Current Efforts in Scheduling and Fault Tolerance

First EnergySFE Workshop

Grenoble, September 1st, 2016

Laércio Lima Pilla

laercio.pilla@ufsc.br

UFSC

ECL

Agenda

Global Scheduling

problems + research + interests

Fault Tolerance

problems + research + interests

Global Scheduling

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

...

UFSC

ECL

Global Scheduling Problems

energy wastage

long computing times

low resource usage

excessive data movement

slow scheduling algorithms

…

Global Scheduling Research

What we want

balanced loads

optimized communications

fast scheduling algorithms

least migrations possible

UFSC

ECL

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

OD = overdecomposition (512 tasks/ 32 PUs)

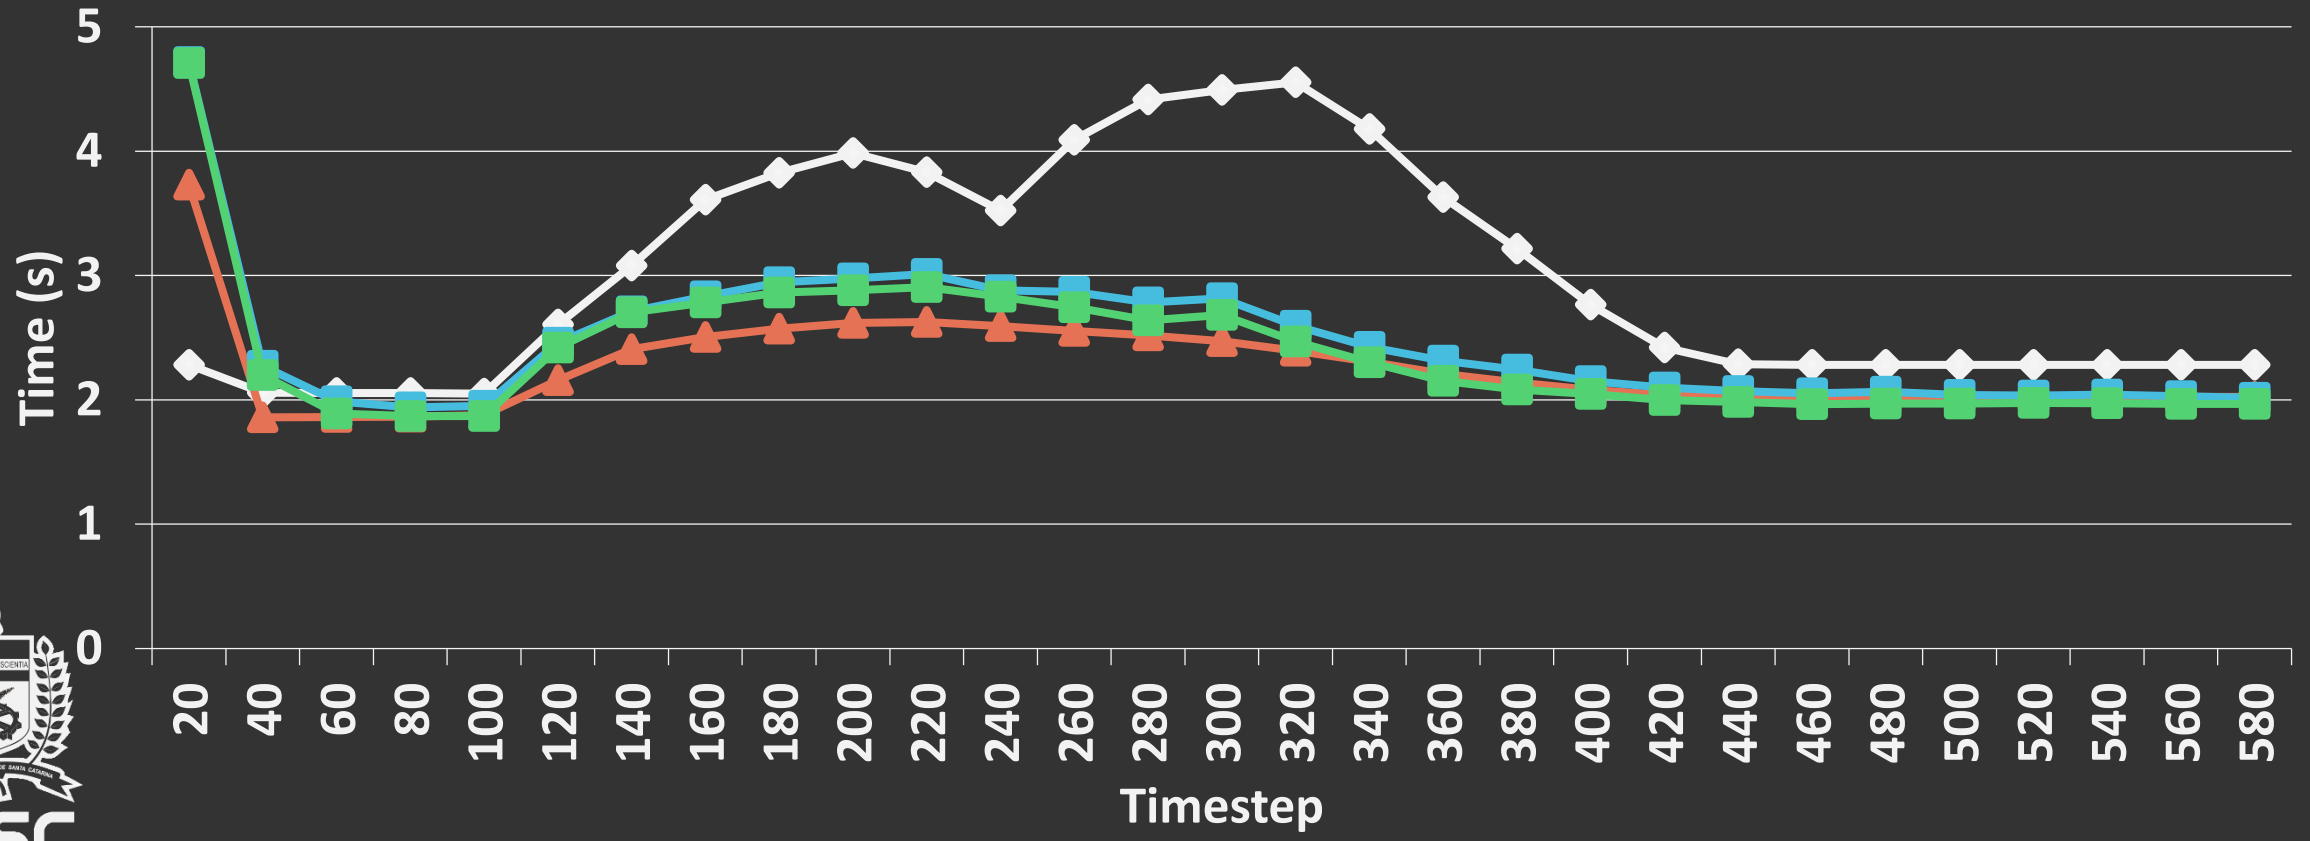

UFSC

ECL

Global Scheduling Research

Application: LeanMD

Weak scalability: Total execution time

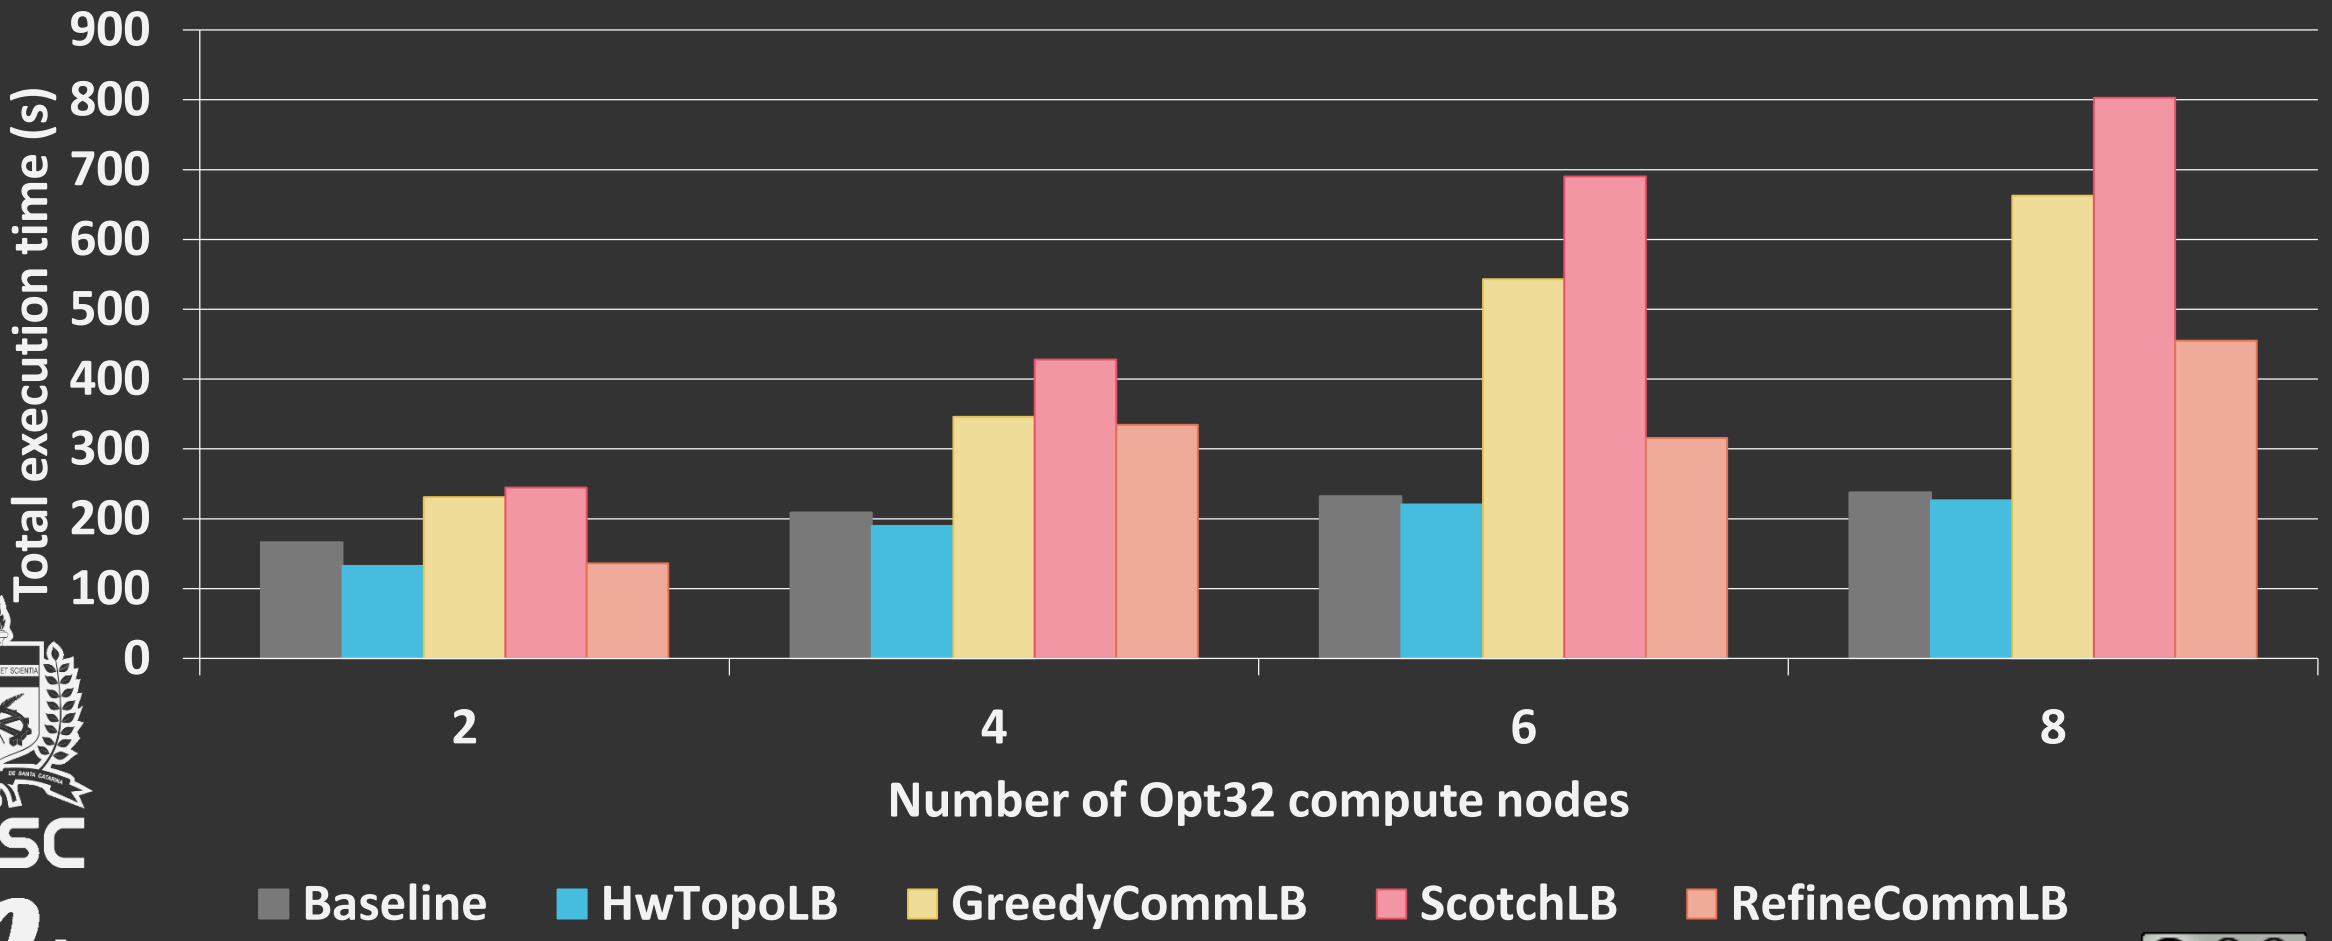

UFSC

ECL

Global Scheduling Research

Strong scalability: LeanMD

Speedup over the baseline on 2 CNs

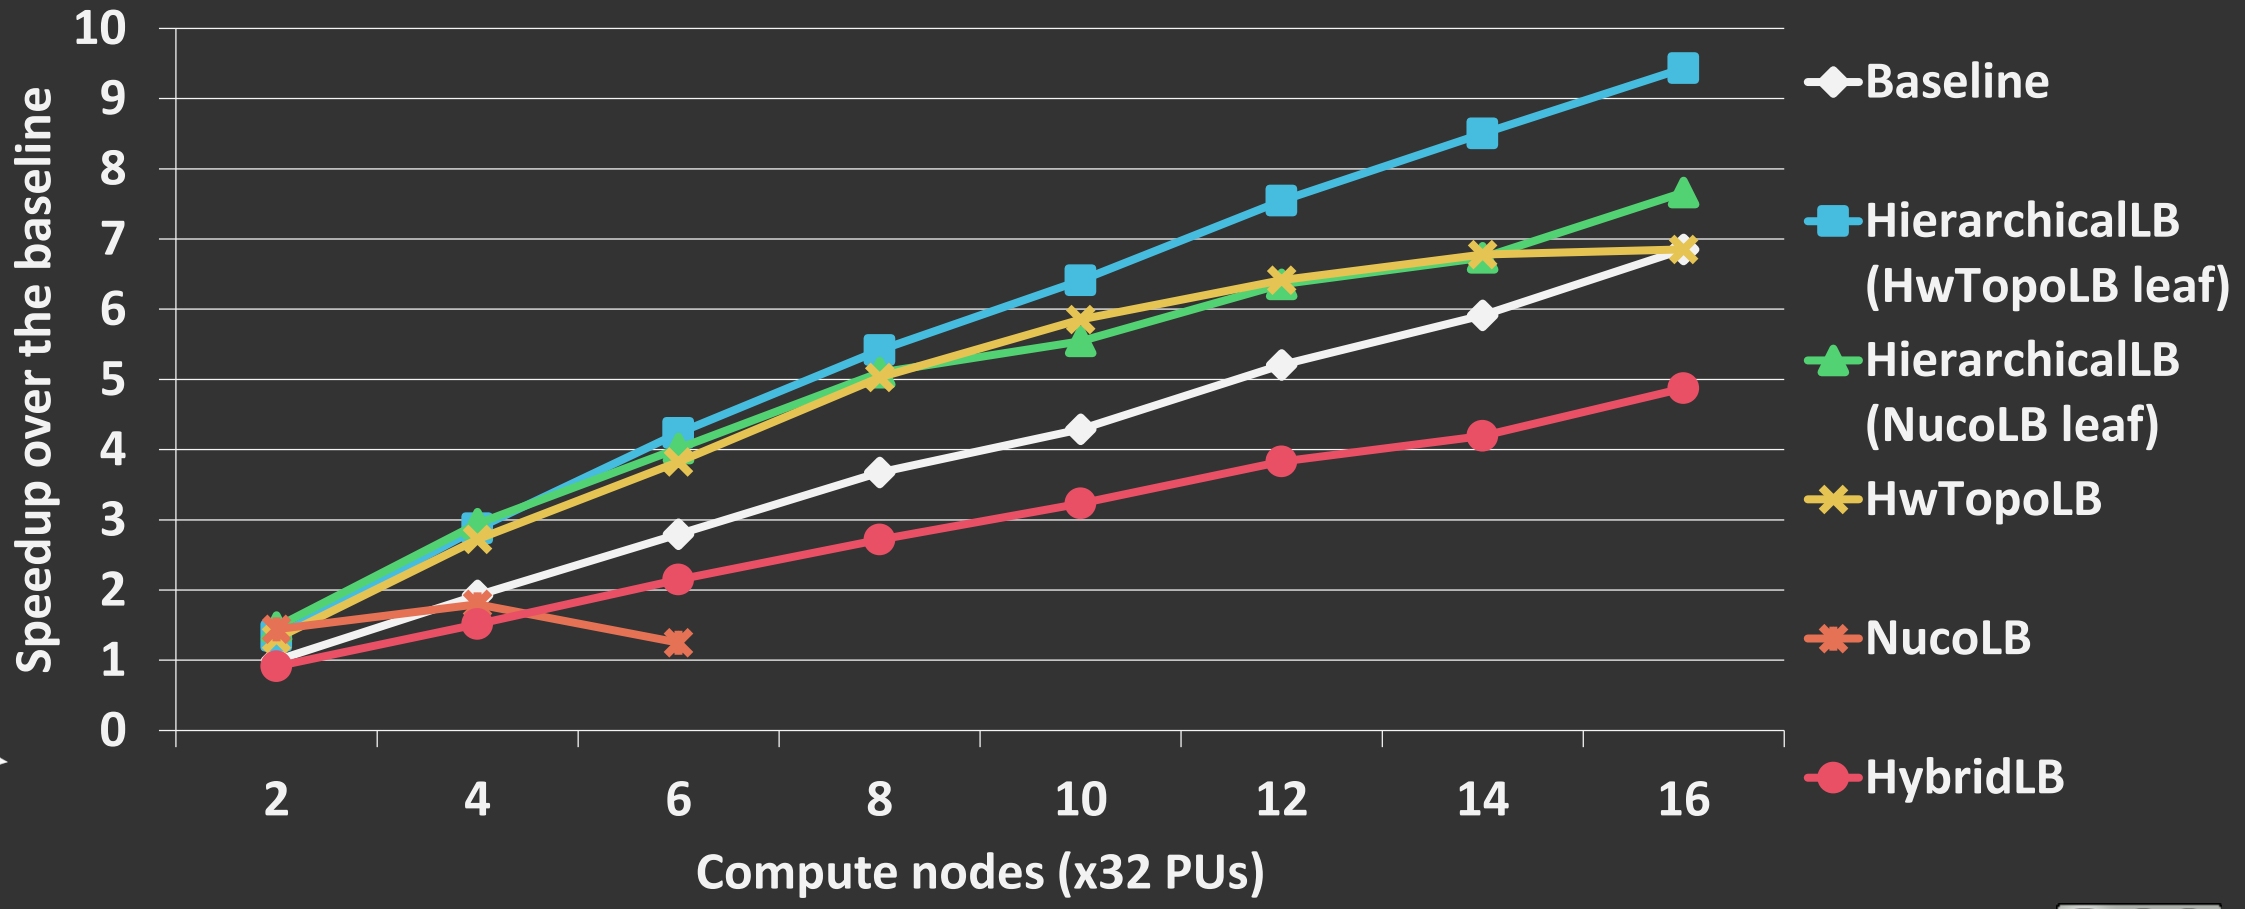

UFSC

ECL

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

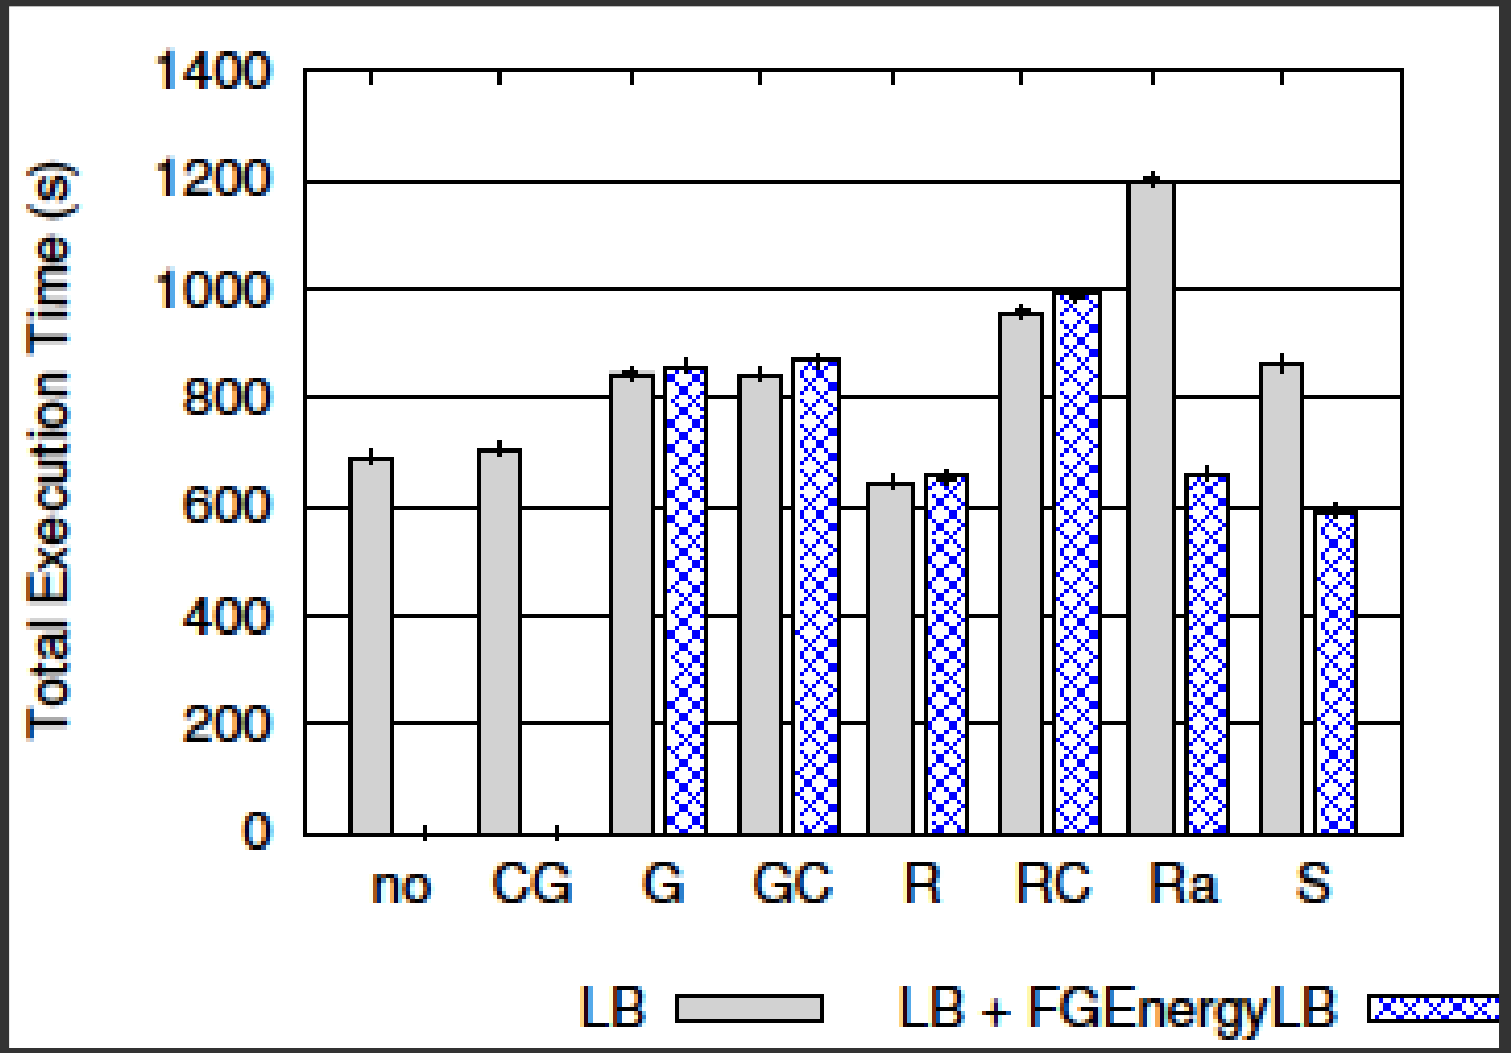

Global Scheduling Research

Example: Lulesh + FG EnergyLB

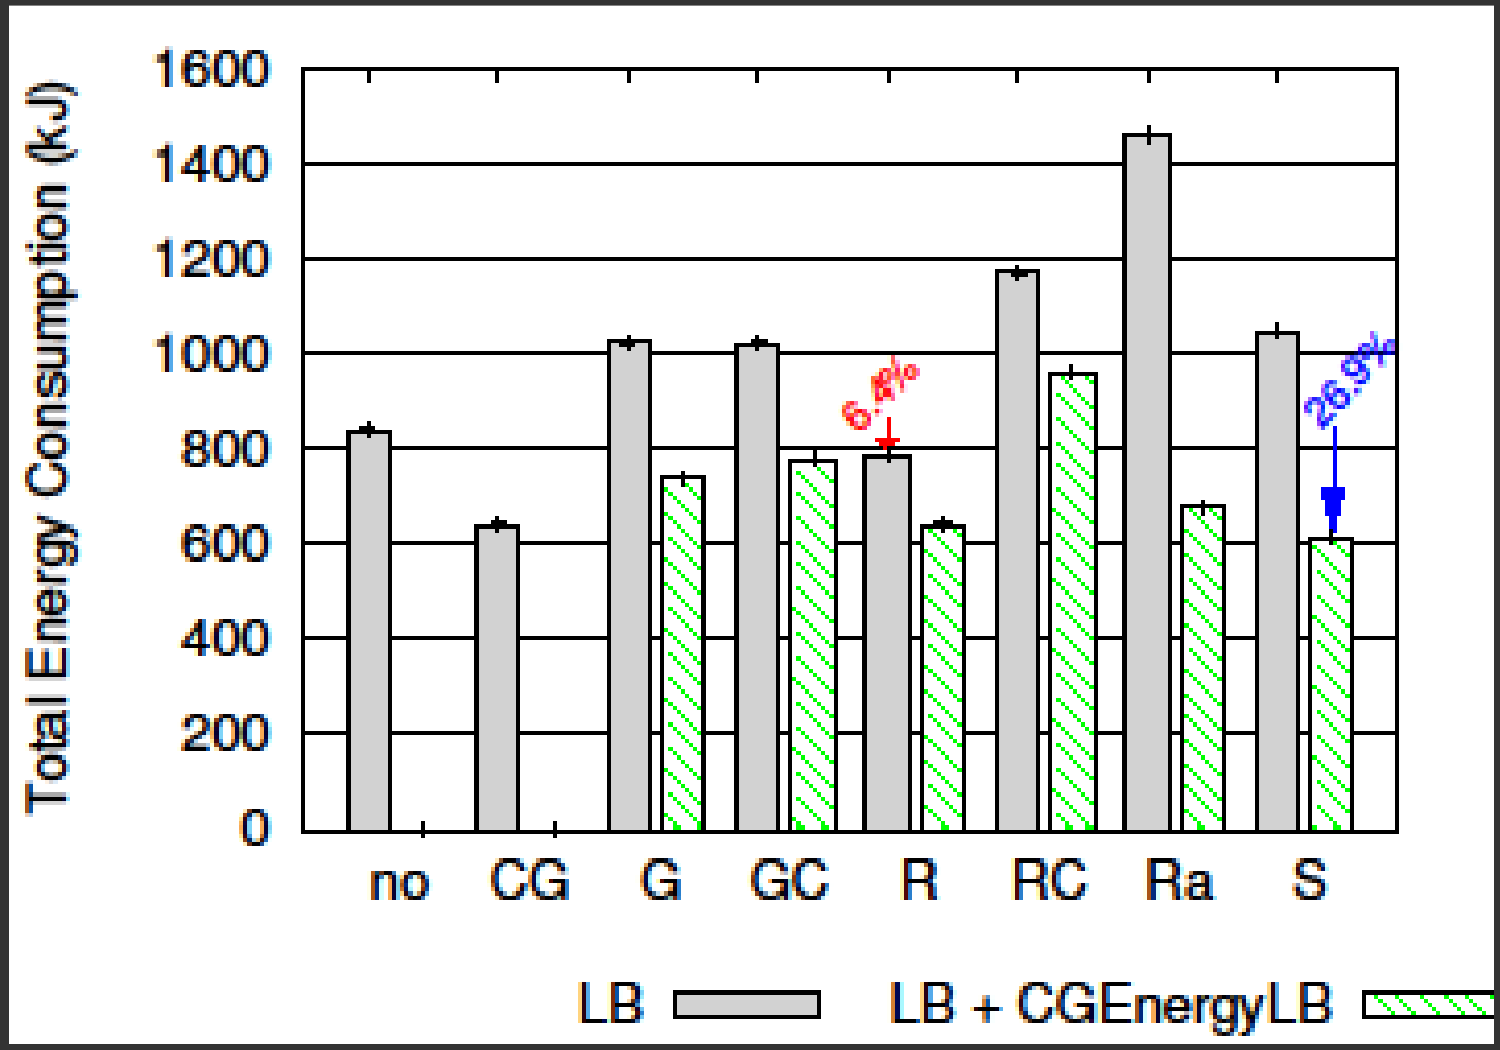

Global Scheduling Interests

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

UFSC

ECL

Fault Tolerance

For more information, we can check with Paolo Rech :-)

Fault Tolerance Problems

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

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

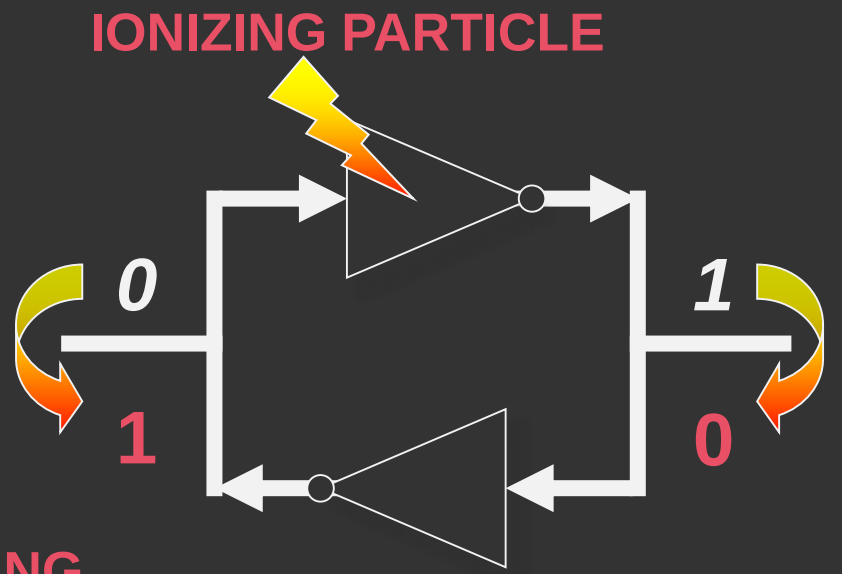

Transient voltage pulse
Single Event Transient (SET)

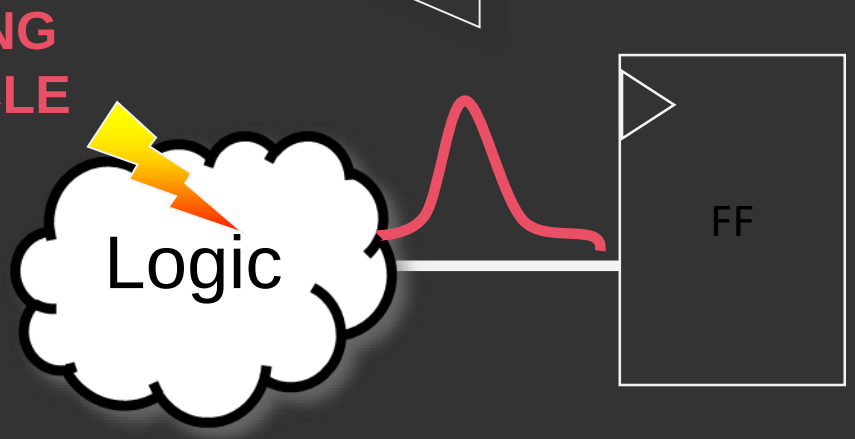

UFSC
ECL

Fault Tolerance Problems

Silent Data Corruption

data caches

register files

ALU

scheduler

Crash

instruction cache

scheduler

PCIe bus

Fault Tolerance Research

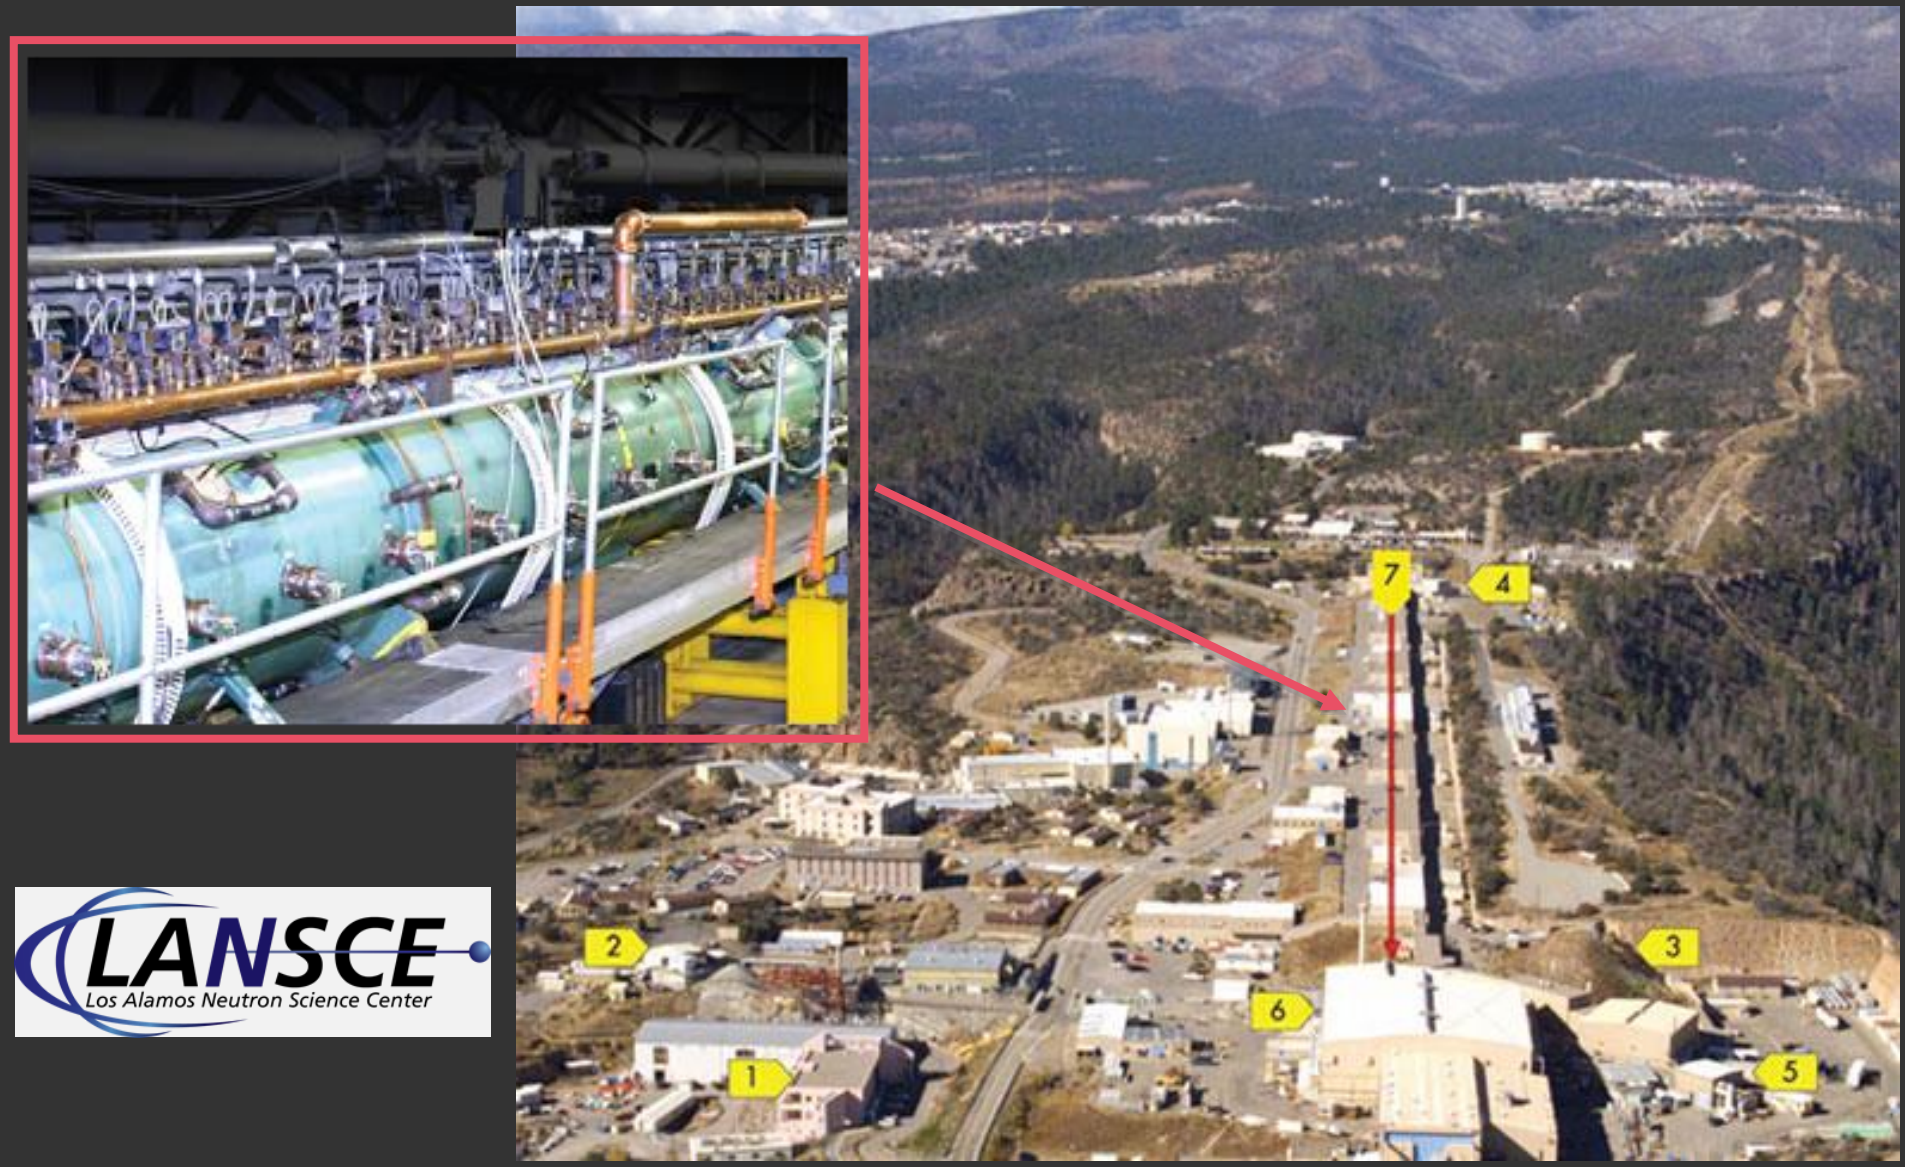

LANSCE
Los Alamos Neutron Science Center

UFSC

ECL

Fault Tolerance Research

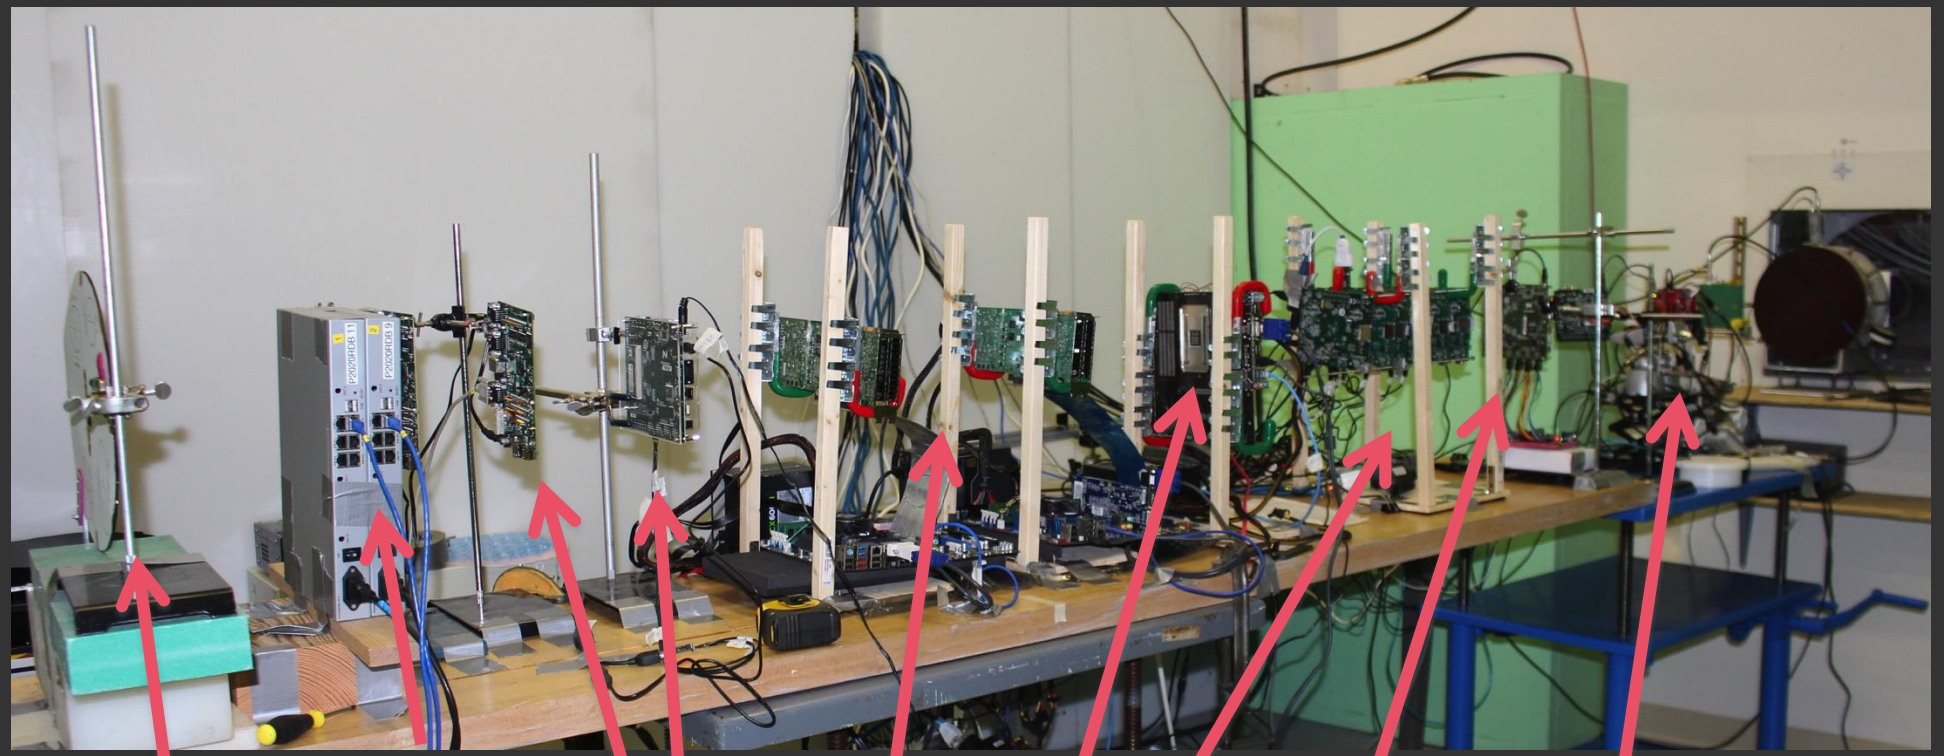

UFSC

ECL

Fault Tolerance Research

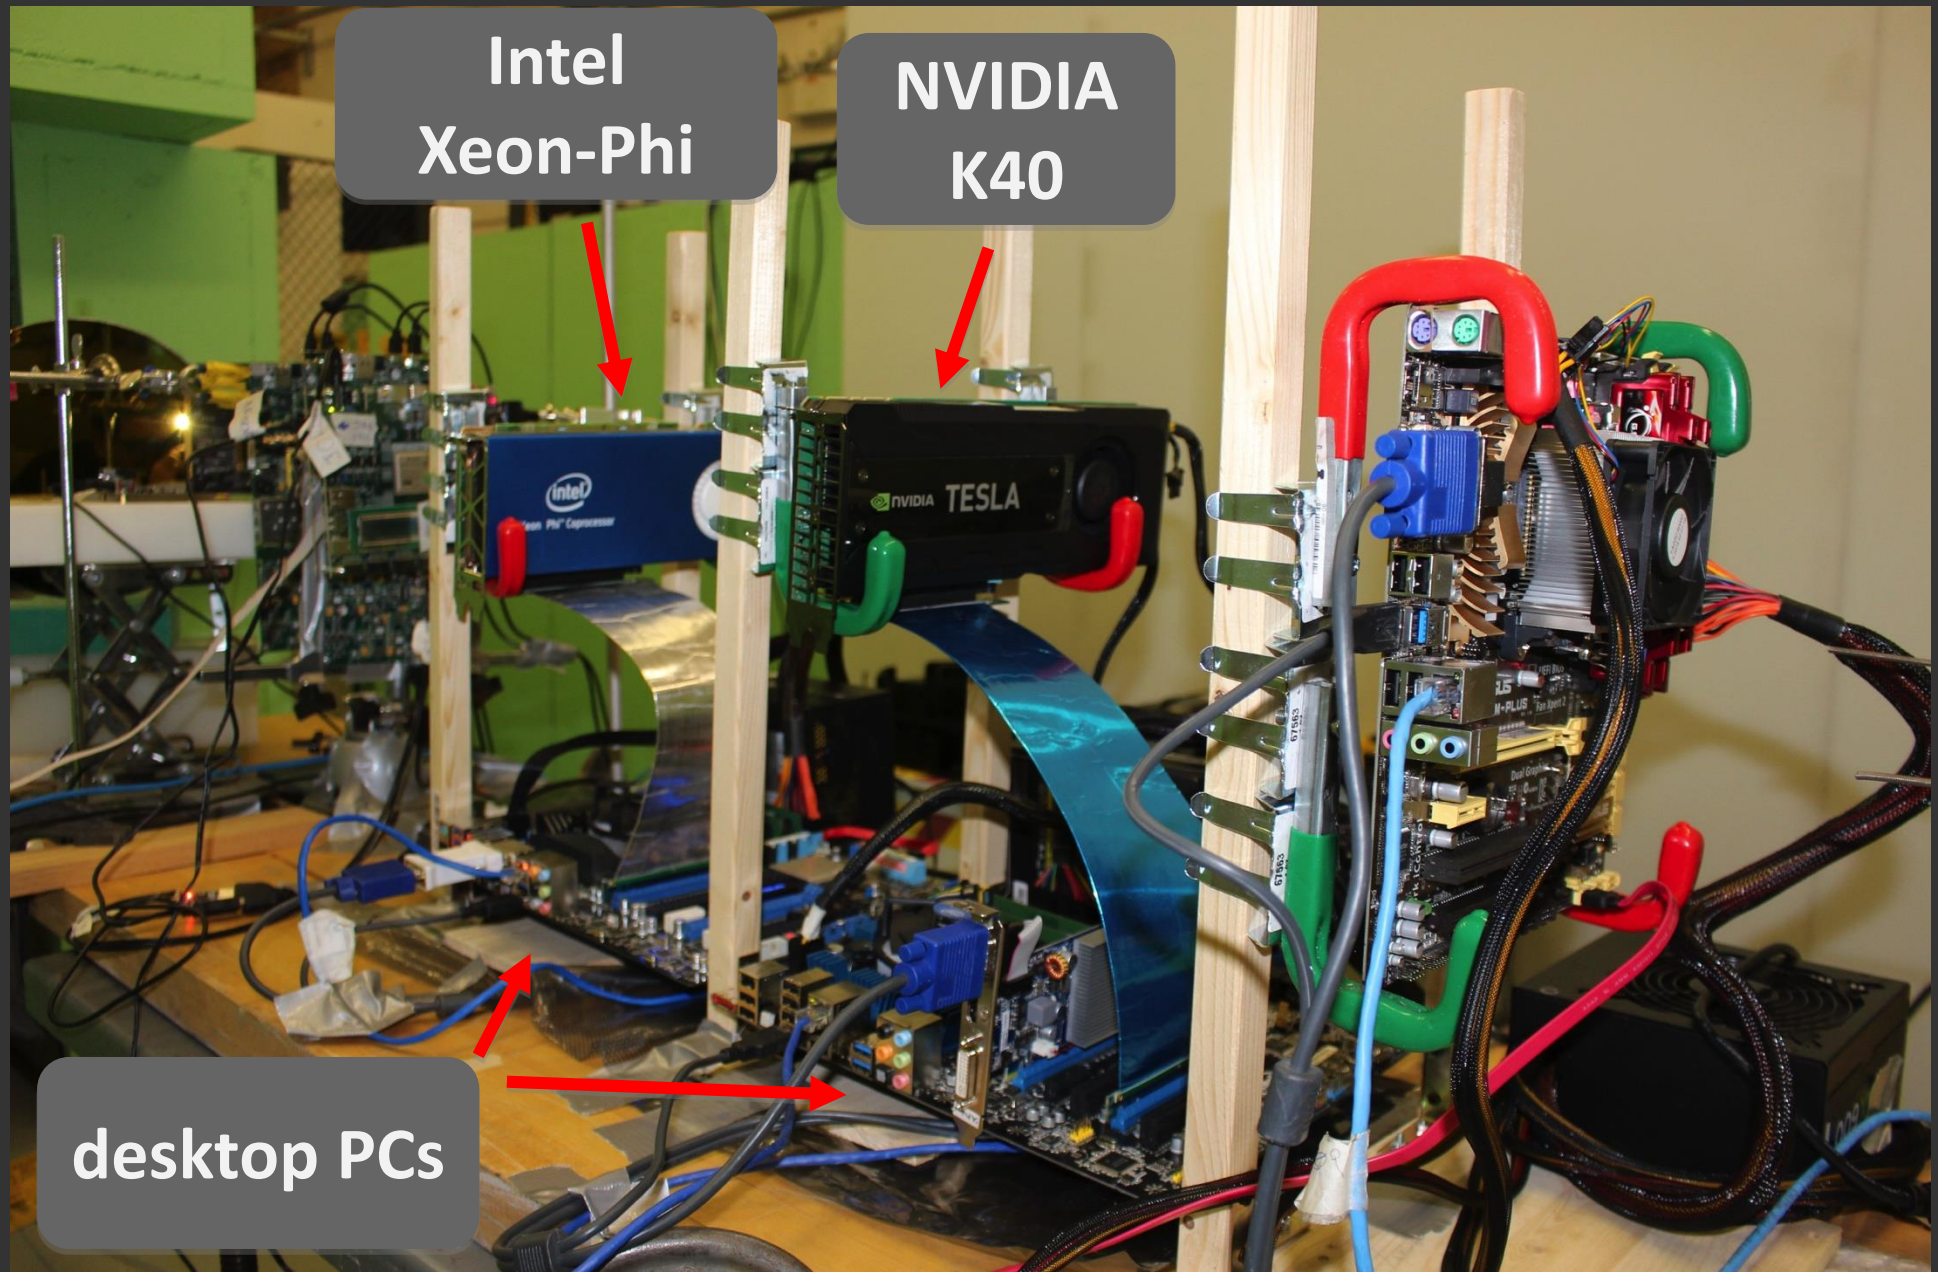

UFSC

ECL

Fault Tolerance Research

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

@NYC 13 n/(cm^2 h)

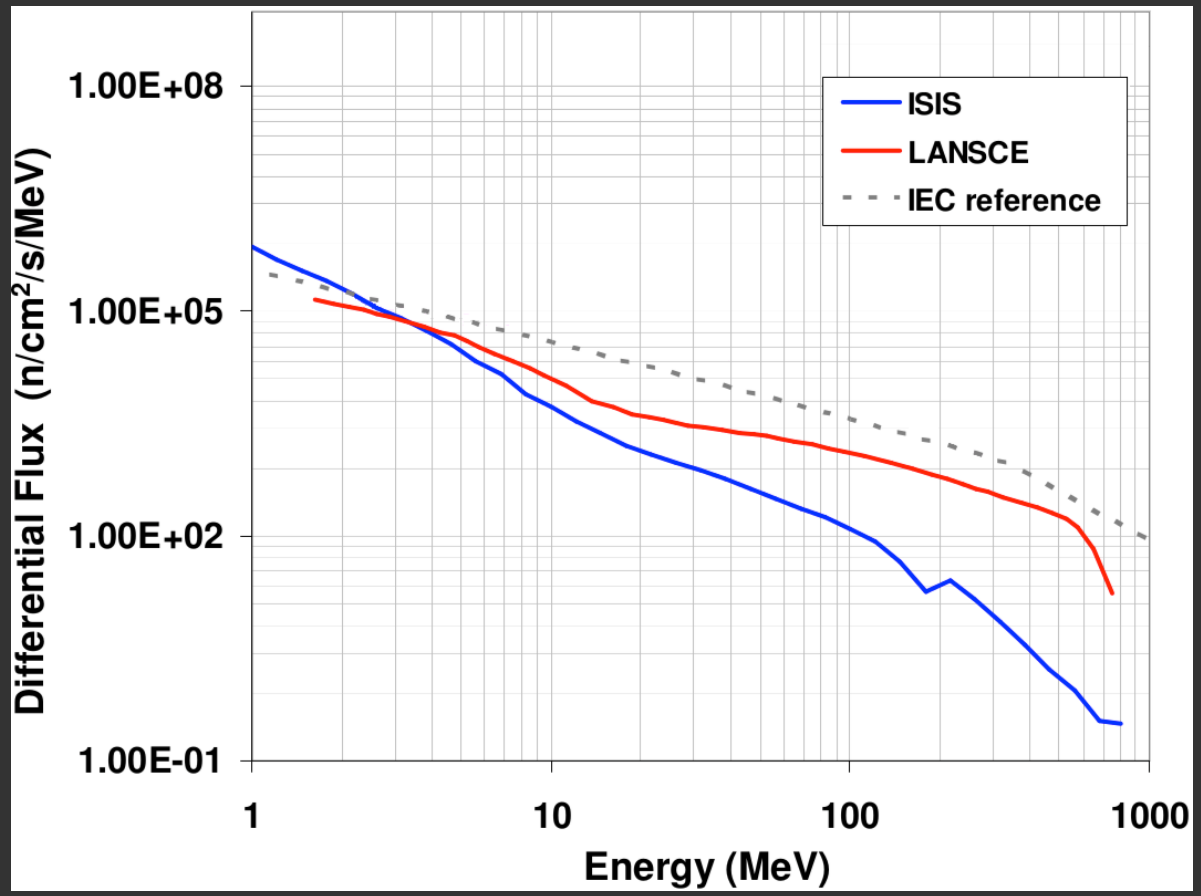

UFSC

ECL

Fault Tolerance Research

Parallel algorithms' reliability

SDC rates vary ~3 orders of magnitude
(details on Oliveira et al. Trans. Comp. 2015)

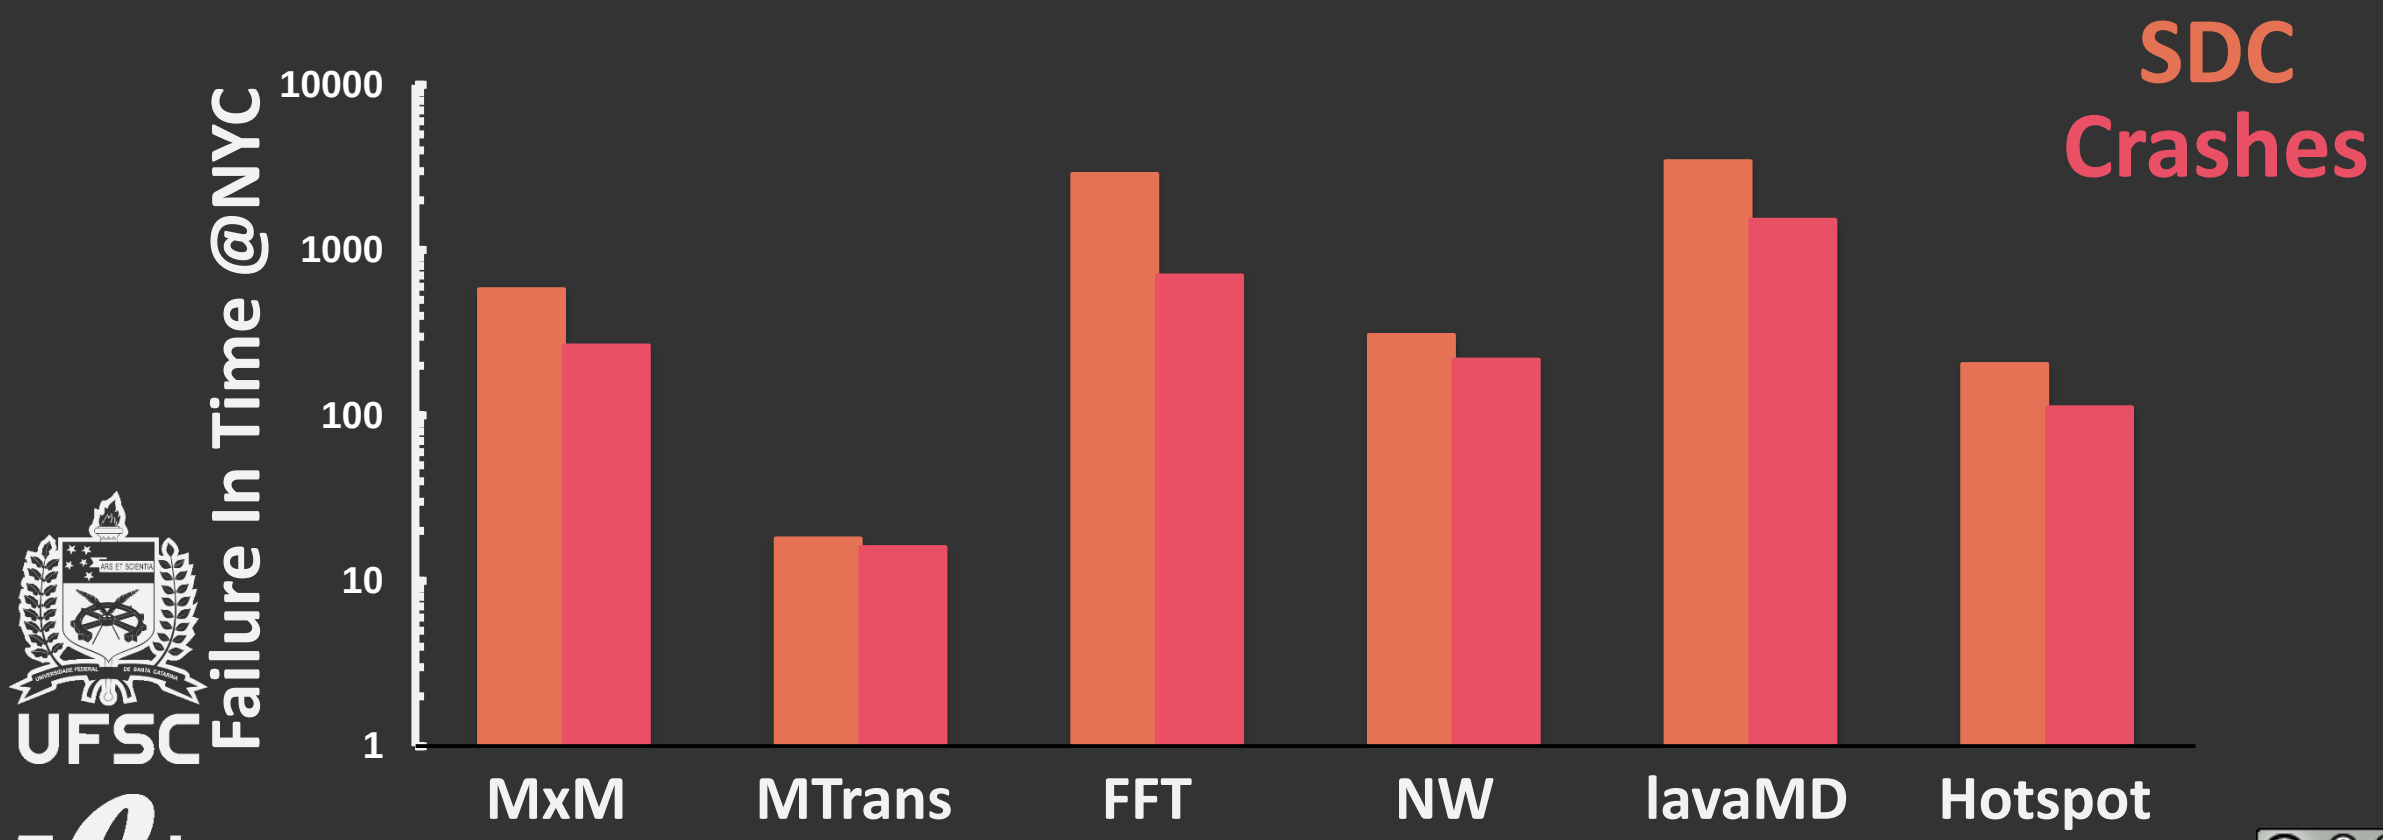

UFSC

ECL

Fault Tolerance Research

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

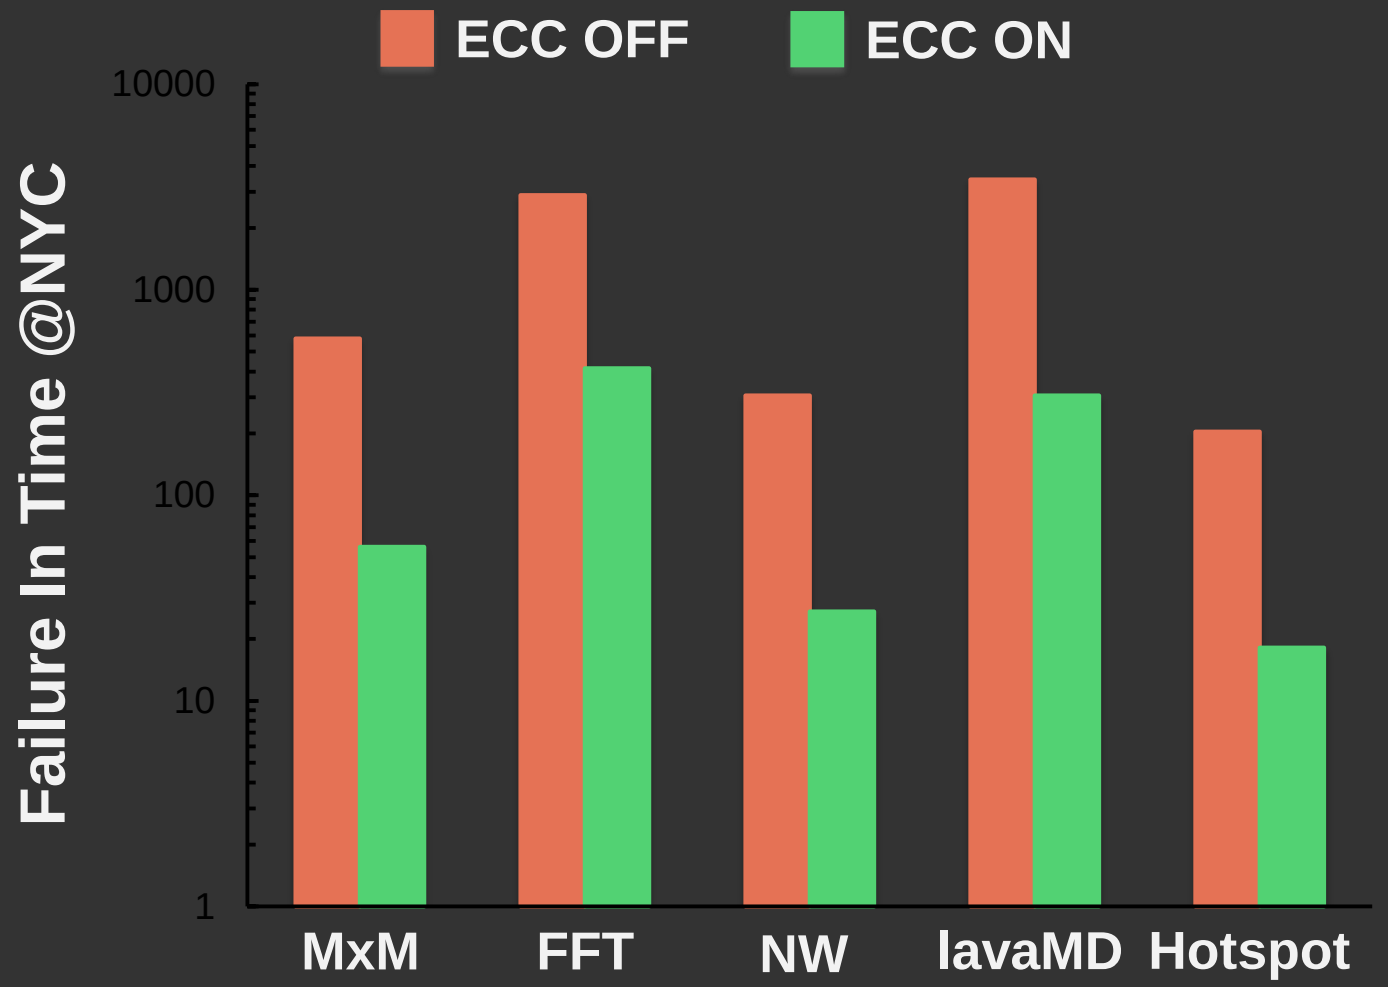

UFSC

ECL

Fault Tolerance Research

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

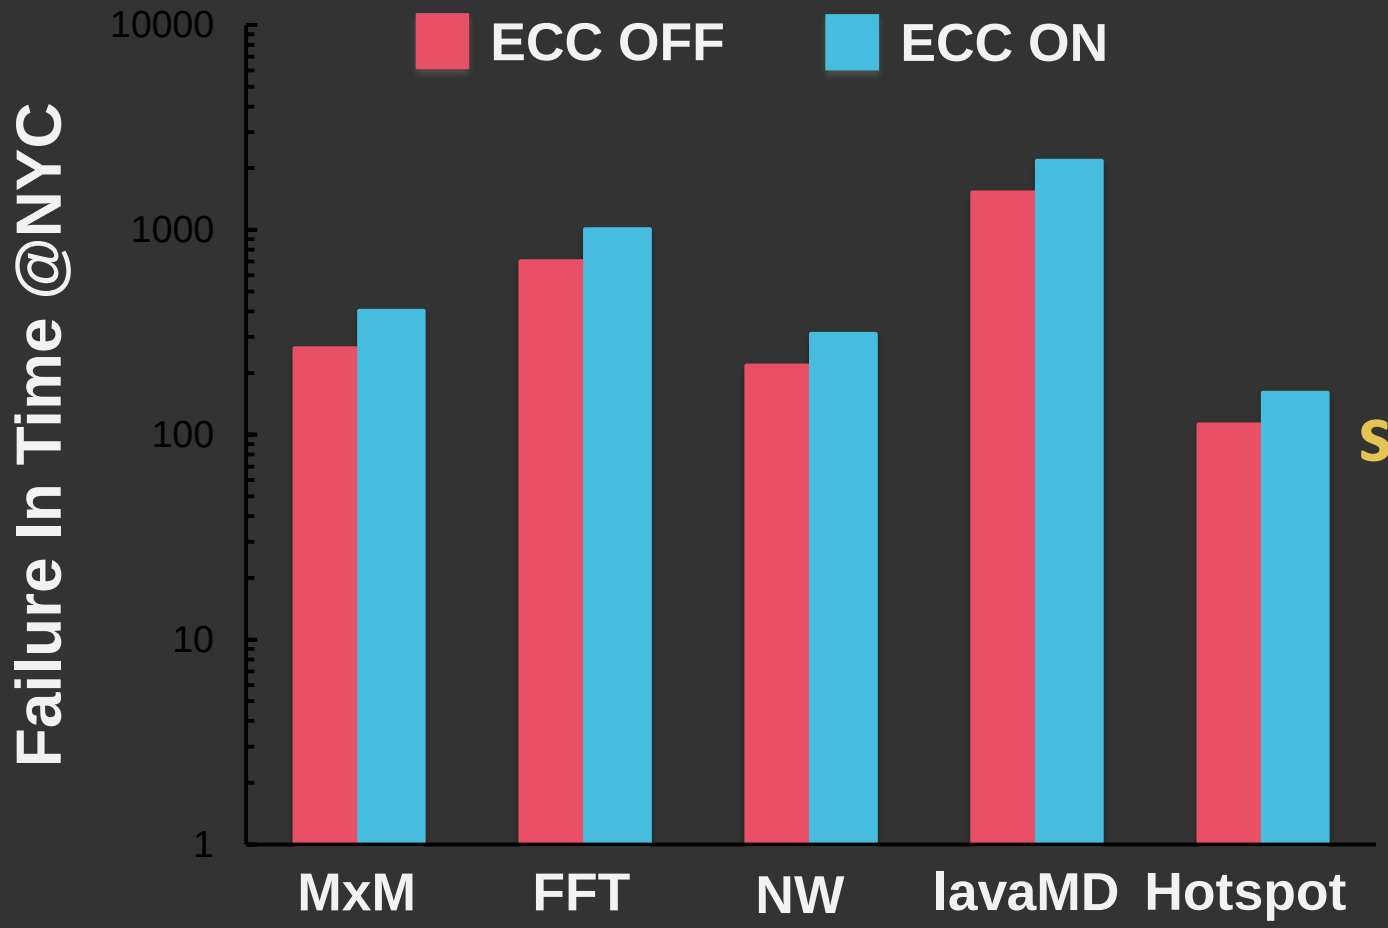

Double Bit Errors cause a crash

scheduler is not protected

UFSC

ECL

Fault Tolerance Research

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

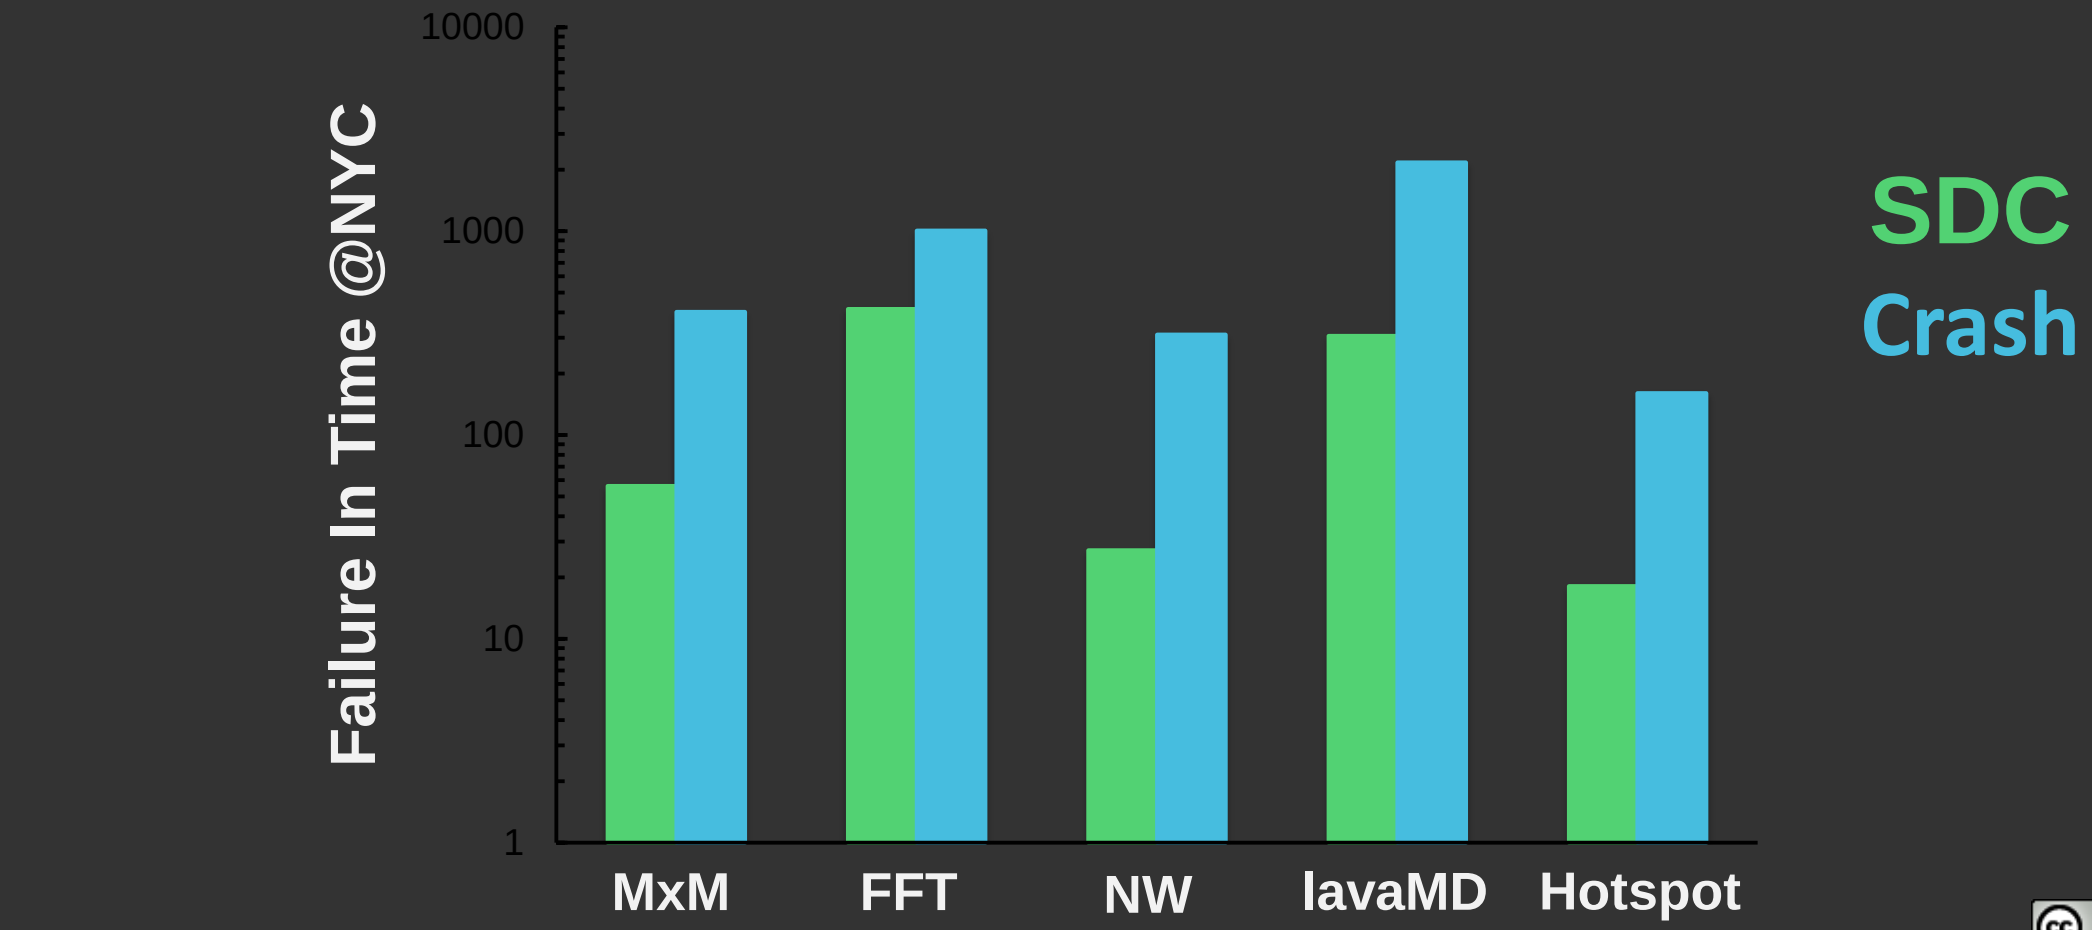

UFSC

ECL

Fault Tolerance Research

ABFT: technique designed specifically for an algorithm.

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

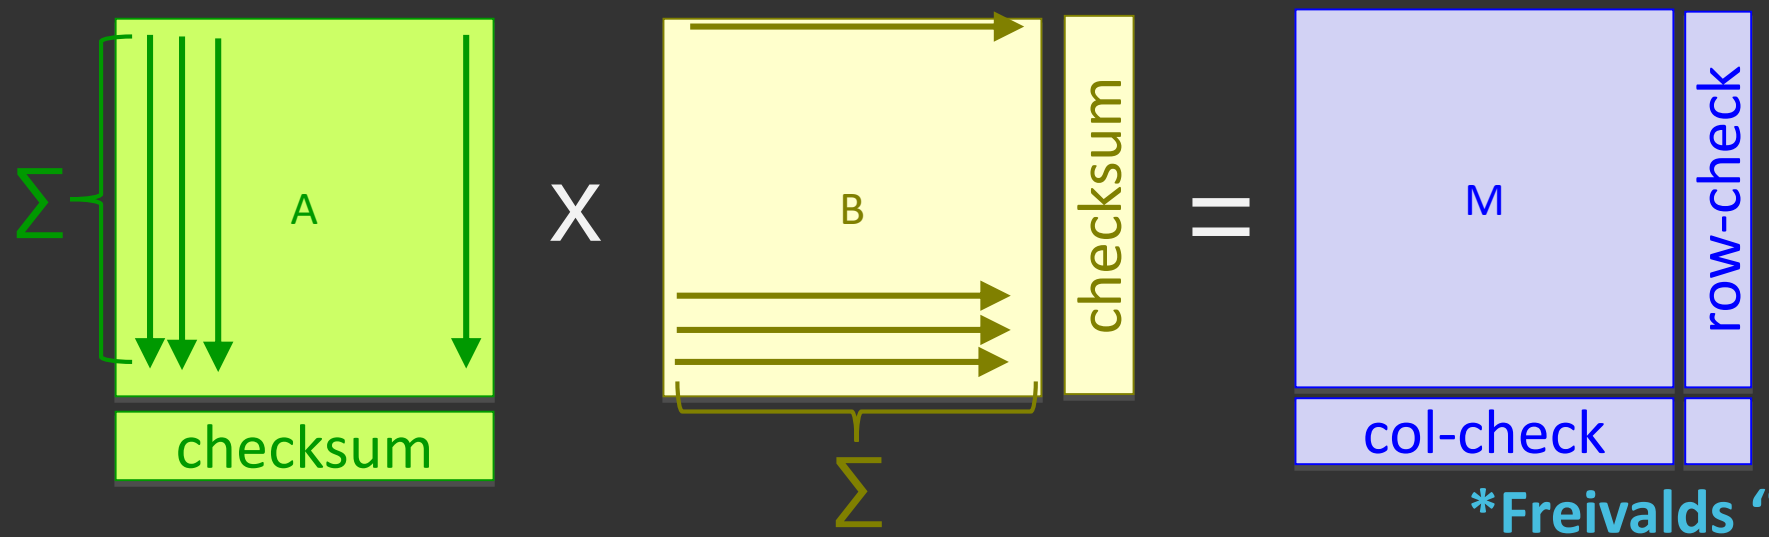

*Freivalds '79

UFSC

ECL

Fault Tolerance Research – ABFT for FFT

*J.Y. Jou and Abraham '88

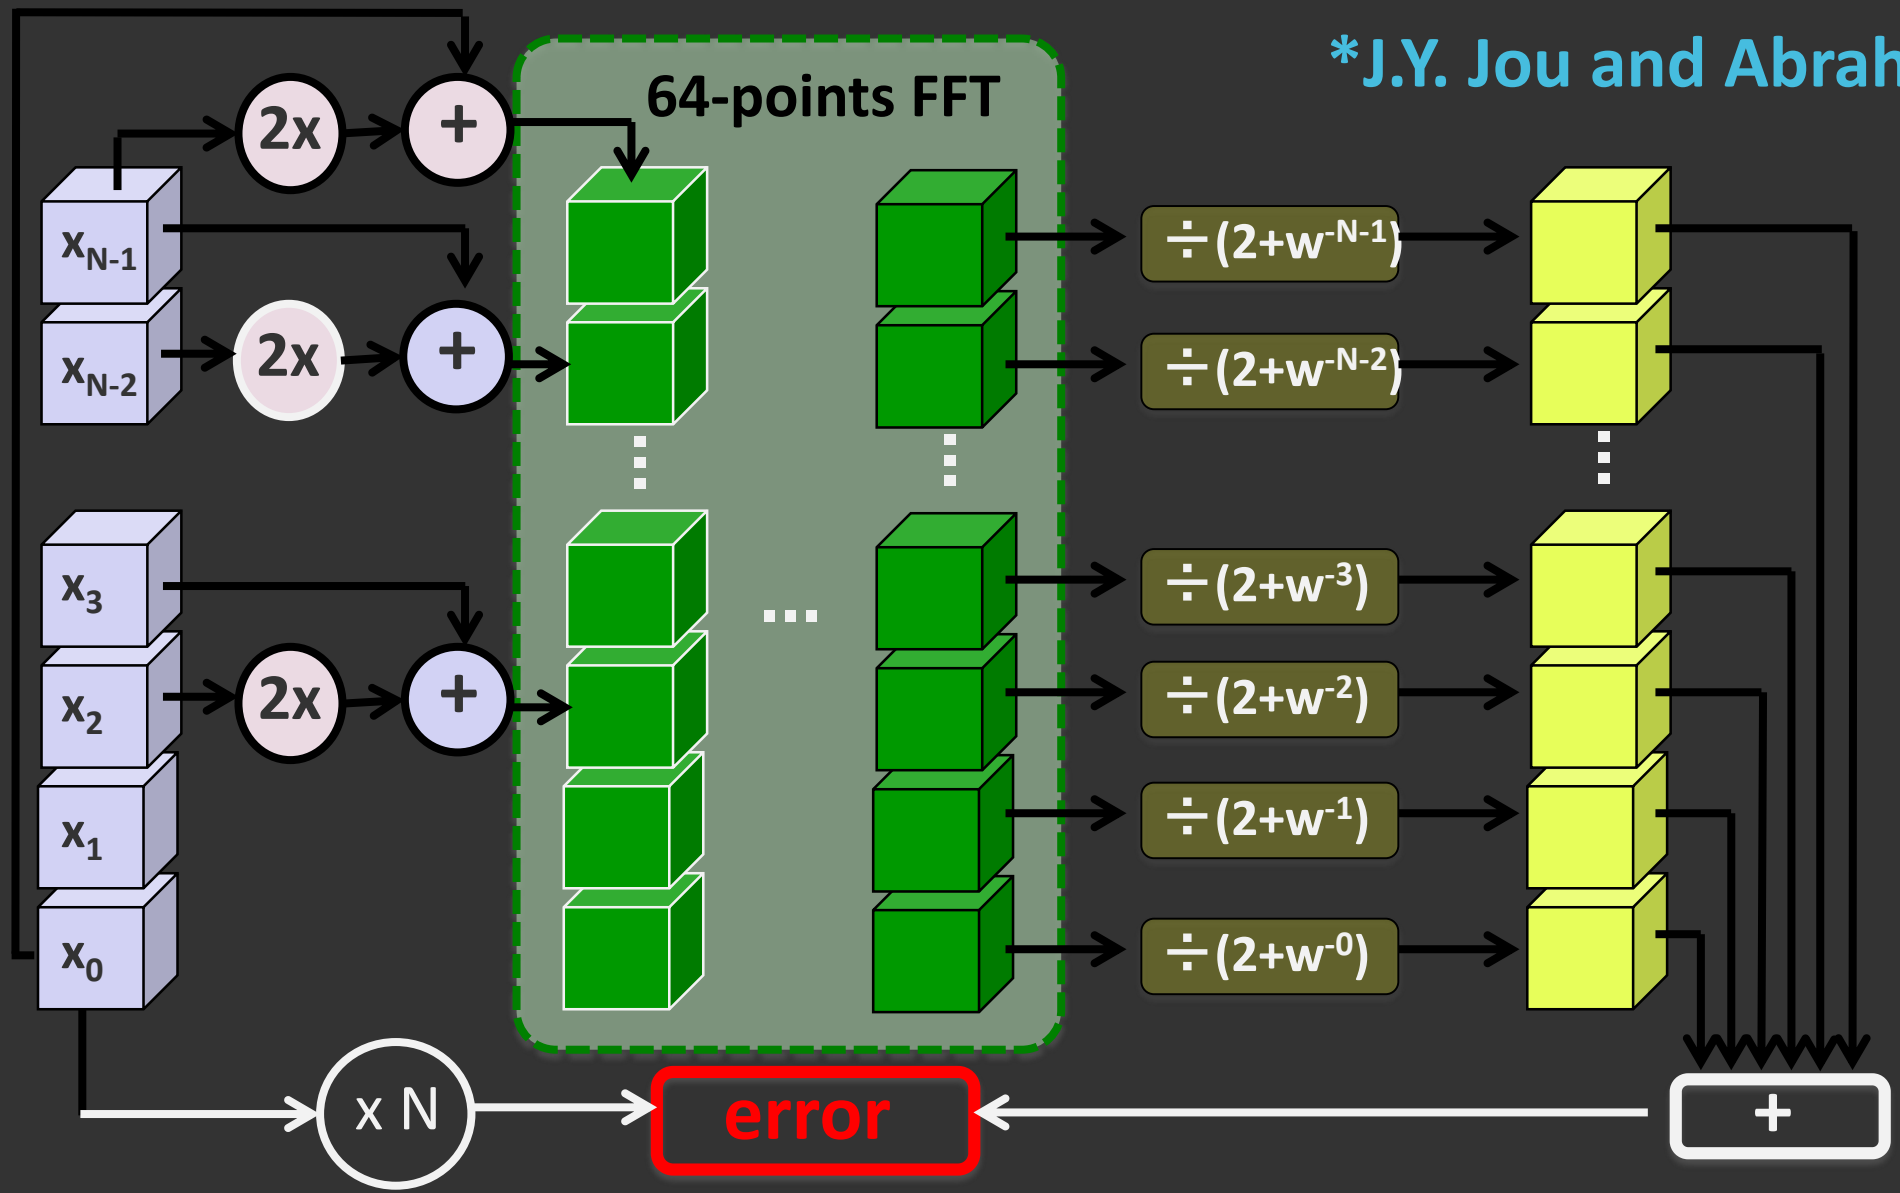

Fault Tolerance Research

ECC x ABFT

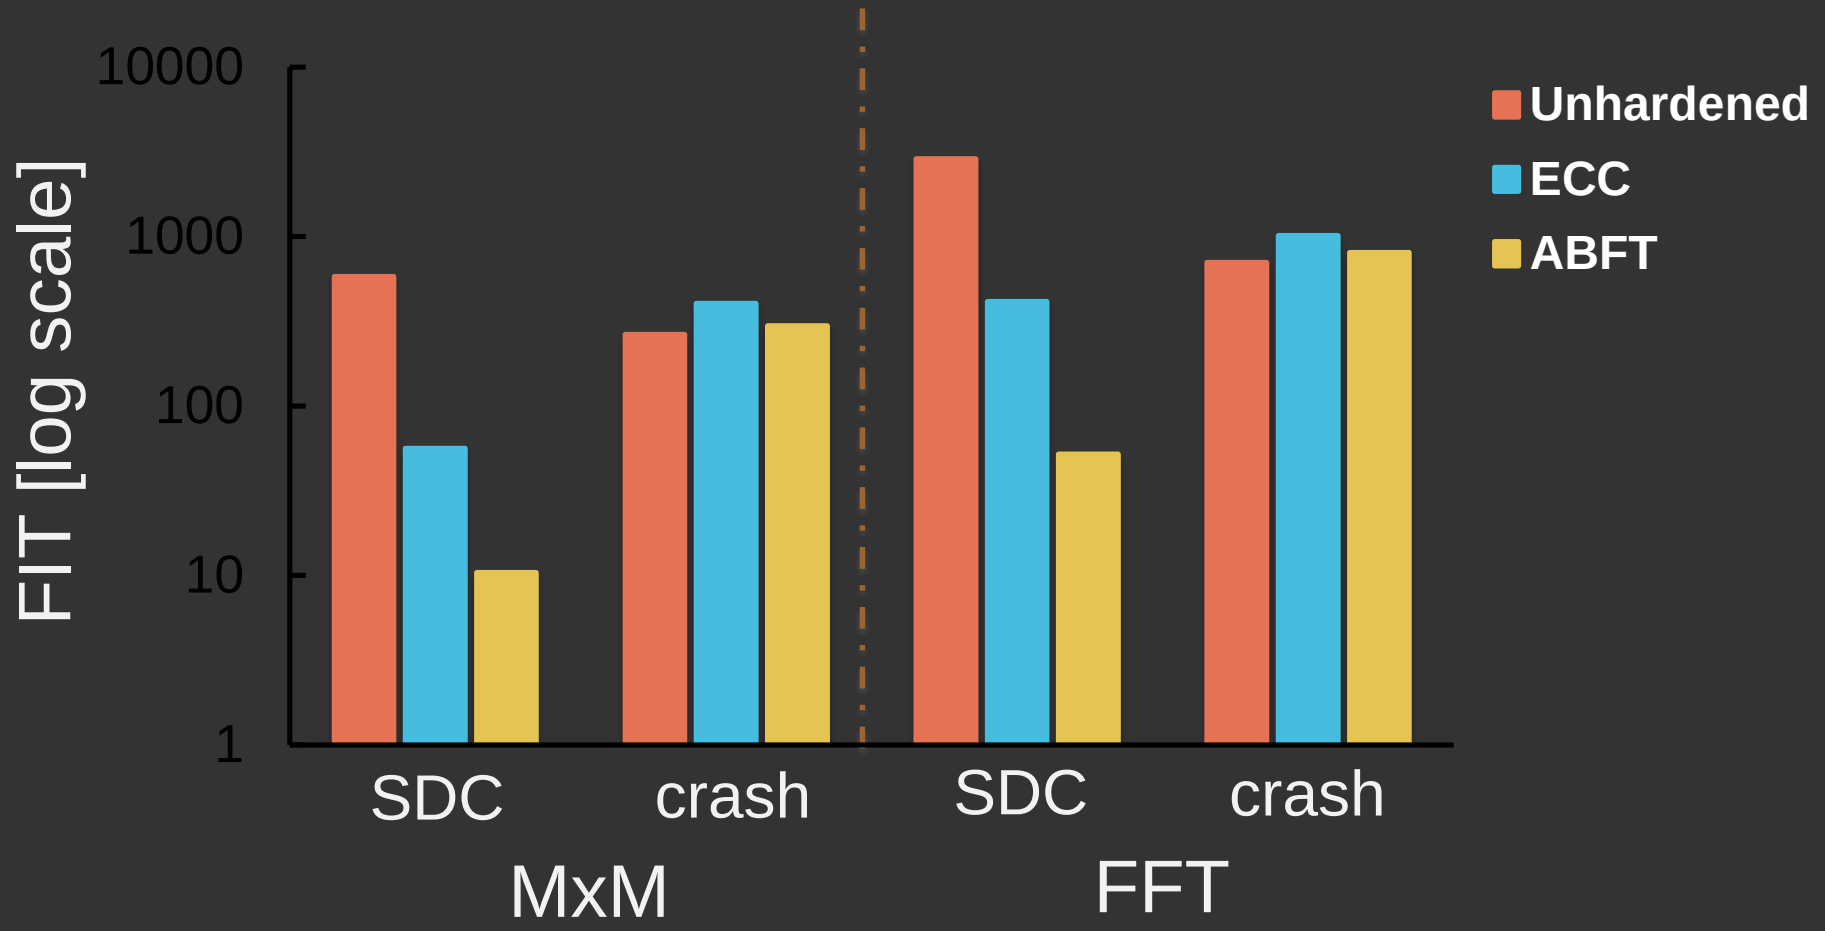

UFSC

ECL

Fault Tolerance Research

ECC x ABFT

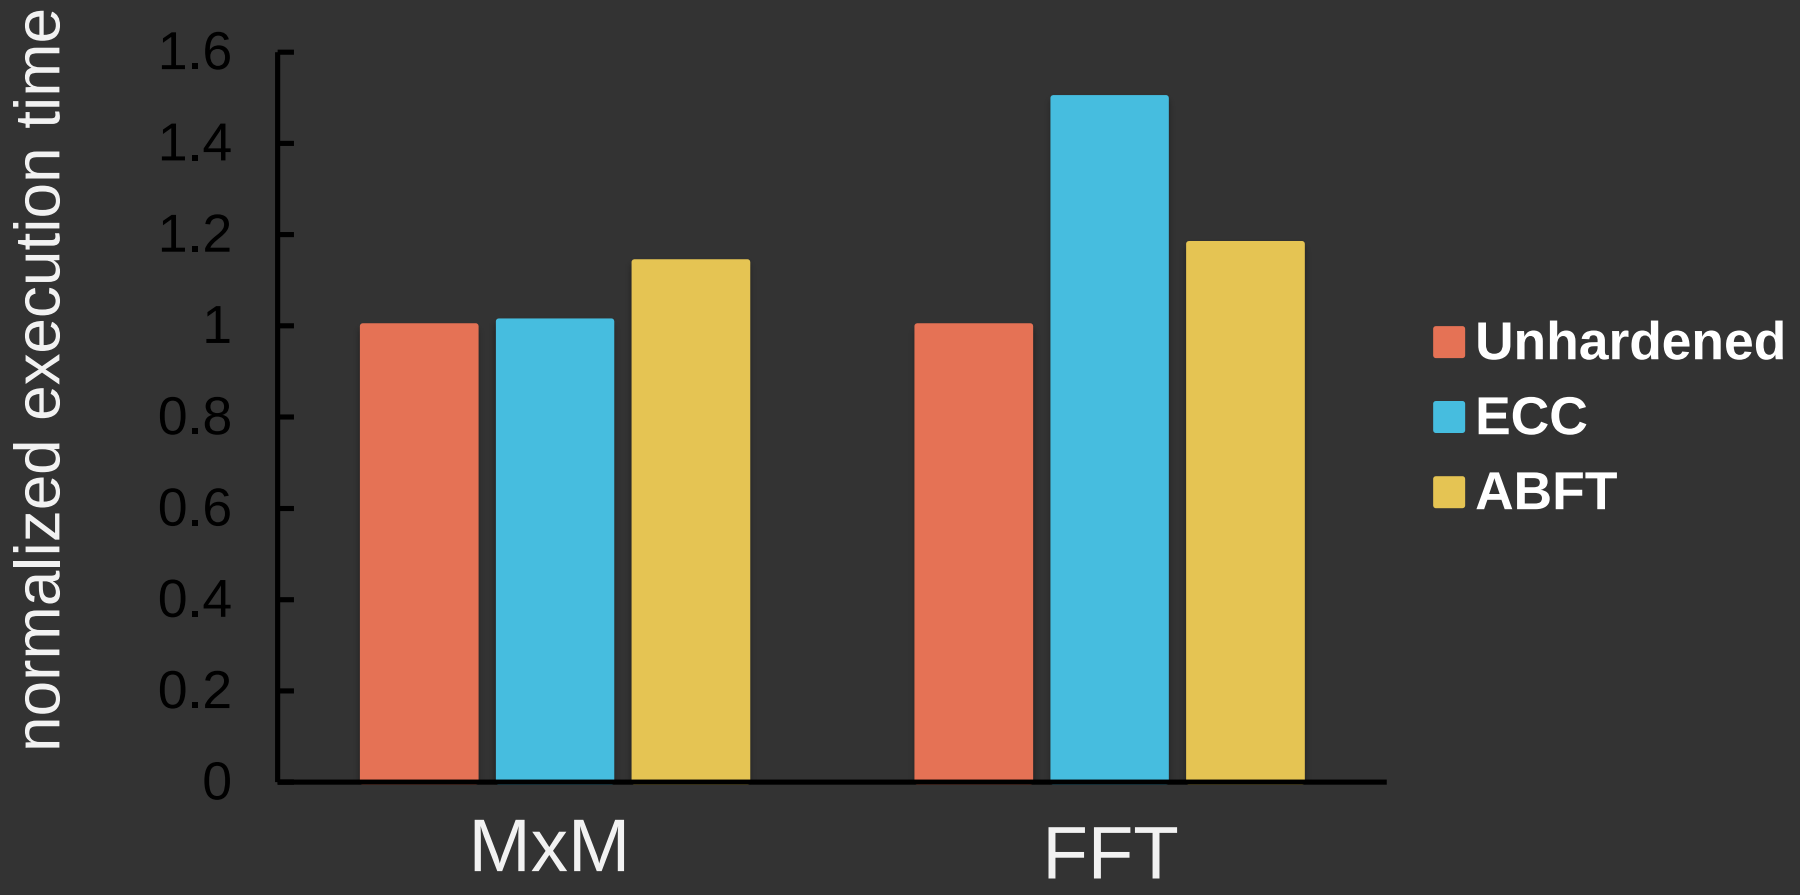

UFSC

ECL

Fault Tolerance Research

Effects of optimizations: MxM

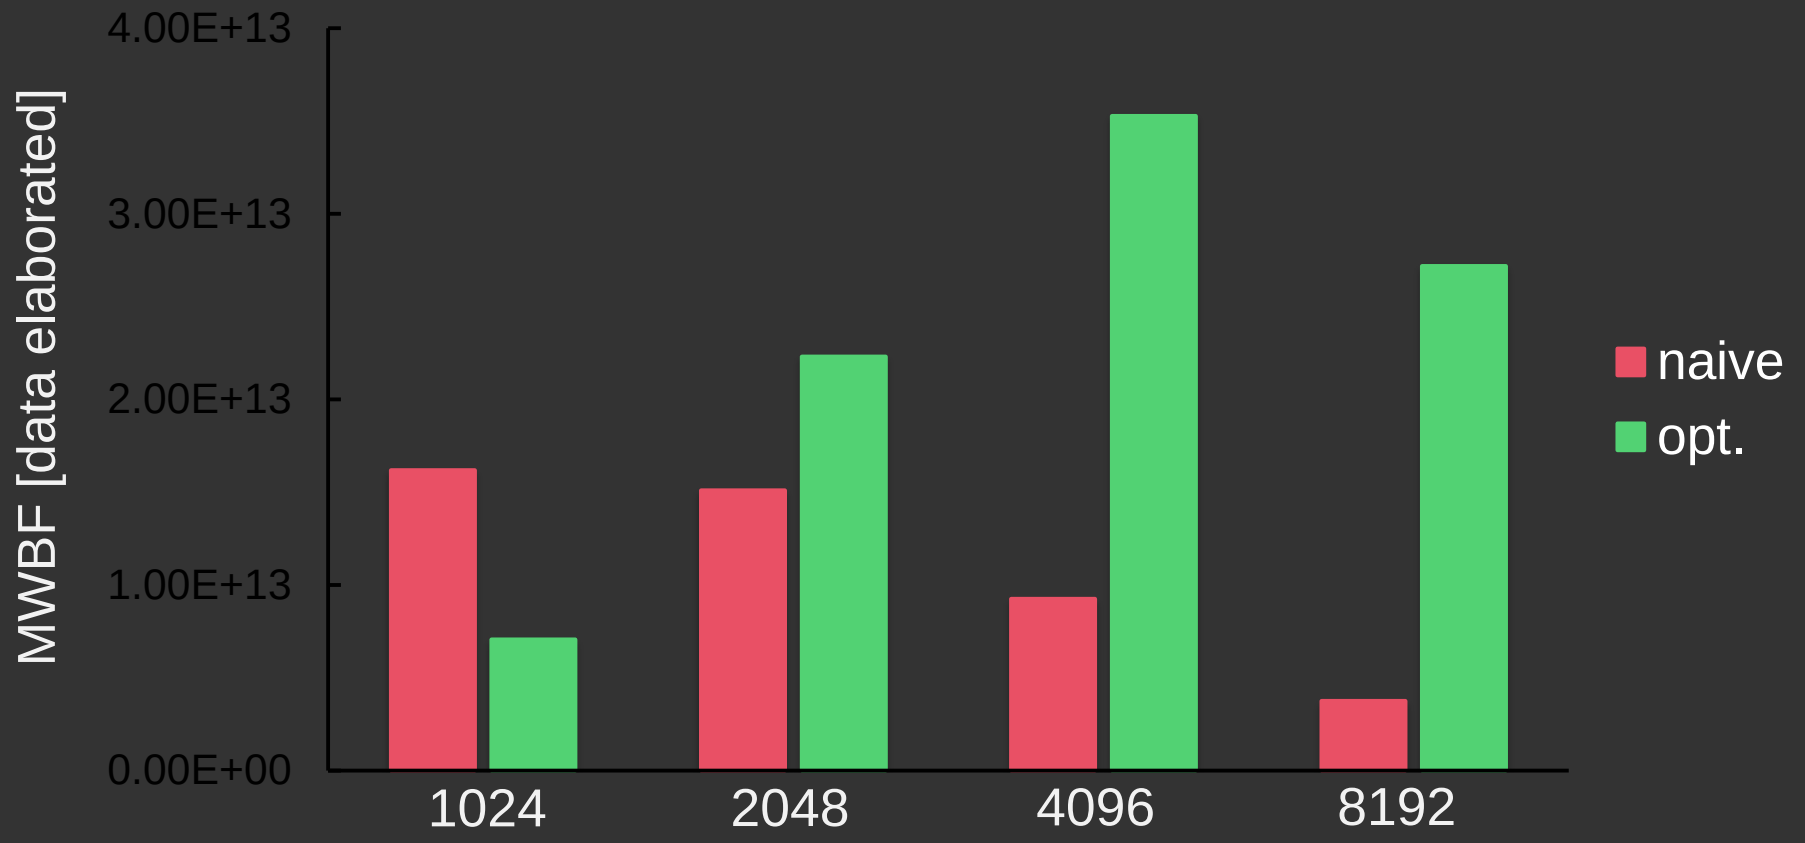

UFSC

ECL

Fault Tolerance Interests

+ applications & kernels for CPU & GPU

+ platforms for energy measurement

+ experiments with radiation

UFSC

ECL

Current Efforts in Scheduling and Fault Tolerance

First EnergySFE Workshop

Grenoble, September 1st, 2016

Laércio Lima Pilla

laercio.pilla@ufsc.br

UFSC

ECL