First Results from Radiation Tolerance Tests on Xilinx Virtex 5

A. Aloisio[1], V. Bocci[2], R. Giordano[1]

1 Physics Dept. - University of Napoli “Federico II”
and INFN Sezione di Napoli, Italy
2 INFN Sezione di Roma 1, Rome, Italy

email: aloisio@na.infn.it, bocci@roma1.infn.it, rgiordano@na.infn.it

Outline

- FPGAs on-detector?
- Benefits of homogeneous FPGA<->FPGA links
- Xilinx Virtex-5 LXT50T facts
- Test bench architecture
- Test facility and conditions
- Results
- Conclusions

Can We Use FPGAs On-Detector ?

- S-RAM FPGAs traditionally excluded from radiation areas
  - Configuration (stored in static RAM) is sensitive to single event upsets (SEUs) => bit-flips
  - A bit flip in the configuration memory can change design functionality
- Mid-range Xilinx FPGAs include tools for configuration error detection and correction

What if FPGAs Could be Used

- The fast link sub-system will be dramatically simplified: we would only have symmetrical links! (FPGA<->FPGA)
- No constraints imposed by line encodings of stand-alone SerDeses (i.e. start/stop bits of DS92LV18 and 8b10b coding TLK2711-A)
- Just one type of links, protocol and line coding customized to ETD requirements
- Fixed-latency proof and thoroughly tested
- Artix-7: cheap Xilinx FPGAs (~ 40$) with ~ 8 embedded SerDeses (GTPs)
  - More than one link per chip
  - Data-rates up to 3 Gbps on all links (including FCTS)
  - (actually SerDes could go even faster : 6.6 Gbps)
  - # of links could be halved or better (when not driven by topology)

Device

- Xilinx Virtex 5 LX50T
  - Includes high-speed SerDeses
  - Mid-range size, yet large enough to fit our needs
  - Tested by Xilinx, NASA, Lawrence Berkeley Lab. and Los Alamos Lab.
  - Embeds configuration CRC with ECC blocks and partial reconfiguration capabilities
  - Two versions: rad-hard (70k$) and industrial (400$) (we focus on the latter, guess why)

Device Facts

- Configuration size 11.37 Mbits
  - CLBs&routing ~ 9 Mbit
  - IOB, DSP, BRAM-interconnect ~ 2.37 Mbit
- Clock cycles per CRC Readback 355,190
- Readback time 7.1 ms (at 50 MHz)

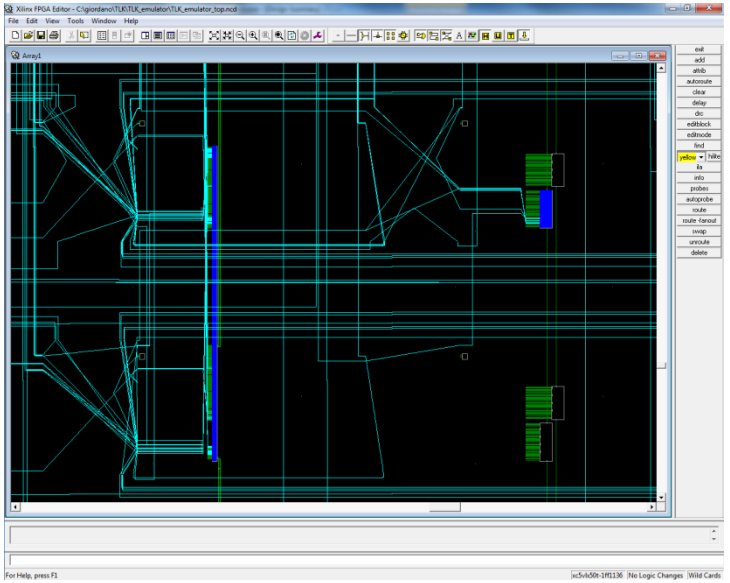

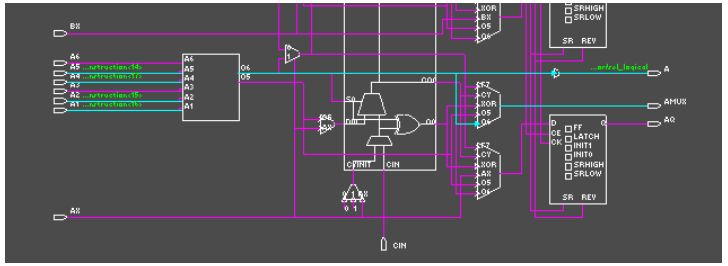

Benchmark Design

- Serial Link running at 2 Gbps (compatible with the TLK2711-A)
  - 16-bit parallel words (18-bit including control bits)
  - 100 MHz parallel clock
  - 8b10b encoding
  - Explicit lock flag
- Dummy Logic on Tx and Rx parallel data-path to observe realistic SEU effects
  - 2x 16 levels of 18-bit registers and combinational logic
- Implemented with Precision Hi-Rel synthesizer
  - Two firmware versions: one with and one without Triple Module Redundancy moderation techniques
- Resource Occupation
  - GTPs : 1 (10%)
  - Slices: 376 (5%)
    - FFs: 926 (3%)
    - LUTs: 980 (3%)
  - BRAM: 1 (1%)
  - PLLs: 1 (16%)
  - Clock Buffers: 6 (18%)
  - IOBs: 47 (10%)

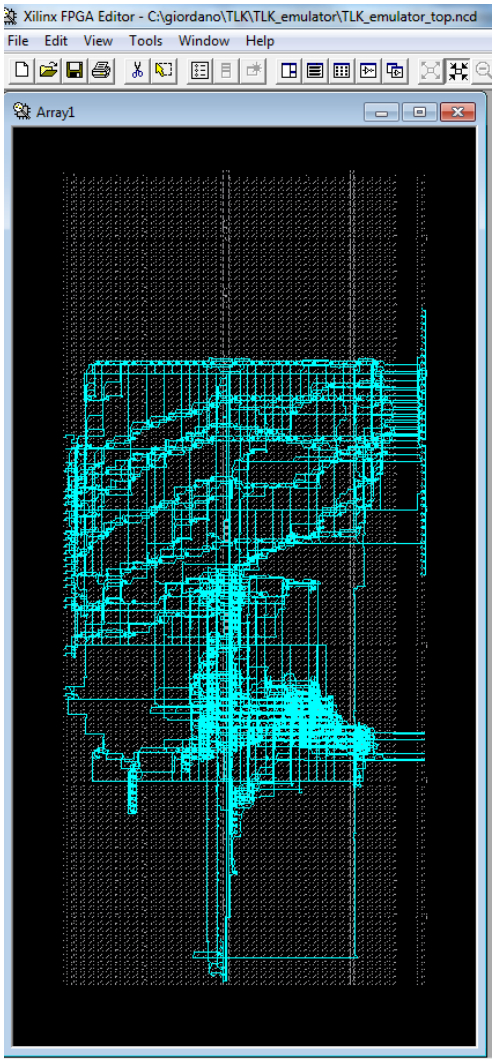

Mentor Precision Hi-Rel Synthesizer

- Aimed at applications in aerospace, medical, high-reliability and safety-critical FPGAs
- Several grades/options for Triple Module Redundancy
  - Local TMR, registers are tripled and voted are inserted
  - Distributed TMR, registers, combinatorial logic and even voters are tripled
  - Global TMR, for highest reliability, see next slide
- Can use power posts of V5 instead of half-latches as a source for logic constants (logic '0' and '1')

Local TMR

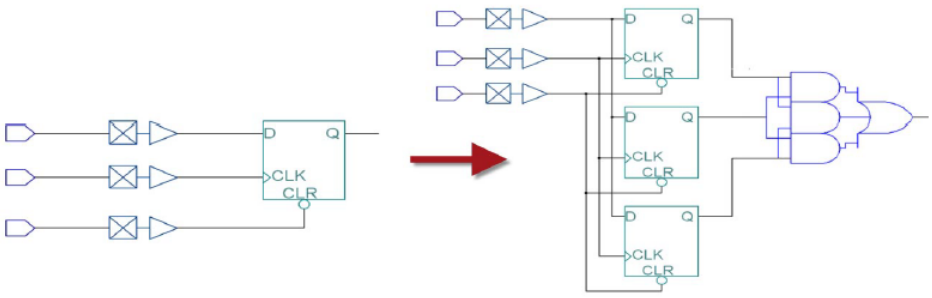

Distributed TMR

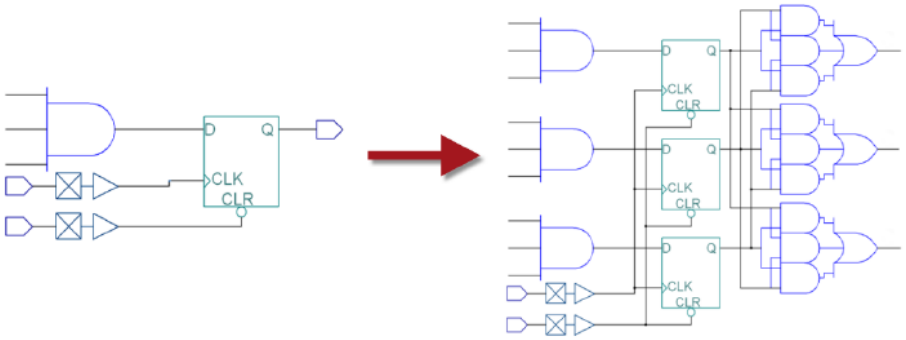

Global Triple Module Redundancy

- We adopted the GTMR option (recommended for S-RAM FPGAs)
  - Combinatorial elements, registers and the pertaining clock routes are tripled and voted out
  - But, in our design, IOs have not been tripled, unfeasible
  - Embedded SerDes (GTP) & PLL not tripled
    - GTP replication would have required Serial IO triplication, unpractical
    - PLL triplication not supported by the tool
- We also used V5 dedicated powerposts

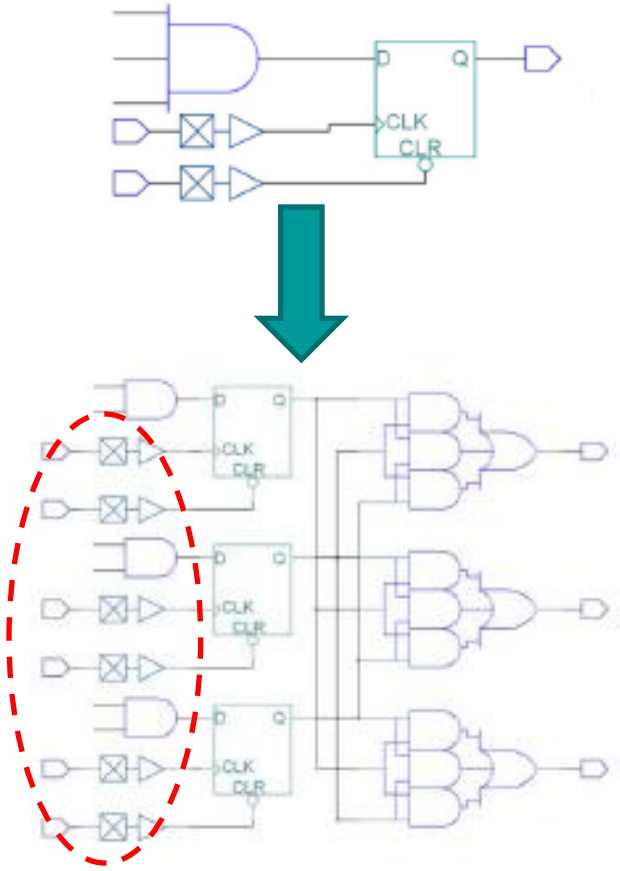

V5LX50T Test Bench

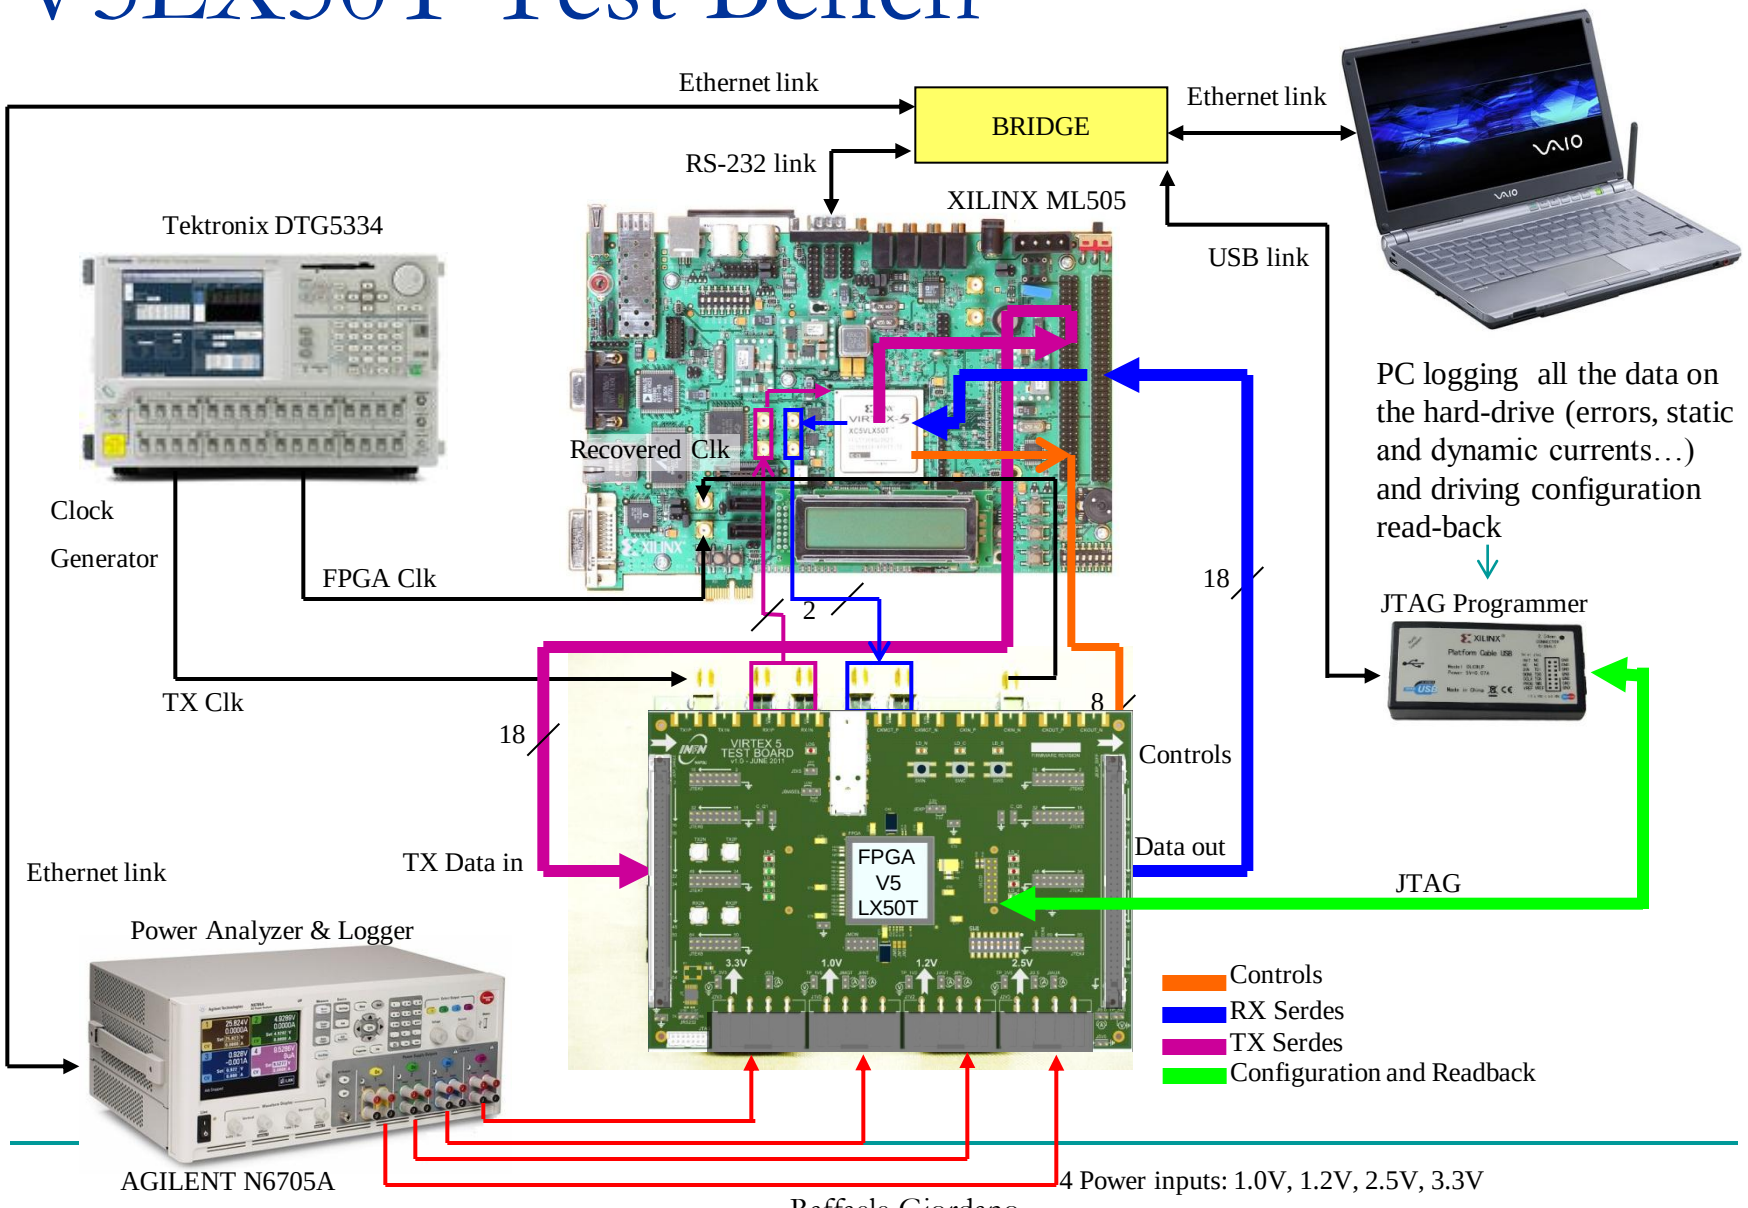

Tester Board

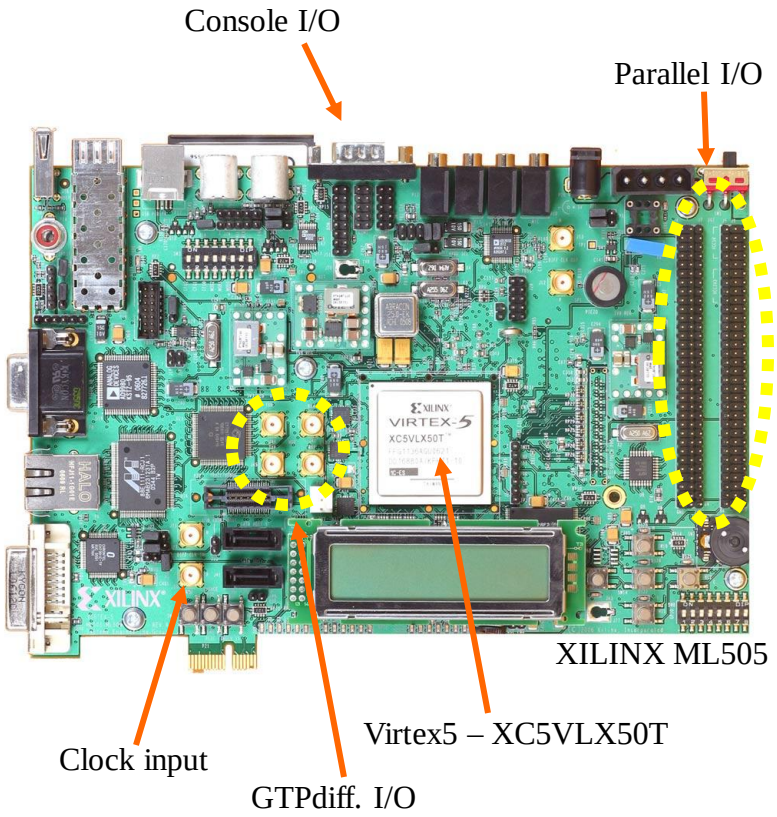

- Same FPGA as DUT board
- Data pattern stored in the FPGA firmware
- TX and RX sections of the benchmark link design are tested independently and simultaneuosly
- Console IO
  - Status and errors are logged on a console handle by an embedded microprocessor

FPGA Board for Beam Test

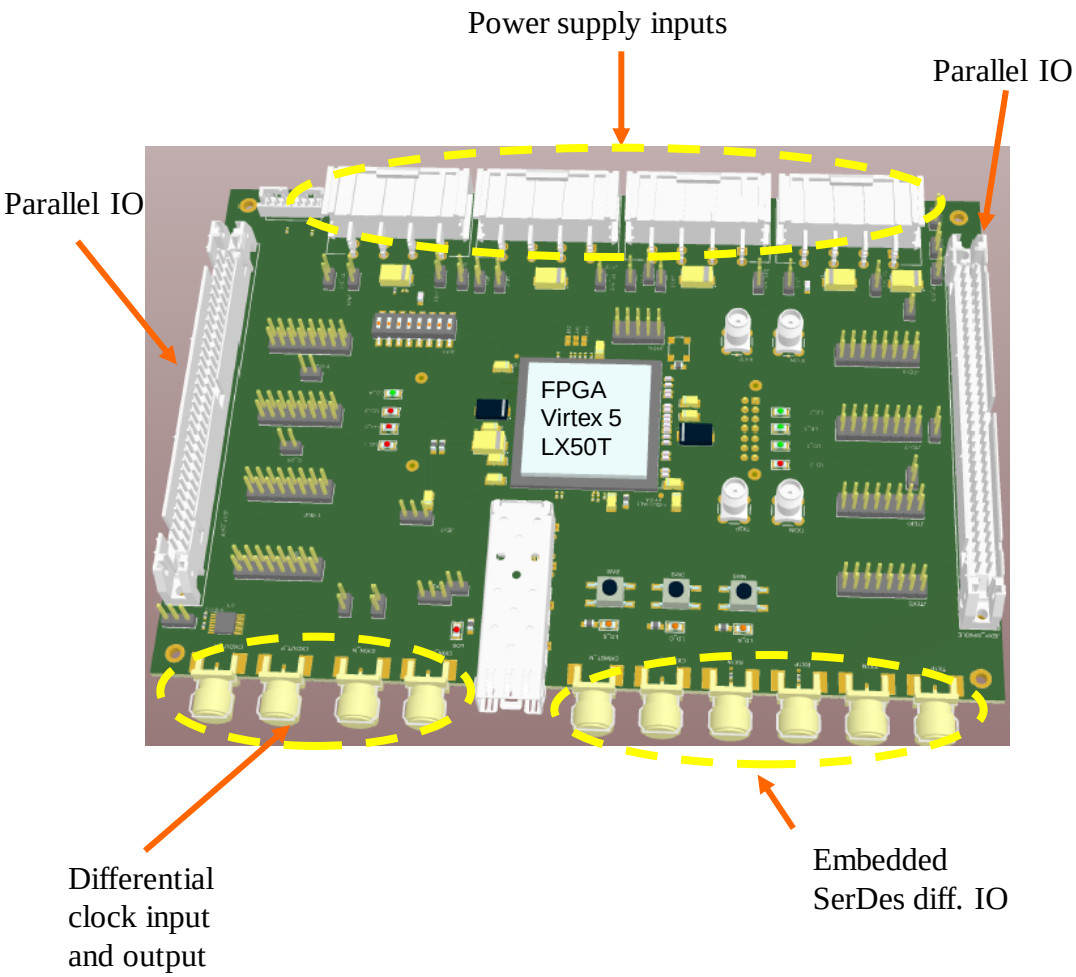

- Compatible with the Xilinx ML505 board used for the presented lab. tests (implements sub-set of features)
- Same firmware used in ML505 for SEU emulation tests (only synthesizer changed)
- SMAs for clocking and serial IO
- No active components
- 4-wires connectors for current sensing power supply
- SFP cage for optional testing of opto-electronics

LNS Facility

LNS accelerator PLAN

Experimental setup

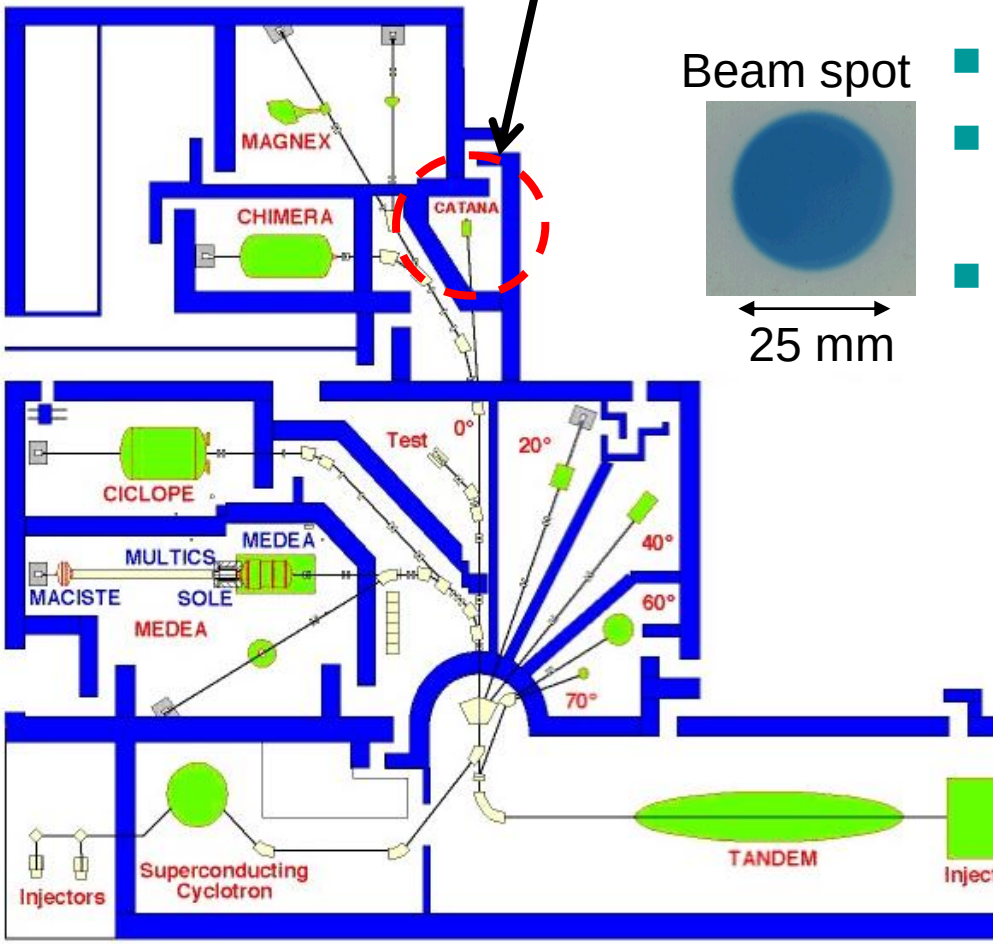

Beam spot

- Superconducting cyclotron at LNS - Catania
- 62-MeV proton beam
- CATANA beam line
- Current tunable up to 300 pA
- Uniform beam intensity on our sample

2D Beam profile

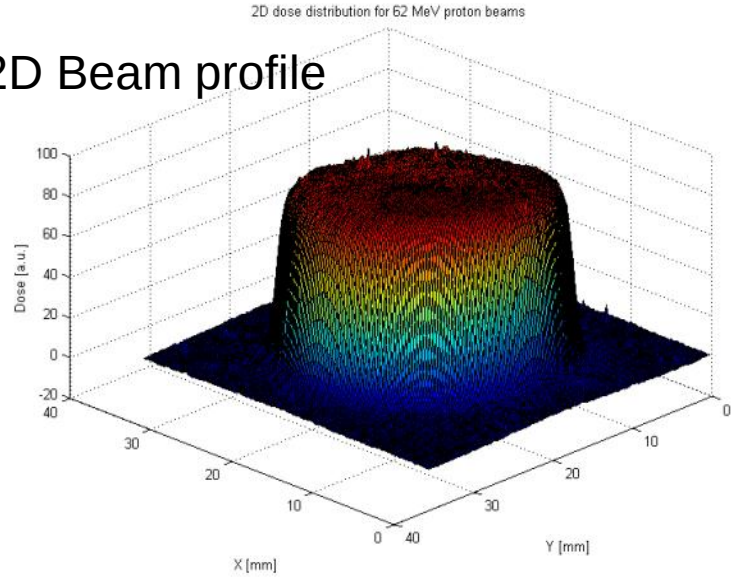

Test-setup

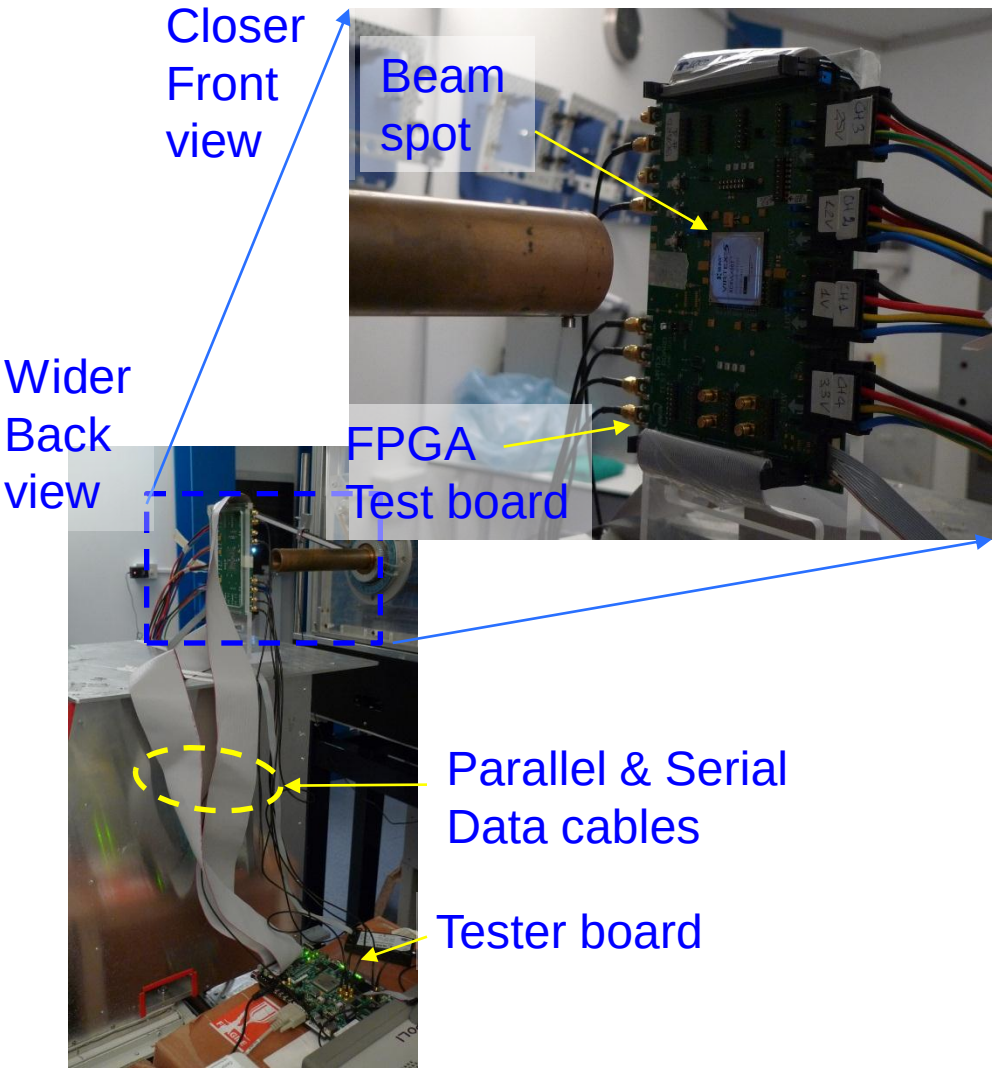

- Test conditions
  - All Vcc set to maximum allowed value (1.05V, 1.32V, 2.63V, 3.45V)
  - f_{clock} = 100 MHz
  - Data rate during test=2 Gb/s
  - Configuration cross section measurement and link failure test
  - Tests split into many runs (~ 50)
  - Configuration x-section, 6 runs, several dose rates from 1 to 20 Gy/min (Si)
  - Tested 2 link firmwares: w/ TMR and w/out TMR (will refer to them as TMR and noTMR), with different dose rates also

NoTMR Firmware: SEU Tolerance

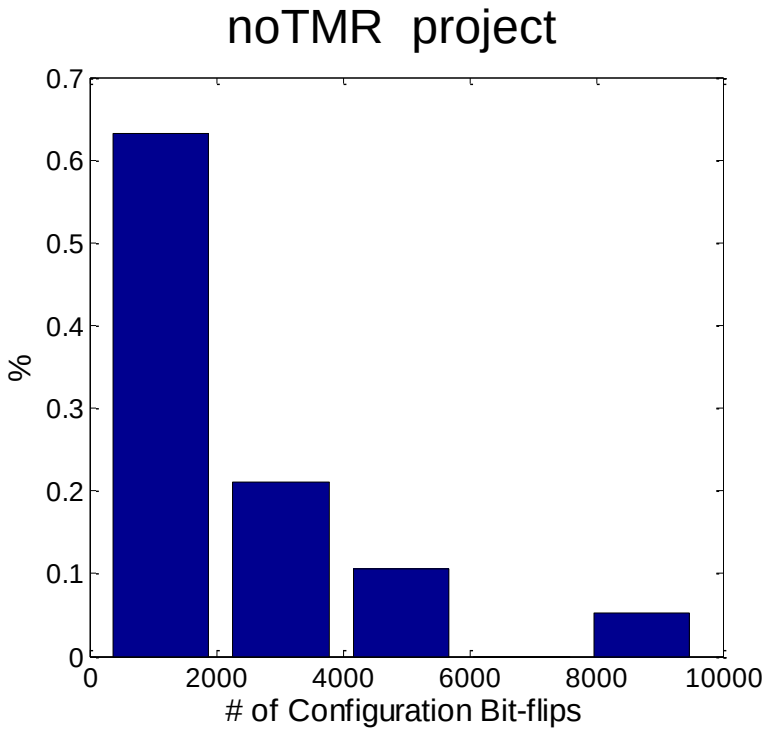

- Good match (<10%) between emulated SEU and real SEUs

- Total of 19 runs
  - 18 runs @ 6 Gy/min (Si)
  - 1 @ 23 Gy/min (Si)
  - Avg. duration 146 s => total dose per run 14 Gy (Si)
- We measured # of configuration bit-flips accumulated before link failure (error burst) (N_{flips})
- On average N_{flips} =2070
- During the SEU emulation tests performed in lab. in May, N_{flips}=2250

TMR Firmware: SEU Tolerance

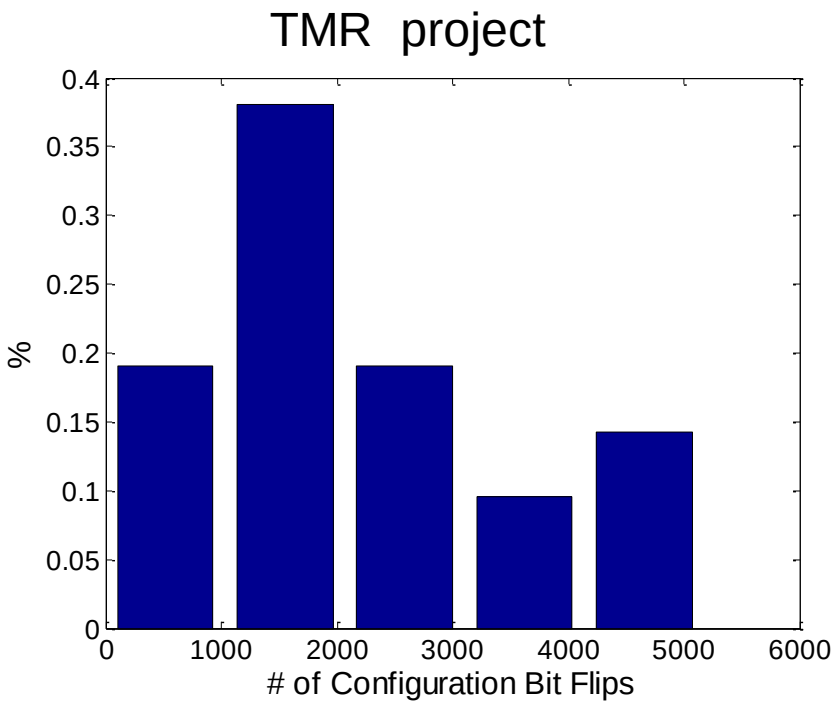

- Total of 24 runs
  - 20 runs @ 6 Gy/min (Si)
  - 4 @ at 23 Gy/min
  - Avg. duration 171 s => total dose per run 19 Gy (Si)
- On average N_{flips} =2170
- Results compatible with the noTMR firmware
- Looks like TMR did not improve SEU-tolerance
  - due to non-tripled modules? (GTP, PLL?)

Configuration Memory Test Methodology

1. Start a new run
2. Program the FPGA
3. Read-back the configuration via JTAG and compare it to the initial one
4. Log number of differences (errored bits)
5. Go back to 3, until total desired dose reached

- The test loop (pts 3,4,5) has been executed at the maximum speed permitted by JTAG, i.e. one readback every 15 seconds

Configuration Memory Error Results

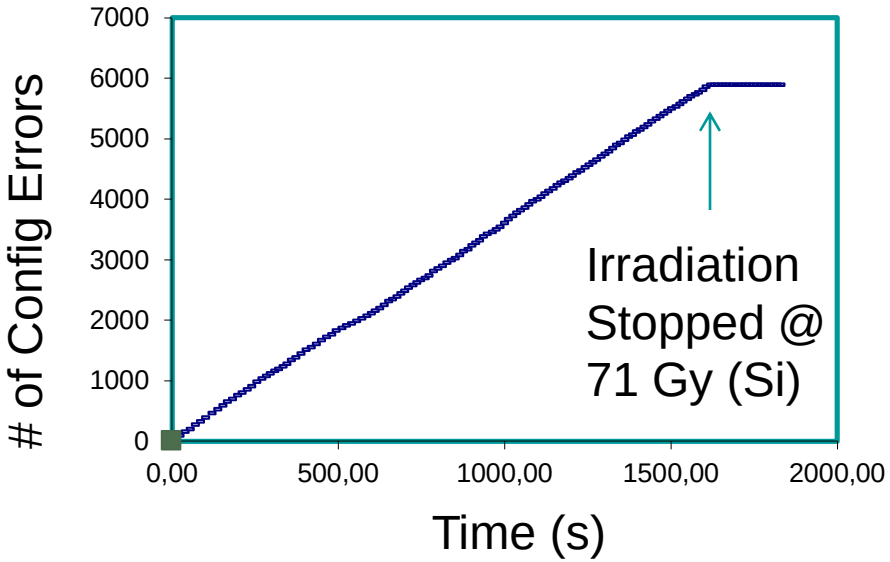

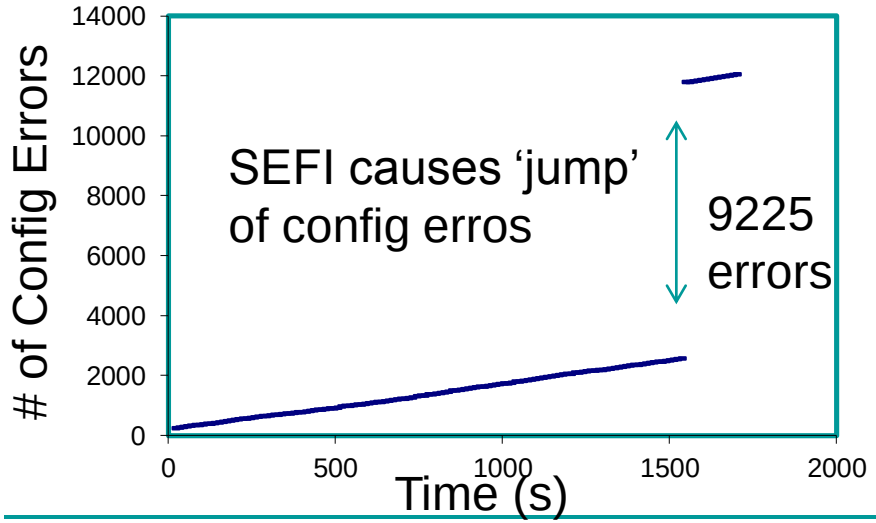

- Accumulation of bit-flips into configuration memory, FPGA working like a «dosimeter»
- Experienced a few SEFIs, a SEU causes failure of readback circuitry => 1 SEU generates many (thousands) configuration errors
- Abrupt jumps of config errors

Configuration Bit-flip Cross-section

- V5LX50T device
  - Die size = 1.4 cm^2
  - p @ 62MeV
    - Measured
      $\sigma = 3{,}5 \cdot 10^{-14}\ cm^2/bit$
    - Published
      $\sigma = 6.4 \cdot 10^{-14}\ cm^2/bit$, see [1]
- Good agreement between measured and published σ (different device, facility, test conditions…)

X-ray image
Package

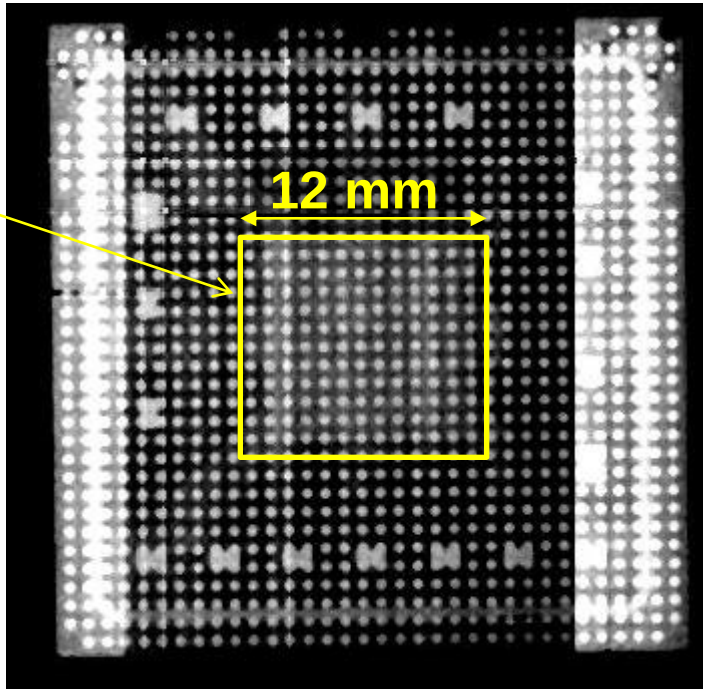

Reference:
[1] Quinn et al., Proceedings of 2007 IEEE Radiation Effetcs Data Workshop, Page(s): 177-184

Link Failure Rates

- Effective x-section to failure is design-dependent (critical bits)
- 0.5 kGy(Si)/year (a) => $4.1 \cdot 10^4$ configuration bit-flips on our design /year (measured)
- ·# of link failures/year = 20 (on average, without ANY recovery strategy)
- Expected one link failure every 18 days (includes both configuration and SEUs in configured logic)

Courtesy of Riccardo Cenci, Elba Meeting, May-Jun. 2011

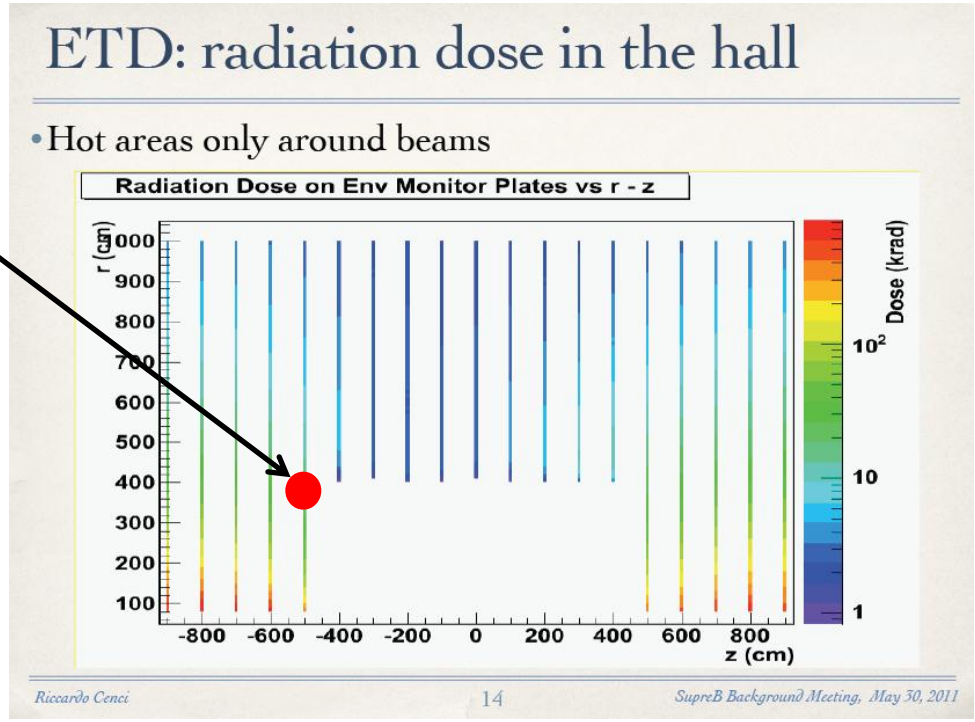

Note:
(a) Estimated by R. Cenci, INFN Pisa

Conclusions

- Very good agreement (within 10%) between test with emulated SEUs and test on beam
- On average ~ 2100 bit-flips needed to have a link failure
- Measured configuration error x-section in agreement with the expected one (within factor 2)
- Test beam data analysis to be completed:
  - Analysis of currents
  - More detailed analysis of link failures, how did it fail?
  - Investigate causes of failures, why did it fail ?
  - Why did the TMR and noTMR designs perform the same?
- Need further testing and updated info on expected radiation in detector area as soon as possible
- Next proton test beam (@LNS, 62-MeV protons) scheduled for Dec. 10th 2011
  - Will test Xilinx Virtex-5 and Virtex-6 FPGA families
- Still unclear if FPGAs are suitable to be used on-detector
- We thank all the LNS staff for the their help and support during the beam test

Back-up Slides

Quick Facts

- ρ_{Si} = 2.3 g/cm^3
- $(dE/dx)_{p@60MeV}$ in Si = 1.8 MeV/mm (or 600 keV in 300 μm)
- $\sigma = (1 / F) * n_{errors}$

- $N_{SEU} = (N_{pbeam} * \sigma_{bit} * N_{bit(eff.)}) / * S_{beam}$
- $N_{sysfail} = N_{SEU} * N_{SEU->failure}$
- 71 Gy(Si) => 5891 config. errors

Single Event Functional Interrupts

- Experienced a few BRAM SEFIs
- Run 30 - This is a known SEFI we call "BRAM SEFI" where the configuration bits (one) get flipped, and the entire BRAM gets its data inverted.
- Another BRAM SEFI is a replacement of a column with the spare column (256 bits in error).

Recovery Strategies

- Scheduled maintenance
  - Reconfigure the FPGA at regular intervals, e.g. once a day, no matter what

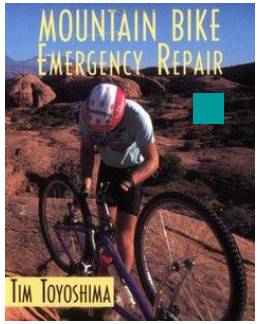

- Emergency maintenance
  - Reconfigure as soon as the service can be interrupted, e.g. exploit any reset or power-cycle of the link to reconfigure

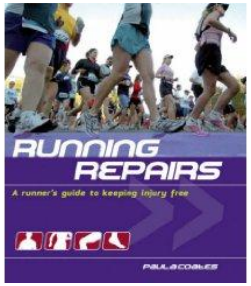

- Repair-while-running
  - Partial reconfigure as soon as the error is detected, even if the link is working => interruptio of service, i.e. dead time

Emulating Configuration SEUs

- Investigate impact of configuration SEUs on design functionality
  - Flip configuration bits by means of internal configuration access port (ICAP)
  - Optional error correction thank to integrated CRC calculator and ECC (FRAME_ECC)
- Programmable integrated controller:
  - SEU generation without correction (error accumulation)
  - SEU generation and on-the-fly correction
- Custom design derivative of a Xilinx core (so called SEU Controller)

Need to Estimate Rad-Tolerance

- Two kind of issues :
  - Total Ionizing Dose (common to all ICs)
  - SEUs
    - Logic errors
    - Configuration errors
      - Lab test (SEU Injector)
      - Beam test
- However, configuration SEUs have been largely over-estimated in the past: they rarely impact the design functionality
- In this talk we will not cover logic errors, they are common to every digital device

SEU Generation&Correction Results

- SEUs generated at 1 Hz (due to limitations of the original Xilinx SEU controller)
- Measured # of (generated&corrected) SEUs before failure
- On average ~ 400 SEUs needed to observe BERT errors
- Very likely the difference is due to reconfiguration=>even correctly working blocks are affected

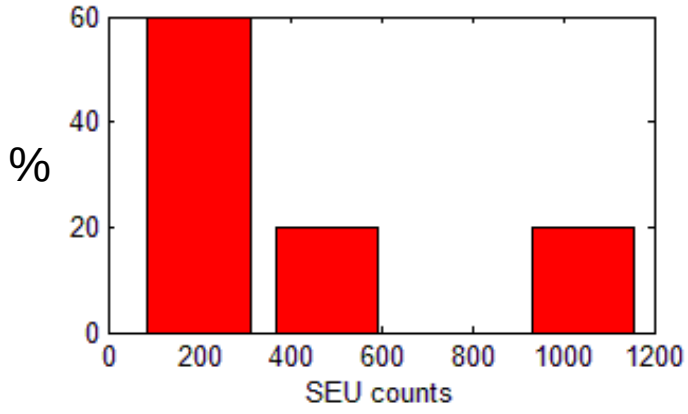

SEU Accumulation Test Results

- SEU generated at 20 Hz
- We measured # of SEUs before failure
- On average 2250 SEUs needed to observe BERT errors