Alliance for Computing at Extreme Scale (ACES)

Petascale System Challenges

SOS-13

March 9, 2008

John Morrison, LANL

Sudip Dosanjh, SNL

Los Alamos NATIONAL LABORATORY EST.1943

Operated by the Los Alamos National Security, LLC for the DOE/NNSA

Sandia is a Multiprogram Laboratory Operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy Under Contract DE-ACO4-94AL85000.

Sandia National Laboratories

Outline

- LANL & SNL partnership
- Zia, next ASC production capability platform in 2010
- Power/Infrastructure
- Concurrency
- Resilience
- Summary

LANL & SNL Partnership

- 3/2008: LANL & SNL Memorandum of Understanding
- ACES: The NNSA New Mexico Alliance for Computing at Extreme Scale
  - Joint design, architecture, development, deployment and operation of production capability systems for NNSA
- Driven by mission needs
- Commitment to the development and use of world class computing
- Continued leadership in high performance computing
- Sharing intellectual capabilities of both laboratories

ACES Alliance Organization Chart

Senior Executives:
Charlie McMillian, LANL / ADWP
Rick Stulen, SNL / CTO & VP-STE

NNSA Defense Programs

Governance Board:
James Peery, SNL / CCIM & ASC
Rob Leland, SNL / 9300
Andy White, LANL / ADTSC
John Hopson, LANL / ASC

ASC Headquarters:
Bob Meisner, ASC Program Director

Co-Directors:
John Morrison, LANL / HPC
Sudip Dosanjh, SNL / CCIM

Advisory Board
(TBD)

Computer Architecture Office
Sited at Sandia
SNL Office Leader, Jim Ang
LANL Deputy Leader, Karl-Heinz Winkler
Joint SNL & LANL team members

Deployment and Operations Office
Sited at Los Alamos
LANL Office Leader, Nick Nagy
SNL Deputy Leader, John Zepper
Joint LANL & SNL team members

User Requirements Office
Mike Haertling, LANL
Ken Alvin, SNL

Zia Goals for the next ASC Capability Platform

- Petascale production capability
  - Capable of running a single application across the entire machine
  - Usage Model will follow the Capability Computing Campaign (CCC) process
- RFP will specify minimum peak, aggregate memory bandwidth and interconnect bandwidth
- 2GB memory per core (minimum)
- Easy migration of existing integrated weapons codes
- Key challenges: Power, reliability, scalability

High-Level Design Targets for Zia

Zia Performance Targets	Specification	Specification[c]
Peak Floating Point (double precision)	> 2 PF	> 1 PF
Total Memory	> 0.5 PB	> 200 TB
Aggregate[a] Memory BW	> 1 PB/s	> 500 TB/s
Aggregate Interconnect Sustained BW	> 1 PB/s	> 200 TB/s
Aggregate Bisection Sustained BW[b]	> 80 TB/s	> 20 TB/s
Aggregate Message Injection Rate	> 100 GMsgs/s	> 50 GMsgs/s
Aggregate I/O BW	> 1 TB/s	> 400 GB/s
Disk Capacity	> 30 PB	> 10 PB
System Power	< 8 MW	< 8 MW
Floor Space	< 8000 ft^2	< 8000 ft^2
Full System Job MTTI	> 25 hours	> 25 hours
System MTBI	> 200 hours	> 200 hours

[a] Total peak bidirectional rate for all nodes.
[b] For an N-dimensional mesh/torus this is defined as the sum of the bandwidths
[c] For a 2PF option, these non-facility performance specs scale accordingly

UNCLASSIFIED

2008

Submit CD0 5/26/2008

CD0 Approval 12/15/2008

Jan 1 | April 1 | July 1 | Oct 1

Develop Draft RFP Requirements

2009

Submit CD1 2/13/09 | CD1Approval 3/16/09 | Submit CD2/3 6/19/09 | C2/3 Approval 7/20/09

Jan 1 | April 1 | July 1 | Oct 1

Draft RFP Issued 1/30/09

RFP Issued 3/30/09

Proposals Due 4/30/09

System Selection 6/10/09

Contract Award (after DOE/NNSA procurement approval) 7/27/09

2010

Jan 1 | April 1 | July 1 | Oct 1

Zia System Delivery

UNCLASSIFIED

2011

Final Acceptance Test CD-4

L2 Milestones, Transition to Production

Los Alamos NATIONAL LABORATORY EST.1943

Operated by the Los Alamos National Security, LLC for the DOE/NNSA

Sandia National Laboratories

World-class facilities support high-performance computing systems

Nicholas C. Metropolis Center for Modeling and Simulation – 2002

43,500 sq ft computer room floor
9.6 MW Rotary UPS power
4800 tons chilled water cooling

2008 Availability 99.38%

14.4 MW power upgrade planned

Los Alamos NATIONAL LABORATORY EST.1943
Operated by the Los Alamos National Security, LLC for the DOE/NNSA
Sandia National Laboratories

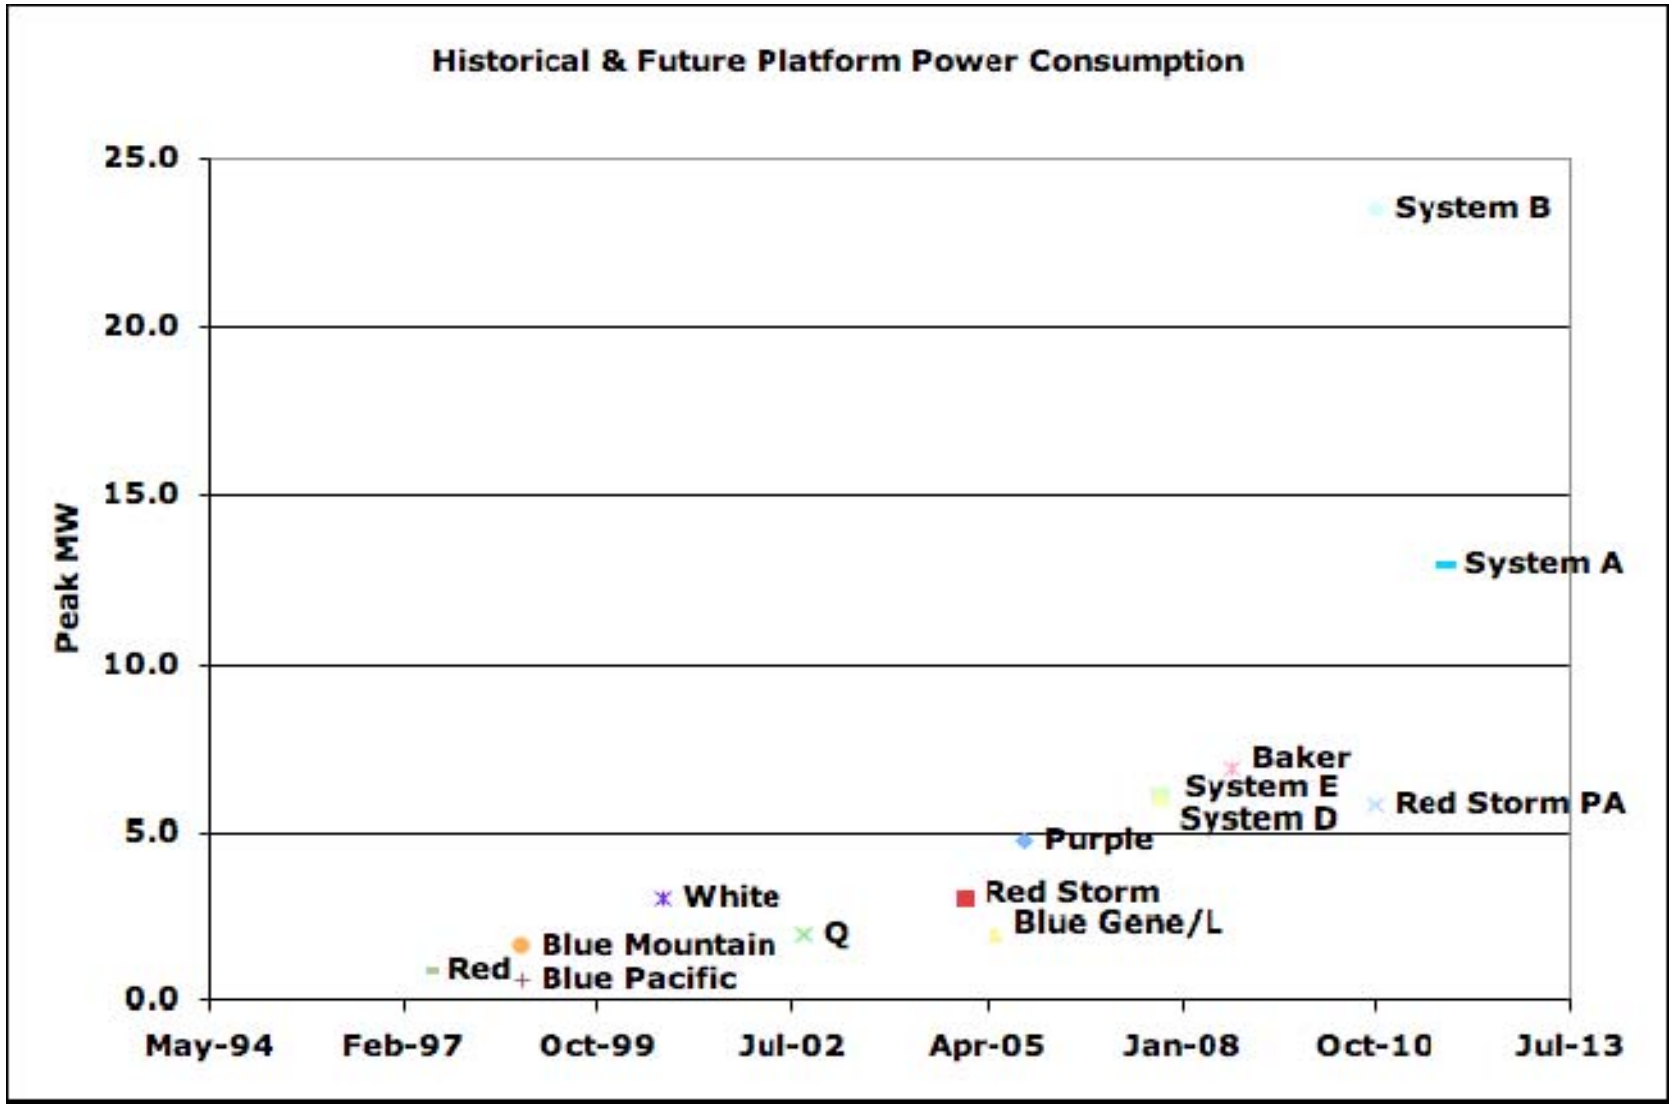
Historical & Future Platform Power Consumption
Peak MW
25.0
20.0
15.0
10.0
5.0
0.0
May-94
Feb-97
Oct-99
Jul-02
Apr-05
Jan-08
Oct-10
Jul-13
System B
System A
Baker
System E
System D
Red Storm PA
Purple
White
Red Storm
Blue Gene/L
Q
Blue Mountain
Red
Blue Pacific

Disturbing Power Trends

DARPA

Level	What	Perf	Power
FPU	FPU, regs,. Instruction-memory	1.5 Gflops	30mW
Core	4FPUs, L1	6 Gflops	141mW
Processor Chip	742 Cores, L2/L3, Interconnect	4.5 Tflops	214W
Node	Processor Chip, DRAM	4.5Tflops	230W
Group	12 Processor Chips, routers	54Tflops	3.5KW
rack	32 Groups	1.7 Pflops	116KW
System	583 racks	1 Eflops	67.7MW

- 1 EF System in 2020 may be 67MW
- With best in class cooling overheads and infrastructure power this is about 80MW
- Power costs/MW will not decrease, currently they are about $1M/MW-year
- Yes Virginia, that is $80M/year
- Even if this is off by a factor of 2-4 this is still a huge concern

FY10 Red Infrastructure Diagram with Zia

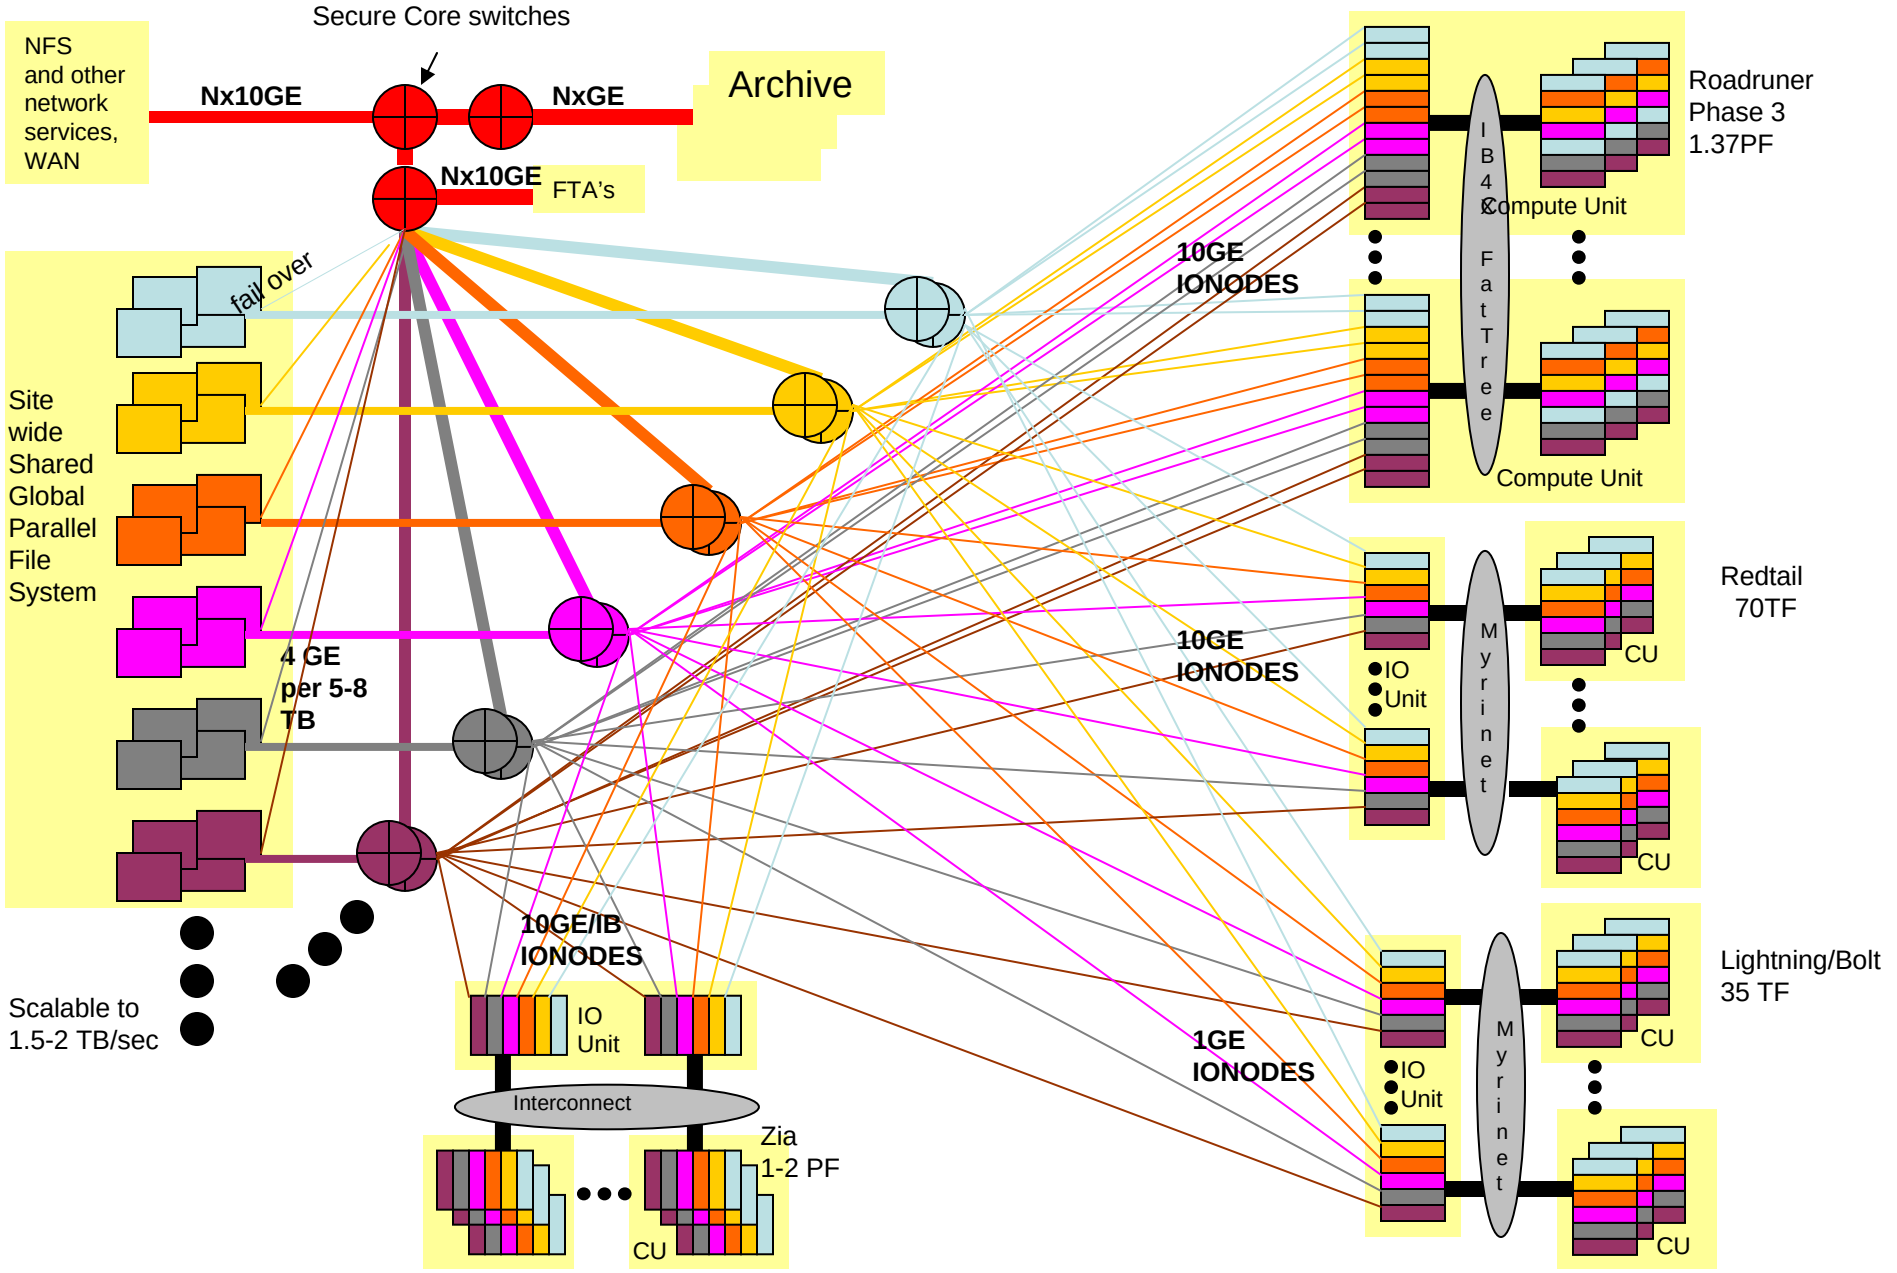

Examples of Disturbing Infrastructure Trends

Technology Capability Increases

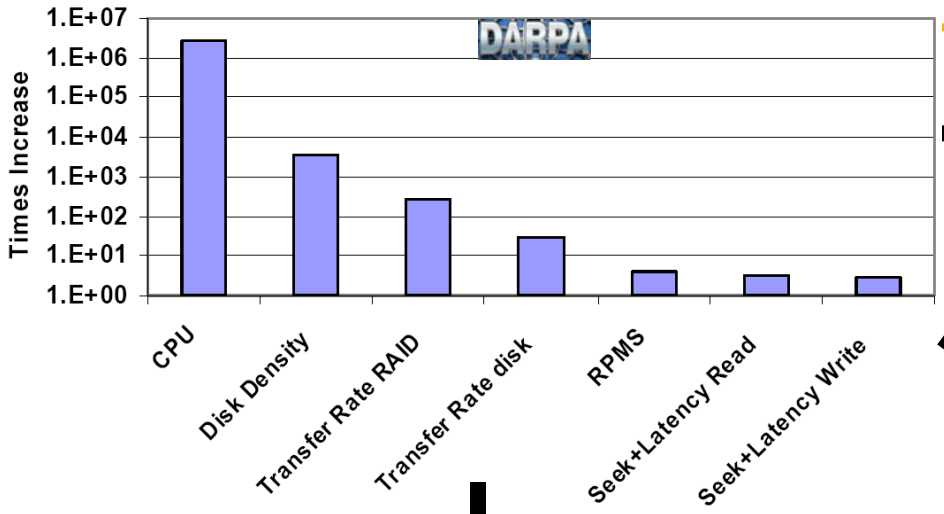

Machine/Environment Size Increases

DARPA	2010 PF	media	2020 EF	media
Memory	0.5 PB		50 PB	
Scratch	20 PB	20k	2 EB	312k
Archive	100 PB	100k	100 EB	1500k

- Gap between CPU and storage technologies continues to grow
- Storage costs have been and are expected to be constant per media part, assuming no miracles
- Implications
  - We could spend 15X current $ on this type of infrastructure or we will have extremely low machine utilization
  - Need to get more serious about resilience and infrastructure R&D
  - Need to expect to spend higher % of total costs on infrastructure in the short term and hope that more serious R&D in the area pays off in the long run
- Even if this is off by a factor of 2-4 this is still a big concern

Projected Application Utilization Accounting for Checkpoint Times

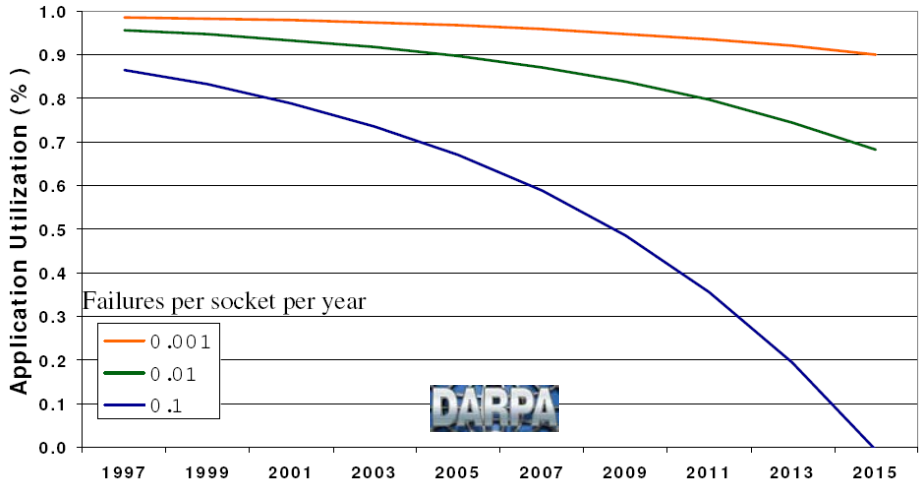

Sandia National Laboratories

Reliability and "Get the Right Answers" are huge issues

- Increased component count and reduced feature size are making problem worse.
- Application initiated checkpoint/restart dominate strategy today
- Resilience
  - Applications apply new strategies to enable computational progress in the face of hardware failures.
    - Will the application developers revolt?
  - Need better tools for identifying failures, perhaps predicting failures.

Concurrency Challenges

- Performance gains driven through increased use of concurrency.
  - Billion “threads” in exascale system!
  - Major impact on applications
  - Integrated multi-physics codes developed over decade time scales
- Approach
  - Engage broader community on programming models
  - Explore new approaches with single physics codes on multiple architectures

SST Preliminary Results

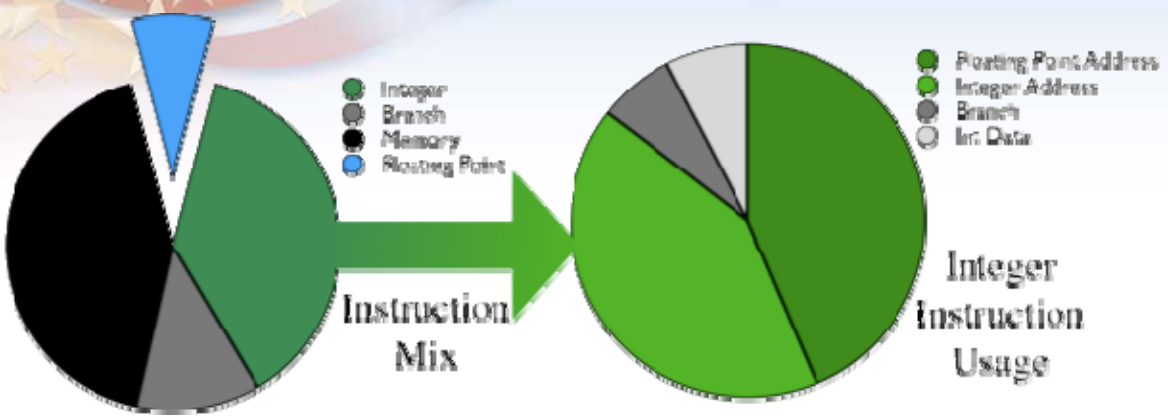

- In the node, Memory performance is key bottleneck
- Even perfect branch prediction and infinite FUs would be less valuable than improving memory latency.
- Prefetching, caches don't help emerging applications
- IEEE Spectrum Article

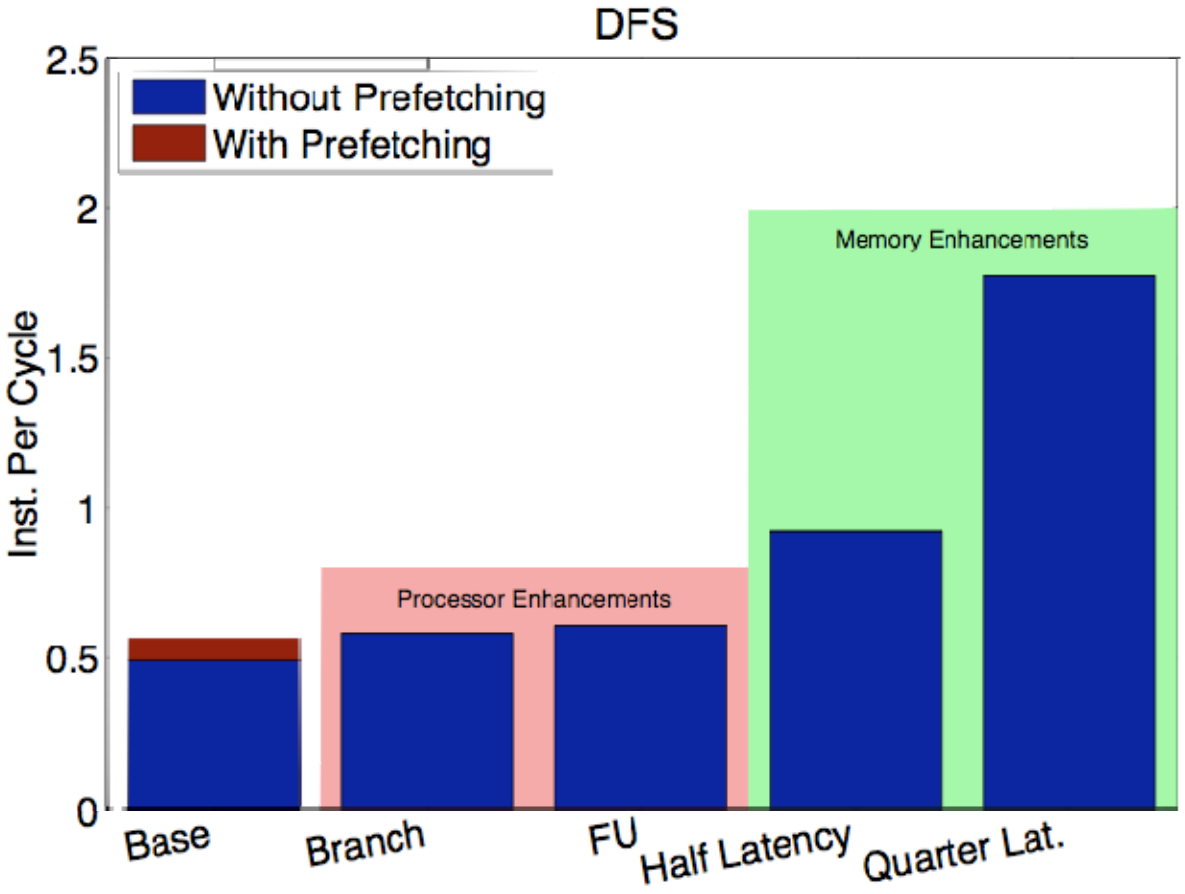

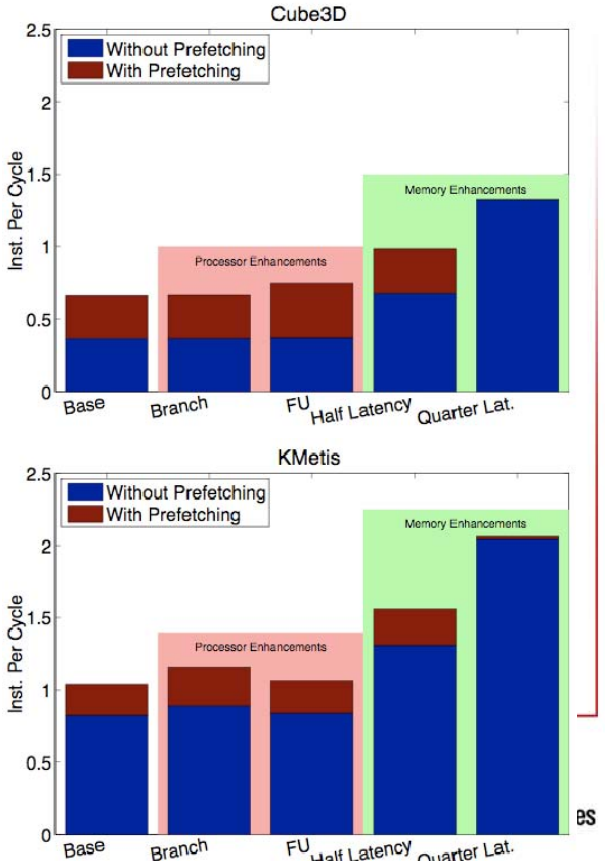

Summary

- Challenges described in Exascale Report are to varying degrees challenges for petascale systems.
- Years of riding clock speed improvements are over and our investment strategy must change if we are to continue to make progress on “simulation science”.
  - Greater emphasis ie funding on dealing with the challenges.
  - Broader collaboration between agencies required to share knowledge and chart the way forward