IBM Deep Computing Team

CAE Performance Trends in 2007

Jeff Zais
zais@us.ibm.com

Challenge

+

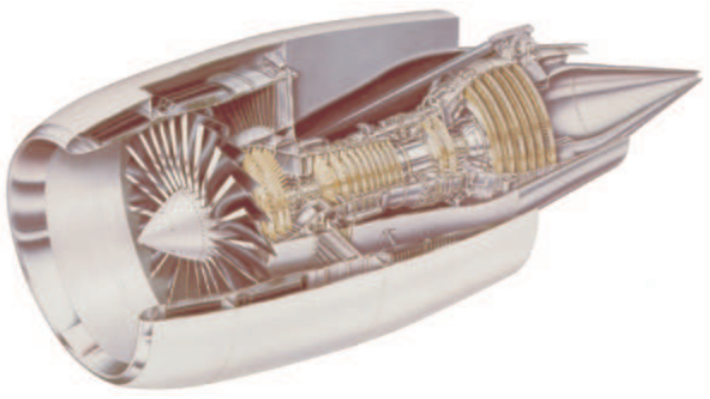

2002: System z + Linux = 97,000

+

=

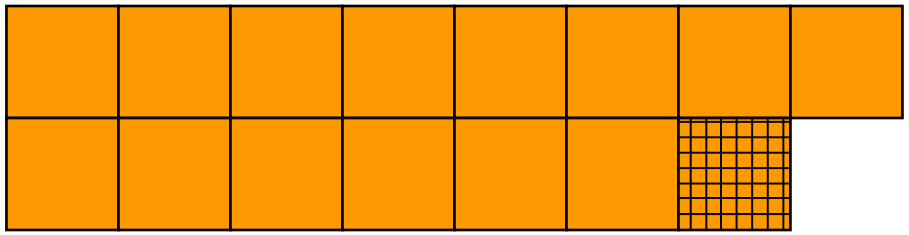

2007: System z + Cell Broadband Engine

+

= **the “GameFrame”**

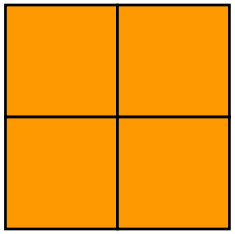

≠

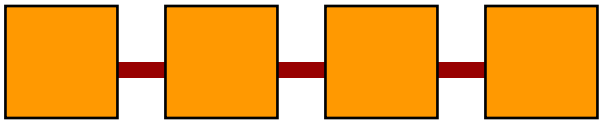

Leading Performance Issues – 2007

1. **“The System x world keeps changing rapidly”**
2. **“We have hit the wall in terms of performance”**
3. **“We are moving to a multicore world”**
4. **“Power consumption is becoming important”**
5. **“Price/Performance is the key buying factor”**

1 - “The System x world keeps changing rapidly”

- **Processor design**
  - single-core
  - dual-core
  - quad-core
- **Performance leadership**
  - Pentium 4 … Opteron … Xeon 5160
- **Interconnect**
  - Myrinet … InfiniBand … SDR … DDR
  - LAM/MPI … Scali … Open MPI

2 - "We have hit the wall in terms of performance"

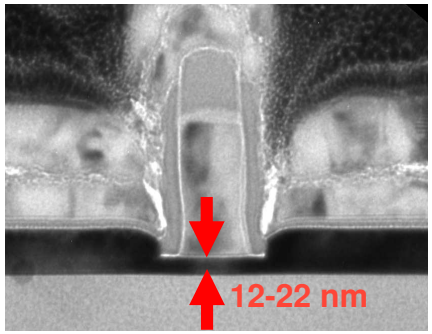

- **As feature size shrinks, passive power grows**
- **Now more important than active power**
- **No longer can we rely on steady clock speed increases to gain more performance**

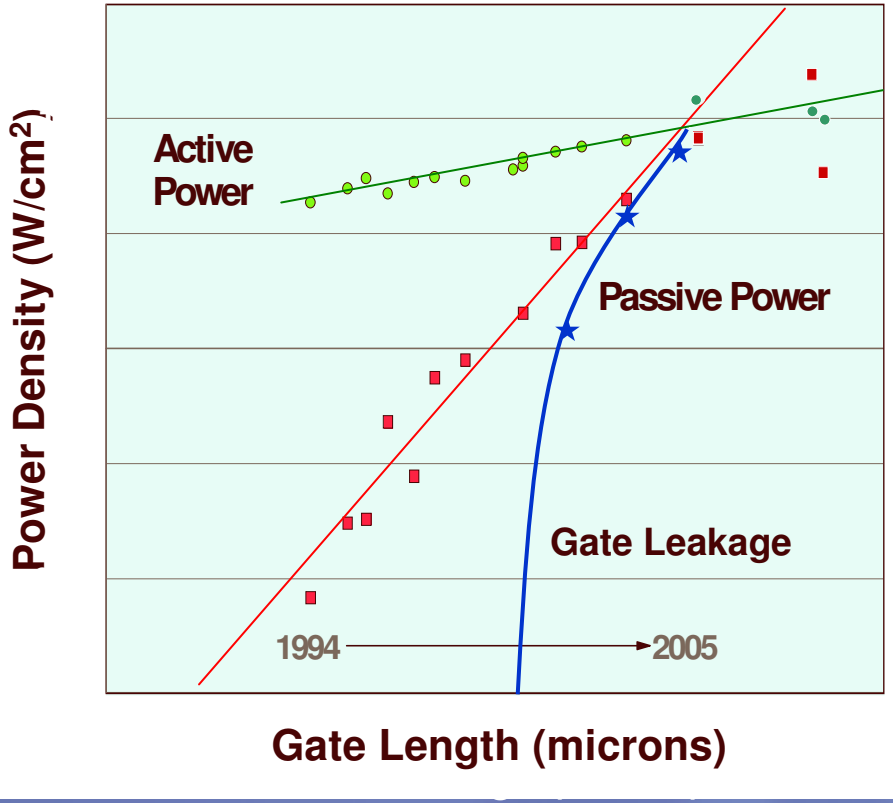

Moving to an Ultra-Low Voltage Operating Point

10X power reduction for only 3X performance penalty

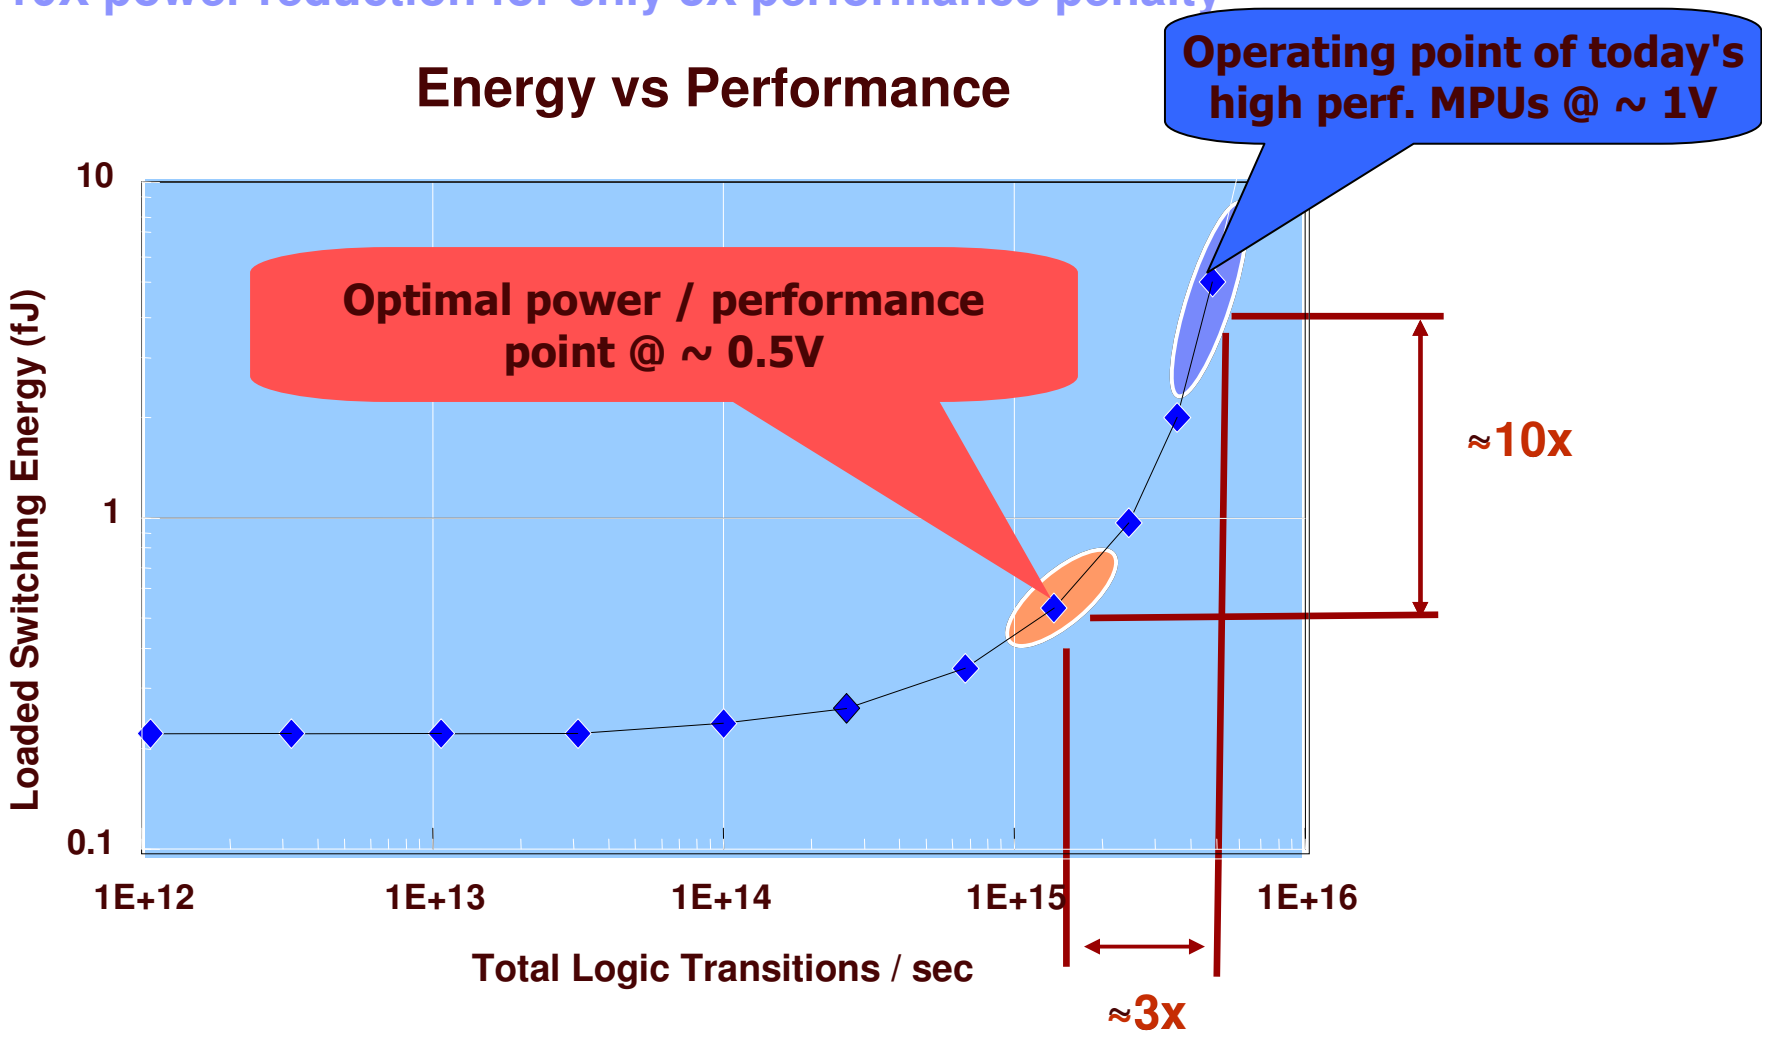

3 - “We are moving to a multicore world”

- **IBM**
  - Cell Broadband Engine, Sony – Toshiba – IBM joint project
  - 1 PowerPC, plus 8 “synergistic processing units”
- **AMD**
  - “Fusion” project – goal of 1 teraFLOP/socket in 2010
  - 48 GPU pipes x 8 FLOPS/cycle x 3 GHz = 1 teraFLOP
- **Intel**
  - project “Larrabee”, due out in 2009
  - 80-core chip reaching 2 teraFLOPS at 191 Watts, or 1 teraFLOP at 46 Watts

4 - “Power consumption is becoming important”

- **Old measures of performance**
  - FLOPS
  - elapsed time to solution
- **New measures of performance**
  - FLOPS per kiloWatt
  - cooling required per rack

Power & cooling spending will exceed new server spending
(Gartner 2006)

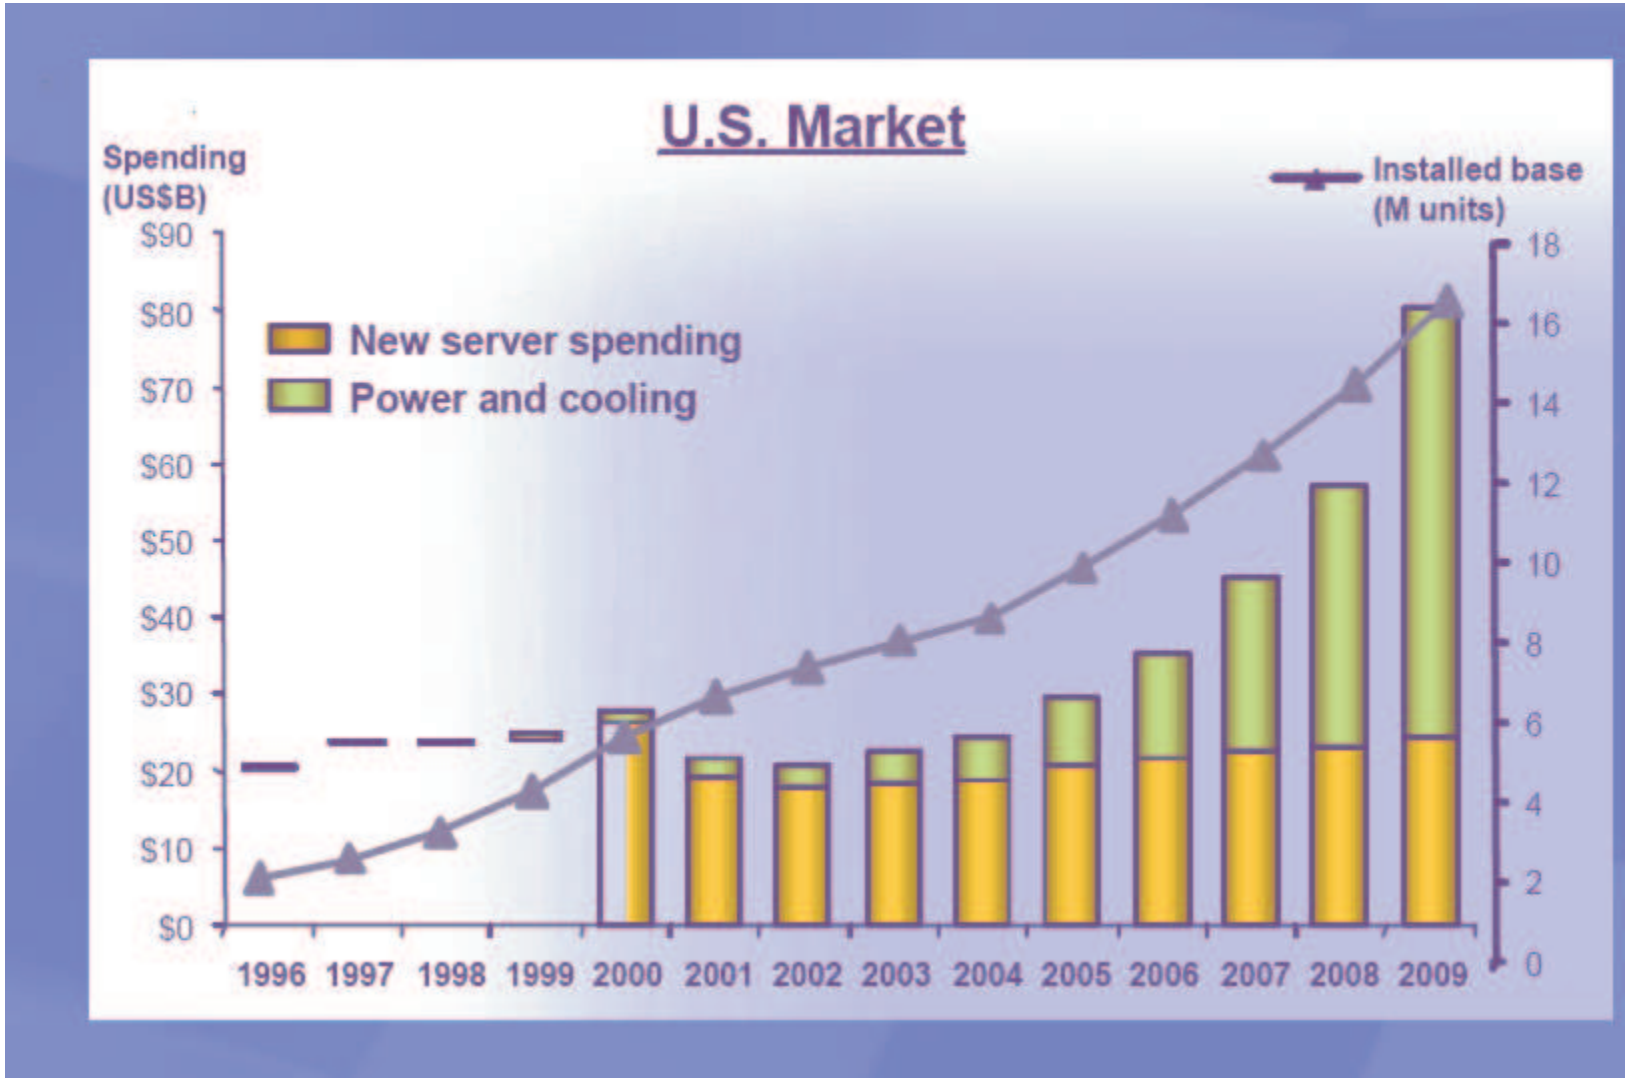

5 - “Price/Performance is the key buying factor”

$$\text{Decision} = \frac{\text{Price}}{\text{Performance}}$$


PIER 17

1 - “The System x world keeps changing rapidly”

2001 Pentium 4

2002 rack-optimized

2003 Myrinet

2004 Opteron

2005 Opteron dual-core

2006 Xeon 5160, quad-core “Clovertown”

2007

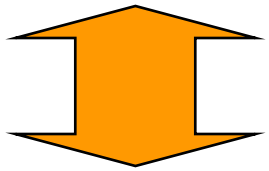

2 - “We have hit the wall in terms of performance”

11 April, 2007: “though-silicon vias”

vertical connections etched through the silicon wafer and filled with metal, allowing multiple chips to be stacked

shortens distance information needs to travel by 1000 times, and allows for up to 100 times more channels

“result of more than a decade of pioneering research at IBM” – Lisa Su, VP of Semiconductor Research

2008 – production status. Starting with wireless communications chips.

more uniform power distribution, putting power closer to cores, so clock speed increases and power consumption reduced up to 20%

Blue Gene chip getting converted to 3-D stacked chip. Also, 3-D technology will fundamentally change the way memory communicates with a microprocessor.

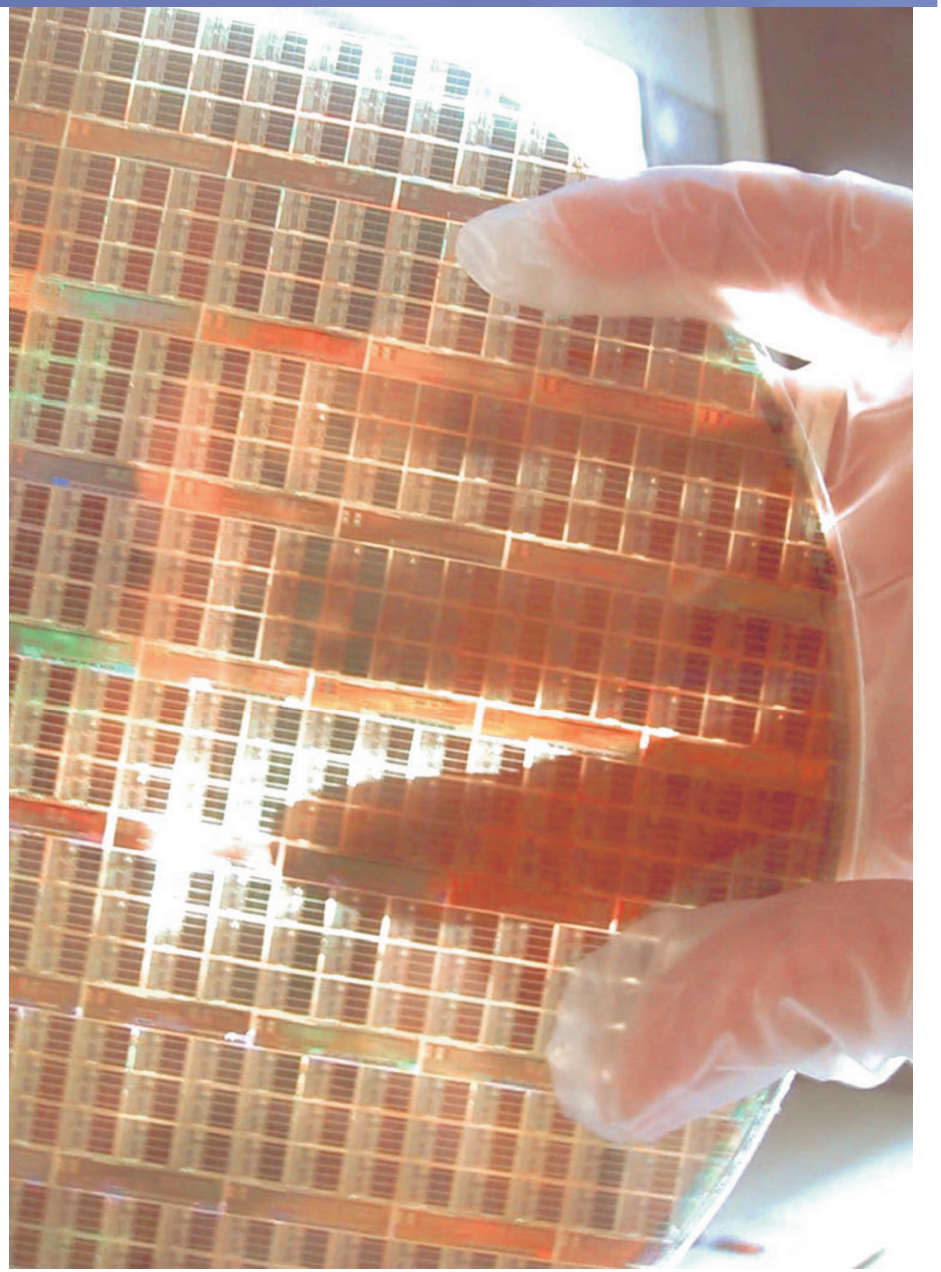

27 January, 2007 innovation from IBM: High-k Metal Gate

New materials can significantly reduce leakage currents

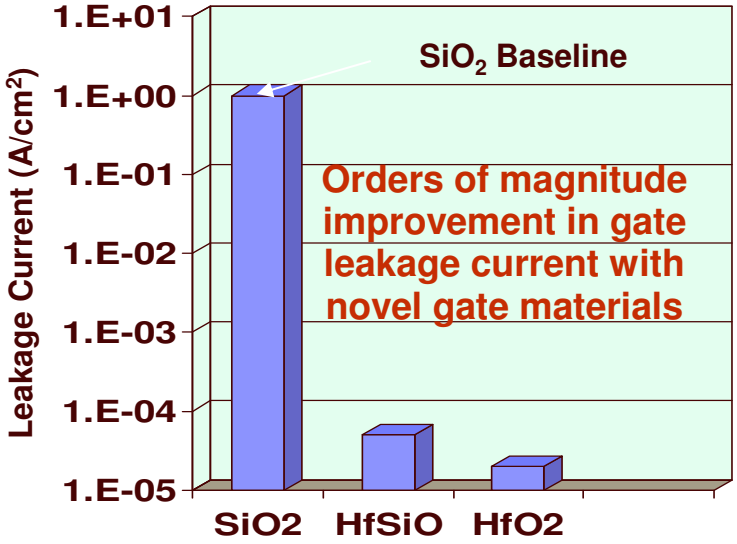

3 May, 2007 – “Using self assembly to create airgap processors

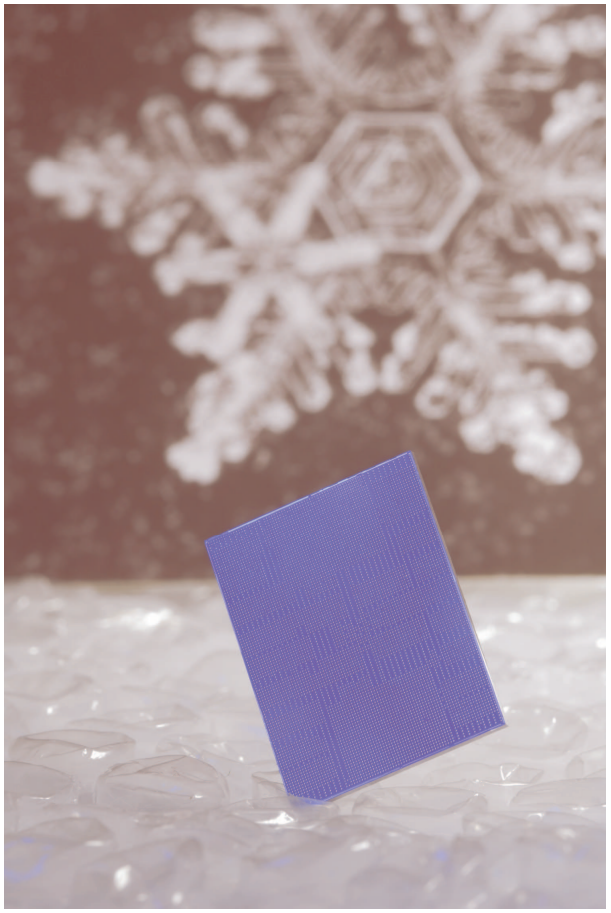

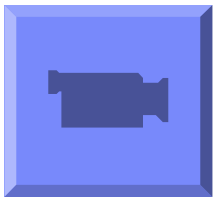

Options (choose 1)

- electrical signals can flow 35% faster on the chip
- chips can consume 15% less energy

10 Years | 10 Breakthroughs

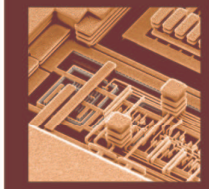

Copper

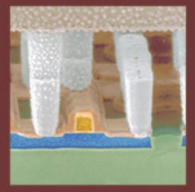

SOI

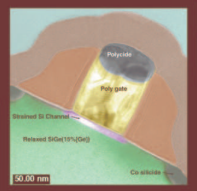

Strained Silicon

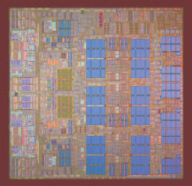

Dual Core

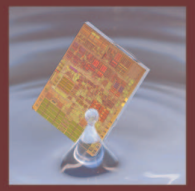

Immersion

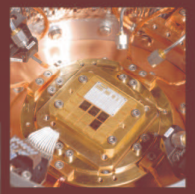

Frozen SiGe Chip

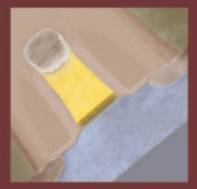

High-k

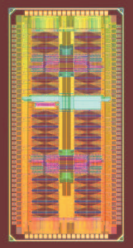

eDRAM

3D Chip Stacking

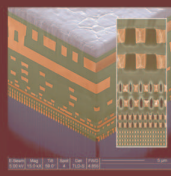

Airgap

IBM

2007

The SIA CMOS Roadmap

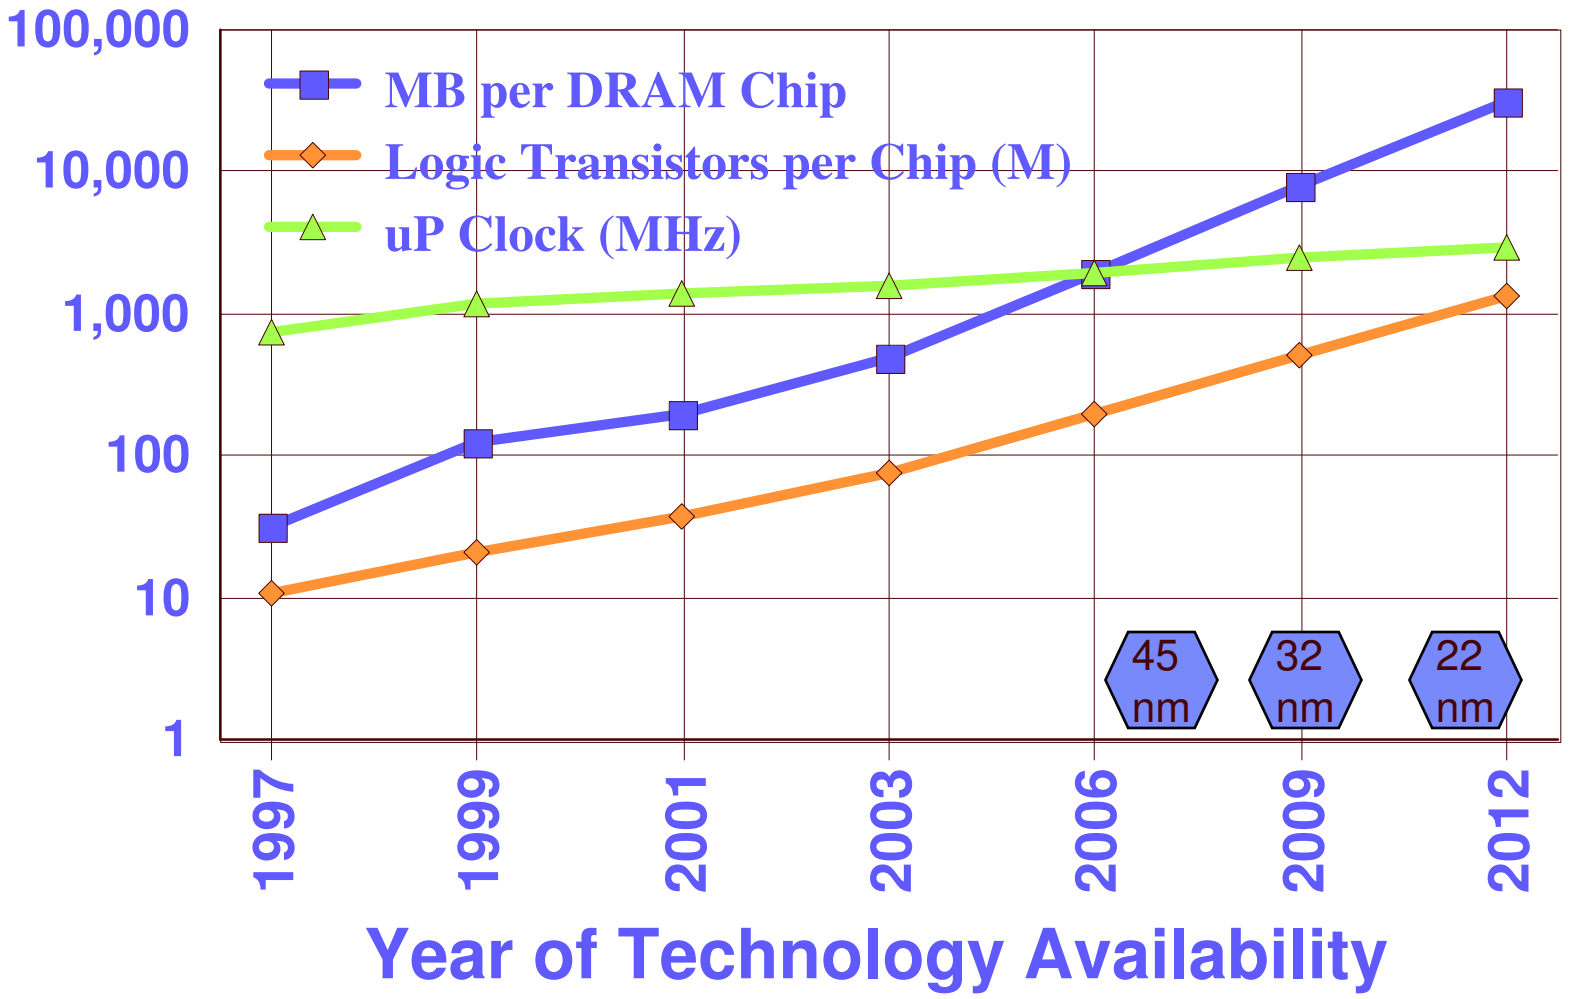

2 - “We have hit the wall in terms of performance”

- **POWER6 should continue IBM tradition of improved performance for CAE codes**

- **6 February, 2006 – International Solid-State Circuits Conference (ISSC) in San Francisco**
  - 65-nm processor
  - clock speed above 4 GHz

3 - “We are moving to a multicore world”

- **In CAE, we moved to multicore in the 1990’s**

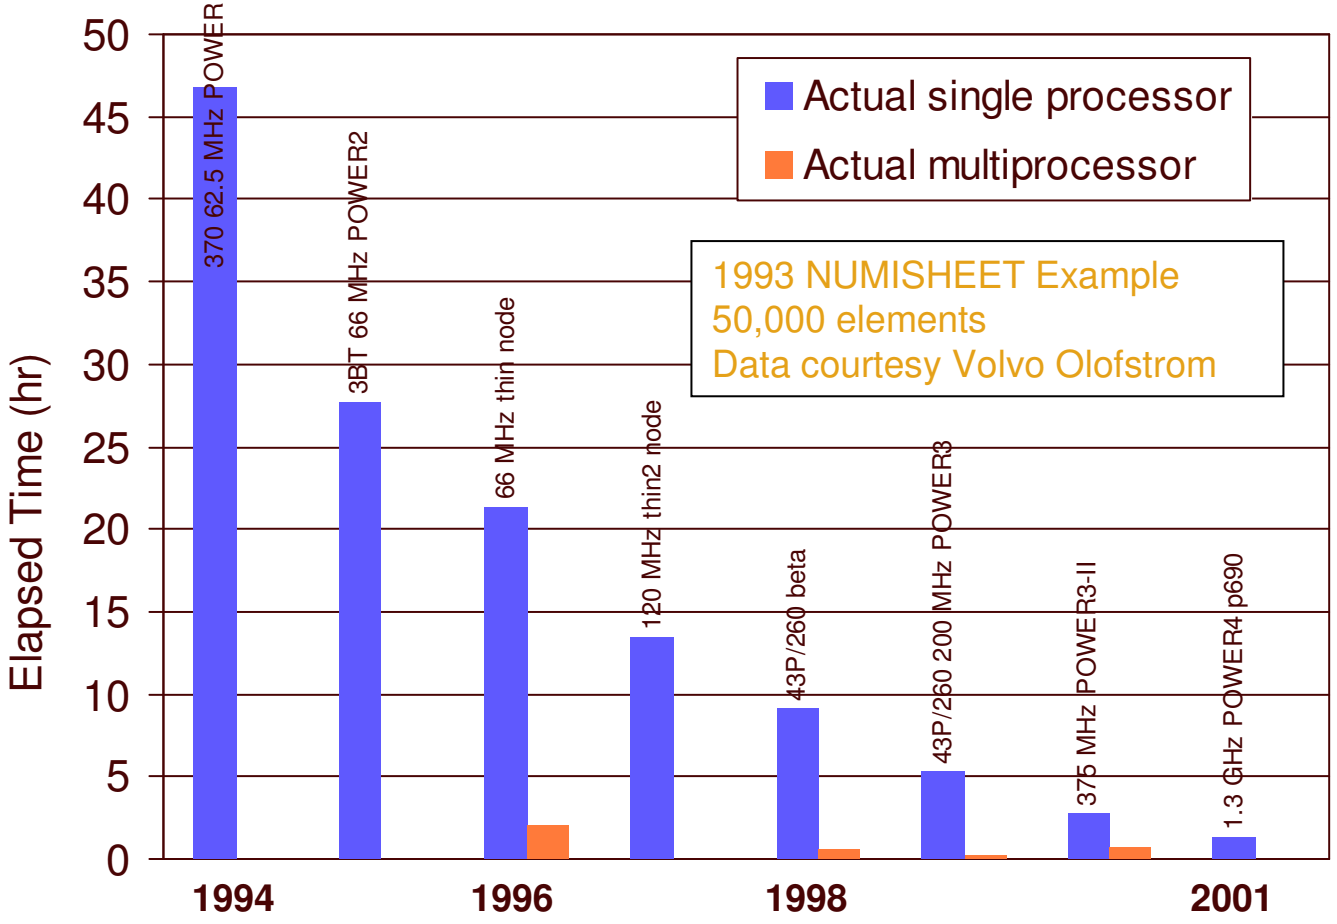

- **Can we catch the next train which leaves the station?**

Can we catch the next train when it leaves the station?

Serial

SMP

MPI
POWER

MPI
x86 Linux

MPI
Blue Gene

Cell
Broadband
Engine

4 - “Power consumption is becoming important”

- **IBM has seen this coming.**

  We’ve been responding on many fronts

BladeCenter Basics:

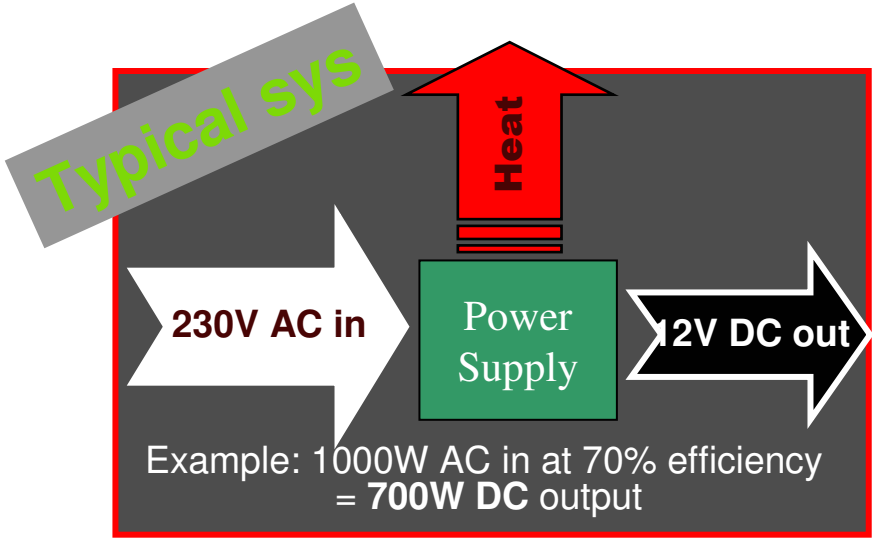

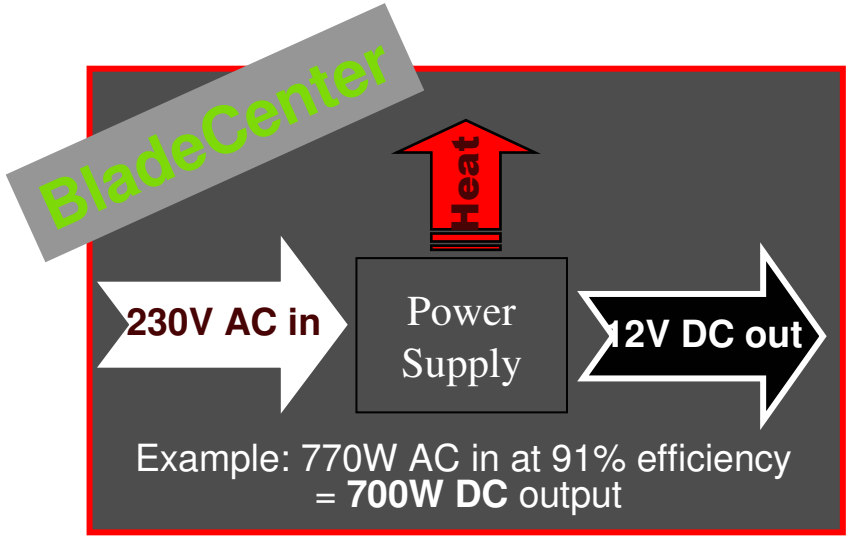

BladeCenter Offerings…

Less Power

Less Cooling

Less Cost

BladeCenter
Announced: Nov. 2002

14 Blades, 7U

BladeCenter T
Announced: Apr. 2004

8 Blades, 8U

Ruggedized Chassis

BladeCenter H
Announced: Feb. 2006

14 Blades, 9U

InfiniBand for data intensive HPC apps

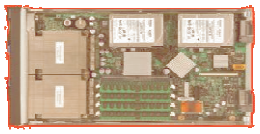

Common Blades, Common Switches

Blade Architecture & OS Flexibility

All blades refreshed to 2nd Generation in 2006

HS2X Blades
Introduced 2002
2-socket
Intel Xeon DP
Up to 14 Per Chassis

HS4X Blade
Introduced 2003
4-socket
Intel Xeon MP
Up to 7 Per Chassis

JS2X Blade
Introduced 2004
2-socket
IBM PowerPC
Up to 14 Per Chassis

QS20 Blade
Introduced 2006
18 cores per blade
Up to 7 Per Chassis

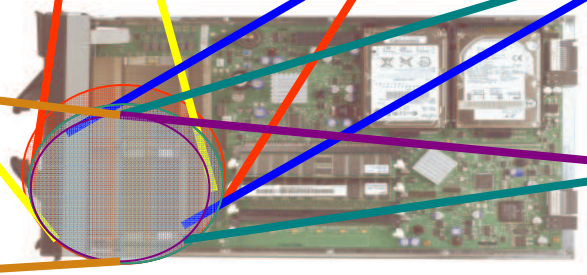

LS2X Blade
Introduced 2005
2-socket
AMD Opteron
Up to 14 Per Chassis

LS4X Blade
Introduced 2006
4-socket
AMD Opteron
Up to 7 Per Chassis

IBM PowerExecutive

PAM-CRASH job running on IBM HS21 XM

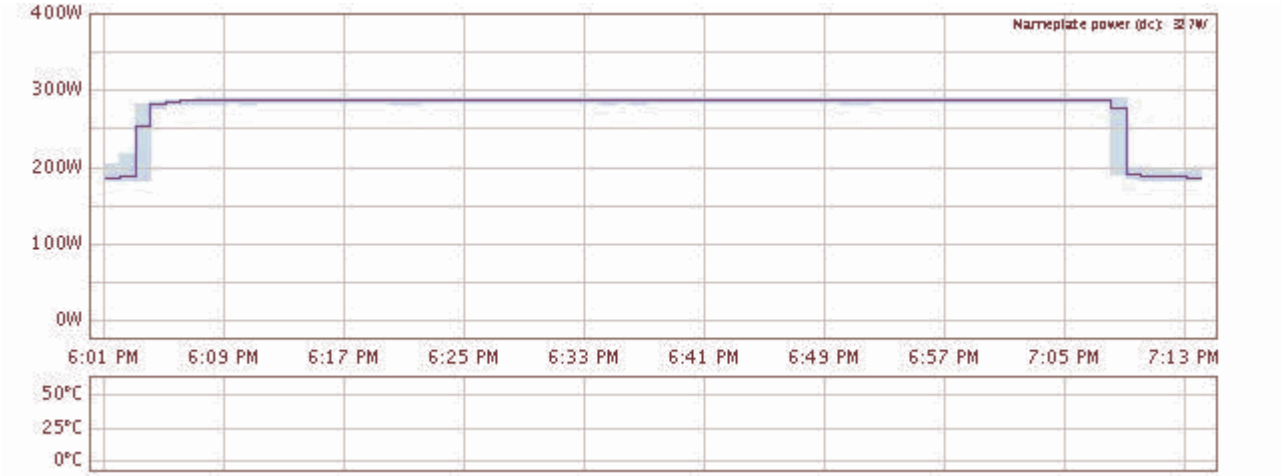

- **tracks actual power usage**
- **monitor installed systems**
- **provides better data for planning future installations**

5 - “Price/Performance is the key buying factor”

$$\text{Decision} = \frac{\text{Price}}{\text{Performance}}$$

- **Both price and performance depend on configuration**
  - single-core or dual-core best?
    - performance of various applications
    - application license costs
    - memory configuration
  - SAS (serial attached SCSI) or SATA (serial ATA)?
  - Interconnect
    - SDR or DDR?

LS-DYNA comparison: dual-core vs. quad-core

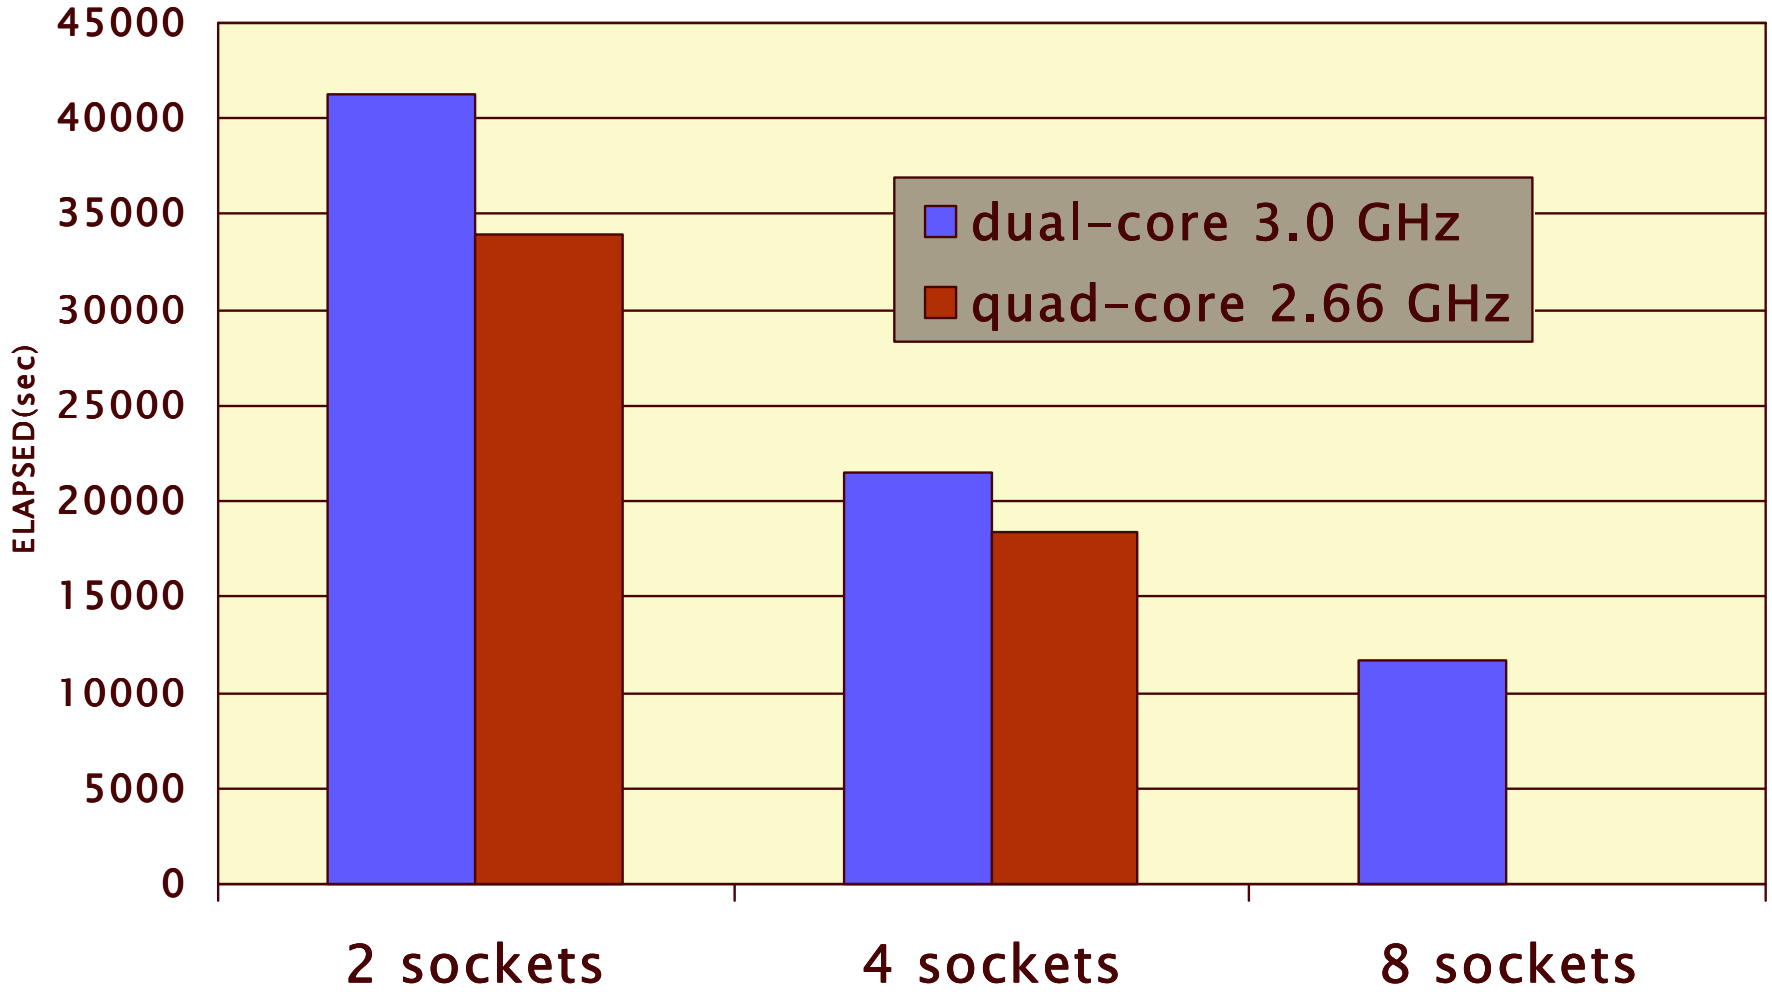

Pitfalls of focusing only on price / performance

- **Examples: difficult-to-quantify costs**
  - application licenses vs. number of cores
  - lifetime costs, including cooling and power consumption
  - …
- **Examples: difficult-to-quantify benefits**
  - applications support
  - stateless computing
    - simplified administration
    - fast booting
  - …

Summary

- **Our IBM team is eager to help customers solve problems! We are working on solutions in many areas**
  - performance of systems based on commodity processors
  - performance of systems based on IBM POWER
  - performance of innovative architectures – Cell and Blue Gene
  - power and cooling
  - finding the best solution in terms of price and performance


Thank you!
Questions / Comments
are welcome