Energy-aware Scheduling and Fault Tolerance Techniques for the Exascale Era

IV Congreso REDU

December 1st, 2016

Laércio Lima Pilla

laercio.pilla@ufsc.br

Federal University of Santa Catarina, Brazil

Agenda

EnergySFE Project

Global Scheduling

problems + research + interests

Fault Tolerance

problems + research + interests

EnergySFE Project

EnergySFE Project

STIC-AmSud Project Grant 99999.007556/2015-02

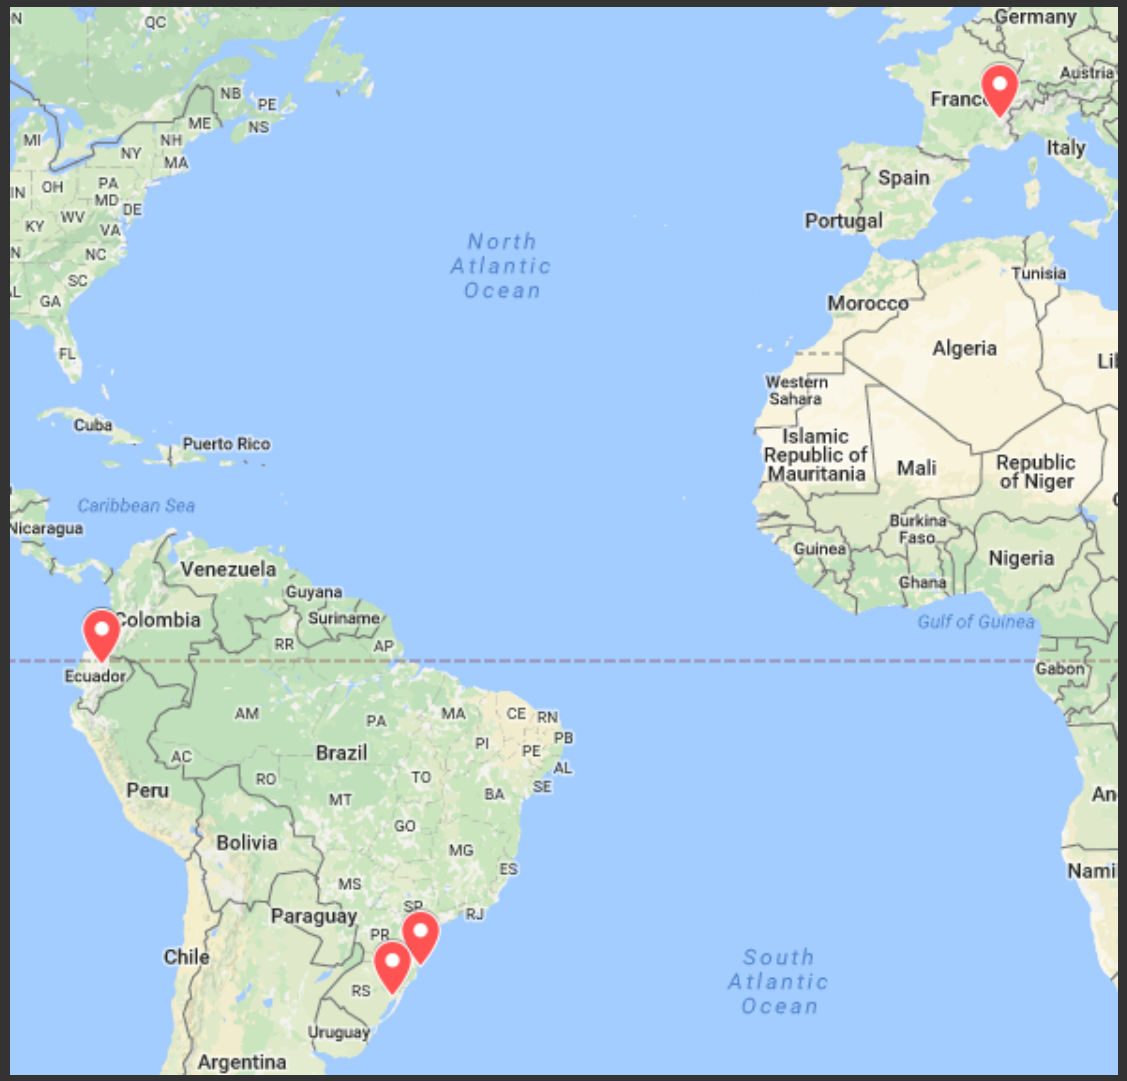

EnergySFE Project

Members

Enrique Vinicio Carrera - EC

François Broquedis - FR

Frédéric Desprez - FR

Jean-François Méhaut - FR

Laércio Lima Pilla - BR

Lucas Mello Schnorr - BR

Márcio Bastos Castro - BR

Mario Antônio Dantas - BR

Pablo Francisco Ramos - EC

Paolo Rech - BR

Philippe O. A. Navaux - BR

Vanessa C. Vargas - EC

UFSC

ECL

http://energysfe.ufsc.br/

EnergySFE Project

Energy-aware Scheduling and Fault Tolerance Techniques for the Exascale Era

Increasing performance with limited energy

Sunway TaihuLight – credits: www.top500.org

EnergySFE Project

Research Questions

1. How to schedule tasks and threads that compete for resources with different constraints while considering the complex hierarchical organization of future Exascale supercomputers?

EnergySFE Project

Research Questions

2. How to tolerate faults without incurring in too much overhead in future Exascale supercomputers?

EnergySFE Project

Research Questions

3. How scheduling and fault tolerance approaches can be adapted to be energy-aware?

UFSC

ECL

Global Scheduling

problems + research + interests

Global Scheduling Problems

objectives

shortest execution time

highest throughput

highest utilization rate

fairness

…

Global Scheduling Problems

load imbalance

poor initial mapping

load dynamicity

platform sharing

heterogeneous platforms

DVFS changes

...

UFSC

ECL

Global Scheduling Problems

communication slowdown

poor initial mapping

complex communication patterns

hierarchical machine topologies

platform sharing

heterogeneous platforms

...

Global Scheduling Problems

energy wastage

long computing times

low resource usage

excessive data movement

slow scheduling algorithms

…

ECL

Global Scheduling Research

What we want with scheduling

balanced loads

optimized communications

fast scheduling algorithms

least migrations possible

Global Scheduling Research

Periodic load balancing

principle of persistence

topology-aware load balancing

energy-aware load balancing

Global Scheduling Research

topology-aware load balancing

NucoLB

greedy algorithm + NUMA node latency

HwTopoLB

refinement algorithm + latency & bw + convergence

HierarchicalLB

LB at node level and at whole machine level

UFSC

ECL

Global Scheduling Research

Application: Ondes3D

irregular loads

dynamic loads

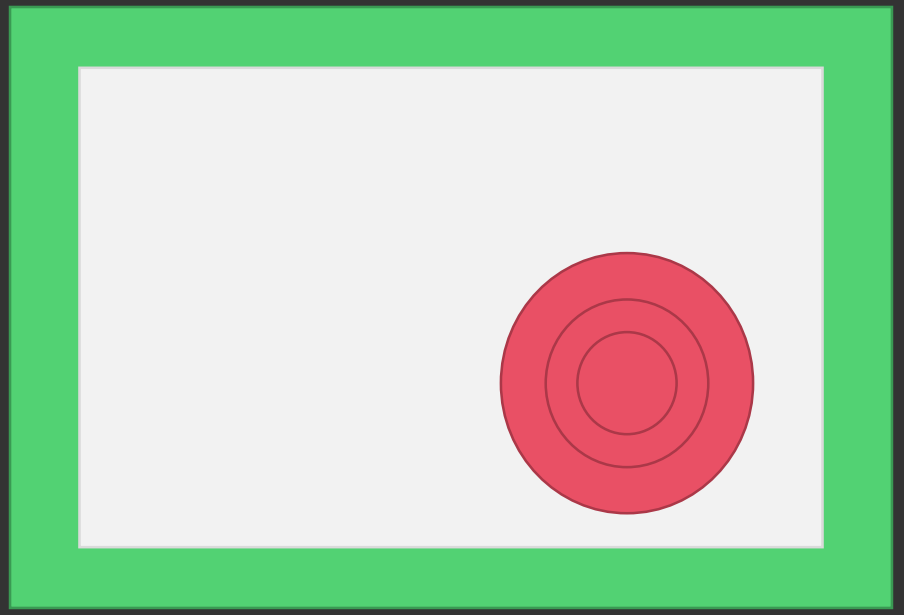

UFSC

ECL

Global Scheduling Research

Application: Ondes3D

irregular loads

dynamic loads

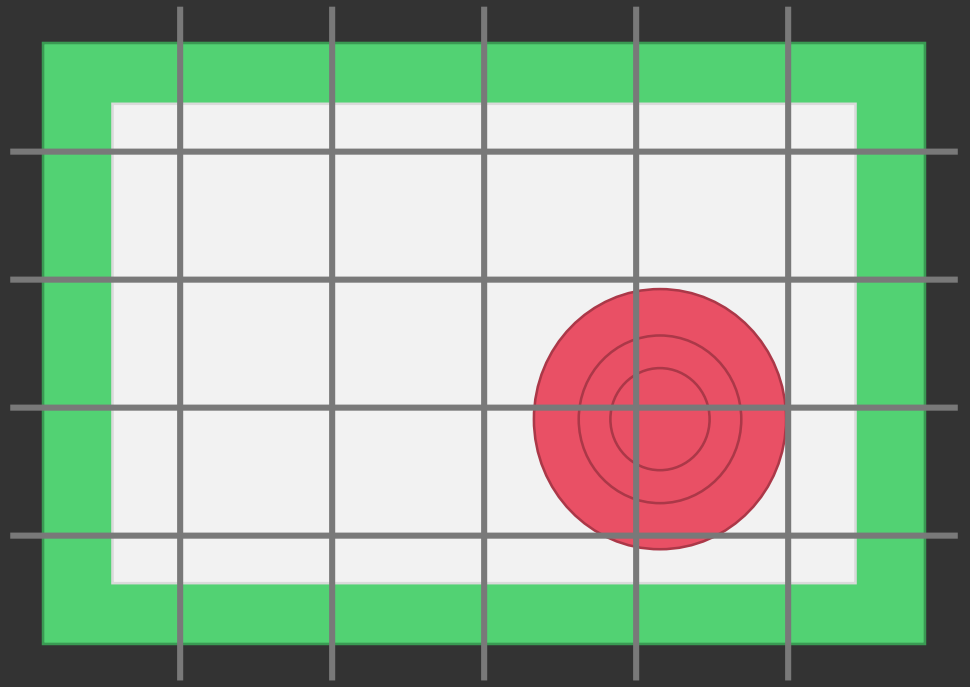

UFSC

ECL

Global Scheduling Research

Application: Ondes3D

OD = overdecomposition (16 tasks / core)

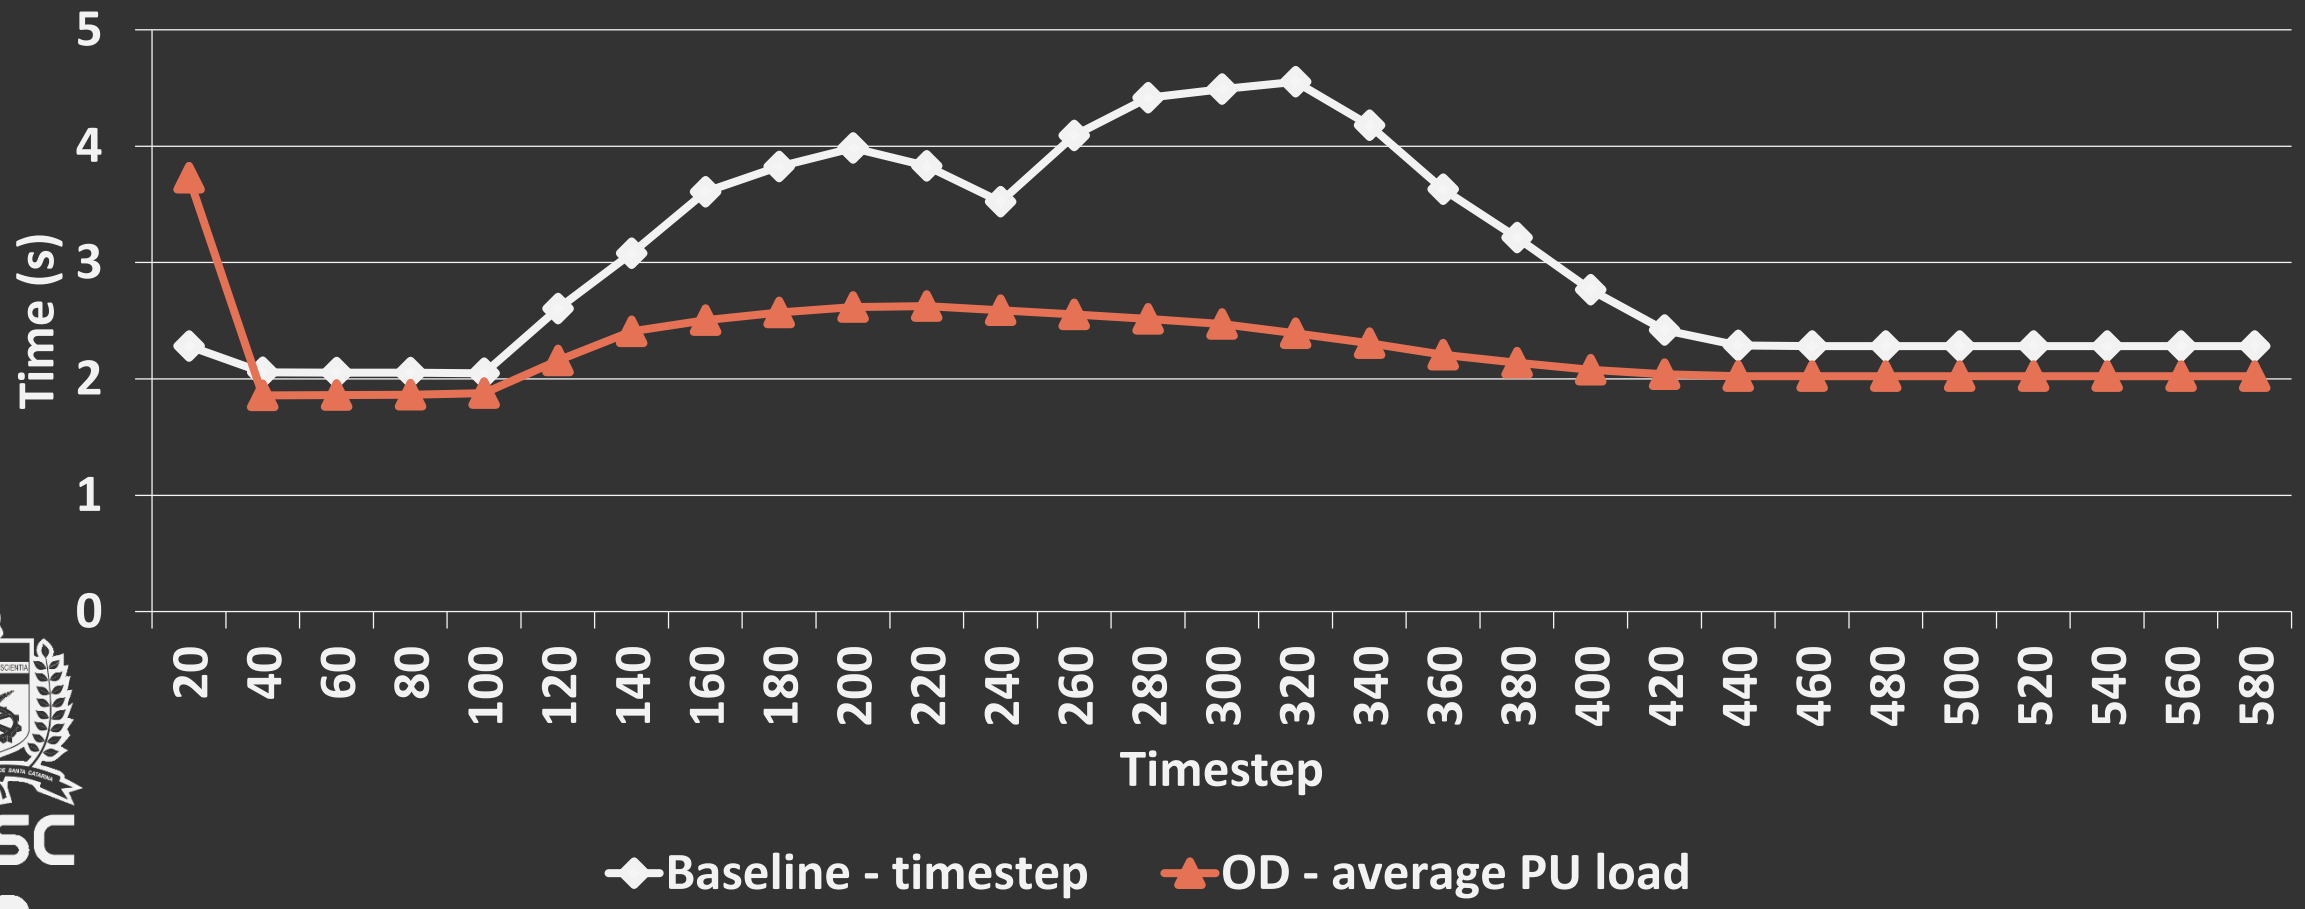

UFSC

ECL

Tesser et al. *Improving the performance of seismic wave simulations with dynamic load balancing*. PDP 2014.

Global Scheduling Research

Application: Ondes3D

OD = overdecomposition (16 tasks / core)

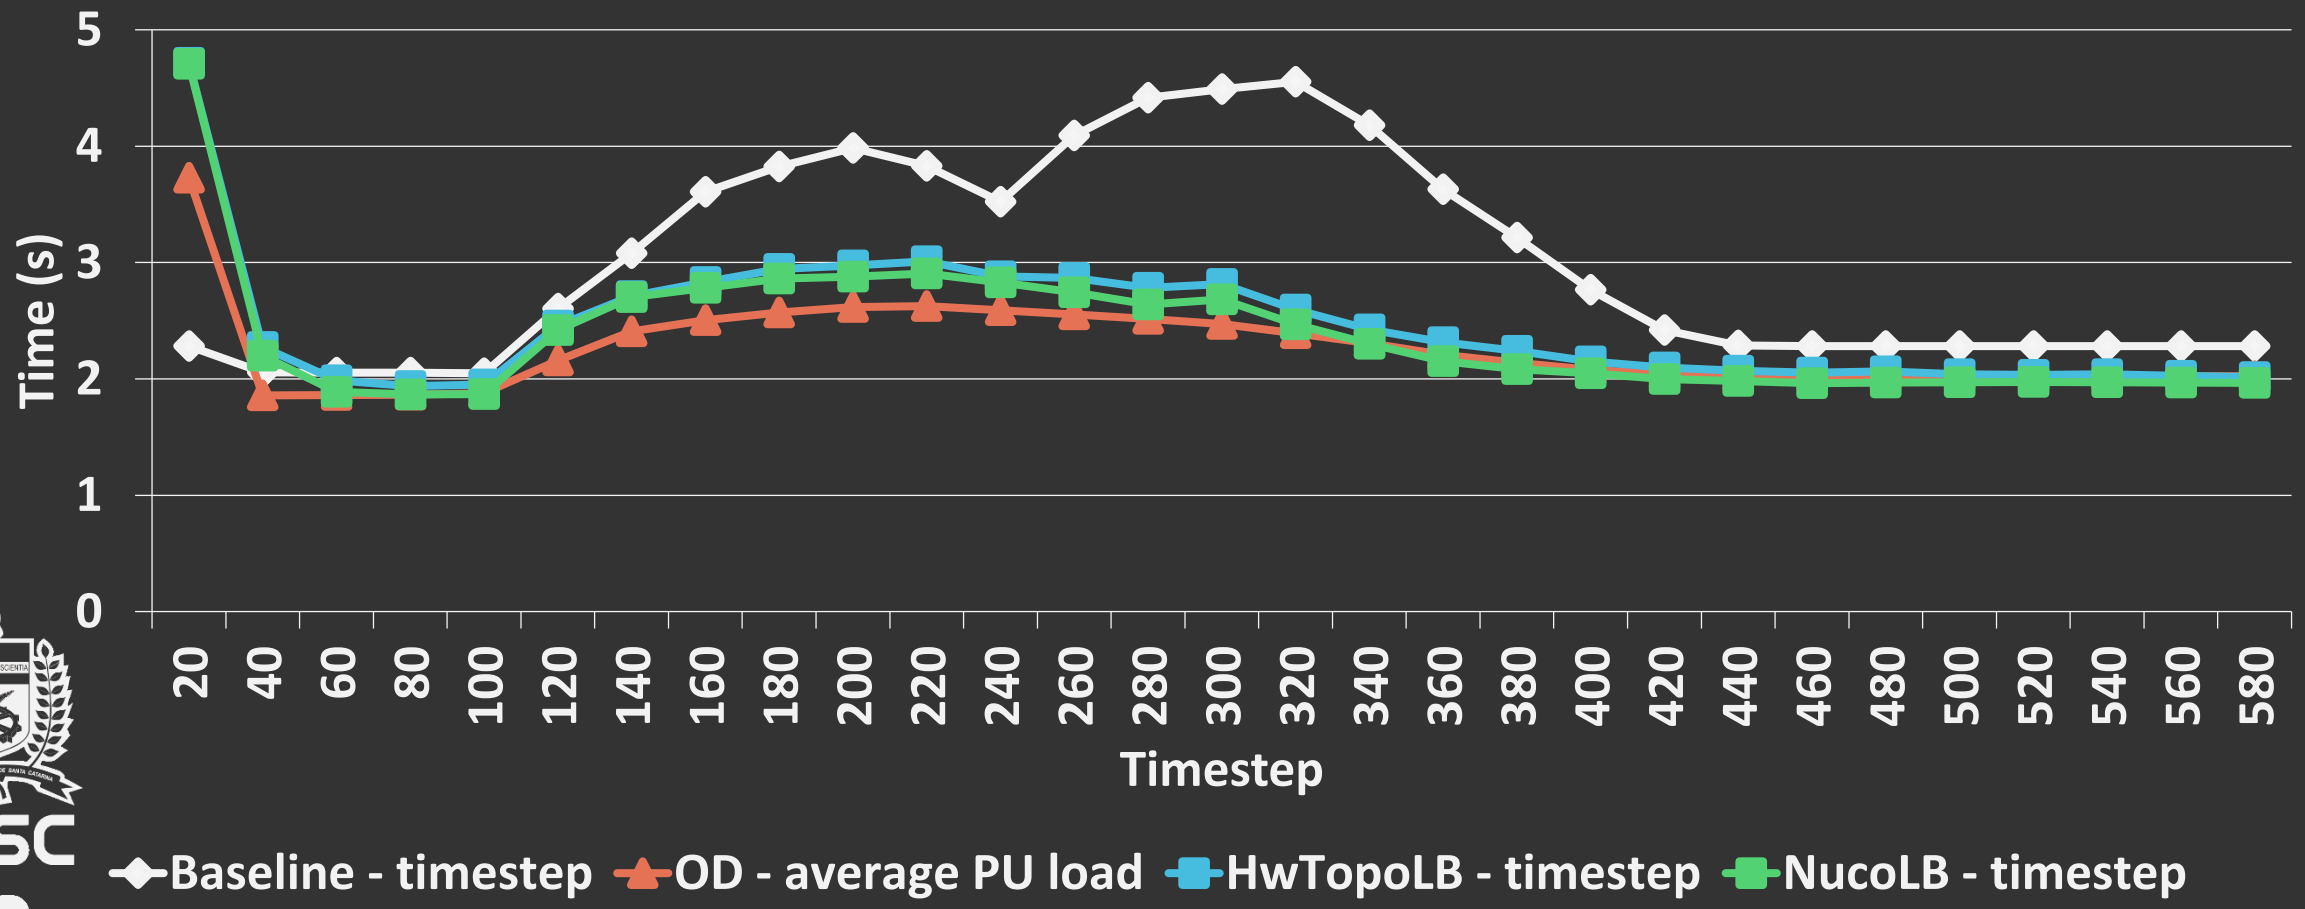

UFSC

ECL

Tesser et al. *Improving the performance of seismic wave simulations with dynamic load balancing*. PDP 2014.

Global Scheduling Research

Application: LeanMD

Weak scalability: Total execution time

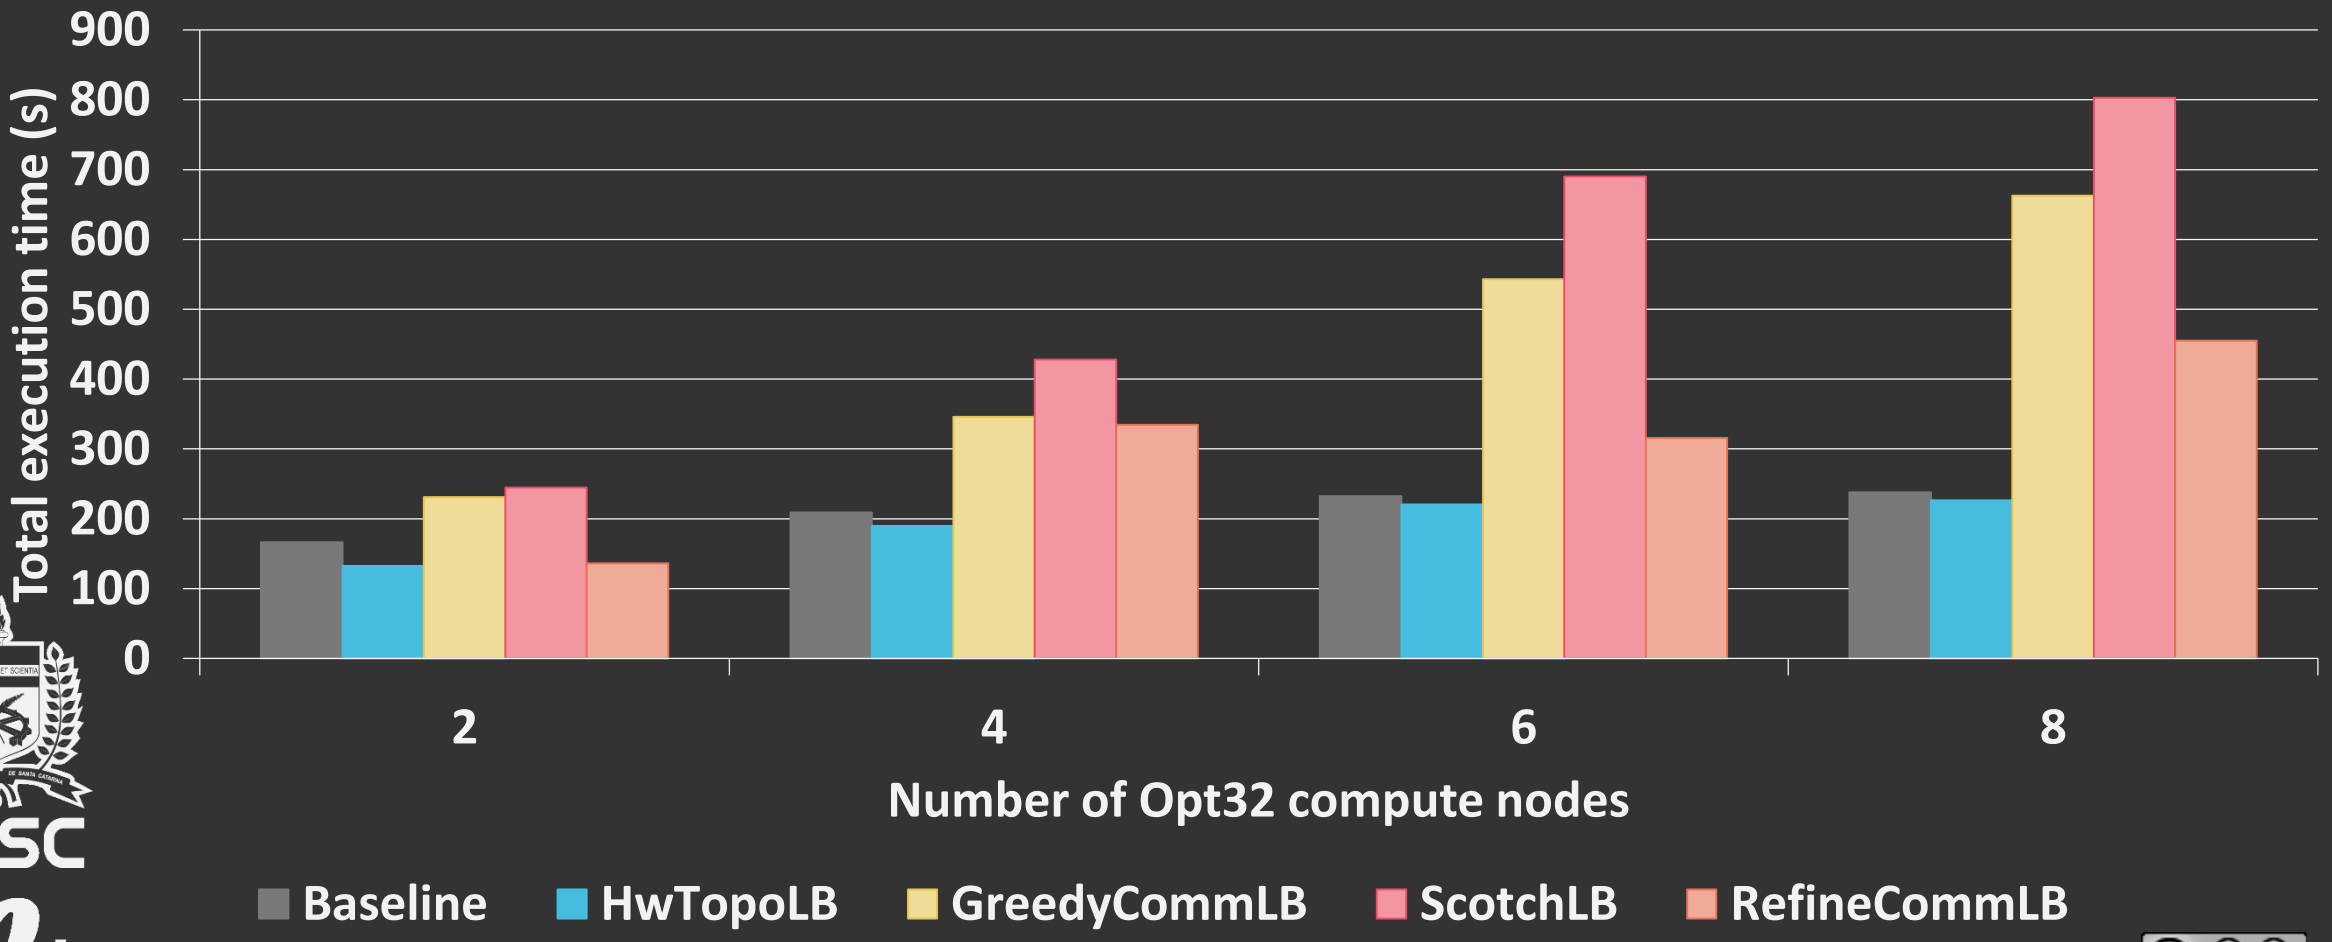

UFSC

ECL

Pilla et al. *A topology-aware load balancing algorithm for clustered hierarchical multi-core machines*. FGCS, 2014.

Global Scheduling Research

Strong scalability: LeanMD

Speedup over the baseline on 2 CNs

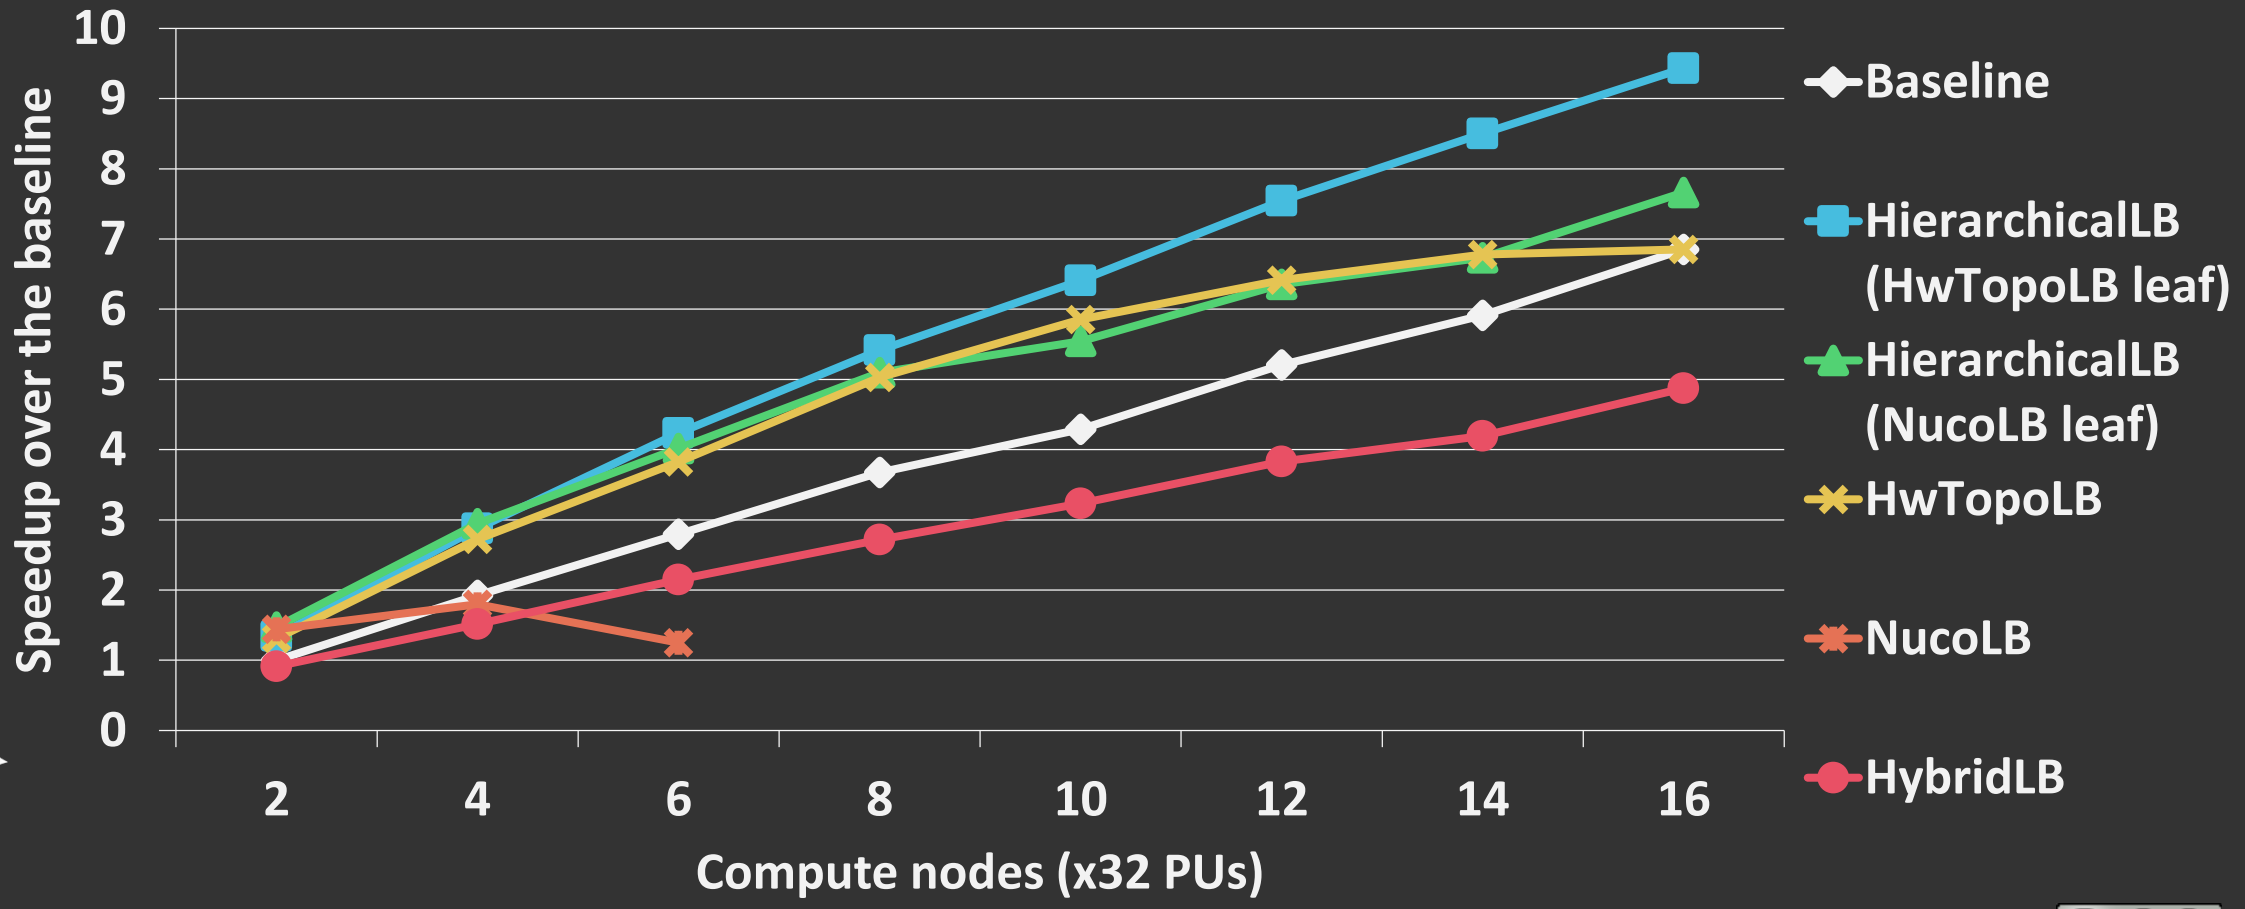

UFSC

ECL

Pilla et al. *A topology-aware load balancing algorithm for clustered hierarchical multi-core machines*. FGCS, 2014.

Global Scheduling Research

energy-aware load balancing

LB + DVFS on underloaded resources

Fine-Grained EnergyLB

per core

Coarse-Grained EnergyLB

per socket + hierarchical algorithm

UFSC

ECL

Global Scheduling Research

Example: Lulesh + FG EnergyLB

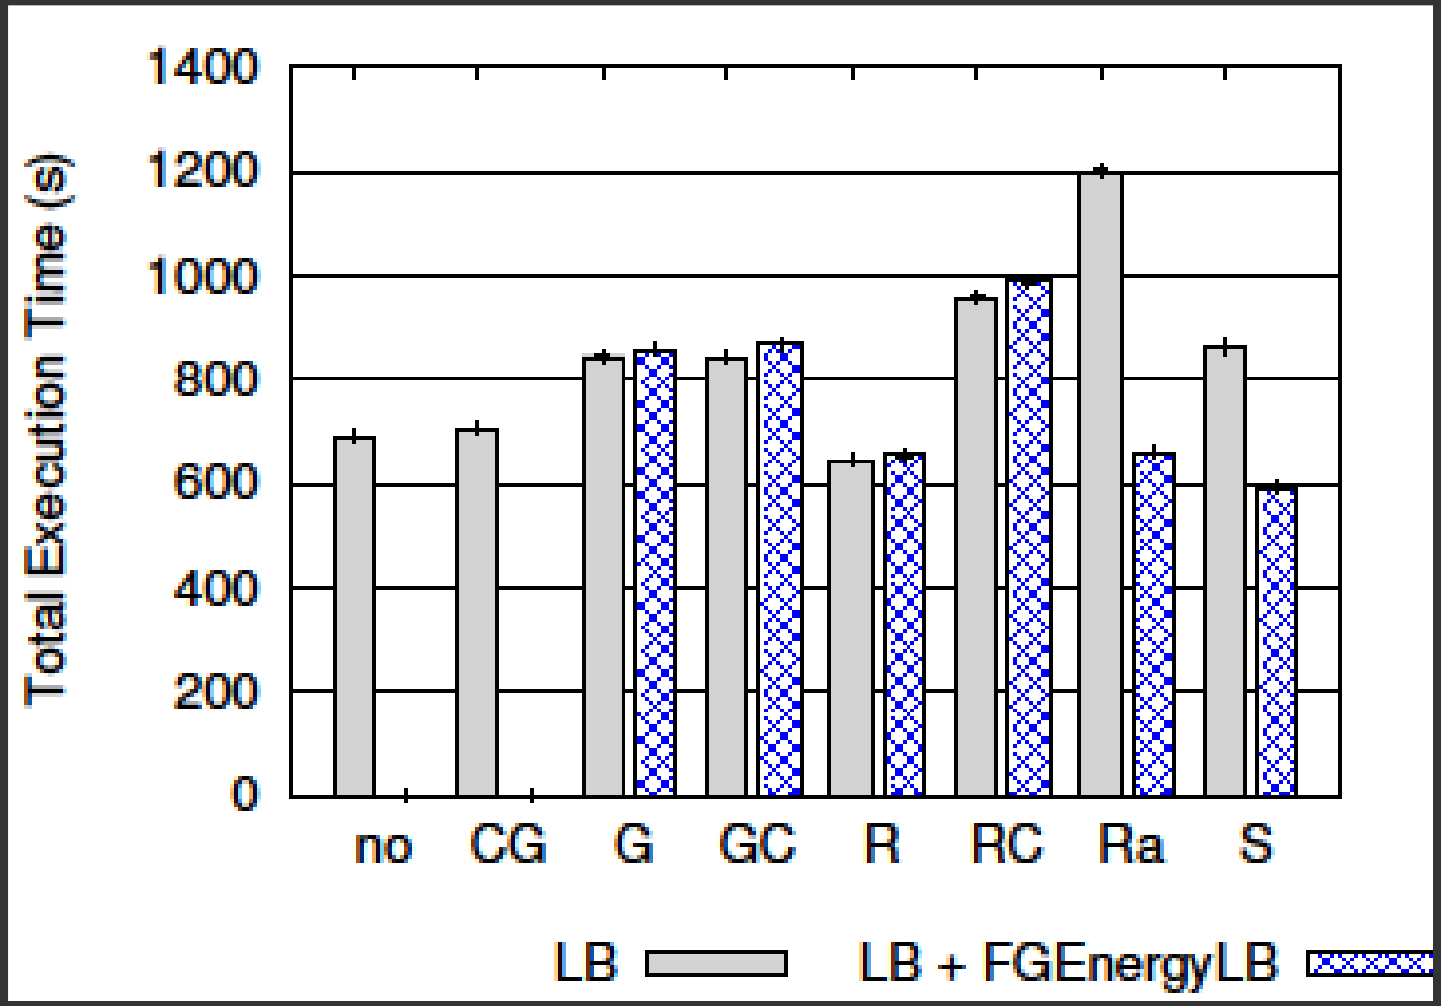

Padoin et al. *Saving energy by exploiting residual imbalances on iterative applications*. HiPC, 2014.

Global Scheduling Research

Example: Lulesh + FG EnergyLB

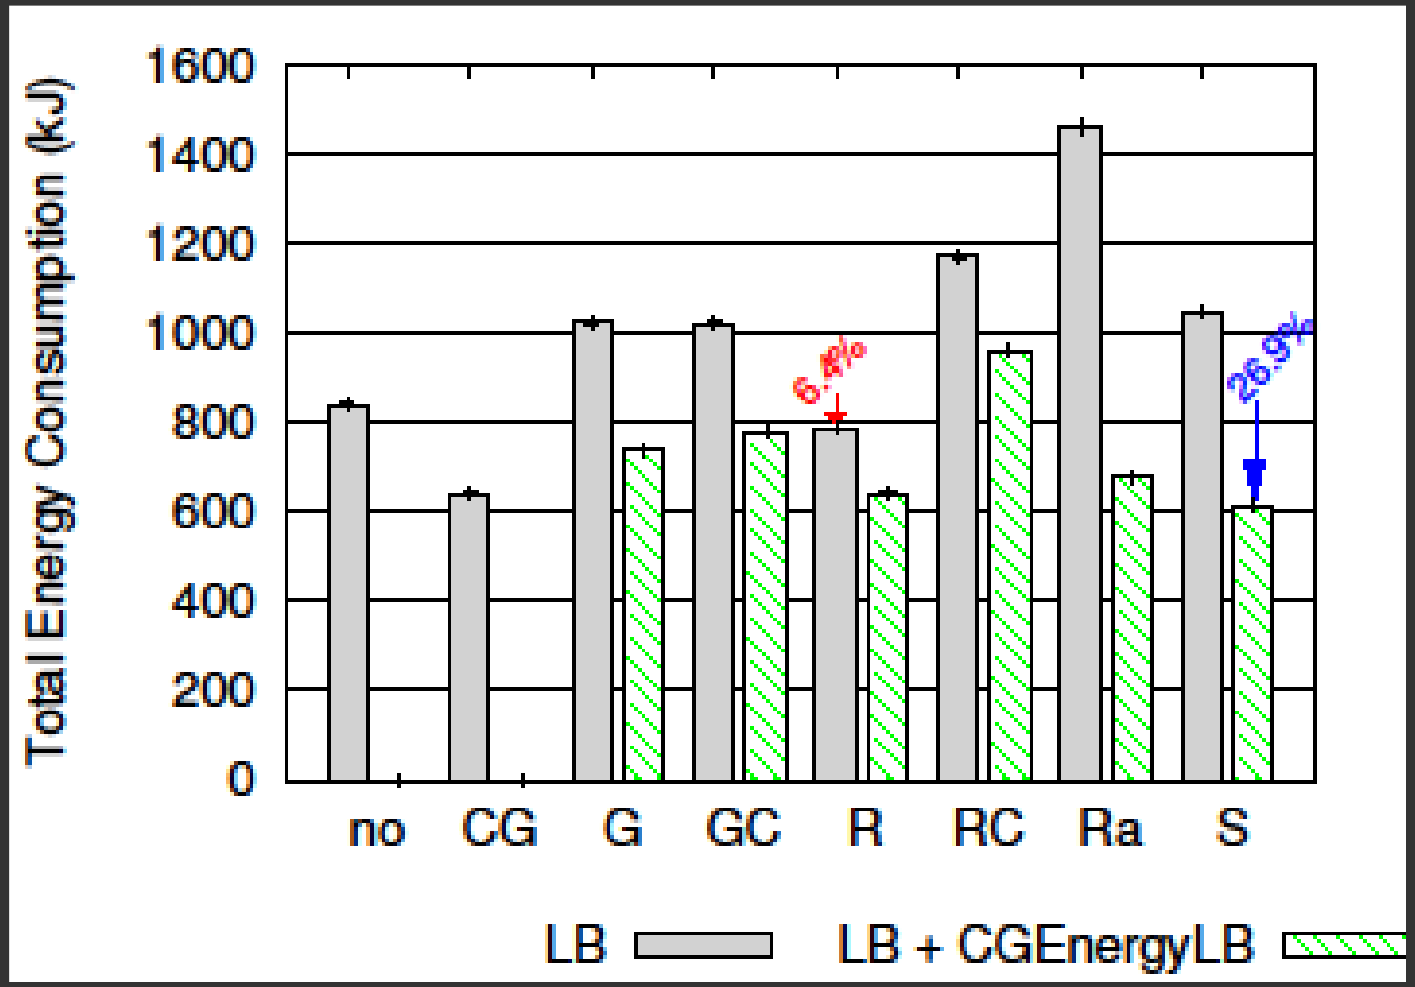

Padoin et al. *Saving energy by exploiting residual imbalances on iterative applications*. HiPC, 2014.

Global Scheduling Interests

+ hierarchical algorithms [+ energy]
+ distributed algorithms [+ energy]
+ platforms for experiments [+ energy]
+ applications

UFSC

ECL

Fault Tolerance

problems + research + interests

Credits: Paolo Rech

Fault Tolerance Problems

failure: systems stops working as expected

fault → error → failure

Fault Tolerance

use redundancy to stop failures from occurring

Fault Tolerance Problems

usual fault tolerance schemes

checkpoint & restart

cannot notice corrupted data

duplication with comparison

expensive in resource (time & energy)

UFSC

ECL

Fault Tolerance Problems

Soft Errors: the device is not permanently damaged, but a particle may generate:

One or more bit-flips
Single Event Upset (SEU)
Multiple Bit Upset (MBU)

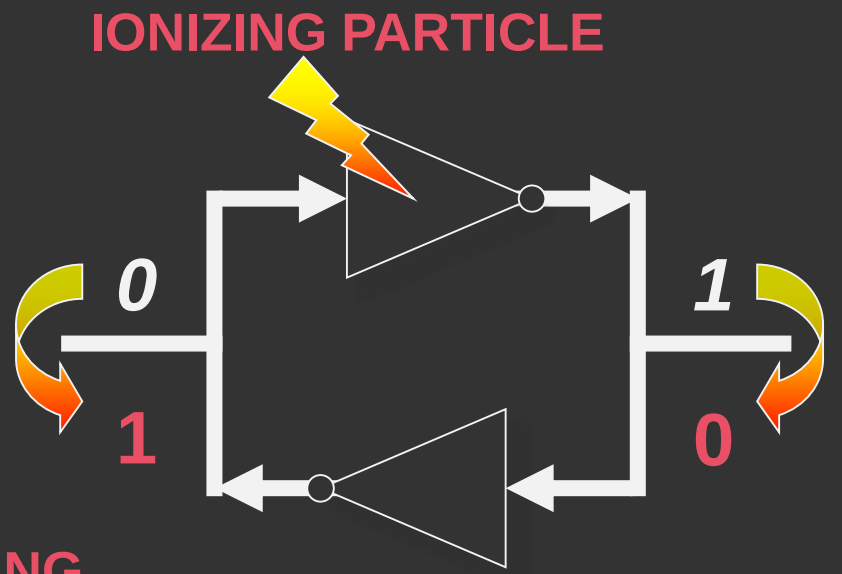

Transient voltage pulse
Single Event Transient (SET)

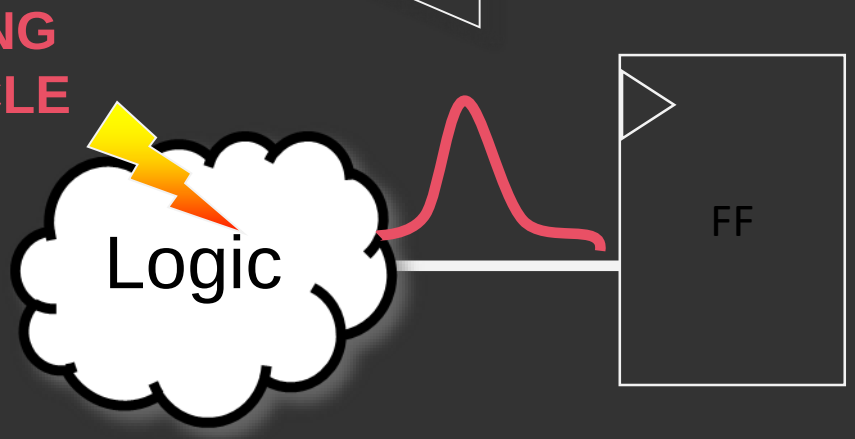

UFSC
ECL

Fault Tolerance Problems

Silent Data Corruption

data caches

register files

ALU

scheduler

Crash

instruction cache

scheduler

PCIe bus

Fault Tolerance Research

LANSCE
Los Alamos Neutron Science Center

UFSC

ECL

Fault Tolerance Research

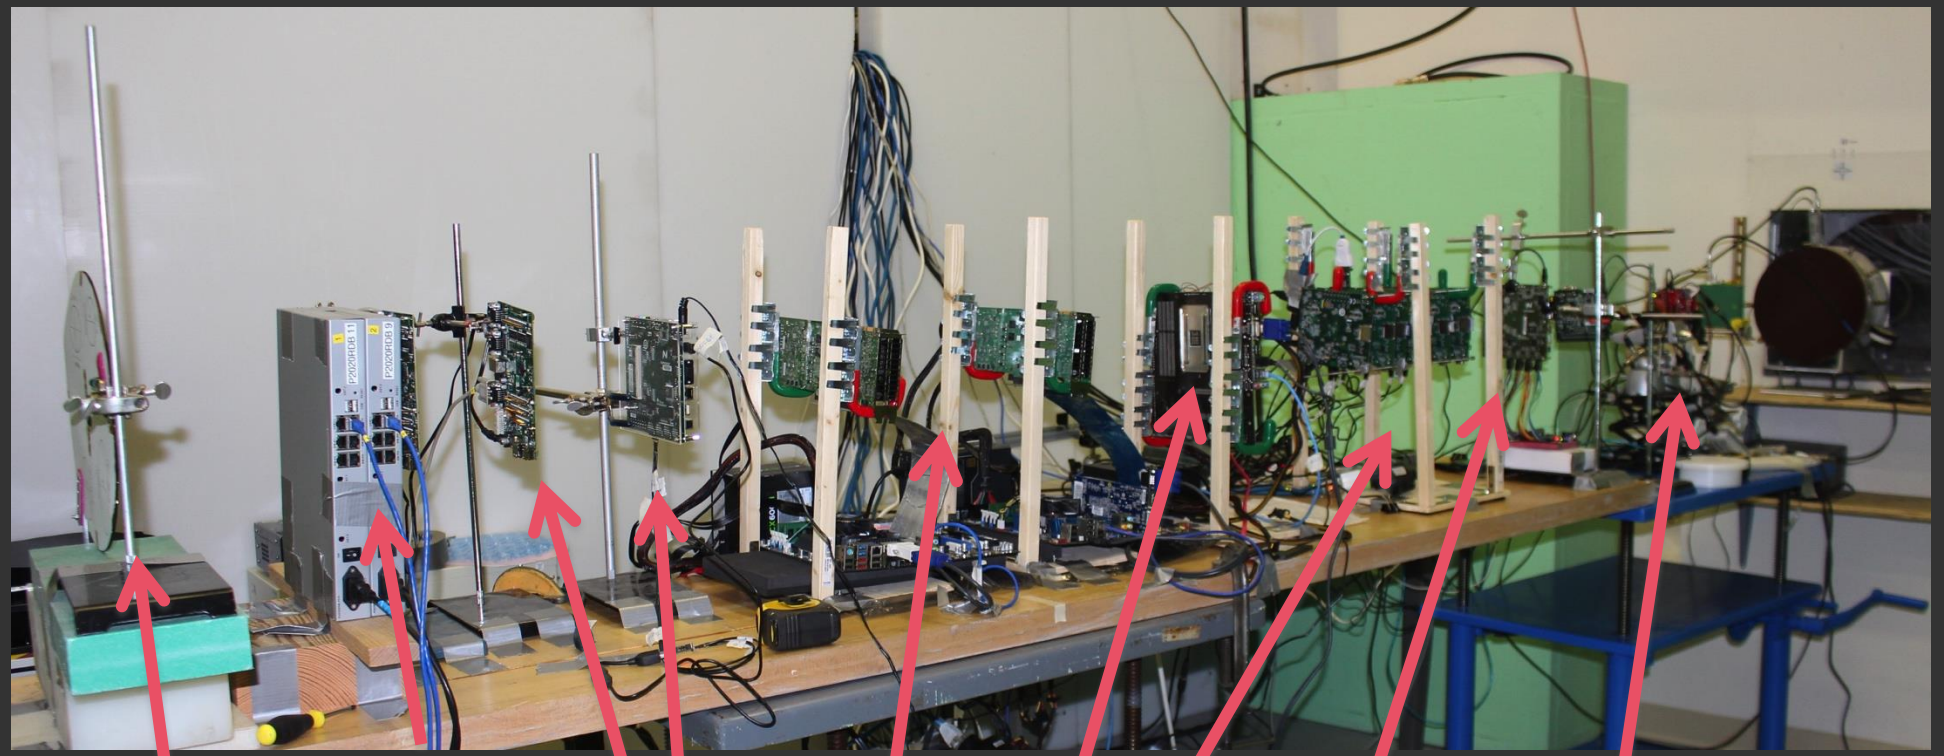

UFSC

ECL

Fault Tolerance Research

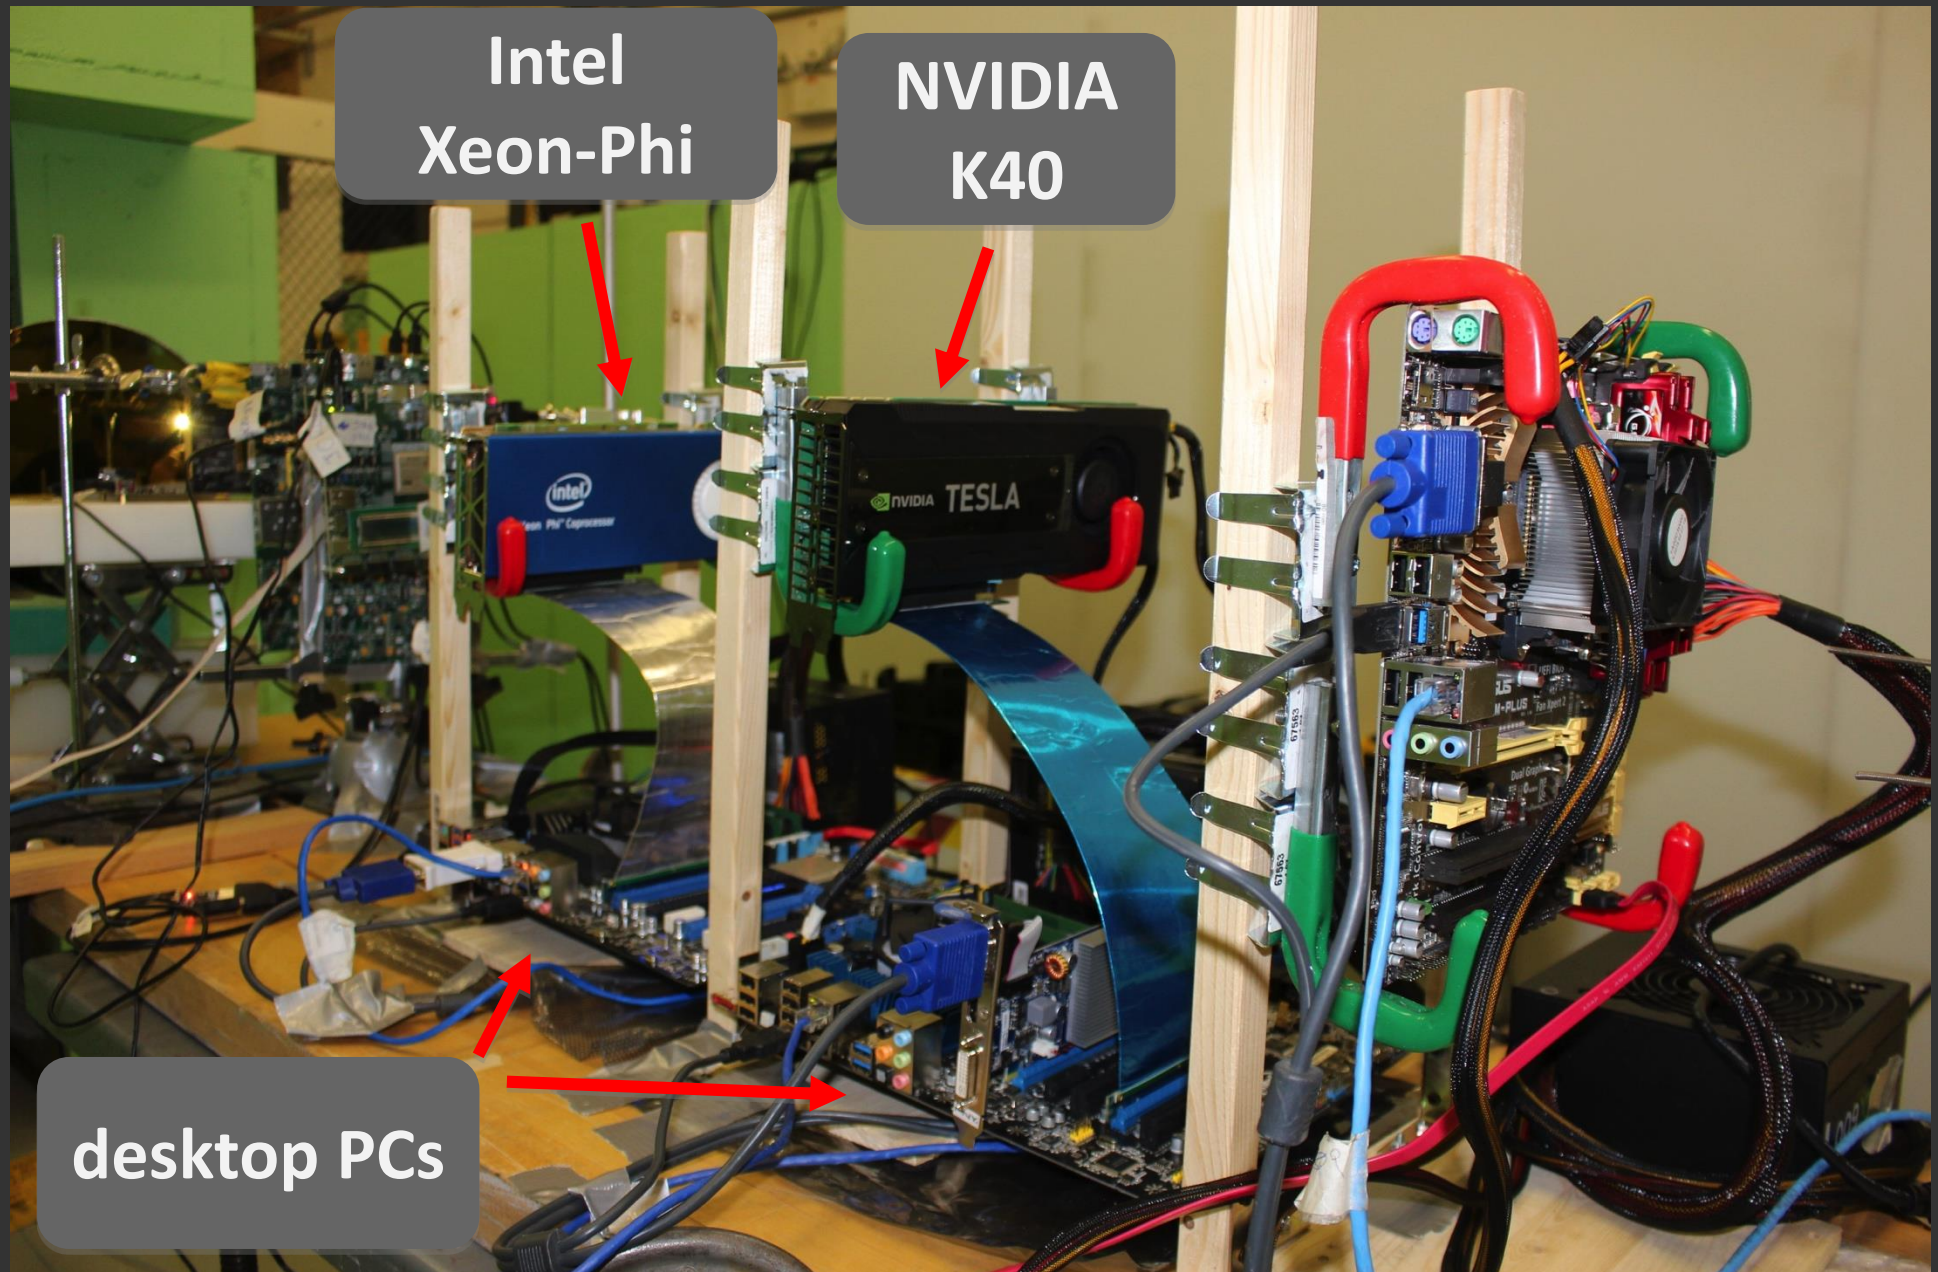

UFSC

ECL

Fault Tolerance Research

@LANSCE 1.8x10^6 n/(cm^2 h)

@NYC 13 n/(cm^2 h)

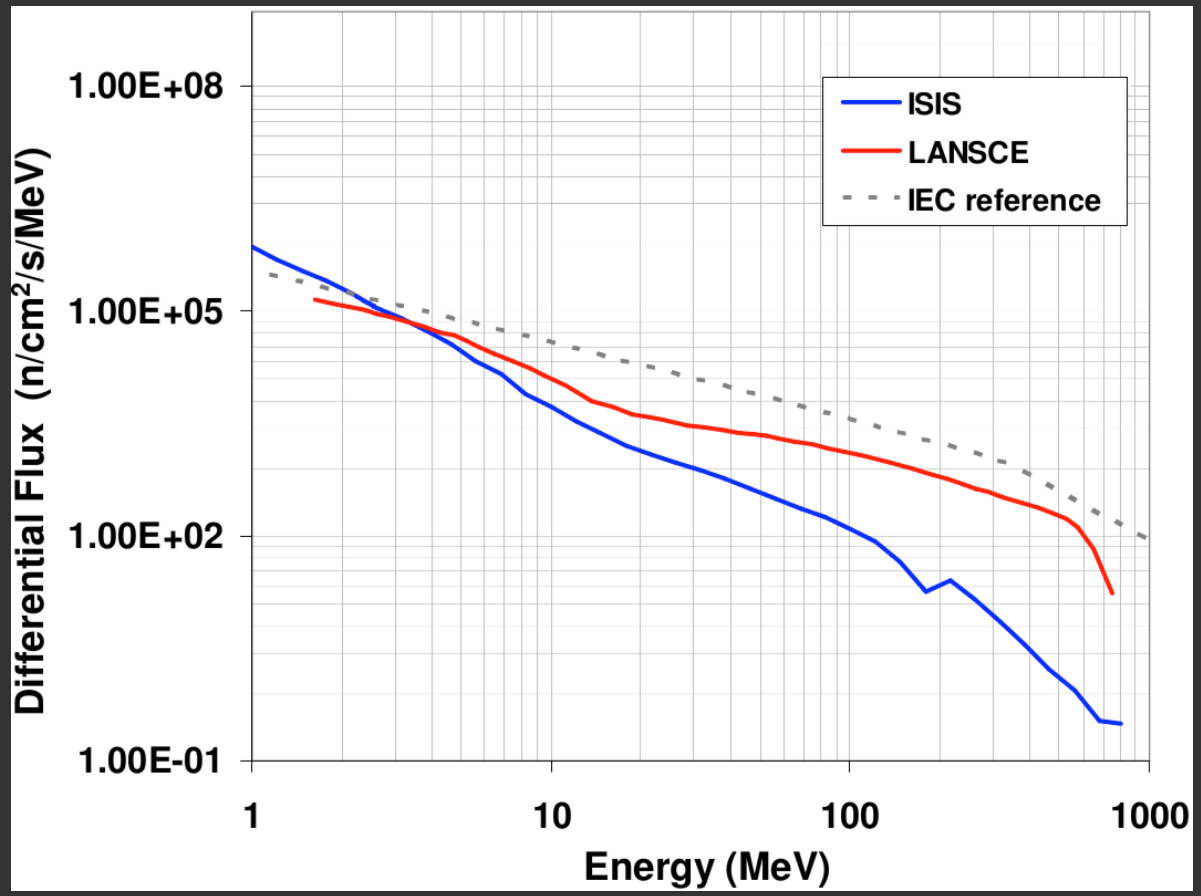

Fault Tolerance Research

Parallel algorithms' reliability

SDC rates vary ~3 orders of magnitude

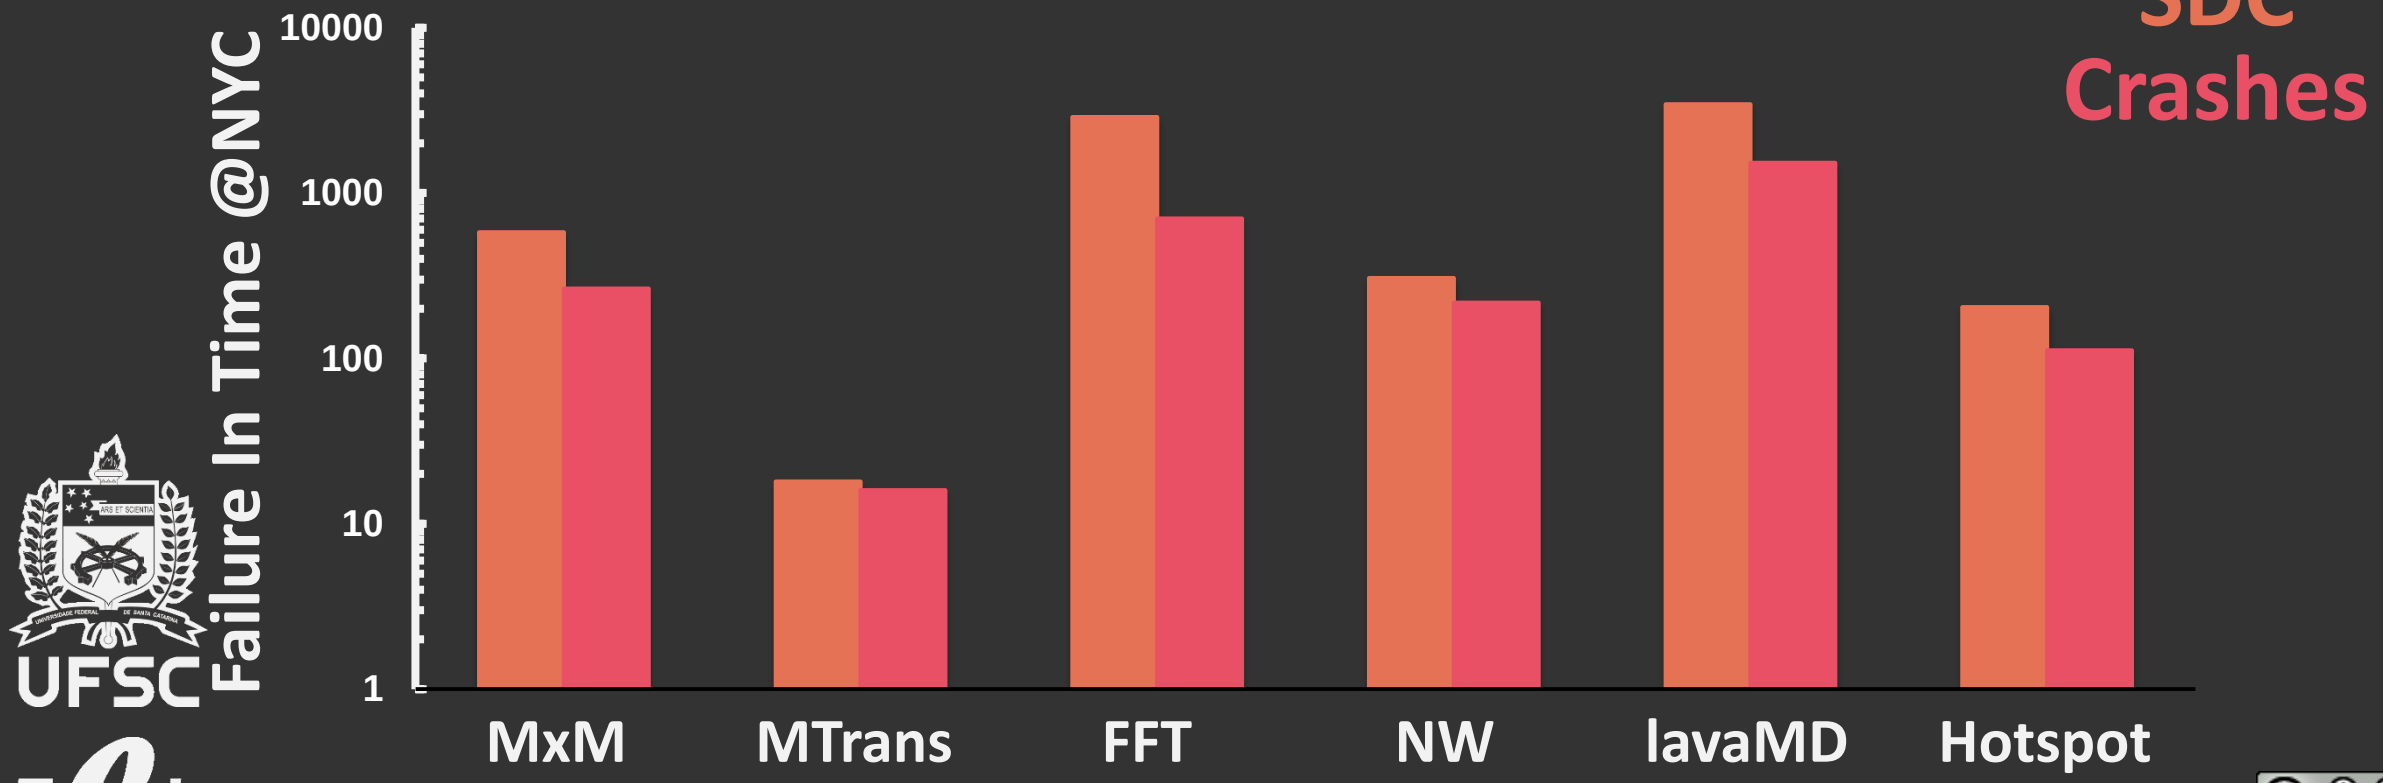

Fault Tolerance Research

ECC reduces the SDC FIT of ~1 order of magnitude (there is almost no code dependence)

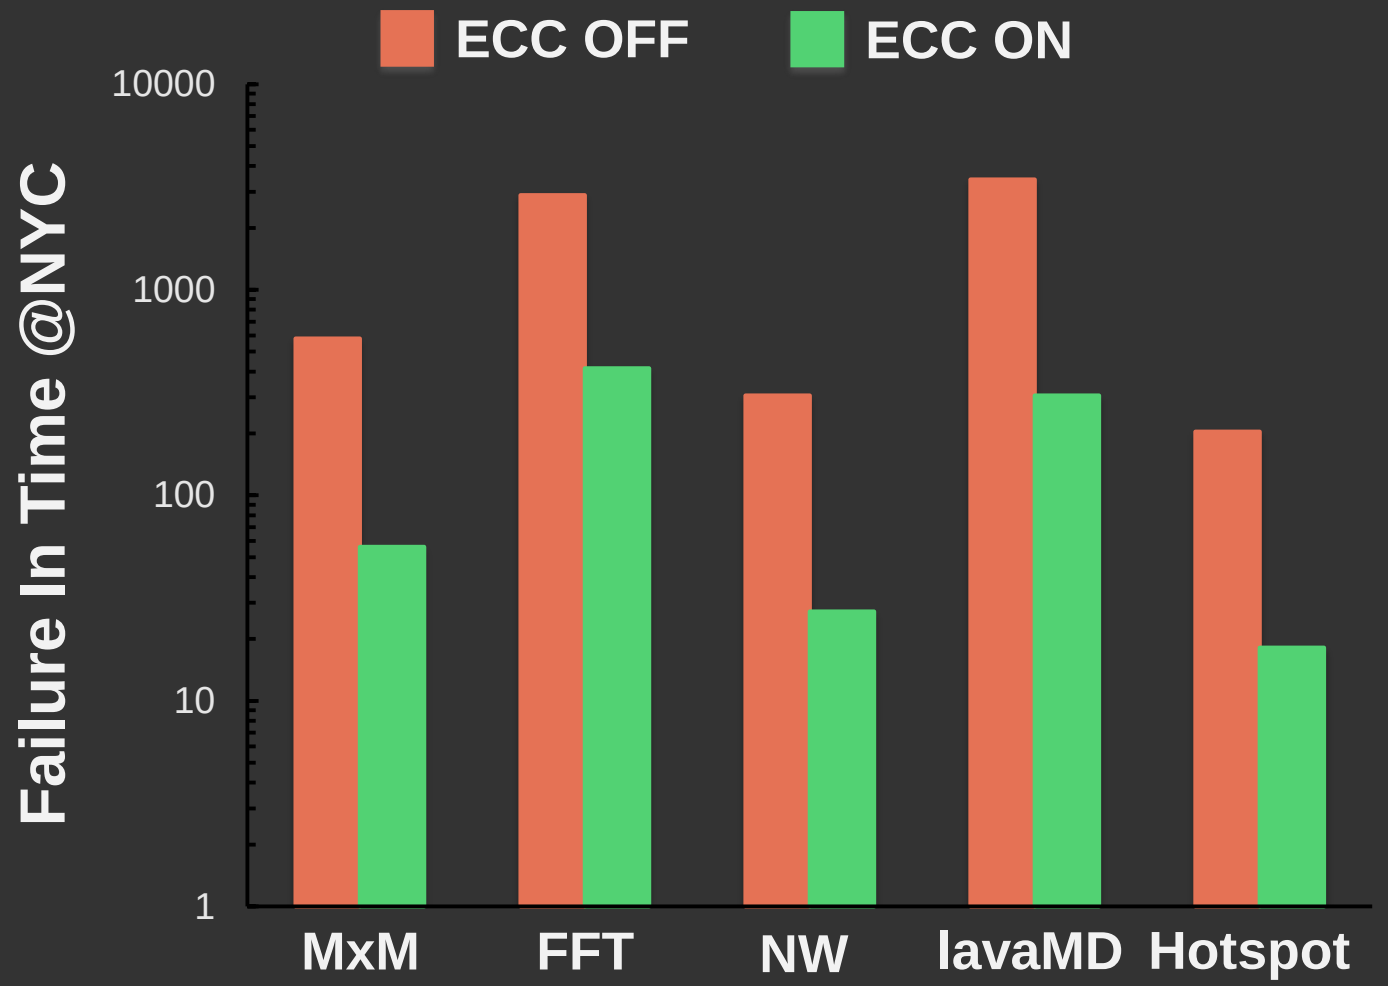

Fault Tolerance Research

ECC increases the Crash FIT of about 50% (there is almost no code dependence)

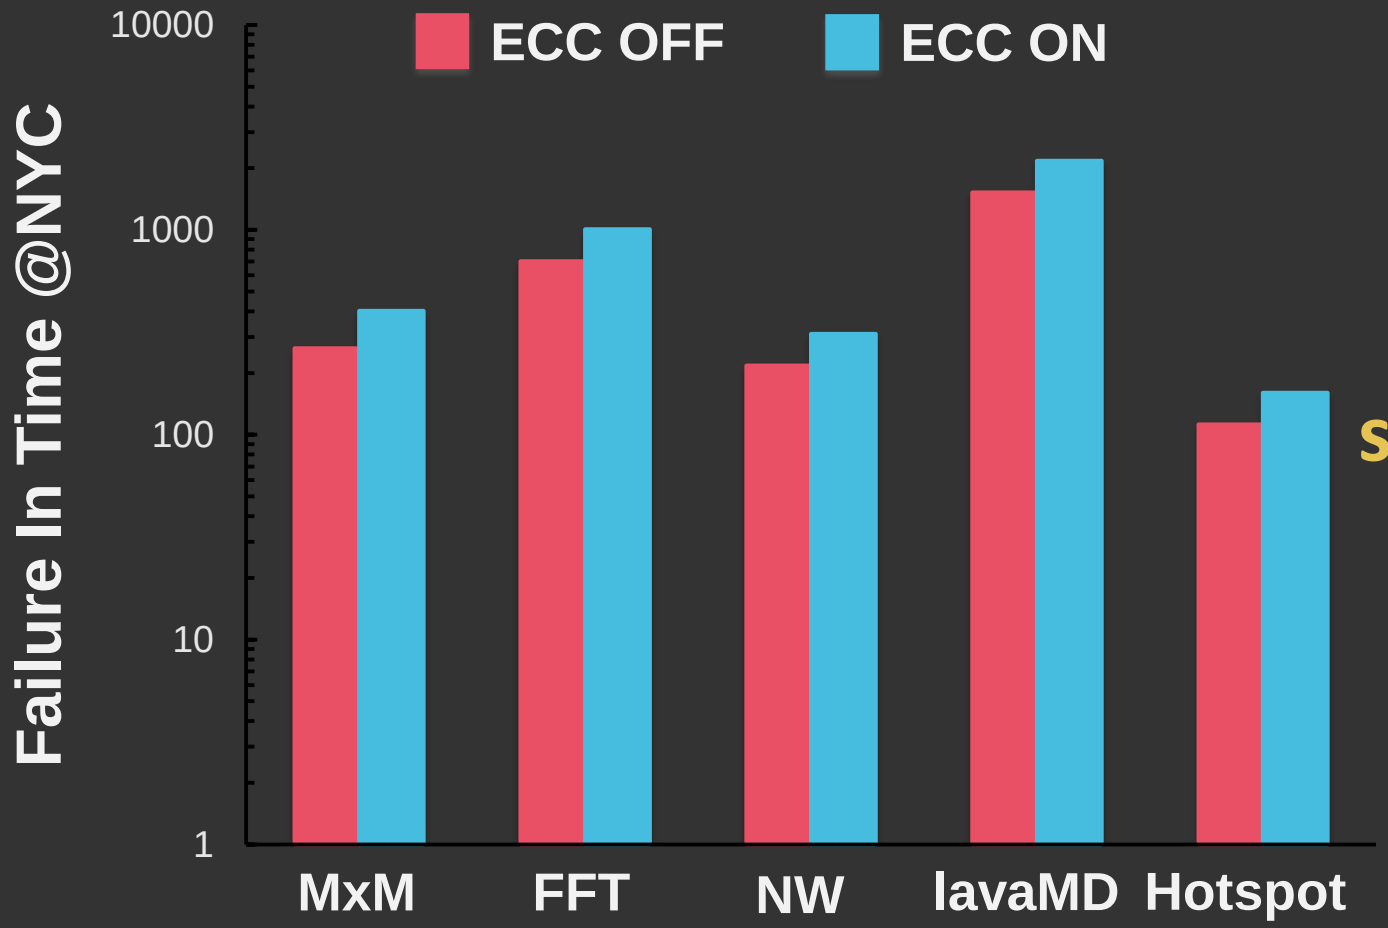

Double Bit Errors cause a crash

scheduler is not protected

UFSC

ECL

Oliveira et al. *Evaluation and mitigation of radiation-induced soft errors in graphics processing units*. Trans. Comp. 2015

Fault Tolerance Research

When the ECC is ON Crashes are more likely to occur than SDCs (this is GOOD for HPC centers!)

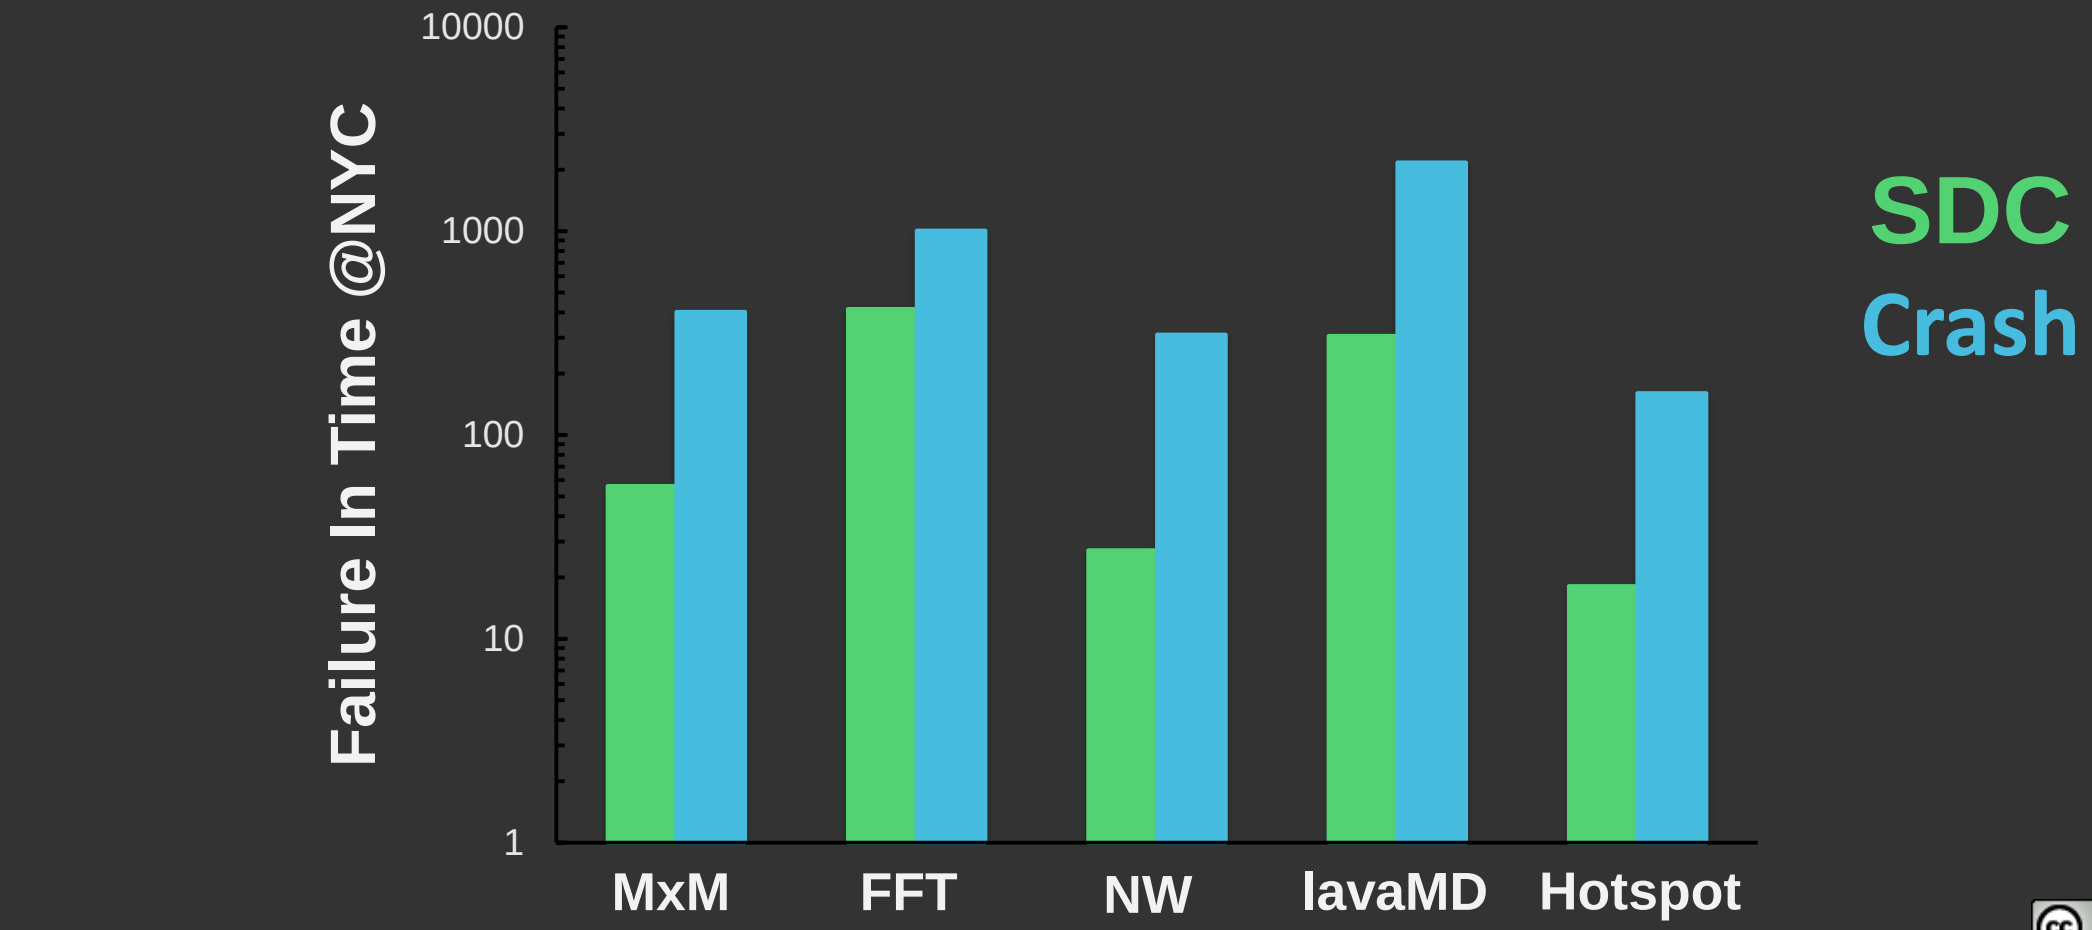

Fault Tolerance Research

ABFT: technique designed specifically for an algorithm.

Usually ABFT requires input coding, algorithm modification, and output decoding with error detection/correction

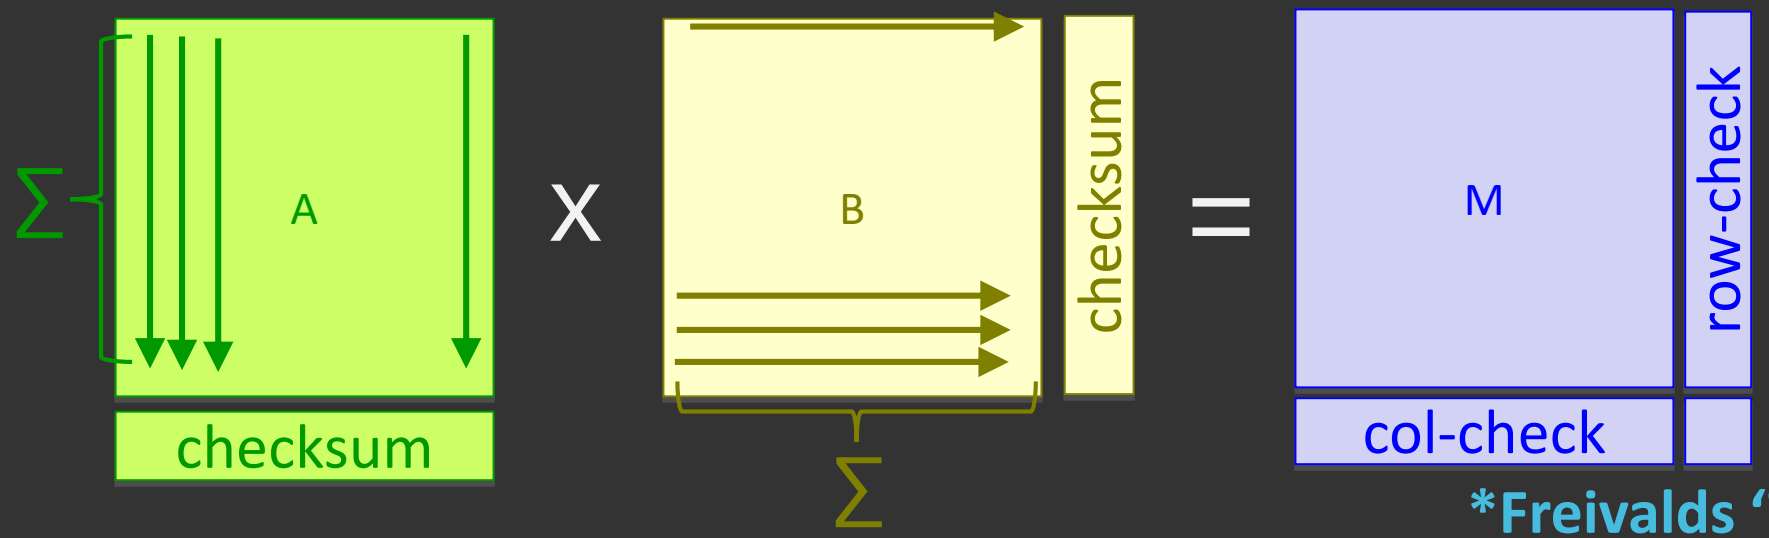

*Freivalds '79

Fault Tolerance Research – ABFT for FFT

*J.Y. Jou and Abraham '88

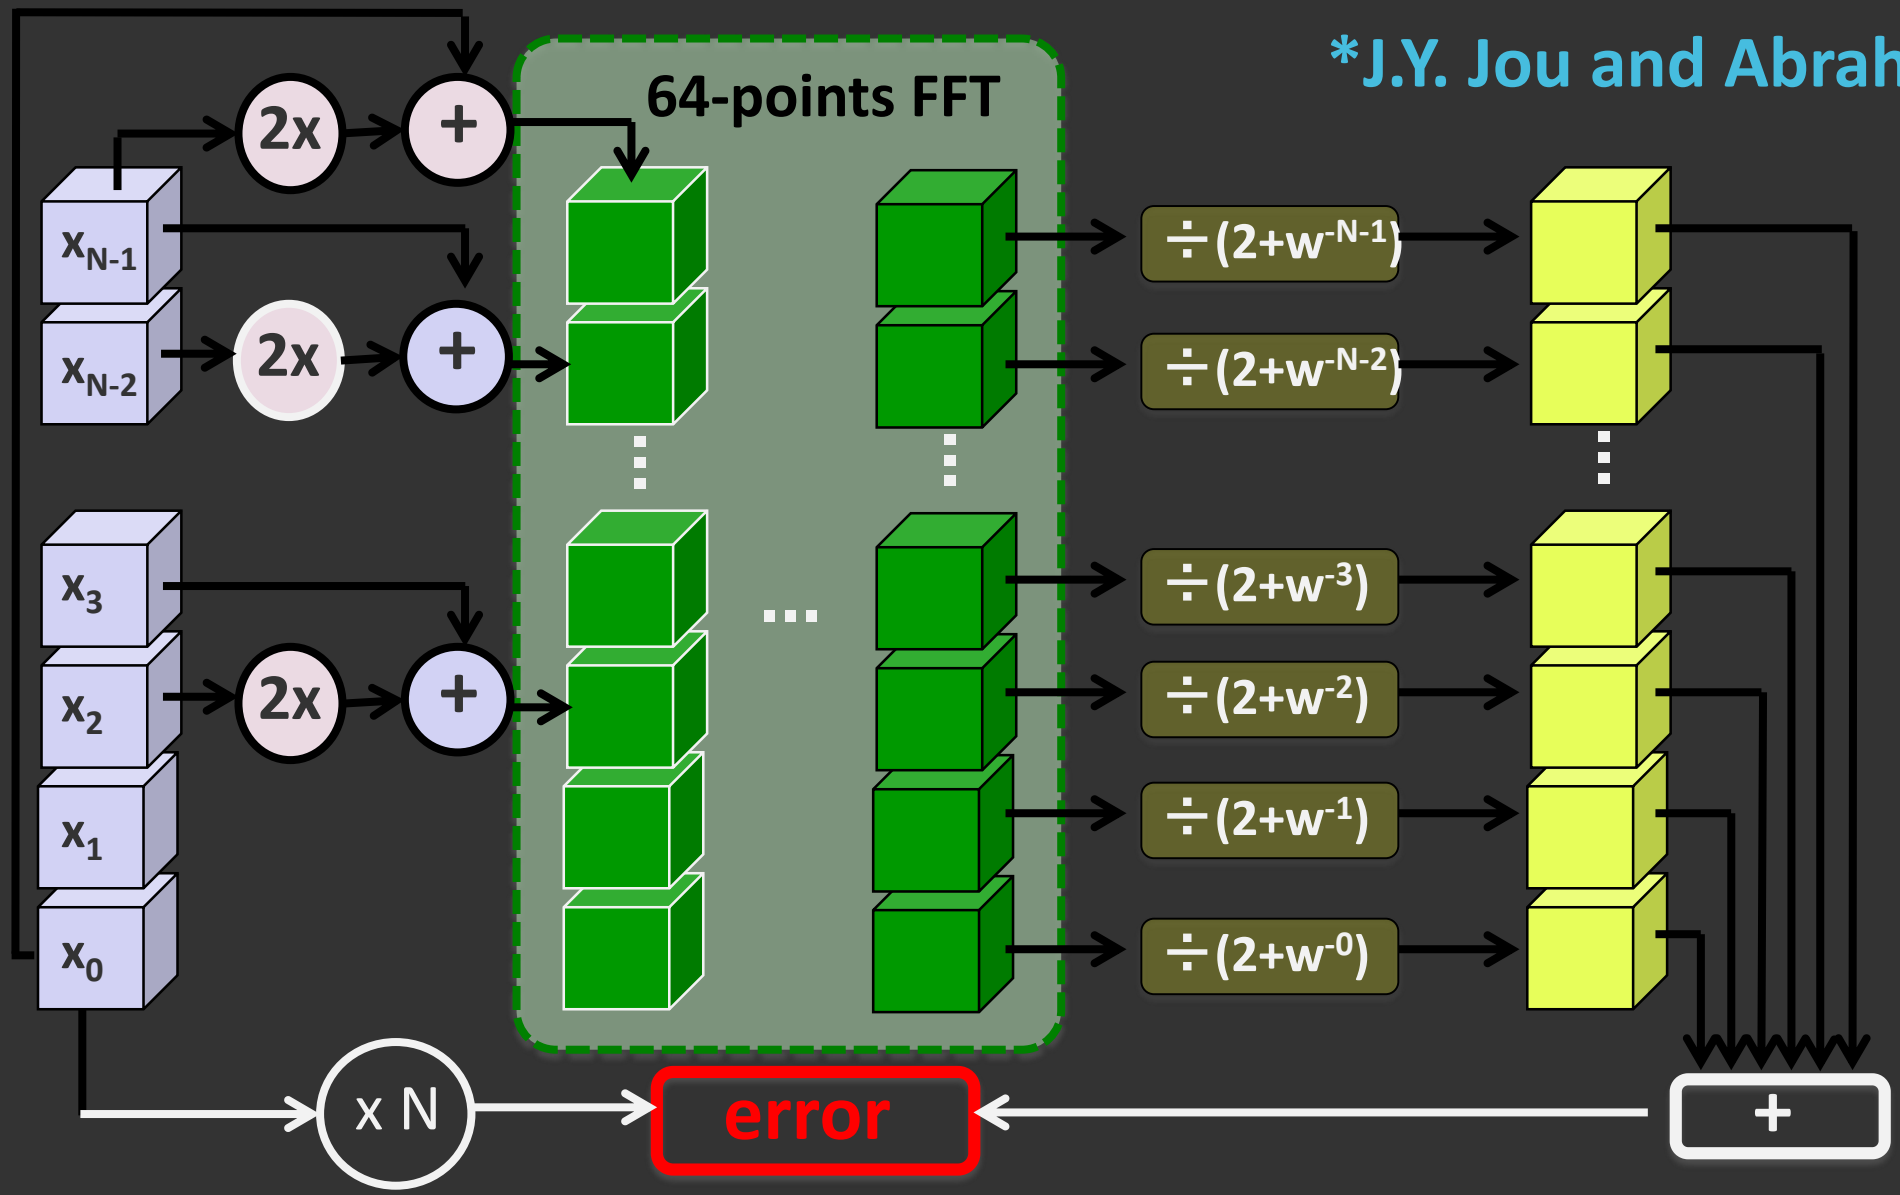

Fault Tolerance Research

ECC x ABFT

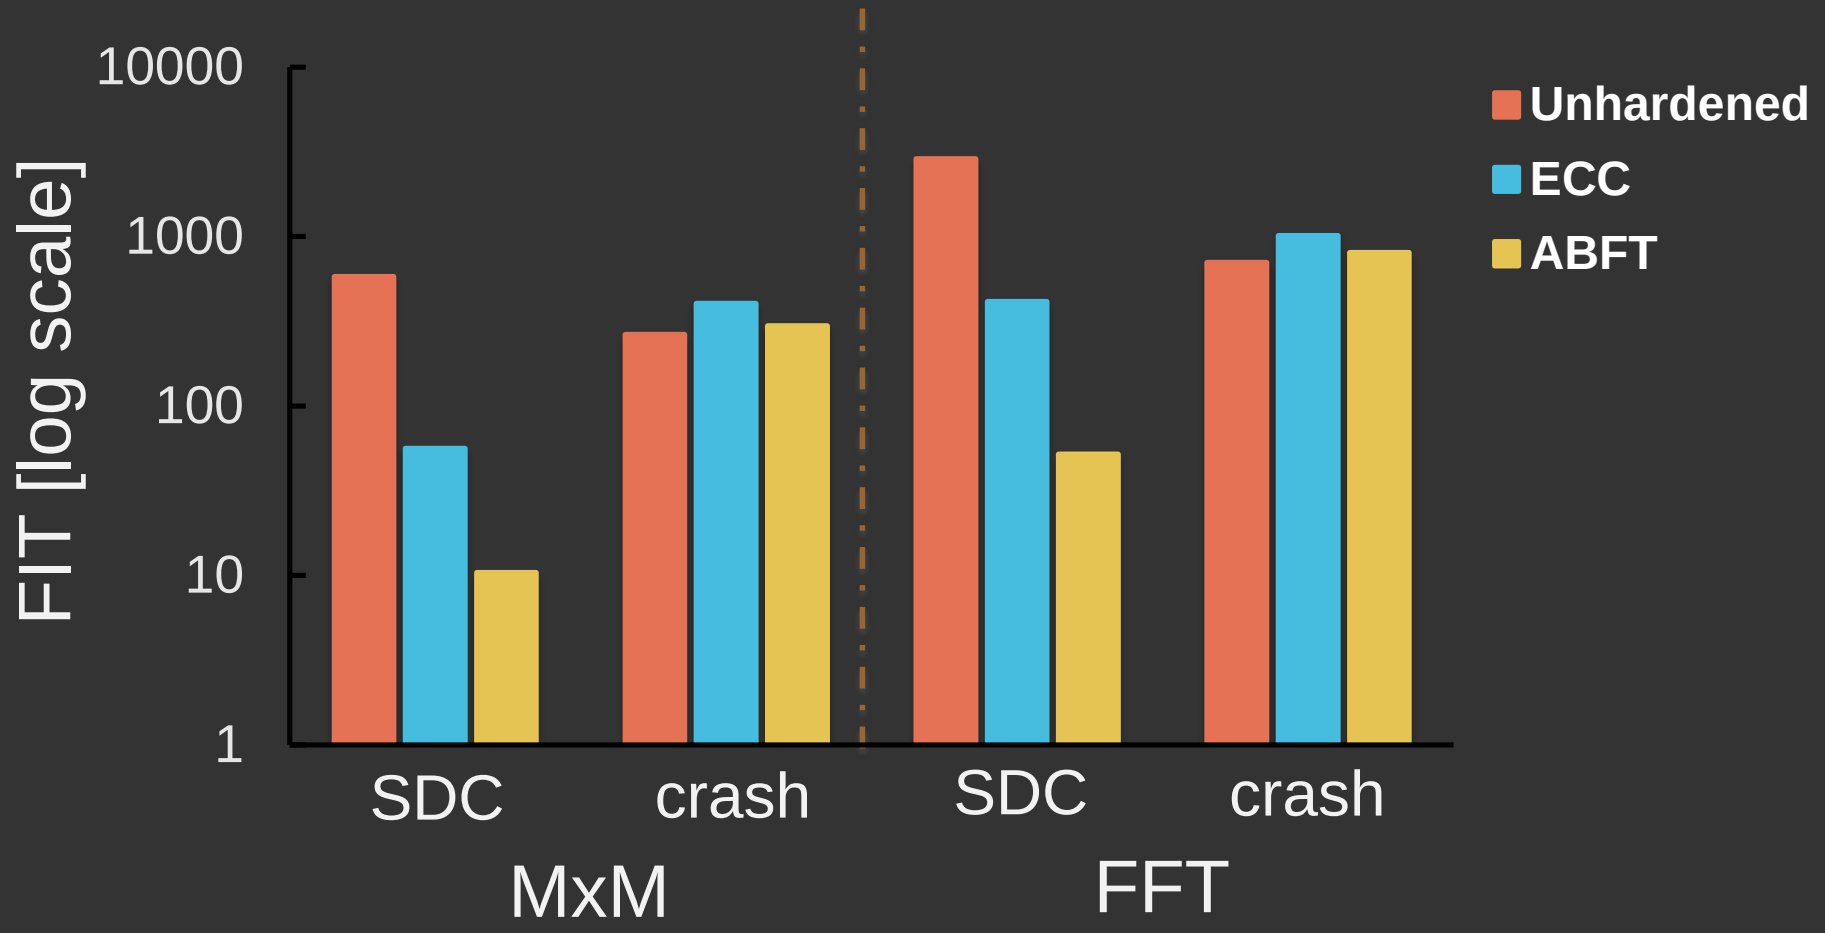

UFSC

ECL

Oliveira et al. *Evaluation and mitigation of radiation-induced soft errors in graphics processing units*. Trans. Comp. 2015

Fault Tolerance Research

ECC x ABFT

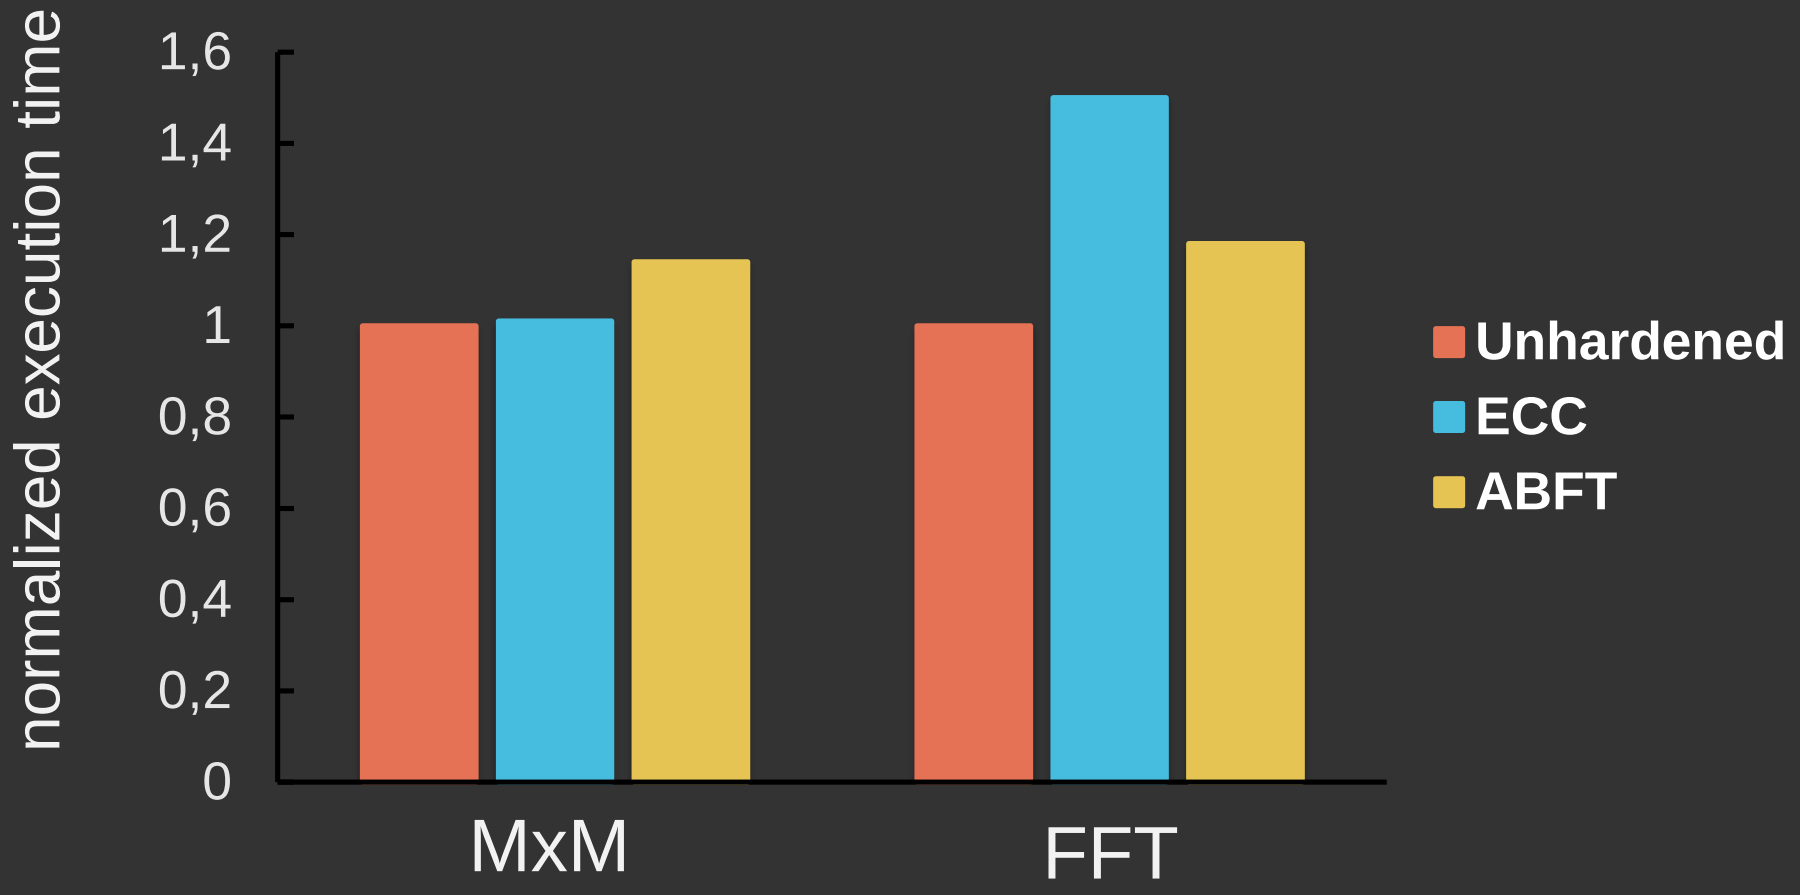

UFSC

ECL

Oliveira et al. *Evaluation and mitigation of radiation-induced soft errors in graphics processing units*. Trans. Comp. 2015

Fault Tolerance Research

Effects of optimizations: MxM

MWBF: Mean Workload Between Failures

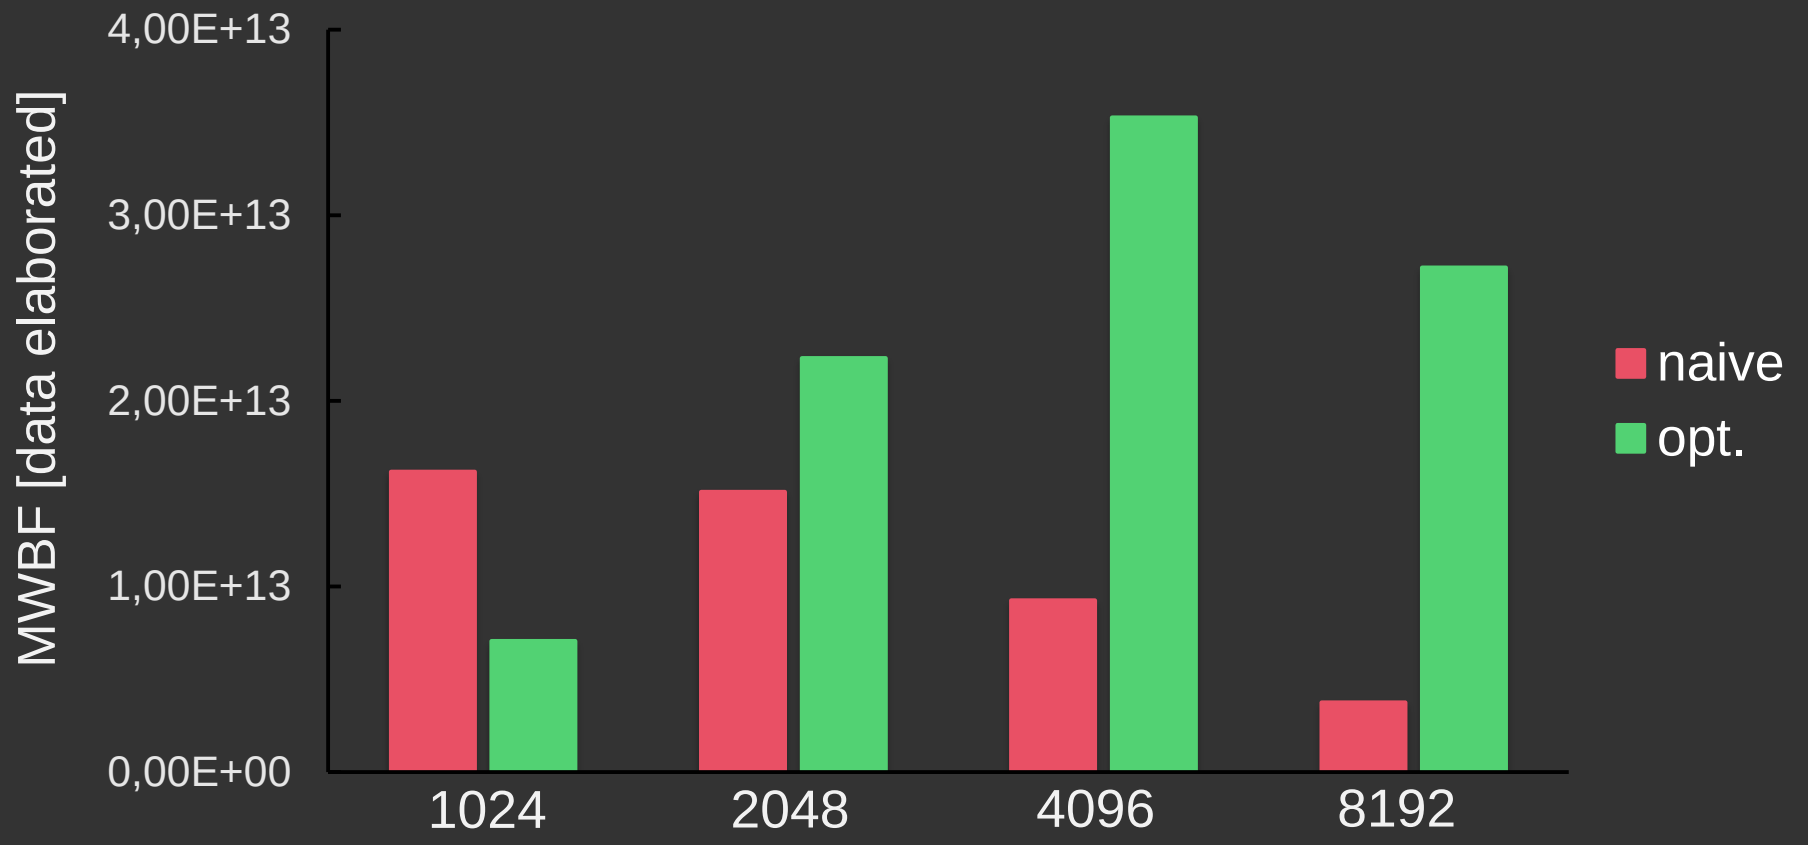

Pilla et al. *Memory Access Time and Input Size Effects on Parallel Processors Reliability*. Trans. Nucl. Sci. 2015

Fault Tolerance Research

Xeon-Phi MWBF decreases significantly with input size.

Even if more prone to be corrupted, Kepler produces more correct data (if parallelism is exploited)

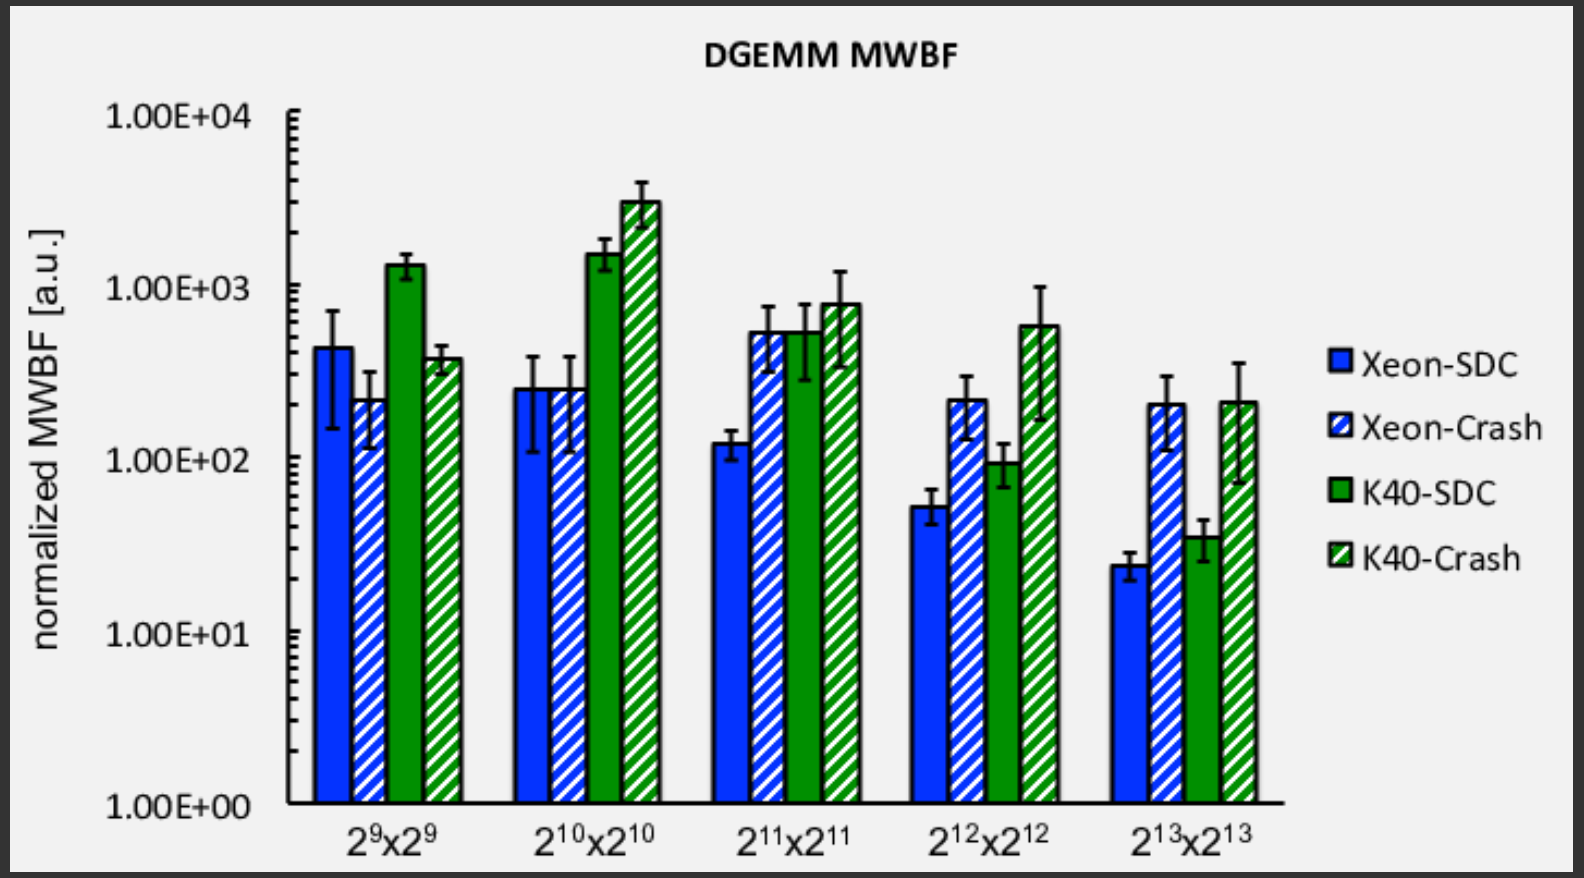

UFSC

ECL

Fault Tolerance Interests

+ applications & kernels for CPU & GPU

+ platforms for energy measurement

+ experiments with radiation

UFSC

ECL

Conclusions

Conclusions

EnergySFE

international collaboration

three research questions

lots of work in scheduling and fault tolerance

lots more to be done

Energy-aware Scheduling and Fault Tolerance Techniques for the Exascale Era

Thank you.

Laércio Lima Pilla

laercio.pilla@ufsc.br

Federal University of Santa Catarina, Brazil