 3, 2011

Agenda

- Technology trends
- Soft error rate (SER) trends
- Servers in 2020
- Server dependability issues
- Dependability issues for other systems

The Largest Scale

- ExaFlops supercomputer (10^18) in 2020

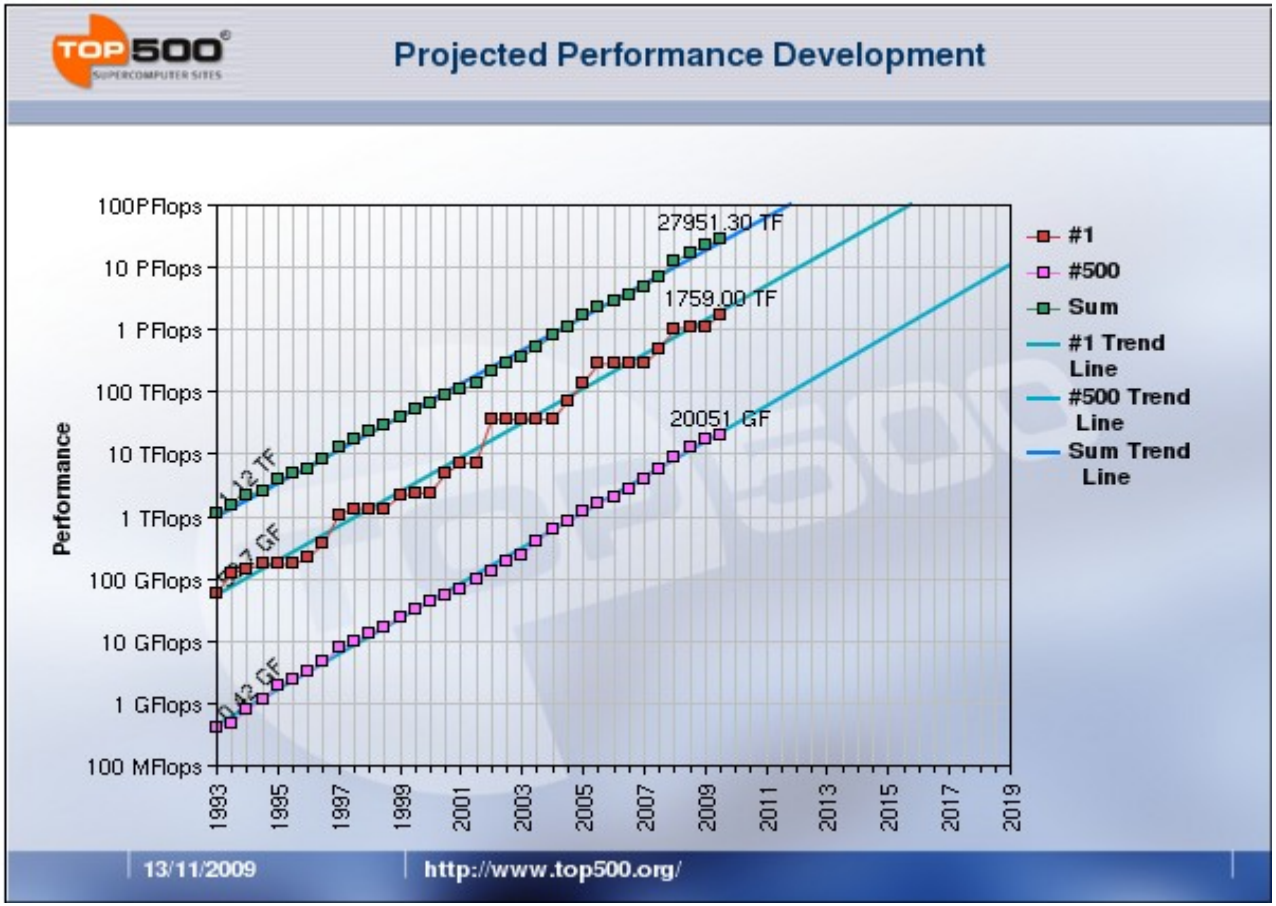

ExaScale Computing Challenges

- Energy – both for base computation and data transport
- Memory and Storage – bandwidth
- Concurrency and Locality – support for a billion parallel threads
- Resiliency - “the ability of a system to continue operation in the presence of either faults or performance fluctuations.”
  - Explosive growth in component count for large systems
  - Advanced technology
    - Lower voltage levels
    - New classes of aging effects

Source: DARPA ExaScale Computing Study

Equivalent Technology Scaling

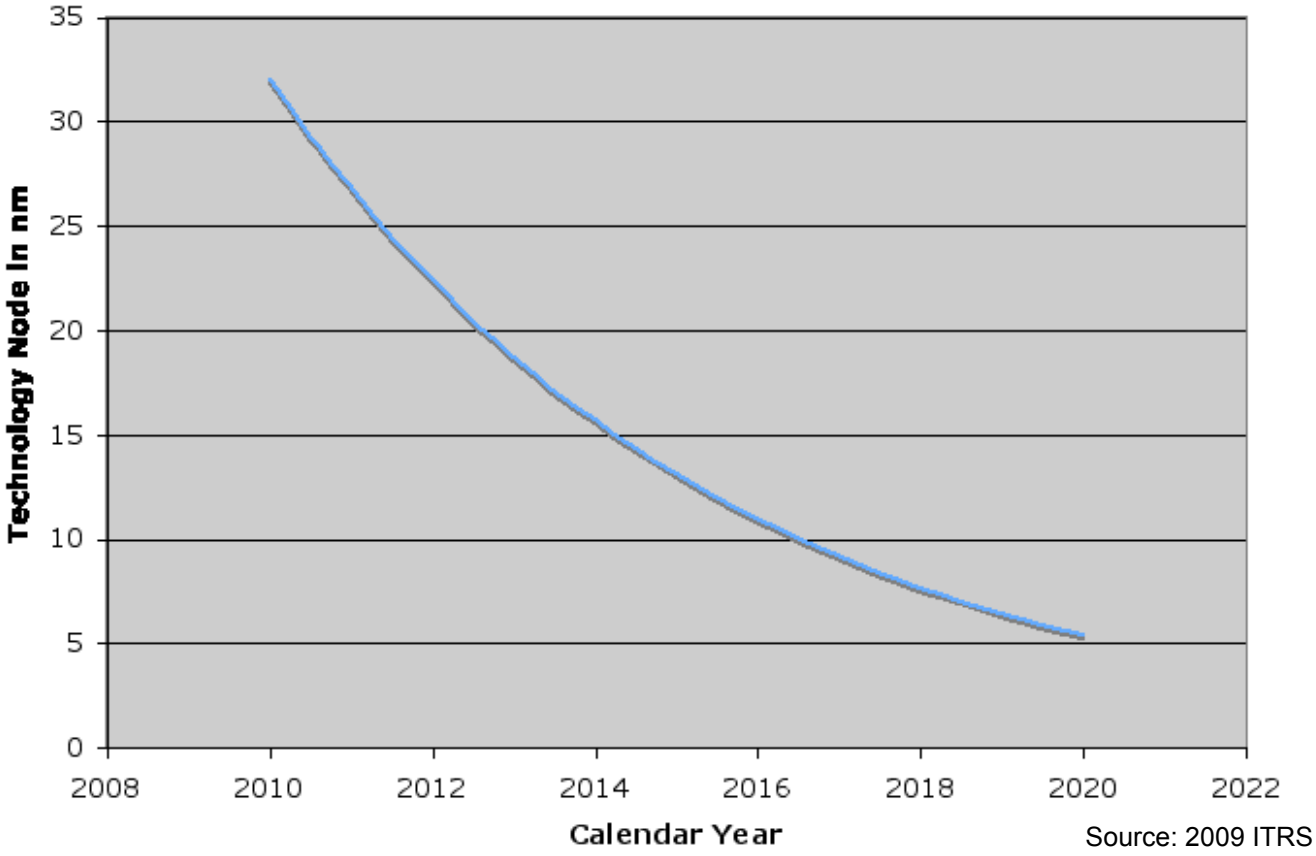

Source: 2009 ITRS

"Equivalent" scaling means the number of functions doubles every 2 years (does not mean half pitch, gate length, feature size)

Feature Size Scaling

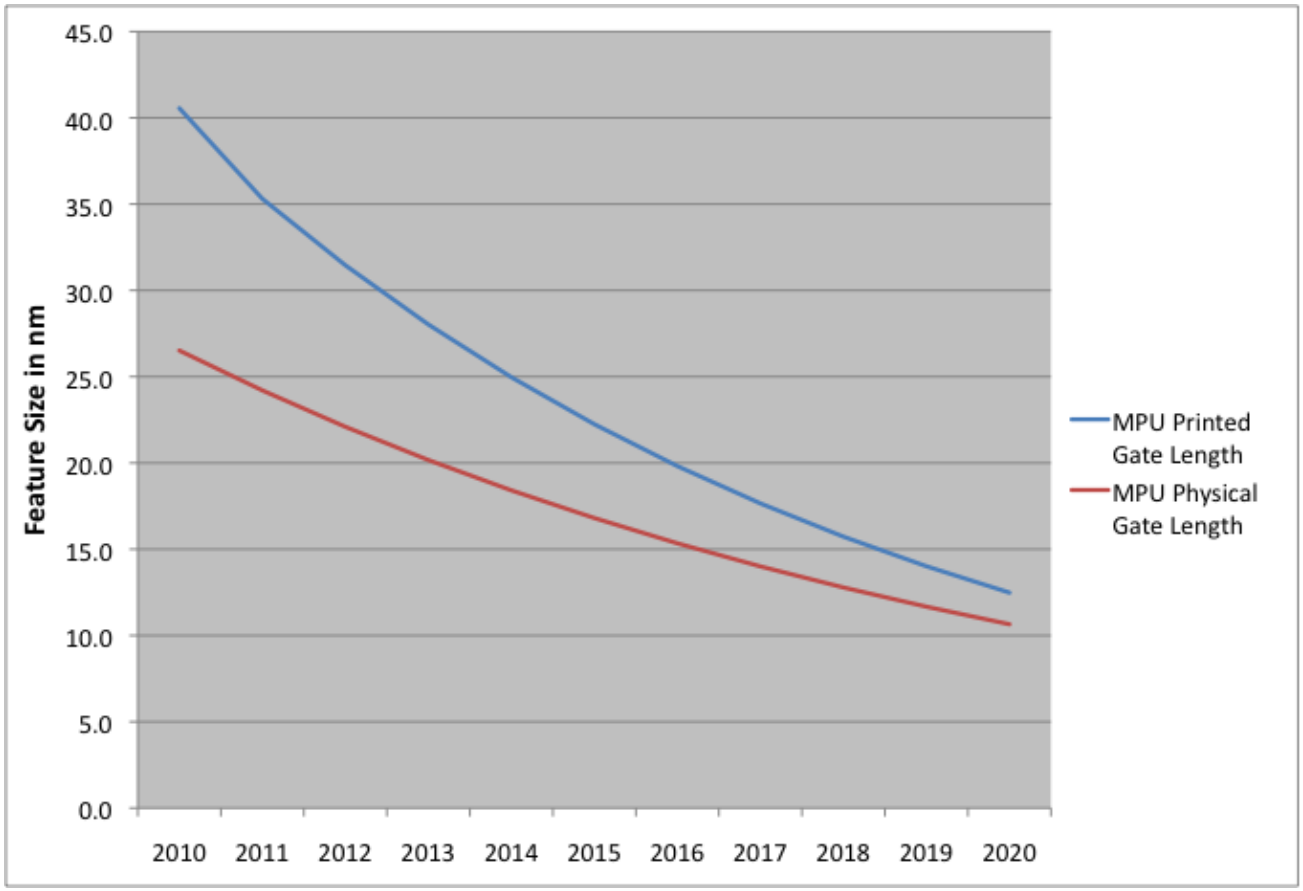

Feature size scaling not quite at Moore's law rate but still worrisome for SER trends

Source: 2009 ITRS

Core Count and Throughput Scaling

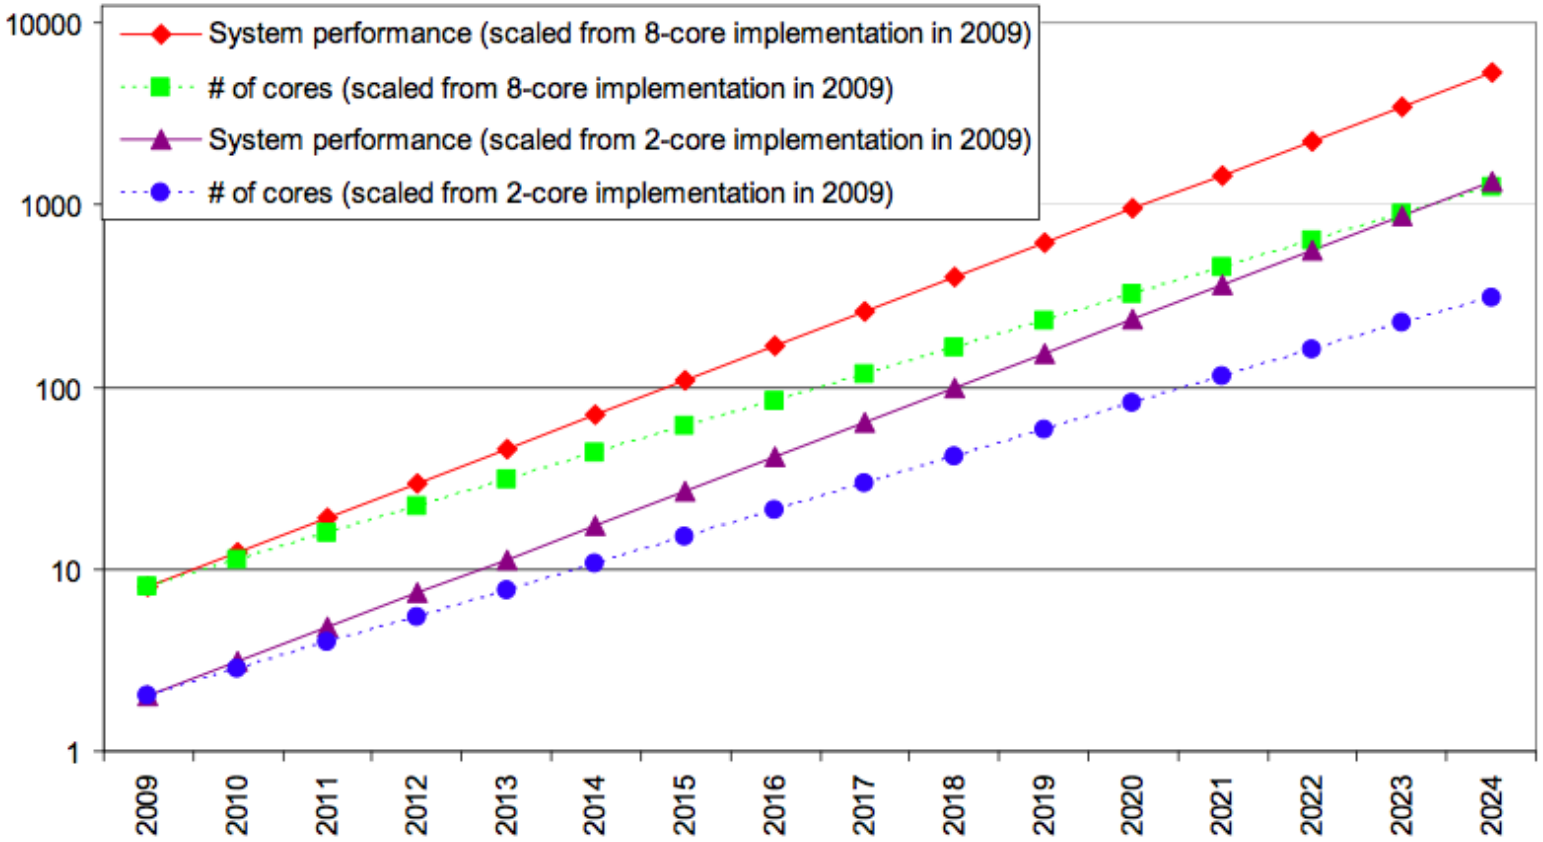

1000x by 2025: Same die size/power

Source: 2009 ITRS

Note: throughput is cores x frequency x instruction per clock

Power and Frequency per Core

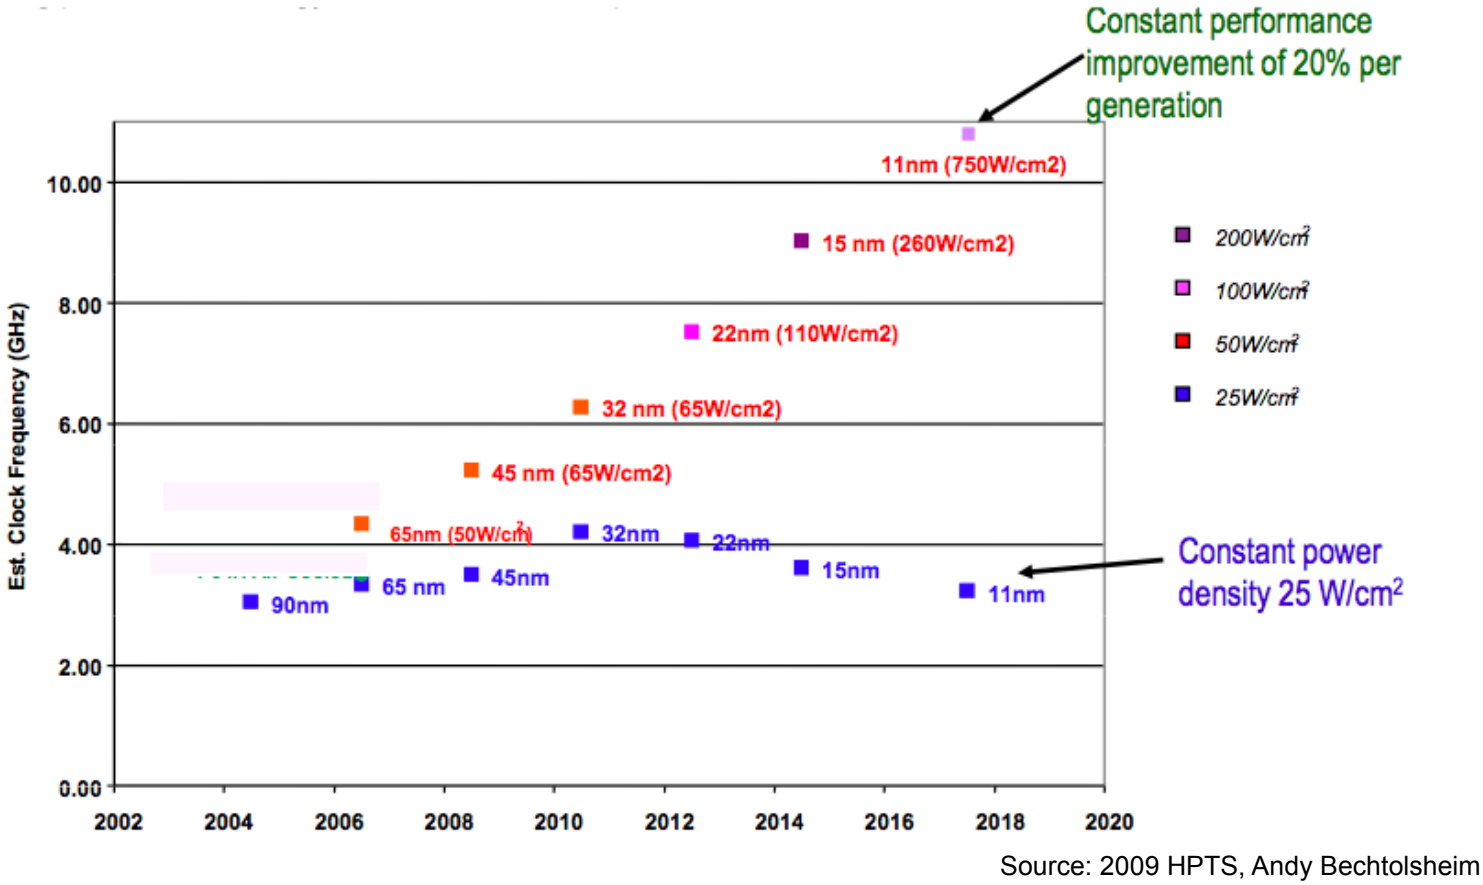

Source: 2009 HPTS, Andy Bechtolsheim

Power is the limiting factor

Bandwidth Scaling – Memory and I/O

- Number of pins/pads not scaling
  - 2x increase by 2025 (ITRS)
  - Need 50-100x Bandwidth by 2020
- 3D integrated MCMs
- Silicon photonics

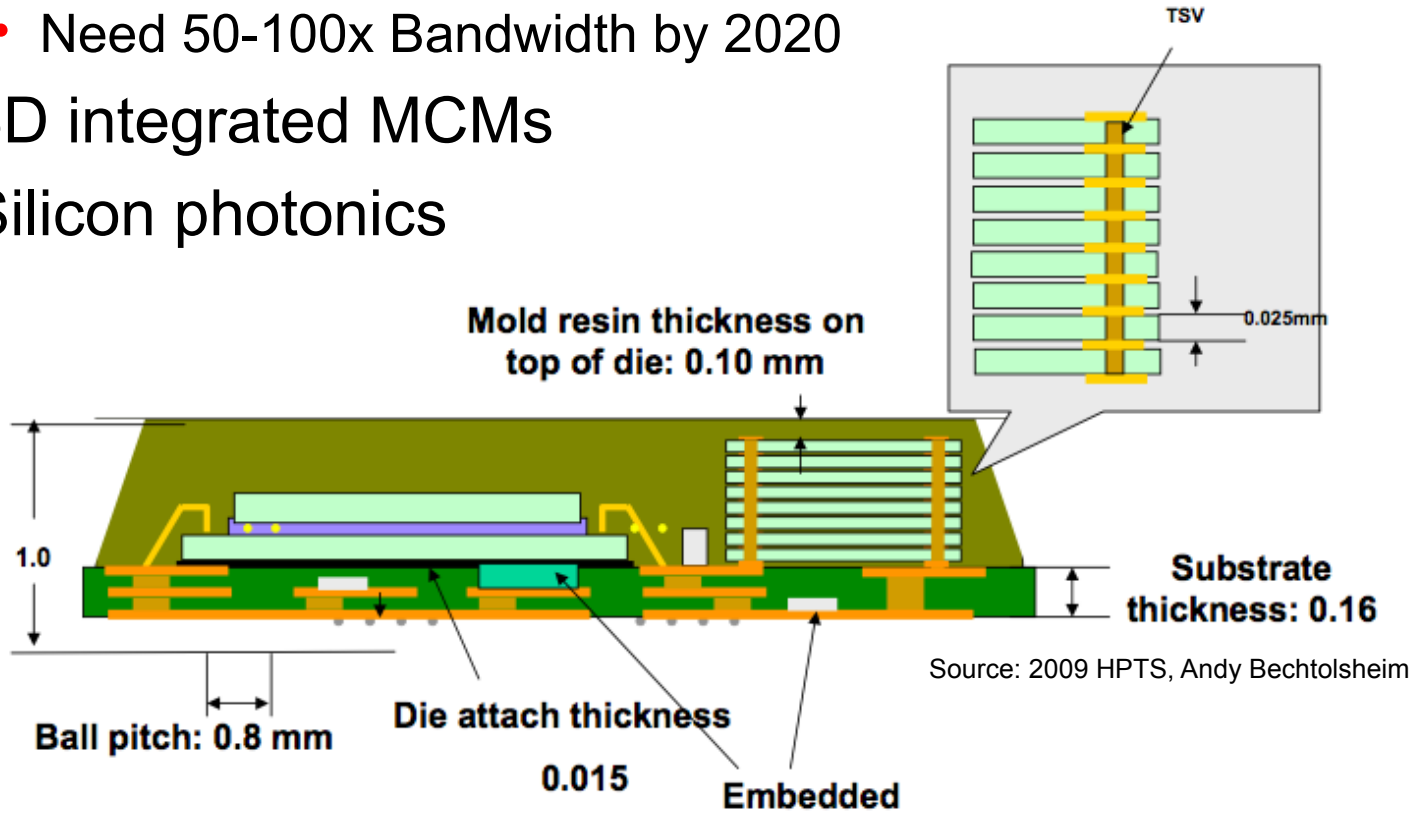

Source: 2009 HPTS, Andy Bechtolsheim

Agenda

- Technology trends
- **Soft error rate (SER) trends**
- Servers in 2020
- Server dependability issues
- Dependability issues for other systems

DRAM SER Trend

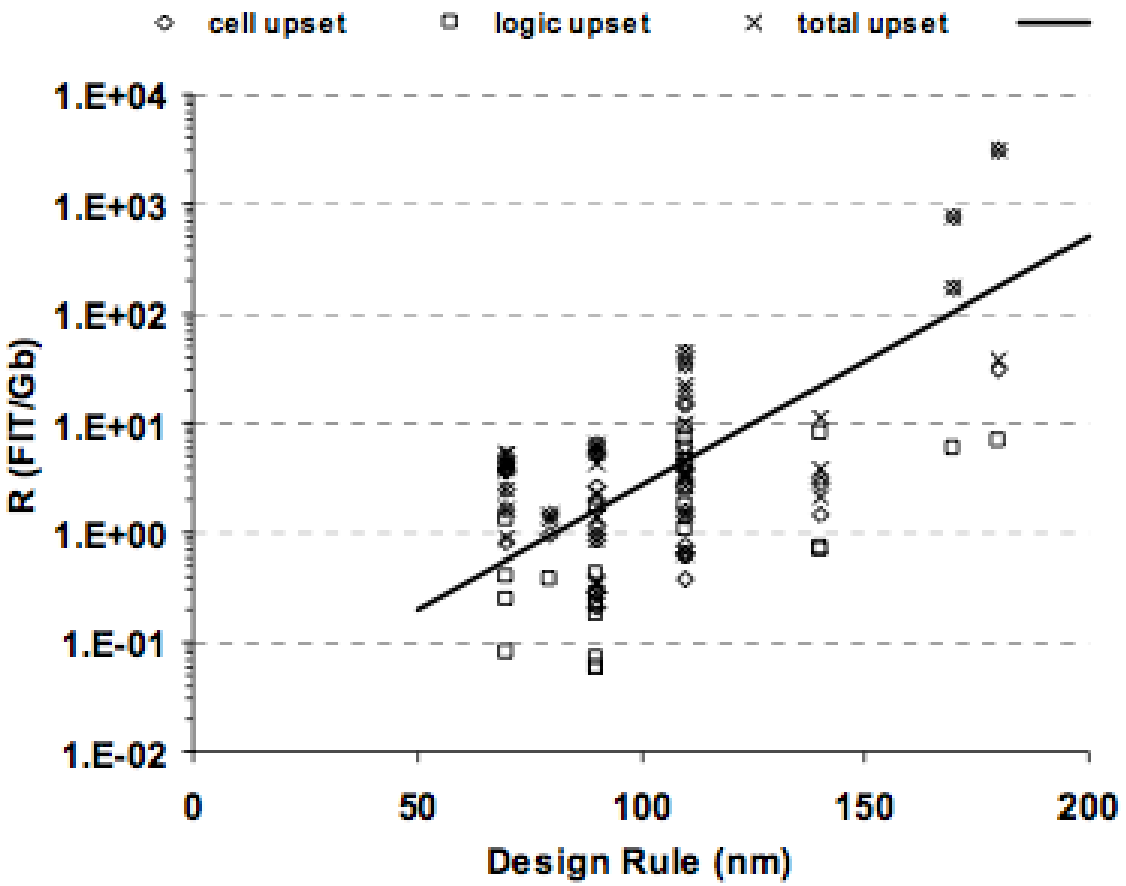

Source: L. Borucki, G. Schindlbeck and C. Slayman, “Comparison of Accelerated DRAM Soft Error Rates Measured at Component and System Level”, IRPS, Phoenix, 2008

SRAM and Logic SER Trend- Oracle

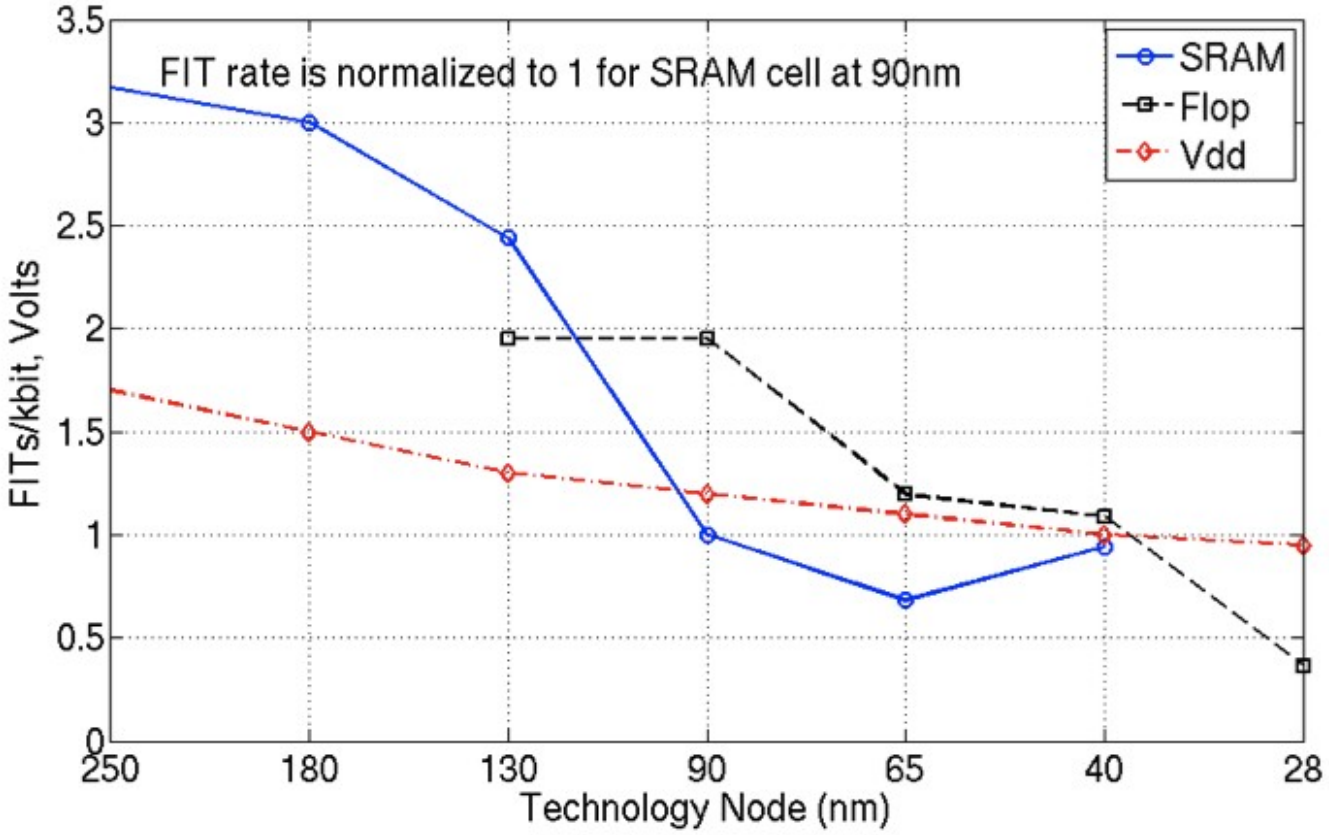

Source: Anand Dixit and Alan Wood, "The Impact of New Technology on Soft Error Rates", IRPS 2011

SRAM SER Trend- AMD

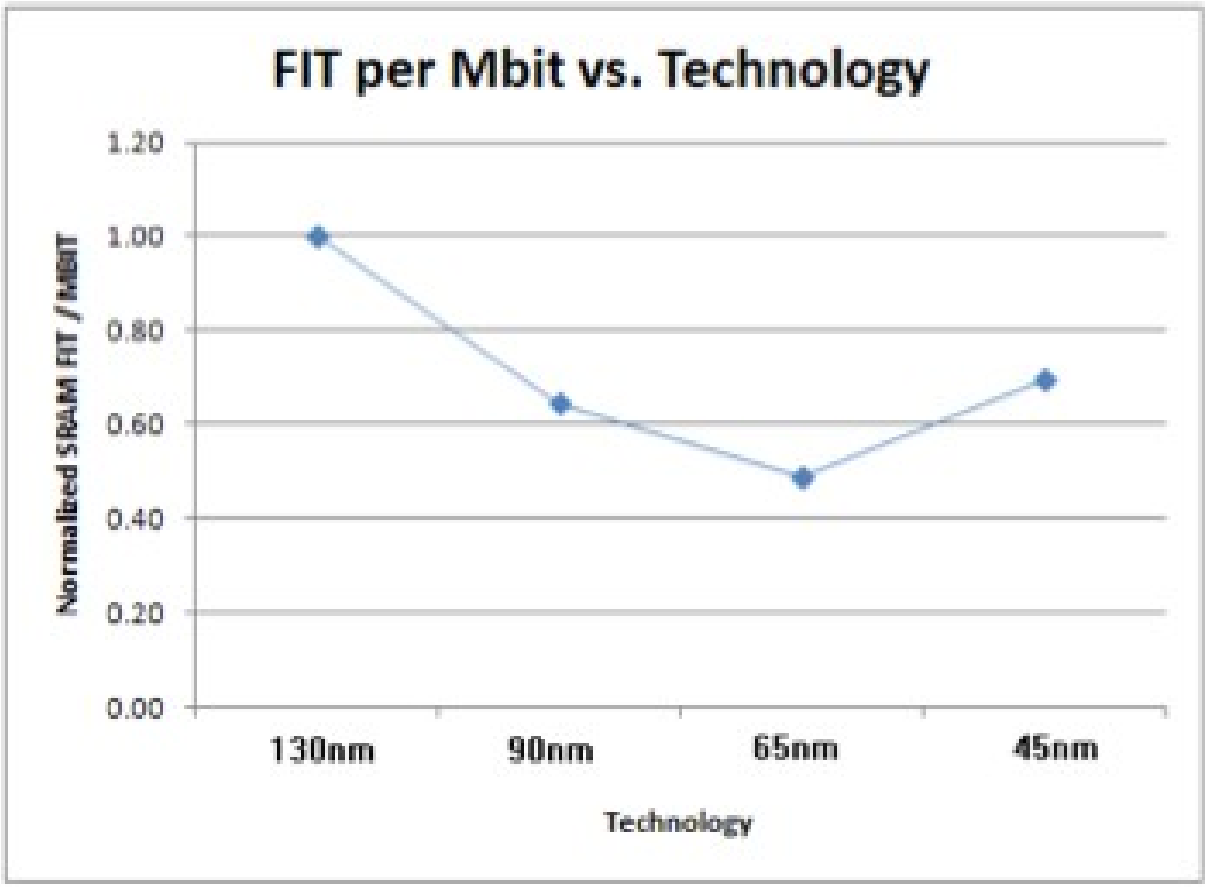

Source: Seth Prejean, “Accelerated Neutron Soft Error Rate Testing of AMD Microprocessors”, SELSE-6, Stanford, 2010

Logic SER Trend as a Function of Voltage

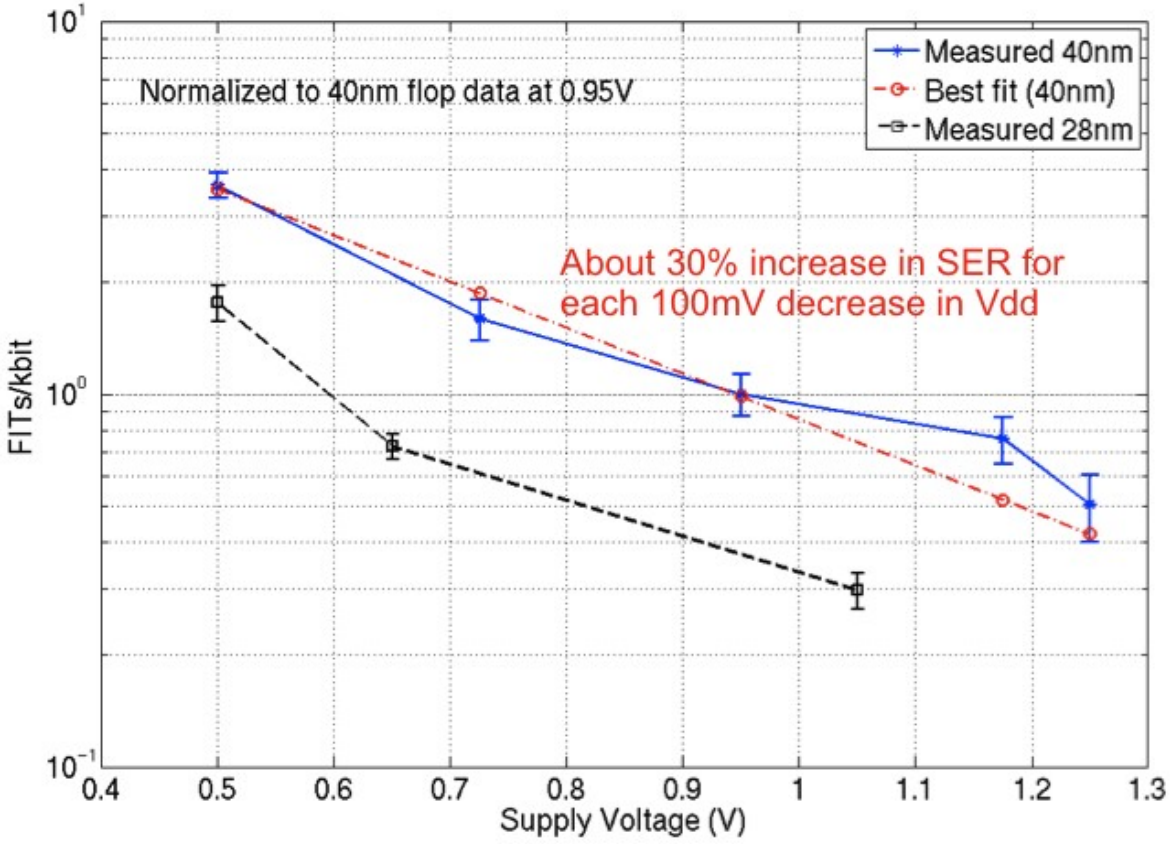

Source: Anand Dixit and Alan Wood, "The Impact of New Technology on Soft Error Rates", IRPS 2011

Microprocessor SER Trend- Oracle

Tech. (nm)	Relative SEU in FITs/kbit	Mbits/Processor	Relative uncorrected SEU/Processor
250	3.2	1.52	5.0
180	3.0	1.52	4.3
130	2.4	3.28	7.9
90	1.0	33.6	33.6
65	0.7	44.3	30.5
40	0.94	71	67.0

Reflects processor design over the years

Source: Anand Dixit, Raymond Heald, and Alan Wood, "The Impact of New Technology on Soft Error Rates", SELSE-6, Stanford, 2010

Agenda

- Technology trends
- Soft error rate (SER) trends
- **Servers in 2020**
- Server dependability issues
- Dependability issues for other systems

Servers in 2020

- Microprocessors
  - ~6-8nm technology (equivalent scaling)
  - ~256 cores per chip
  - ~16 Billion transistors per chip
  - Mostly SOCs
  - CMOS replacement?
- Memory
  - Stacked or embedded (no DIMMs)
  - Flash important part of memory hierarchy
  - New technologies (PCRAM, NRAM, …)

Servers in 2020 - 2

- Storage
  - Flash (SSDs) everywhere
  - New technology (holographic)?
- I/O
  - Fiber everywhere
  - Silicon photonics?
- Packaging
  - 3D
  - Liquid cooling
    - Including on-chip, e.g., heat pipes
  - Free-space optics?

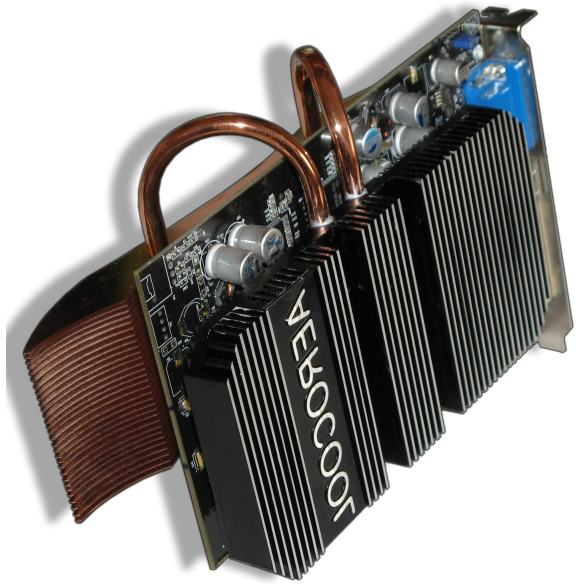

Server Reliability Environment in 2020

- Data centers operate at higher temperature and humidity
- Increased power and temperature cycling
  - From turning off equipment
  - From throttling performance
- Hopefully less vibration due to liquid cooling
- More operating and maintenance software
  - Reduced power states
  - Software control of hardware
- Parallel programming for 100,000+ threads
  - Instruction set parallelism
  - Functional parallelism within node
  - Data parallelism across nodes

Server Reliability in 2020

- Most chips will have manufacturing defects
  - Increased process variability
  - Decreased margins
- Errors/failures will be common
  - Logic soft error rates at least an order of magnitude higher
  - Also more susceptible to NBTI, HCI, other failure modes
  - Need to compute through failure
- On-chip sparing will be common
  - For both initial configuration and dynamic sparing
  - e.g., 132 cores on a microprocessor, ship with 128
  - Map out bad cache/memory areas, even I/O controllers

Dependability Concepts in 2020

- Focus on techniques that minimize power
  - Information redundancy (on larger blocks of data)
  - Self-checking logic, e.g., state machines
  - Control flow checking, data flow checking, assertion checking
- Tradeoff power consumption with recovery latency
  - Retry instead of active redundancy/in-line recovery
- (Re)configurable fault-tolerance
  - Application chooses
  - Reconfiguring around faults (or not, depending on application)
- Can software save hardware?

Agenda

- Technology trends
- Soft error rate (SER) trends
- Servers in 2020
- **Server dependability issues**
- Dependability issues for other systems

Trends and Dependability Impact

- Technology Scaling
  - (- - -) More transient errors
  - (-) Increased prevalence of some aging effects
  - (-) Design and test complexity
  - (+) Opportunity to use extra transistors for dependability
- Increasing cores/threads per CPU
  - (+) Better heat distribution across chip
  - (+) Repeated structures are easier to design and test
  - (+) Opportunity to have a dedicated diagnostic thread
    - (but be careful about energy)
  - (+) Redundant computations, e.g., threads at different precision for sanity checking (Los Alamos)
    - (but be careful about energy)

Trends and Dependability Impact - 2

- Memory/Storage
  - (+) Solid-state reliability/durability/power
  - (?) New technology
  - (+) No DIMM replacement
- Interconnect - optics
  - (-) Fiber attach/connector failures
  - (+) Less electrical noise
  - (+) Fiber handling/reliability (weight)
- Software
  - (-) Increased complexity

Trends and Dependability Impact - Packaging

- Liquid cooling
  - (+) Better temperature control
  - (+) Fewer fans
  - (-) New failure modes/disaster potential
- 3d MCMs
  - (-) New failure modes

Trends and Dependability Impact – Data Center

- Energy-efficiency
  - (-) Higher operating temperature and humidity
  - (+) Less vibration
  - (--) Redundancy may be limited
  - (-) Lower voltage/power increases soft error vulnerability
  - (?) Increased power and temperature cycling
    - From turning off equipment
    - From throttling performance
  - (-) Many new states to verify
    - Reduced power states
    - Software control of hardware
  - (?) Derating, hot spots, wearout, ... may all change

Agenda

- Technology trends
- Soft error rate (SER) trends
- Servers in 2020
- Server dependability issues
- **Dependability issues for other systems**

Societal Trend is Less Dependable Applications

- Computers were originally for science, then business, then commercial
- Now growth is in mobile, social
  - 350,000 iPhone apps
  - 500M+ Facebook users
- Means hardware can be less dependable
- Or does it?

New Mobile Financial Apps

Credit Card

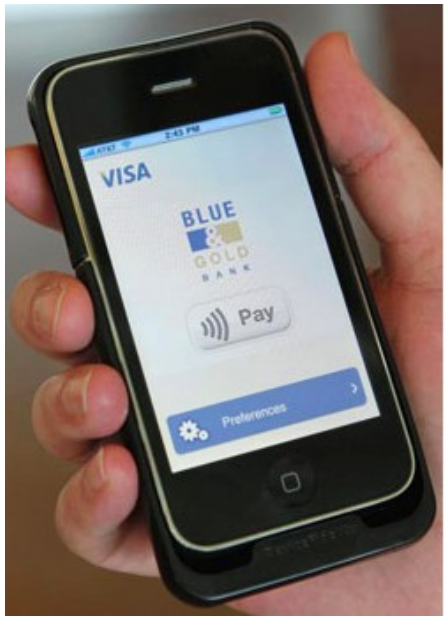

Bank Deposits

By 2020, you will be able to drive your car with your cell phone (my prediction)

Mobile Dependability Issues (IMHO)

- Wearout not important
- High-volume yield considerations will reduce infant mortality
- SER a big issue
- Security a huge issue
  - IEEE Spectrum, May, 2011: “Several programs available via Google's Android Market early this year appeared to be legitimate software, but hackers had actually added Trojans to them.”

Questions?

DRAM SER Trend Explanation

- DRAM memory cell SER has decreased by 2-3 orders of magnitude in the last 10 years
- Memory Cells
  - Basic DRAM cell has not changed much, so cell capacitance has not changed much, so Qcrit has not changed much
  - Charge collection area decreased by a factor of 2 with each generation
- DRAM Logic
  - Charge collection area has decreased, but decreases in voltage and different circuit designs has significantly decreased Qcrit, which significantly increases SER

SER Trend Explanation

Vdd ↓ → Critical Charge ↓ → SER ↑

Area ↓ → Sensitive depletion region ↓ → SER ↓

SER α Area * exp(-Qcrit/Qcoll)

Linear with Area; Exponential with Vdd

Source: Anand Dixit, Raymond Heald, and Alan Wood, “The Impact of New Technology on Soft Error Rates”, SELSE-6, Stanford, 2010