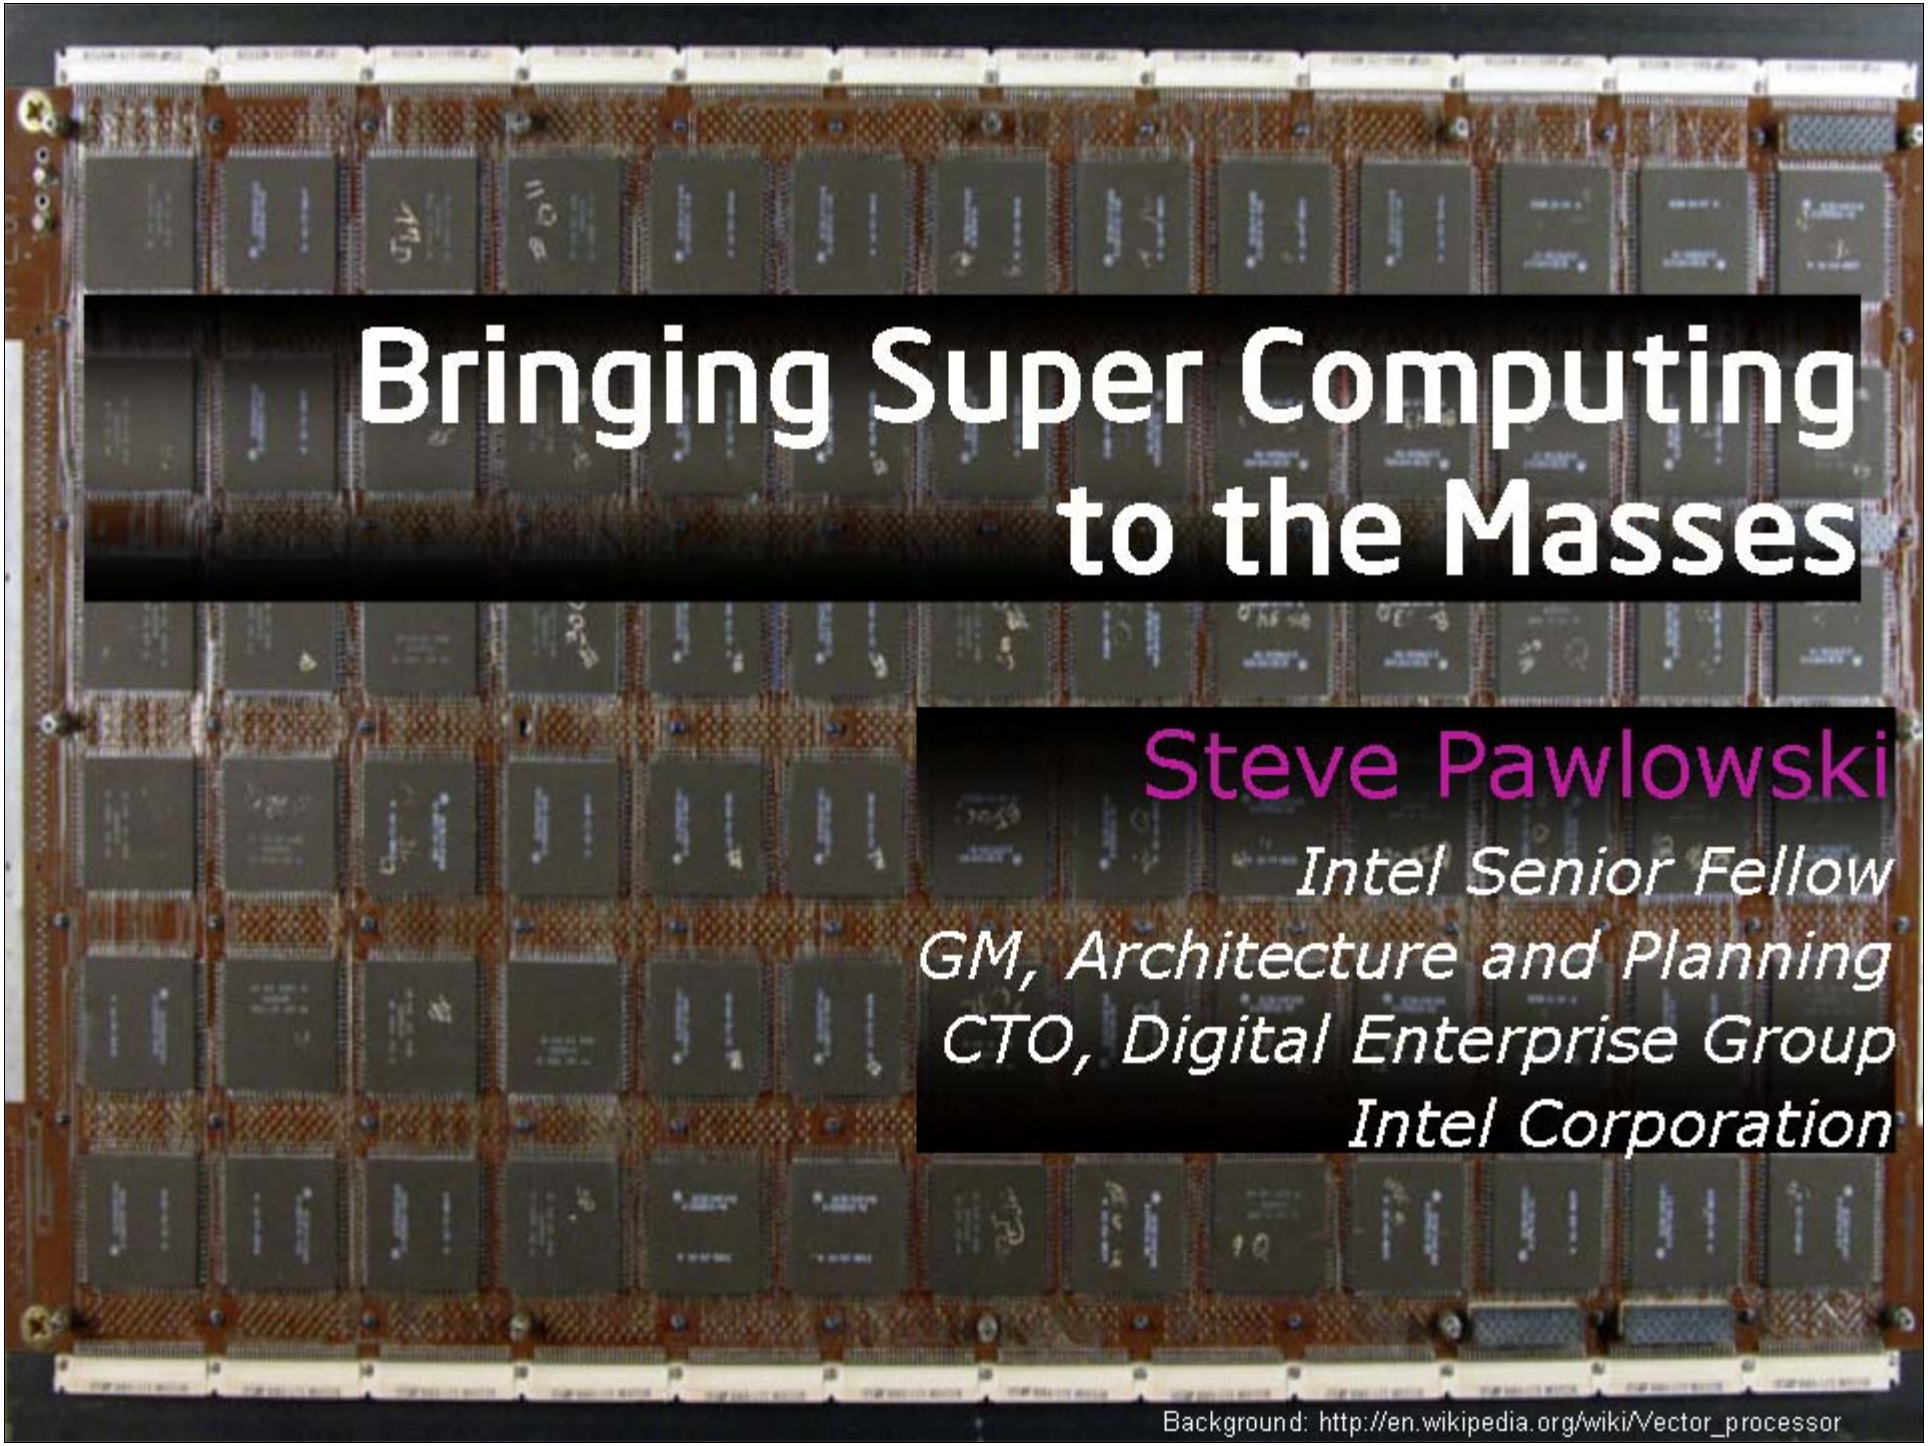

# Bringing Super Computing to the Masses


Steve Pawlowski

*Intel Senior Fellow*

*GM, Architecture and Planning*

*CTO, Digital Enterprise Group*

*Intel Corporation*


Background: http://en.wikipedia.org/wiki/Vector_processor

# Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control or safety systems, or in nuclear facility applications.

Intel products may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel may make changes to dates, specifications, product descriptions, and plans referenced in this document at any time, without notice.

This document may contain information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Wireless connectivity and some features may require you to purchase additional software, services or external hardware.

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

Intel, the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.





# IDC Prediction of HPC Market Growth

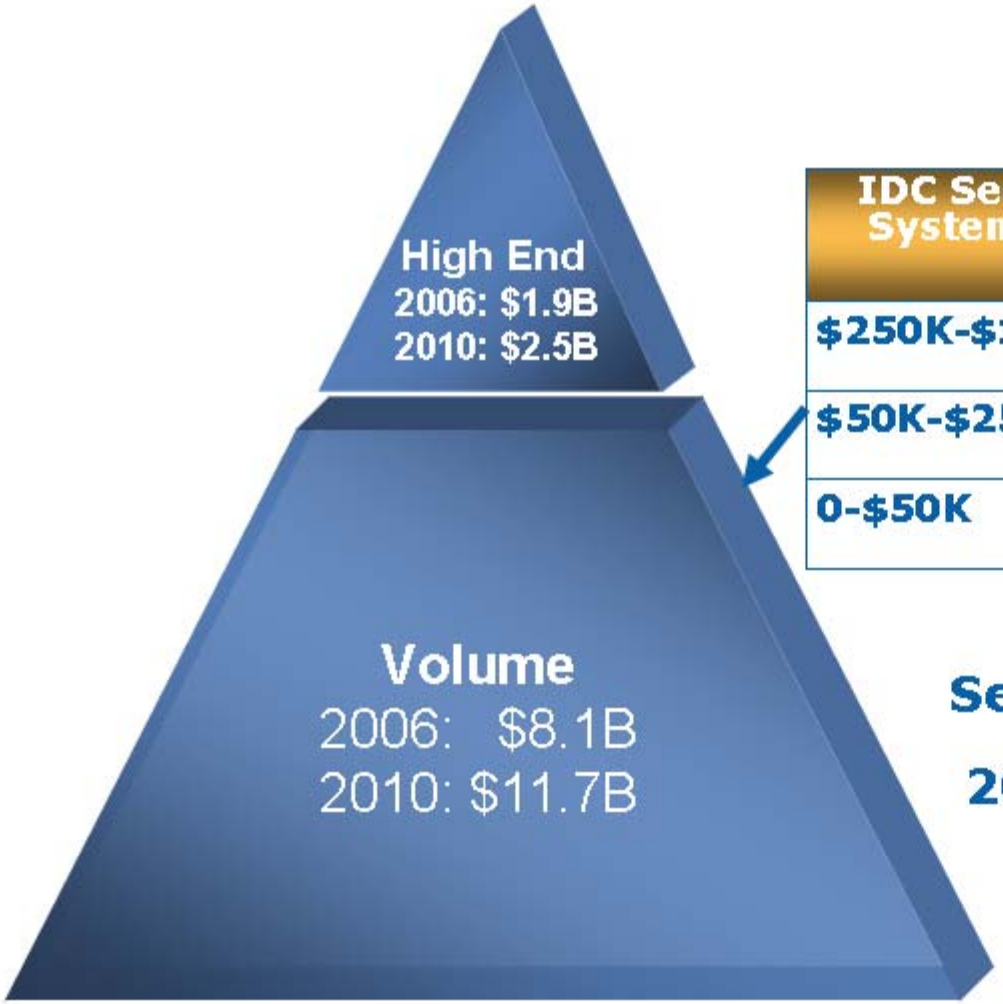


| IDC Segment System Size | 2006 | 2010 | CAGR |
|---|---|---|---|
| $250K-$1M | $2.1B | $3.4B | 11.8% |
| $50K-$250K | $3.6B | $4.9B | 10.7% |
| 0-$50K | $2.5B | $3.4B | 9.6% |

**Segment Total 2006: $10.0B**

**2010 forecast: $14.3B**

- Source: IDC Worldwide Technical Computing Systems Revenue 2006-2010 #201733 May 2006.
- Source: IDC Technical Server QView February 2007 for 2006 actual data

# Yesterday's Supercomputer is Today's PC

CDC 6600 – First successful Supercomputer 9MFlops

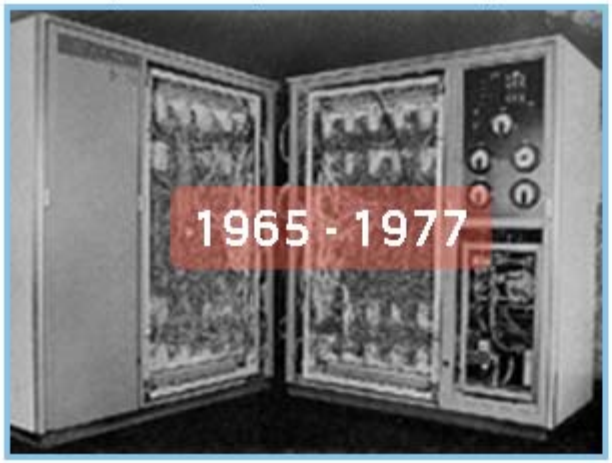



ASCI Red
(world fastest on top500 till 2000)
First Teraflop Computer,
9298 Intel Pentium® II Xeon Processors

Yesterday

2005

XBOX 360
1 TFlops

1TFlops with just a game Console

Today



*Bring supercomputing to the masses: Why do we care?*

intel


VIRTUAL WORLDS
THE RISE OF THE 3D INTERNET
Justin Rattner
Intel Developer Forum
October 2007

Accelerating Technology and Capacity

# GROWTH CURVES


Source: metaverseraodmap.org

# The Transition to the 3D INTERNET

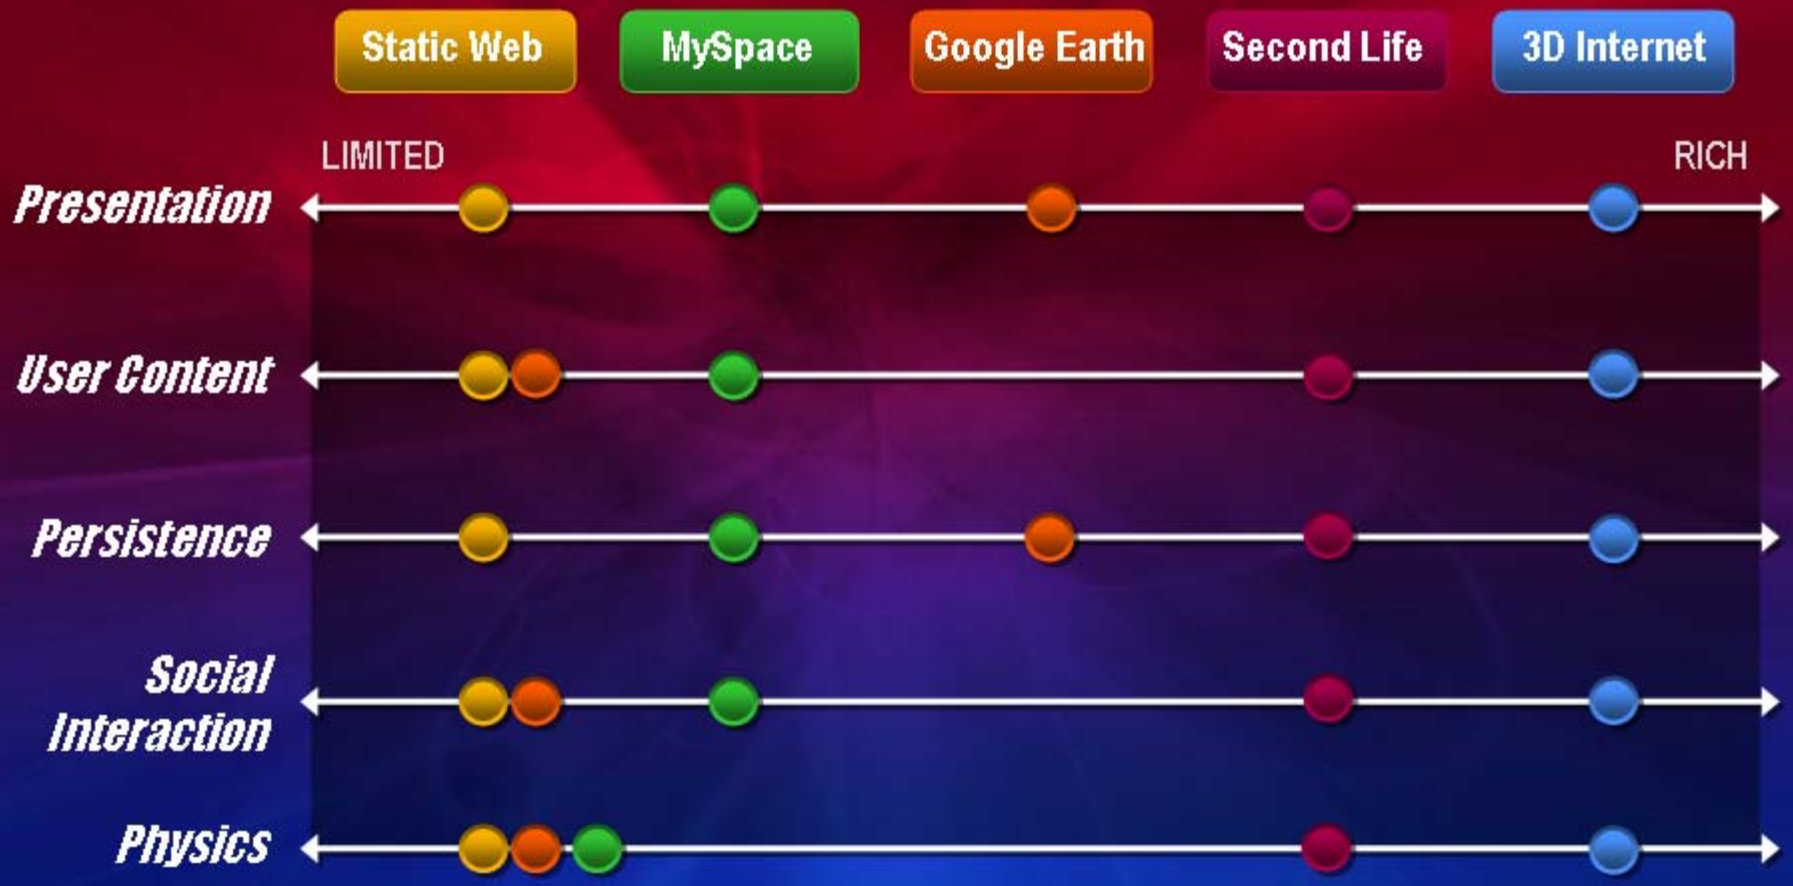


Virtual Worlds Provide The Richest Space In Which To Understand The 3D Internet

# Flavors of VIRTUAL WORLDS

WoW

**MULTI-PLAYER GAMES**

Croquet

Second Life

**METAVERSES**

Fully Immersive
3D Virtual Spaces

# Flavors of VIRTUAL WORLDS


**PARAVERSE**
Real + Virtual Mashups

**MACHINIMA**
Interactive Cinema

**Virtual Surgery**

**Mscape**

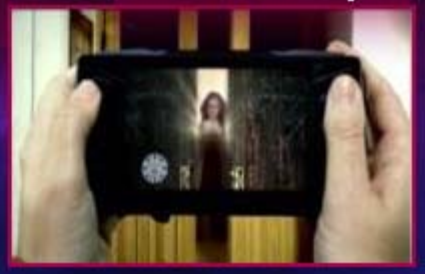

**globalkids.org**

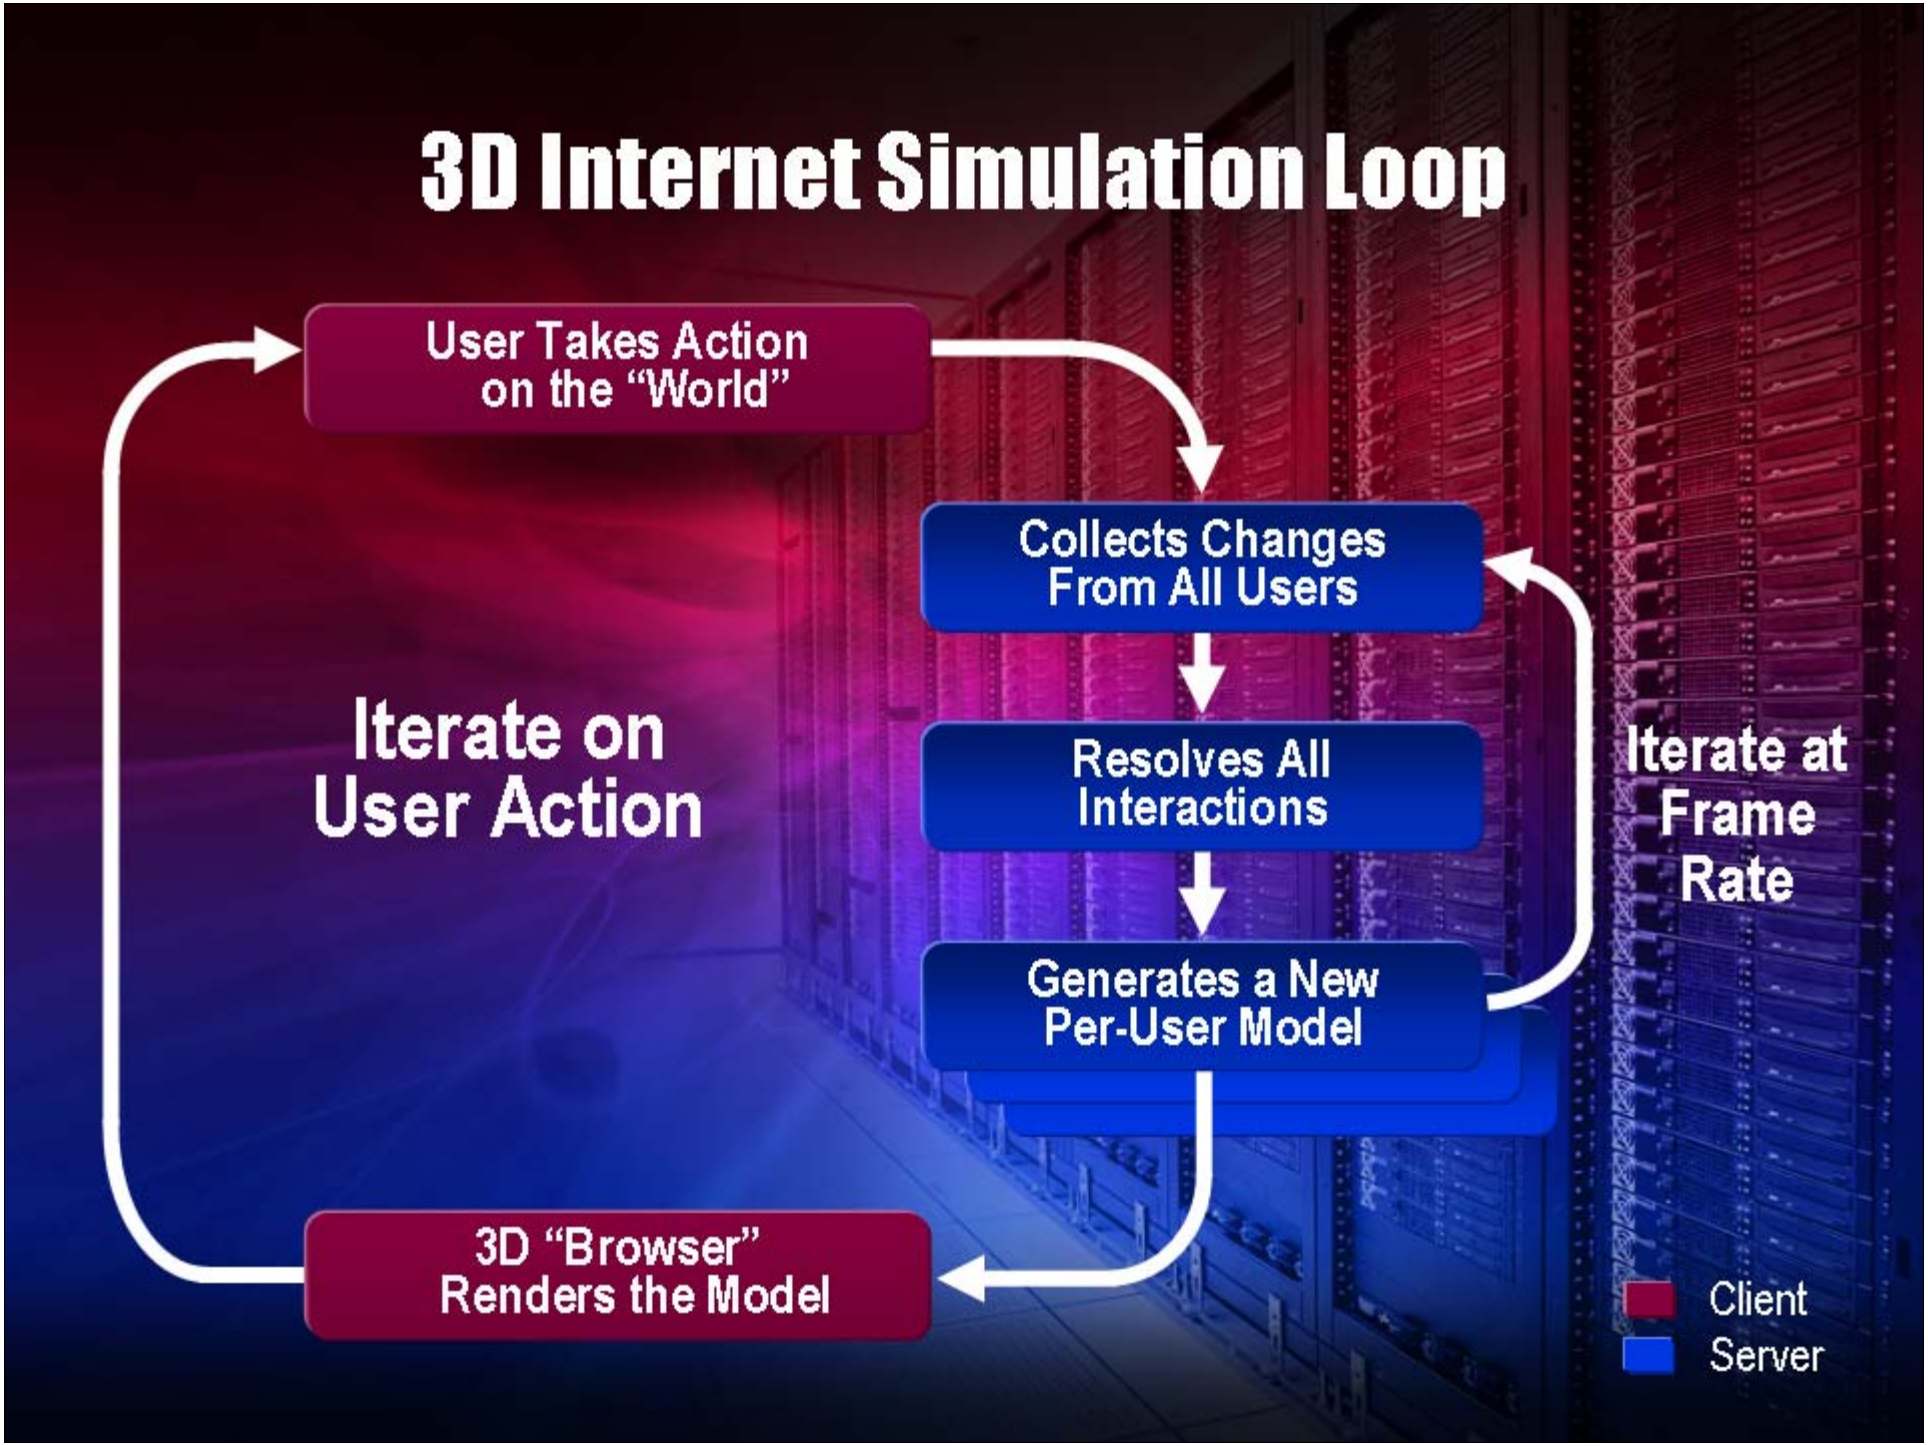
3D Internet Simulation Loop
User Takes Action on the "World"
Collects Changes From All Users
Resolves All Interactions
Generates a New Per-User Model
3D "Browser" Renders the Model
Iterate on User Action
Iterate at Frame Rate
Client
Server

# DEMANDS on Server Performance

**75%+**
**Time Spent in Compute-Intensive Components**

**Unlike 2D Internet Server Side:**

- Floating Point, Matrix Intensive
- Quadratic Growth

**Server Performs an Order of Magnitude More Work**

| TYPE | SOFTWARE | MAXIMUM CLIENT/ SERVER |
|---|---|---|
| MMORPGS | Eve Online | 34420 |
| | WoW | 2500 |
| VWs | Second Life | 160 |

**VW's Require 10-100x More Computation Per Client Than MMORPGS**

# DEMANDS on Client Performance

**65%+**
**Time spent in Compute-Intensive Components**

**Unlike 2D Internet Client Side:**

- Floating Point, Matrix Intensive
- Quadratic Growth in Compute

**Client CPU/GPU Performance**
*2D Internet Vs. 3D Internet*

| APPLICATION | % CPU UTILIZATION | % GPU UTILIZATION (nVidia G80) |
|---|---|---|
| 2D Websites | 20 | 0-1 |
| Google Maps | 60 | 3-5 |
| Google Earth | 50 | 10-15 |
| Second Life | 70 | 35-75 |

**VW's Need At Least 3x CPU, 20x GPU Processing Compared To 2D**

# DEMANDS on Network Performance

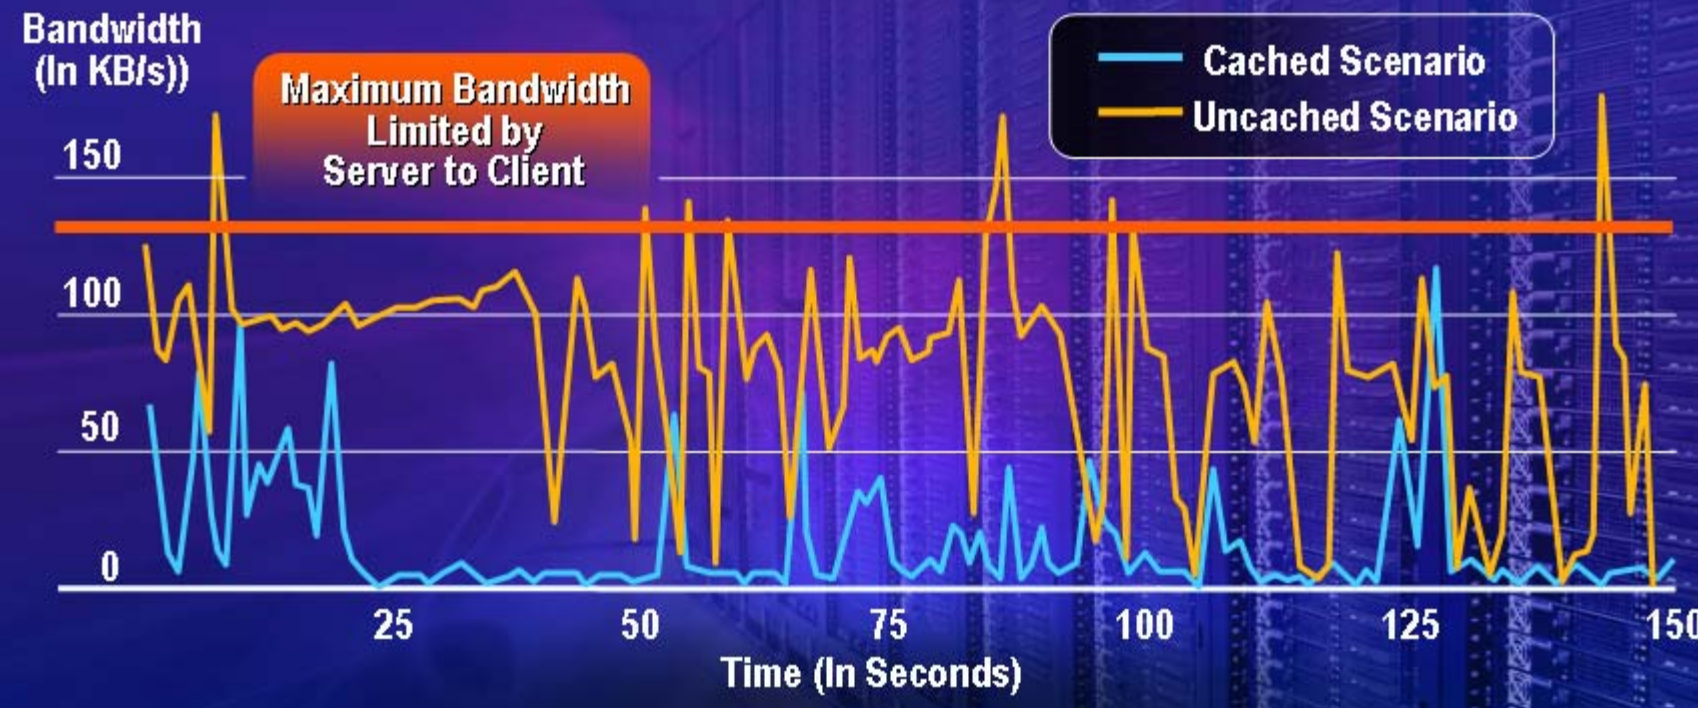


## VW's Need At Least 100x Bandwidth Compared To MMORPGS

# Cinematic Quality Realism*

**Reflection**

**Nature****

**Shadow**

**www.wallis-eck.de *www.q4rt.de

# Progress in Ray Tracing

## IDF 2004

50 Intel® Xeon™ Processors

4 Frames per Second
640x480

## Games Convention 2007

Yorkfield
(45nm Quadcore)

~90 frames per Second
768x768

## Fall IDF 2007

Dual-X5365
(total: 8 cores)

~90 frames per Second
**1280x720**

# More Engaging

Human-Computer Interface

## Haptics

## Facial Tracking

## 3D Mouse

## Open GL Printers

## Screen Technologies

# High Performance Computing - Serving People from Grand Scientific Problems Today to Our Daily Lives Tomorrow




*"Don't be encumbered by past history. Go off and do something wonderful." – Robert Noyce, Intel, 1970*

# Exploding Medical Imaging Data Processing

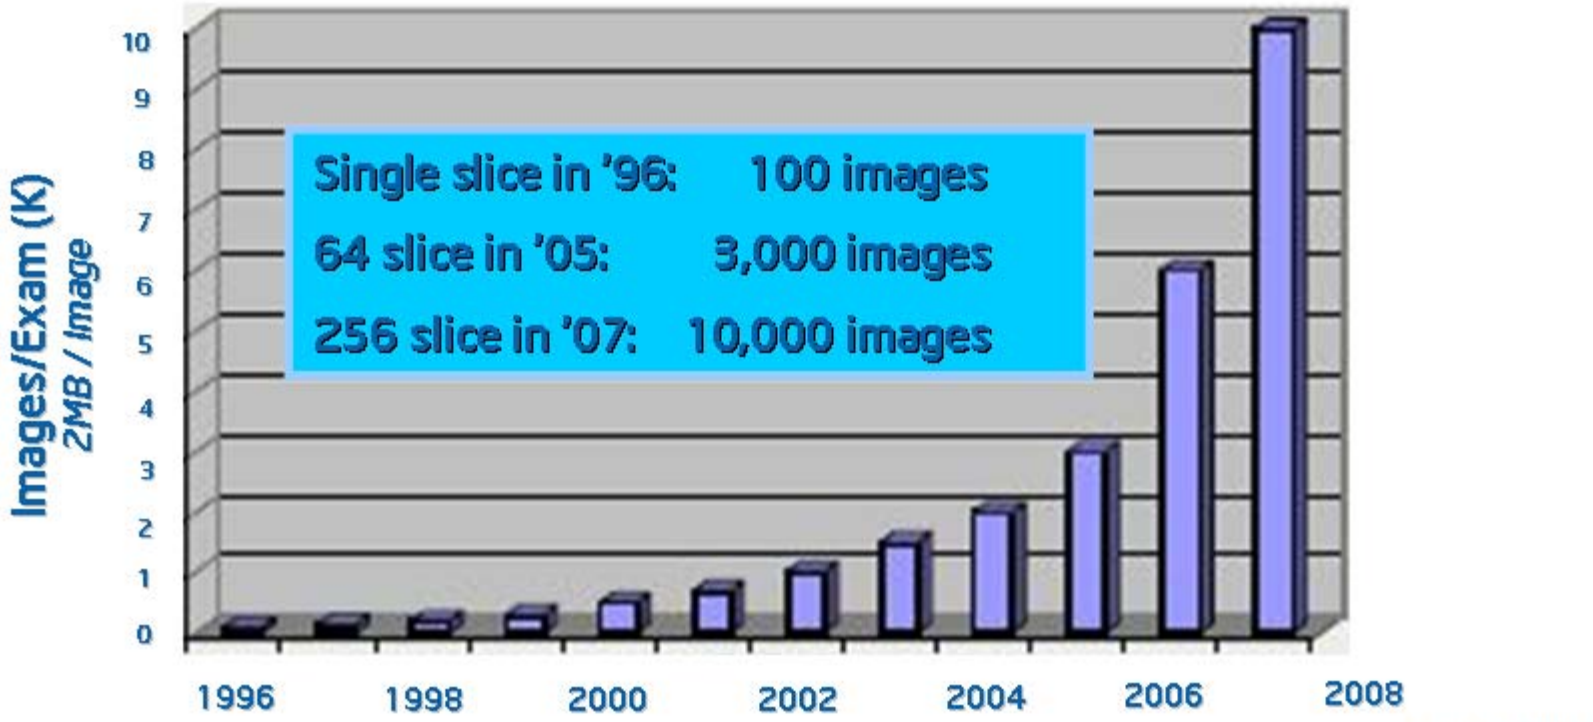


Source: Intel

- 3-D renderings of the images
- Computer aided diagnostic algorithms
- Fusions of images from different modalities
  - MRI, CT, PET, and SPECT
- Real-time applications are appearing

*Full Body CT – 256 slice/10,000 images: a 20GB file*

# Global Healthcare will Reach the Masses

## *with Ever Increasing Accuracy*

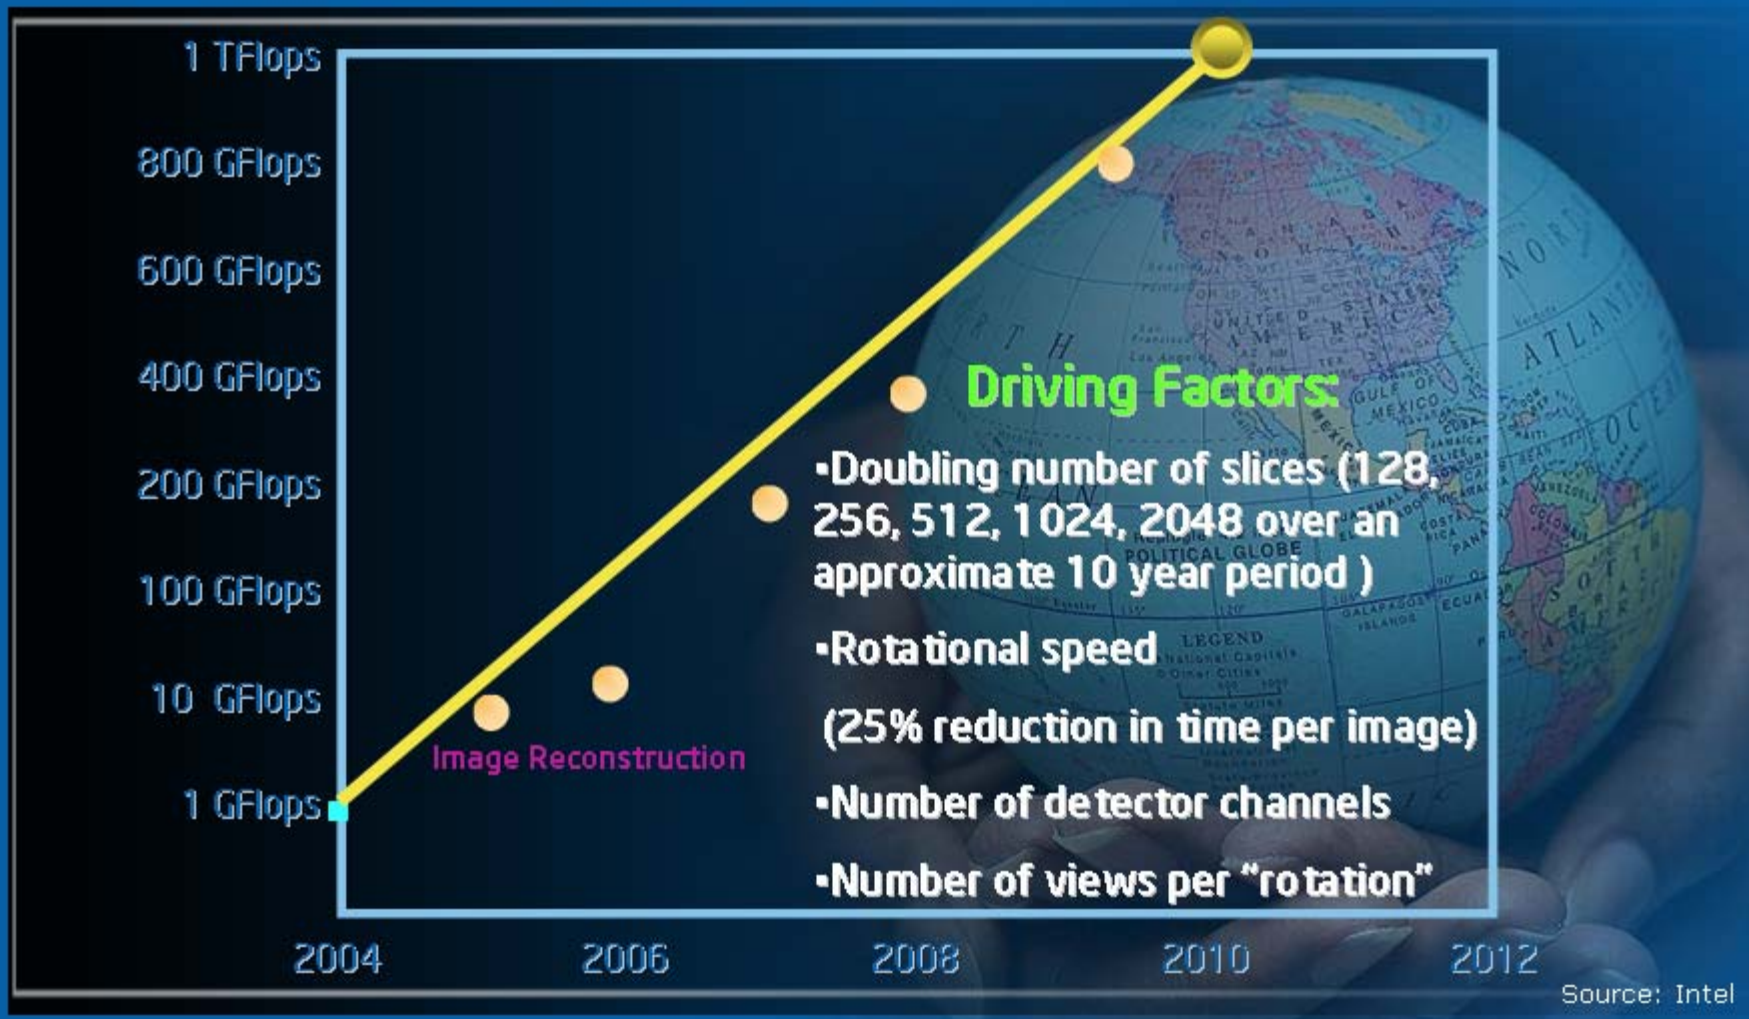


*Medical Imaging: Needs 1 TFlops by 2010*

# Data changes the game in Client Computing


[Kasparov vs. Deep Blue]

Rule-based system exceeds human performance in a structured, deterministic domain

↓

Statistical inference (not rules)

- 100s of TB of training data
- Racks of computation




*doubling training data → constant accuracy improvement*

# The Top500 Illustration

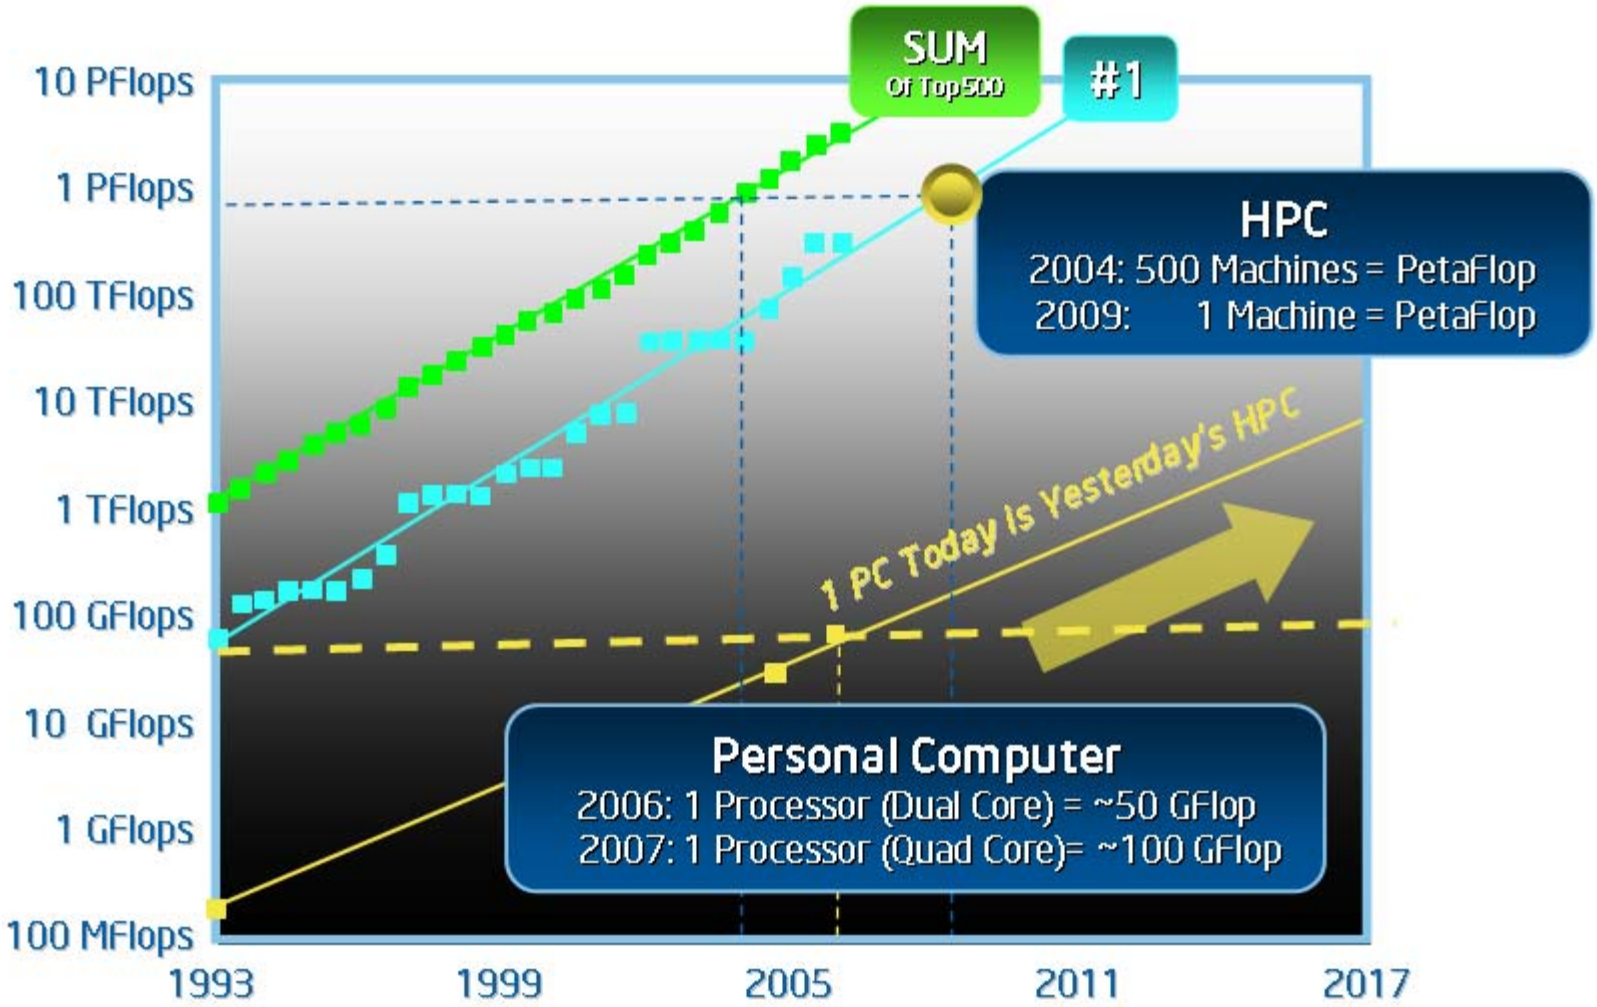




Source: HPC - www.top500.org, June 2006; PC - Intel

# Insatiable Future Super Computing Demand


Source: Dr. Steve Chen, "The Growing HPC Momentum in China," June 30th, 2006, Dresden, Germany

# Applications Demand Petascale

- Many believe that global warming will produce more extremes weather (drought/flooding).
- Current models are too coarse for predicting climate change at the national level.
- To predict regional climate change:
  - Community climate model resolution goal is *10 km*
  - Currently can simulate *50 days/day* on Red Storm at 10 km using NCAR/SNL SEAM
  - Typical climate simulation is for *100 yrs.*

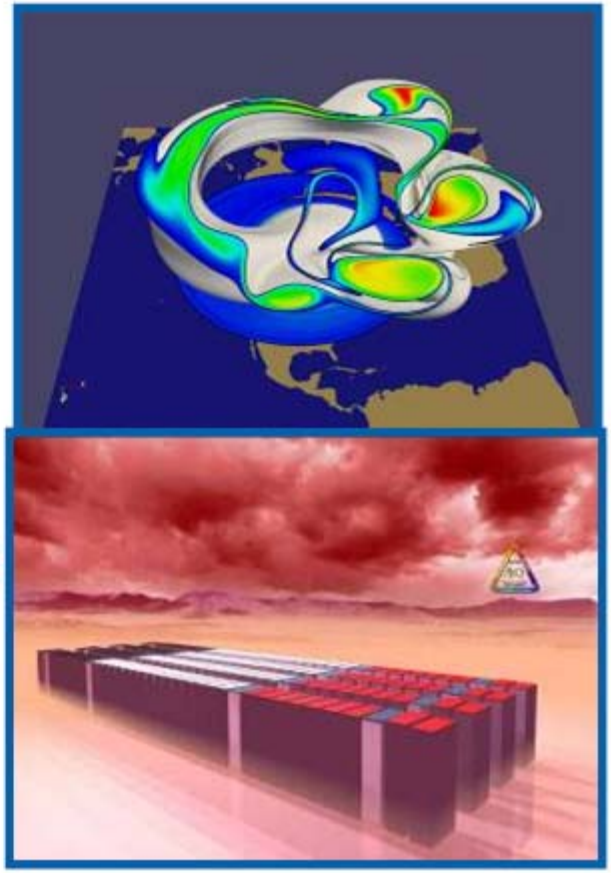

*To simulate 100 year climate:*
*1.6 PFlops in 6 month of computing; 40 PFlops per week*

# Computing that is Impossible Today

100 EFlops
1 ZFlops
10 EFlops
1 EFlops
100 PFlops
10 PFlops
1 PFlops
100 TFlops
10 TFlops
1 TFlops

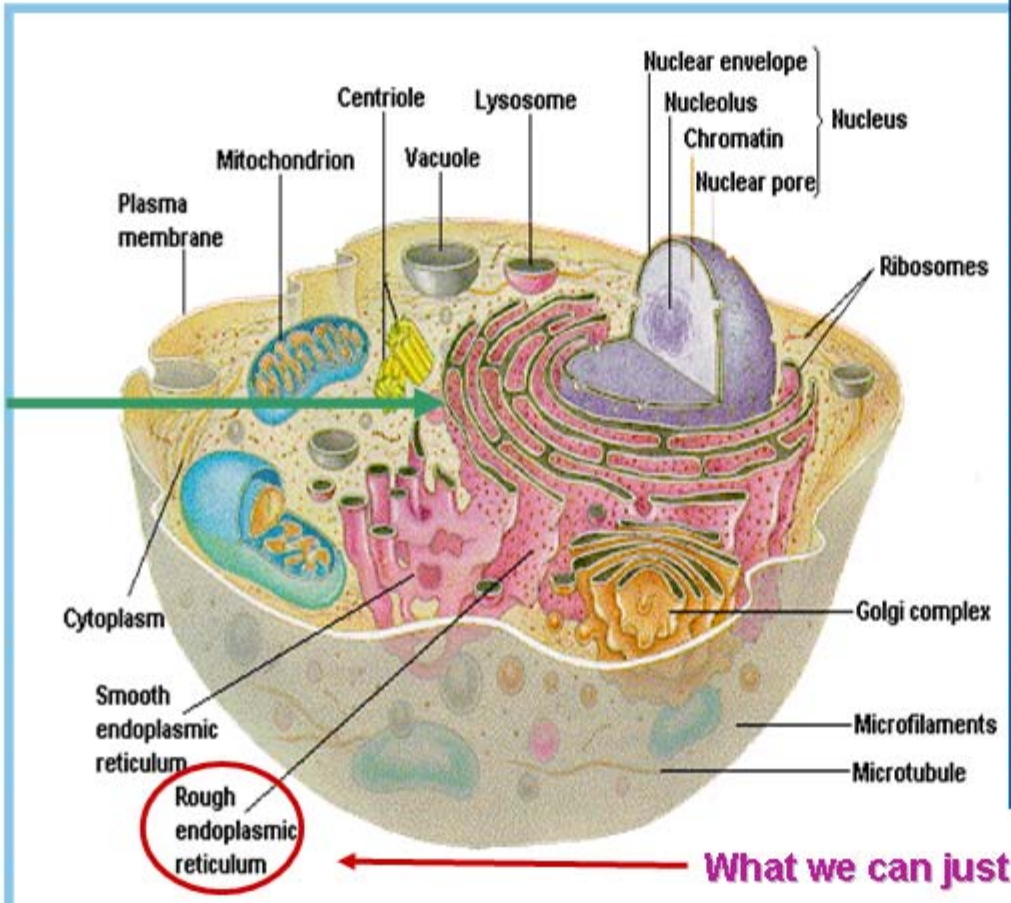


Taken from *Human Biology* by Daniel Chiras

**Organelles**
4 Million Ribosomes
30,000 Proteasomes
Dozens of Mitochondria

**Macromolecules**
5 Billion Proteins
5,000 to 10,000 different species
1 meter of DNA with Several Billion bases
60 Million tRNAs
700,000 mRNAs

**Chemical Pathways**
Vast numbers
Tightly coupled

**What we can just model today with <100TF**

1993 1999 2005 2011 2017 2023 2029



*We have a way to go ...*

intel

# What can we expect!

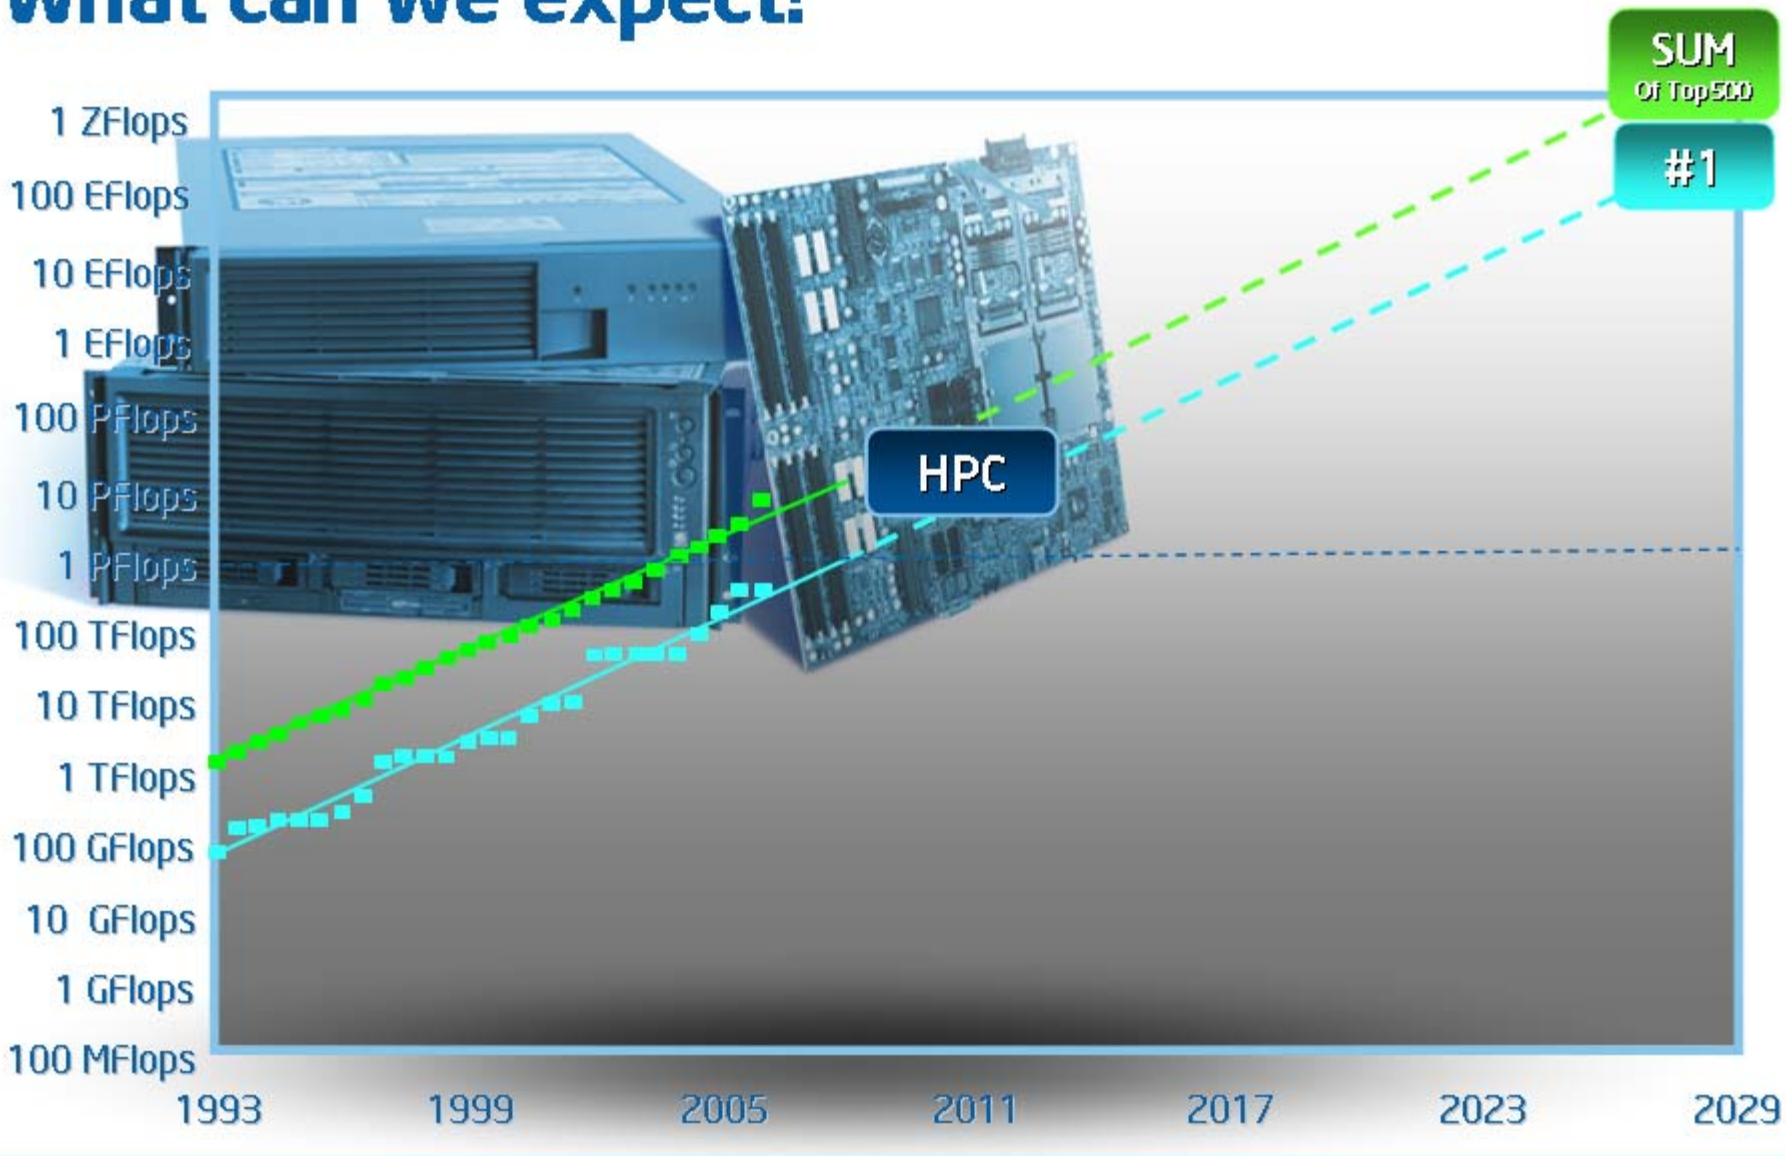

# Performance Acceleration

## *Marching Towards Petascale and Beyond*

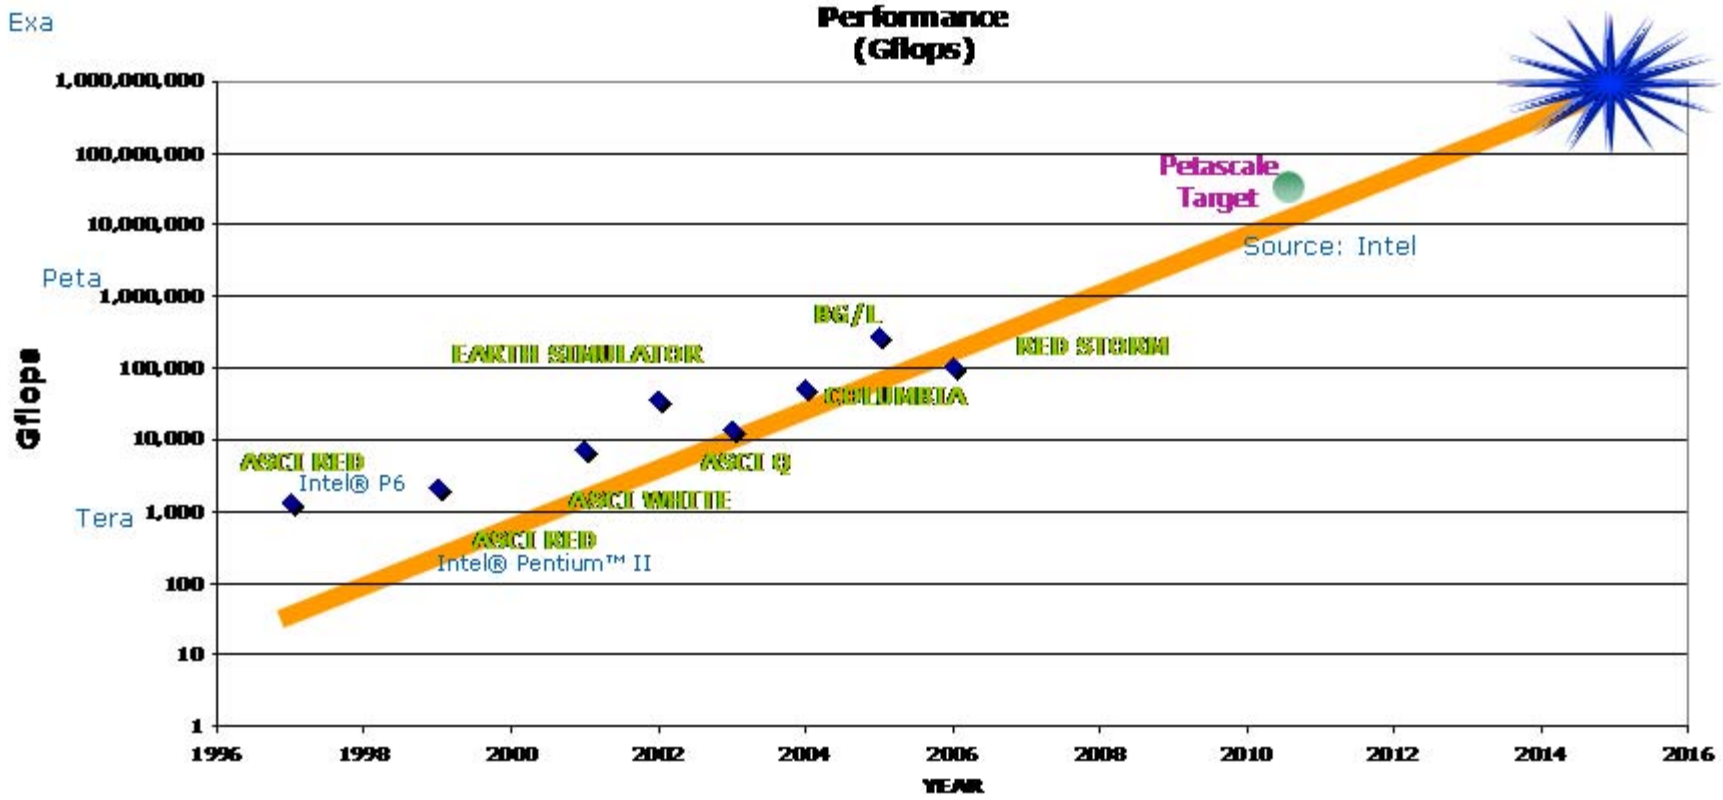


Source: Exascale Study Group

# Technologies Improve Performance vs. Area

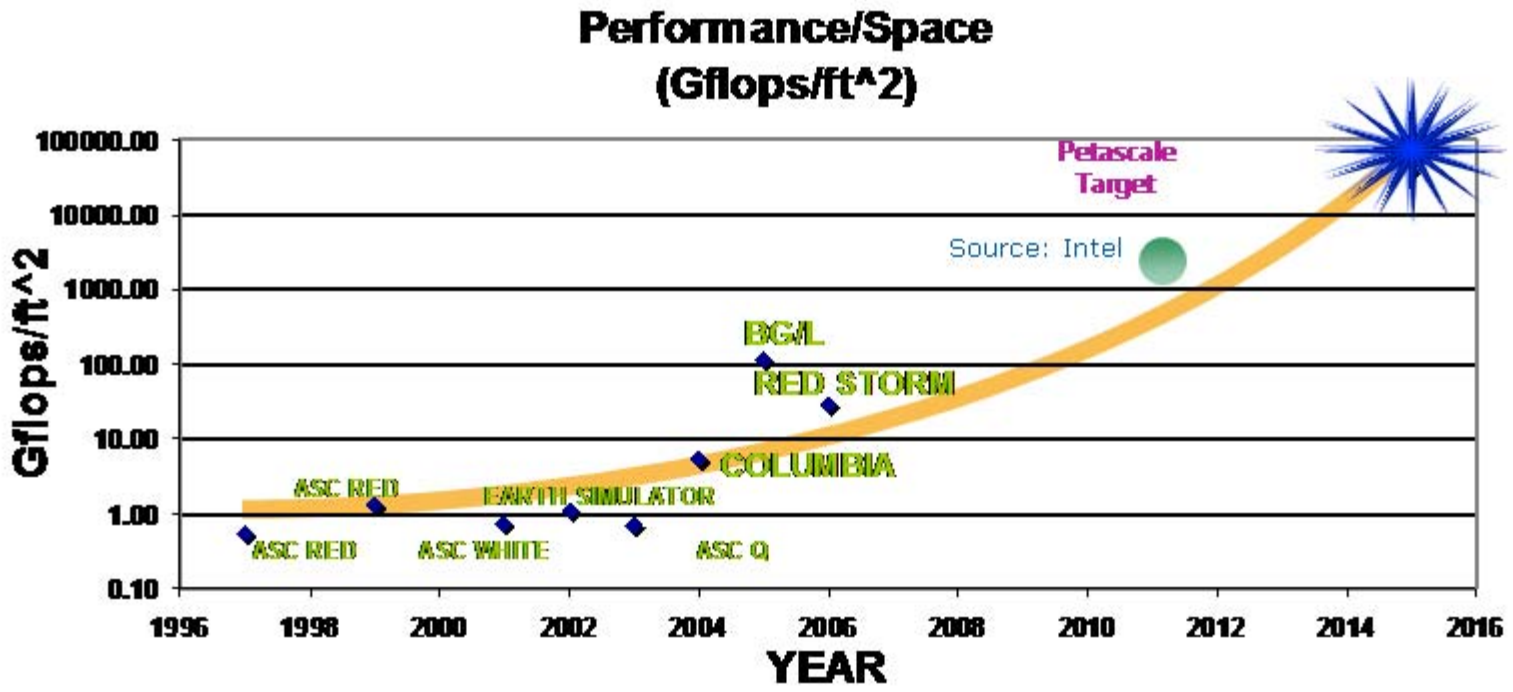


Source: Exascale Study Group

CDC 6600 filled ~500 ft$^2$ of floor space
1965

The CRAY-XMP occupied ~100 ft$^2$
1984

# Performance vs. Power

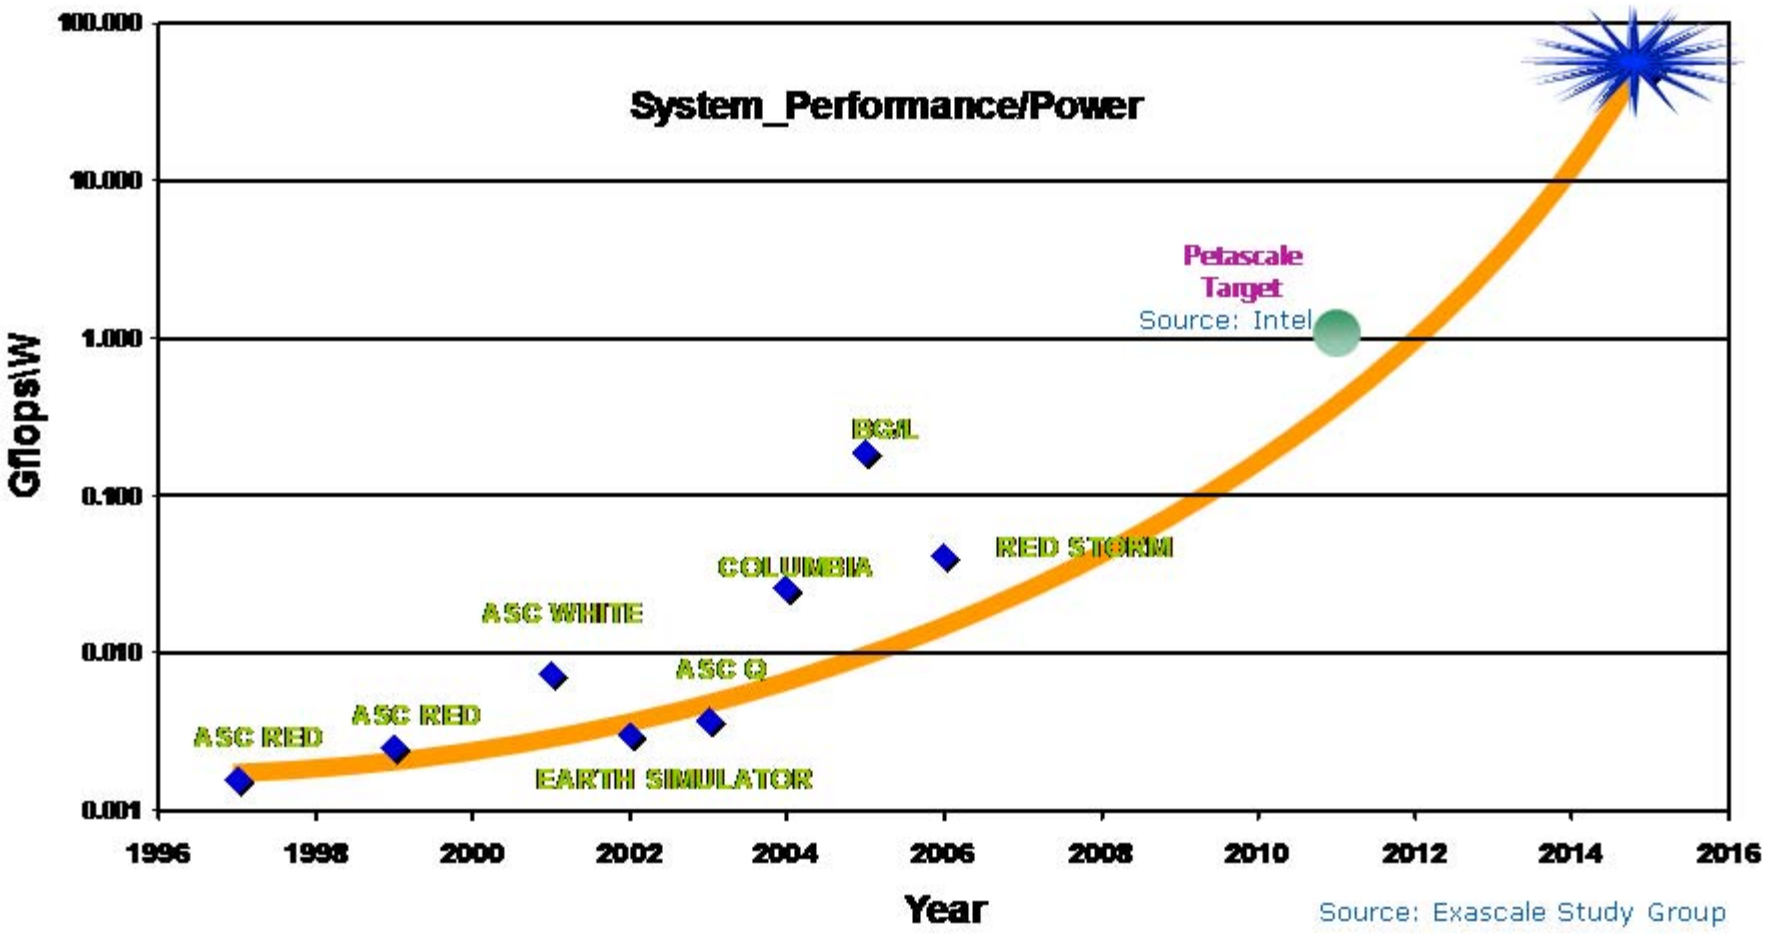


Source: Exascale Study Group



*Power reductions made possible of personalizing past HPC*

intel

# Grand Challenge Problems Break All Rule of Thumb

~~$10-20 Million?~~ → > $100 Million


~~Watt~~ → Mega Watt

~~>100 sqft~~ → >Thousands of sqft

~~Countable Processors~~ → Millions of Cores and Sockets



*Technologies push the envelope…*

intel

# What Brings HPC to the Masses?

# Technology Scaling

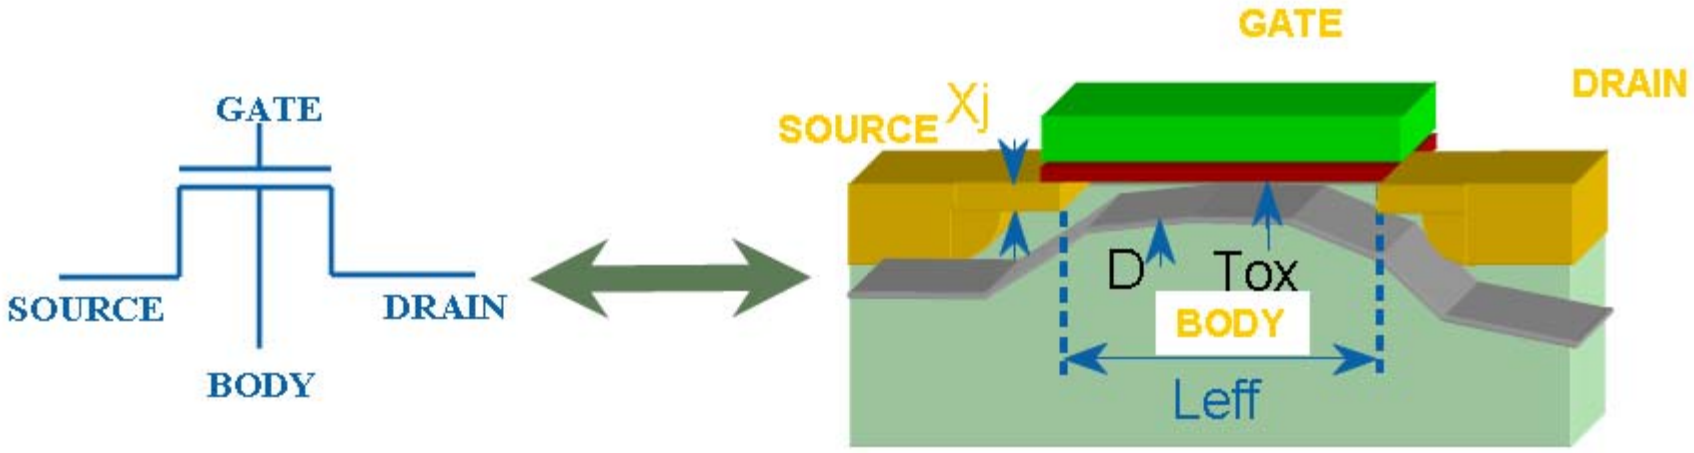


| | |
|---|---|
| Dimensions scale down by 30% | Doubles transistor density |
| Oxide thickness scaling | Faster transistor, higher performance |
| Vdd & Vt scaling | Lower active power |

# Computing Affordability

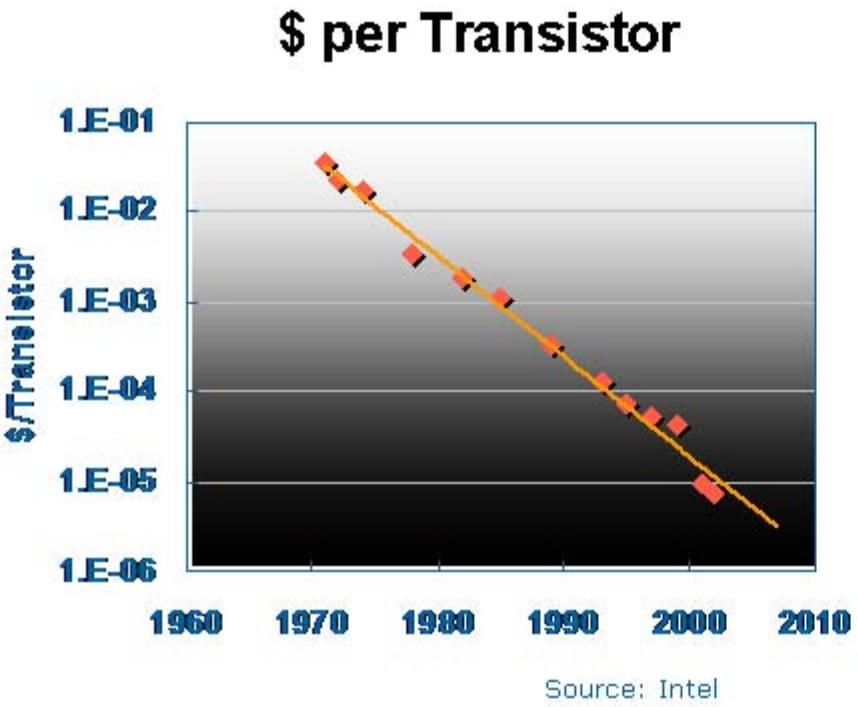


Source: Intel

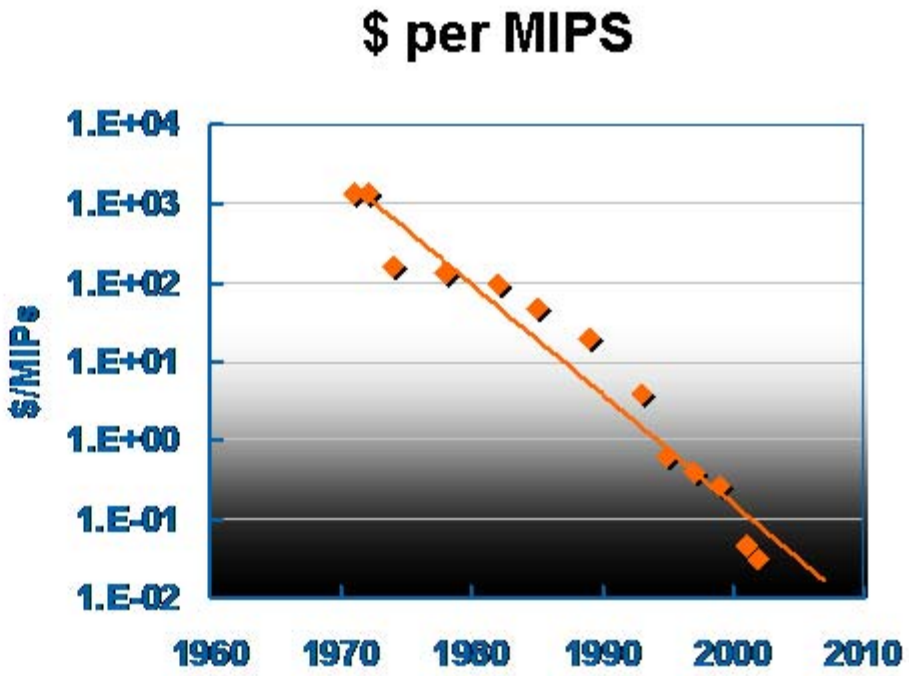


"Supercomputer prices range from about $1 million to $20 million, but in terms of constant dollars, the $10 million average is a useful rule of thumb over about 3 decades." - source: "Characteristics of High-Performance Computers," 1989, Jack Worlton



*Xbox 360's GPU has 332 Million Transistors, system cost ~$350*
*CRAY X-MP had <1M Transistors, system cost $15 Million in 1984*

# Accelerating Multi- and Many-core

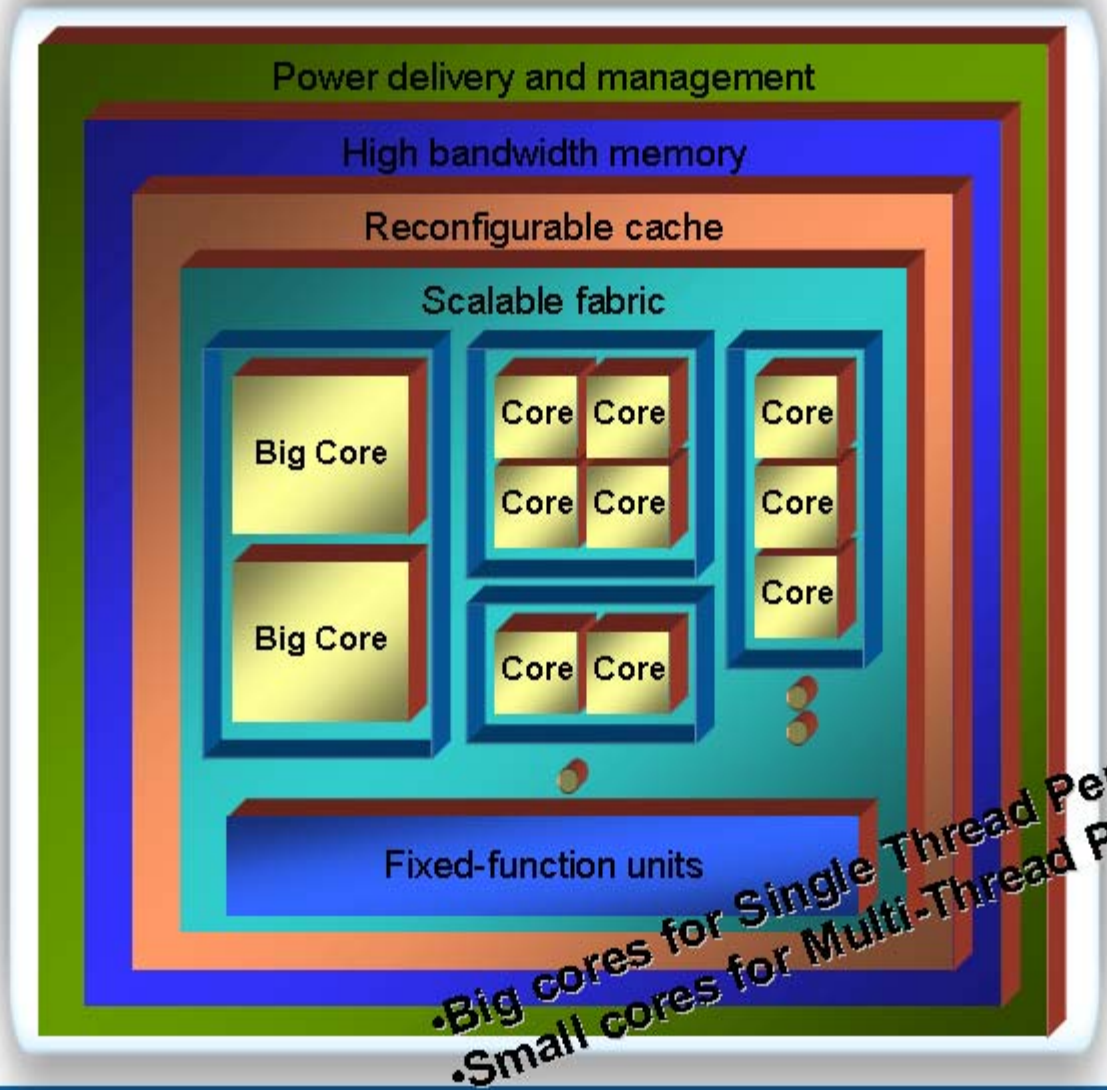


**Performance Through Parallelism**



intel

# Future Silicon Photonics at the Cost of Silicon

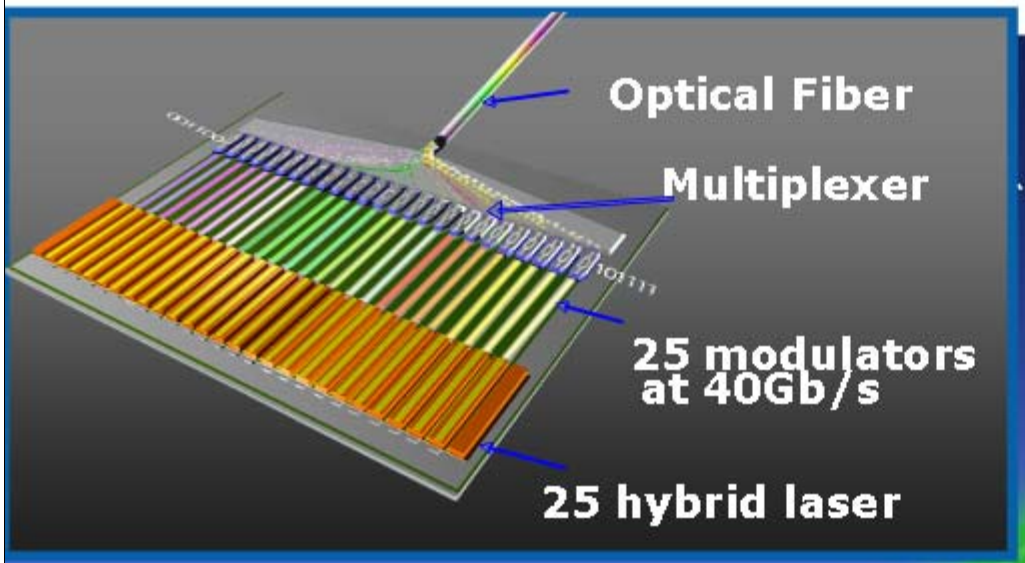


Pioneering work by Dr. Richard Soref (early 1980's)

2002
- Integrated APD+TIA — UT
- Inverted Taper — NTT, Cornel

2003
- DGADC — Surrey
- PBG WG <25dB/cm — IBM

2004
- 30GHz Si-Ge Photodetector — IBM
- PBG WG <7dB/cm — IBM, FESTA, NTT

2005
- QCSE in Si — Stanford
- Stim-Emission — Brown
- CW Raman Laser — Intel
- 10Gb/s Modulator — Intel, Luxtera
- 1.5Gb/s Ring Mod. — Cornell
- 39GHz Si-Ge Photodetector — Univ. Stuttgart
- PBG WG <3db/cm — NTT

2006
- Hybrid Silicon Laser — Intel - UCSB
- Broadband Amplification — Cornell
- E-O Effect Strained-Si — DTU
- 10Gb/s SiGe PIN Commercial Quality — Intel

2007
- 40G Raman Amp & λ Conv. — Intel
- 40Gb/s Modulator — Intel

Note: Not an exhaustive list



intel

# Increasing Signaling Rate

## More Bandwidth & Less Power

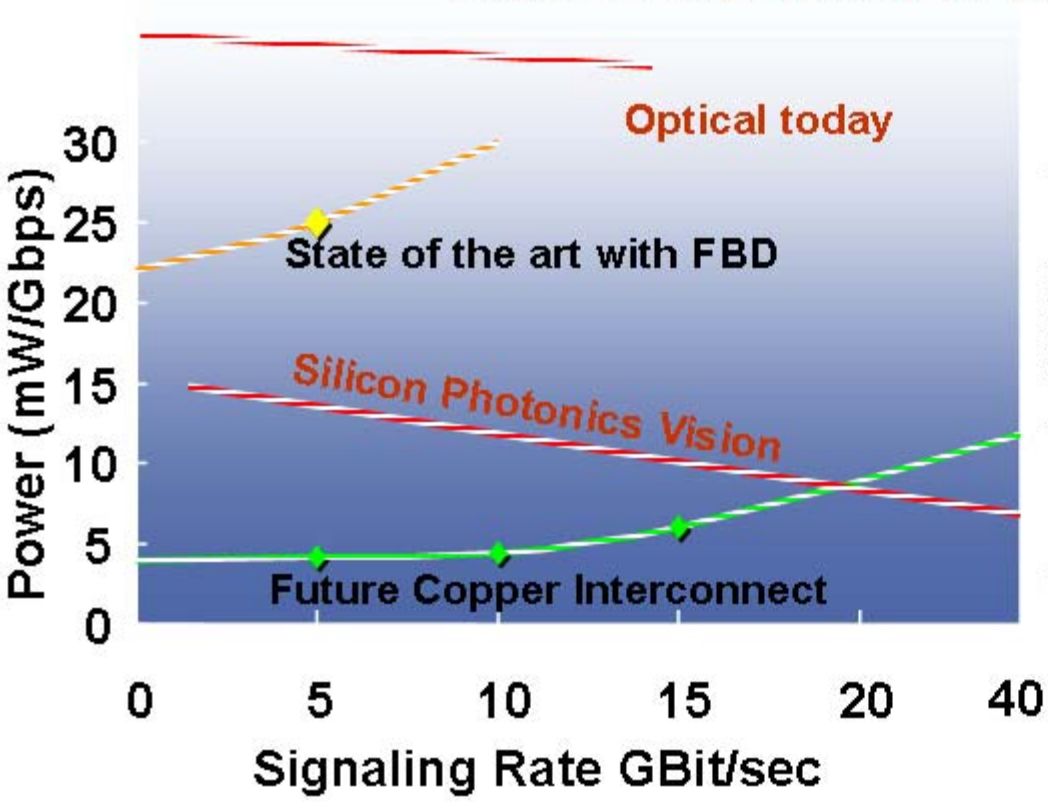


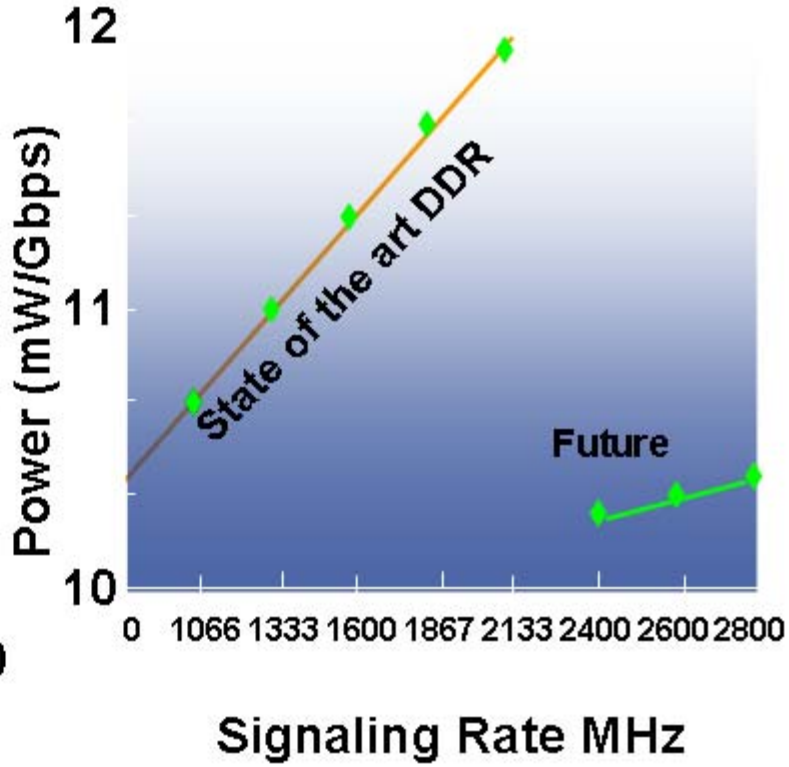


Differential Copper Interconnect & Silicon Photonics

Single-ended Interconnect

# Technologies for Smaller and Better Flash Memory

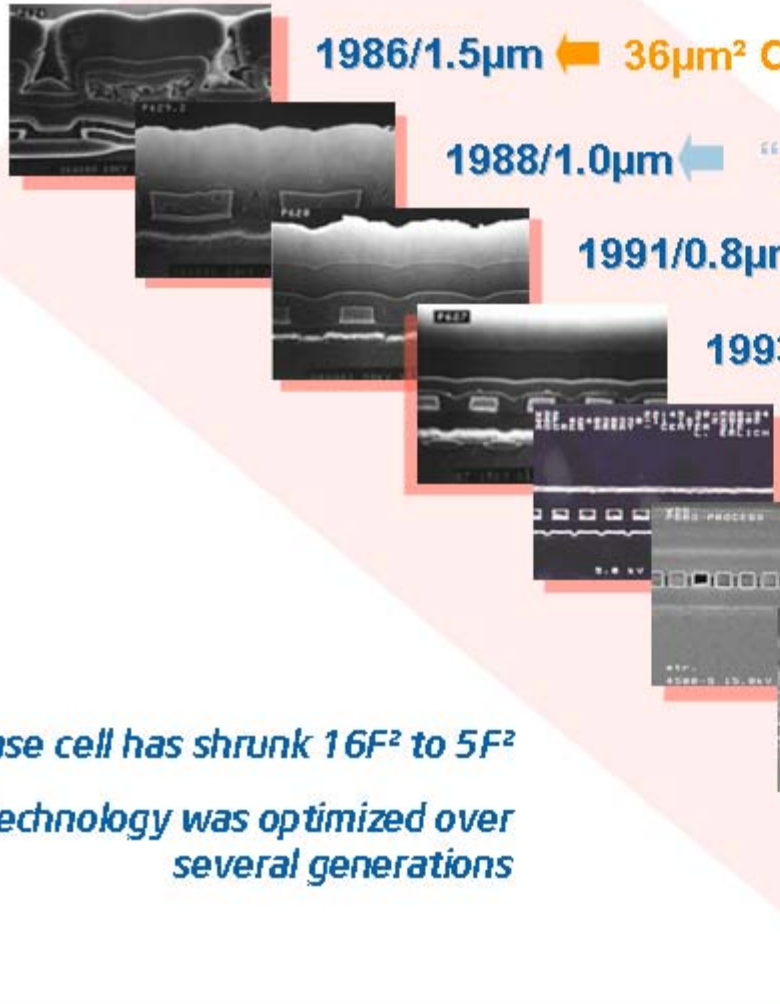

1986/1.5µm ← 36µm² Cell = 16F²

1988/1.0µm ← "Boot Block"

1991/0.8µm ← Intel FlashFile™ Memory

1993/0.6µm ← 10F² cell, Flash Takes Over EpROM

1996/0.4µm ← First MLC

1998/0.25µm ← Wireless Boom

2000/0.18µm ← 4.9F² eff. cell

2002/0.13µm ← 1.8V ISF

2004/90nm ← 4.7F² effective cell

*Base cell has shrunk 16F² to 5F²*

*Technology was optimized over several generations*

# Personalize Bigger and More Powerful Engine

## Tomorrow's Personal Computing

**Today's HPC Performance**
**Vs.**
**Efficient Power + Small Space**

CERN OpenLab, Intel HPC Roundtable '06

Becomes this?

# Driving to the Future:

## *New Generation of Mass Computing*

Wall-street, Google, Pixar/ILM, SecondLife…

| Computer Vision | Rendering | Physical Simulation | Financial Analytics | Data mining |
|---|---|---|---|---|

## *Grange Challenge Problems*

- Field containment for Nuclear fusion
- Multi modal fluid structures problems
- Modeling materials' behaviors in phase change
- Many more…



intel