



TMR in Virtex-4 and RHBD in Virtex-5

Current Status of Two Approaches to Attaining Robustness
of Reconfigurable FPGA Applications in Upset Environments

Gary Swift
and the Xilinx Radiation Test Consortium

Three Options for Dealing with Upsets

1. Do nothing - live with intrinsic upset rate

For rFPGA's not all config upsets yield errors

However, nth error 'breaks' design ($n \approx 10$)

2. Upset mitigation - upsets \neq errors

Prevent a single upset from causing an error

Prevent upset accumulation

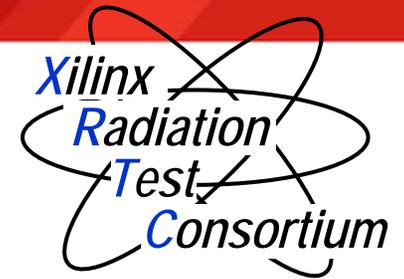
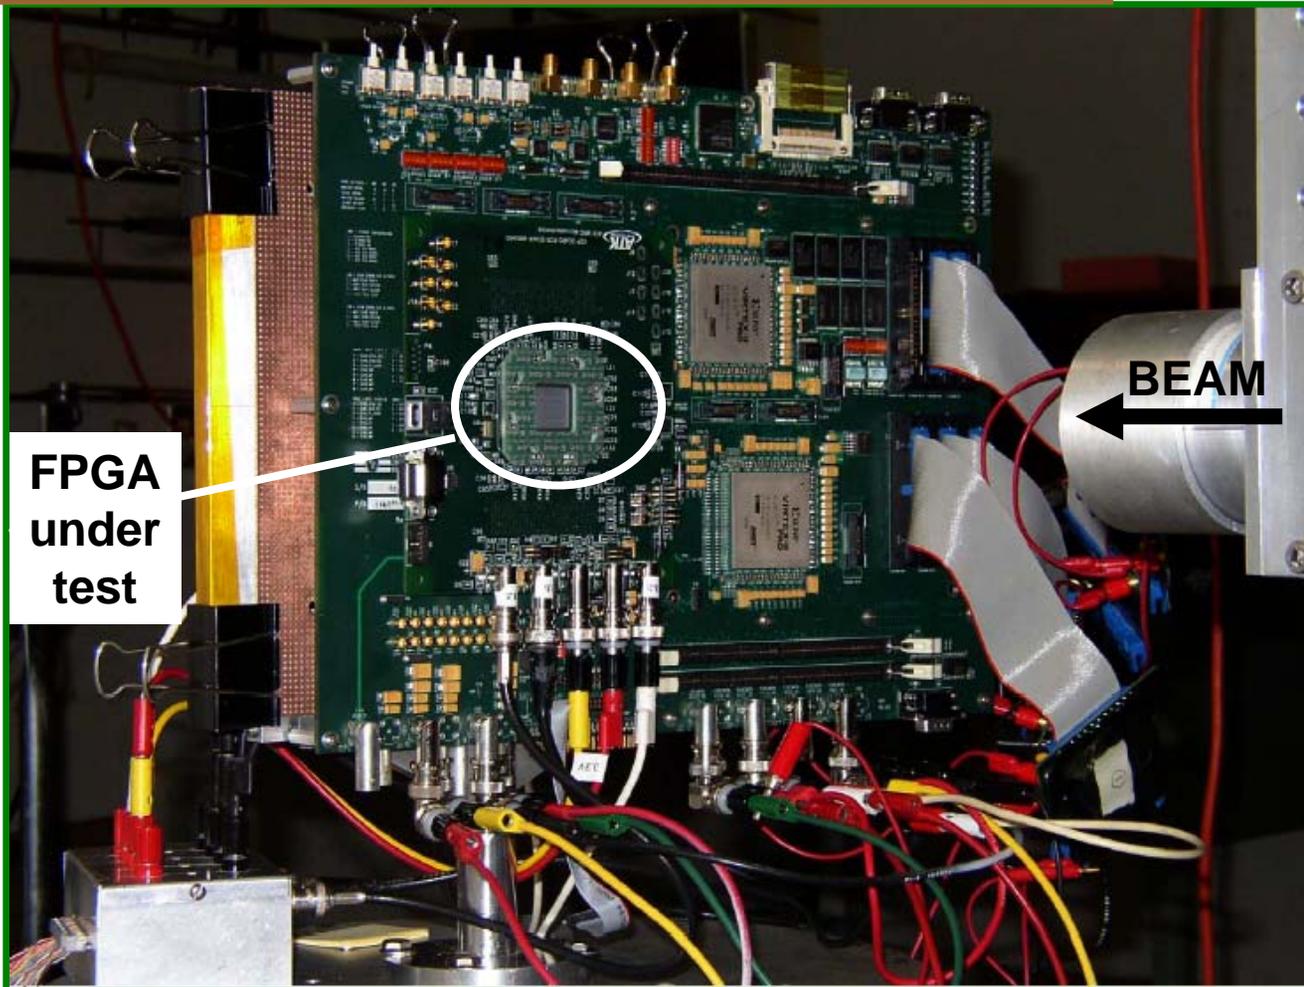
3. Harden to upset - no upset = no error

Outline

- **XRTC (Xilinx Radiation Test Consortium) Background**
 - Voluntary membership
 - Test types: static, dynamic, and mitigation
- **Design Robustness**
 - TMR (triple modular redundant) designs plus config management
 - RHBD (rad-hard by design) fabric
- **Calculating Upset and Error Rates**
 - Ordinary (single-node) - SEFI example
 - TMR case - application example
 - Dual-node case - data-based example

XRTC Apparatus in Action

Testing at Texas A&M Cyclotron Institute in Air



Boeing



XRTC Beam Tests

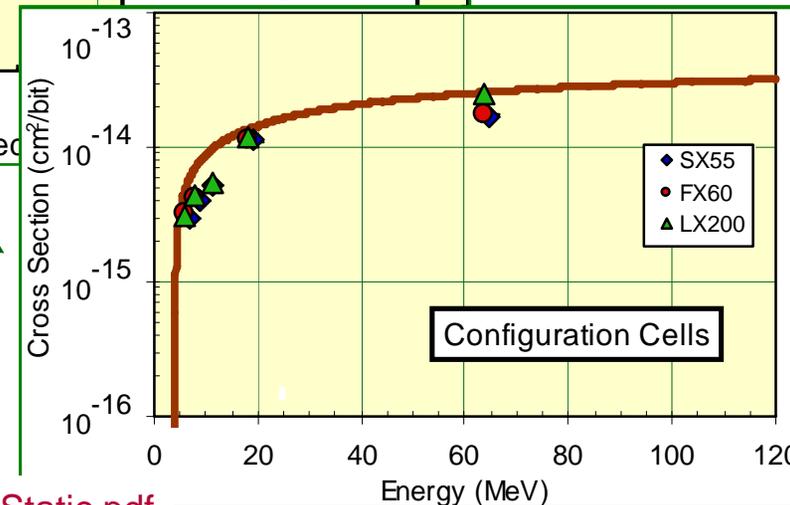
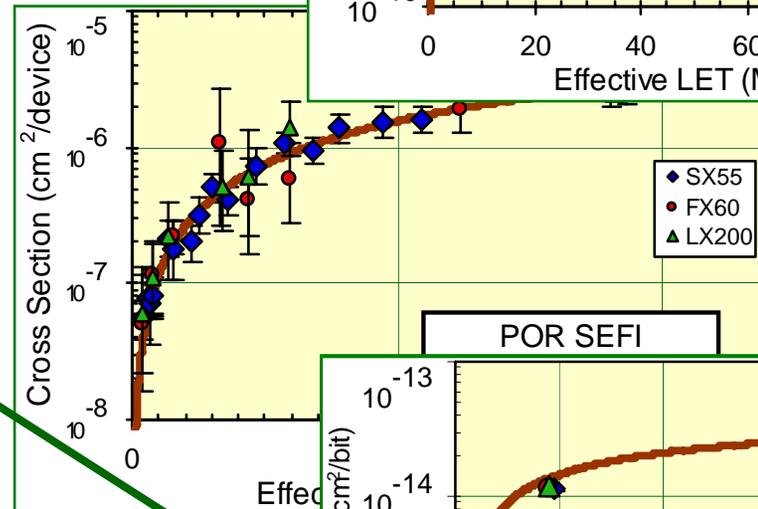
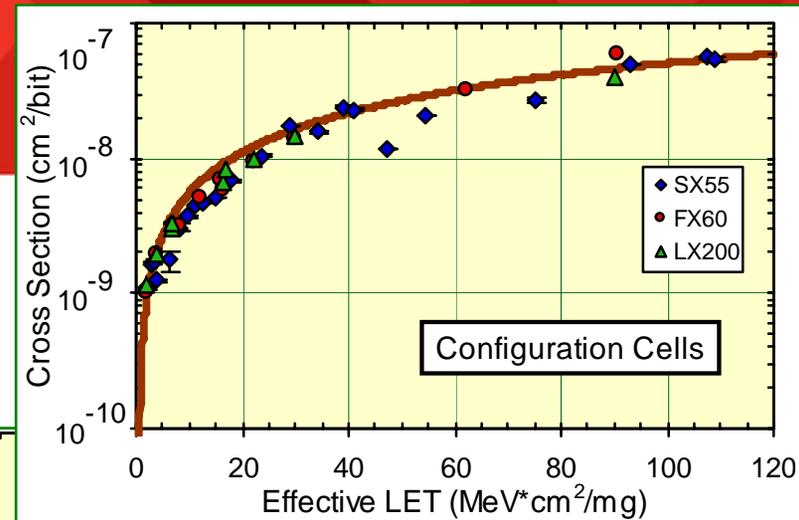
Static Results on V4-QV

- Config cells
- User BRAM & FFs
- Functional Upsets (aka SEFIs)
- Both Heavy Ions & Protons

Dynamic Results

- Digital Clock Managers
- Half Latches

Mitigation Campaign, Ongoing



Upset Mitigation

Redundancy -

Extra information (bits) prevents all upsets from yielding system errors.

Scrubbing required –

Accumulation of errors rapidly kills mitigation effectiveness.

Effective –

Most spacecraft now fly large arrays of upset-soft memories with few or no errors.

Typically, uncorrectable errors are detectable.

Upset Hardening – Two Basic Approaches

Both Approaches -

- Add circuit elements to basic storage cell
- Increase storage cell stability

Approach 1 - Increase “critical charge” to upset

- Add passive element(s) into cell feedback path.
- Cell size increase may be small, but it's slower
- Standard upset rate calculation does work

Approach 2 - Require charge in two nodes

- Add geometrically separated active elements.
- Standard upset rate calculation doesn't work

Space Upset Rate Calculation

Involves three basic elements:

1. Upset susceptibility measurements

cross section vs. “effective” LET

2. Environment specification

integral flux vs. LET

3. Angular response model

RPP[†] chord length distribution

one adjustable parameter:

charge collection depth

(aka funnel length)

[†] RPP = rectangular parallel piped, i.e. a 3-D box

Simplifying concepts (or useful fictions)

Critical charge:

If a node collects more charge than the critical amount, then the cell upsets.

Effective LET:

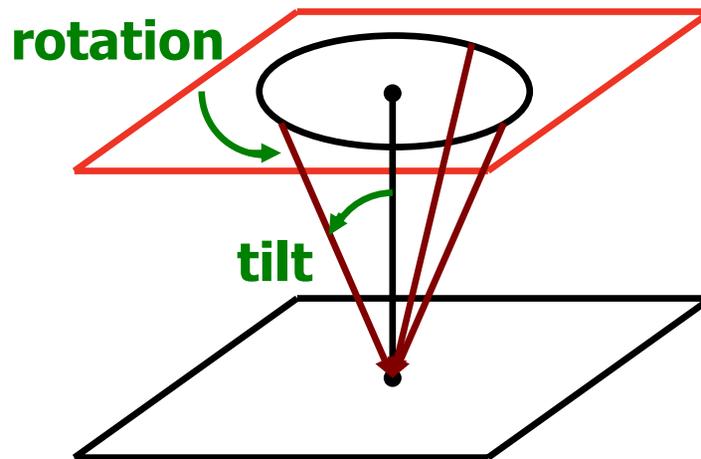
An ion's "effective" energy (or charge) deposition is related to the cosine of the tilt angle (off normal incidence) that it strikes.

RPP charge collection volume:

All charge deposited in RPP goes to node, while all charge outside does not.

Inherently, this is a “single node” calculation

Although a cell may contain multiple charge collection nodes capable of upsetting the cell, the charge collected is only dependent on the “tilt” angle and not the rotational orientation:



Several ion trajectories all with same tilt angle, but various rotation angles.

Results for Virtex-4QV FPGAs in GEO

Configuration upsets:

Less than twelve per day

SEFIs:

About one per century

Processor Upset Rates – Mild Environment

for GEO:

	<i>Hardening</i>	<i>Upset Rate</i>	<i>Ratio</i>
BAE RAD750 (estimated)	RHBD	2.2	x1
Maxwell SCS750	TMR*	1.1×10^{-5}	÷200,000
Virtex II-Pro ePPC405	none	13	x6

per year

RHBD = radiation (upset) hardened by design

* Processor-level TMR, scrubbed ten times per second, with ~3% performance hit

Notes:

Assumes 100% duty cycle on all bits (registers and L1 caches)

Environment = Galactic Cosmic Ray (GCR) background at solar minimum

Shielding = 100 mil Aluminum-equivalent

Limits of Upset Mitigation

Common sense says -

At some point, upsets will occur too rapidly and the mitigation will be “overwhelmed.”

In fact, Edmonds approx. equation says –

There’s not really a “cliff.”

The relationships are known; the error rate:

- (1) increases with the square of upset rate
- (2) decreases linearly with faster scrub rates
- (3) is directly proportional to EDAC word size[†]

[†] EDAC word size = data bits + check bits ; EDAC=error detection and correction

Edmonds TMR Equation

- Approximation when r (upset rate) is small:

$$R \approx 3M T_C (\mathcal{M}_2 r)^2$$

System Error Rate (points to R)

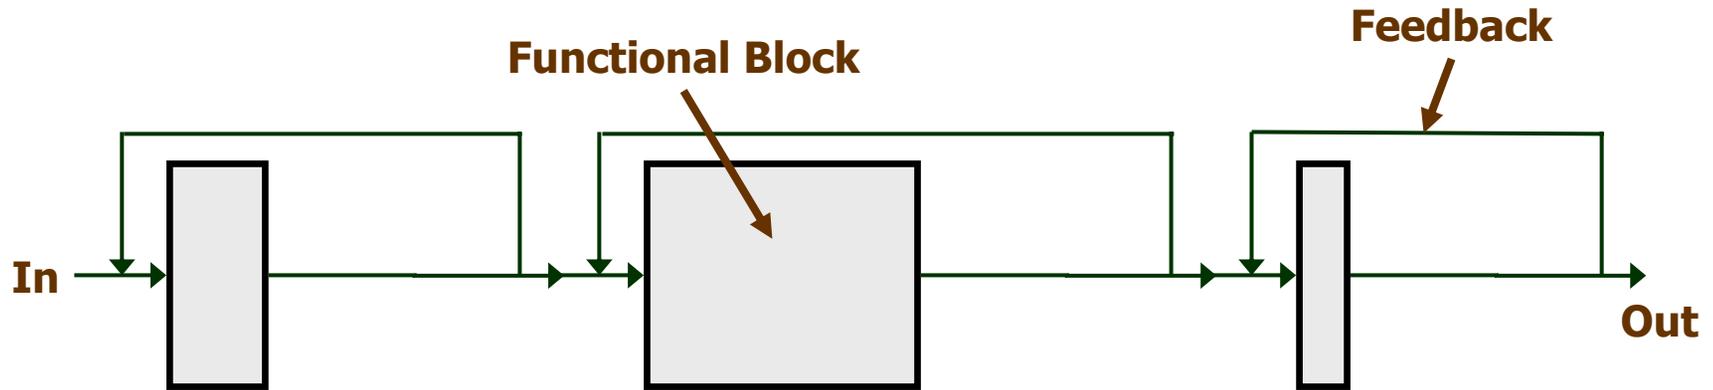
Total Modules (points to $3M$)

Scrub Time (points to T_C)

Underlying Upset Rate (points to r)

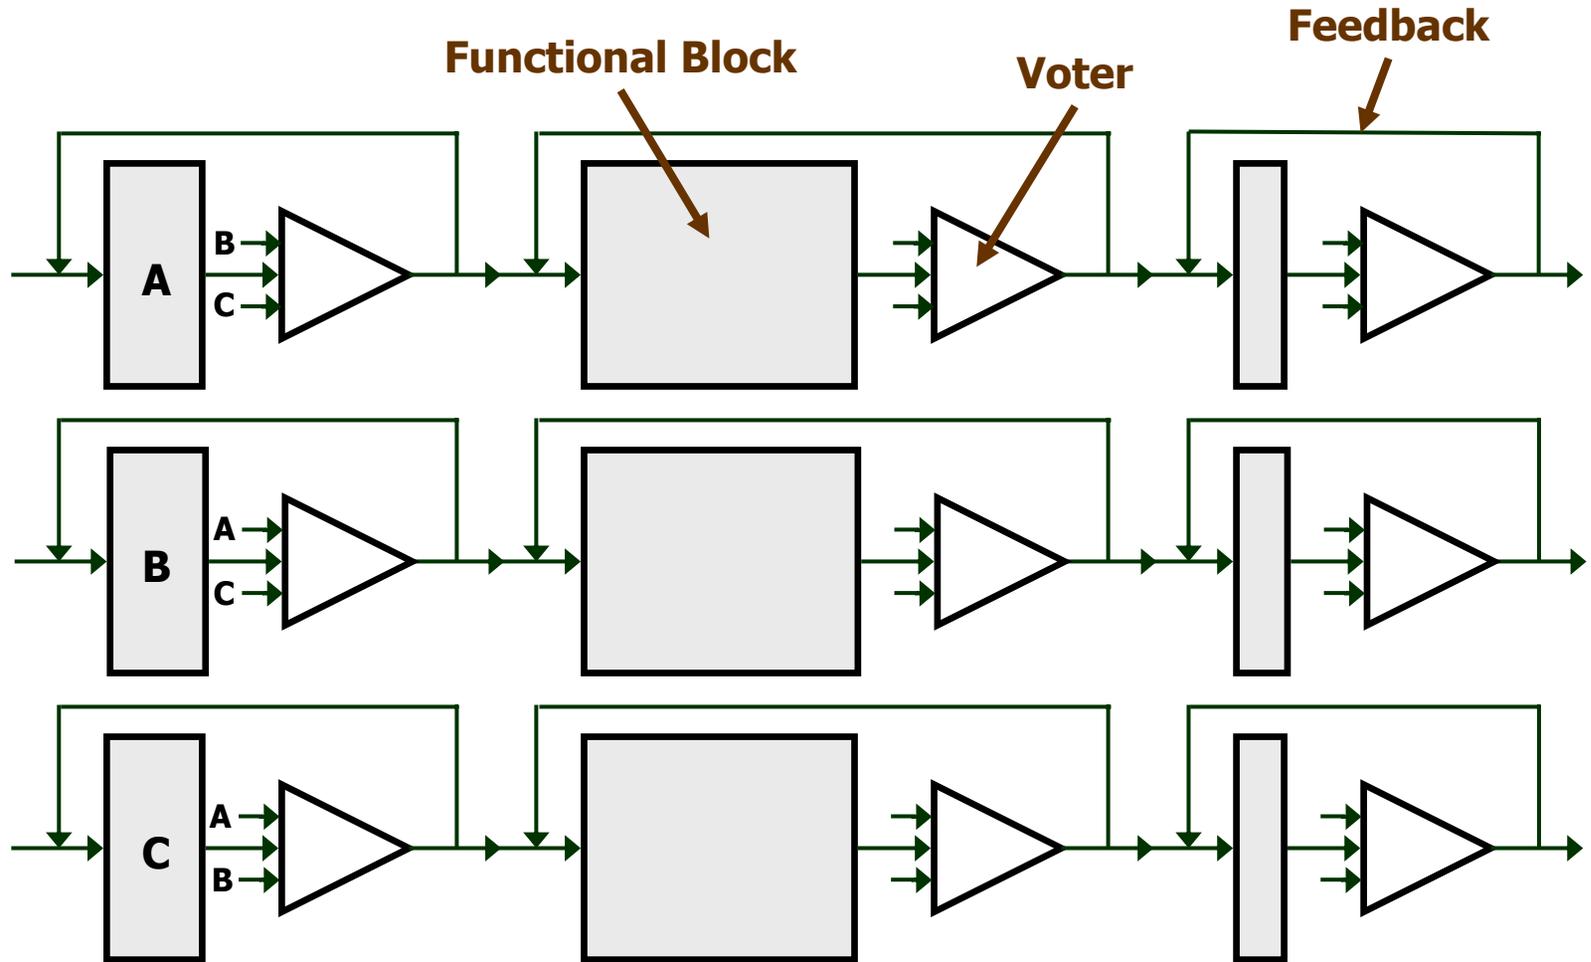
“Fitting” Parameter
is
root mean square module size

Single-String Design



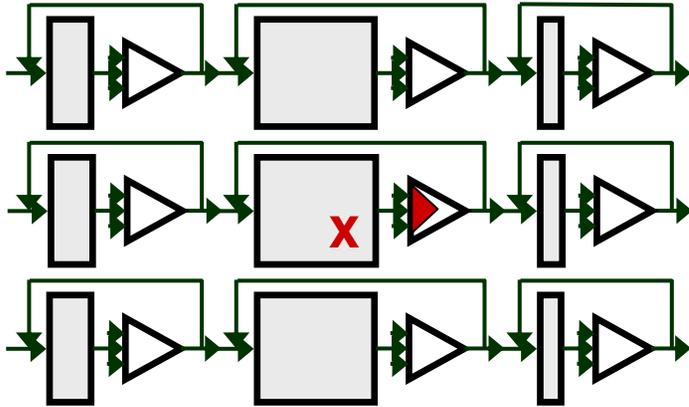
Conceptually, a design is a string of logic blocks (sequential or combinational) bounded by feedback loops.

TMR Design

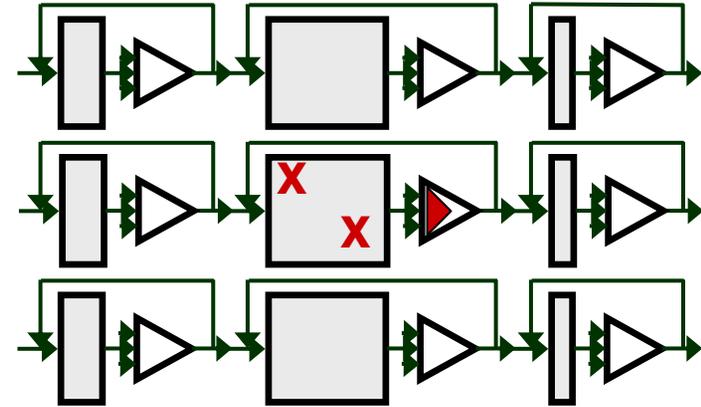


Feedback from the voters corrects state errors inside blocks

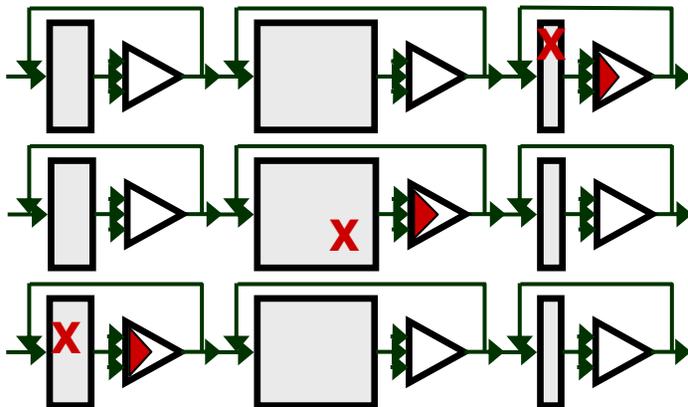
TMR prevents almost all errors



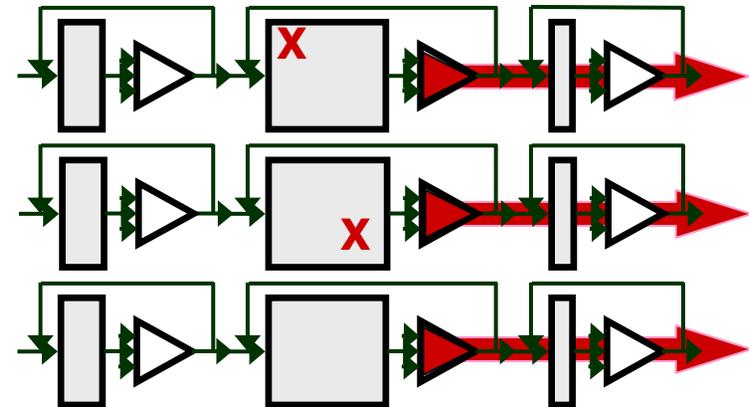
Single upsets cannot cause errors



Multiple upsets but no error

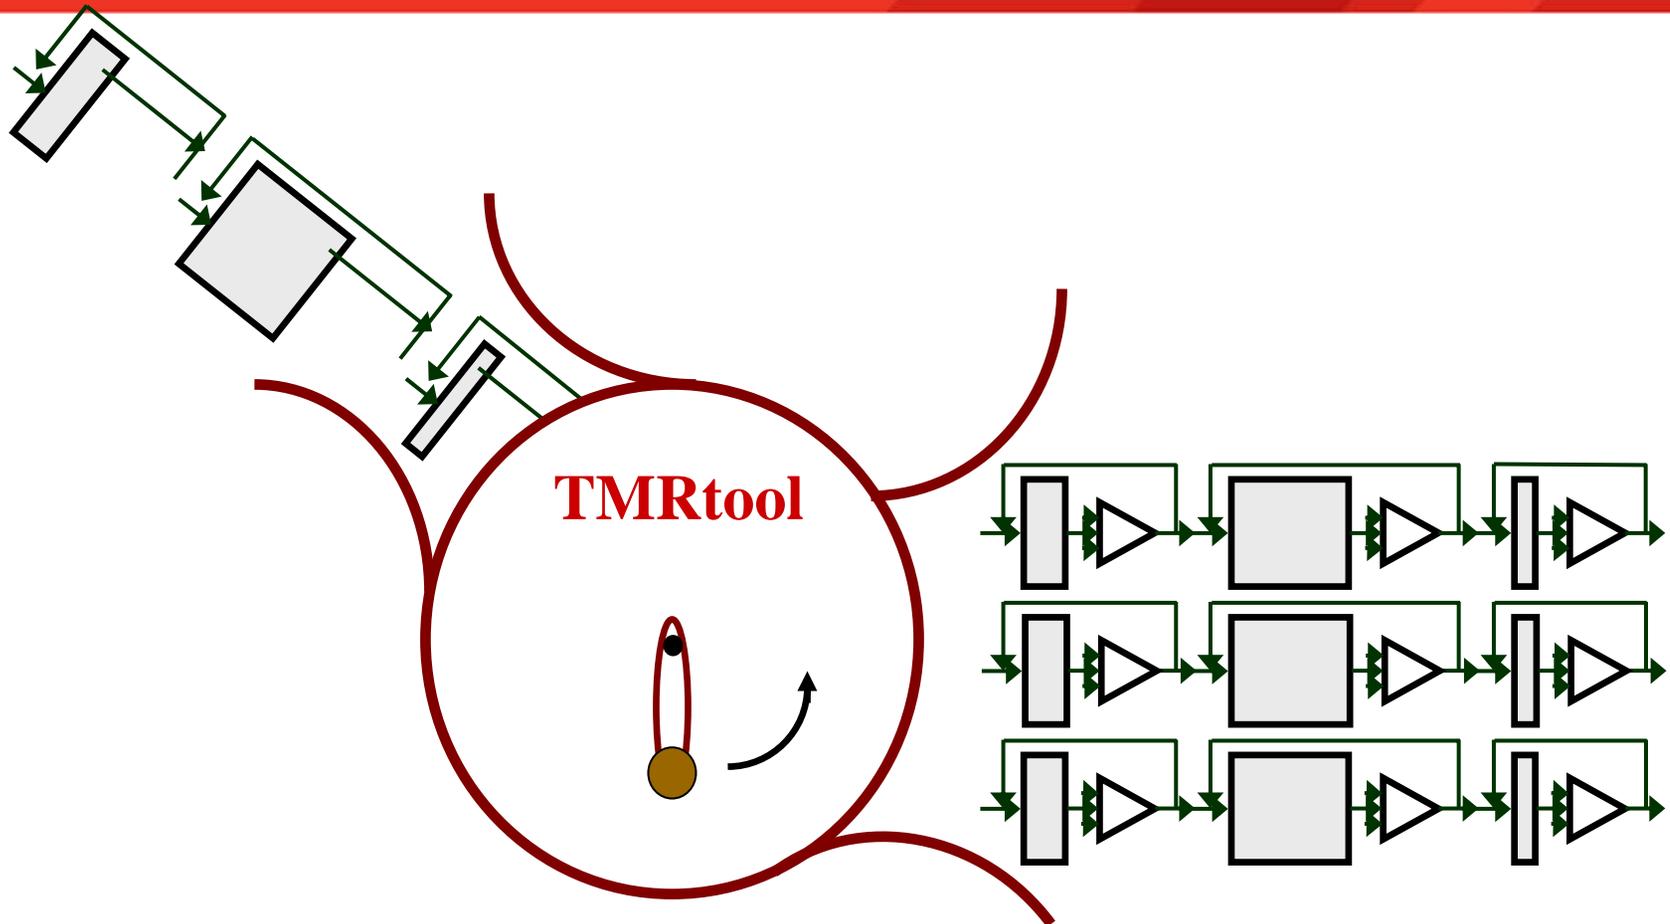


Multiple upsets but no error



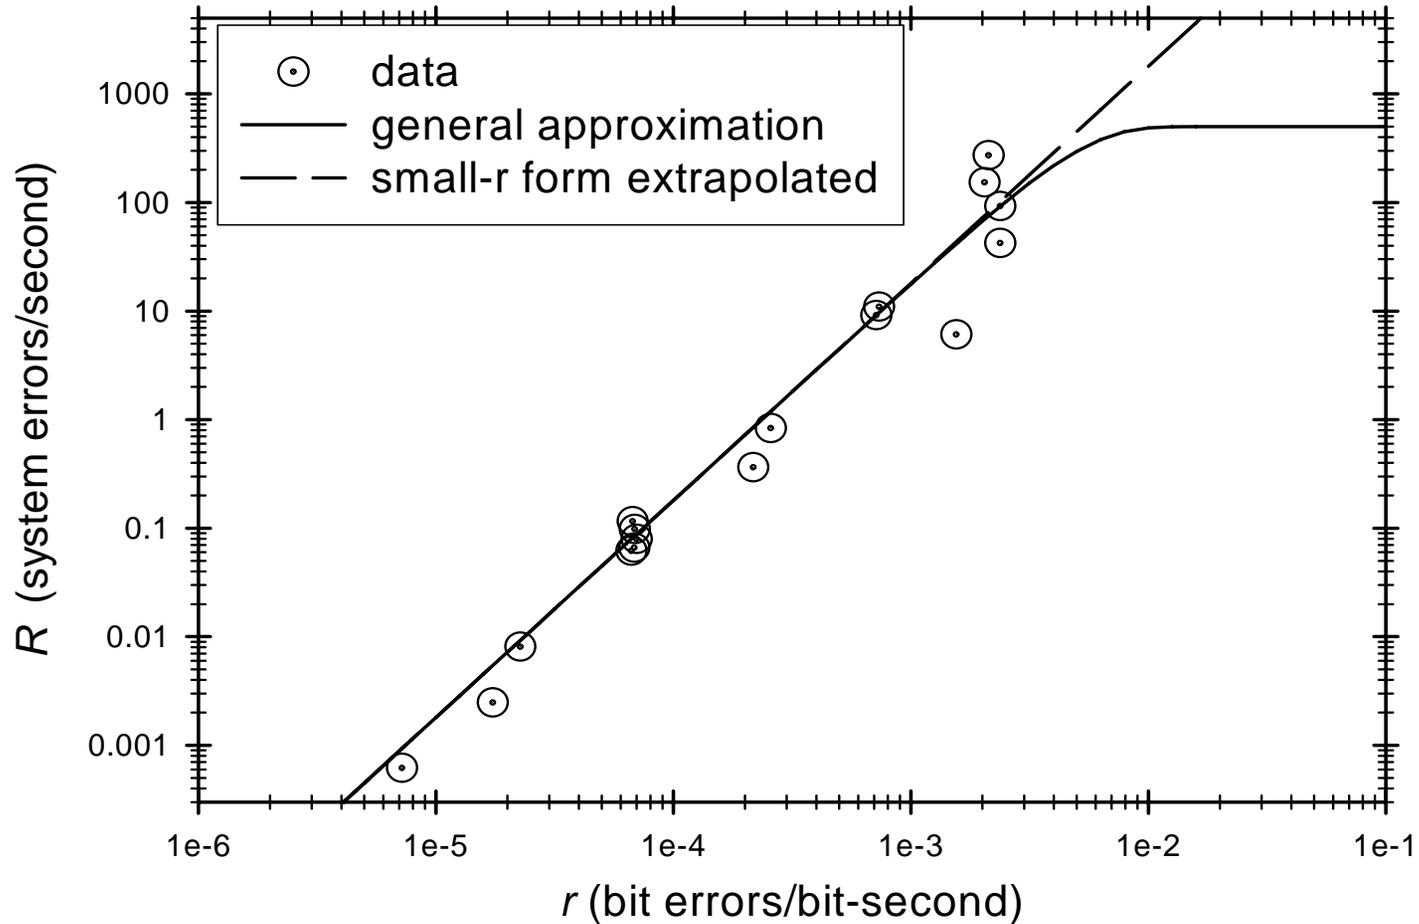
Error propagation requires upsets in two parallel modules (within a scrub cycle).

Designer's TMR "Burden"



Run the working single-string design through the TMRtool to obtain the correct Xilinx-style triplicated and voted design.

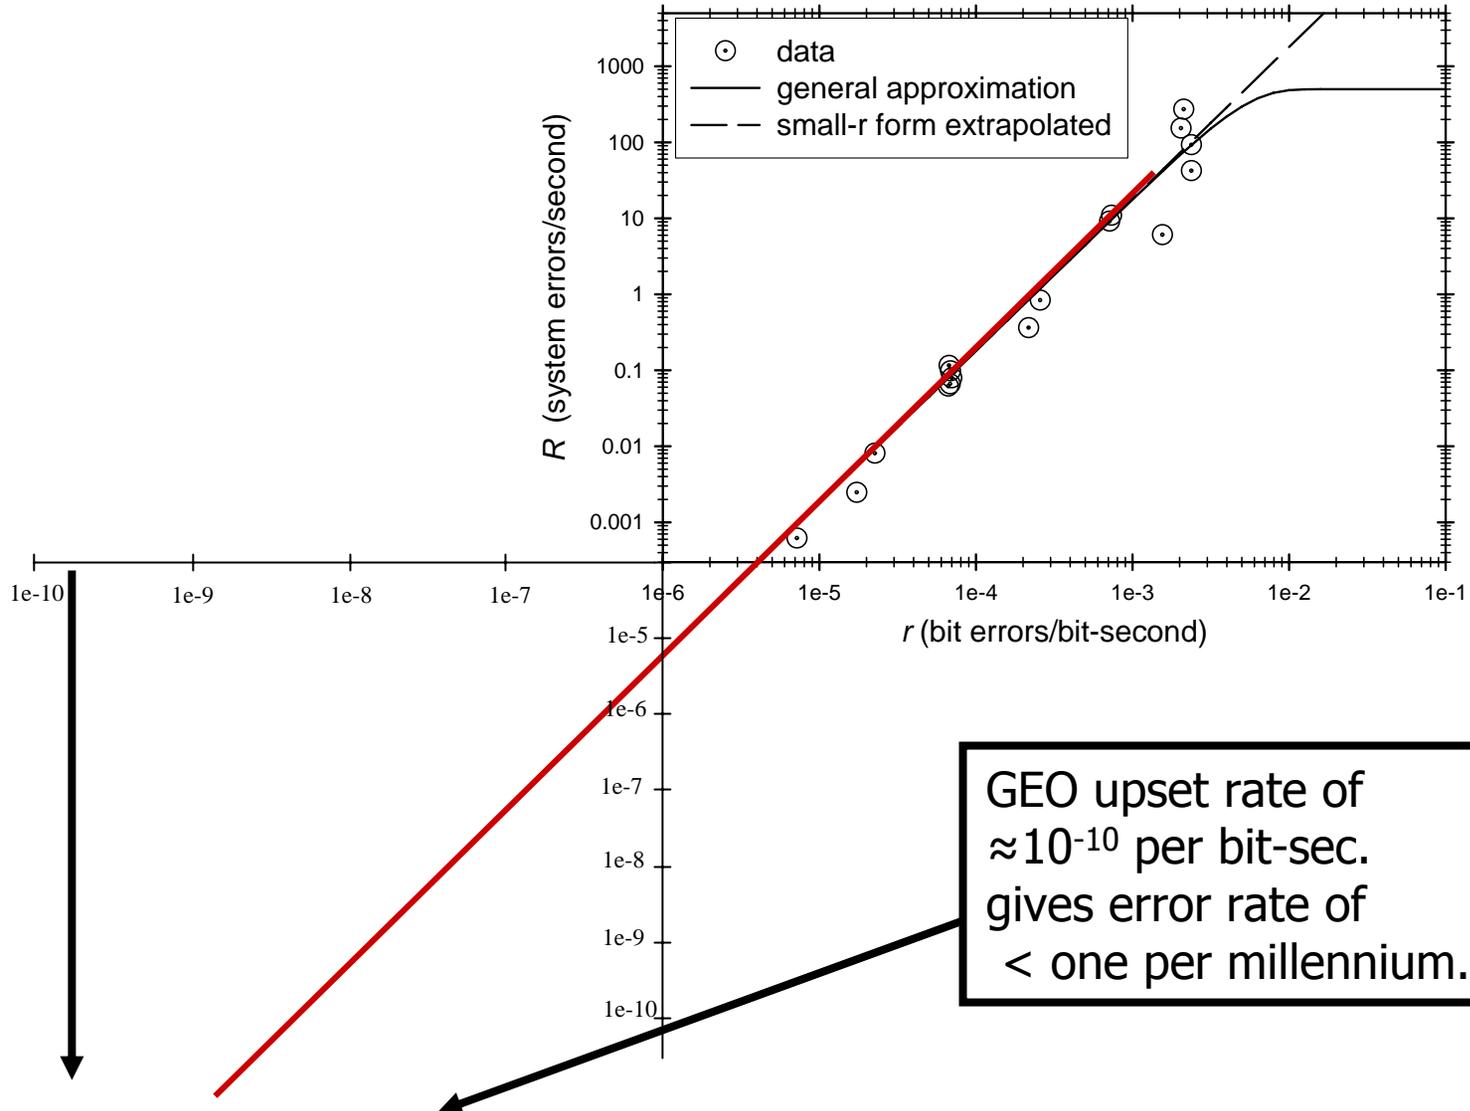
Example App - XQR2V6000 BRAM Scrubber



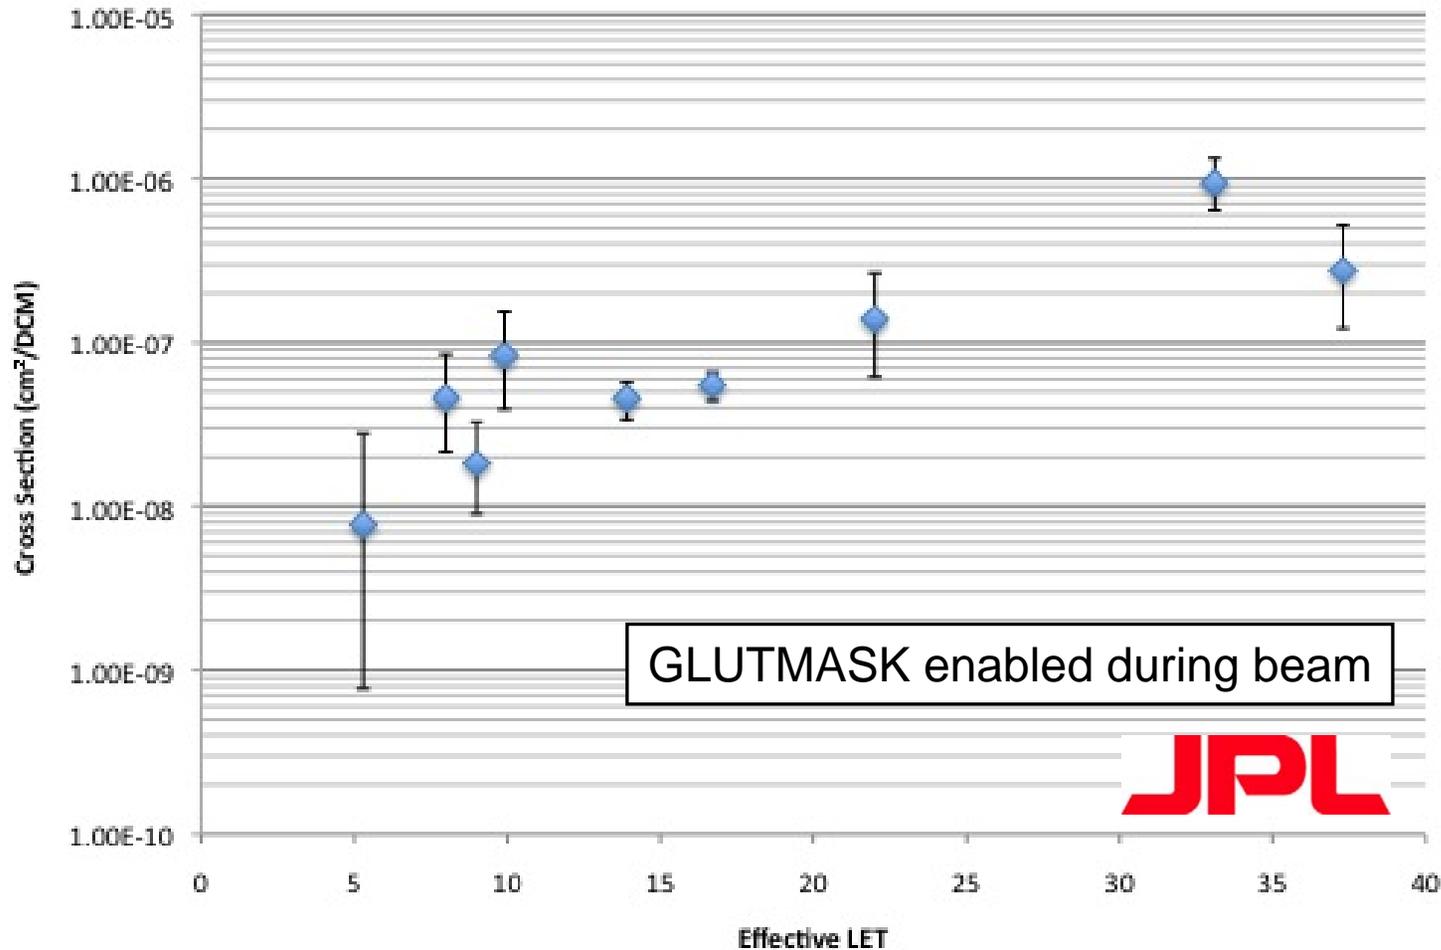
Given parameters: $T=2$ ms, $M=48000$

Fit parameter: $M2=M3=M4= 250$

Extrapolating to Space Rates



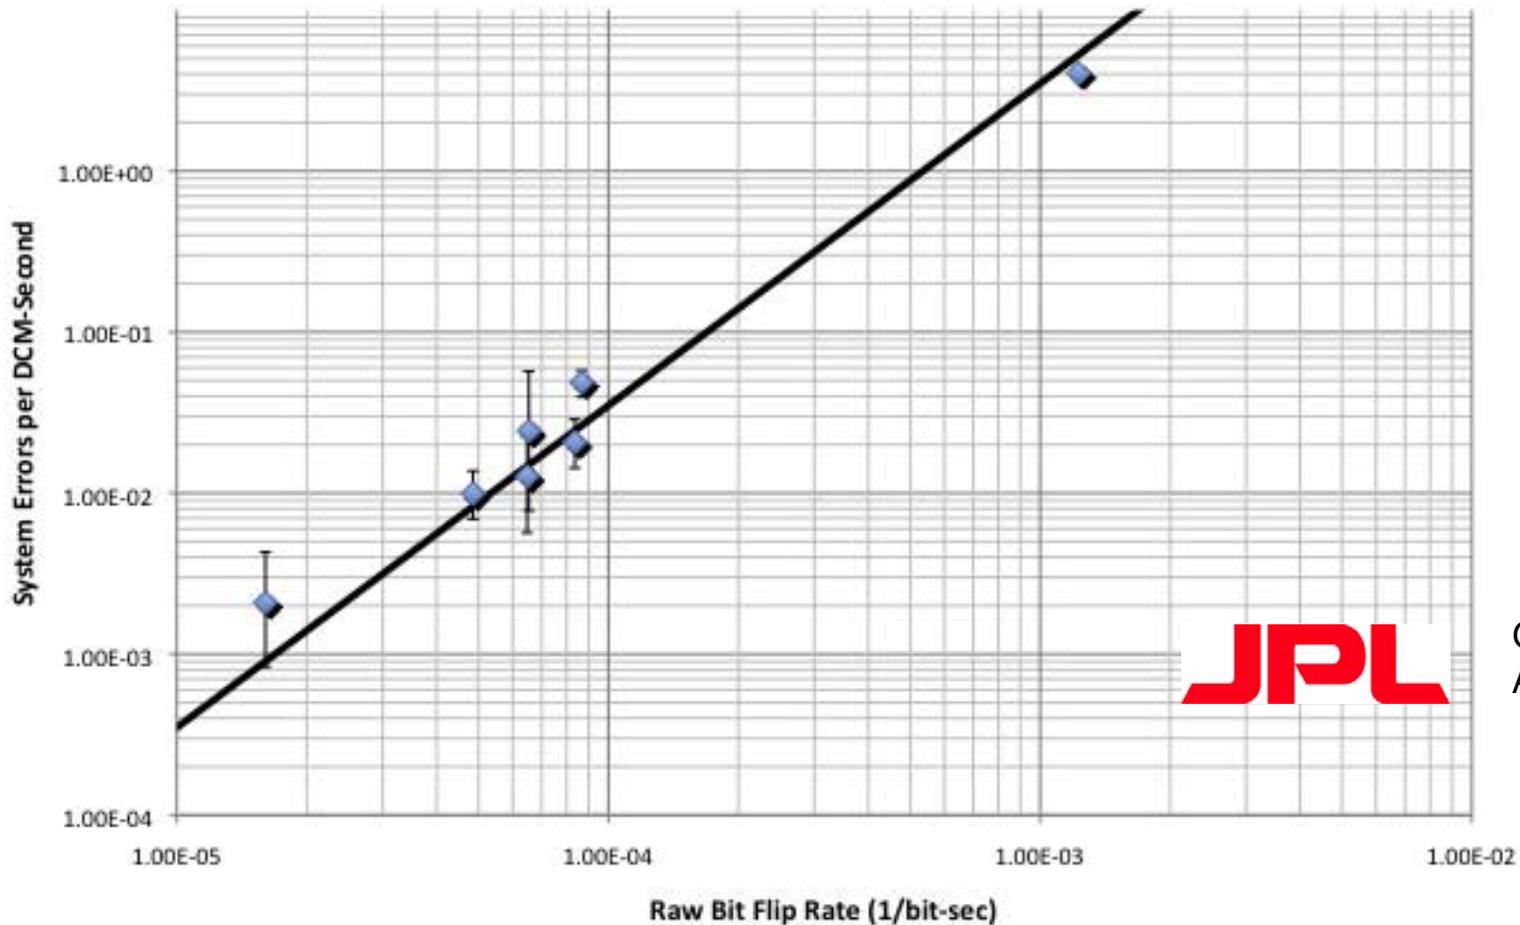
V4 DCM Dynamic Results



Greg Allen

All DCM fails fixed by either DCM reset, re-writing settings through the DRP, or scrubbing with GLUTMASK disabled.

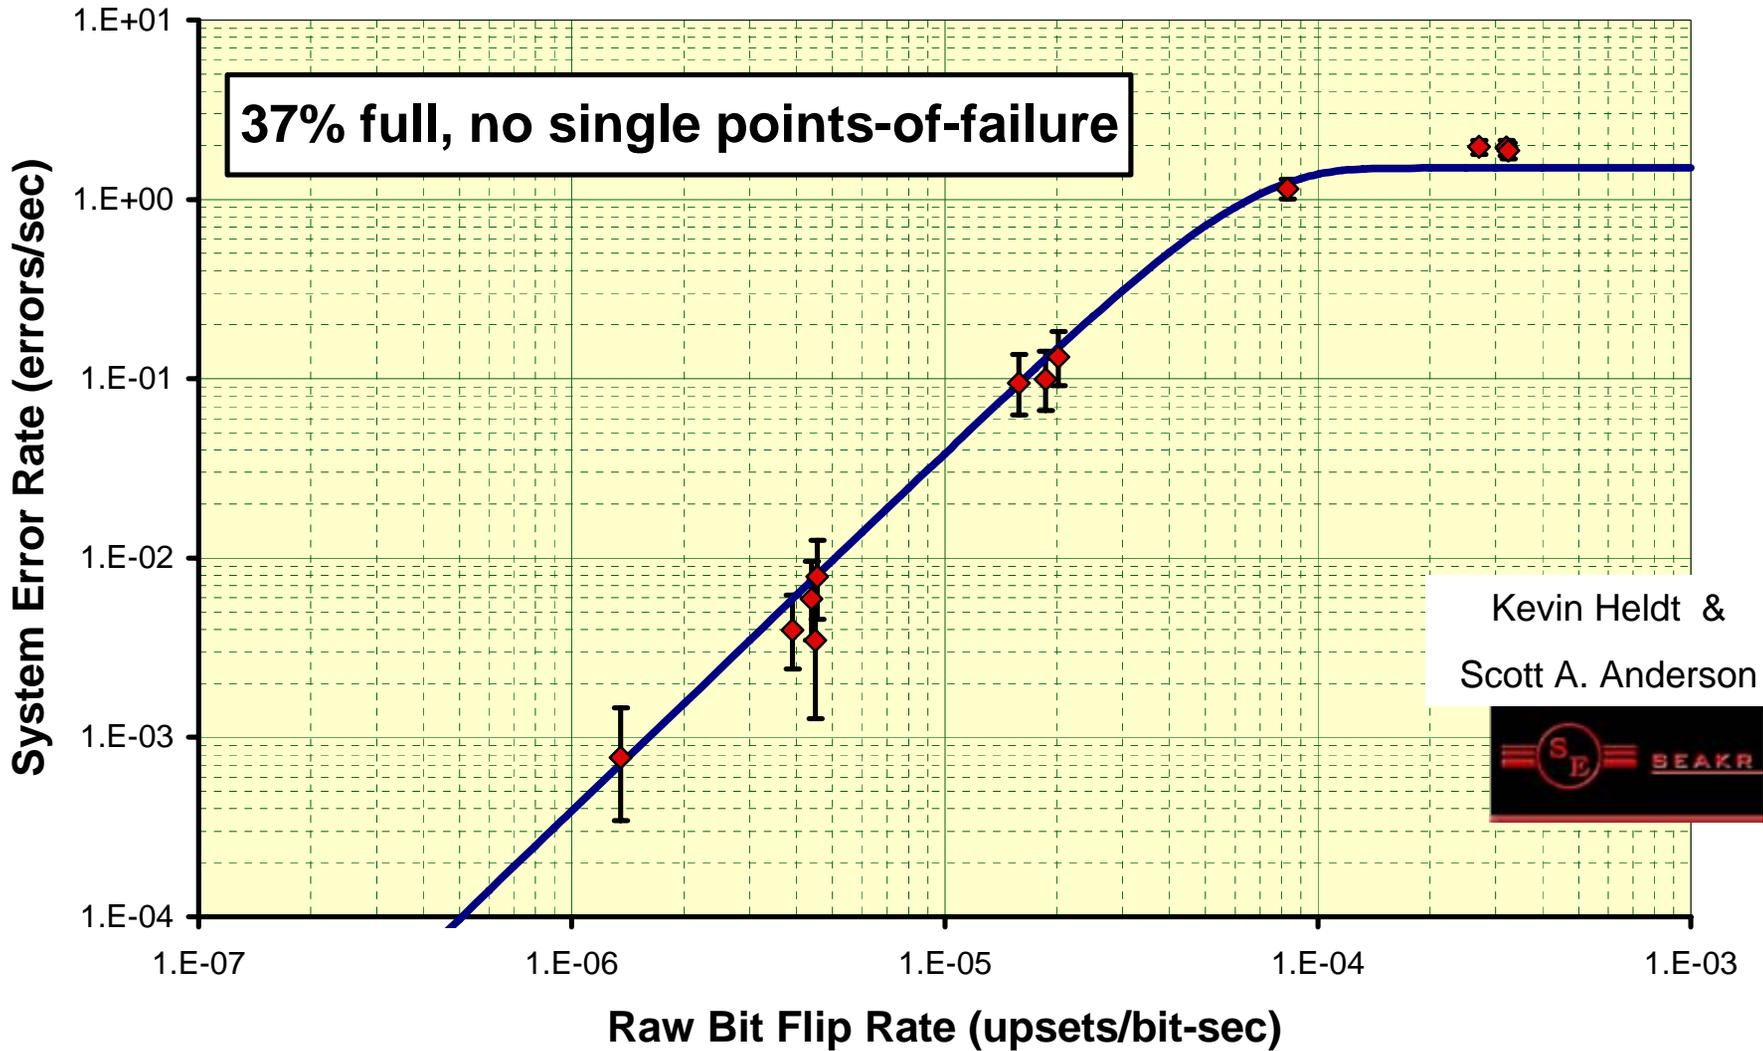
V4 DCM Mitigation Results



Greg Allen

“per DCM” means “per DCM triplet”

V4-QV TMR-Counter Results

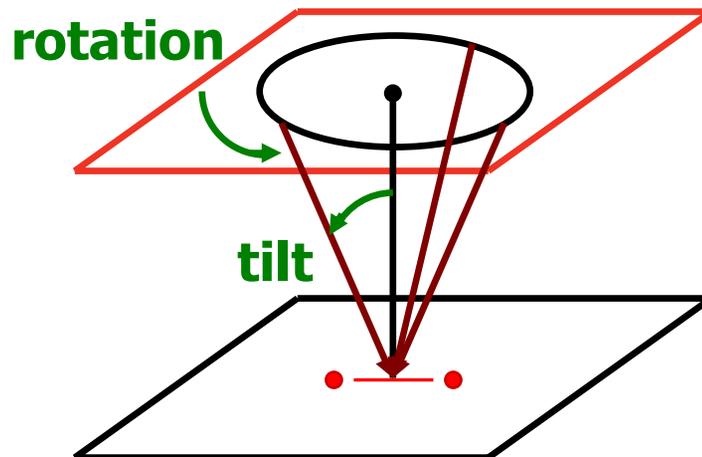


Geometrical RHBD is *two-node* problem

- To upset a cell requires some charge collection at *both* of a pair of nodes, that is,
if one node collects no charge, the cell will not upset no matter how much charge is collected at the other of the pair.
- A cell may contain one or several such pairs, but the two nodes of a given pair must be as widely separated as possible.

Two-node case makes rotation important

- The more an ion trajectory aligns with the line defined by the two nodes, the more likely it is to be able to cause an upset:



For a given tilt, different rotation angles give more or less alignment with line through the nodes.

Model to Guide Data Fits

Model necessary because ‘brute force’ :

requires too much data.

needs extrapolation to impossible tilts (90°).

Model assumes existence of a charge collection efficiency function with ellipsoidal volumes (like rounded RPPs).

Many (8) fitting parameters in current model:

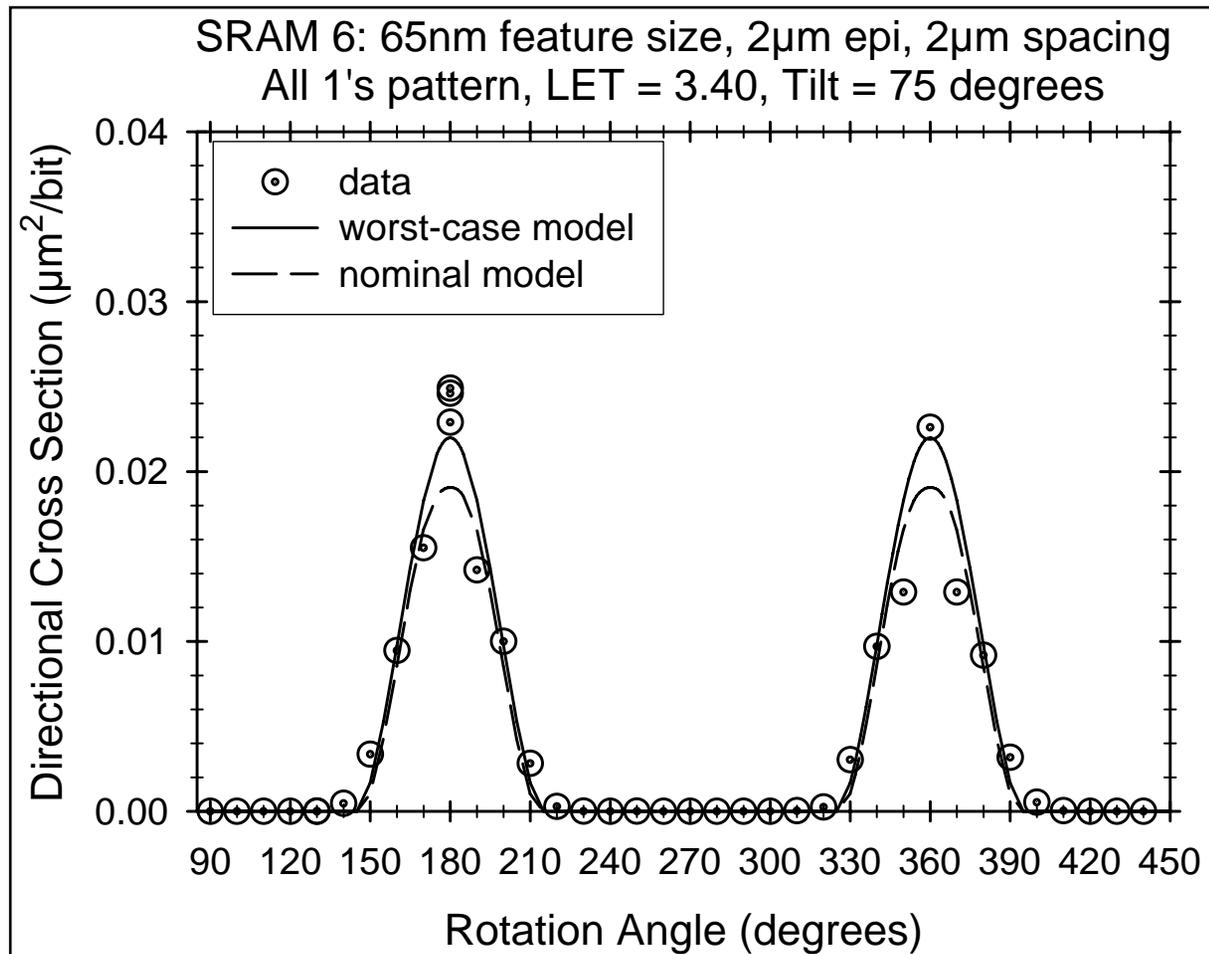
two (A, B) relate to ellipsoid shape

four – LET threshold and sat. cross section per node

plus two others

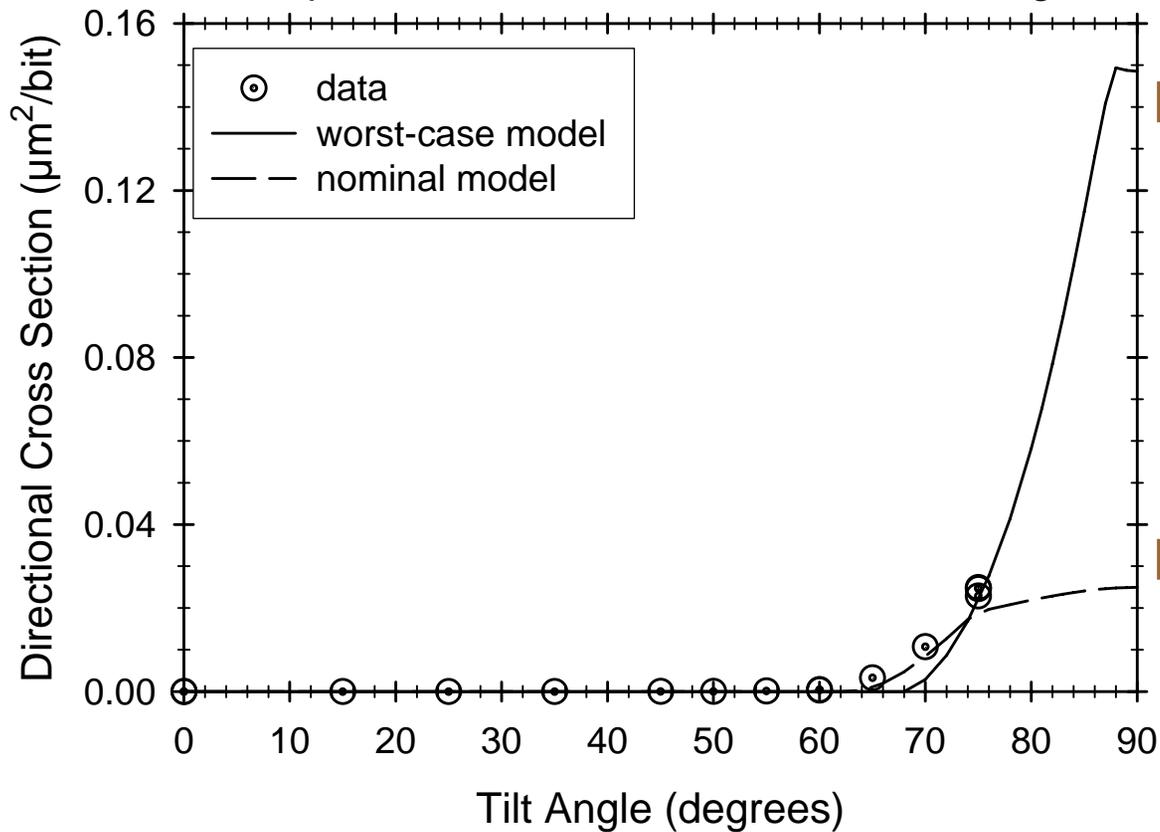
Directional Upset Response

Clearly there is a strong dependence on rotation angle:



Necessary Extrapolation

SRAM 6: 65nm feature size, 2 μ m epi, 2 μ m spacing
All 1's pattern, LET = 3.40, Rotation = 180 degrees



Note
factor of
4 or 5

Model Results

GEO rate for ones is $<9E-10$ upsets per bit-day.

GEO rate for zeros is $<9E-11$ upsets per bit-day.

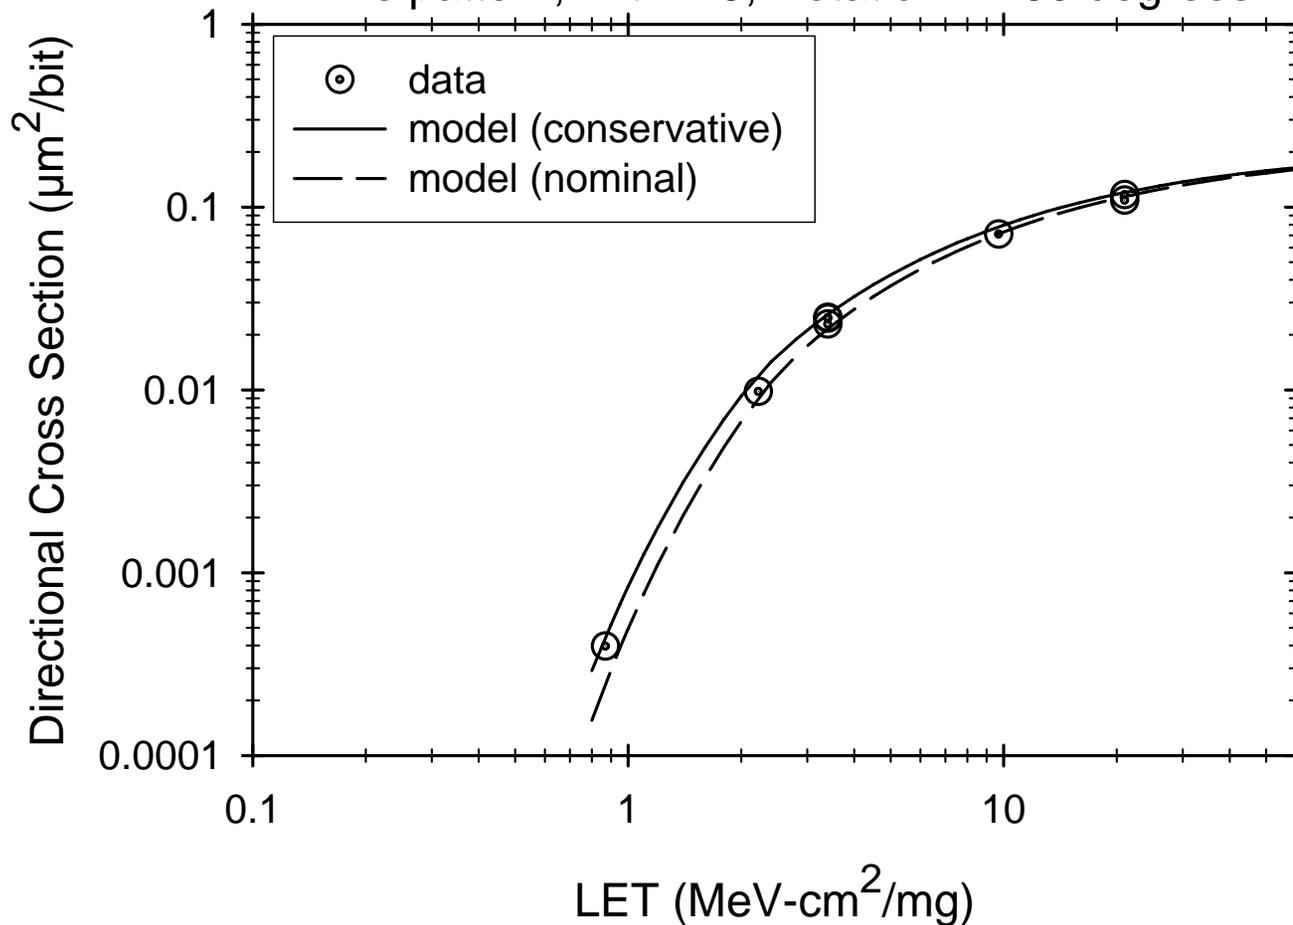
Typical design has more than 90% zeros and takes about ten (or more) upsets to cause an error:

**GEO rate for typical design is
 $<2E-11$ errors per bit-day
or approx. one every 2 years.**

Data & Model for an LET Sweep

Good agreement at worst rotation:

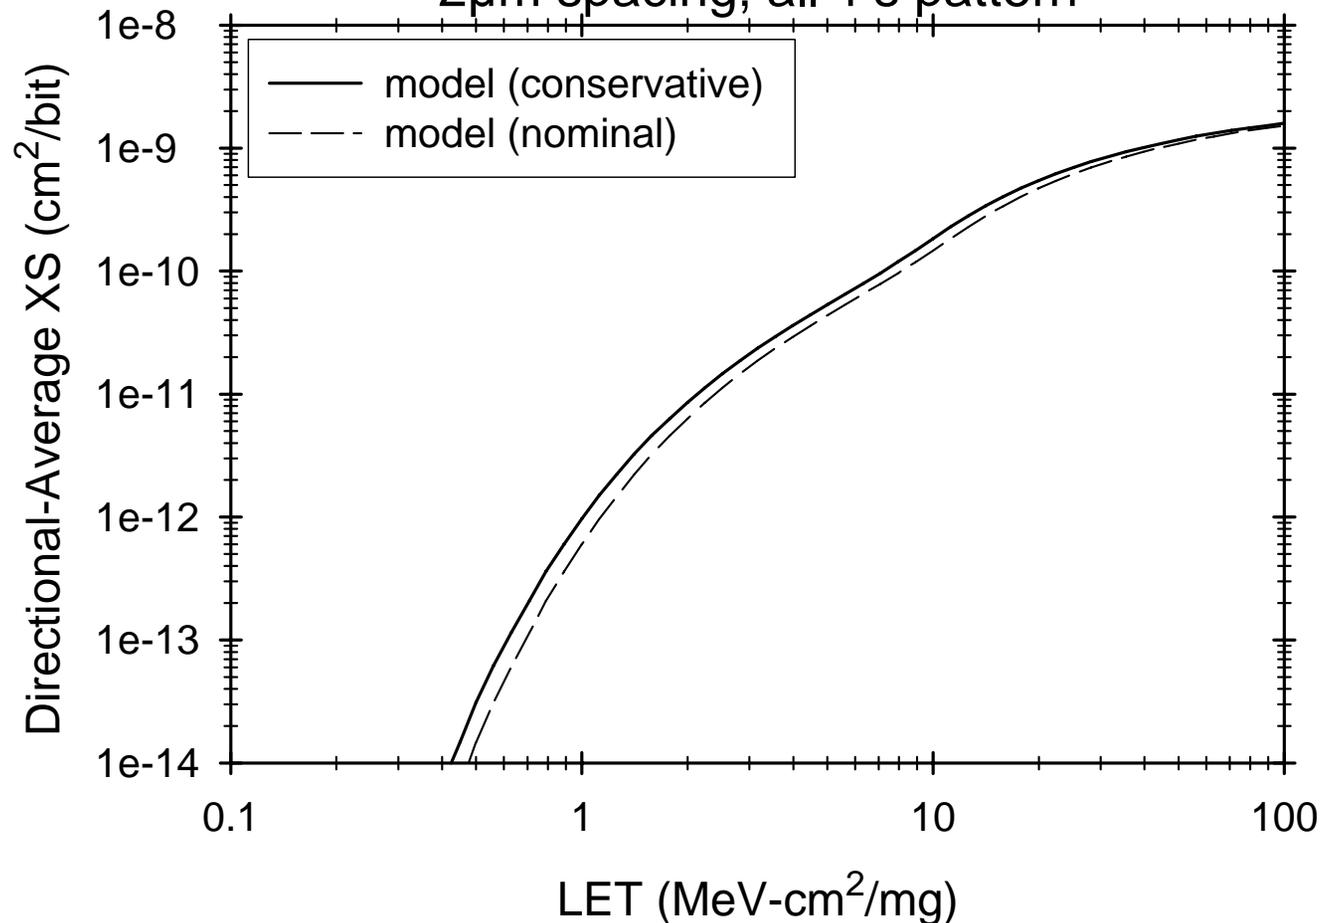
SRAM 6: 65nm feature size, 2 μ m epi, 2 μ m spacing
All 1's pattern, Tilt = 75, Rotation = 180 degrees



Average Cross Sections

... are useful for 'estimating' rates via standard calculation

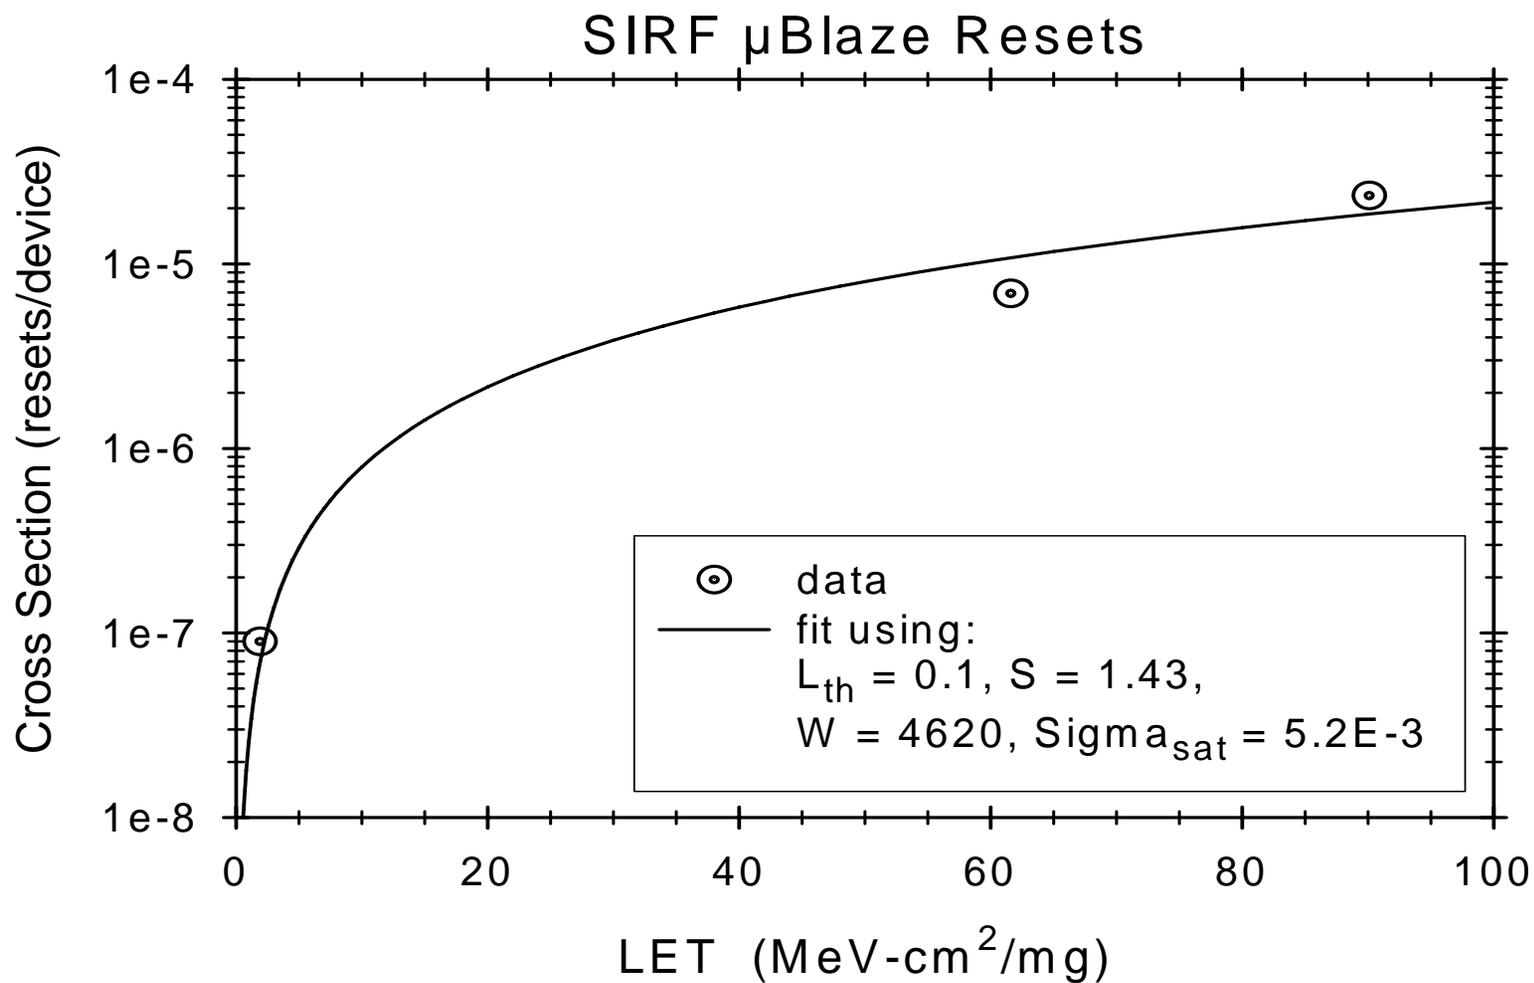
SRAM6: 65nm feature size, 2 μ m epi,
2 μ m spacing, all 1's pattern



Preliminary Results

Energy (MeV/u)	Ion	Eff. LET (Mev-cm ² /mg)	Flux (ions/cm ² /s)	Fluence (ions/cm ²)	Resets (events/device)	Runaways (events/device)
15	Au	90.1	1.340E+03	7.271E+05	17	1* - due to SEFI
24.8	Xe	61.6	9.870E+03	5.500E+06	38	0
24.8	Ne	1.9	1.000E+05	2.000E+08	18	0

Preliminary Results – Resets



Note : Weibull Fit is just a guide for the eye

MicroBlaze Results

TMRed in V4

Analysis in progress

Mike
Pratt



Theoretically, TMRed MicroBlaze in V4 will extrapolate to a lower system error rate in space than single-string in RHBD V5, but SEFI performance makes RHBD V5 better overall.

Maximum Robustness Conclusion

- **Single-FPGA design robustness is limited by the SEFI rate:**
 - Approx. 1 per century in GEO for V4
 - Approx. 1 per 100 centuries in GEO for RHBD V5
- **Properly TMRed Virtex 4-QV designs, i.e. having no single points of failure, extrapolate to an upset-induced system error rate lower than the SEFI rate**
 - 100 bits that are single points of failure translate to a system error rate of about one per century in GEO
- **Not good enough? Fly through SEFIs by using three FPGAs**
 - See XAPP987

**Assuming a SEFI outage lasts one second,
then one per century is better than 10 nines of availability.**

Conclusions - RHBD vs. TMR

Both can yield good system robustness.

- - TMR
 - Requires designer involvement
 - Costs times 3+ in gates and power
 - Extrapolation required for space error rates
 - RHBD
 - Transparent to the designer
 - Requires extra silicon area
 - Extrapolation required for space error rates
 - Potentially more robust in “extreme” environments

BACKUP MATERIAL

Processor Upset Rates – Severe Environment

for JPL Design Case Flare (DCF):

	<i>Upset Rate</i>	<i>Ratio</i>
BAE RAD750 (estimated)	6.6*	x3
Maxwell SCS750	0.36**	÷ 6
Virtex II-Pro ePPC405	85*	x40

per flare

* Upsets from heavy ions only; proton upsets insignificant or neglected

** Includes 0.14 /flare from protons

Notes:

Assumes 100% duty cycle on all bits (registers and L1 caches)

Environment ≅ actual events in October 1989 and January 1972

Shielding = 100 mil Aluminum-equivalent

TMR System Model

