Energy Efficient Computing

Lennart Johnsson
Advanced Computing Research Lab

UNIVERSITY of HOUSTON

What we do?

- Seek to maximize impact by working with companies, labs, institutes and other Universities to enhance/advance architectures and platforms for HPC
- Validation of designs through energy efficiency and performance benchmarks
- Develop algorithms and software tools

Why is energy efficiency important?

Rule of thumb: 1 MW = $1M/yr in electricity cost

A large data center (Google, Microsoft, Facebook, ….) consumes 100+ MW! (Equiv. of ~25,000 homes)

Energy Challenges

- Of the Total Cost of Ownership (TCO) more than 50% is due to energy for operations and cooling
- For Exa-scale, the next supercomputing platform challenge
  - Business as usual: ~200 MW (not acceptable)
  - Acceptable/Target: 13 MW (embedded)
    20 MW (data center)

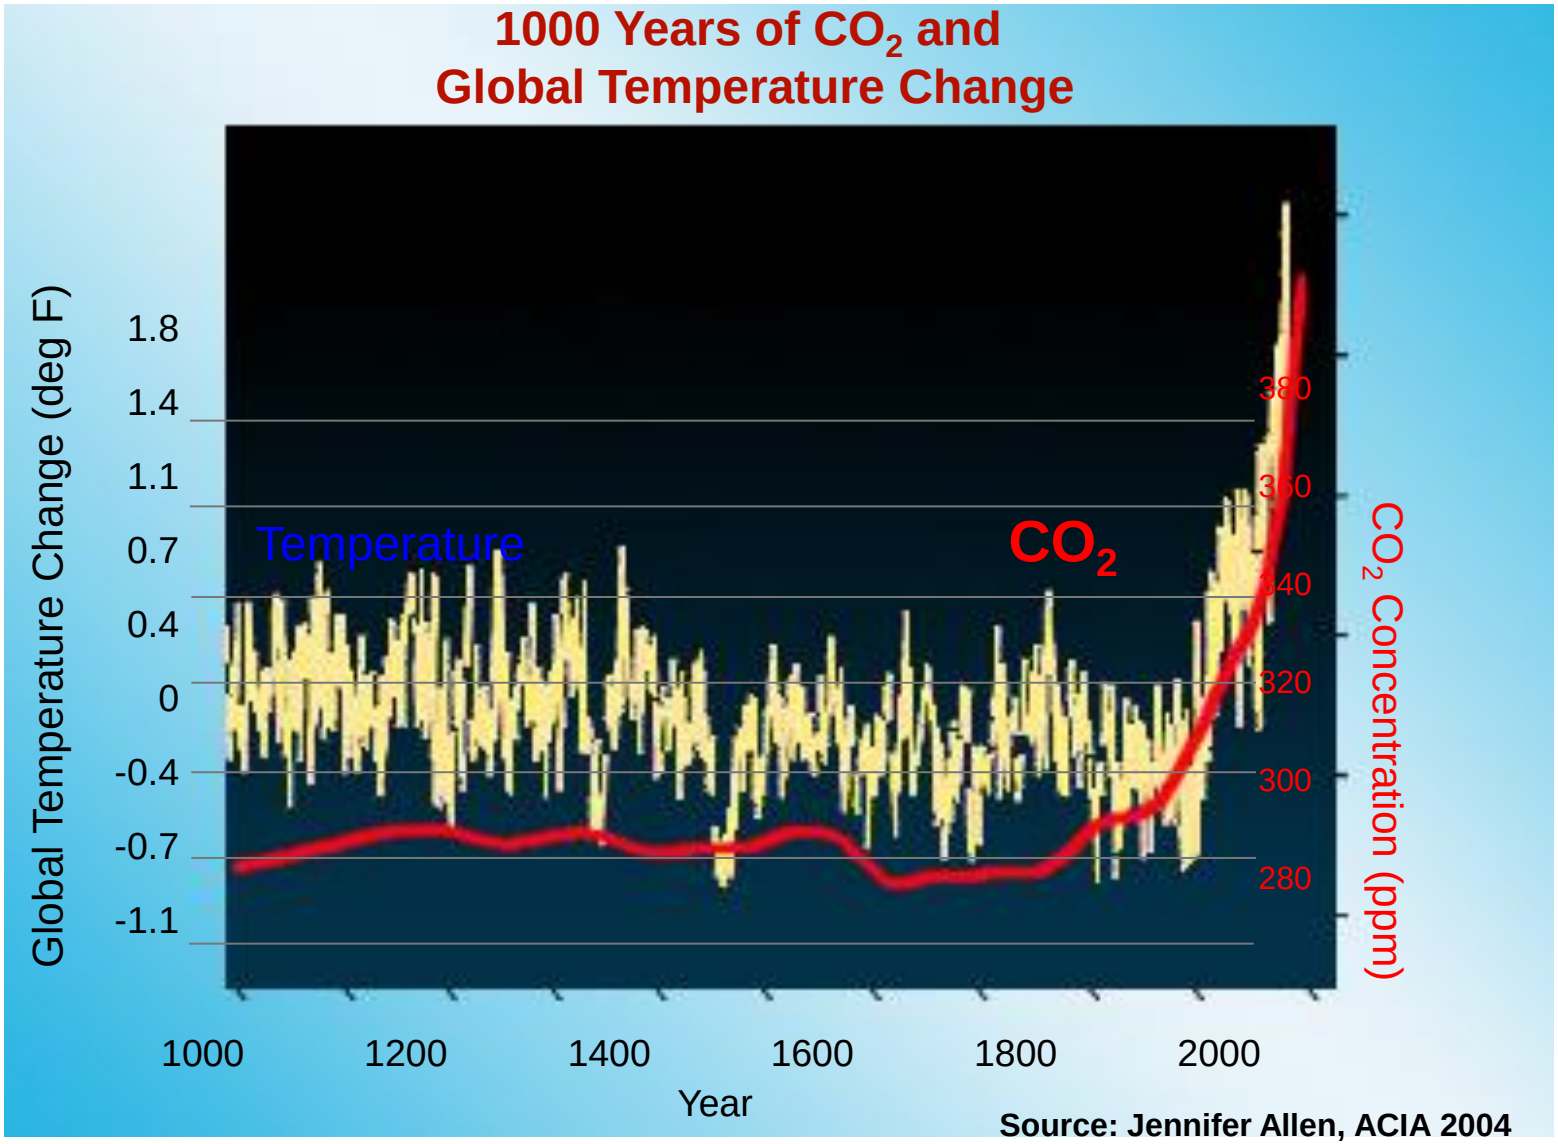

Source: http://alaskaconservationsolutions.com/acs/images/stories/docs/AkCS_current.ppt

Energy Consumption

"We are on the Wrong side of a Square Law" Fred Pollack 1999

New goal for CPU design: "Double *Valued Performance* every 18 months, at the same power level", Fred Pollack

Pollack, F (1999). *New Microarchitecture Challenges in the Coming Generations of CMOS Process Technologies.* Paper presented at the Proceedings of the 32nd Annual IEEE/ACM International Symposium on Microarchitecture, Haifa, Israel.

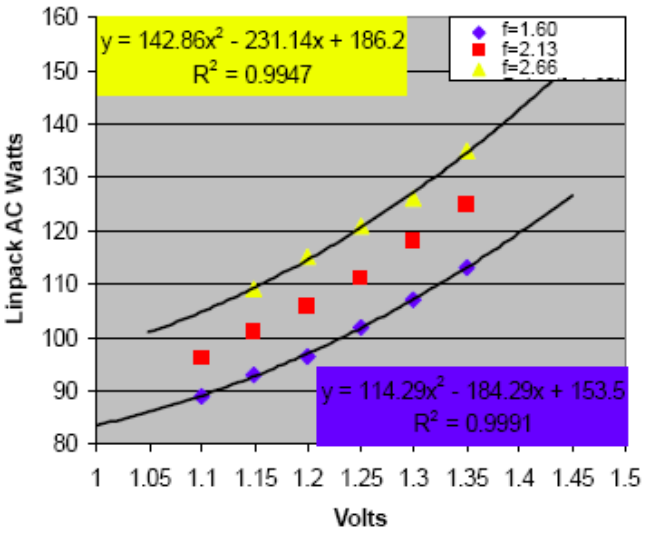

Linpack: $15f(V-0.2)^2+45V+19$

STREAM: $5f(V-0.2)^2+50V+19$

Product	Normalized Performance	Normalized Power	EPI on 65 nm at 1.33 volts (nJ)
i486	1.0	1.0	10
Pentium	2.0	2.7	14
Pentium Pro	3.6	9	24
Pentium 4 (Willamette)	6.0	23	38
Pentium 4 (Cedarmill)	7.9	38	48
Pentium M (Dothan)	5.4	7	15
Core Duo (Yonah)	7.7	8	11

Ed Grochowski, Murali Annavaram Energy per Instruction Trends in Intel® Microprocessors. http://support.intel.co.jp/pressroom/kits/core2duo/pdf/epi-trends-final2.pdf

What type of Architecture?

NERSC
National Energy Research Scientific Computing Center

Reducing Waste

Mark Horowitz 2007: *"Years of research in low-power embedded computing have shown only one design technique to reduce power: **reduce waste**."*

Seymour Cray 1977: *"Don't put anything in to a supercomputer that isn't necessary."*

Office of Science
U.S. Department of Energy

Berkeley Lab

What type of Architecture?

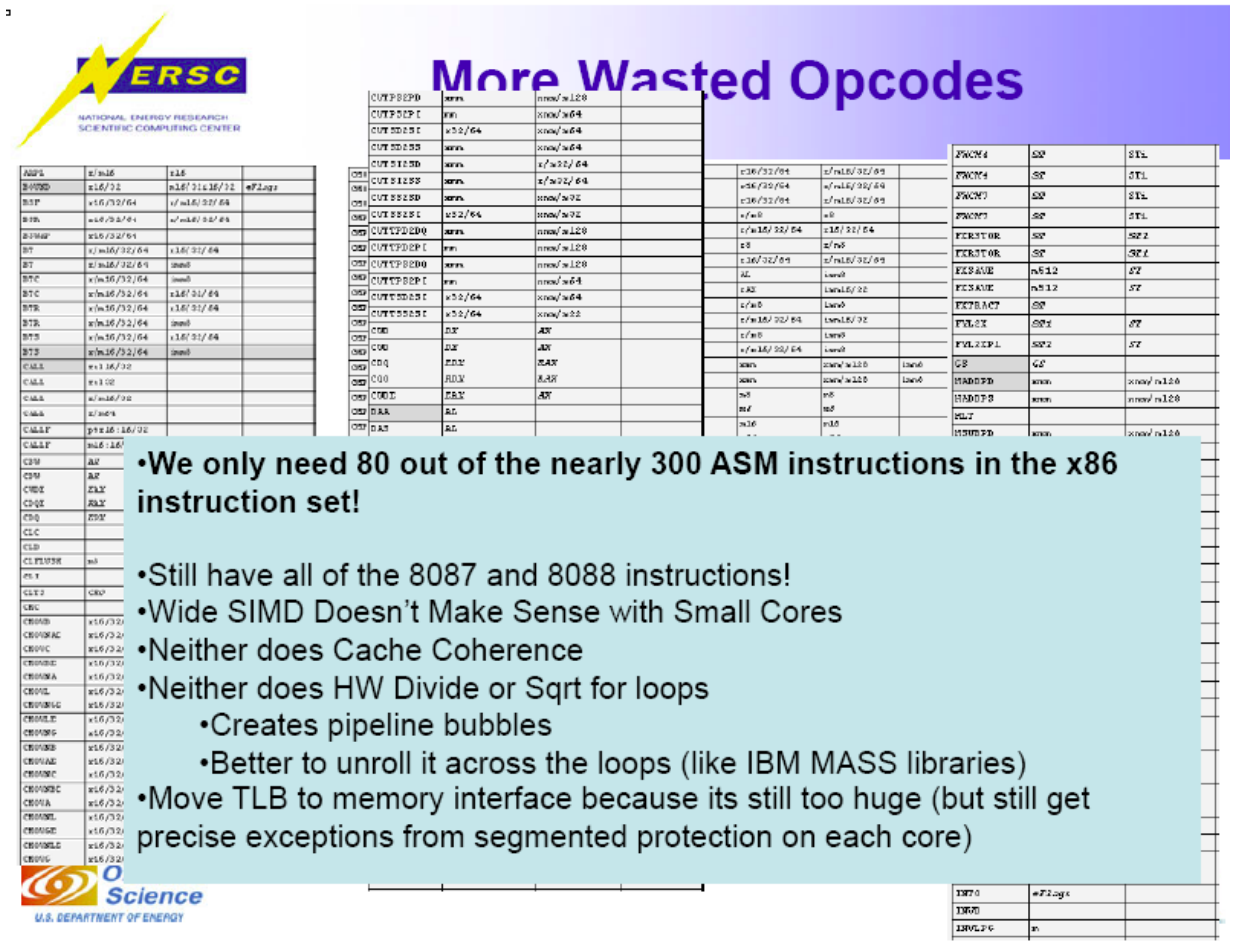

Exascale Computing Technology Challenges, John Shalf
National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory
ScicomP / SP-XXL 16, San Francisco, May 12, 2010

What kind of architecture (core)

How Small is Small

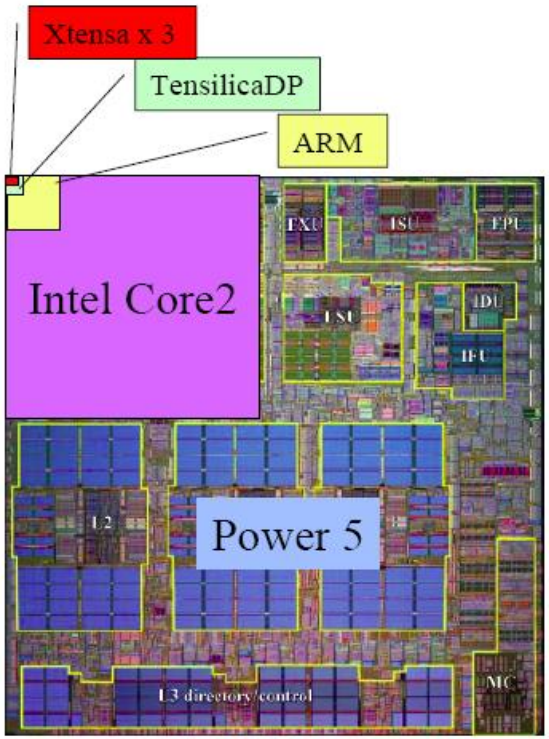

- Power5 (server)
  - 389mm^2
  - 120W@1900MHz
- Intel Core2 sc (laptop)
  - 130mm^2
  - 15W@1000MHz
- ARM Cortex A8 (toaster oven)
  - 5mm^2
  - 0.8W@800MHz
- Tensilica DP (cell phones)
  - 0.8mm^2
  - 0.09W@600MHz
- Tensilica Xtensa (Cisco Rtr)
  - 0.32mm^2 for 3!
  - 0.05W@600MHz

Each core operates at 1/3 to 1/10th efficiency of largest chip, but you can pack 100x more cores onto a chip and consume 1/20 the power

- Cubic power improvement with lower clock rate due to V^2F
- Slower clock rates enable use of simpler cores
- Simpler cores use less area (lower leakage) and reduce cost
- Tailor design to application to reduce waste

http://www.csm.ornl.gov/workshops/SOS11/presentations/j_shalf.pdf

Energy Cost of Operations

Operation	Energy (pJ)
64b Floating FMA (2 ops)	100
64b Integer Add	1
Write 64b DFF	0.5
Read 64b Register (64 x 32 bank)	3.5
Read 64b RAM (64 x 2K)	25
Read tags (24 x 2K)	8
Move 64b 1mm	6
Move 64b 20mm	120
Move 64b off chip	256
Read 64b from DRAM	2000

http://www.lbl.gov/cs/html/Manycore_Workshop09/GPU%20Multicore%20SLAC%202009/dallyppt.pdf

New cost metric

operations inexpensive

moving data expensive!

SNIC/KTH PRACE Prototype I

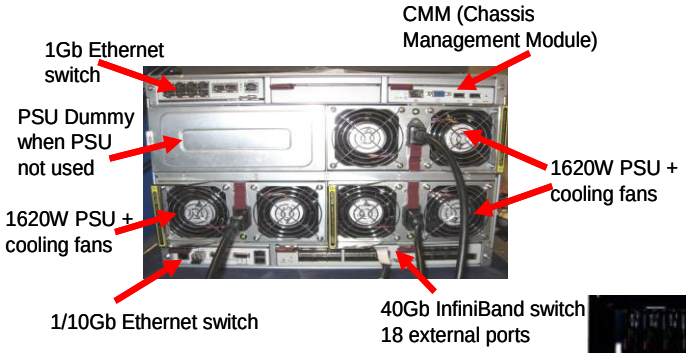

- New 4-socket blade with 4 DIMMs per socket supporting PCI-Express Gen 2 x16
- Four 6-core 2.1 GHz 55W ADP AMD Istanbul CPUs, 32GB/node
- 10-blade in a 7U chassis with 36-port QDR IB switch, new efficient power supplies.
- 2TF/chassis, 12 TF/rack, 30 kW (6 x 4.8)
- 180 nodes, 4320 cores, full bisection QDR IB interconnect

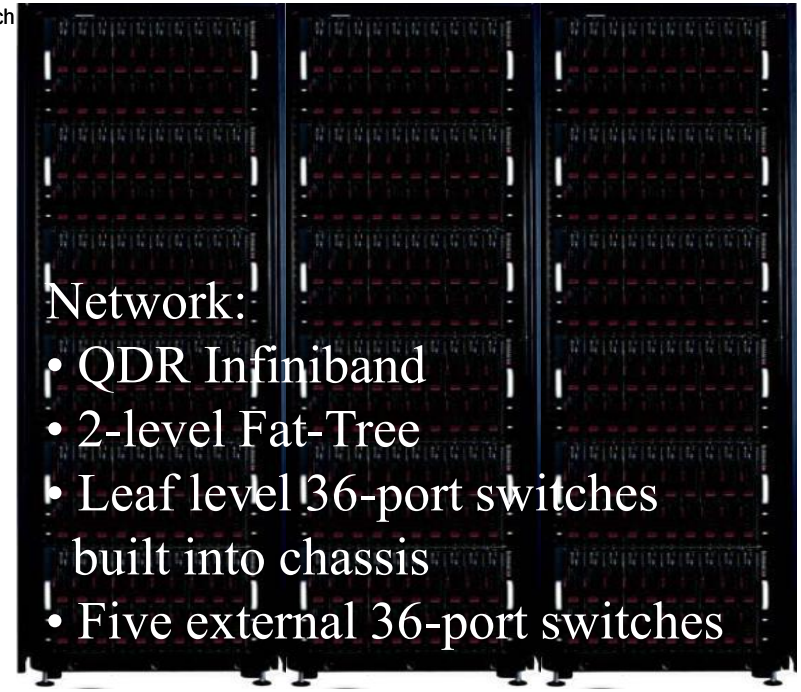

Became Supermicro product!

SNIC/KTH/PRACE Prototype I

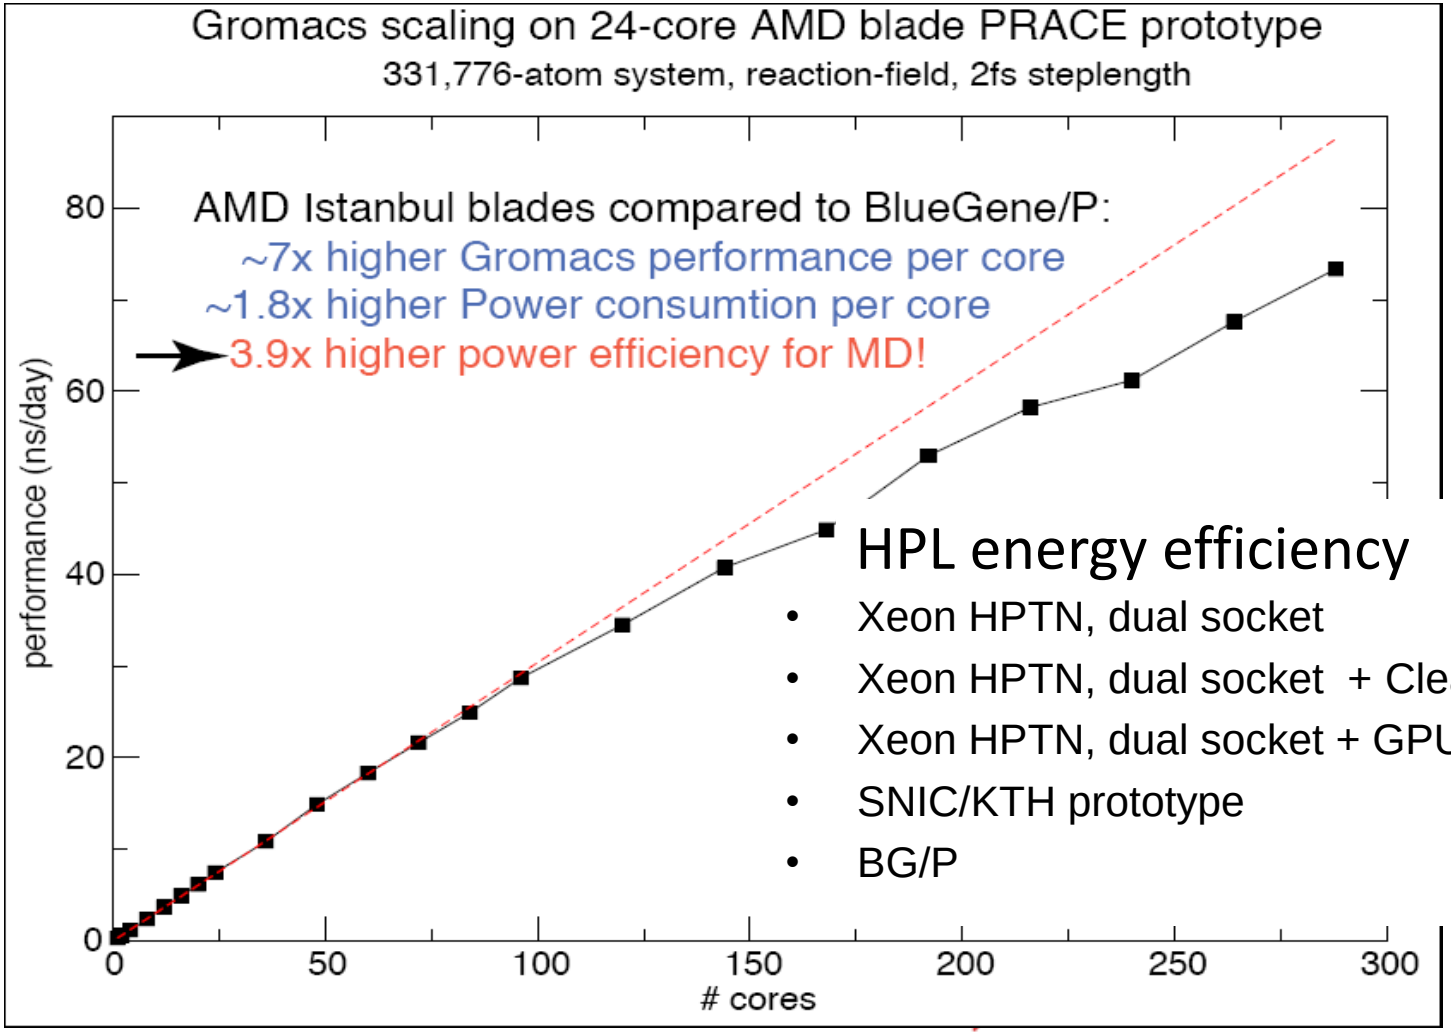

HPL energy efficiency

- Xeon HPTN, dual socket 240 MF/W
- Xeon HPTN, dual socket + Clearspeed 326 MF/W
- Xeon HPTN, dual socket + GPU 270 MF/W
- SNIC/KTH prototype 344 MF/W
- BG/P 357 MF/W

KTH/SNIC/PRACE DSP HPC node

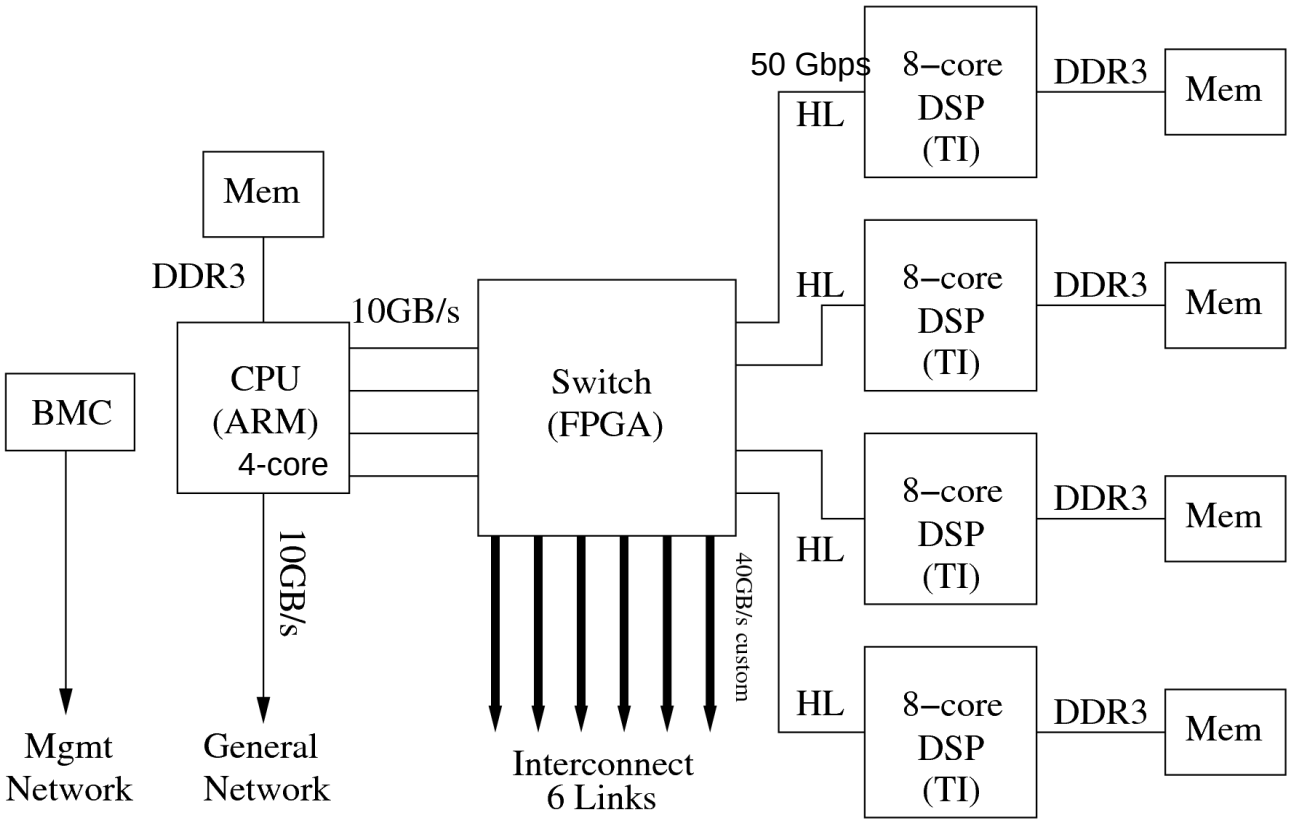

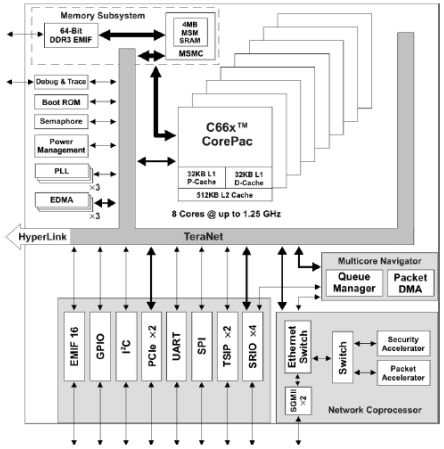

Target:
15 – 20W
32 GB
2.5 GF/W Linpack

TEXAS INSTRUMENTS

UNIVERSITY of HOUSTON

Linpack on TI 6678 EVM

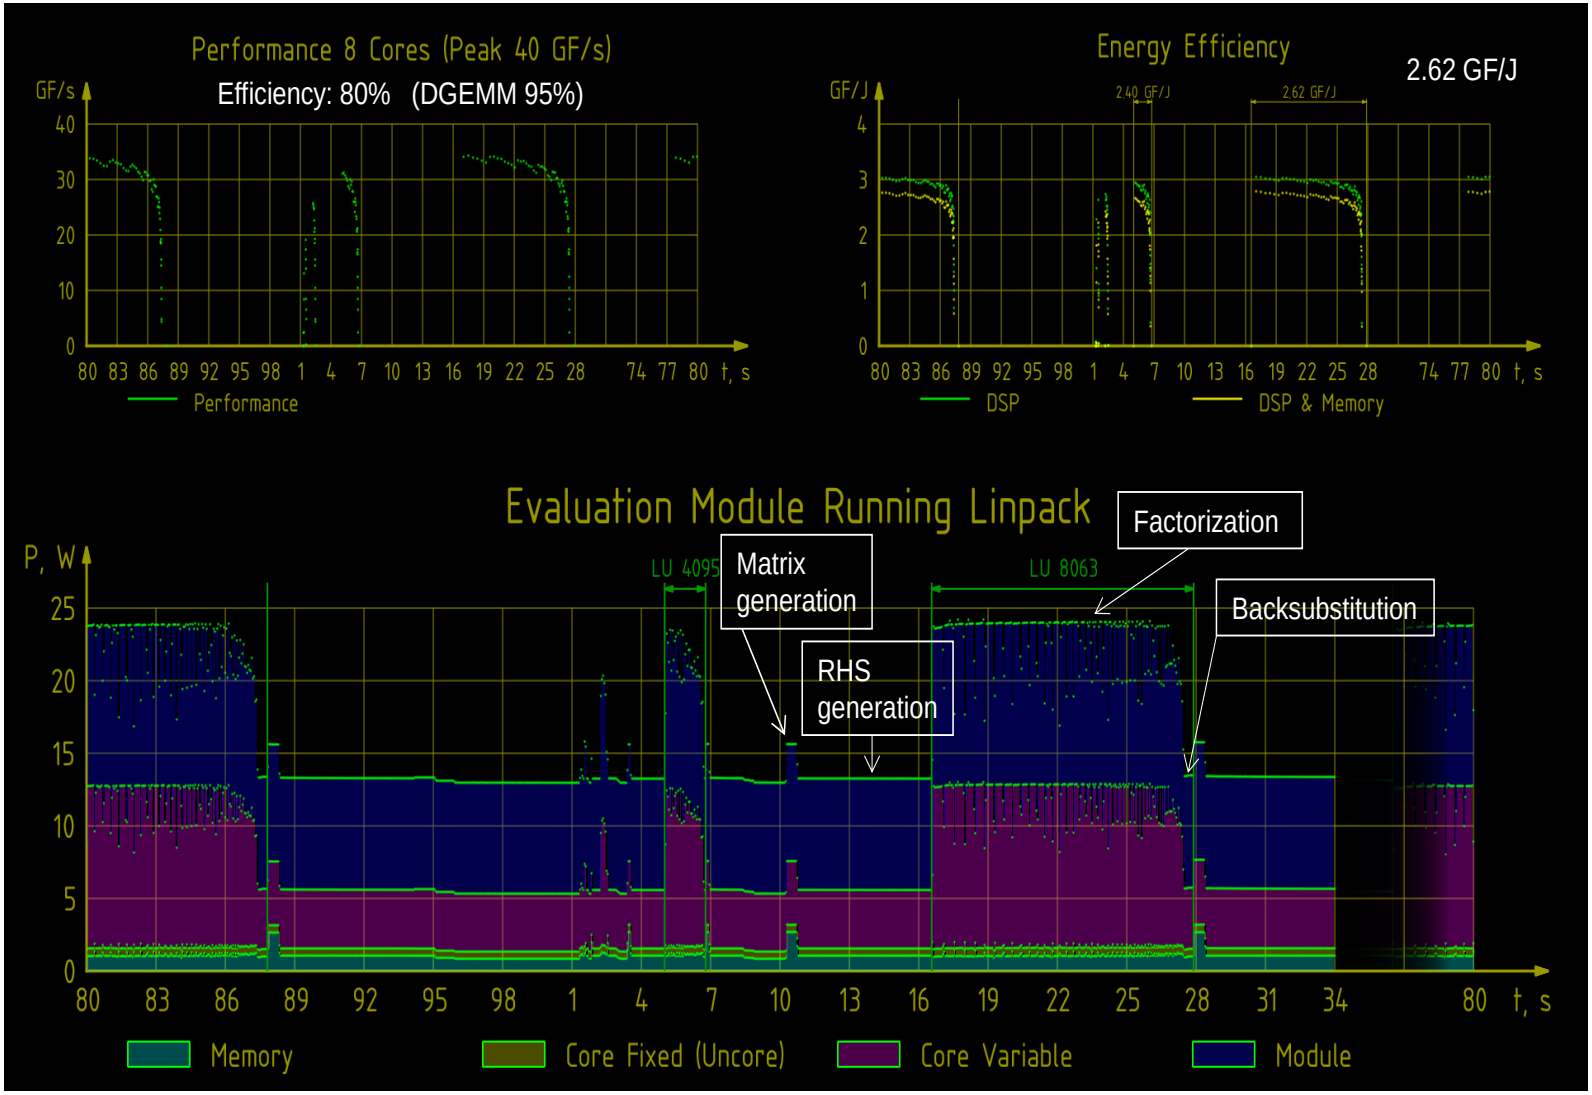

Imagine the impact… TI's KeyStone SoC + HP Moonshot

2013-04-19. Last week, market leader Hewlett Packard announced a huge change in the server landscape with its recent Moonshot announcement. …..

….. "TI's KeyStone II-based SoCs, which integrate fixed- and floating- point DSP cores with multiple ARM® Cortex™A-15 MPCore processors, packet and security processing, and high speed interconnect, give customers the performance, scalability and programmability needed to build software-defined servers."

HP Project Moonshot is dedicated to designing extreme low-energy server technologies. HP expects data center efficiencies to reach new heights for select workloads and applications, consuming up to 89% less energy .

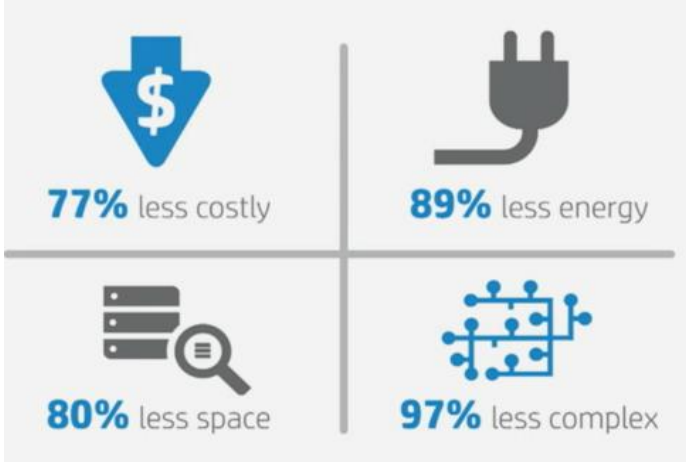

We are pursuing HPC cartridges with HP and TI …..

Next? – Enhanced Mobile Video CPU

Fragrak 28nm Platform by Movidius

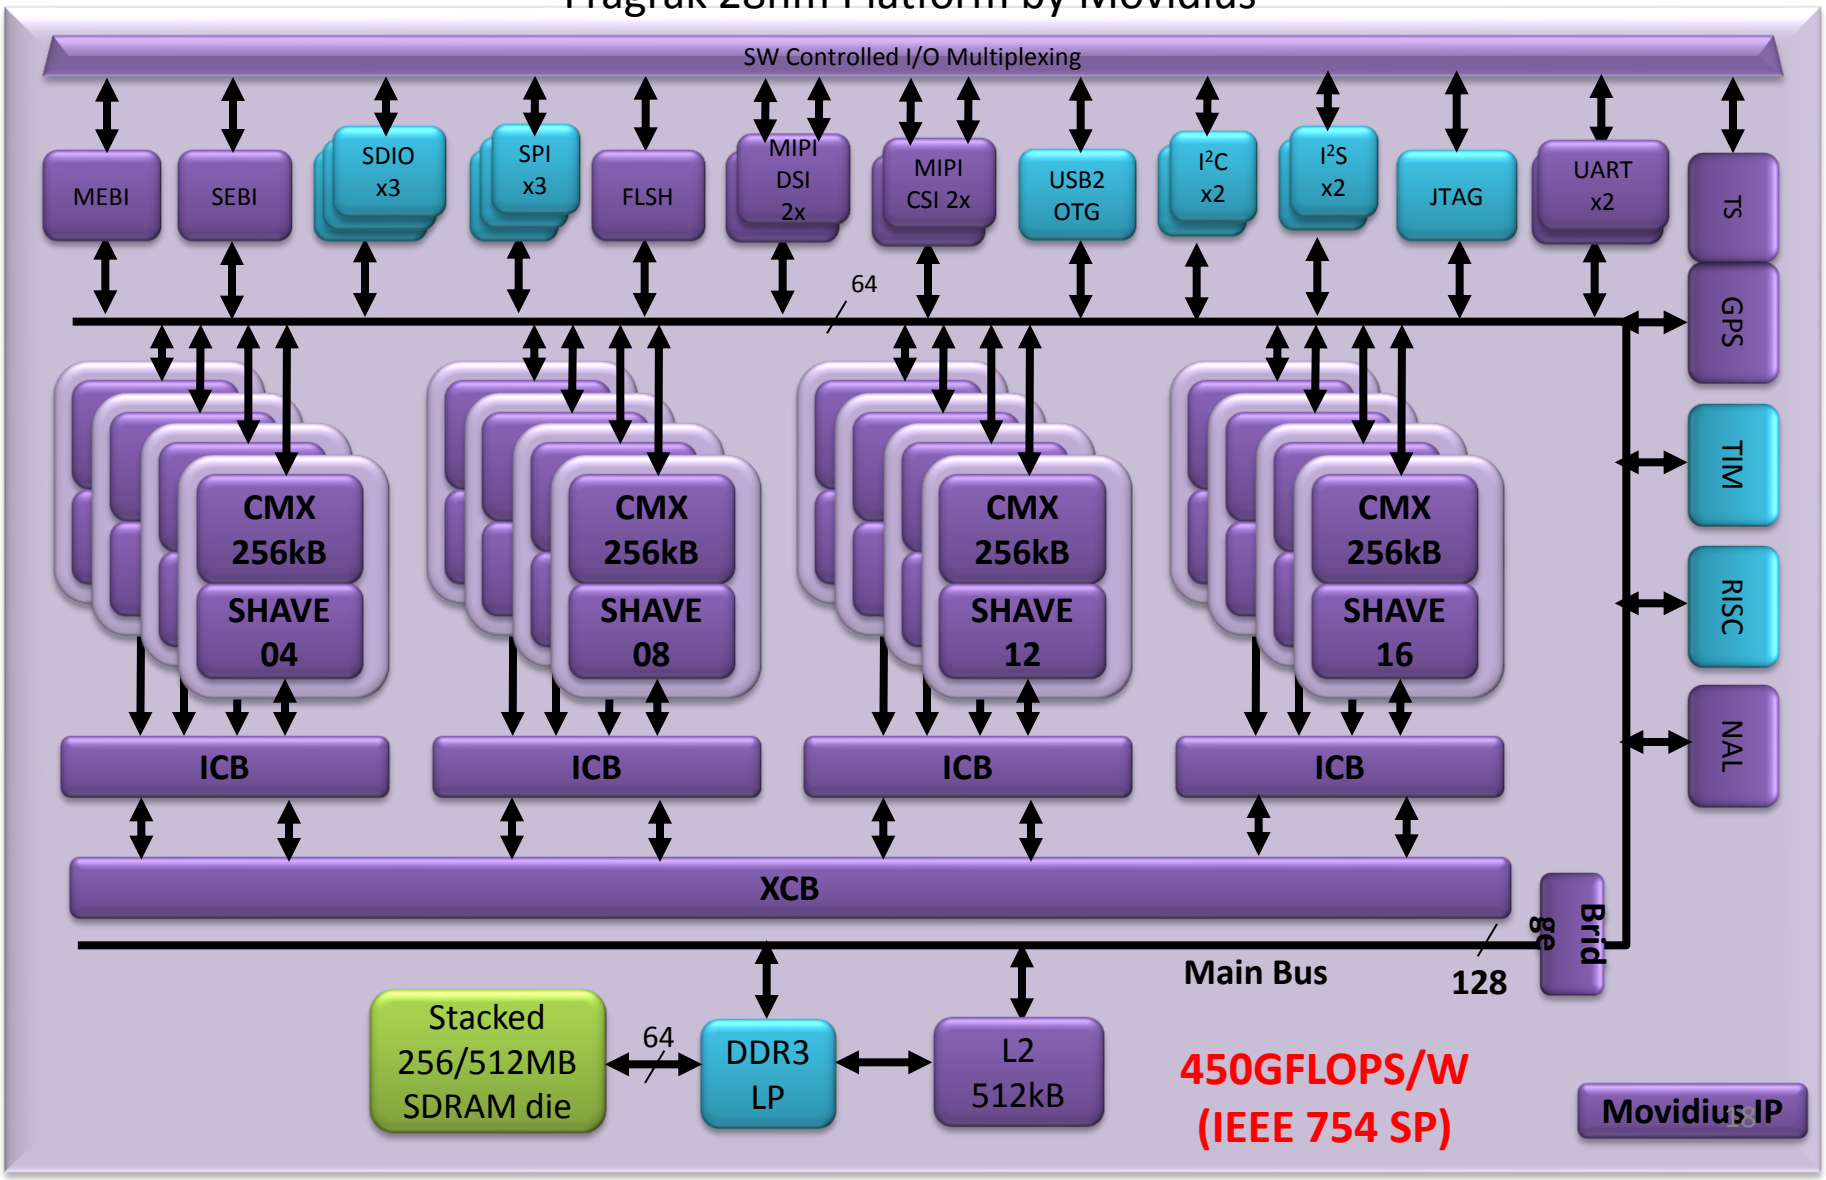

New Students Welcome!!!

Past students employment

Purdue Univ, Prof
Indian Institute of Technology, Prof
MIT, Research Scientist (HCI)
IBM Almaden Research Center, (Big Data)
HP (Security)
HP Research Lab (Grids and Clouds)
Intel, Chief Architect
Microsoft
Microsoft Research
…….
…….