Современные тенденции разработки высокопроизводительных приложений

Игорь Одинцов
Intel

Архангельск, 9 февраля 2015 г.

Учиться параллельным вычислениям

Летняя интернатура в офисах Intel

Краткосрочные тренинги в офисах Intel (курсы Дельта)

Молодежные школы и тренинги в университетах

Континуум образовательных программ и ресурсов

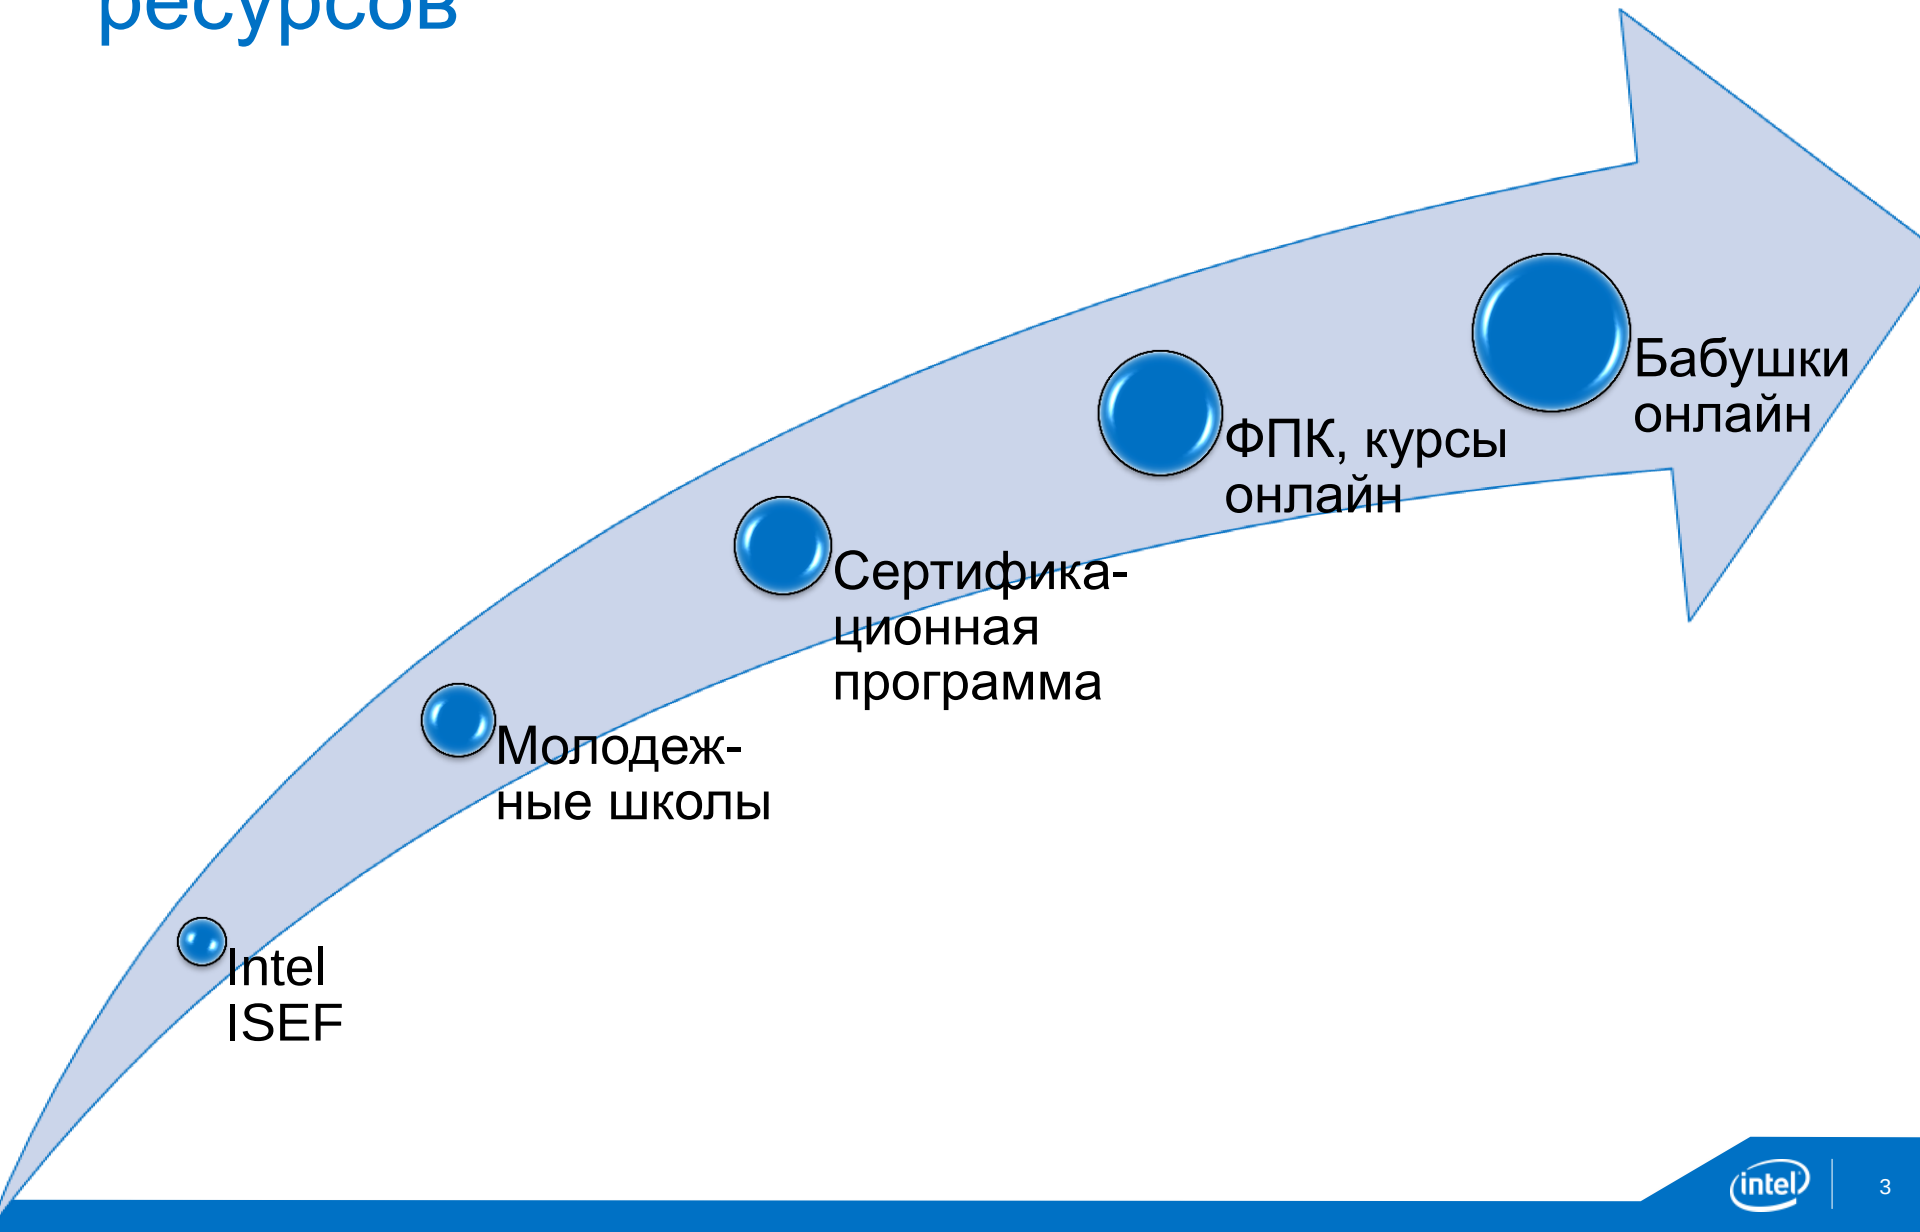

Насколько сложна разработка программ с параллельными вычислениями?

Тенденции

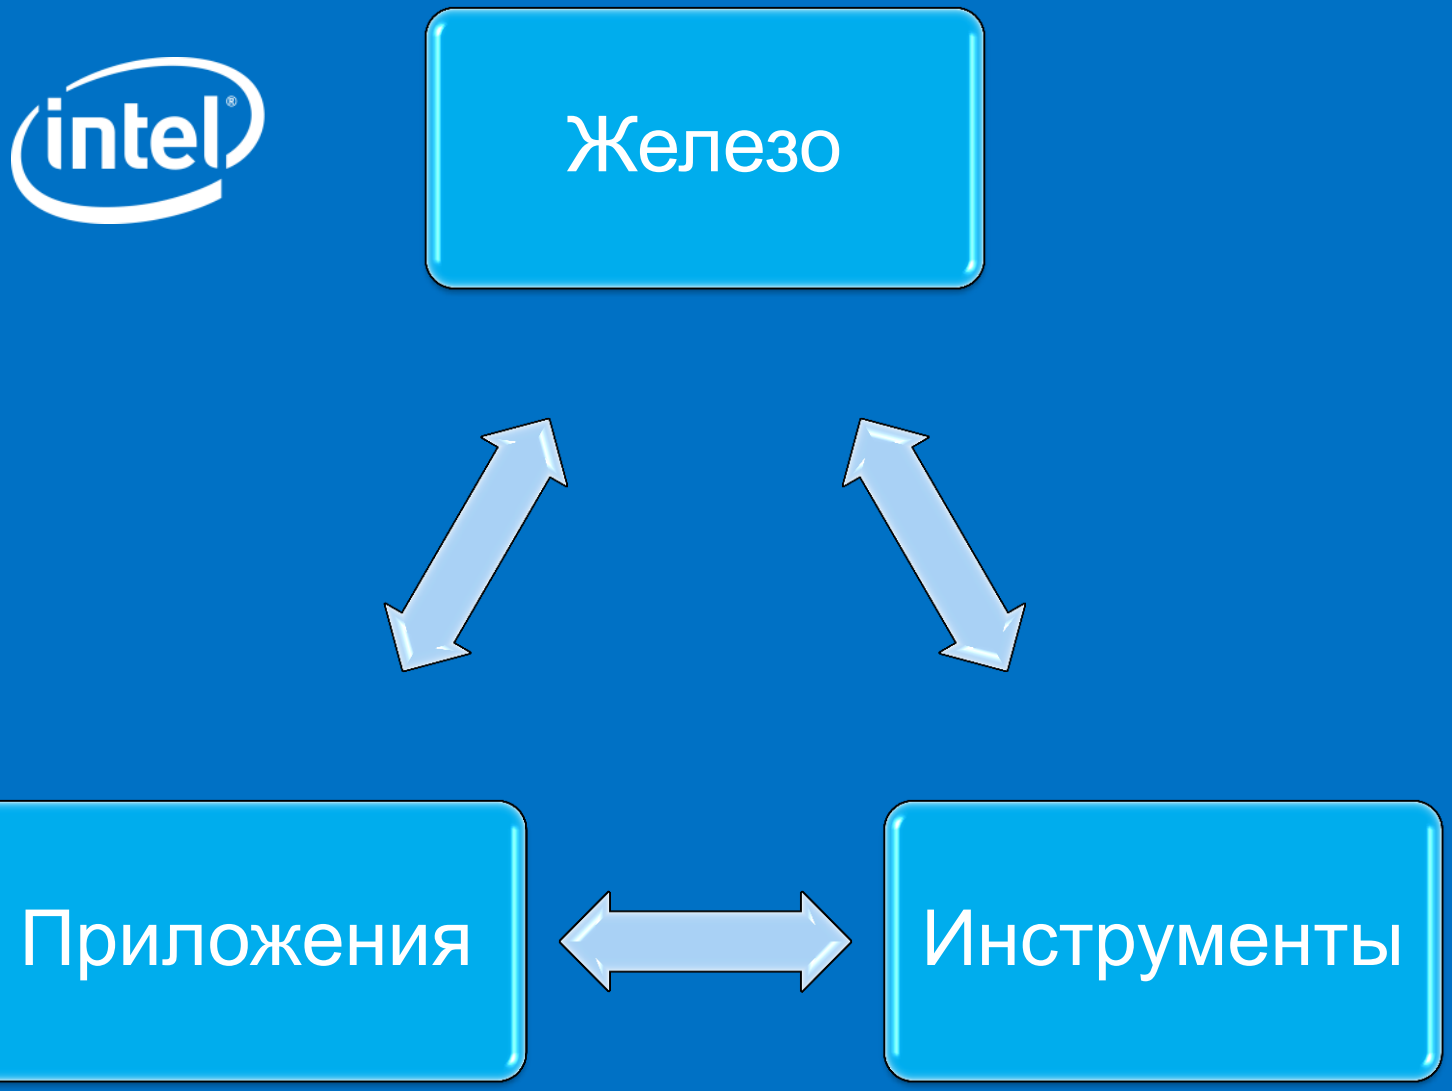

Unveiling Details of Knights Landing

(Next Generation Intel® Xeon Phi™ Products)

2nd half '15
1st commercial systems

3+ TFLOPS[1]
In One Package
Parallel Performance & Density

Platform Memory: DDR4 Bandwidth and Capacity Comparable to Intel® Xeon® Processors

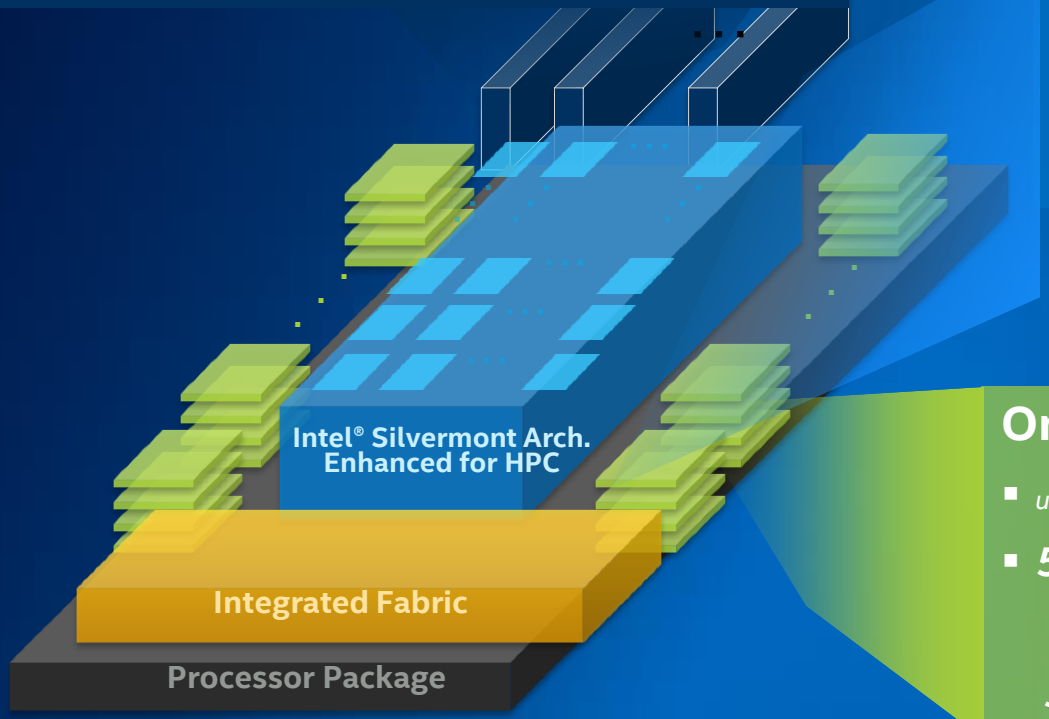

Compute: Energy-efficient IA cores[2]

- *Microarchitecture enhanced for HPC[3]*
- ***3X** Single Thread Performance vs Knights Corner[4]*
- *Intel Xeon Processor Binary Compatible[5]*

On-Package Memory:

- *up to **16GB** at launch*
- ***5X** Bandwidth vs DDR4[7]*
- ***1/3X** the Space[6]*
- ***5X** Power Efficiency[6]*

Jointly Developed with Micron Technology

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. [1]Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expectations of cores, clock frequency and floating point operations per cycle. FLOPS = cores x clock frequency x floating-point operations per second per cycle. . [2]Modified version of Intel® Silvermont microarchitecture currently found in Intel® Atom™ processors. [3]Modifications include AVX512 and 4 threads/core support. [4]Projected peak theoretical single-thread performance relative to 1st Generation Intel® Xeon Phi™ Coprocessor 7120P (formerly codenamed Knights Corner). [5] Binary Compatible with Intel Xeon processors using Haswell Instruction Set (except TSX) . [6]Projected results based on internal Intel analysis of Knights Landing memory vs Knights Corner (GDDR5). [7]Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory only with all channels populated.

Conceptual—Not Actual Package Layout

Intel® Xeon Phi™ Product Family

Industry and User Momentum

1 TFLOPS[1]

Knights Corner

Intel® Xeon Phi™ Coprocessor – Applications and Solutions Catalog

3+ TFLOPS[2]

-Bootable Processor
-On-Pkg, High BW Memory
-Integrated Fabric

Knights Landing

2H'15 *First Commercial Systems*

>50 systems providers expected[3]

+

many more *card-based systems*

>100 PFLOPS *customer system compute commits to-date[3]*

Announcing

Knights Hill

3rd Generation Intel® Xeon Phi™ Product Family

2nd Generation Intel Omni-Path Architecture

10nm process technology

[1] Claim based on calculated theoretical peak double precision performance capability for a single coprocessor. 16 DP FLOPS/clock/core * 61 cores * 1.23GHz = 1.208 TeraFLOPS
[2] Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expectations of cores, clock frequency and floating point operations per cycle. FLOPS = cores x clock frequency x floating-point operations per second per cycle. [3] Intel internal estimate

Intel® Omni Scale—The Next-Generation Fabric

- **Designed for Maximum Scalability**
- **Rich Set of Programming Models**
- **Flexible Configurations**
- **End-to-End Solution**

INTEGRATION

Intel® Omni Scale Fabric

Starting with Knights Landing

Intel® Omni Scale Fabric

Future 14nm generation

Coming in '15

- √ PCIe Adapters
- √ Edge Switches
- √ Director Systems
- √ Intel Silicon Photonics
- √ Open Software Tools*

Intel® True Scale Fabric Upgrade Program *Helps Your Transition*

*OpenFabrics Alliance

AMBER WRF VISIT VASP UTBENCH SU2 SG++ SeisSol ROTOR SIM R Quantum Espresso Optimized integral OPENMP/ MPI Openflow NWChem

Modernizing HPC Community Codes

AVBP (Large Eddy) NEMO5

Breakdown of Application Hours
NERSC - Hopper 2012[1]

MPAS

Blast Mardyn

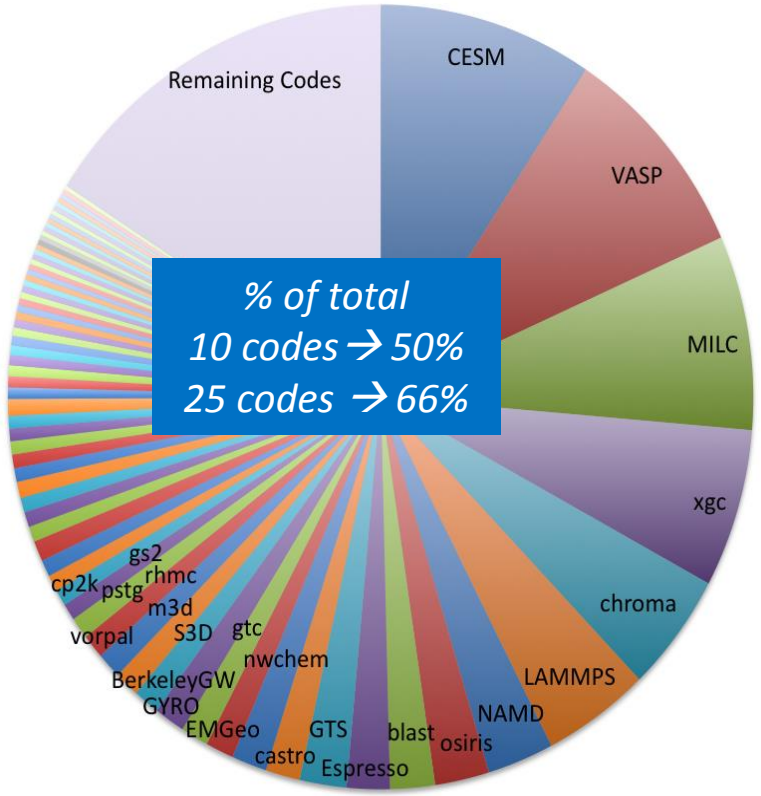

BUDE MACPO

CAM-5 Ls1

CASTEP Harmonie

CESM GTC

CFSv2 GS2

CIRCAC Gromacs

GPAW

Intel® Parallel Computing Centers

Collaborating to accelerate the pace of discovery

>40 Centers
13 Countries
>70 Codes
2 User Groups

https://software.intel.com/en-us/ipcc

CliPhi (COSMOS) COSA Cosmos codes DL-MESO DL-Poly ECHAM6 Elmer FrontFlow/Blue Code GADGET GAMESS-US

[1]*Source: NERSC*

[2] www.ihpcc2014.com

Create the applications that shape and enable innovation

Software

Intel® Parallel Studio XE

Technical Computing Enterprise, and HPC Software

Intel® System Studio

Embedded Systems/Device Development Software

Intel® Media Server Studio

Media Development Software

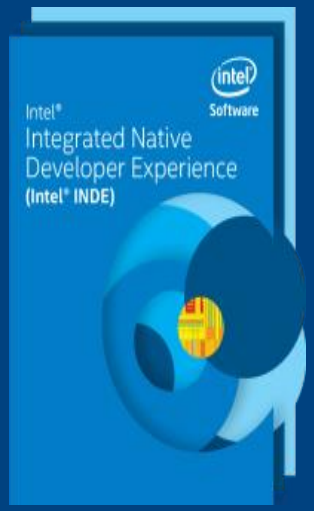

Intel® Integrated Native Developer Experience

Native App Development Software

Intel® XDK

Web and hybrid HTML5 App Development Software

Высокопроизводительные вычисления и компьютерный континуум

суперкомпьютеры

центры обработки данных

облачные среды

ЭНЕРГОЭФФЕКТИВНАЯ ПРОИЗВОДИТЕЛЬНОСТЬ

БЕЗОПАСНОСТЬ И ЗАЩИТА ДАННЫХ

ИНТЕРНЕТ-ДОСТУП И СЕТЕВАЯ СРЕДА

десктопы

ноутбуки

нетбуки

планшеты

смартфоны

«умные» телевизоры

встроенные устройства

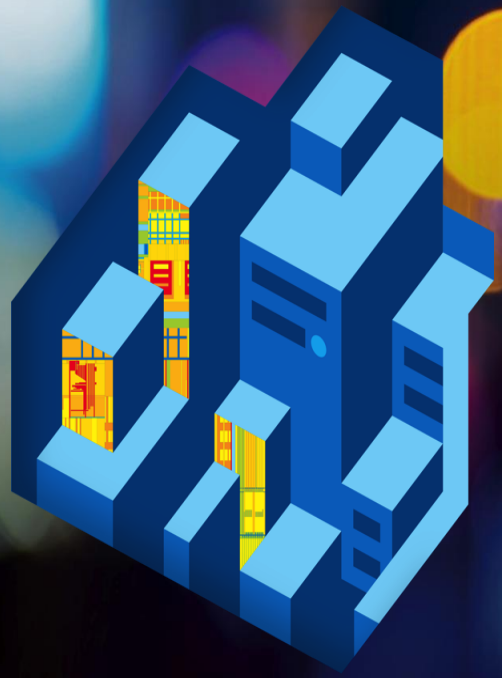

Intel® Parallel Studio XE

ACCELERATE

improve application performance, scalability and reliability.

TRANSFORM YOUR CODE >

Technical Computing Enterprise, and HPC Software

Intel® Parallel Studio XE

Improve application performance, scalability, reliability

Industry leading performance

- Advanced compilers and libraries
- Linux*, Windows*, and OS X

Parallel programming models

- Task, Data Parallel Performance
- Distributed Performance

Insightful analysis tools

- Parallelism Assistant
- Thread and Memory Analysis
- Performance Analysis

Create → Build → Verify → Tune

Faster Code Faster

Intel® Parallel Studio XE 2015

- Simplifies building, debugging and tuning parallel code
  - Integrated C++ and Fortran tool suite
  - Drops into development environment
    - e.g., Visual Studio*
  - Windows*, Linux* & OS X*

Code Faster

Compilers with high level parallelism features including OpenMP* 4.0

Parallelism prototyping assistant

Advanced parallel models and libraries, simple update with relink

Graphical profilers visualize bottlenecks

Memory, thread and MPI error checkers help remove errors

Faster Code

Performance without compromise through optimizations for current and future processors
- Compilers
- Libraries

Profilers simplify tuning parallel code for best performance

How Intel® Parallel Studio XE 2015 helps make *Faster Code Faster* for HPC

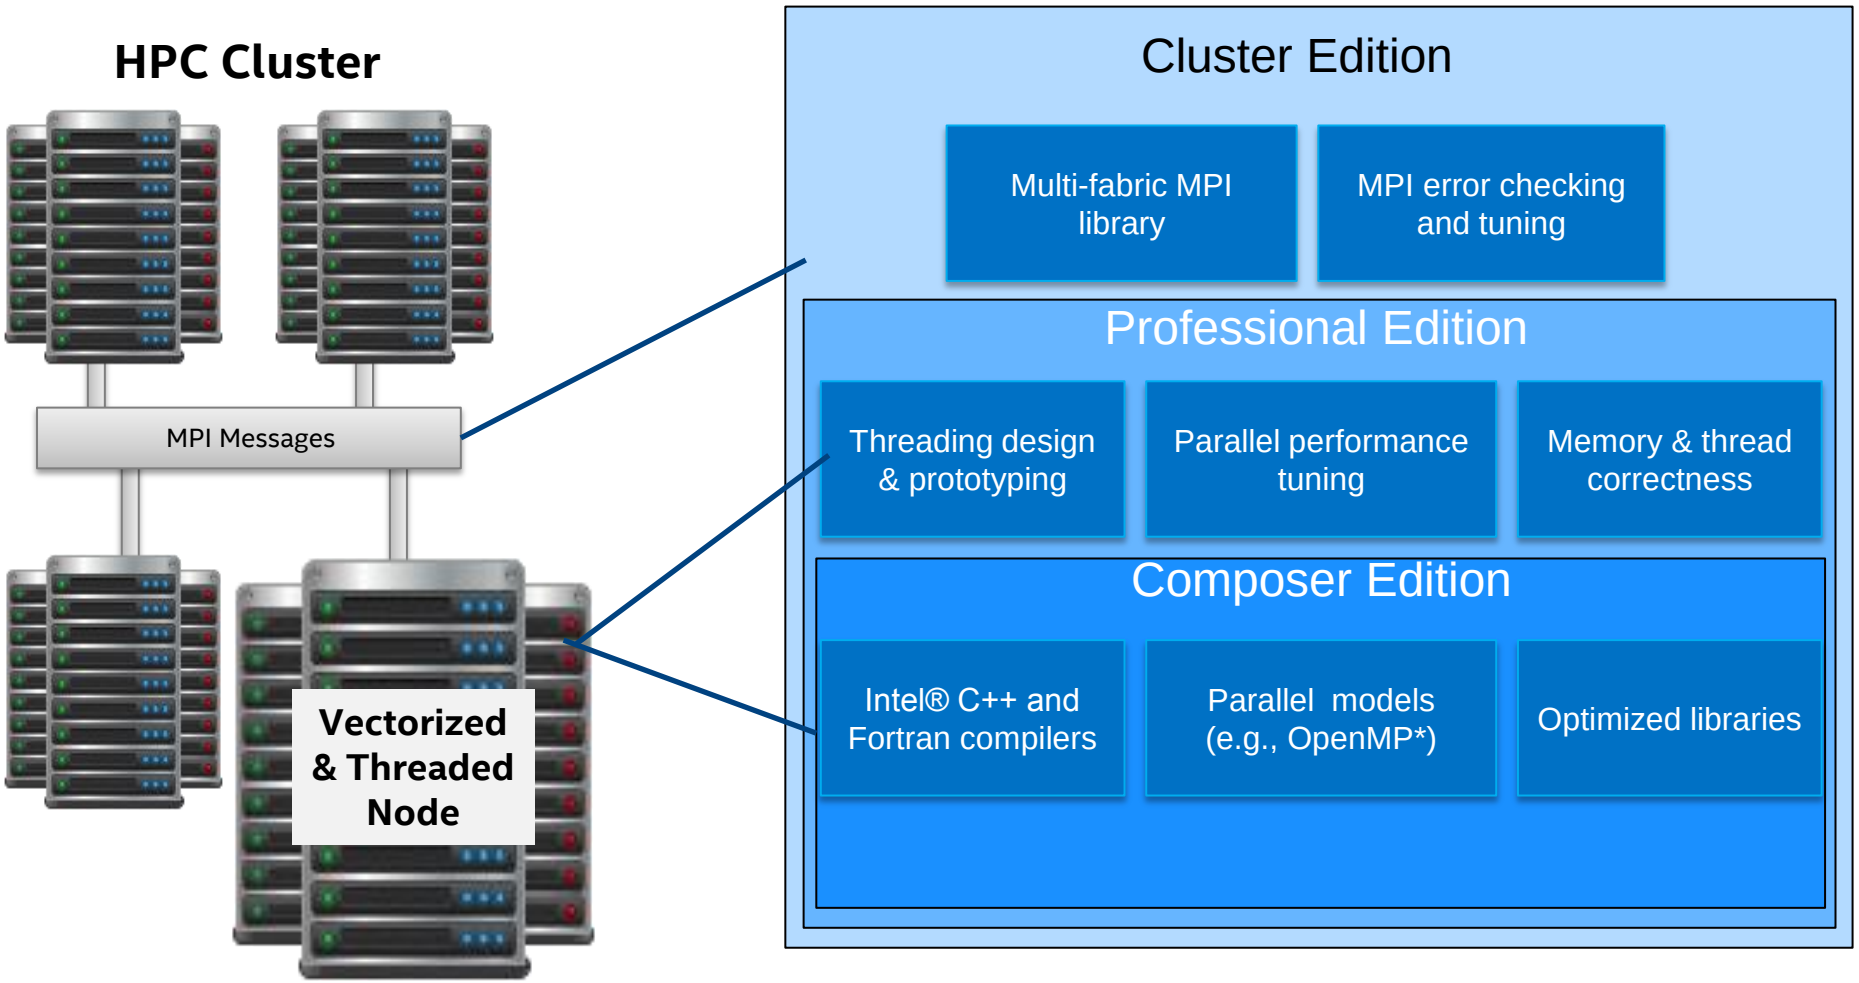

Intel® Parallel Studio XE 2015

Phase	Product	Feature	Benefit
Build	**Intel® Composer XE**	Compilers, Performance and Threading Libraries	Out of the box performance
	Intel® MPI Library†	High Performance Message Passing (MPI) Library	Interconnect independence
	Intel® Advisor XE	Threading Prototyping Tool (Studio XE products only)	Simplifies parallel application design
Verify & Tune	**Intel® VTune™ Amplifier XE**	Performance Profiler	Find performance bottlenecks
	Intel® Inspector XE	Memory & Threading Dynamic Analysis	Code quality, improved security
	Intel® Trace Analyzer & Collector†	MPI Performance Profiler	Find performance bottlenecks in cluster-based applications

Efficiently Produce Fast, Scalable and Reliable Applications with Intel Tools

Intel® C++ and Fortran Compilers

Boost application performance on Windows* and Linux*

Boost C++ application performance on Windows* & Linux* using Intel® C++ Compiler (higher is better)

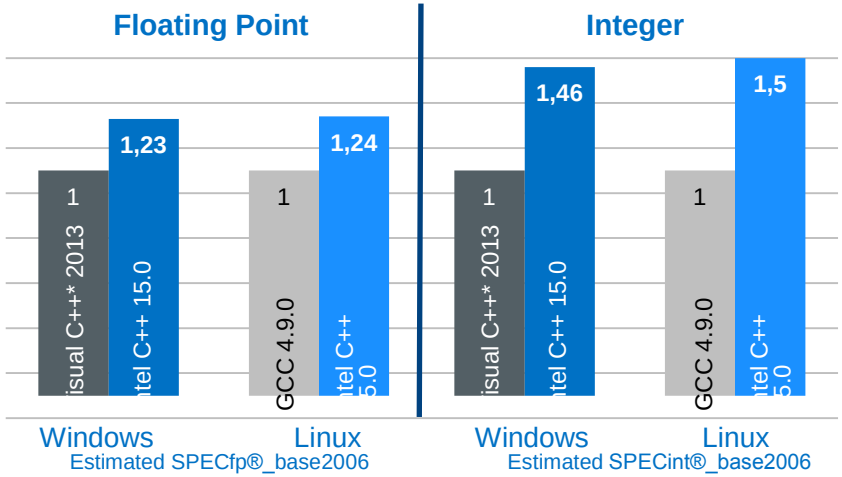

Relative geomean performance, SPEC* benchmark - higher is better

Boost Fortran application performance on Windows* & Linux* using Intel® Fortran Compiler (lower is better)

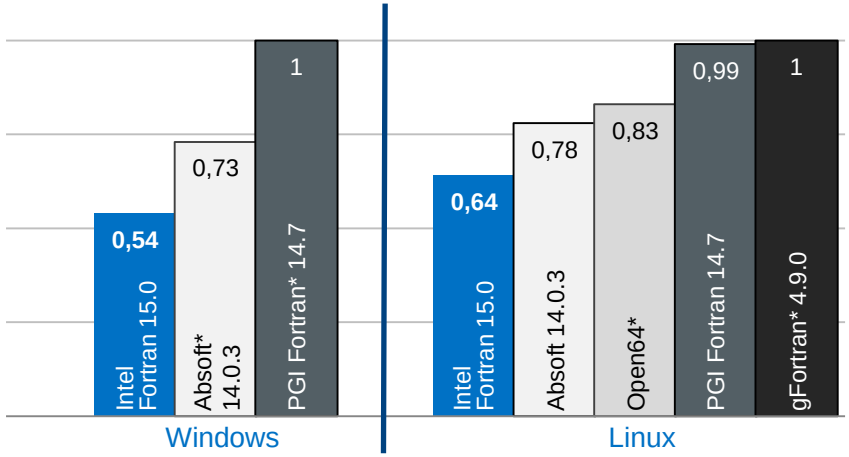

Relative geomean performance, Polyhedron* benchmark– lower is better

Configuration: Hardware: HP ProLiant DL360p Gen8 with Intel® Xeon® CPU E5-2680 v2 @ 2.80GHz, 256 GB RAM, HyperThreading is on. Software: Intel C++ compiler 15.0, Microsoft Visual C++ 2013, GCC 4.9.0. Linux OS: Red Hat Enterprise Linux Server release 6.5 (Santiago), kernel 2.6.32-431.el6.x86_64. Windows OS: Windows 7 Enterprise, Service pack 1. SPEC* Benchmark (www.spec.org).

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. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804 .

Configuration: Hardware: Intel® Core™ i7-4770K CPU @ 3.50GHz, HyperThreading is off, 16 GB RAM. . Software: Intel Fortran compiler 15.0, Absoft*14.0.3,. PGI Fortran* 14.7, Open64*, gFortran* 4.9.0,. Linux OS: Red Hat Enterprise Linux Server release 6.4 (Santiago), kernel 2.6.32-358.el6.x86_64. Windows OS: Windows 7 Enterprise, Service pack 1. Polyhedron Fortran Benchmark (www.polyhedron.com). Windows compiler switches: Absoft: -m64 -O5 -speed_math=10 -fast_math -march=core -xINTEGER -stack:0x80000000. Intel® Fortran compiler: /fast /Qparallel /link /stack:64000000. PGI Fortran: -fastsse -Munroll=n:4 -Mipa=fast,inline -Mconcur=numa. Linux compiler switches: Absoft -m64 -mavx -O5 -speed_math=10 -march=core –xINTEGER. Gfortran: -Ofast -mfpmath=sse -flto -march=native -funroll-loops -ftree-parallelize-loops=4. Intel Fortran compiler: -fast –parallel. PGI Fortran: -fast -Mipa=fast,inline -Msmartalloc -Mfprelaxed -Mstack_arrays -Mconcur=bind. Open64: -march=bdver1 -mavx -mno-fma4 -Ofast -mso –apo.

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. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804 .

Scalability and Productivity - TBB

Intel® Threading Building Blocks

Scalability on Intel® Xeon Phi™ Coprocessor

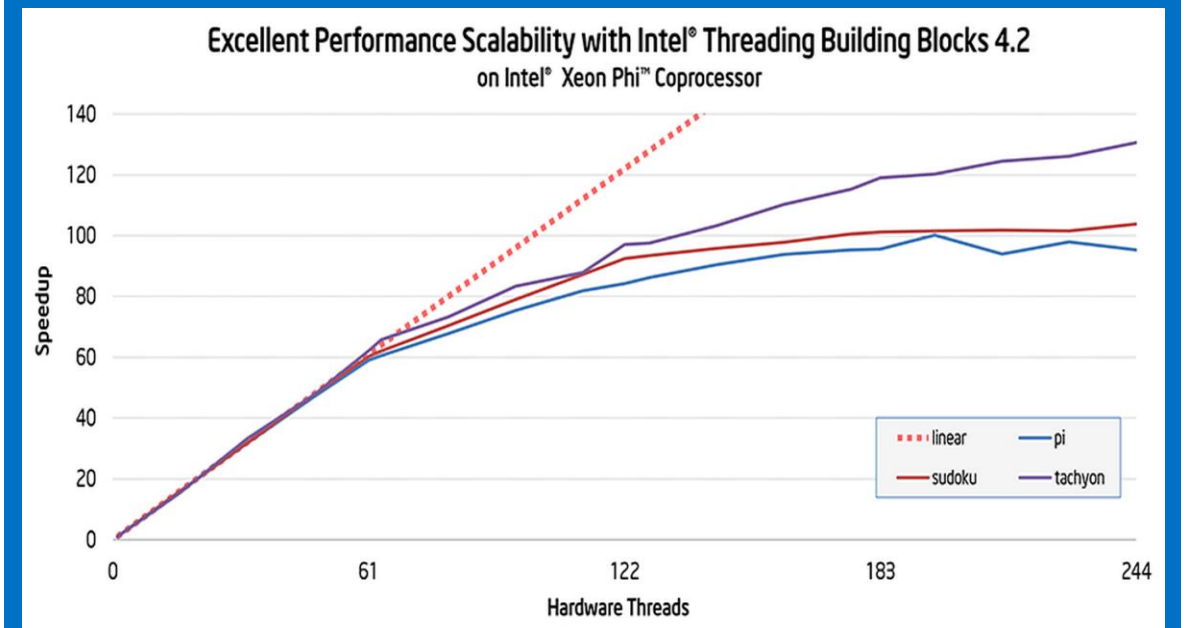

Configuration Info - SW Versions: Intel® C++ Intel® 64 Compiler, Version 14.0, Intel® Threading Building Blocks (Intel® TBB) 4.2; Hardware: Intel® Xeon Phi™ Coprocessor 7120X (16GB, 1.238 GHz, 61C/244T); MPSS Version: 2.1.6720-13; Flash Version: 2.1.02.0386; Host: 2x Intel(R) Xeon(R) CPU E5-2680 0 @ 2.70GHz (16C/32T); 64GB Main Memory;OS: Red Hat Enterprise Linux Server release 6.2 (Santiago), kernel 2.6.32-220.el6.x86_64; Benchmarks were run on Intel® Xeon Phi™ Coprocessor. Benchmark Source: Intel Corp. Note: sudoku and tachyon are included with Intel TBB

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, refer to www.intel.com/performance/resources/benchmark_limitations.htm.

* Other brands and names are the property of their respective owners.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User

Developer Productivity

"Intel® TBB provided us with optimized code that we did not have to develop or maintain for critical system services.
I could assign my developers to code what we bring to the software table."

Michaël Rouillé, CTO, Golaem

Intel® Integrated Performance Primitives Overview

A software developer's competitive edge

Optimized performance focused on the compute-intensive tasks that matter to you

Easy to use building blocks to create high performance workflows in String Processing, Data Compression, Image Processing, Cryptography, Signal Processing, & Computer Vision

Unleash your potential through access to silicon

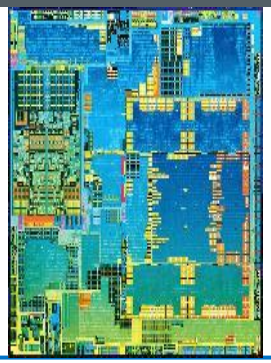

Consistent C APIs span multiple generations of Intel's processors and SoC solutions, removing the need to develop for specific architecture optimizations

Included in Intel® Parallel Studio XE Suites

The optimizations you need, available where you need them

Intel® Math Kernel Library is a Computational Math Library

Mathematical problems arise in many scientific disciplines


Energy


Signal Processing

Financial Analytics


Engineering Design


Digital Content Creation


Science & Research

These scientific applications areas typically involve mathematics …

- Differential equations
- Linear algebra
- Fourier transforms
- Statistics

$$-\frac{\partial u^2}{\partial x^2}-\frac{\partial u^2}{\partial y^2}+q\,u=f(x,y)$$

Intel® MKL can help solve your computational challenges

Tune Applications for Scalable Multicore Performance

Intel® VTune™ Amplifier XE Performance Profiler

Is your application slow?

Does its speed scale with more cores?

Tuning without data is just guessing

- Accurate CPU, GPU[1] & threading data
- Powerful analysis & filtering of results
- Easy set-up, no special compiles

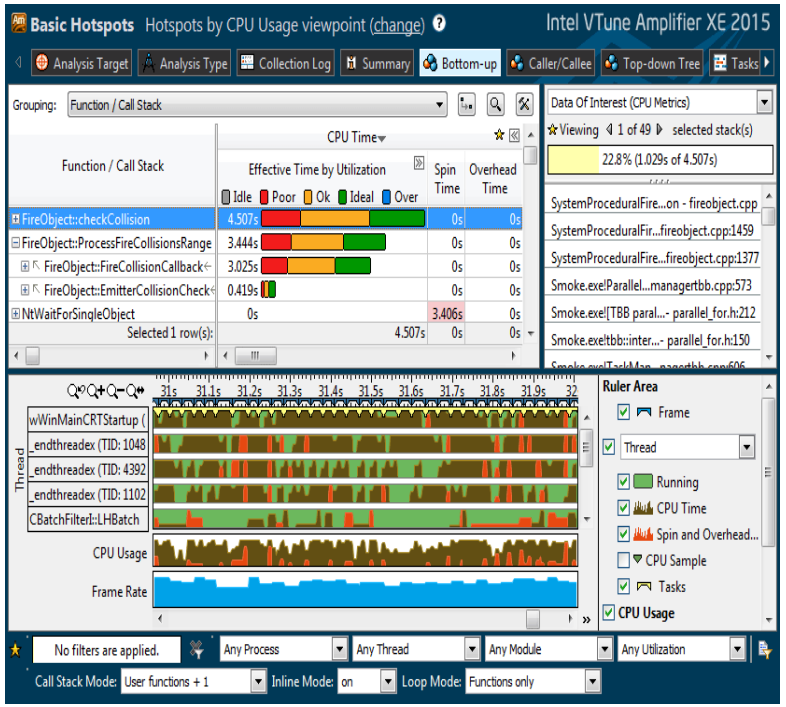

For Windows* and Linux* From $899
(GUI only now available on OS X*)

"Last week, Intel® VTune™ Amplifier XE helped us find almost 3X performance improvement. This week it helped us improve the performance another 3X."

Claire Cates
Principal Developer
SAS Institute Inc.

[1] Windows* only.

http://intel.ly/vtune-amplifier-xe

Design Then Implement

Intel® Advisor XE Thread Prototyping

Design Parallelism

- No disruption to regular development
- All test cases continue to work
- Tune and debug the design before you implement it

1) Analyze it.

2) Design it.
(Compiler ignores these annotations.)

3) Tune it.

4) Check it.

Implement Parallelism

5) Do it!

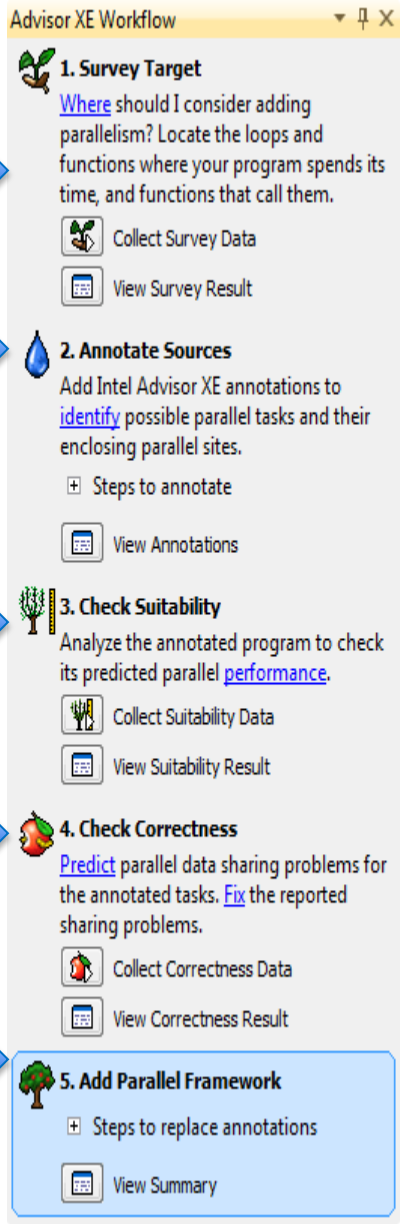

Less Effort, Less Risk, More Impact

История успеха

Межведомственный Суперкомпьютерный Центр РАН

intel Software

МСЦ РАН JSCC

РСК

Intel Cluster Ready

Оптимизацию скорости работы приложений в распределенных и гибридных средах на мощном и энергоэффективном суперкомпьютере МВС-10П с сопроцессорами Intel® Xeon Phi™, созданном в МСЦ РАН группой компаний РСК на базе архитектуры «РСК Торнадо», обеспечивает программный пакет Intel® Cluster Studio XE 2013. Например, это приложения с методами Monte-Carlo, приложения квантовой хромодинамики, а также при решении задач расчета прогноза погоды WRF и др.

История успеха

Лаборатория I-SCALARE в МФТИ

МФТИ

РСК

Intel Cluster Ready

Оптимизацию скорости работы приложений на энергоэффективном суперкомпьютере на базе архитектуры «РСК Торнадо» в Лаборатории суперкомпьютерных технологий для биомедицины, фармакологии и малоразмерных структур в МФТИ обеспечивает программный пакет Intel® Cluster Studio XE 2013. Тем самым, российские ученые получили возможность проводить более масштабные исследования, что позволяет достичь очередных успехов в моделировании поведения вирусов и в создании в будущем новых лекарств для борьбы со многими опасными заболеваниями.

8 https://software.intel.com/en-us/articles/icp-case-studies

История успеха

Федеральная служба по гидрометеорологии и мониторингу окружающей среды Российской Федерации (Росгидромет)

intel Software

РСК

Intel Cluster Ready

Новый энергоэффективный и компактный вычислительный кластер Росгидромета с пиковой производительностью 35 ТФлОПС благодаря инновационной архитектуре «РСК Торнадо» с жидкостным охлаждением, созданный на базе процессоров Intel® Xeon® E5-2690 и серверных плат Intel® S2600JF, стал рабочим инструментом российских метеорологов для дальнейшего развития оперативных технологий с целью повышения точности, заблаговременности и детализации прогнозов погоды. Например, с его помощью решались задачи оперативного моделирования и прогнозирования погоды в районе Сочи во время проведения Зимних Олимпийских и Паралимпийских Игр в марте 2014 г.

Конференция Intel для разработчиков

Москва 19 сентября

Иркутск 17 сентября

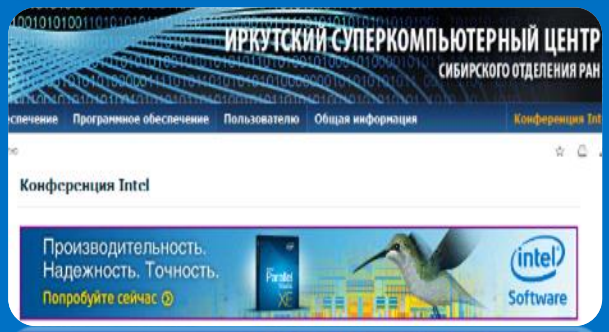

Vladivostok RSAP 2013

REMOTE SENSING IN ASIA PACIFIC

Международная конференция

Дистанционное зондирование окружающей среды: научные и прикладные исследования в Азиатско-Тихоокеанском регионе (RSAP2013)

Владивосток, Россия, 24-27 сентября 2013 г.

ИФА РАН ДВФУ ДВО РАН

Официальный спонсор

intel Software

История успеха: ЮУрГУ

Механизм «одевания» виртуального платья на манекен

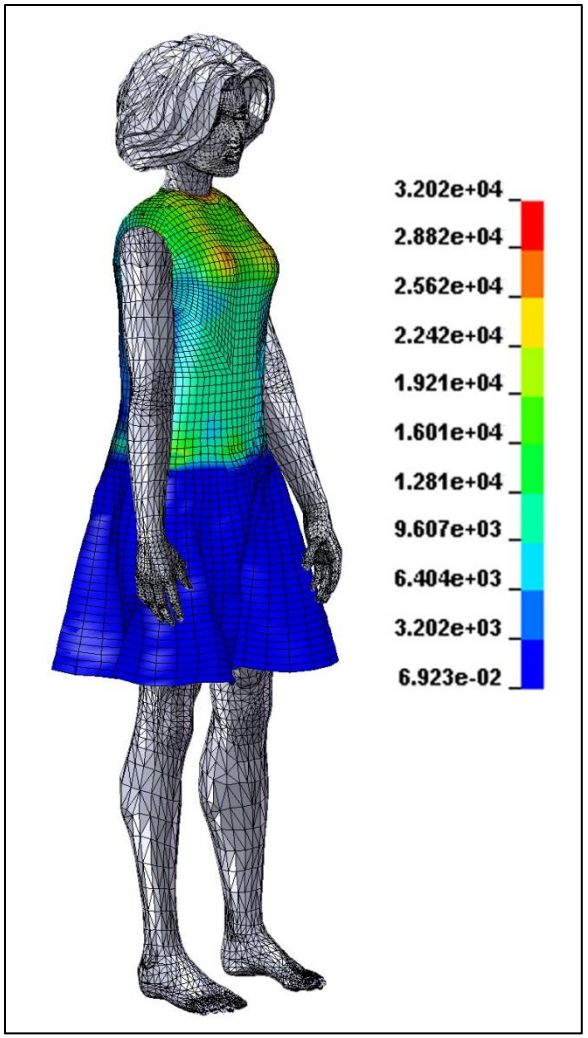

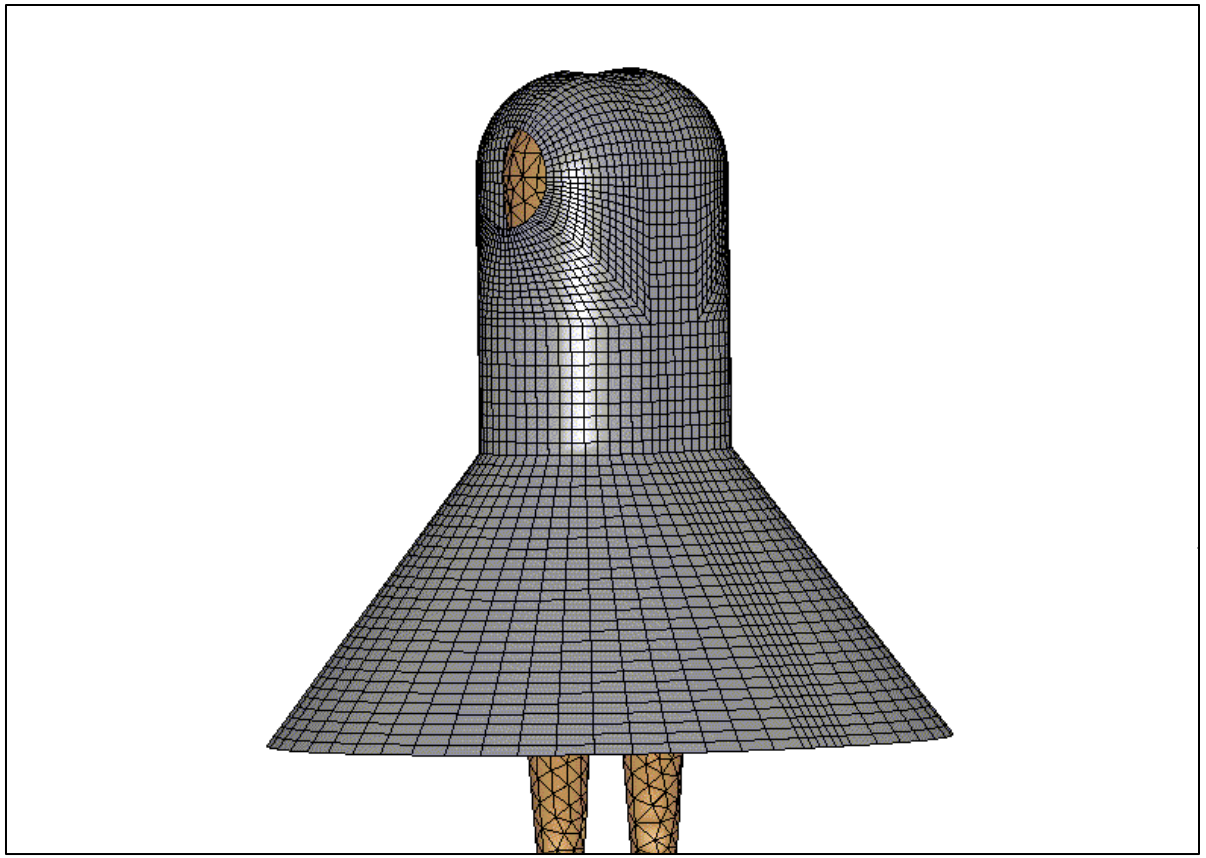

Процесс динамического взаимодействия баллистической ткани с пулей

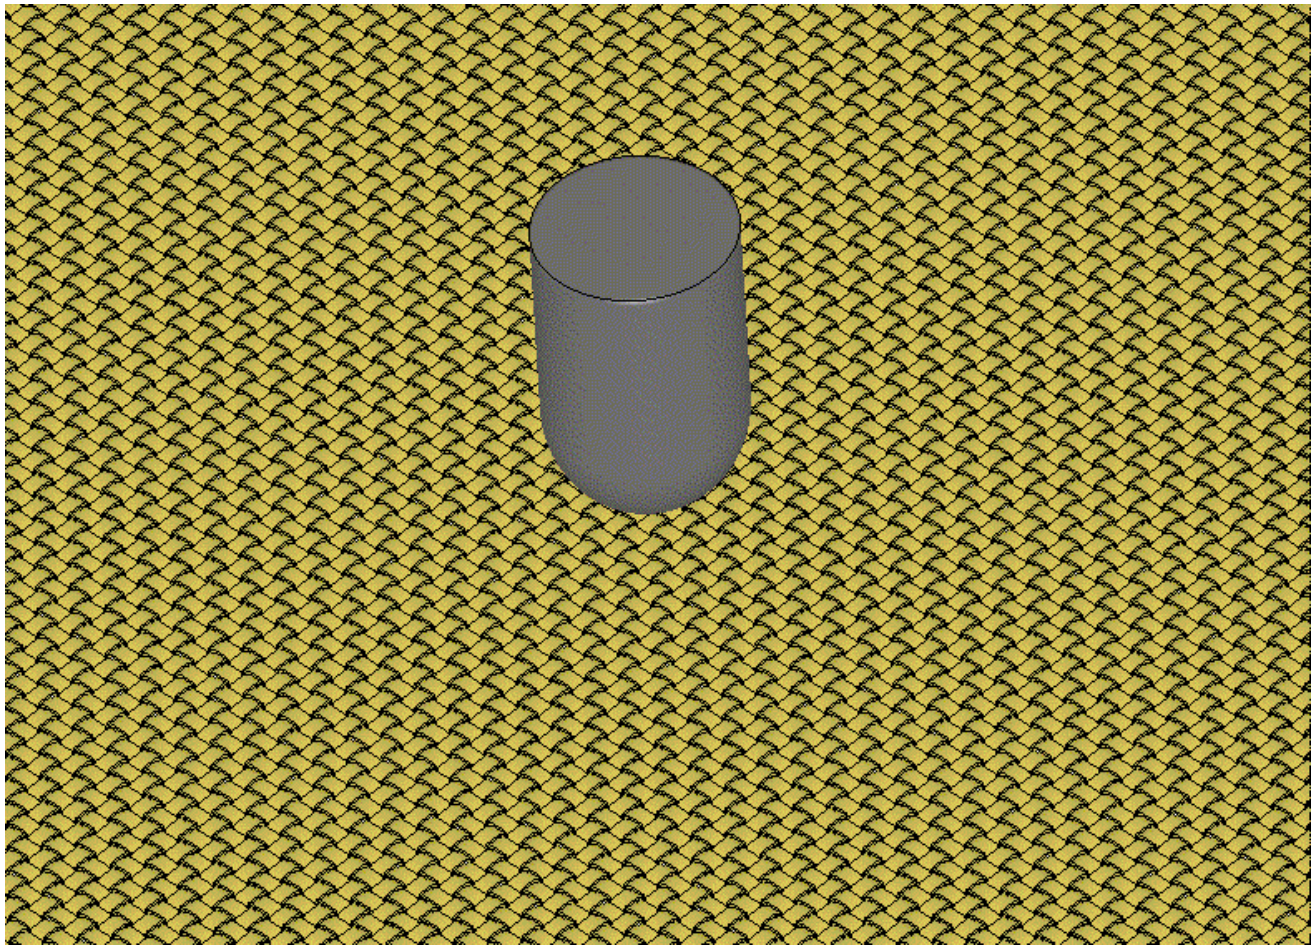

How can Intel Tools help Oil & Gas applications?

Task	Typical Problems
Seismic Exploration	Very Large Jobs, Fault tolerance, Complex FFTW's
Reservoir Modeling	Memory Bound Apps, Complex Fluid Mechanics, Large Memory Footprint
Field Operations	Optimizing oil distribution in the field, Signal Processing

"Both Shell and Intel have made significant investments in engineer-to-engineer collaboration and training to optimize Shell's algorithms and applications for the Intel processors and to enable our codes to fully utilize rising core counts. We believe **our optimization efforts are a key to our HPC success**.

We utilize the Intel® C++ and Fortran compilers and Intel® Math Kernel Library, as well as tools such as Intel® VTune™ Performance Analyzer."

Oil & Gas: Additional Reading
Article: Optimize Seismic Image Processing on Intel MIC
Article: Developing seismic imaging code for Intel Xeon Phi
White Paper: 3D Finite Differences on Multi-core Processors
White Paper: Shell Drill Downs on Hundredfold Improvements in HPC
Tips and Tricks for finite difference

Intel Tools are used by majority of leading oil and gas companies

Oil & Gas: Reservoir Modeling

Problem	Memory bound apps, complex fluid mechanics, large memory footprint

1: Identify bottlenecks w/ Intel® VTune™ Amplifier XE, ensure your software is utilizing all CPU cores.

2: Increase app performance by reusing optimized Linear Algebra functions in Intel Math Kernel Library

Linear Algebra
- BLAS
- LAPACK
- Sparse Solvers

CLUSTER
- ScaLAPACK

Where is my application...

Spending Time?

Function - Call Stack	CPU Time
algorithm_2	3.560s
do_xform	3.560s
algorithm_1	1.412s
BaseThreadInitTh	0.000s

Wasting Time?

Line		MEM_LOAD... LLC_MISS
475	float rx, ry, rz =	
476	float param1 = (A	30,000
477	float param2 = (A	
478	bool neg = (rz < 0	

Waiting Too Long?

Wait Time	Wait Count
Idle / Poor / Ok / Ideal	
176.504s	18,277
84.681s	5,499
84.612s	5,489

3: Use Intel® Inspector XE to identify memory and threading errors

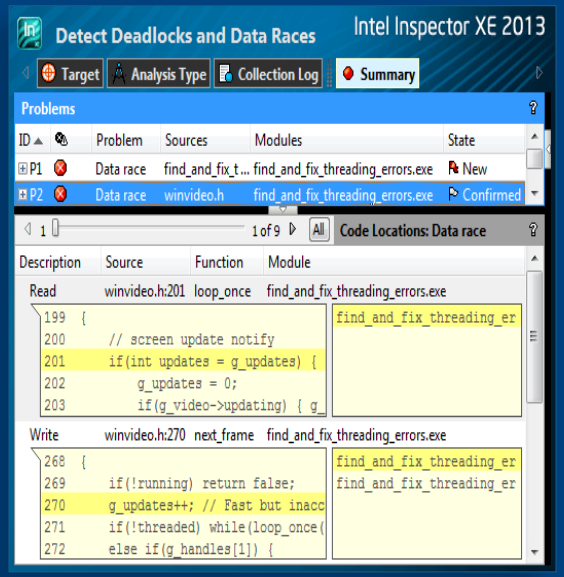

Oil & Gas: Seismic Exploration

Problem	Very large jobs, apps require fault tolerance, complex FFTW's

1: Intel MPI Library scales to 120k processes and is interconnect independent. Intel Trace Analyzer and Collector scales to 6k processes, can trace communication type and identify user code

2: Automatic job restart via support for Berkeley Labs Checkpoint Restart

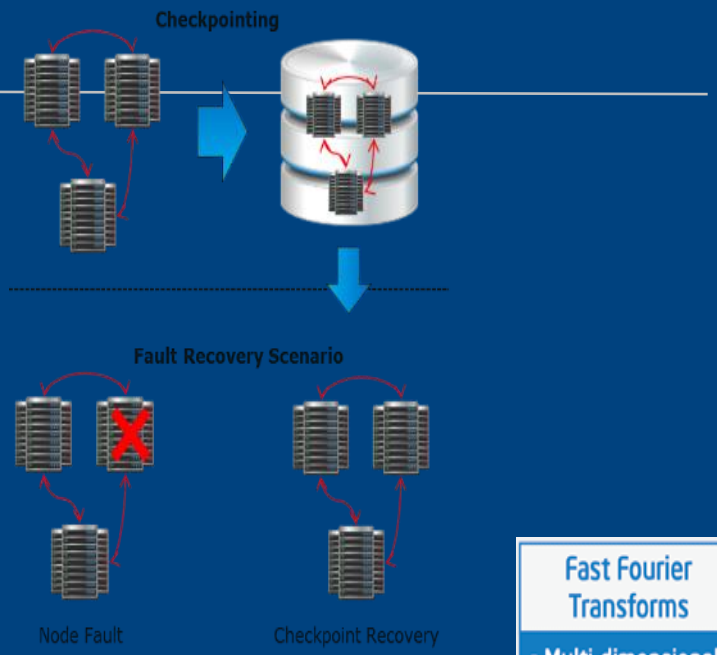

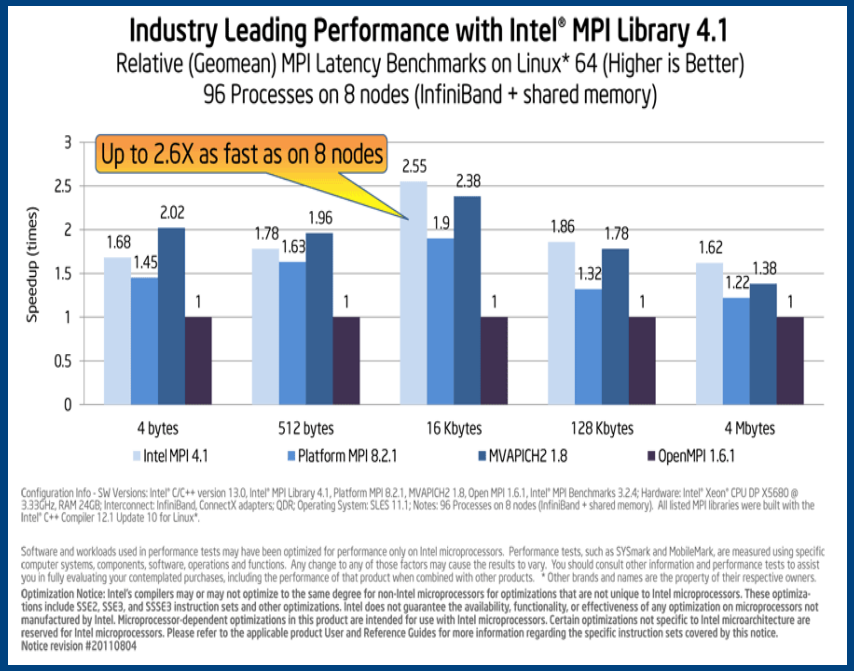

3: Improve performance using pre-optimized FFT functions in Intel Math Kernel Library

Fast Fourier Transforms

- Multi-dimensional (up to 7D)
- FFTW* Interfaces

CLUSTER

- Cluster FFT

EDC North America Development Survey 2011, Volume II

libraries users rely on Intel's Math Kernel Library

"Intel MKL is indispensable for any high-performance computer user on x86 platforms." —PROF. JACK DONGARA, INNOVATIVE COMPUTING LAB, UNIVERSITY OF TENNESSEE, KNOXVILLE

"By adopting the Intel® MKL DGEMM libraries, our standard benchmarks timing improved between 43 percent and 71 percent…"
MATT DUNBAR, SOFTWARE DEVELOPER, ABAQUS, Inc.

A Very Good Kitty, Indeed

DreamWorks Animation's *Puss in Boots* Uses Intel® Math Kernel Library to Help Create Dazzling Special Effects

BY GARRET ROMAINE

"We want solid building blocks that we know will be robust and have optimal performance. Intel MKL provides that. We can start understanding the artistic benefits of a complex algorithm more quickly if we don't have to build every component of a system from scratch." —RON HENDERSON, SR.MANAGER, R&D, DREAMWORKS ANIMATION

Case study: Engineering analysis

- Intel® Parallel Studio helped identify conventionally difficult errors in the parallel version of the SIMULIA software. Simulation time reduced by more than 60 times when running on 128 cores.

Multicore Intel® Architecture	Head gasket simulation: Sparse linear equations with 5.3M variables, 2.6E13 Flops
1 core	3 days
16 cores	06:40:40
32 cores	03:20:20
64 cores	01:40:40
128 cores	01:06:06

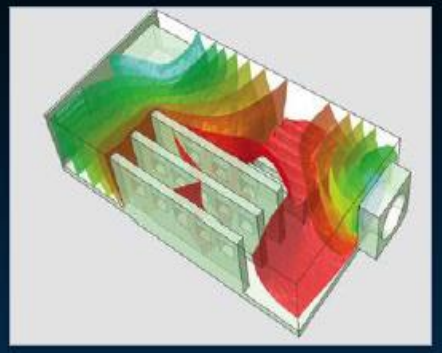

Тенденции

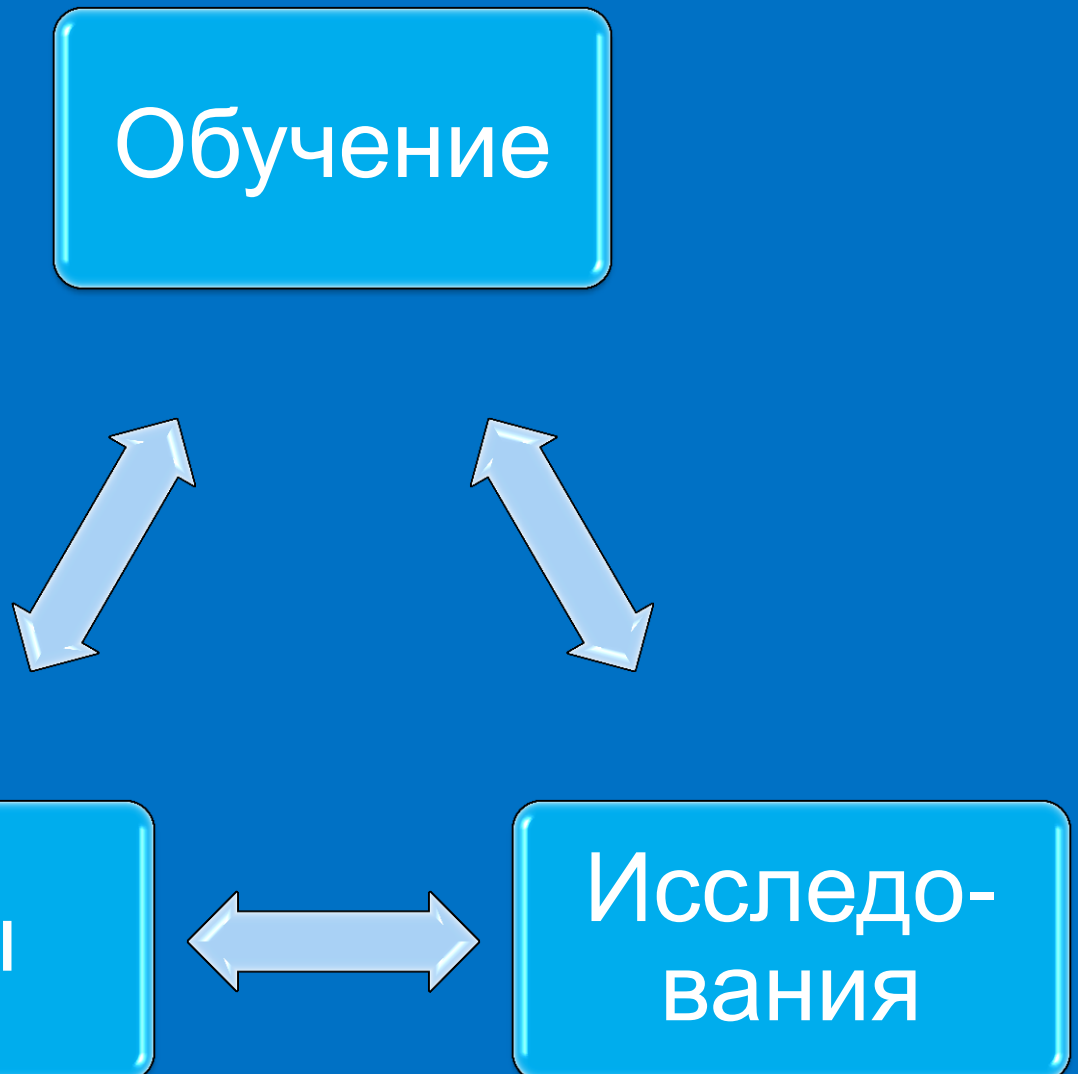

Тенденции: IT-обучение

IT-индустрия все больше вовлекается
в систему образования и участвует
в формировании образовательных программ

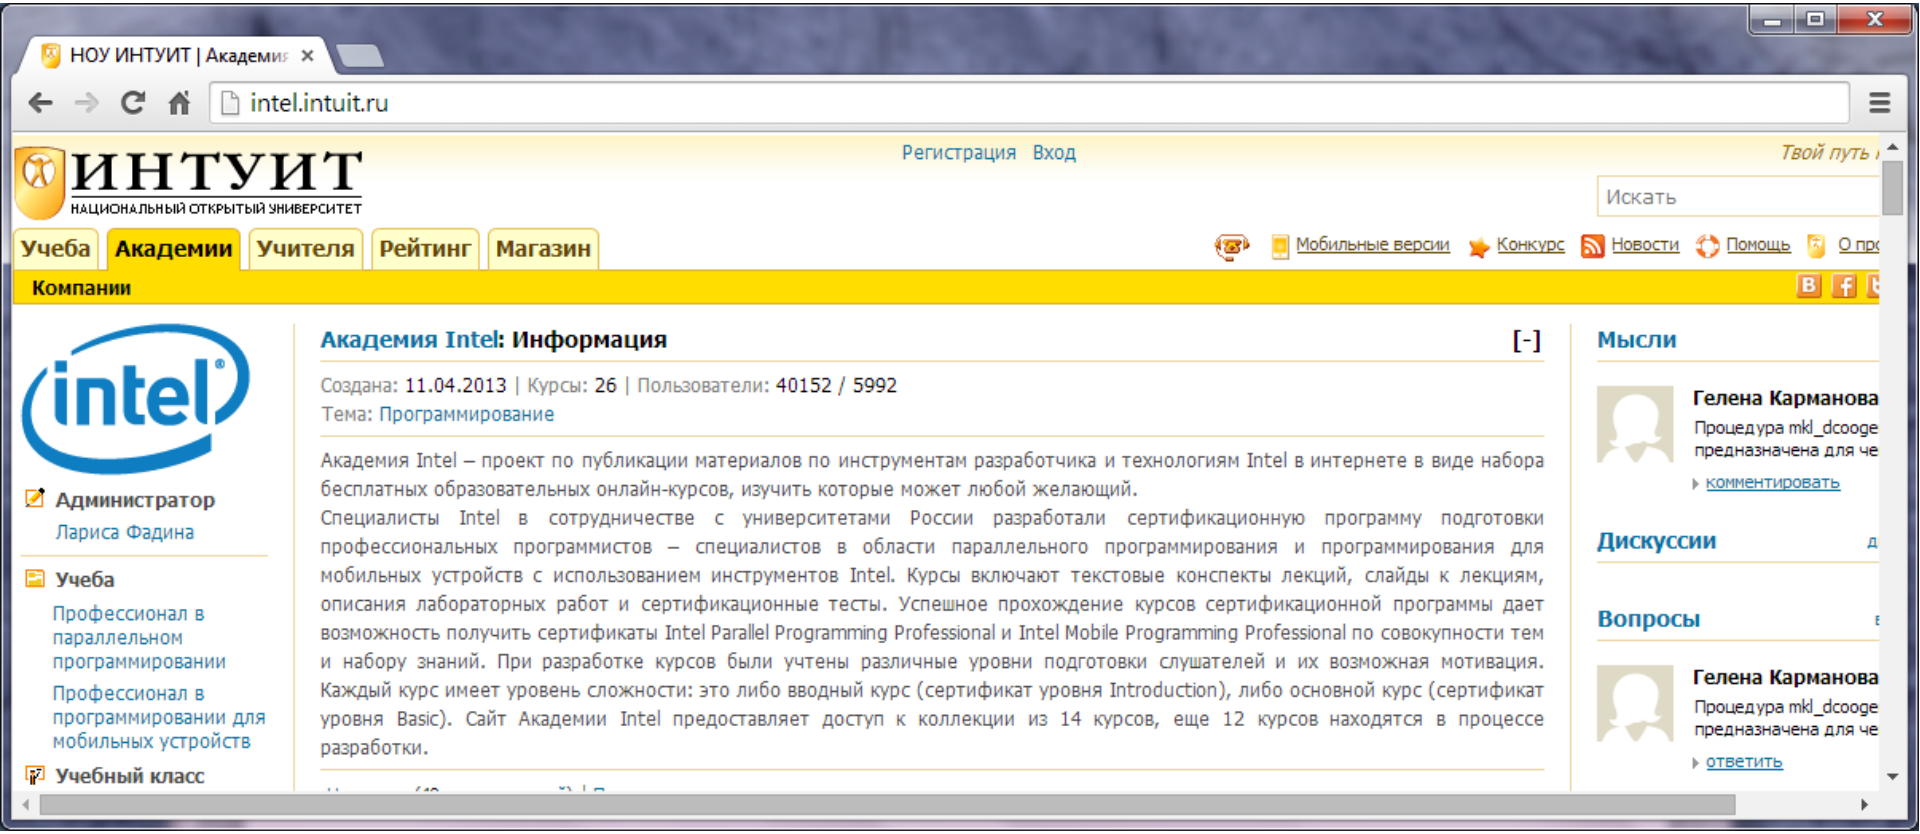

Тенденции: IT-кадры

IT-индустрия создает внешние информационные порталы для студентов и внутренние корпоративные университеты

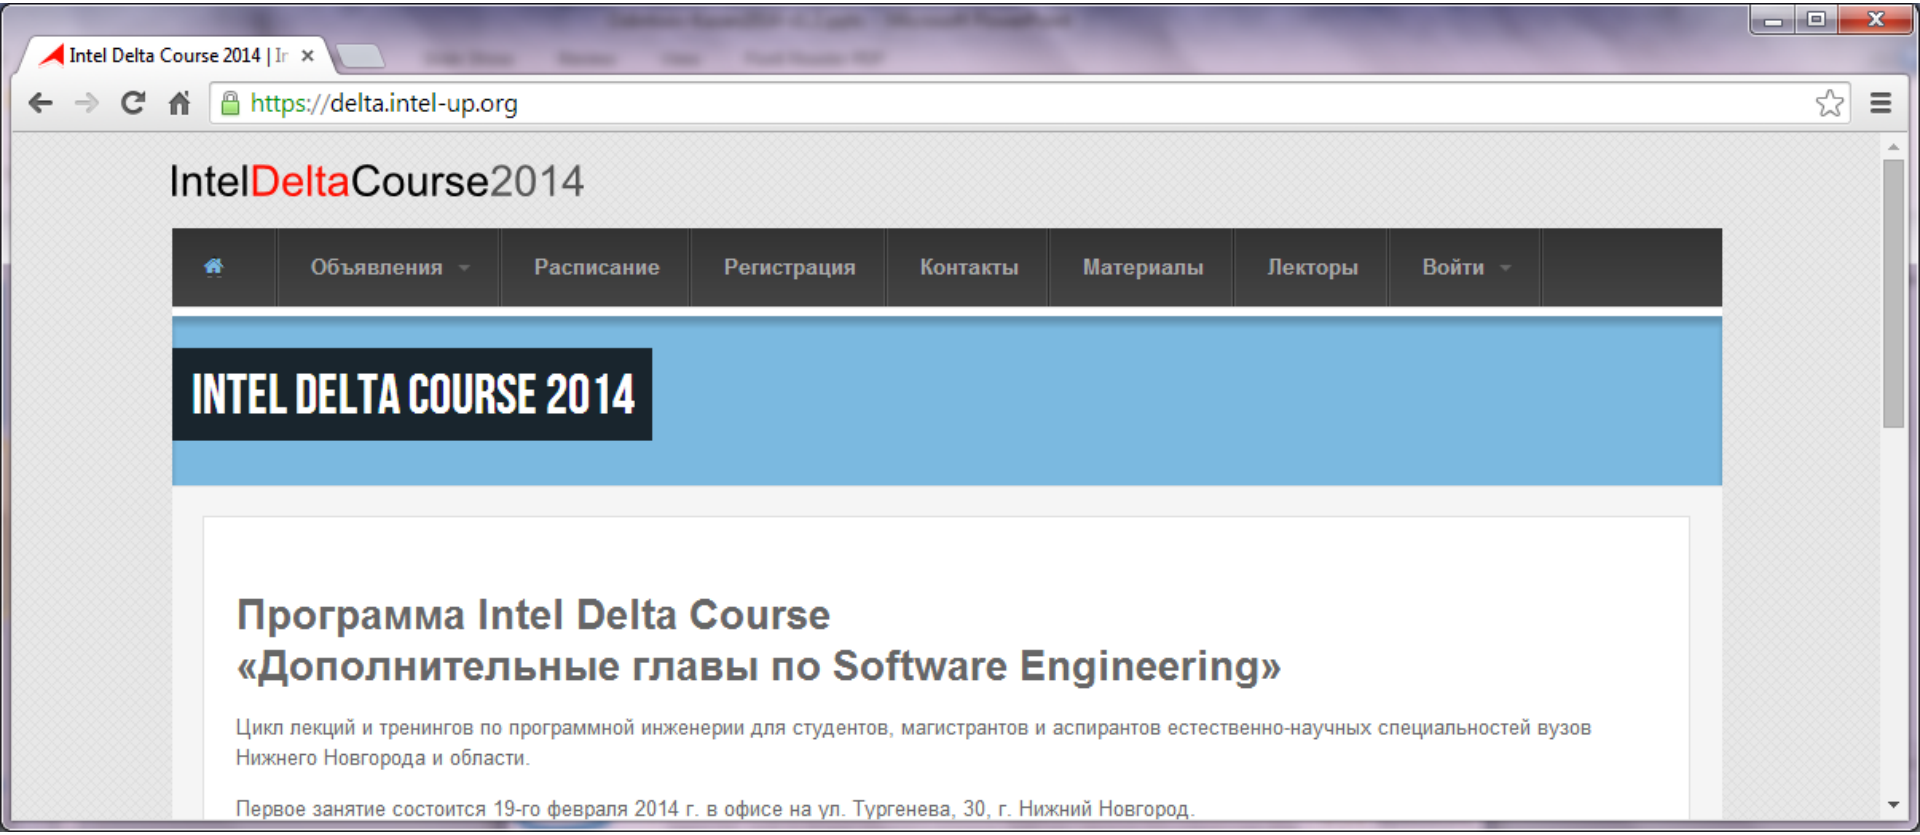

Тенденции: IT-исследования

IT-индустрия говорит науке, что надо исследовать

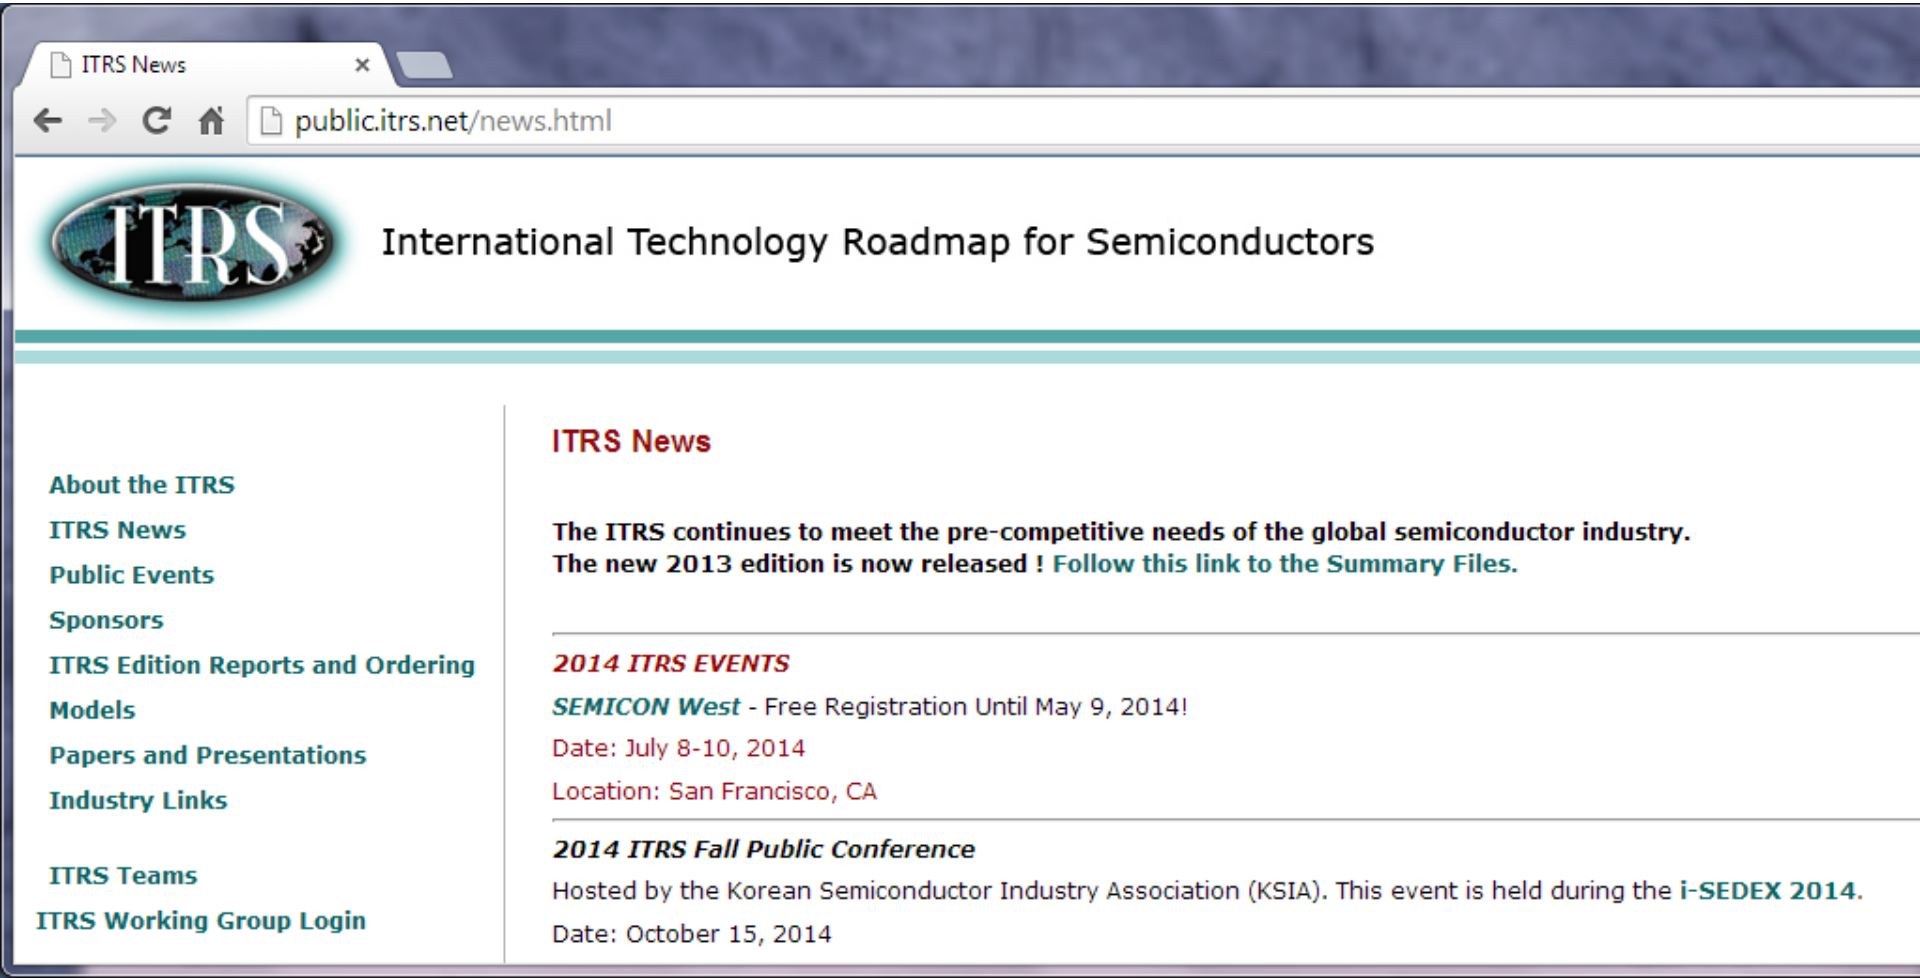

Современная IT-индустрия становится лидером образовательных активностей в IT-областях

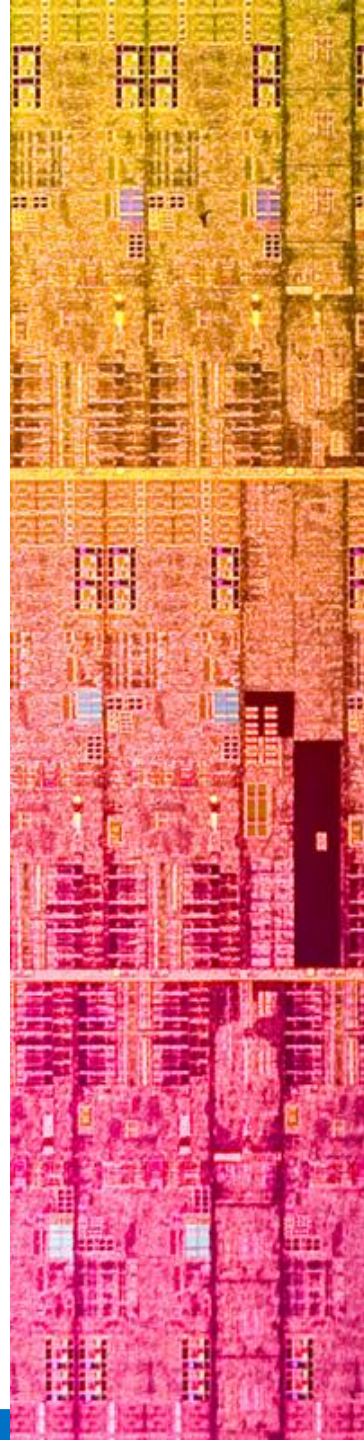

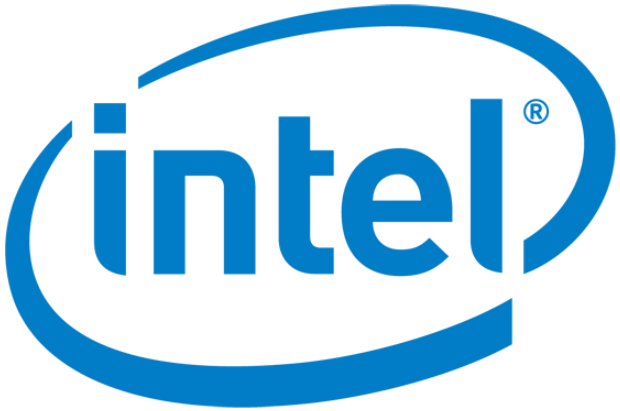
intel