


Transient Faults Effects on HPC Applications

SC17: Experimental and Analytical Study of Xeon Phi Reliability

Daniel Oliveira, Laércio Pilla, Nathan DeBardeleben, Sean Blanchard, Heather Quinn, Israel Koren, Philippe Navaux, Paolo Rech

- **Motivation**
- Radiation Experiments
- Fault Injection
- Conclusion


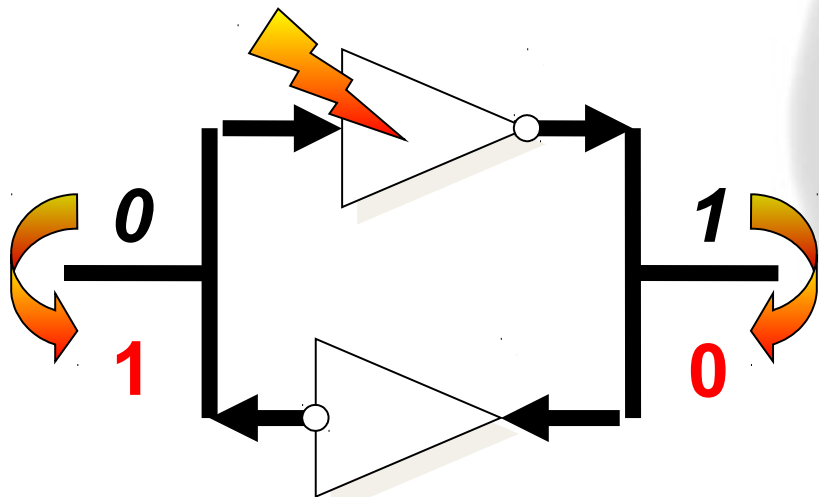
Terrestrial Radiation Environment



Galactic cosmic rays interaction with atmosphere generates **neutrons**.

13 n/(cm²*h) @sea level

Soft Errors: the device is not permanently damaged, but the particle may generate **bit-flips** or **logic errors**



HPC Reliability Importance



Follow on: [f](#) [t](#) [in](#) [+](#) [m](#)

Engineering Topics ▾ Special Reports ▾ Blogs ▾ Multimedia ▾

Advertisement 

Feature | Computing | Hardware

How To Kill A Supercomputer: Dirty Power, Cosmic Rays, and Bad Solder

Will future exascale supercomputers be able to withstand the steady onslaught of routine faults?

By Al Geist

Posted 23 Feb 2016 | 16:00 GMT



As a child, were you ever afraid that a monster lurking in your bedroom would leap out of the dark and get you? My job at Oak Ridge National



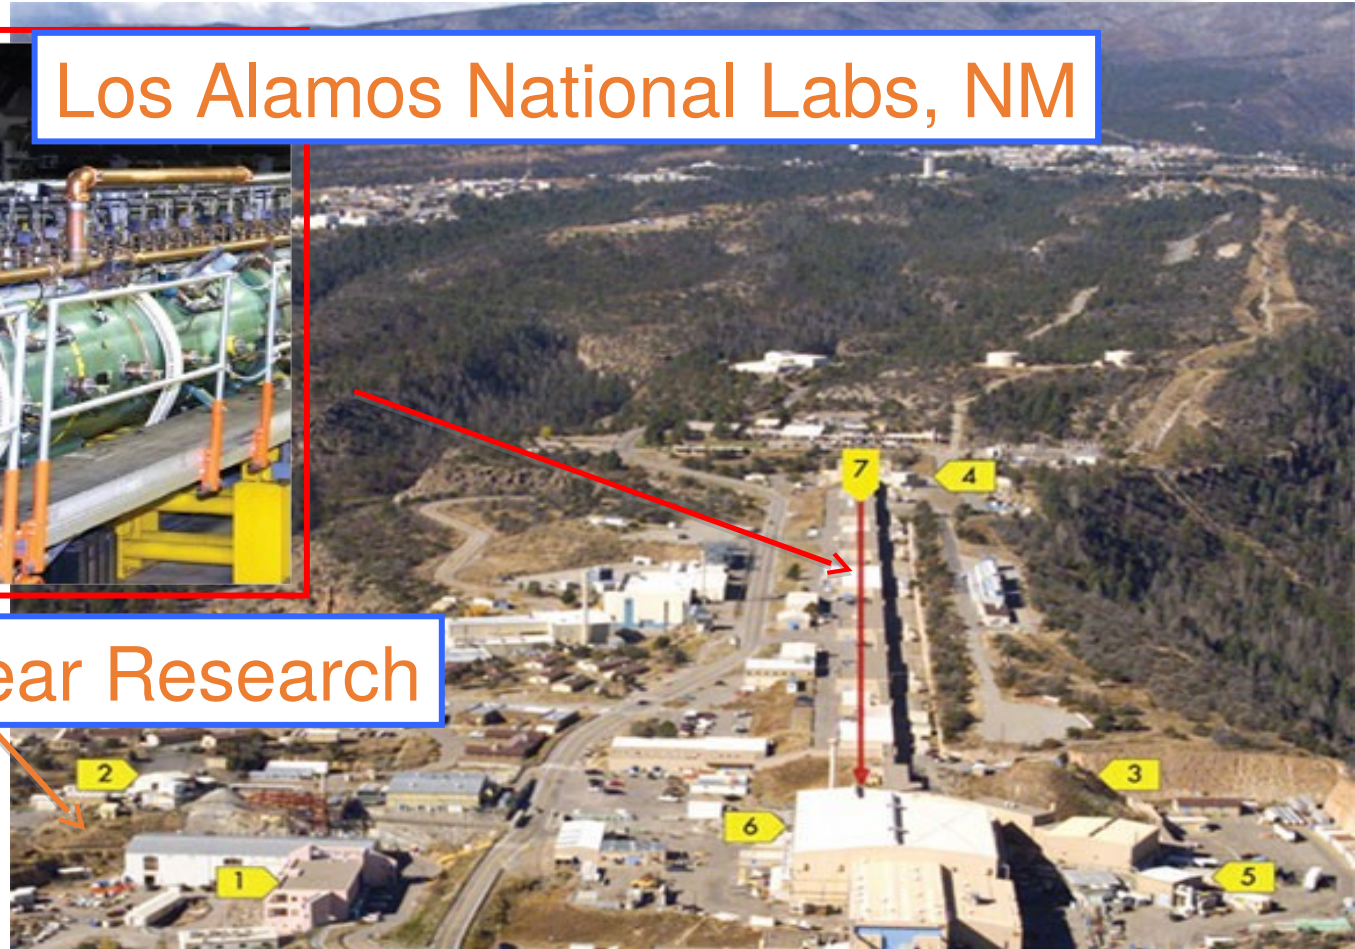
- Motivation
- **Radiation Experiments**
- Fault Injection
- Conclusion

Radiation Experiment Facilities

Los Alamos National Labs, NM




Weapon Nuclear Research




Quantify and Qualify SDCs


Number of incorrect elements




Spatial Locality



square




random




Relative Error

how different the error is from the expected value

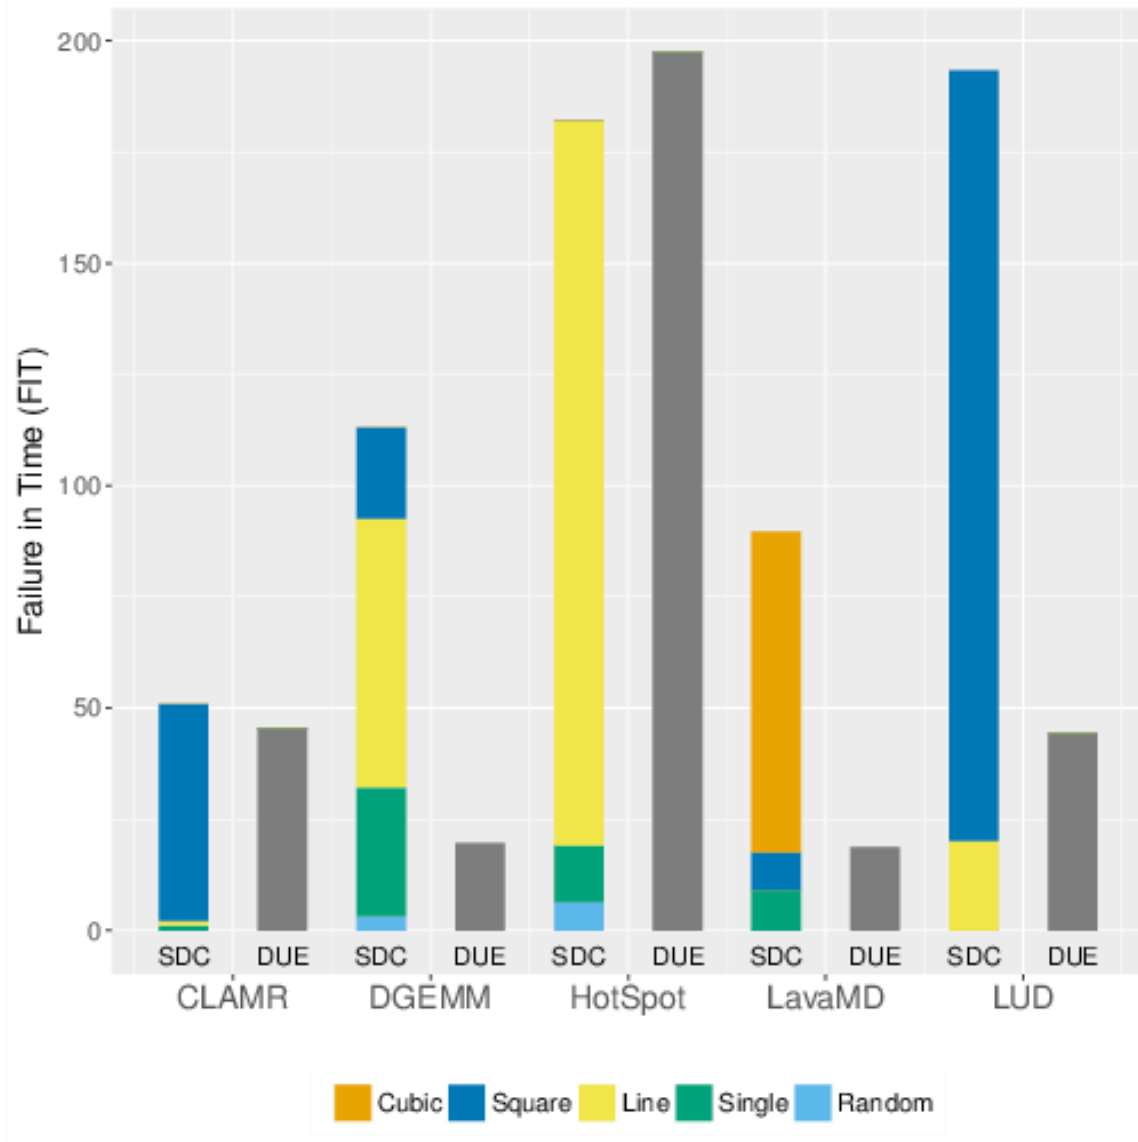


Potentially Masked Errors

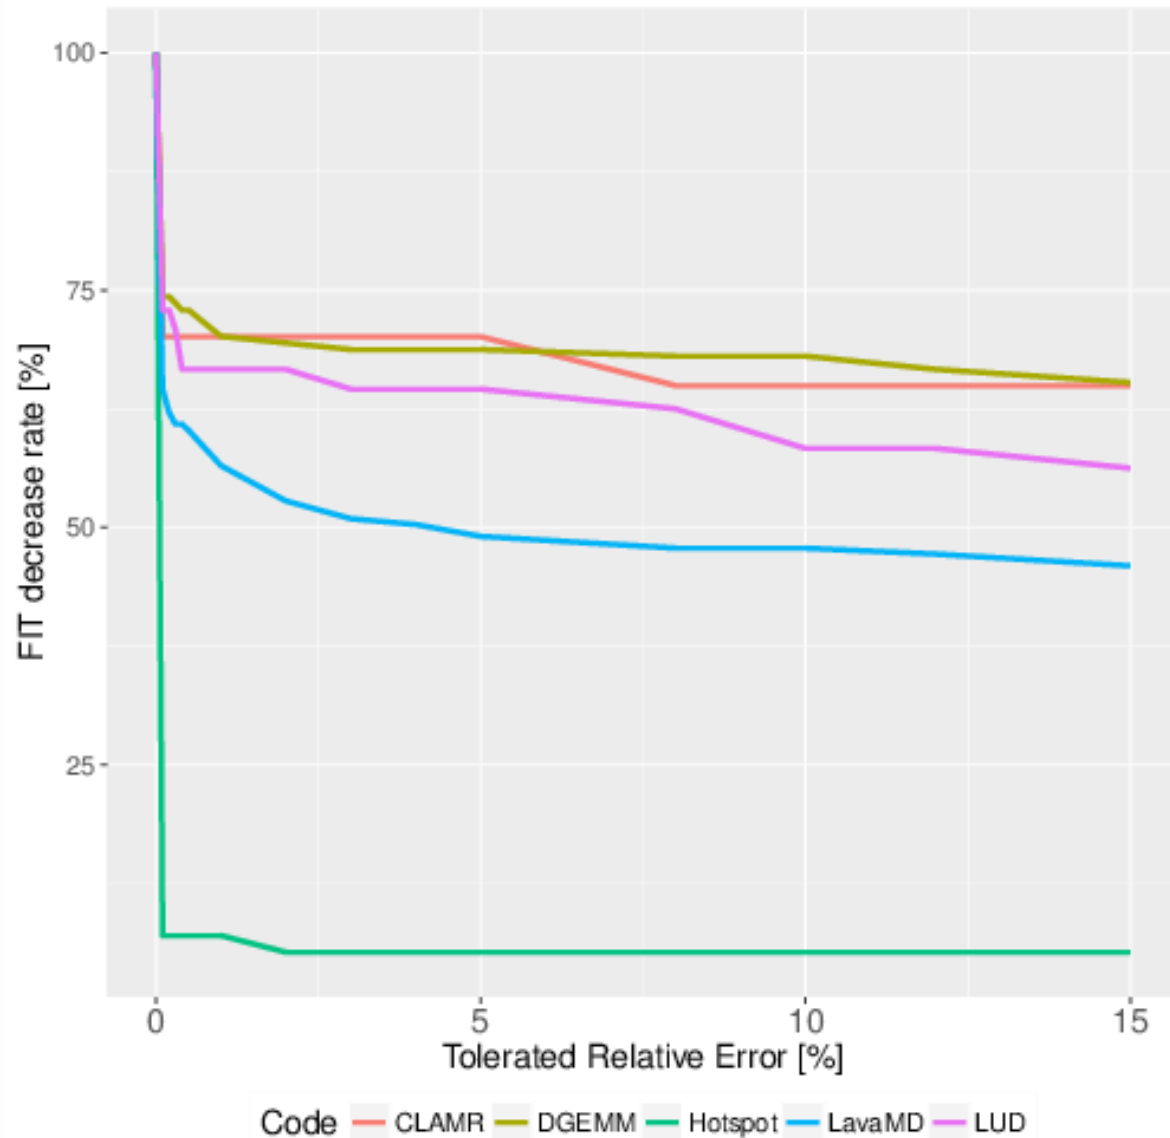
relative error $< 2\%$ is tolerable



Radiation Experiment Results



Relative Error Impact




- Motivation
- Radiation Experiments
- **Fault Injection**
- Conclusion


- High-Level Fault Injector
 - Modify memory content that is currently allocated by the system
 - Debug information used to correlate allocated memory portion with source code's variables.
 - Fault model can be adapted (single, double, random, zeros, and LSE implemented)
 - Overhead up to 8x (4x on average)

CAROL-FI Results

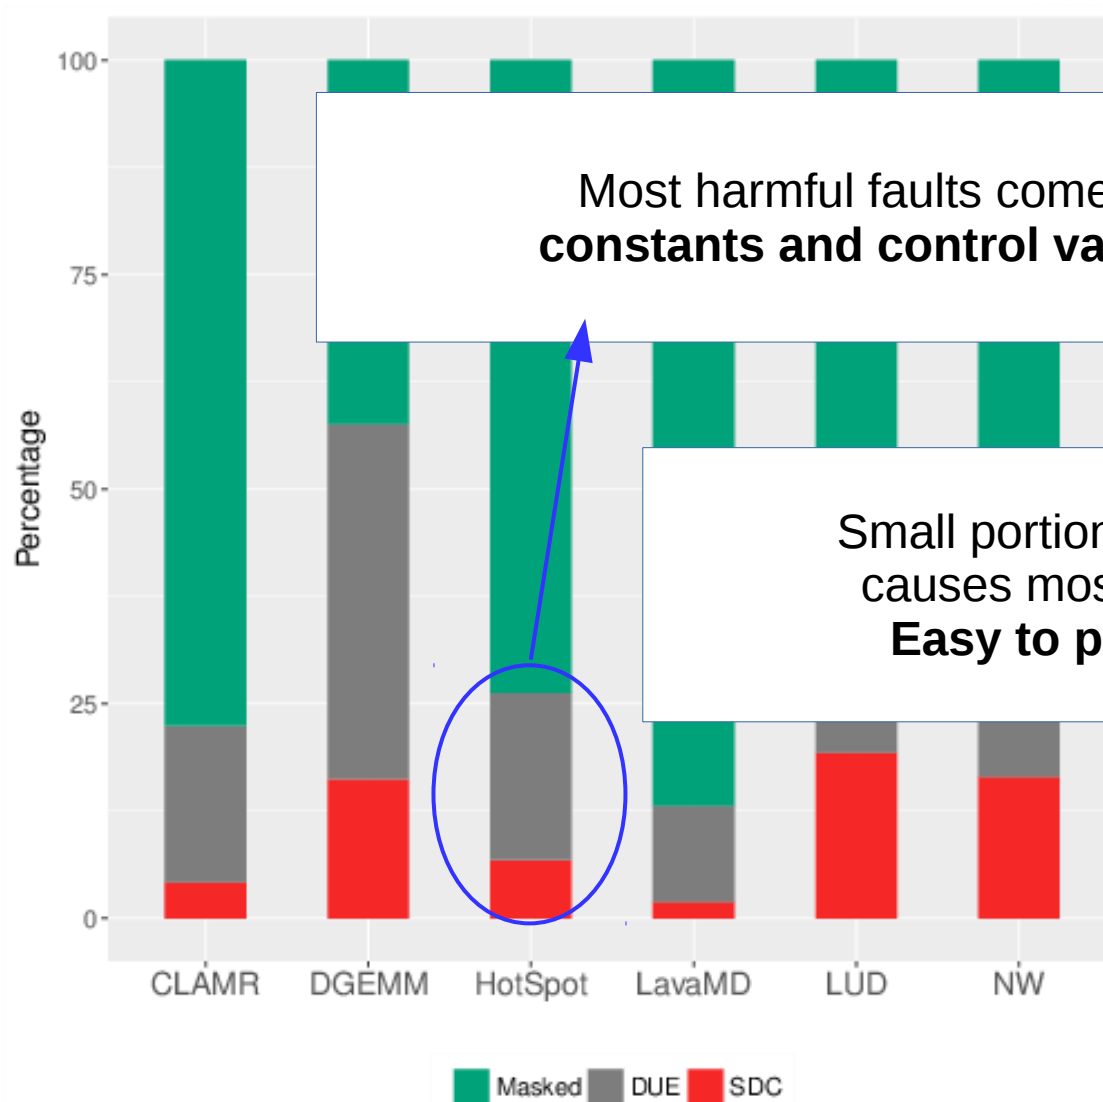
Fault injection campaign with more than 10,000 faults into each benchmark



DGEMM



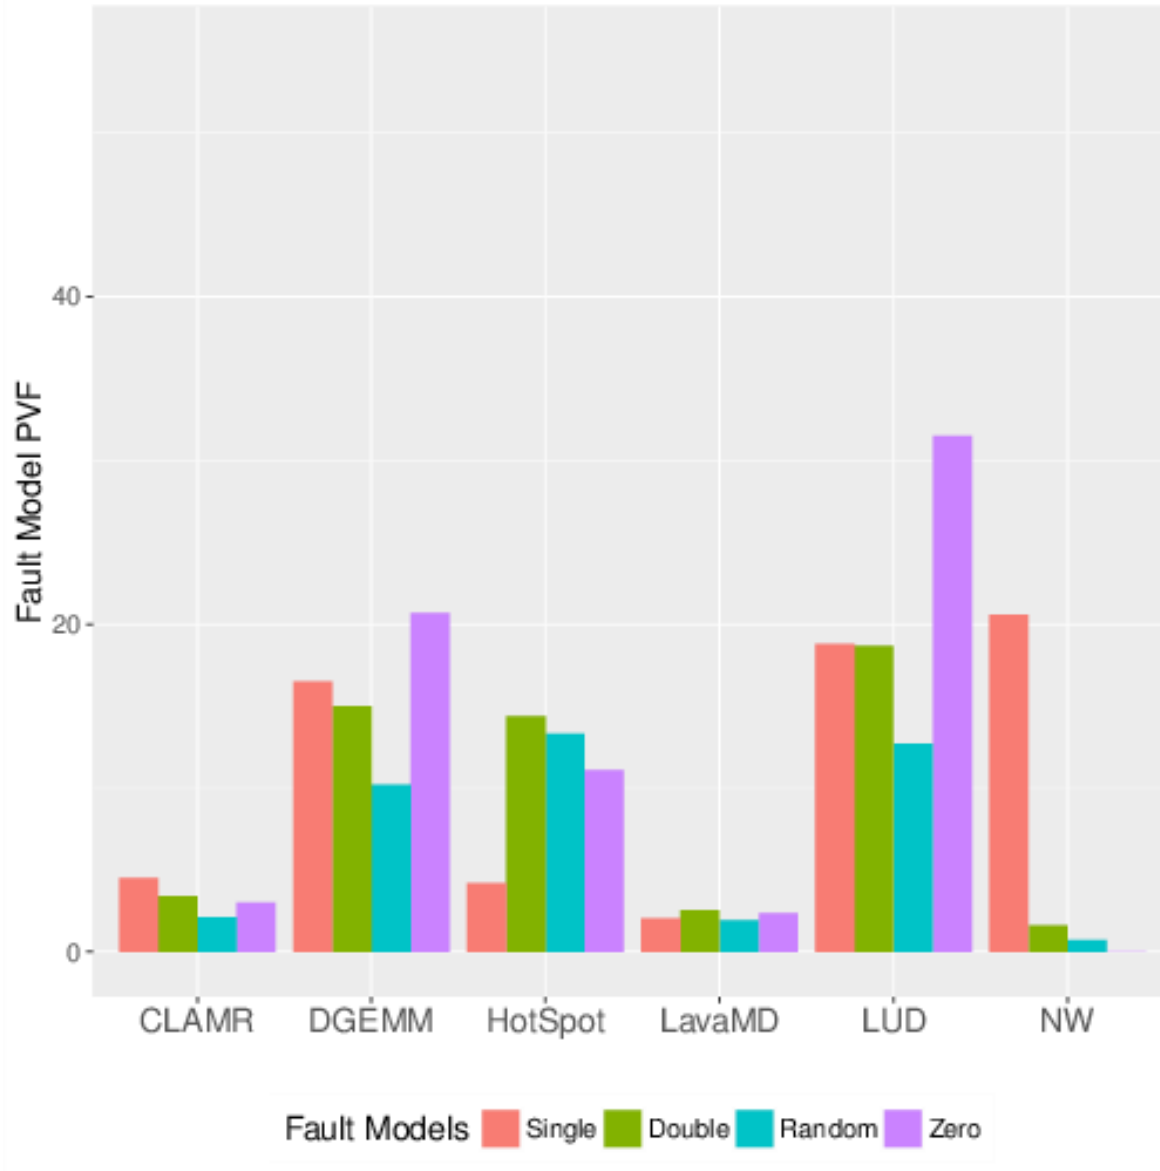
Hotspot



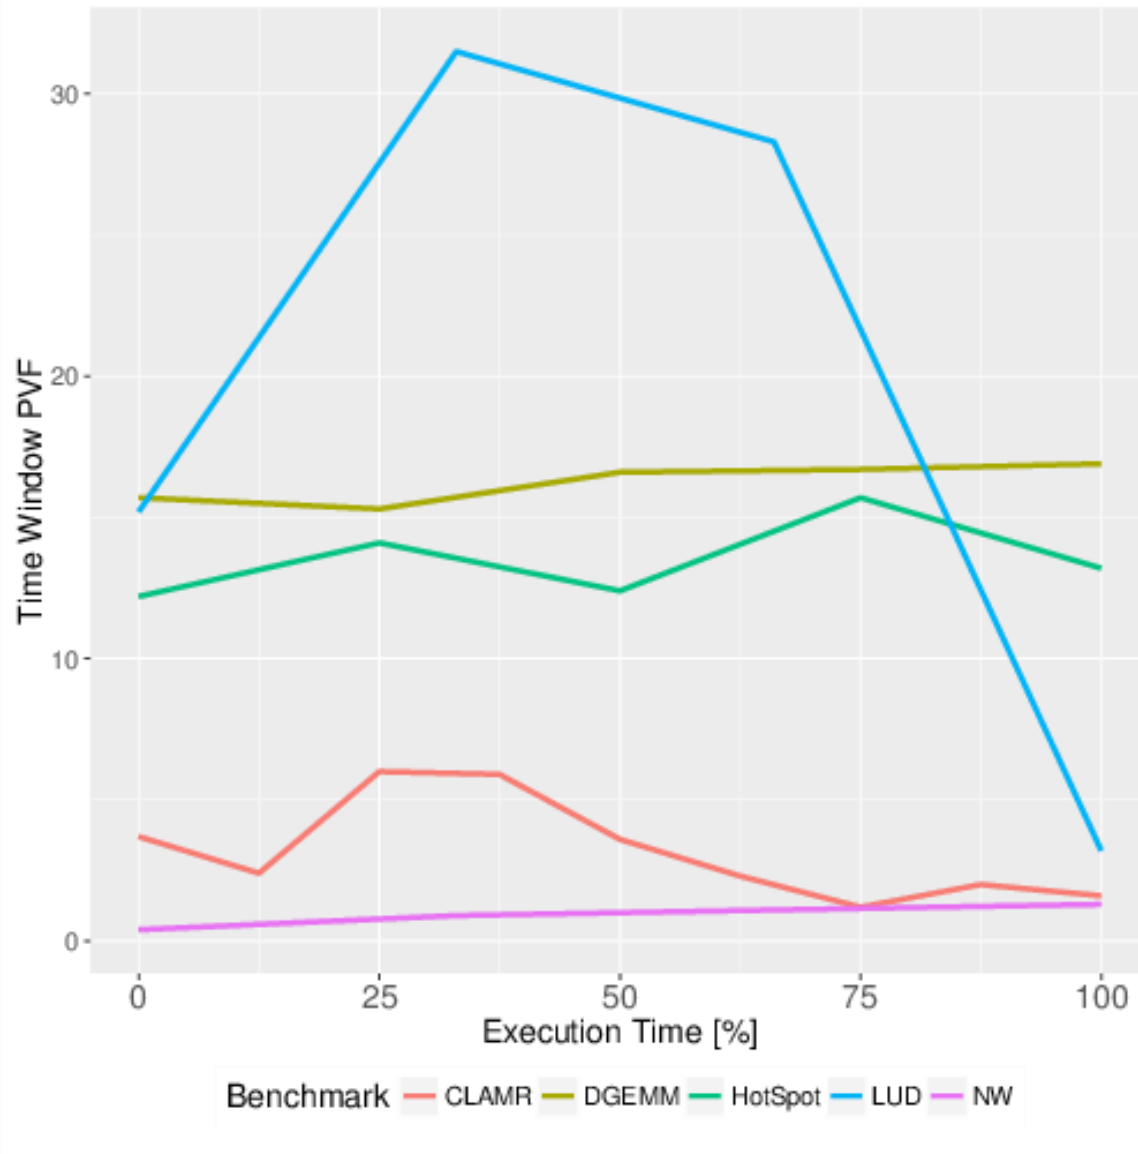
Most harmful faults come from **constants and control variables**

Small portion of data causes most harm:
Easy to protect

Fault Model Impact - SDC



Time Window Impact - SDC



- Motivation
- Radiation Experiments
- Fault Injection
- **Conclusion**

- Realistic evaluation of how the errors spread and the severity of SDCs
 - Output pattern may can be beneficial to mitigation techniques
 - Imprecise computation impact on HPC
- Deeper understanding using CAROL-FI
 - Codes portion that are more sensitive
 - Fault model and time window impact

Transient Faults Effects on HPC Applications

SC17: Experimental and Analytical Study of Xeon Phi Reliability

Daniel Oliveira, Laércio Pilla, Nathan DeBardeleben, Sean Blanchard, Heather Quinn, Israel Koren, Philippe Navaux, Paolo Rech

Thanks