HPC in the Multicore Era

-Challenges and opportunities -

David Barkai, Ph.D.
Intel HPC team

14th Workshop on the Use of High Performance Computing in Meteorology
ECMWF, Shinfield Park, Reading, UK
1-5 Nov., 2010

High Performance Computing

Outline

- Moore's law is alive
- Heterogeneous computing – MIC
- NWS work at Intel
- Challenges for large (up to Exascale) systems

MOORE'S LAW – ALIVE AND WELL

Even as we design with multicore and accelerators..

Moore's Law: Alive and Well at Intel

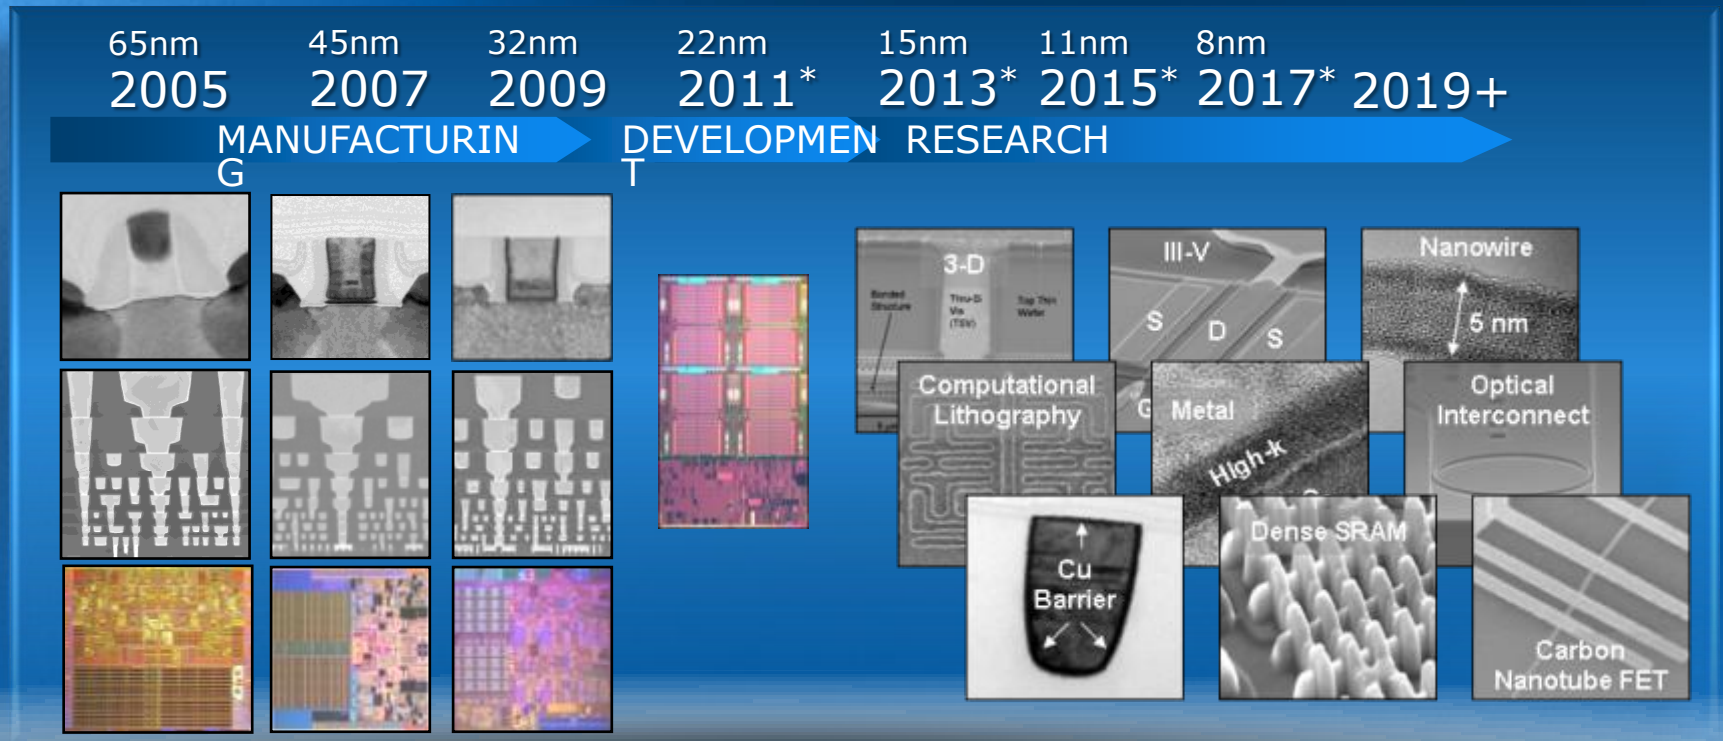

Intel Innovation-Enabled Technology Pipeline is Full

High Performance Micro-Architecture for PetaScale Deployments

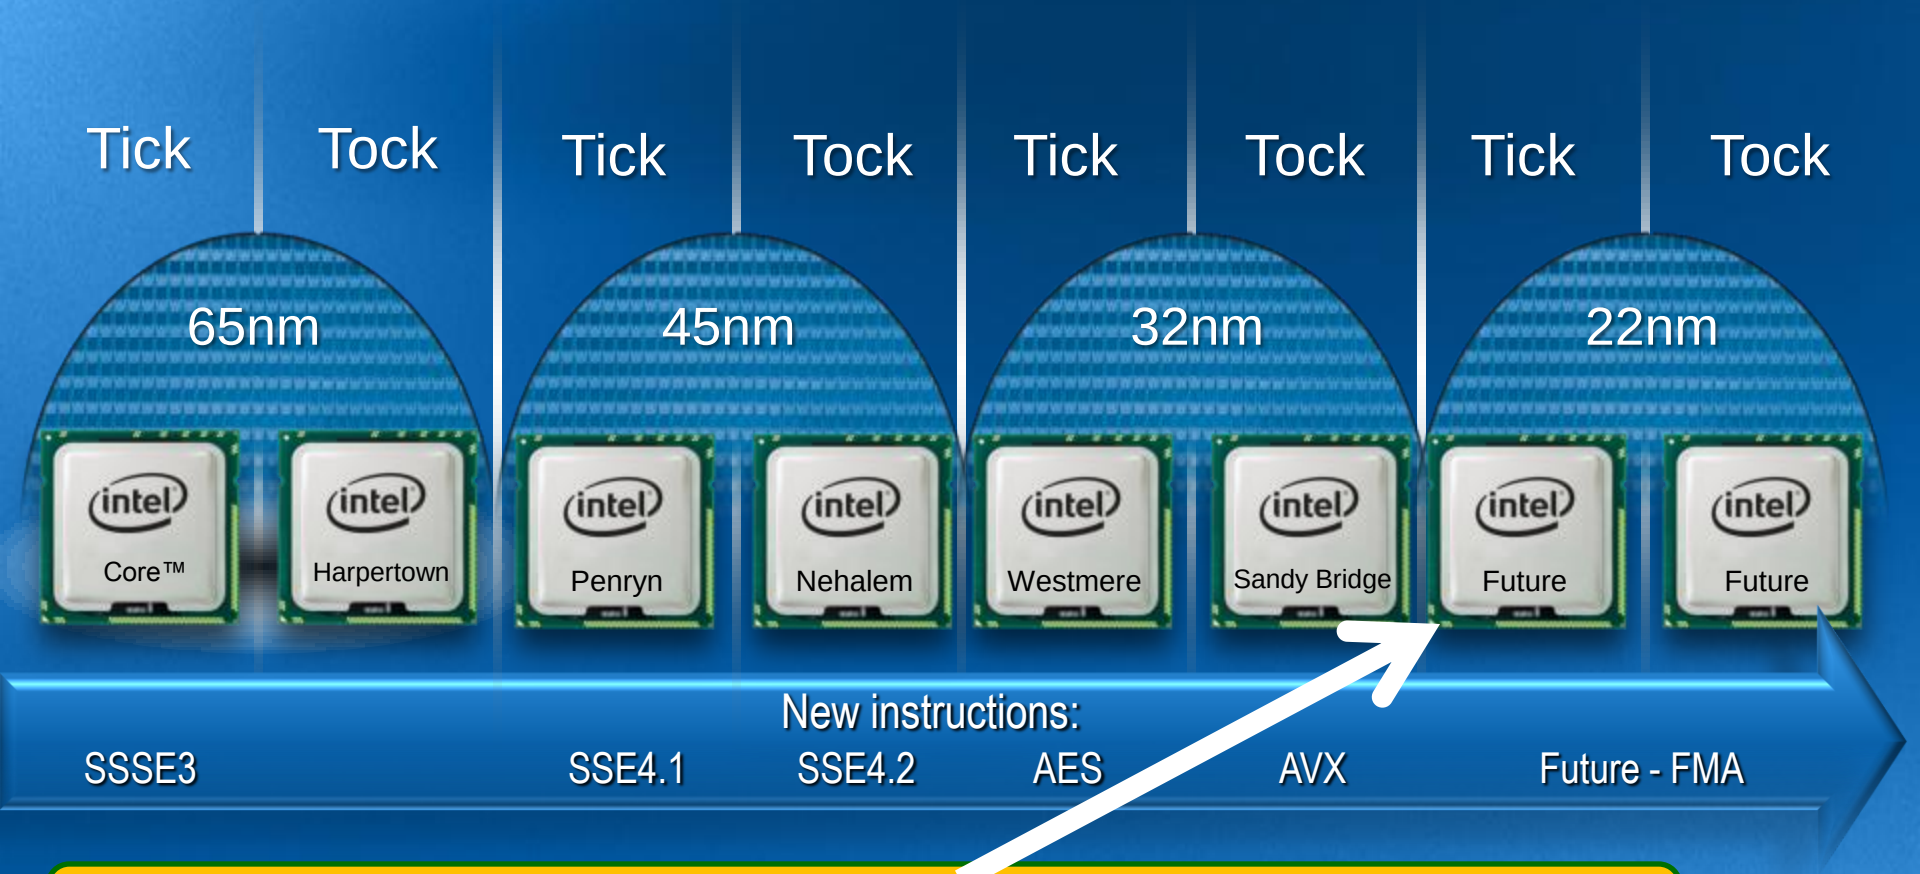

Intel just announced $6-8B investment in two 22nm fabs

COMING SOON

Westmere-EX Overview
New Levels of Scalabilities for the Enterprise

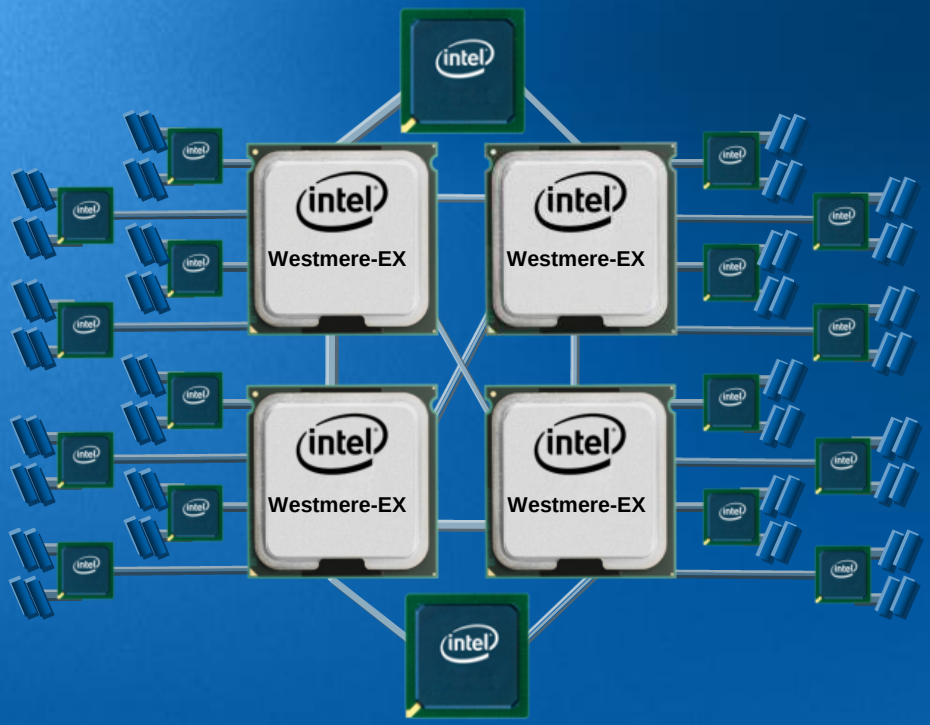

- Up to 10 Cores/ 20 Threads
- 2X the Memory Capacity*
- Enhanced Security
- Second Generation High-k Metal Gate (32nm)
- Target Availability 1H'11

* WSM-EX supports DIMM densities of up to 32GB; this is double the 16GB DIMM maximum supported by Xeon® 7500

Evolution of the Intel® Instruction Set

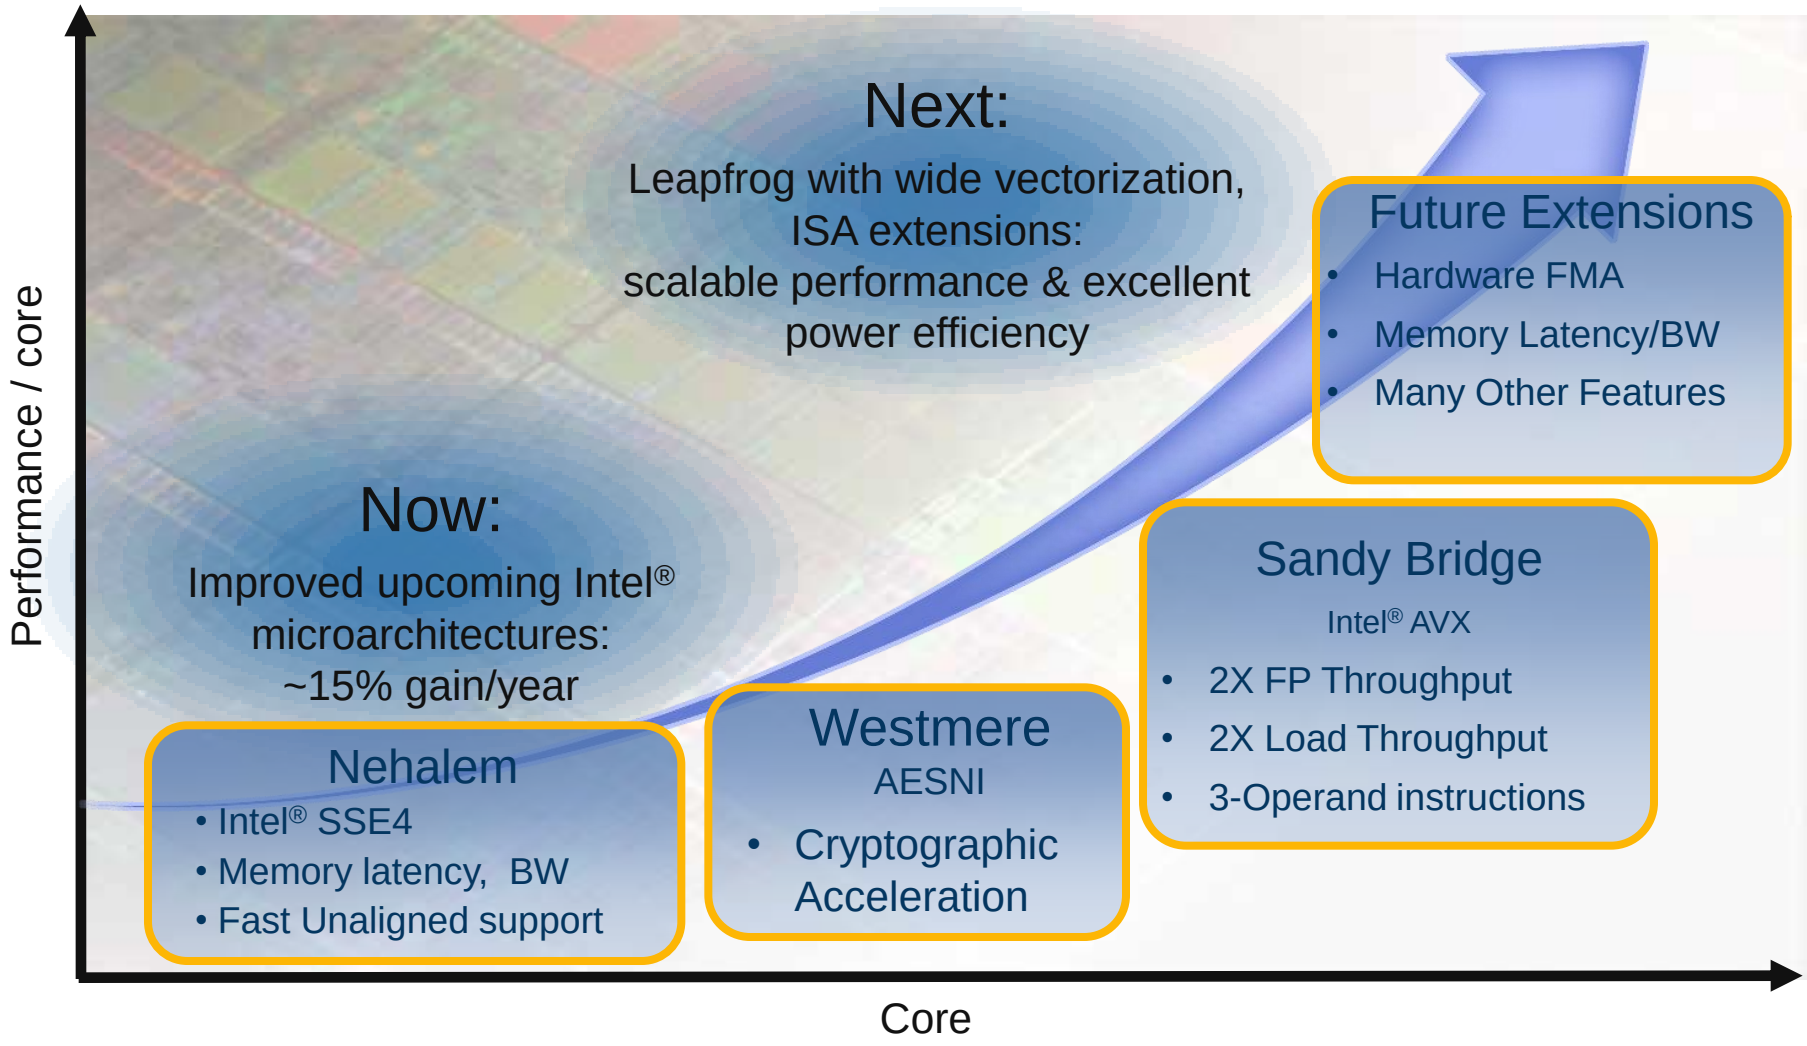

^AESNI - Advanced Encryption Standard New Instruction, Intel® AVX – Intel® Advanced Vector Extensions, Intel® SSE4 – Intel® Streaming SIMD Extensions 4.0, Intel® microarchitecture, codename Nehalem, Intel® microarchitecture, codename Westmere, Intel® microarchitecture, codename Sandy Bride

Intel® Advanced Vector Extension (AVX)

256-bit vector extension to SSE for FP intensive applications

KEY FEATURES	*BENEFITS*
Wider Vectors Increased from 128 bit to 256 bit	Up to 2x peak FLOPS output
Enhanced Data Rearrangement Use new 256 bit primitives to broadcast, mask loads and do data permutes	Organize, access and pull only necessary data more quickly and efficiently
Three Operand, Non Destructive Syntax Designed for efficiency and future extensibility	Fewer register copies, better register use, more opportunities for parallel loads and compute operations, smaller code size

Key AVX Advantages:
Most apps written with intrinsics need only recompile
Straight forward porting of existing SSE to AVX/GSSE 256 with Intel libraries, IPP, etc..
All SSE/2 instructions are extended via simple prefix (“VEX”)

http://www.intel.com/software/avx

In next year's new processor:
New Intel® Advanced Vector Extension (AVX) Instructions

New instructions that delivers up to 2x peak FLOPS output for Technical Computing

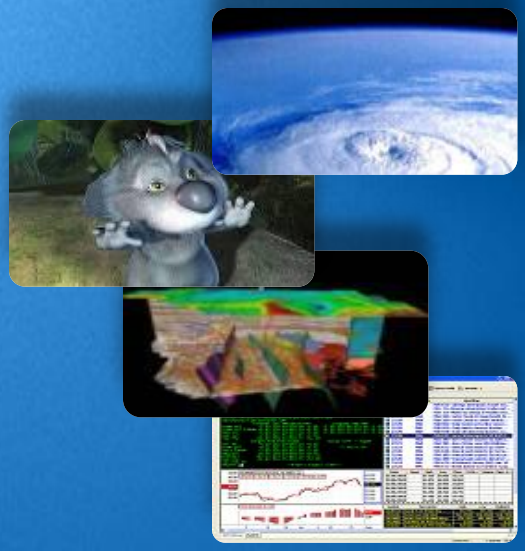

Performance:

- Vectors increased from 128- to 256-bit for increased performance on FP-intensive apps

Power Efficient:

Significantly higher performance for small incremental power

Extensible:

Backward compatible with existing apps, and designed for future ISA extension

http://www.intel.com/software/avx to get more info

Tools available for ISV enabling

Intel® AVX improves floating point & vector computation HPC applications

Trends in Multicore Evolution

- Not just adding cores
- Instead, optimum evolution of the socket architecture involves a modest increase of core count combined with architectural innovation that enhances the performance per core
  - # of ops, cache, registers, new instructions, ..
- Arguably the biggest challenge is keeping the reasonable balance we have today between memory bandwidth and flops

In addition, we have plans for compute accelerators

MANY INTEGRATED CORE (MIC) ARCHITECTURE

Intel's Many-Core Research Program

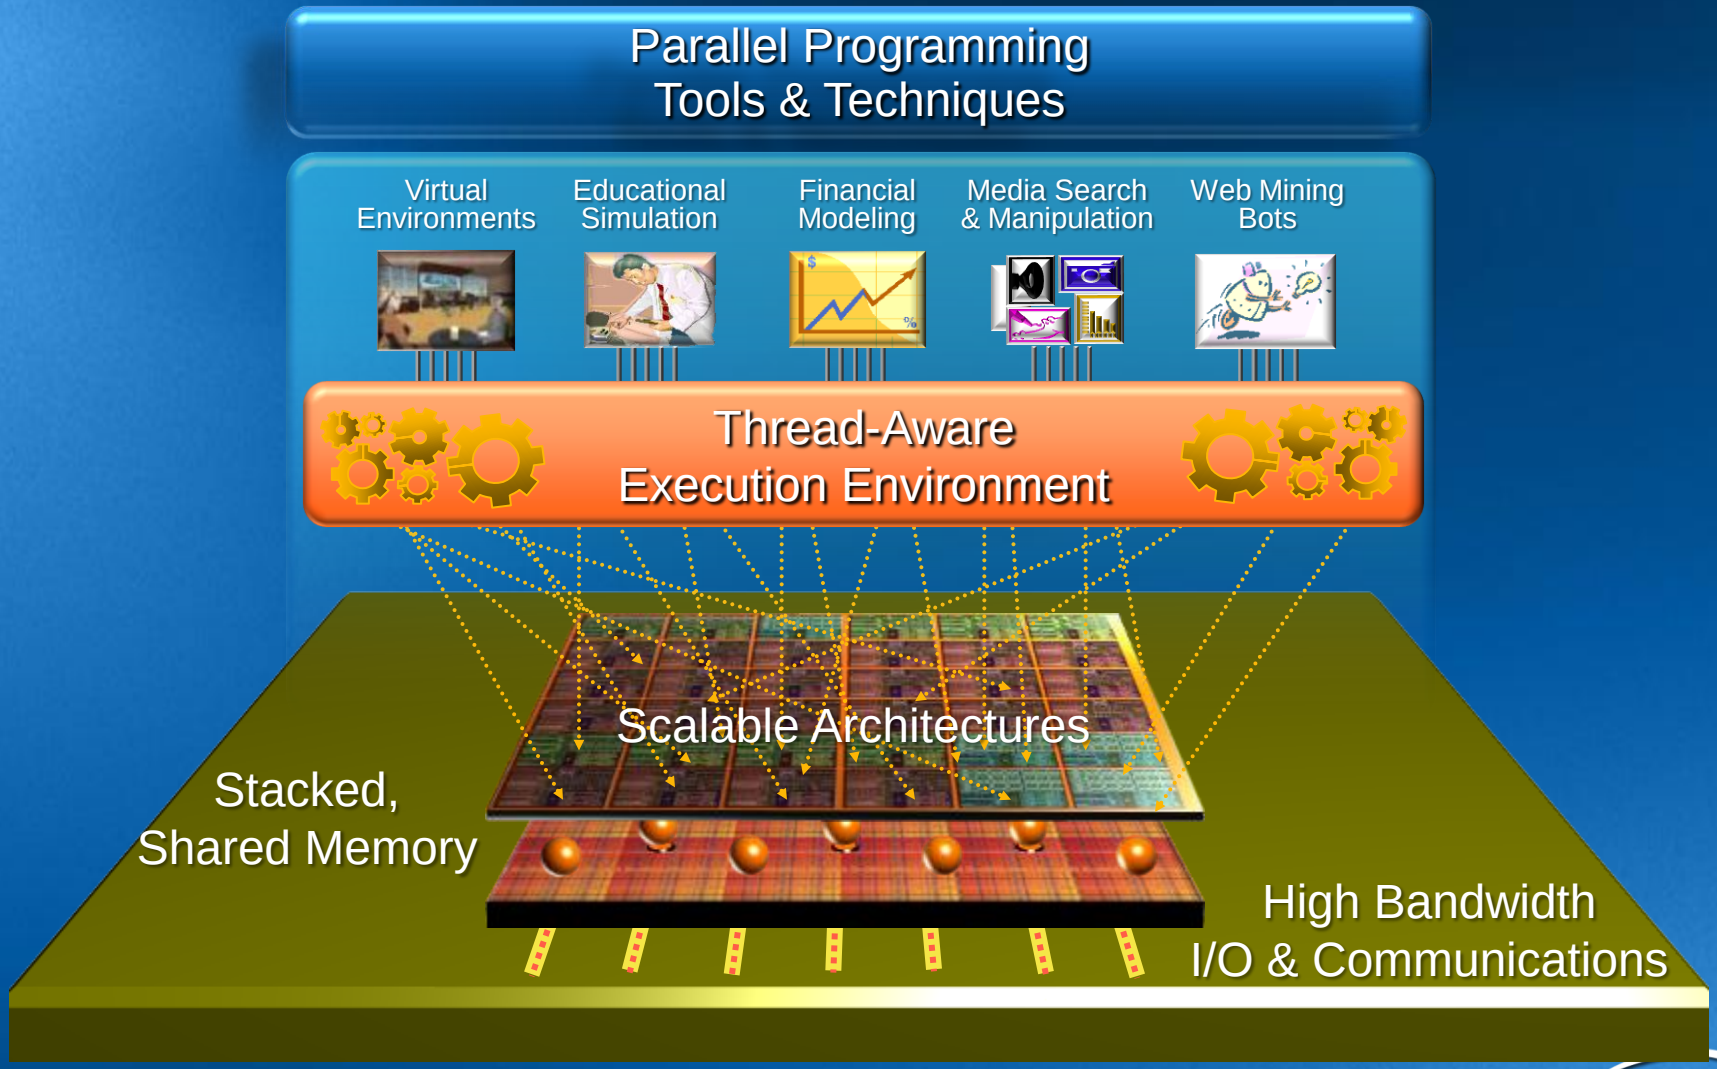

From Research to Realization.

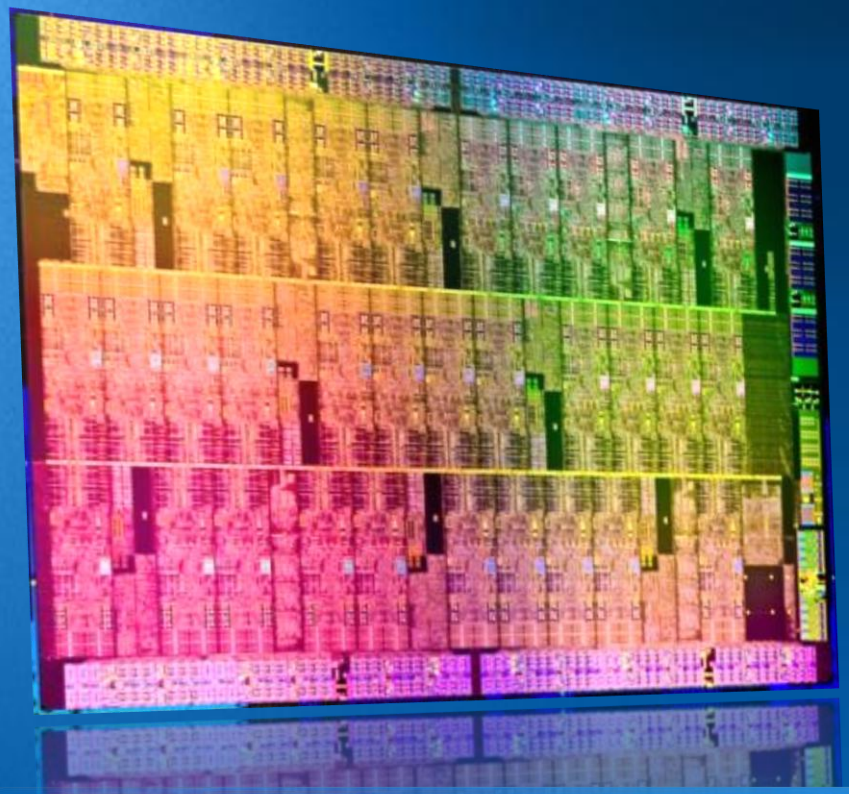

Intel®
Many
Integrated
Core
Architecture

The Newest Addition to the Intel Server Family.
Industry's First General Purpose Many Core Architecture

Application-Driven Architecture Research

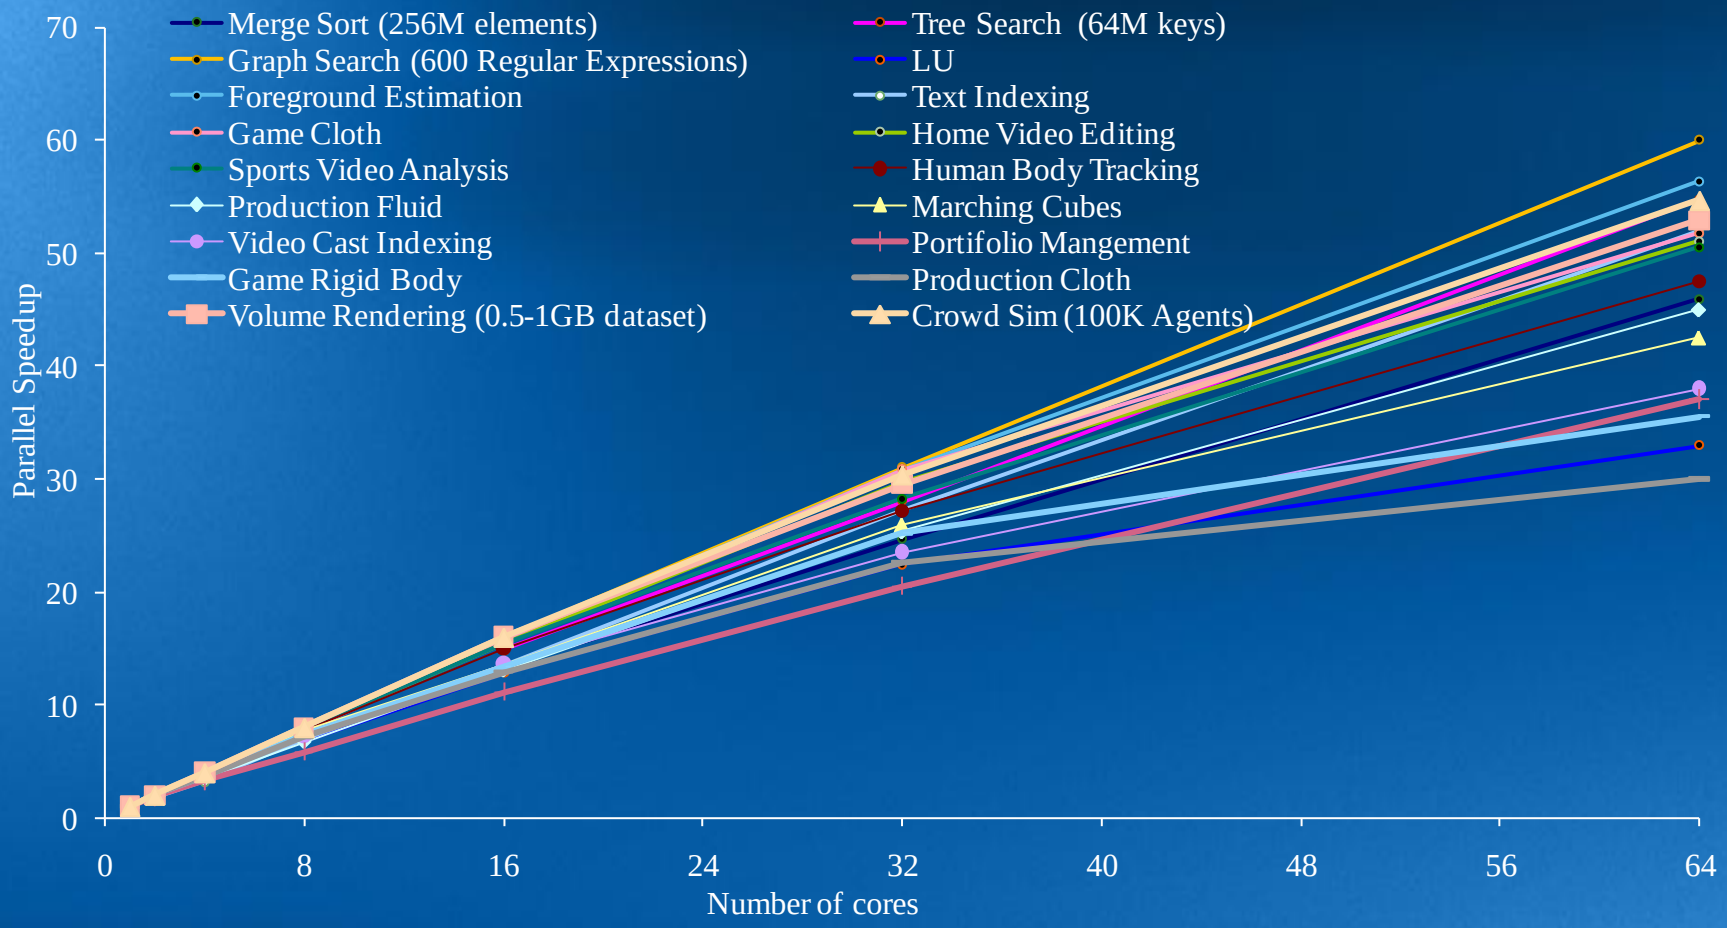

Constantly Evaluating Options for All Workloads

Our Strategy: Computing For Highly Parallel Workloads

Xeon® serves Most HPC Workloads

Xeon right for most workloads

Extending instruction set arch (AVX, etc.)

Common tools suite and programming model

Tools become architecturally aware

Ct programming model

Intel® MIC Products for highly parallel applications

First intercept 22nm process

Software development platforms available this year

Common Tool Suites insures Application Development Continuity, and Fast Time to Performance

Intel® MIC Architecture:
An Intel Co-Processor Architecture

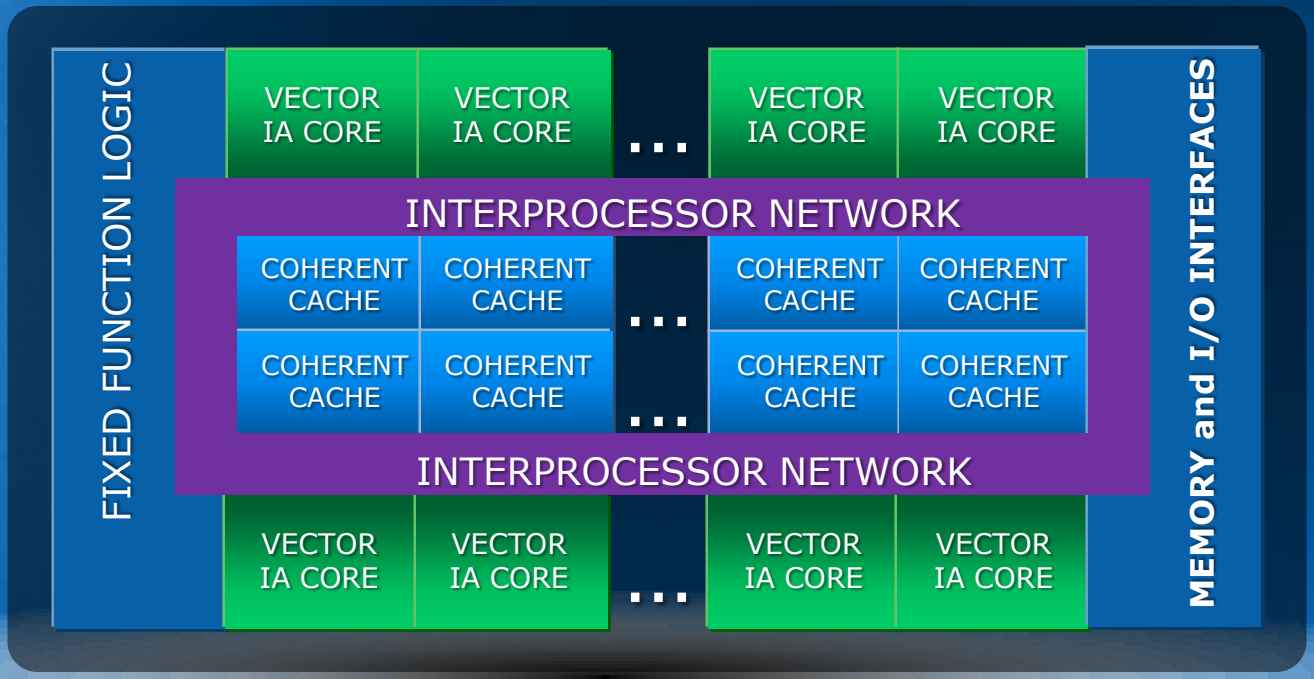

Many cores and many, many more threads

Standard IA programming and memory model

Knights Ferry

- Software development platform
- Growing availability through 2010
- 32 cores, 1.2 GHz
- 128 threads at 4 threads / core
- 8MB shared coherent cache
- 1-2GB GDDR5
- Bundled with Intel HPC tools

Software development platform for Intel® MIC architecture

The Knights Family

Future Knights Products

Knights Corner

1st Intel® MIC product

22nm process

>50 Intel Architecture cores

Knights Ferry

Intel® MIC Architecture Programming

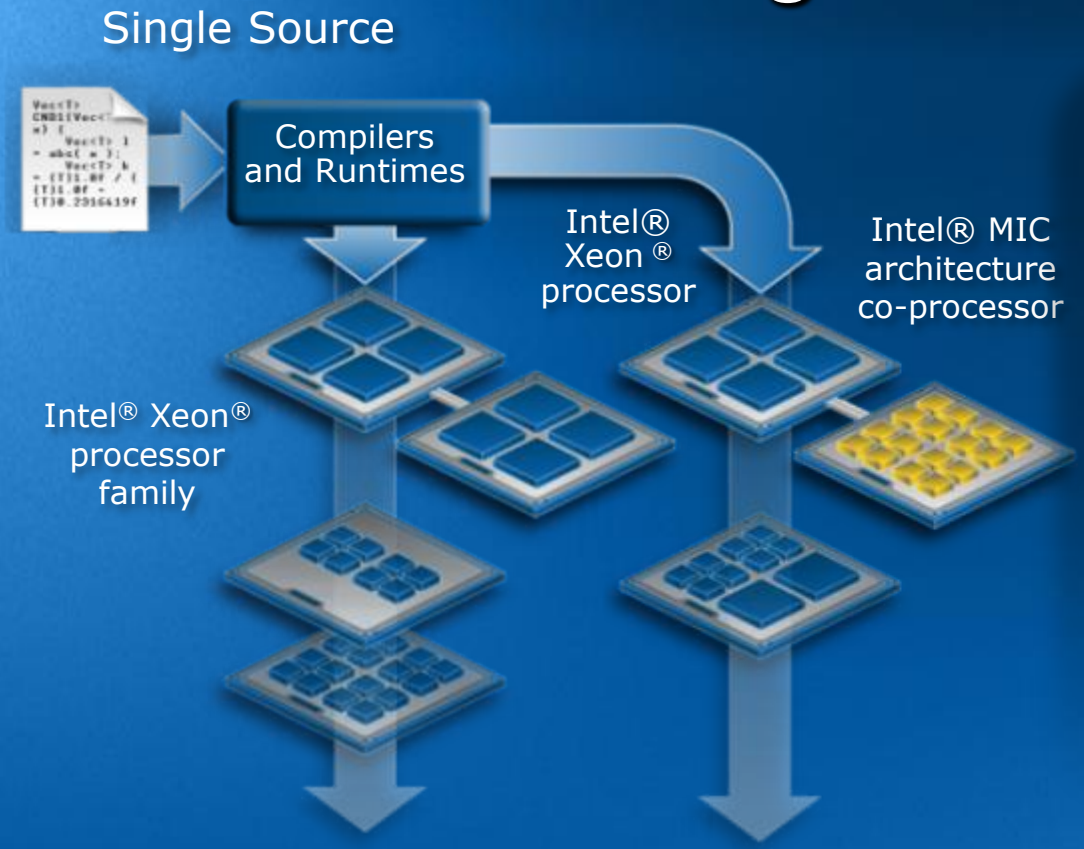

Common with Intel® Xeon®

- Languages
- C, C++, Fortran compilers
- Intel developer tools and libraries
- Coding and optimization techniques
- Ecosystem support

Eliminates Need for Dual Programming Architecture

Multicore and manycore mean more work for software and application developers..

NWS

HPC/Petascale Programming Challenges

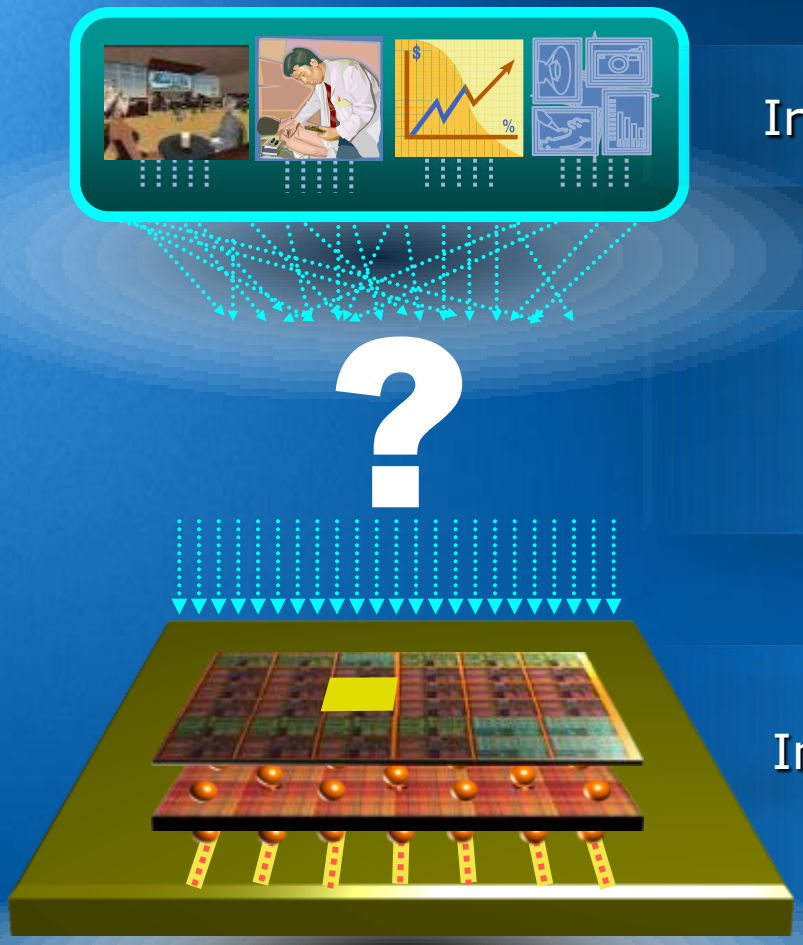

Irregular Patterns and Data Structures

Scale to Multi-Core → Hard
Scale to Many-Core→ Harder

Increasing number of cores & threads
Vector instructions

Models Supported

Application engineers assigned to track and tune major models used by the community (one person per model):

- WRF
- CESM (was CCSM)
- HOMME
- RAPS/IFS
- UM
- 4DVAR
- ECCO
- HARMONIE

Areas to investigate

- Addressing scalability to large clusters
- Hybrid programming implementations
- Use of accelerators/co-processors?

Presence in Op Centers

- In at least 8 European countries
- In major Asia countries
- Middle East
- Australia
- US agencies

  Engaged in future plans of the largest operational centers

  Much of the development of techniques and models now done on x86 clusters

Looking further out .. Up to circa 2018

CHALLENGES FOR FUTURE LARGE SYSTEMS (UP TO EXASCALE)

Meeting Today's HPC Challenges

- Performance
- Ease of Development
- Confidence/Trust

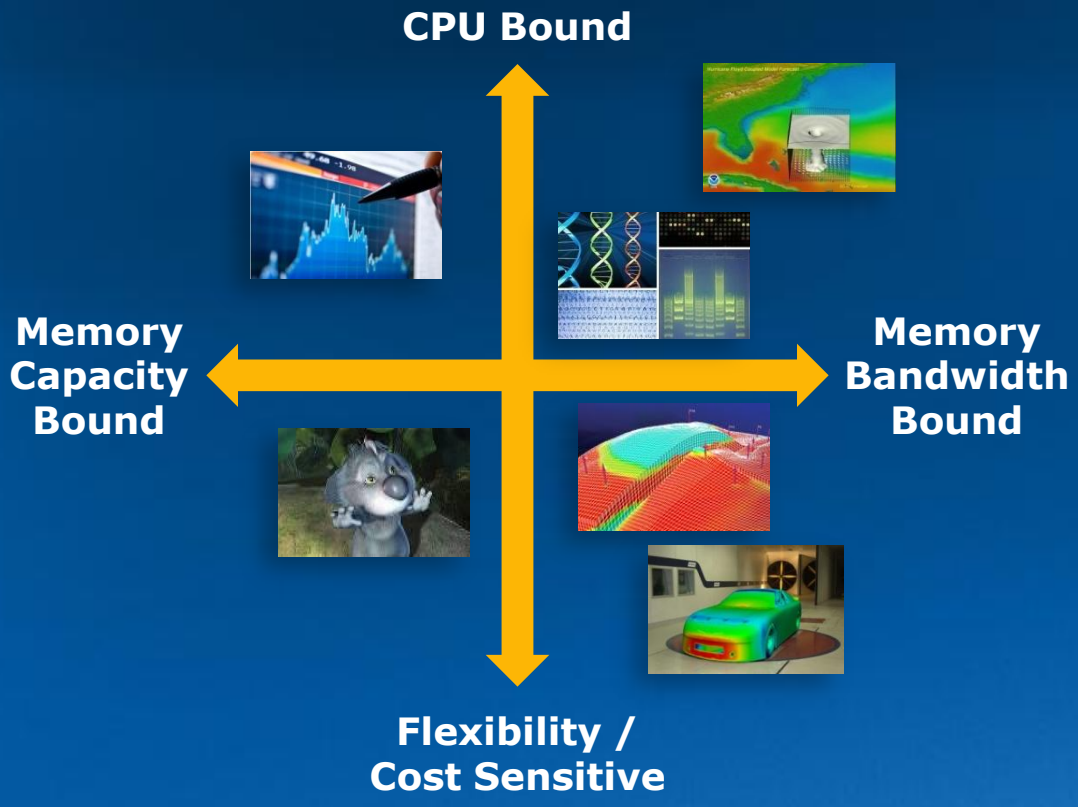

For demonstration purposes only

Versatile, scalable solutions to enable innovation

Still an Insatiable Need for Computing

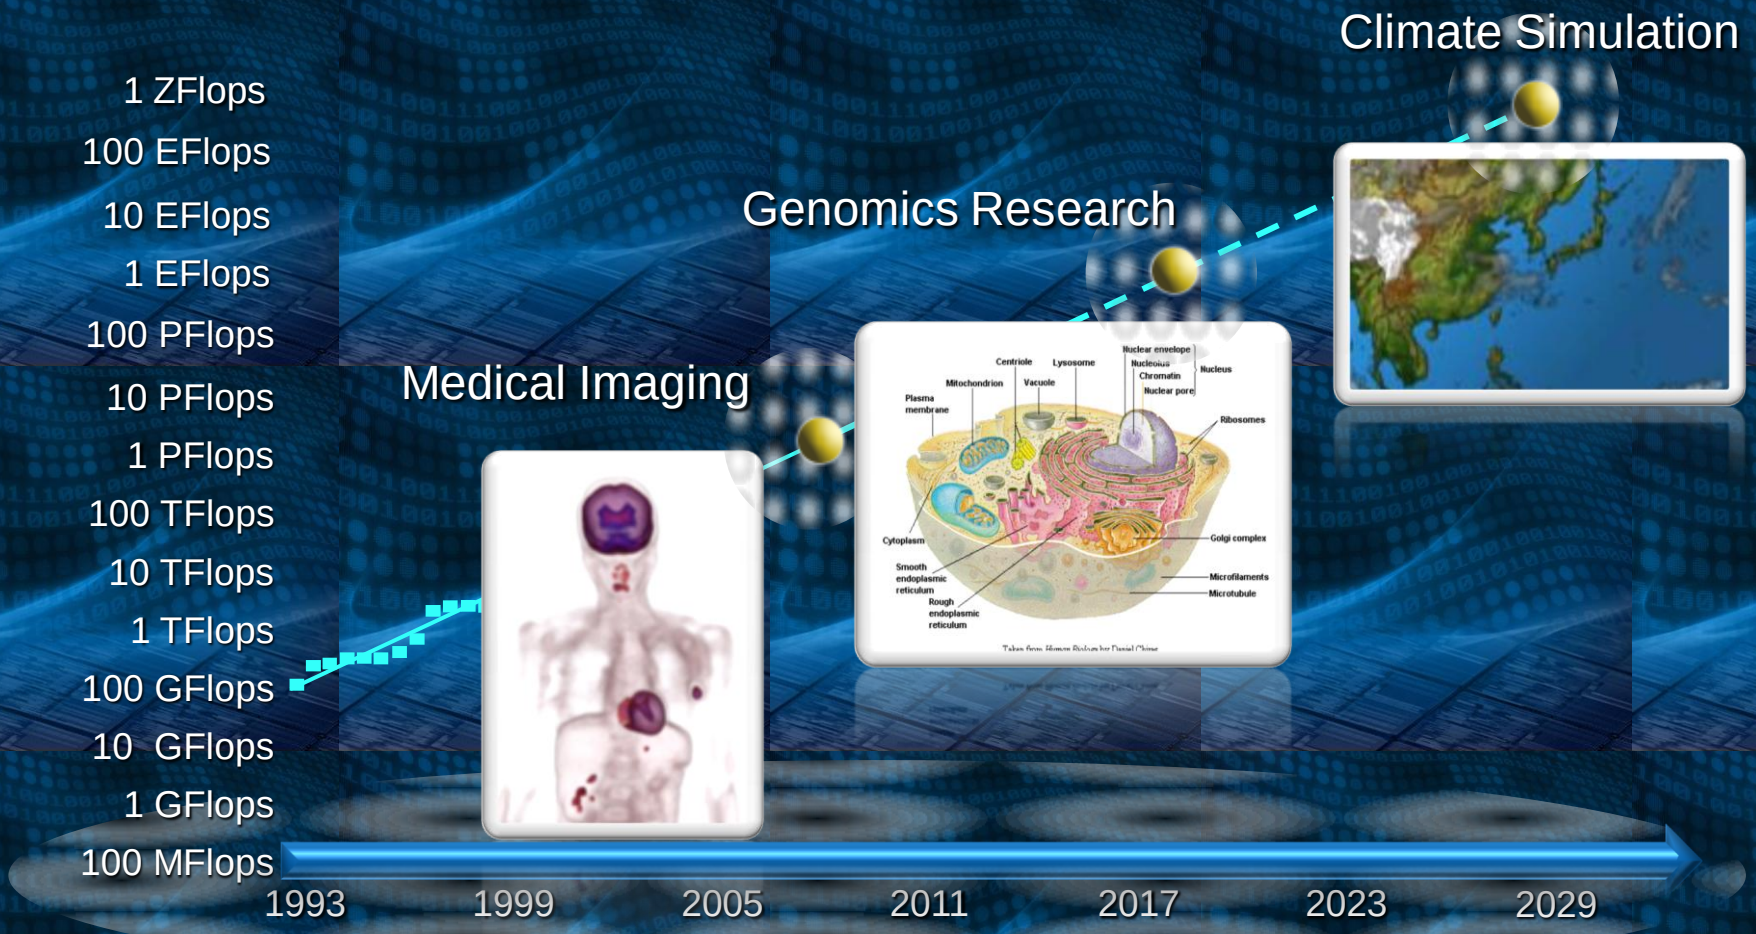

Source: www.top500.org

Programming at Exascale

- New languages needed?
- First: actually support and use the ones we have
- New Runtime Systems and OS approaches are needed
- Resiliency cannot be the primary focus of programming
- Get rid of Bulk Synchronous Programming
  - It is wasteful of resources and power
- Do algorithms need to inform and be informed by system SW and HW (e.g., resources going south, jitter, …)?
- We may be approaching a crossover point where probabilistic methods come into their own

Resiliency

- Combine
  - new approaches to classic reliability and resiliency
    - Advanced multi-scale/multi-frequency checkpointing
    - Pervasive RAS monitoring by RTS
    - Redundancy or elimination of high-failure parts
  - With
    - Resilient circuit technologies
    - Ckt level detection and correction
    - Residue checking and redundant multi-threading
    - Abstract Machine to virtualize resources scalably

30

The Path Forward

Research Needed to Achieve Exascale Performance

- Extreme voltage scaling to reduce core power
- More parallelism 10x – 100x to achieve speed
- Re-architecting DRAM to reduce memory power
- New interconnect lower power and distance
- NVM to reduce disk power and accesses
- Resilient design to manage unreliable transistors
- New programming tools for extreme parallelism
- Applications built for extreme parallelism

12

31

Intel Co-Sponsored HPC Labs in Europe

ExaTec Lab, Paris

Performance and scalability of Exascale applications

Programming models
Performance analysis
Applications performance

ExaCluster Lab, Jülich

Exascale cluster scalability and reliability

System management
Software tools
Simulation

ExaScience Lab, Leuven

Exascale Science power and reliability

Space weather app
Programming models
Power & reliability

Advancing Exascale Computing on Intel Architecture

Co-design is Essential

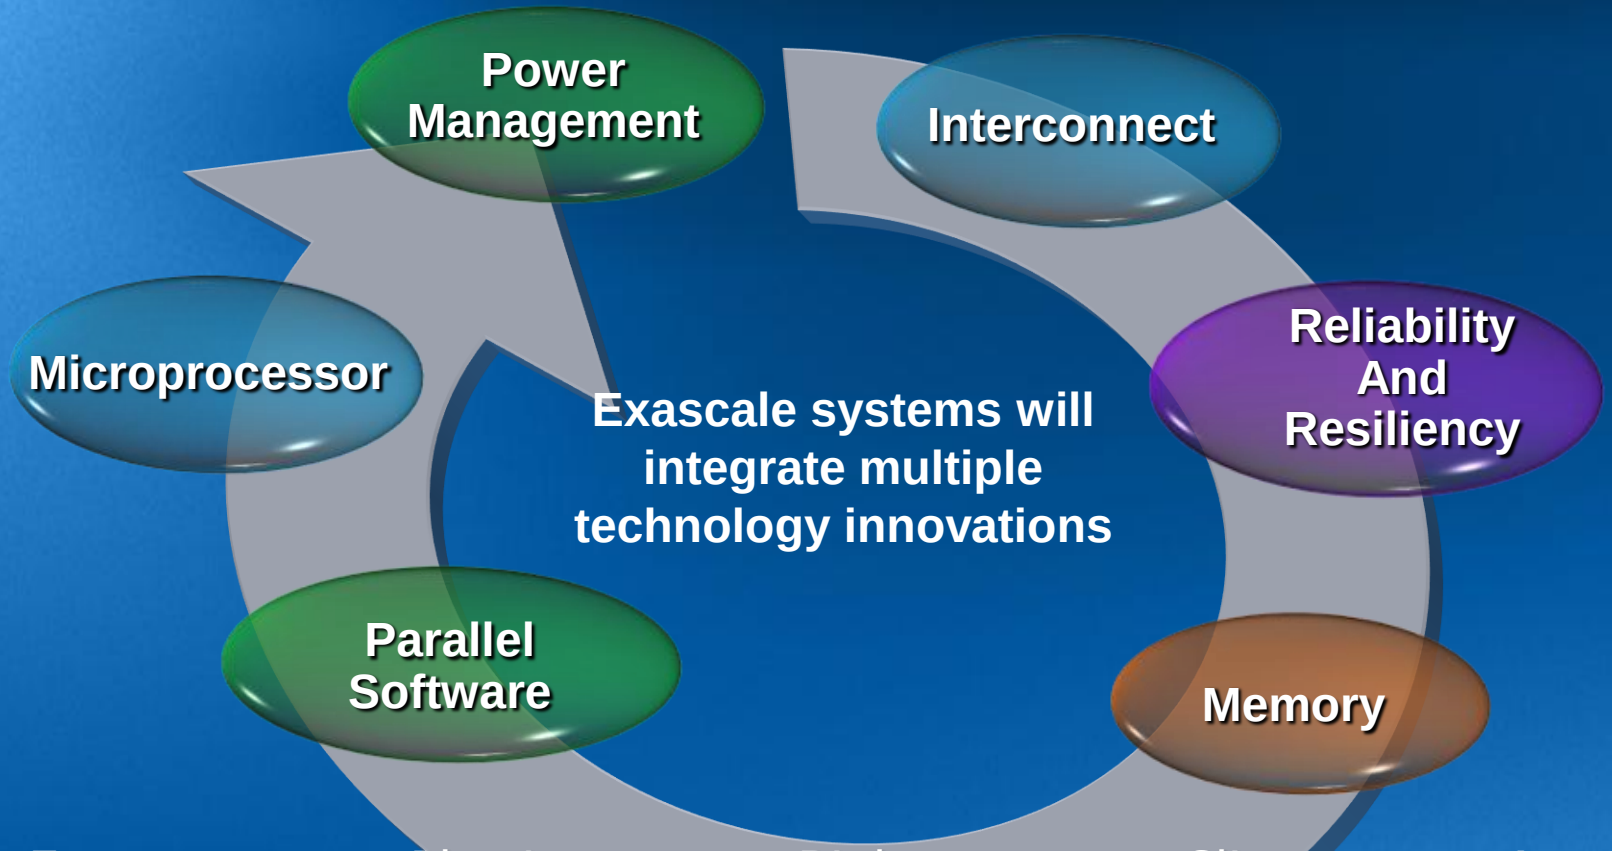

Energy | Physics | Biology | Climate | Astrophysics

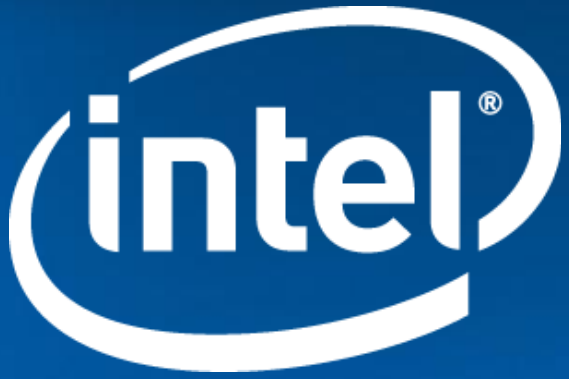
intel


intel