

SAND2008-5040P

Radiation Effects Issues and Trends for High Performance Space Computing

Paul E. Dodd

Sandia National Laboratories

pedodd@sandia.gov

(505) 844-1447

Fault-Tolerant Spaceborne Computing Employing
New Technologies Workshop

5/28/2008


This work was supported by the United States Department of Energy. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.



Classes of Radiation Effects

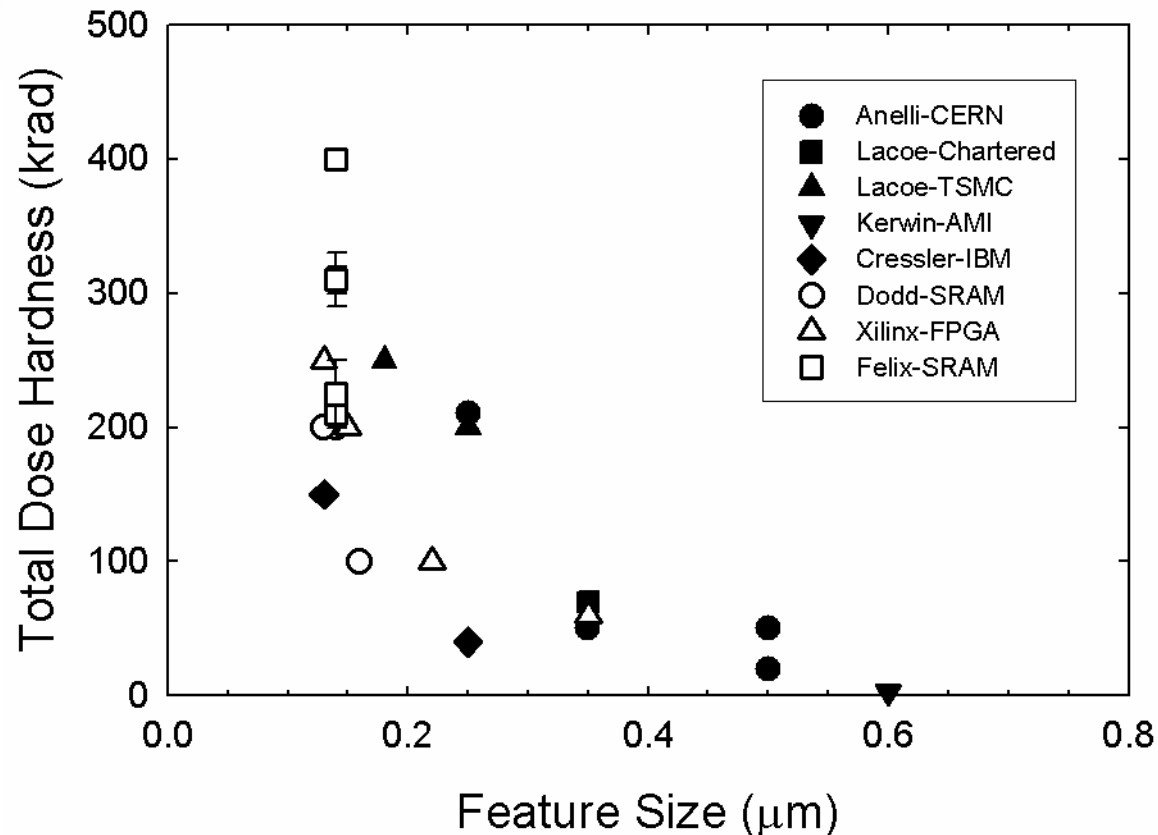
- Total Ionizing Dose (TID)
 - space and weapons environments
 - accumulated effects of trapped charge in oxides
- Single-Event Effects (SEE)
 - space and high-altitude environments, increasingly a concern at ground level
 - results from single energetic particle strike
- Transient (Dose-Rate)
 - nuclear weapons environment
 - burst of charge generation globally within devices
- Displacement Damage
 - space and weapons environments
 - accumulated effects of lattice damage in minority carrier devices

TID Hardness of Commercial Technologies is Limited by Field Oxide Leakage



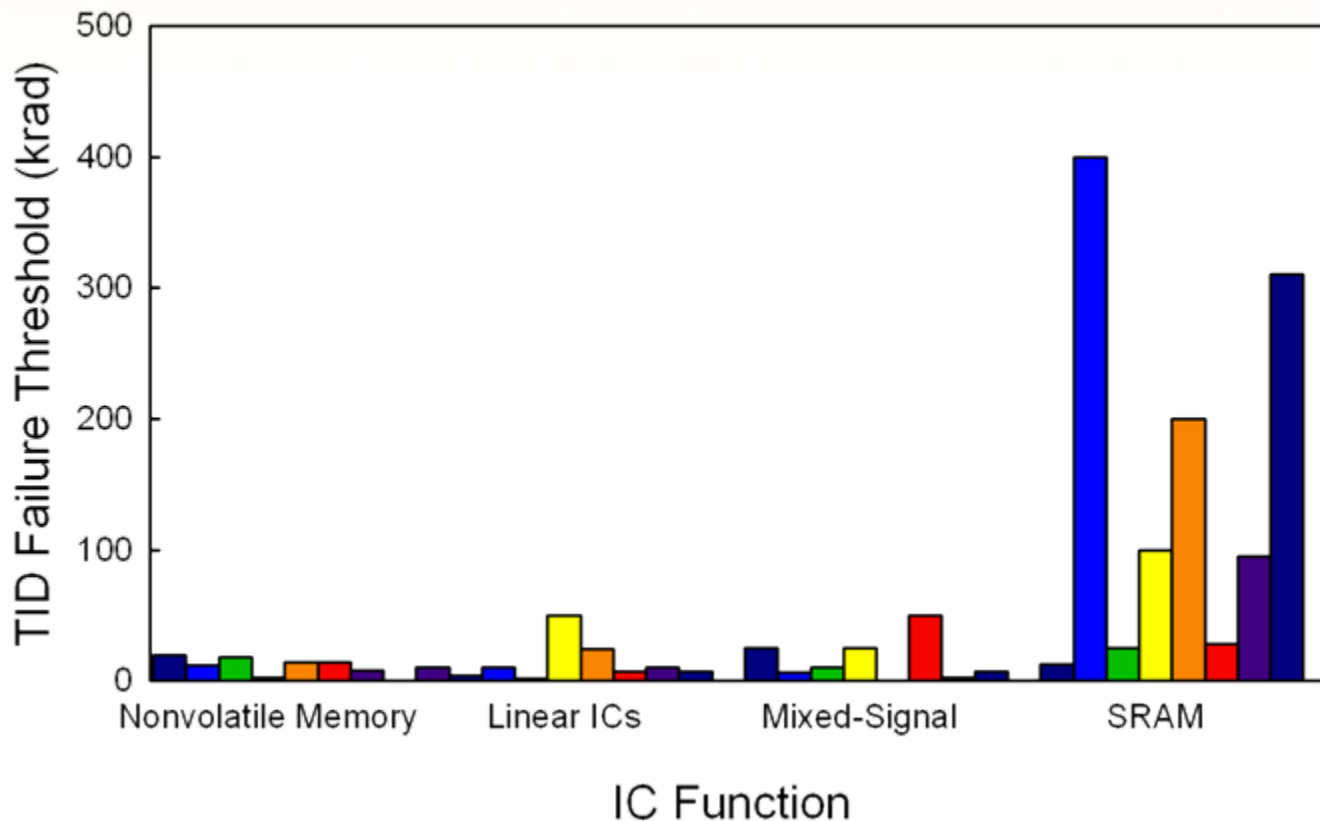
The migration to shallow trench isolation (STI) at 0.35 μm has generally led to improved TID hardness.

TID Hardness of Commercial CMOS Technologies is Generally Improving with Scaling



Overall trend is toward increased hardness with technology scaling, however, this serendipitous hardness is usually neither monitored nor controlled.

TID Hardness of Commercial Technologies Varies Widely Between Classes of ICs



Some classes of ICs (such as nonvolatile memory) are historically sensitive to total-dose radiation. Even within a class of IC, total dose hardness can vary widely between/within manufacturer.

Real Systems Use a Wide Variety of Microelectronic Technologies

- Even advanced systems contain a wide mix of technologies and generations
 - CMOS, bipolar, compound semiconductor
 - Digital, analog, mixed-signal, radio frequency
 - ICs, hybrids, and discretetes
 - 90 nm to several micron geometries
- Radiation response of these technologies can vary over a wide range
 - Different part types or technologies may be vulnerable to different components of a given radiation environment
 - System radiation response is limited by weakest component

Some Classes of Single-Event Effects (SEE)

Soft Errors (no permanent damage)

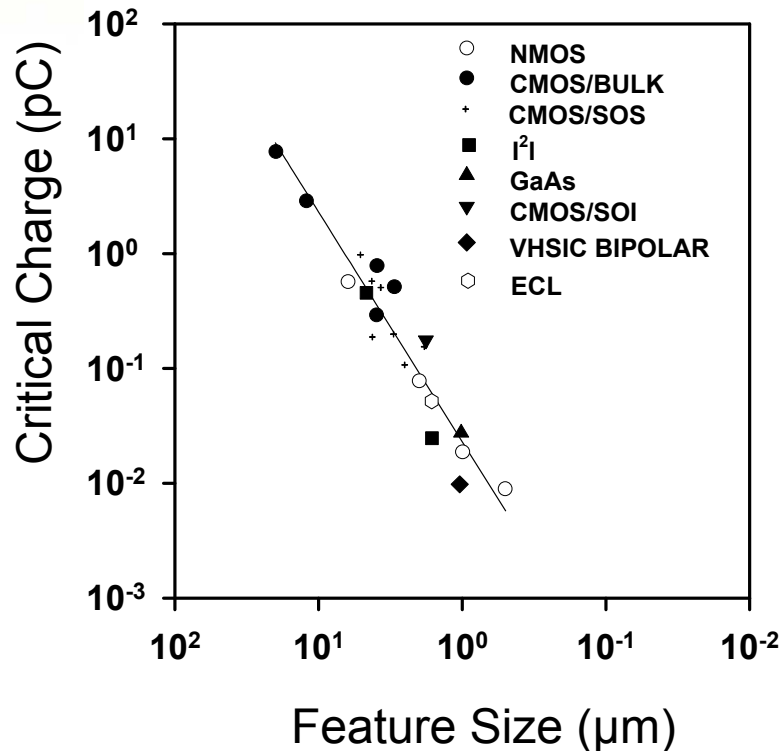
- Single-Event Upset (SEU)
- Single-Event Functional Interrupt (SEFI)

Hard Errors (permanent damage to device/circuit)

- Single-Event Latchup (SEL)
- Single-Event Burnout (SEB)
- Single-Event Gate Rupture (SEGR)

This is not a complete listing of all possible single-event effects!

Technology Trends Lead to Increased SEE Sensitivity



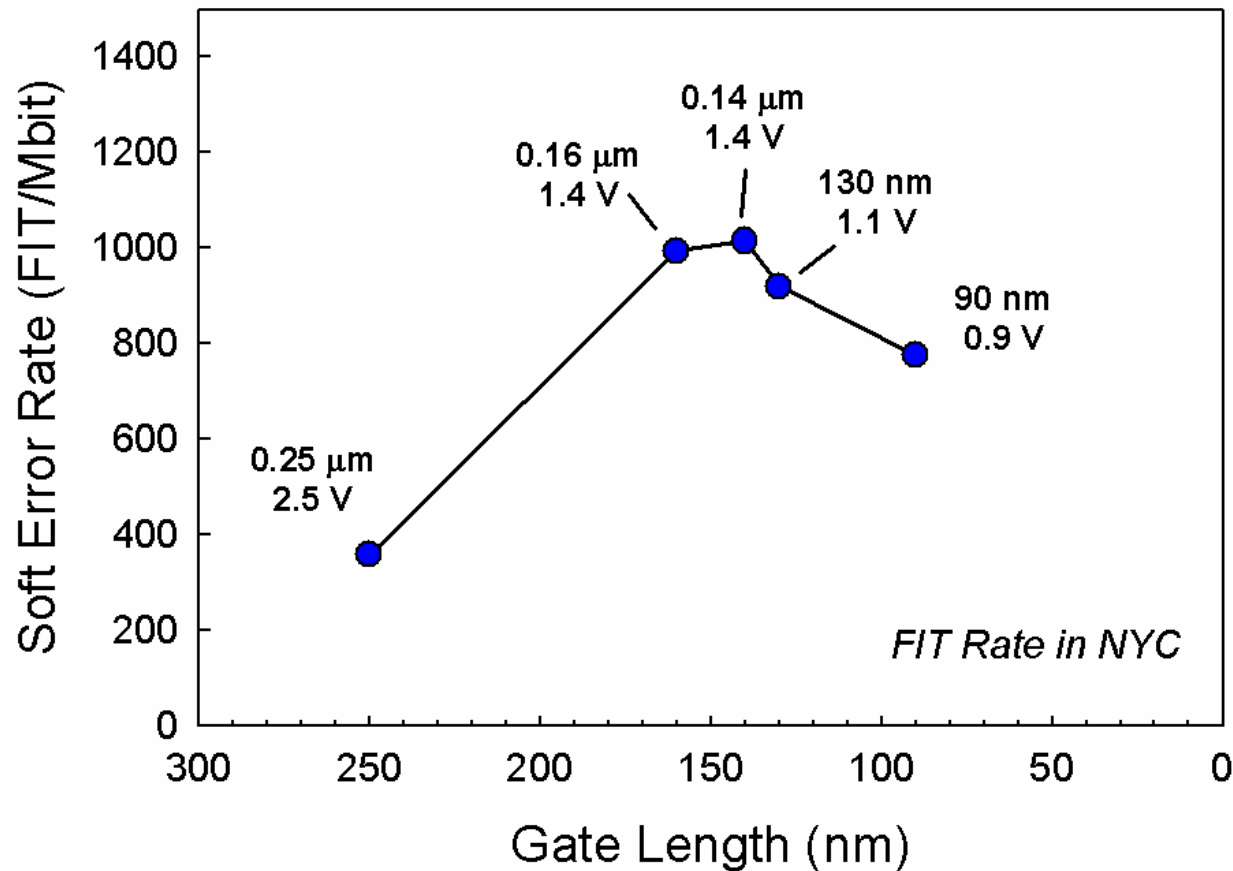
- decreasing feature size
 - decreased nodal capacitance
- lower voltages
- higher speed
 - more transients propagated as errors

The trend line is remarkably consistent across a wide variety of technologies. Does this 20-year old curve still apply?

After Petersen, et al., J. Rad. Eff. R&E 6, p. 6, 1988.

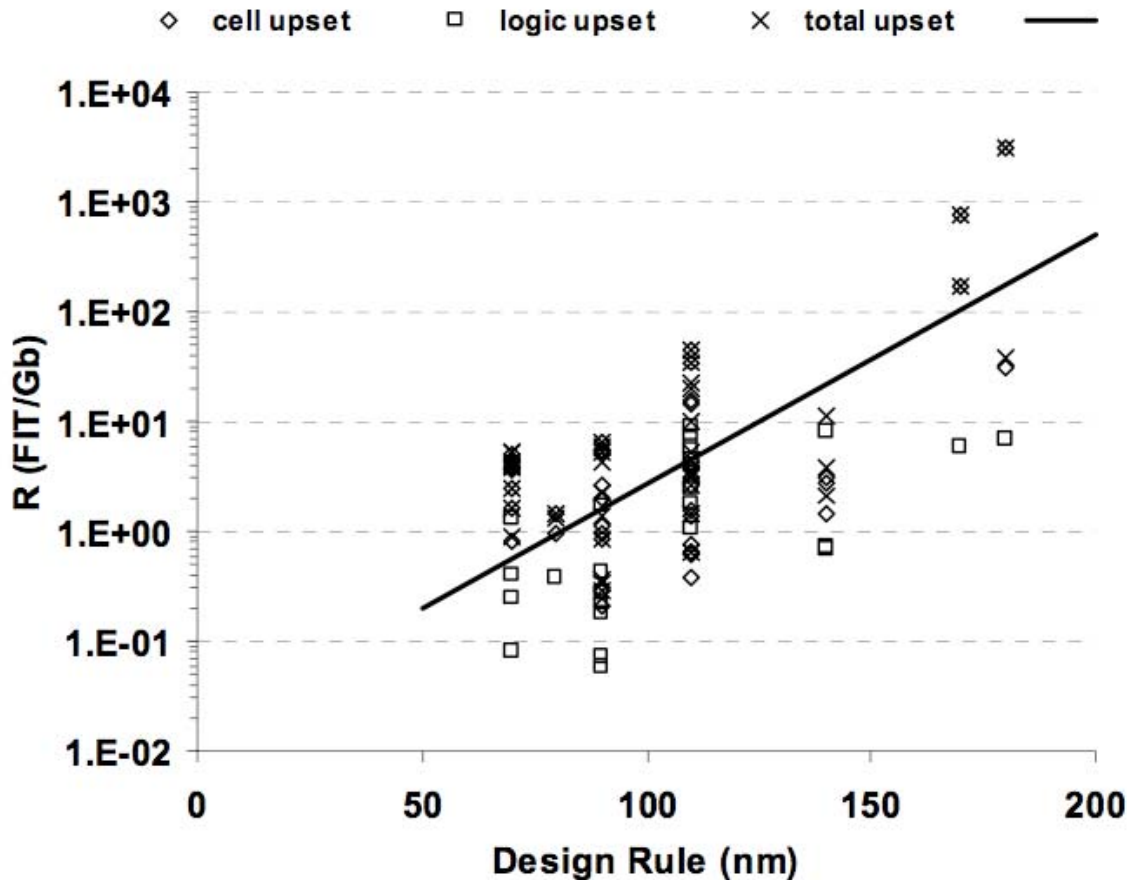
SRAM Terrestrial Soft Error Rate Scaling Trend

6T SRAM Bit nSER Trend
(terrestrial cosmic ray neutrons only)



After Dodd, et al., Proc. IRPS, 2003 and unpublished 90-nm data.

DRAM Terrestrial Soft Error Rate Scaling Trend



DRAM SER at 90 nm:
2 FIT/Gbit

(~400,000 times less
than SRAM SER)

After Borucki, et al., Proc. IRPS, pp. 482-487, 2008.

Memories are the Most SEU-Sensitive Element of Electronic Systems

- **SRAM**
 - By far the most sensitive type of memory, especially in highly-scaled CMOS technologies
- **DRAM**
 - Also sensitive, but far less so than SRAM because stored charge has not changed much with technology scaling
- **FPGAs**
 - SRAM-based reconfigurable FPGAs inherit the SEU sensitivity of the SRAMs on which they are based
- **Microprocessors**
 - Logic cells are sensitive to upset and single-event transients and are becoming a growing concern
 - Large embedded (or even off-chip) SRAM caches can be the biggest contributor to error rates in high-performance computers

Terrestrial SEU Rate Case Study: SGI Altix 4700 Supercomputer

Processing blades	Intel Itanium 2 dual core processors		
	L2 cache SRAM	2 Mbyte	per processor
	L3 cache SRAM	24 Mbyte	per processor
	Blade DRAM	16 Gbyte	per processor
	processor density	2 processors	per blade
	Total SRAM per blade	416 Mbit	
	Total DRAM per blade	256 Gbit	
RASC blades	Reconfigurable processing nodes with Xilinx V4s		
	Xilinx V4 LX200 CLB	1.36 Mbit	per V4
	Xilinx V4 LX200 BRAM	5.91 Mbit	per V4
	FPGA density	2 FPGAs	per blade
	RASC QDR SRAM	80 Mbyte	per blade
	Total CLB RAM per blade	2.72 Mbit	
	Total BRAM per blade	11.82 Mbit	
	Total SRAM per blade	640 Mbit	

Note: This configuration is based on a system SGI recently used to break a bioinformatics computing benchmark record.

Terrestrial SEU Rate Case Study: SGI Altix 4700 Supercomputer

Blade configuration	# processing blades	32 blades	total
	# RASC blades	35 blades	total
	Total system SRAM	35712 Mbit	
	Total system DRAM	8192 Gbit	
	Total system FPGA CLB	9.52 Mbit	
	Total system FPGA BRAM	413.7 Mbit	
	SRAM SER	27748224 FIT	
	DRAM SER	16384 FIT	
	FPGA CLB SER	2342 FIT	
	FPGA BRAM SER	145622 FIT	
	Total SER	27912572	
	MTTF	35.8 hours	
		1.5 days	

Processor cache SRAM and high-speed SRAM on the FPGA boards completely dominate the predicted terrestrial soft error rate.

SEU Mitigation Techniques: Device Level

- Reduce cell recovery time
 - oversized transistors with increased current drive
 - × area penalty
- Increase cell feedback time
 - feedback resistors or transistors
 - additional node capacitance
 - × speed & area penalties, temperature dependence, process complexity
- Reduce charge collection
 - thin epi or neutron-irradiated substrates
 - SOI
 - multiple doping wells
 - scaled sensitive volumes in advanced technologies
 - × cost and process/design complexity

Classical Radiation-Hardened SRAMs Use Cross-Coupled Feedback Resistors

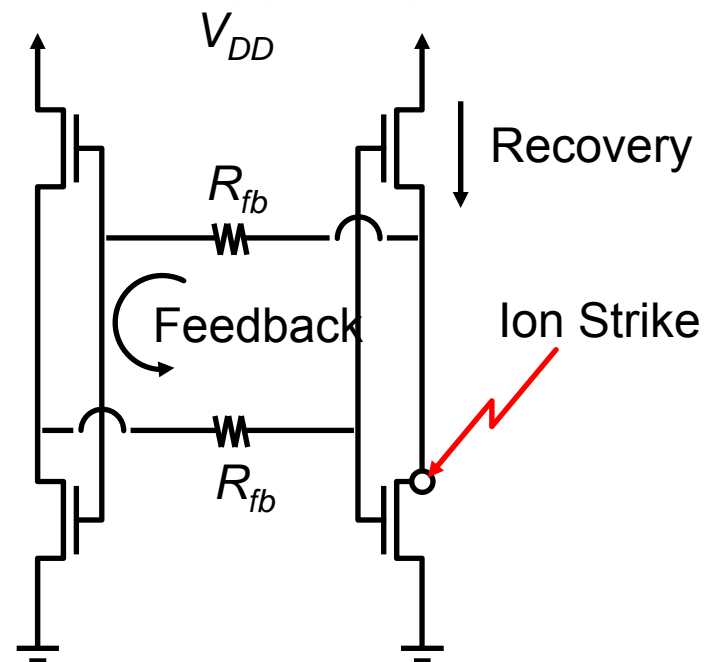
Slows feedback down so that recovery “wins the race”

But:

increases write time



increases process complexity

results in temperature dependence
of write performance and SEU
immunity



*Other feedback elements (capacitors, diodes, or transistors) can also be used.
Resistive hardening is not scalable to high-performance technologies.
Or is it?*

Silicon-on-Insulator Has Greatly Reduced Charge Collection Volume



In bulk technologies, charge deposited several μm deep can be collected.

In thin-film SOI, the top silicon layer is typically less than 300 nm thick.

In practice, floating body effects can limit the SEU hardness of partially-depleted SOI technologies. Consequently, radiation-hardened SOI technologies use body ties to reduce floating body effects.


SEU Mitigation Techniques: Circuit Level

- On-chip error detection and correction (EDAC)
 - error-correcting codes of varying complexity/overhead
 - ✗ area and performance penalties, what about logic?
- Remove dynamic and quasi-static logic
 - ALL circuits should be replaced with fully-static logic
 - no circuits requiring refresh
 - ✗ area and performance penalties
- On-chip redundant circuit designs
 - triple modular redundancy (e.g., Xilinx TMRTTool)
 - spatially redundant latches (6T SRAM → 12-16T SRAM)
 - temporally redundant designs to prevent SEU and SET
 - ✗ area and performance penalties, increased design complexity


Recent Data and Simulations Indicate DICE Protection can Fail for Scaled Technologies

DICE Cell SEU Simulations: 10° Vertical Angle through sensitive transistor pair


Electrons and Potential: 25 fs



Electrons and Potential: 10 ps



Electrons and Potential: 100 ps



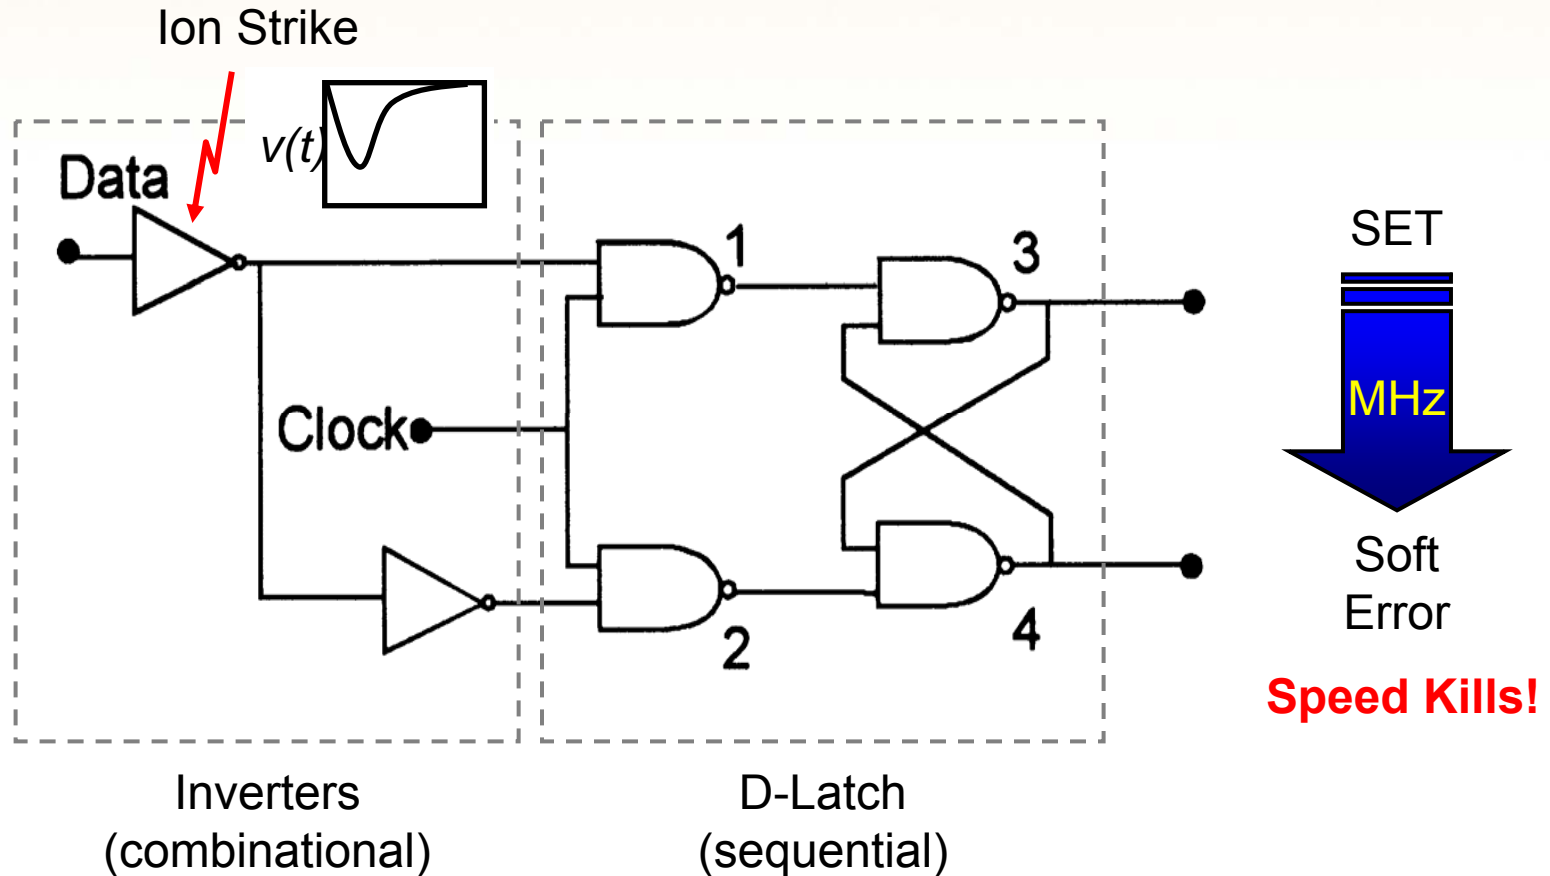
LET = 10 MeV-cm²/mg (Cell upsets)

Multinode charge collection limits the minimum spacing of sensitive node pairs in internally redundant memory and logic designs.

SEU Mitigation Techniques: System Level

- Shielding
 - add shielding or bury system deep in spacecraft
 - × not effective against high-energy cosmic rays
 - × weight penalty (weight = \$\$\$\$)
- System-level EDAC
 - use SEU-hardened chips to process EDAC
 - × volume/weight and performance penalties
- System-level redundancy
 - triple modular redundancy with voting subsystems
 - × volume/weight penalties, increased system complexity
- Fault-tolerant hardware and software architectures
 - application-oriented software fault tolerance
 - watchdog co-processing
 - × difficult to detect transient faults
 - × overhead and performance loss during return to known good state

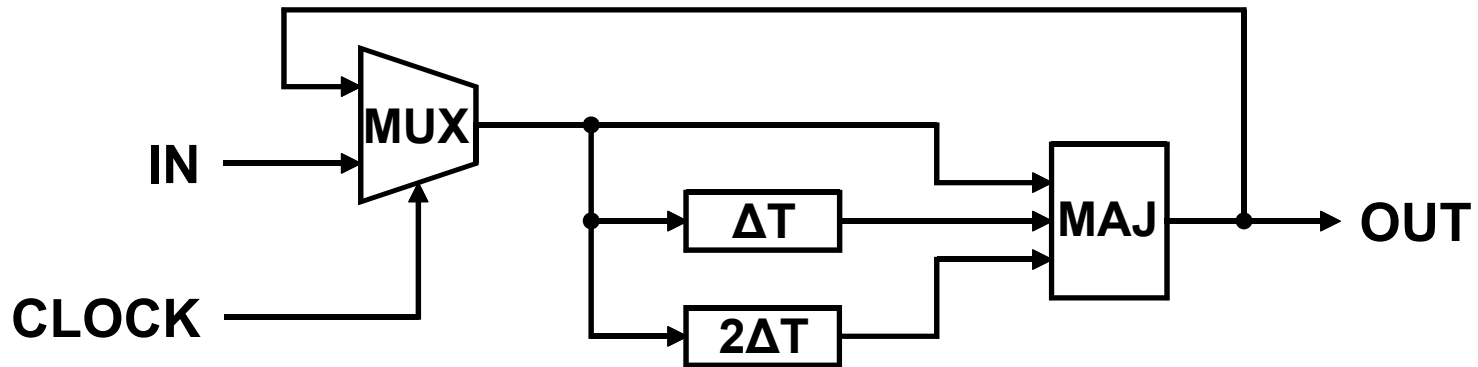
Single-Event Transients (SETs) Are An Increasing Problem in High-Speed Circuits



Transients generated in combinational logic can become errors if latched by subsequent circuitry. Clock transients are even worse!

After Buchner, et al., *IEEE Trans. Nucl. Sci.* 44, p. 2209, 1997.

Temporal Latch for Eliminating SETs



Latch filters SETs by voting temporally redundant outputs, but requires delay element (ΔT) larger than expected transient width.

Impact of Technology Trends on Mitigation Techniques

- Device-level techniques such as feedback elements do not scale well with technology, leading to stagnant performance
- SOI may buy an improvement in SER, but only temporarily
- Techniques that address SEU and not SET will become less effective as circuit speeds continue to increase
- Plus new failure modes such as functional interrupt (SEFI) in today's complex ICs
- Redundant circuit techniques were promising, but require area to implement and are becoming less effective with scaling
 - Some rad-hard vendors are returning to RC feedback delay hardening
- Will System-on-chip become 3 systems-on-chip?
- Fault-tolerant architectures will take on increasing importance

Opportunities Exist for “Beyond CMOS” Devices

- Continued CMOS scaling will only exacerbate SEU issues
 - Smaller devices → smaller critical charge
 - Higher speeds → more propagated transients
- New technologies may break the current trends
 - Carbon nanotubes, molecular electronics, phase change memories, ferroelectrics, etc.
 - Many of these technologies have shown promising initial radiation hardness
 - Much work remains to be done to bring these technologies to mass market and to qualify them for use in space radiation environments
- New fault-tolerant computing architectures will also be key
 - Can multi-core parallel processors with built-in fault tolerance provide some measure of redundancy while retaining computational efficiency?

Conclusions

- Radiation effects continue to be a significant concern for advanced spaceborne computing platforms
 - Technology trends favor improved TID hardness, but have led to increased SEU sensitivity with scaling
- A variety of techniques exist for mitigating soft errors in mission-critical systems
 - device, circuit, and system-level solutions
- Unfortunately, most (all?) of these techniques require tradeoffs that are difficult to accept
- Fault tolerant architectures and new device concepts offer promise for breaking the downsides of Moore's Law