Accelerate to Exascale: Hardware

Stephen S. Pawlowski
Intel Senior Fellow
General Manager, Architecture and Planning
Chief Technology Officer, Digital Enterprise Group
Intel Corporation

September 25th, 2009

Legal Disclaimer

Notice: This document contains 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. Contact your local Intel sales office or your distributor to obtain the latest specification before placing your product order.

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, or life sustaining applications. Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

All products, dates, and figures are preliminary for planning purposes and are subject to change without notice.

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.

The (Intel products discussed herein) 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.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at http://www.intel.com.

Intel® Itanium®, Xeon™, Pentium®, Intel SpeedStep® and Intel NetBurst® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Copyright © 2009, Intel Corporation. All rights reserved.

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

Agenda

Exascale Hardware Challenges

Exascale Goal

Solution Space

Optimal Solution Space

Tera->Petaflop (12 years); An Exaflop is <u>much</u> harder

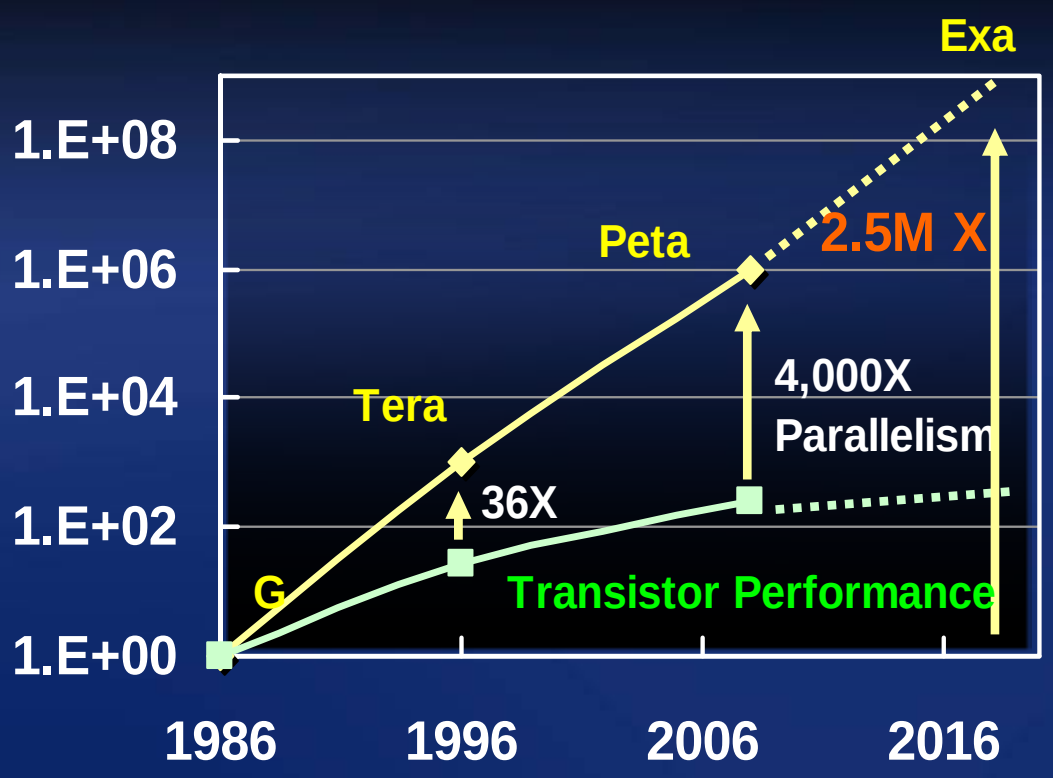

Source: Intel

Relative Performance and Power (GFlops as the base)

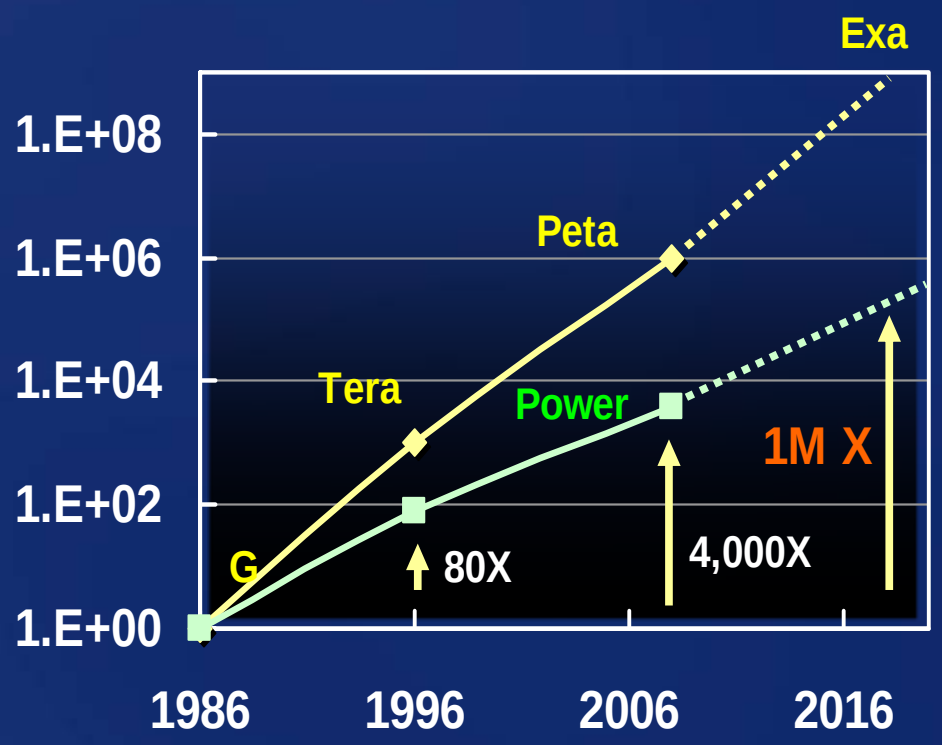

Business as usual is not an option.

Computing is driven by power efficiency and scalability.

The Fact of Voltage Scaling

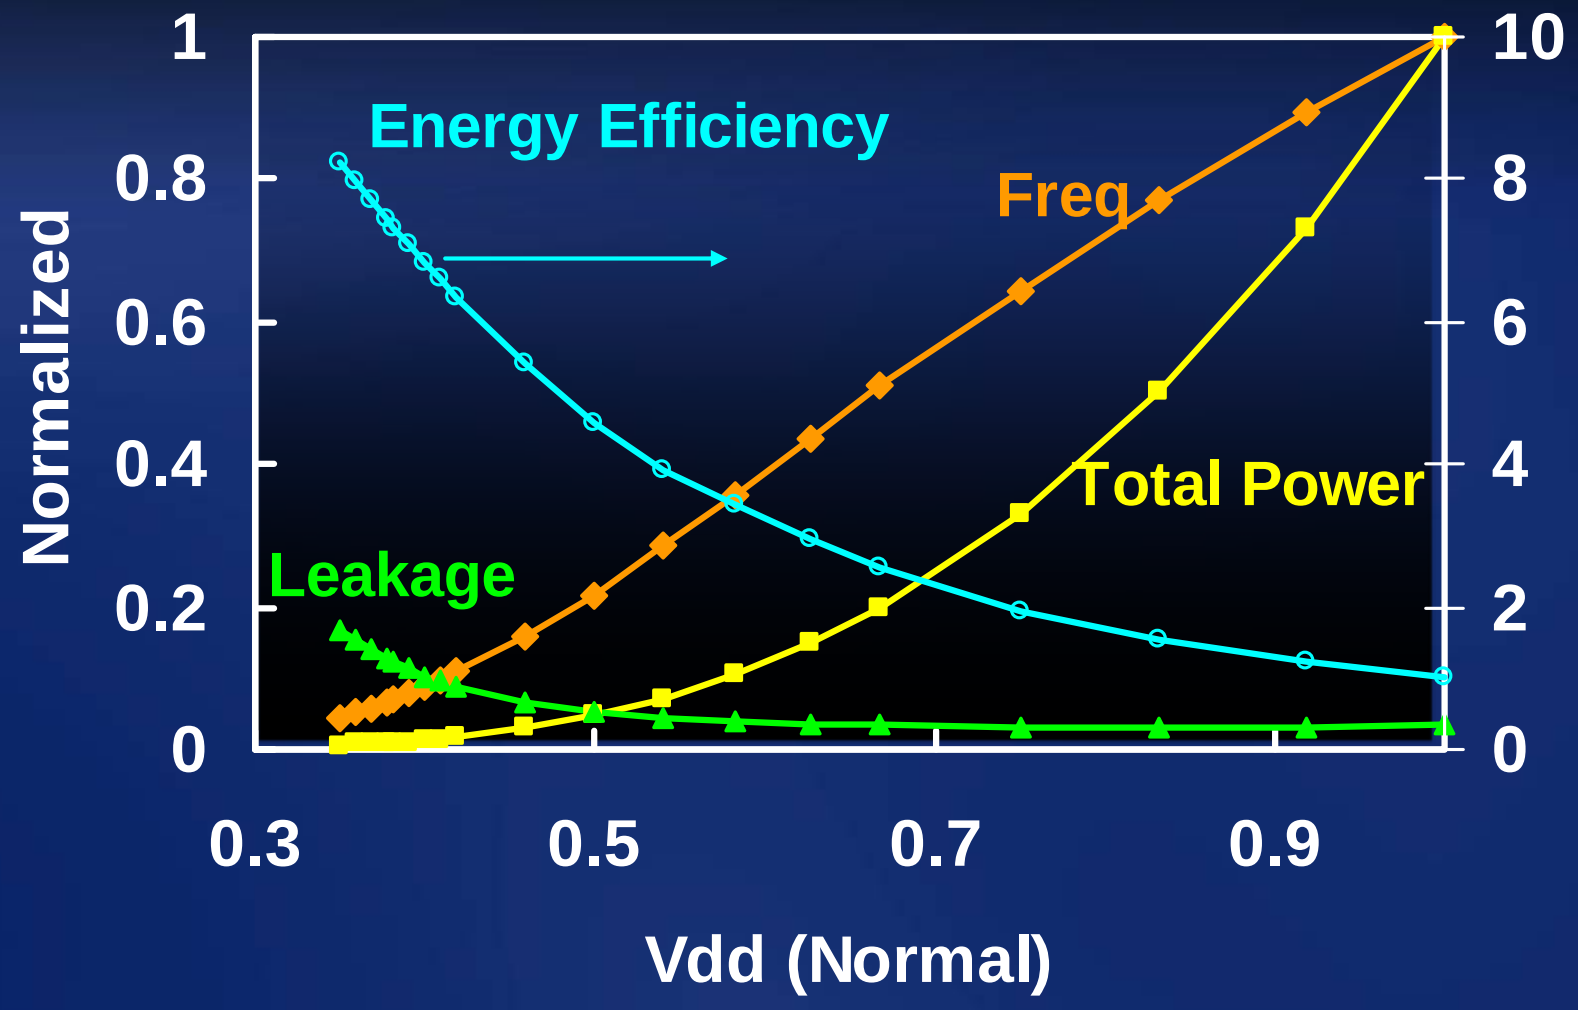

Balancing between voltage and frequency scaling for energy efficient performance

Future Scaling Outlook / Assumption

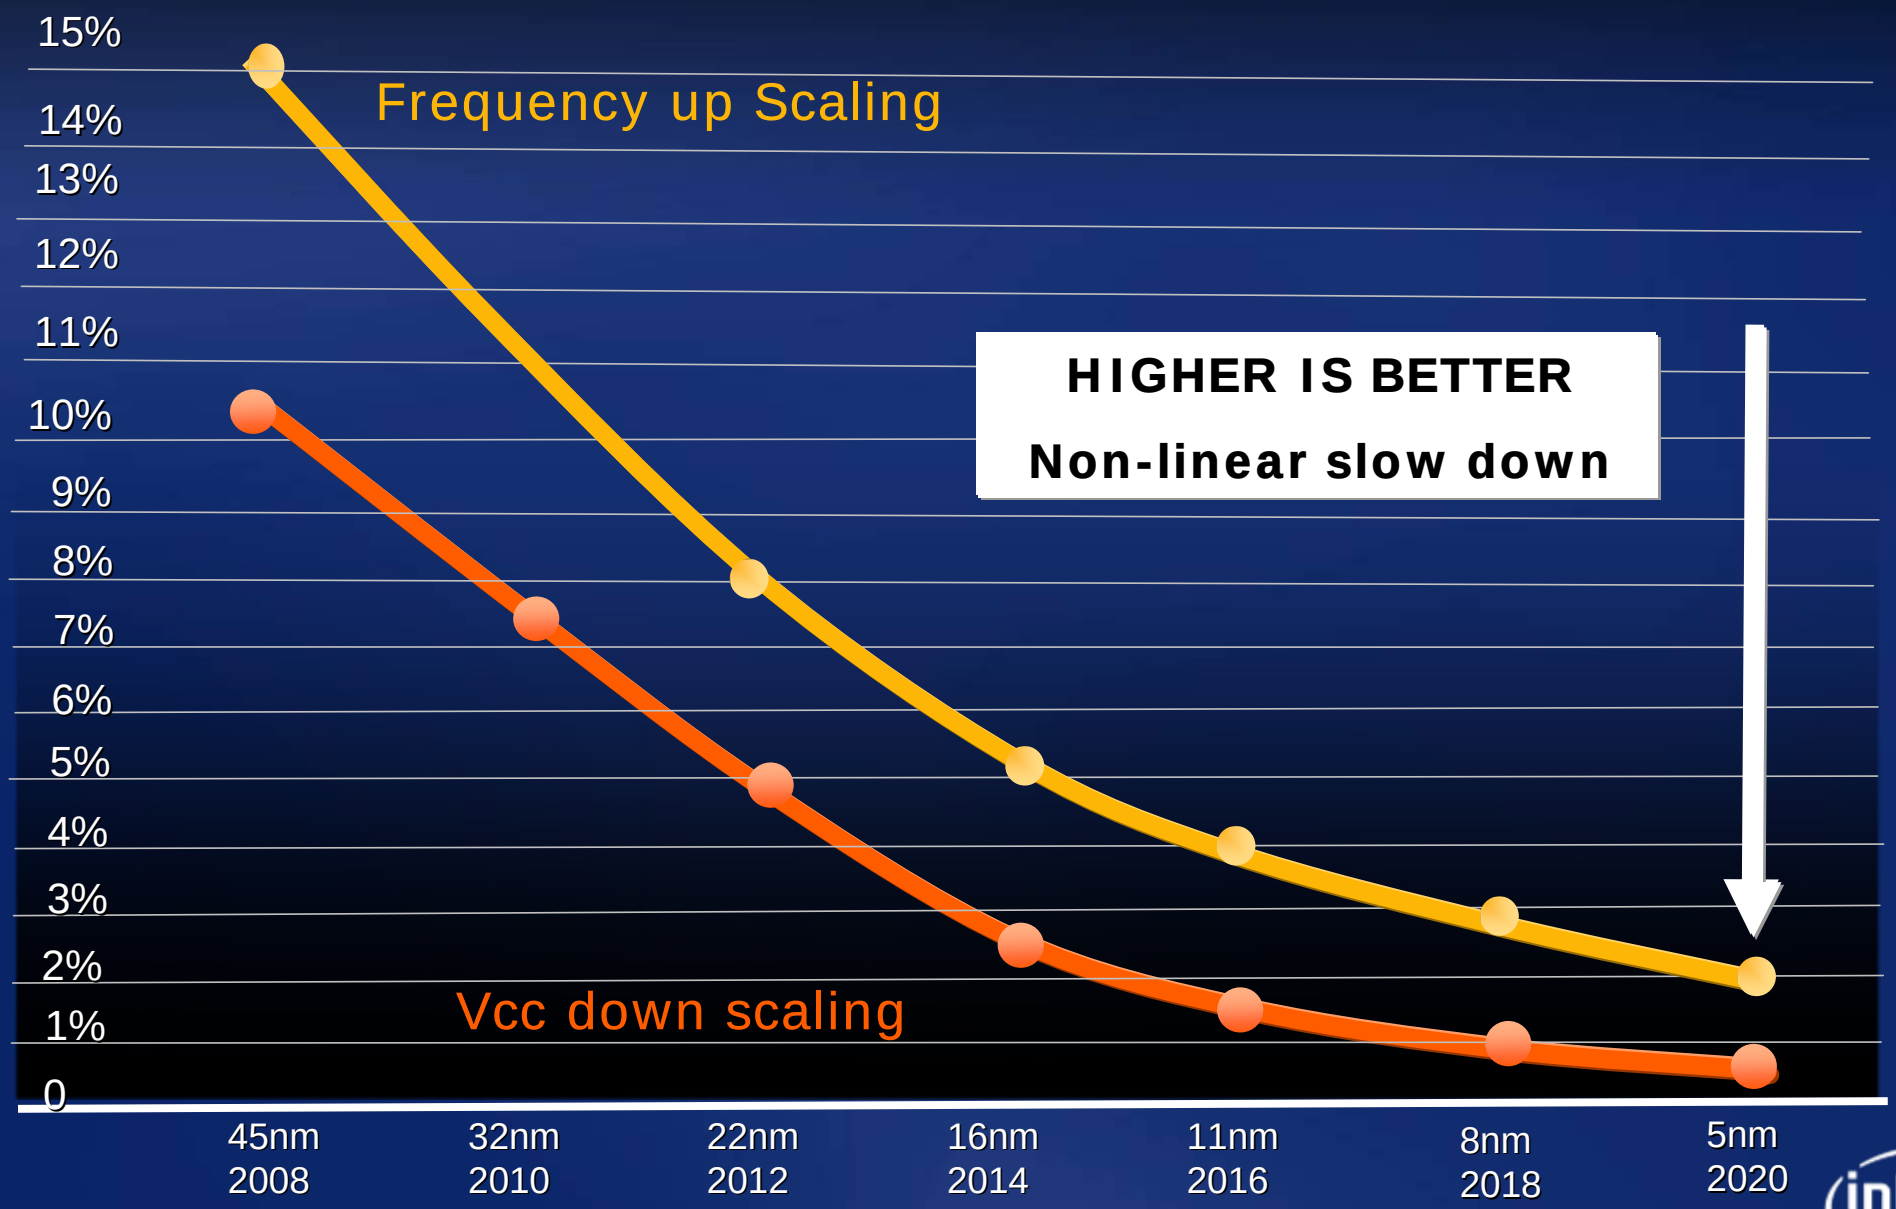

Source: Intel

Is An 8nm Core Good Enough for Exa?

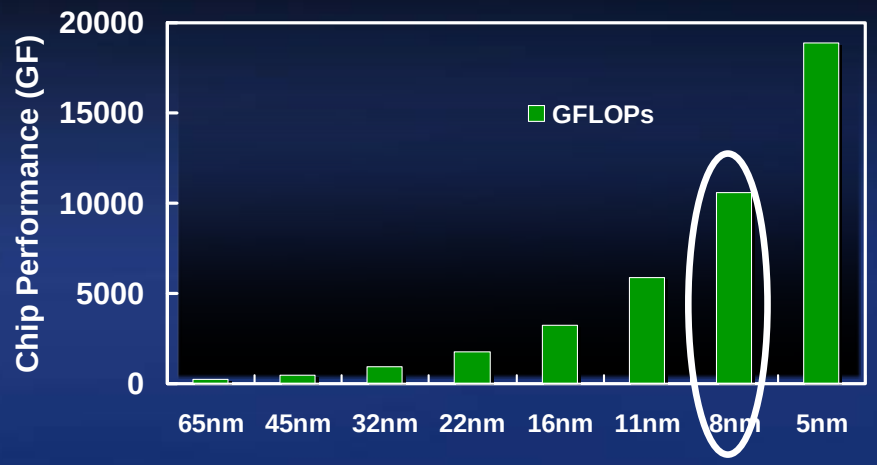

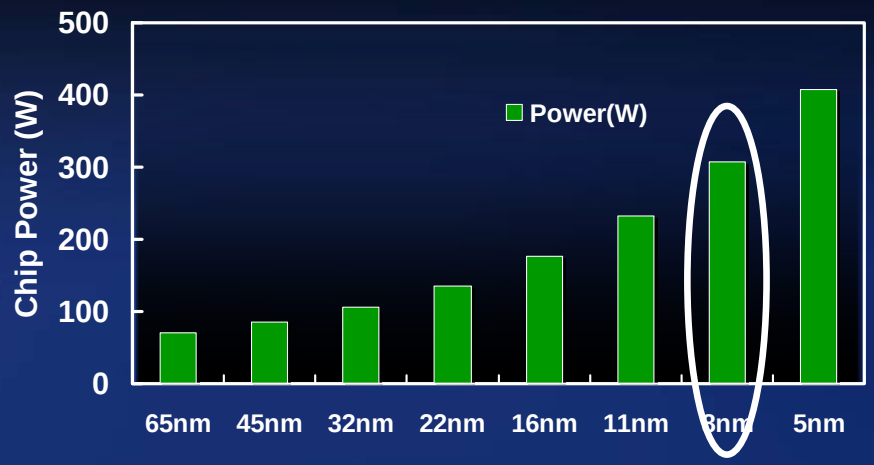

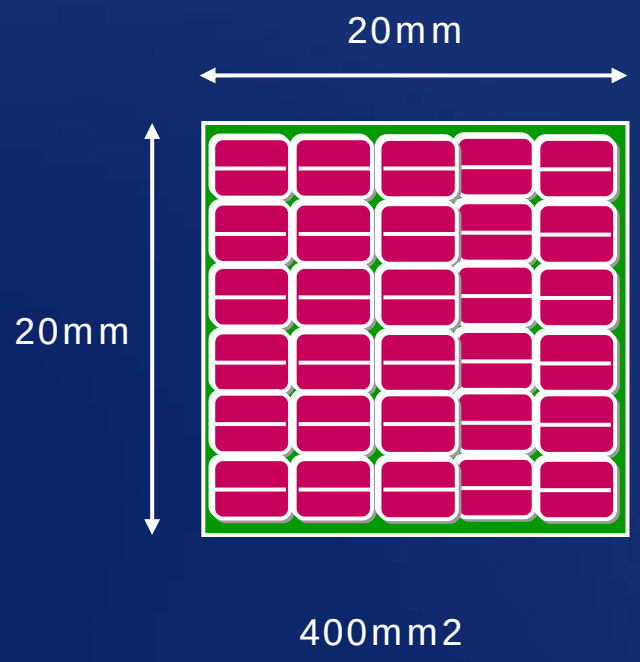

2018, 8nm computing node/socket

Cores	1150
Total Local Memory	400 MB
Frequency	4.61 GHz
Peak performance	10.6 TF
Power	300 - 600W
Energy efficiency	34 - 18 GF/Watt

30-60 MW Peak Exa Flop Just to Compute

The Exascale Hardware Challenges

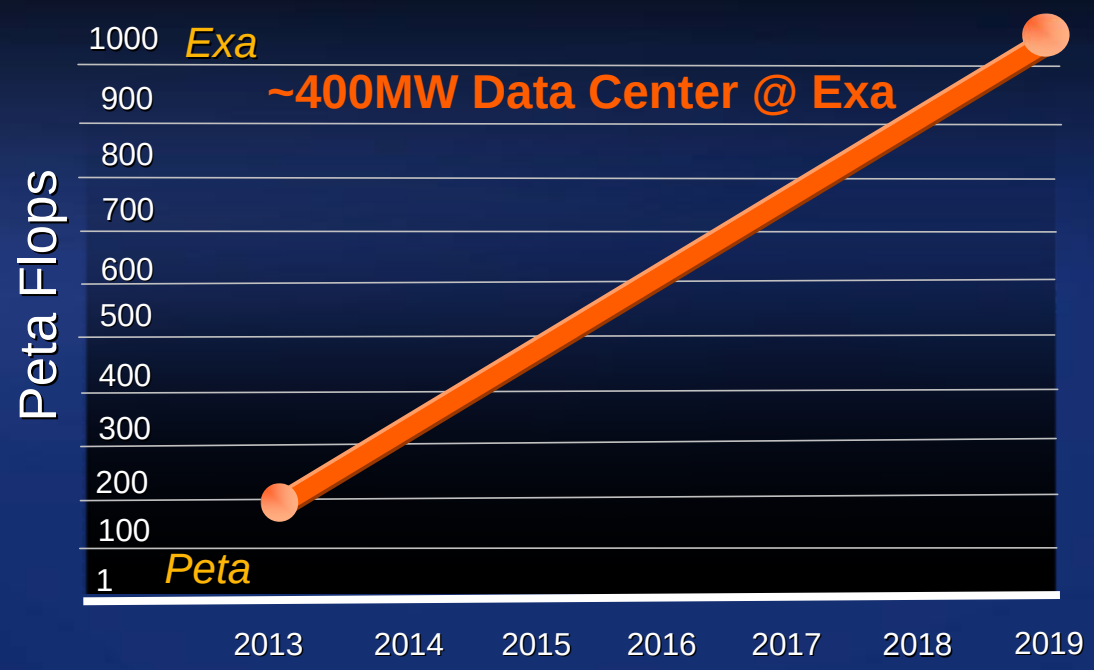

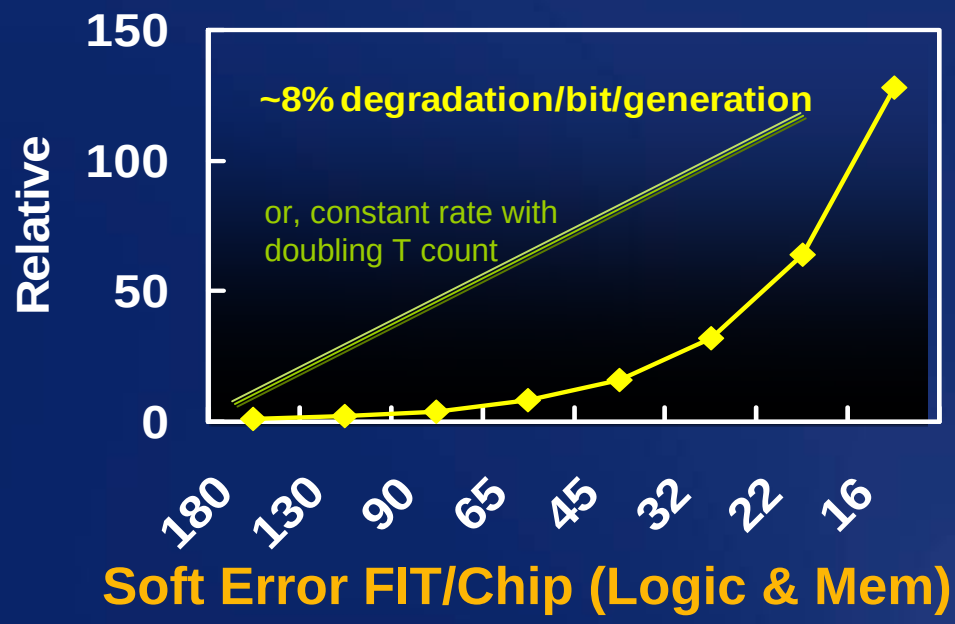

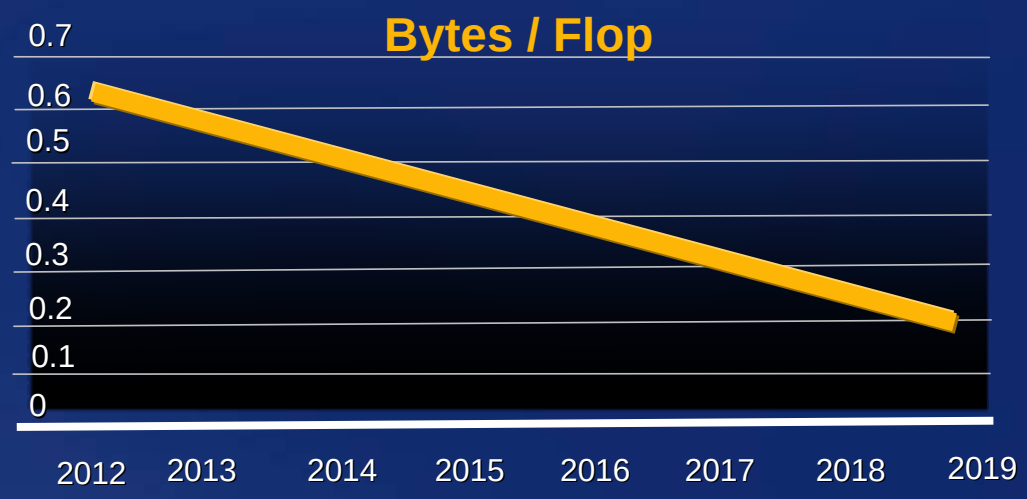

Exascale Goal

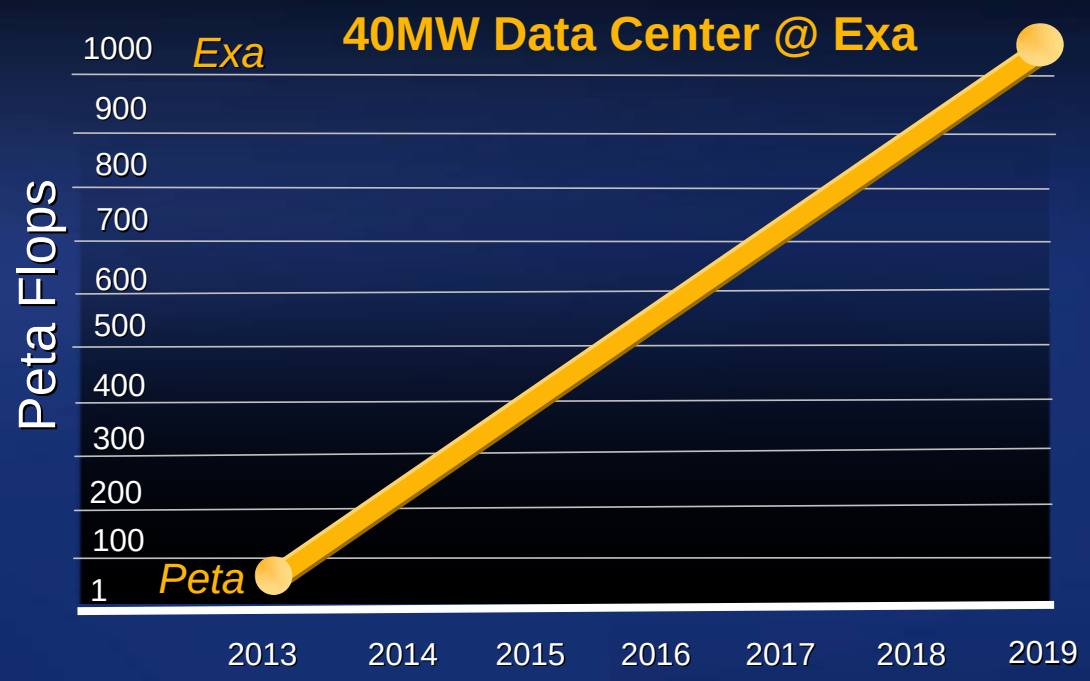

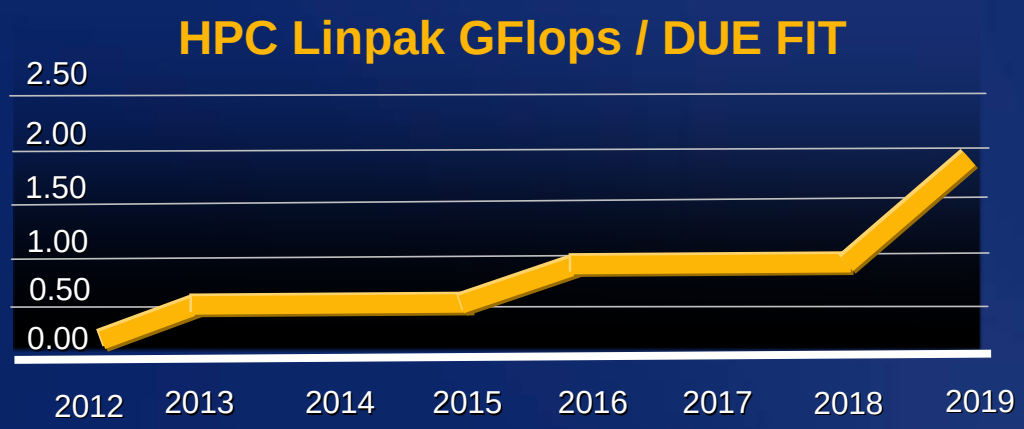

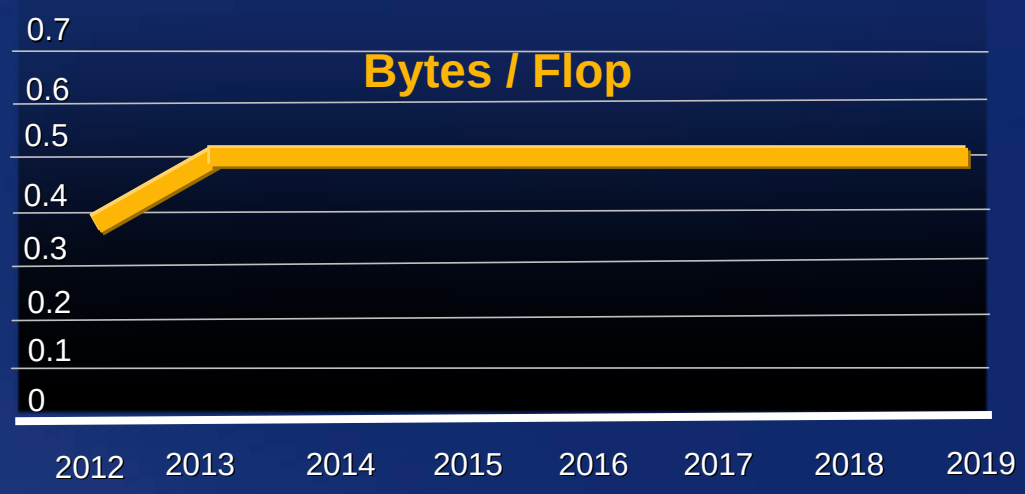

Source: Intel

Solution Space

Hybrid Heterogeneous Processors

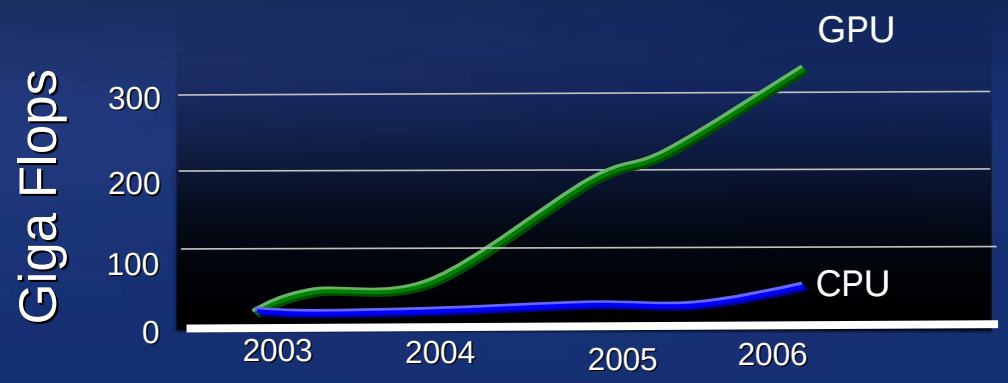

Source: Tom's Hardware "Nvidia's CUDA: The End of the CPU?"

GPU – Graphics Processing Unit

Great Potential for Future HPC

GPGPU – General Purpose GPU

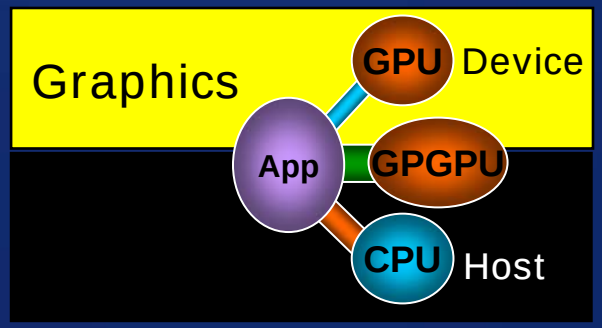

A thread on the GPU is a basic element of the data to be processed. Unlike CPU threads, GPU threads are extremely "lightweight," meaning that a context change between two threads is not a costly operation.

Source: Tom's Hardware "Nvidia's CUDA: The End of the CPU?"

Problem with Current CPU+GPU

Computing tailored to limitations of current GPUs

- Use simple data structures like vectors, matrices
- Very difficult to program, debug, impacts performance
- Coarse-grain interaction (offload & forget)

Programmers are used to the x86 look and feel.
We need to provide convergence with homogeneity

Why the Convergence?

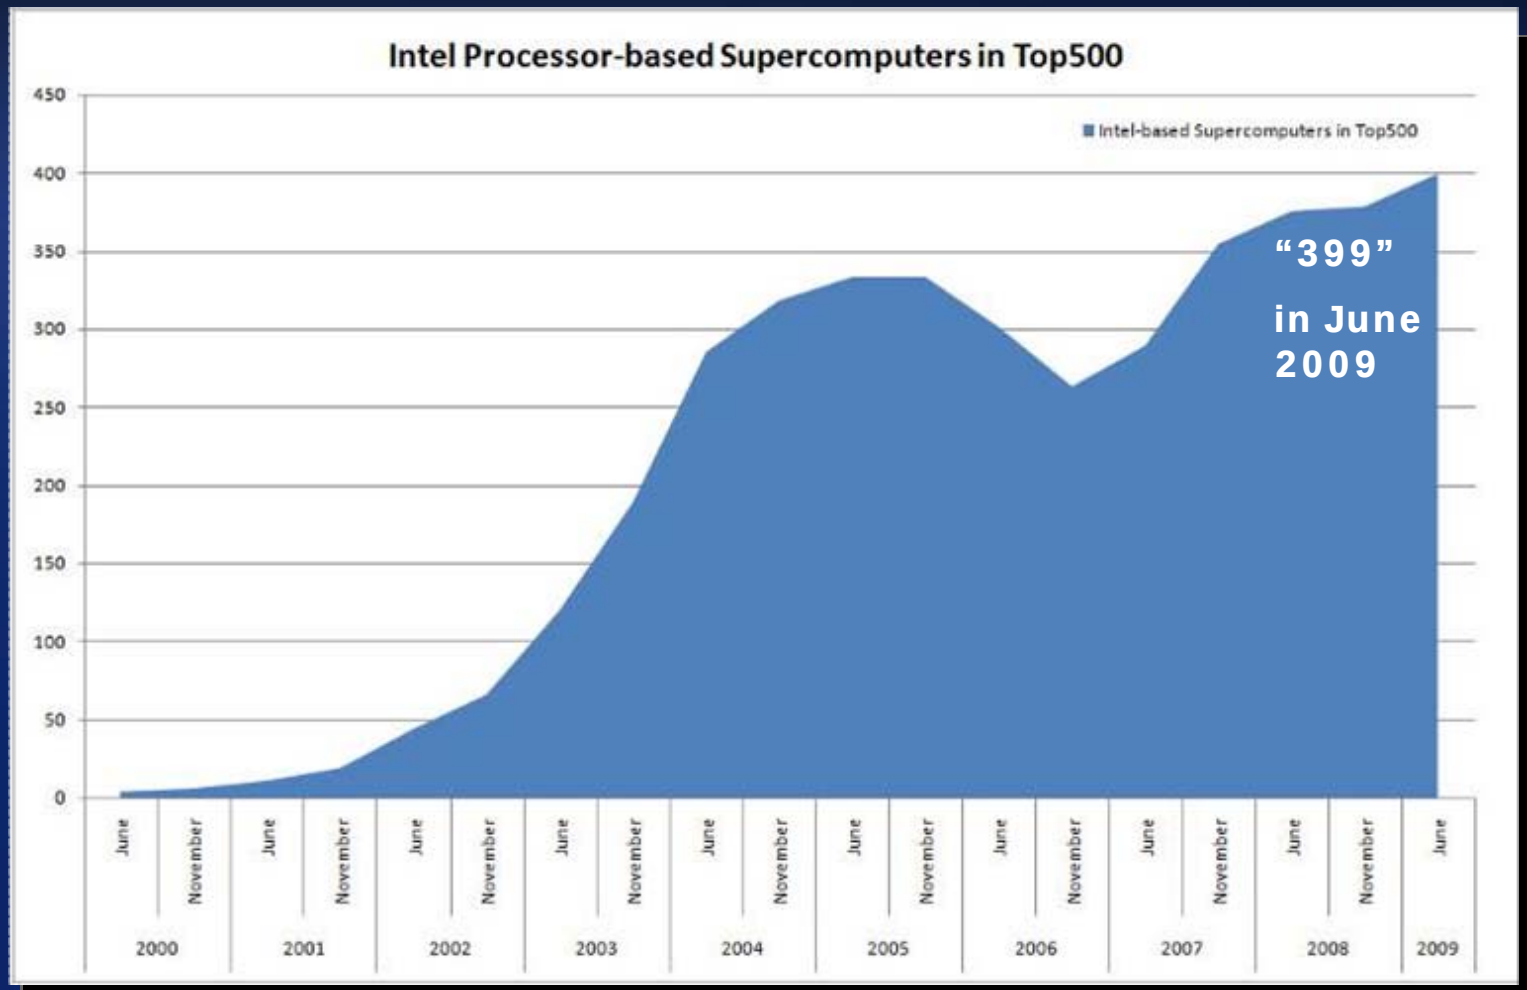

x86 Architecture for Industry-quality code with Ease of Optimization, Programmability, Reusability and Compatibility

Optimal Solution Space #1

Convergence with Homogeneous Computing

Very "Lightweight" Near Threshold logic

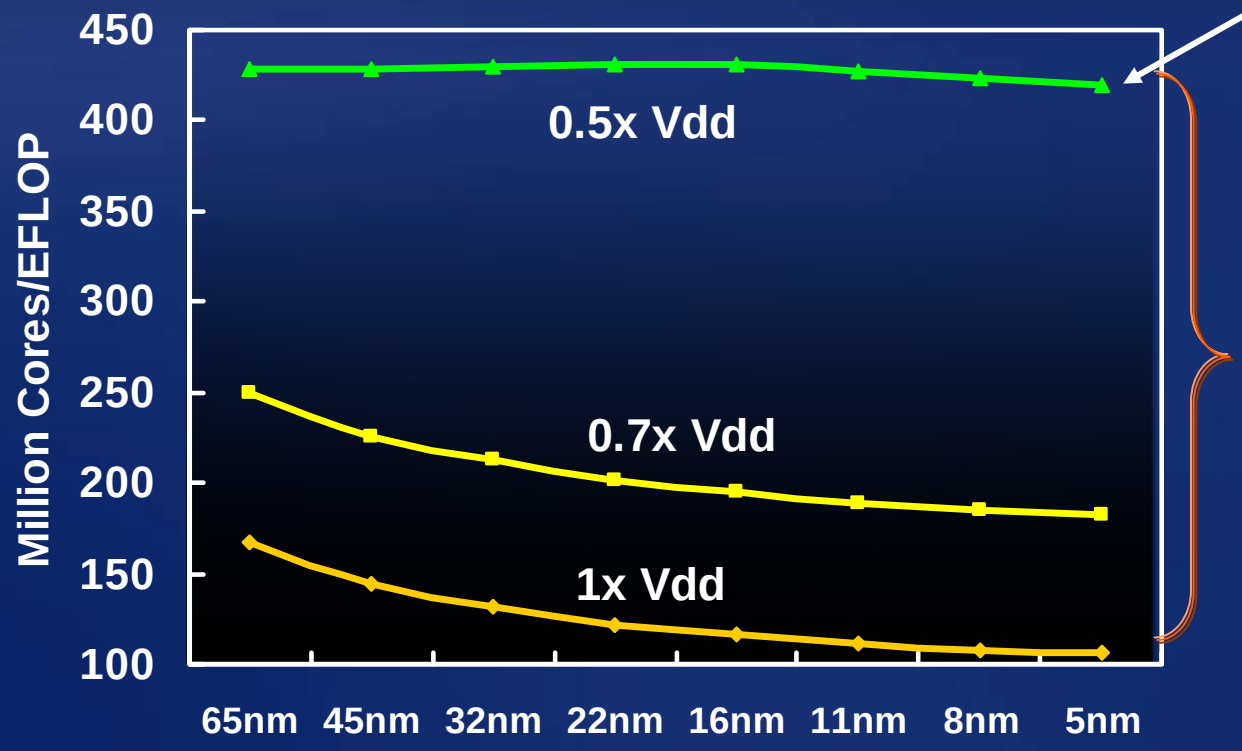

Almost flat because Vdd close to Vt

4X increase in the number of cores (Parallelism)

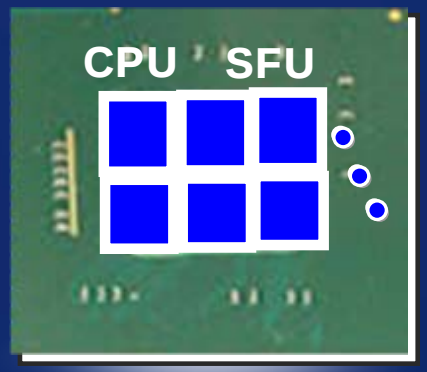

X86 Instruction Set Processor

Optimal Solution Space #2

Convergence with Heterogeneous Computing

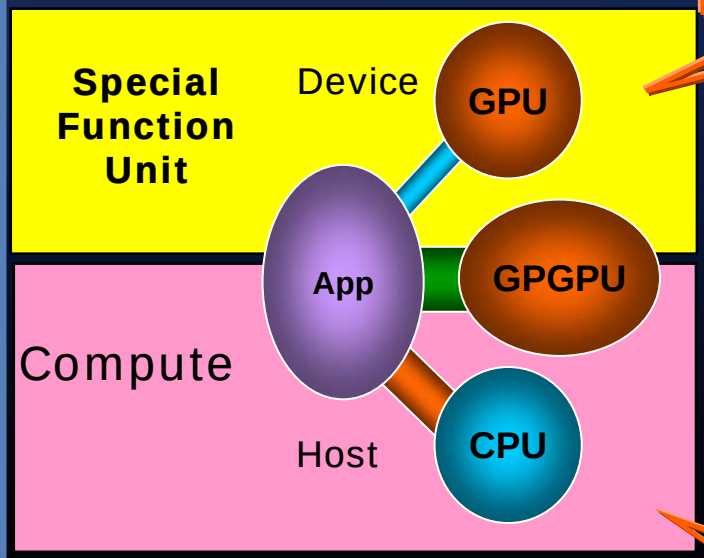

Very "Lightweight" x86 Programmable Special Function Unit (SFU)

+

Uniform Shared Memory Model between SFU and CPU

- Share data structures and pointers
- Unified memory semantics

+

Very "Lightweight" CPU threads

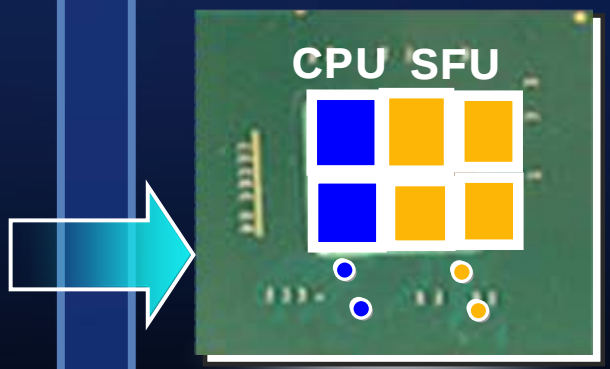

Intel Architecture Instruction Set CPU + SFU

Integrated

Use and share complex data structures, split application between CPU and SFU

Uncompromised Reliability & Resiliency

Billions of Transistors
Hundreds of Cores

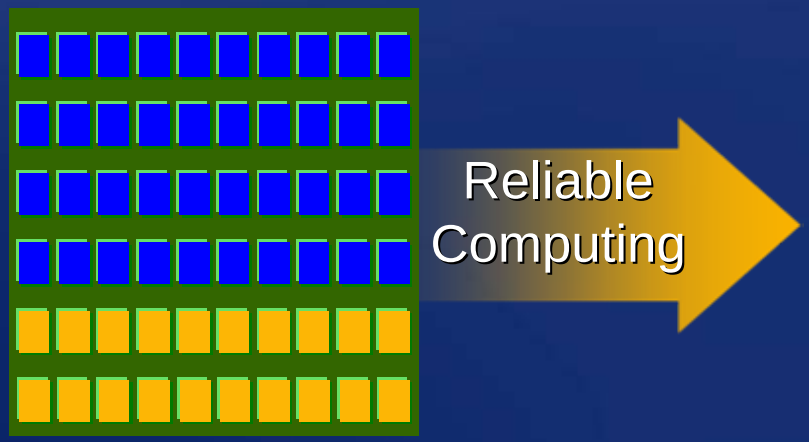

Transient errors - *Soft Errors*
Gradual errors - *Variations*
Time dependent - *Degradation*

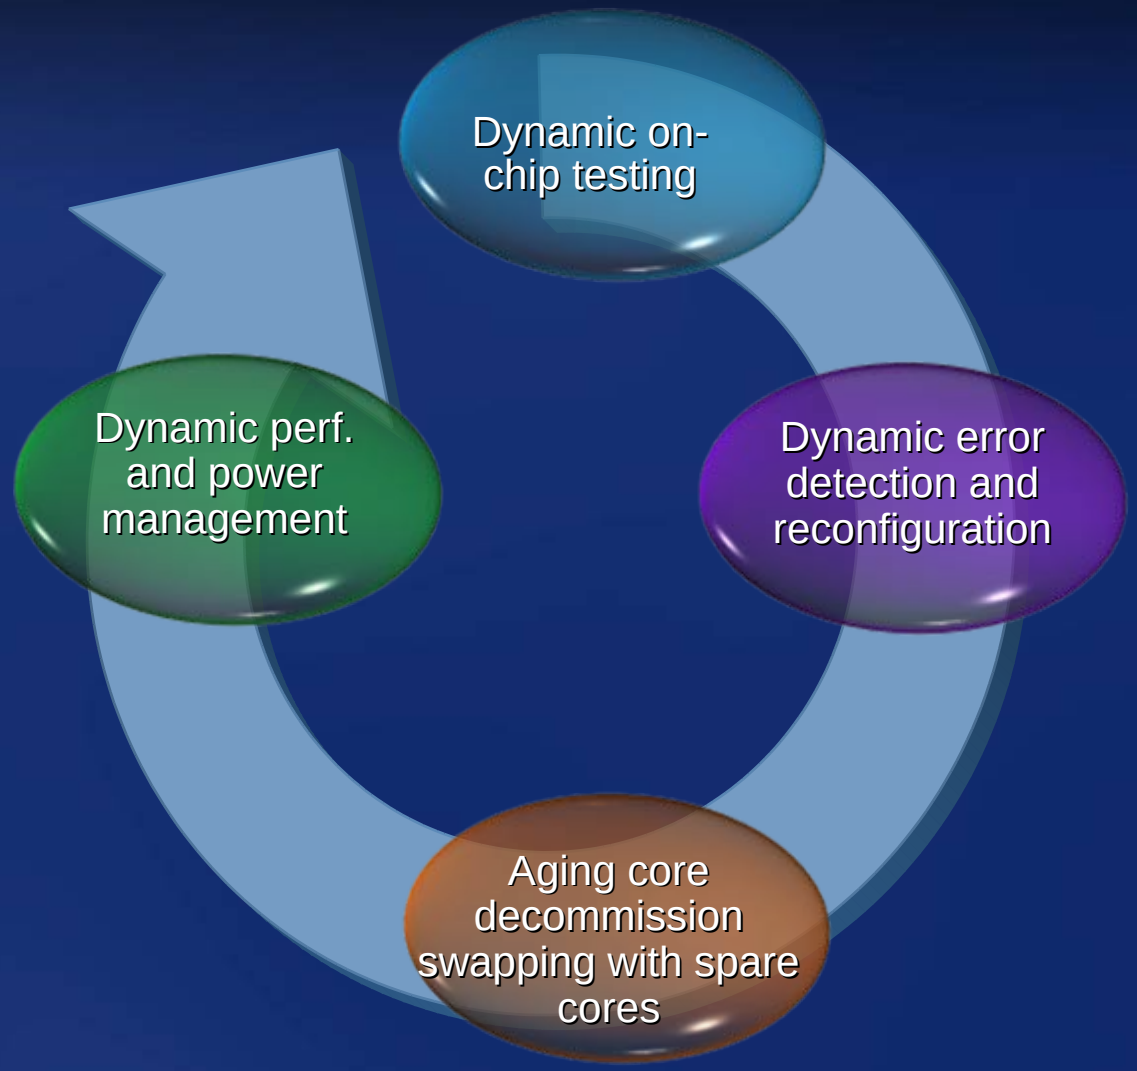

Summary

- The industry is facing great challenges to reach Exascale in both hardware and software
- Business as usual is not an option, driven by energy efficiency and scalability
- CPU and GPU heterogeneous computing has great potential for HPC
- But, the optimal solution is in the ease of use by the end users - ease of optimization, programmability, reusability and compatibility

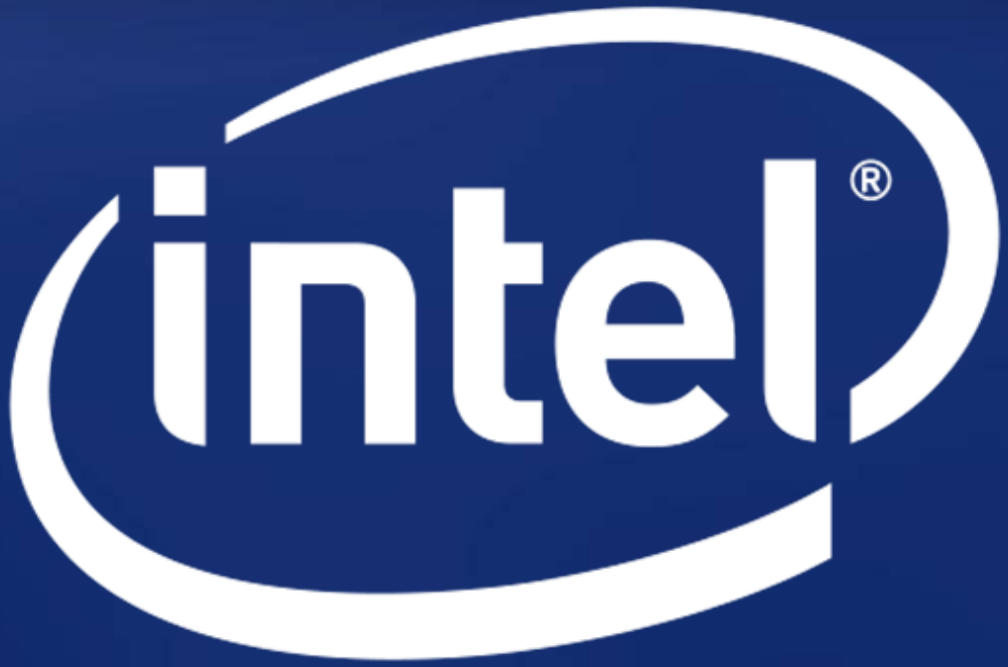
intel