Panel Session:
Multicore Meltdown?

James C. Anderson
MIT Lincoln Laboratory

HPEC07
Wednesday, 19 September 2007

This work was sponsored by the Department of the Air Force under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author, and are not necessarily endorsed by the United States Government.

Objective & Schedule

- **Objective: assess the impact of multicore processors on past and present DoD HPEC systems, and project their importance for the future**

- **Schedule**
  - **1545-1605: Overview**
  - **1605-1615: Guest speaker Mr. Kalpesh Sheth**
  - **1615-1620: Introduction of the panelists**
  - **1620-1650: Previously-submitted questions for the panel**
  - **1650-1710: Open forum**
  - **1710-1715: Conclusions & the way ahead**

On-chip Performance for Representative COTS Single- & Multi-core Processors (2Q98)

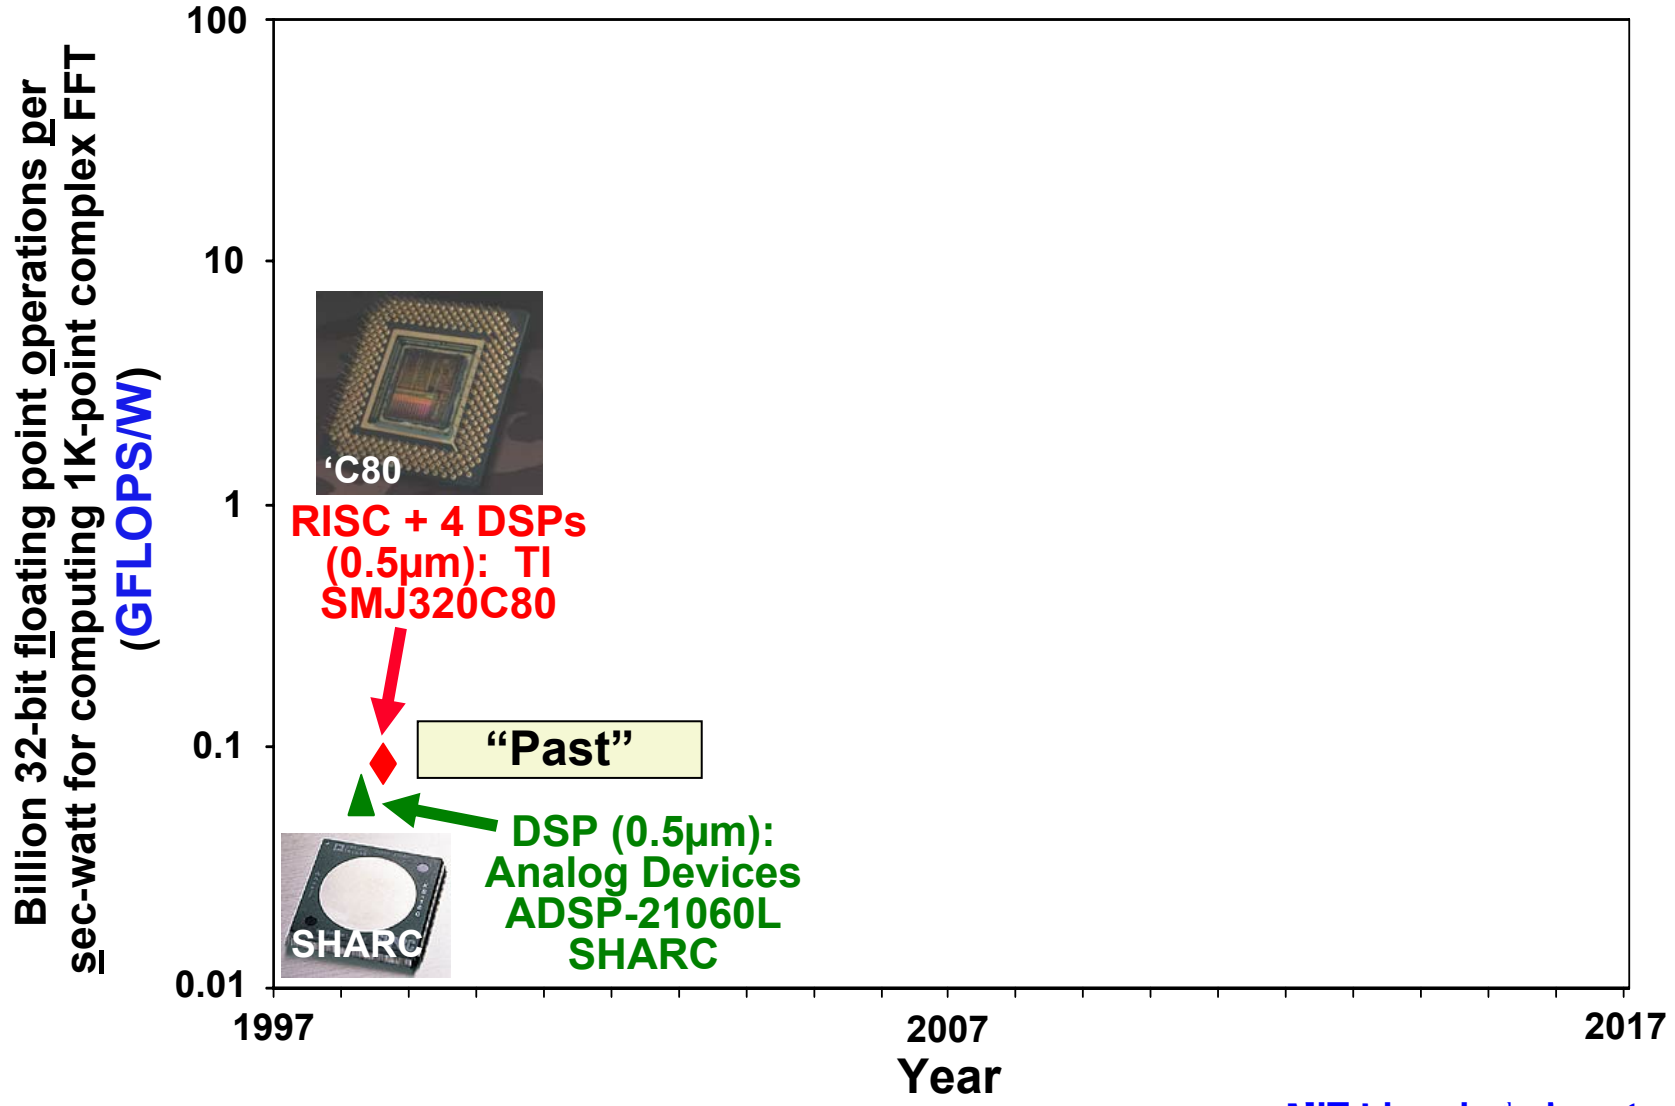

On-chip Performance for Representative COTS Single- & Multi-core Processors (3Q07)

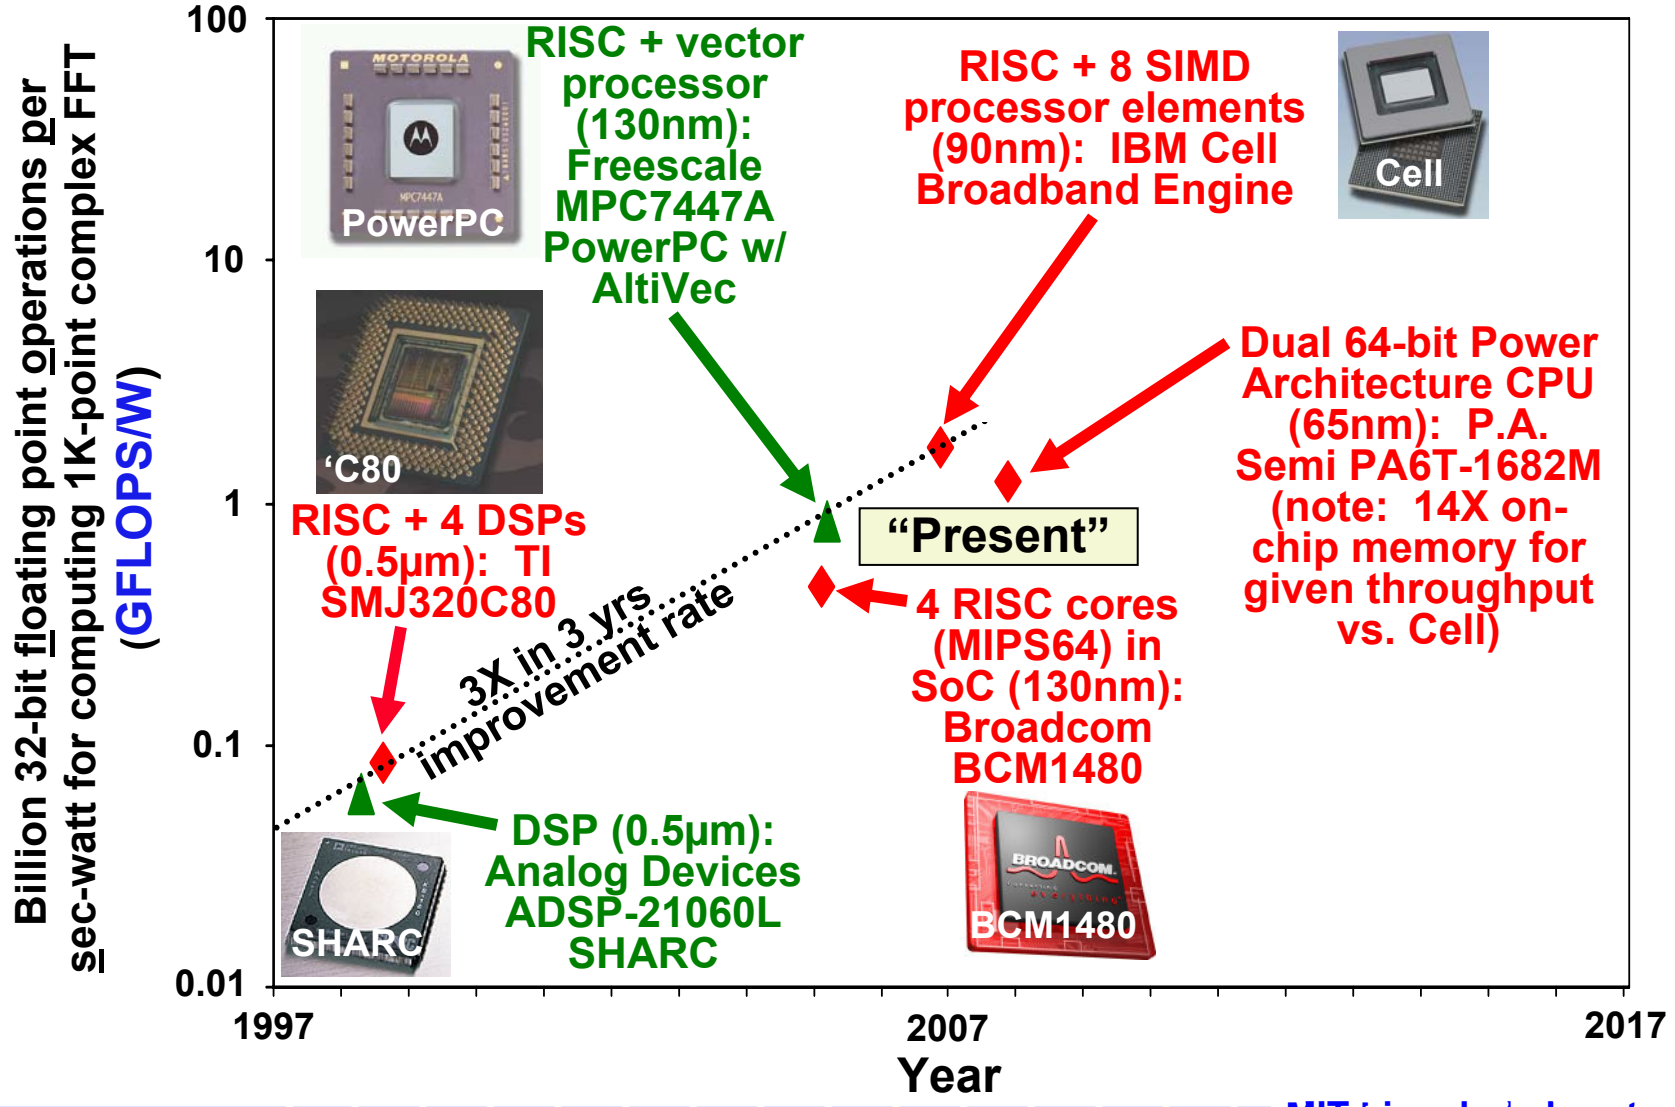

On-chip Performance for Possible Future COTS Single- & Multi-core Processors

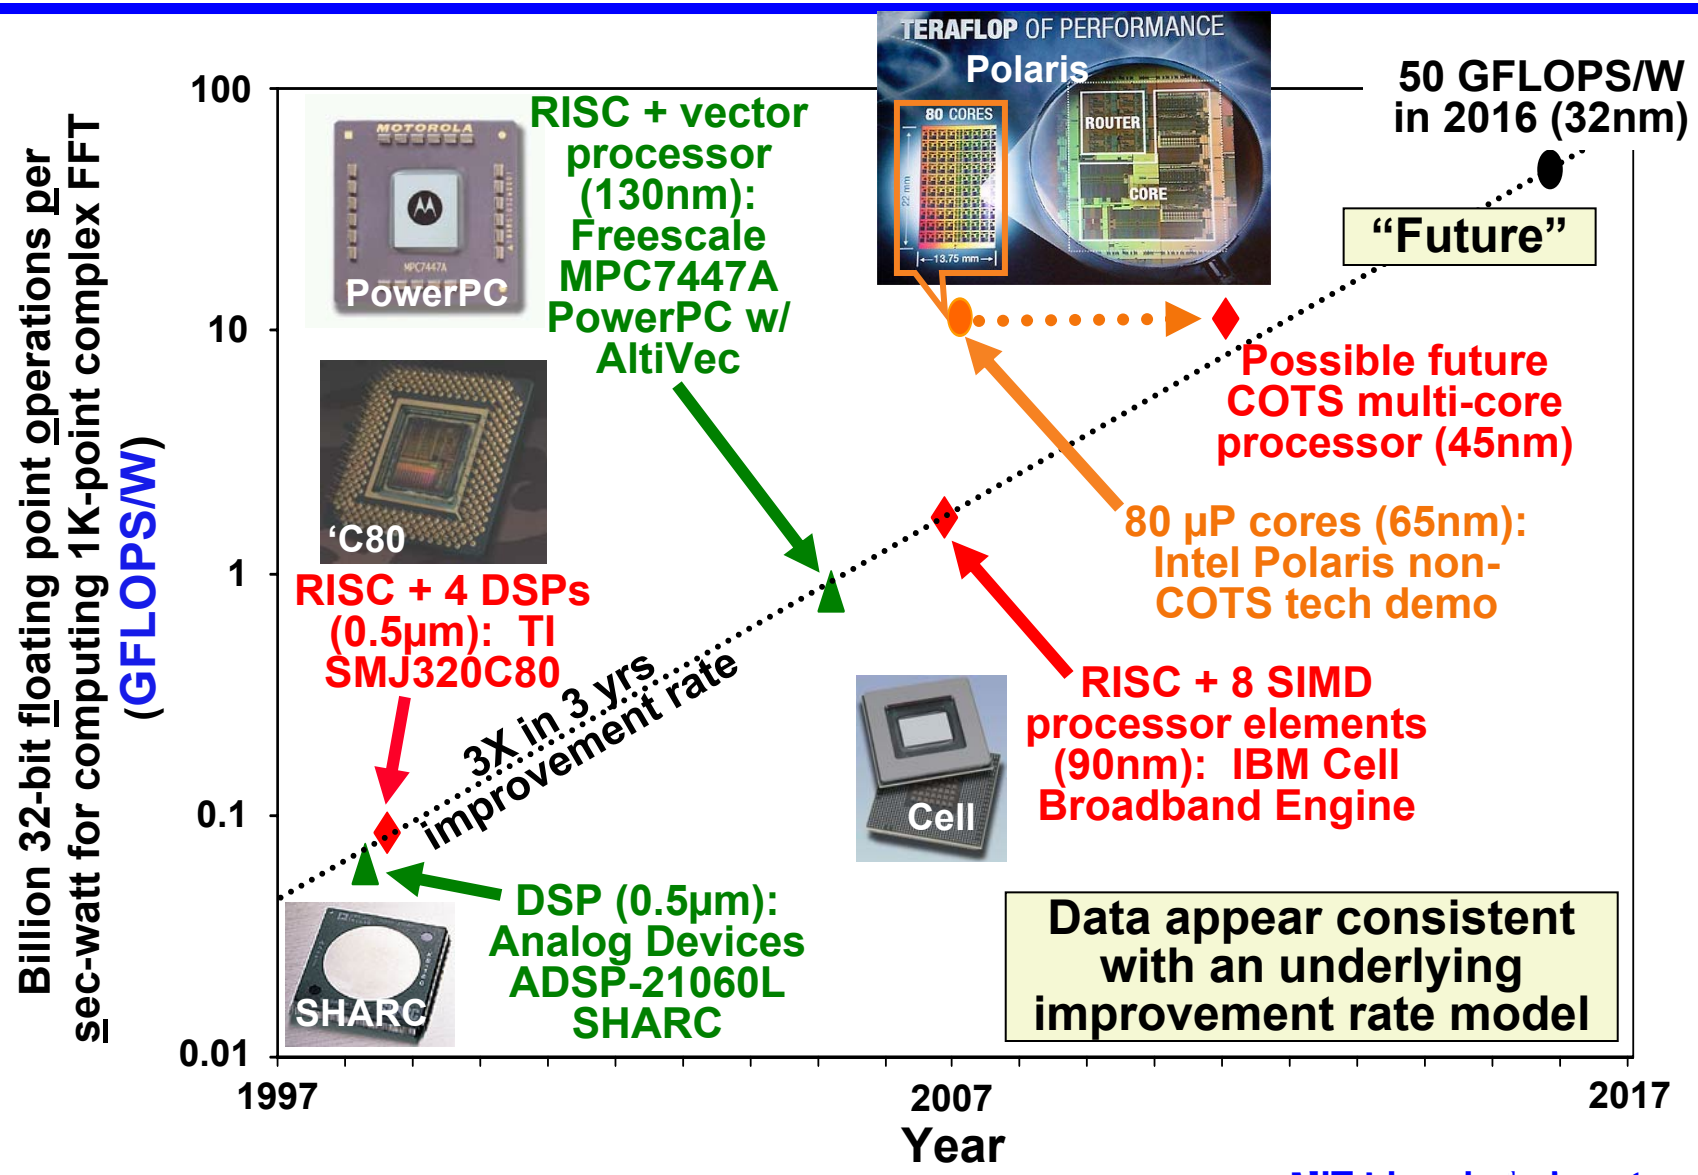

Single-core Microprocessor Family Improvement Rate Model

- Derived from ITRS06 projections (2005-2020)
  - $1/\sqrt{2}$ geometry reduction every 3 yrs
  - Normalized for constant die size & power

1). Present-generation uni-processor

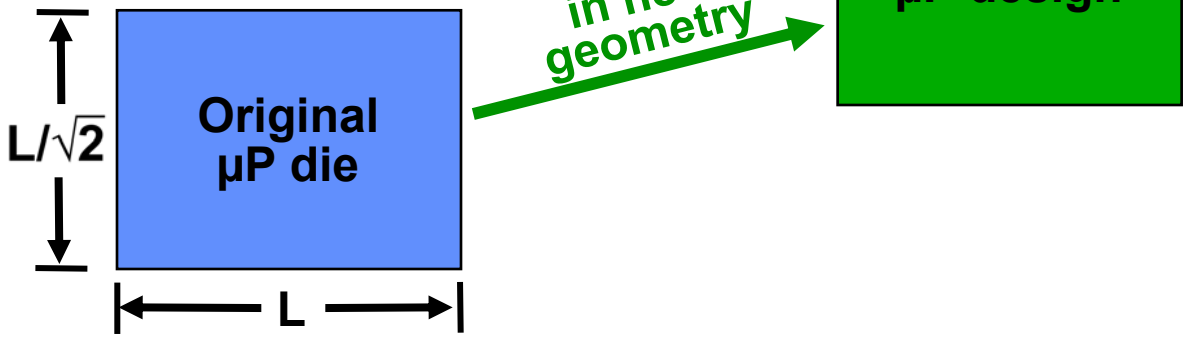

2a). Next-generation uni-processor

- 2X throughput (2X transistors @ similar speed, bytes/OPS unchanged)
- Compatible HW/SW

Single- & Multi-core Microprocessor Family Improvement Rate Model

- Derived from ITRS06 projections (2005-2020)
  - $1/\sqrt{2}$ geometry reduction every 3 yrs
  - Normalized for constant die size & power

1). Present-generation uni-processor

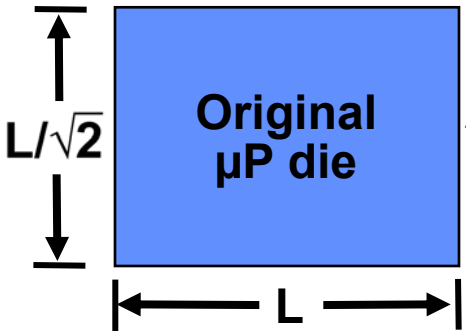

Re-design in new geometry

2a). Next-generation uni-processor

- 2X throughput (2X transistors @ similar speed, bytes/OPS unchanged)
- Compatible HW/SW

Re-implement in new geometry

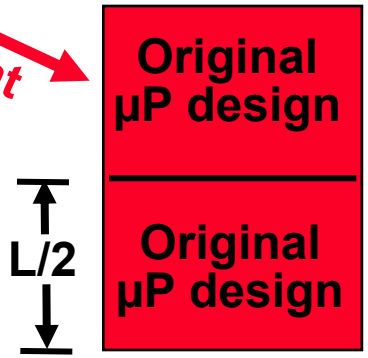

2b). Next-generation dual-core processor

- 3X throughput (2X transistors @ 1.5X speed, but 1/3 fewer bytes/OPS)
- Incompatible HW/SW (2X pins, not a uni-processor)

Note: 3X in 3 yrs improvement rate also projected for COTS ASICs (e.g., Cell) & SRAM-based FPGAs

"Multicore meltdown" occurs when core-based design methods can no longer provide substantial improvements

Notional Model Projection Based On IBM's *Cell Broadband Engine*

- 90nm geometry baseline ca. 2006
  - 1 mm pitch 1320 (41^2-19^2) BGA (ball grid array)
  - 234M transistors on 221 mm^2 die
    - L = 17.7 mm
    - W = L/$\sqrt{2}$ = 12.5 mm
  - 1333 pads (43x31) possible using 400µm pitch
  - 100W (est) @ 3.2 GHz (170 GFLOPS sustained for 1.7 GFLOPS/W)
  - 2592 Kbytes on-chip memory (66 KFLOPS/byte)

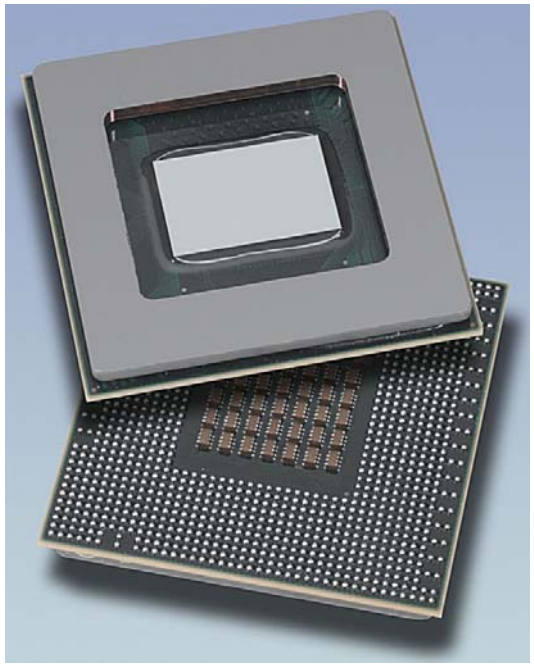

- 22nm geometry ca. 2018 (16X transistors & 5.1X speed @ constant power & die size)
  - 0.25 mm pitch 21,120 (164^2-76^2) FBGA (fine pitch BGA, with 0.15 mm feasible ca. 2010)
  - 21,328 pads (172x124) possible using 100µm pitch (typical ca. 2015)
  - 139 GFLOPS/W & 334 KFLOPS/byte

Core-based design methods appear *feasible* for the coming decade, *but how useful are the resulting devices given that memory & I/O shortcomings are compounded over time?*

Timeline for Highest Performance COTS Multiprocessor Card Technologies (3Q06)

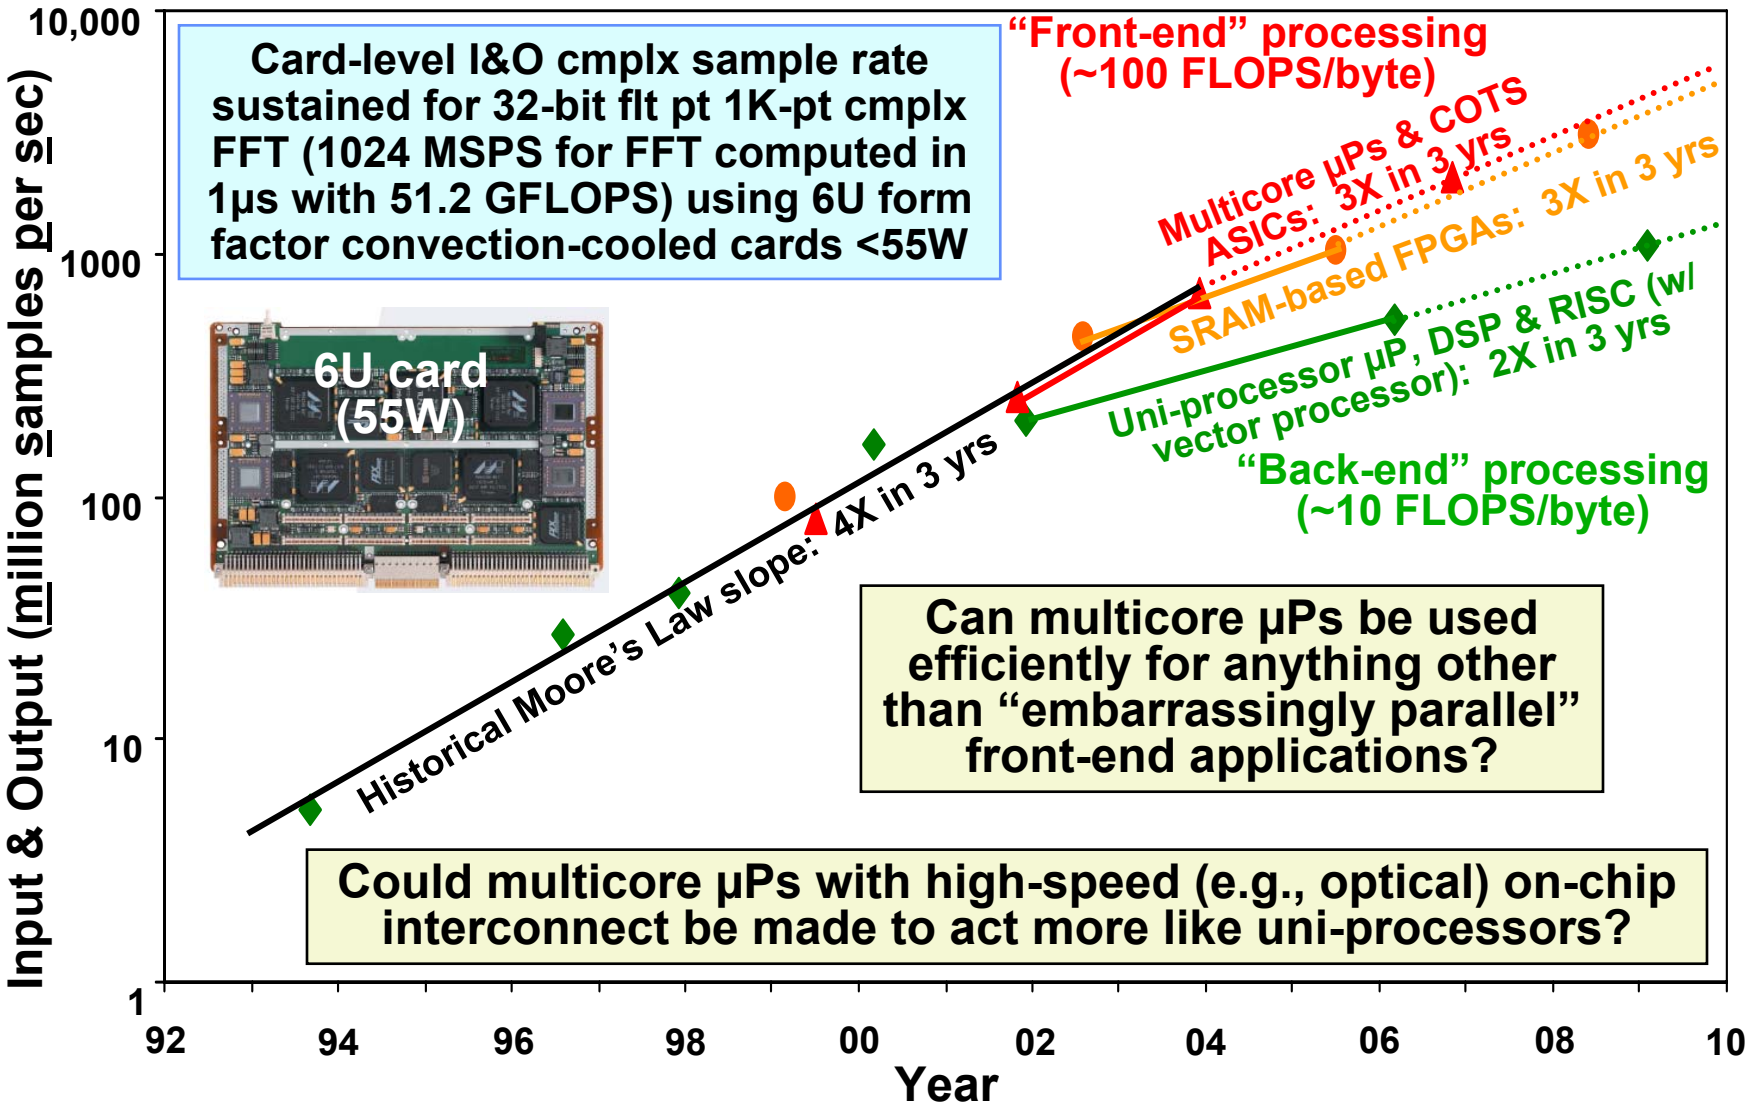

Timeline for Highest Performance COTS Multiprocessor Card Technologies (4Q06)

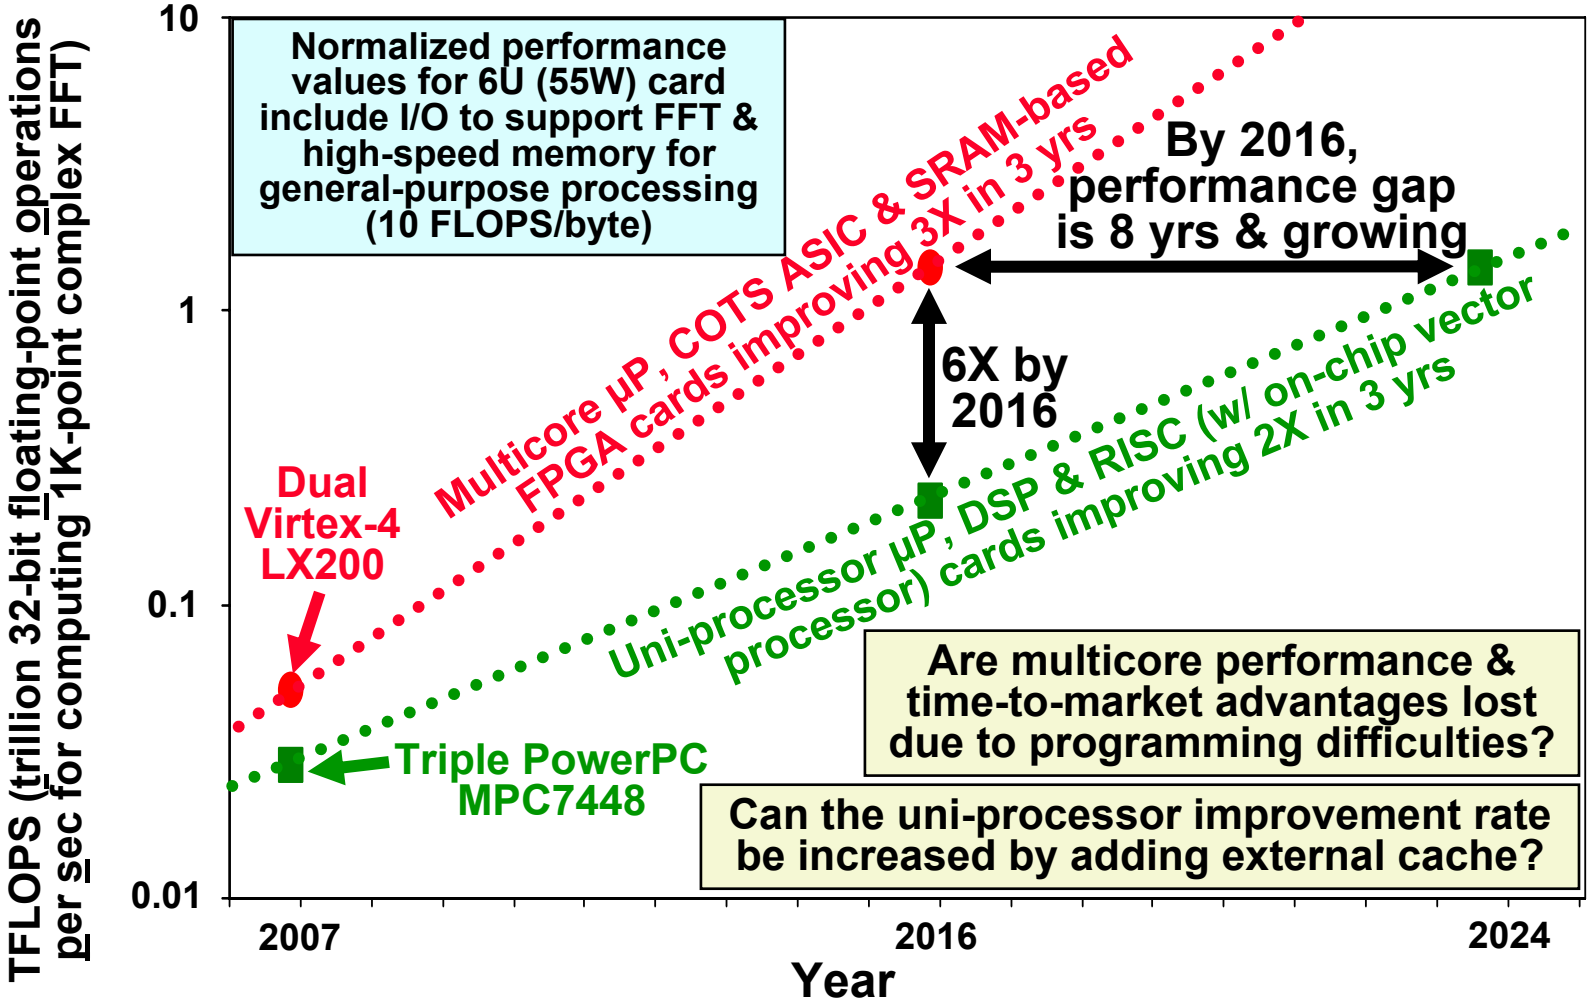

Improvements in COTS Embedded Multiprocessor Card Technologies (4Q06)

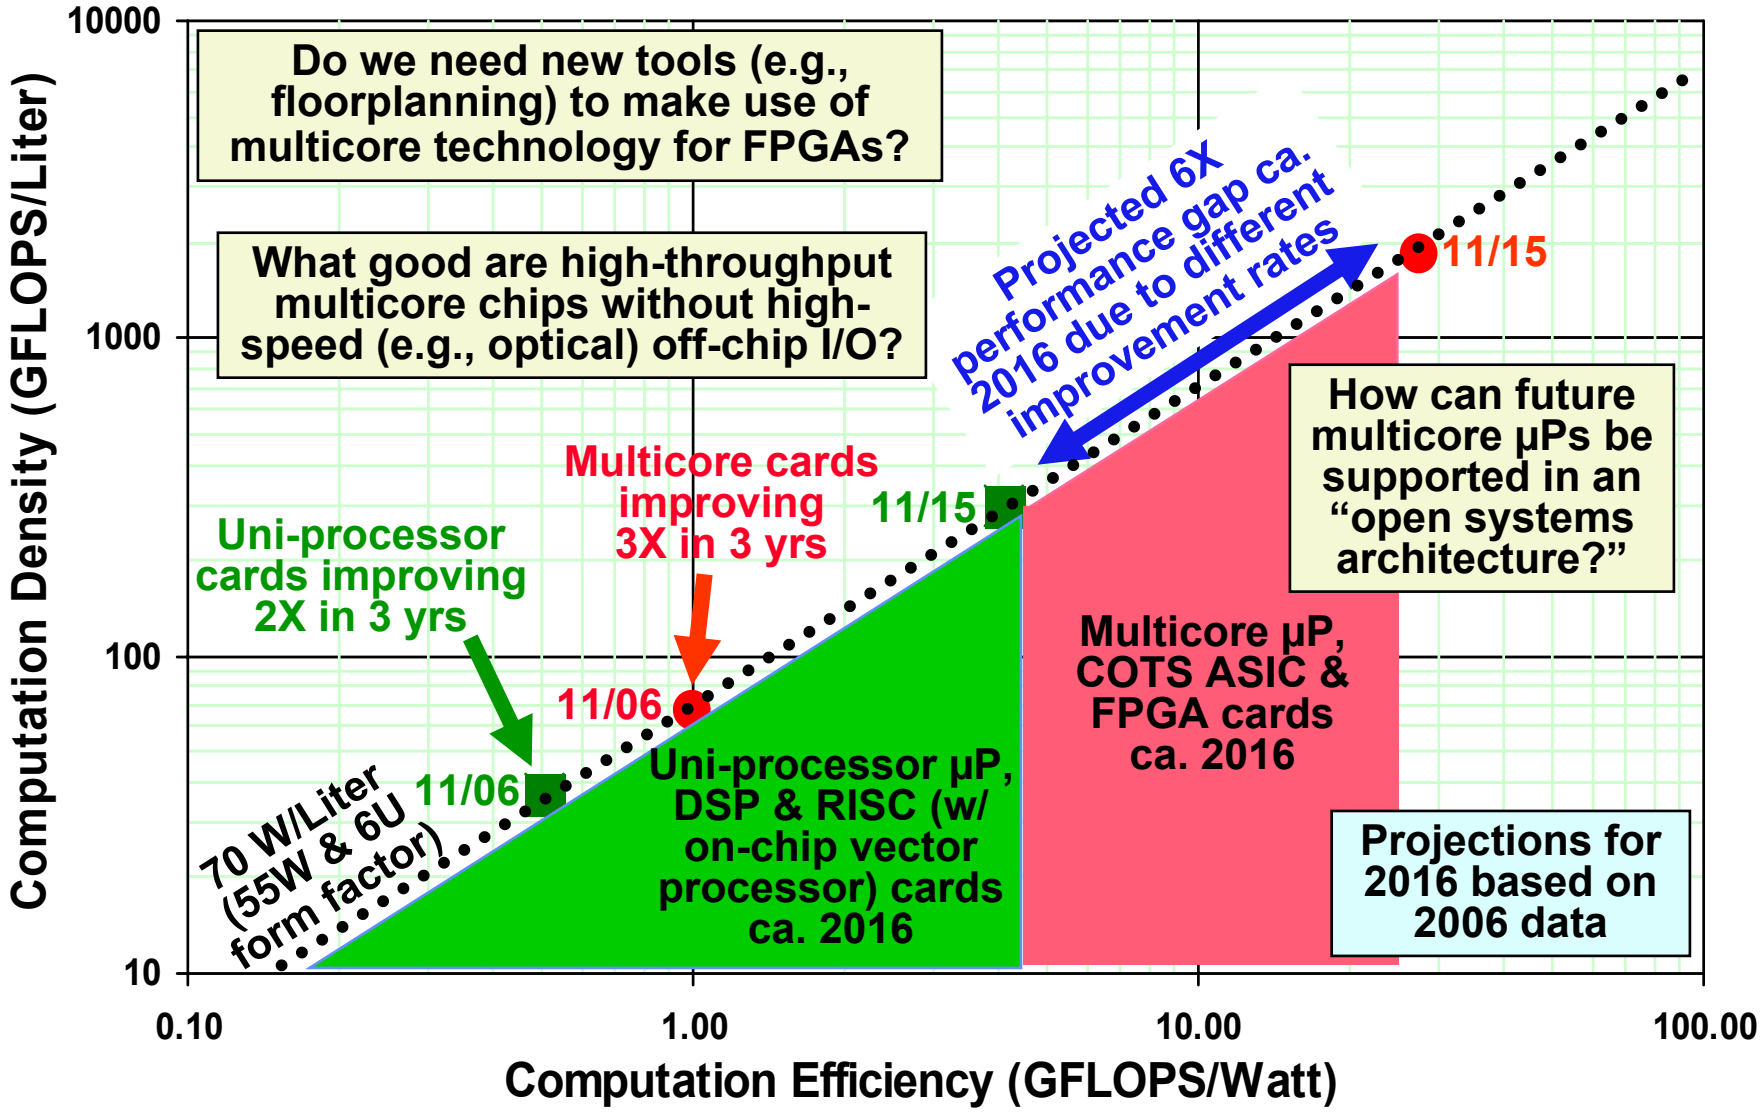

Objective & Schedule

- **Objective: assess the impact of multicore processors on past and present DoD HPEC systems, and project their importance for the future**

- **Schedule**
  - **1545-1605: Overview**
  - **1605-1615: Guest speaker Mr. Kalpesh Sheth**
  - **1615-1620: Introduction of the panelists**
  - **1620-1650: Previously-submitted questions for the panel**
  - **1650-1710: Open forum**
  - **1710-1715: Conclusions & the way ahead**

Panel Session:
"Multicore Meltdown?"

Moderator: Dr. James C. Anderson, MIT Lincoln Laboratory

Dr. James Held, Intel Corp.

Mr. Markus Levy, The Multicore Association & The Embedded Microprocessor Benchmark Consortium

Mr. Greg Rocco, Mercury Computer Systems

Mr. Kalpesh Sheth, Advanced Processing Group, DRS Technologies

Dr. Thomas VanCourt, Altera Corp.

Panel members & audience may hold diverse, evolving opinions

Objective & Schedule

- **Objective: assess the impact of multicore processors on past and present DoD HPEC systems, and project their importance for the future**

- **Schedule**
  - **1545-1605: Overview**
  - **1605-1615: Guest speaker Mr. Kalpesh Sheth**
  - **1615-1620: Introduction of the panelists**
  - **1620-1650: Previously-submitted questions for the panel**
  - **1650-1710: Open forum**
  - **1710-1715: Conclusions & the way ahead**

Conclusions & The Way Ahead

- Embedded processor hardware still improving exponentially
  - Throughput, in FLOPS (floating-point operations per second)
  - Memory, in bytes
  - Standard form factors with constant SWAP (size, weight and power) still required for upgradeability in DoD HPEC systems to support an "open systems architecture"
- At system level, ability to effectively utilize hardware improves slowly, leading to ever-worsening "performance gap"
  - Multicore processors, COTS ASICs (e.g., Cell) & FPGAs improving 3X every 3 yrs, but some performance advantages may be lost due to programming difficulties (relatively small memory) & I/O bottlenecks
  - Traditional uni-processors easier to program, but improving at slower rate (2X every 3 yrs)
- Several methods may help "narrow the gap"
  - Improved processor-to-processor & processor-to-memory communication technology (e.g., optical on-chip & off-chip) may allow multi-processors to behave more like uni-processors
  - Control logic & 3D interconnect to attach more high-speed memory for a given throughput may improve multicore processor utility (and/or allow uni-processors to improve at a faster rate)

"Multicore meltdown" avoidable for the coming decade, but taking full advantage of performance benefits will be challenging

Backup Slides

Intel Benchmark Performance

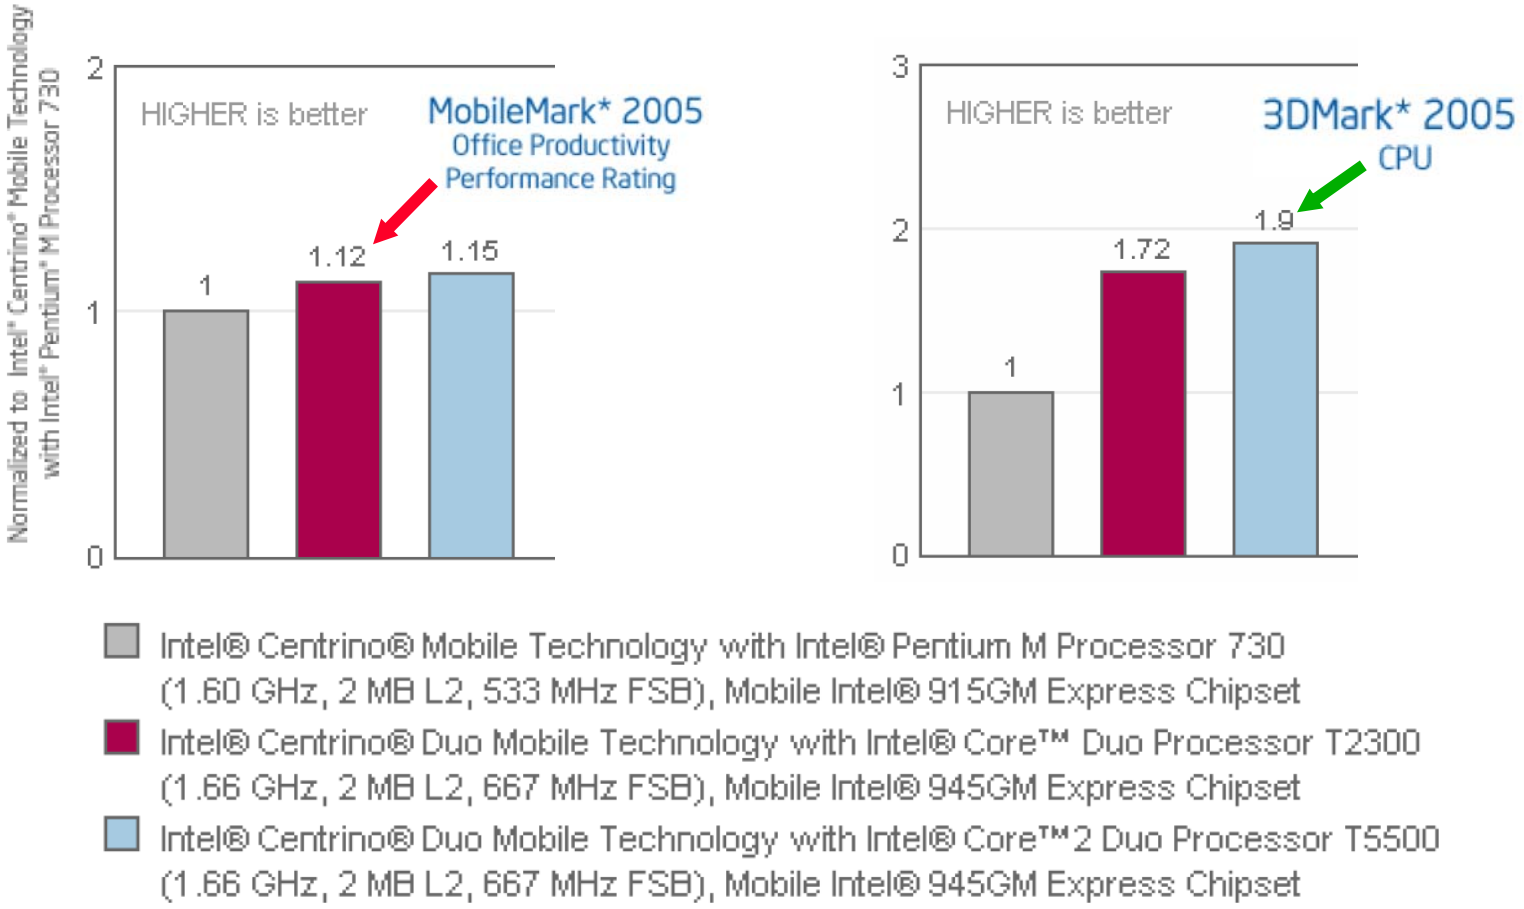

Depending on task, 2 CPU cores perform 1.12 - 1.9X as fast as one

6U Cards Feasible (but not built) using 90nm COTS Devices (4Q06)

- µP: Freescale MPC7448 PowerPC
  - 10W @ 1.4 GHz
  - 9.4 GFLOPS sustained for 32-bit flt pt 1K cmplx FFT (83.6% of peak)
    - 2.3W budget for external 1.5 Gbytes/sec simultaneous I&O
    - 2W budget for 1 Gbyte external DDR2 SDRAM (~10 FLOPS/byte)
    - 2.5W budget for misc. logic (host bridge, clock, boot flash)
    - 1.6W budget for on-board DC-to-DC converter (91% efficient)
  - Triple CN (compute node) card
    - Easier component placement vs. quad CN
    - 28.2 GFLOPS sustained for 55W
    - 0.51 GFLOPS/W & 37 GFLOPS/L

- FPGA: Xilinx Virtex-4 LX200
  - 41W (est.) for a legacy card re-implemented using 90nm devices
    - Card includes DC-to-DC converter
    - 2 FPGAs w/ 225 MHz core & 375 MHz I/O speeds
    - 528 Mbytes DDR-SRAM
    - 8 Gbytes/sec I&O
    - 51.2 GFLOPS sustained for 32-bit flt pt 1K cmplx FFT
  - 8W budget for additional 4 Gbytes DDR2 SDRAM (~10 FLOPS/byte)
  - Dual CN card
    - 51.2 GFLOPS sustained for 50W (1W added for increased converter dissipation)
    - 1.0 GFLOPS/W & 67 GFLOPS/L