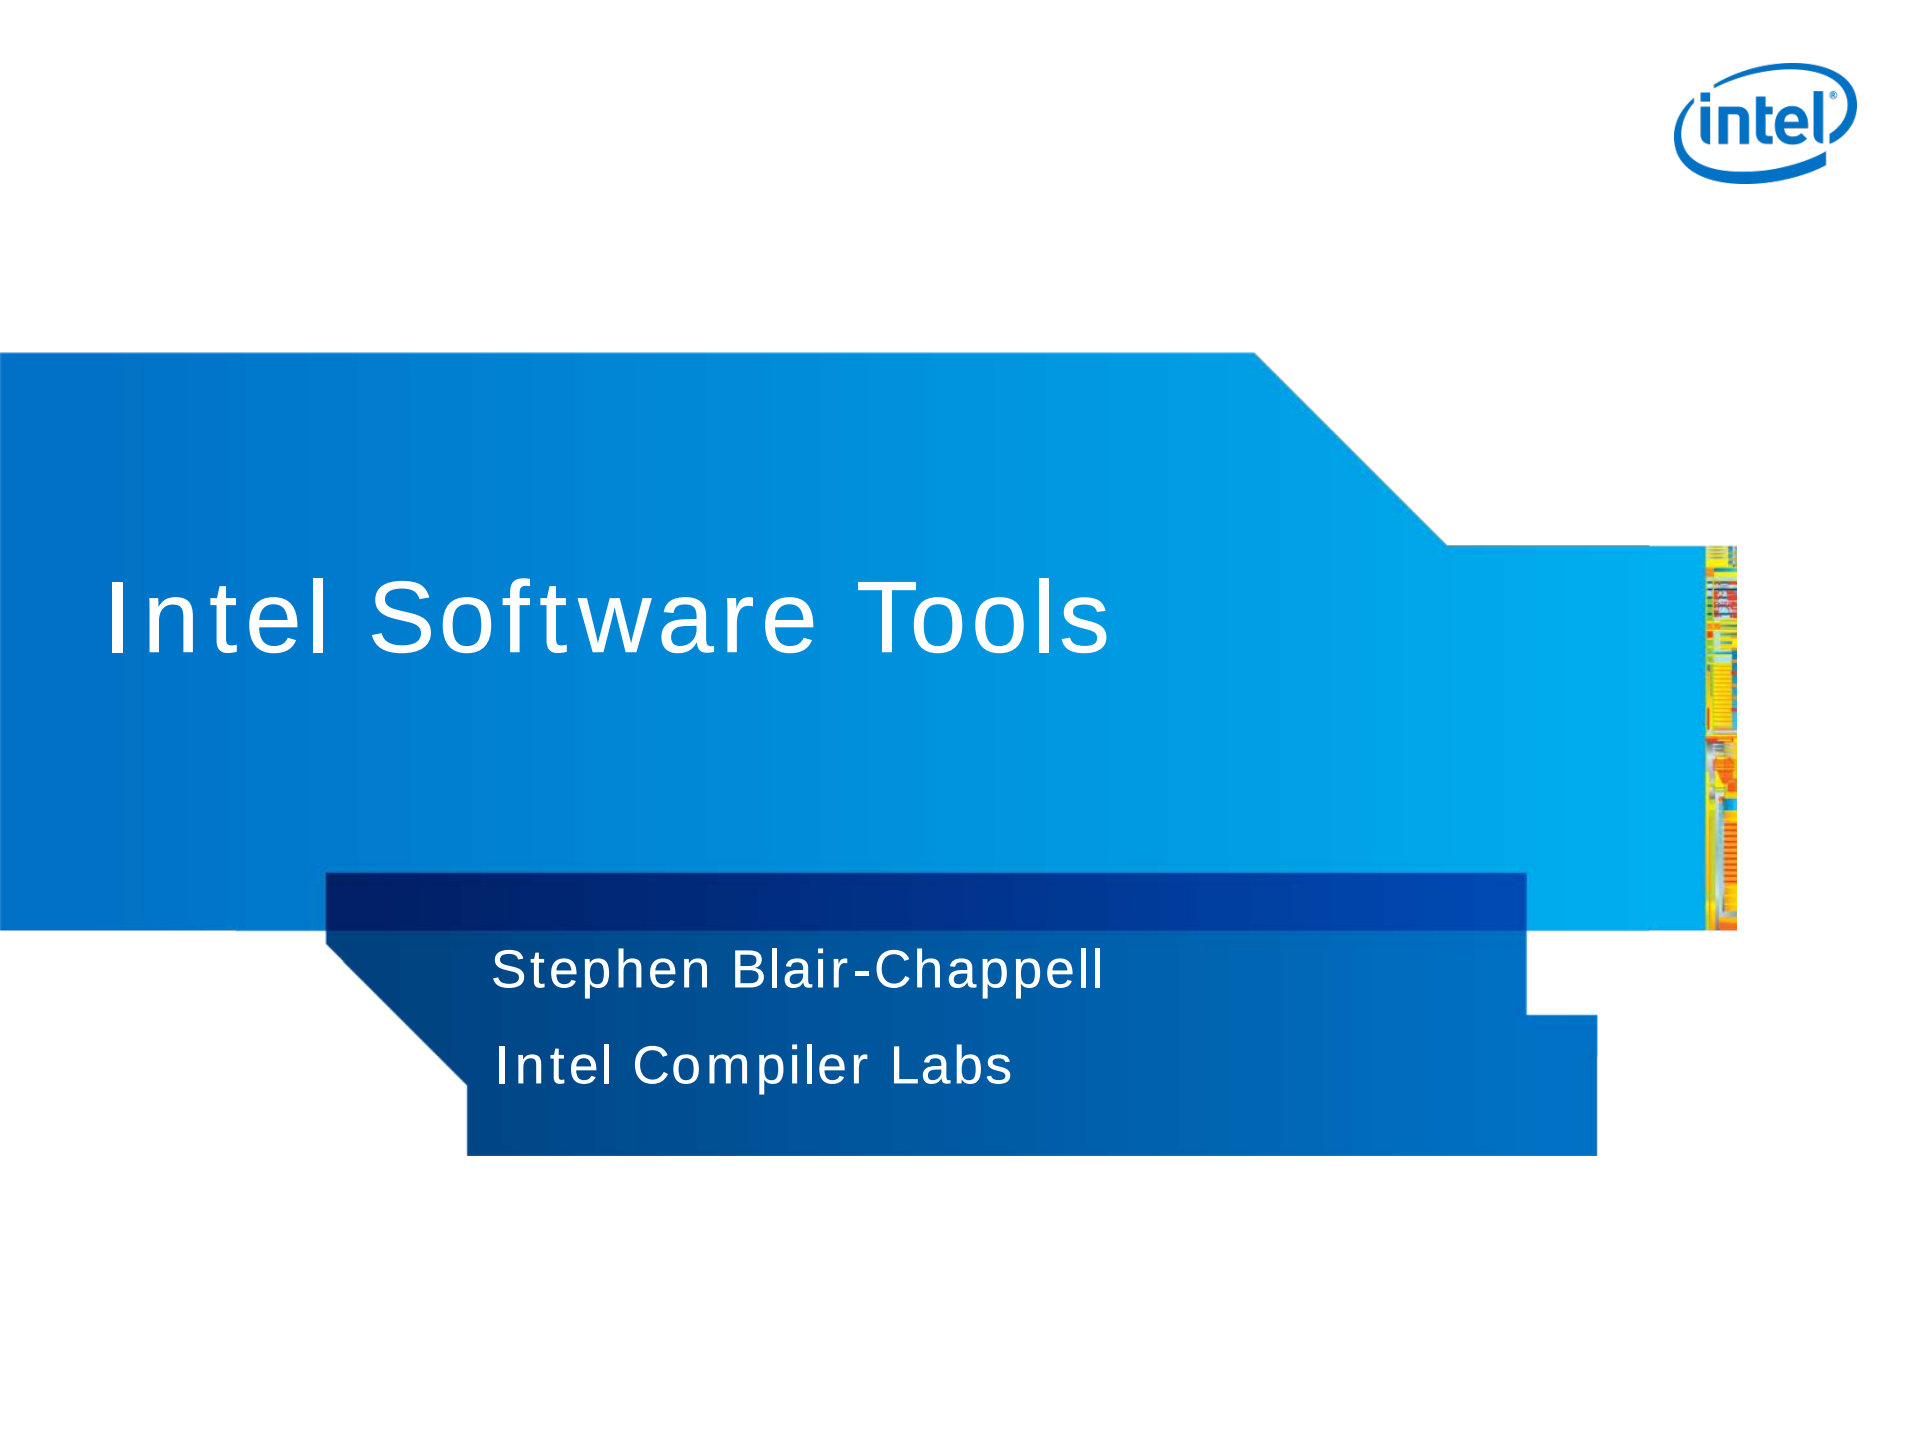

Intel Software Tools

Stephen Blair-Chappell

Intel Compiler Labs

Intel® Parallel Studio XE 2013 and Intel® Cluster Studio XE 2013

Helping Developers Efficiently Produce Fast, Scalable and Reliable Applications

More Cores. Wider Vectors. Performance Delivered.

Intel® Parallel Studio XE 2013 and Intel® Cluster Studio XE 2013

More Cores

Multicore Many-core

Xeon

Xeon Phi

50+ cores

Wider Vectors

128 Bits

256 Bits

512 Bits

Scaling Performance Efficiently

Serial Performance

Task & Data Parallel Performance

Distributed Performance

- Industry-leading performance from advanced compilers
- Comprehensive libraries
- Parallel programming models
- Insightful analysis tools

Intel® Parallel Studio XE 2013 and Intel® Cluster Studio XE 2013 †

Phase	Product	Feature	Benefit
Build	**Intel® Advisor XE**	Threading design assistant (Studio products only)	• Simplifies, demystifies, and speeds parallel application design
	Intel® Composer XE	• C/C++ and Fortran compilers • Intel® Threading Building Blocks • Intel® Cilk™ Plus • Intel® Integrated Performance Primitives • Intel® Math Kernel Library	• Enabling solution to achieve the application performance and scalability benefits of multicore and forward scale to many-core
	Intel® MPI Library†	High Performance Message Passing (MPI) Library	• Enabling High Performance Scalability, Interconnect Independence, Runtime Fabric Selection, and Application Tuning Capability
Verify & Tune	**Intel® VTune™ Amplifier XE**	Performance Profiler for optimizing application performance and scalability	• Remove guesswork, saves time, makes it easier to find performance and scalability bottlenecks
	Intel® Inspector XE	Memory & threading dynamic analysis for code quality Static Analysis for code quality	• Increased productivity, code quality, and lowers cost, finds memory, threading , and security defects before they happen
	Intel® Trace Analyzer & Collector†	MPI Performance Profiler for understanding application correctness & behavior	• Analyze performance of MPI programs and visualize parallel application behavior and communications patterns to identify hotspots

Efficiently Produce Fast, Scalable and Reliable Applications

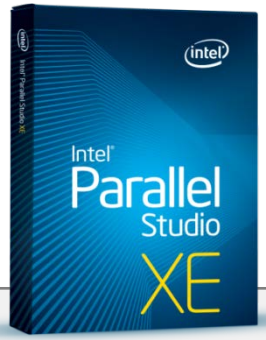

Top New Features

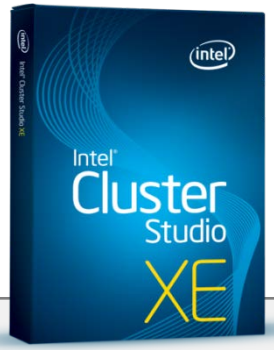

Performance	Performance Profiling	Reliability	Reproducibility	Standards	Parallelism Assistance
Improved compiler and library performance + Ivy Bridge microarchitecture + Haswell microarchitecture + Intel® Xeon Phi™ coprocessor	A dozen new analysis features Low overhead Java* profiling CPU Power Analysis	Pointer checker Heap growth analysis Improved MPI fault tolerance†	Conditional numerical reproducibility	Expanded C++ 11 Expanded Fortran 2008 MPI 2.2†	Analysis extended to include Linux*, Fortran and C# (in addition to Windows* and C/C++)

†Intel® Cluster Studio XE

Efficiently produce fast, scalable and reliable applications running on Windows* and Linux*

A Story …

A Bank near you!

The Reason Why

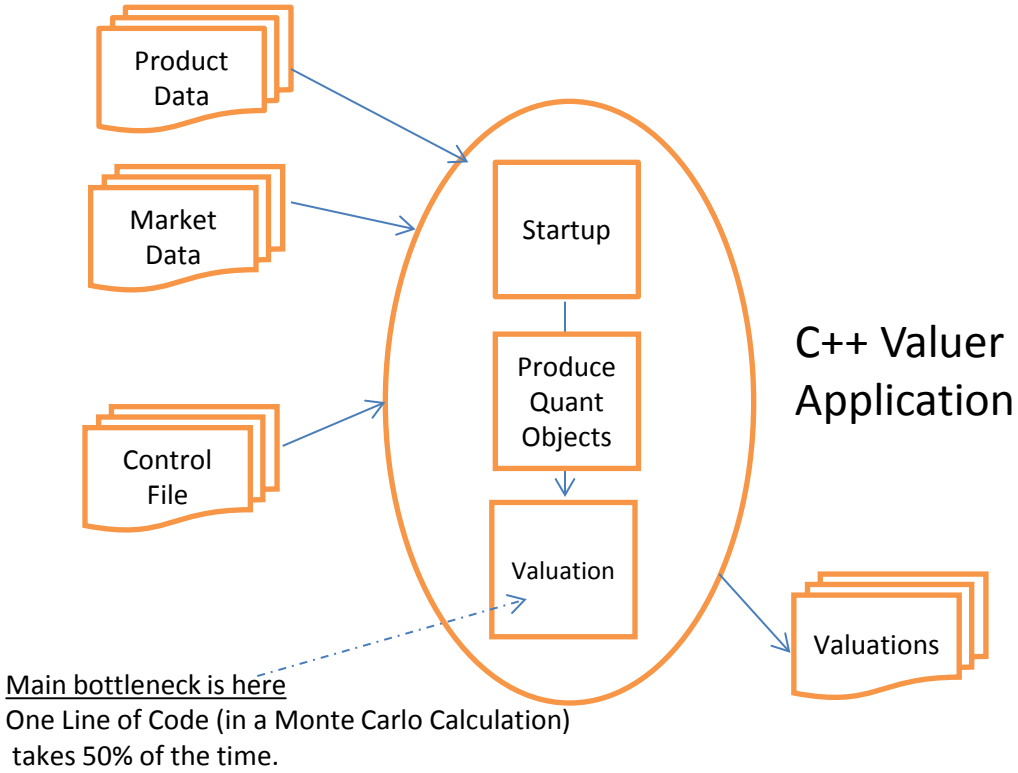

- Long overnight runtime
- Cost of renting space in data centres
- Power Consumption

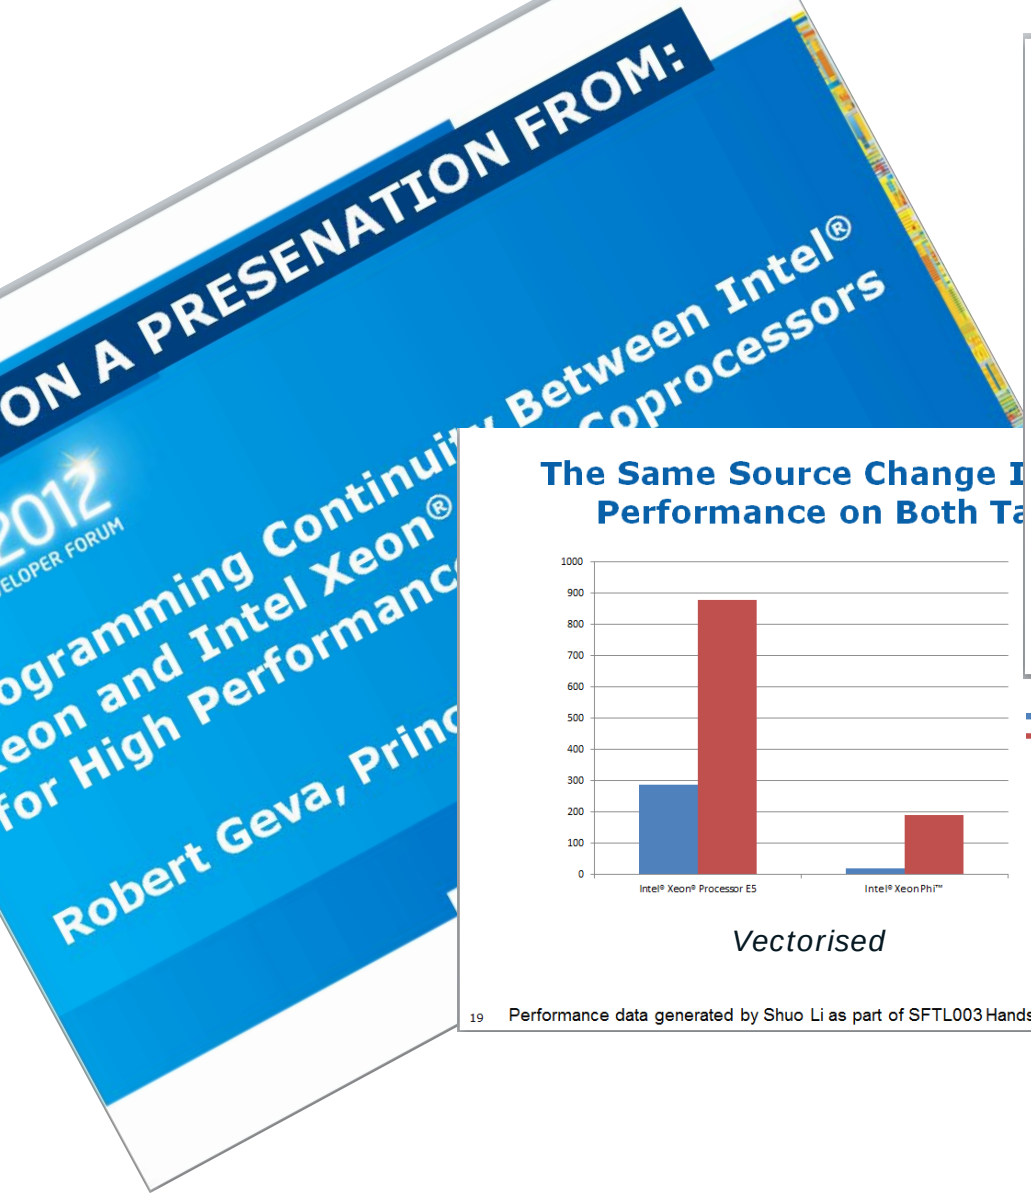

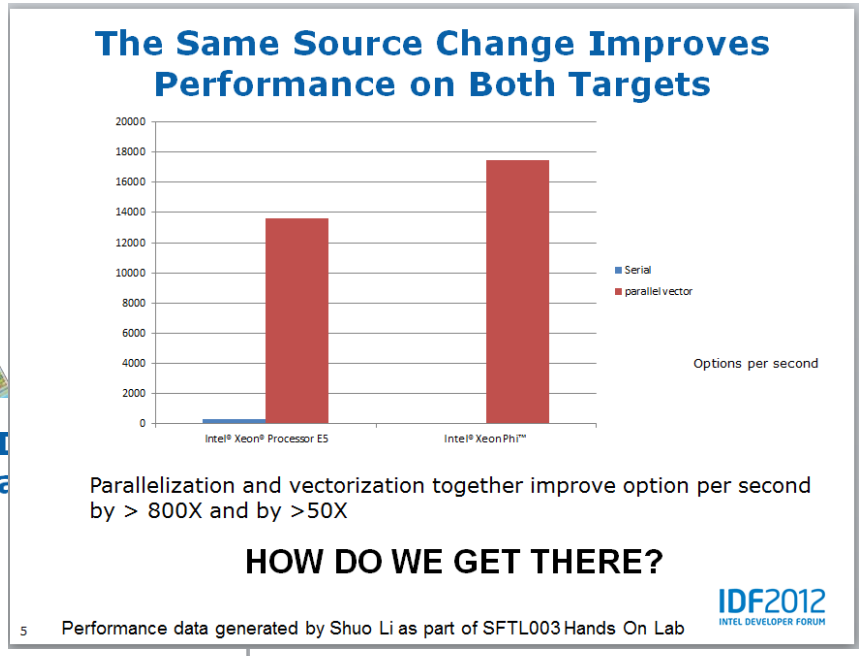

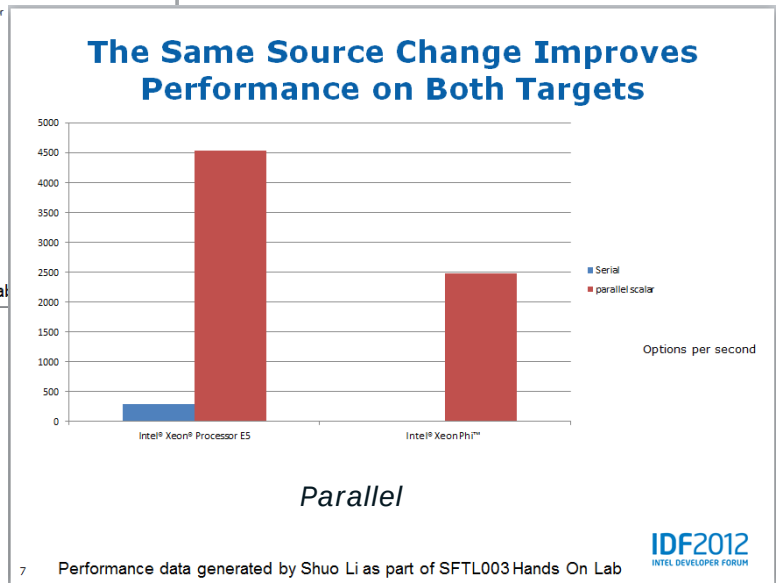

On the graphs, bigger is better

Timing Summary

	initialization	calculation	total
MS VS 10 - CL , base line	1324	627	1951
ICL , base line	1172	487	1659
ICL, vectorized	1161	202	1363
ICL, vectorized + OMP threading	612	105	717
ICL, vectorized + Cilk tasking	608	117	725
ICL, vectorized + OMP + MKL VSL	99	103	202

Source code freely available : Please contact presenter in case you want to have the source code and build scripts to reproduce the measurements

ArraySection version available from software.intel.com (search for Black-Scholes)

Three Common Requests

"How can I make my program run faster?"

"How can I make my program parallel?"

"Will my code run on any CPU? - compatibility"


Programs on Xeon Phi . . .
Code must be
highly Parallel
effectively Vectorised

Application Performance: Intel® Xeon Phi™ Coprocessor

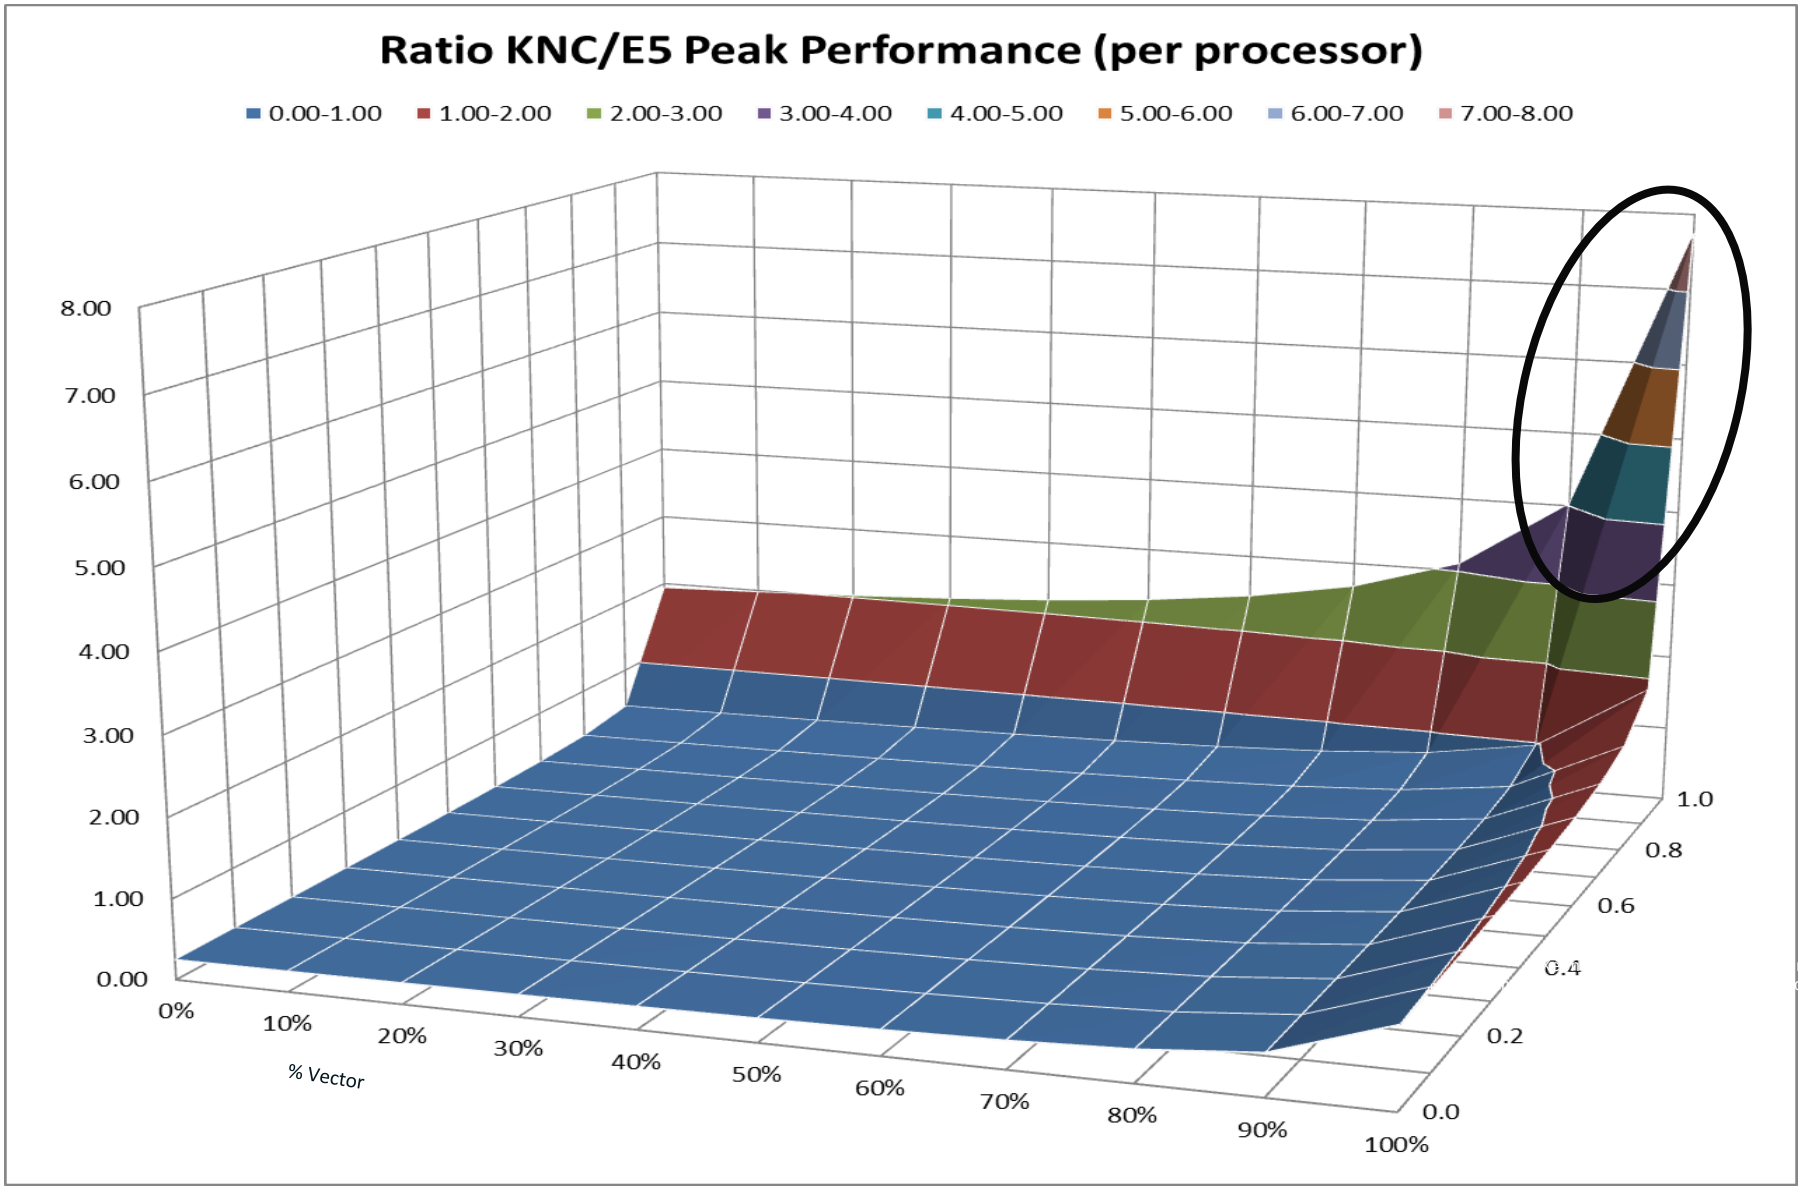

For more information go to http://www.intel.com/performance

And three more questions of late . . .

Will my code run on a Xeon Phi?

Do I have to change my code to much?

What performance will I get?

Intel® Parallel Studio XE

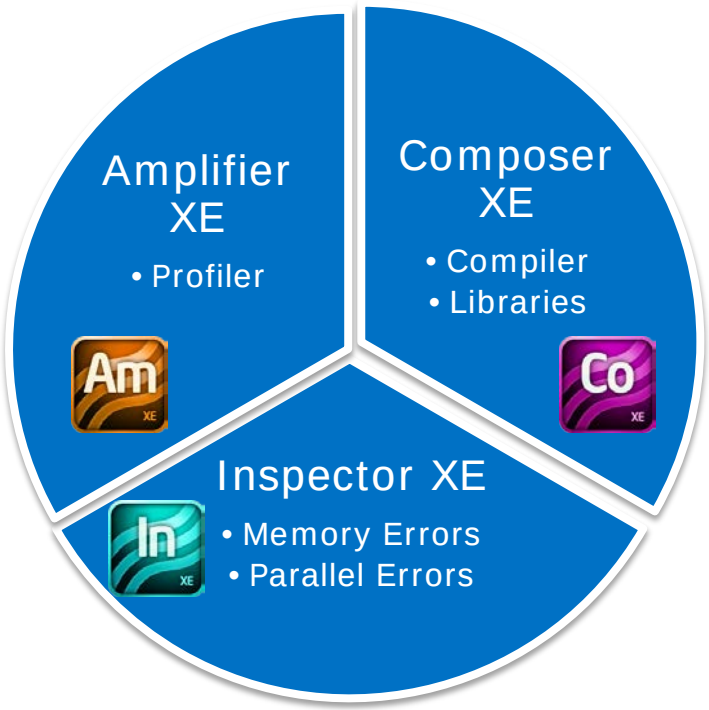

+ Advisor 

- **Intel® Parallel Advisor**
  - Use to model parallelism in your existing applications
- **Intel® Composer XE**
  - Use to generate fast, safe, parallel code (C/C++, Fortran)
- **Intel® VTune™ Amplifier XE**
  - Find hotspots and bottlenecks in you code.
- **Intel® Inspector XE**
  - Use to find memory and threading errors

Four Components

Three Common Requests

"How can I make my program run faster?"

"How can I make my program parallel?"

"Will my code run on any CPU? - compatibility"

The compiler uses many optimisation techniques

Faster Code

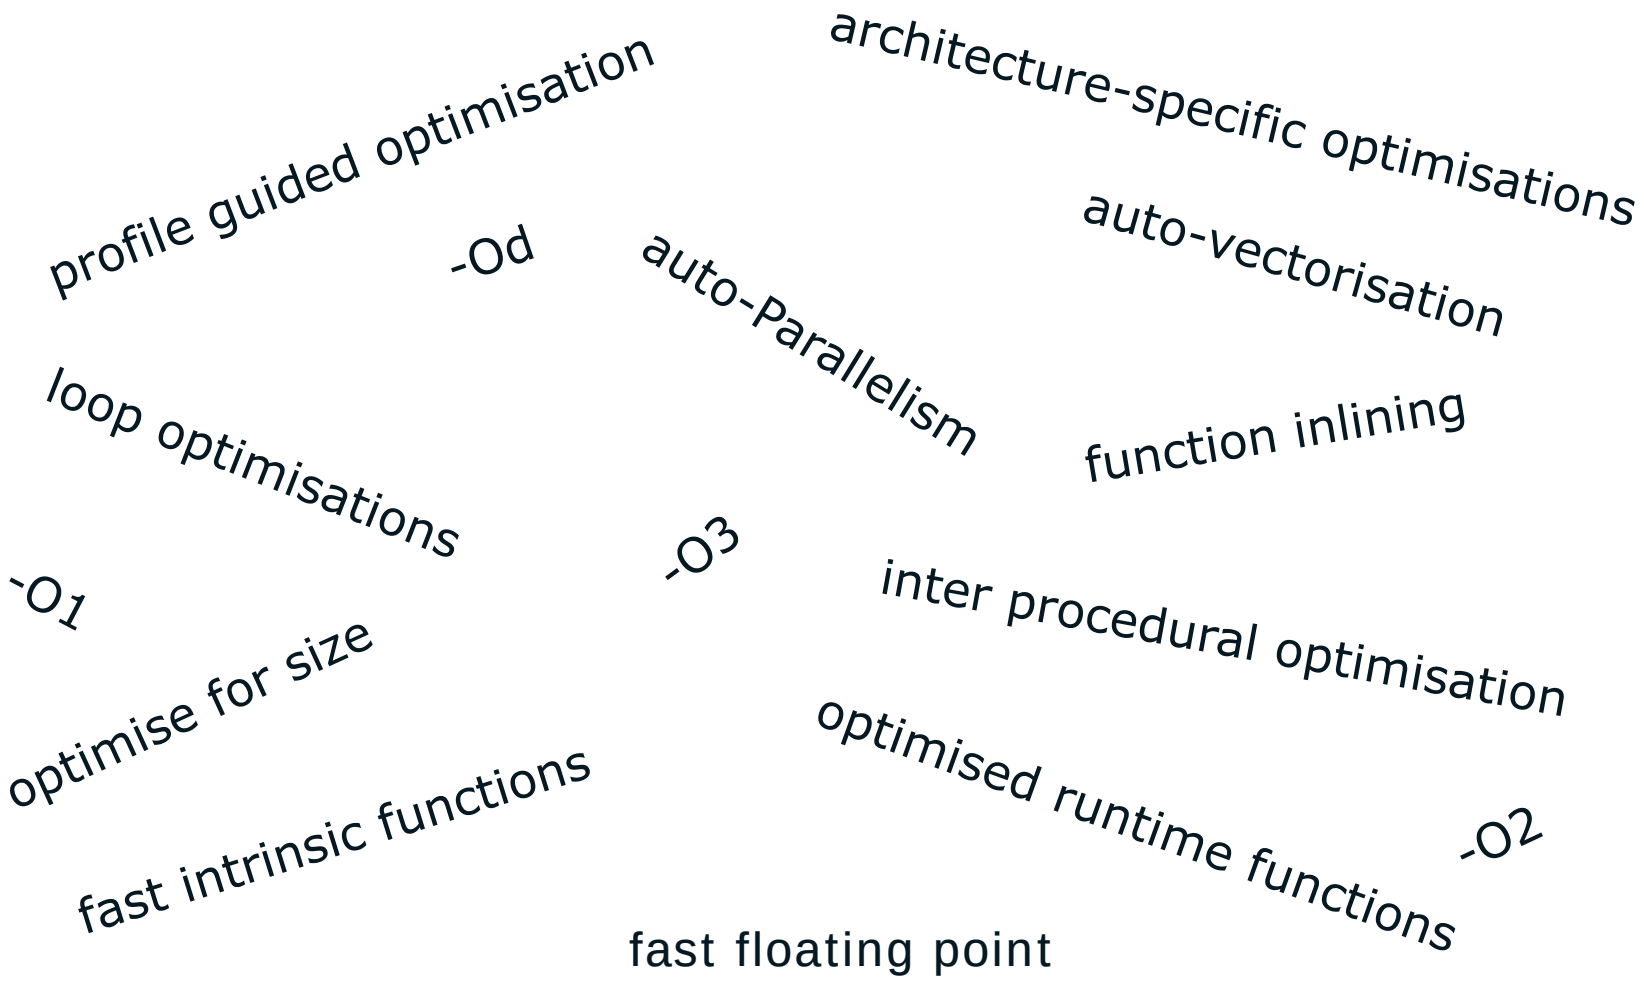

http://www.intel.com/content/www/us/en/architecture-and-technology/64-ia-32-architectures-optimization-manual.html
http://software.intel.com/sites/products/collateral/hpc/compilers/compiler_qrg12.pdf

Often we are happy with out-of-the-box experience

When was the last time you looked at some documentation?

SIMD Instruction Enhancements

1999	2000	2004	2006	2007	2008	2009	2011	2012\2013	2012
SSE	SSE2	SSE3	SSSE3	SSE4.1	SSE4.2	AES-NI	AVX	AVX2	MIC
70 instr Single-Precision Vectors Streaming operations	144 instr Double-precision Vectors 8/16/32 64/128-bit vector integer	13 instr Complex Data	32 instr Decode	47 instr Video Graphics building blocks Advanced vector instr	8 instr String/XML processing POP-Count CRC	7 instr Encryption and Decryption Key Generation	~100 new instr. ~300 legacy sse instr updated 256-bit vector 3 and 4-operand instructions	Int. AVX expands to 256 bit Improved bit manip. fma Vector shifts Gather	512-bit vector

```
for (i=0;i<MAX;i++)
   c[i]=a[i]+b[i];
```

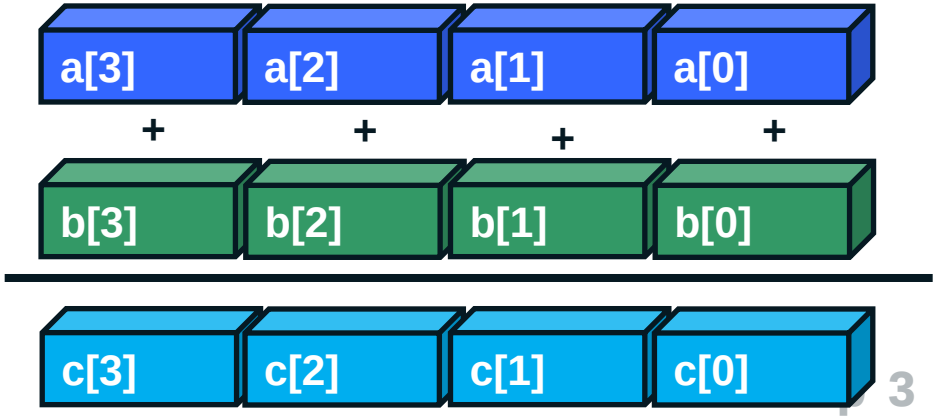

Other Ways of Inserting Vectorised Code

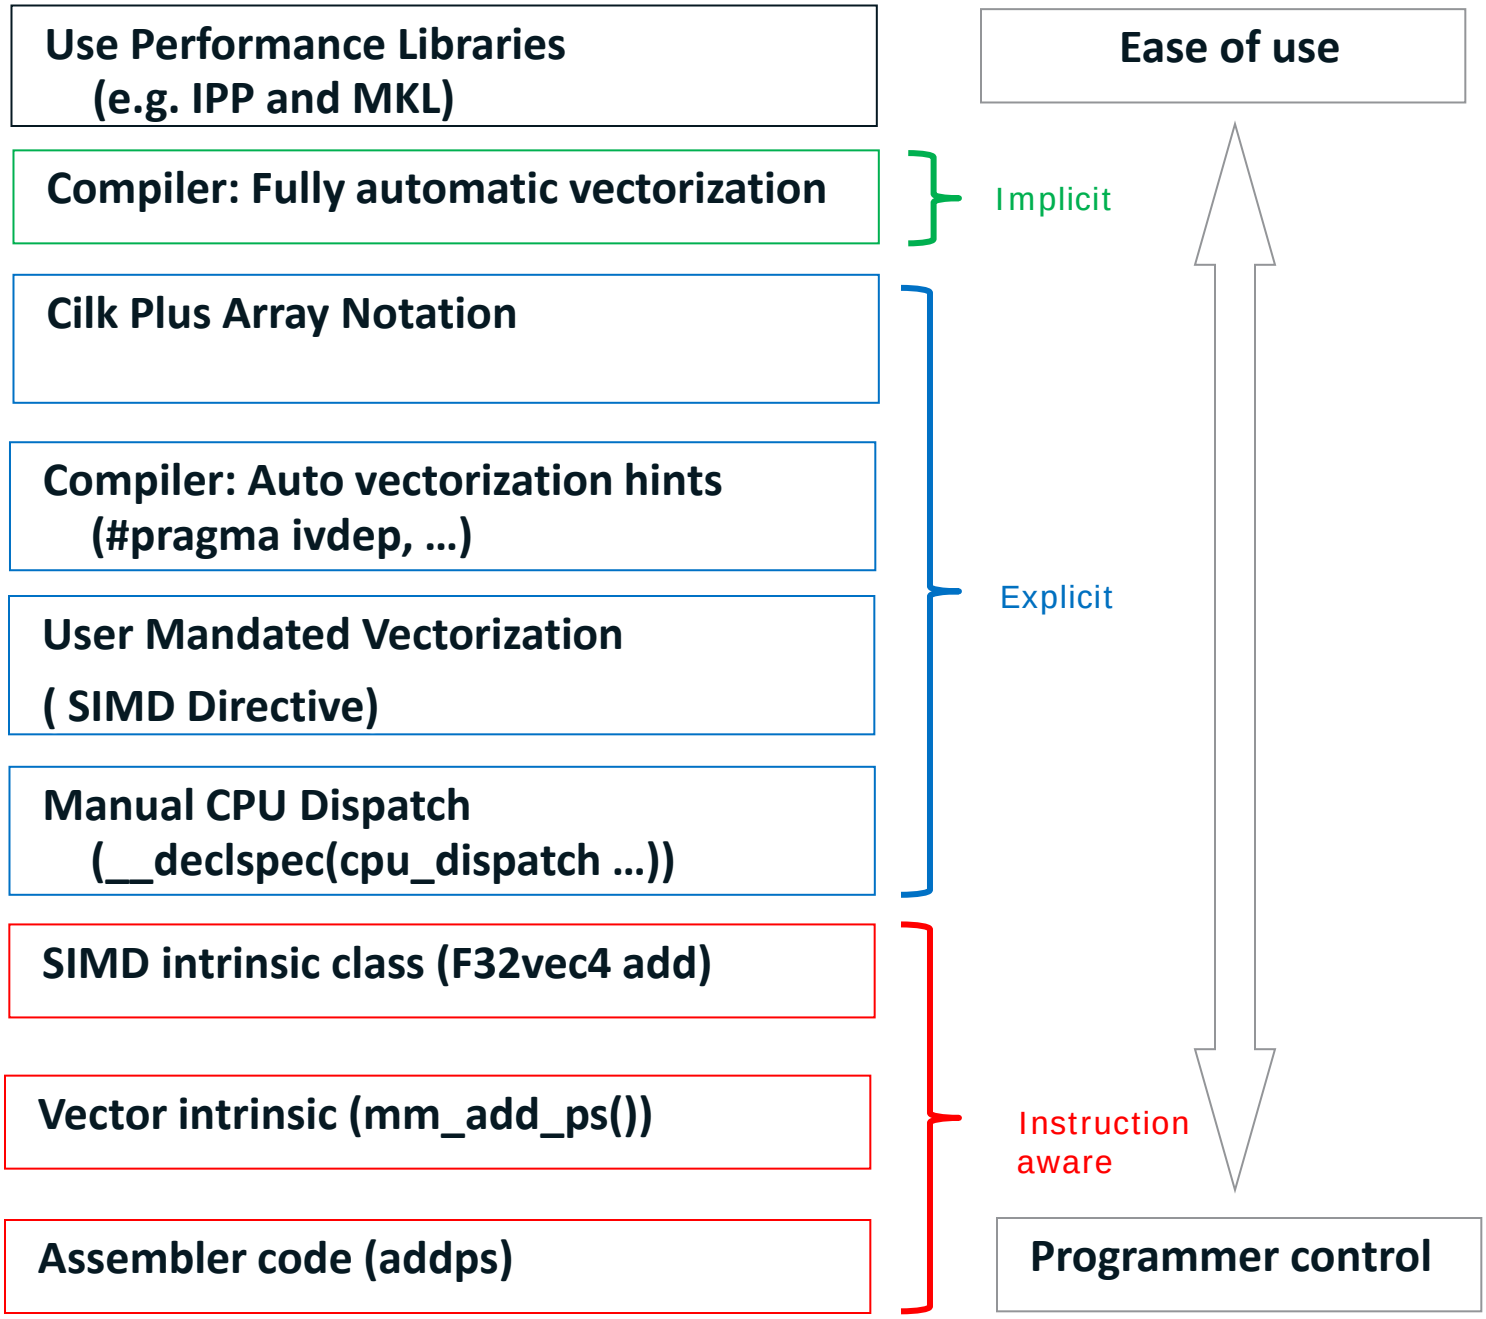

An example

Faster Code

Speedup by upgrading silicon

CPU	No Auto-Vectorisation	With Auto-Vectorisation	Speedup
P4	39.344	21.9	**1.80**
Core 2	5.546	0.515	**10.77**
Speedup	**7.09**	**45.52**	**76**

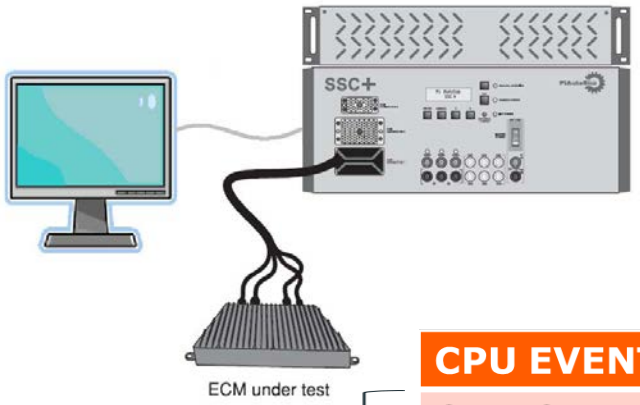

Speedup by swapping compiler

Verified using VTune

CPU EVENT	Without Vect	With Vect
CPU_CLK_UNHALTED.CORE	16,641,000,448	1,548,000,000
INST_RETIRED.ANY	3,308,999,936	1,395,000,064
X87_OPS_RETIRED.ANY	250,000,000	0
SIMD_INST_RETIRED	0	763,000,000

Full paper available here:
http://edc.intel.com/Link.aspx?id=1045

Three Common Requests

"How can I make my program run faster?"

"How can I make my program parallel?"

"Will my code run on any CPU? - compatibility"

Speedup using parallelism

Parallel Code

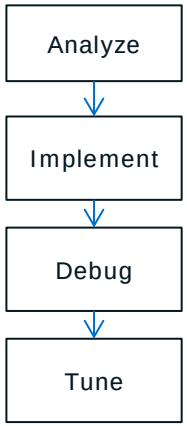

1 **Analyze**

- Amplifier XE
  - Hotspot
  - EBS (XE only)

2 **Implement**

- Composer XE
  - Compiler
  - Cilk Plus
  - OpenMP
  - Libraries
    - MKL
    - TBB
  - IPP

3 **Debug**

- Inspector XE
  - Threads
  - Memory

4 **Tune**

- Amplifier XE
  - concurrency
  - Locks & waits

Four Step Development

Language to help parallelism

Parallel Code

Intel® Cilk™ Plus

OpenMP

```
#pragma omp parallel for
for(i=1;i<=4;i++) {
    printf("Iter: %d", i);
}
```

Intel® Threading Building Blocks

Intel® MPI

Fortran Coarrays

OpenCL

Native Threads

```
cilk_for (int i = 0; i < max_row; i++)
{
    for (int j = 0; j < max_col; j++ )
    {
        p[i][j] = mandel( complex(scale(i), scale(j)));
    }
}
```

An example …

Parallel Code

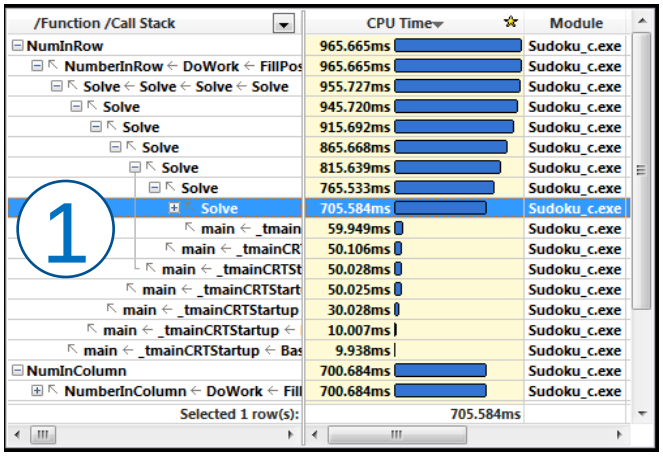

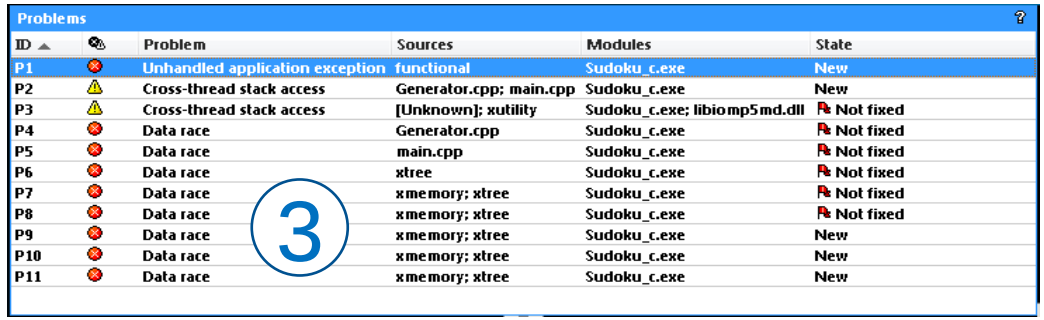

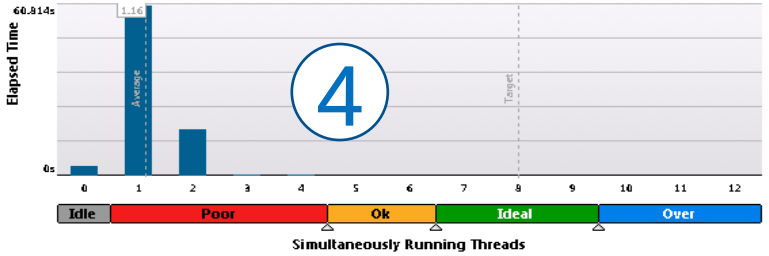

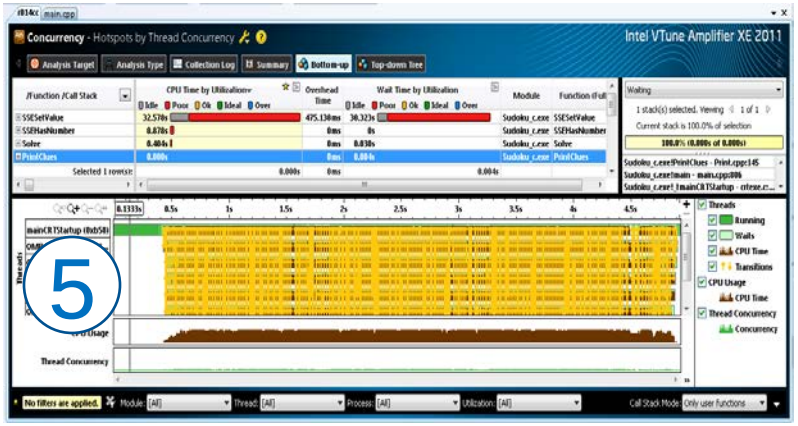

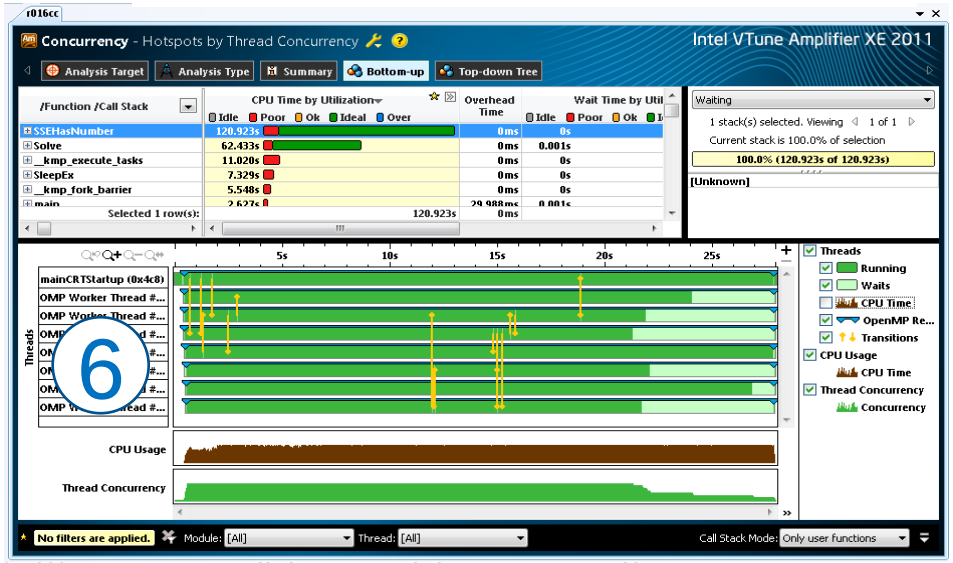

1. Hotspot Analysis
2. Implement
3. Find Threading Errors

4,5,6. Tune Parallelism

https://makebettercode.com/parallel_landing_required/lib/pdf/5373_IN_ParallelMag_Sudoku_060911.pdf

Three Common Requests

"How can I make my program run faster?"

"How can I make my program parallel?"

"Will my code run on any CPU? - compatibility"

Will my program run on any CPU?

Compatible Code

Compatibility

- run?

OS-agnostic
CPU-agnostic
Language / Standards
Tools
Scalability

Future Proofing

- build?
- Performance?

Thank You

Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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.

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

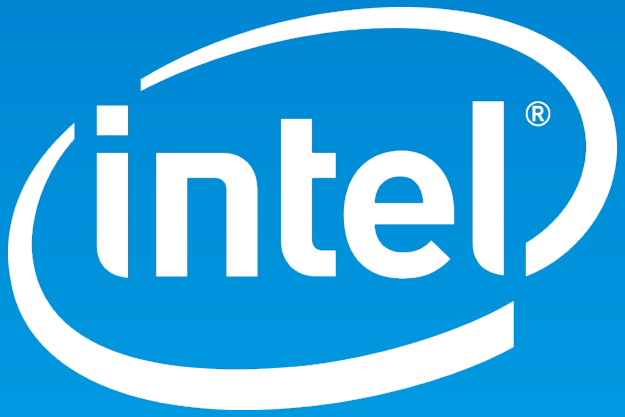
intel