Tera Tera Tera

Ed Davis – Chief Architect

AGENDA

- Supercomputing Roadmap
- Quad-core, Eight-core, 10-100 core
- Compute / Networking / Storage / Software
- Tera-Scale Computing
- Software Solutions & Challenges
- Ongoing Research

intel Leap ahead

Pat Gelsinger laying out the future roadmap hurdles

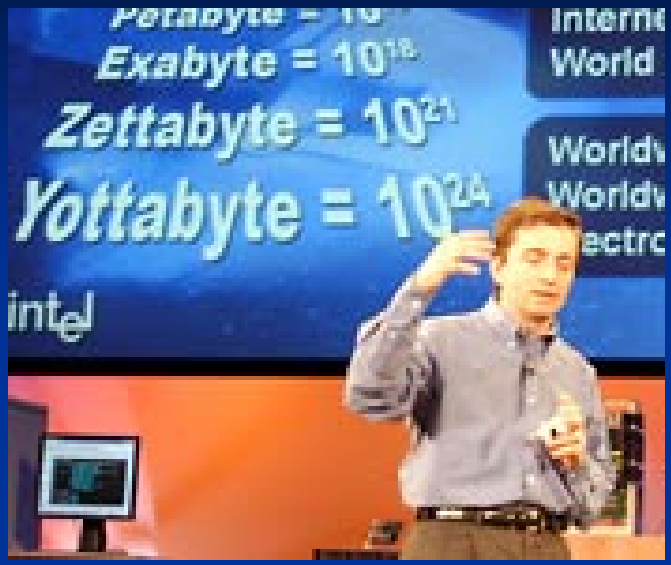

- Xonabyte = 10^{27}
- Wekabyte = 10^{30}
- Vundabyte = 10^{33}
- Udabyte = 10^{36}
- Tredabyte = 10^{39}

Currently Google uses 2 PetaBytes of disk space

intel Leap ahead

Supercomputing Roadmap (PetaFLOPs +)

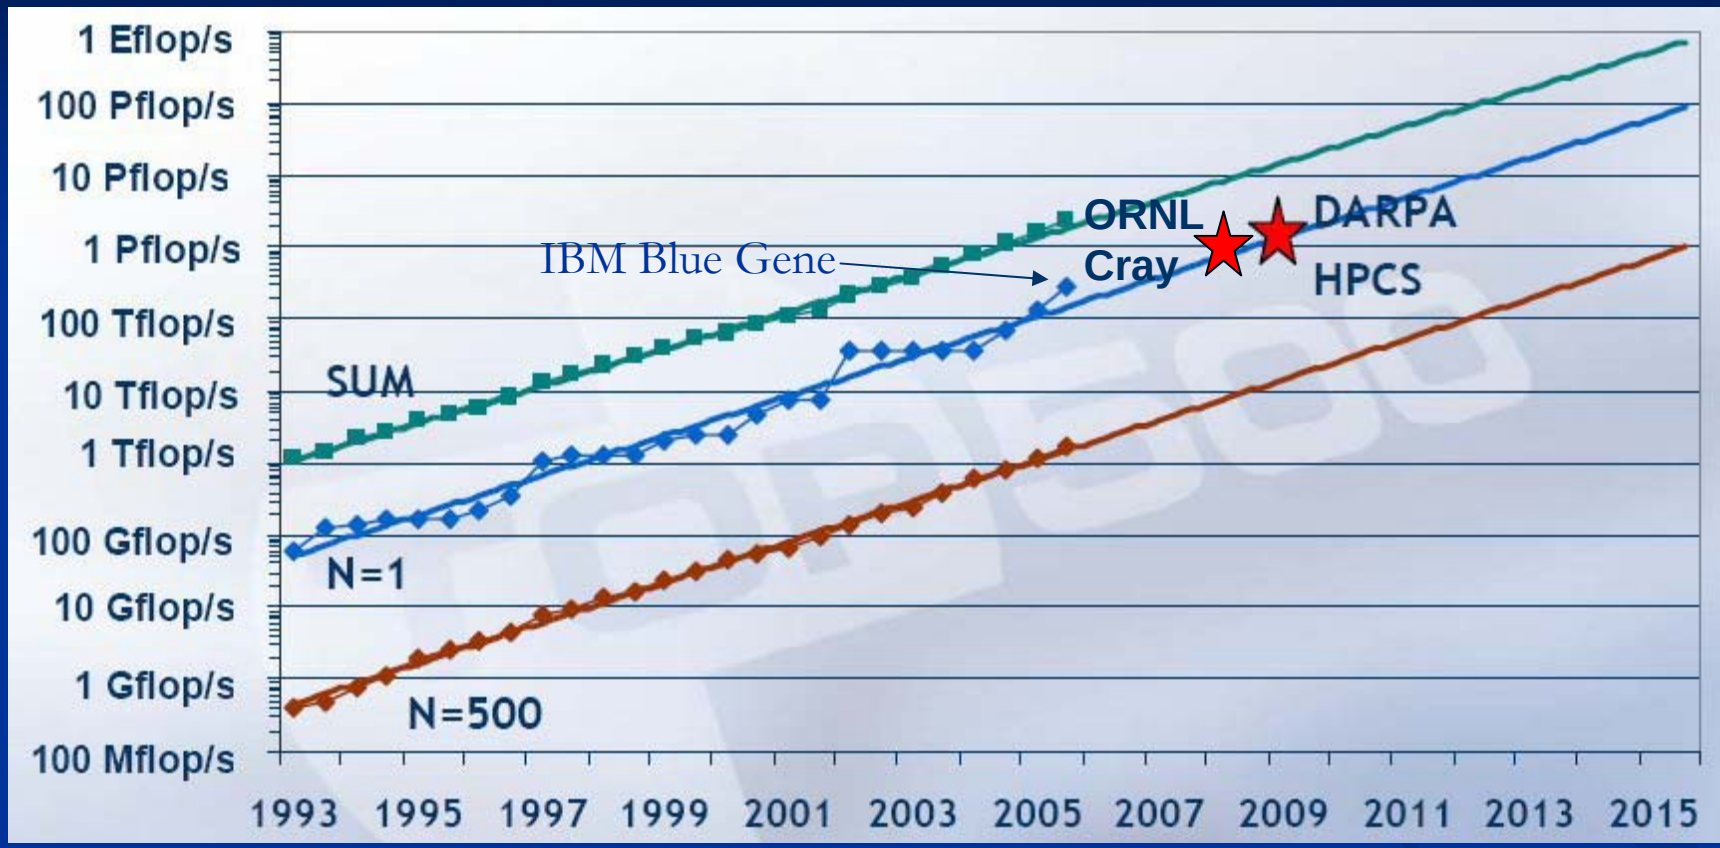

We've already achieved a PetaFLOP in aggregate now the race is on for World's first PetaFLOP system

intel Leap ahead

Quadcore Launch at CERN

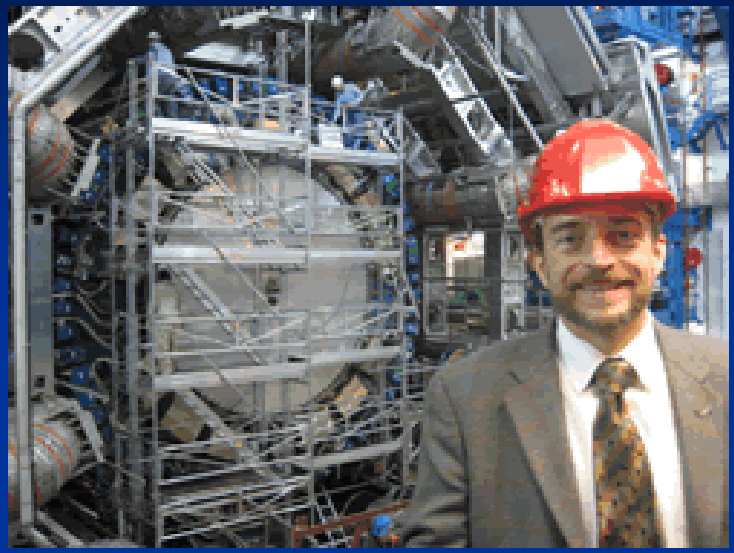

- Pat Gelsinger Sr. V.P. and GM Digital Enterprise Group Launches Quad-core “Clovertown” at CERN

intel Leap ahead

QUAD-CORE: 65nm

Gigabyte GA-965P-S3

1066/1333 MHz

Clovertown (Server)

Kentsfield (Desktop)

+ **Faster TTM (Out ship AMD 20:1 in the first 3 Qtrs)**
+ **Simple Coherency Model**
+ **Flexible Die Selection**
+ **Utilizes Capacity**

Up to 70% better perf. over Core2 Duo*

Intel's Quad-Core solution is a careful choice of performance, schedule, and cost.

*SPECint_rate

intel Leap ahead

WHAT THE INDUSTRY IS SAYING

"Intel Enhances Its Best Processors"

- Computing News, Nov 2006

"If you want quad-core today, Intel is the only game in town"

- Nathan Brookwood quoted in the Mercury News, Nov 14, 2006

"A total of eight cores offering unmatched server performance"

- Koen Crijns, Hardware.Info, Nov 13, 2006

"With no price premium, Intel has laid the gauntlet down for AMD."

- Charlie Demerjian, Inquirer Nov 14, 2006

"Woodcrest offers better performance than [Opteron] "Rev F", but Clovertown takes performance to a new level"

- Thomas Weisel Partners, EQUITY RESEARCH, Nov 14, 2006

"These new Xeon chips fit into the same power envelope as their dual-core predecessors, but pack twice the processor cores"

- *Geoff Gasior, The Tech Report, Nov 13, 2006*

intel Leap ahead

GOING MAINSTREAM

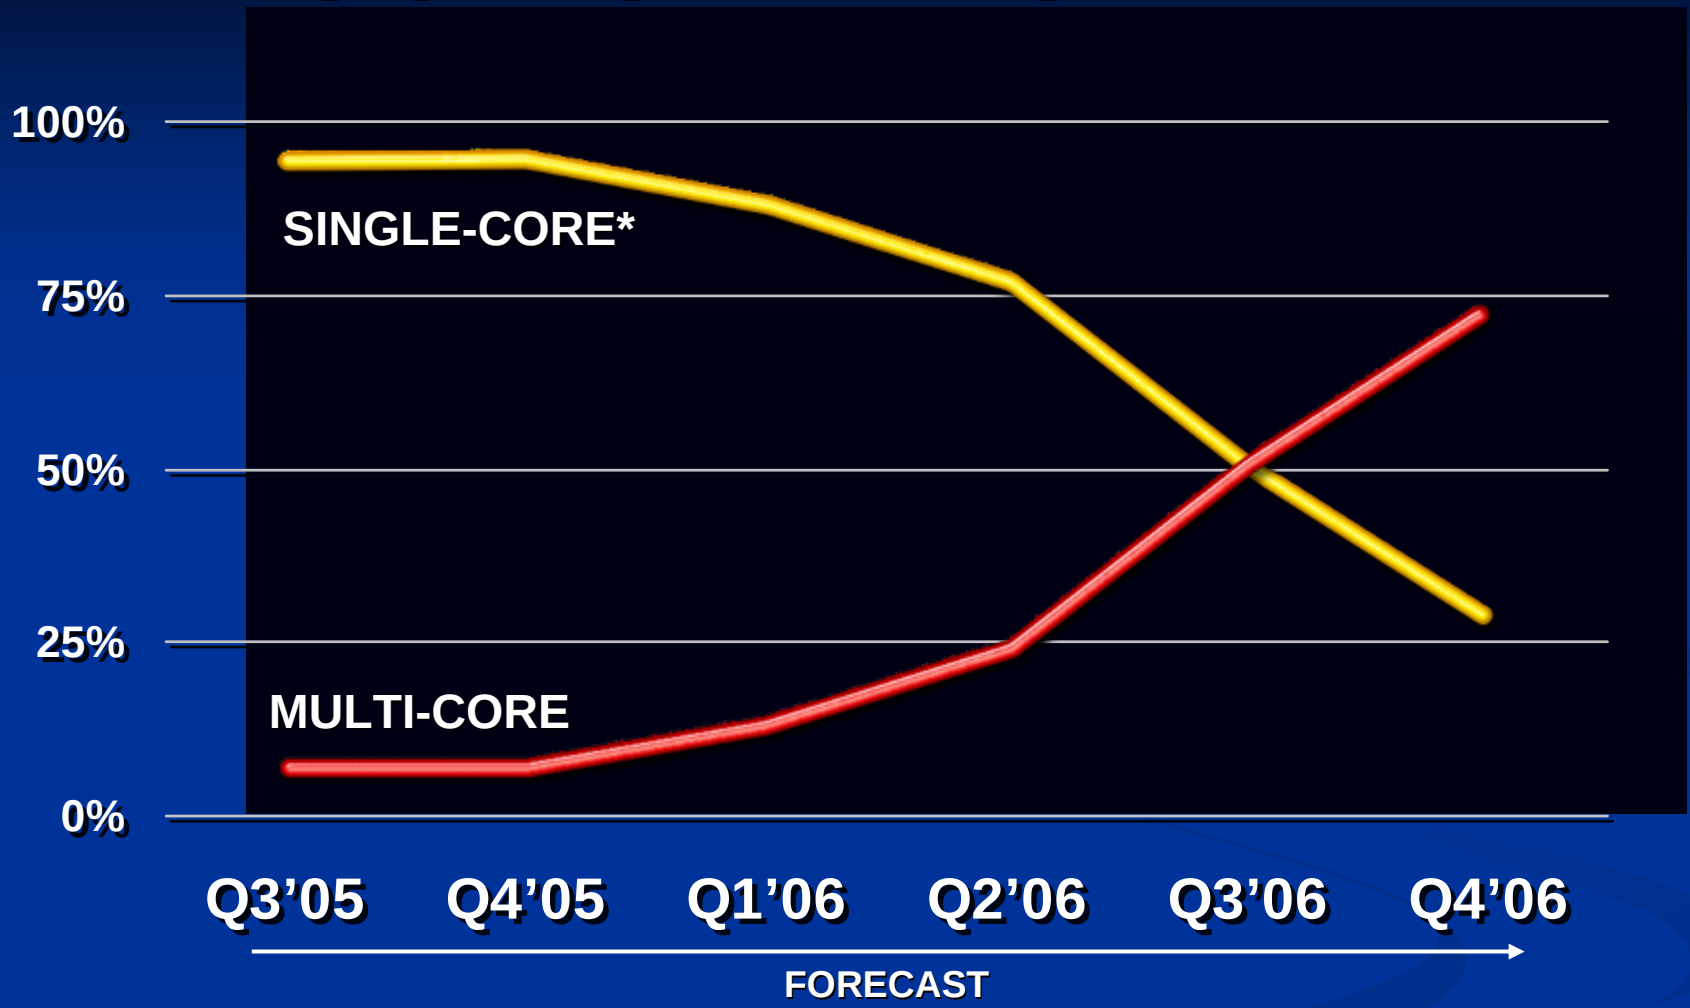

Intel will ship greater than 60 million multi-core CPUs by EOY 2006

Source – Intel Corporation

intel Leap ahead

Advancing to Terascale Applications

Source: electronic visualization lab
University of Illinois

Entertainment

Learning & Travel

Next Gen. Particle Accelerators LHC, ILC,..

Personal Media Creation and Management

Health

Performance

TIPS

GIPS

MIPS

KIPS

RMS

3D & Video

Mult-Media

Text

Tera-scale

Multi-core

Single-core

Kilobytes Megabytes Gigabytes Terabytes

Dataset Size

intel Leap ahead

Compute, Storage, Networking & Software

Intel Confidential

Intel® Blade Server Family

Consolidate – Compute, Networking, Storage, & Management

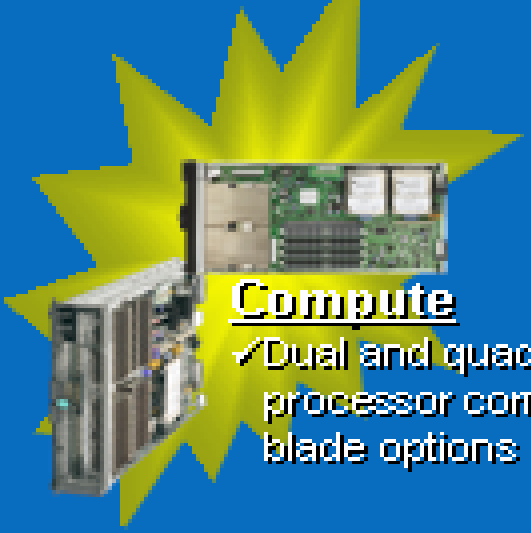

Compute

✓Dual and quad processor compute blade options

Networking

✓Ethernet switching & other fabric options

Storage

✓SAN & NAS connections via Fibre Channel & Ethernet

Integration of resources in a single package for improved utilization

Management

✓Complementary hardware and software simplifies daily tasks

Copyright © 2006, Intel Corporation. All rights reserved. *Other names and brands may be claimed as the property of others.

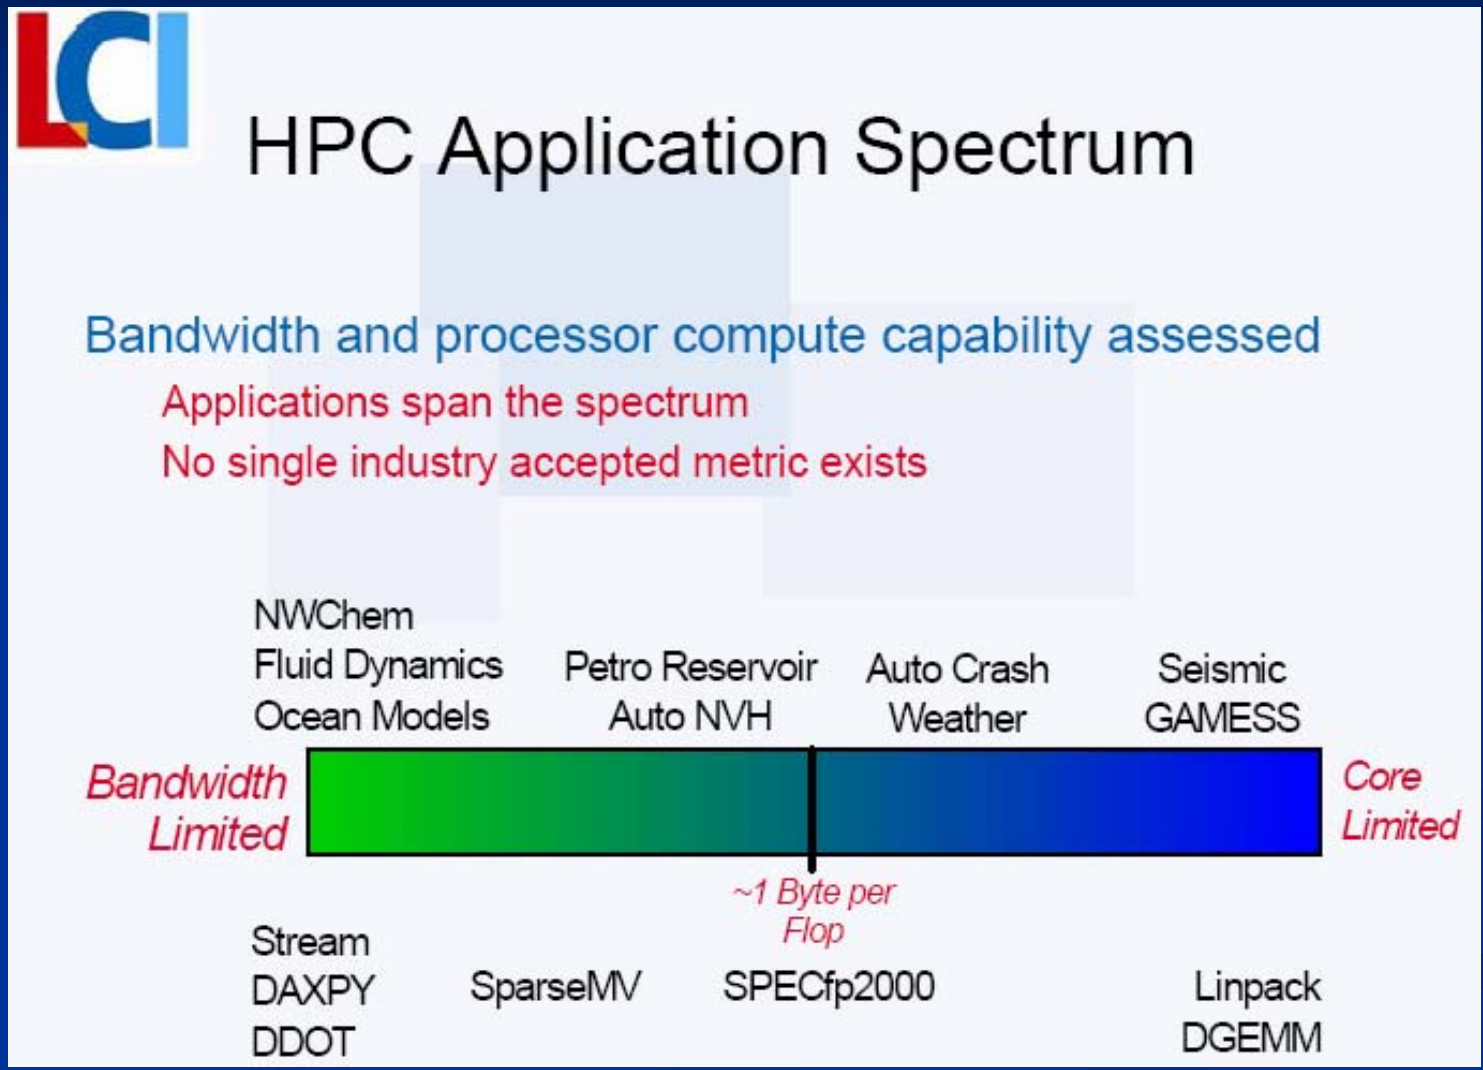

http://www.linuxclustersinstitute.org/Linux-HPC-Revolution/Archive/PDF05/presentations/Pase_D.pdf

intel Leap ahead

Compute: TERAFLOP OF PERFORMANCE

80 Cores
98 Watts
256 GB/s bisection

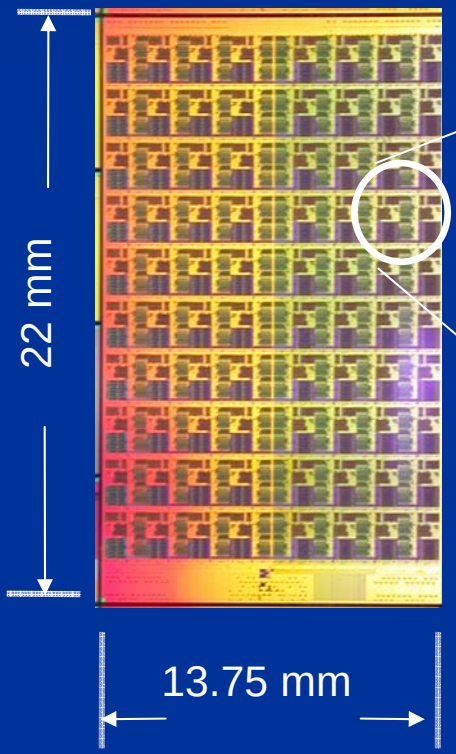

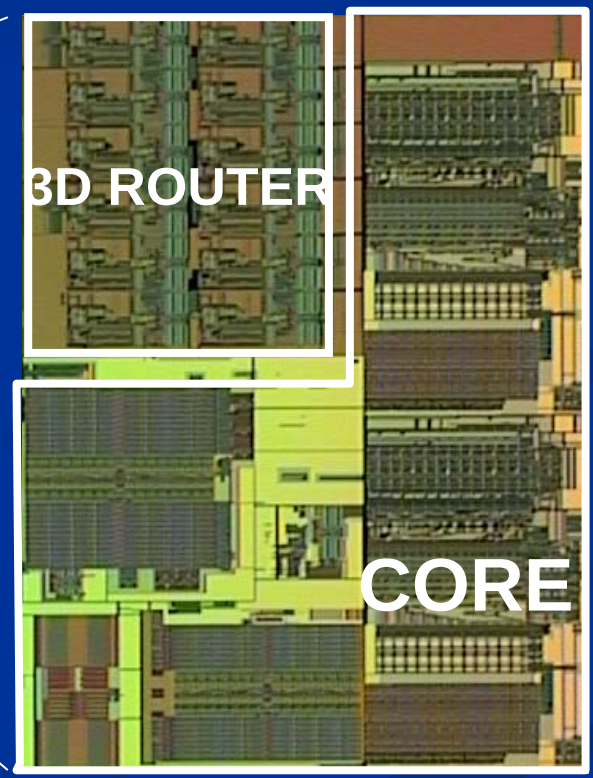

Source: Intel

intel Leap ahead

Storage: TERABYTES OF BANDWIDTH

256 KB SRAM per core
4X C4 bump density
8490 thru-silicon vias

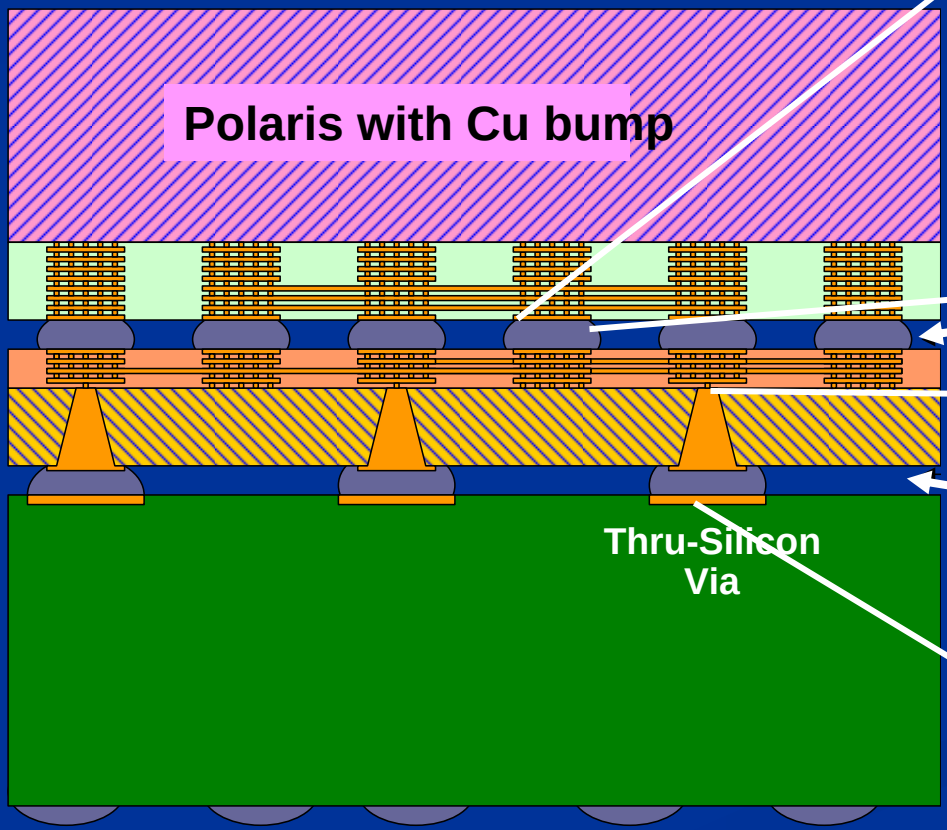

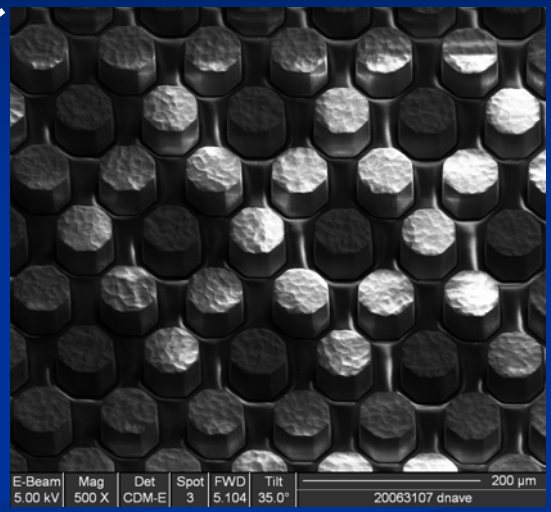

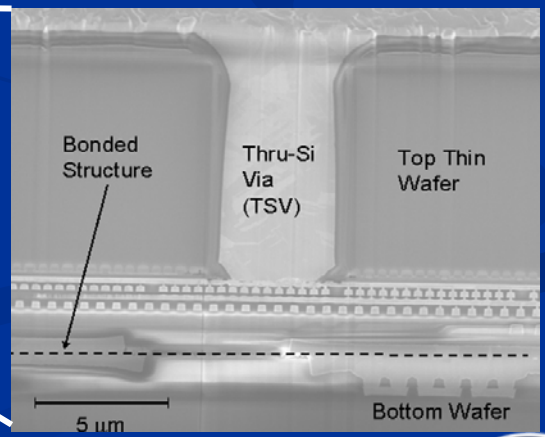

Source: Intel

intel Leap ahead

Networking/Comms: TERABITS OF I/O

TERABITS OF I/O

Target: 1 Tbps chip-to-chip link

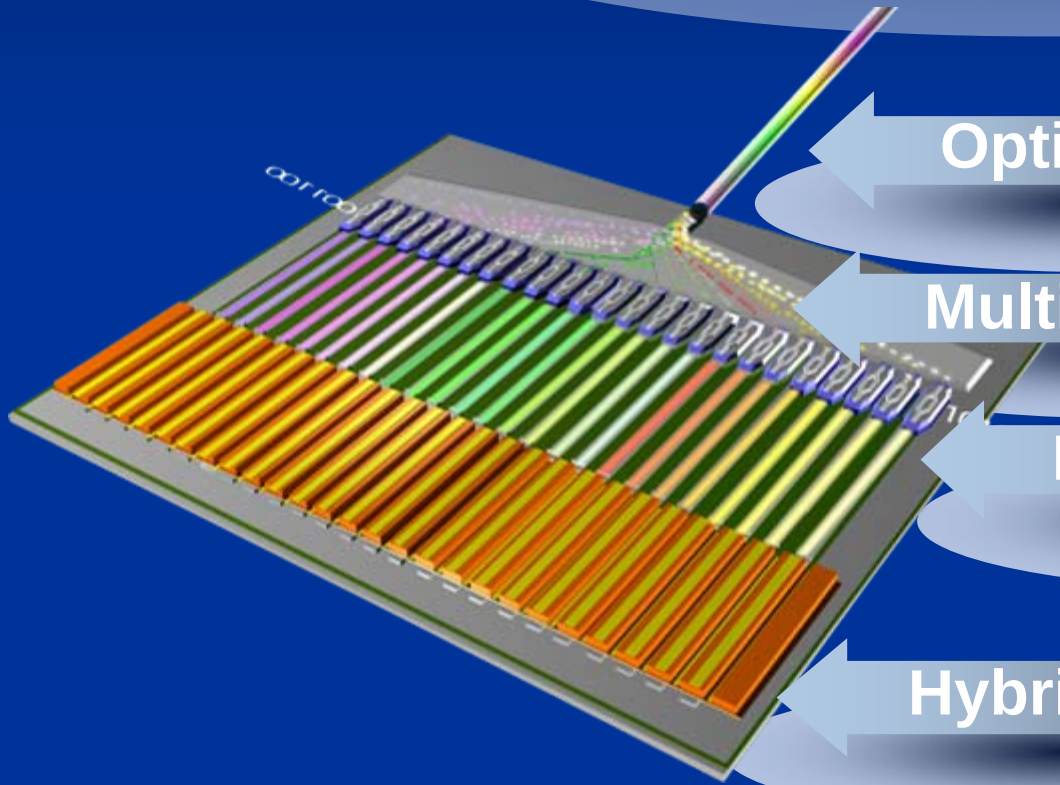

Optical Fiber

Multiplexer

Modulators

Hybrid lasers

30 Gb/s modulator

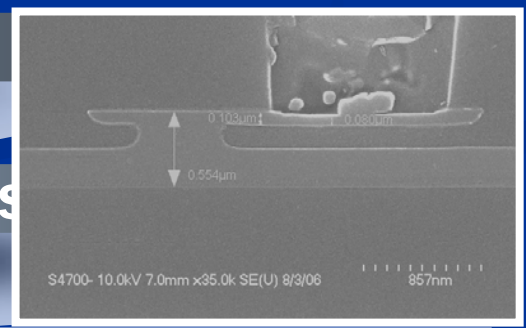

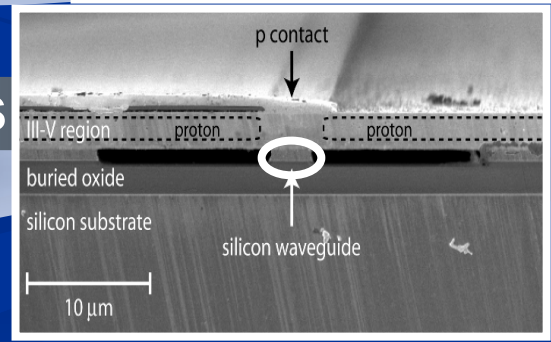

Source: Intel

intel Leap ahead

PARALLEL PROGRAMMING CHALLENGES

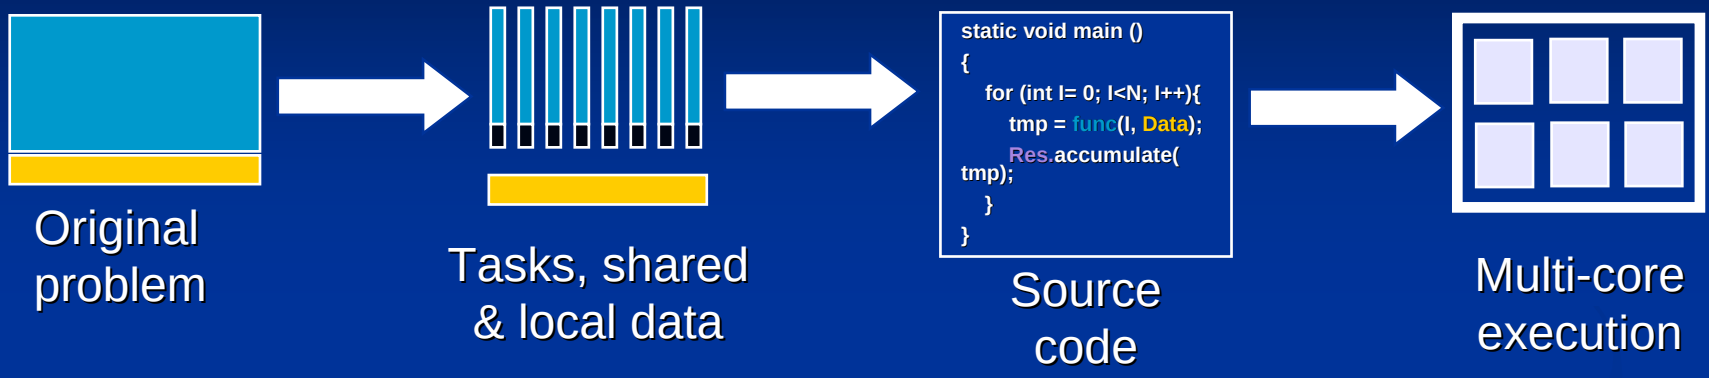

- Extracting concurrency
  - **Tools**
- Expressing concurrency
  - **Programming abstractions & languages**
- Exploiting concurrency
  - **Compilers, runtimes & hardware support**

KEY IS TO IDENTIFY (algorithmic, manual) & MANAGE (locks, conditionals) PARALLELISM

intel Leap ahead

Chip and System Interconnects

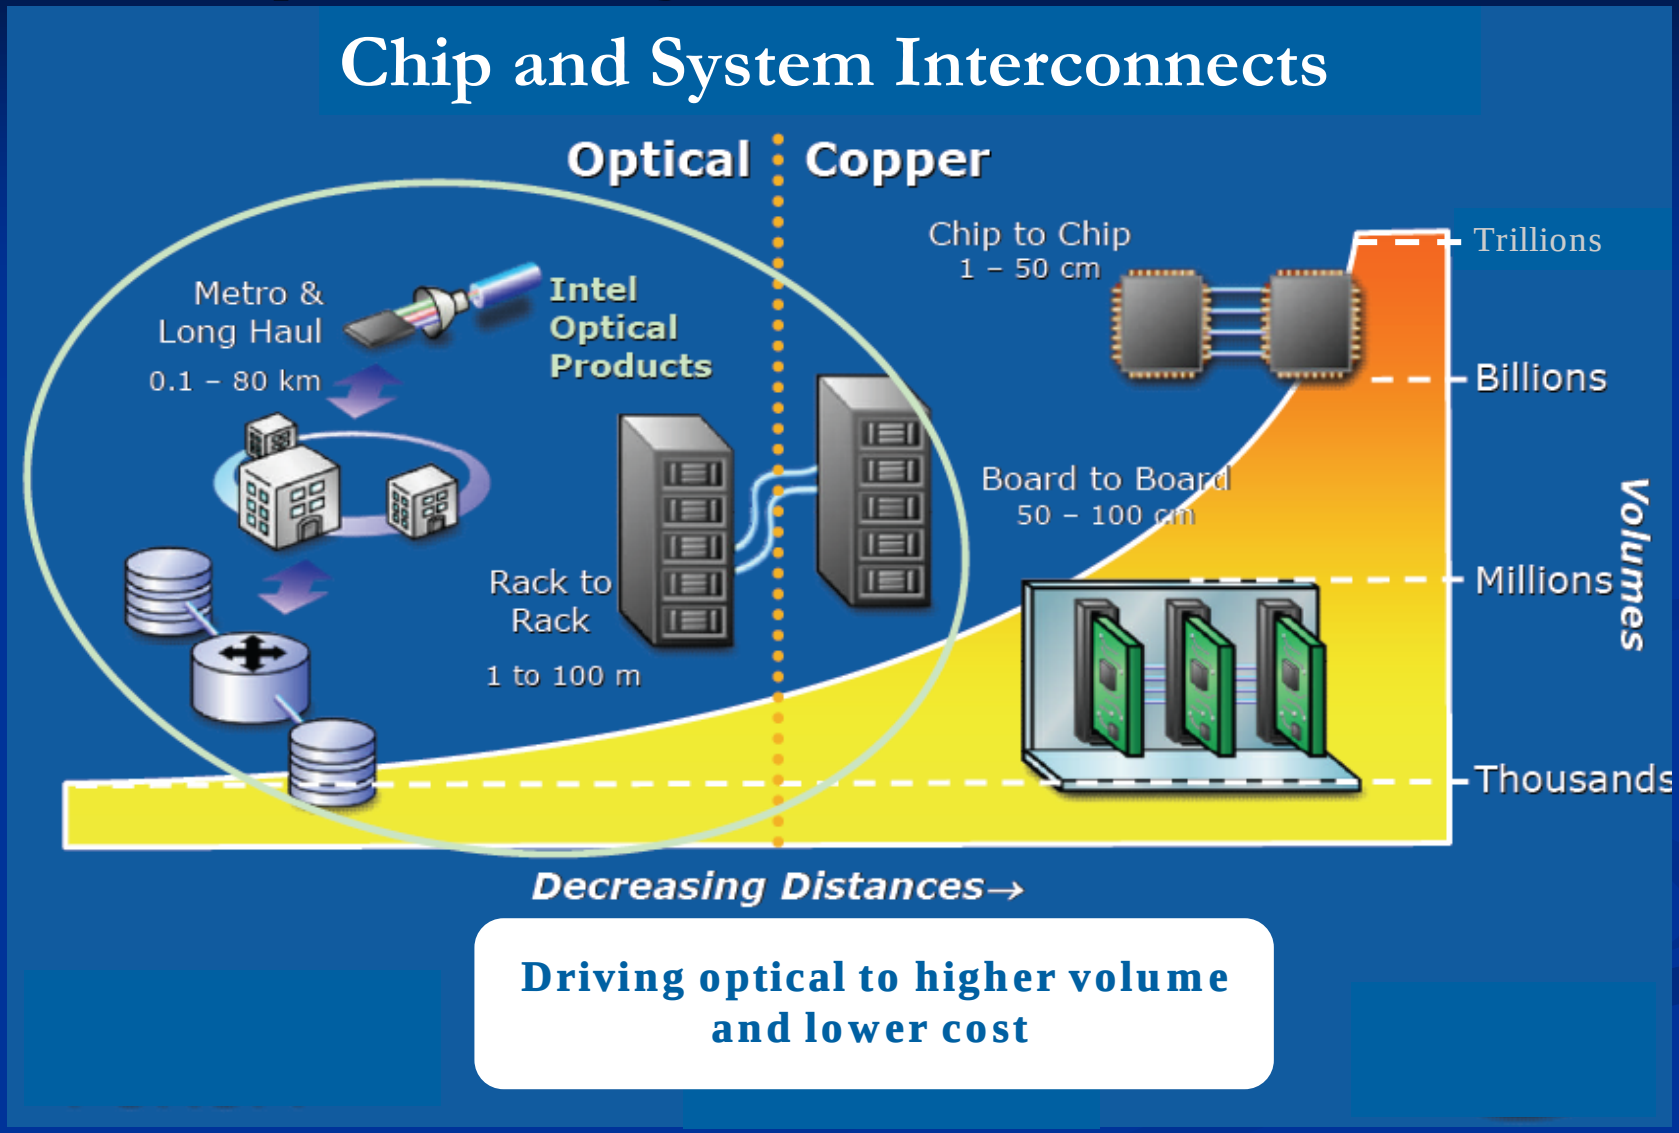

Driving optical to higher volume and lower cost

TRANSACTIONAL MEMORY

Transactional memory is a technique for coordinating how multiple threads access the same memory

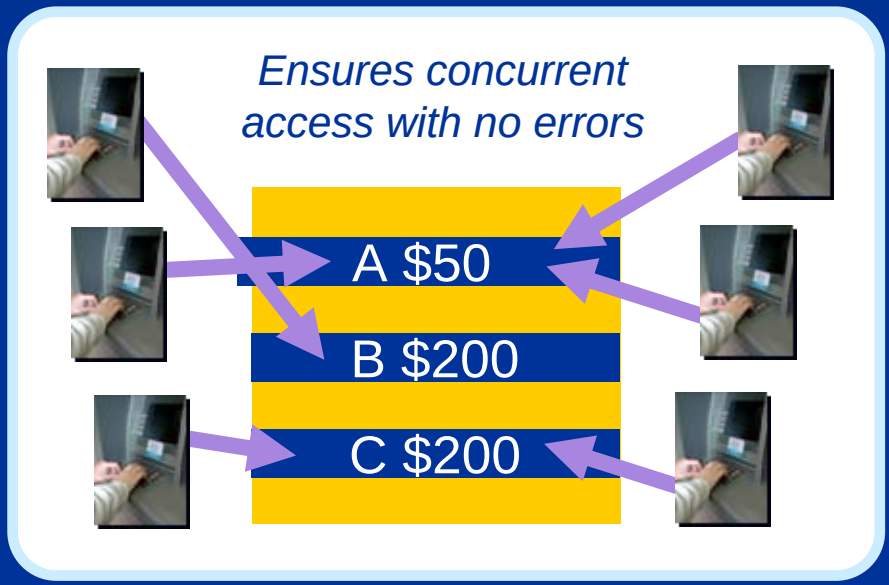

- Greater performance due to concurrent execution
- Eases the writing of parallel programs that work and scale
- Eliminates deadlocks
- Provides a failure recovery mechanism

Java based proof-of-concept shows 3-4x performance gain on a 16-way Xeon® Processor-based system

intel Leap ahead

Robson Technology

- Over the past 10 years alone, processor performance has increased by over 30X hard-drive performance has increased by only 1.3X. And, the gap will continue to grow
- factor of 100,000 difference between DRAM and HDD

  (random read latency of 150 nanoseconds vs. 15 milliseconds)
- About two orders of magnitude in cost per bit for equivalent capacity
- NAND Flash could go between L2 and Disk or in some cases to replace disk altogether.
- Robson's smart controller uses clever write-leveling algorithms to spread the block erasures evenly across the array giving the NAND flash memory a service life consistent with the rest of the platform.
- http://blogs.zdnet.com/OverTheHorizon/?p=11

intel Leap ahead

UNLOCKING PARALLELISM

Must carefully control how multiple threads access common memory

Today we “lock” memory for one thread at a time

- Other threads must wait, reducing multi-core benefit
- Locking code scales poorly, must re-do for more threads
- Can cause critical software deadlocks and errors

Account locked during access

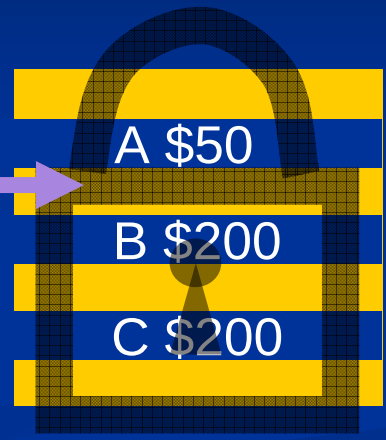

2003 Northeast blackout

Courtesy NASA/JPL-Caltech

Mars rover problem

intel Leap ahead

PROGRAMMING RESEARCH

ENABLE PARALLELISM EVERYWHWERE WITH H/W & S/W

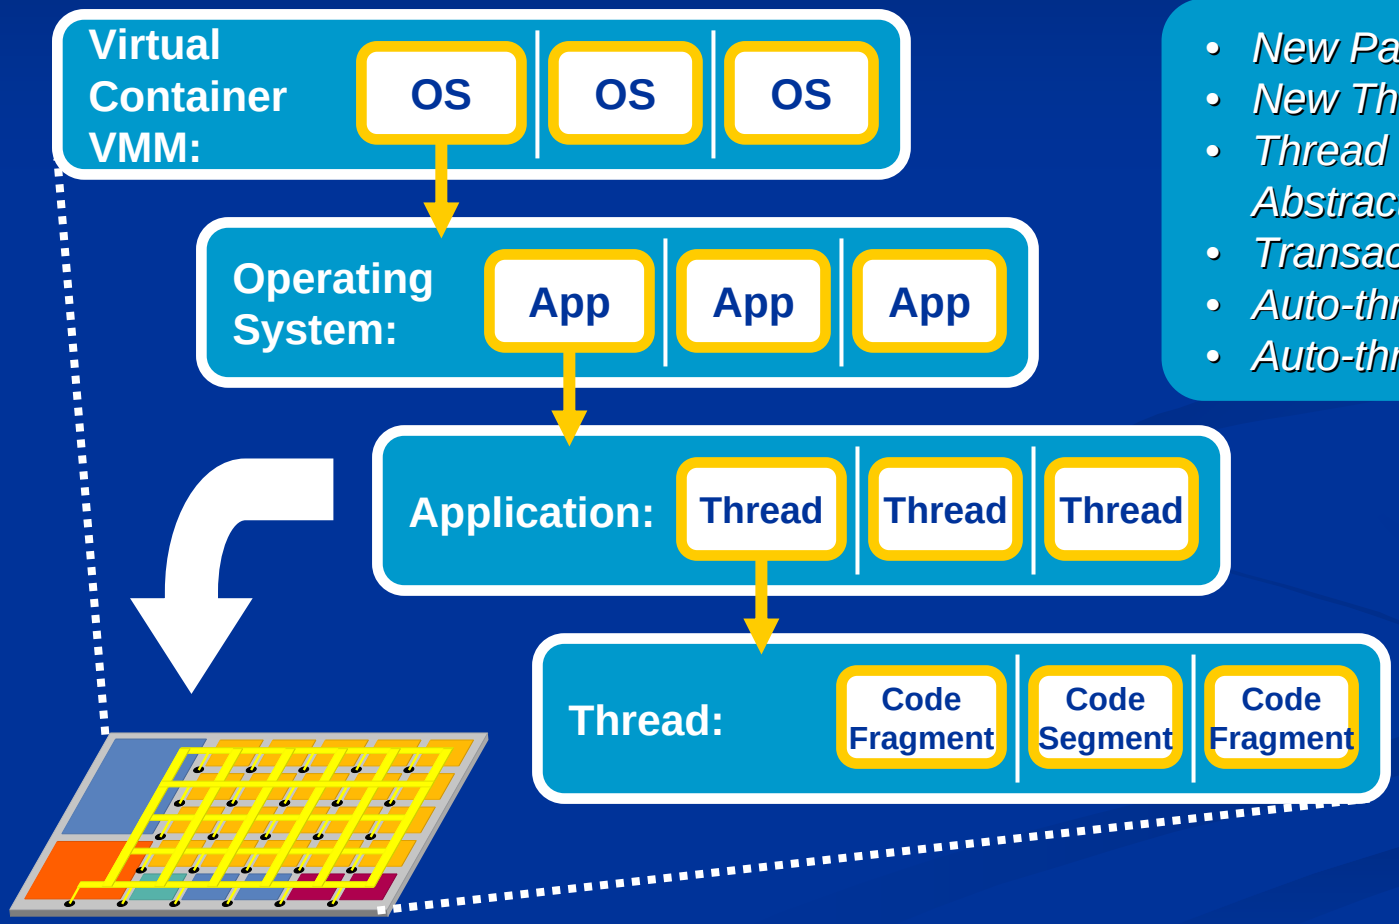

- *New Parallel Languages*
- *New Threading tools*
- *Thread Management & Abstraction layers*
- *Transactional memory*
- *Auto-threading compilers*
- *Auto-threading hardware*

intel Leap ahead

ENABLING MULTI-THREADED S/W

Automatic Loop Parallelization

Optimized MKL and IPP Libraries

Load Balancing w/ VTune™

Critical Path Analysis w/ Thread Profiler

Media and Music	Digital Home	Digital Office	Games	Content Creation	Server
Real Networks, MusicMatch, Sony, Thompson, Windows Media	Roxio, Microsoft, Pinnacle, Cakewalk, Steinberg	Microsoft Office, UGS	EPIC, Id, Square Enix, Activision	Macromedia, Adobe, Discreet, Pixar, Alias, Canopus	Oracle, SQL, DB2, Exchange, SAP, Lotus, BEA, SAS

intel Leap ahead

TERA-SCALE RESEARCH

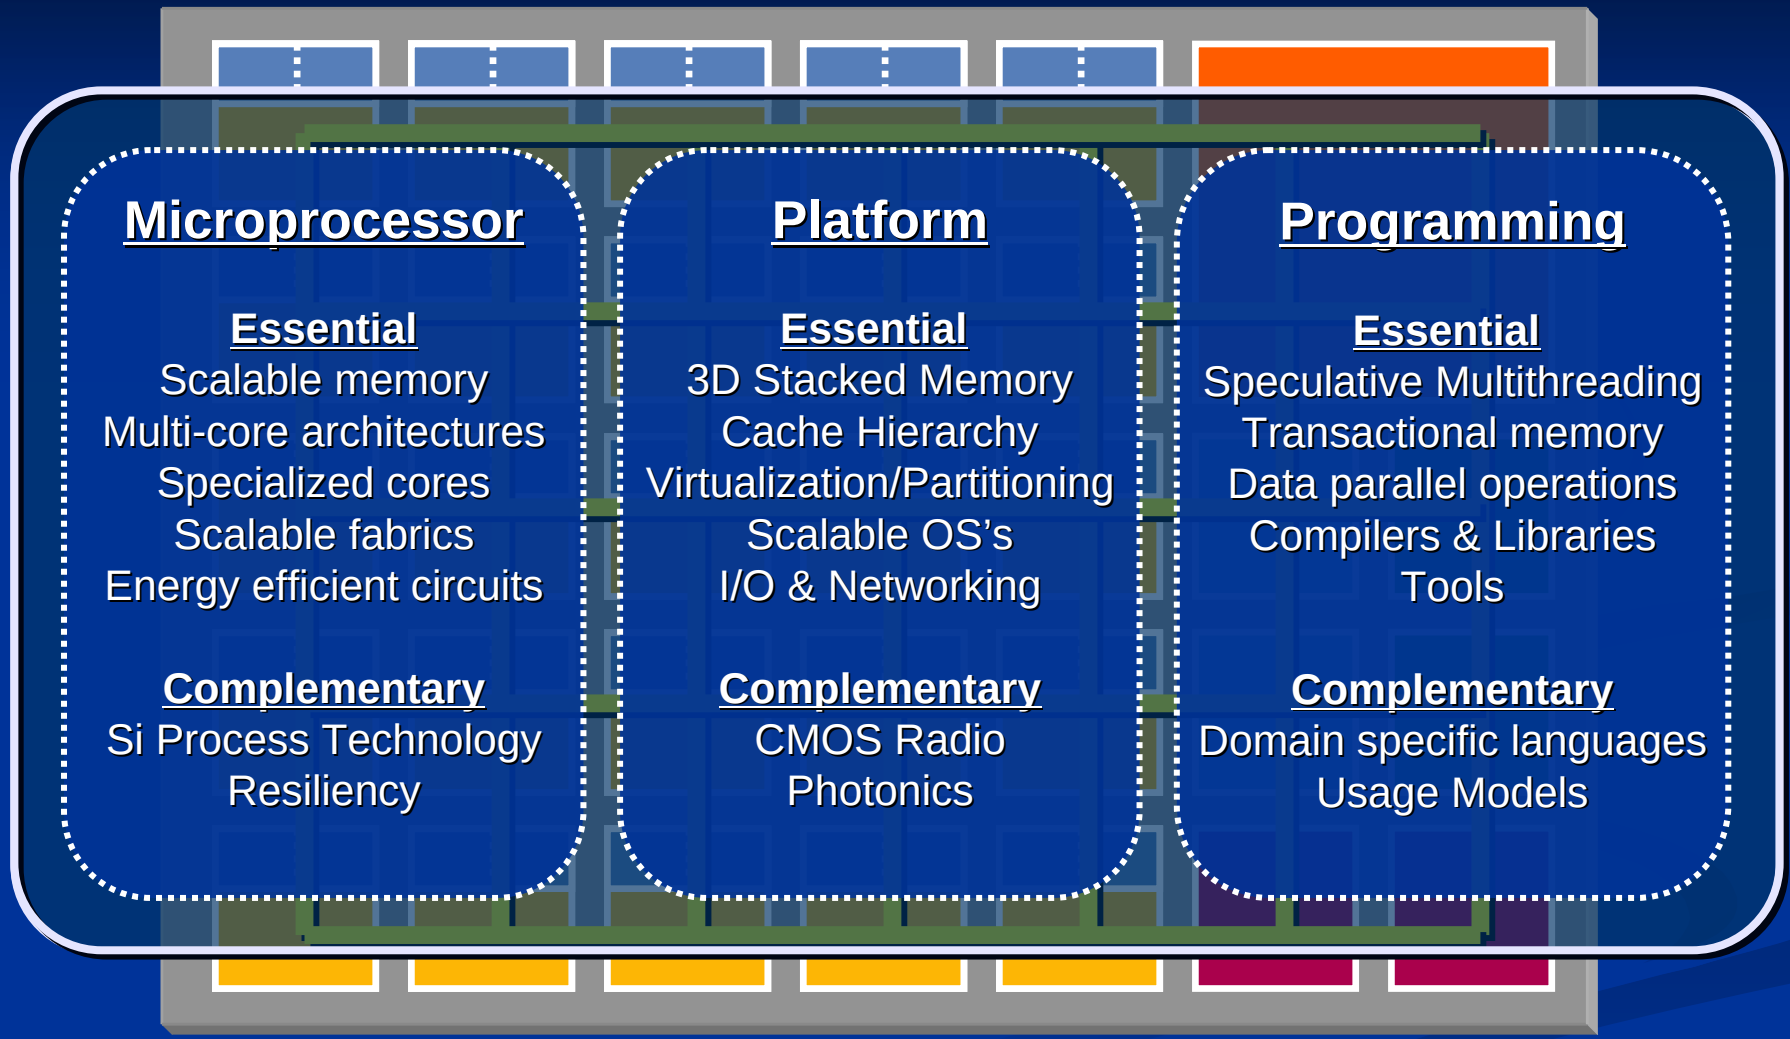

100+ Projects Worldwide

intel Leap ahead

PUTTING IT ALL TOGETHER

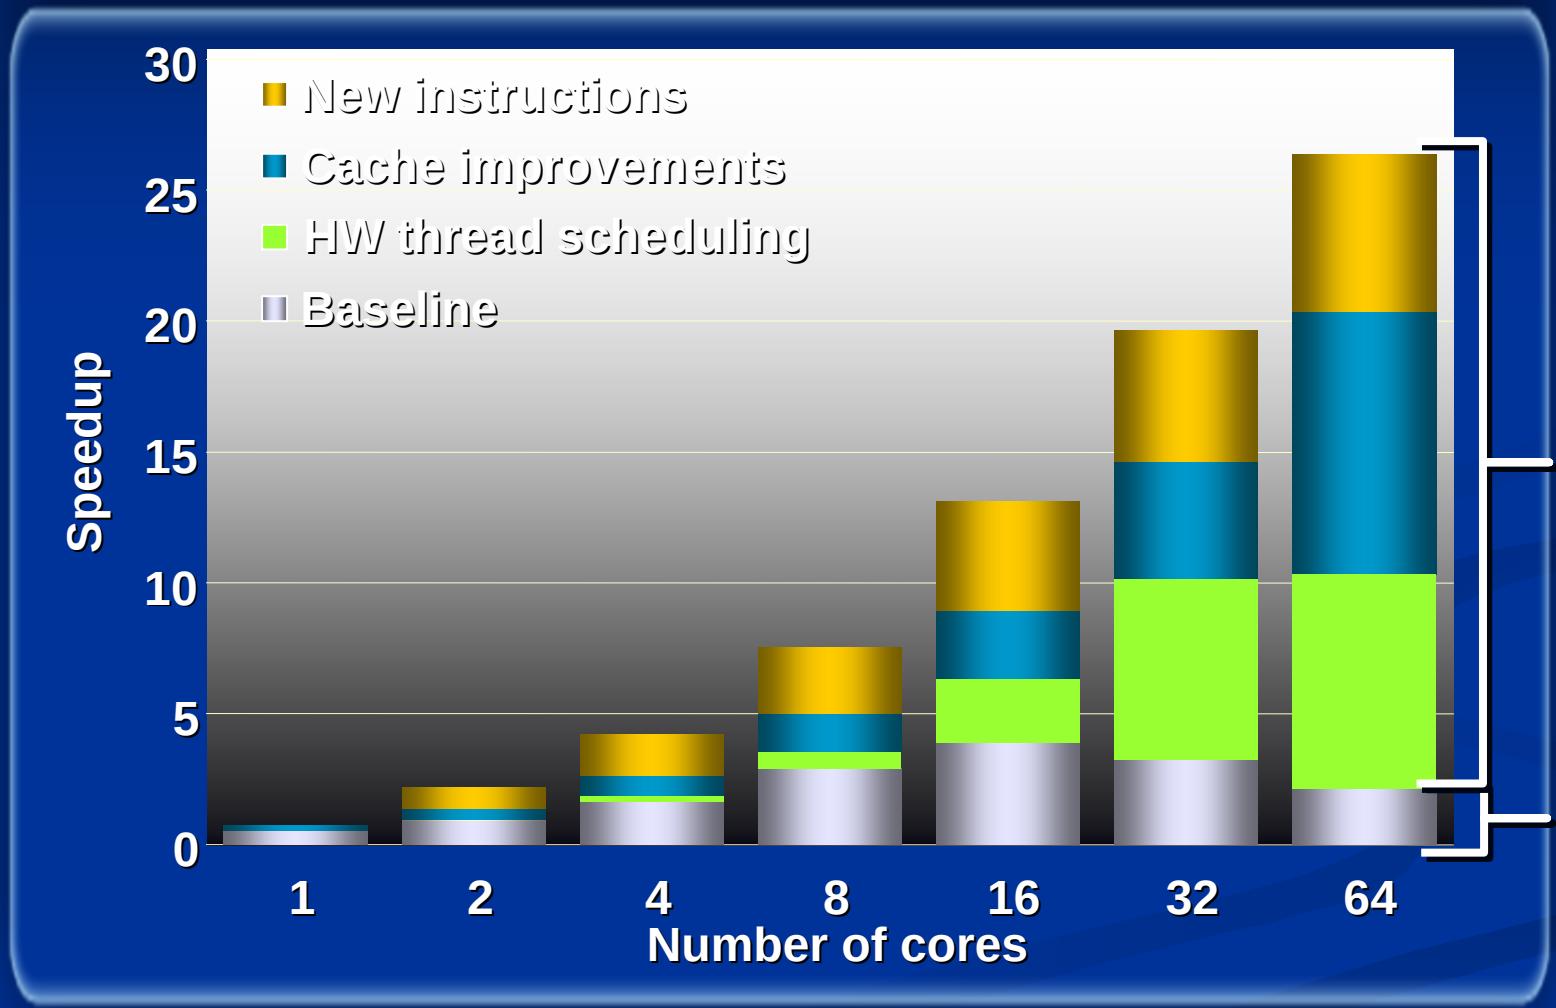

1.7 Billion Transistors per chip already

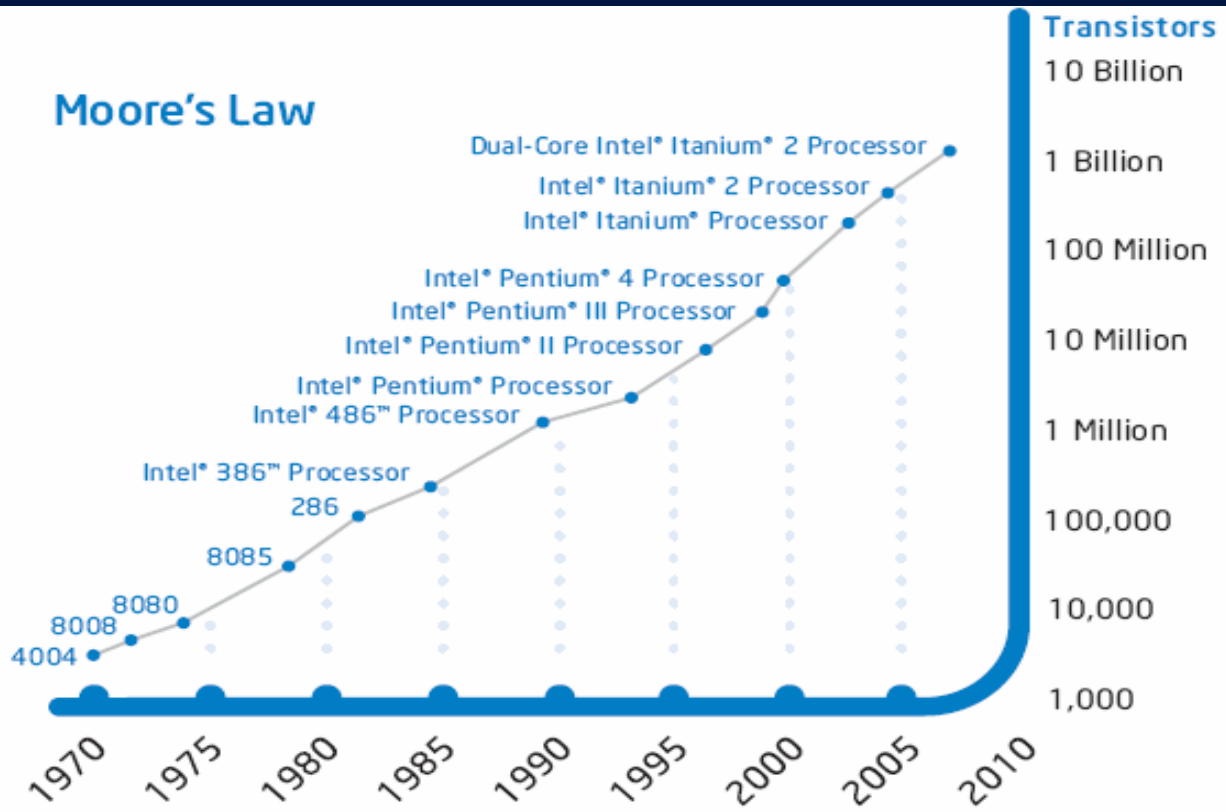

Figure 1. Scaling transistors. The number of transistors is expected to continue to double about every two years, in accordance with Moore's Law. Over time, the number of additional transistors will allow designers to increase the number of cores per chip.

A Multiplier is about 150K gates, roughly 7 per million gates
About 1200 will fit, running at 2 GHz is 2.4 TeraFLOPS!

AND IT GETS BETTER!!! >>>

intel Leap ahead

Intel – New technology generation every 2 years

45 nm Logic Process on Track for Delivery in 2007

Process Name	P1262	P1264	P1266	P1268
Lithography	90 nm	65 nm	45 nm	32 nm
1st Production	2003	2005	2007	2009

Moore's Law continues!

Intel continues to develop a new technology generation every 2 years

intel Leap ahead

Tri-Gate Transistor

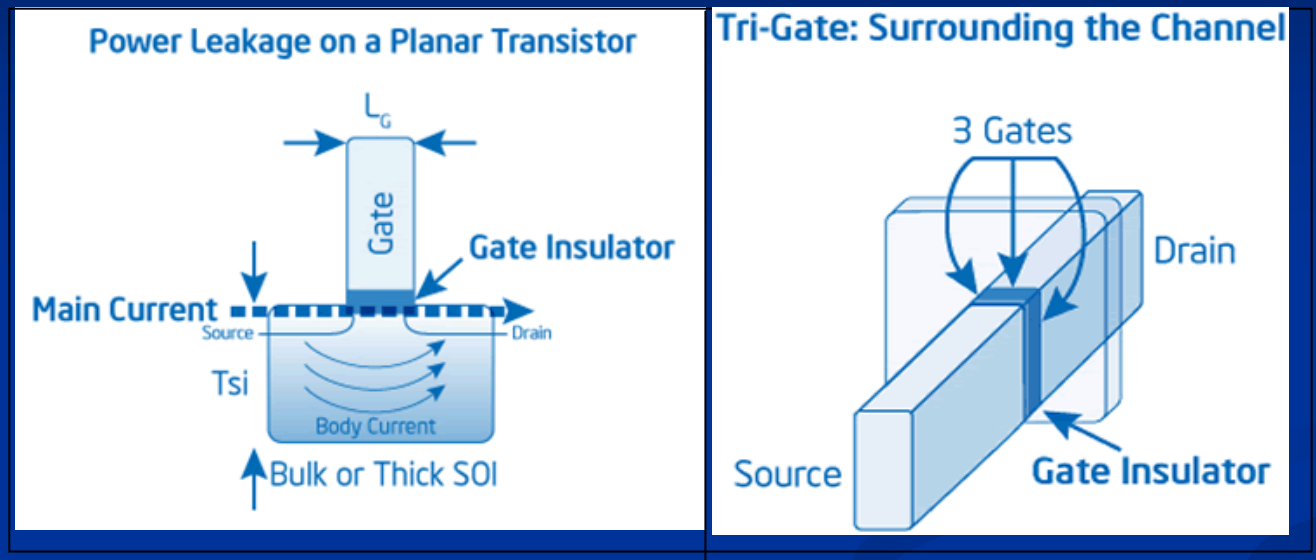

- Shrinking geometries rapidly is a huge win, but it's not the only way to win! More efficient use of the silicon is equally important

Intel will continue to Lead the Way to Exa, Yotta, Zetta,...!!!

intel Leap ahead

Comparison AMD, Xeon, Itanium Memory Speed, etc.
http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8799,00.html

intel Leap ahead


intel
Leap ahead™

MULTI-CORE ENABLING INITIATIVES

USE INNOVATION	Joint effort with SSG, CTG and DEG focused on driving MC uses into BUs
DEV. OUTREACH	64 classes, 1587 Students, Webcasts and Tutorials : 14,239, Downloads: 19, 356
ENGR. METHODS	Tools: Compiler, Libraries, Thread Checking, performance analyzers;
RESEARCH AND EDUCATION	30 3-YR grants awarded, Apps, Algo., Tools, 20 Univs. offering courses
GTM PROGRAMS	GTM Program with MC-enabled ISVs Reseller Matchmaking program for ISVs

intel Leap ahead

SYSTEM CONFIGURATION

1. SPECjbb™2005 Configuration Details: Published/measured results as of Oct 09, 2006. System power was measured during the steady-state of window of the performance run.
   - Intel Xeon Processor 3.60 GHz based platform details: Intel preproduction Server platform with two 64-bit Intel® Xeon™ processors 3.60 GHz with 2MB L2 Cache and 800 MHz system bus and 8GB (8x1024 MB) DDR2-400 memory, Microsoft Windows Server x32 Enterprise Edition. Results published by Principled Technologies at http://www.principledtechnologies.com/clients/reports/Intel/WSPECjbb2005_0506.pdf as of May 23, 2006.
   - Dual-Core Intel Xeon Processor 5160 based platform details: Intel Server pre-production platform with two Dual-Core Intel Xeon Processor 5160, 3.00 GHz with 4M L2 Cache, 1333 MHz system bus, 8GB (8x1GB) 667MHz FB-DIMM memory, Windows 2003 Enterprise Edition. BEA JRockit(R) 5.0 P26.4.0. Run with two jvm instances
   - Quad-Core Intel Xeon Processor based platform details: Intel Server pre-production platform with two Quad-Core Intel Xeon Processor X5345, 2.33 GHz with 2x4M L2 Cache, 1333 MHz system bus, 8GB (8x1GB) 667MHz FB-DIMM memory, Windows 2003 Enterprise Edition. BEA JRockit(R) 5.0 P26.4.0. Run with two jvm instances
2. Performance tests and ratings are measured using specific systems and/or components and reflect approximate performance of Intel products as measured by those tests. Any difference in system hardware, software, or configuration may affect actual performance. Buyers should consult other sources of information to evaluate system or component performance they are considering purchasing. For information on performance tests and performance of Intel products, visit http://www.intel.com/performance/resources/limits.htm

intel Leap ahead

RECORD SETTING TPC-C PERF. TODAY

Intel® Xeon® 7100 Series 4-Way Servers

- **Configurations Published:**
  - Database Performance on TPC-C*; represents the transaction throughput of a database server in an OLTP client/server environment.
  - **AMD Opteron* platform:** Hewlett-Packard* ProLiant* DL585-G2 with 4x AMD Opteron* processor 8220SE (4 Processors/8 Cores/8 Threads), 32x 4 GB DDR2, Microsoft* SQL Server* 2005 Enterprise Edition x64 SP1 database, Microsoft* Windows* Server 2003 Enterprise Edition x64 SP1. Results at http://www.tpc.org/results/individual_results/HP/HP_DL585G2_2.8DC_es.pdf. Referenced as published: 262, 989 tpmC; $2.09/tpmC; Availability Date March 26, 2007.
  - **Intel Xeon 7140M platform** IBM* System x3950 Server System using 4x Dual-Core Intel® Xeon® processor 7140N (3.33 GHz, 16 MB L3 cache, 4 Processors/8 Cores/16 Threads), 32x 4 GB (128 GB) memory, DB2* 9 database, SuSe* LINUX Enterprise 10. Results at http://www.tpc.org/results/individual_results/IBM/ibm_x3950_DB2_Linux_es.pdf Referenced as published: 314,468 tpmC; $4.75/tpmC; Availability Date Nov 30, 2006.
  - **Intel Xeon 7140N platform** HP ProLiant ML570G4 Server System using 4x Dual-Core Intel® Xeon®processor 7140M (3.40 GHz, 16 MB L3 cache, 4 Processors/8 Cores/16 Threads), 16x4 GB (64 GB) memory, Microsoft SQL Server 2005 Enterprise Edt (x64), Microsoft Windows Server 2003 Enterprise x64 Edition SP1. Results at http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=106101901 Referenced as published: 318,407 tpmC; $1.88/tpmC; Availability Date April 19, 2007.

intel Leap ahead

MULTI-CORE MOTIVATION

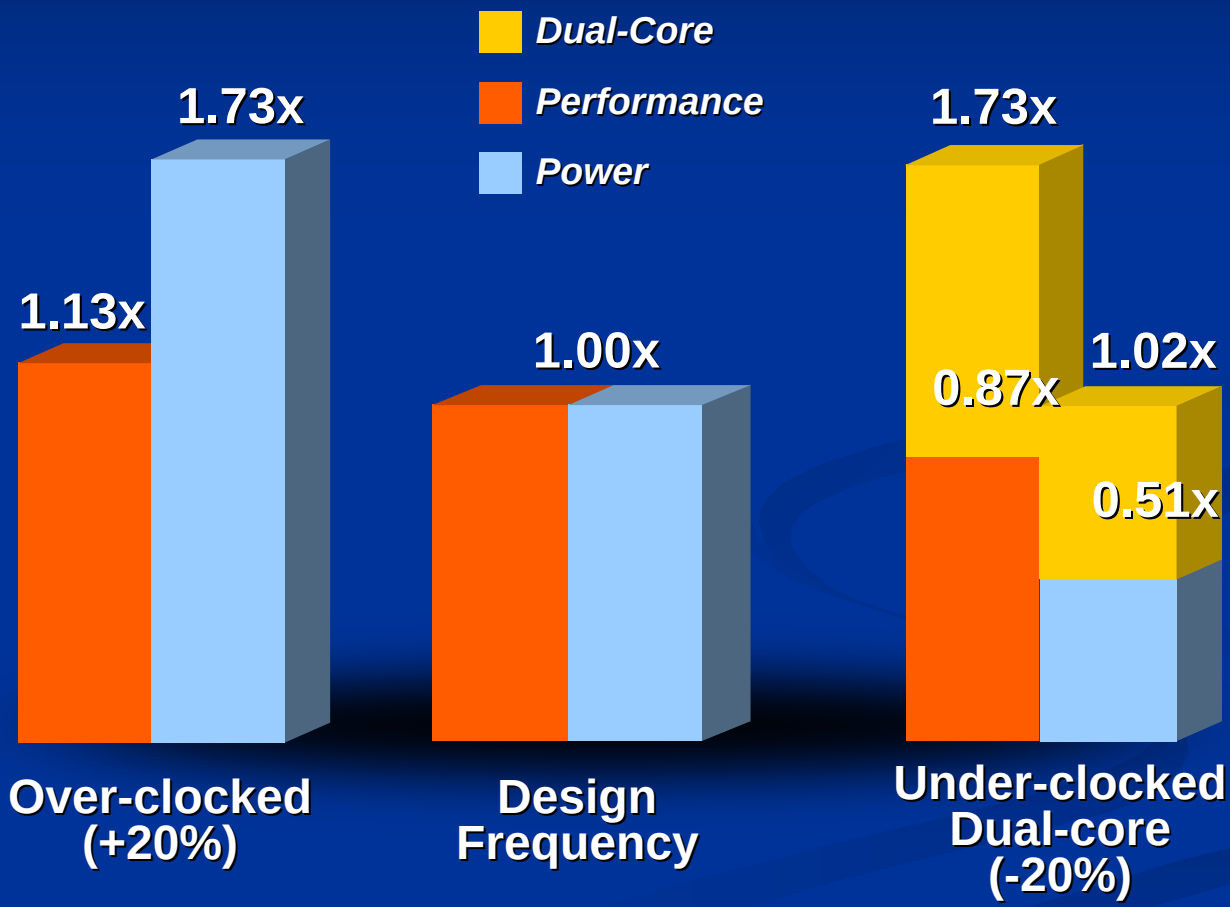

Relative single-core frequency and Vcc

intel Leap ahead

MULTI-CORE POWER EFFICIENCY

System Power
(measured system Watts)

Lower is better

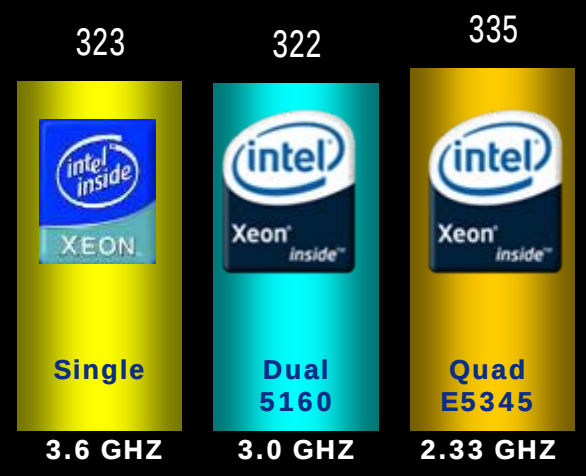

Performance
(SPECjbb2005*)

Higher is better

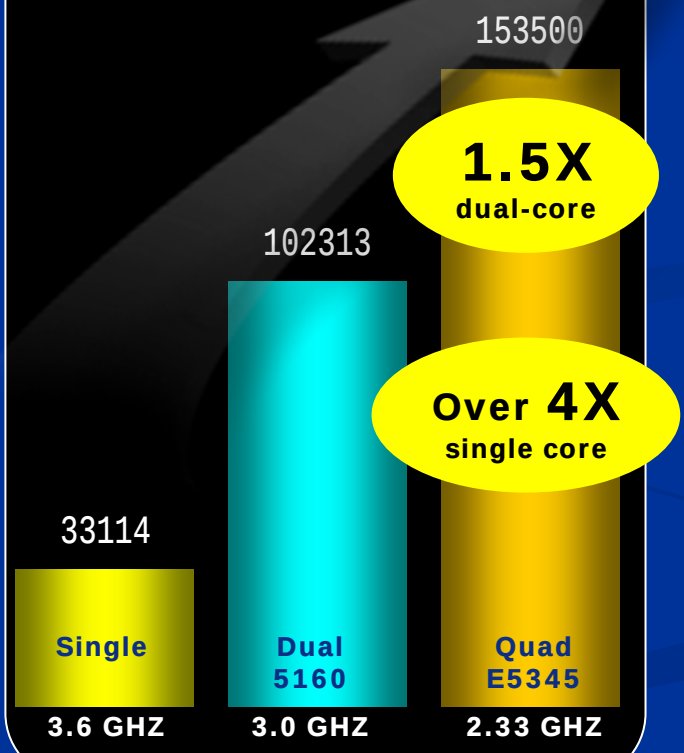

Exceptional Software Scaling**

Same Power Envelope

Outstanding Performance per Watt

* http://www.principledtechnologies.com/
** For most software applications

EMERGING ‘KILLER’ APPS

Recognition

Feature Extraction
Cancer Cell Detection

Mining

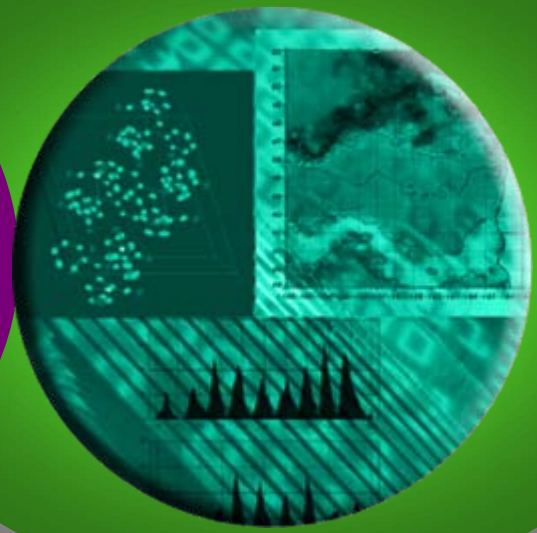

Financial Analytics
Media Search &
Retrieval

Synthesis

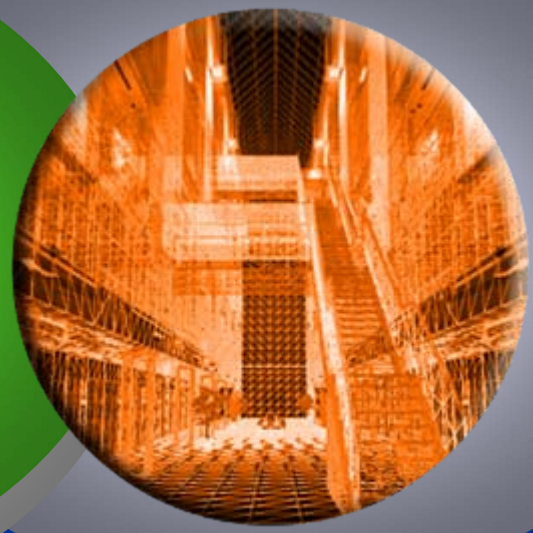

Ray Tracing
Physical Simulation

Users Benefit + Highly Parallel = Tera-scale Application

intel Leap ahead