

System-Level Radiation Hardening

Ray Ladbury

NASA Goddard Space Flight Center

Acronyms and Abbreviations



- ASIC—Application Specific Integrated Circuit
- C_f —Cost/consequence of failure
- CL—Confidence Level
- CMOS—Complementary Metal-Oxide-Semiconductor
- COTS—Commercial Off The Shelf
- DDD—Displacement Damage Dose
- DDR—Double-Data Rate
- DSEE—Destructive SEE
- EDAC—Error Detection and Correction
- ELDRS—Enhanced Low Dose Rate Sensitivity
- EOL—End of Life
- F—Fluence
- FPGA—Field Programmable Gate Array
- GCR—Galactic Cosmic Rays
- LET—Linear Energy Transfer
- LRO—Lunar Reconnaissance Orbiter
- MBU—Multi-Bit Upset
- MCU—Multi-Cell Upset
- MOSFET—Metal-Oxide-Semiconductor Field Effect Transistor
- N—Number of events
- P_f —Probability of failure
- P_s —Probability of success
- RHA—Radiation Hardness Assurance
- RLAT—Radiation Lot Acceptance Test
- SDRAM—Synchronous Dynamic Random Access Memory
- SEB—Single-Event Burnout
- SEE—Single-Event Effect
- SEFI—Single-Event Functional Interrupt
- SEGR—Single-Event Gate Rupture
- SEL—Single-Event Latchup
- SET—Single-Event Transient
- SWAP—Size, Weight and Power
- TID—Total Ionizing Dose
- VDS—Drain-to-Source Voltage
- VGS—Gate-to-Source Voltage
- WC—Worst case
- w/o—Without
- σ —Cross section

System-Level Radiation Hardening: Why not and Why?



- System-level hardening can be expensive on many different levels
 - Cost—System designed to compensate for component weaknesses
 - Schedule—Design and verification of hardening are time consuming
 - Performance—Most radiation mitigations require compromising performance
 - Redundancy means greater Size, Weight And Power (SWAP)—Important for space systems
- So why do it?
 - Commercial components may offer greatly improved performance even after hardening—not obtainable by other means
 - Speed—Processors, Field-Programmable Gate Arrays (FPGAs), data converters, memories, etc.
 - Memory density—Double-Data Rate (DDR) Synchronous Dynamic Random Access Memories (SDRAM) and FLASH >1000x denser than radiation hardened counterparts
 - Precision—Commercial data converters ~1000x more precise than rad hard devices
 - Smaller, more efficient chips mean hardened systems can still win the SWAP tradeoff
 - Sometimes system-level hardening implemented en lieu of complete characterization



Hardening is Expensive: Must Be Driven by Requirements

- Good Requirements must be:
 1. Clear to all affected parties
 2. Relevant to mission objectives (not “desirements”)
 3. Verifiable by test or analysis (preferably before spacecraft launch)

General; e.g. specifies mission duration and environment

Top Level Requirement
The System shall operate for five (5) years in a geostationary environment at 175° W Longitude.

More specific; relevant to design and parts engineers

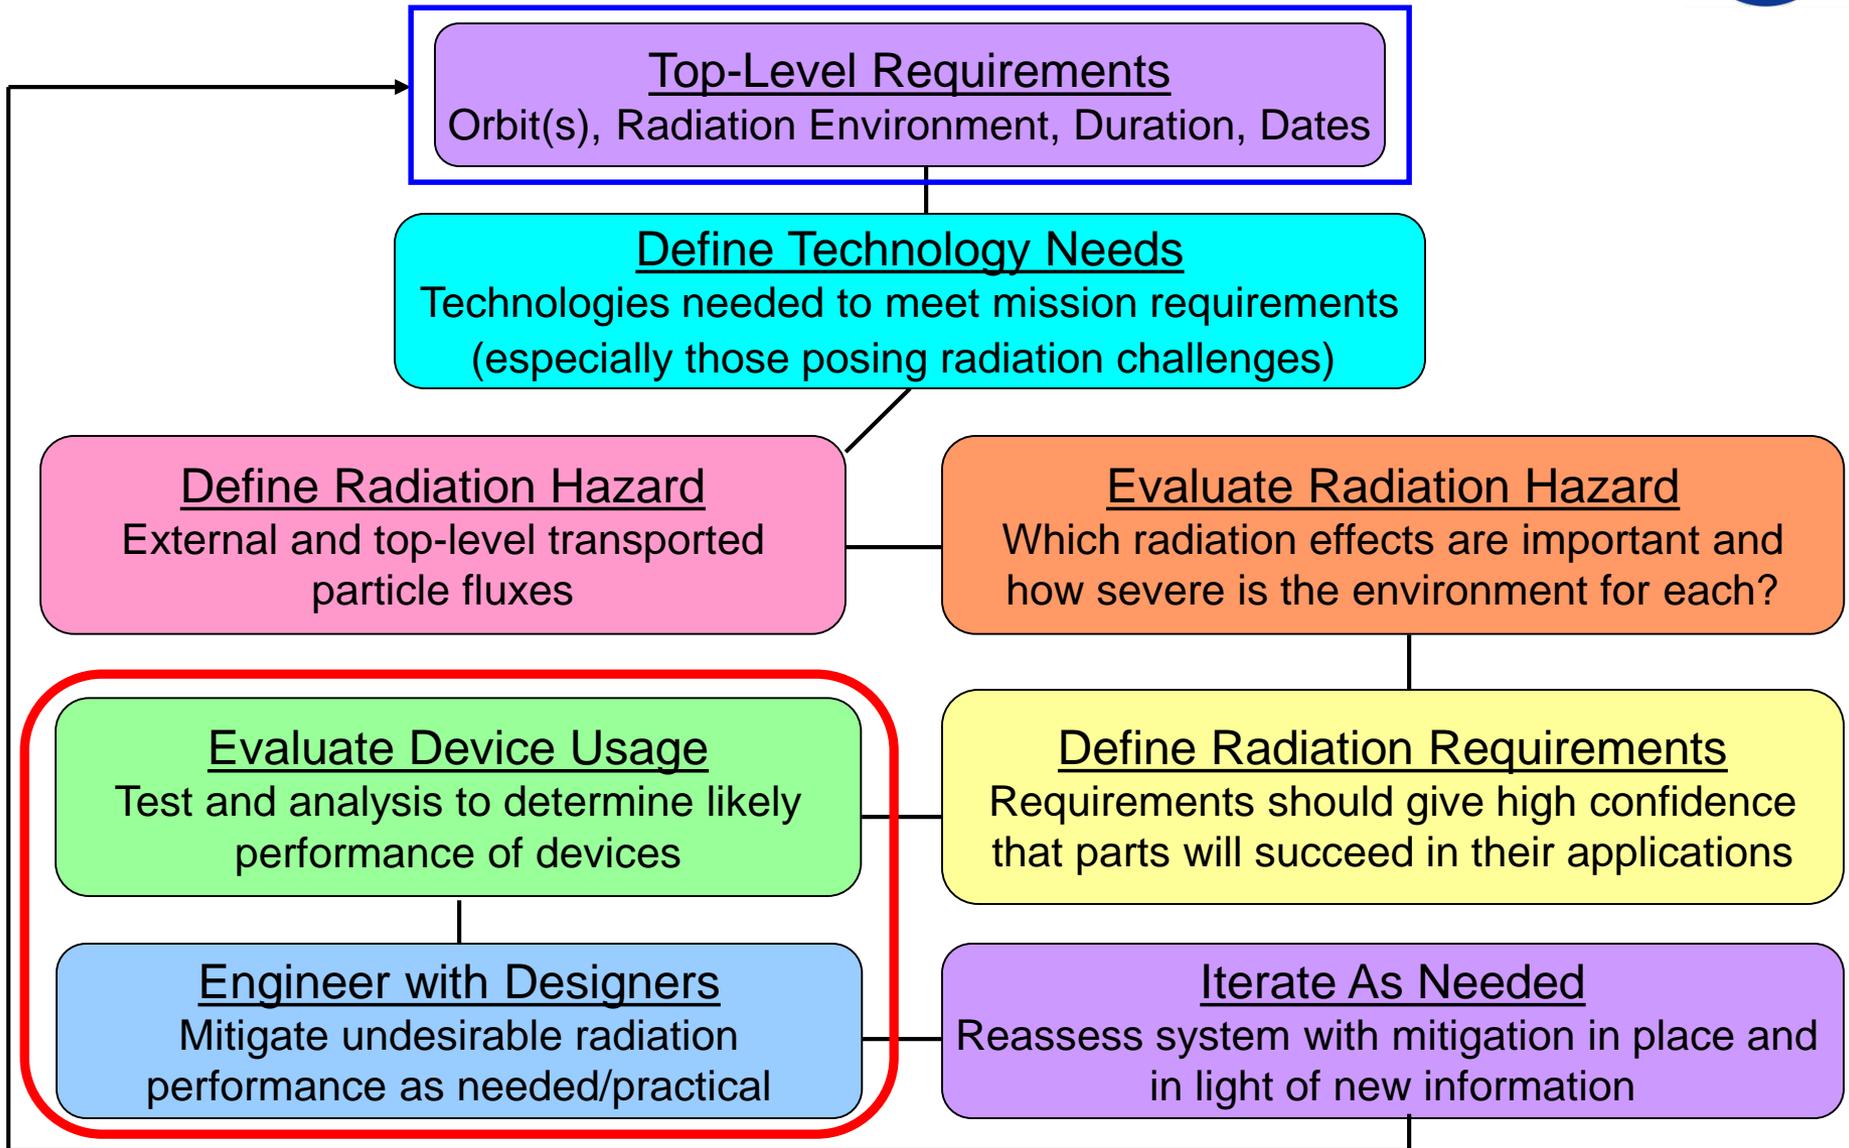
Second Level Requirement
Parts used in the system shall be immune to destructive single-event effects (DSEE).

Specific and verifiable; gives guidelines for verification

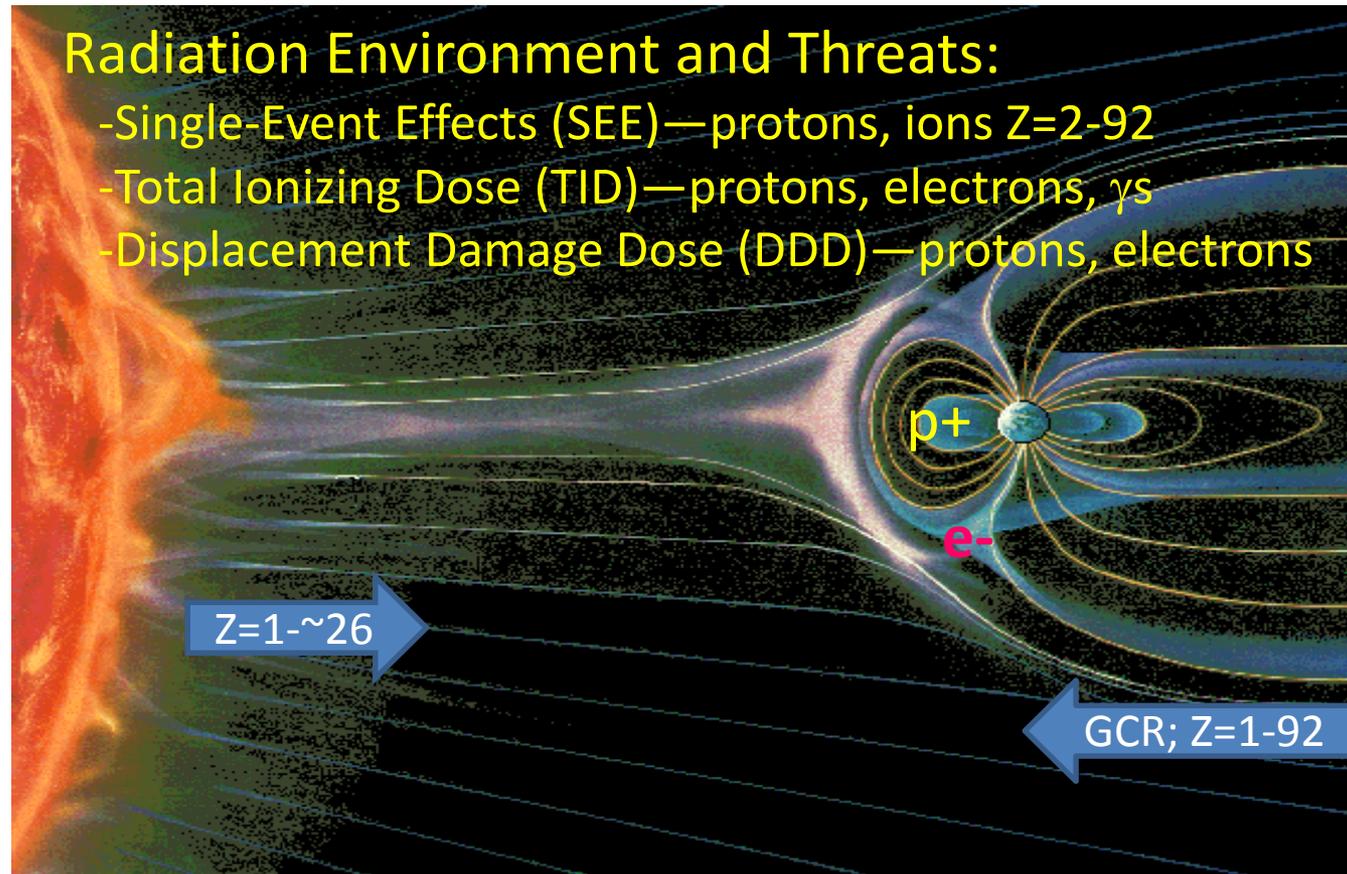
Third Level Requirement
Parts shall be defined immune to DSEE if they survive 10^7 ions/cm² with LET* > 60 MeV·cm²/mg

*LET=Linear Energy Transfer

Radiation Hardness Assurance (RHA) Process at NASA



Radiation Threat Environments



Adapted from K. Endo,
Nikkei Science, Japan

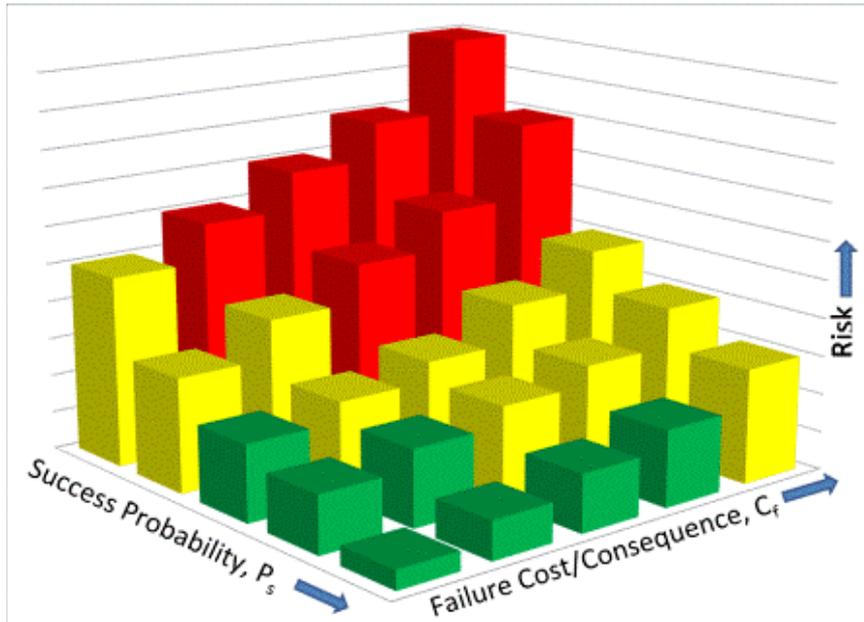
- SEE are prompt effects due to passage of ionizing particles—Poisson processes
 - May be nondestructive (e.g., single-event upset—SEU) or destructive (e.g. single-event latchup—SEL)
- TID and displacement damage are cumulative processes

Three Types of Radiation Effects



- **Destructive SEE—Poisson process, constant rate, affects single die; redundancy effective as mitigation, but very costly**
 - SEL—Single-Event Latchup (Complementary Metal Oxide Semiconductor-CMOS)
 - SEGR—Single-Event Gate Rupture (MOS Field Effect Transistors-MOSFETs)
 - SEB—Single-Event Burnout in discrete transistors
 - Others—Stuck Bits, Snapback (Silicon on Insulator), Single-Event Dielectric Rupture
- **Nondestructive SEE—Poisson process, const. rate, single die, recoverable**
 - SEU—Single-Event Upset in digital device (or portion of device)
 - MBU/MCU—Multibit/Multi-Cell Upset in digital device (or portion)
 - SET—Single-Event Transient in digital or analog device
 - SEFI—Single-Event Functional Interrupt (full or partial loss of functionality)
- **Degradation Mechanisms—cumulative, end-of-life, affects most die as mission approaches mean failure dose. Redundancy ineffective**
 - TID—Total Ionizing Dose (degradation due to charge trapped in device oxides)
 - DDD—Displacement Damage Dose (degradation from damage to semiconductor)

Risk and Mitigation



- $Risk = (1 - P_s) \times C_f$
 - $P(\text{Fail}) = 1 - P_s$
 - C_f can be quantitative (e.g. \$,€) or qualitative; may include intangibles > cost of mission

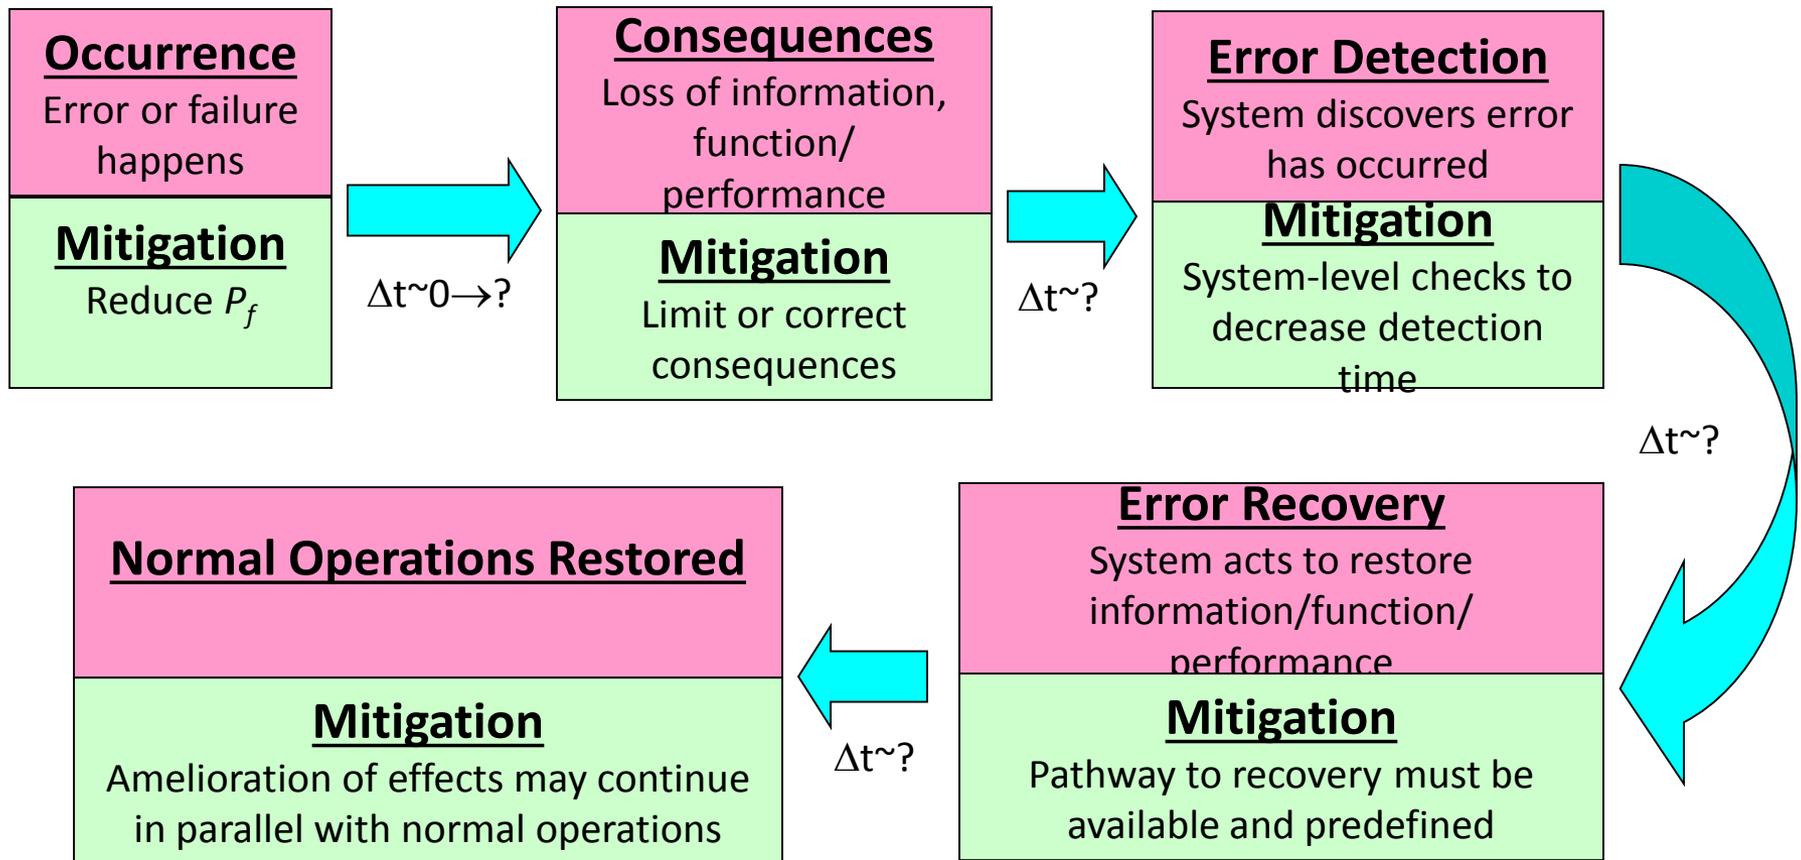
- Mitigation can reduce risk by
 - Reducing probability of error/failure
 - Replace with hardened part
 - Derate operating conditions
 - Limit time in vulnerable condition
 - Add shielding to limit TID to component
 - Reducing consequences of error/failure
 - Implement event detection/circumvention
 - Capacitive filtering of transients
 - Add cold spare to replace failed unit
 - Error Detection and Correction for SEUs
 - Or both
 - Triplicate voting and temporal voting mean no single error causes failure
 - System failure probability = $3P_f^2 + P_f^3$
 - Temporal voting similar for transients
 - Some operating conditions reduce both transient rate and duration



Mitigation Usually Focuses on Effects

- Radiation effects can result in a range of impacts
 - Permanent Loss of Functionality/Capability (TID/DDD failures, destructive SEE)
 - Metric affected is System Reliability=Probability of meeting requirements at end of life
 - Reliability P_s of System composed of components A, B, C...Z = $P_s(A) P_s(B) P_s(C)...P_s(Z)$
 - Survivability is a related concept = Ability to remain mission capable in the face of threats
 - Mitigation to improve reliability requires redundant subsystems or higher P_s values
 - Temporary Loss of Functionality/Capability (SEFI, spurious reset, etc.)
 - Metric affected is System Availability = (Up time)/(Up time + Down time)
 - If time to failure = T_F , Time to Recover= T_R , Availability = $T_F/(T_F+T_R)$
 - Mitigating availability: Increase T_F or decrease T_R
 - Multiple systems functioning simultaneously voted, averaged, etc.
 - Data Accuracy (bit errors from SEU, MBU, stuck bits, SEFI)
 - Example of a metric: Bit error rate = #bits in error/total bits processed
 - Mitigation—redundant bits for EDAC/voting
 - Performance degradation—(from TID/DDD)
 - Examples: system speed, decreased resolution
 - Mitigation: redundancy ineffective; sometimes application conditions can affect degradation; usual approach is to keep dose well below failure dose. Also programmatic.

Stages of a Failure (and their mitigation)



In these 5 steps, only one of the mitigations involves reducing failure probability. The rest involve limiting consequences or hastening restoration of normal operations.



Types of Mitigations

Threat Avoidance/Reduction

Reduce stress to reduce vulnerability

Effective for: destructive SEE, degradation

Example: shielding for TID/DD; safe voltages for SEGR

Rely On Infrastructure

Rely on system to correct errors

Effective for: all

Example: watchdog; event detection

Performance Matching or Derating

Avoid over-performing conditions

Effective for nondestructive SEE

Example: filter SETs, limit speed

General Mitigation Strategies

Opportunistic Strategies

Use failure characteristics to limit risk

Effective for: depends on part characteristics

Example: bit interleaving; select best op. cond.

Redundancy

Additional hardware/info to limit failures

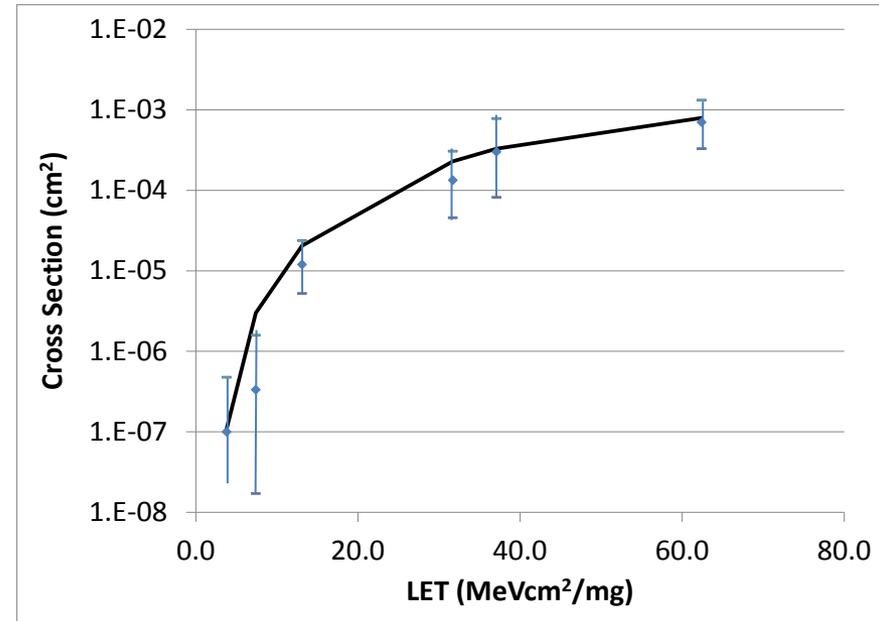
Effective for: destructive/nondest. SEE

Example: cold spares; EDAC, voting...



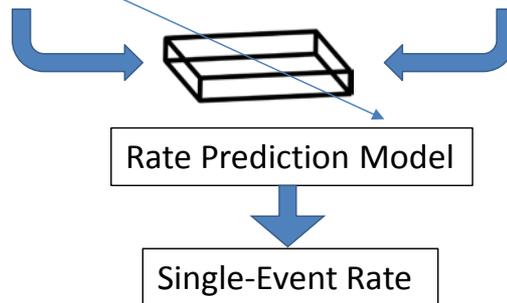
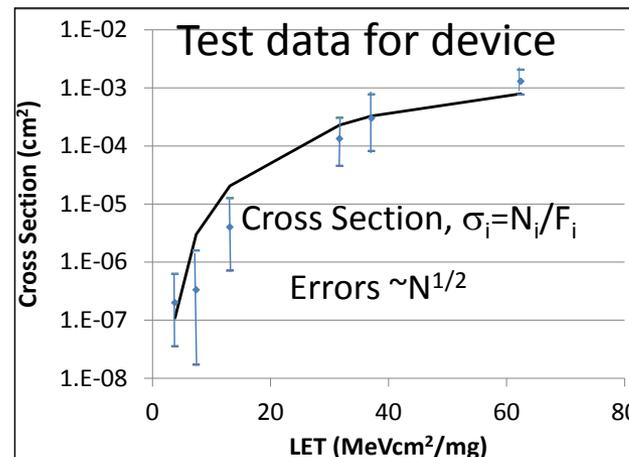
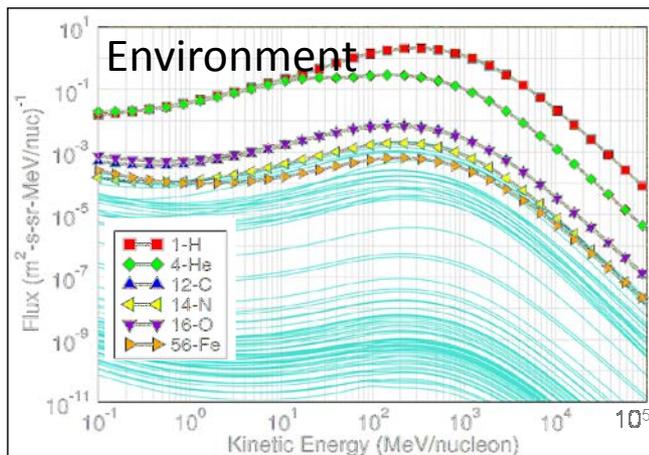
SEE Hardness Assurance

- SEE testing definitions:
 - LET_{∞} ion track charge density
 - Events observed = N_i @ LET_i
 - Fluence @ LET_i , $F_i = \# \text{ ions/cm}^2$
 - Cross section @ LET_i , $\sigma = N_i / F_i$
- Estimate rates with σ vs. LET curve and radiation environment model
 - Poisson errors on $N_i \sim N_i^{1/2}$
- Mitigate if risk too high
 - Nondestructive SEE—use redundancy
 - Destructive SEE—threat avoidance (preferred); cold sparing or detection/circumvention circuit in some cases
- Significance of rare events
 - Rare events means few observed, so errors on cross sections are large
 - May be expensive and time consuming to accumulate statistics



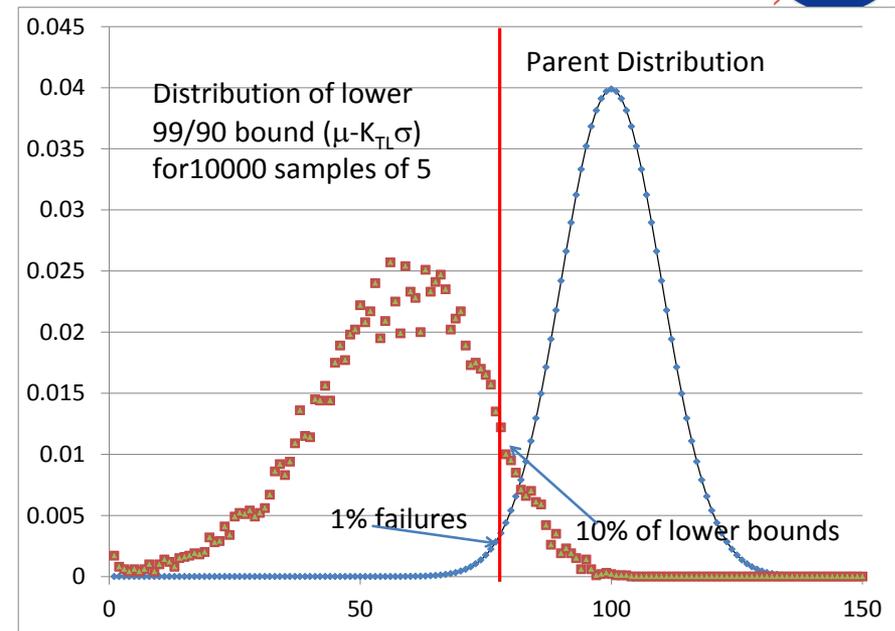
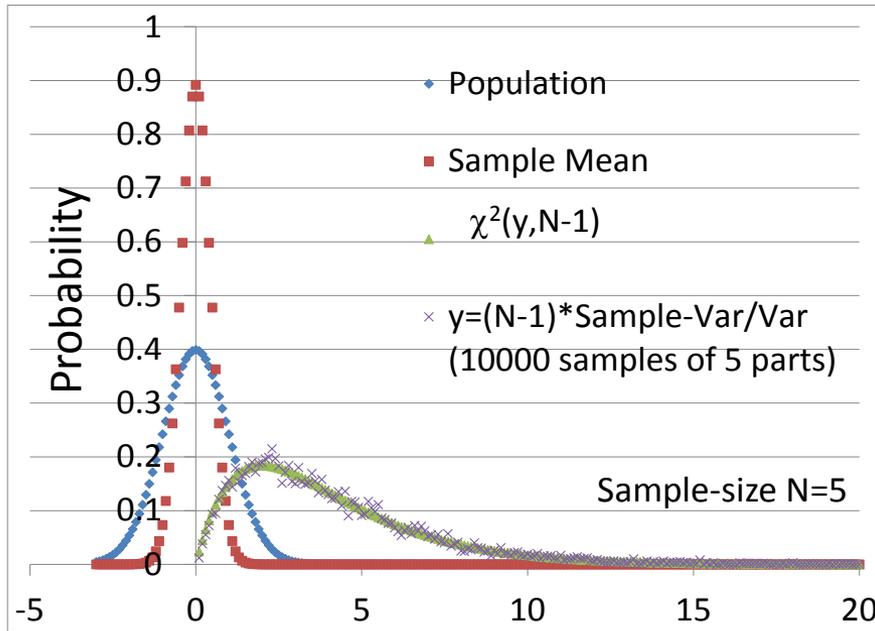
LET	Fluence, F	Evt. Obs., N	Cross section $\sigma=N/F$
3.8	1E+07	5	5.E-07
7.4	3E+06	1	3.E-07
13.1	5.E+05	5	1.00E-05
31.7	3.E+04	4	1.3E-04
37.1	1.E+04	2	2.0E-04
62.4	1.E+04	13	1.3E-03

SEE Hardness Assurance



- SEE testing exposes device to ions with given Linear Energy Transfer, LET_i
 - N_i events observed during fluence $F_i \text{ cm}^{-2} \rightarrow \sigma_i = N_i / F_i$
- If rate is too high, mitigation of SEE consequences is required
- Rare Events can dominate risk—poor statistics but severe consequences

TID Hardness Assurance

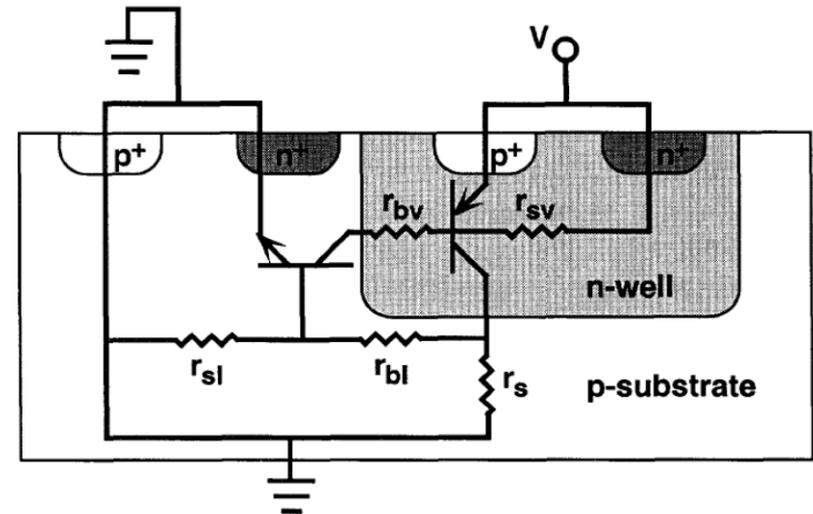


- Conventional TID RHA assumes “lot” is well behaved and unimodal
 - Wafer diffusion lot usually (not always) shows homogeneous performance
 - Assuming Normal distribution sample mean and variance converge to population values rapidly
 - Can use one-sided tolerance limits— $KTL(P_s, CL, n)$ —to bound performance and ensure success probability P_s with confidence CL for a test sample of size n
- Rare events for TID usually mean poorly behaved or “maverick” parts
 - Deviate from “well-behaved” assumption and from behavior typical for part
 - If mavericks are a small subsample of the lot, difficult to detect with a typical small sample test

Destructive SEE: Single-Event Latchup (SEL)



- Mitigation: threat avoidance/reduction
 - Reduce probability of SEL
 - Select part that is not susceptible
 - Use part at lower voltage if possible and take the performance hit
 - Reduce consequences of SEL
 - Cold sparing → complicated, heavy system
 - SEL detection and circumvention may not be feasible and results in spurious resets
 - Current limiting limits performance and may not be feasible
- Other complicating factors
 - Nondestructive SEL gets treated like a SEFI requiring a power cycle for recovery
 - Latent damage compromises part's reliability and possibly its SEL response
 - Mitigation needs to demonstrate it works against latent damage as well as SEL
 - SEL data sets often include several parts
 - Have to separate variability & Poisson errors

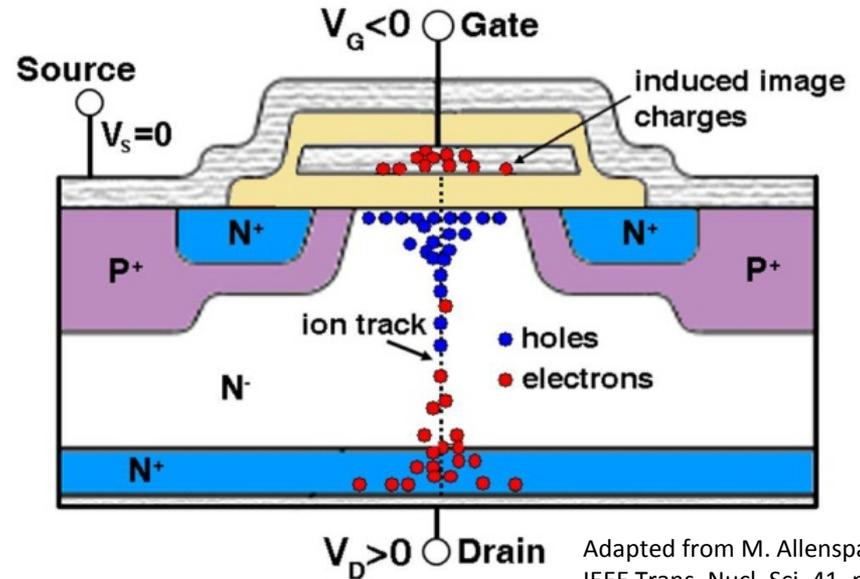


- SEL is a parasitic bipolar effect
 - Mechanism understood, but sensitive volume complicated
 - Depends on ion LET AND energy
 - Likely more variable than nondestructive SEE
 - Statistics poor due to disruptive nature of effect in testing

Destructive SEE: Single-Event Gate Rupture (SEGR)



- Mitigation: threat avoidance
 - Testing finds “safe” VDS and VGS for each test ion; often derated further
 - Mitigating consequences ineffective
 - If SEGR, failure cannot be circumvented
 - SEGR “precursor” events leave part functional but unreliable
- Rate estimation is complicated
 - Mechanism only partly understood
 - Poor statistics
 - No standard rate estimation method
- Risk avoidance gets more difficult
 - New vendors and commercial hardware being used in space
 - Lack of risk metric complicates comparing different parts/designs

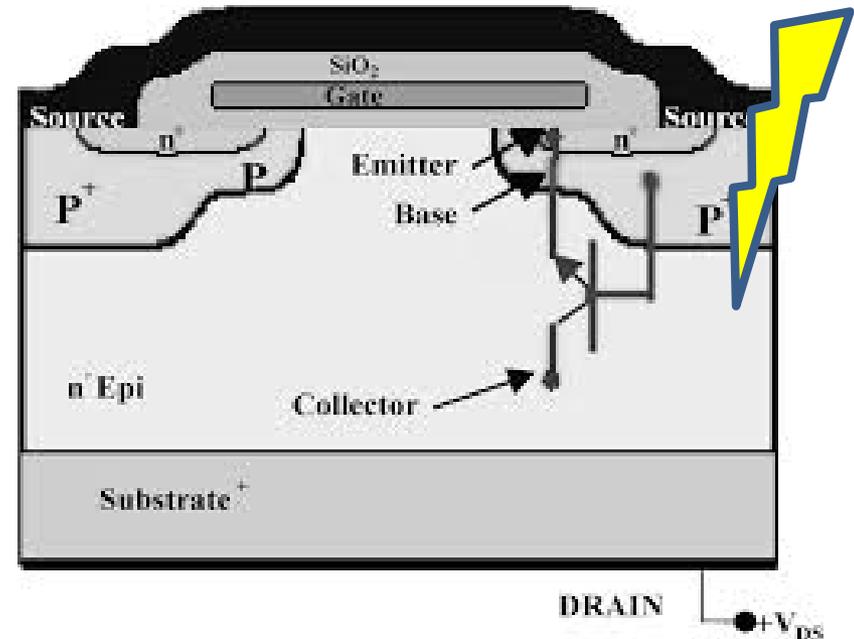


Adapted from M. Allenspach, IEEE Trans. Nucl. Sci. 41, pp. 2160-2166.

- SEGR: ion weakens gate oxide, then charges pile up under gate
 - Always destructive—testing costly and produces poor statistics
 - Depends on ion Z, energy and angle, as well as $V_{DS}=V_D-V_S$, and $V_{GS}=V_G-V_S$.
 - Similar failures in flash memory

Destructive SEE: Single-Event Burnout (SEB)

- Mitigation: threat avoidance/reduction
 - Testing finds “safe” VDS and VGS for each test ion, then voltage further derated
 - SEE-rad hard MOSFETs control SEB by design, but commercial parts may be at risk
 - For discrete transistors, failure can be circumvented by current limiting
 - Circumvention not possible for flash memories or bipolar microcircuits
- Rate estimation is difficult
 - Mechanism complicated
 - Statistics slightly better than SEGR/SEL
- Risk avoidance gets more difficult
 - New vendors and commercial hardware being used in space
 - Lack of risk metric complicates comparing different parts/designs



- SEB is parasitic bipolar effect
- Susceptibility complex—depends on
 - Ion Z, Energy and angle, application voltages
- May not be destructive if current limited
- Similar effects in Schottky diodes and some bipolar microcircuits—always destructive



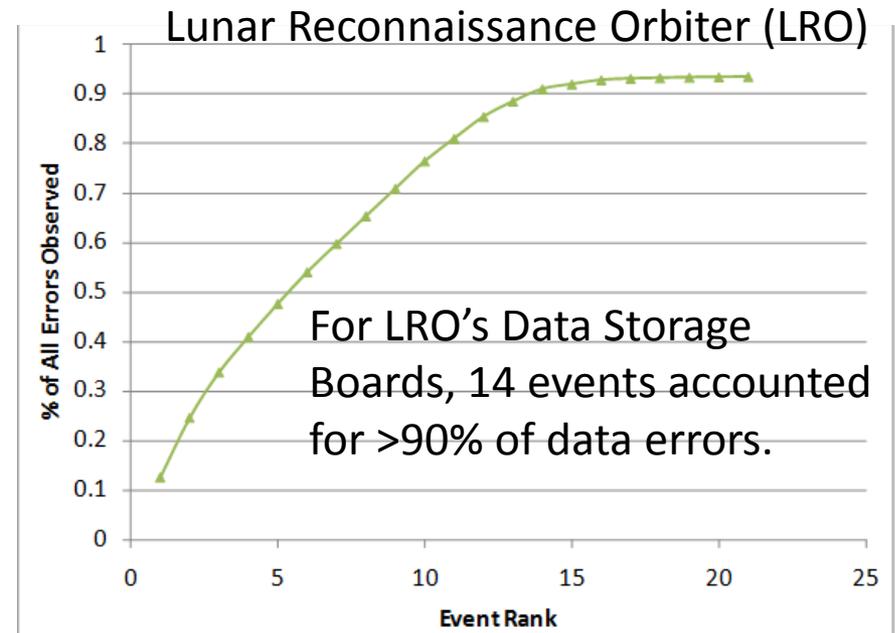
Stuck Bits

- Any digital device can conceivably be susceptible to stuck bits
- Mechanism: Charge trapped in device oxides causes threshold voltage shift and high leakage current, leading to inability to store one parity
- May exhibit high variability
 - Recent test of DDR3 SDRAM showed difference of >100x part to part in stuck bit rate; previous die revision did not show stuck-bit susceptibility at all
- Testing for stuck bits is complicated
 - Sometimes bits merely “leaky” and may not be stuck in all runs
 - May appear to be SEUs unless errors in each run compared with previous runs
 - Stuck bits may anneal over times from minutes to weeks; some do not anneal
- Mitigation options limited
 - Avoidance: Test candidate parts and select part with lowest stuck bit rate
 - EDAC/voting will correct stuck bit, but decreases protection for that word
 - Argues for a robust EDAC scheme to ensure reliability at end of life (EOL)
- Stuck bit susceptibility to date has been low enough not to drive design

Nondestructive SEE: Single-Event Functional Interrupt



- Possible interventions for functionality
 - Soft reset/re-initialization
 - Hard reset (may include clock)
 - Power cycle
 - Need to verify bus contention does not cause latent damage
- May also need to correct lost data
 - Data correction always uses redundancy
 - Requires EDAC or voting that is not overwhelmed by massive errors
- SEFI mitigation often complicated and can drive system design
- SEFI rates often based on poor statistics
 - Both SEFI and recovery are disruptive, limiting # observed during test
 - May have multiple SEFI modes
- SEFI susceptibility likely getting worse for ever more complex parts

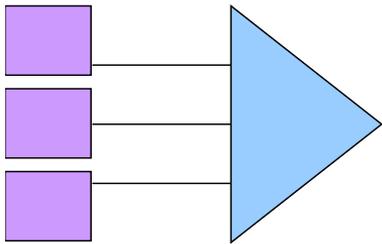


- SEFI: broad SEE category in complex devices
 - Processors, memories, data converters, etc.
 - Interrupts normal device functionality
 - Requires intervention to restore operations
 - May result in massive data corruption/loss
 - High-current bus contention may damage part
 - Effects limited to single die

Data Errors: SEUs, MCUs, MBUs and More

- Mitigation is based on redundancy
 - Requires probability of multiple correlated errors (MBU, SEFI...) to be small
 - Other measures can ensure noncorrelation

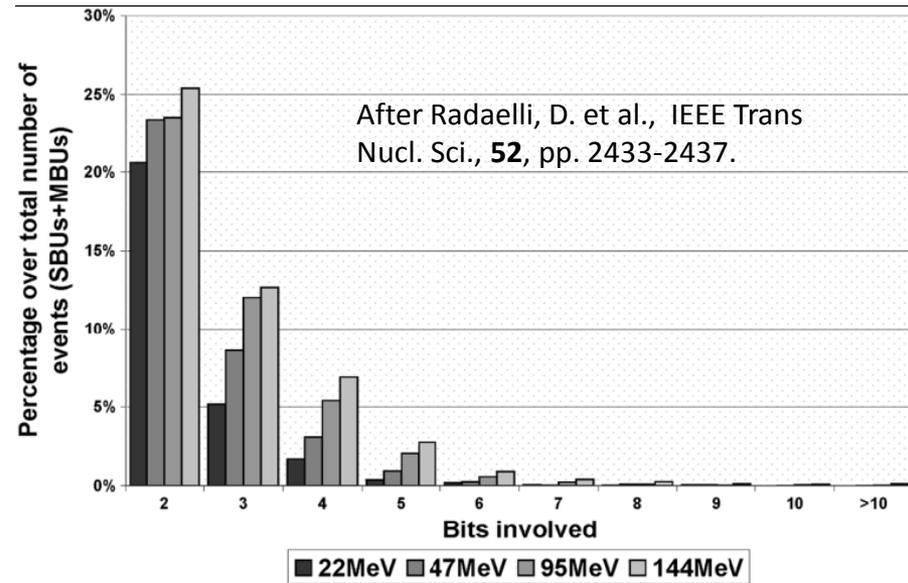
Triplicate Voting



- Reliable if common elements hardened
 - Costly in size, weight and power
 - Can also duplicate vote w/ retry

- EDAC uses redundant bits to correct data errors up to a maximum size
- To avoid overwhelming EDAC
 - Scrub data to avoid error accumulation
 - Interleave bits in data word across die

Data Bits					Correction Bits				
D0	D1	D2	...	DN	C0	C1	C2	...	CM

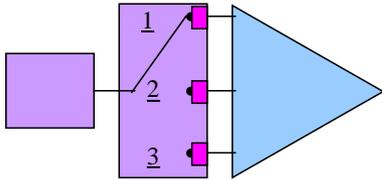


- Single-event upset (SEU) flips one bit
 - Sometimes called single bit upset (SBU)
- Multi-bit upset (MBU) flips multiple bits in the same data word
- Multi-cell upset (MCU) flips multiple bits, but in different data words
- Block errors corrupt large data blocks
- SEFI & SEL can corrupt all data on chip

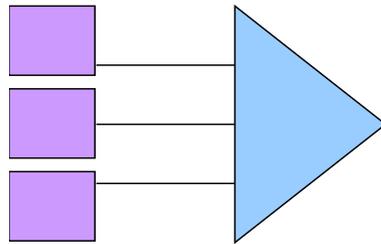
Nondestructive SEE: Single-Event Transients (SET)



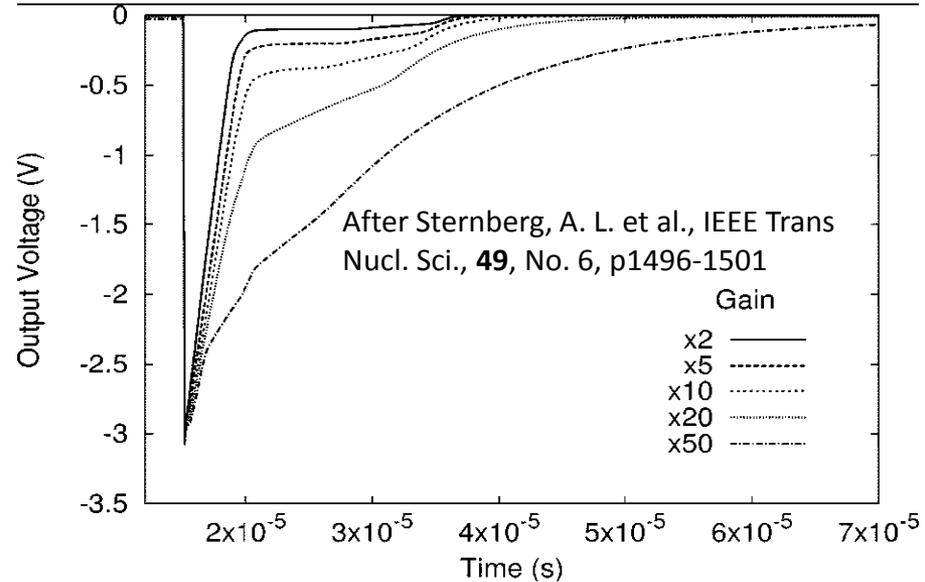
Temporal Voting



Spatial Voting



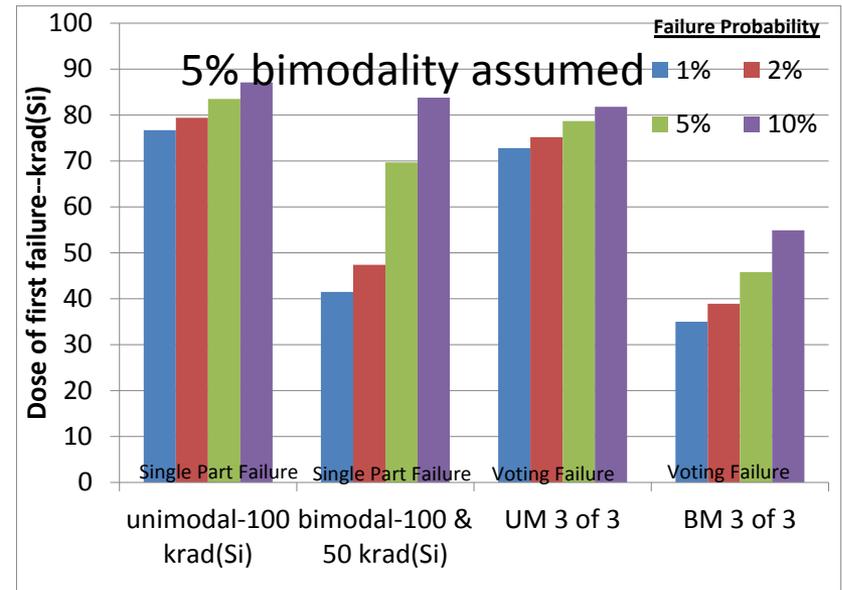
- SETs are limited temporally and spatially
- Temporal mitigation
 - Capacitive filtering
 - Multiple sampling with $\Delta t >$ transient width
 - Vote or average samples
 - Mitigation defeated if transients too long
 - Significant time penalties
- Spatial voting—transients in spatially separated instantiations unlikely
 - If nodes too close, multiple error can occur
- All mitigations penalize component speed
 - Some may have to live with transients



- SET response can be highly variable
 - Application affects duration and probability
 - Effect depends on circuitry downstream
- Analog SETs are characterized by amplitude and duration
- Digital SETs are characterized by duration only
 - Very important for ASICs, FPGAs, other devices with complex logic

Total Ionizing Dose Degradation

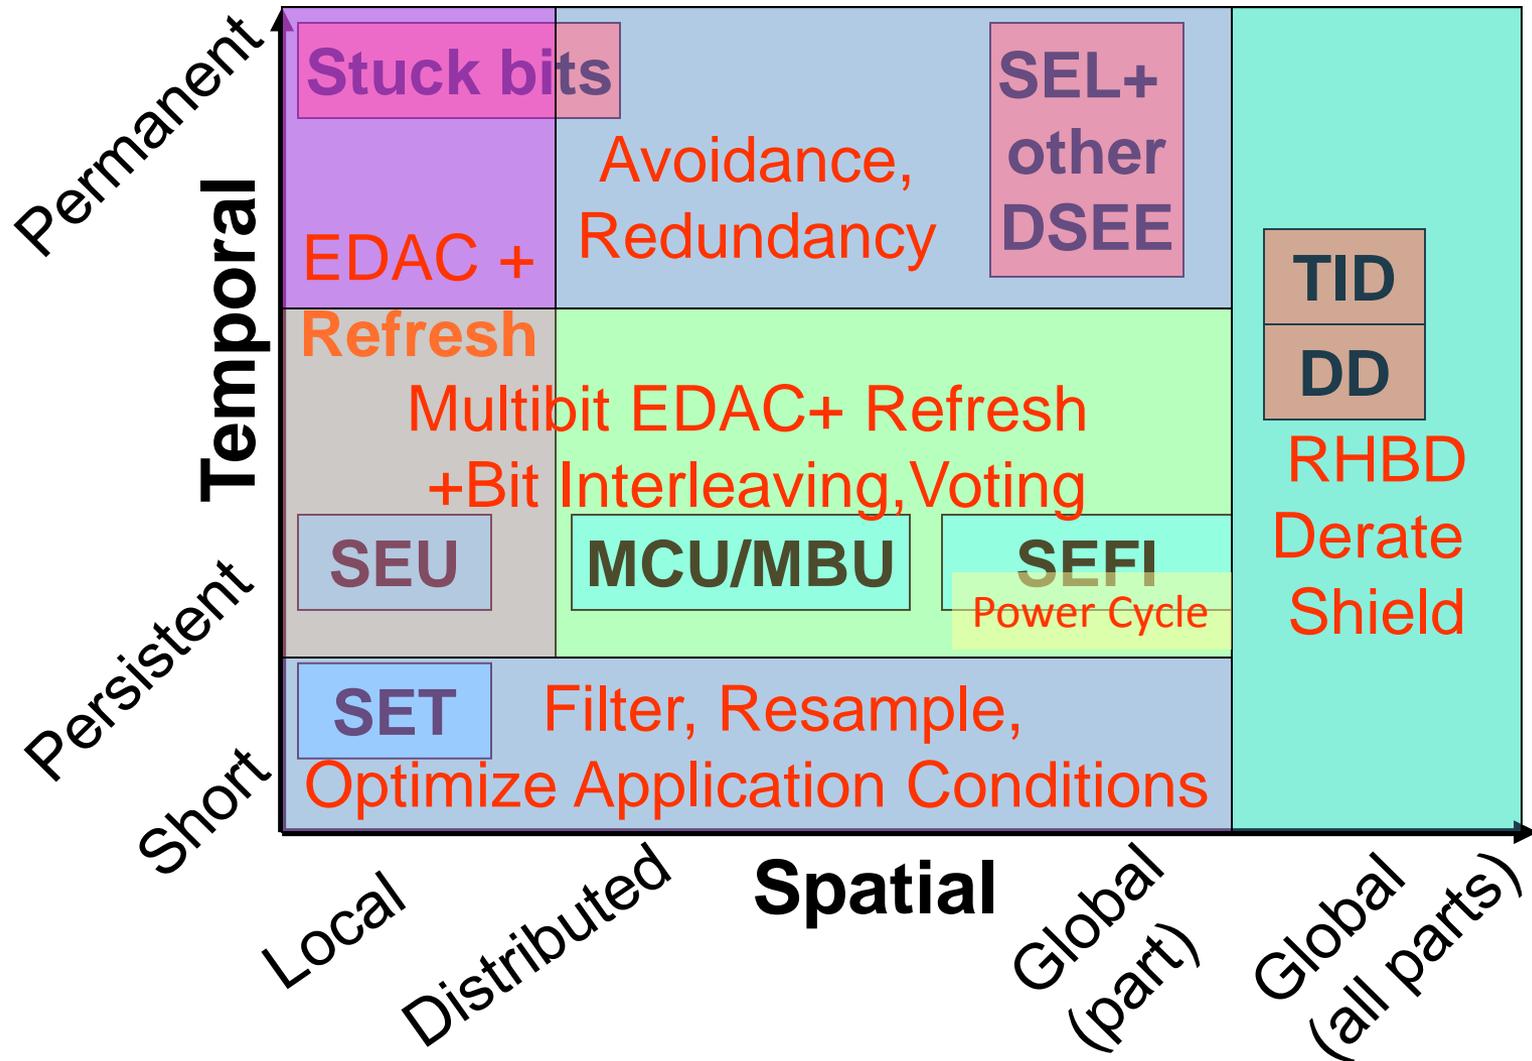
- Caused by charge trapped in device oxides
 - Degrades part functionally and parametrically followed by eventual failure
- Failures tend to cluster around the mean
 - Single wafer diffusion lot usually follows well-behaved distribution
- Tendency of failures to cluster means mitigation must be by threat avoidance
 - Increased shielding to reduce dose
 - Selecting TID hard parts
 - Increase Radiation Design Margin (RDM) to ensure parts far from failure dose
 - $RDM = \frac{\text{Failure Dose}}{\text{Mission Dose}}$
 - Varying application not effective
 - Failure Dose often taken as 99/90 level assuming Normal or lognormal
- Displacement damage is similar



- TID RHA usually assumes well-behaved lot
- If the lot not well-behaved RHA undermined
- Example above compares 2 lots
 - Unimodal Normal w/ mean failure @100 krad(Si)
 - Bimodal w/ 95% having mean 100, 5% mean 50
 - Bimodal lot can compromise RHA for a part, but it is very likely to compromise a triplicate voting scheme by causing one of three parts to fail



Combining Mitigations





Validation

- Once we implement mitigation, we have to show it works
- Validation may be by test or by analysis
 - Types of mitigation most likely to require validation by test are those most difficult to model
 - SEL detection and protection and/or cold sparing
 - SET mitigation
 - Compensation circuitry for TID
 - Often, mitigation is already based on the best test data available and additional testing would not improve confidence, so validation is done by modeling and analysis
 - Example I: Analysis of RLAT TID data for SDRAMs shows that dose levels will be less than half the 99/90 dose level.
 - Example II: Analysis shows scenario 2 mitigation reduces outages by >99.7%
 - Fault injection can be a very valuable tool
 - Useful for test planning as well as validation
 - May be essential for validating complicated systems with multiple mitigation layers
 - Often impossible to simulate all possible errors from all possible states
 - Some applications require nothing less than system test under irradiation

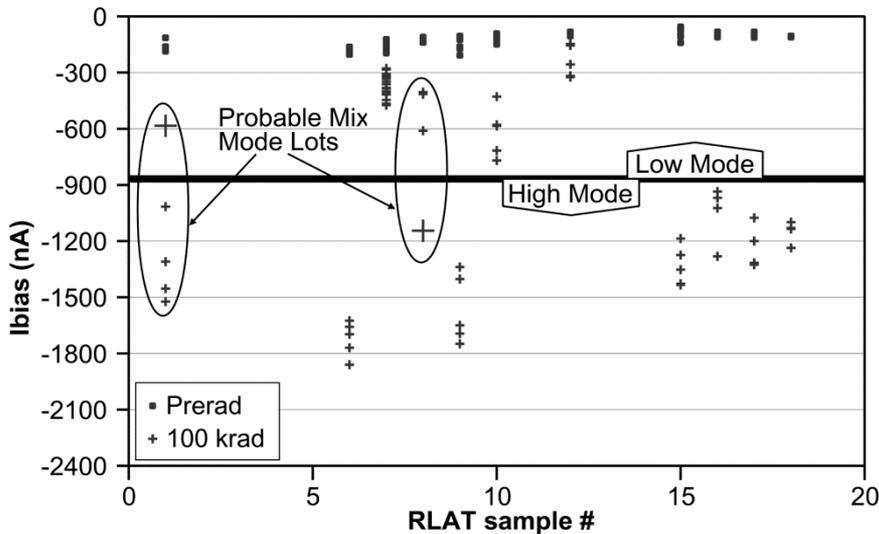
Successful Mitigation Predicated on Assumptions



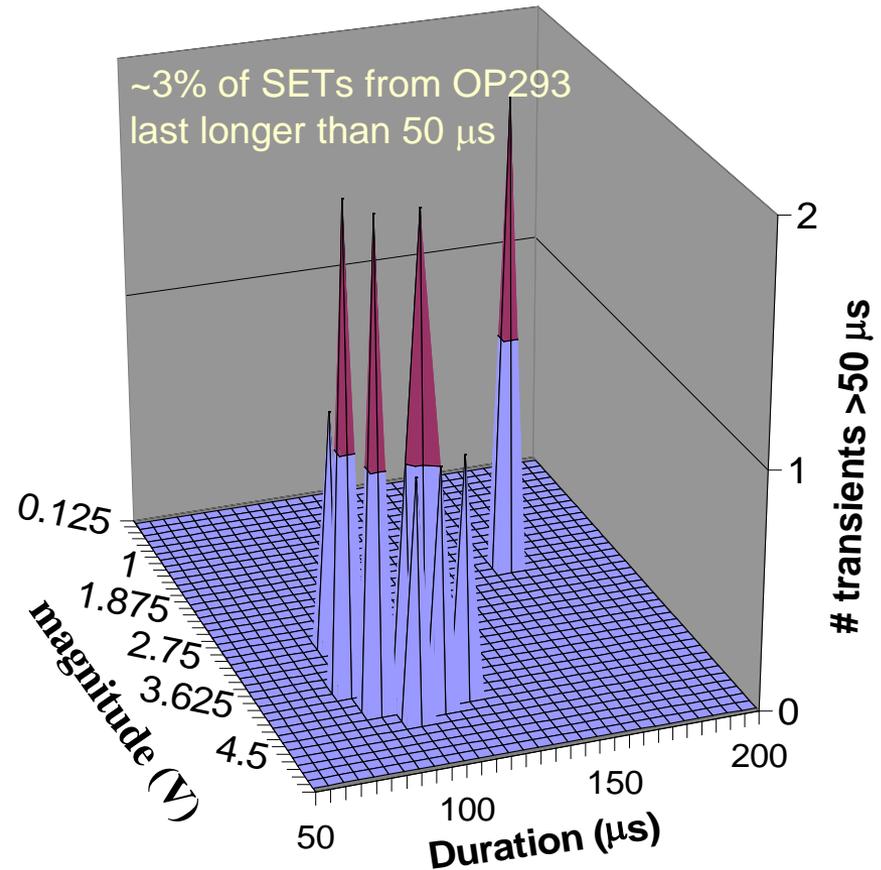
- RHA testing and analysis are expensive; some assumptions make it cheaper
 - 1) It is possible to define a “lot” from which test parts will have radiation behavior similar to flight parts
 - a) A lot for TID/DDD is a wafer diffusion lot; assumed to follow well behaved distribution
 - b) A lot for non-destructive SEE includes parts with same mask set and process
 - c) A lot for destructive SEE may be a wafer diffusion lot, but may be more general
 - 2) We can use our understanding of the radiation effect mechanism to limit testing
 - a) Example: only CMOS devices are susceptible to SEL
 - b) Only bipolar devices show significant Enhanced Low Dose Rate Sensitivity (ELDRS)
 - c) SETs are expected to be less than some maximum duration (e.g. 100 μ s)
 - d) Testing may deviate from 100% fidelity as long as we get bounding results
 - 3) Different radiation threats can be considered independently (no synergy)
 - 4) Testing detects all significant error/failure modes
 - a) Rare SEE modes with severe consequence can still dominate risk
 - b) Small TID samples may not detect bad parts if distributions are pathological
- When assumptions are violated, RHA and system level mitigation can fail

Violation of Assumptions

OP484 Ibias Prerad and 100 krad(Si)



- LM111 and OP484 show bimodal TID response within single wafer lot
 - Invalidates use of small samples
 - May require binomial sampling

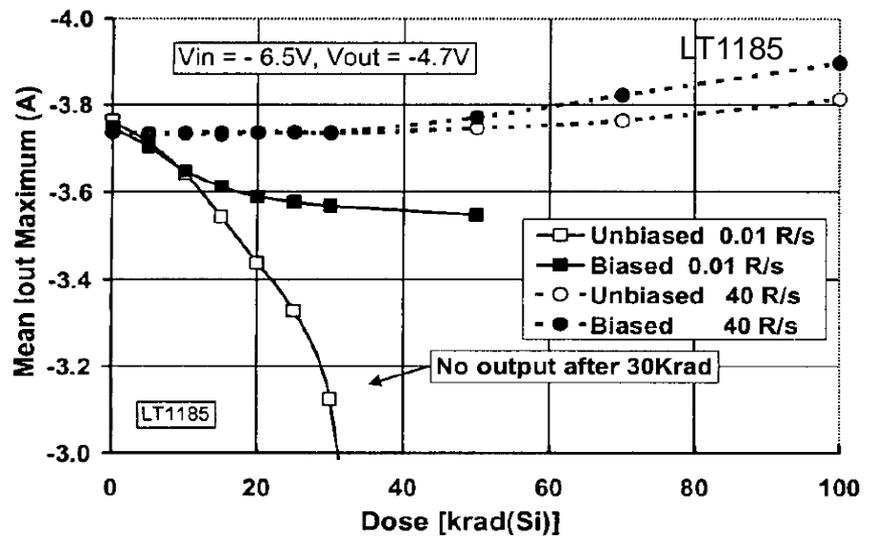
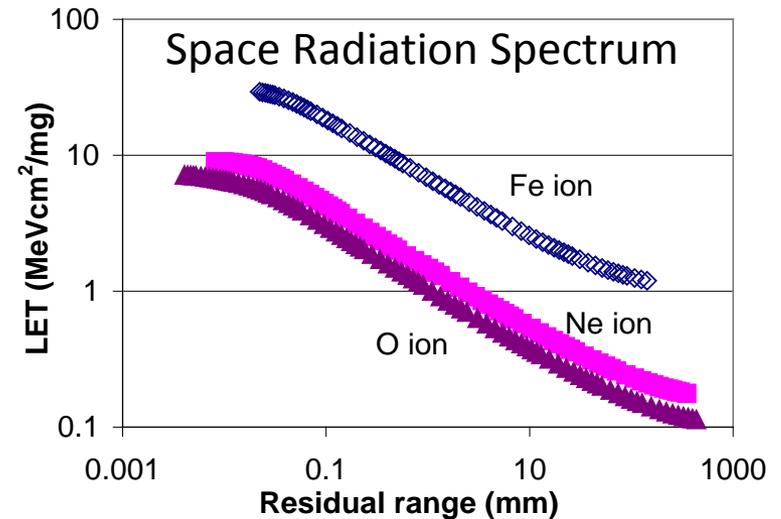


- Pre-2003, SET duration assumed $<100\mu$ s
 - OP293 showed $\gg 100\mu$ s SETs
 - LM6144 showed >1 ms SETs

Test Fidelity

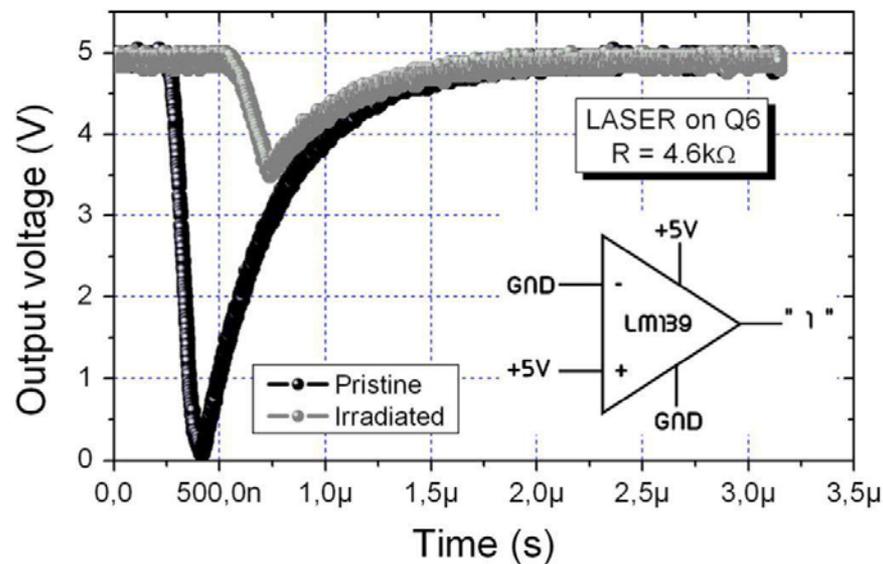
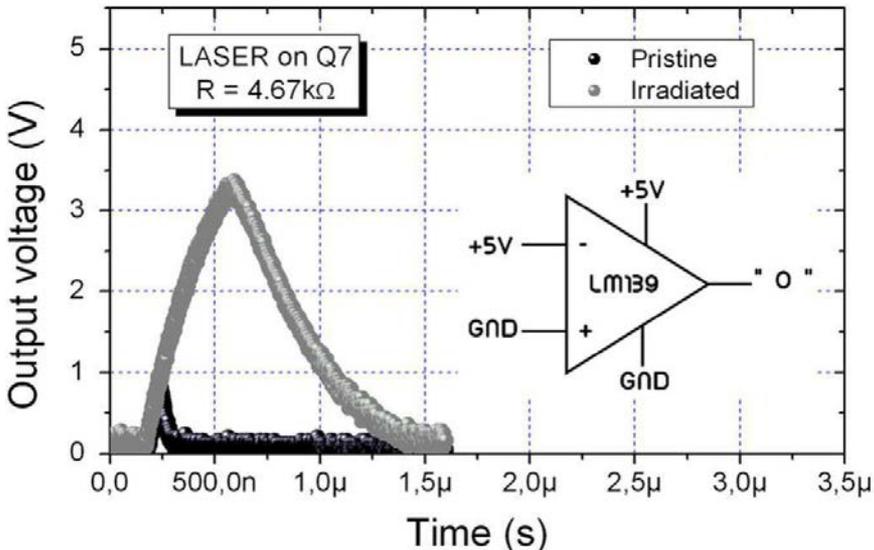


- Terrestrial heavy ions have much lower energies than those in space
 - Fe ion range @ GCR peak >14 cm
 - May affect efficacy of MBU hardening, etc.
- Enhanced Low Dose Rate Sensitivity (ELDRS) and bias effects can compromise test fidelity
 - Some ELDRS effects persist down to at least 0.5 mrad(Si)/s (ultra-ELDRS)
 - At high and low dose rates parts can degrade in opposite directions
 - Bias and ELDRS combined can drastically alter response
- Testing should be done under worst-case (WC) conditions
- ELDRS behavior can be modified significantly by the presence of hydrogen and thermal history

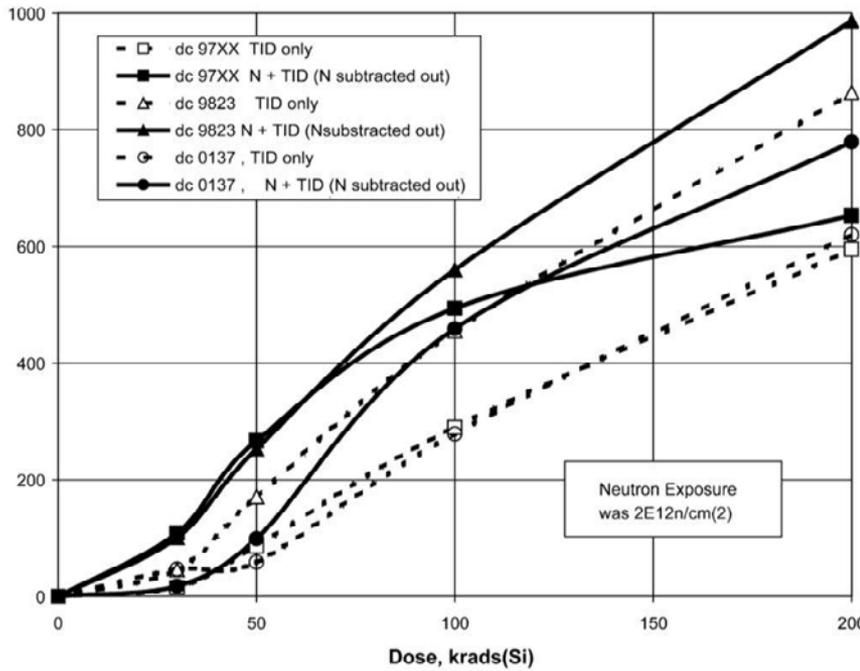


After McClure et al., IEEE REDW, p. 100, 2000.

Synergistic Effects



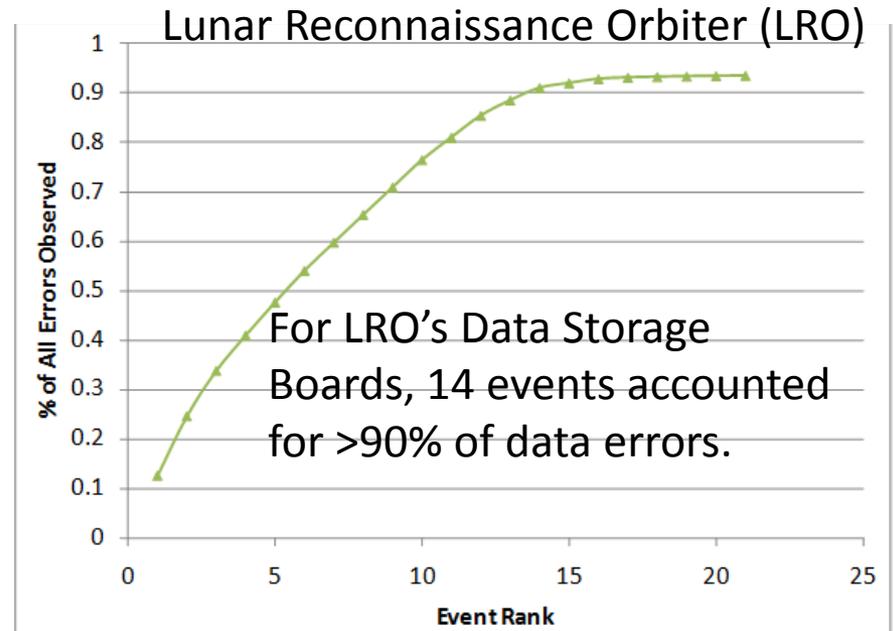
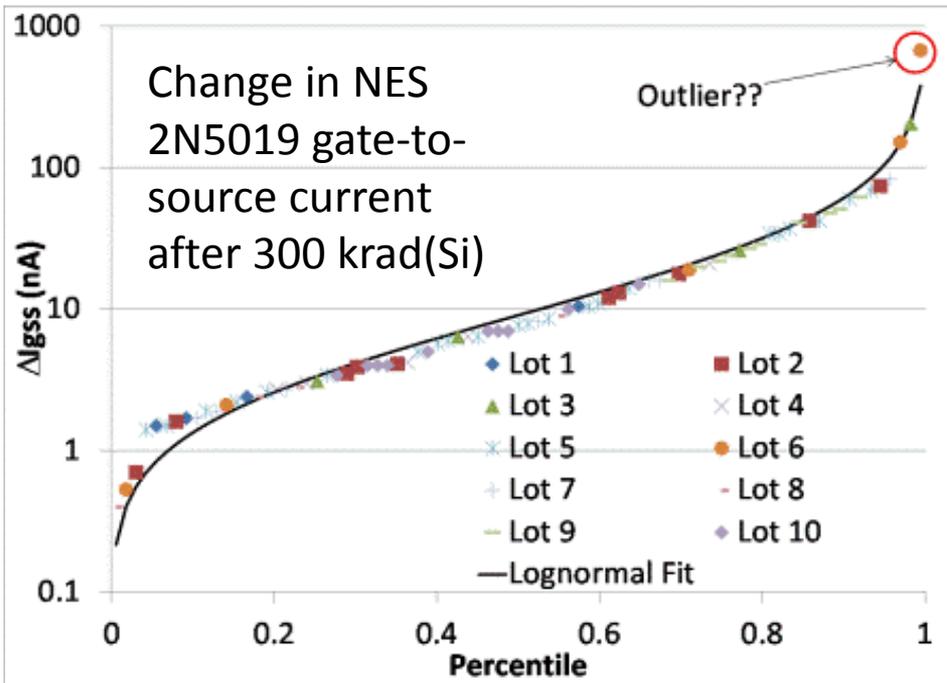
After Bernard, M., IEEE Trans. Nucl. Sci. **54** pp. 2534-2540



After Gorelick, J., IEEE Trans. Nucl. Sci. **51** pp. 3679-3685

- Susceptibility can change depending on past radiation history
 - TID affects SET duration, amplitude, rate
 - Displacement damage affects TID damage
 - Latent damage due to SEL can change SEL susceptibility in a test or mission

Rare Events Can Drive Risk



- For TID, is our sample representative?
- Not all wafer lots are well-behaved
- New England Semiconductor (NES) 2N5019 FET shows very thick tail in ΔI_{gss}
 - Lognormal fit to multiple lots underestimates WC part by $\sim 2x$. Is it an outlier or pathological part?

- SEFI—rare but severe consequences
 - May detect only a few WC SEFIs in SEE testing—very poor statistics
 - Difficult to estimate rate
- Situation similar for destructive SEE
- Risk driven by most poorly known SEE



A Few More Unpleasant Surprises

Rule of Thumb	Unpleasant Exception
Diodes immune to SEE	DSEE in Schottky diode (2011)
SEB depends only on ion LET and range/energy	Susceptibility depends on ion atomic number Z as well, complicating testing and rate estimation (2011)
WC for SEL is high temperature	Discovery of SEL at cryogenic temperature (2010)
SEE caused only by heavy ions and high-energy protons	SEUs seen in SRAMs due to direct ionization by low-energy protons and even muons (2007-11)
Digital SETs are short, so not as severe a problem as analog SETs	Has assumed increasing importance for hardening since 2005
SELs that leave part functional are nondestructive	Latent damage seen in SEL susceptible parts (2005)
MOSFETs immune to SEGR if $V_{DS} < 30\%$ of rated value	IRF640 commercial 200 V MOSFET fails @22% of rated value (2003)
SEGR depends only on ion LET and energy	Susceptibility depends on ion atomic number Z as well, complicating testing and rate estimation (2011)
TID independent of dose rate	ELDRS in bipolar (1994); in CMOS (sort of, 2005)



Other Threats to Affordable RHA

- Testability of complex state-of-the-art devices
 - Packaging may preclude irradiating sensitive volumes with heavy ions from conventional accelerators; Ultrahigh energy accelerators cost ~\$5K/hr
 - Complex state space and application dependence may preclude complete testing
- Commercial Off The Shelf (COTS) vs. radiation-hardened technologies
 - COTS outperform Rad Hard technologies by > 2 generations, often much more
 - Can be difficult to find rad hard parts to support COTS
 - If COTS support parts needed, they may also require significant system-level hardening
 - COTS lifetime <24 months; governed only by manufacturer—can change w/o notice
 - COTS qualification can take >1 year—especially for complex devices
- Proliferation of new technologies and imperfect mechanism models
 - End of conventional CMOS scaling in ~2000 led to many new technologies entering market
 - Radiation studies on many of these new technologies are very limited
 - Models imperfect even for simple technologies (e.g. SEGR/SEB in MOSFETs)
- Trends likely to worsen in the future
 - System-level hardening can facilitate reliability in the face of uncertainty



Conclusions (1)

- System-level hardening can be costly
 - \$\$\$ to design a system to mitigate error/failure modes in a device
 - Performance often compromised by mitigation efforts
 - System will often be more complicated, consume higher power and be heavier
 - Nonetheless, often the only way to realize benefits of COTS technologies
 - Because it is costly, system-level hardening should be driven by requirements
- System-level hardening in 5 categories:
 - Threat reduction/avoidance—especially for TID, DDD and destructive SEE
 - Infrastructure—Watchdog timers for SEFI, SEL circumvention circuits, system architecture and scrubbing to avoid multiple errors, etc.
 - Performance matching/derating—important for SETs and destructive SEE
 - Opportunistic mitigations—can be effective if part allows (e.g. interleaving for MBU, tailoring application conditions to minimize SET rate)
 - Redundancy—can be effective against data errors (but requires reliable parts); cold sparing can be effective for destructive SEE, but added weight is a problem
 - Techniques often combined to achieve greater hardness—need to be validated



Conclusions (2)

- Successful mitigation is predicated on the same assumptions as RHA
 - Lot traceability and test parts representative of flight parts
 - Mechanisms for radiation effects understood, allowing deviations from 100% environment/application fidelity
 - Testing can detect and illuminate all error/failure modes
 - When assumptions violated, hardening can fail
 - Rare events can drive risk
- New Technologies pose challenges for radiation hardening
 - Complex parts and packaging pose significant test challenges
 - New technologies emerging with little understanding of radiation threats
 - COTS pose challenges for traceability, availability, but performance is attractive
- System-level hardening likely to assume increased importance, due to
 - Increased uncertainty over radiation response of complex parts
 - Increased attractiveness of COTS due to their higher performance
- Challenge: balancing reliability and cost of assurance