


SDRAM Testing: Lessons Learned or How to test an SDRAM in Less than 7 years

Ray Ladbury


Radiation Effects and Analysis Group
NASA Goddard Space Flight Center



A Rad Hard Memory?


- Hard to 50 krad(Si)
- $BER < 10^{-11}$ errors/bit-day (97.6% of the time)*
 - *New error mode: SERBD
 - Single-Event Really Bad Day (2.4% of days)
 - BER up to 10^{-6} errors/bit-day
- That's SDRAMs
 - >90% of the time, they're great
 - But when they fail, it's spectacular!





And It's Not Just Data Loss

- Destructive Failures
 - Single-Event Latchup (SEL)
 - Others? Not yet, but do you feel lucky?
- Nondestructive Errors
 - Single-Event Functional Interrupt (SEFI) requiring power cycle to restore functionality—SEFI-PC
 - Recoverable SEFI (SEFI-rec), Block and Logic, Multi-Bit/Cell Upset (MBU), Stuck and SEU
- Error rates may depend on
 - Operating mode (dozens possible), frequency, supply voltage...
 - Beam Daddy estimates ~7.5 years for an exhaustive test!
 - Beam Daddy usually gives me 12 hours
- Result: All SDRAM tests are application specific
 - Test plan must consider not just application conditions, but also possible mitigation for the application
 - Hey, I'm trying to save you 7.5 years of testing

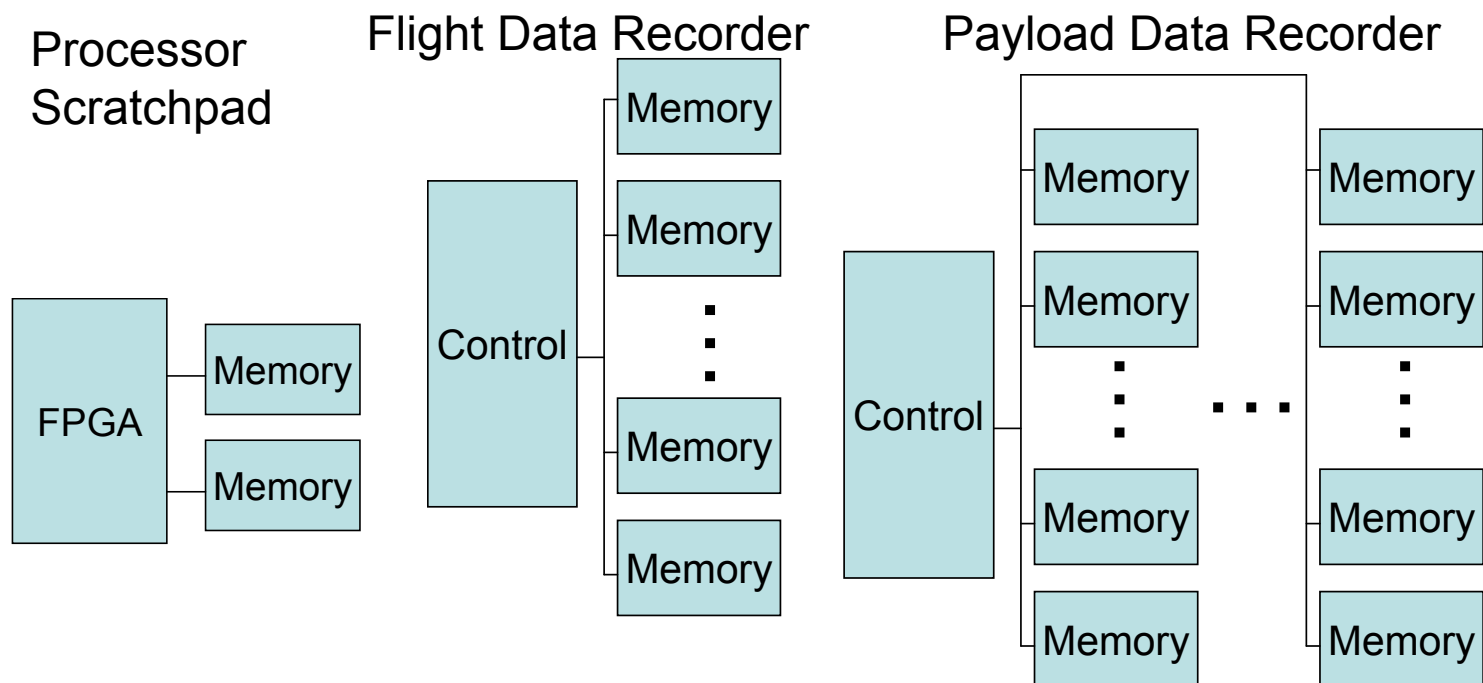


Outline

- I. Typical SDRAM Applications, Requirements and Mitigation
- II. Error occurrence and Propagation in SDRAM Applications
- III. Test Planning and Setting Priorities
- IV. Test Execution: Time Management and Statistics
- V. Data Analysis and Rate Calculation




The Application Drives Everything



Size	<10 Gbit	>100 Gbit	>500 Gbit
Mitigation	Limited	Good	Good
Reliability	High	High	High
Availability	High	High	High
Integrity	High	Moderate to High	Moderate
Retention	Short	Long	Long


Mitigation: What We Can and Can't Do




Not So Great

Pretty Darned Good

Destructive Failure
Consequences:
Permanent loss of
one memory die
Mitigation:
1) Find Immune Part
2) Redundancy



SEFI
Consequences:
Loss of functionality
of one memory die
Mitigation:
1) Find Immune Part
2) Reload Mode Reg.
3) Cycle Power



Data Corruption
Consequences:
Loss of up to all data on a single memory die
Mitigation:
1) EDAC + bit interleaving + Error Scrubbing
Hamming Code—SECDED; 20% overhead
Reed-Solomon— ≤ 2 nibbles; 50% overhead
2) Triplicate voting + Error Scrubbing
Overhead: 200%

Bonus
Data loss also corrected for SEU, MBU, and
even stuck bits

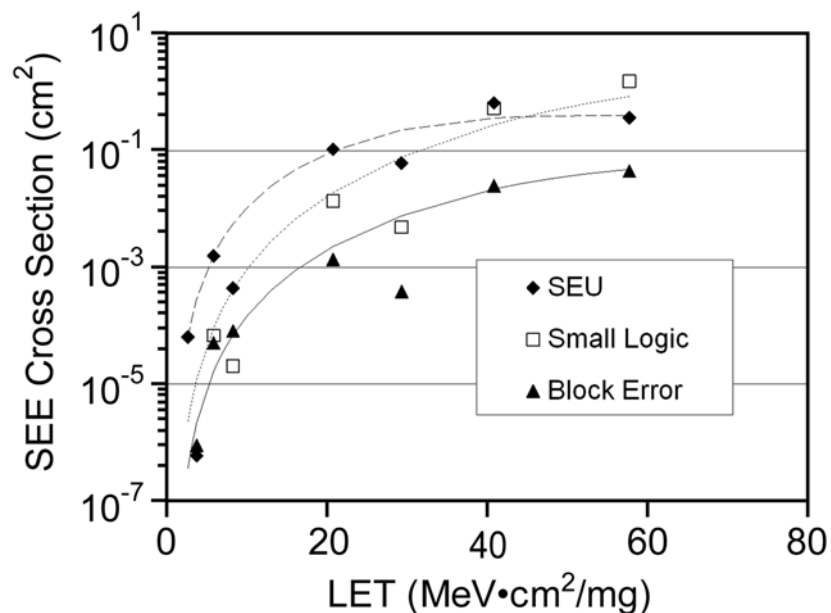


How Errors Occur

LRO Data Summary

- 248 days × 864 die =
 - 214272 device days
 - 87264 error free (41%)
 - 58176 singles (27%)
 - ≤10 errors (97%)
 - ER = 1.12E-8/day
 - WC ER = 1.2E-7/day
 - ER excluding 6 worst days = 6.25E-10/day

GCR Rates	Bits in Error	Predicted	Observed (LRO)*
SEU (bit ⁻¹ *day ⁻¹)	1	2 × 10 ⁻¹¹	6.8 × 10 ⁻¹²
Logic Error (dev ⁻¹ *day ⁻¹)	2-20	1 × 10 ⁻²	5 × 10 ⁻⁴
Block Error (dev ⁻¹ *day ⁻¹)	21-10 ⁴	3.5 × 10 ⁻³	4.2 × 10 ⁻⁵
SEFI (dev ⁻¹ *day ⁻¹)	5.12 × 10 ⁸	~8 × 10 ⁻⁶	< 5 × 10 ⁻⁶



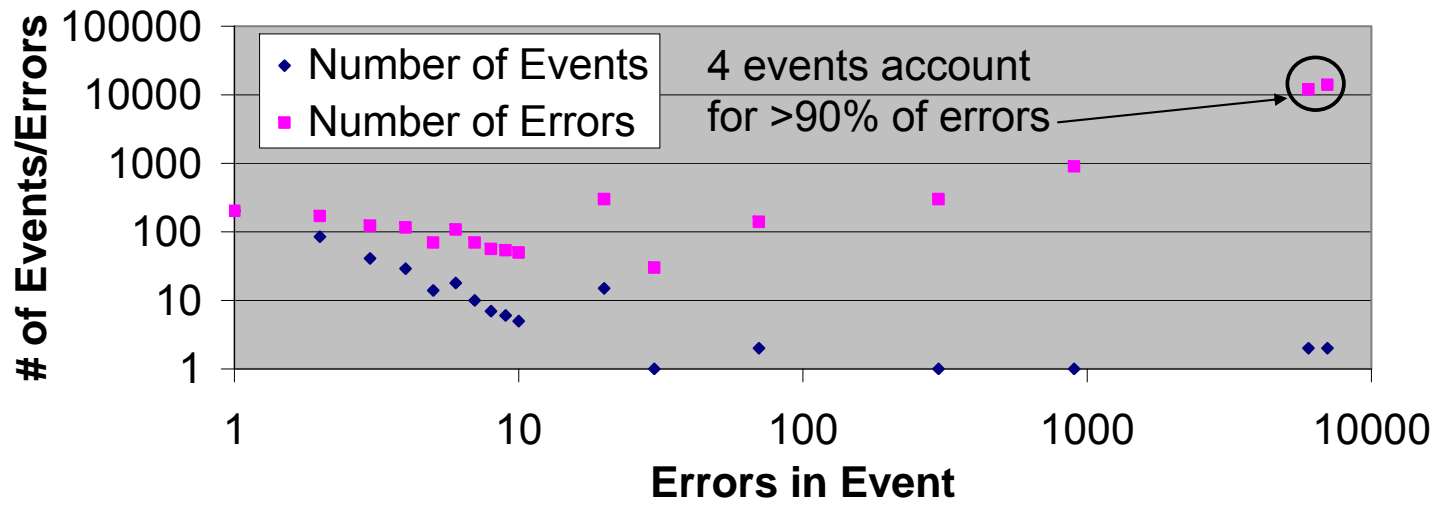
- SEU rate ~3.5E-3 per die per day
- Logic + Block Error ~5.4E-4 per die per day
 - Average ~160 words/error corrupted

Questions

- How important is it not to have a bad day?
- How much is bullet-proofing worth?



Error Occurrence and Propagation



On-Orbit Data
from LRO for
Elpida 512-Mbit
SDRAM

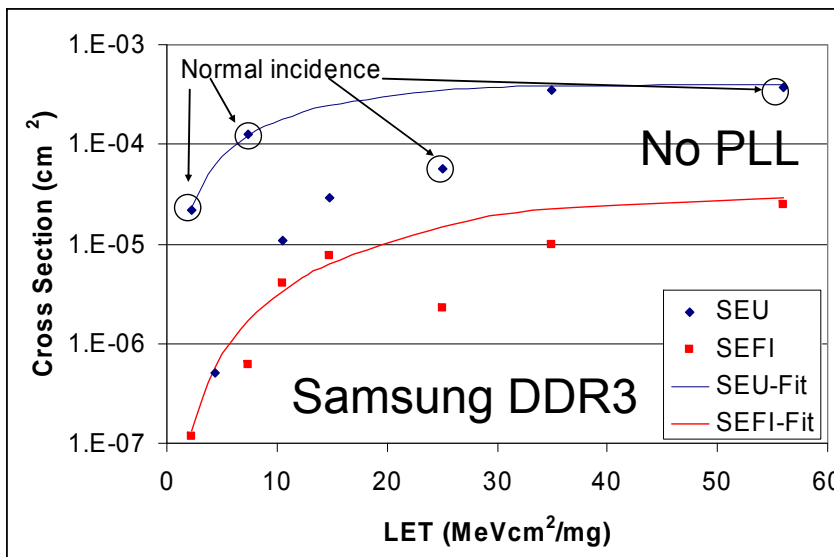
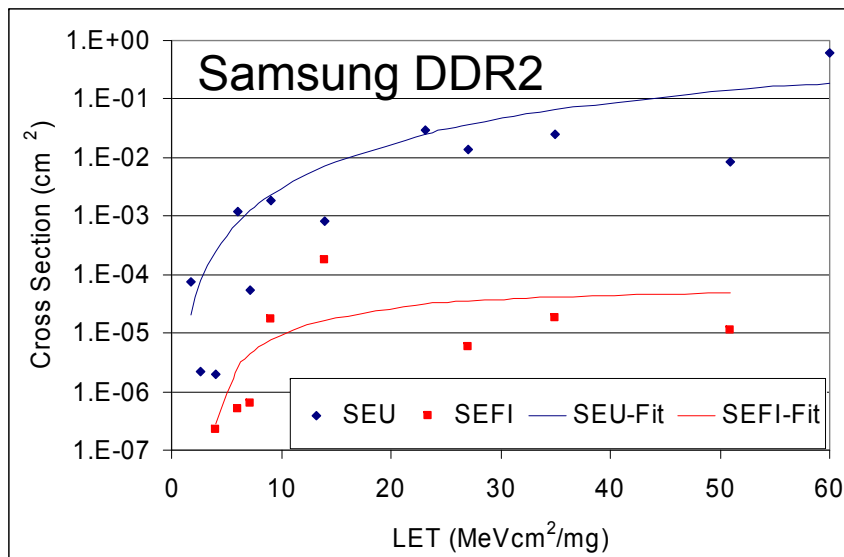
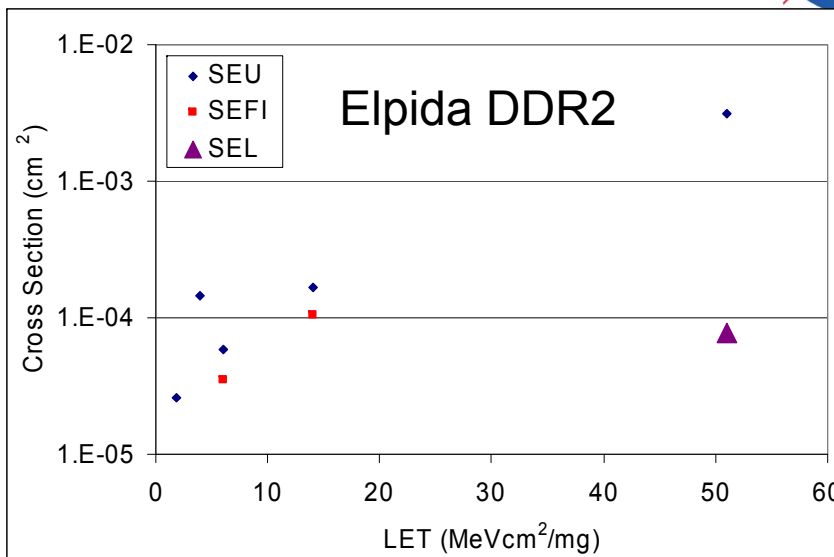

- If it takes $n+1$ errors (occurring with rate r) in N devices during refresh period t to cause an error at system level:


$$P(rt, n) = 1 - \sum_{i=0}^n \binom{N}{i} (rt)^i (1 - rt)^{N-i} \cong \binom{N}{n+1} (rt)^{n+1}$$

- If rate r is in error by a factor x , system level error factor $\sim x^{n+1}$
- Usually, dominant error modes for SDRAMs are SEFIs or Block Errors, which usually have poor statistics.



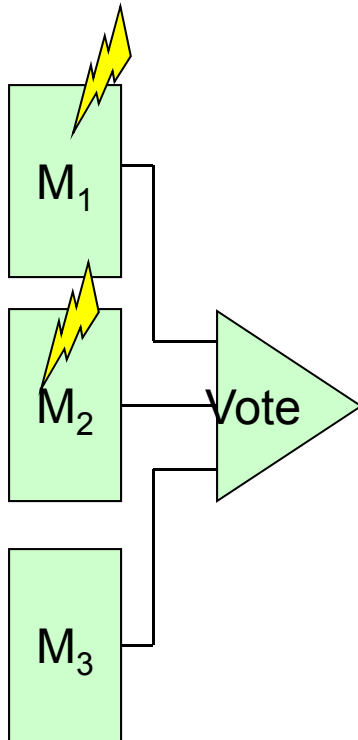
Other Recent Results



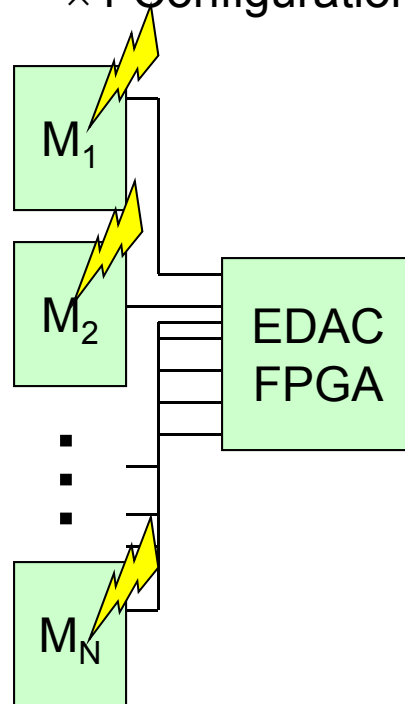


Organization and Error Propagation

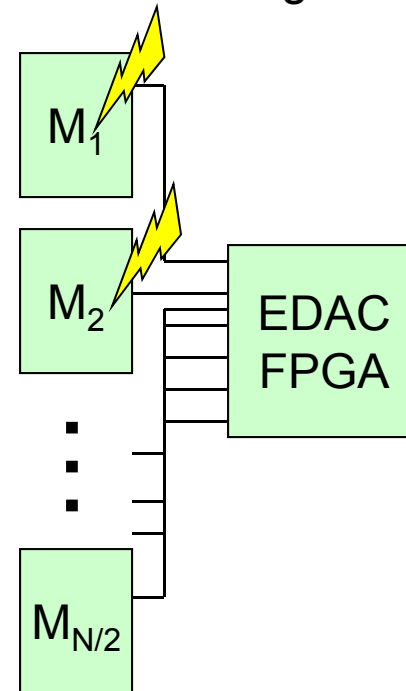
Triplicate Voting



R-S Double-Nibble ×4 Configuration



R-S Double-Nibble ×8 Configuration



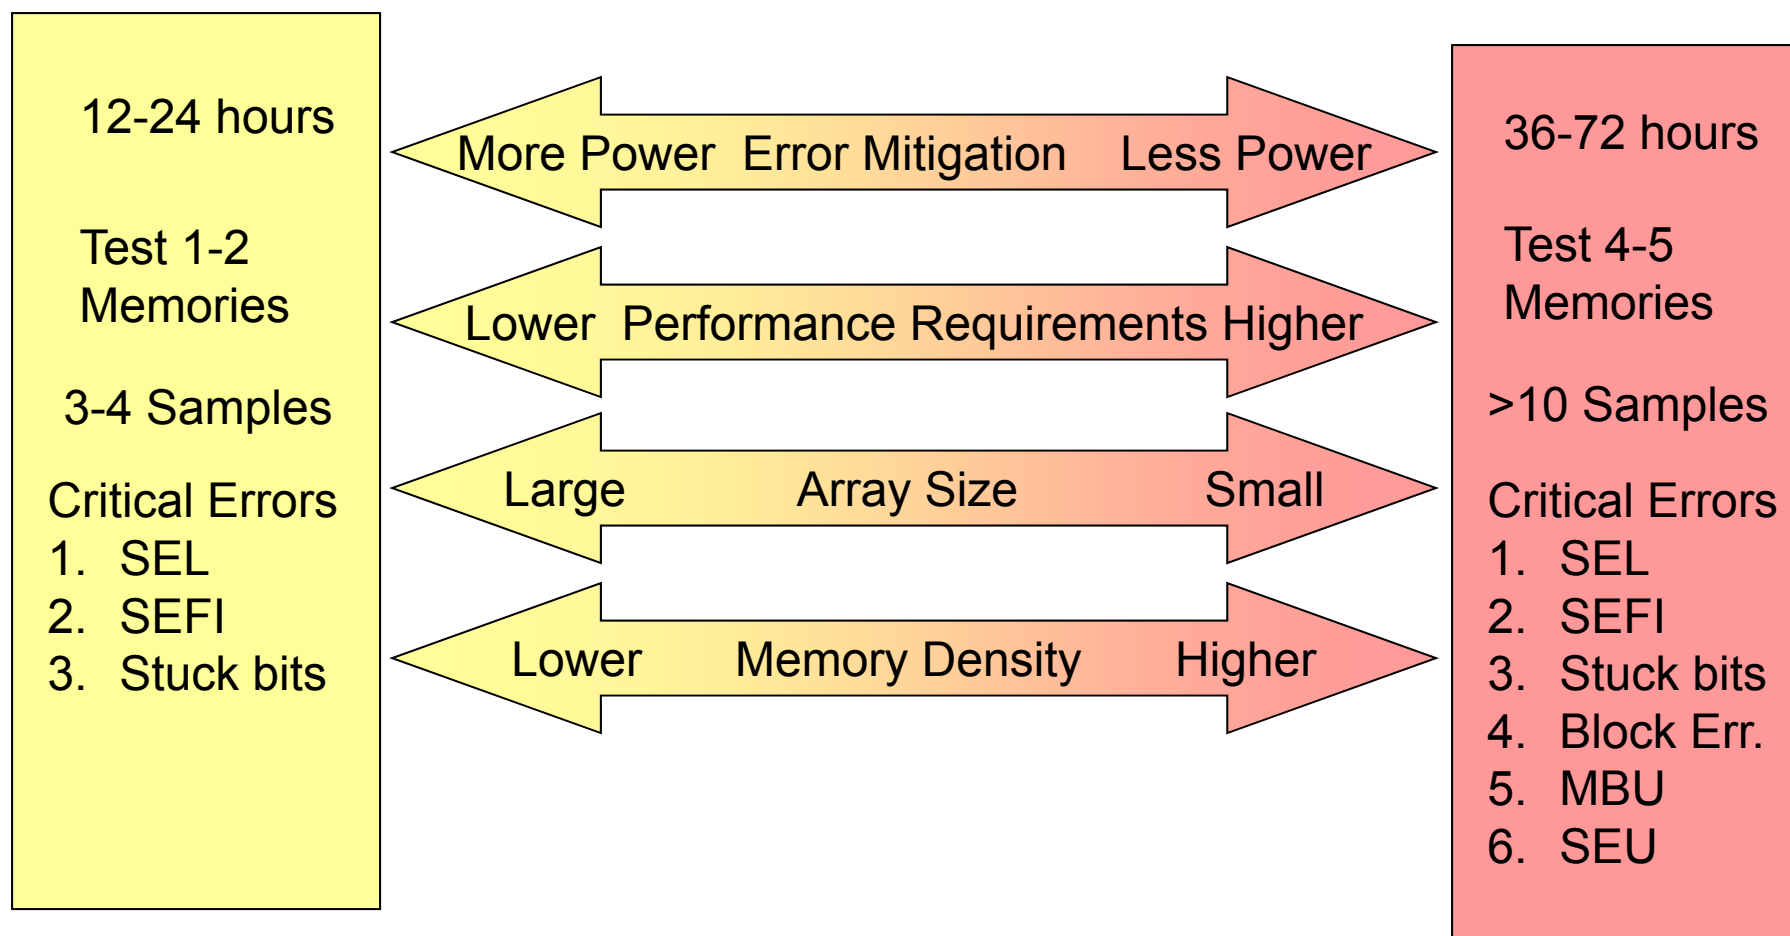


Testing Drivers

“Easy” Testing

Determinants of Testing Difficulty

Hard Testing







Test Planning: Where does all the time go?

- Test Priorities
 - 1) Establish (effective) immunity to destructive SEE (e.g. SEL)
 - Requires high LET, fluence, but does not require normally functioning DUT
 - Single high-flux/fluence run may be sufficient
 - 2) Establish immunity or susceptibility to SEFI-PC (requiring power cycle)
 - Problem: Susceptibility may depend on state of DUT
 - Need many low-flux/fluence runs totaling high fluence w/ DUT in known state
 - 3) Other Errors (SEFI-rec, Block Errors, MBU, SEU, stuck bits)
 - Susceptibility may be state dependent, but cross sections are large
 - Runs should be low flux/fluence, but number of runs is manageable
- Need to determine SEFI-PC susceptibility drives test requirements
 - Need high total fluence to ensure rare events are seen, but low-fluence runs to avoid having DUT in an unknown state for majority of the test
 - Long runs with frequent reset have large deadtime
 - SEFI-rec and Block Errors are common and extremely disruptive
- Can we screen for SEFI-PC susceptibility another way?



Laser Testing to Complement Heavy-Ion Testing

- Laser testing can establish
 - SEL susceptibility
 - Susceptibility to SEFI-PC
 - Whether multiple errors result in a spurious SEFI-PC
 - Relative susceptibilities to SEU, Block Errors, SEFI-rec, etc.
- Validate laser test w/ Heavy-Ion test
- Challenges
 - Die need to be polished for laser tests
 - Places added stress on die
 - Laser spot size $>1 \mu\text{m}^2$
 - Difficult to translate luminosity to LET
- For 512 Mbit Elpida SDRAM
 - TPA laser test revealed SEFI-PC susceptibility and state dependence





The Poor-Man's Laser Test

- Xe strobe of a camera can induce SEE in microelectronics
- **WARNING: THIS MECHANISM IS DISTINCT FROM SEU**
 - Floods substrate with charge
 - Cannot calculate critical charge
- **BUT**
 - Can verify test set-up working
 - Can give qualitative idea of susceptibility to
 - SEL, SEFI
 - Vary charge by varying distance
 - Parts that show little susceptibility to photoflash SEL or SEFI MAY be more immune to HI as well.





Test Execution: Who Needs a Schedule?

- Thinned parts can fail any time
 - ~50 % failed before or just after testing
 - Schedule high-priority tests first
 - Minimize mounting/demounting
 - Bring extra die!!!
- OK, maybe the results are interesting
 - But are they interesting in the application?
 - Is it potentially destructive?
 - Is it severely disruptive?
 - Will it overwhelm error correction?
- Remember:
 - For Elpida 512 Mbit SDRAM, errors on SEFI, Block Error rates were large
 - These modes dominate data loss
 - Statistics are key to improving rates
 - But






Interference between Error Modes


Near Saturation, SEUs and Block Errors have nearly the same σ

LET	Avg. Fluence to 1st SEU	Avg. Fluence to 1st Block Err.	Avg. # SEU@ 1st Block Err.	% Read before 1st Block Err ($100 \text{ cm}^{-2}\text{s}^{-1}$)
2	33333	2500000	75	All
22	10	667	67	~7%
42	3	63	25	<1%

- Fluxes $< 100 \text{ particles cm}^{-2}\text{s}^{-1}$ are impractical due to dosimetry
- Once Block Error or SEFI occurs, further SEUs are likely lost
 - For LET=42 MeVcm²/mg, likely to read Block Error before 1st SEU!
- Runs ~1 sec. @ Fluxes $< 100 \text{ particles cm}^{-2}\text{s}^{-1}$ can yield clean data
 - For LET=42 MeVcm²/mg, 20% of 100-ion runs will be block error free with ~33 SEUs/run
 - Same is true for SEL or SEFI in susceptible devices.




Test Execution: How many Events are Enough?



- σ_{sat} errors affect rate ~linearly
 - W and σ_{sat} are correlated
 - ~16 events at high LET correlates to ~25% error
- LET_0 error δ affects rate as roughly $2\delta / LET_0$
 - A few events is enough
- Error on s can have a big effect on rate for small LET_0 and σ_{sat}
 - Many points in σ vs. LET curve are helpful, but effective LET doesn't hold
- Each event require several minutes to assess recovery and reinitialize part

$$\sigma(LET) = \sigma_{sat} * \left(1 - \exp\left(-\left(\frac{LET - LET_0}{W} \right)^s \right) \right)$$

- Curves onset behavior determines LET_0
- Saturation behavior determines σ_{sat} and W
- s determines how “s” shaped the curve is



Data Analysis and Rate Calculation

- Data analysis begins with categorizing events
 - SEU cross sections estimated from “clean” runs
 - Multi-bit upsets may contaminate SEUs, elevating σ at high LET
 - Block errors averaged over several runs
 - May want to break block errors into categories by #bits affected
 - Block errors should be recoverable
 - Single-Event Functional Interrupt
 - Recoverable SEFIs may self-recover or may require mode register refresh
 - Non-recoverable SEFIs require a power cycle or other disruptive measure
 - Massive data loss and unavailability of system during recovery
 - Other modes: SEL, Stuck bits
- Rate Calculation is challenging
 - Effective LET may be invalid
 - Large errors necessitate a conservative fit
 - With good statistics, rates good to better than 10x



Conclusions

- SDRAM are the Jekyll and Hyde of space application and testing
 - Testing necessarily application dependent, because
 - Different applications have radically different requirements
 - Different mitigation options are feasible for different applications
- Testing is demanding and time consuming due to high cross sections for disruptive error modes
 - Disruptive error modes drive data loss, outages and test requirements
 - Complementary testing can be invaluable
 - Mix of high-fluence and low-flux/fluence runs needed to achieve adequate statistics for all error modes.
- CREME96 can still give adequate results despite severe challenges
- Possible improvements
 - New ion beams at high energy to fill in LET points on σ vs. LET curve
 - More rapid automated recovery from disruptive errors in testing
 - Better integration of complementary test techniques