

# MIC, Intel and Rearchitecting for Exascale



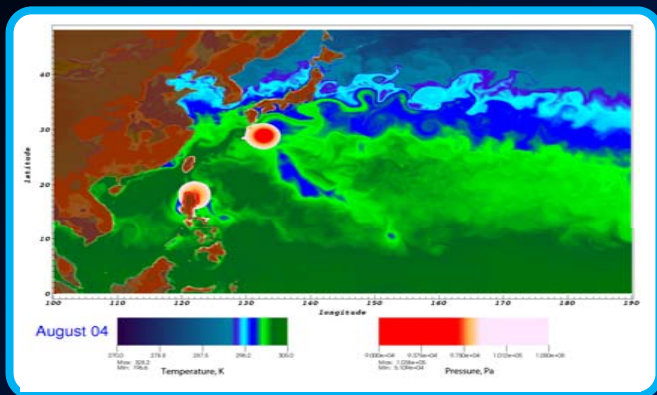
John Hengeveld  
Director of Marketing, HPC Evangelist  
Intel Data Center Group

Dr. Jean-Laurent Philippe, PhD  
Technical Sales Manager & Exascale Technical  
Lead  
Intel EMEA Sales and Marketing Group




# Exascale Answers Mankind's Challenges In...


*Weather / Climate*




*Healthcare*



*New Forms of Energy*



# An Insatiable Need For Computing



*Exascale Problems Cannot Be Solved Using the  
Computing Power Available Today*

# *Intel Commitment To Exascale*

Efficient  
Performance



Programming  
Parallelism



Extreme  
Scalability



*Intel Exascale Commitment:  
>100X Performance Of Today At  
Only 2X The Power of Today's #1 System  
Scaling Today's Software Model*




# What will Exascale Workloads Look like?

- Massively Parallel (assume 100x threads)
  - Assume low growth to total HPC SW talent (maybe up 50%)  
 $100x \text{ threads} / 1.5x \text{ engineers} = 66x \text{ today's threads per head}$
  - SW methodology must absorb this productivity requirement
- Memory Intensive – some will be... some will not be
  - What Bytes per flop ratio is sufficient?
  - The cache debate is in full swing.
  - Alternative memory architectures?




# Process Technology: First 16x of the 100x – 6x to go....

2007  
45 nm




Invented  
Gate-Last  
High-k  
Metal Gate

2009  
32 nm



2nd Gen.  
Gate-Last  
High-k  
Metal Gate

2011  
22 nm



First to  
Implement  
Tri-Gate

STRAINED SILICON

HIGH-k METAL GATE

TRI-GATE

*With Trigate...  
Skyscraper  
Versus  
Ranch Houses*

*Semiconductor Likely  
Scales through Exascale*

*2x Performance each  
2 years (roughly)  
8 years... 16x*

*The foundation for all computing*






# Breakthroughs Required to Get There

	Delivered Performance in 8 years	Issue
Moore's Law	16x	
Increased Node Scaling	(1.5x?)	Interconnect Performance Scales?
HW architecture – Performance per Watt improvement	5x?	Memory Performance Scales? Storage Performance Scales?
SW Complexity	Probably increases 100x	Methodology gap?




# Many Core and Multi-Core (~4x)

Many Integrated Core Aubrey Isle at 1-1.2 GHz



Multi-core Intel® Xeon® processor at 2.26-3.5 GHz



Die Size not to scale

In Intel® MIC architecture, each core is smaller, has lower power limit, has lower single thread performance, but higher aggregate performance

Many core relies on a high degree of parallelism to compensate for the lower speed of each individual core





# Intel Labs & HPC (lots of other Issues)

## Strong Research Partnerships

### Industry



### Government




### Universities



## World Class Research in HPC

### Memory Stacking & Technologies



### Silicon Photonics




### Security



### Programmability



### Interconnect Technologies



### Power Reduction



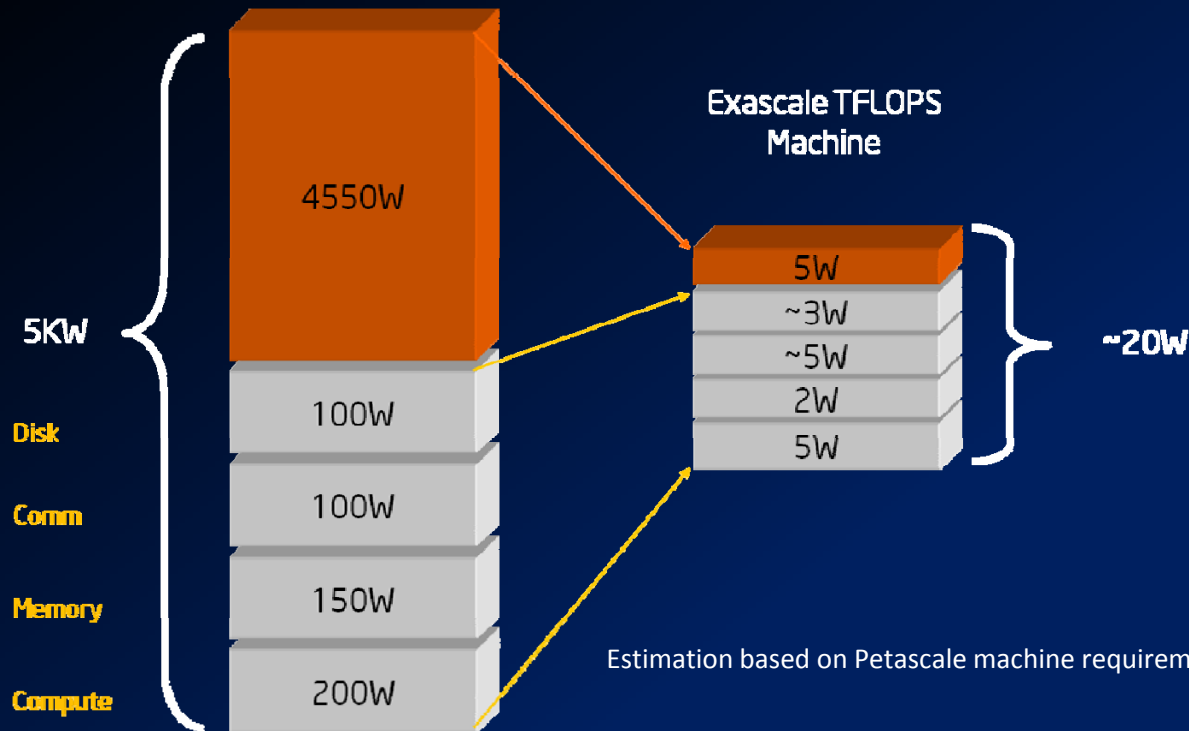
*Delivering Breakthrough Technologies to Fuel Innovation*



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

# Exascale Requirements

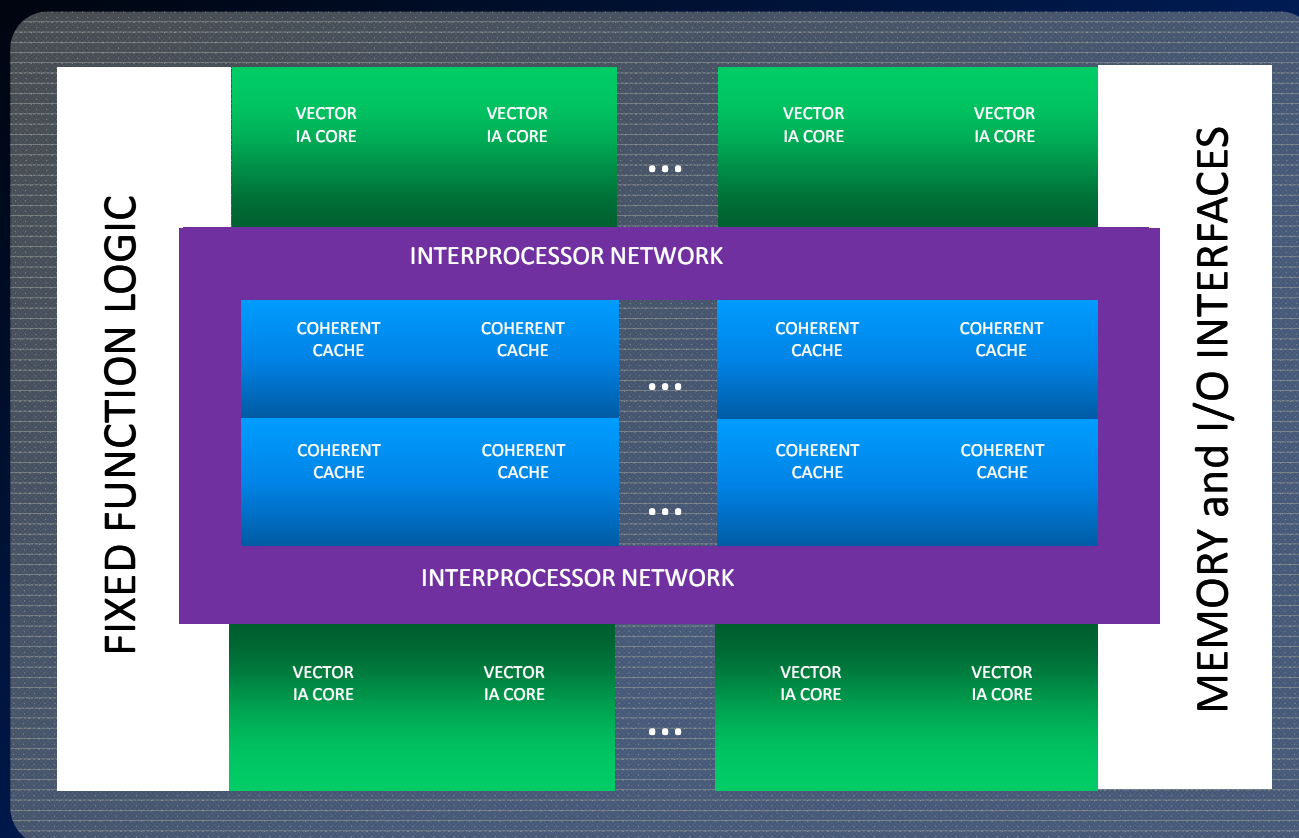
Petascale Machine of 2010: TFLOP of Compute



Compute 40x  
Memory 75X  
Comms 20x  
Disk/Storage 33x  
Other 900x

*Visceral Focus on System Power Efficiency Improvement*

# Aubrey Isle Co-Processor Architecture



Multiple Fully Functional x86 cores

- In-order, short pipeline
- Multi-thread support

16-wide vector units (512b)

- Extended instruction set
- Fully coherent caches

1024-bit ring bus

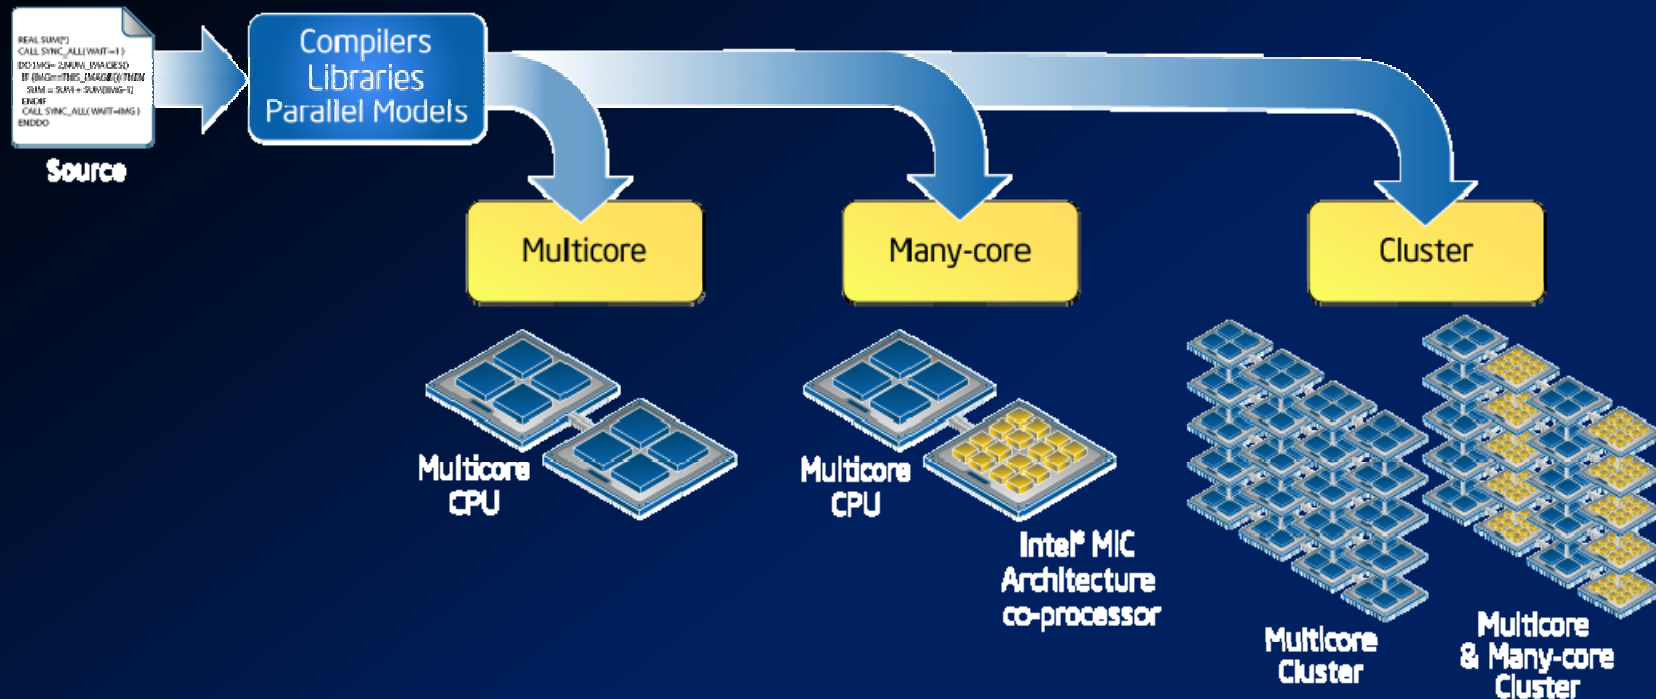
- GDDR5 memory
- Supports virtual memory

## Standard Intel Architecture Programming and Memory Model

For illustration only.  
Future options subject to change without notice.



# Scaling Programmability



*One Programming Model Democratizes Usage  
...Avoid Costly Detours*



# Knights Corner (22nm MIC product)

## *Intel® Many Integrated Core (Intel® MIC) Architecture*

### *Delivered Performance*


Launching on 22nm with >50 cores to provide outstanding performance for HPC users

### *Performance Density*

The compute density associated with specialty accelerators for parallel workloads

### *Programmability*

The many benefits of broad Intel CPU programming models, techniques, and familiar x86 developer tools



*A Step Forward In Dealing With  
Efficient Performance & Programmability*





# MIC On Track: ISC Demonstrations<sup>1</sup>



## Hybrid LU Factorization

*Leverages compute power of both Intel® Xeon® CPUs and Intel® MIC  
Delivers optimal performance by dynamically balancing large and small  
matrix  
Computations between Intel® Xeon® and Intel® MIC*

Up to 772  
GFLOP



## Hybrid Computing – SGEMM with Intel® MKL

*High performing SGEMM with just 18 lines of code – common between Intel® Xeon® CPUs  
and Knights Ferry  
Uses Intel® MKL in current version of Alpha stack/tools on Knights Ferry*

1+ TFLOP



## 7.4 TFLOP SGEMM in a node

*Simultaneous execution of SGEMM on 8 Knights Ferry cards to deliver 7.4 TFLOPS  
in 1 4U server*

7.4 TFLOP

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. Source: Intel measured results as of March 2011. See backup for details.  
For more information go to <http://www.intel.com/performance>

<sup>1</sup> Refer to backup material for system configurations

Optimized MIC Software Development Platform Performance





# MIC Partners at International Supercomputing 2011

*Showcasing  
Applications*



*Showcasing  
Platforms with  
Knights Ferry*



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





# Preface



- Programming models are the key to harness the computational power of massively parallel devices.
- Obviously, Intel has realized this trend and substantially supports open standards and invests in innovative programming models.
- LRZ and TUM are using Intel hard- and software for many years and know the tool chain by heart.
- We expect: A hardware product that delivers good performance (and energy-efficiency) without losing programmability.



# Advantages of the MIC Architecture



- Is a standard x86 architecture!
  - Allows many different parallel programming models like OpenMP, MPI and Intel Cilk!
  - Offers standard math-libraries like Intel MKL!
  - Supports whole Intel tool chain, e.g. Compiler & Debugger!
- ***Writing MIC-accelerated code with minimal effort and great performance***



# Workloads under Investigation




- Euroben Kernels (7 dwarfs of HPC)
- Data Mining
- TifaMMY – Matrix Operations (Demo at ISC'11!)
- Further Linear Algebra and Simulation Codes

- Selected micro-benchmarks used in PRACE for the evaluation of accelerator hardware & new languages:

<http://www.prace-project.eu/documents/public-deliverables/d6-6.pdf>

- Example: mod2am: dense matrix-matrix multiplication (MxM)



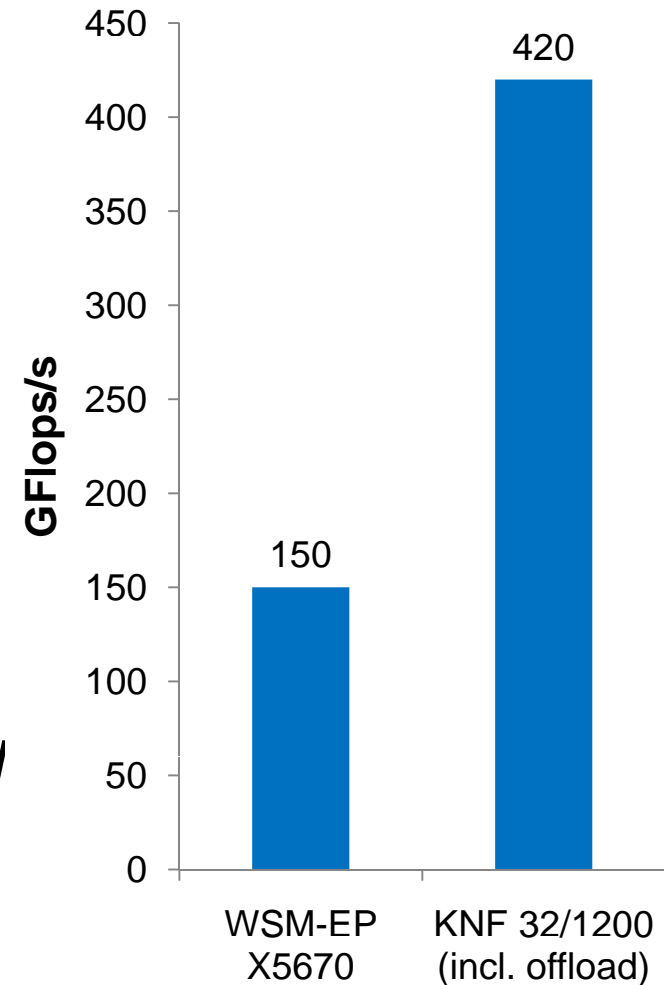
Performance evaluation of mod2am on KNF with 30 cores @1050 MHz using Intel's Offload Compiler, single precision, data transfer times excluded



# Data Mining with Adaptive Sparse Grids

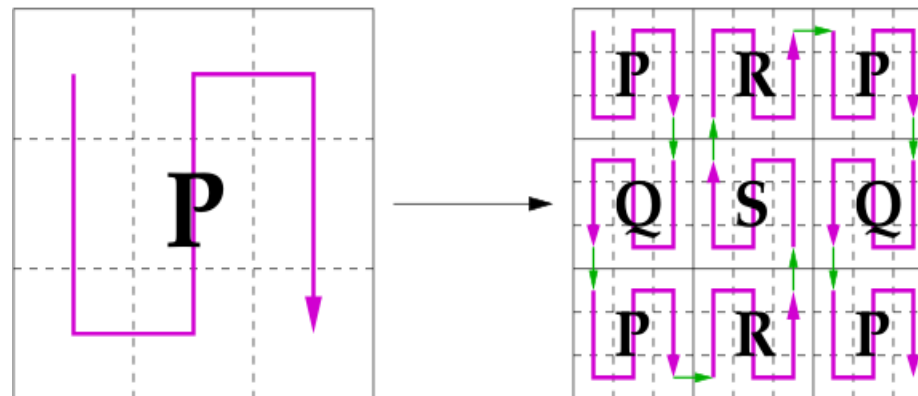


- Machine learning algorithm
  - Learning function from a training dataset
  - Important workload for classification and regression of huge datasets
- *MIC-Execution: Straightforward*
- *First version within a few hours*
  - *Optimized version took **2** days*






- TifaMMMy: self-adaptive and cache-oblivious framework for matrix operations optimized on fat x86 cores
  - This is done by nested recursions and vectorized kernels
    - On MIC, only the kernels were changed, MIC's x86 cores are able to tackle nested recursions!
- *Parallelization scheme employing OpenMP can be reused*
- *Having SSE kernels, bringing code to MIC is nearly for free*





# TifaMMMy – Performance Matrix Multiplication





# Advantages of the MIC Architecture



- Is a standard x86 architecture!
  - Allows many different parallel programming models like OpenMP, MPI and Intel Cilk!
  - Offers standard math-libraries like Intel MKL!
  - Supports whole Intel tool chain, e.g. Compiler & Debugger!
- ***Pre-release MIC-accelerated code for a typical scientific workload (e.g. Data Mining, TifaMMY) can reach up to 50% of peak performance!***

# How Intel® Delivers its Commitments:

*Intel Exascale Commitment:  
>100X Performance Of Today At  
Only 2X The Power Of Today's #1  
Scaling Today's Software Model*

Committed roadmap now and in the future

Flexible, open and scalable programming models

Collaborating with others to ensure the exascale future



