Manycore Computing and MIC

OFA
2011 International Monterey Workshop

Pradeep K. Dubey

IEEE Fellow and Director of Parallel Computing Lab
Intel Corporation

Parallel Computing Applications

Illustrative Parallel Computing Apps

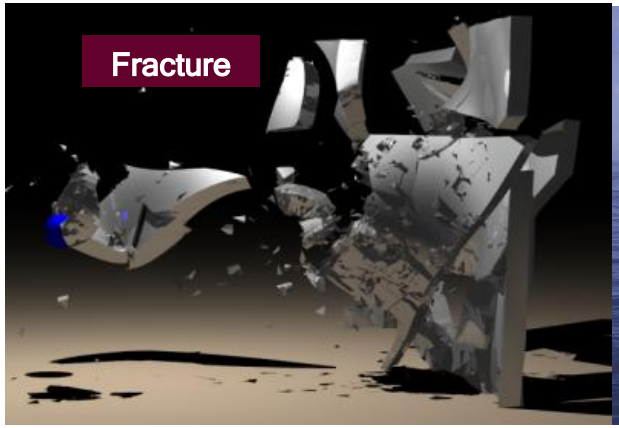

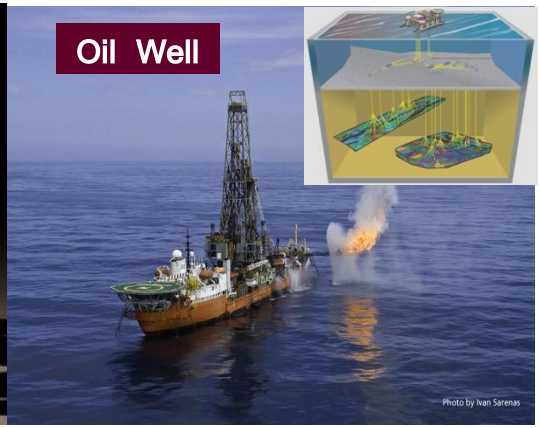

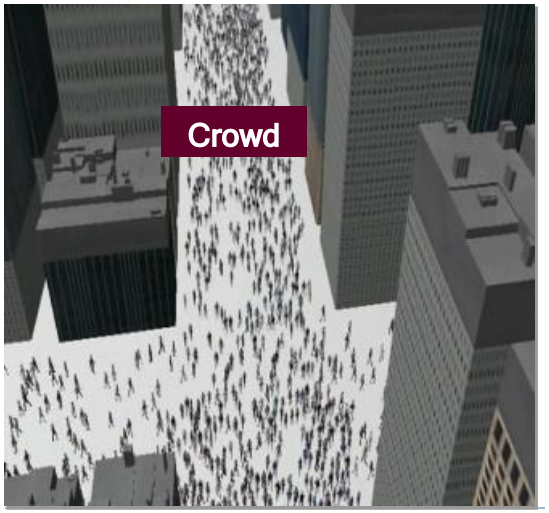

More the better …

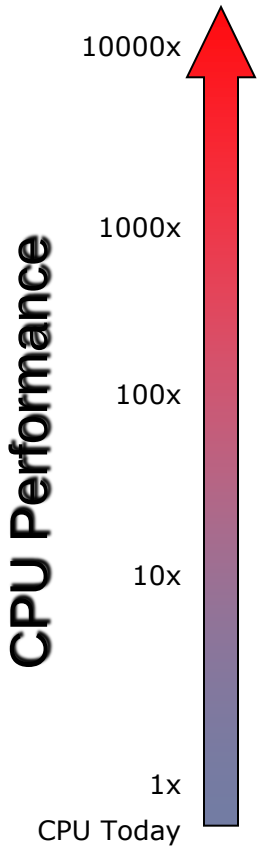

- Haptic dynamics in haptic training apps
  - System fully usable in the operating room
- Real-time implicit simulation for 100K elements
  - For real-time training tools & very accurate prediction
- Interactive quasi-statics simulations of 100K elements
  - Good usability for planning and prediction
- Offline dynamic Simulations of 100K elements
  - Limited usability for prediction
- Offline quasi-statics simulations of 10K elements
  - Impractical for use in clinical environments

100K - 1M elements

10 - 100K elements

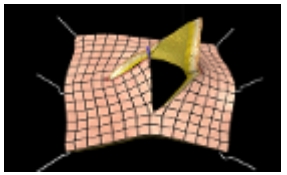

1 - 10K elements

Force simulations for visual rendering: 10s of Hz
Force simulations for haptic rendering: KHz or more

Entertainment to Interventional Medical Imaging: Physics plays a critical role and drives compute!

Estimating the Exaflood, Discovery Institute, 1/08; *Amassing Digital Fortunes, a Digital Storage Study*, CEA, 3/08

A Wealth of Data to Move

Personal Media	Business	Medical	Social Media	Science
				
Ave. Files on HD 54GB	Retail Customer DB 600 TB	Clinical Image DB ~1PB	HD video forecast 12 EB/yr	Physics (LHC) 300 EB/yr

More than 15B connected devices by 2015

Massive Data & Ubiquitous Connectivity

- Data-driven models are now tractable and usable
  - We are not limited to analytical models any more
  - No need to rely on *heuristics* alone for unknown models
  - Massive data offers new algorithmic opportunities
    - Many traditional compute problems worth revisiting
- Web connectivity significantly speeds up model-training
- Real-time connectivity enables continuous model refinement
  - Poor model is an acceptable starting point
  - Classification accuracy improves over time

Nested RMS

Recognition	Mining	Synthesis
What is …?	Is it …?	What if …?

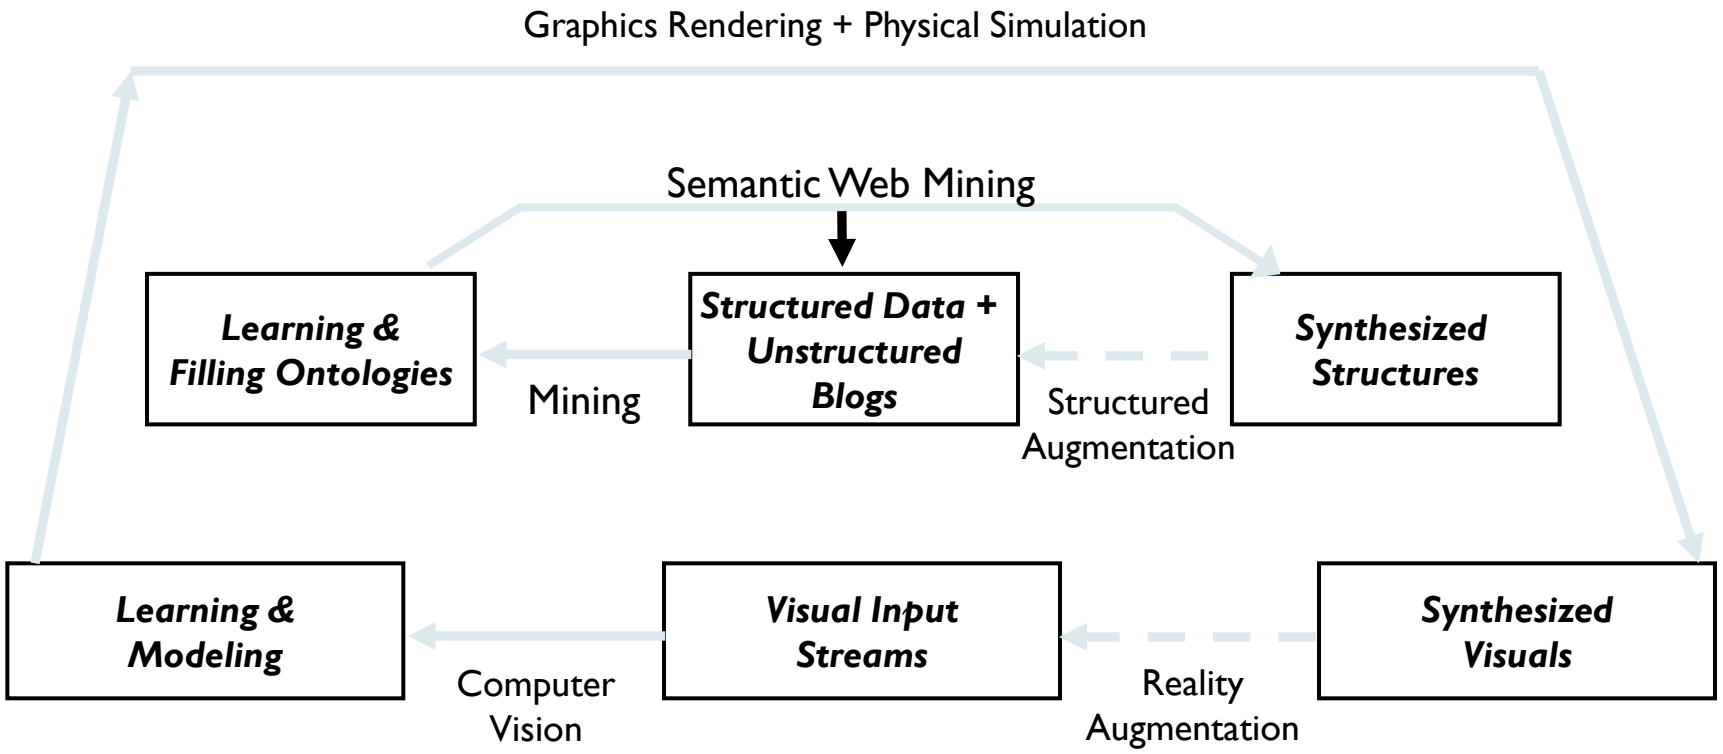

Pradeep.Dubey@intel.com

Where is my computer ☺

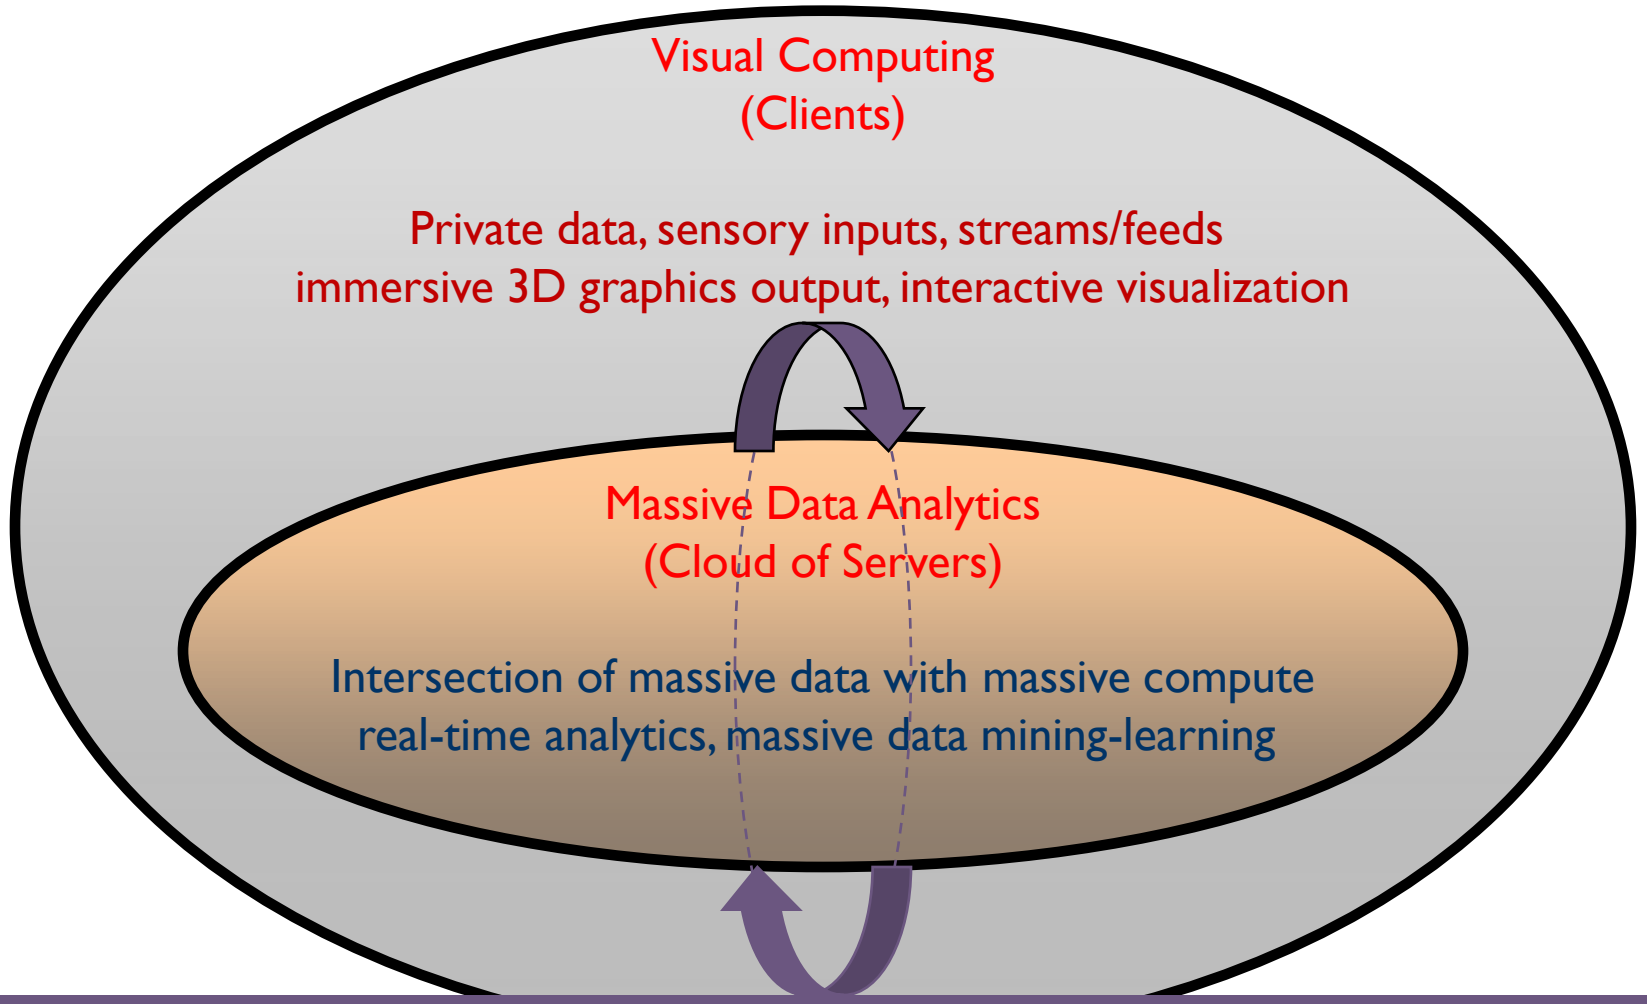

Architectural Implications Are Radical!

Pradeep.Dubey@intel.com

MIC Architecture and Tools

Multicore Versus Manycore

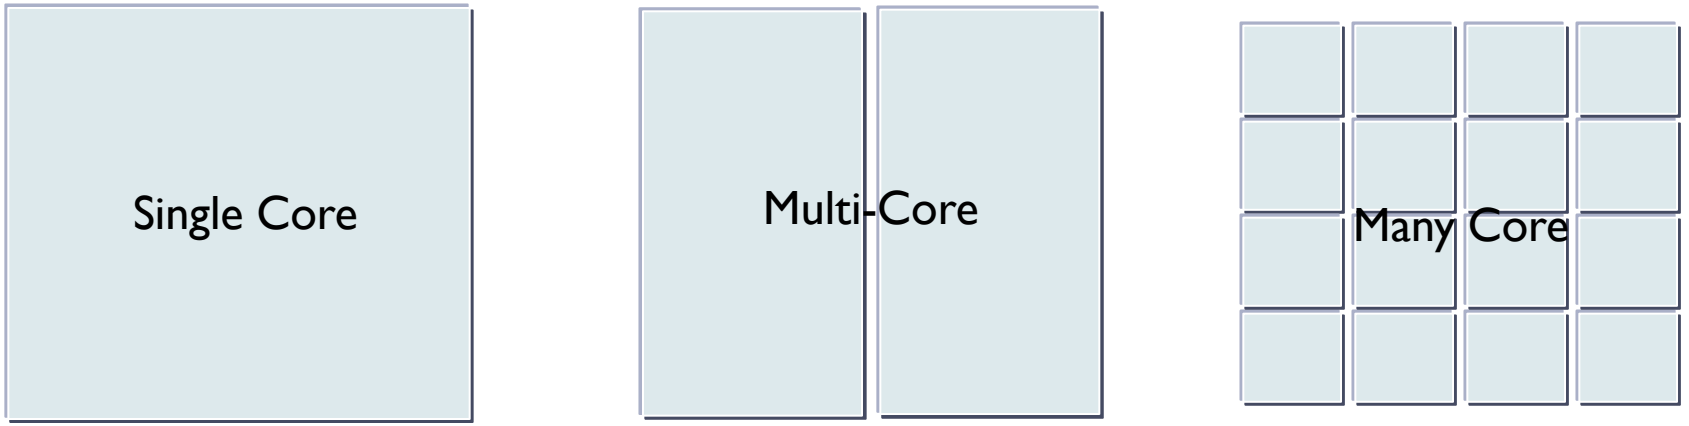

Many Core makes sense for workloads with high enough "P "- parallel component - for simplicity, we call these Highly Parallel

$$S = \frac{1}{(1-P) + \frac{P}{N}}$$

$$S = \frac{1}{(1-P)K_N + \frac{P}{N}}$$

$$For\ S \geq 1,\ P \geq \frac{N(K_N - 1)}{NK_N - 1}$$

S = speedup, P = parallel fraction, # of Cores = N, Kn = single thread performance (single core/multicore)

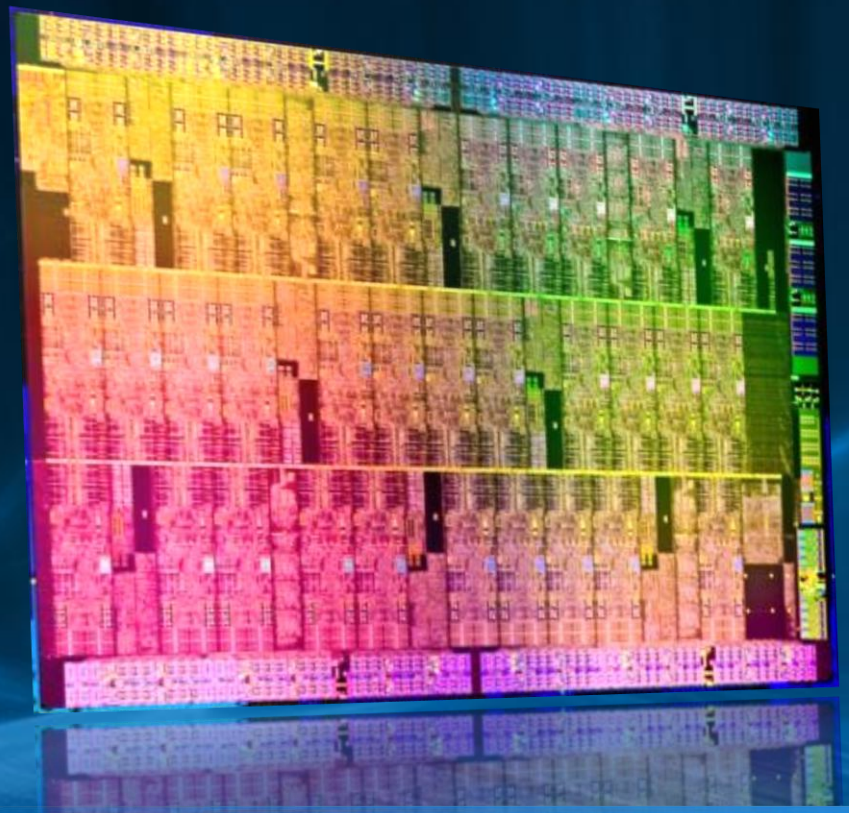

Intel® Many Integrated Core Architecture

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

intel

Intel® Many Integrated Core (Intel® MIC) Architecture

- Up to 32 Intel coherent Intel processor cores on 1 silicon die
- Implements all four salient architectural features of Intel® CPUs
  - x86 Cores, Coherent caches, SIMD, SMT threads
- Enables developers to *scale* applications forward to future Intel® MIC products

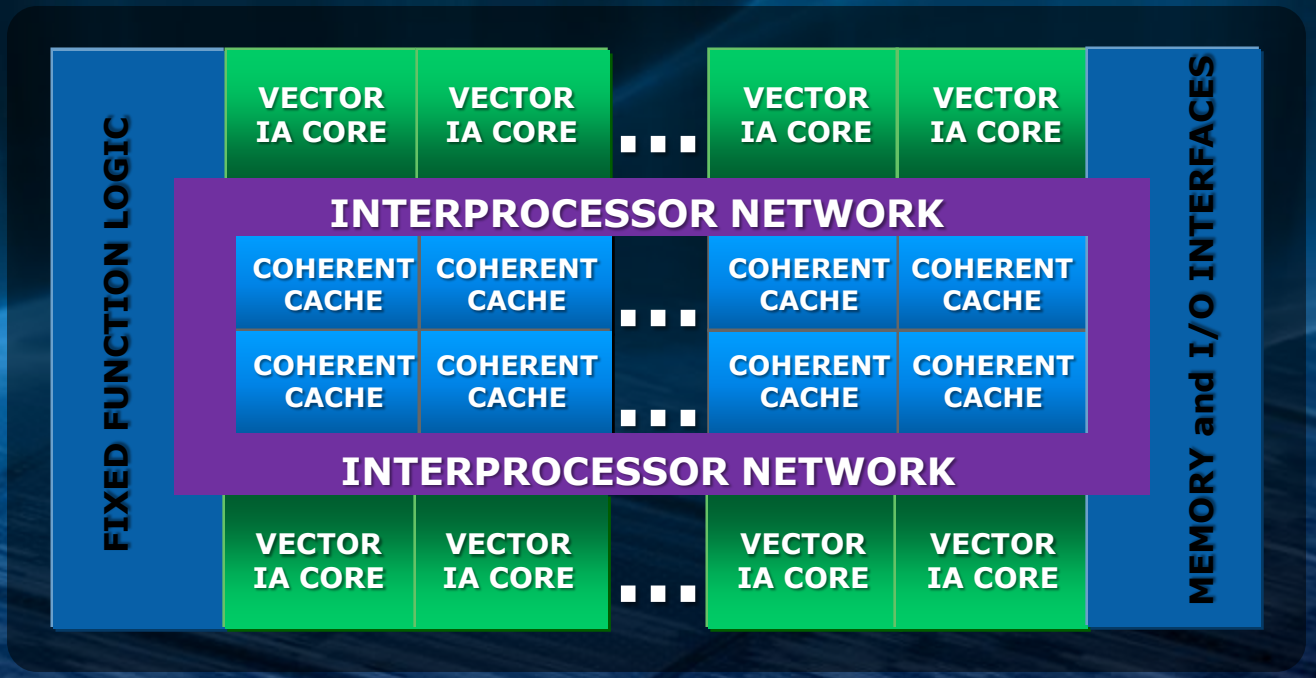

Knights Ferry

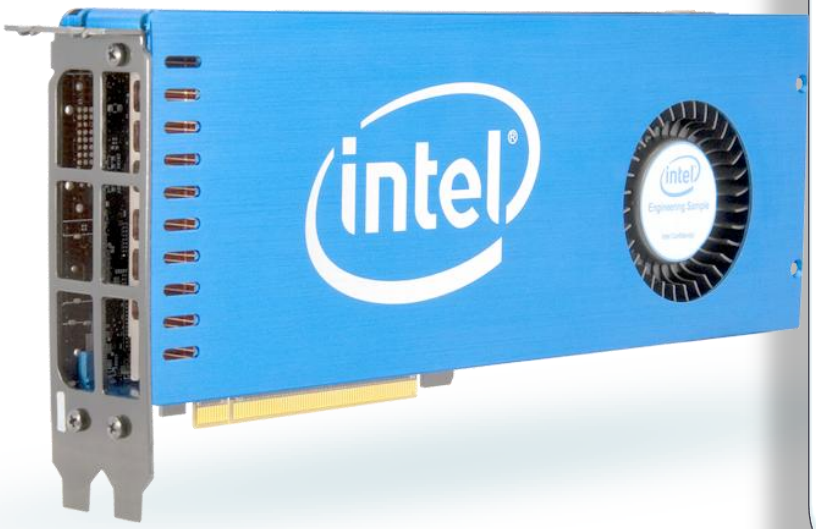

Software development platform for Intel® MIC Architecture includes:

PCIe 2.0 Adapter Card

- 1 or 2 Gigabytes of GDDR
- Aubrey Isle Intel® MIC processor with up to 32 cores and 128 threads, and 8 megabytes of coherent L2 cache
- Designed for systems that support 300W PCIe design guidelines

Intel Software Development tools

Availability beginning in 2H 2010

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's Development Tools for HPC

Leading developer tools for performance on nodes and clusters

Advanced Performance

C/C++ and Fortran Compilers, Intel® Math Kernel Library (Intel® MKL)/Intel® Integrated Performance Primitives (Intel® IPP) Libraries & Analysis Tools for Windows*, Linux*, Mac OS* developers on Intel® architecture multi-core node

Distributed Performance

MPI Cluster Tools, with C++ and Fortran Compiler and Intel MKL Libraries, and analysis tools for
Windows, Linux developers on Intel architecture clusters

Intel® MIC Architecture Programming

Single Source

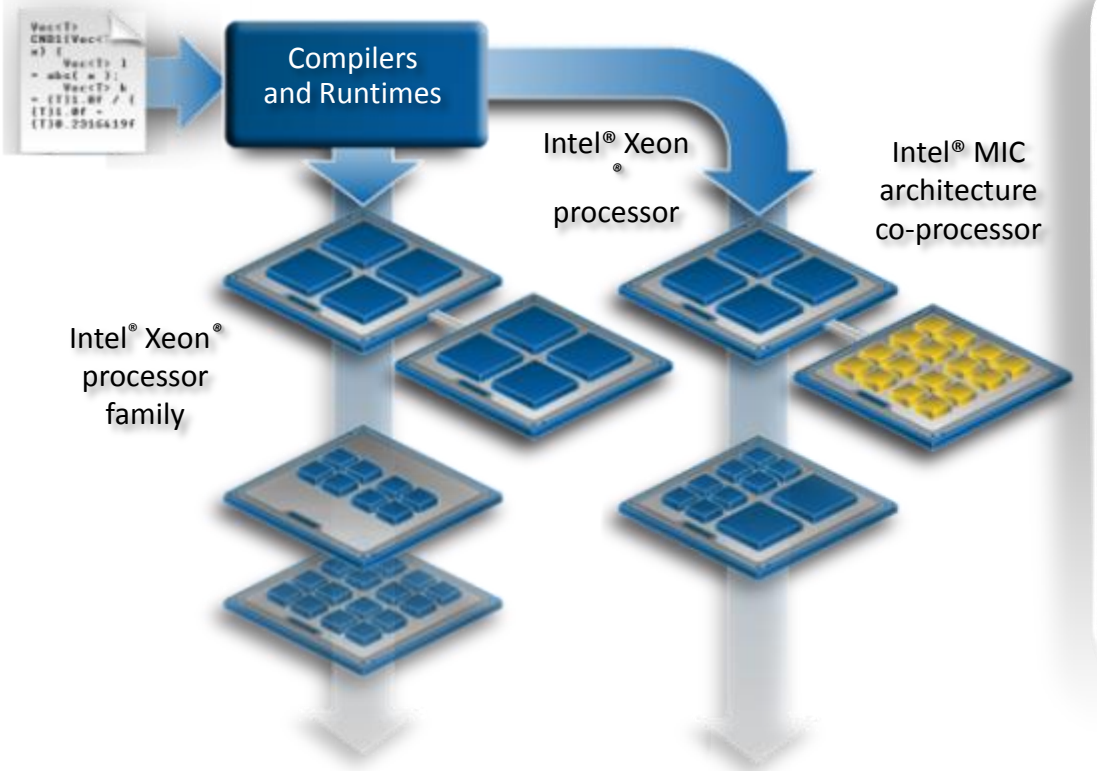

Common with Intel® Xeon®

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

Eliminates Need for Dual Programming Architecture

Sample Application: Compressed Sensing

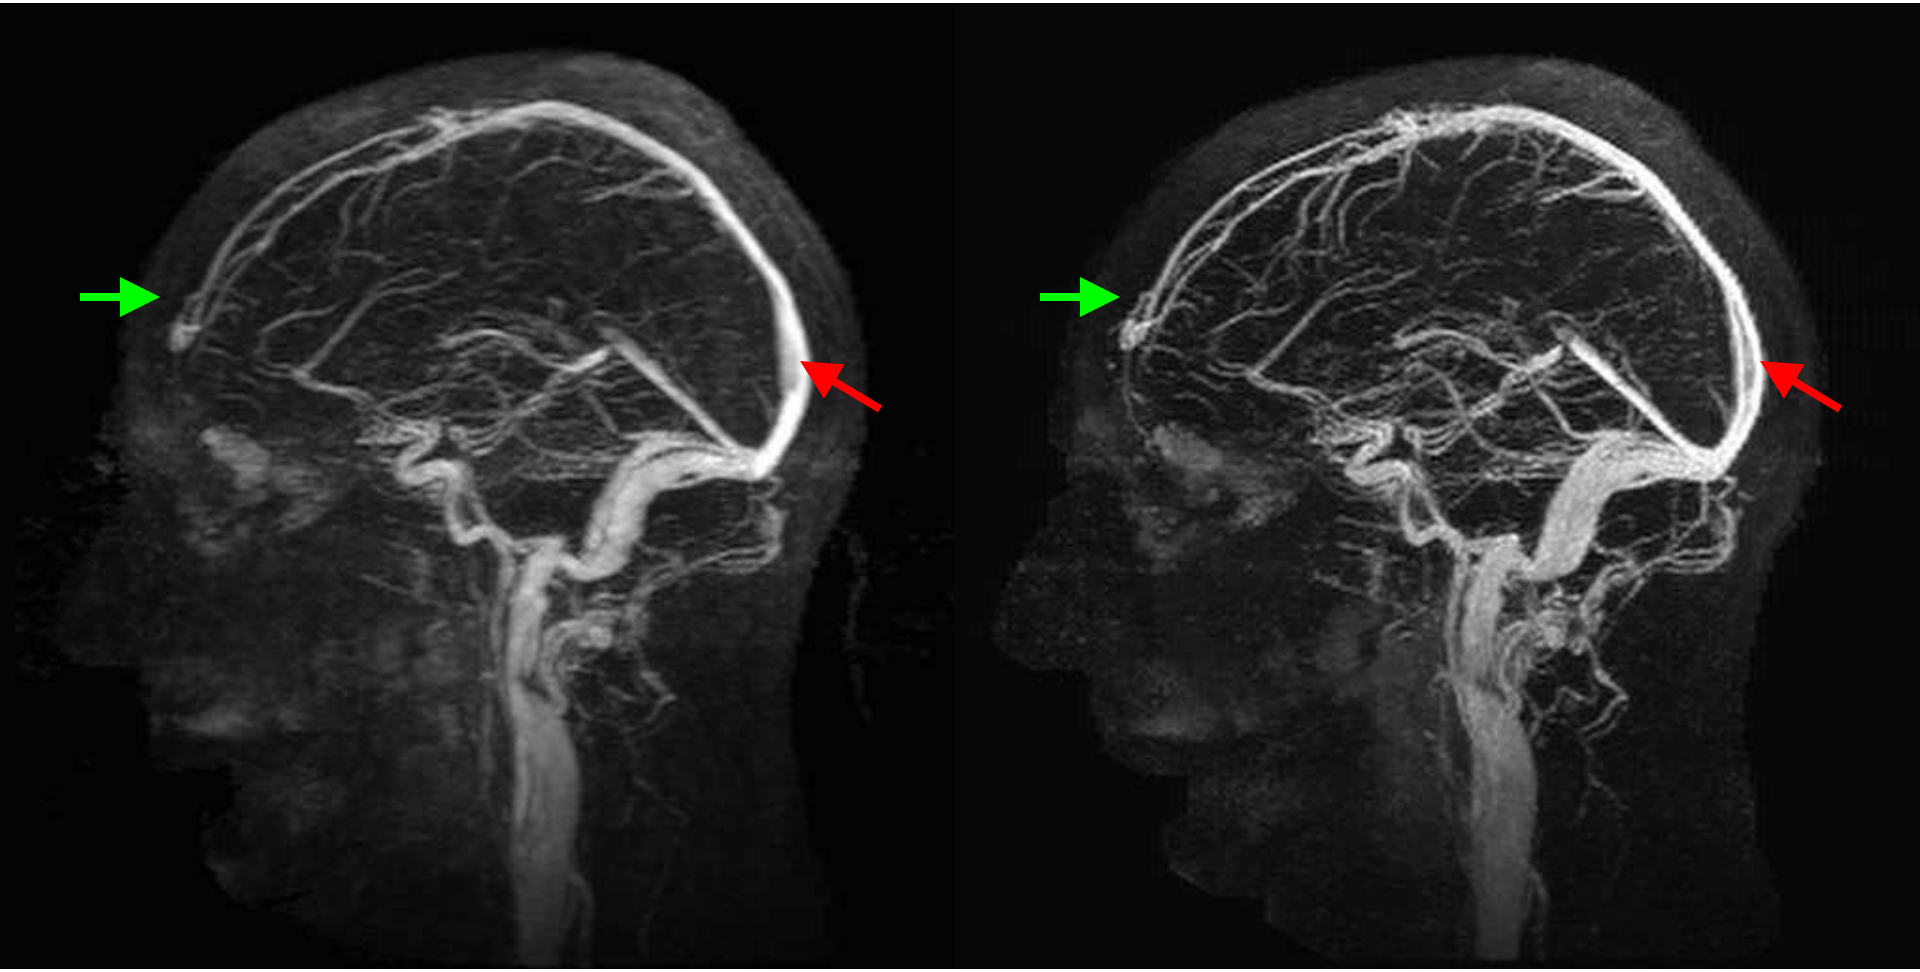

Current Clinical (SENSE+PF)

Compressive Sensing

Single time frame from a CAPR CE-MRA exam (8-channel, R=19x, 256x160x80) [Trzasko2010]

*"High-Performance 3D Compressive Sensing MRI Reconstruction" International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS 10)., Daehyun Kim, Joshua D. Trzasko, Mikhail Smelyanskiy, Clifton R. Haider, Armando Manduca, and Pradeep Dubey.

Backup

Summary

To Be Added

- Performance variability on the rise with parallel architectures
- *Feeding the Beast: increasingly a performance bottleneck*
- Programmer productivity key to market success