Aim High

Intel Technical Update

Teratec '07 Symposium

June 20, 2007

Stephen R. Wheat, Ph.D.
Director, HPC
Digital Enterprise Group

Risk Factors

Today's presentations contain forward-looking statements. All statements made that are not historical facts are subject to a number of risks and uncertainties, and actual results may differ materially. Please refer to our most recent Earnings Release and our most recent Form 10-Q or 10-K filing available on our website for more information on the risk factors that could cause actual results to differ.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations (http://www.intel.com/performance/resources/limits.htm).

Rev. 7/19/06

Real World Problems Driving Petascale & Beyond

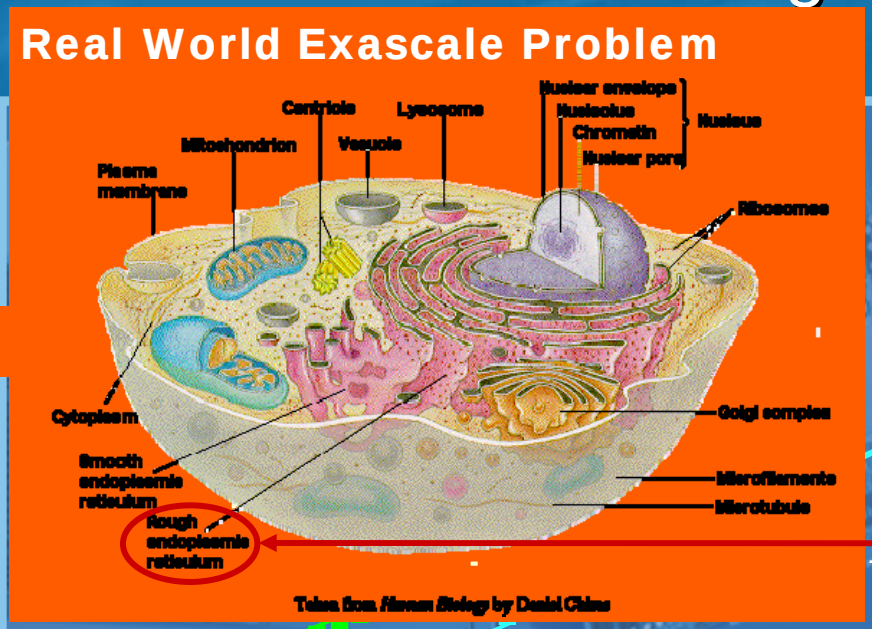

SUM
Of Top500

#1

What we can model today with <100TF

Example real world challenges:

- Full modeling of an aircraft in all conditions
- Green airplanes
- Genetically tailored medicine
- Understand the origin of the universe
- Synthetic fuels everywhere
- Accurate extreme weather prediction

1 ZFlops
100 EFlops
10 EFlops
1 EFlops
100 PFlops
10 PFlops
1 PFlops
100 TFlops
10 TFlops
1 TFlops
100 GFlops
10 GFlops
1 GFlops
100 MFlops

1993 1999 2005 2011 2017 2023 2029

Germany

Silicon Future

90nm
2003

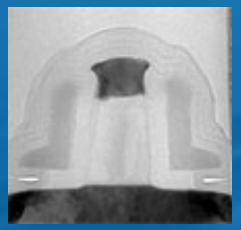

65nm
2005

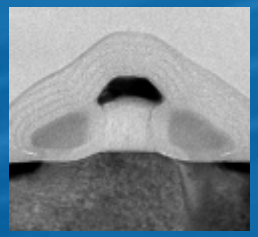

45nm
2007

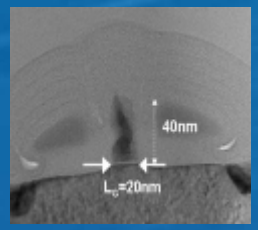

2005 – 2012

32nm
2009

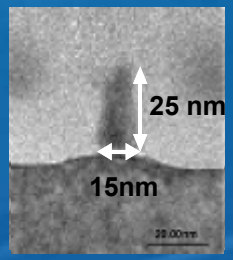

22nm
2011

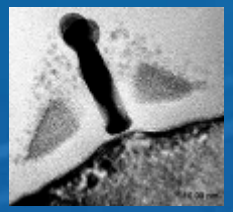

2013 – 2017

16nm
2013

11nm
2015

8nm
2017

New Intel technology generation every 2 years
Intel R&D technologies drive this pace well into the next decade

Roadmap

Research

Intel Design & Process Cadence

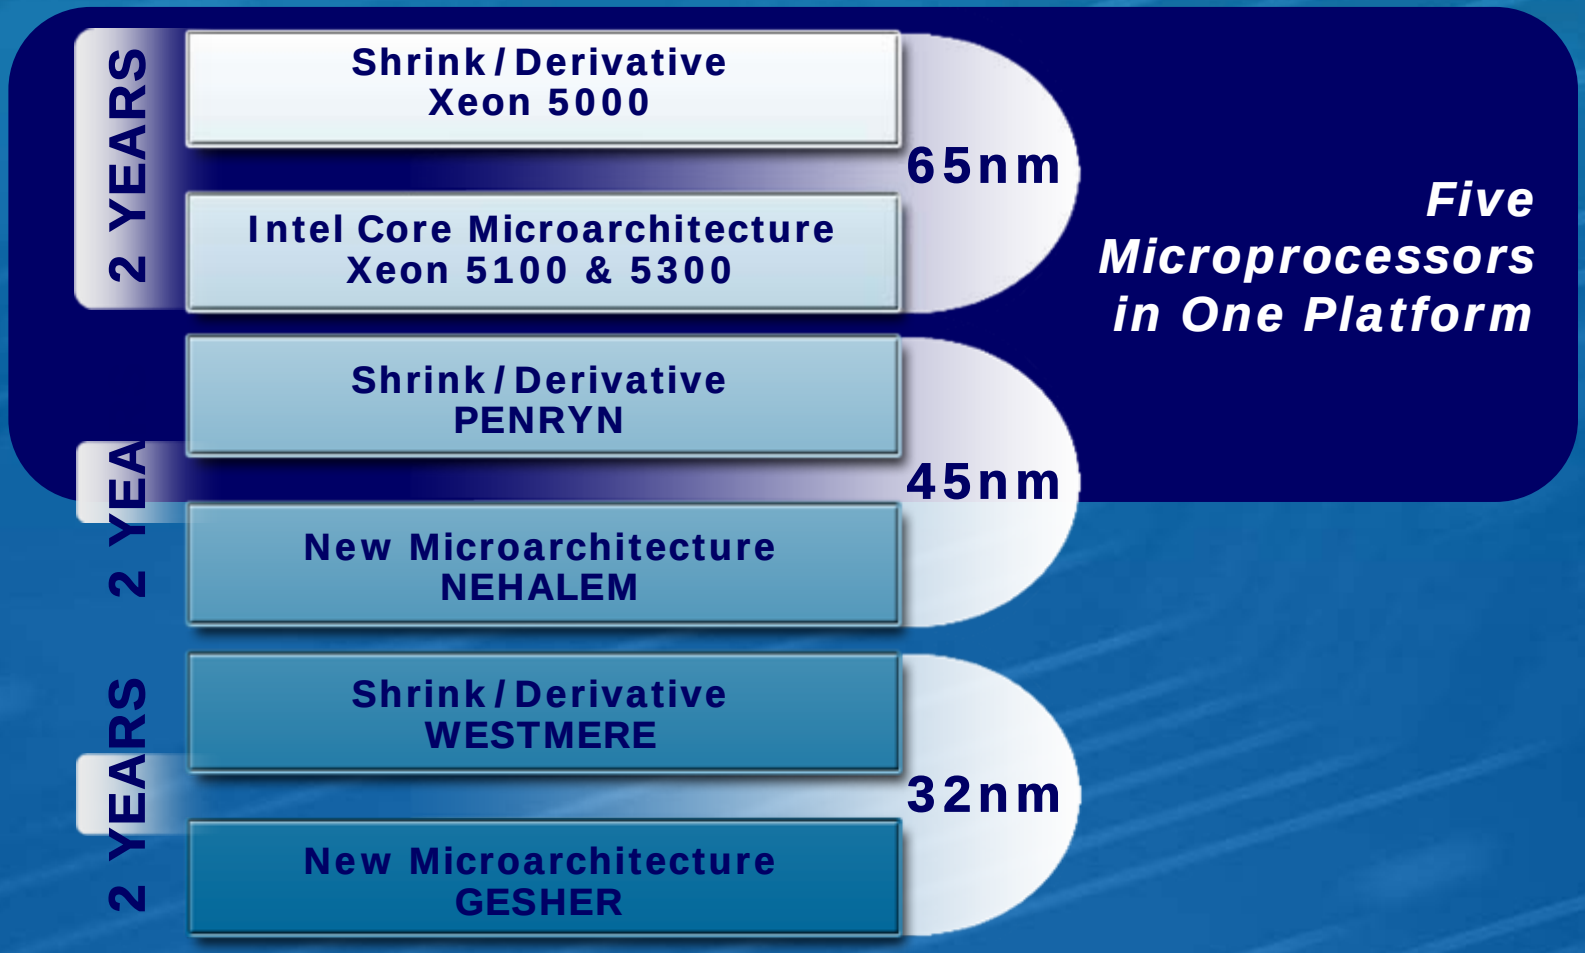

Processor Performance

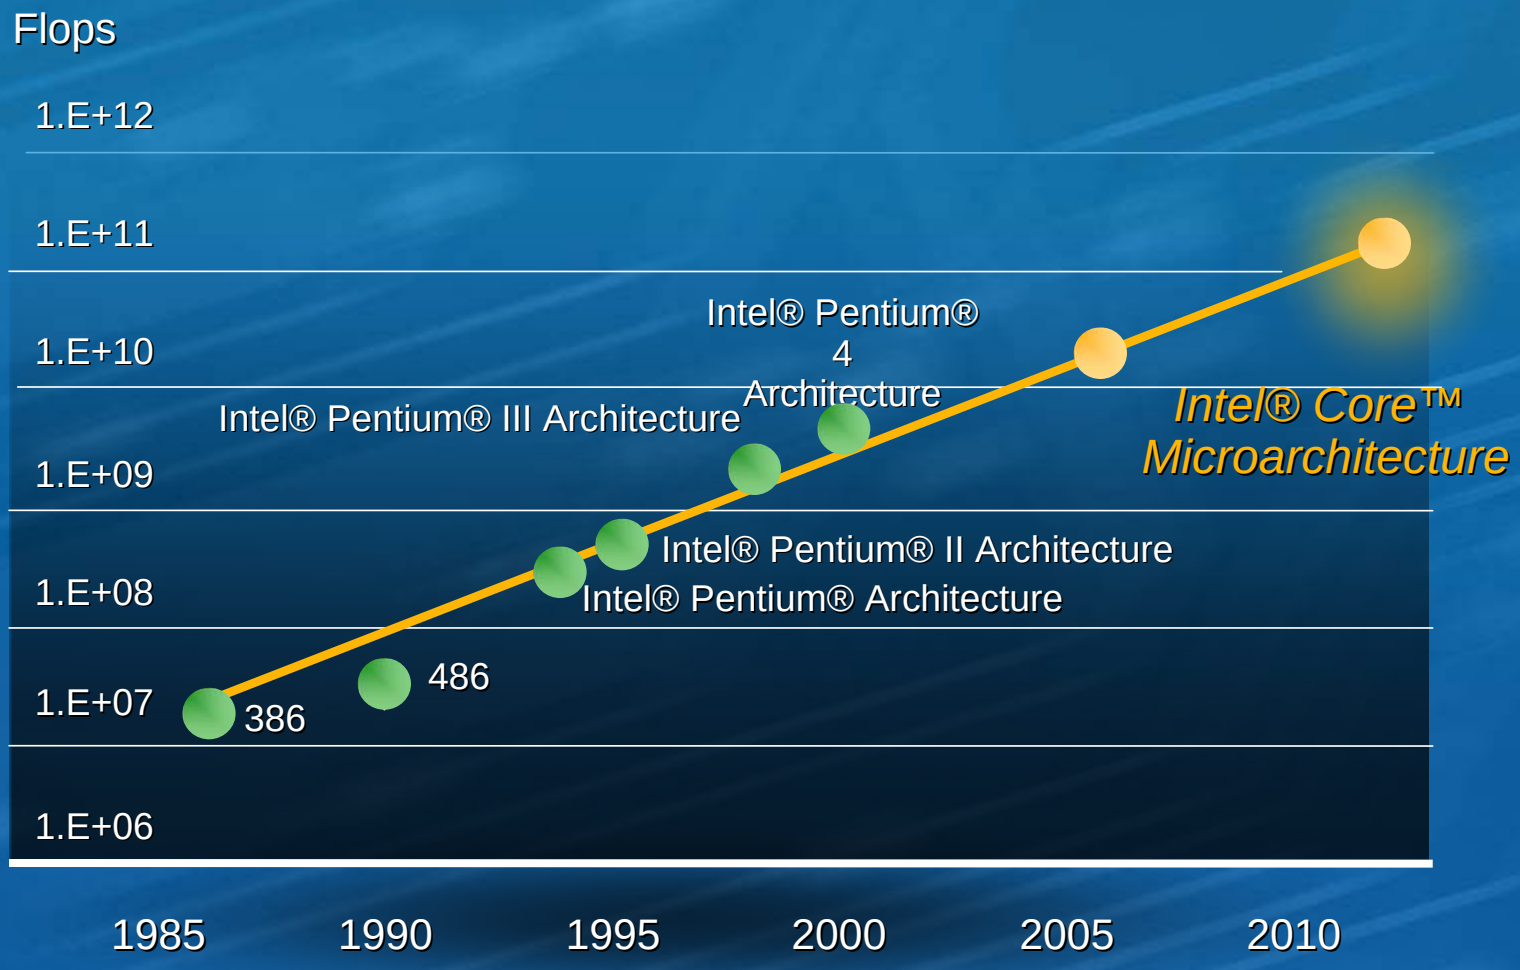

*Reaching Petascale with ~100,000 Processors in 2010**

Assuming approx. 100Glops processors

* Petascale assumes 10's of PF Peak Performance and 1PF Sustained Performance on HPC Applications.

Why Multi-Core?

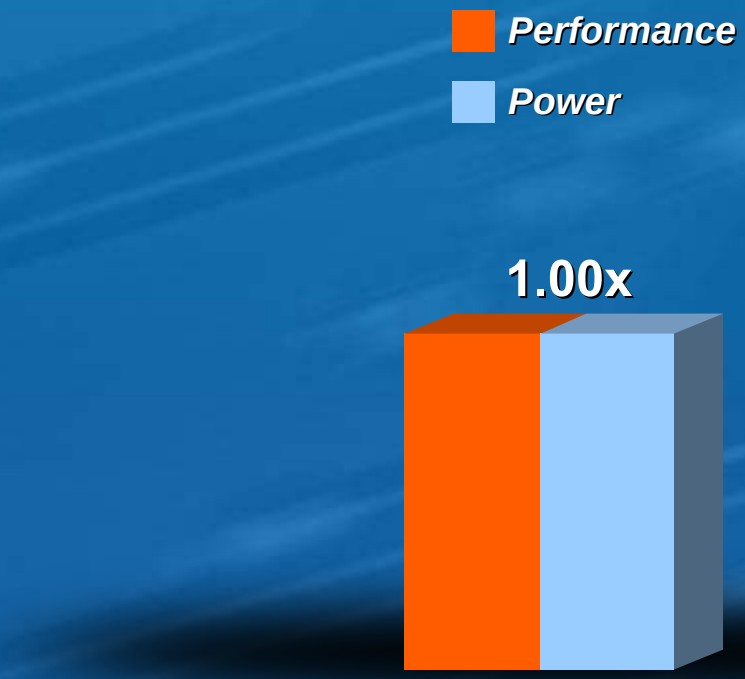

Relative single-core frequency and Vcc

Over-clocking

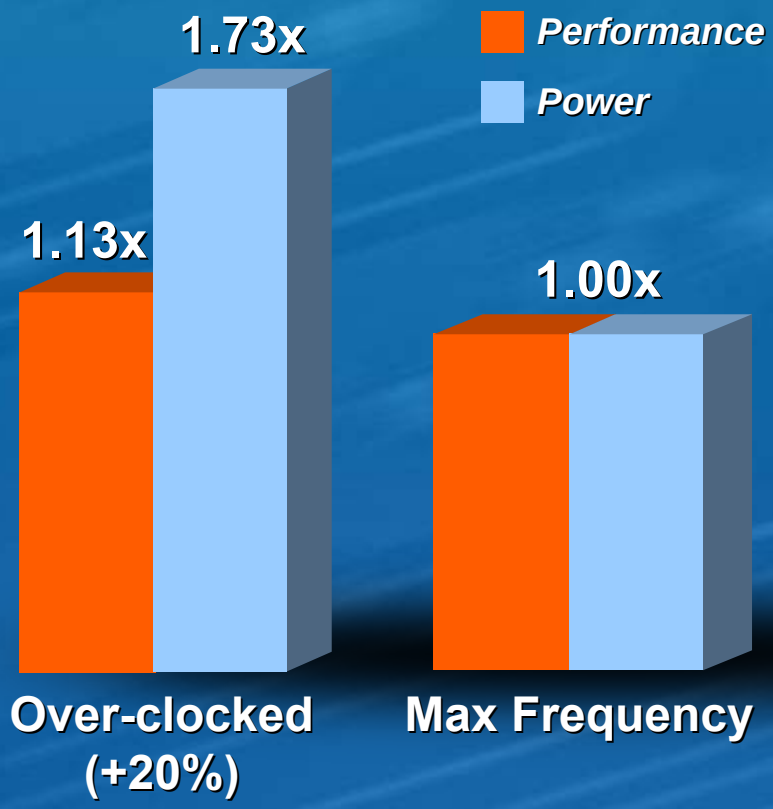

Relative single-core frequency and Vcc

Under-clocking

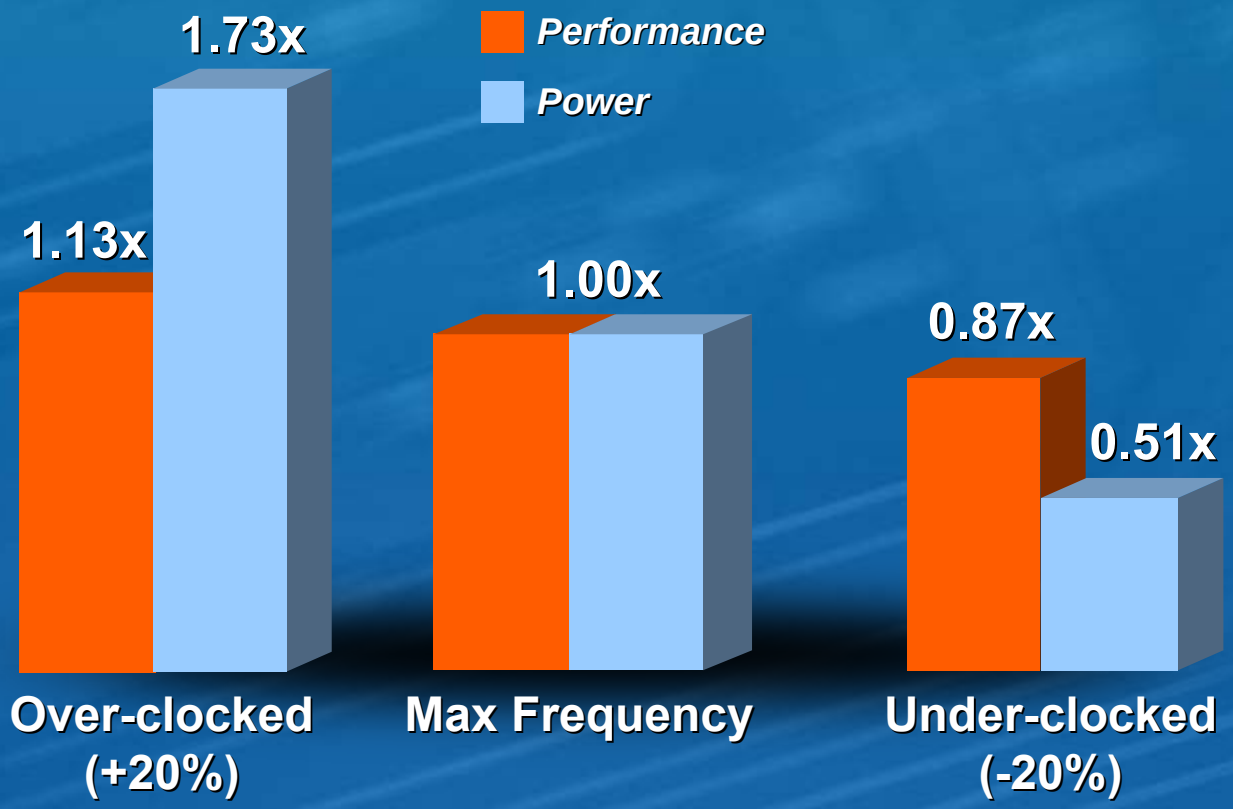

Relative single-core frequency and Vcc

Multi-Core
Energy-Efficient Performance

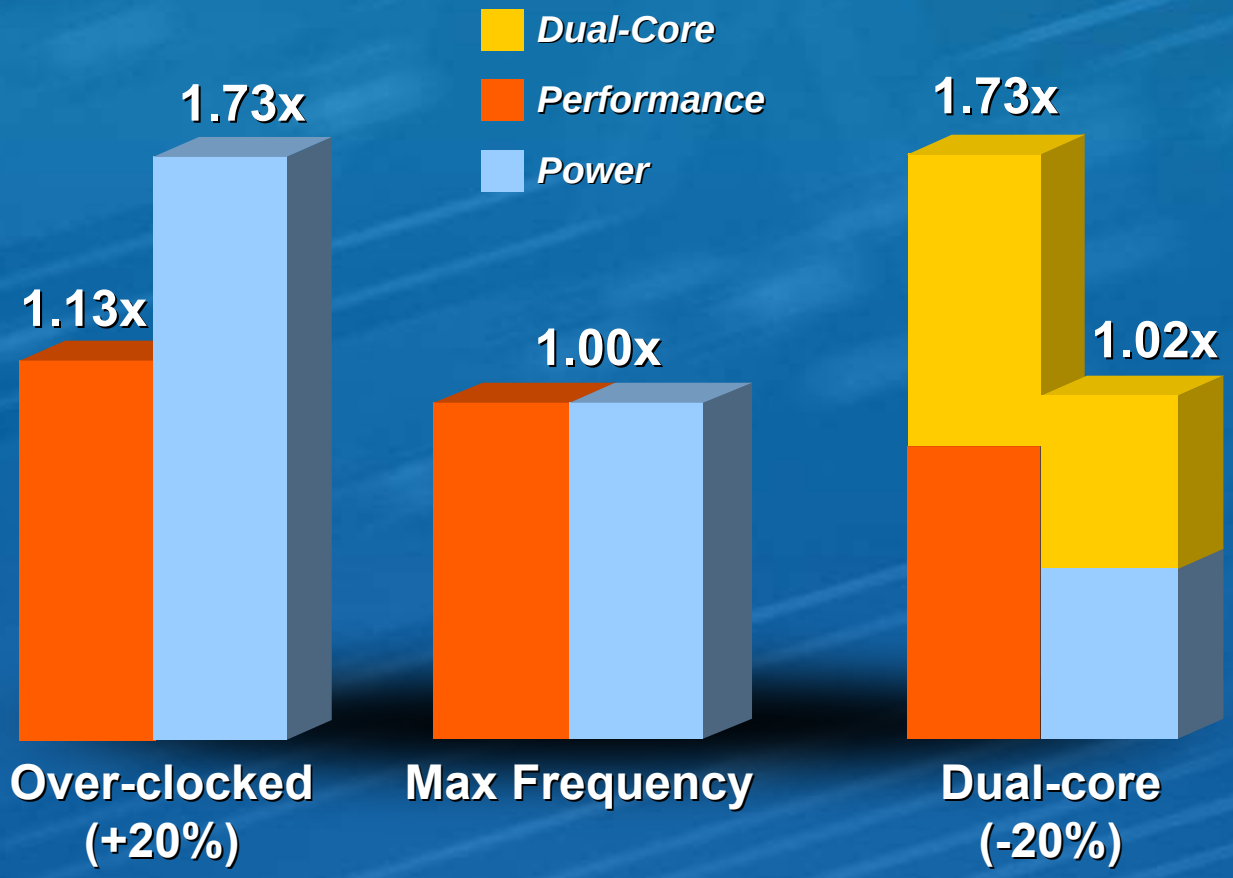

Relative single-core frequency and Vcc

Multi-threaded Cores

All Large Core

Mixed Large and Small Core

All Small Core

Goal: Energy Efficient Petascale with Multi-threaded Cores

Note: the above pictures don't represent any current or future Intel products

intel

Increasing Throughput through Parallelism

Amdahl's Law: Parallel Speedup = 1/(Serial% + (1-Serial%)/N*)

12 Cores

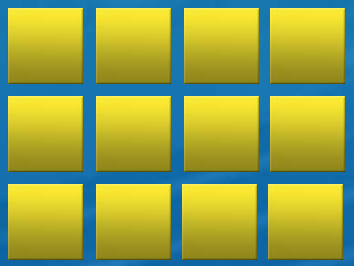

48 Cores

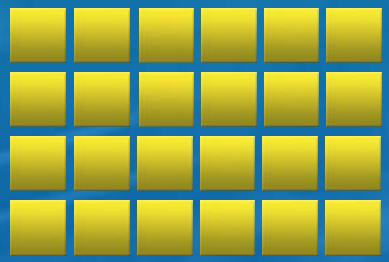

144 Cores

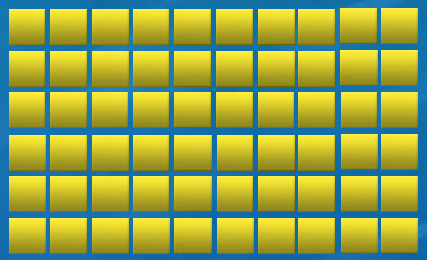

Single Core Performance

Relative Performance

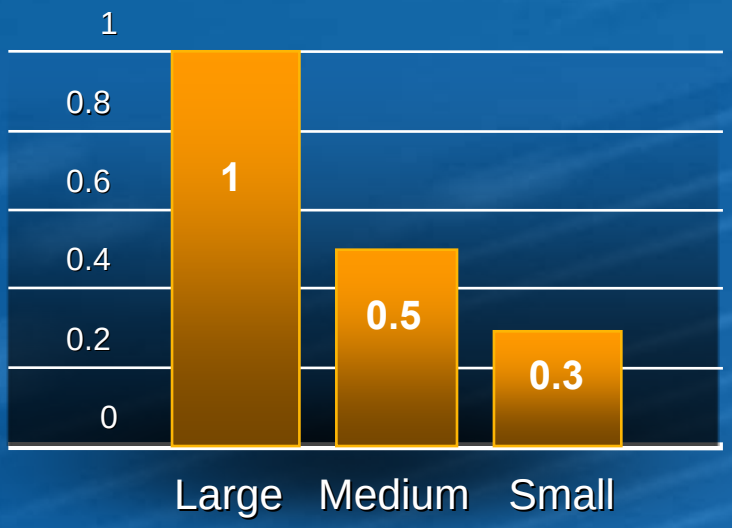

System Performance

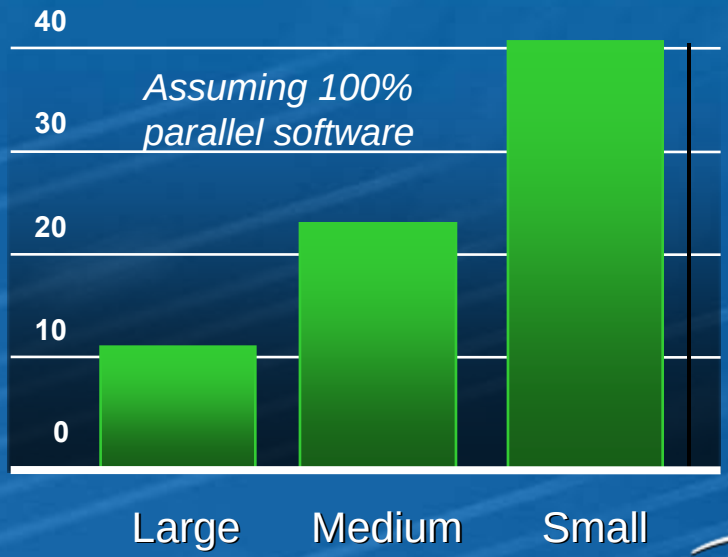

* N = number of cores

Teraflops Research Chip

100 Million Transistors • 80 Tiles • 275mm²

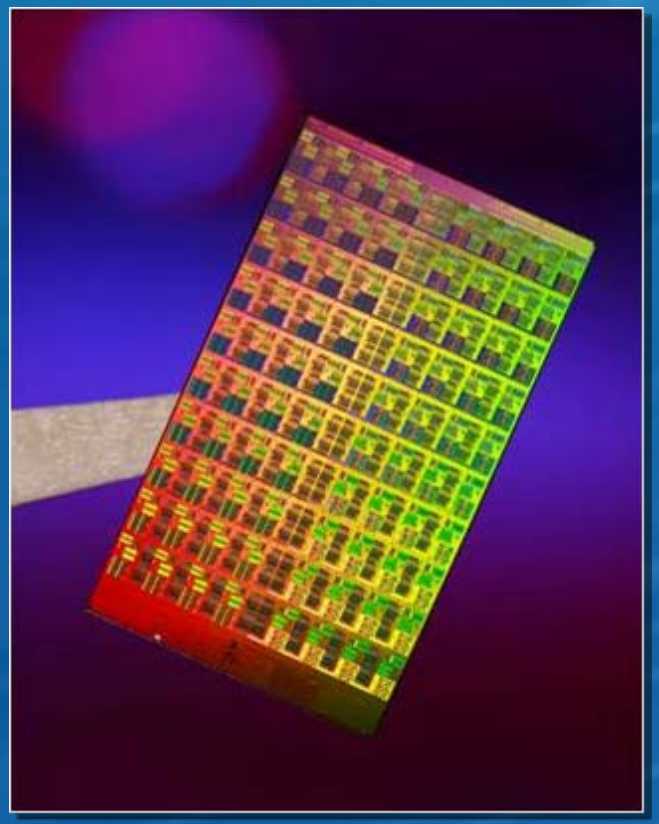

First tera-scale programmable silicon:

- Teraflops performance
- Tile design approach
- On-die mesh network
- Novel clocking
- Power-aware capability
- Supports 3D-memory

Not designed for IA or product

What is Tera-scale?

Teraflops of performance operating on Terabytes of data

Entertainment

Learning & Travel

Personal Media Creation and Management

Health

Performance

TIPS

GIPS

MIPS

KIPS

RMS

3D & Video

Multi-Media

Text

Single-core

Multi-core

Tera-scale

Kilobytes

Megabytes

Gigabytes

Terabytes

Dataset Size

Tera-scale Introduction

- Represents significant Intel transition from “large” cores to 32+ low-power, highly-threaded IA cores per die
- Motivations for a new architecture
  - Enable emerging workloads and new use-models
  - Low Power IA cores provide 4-5X greater performance-power efficiency
  - Scaling beyond the limits of Instruction level parallelism and single-core power
- Tera-scale is *NOT* simply SMP-on-die
  - Will require complete platform and software enabling

Parameter	SMP	Tera-scale	Improvement	Optimizations
Bandwidth	12 GB/s	~1.2 TB/s	~100X	Massive bandwidth between cores
Latency	400 cycles	20 cycles	~20X	Ultra-fast synchronization

Intel Tera-scale Research

100+ Research Projects Worldwide

Microprocessor

Examples:

Scalable memory
Multi-core architectures
Specialized cores
Scalable fabrics
Energy efficient circuits

Platform

Examples:

3D Stacked Memory
Cache Hierarchy
Virtualization/Partitioning
Scaleable OS's
I/O & Networking

Programming

Examples:

Speculative Multithreading
Transactional memory
Workload analysis
Compilers & Libraries
Tools

ACCELERATE TRANSITION TO PARALLEL PROGRAMMING

intel Software Products

www.intel.com/software/products

University Outreach
Intel ® Press
Intel® Software College

intel

Expected Tera-scale Insights

- Power management of many cores
  - Research prototype enables extensive studies on fabric and core power consumption & management
- Physical implementation challenges of high speed fabric and multiple cores
- 3D stacked silicon technology
- On-chip bandwidth and latency impact

Tiled Design & Mesh Network

Repeated Tile Method:

- Compute + router
- Modular, scalable
- Small design teams
- Short design cycle

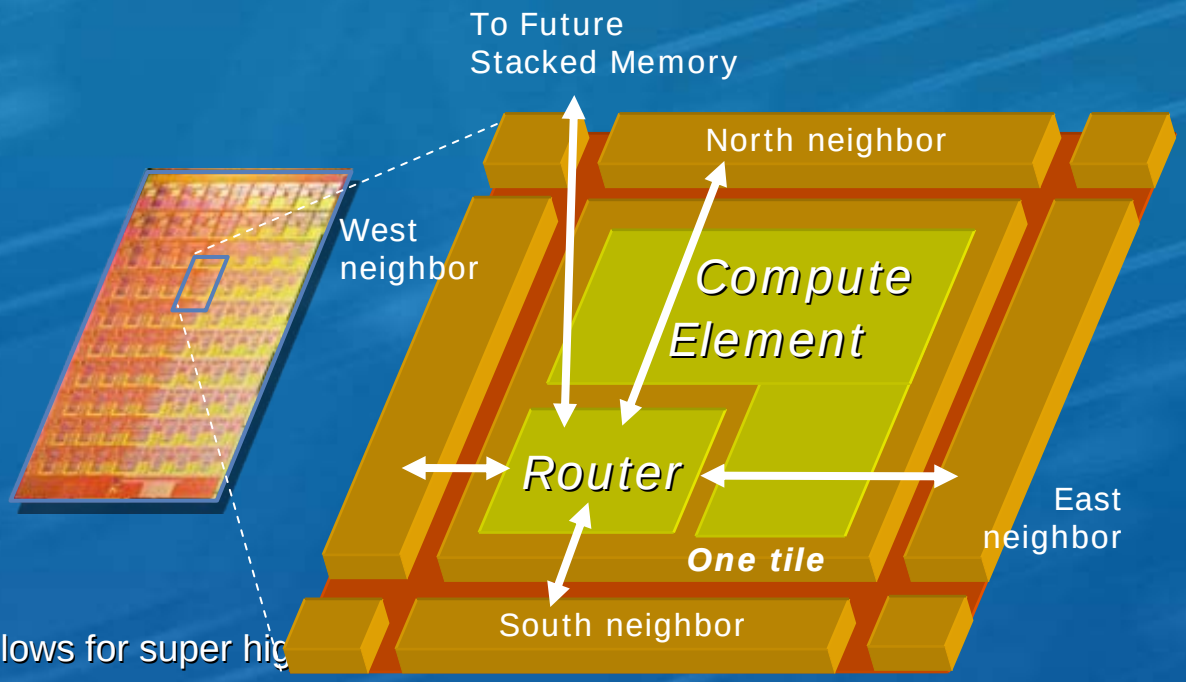

Mesh Interconnect:

- "Network-on-a-Chip"
  - Cores networked in a grid allows for super hig
    communications in and between cores
- 5-port, 80GB/s* routers
- Low latency (1.25ns*)
- Future: connect IA/or and special purpose cores

* When operating at a nominal speed of 4GHz

Fine Grain Power Management

- Novel, modular clocking scheme saves power over global clock
- New instructions to make any core sleep or wake as apps demand
- Chip Voltage & freq. control (0.7-1.3V, 0-5.8GHz)

Dynamic sleep

STANDBY:
- Memory retains data
- **50%** less power/tile

FULL SLEEP:
- Memories fully off
- **80%** less power/tile

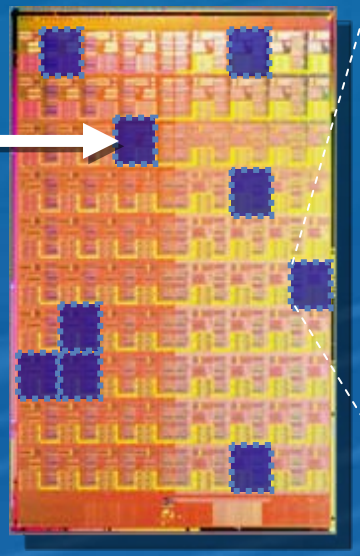

21 sleep regions per tile *(not all shown)*

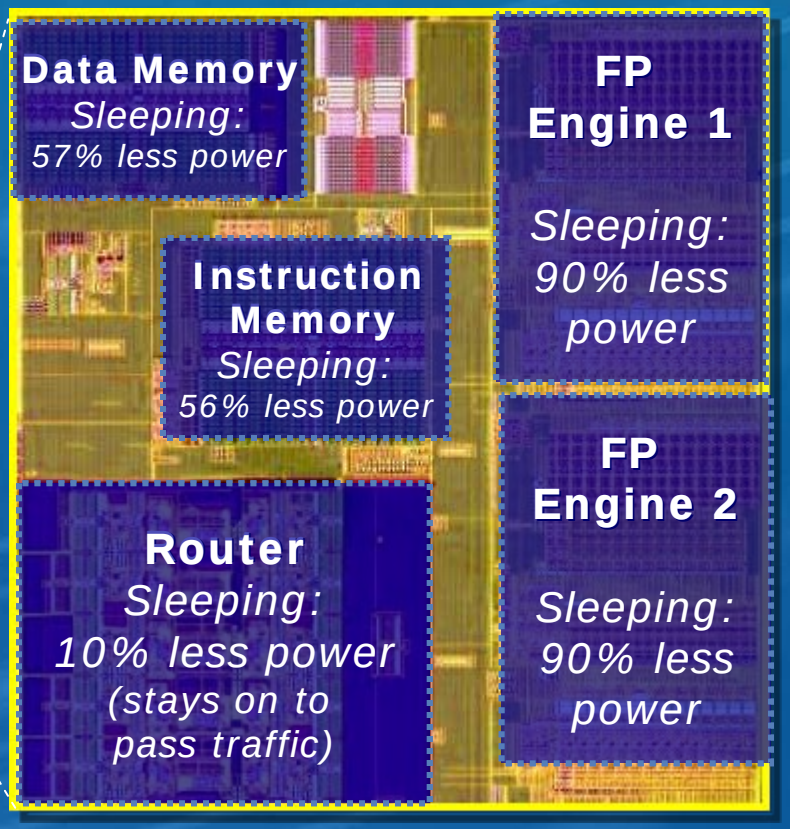

Industry leading energy-efficiency of 16 Gigaflops / Watt

intel

Research Data Summary

Frequency	Voltage	Power	Bisection Bandwidth	Performance
3.16 GHz	**0.95 V**	**62W**	**1.62 Terabits/s**	**1.01 Teraflops**
5.1 GHz	1.2 V	175W	2.61 Terabits/s	1.63 Teraflops
5.7 GHz	1.35 V	265W	2.92 Terabits/s	1.81 Teraflops

1.01 Teraflops
62 Watts

More than the Cores

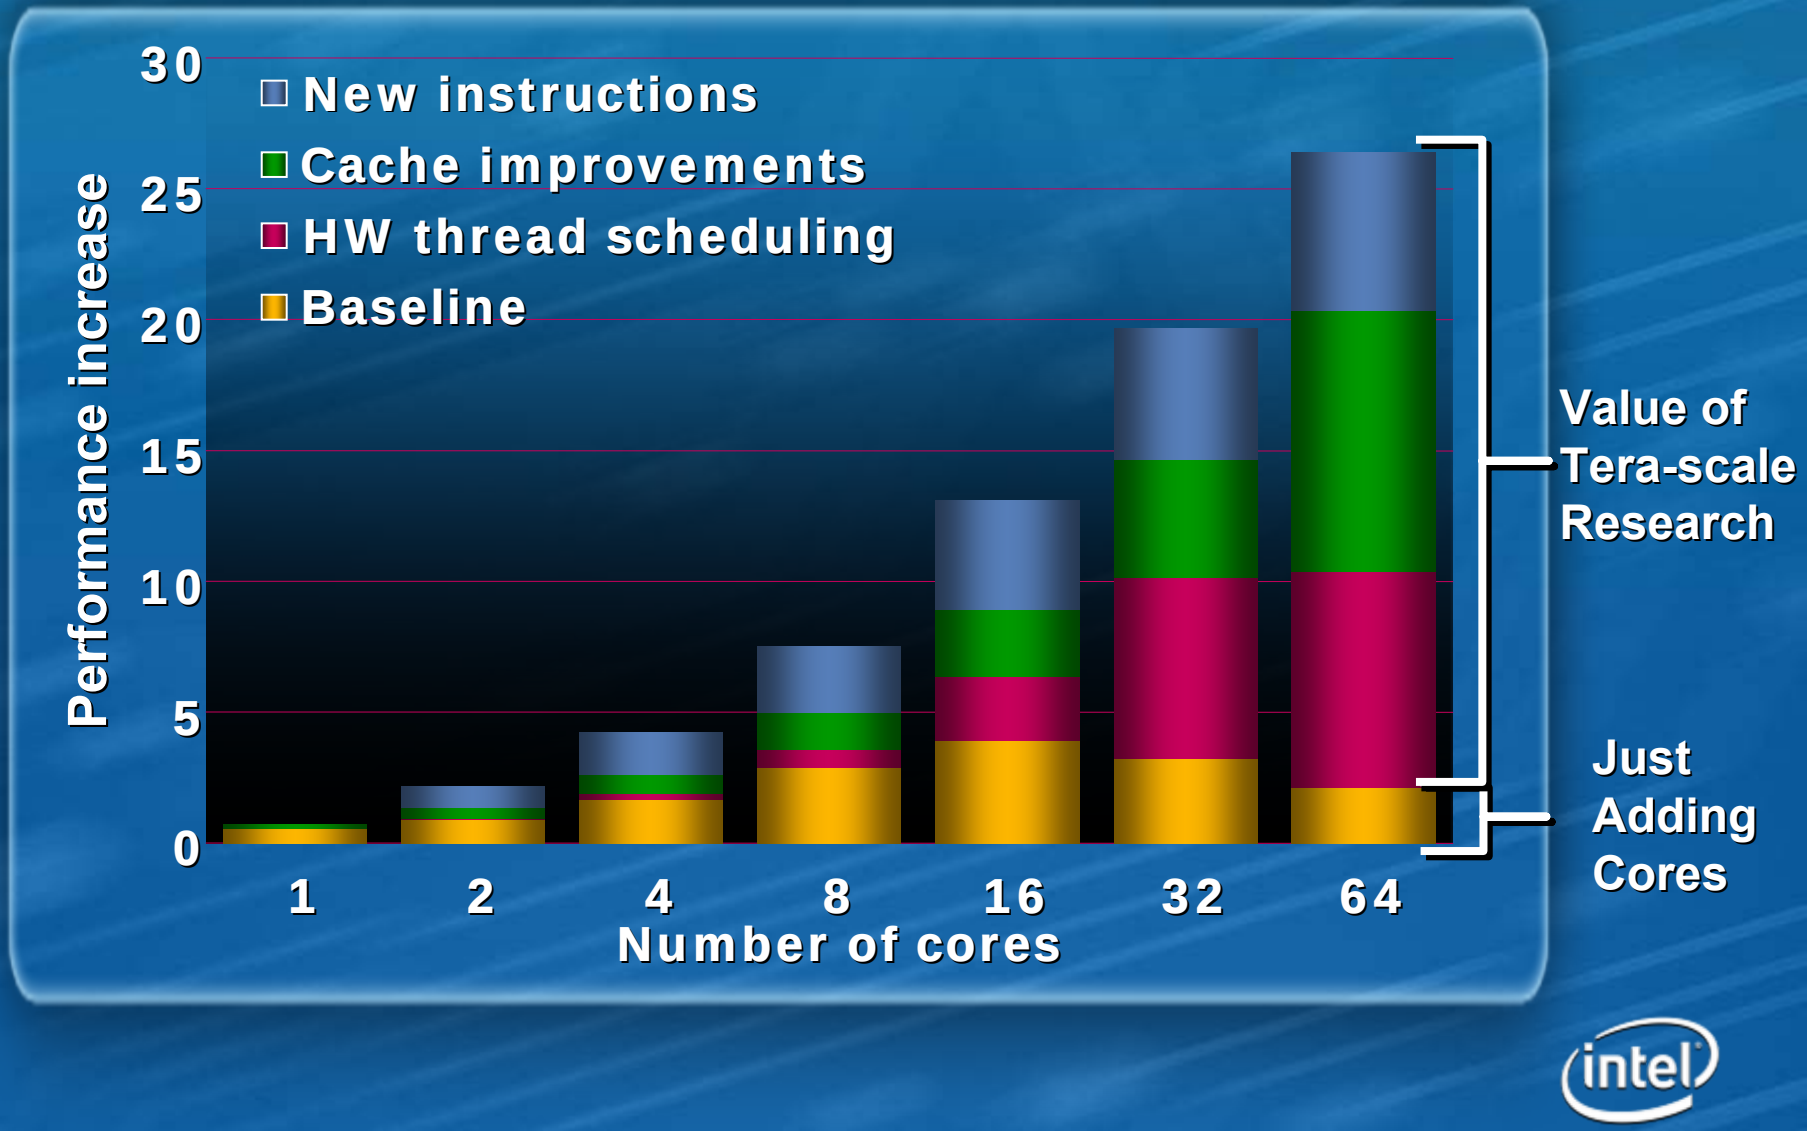

Processor Performance

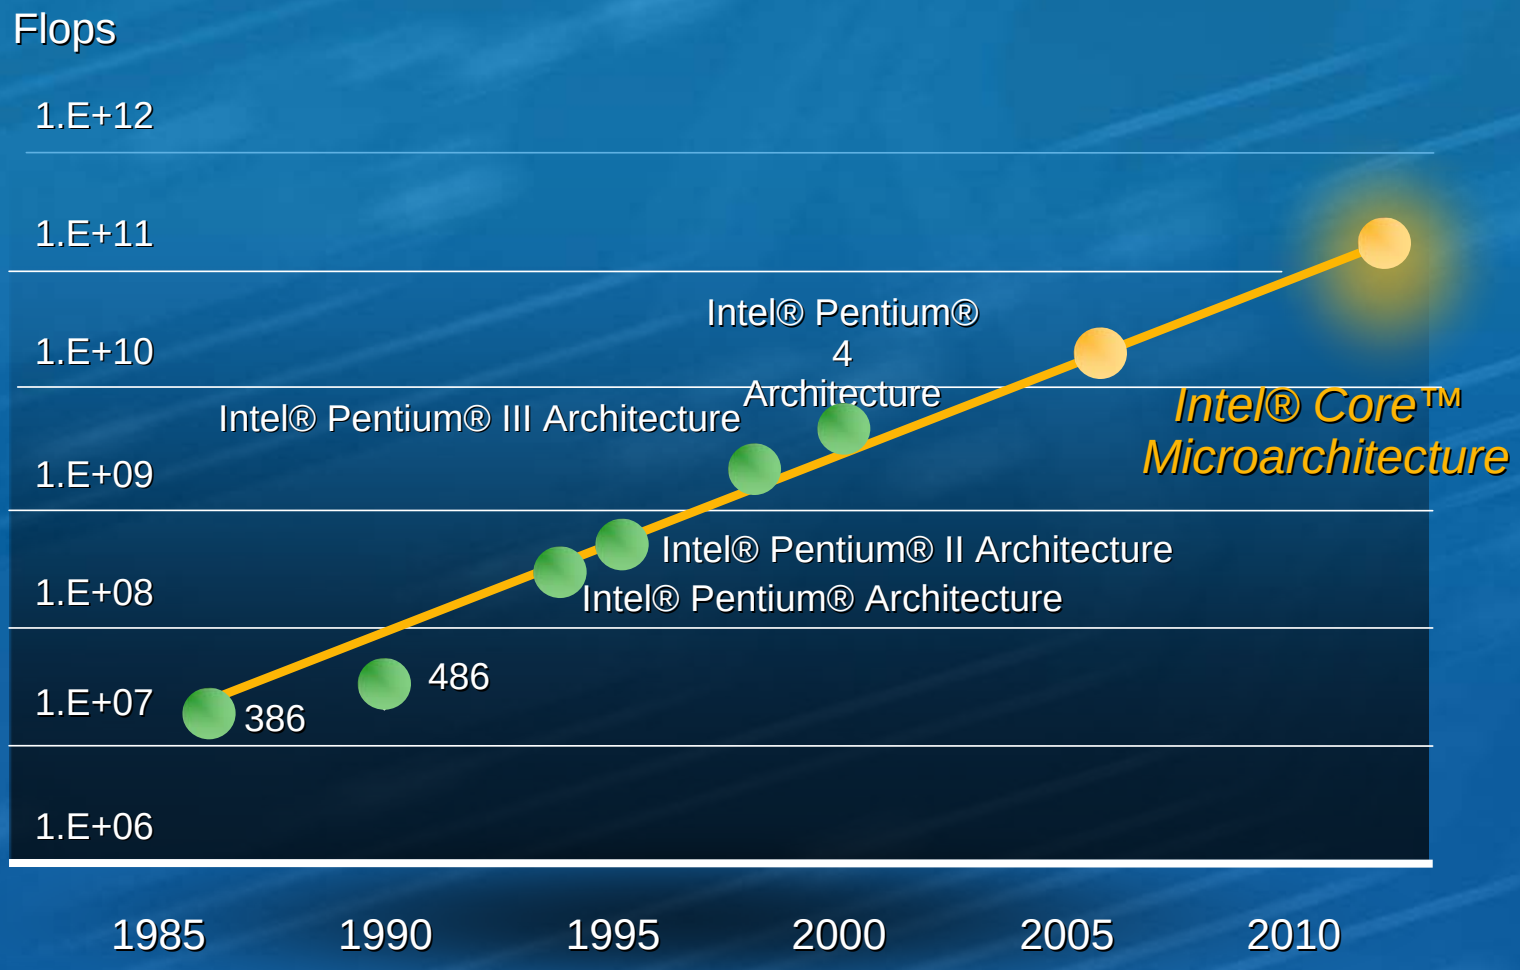

*Reaching Petascale with ~100,000 Processors in 2010**

Assuming approx. 100Glops processors

* Petascale assumes 10's of PF Peak Performance and 1PF Sustained Performance on HPC Applications.

Increasing Processor Performance

Through Multi-threaded Cores

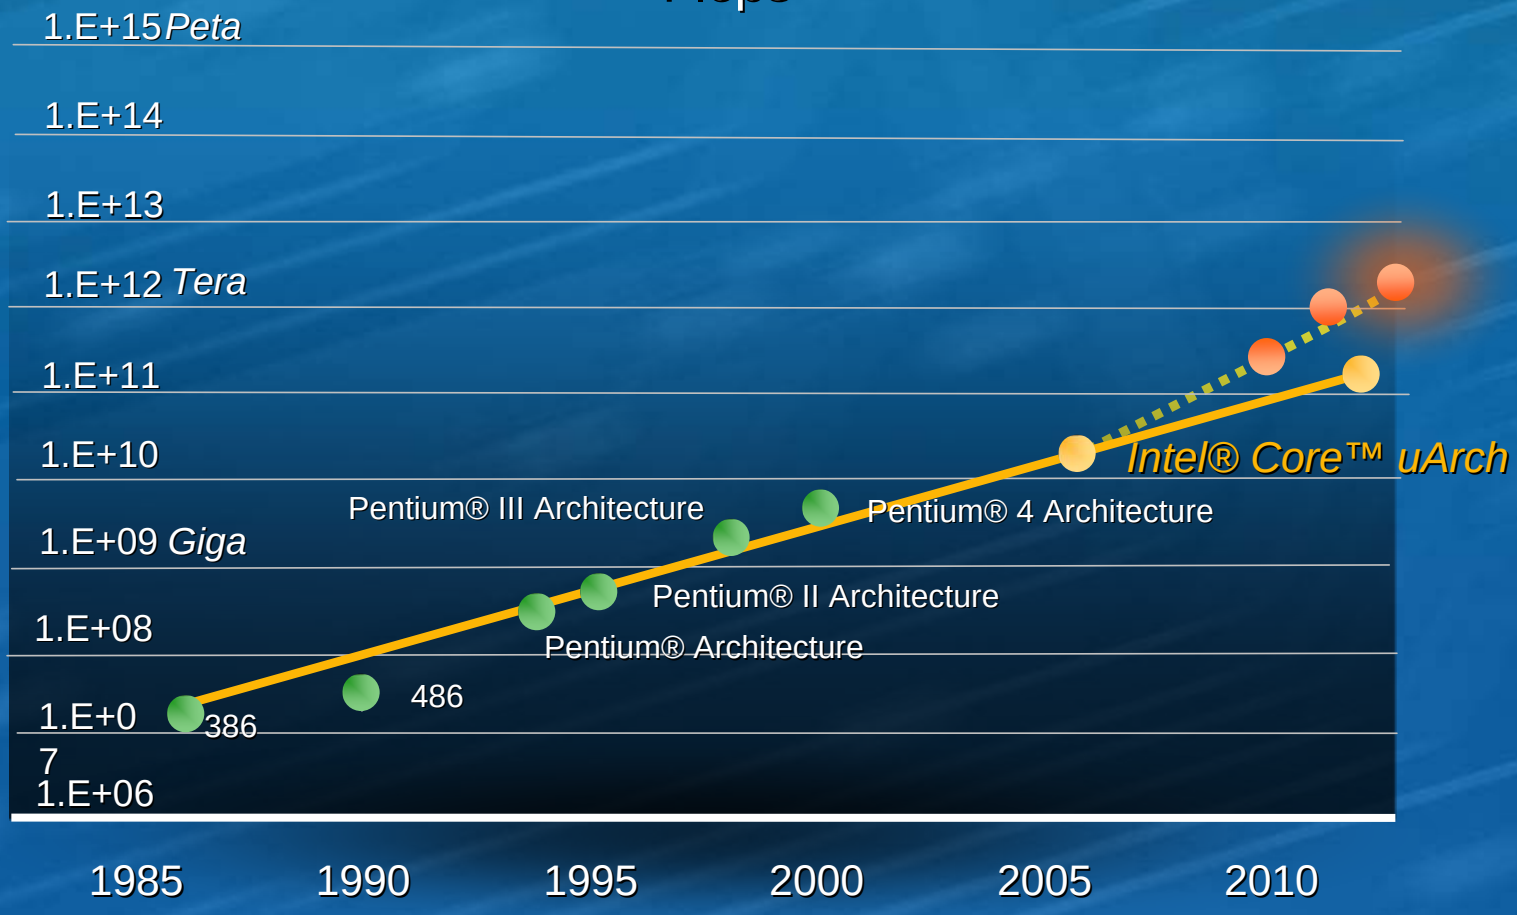

Reaching Petascale with ~5,000 Processors

Increasing I/O Signaling Rate to Fill the Gap

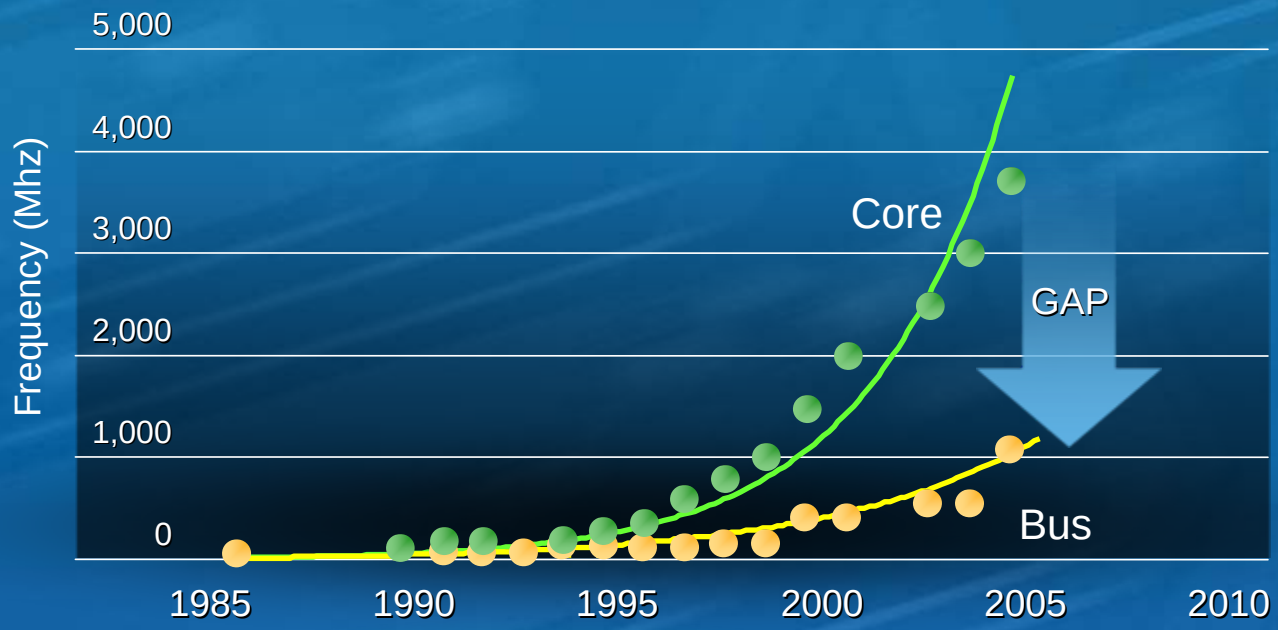

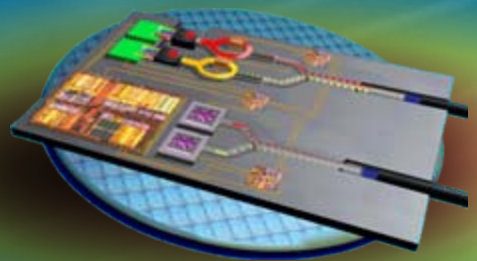

Silicon Photonics

Source: Intel

Increasing Memory Bandwidth
to Keep Pace

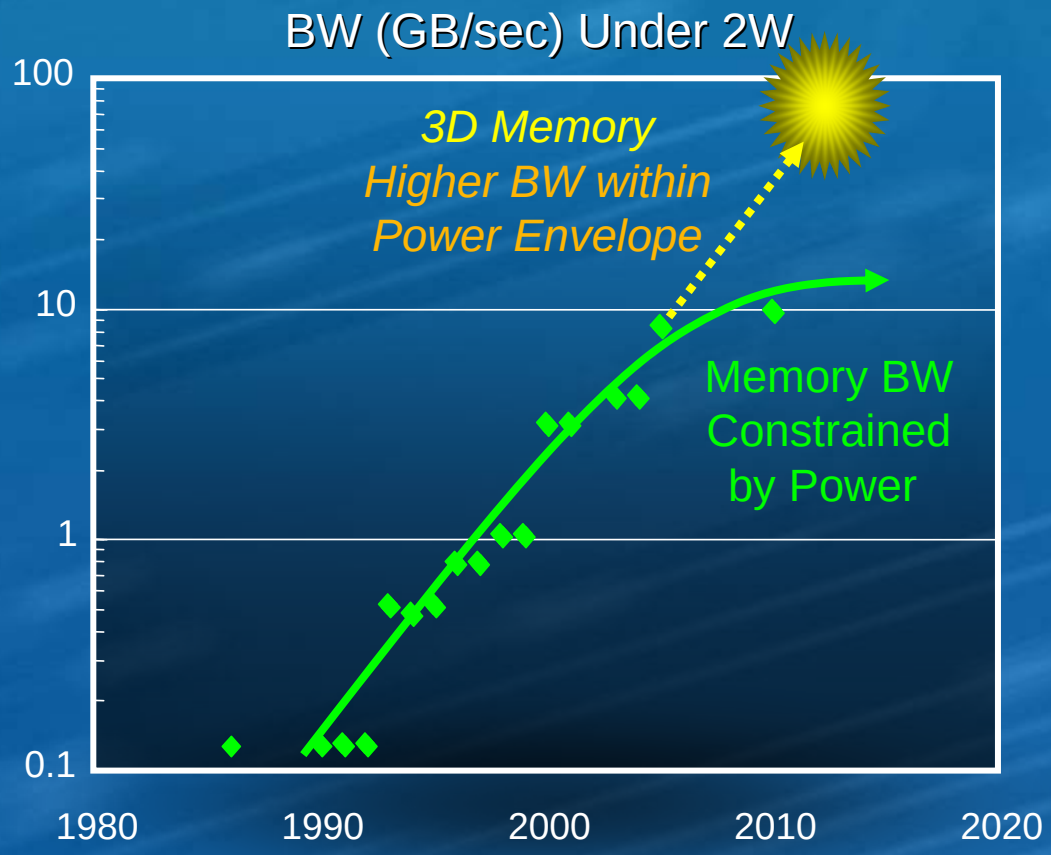

3D Memory Stacking

Power and IO Signals Go Through DRAM to CPU

Thin DRAM Die

Through DRAM Vias

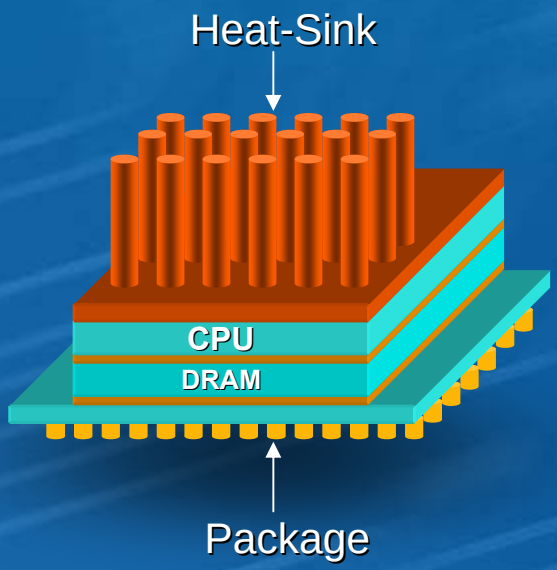

Source: Intel

What can we expect!

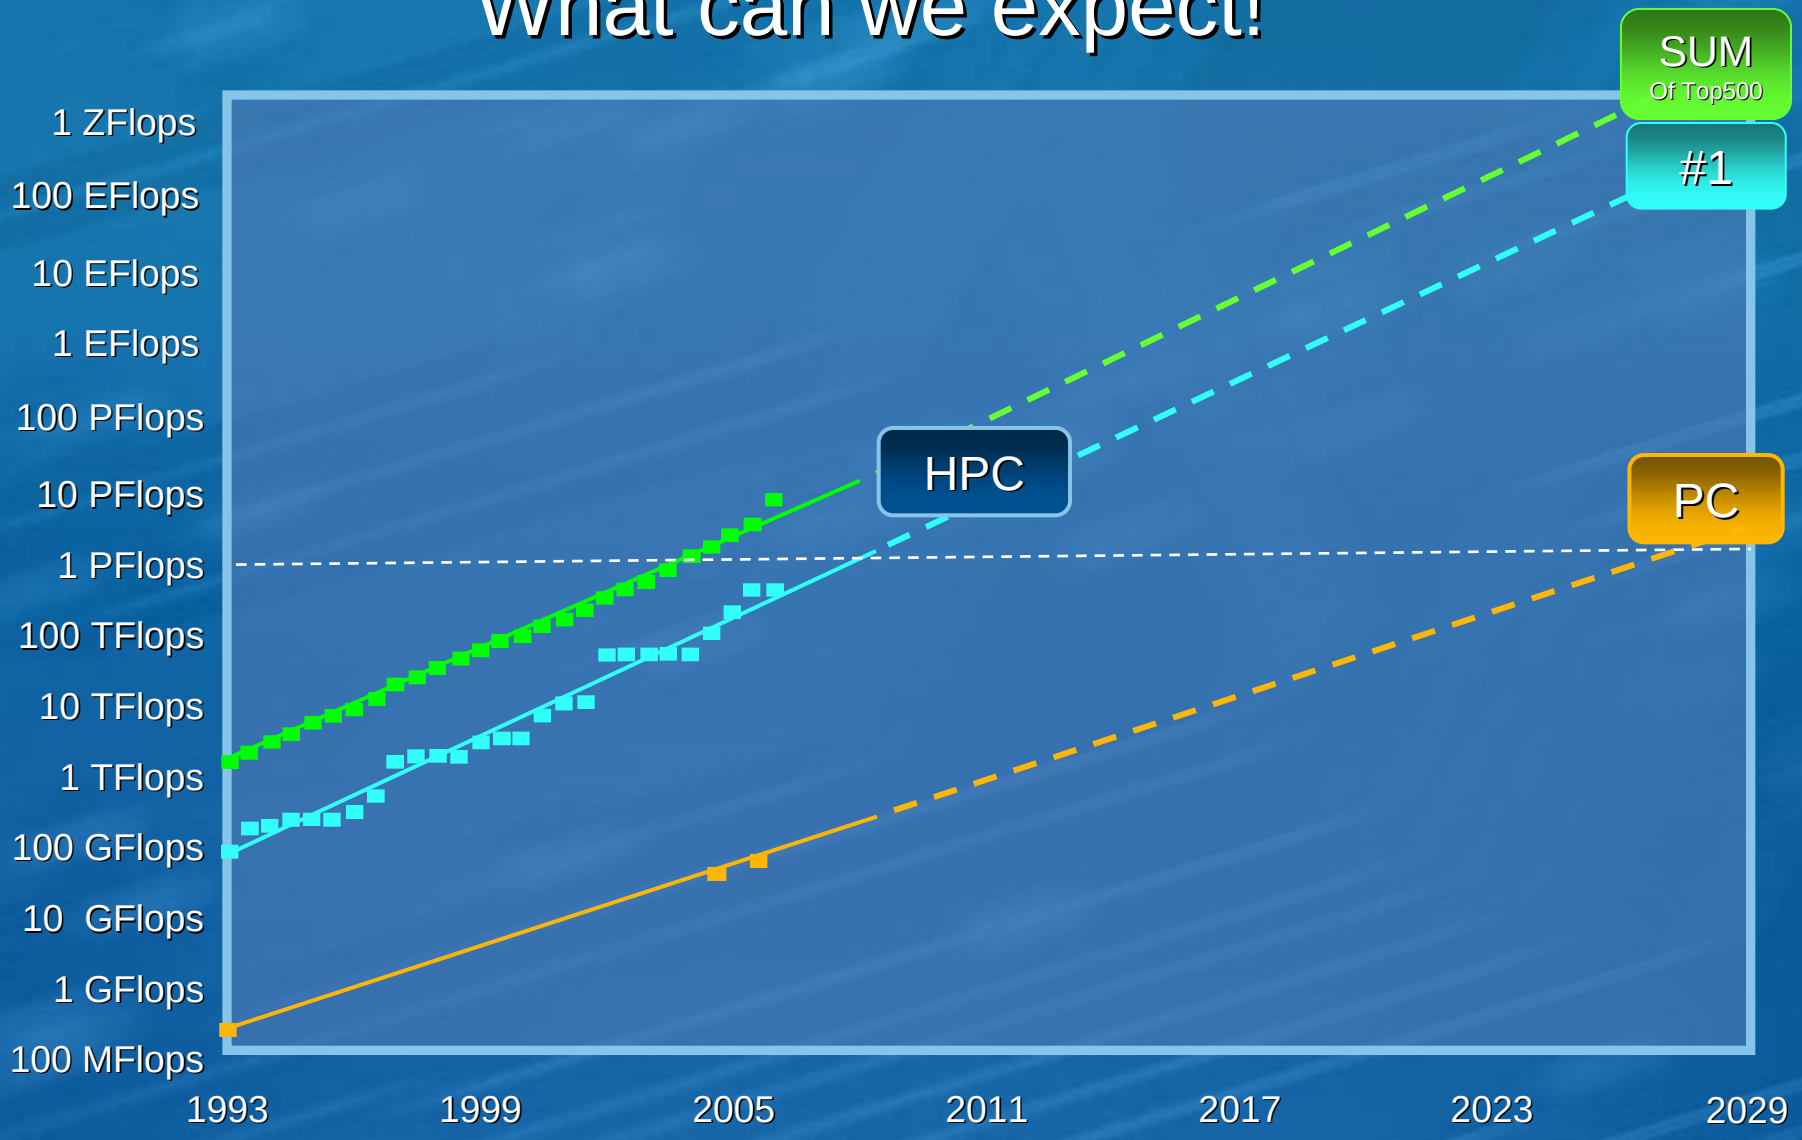

Source: HPC - www.top500.org, June 2006, Intel


intel
Leap ahead™
HPC @ Intel