Designing Reliable, Power-Efficient Systems

Mary Jane Irwin (www.cse.psu.edu/~mji)

Vijay Narayanan, Mahmut Kandemir, Yuan Xie

CSE Embedded and Mobile Computing Center (emc^2)
Penn State University

Outline

- Motivation
- Tradeoffs
- Designing Reliable, Power-Efficient Interconnect
- Soft Error Design Issues

Tradeoffs

emc^2

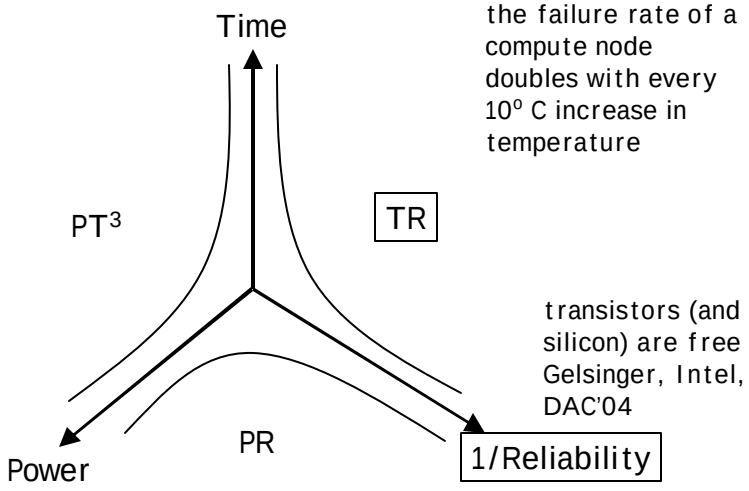

the failure rate of a compute node doubles with every 10^0 C increase in temperature

transistors (and silicon) are free Gelsinger, Intel, DAC'04

Threats to System Reliability

emc^2

- Noise: An unwanted signal or a disturbance in an electronic device [Webster] that makes signals deviate from their intended or ideal value.
  - The deviation is transient (temporary) and intermittent
  - Can be proportional (to the signal swing) or independent
    - With scaling, noise becomes more significant
- Sources of noise
  - **Crosstalk**
  - Process variation
  - Power supply noise; Substrate noise; Thermal noise
  - **Soft errors**
  - Electromagnetic interference

emc^2

What is Crosstalk?

- Crosstalk is the interaction (due to capacitive coupling) between signals on two different nets. The noise on one wire is induced by switching activity on neighboring wires.
  - Can cause a propagation delay – crosstalk delay
  - Can cause a voltage spike – crosstalk glitch
- Capacitive coupling is becoming more significant.
  - Wire spacing shrinks faster than wire height
  - Clock frequency increases – so delay is a more critical issue

"Crosstalk due to coupling capacitance between adjacent interconnect lines, in 0.18 micron and below, has become a major performance limiting factor that can cause both noise injection and signal timing deviation."

-Bruno Franzini, STMicroelectronics, SNUG Europe 01.

emc^2

Capacitive Coupling Example

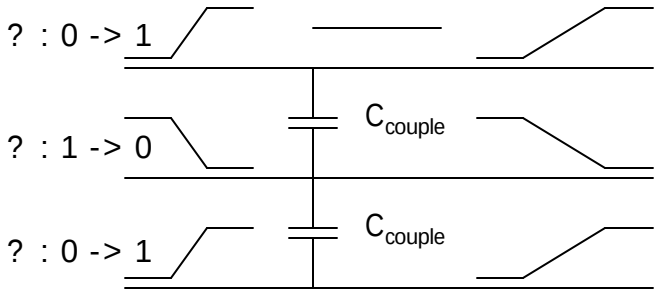

When adjacent signals transition in opposite directions, the crosstalk delay is longest.

In this case: $C_{total} = 4*C_{couple} + C_{ground}$

Solutions for Crosstalk

emc^2

- Crosstalk reduction techniques
  - **Shielding**
  - Spacing
  - **Crosstalk aware signal coding and transmission**
  - Buffer insertion
  - Wire ordering
- Crosstalk tolerance techniques
  - On-line detection
  - Bus guardian
- Combination of low power and reliability techniques
  - Low-power error resilient encoding
  - Adaptive low power transmission schemes
  - **Adaptive error protection**

Crosstalk Reduction Techniques

emc^2

- Shielding: grounded wires between signal wires

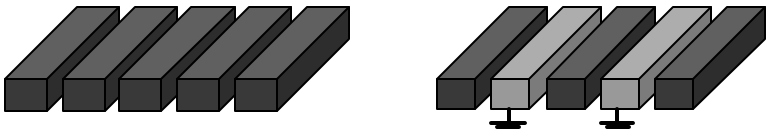

  - Removes the opposite-direction transitions in adjacent wires
  - Drawback – wiring area doubles
- Shielding can be thought of as a data encoding where two wires are required for every signal bit
  - "1" is encoded as "10" and "0" is encoded as "00"
- Crosstalk aware signal coding: Are there other encodings that can prevent adjacent wires on a bus from transitioning in opposite directions?
  - Successive codewords cannot allow a ***rising bit transition*** next to a ***falling bit transition***.

Crosstalk Coding Overhead

- Need 40%~50% extra bit lines
- Need an encoder and decoder
- Designing the encoder and decoder for a large number of bits may be impractical; encode *n*-bit blocks and put ground shields between them.

Bits of symbol	Bits of codeword	Extra bits
3	4	33%
4	6	50%
8	12	50%
16	23	43.8%
32	46	43.8%

Victor & Keutzer, Bus Encoding to Prevent Crosstalk Delay , ICCAD'01.

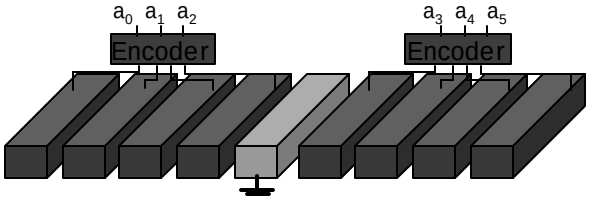

Crosstalk Aware Interconnect

- Instead of signal coding allow a variable number of transmission cycles depending on successive code words and the resulting crosstalk delay

Lin, et.al., A Crosstalk Aware Interconnect with Variable Cycle Transmission, DATE'04

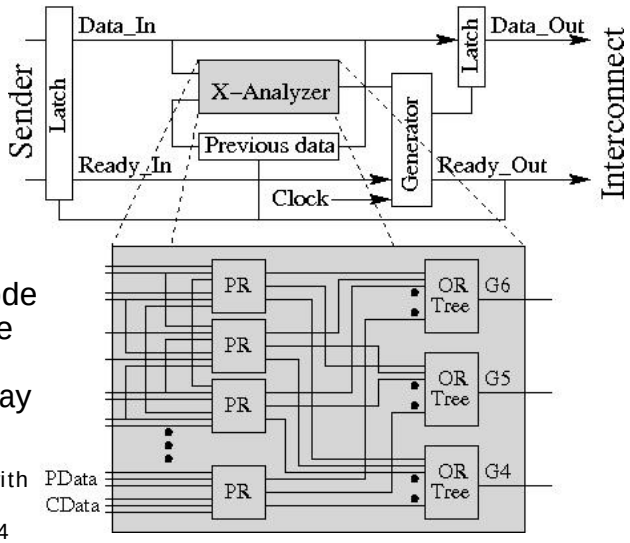

Crosstalk Patterns

- Different transmission patterns have different C_{total} and thus have different delay

? : 0 -> 1

? : 1 -> 0

- : no change

C_{total} of line k	Patterns (lines k-1, k, k+1)			
0	- - -			
	- - ?	- - ?	? - -	? - -
	? - ?	? - ?	? - ?	? - ?
C_{ground}	? ? ?	? ? ?		
$C_{couple} + C_{ground}$	- ? ?	- ? ?	? ? -	? ? -
$2*C_{couple} + C_{ground}$	- ? -	- ? -		
	? ? ?	? ? ?	? ? ?	? ? ?
$3*C_{couple} + C_{ground}$	- ? ?	- ? ?	? ? -	? ? -
$4*C_{couple} + C_{ground}$	? ? ?	? ? ?		

Sotiriadis & Chandrakasan, Reducing bus delay in submicron technology using coding, ASP-DAC'01

Crosstalk Pattern Analysis

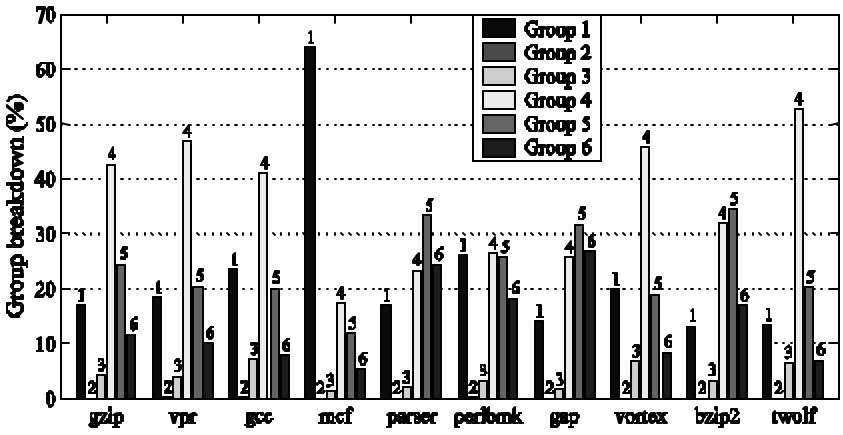

The average distributions are 22.64%, 0.05%, 4.06%, 35.4%, 24.2%, and 13.7% for Group 1 through Group 6.

emc^2

Variable Cycle Transmission

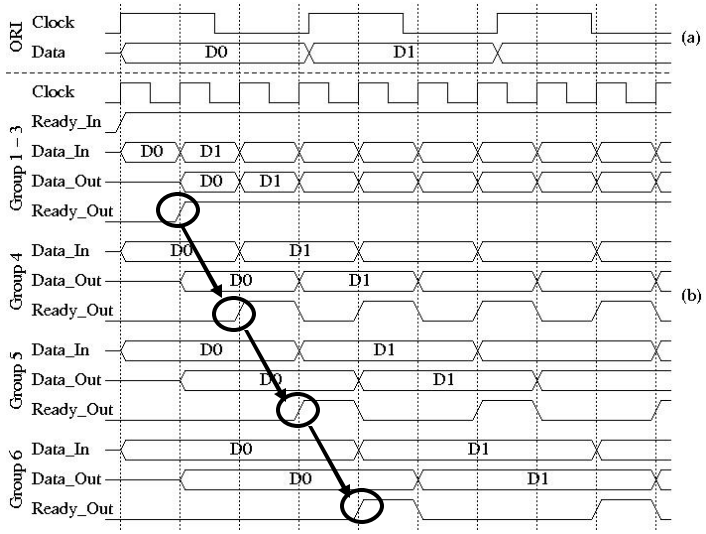

emc^2

Parameters and Area Overheads

ORI: Original interconnect

CPC: Crosstalk Prevention Coding

DYN: Variable cycle crosstalk detection

DBS: Double spacing

SHD: Shielding

	C_{ground} (fF/mm)	C_{couple} (fF/mm)	# of wires	Normalized cycle time
ORI	36.3	115.1	32	3.28
CPC	36.3	115.1	53	2.76
DYN	36.3	115.1	33	1.00
DBS	53.1	60.4	32	1.95
SHD	36.3	115.1	63	1.76

Normalized area for each scheme (including bus and codec)

	2mm	5mm	10mm
ORI	100	100	100
CPC	174	170	168
DYN	132	113	106
DBS	149	149	149
SHD	198	198	198

Performance Results

emc^2

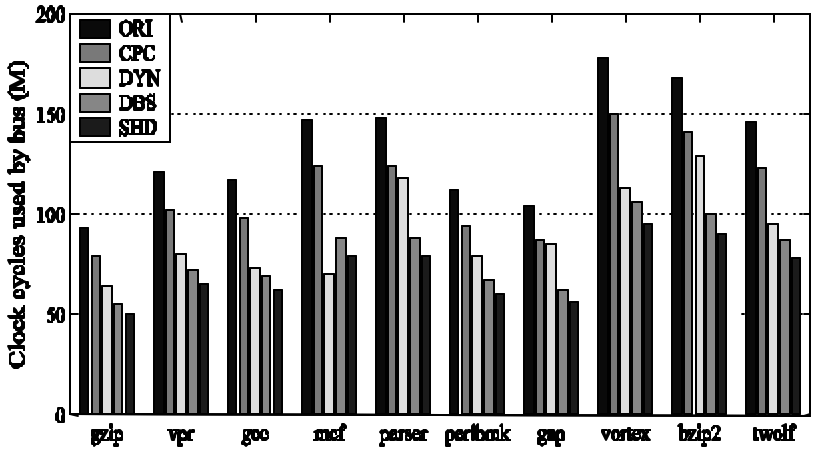

DYN provides an average of 31.5% performance improvement over ORI.

Energy Consumption Results

emc^2

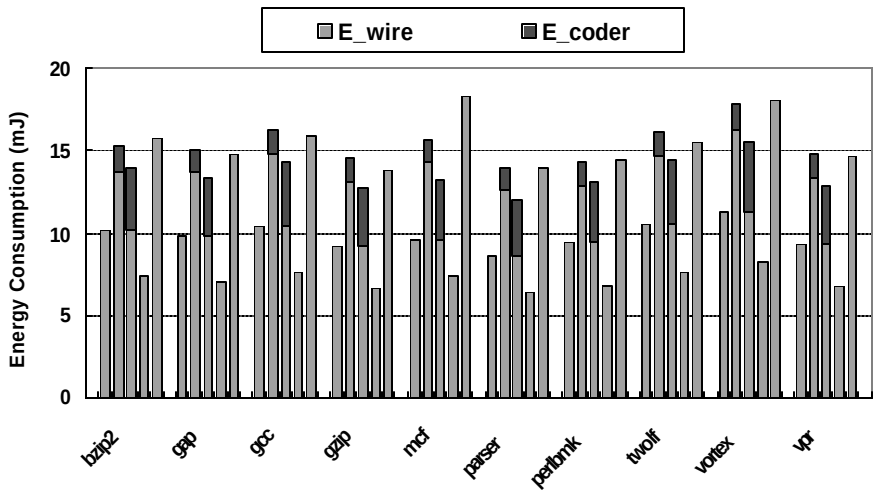

From left to right are ORI, CPC, DYN, DBS, SHD

Crosstalk Protection Techniques

emc^2

- Error detection coding + correction or retransmission
- Different coding methods have different error detection capabilities and different energy overheads
  - Parity (PAR): 1 extra bit, detects all odd number errors
  - Double Error Detection (DED): (38,32) Hamming code
  - Triple Error Detection (TED): (38,32) Hamming code + Parity

	Encoder	Decoder	Total
PAR	1.0	1.0	1.0
DED	2.1	2.4	2.3
TED	2.7	3.5	3.1

- Correction versus retransmission
  - Retransmission: More delay and more bus transitions but simpler codec and more detection capability
  - Correction: Smaller delay and fewer bus transitions but more complex codec

Adaptive Error Protection

emc^2

- Designing an error protection scheme for the worst case scenario may not be energy efficient
  - The more powerful an error protection scheme, the more energy it consumes.
  - Noise behavior varies over time due to environmental factors and operational conditions

→ An adaptive error protection scheme for on-chip interconnect that adapts the strength of the error detection scheme dynamically based on the noise behavior observed.

  - Detecting the variation in noise behavior
  - Identifying the protection scheme to employ for the observed noise behavior

Lin, et.al., Adaptive Error Protection for Energy Efficiency, ICCAD'03

Detecting the Variation of Noise

emc^2

- Detected error rates are an indicator of the variation in the noise
  - A small variation in detectable error rates indicate huge variation in undetectable error rates

- Victim bus line
  - uses half the voltage swing of the normal bus lines
  - so is more susceptible to variations in noise
  - amplifies the number of detectable errors

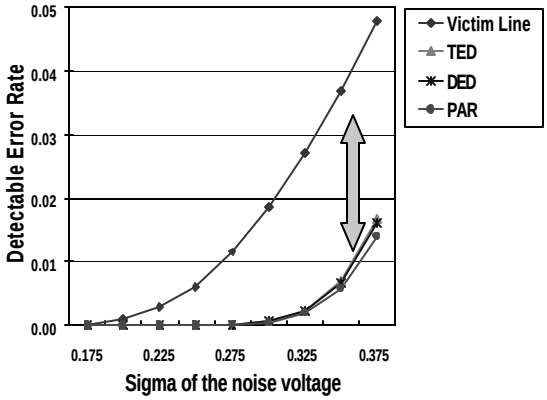

An Adaptive Scheme

emc^2

Keep the undetectable word error rate below a preset value (threshold), e.g., 1e-10.

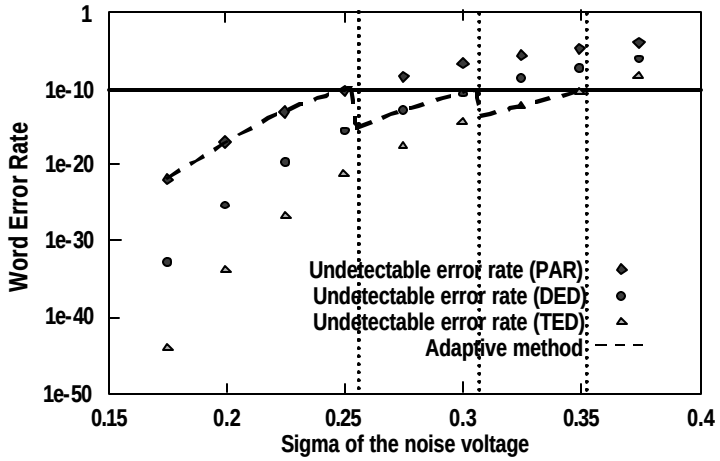

emc^2

Energy Consumption (Noise Profile 1)

Noise profile 1: Slow increase of the noise followed by a phase of slow decrease.

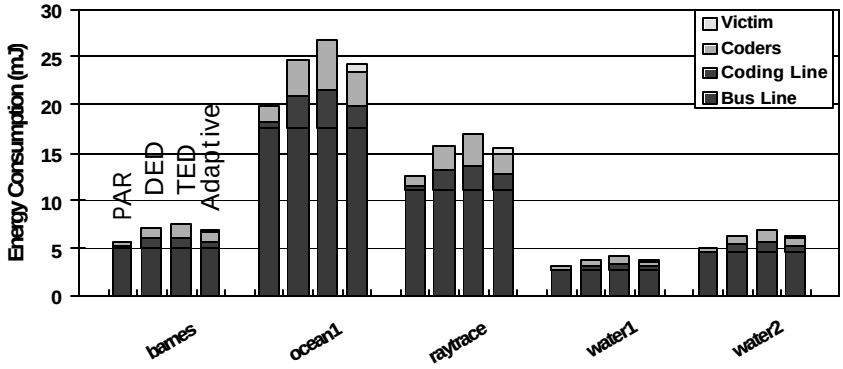

Adaptive scheme has 8% energy savings over TED.

emc^2

State Breakdown (Noise Profile 1)

Cumulative cycles spent in each scheme

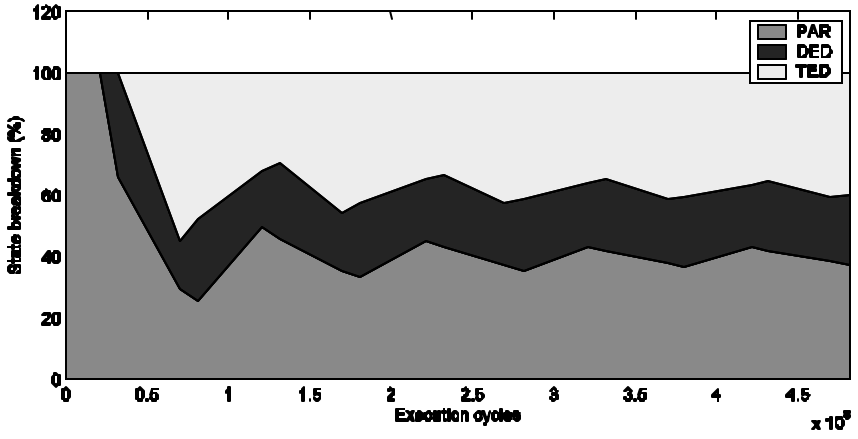

PAR and DED are used around 60% of the time.

Outline

- Motivation
- Tradeoffs
- Designing Reliable, Power-Efficient Interconnect
- Soft Error Design Issues

What are Soft Errors ?

- Soft Errors (or single event upsets – SEUs) are when the internal states of nodes are flipped due to excess charge carriers induced primarily by external radiation.
  - These errors cause an upset event but the circuit itself is not damaged.

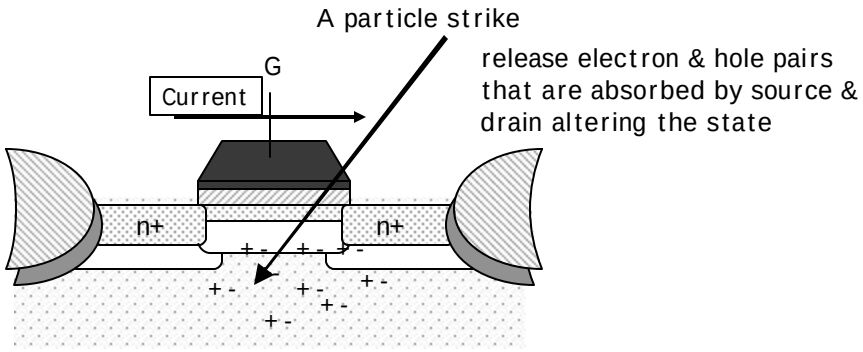

Problems Caused by SEU

emc^2

- Soft Errors can cause problems in different ways
  - Change the data value in the caches and memory
  - Corrupt the execution of an instruction due the flip of data in the pipeline registers
  - Change the character of a SRAM-based FGPA circuit (Firm Error)

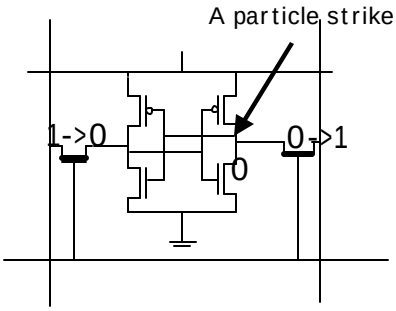

  - Datapath (combinational) logic SET (Single Event Transient) caught by registers/memory

What causes Soft Errors?

emc^2

- At ground level, there are three major contributors to Soft Errors
  - **Cosmic Ray induced neutrons** cause errors due the charge induced by Silicon Recoil
    - The upset rate increases with altitude by a factor of 2.36 every 1K meters
  - **Alpha particles** emitted by decaying radioactive impurities in packaging and interconnect materials
    - plastic packages are the worst - Ceramic, HyperBGA, Flip-chip PBGA
  - **Neutron induced ^{10}B fission** which releases a Alpha particle and ^{7}Li

emc^2

Soft Error Rate (SER)

- For a soft error to occur, the collected charge Q at a node should be more than $Q_{critical}$

$$SER \propto N_{flux} * CS * \exp\left(-\frac{Q_{critical}}{Q_s}\right)$$

  - N_{flux} : intensity of the neutron flux
  - CS : the area of the cross section of the node
  - $Q_{critical}$: critical charge necessary for a bit flip (proportional to the node capacitance and the supply voltage)
  - Q_s : charge collection efficiency (dependent on doping)
- As CMOS device sizes decrease, the charge stored at each node decreases (due to lower nodal capacitance and lower supply voltages) but the collection area also decreases
- So do SERs go up or down as technology scales?

emc^2

Impact of Scaling on SER (Bohkar, DAC)

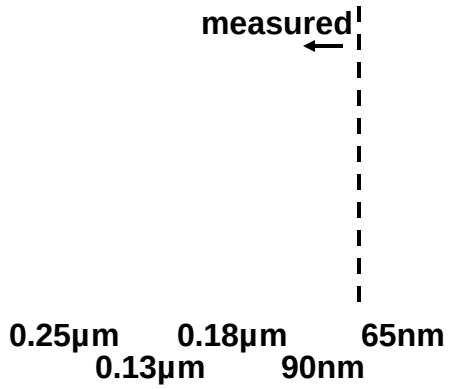

8% increase in SER/bit per generation

Leakage Vs Soft Error Susceptibility

emc^2

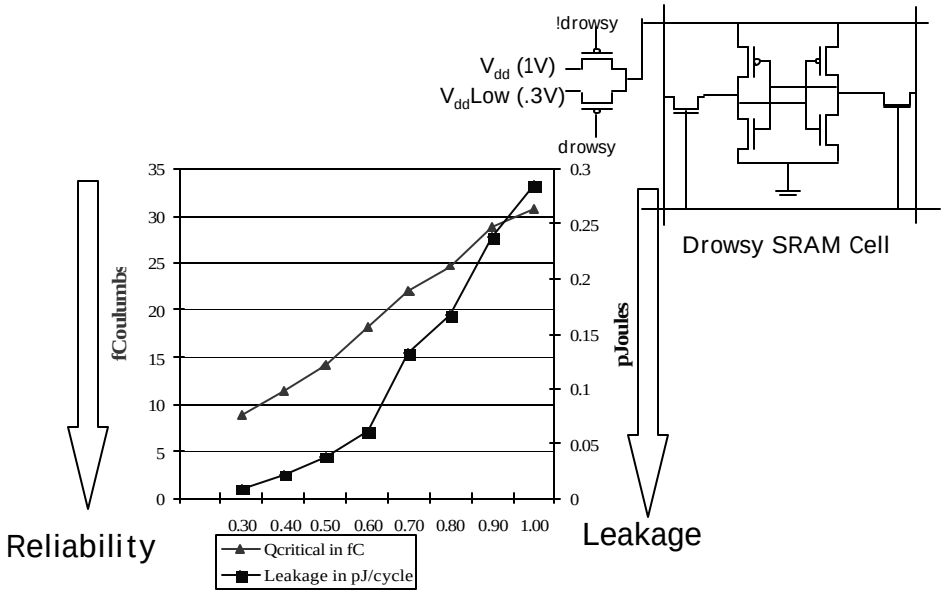

Supply Scaling Impacts on SEUs

emc^2

- SEU exposure results (Cypress CY7C128A: 2Kx8 SRAM)
  - 10e7 particles/sec beam intensity
  - 5V supply, 1 hour exposure ? 2 single SEU events
  - 3V supply, 1 hour exposure ? 13 single and 2 double SEU events

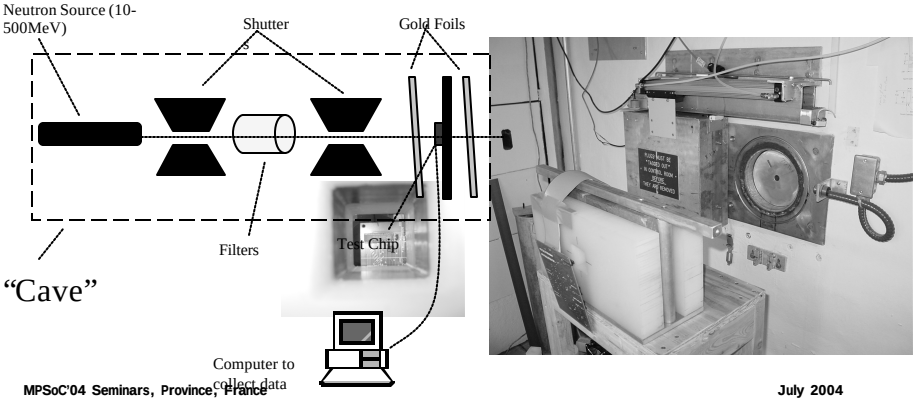

SER Solution Approaches

emc^2

- Physical solutions (may be hard!)
  - Shielding?
    - No practical absorbent (e.g., approximately > 10 ft of concrete)
  - Radiation-hardened cells?
    - 10x improvement possible with significant penalty in performance, area, cost
    - 2-4x improvement may be possible with less penalty
  - SOI
- Error detection and correction in memories
  - Adds datapath delay for ECC calculations
  - Cache scrubbers
- Circuit modifications to increase the node capacitance
  - Replace diffusion capacitance with gate capapcitance
  - Add active device (restorer inverter + pull-up transistor)

SER Solution Approaches, con't

emc^2

- Spatial redundancy
  - Hardware duplication and voting
- Temporal redundancy
  - Redundant multithreading
- Software techniques
  - ABFT-Algorithmic based fault-tolerance (CRC-Stanford) & Abraham (UT)
  - Procedure call duplication
    - Duplicate instructions but with different registers and variables
    - A *master* original instruction and a *shadow instruction* in the duplicated code
    - General purpose registers and memory are partitioned into two groups for master and shadow instructions
- ???

Thank You

Much of the research presented supported by the GSRC, a MARCO Focus Center