

# **Silent Data Corruption and Other Anomalies**

**Sarah Michalak**

**Statistical Sciences Group**

**Los Alamos National Laboratory**

**michalak@lanl.gov**

# Silent Data Corruption (SDC)


---

- **“SDC occurs when incorrect data is delivered by a computing system to the user without any error being logged” C. Constantinescu (AMD)**
- **SDC can have many causes** [Constantinescu 2008]
  - Environmental (temperature/voltage fluctuations; particles), manufacturing residues, oxide breakdown, electro-static discharge
- **New technologies may experience more data corruptions** [Borkar 2012]
  - Increasing number of bits/latches, noise levels; decreasing voltages
- **Can affect desktops/laptops and larger-scale systems**
- **Greater issue with increasing scale**
  - ECC/parity provide protection, but perfect systems can't be assumed
- **Impact is hard to quantify precisely: SDC is silent and rare**
  - Different applications may have different susceptibility

# LANL SDC Research

## ■ Platform Testing:

- 9 platforms tested: all with HPL; 5 with HPL and Crisscross interconnect tester
- HPL SDC observed on 4 platforms post-decommissioning



*A single node participated in 70 incorrect Linpack calculations*

## ■ Laboratory Testing:

- Dual-core 65nm processor on a high-performance overclocking MOBO
  - Manipulate frequency, voltage, fan speed/temp while running Linpack
- Multiple forms of SDC outside of nominal conditions
  - Incorrect Linpack results, erroneous timestamps/environmental data
- Other Errors: program termination with no error, program crash, system crash

## ■ Neutron Beam Testing of Roadrunner (1<sup>st</sup> petaflop/s supercomputer):

- 2 SDCs observed on Cell processor; 2 SDCs observed on Opteron processor
- Estimate 1 cosmic-ray-neutron-induced SDC every 1.5 months of operation

# Other Anomalies

---

- **Data parsing (grep?) issue: different counts of failed HPL calculations**
  - 5 false positives in 10M lines parsed; rate replicated on second cluster
  - Appears related to a binary failing to link to a required dynamic library
    - NOT silent if you capture all messages; how many people do this?
- **R duplicated lines reading in file using read.table command**

***Computational Reproducibility Must  
Account for SDC and Other Anomalies***

# Collaborators

---

- Sean Blanchard
- Andy DuBois
- Dave DuBois
- Andrea Manuzzato
- Dave Modl
- Heather Quinn
- Bill Rust
- Ruben Salazar
- Curt Storlie
- Bruce Takala
- Steve Wender
- And many others...