Computer Science
UNIVERSITY OF TORONTO

TEMPERATURE MANAGEMENT IN DATA CENTERS: WHY SOME (MIGHT) LIKE IT HOT

Nosayba El-Sayed, Ioan Stefanovici, George Amvrosiadis,
Andy A. Hwang, Bianca Schroeder

{nosayba, ioan, gamvrosi, hwang, bianca}@cs.toronto.edu

ACM SIGMETRICS/Performance 2012

DATA CENTERS: POWER, EMISSIONS AND $$

- **Data centers are major energy hogs**
  - 30,000 square ft data center; consuming 10MW
  - Annual cooling cost = **$4-8 million**
- **Greenhouse Gas Emissions**
  - 2008: as much CO_2 as Argentina

    (McKinsey & Co., July 2008)

- **Where does all the power go?**

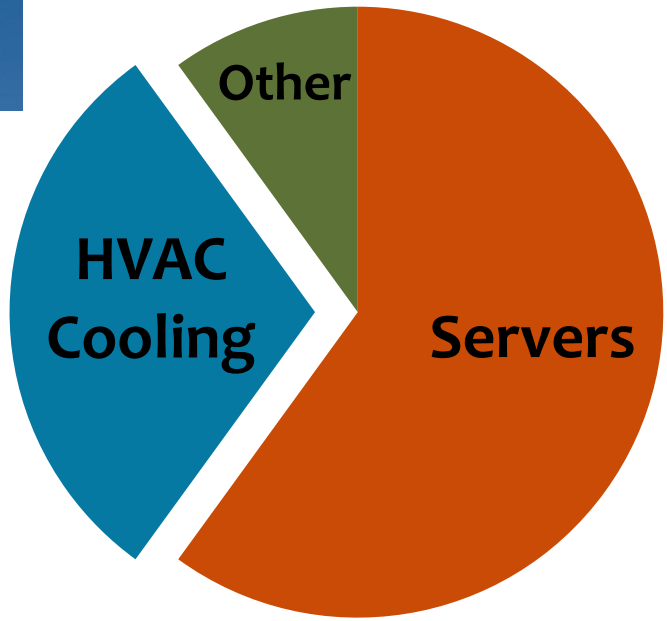

How do we reduce the power spent on cooling?

WHAT CAN BE DONE TO REDUCE COOLING POWER?

- **Improve air-flow management**
  (Sullivan, 2000), (Patel, 2003)
- **Load Balancing and temperature-aware workload placement**
  (Pinheiro, 2001) , (Bradley, 2003), (Rajamani, 2003), (Sharma, 2005)
- **Power reduction features in servers**
  (Flautner, 2002), (Gandhi, 2009)
- **Move to the Arctic Circle..**
  (Facebook, 2011)

- **Making data centers warmer!**
  - **1°C** increase in setpoint temperature
    → reduce energy consumption **2-5%**

Warm data centers: What can go wrong?

UNDERSTANDING THE IMPACT OF TEMPERATURE

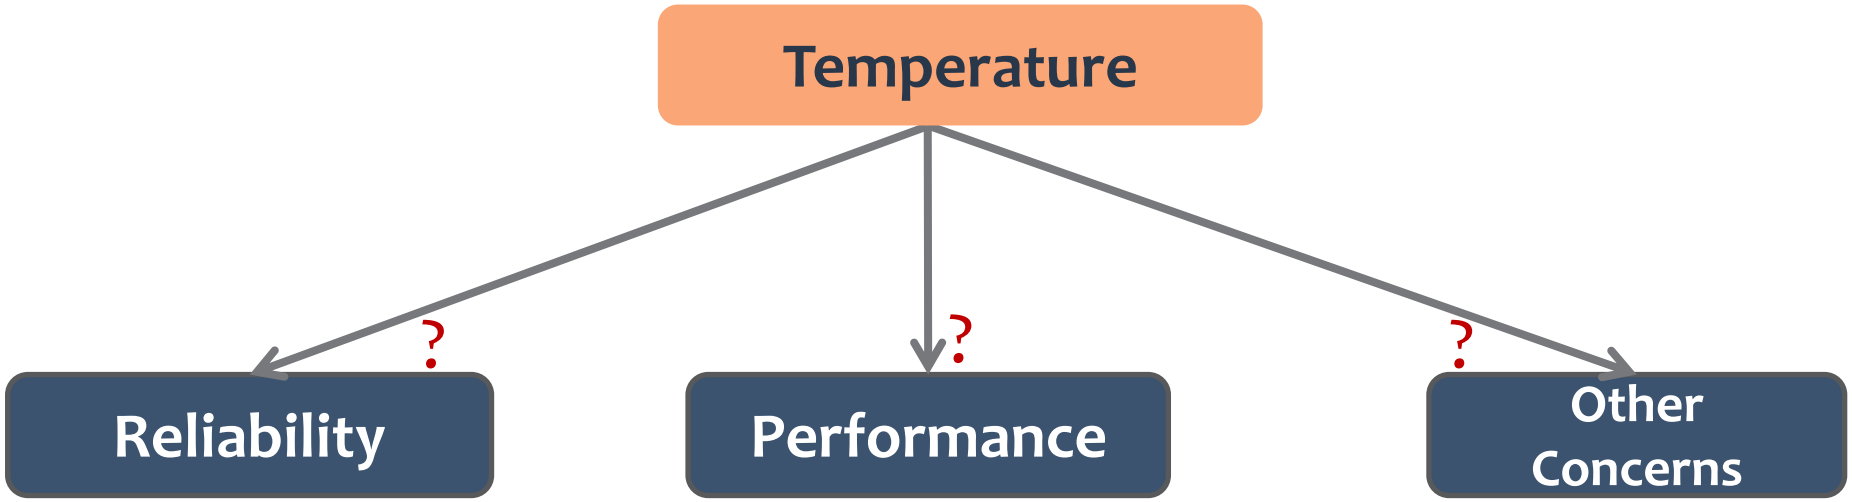

UNDERSTANDING THE IMPACT OF TEMPERATURE

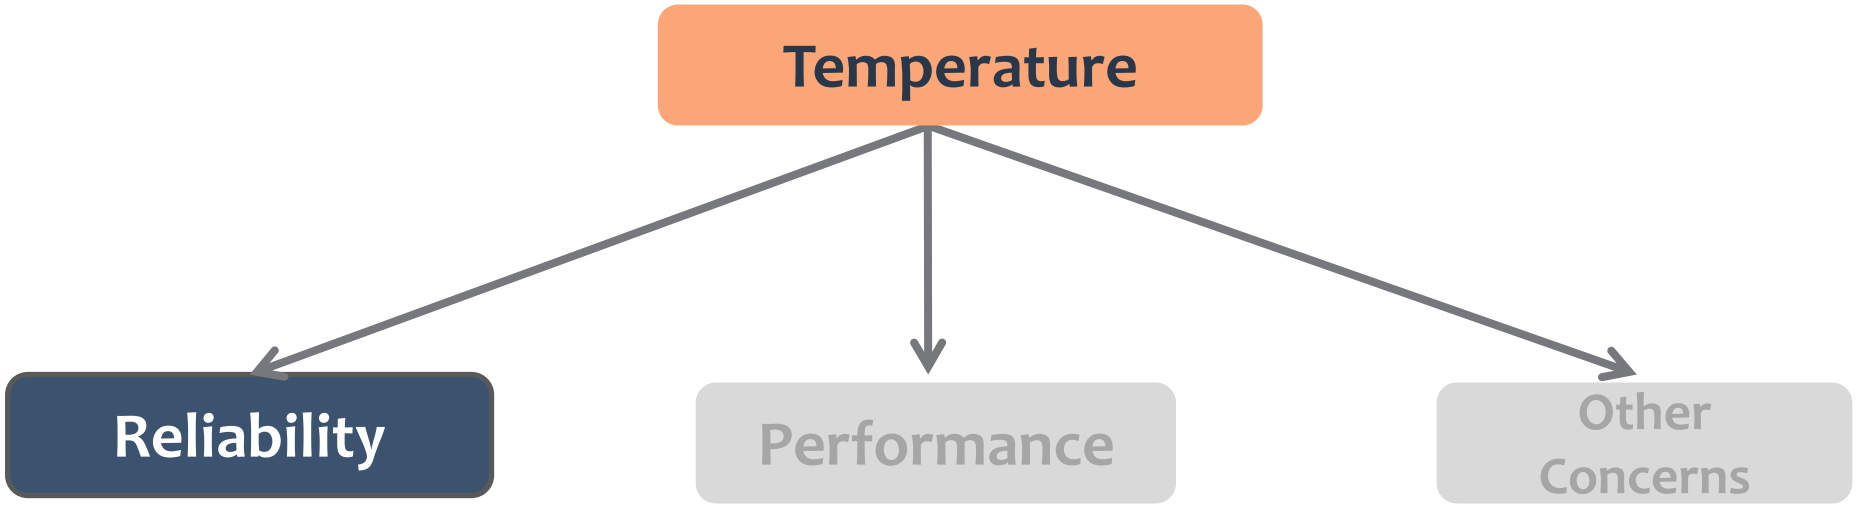

EXISTING MODELS: ARRHENIUS EQUATION

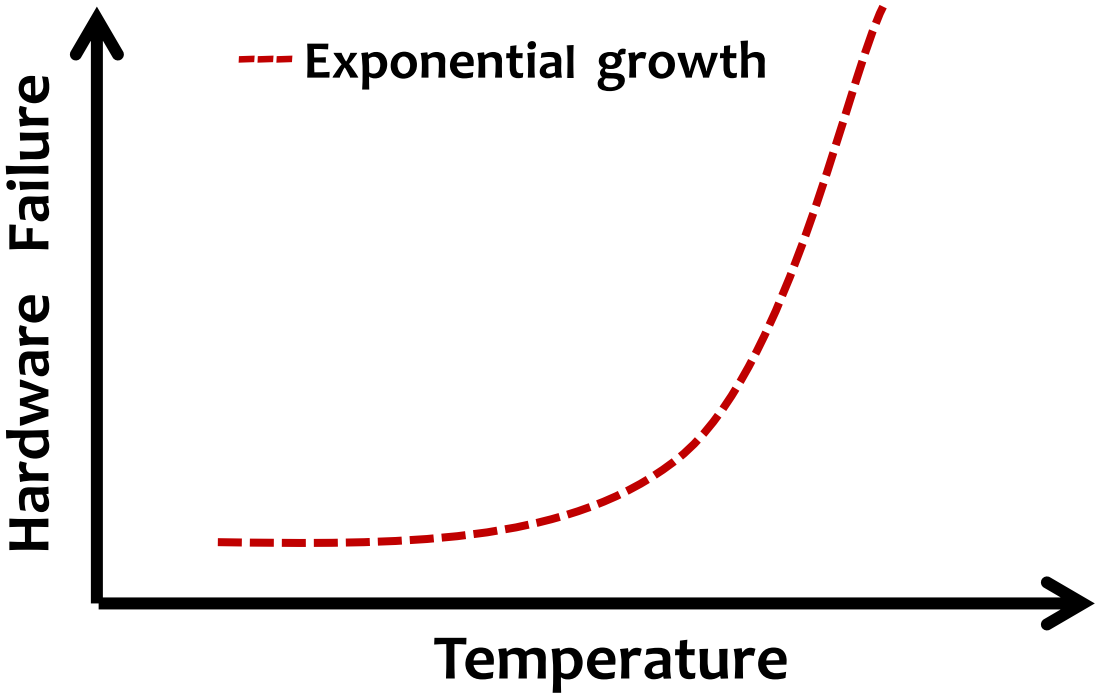

The real world: difficult to obtain empirical data!

TEMPERATURE AND RELIABILITY

1. Temperature and Hard Disk reliability
2. Temperature and DRAM reliability
3. Temperature and overall system reliability

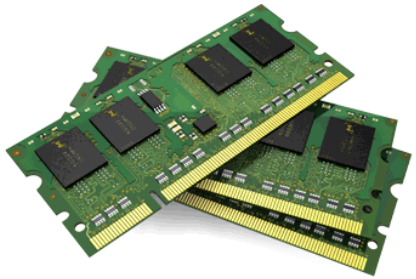

TEMPERATURE AND RELIABILITY

1. **Temperature and Hard Disk reliability**
   - **Disk Replacements**
   - **Latent Sector Errors (LSEs)**
2. Temperature and DRAM reliability
3. Temperature and overall system reliability


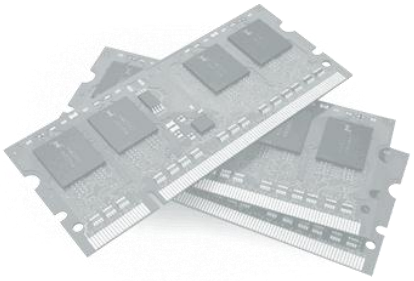


TEMPERATURE AND DISK FAILURES

- **Disk Replacements**
  - Typically **1-5%** of drives per year
    (Schroeder and Gibson, FAST 2007), (Pinheiro et al., FAST 2007)

- **Impact of temperature?**
  - **Google:** higher failure rates in **colder** temperatures!

(Pinheiro et al., FAST 2007)

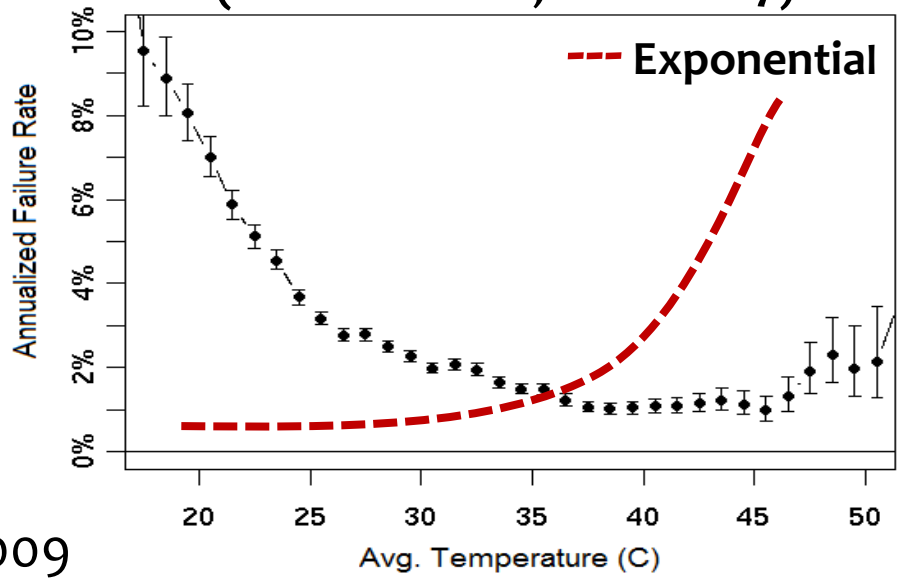

- **Our data: Google**
  - Sample: January 2007 - May 2009
  - 200,000 disks: 5 drive models; 19 data centers
  - Average Internal Temperature; Disk Age; Disk Utilization; Replaced?

HOW DOES TEMPERATURE IMPACT DISK FAILURES?

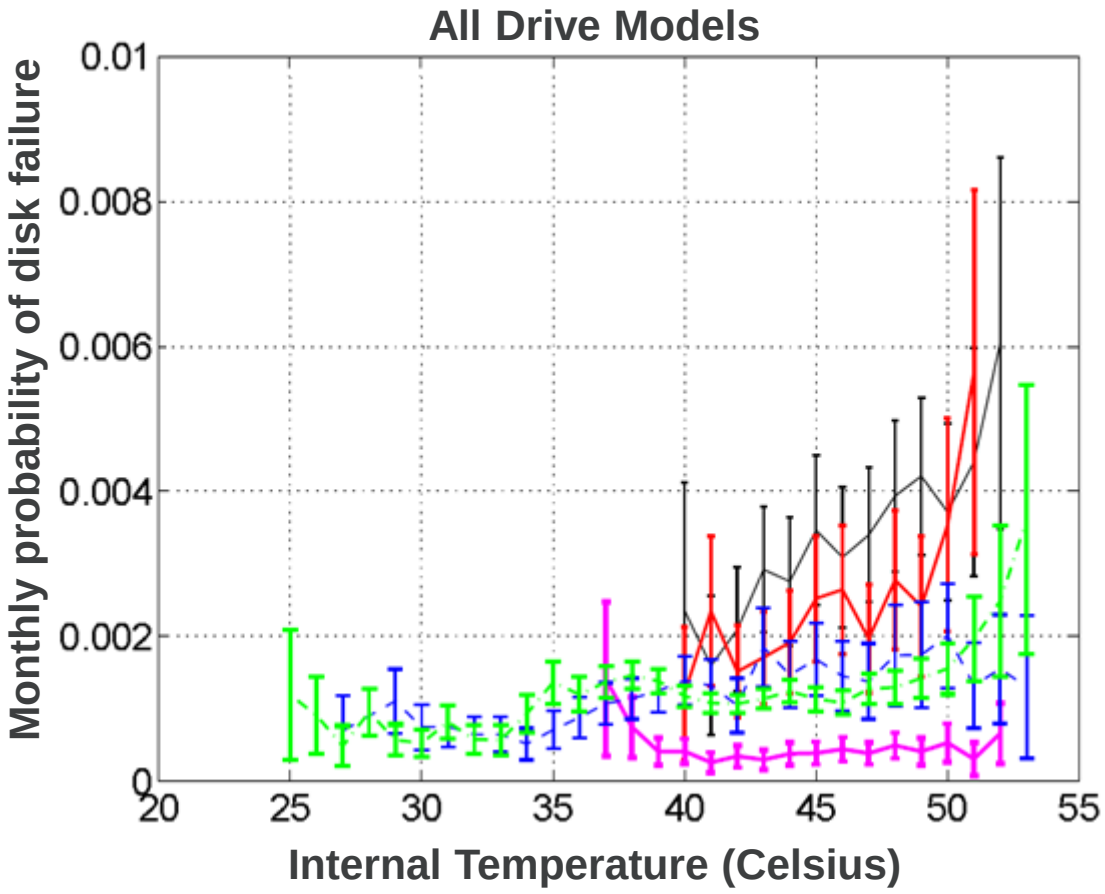

- Curve fitting: Linear model provided a better fit than exponential model for all lines!
- Contradiction with the Google 2007 study?

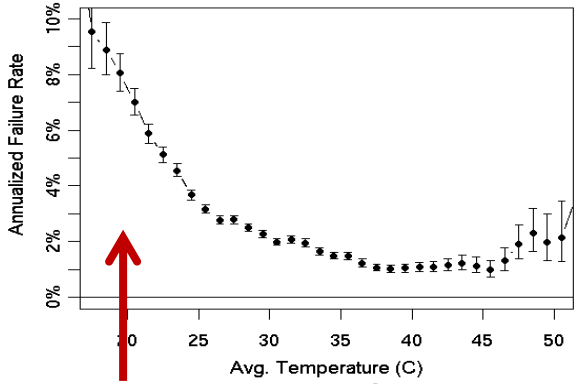

Aggregation of all models

Disk failures grow more slowly with temperature than expected!

TEMPERATURE AND LATENT SECTOR ERRORS

- **Latent Sector Errors (LSEs)**
  - Individual sectors on disk becoming inaccessible → data loss
  - Common failure mode: **3-4%** of disks
    (Bairavasundaram et al., SIGMETRICS 2007)
  - No prior work on how temperature affects LSE rates

- **Our data: Google**
  - Sample: January 2007 - May 2009
  - 70,000 disks: 3 drive models; 7 data centers
  - Average Internal Temperature; LSE Counts; Disk Age; Disk Utilization

HOW DOES TEMPERATURE IMPACT LSE PROBABILITY?

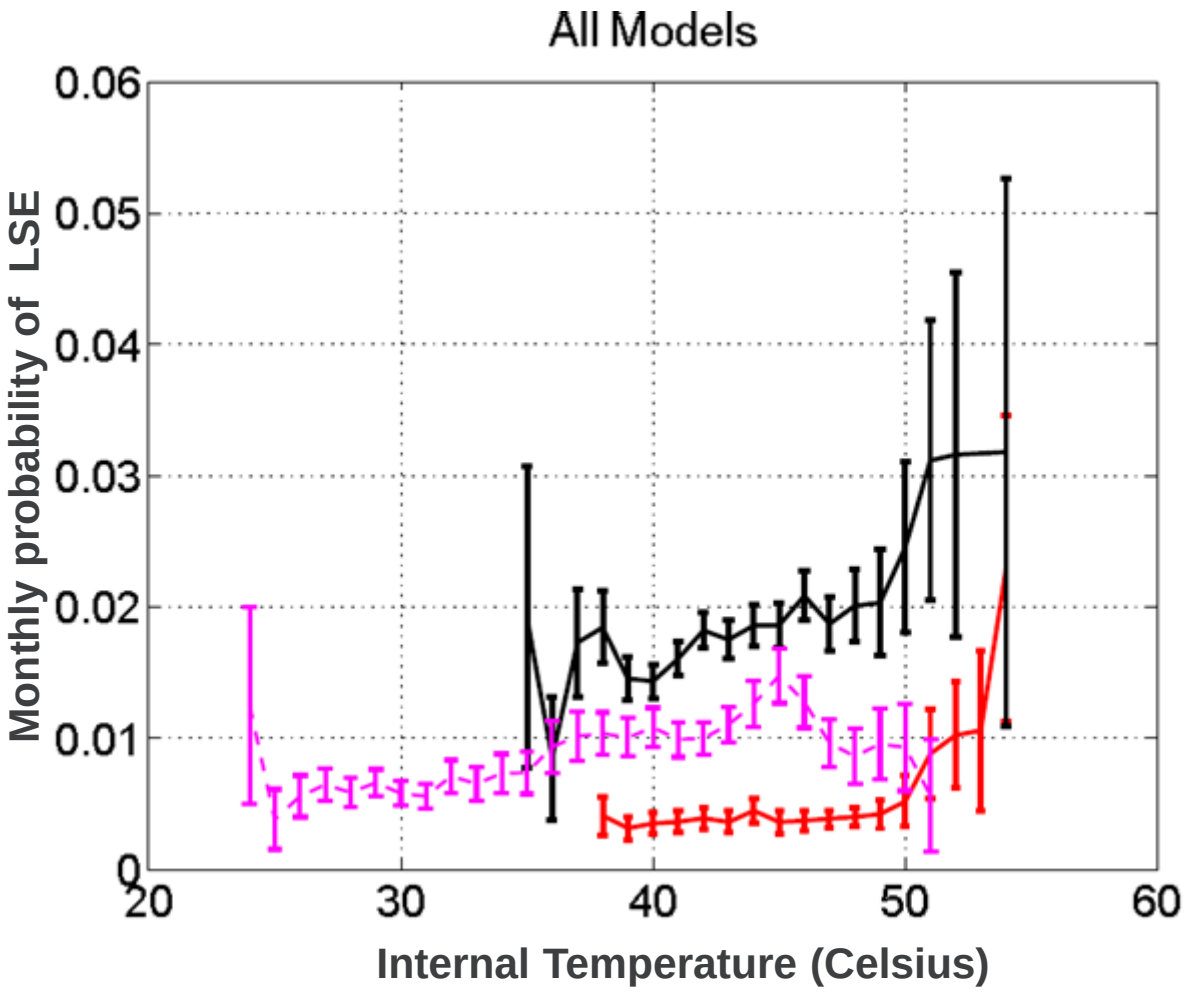

- **Curve fitting: Linear provided comparable and sometimes better fit than exponential**
- **Data center specific factors?** *(humidity, handling procedures, etc)*
- → **LSE probability varied across data centers** (more than **2x** difference)

- **LSEs increase more slowly with temperature than expected!**
- **Other data center specific factors seem to matter more**

LSEs AND TEMPERATURE: OTHER FACTORS

1. Age

Older disks are not more likely to develop LSEs under temperature!

2. Utilization

Disks with higher utilization are not more sensitive to temperature!

3. Temperature Variability

Impact of temperature variability was found to be stronger and more consistent than average temperature!

4. LSEs Frequency

Once LSEs have developed: higher temperatures did not increase LSE frequency!

TEMPERATURE AND RELIABILITY

1. Temperature and Hard Disk reliability
2. **Temperature and DRAM reliability**
3. Temperature and overall system reliability

TEMPERATURE AND DRAM RELIABILITY

- **Memory: what could go wrong?**
  - Corruption of one or multiple bits
    1. **Correctable errors**
    2. **Uncorrectable errors:** **cannot** be corrected with memory Error Correcting Codes (ECC) → **DIMM replacement** / **server crash!**
- **Our data**

Google

Uncorrectable Errors

SciNet

DIMM Replacements

Los Alamos NATIONAL LABORATORY EST. 1943

Server outages due to DRAM problems

TEMPERATURE AND DRAM RELIABILITY

- **Google**
  - Counts of Uncorrectable Errors (UEs); Internal temperatures
  - Five H/W platforms

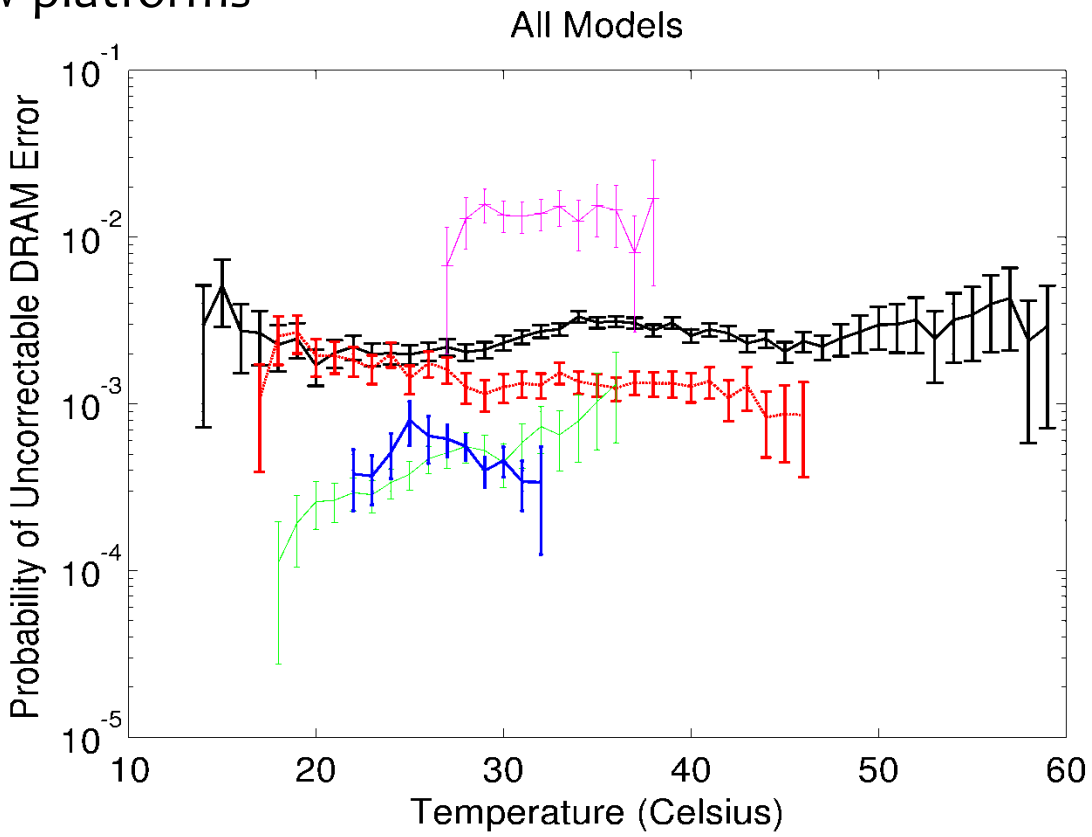

- **No evidence of increasing UEs in higher temperatures!**
- **Similar observation for DIMM replacements (SciNet), and node outages due to DRAM failures (LANL)**

TEMPERATURE AND RELIABILITY

1. Temperature and Hard Disk reliability
2. Temperature and DRAM reliability
3. **Temperature and overall system reliability**


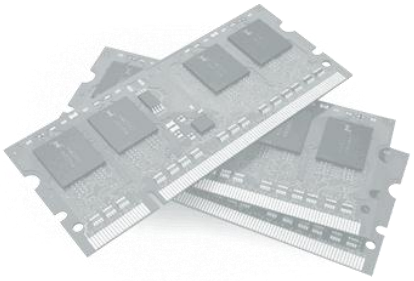


TEMPERATURE AND OVERALL RELIABILITY

- **What is the impact of temperature on overall system reliability and availability?**

- **Data:**

Los Alamos NATIONAL LABORATORY EST. 1943

Server outages due to any H/W problem

SciNet

H/W Replacements

TEMPERATURE AND NODE OUTAGES

- **Los Alamos National Lab (LANL)**
  - 13 HPC clusters (4384 nodes; 24,208 processors)
  - Node outages due to **hardware** problems (2002-2008); **rack** positions
  - One cluster (256 nodes): ambient **temperature** data (motherboard sensors)

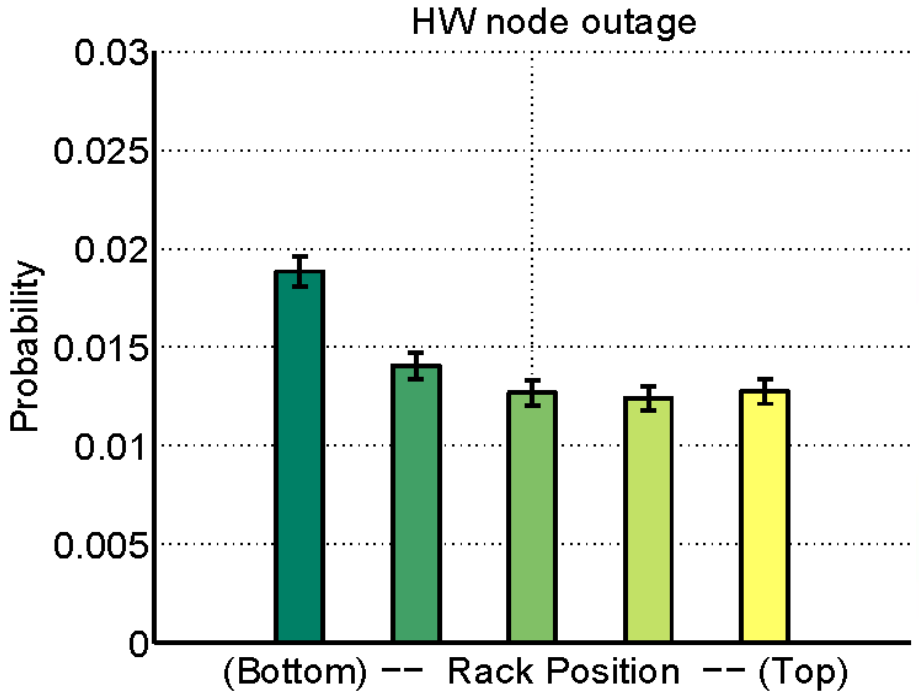

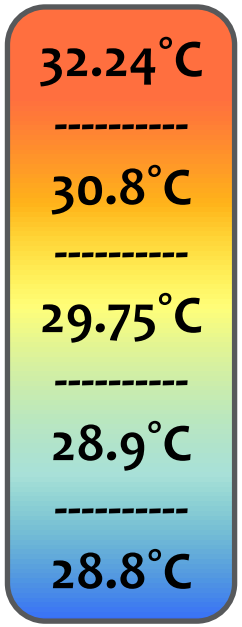

- **No evidence of increasing node outages in warmer rack positions!**
- **Similar observation found in H/W replacements in SciNet clusters!**

UNDERSTANDING THE IMPACT OF TEMPERATURE

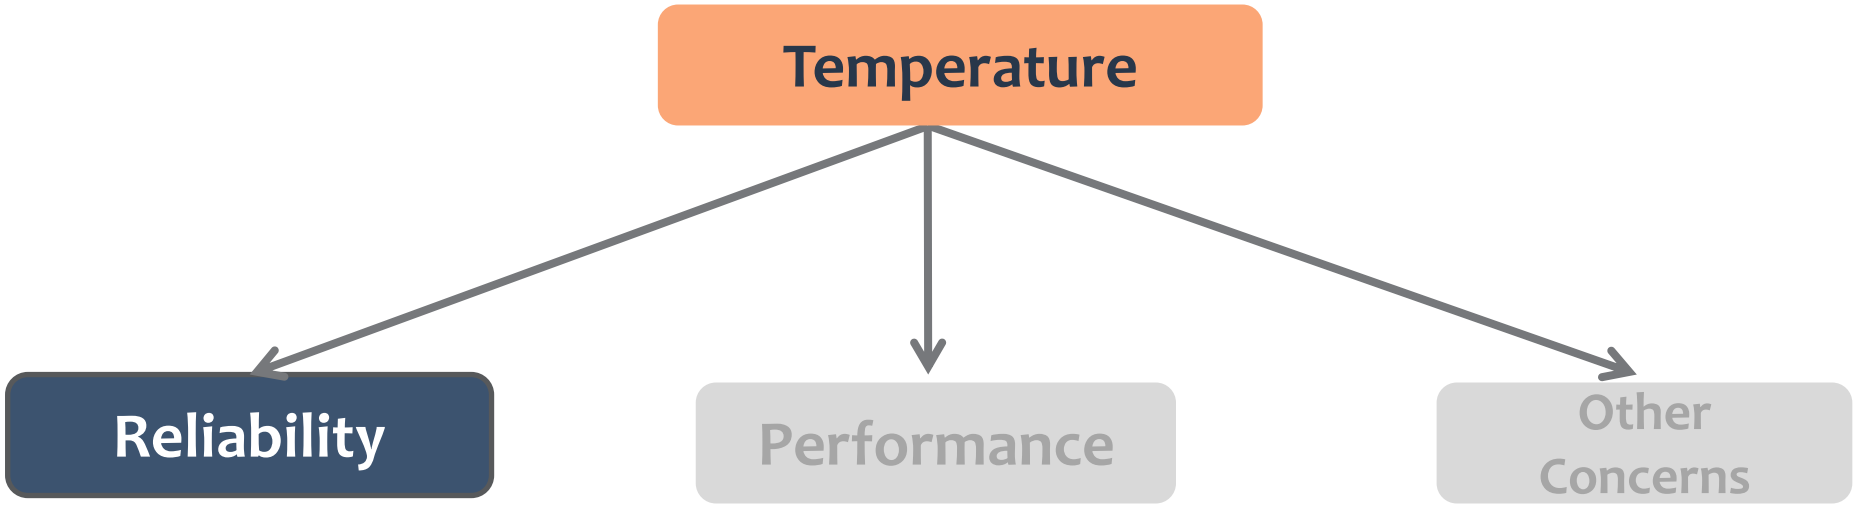

UNDERSTANDING THE IMPACT OF TEMPERATURE

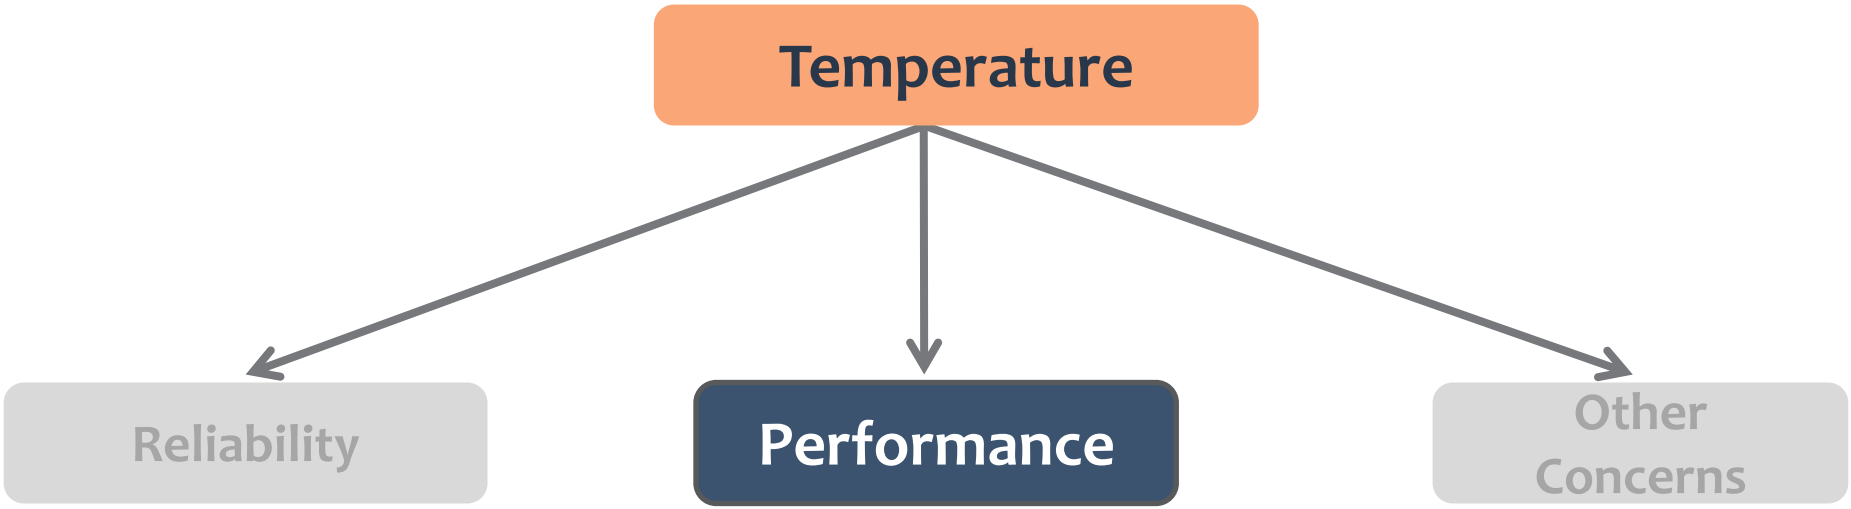

TEMPERATURE AND PERFORMANCE

- **What could go wrong?**
- Protection mechanisms at certain **temperature** thresholds
  ➔ Performance overhead
- **Problem: features and associated parameters not well-documented!**

- **Experimental Study**
  - Thermal chamber (10˚C to 55˚C)
  - Server: Dell PowerEdge R710
  - 2 synthetic workloads; 4 microbenchmarks
    6 macrobenchmarks
  - Variety of disk models (3 SATA, 4 SAS)

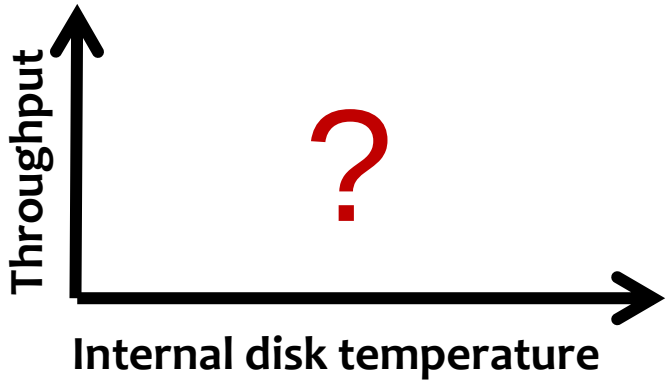

Graduate student

TEMPERATURE AND DISK PERFORMANCE

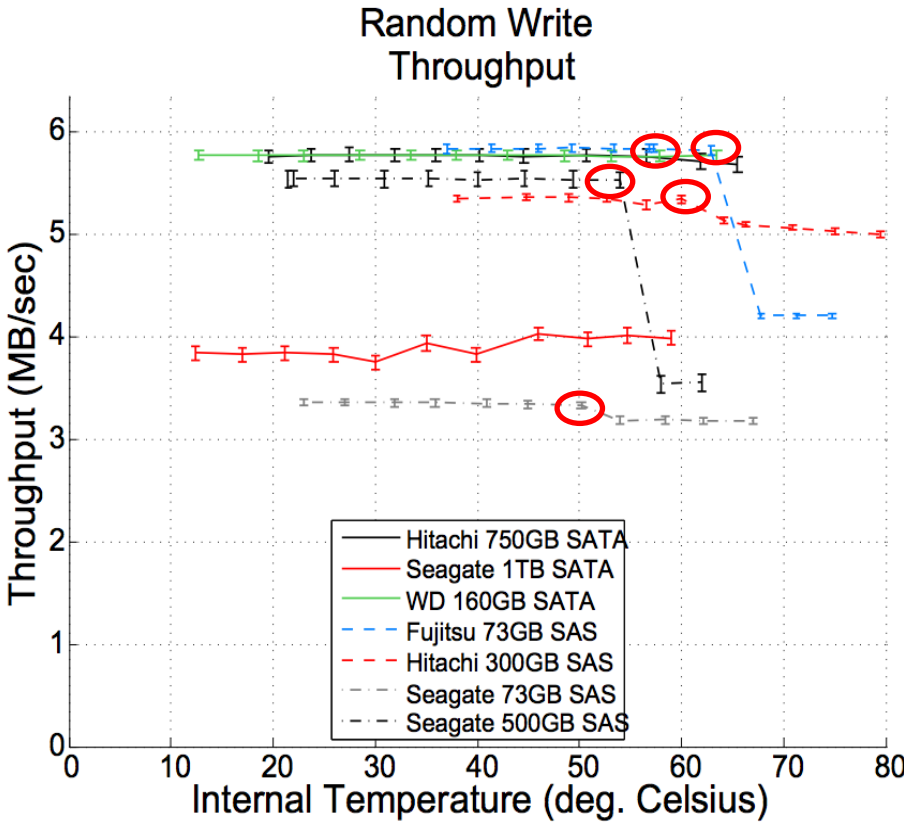

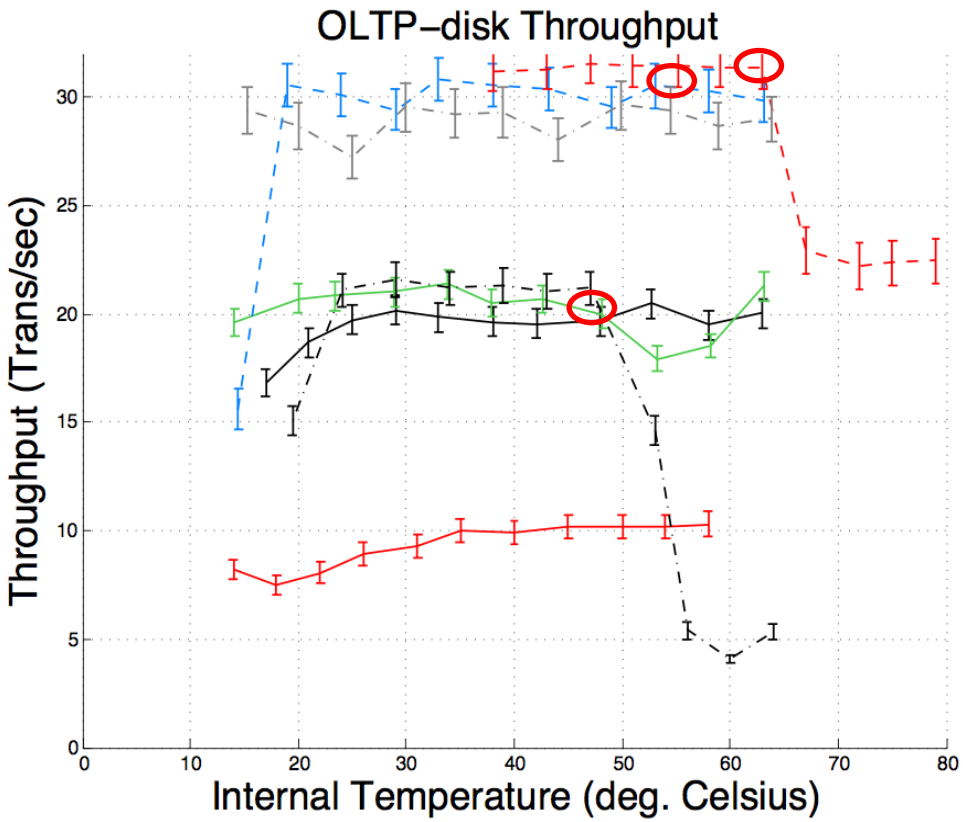

- 5 of 7 drives had throughput drops: can go up to 30% (synthetic), 80%(macro)
- For each drive: throughput drops observed at same temperatures across workloads! (due to enabling of protection mechanisms?)

MEMORY CONFIGURATIONS AND PERFORMANCE

- **Memory configuration options**
  - **Memory protection schemes:**
    - **S**ingle **E**rror **C**orrection-**D**ouble **E**rror **D**etection **(SEC-DED)**
    - Advanced ECC (Chipkill): detection and correction of multi-bit errors
    - Mirroring
  - **Memory bus speeds:** 800MHz, 1066MHz

manually enabled (by sysadmin)

dynamically activated (e.g.: at high temperatures)
→ None observed in our experiments

- **Experiment: manually configure server and observe impact**
  - Different bus speeds, ECC schemes
  - 7 different workloads (CPU-bound, mem-bound, macrobenchmarks)

Significant throughput drops (up to 40%) for memory-bound microbenchmarks when activating protection mechanisms!

WARM DATA CENTERS: WHAT COULD GO WRONG?

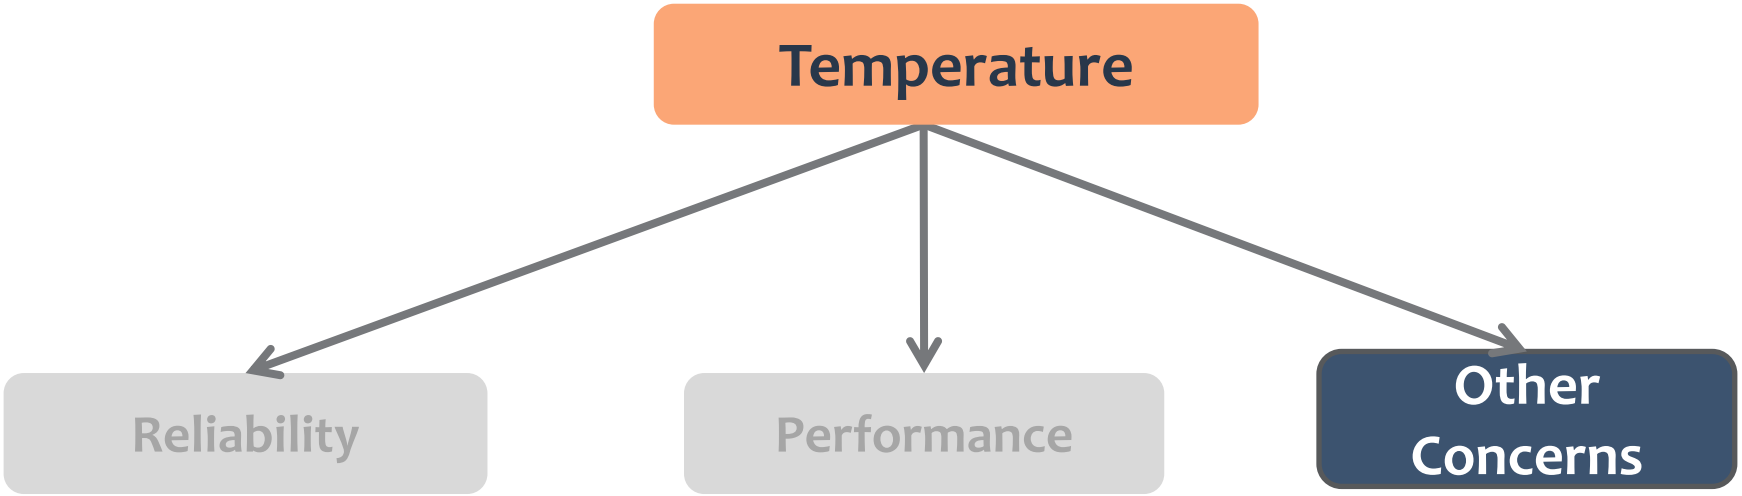

- Server Power Consumption
- Hotspots

TEMPERATURE AND POWER CONSUMPTION

- **Server power: what could go wrong in high temperatures?**
  - Increased leakage power
  - Increased server-fan speeds

  } increased server power

- **Quantify increase in power consumption under ambient temperature**

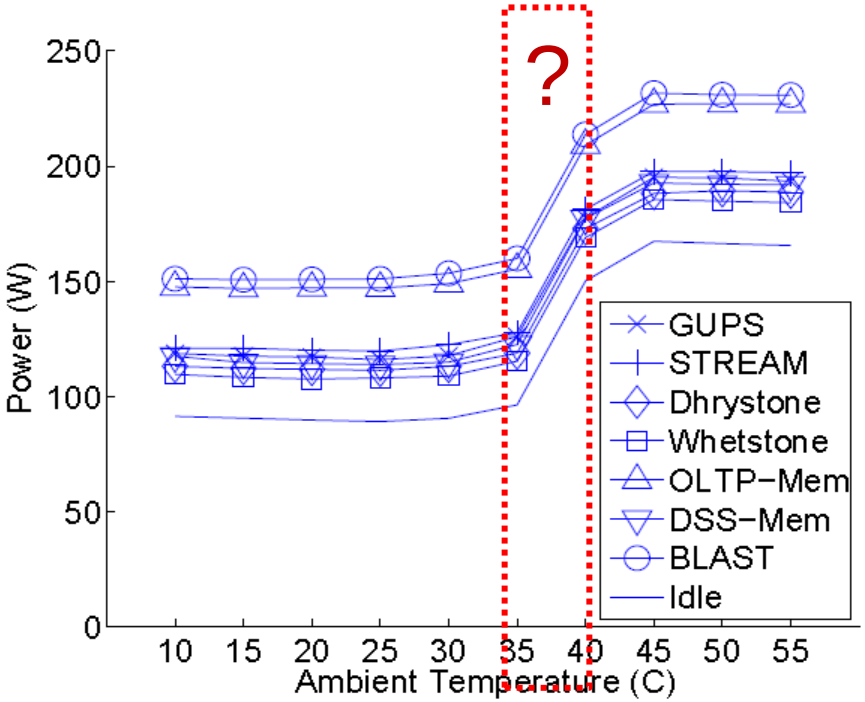

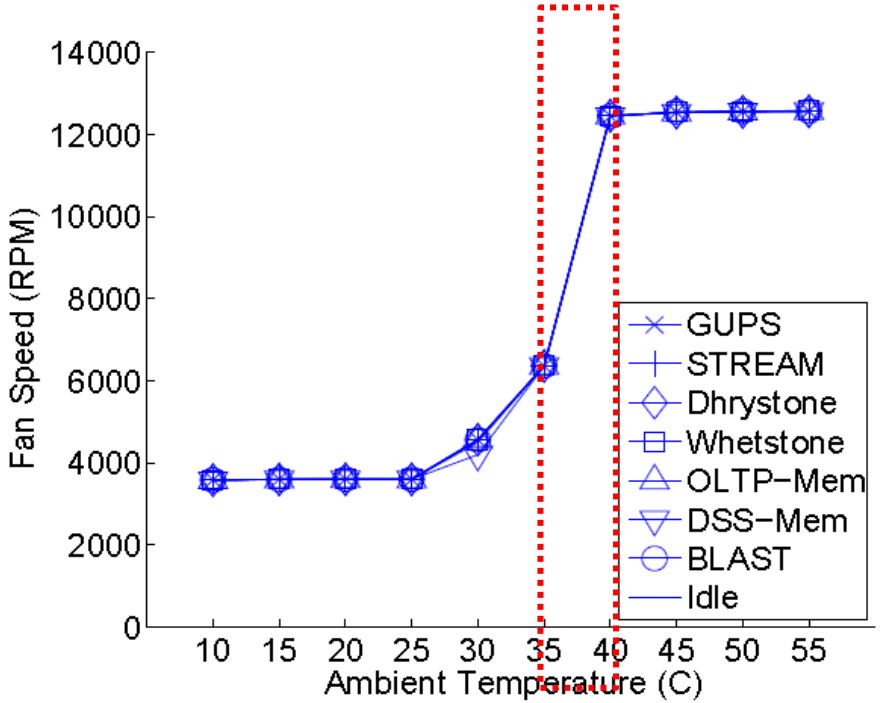

- **Evidence that power increase could be dominated by fan power**
- **Need more sophisticated fan control!**

REDUCED SAFETY MARGINS: HOT SPOTS

- **Hot Spots in data centers**
  - Significantly hotter than average room temperature
- **Raising setpoint temperatures**
  → Even hotter hot spots?

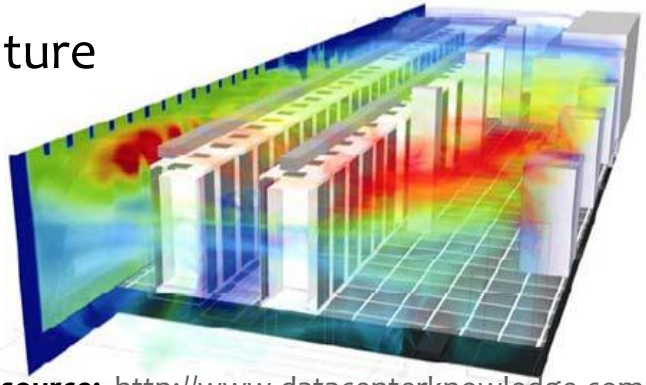

source: http://www.datacenterknowledge.com

- **Understand temperature imbalances and variation across nodes in different data centers**
  - **hottest 5%** nodes/disks: **5˚C** higher than median
  - **hottest 1%** nodes/disks: **8-10˚C** higher than median

Google Los Alamos NATIONAL LABORATORY EST. 1943

The degree of temperature variation across nodes was similar for different facilities managed by independent organizations!

SUMMARY AND IMPLICATIONS

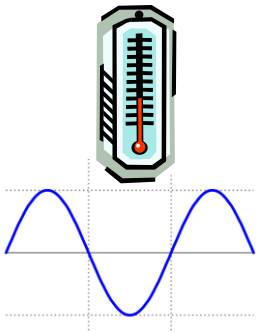

- **Temperature and reliability: impact smaller than assumed**
  - ➔ Consider raising setpoint temperature
- **More attention to temperature variability than average**
- **No correlation between DRAM failures and temperature**
  - ➔ Avoid performance overhead by disabling protection mechanisms
- **Disks and high temperatures: expect to deal with increasing errors (LSEs) more so than full disk failures!**
  - ➔ Consider periodic "scrubbing" to proactively detect LSEs

- **Need for smarter and more sophisticated fan controllers**

- **Reduced safety margins: keep in mind impact of raising temperature on hottest nodes in the facility!**
  - ➔ More detailed monitoring to react quickly to thermal shutdowns

TEMPERATURE MANAGEMENT IN DATA CENTERS

Questions?

nosayba@cs.toronto.edu

{ioan, gamvrosi, hwang, bianca}@cs.toronto.edu

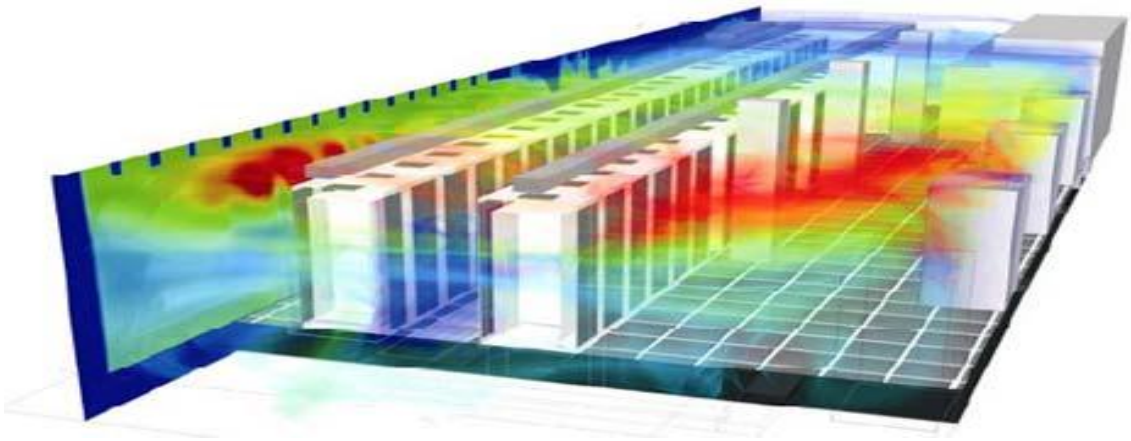

Computer Science
UNIVERSITY OF TORONTO

TEMPERATURE MANAGEMENT IN DATA CENTERS

Backup Slides

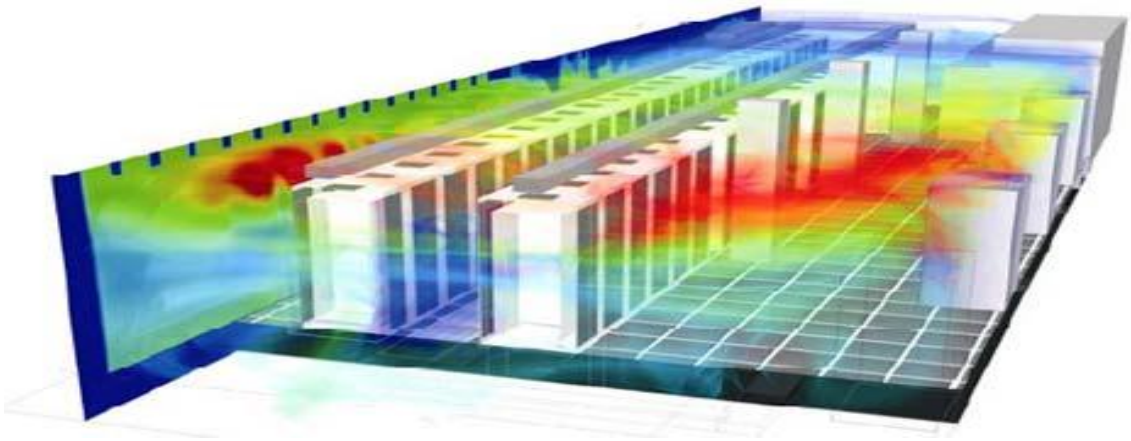

TEMPERATURE AND DISK FAILURES

(Pinheiro et al., FAST 2007)

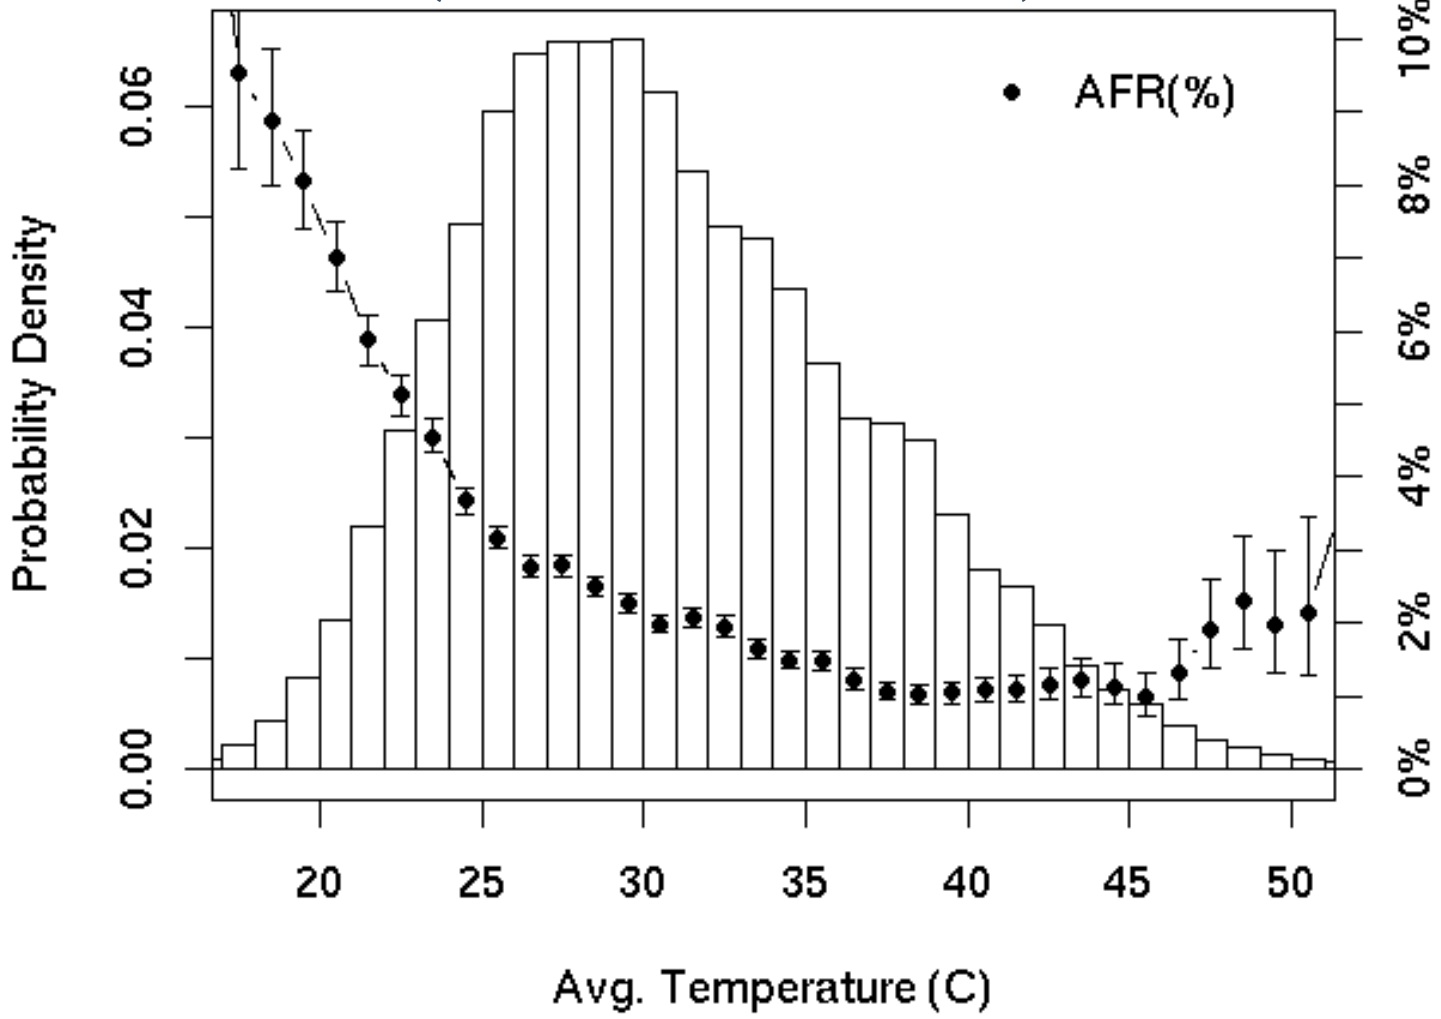

HOW DOES TEMPERATURE VARIANCE IMPACT LSEs?

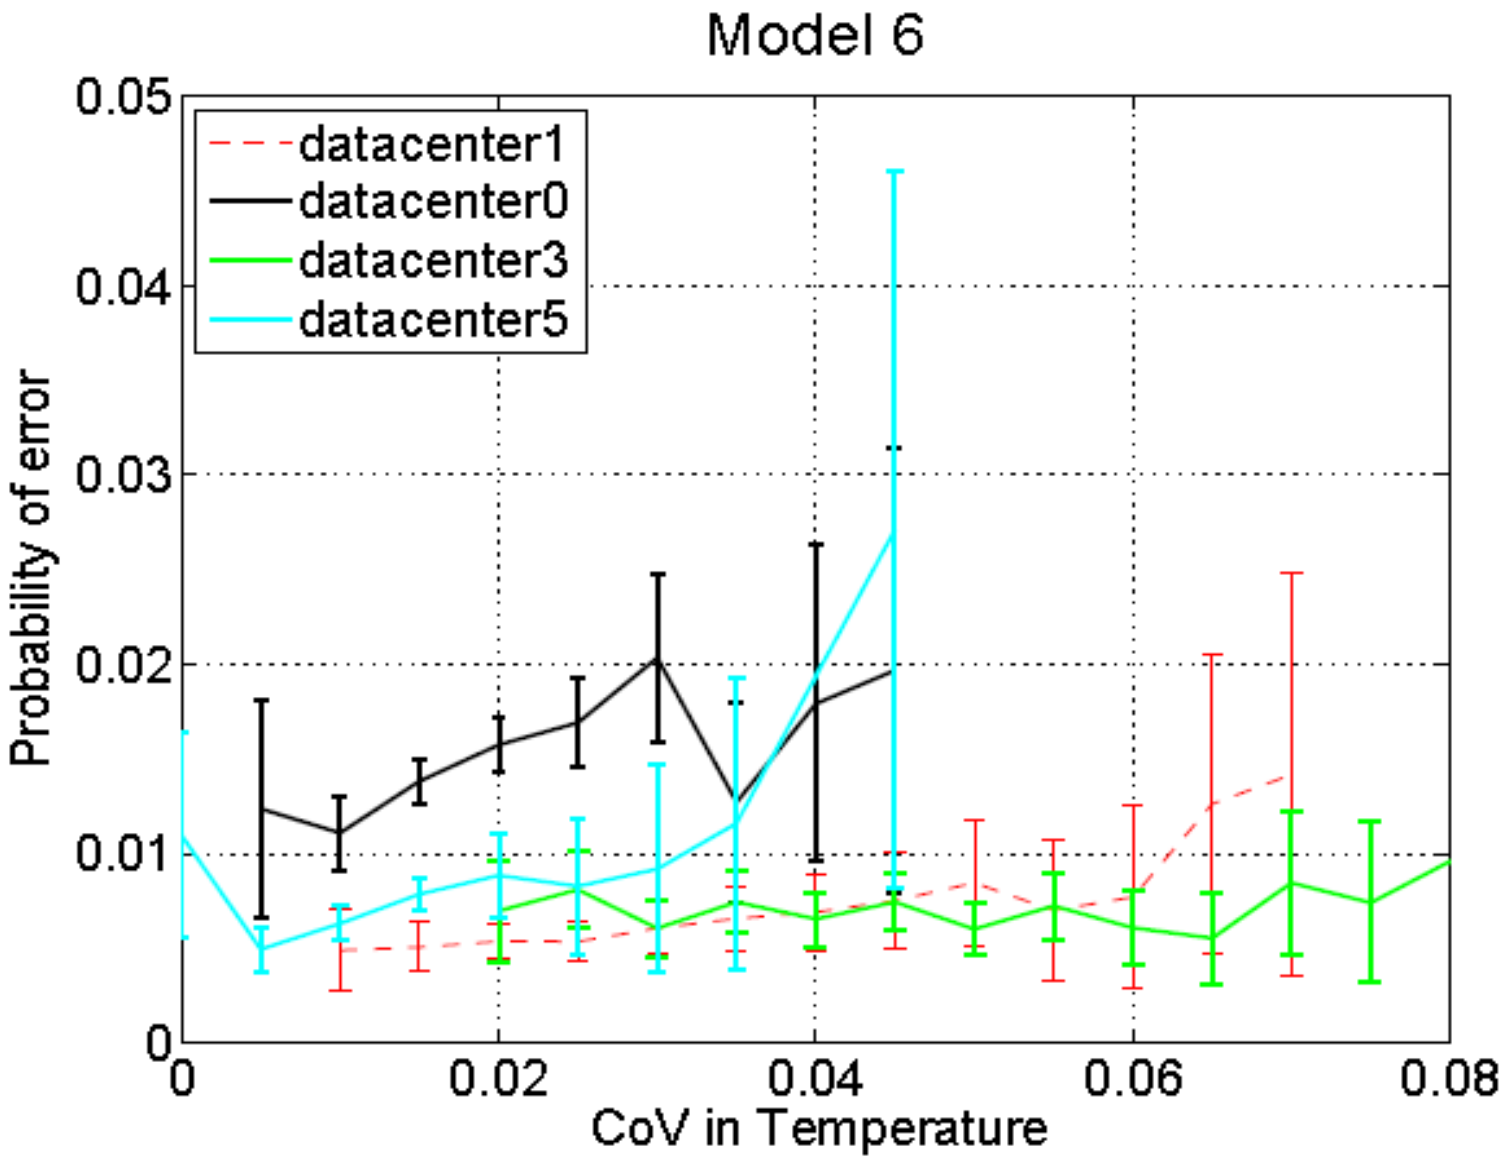

TEMPERATURE VARIANCE IMPACT ON NODE-OUTAGES

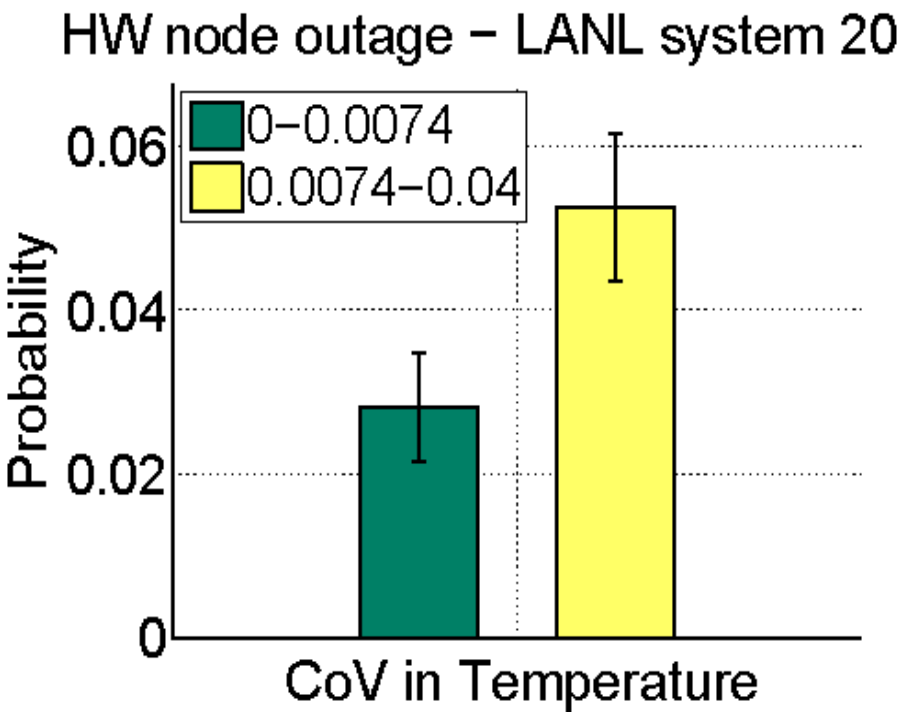

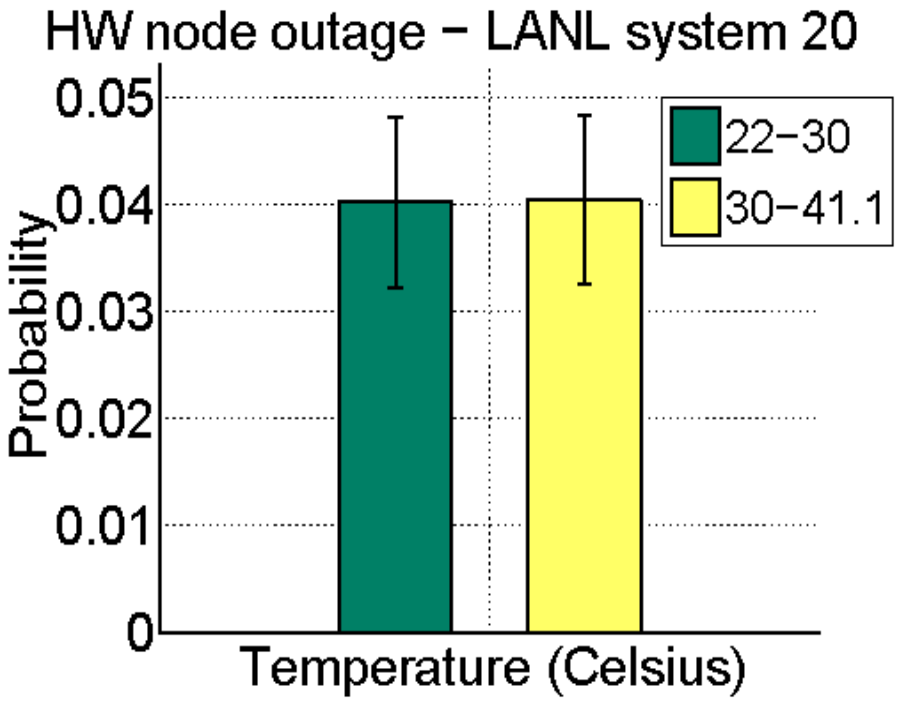

Variability in temperature has stronger effect on node reliability than average temperature

DO HIGHER TEMPS LEAD TO HIGHER NUM OF LSEs?

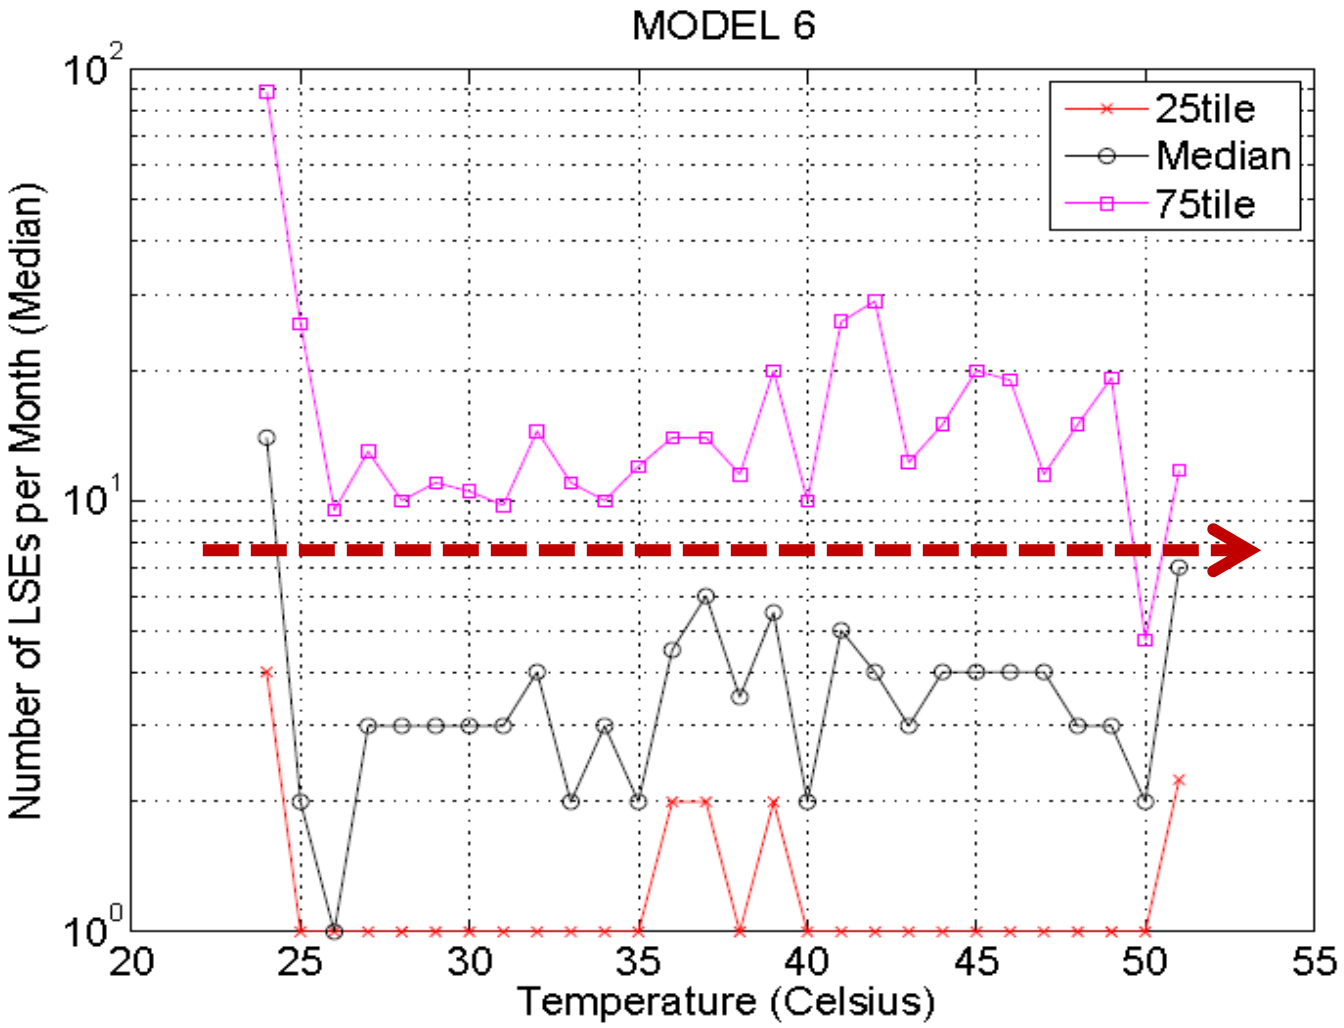

Hotter drives with errors don't develop more LSEs than colder drives!

HOW DOES AGE IMPACT LSEs?

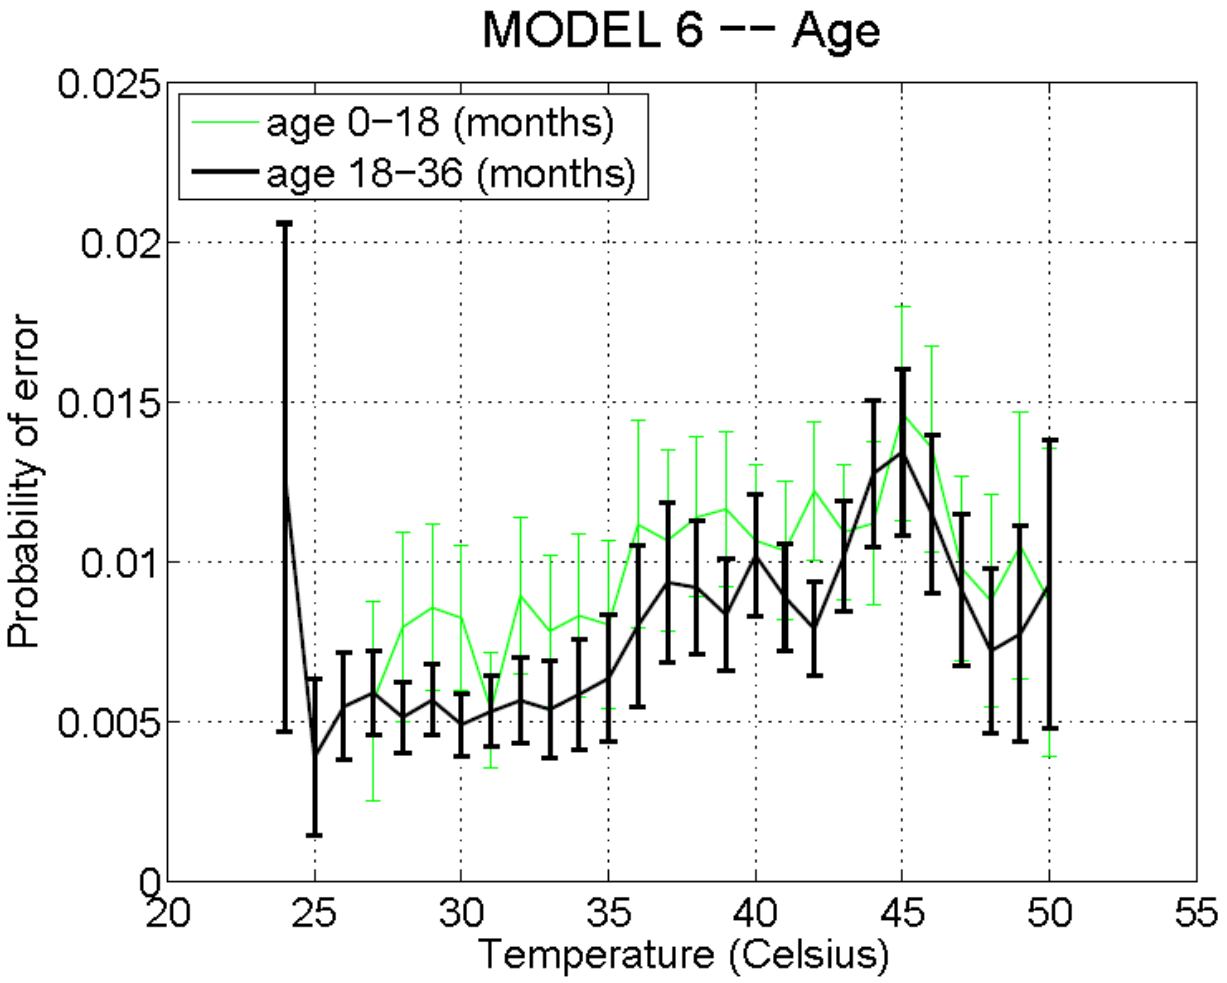

- Divide drives into two groups:
  - 0-18 months
  - 18-36 months
- No evidence that older drives are more sensitive to temperature than younger drives!
- Confirmed by ANOVA test

HOW DOES UTILIZATION IMPACT LSEs?

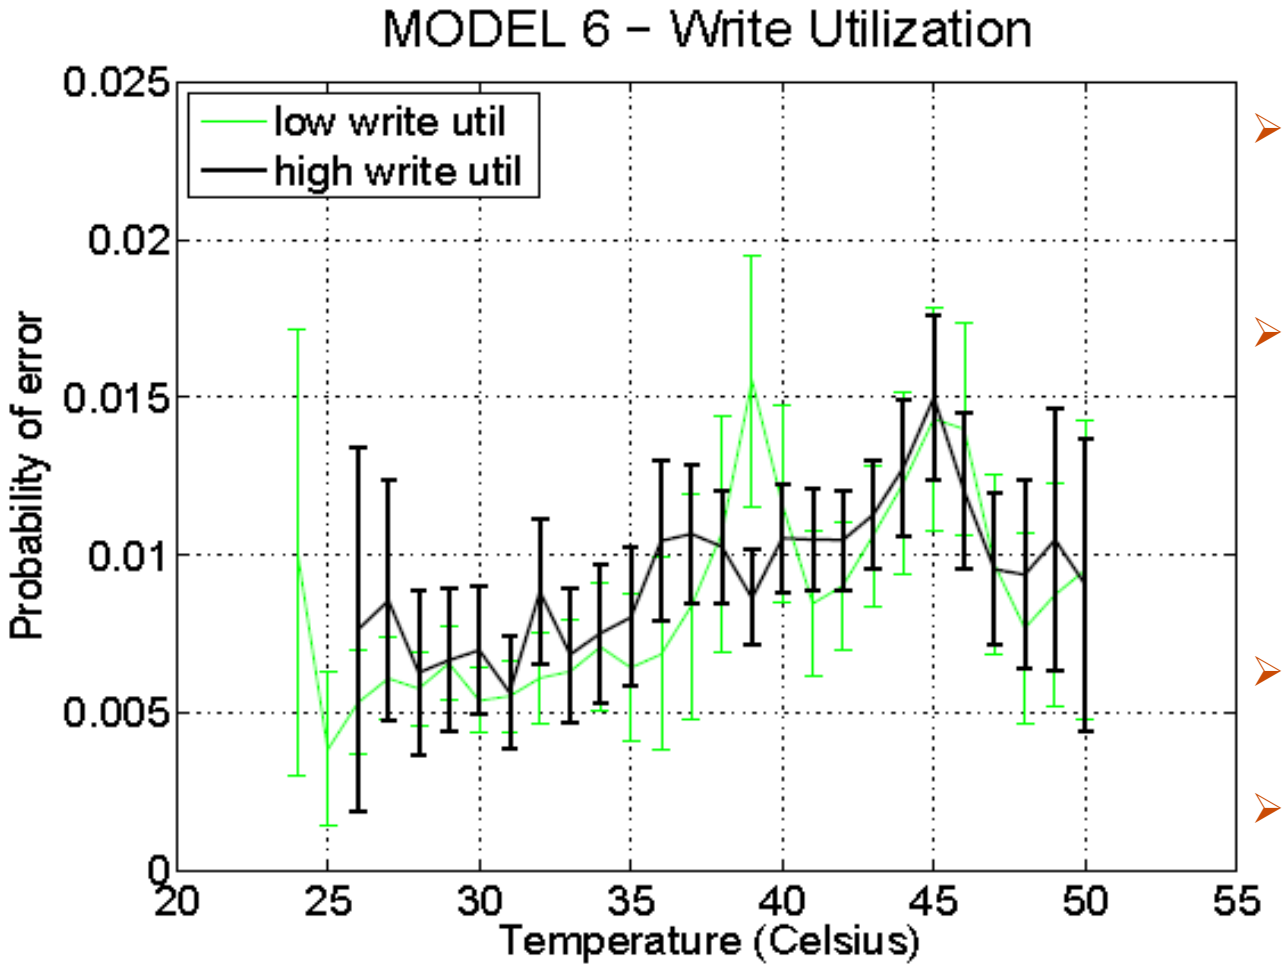

- Write util: # write operations per month
- No evidence that drives with higher util are more sensitive to temperature
- Similar for read util
- Confirmed by ANOVA test

TEMPERATURE AND PERFORMANCE

- **Experimental Study**
- **Thermal chamber (10˚C to 55˚C)**

- **Server: Dell PowerEdge R710**
  - Quad-core 2.26 GHz Intel Xeon
  - 8MB L3, 16GB DDR3 ECC
  - Ubuntu 10.04 Server (2.6.32-28-server Linux Kernel)

- **Workloads**
  - Synthetic microbenchmarks; macrobenchmarks
  - Stress different components **(disk, CPU, memory)**

TEMPERATURE AND DISK PERFORMANCE

- **Variety of hard disk drives**

Manufacturer	Model	Interface	Capacity	RPM
Hitachi	Deskstar	SATA	750GB	7200
Western Digital	Caviar	SATA	160GB	7200
Seagate	Barracuda	SATA	1TB	7200
Seagate	Constellation	SAS	500GB	7200
Seagate	Cheetah	SAS	73GB	15000
Fujitsu	MAX3073RC	SAS	73GB	15000
Hitachi	Ultrastar	SAS	300GB	15000

- **Disk-bound workloads**
  - Random-Read/Write
  - Sequential-Read/Write
  - Postmark
  - OLTP (TPC-C-based)
  - DSS (TPC-H-based)

Synthetic

Macrobenchmarks

TEMPERATURE AND MEMORY PERFORMANCE

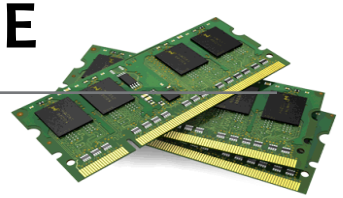

Performance

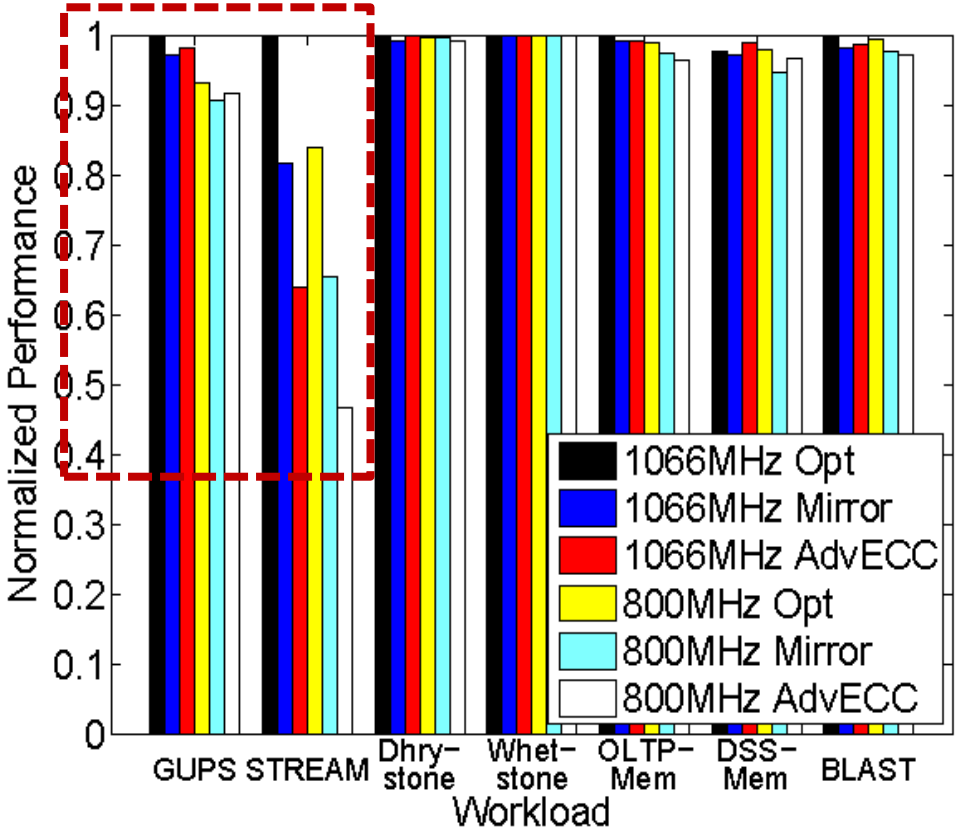

Power Consumption

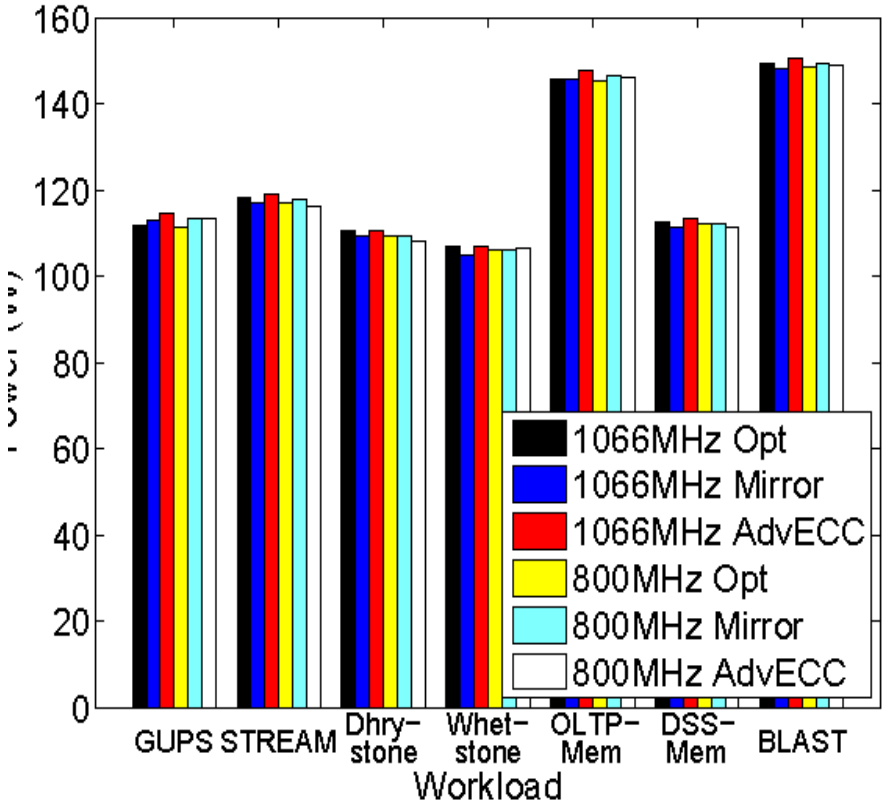

Significant throughput drops (up to 40%) for memory-bound microbenchmarks when activating protection mechanisms

AMBIENT vs INTERNAL TEMPERATURE: DISKS

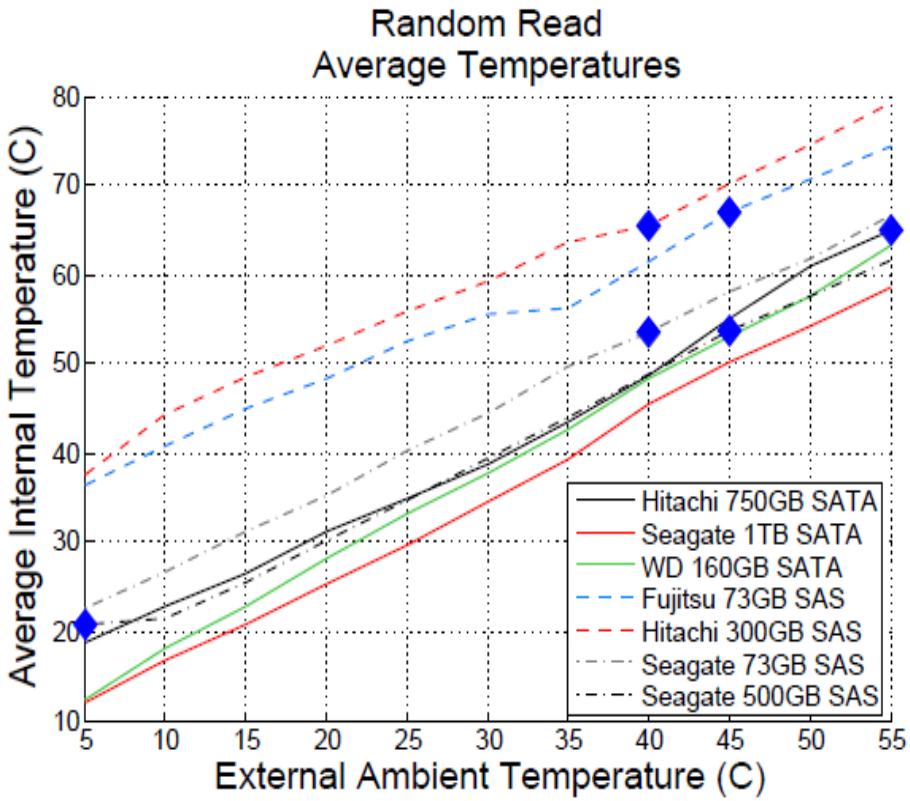

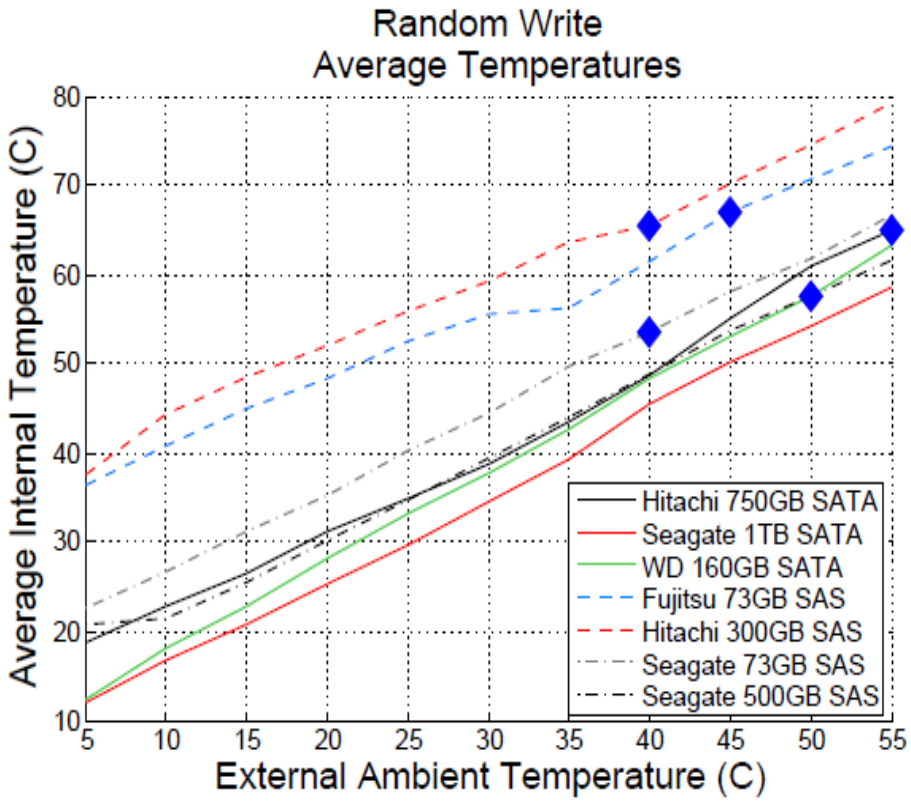

AMBIENT vs INTERNAL TEMPERATURE: CPU/MEMORY

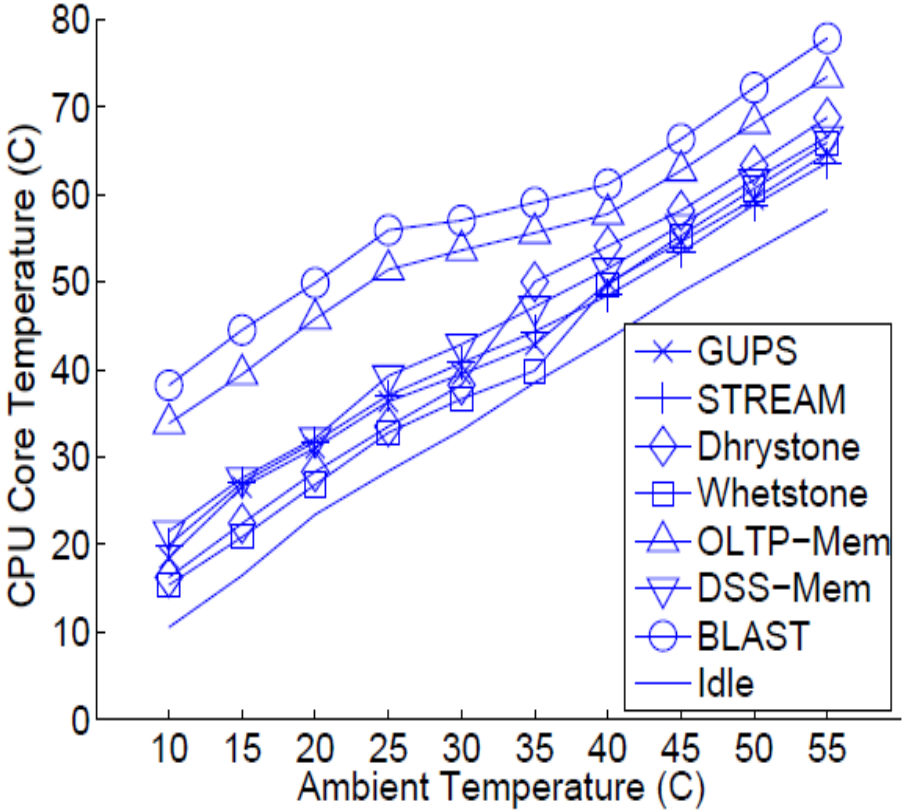

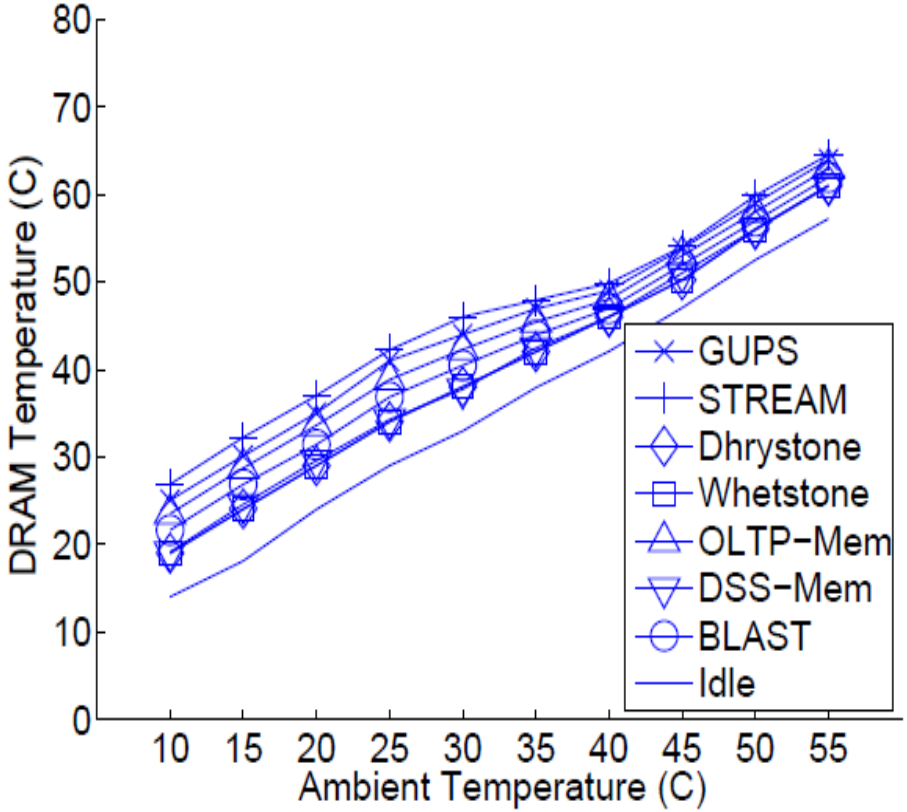

REDUCED SAFETY MARGINS: HOT SPOTS

- **hottest 5%** nodes: **5˚C** higher than median;
- **hottest 1%** nodes: **8-10˚C** higher than median

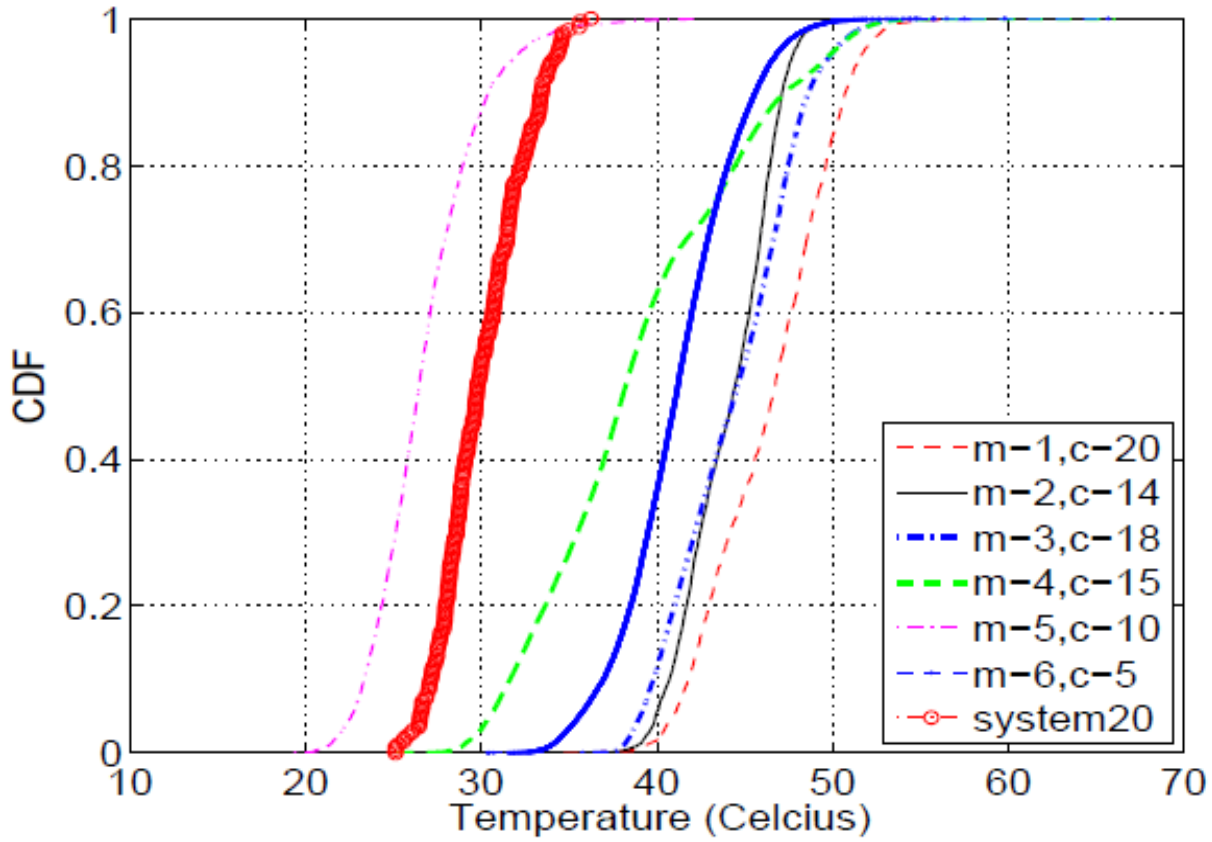

DELL POWEREDGE FAN – POWER PROFILE