



INTEL® PARALLEL STUDIO XE 2016

WHAT'S NEW

September 2015

Intel Software and Services Group

Intel® Parallel Studio XE 2016 Suites

Vectorization – Boost Performance By Utilizing Vector Instructions / Units

- Intel® Advisor XE - **Vectorization Advisor** identifies new vectorization opportunities as well as improvements to existing vectorization and highlights them in your code. It makes actionable coding recommendations to boost performance and estimates the speedup.

Scalable MPI Analysis– Fast & Lightweight Analysis for 32K+ Ranks

- Intel® Trace Analyzer and Collector add **MPI Performance Snapshot** feature for easy to use, scalable MPI statistics collection and analysis of large MPI jobs to identify areas for improvement

Big Data Analytics – Easily Build IA Optimized Data Analytics Application

- Intel® Data Analytics Acceleration Library (Intel® DAAL) will help data scientists speed through big data challenges with optimized IA functions

Standards – Scaling Development Efforts Forward

- Supporting the evolution of industry standards of **OpenMP***, **MPI**, **Fortran** and **C++** Intel® Compilers & performance libraries

DETAIL SLIDES

Intel® C/C++ and Fortran Compilers 16.0

Get best performance with latest standards

- More of C++14, generic lambdas, member initializers and aggregates
- More of C11, `_Static_assert`, `_Generic`, `_Noreturn`, and more
- `BLOCK_LOOP` Pragma/Directive Syntax
- Intel® Cilk™ Plus Combined Parallel & SIMD loops
- OpenMP 4.0 C++ User Defined Reductions, Fortran Array Reductions
- Vectorization enhancements
- OpenMP 4.1 asynchronous offloading, `simdlen`, `simd` ordered
- Offload Feature Enhancements for Intel® Xeon Phi™
- F2008 Submodules, Impure Elemental Functions
- F2015 `TYPE(*)`, `DIMENSION(..)`, `RANK` intrinsic, attributes for args with `BIND`

C++14 Standard Support

Generic Lambdas



```
auto glambda = [ ] (auto a) { return a; };
```

Generalized lambda captures



```
int x = 4;  
int z = [&r = x, y = x+1]  
{  
    r += 2;        // set x to 6; "R is for Renamed Ref"  
    return y+2;   // return 7 to initialize z  
}();             // invoke lambda
```

Digit Separators

[[deprecated]] attribute

Function return type deduction

Member initializers and aggregates

Feature Test macros



- Test for header files existence

- Test for compiler features

```
#if __has_include("shared_mutex") // use standard header  
#elif __has_include("boost/shared_mutex.h") // use from BOOST  
#endif  
  
...  
#ifndef __cpp_constexpr           // no constexpr available  
#elif __cpp_constexpr == 200704  // c++11 constexpr available  
#else                             // c++14 constexpr available  
#endif
```

C11 Standard Support

New Feature Support

- Unicode strings
- C11 anonymous unions

New Keyword Support

| | |
|-----------------------------|----------------------------|
| <code>_Alignas</code> | <code>_Alignof</code> |
| <code>_Static_assert</code> | <code>_Thread_local</code> |
| <code>_Noreturn</code> | <code>_Generic</code> |

`_Generic` : template like functionality in C w/o template support in the language
defines different actions based on different arguments types

```
#define pow(X) _Generic((X), long double: powl, \  
                        default: pow, \  
                        float: powf)(X)
```

when `pow()` is called with a long double argument it calls `powl()`
when `pow()` is called with a float, it calls `powf()`,
when `pow()` is called with other types it calls `pow()`

BLOCK_LOOP Pragma/Directive Syntax

C++:

```
#pragma block_loop [clause[,clause]...]
#pragma noblock_loop
```

Fortran:

```
!DIR$ BLOCK_LOOP [clause[[,] clause]...]
!DIR$ NOBLOCK_LOOP
```

clause → *factor*(*expr*), *level*(*levels*)

- Enables or disables loop blocking for the immediately following nested loops
- Works seamlessly with other directives including SIMD directive

```
#pragma block_loop factor(256) level(1:2)
for( j=1; j<n ; j++){
    f = 0;
    for( i=1; i<n; i++){
        f = f + a[i] * b[i];
    }
    c[j] = c[j] + f;
}
```



```
for( jj=1 ; jj<n/256+1; jj++){
    for( ii=1; ii<n/256+1; ii++){
        for( j=(jj-1)*256+1; min(jj*256, n); j++){
            f = 0;
            for ( i=(ii-1)*256+1; i<min(ii*256,n); i++){
                f = f + a[i] * b[i];
            }
            c[j] = c[j] + f;
        }
    }
}
```

Intel® Cilk™ Plus

Combined Parallel & SIMD loops

- New combined Parallel + Vector loop

```
_Cilk_for _Simd (int i = 0; i < N; ++i)  
    // Do something
```

- or

```
#pragma simd  
_Cilk_for (int i = 0; i < N; ++i)  
    // Do something
```

- Combined loop gives both thread-parallelism and vectorization
- Behaves approximately like this pair of nested loops

```
_Cilk_for (int i_1 = 0; i_1 < N; i_1 += M)  
    for _Simd (int i = i_1; i < i_1 + M; ++i)  
        // Do something
```

- The chunk size, M, is determined by the compiler and runtime

Optimization Notice

Improvements in Vectorization

Intel® Cilk™ Plus and OpenMP* 4.0

simdlen (i.e. vectorlength) and safelen for loops

- Usable with #pragma simd (Intel Cilk™ Plus) and omp simd (OpenMP*)

Array reductions

- Fortran only

User-defined reductions

- Supported for parallel in C/C++ for POD types. No support for Fortran, SIMD, or non-POD types (C++)

omp-simd collapse(N) clause

FP-model honoring for simd loops

- No more undesirable speculations and transformations

Adjacent gathers optimization

- Replaces series of gathers with series of vector loads & sequence of permutes

Ordered blocks in SIMD context

- ordered with simd specifies structured block simd loop or SIMD function

```
#pragma simd safelen(9) simdlen(4)
for (i = 0; i < N; i++) {
    a[i] = a[i] + a[i-9];
}
```

```
#pragma omp simd simdlen(4,8)
for (i = 0; i < N; i++) {
    foo(a[i]); //versions for VL=4,8
}
```

```
#pragma omp simd collapse(2)
for (i = 0; i < N; i++)
    for (j = 0; j < 4; j++) {
        a[5*i+j] += da[i]*cf[j];
    }
```

```
#pragma omp ordered [simd]
    structured code block
#pragma simdoff
    structured code block
```

OpenMP* 4.1 Extensions

Support for Features in OpenMP* 4.1 Technical Report 3

- Non-structured data allocation
 - `omp target [enter | exit] data`
 - Begins/ends data region for the target
 - `map` specifies data to be mapped/unmapped with motion
- Asynchronous offload
 - `nowait` clause on `omp task` and `target`
 - Enables asynchronous offload using existing tasking model with `nowait` clause
- Dependence (signal)
 - `depend` clause on `omp task` and `target`
 - Enables offload with dependence
- Map clause extensions
 - Modifiers `always` and `delete`
- Available for C/C++ and Fortran

Offload Feature Enhancements for Intel® Xeon Phi™

New Features available for C/C++

- Offload of structures with pointer members
 - Enabling offload inside member function
 - 'pointer->field' is now allowed in offload constructs
 - Structured objects are copied without pointer fields. This makes them bitwise copyable
 - Pointer fields transfer must be pointed out in the pragma explicitly
- Offload with MIC-only memory allocation using new modifiers
 - targetptr for MIC-only memory allocation by offload runtime
 - preallocated targetptr for informing offload runtime about user allocated memory on MIC
- Offload using Streams - new stream clause and associated APIs
 - Offload multiple concurrent computations to a coprocessor from a single CPU thread
 - stream is a logical queue of offloads
 - Offloads in any one stream complete in the order they were issued to the stream
- Performance Improvements for C/C++ and Fortran
- Asynchronous offload
- Memory Allocation and Data Transfers

Optimization Notice

Submodules (F2008) – The Problem

```
module bigmod
...
contains
subroutine sub1
...<implementation of sub1>

function func2
...<implementation of func2>

subroutine sub47
...<implementation of sub47>

...
end module bigmod
```

```
! Source source1.f90
use bigmod
...
Call sub1
```

```
! Source source2.f90
use bigmod
...
x = func2(...)
```

```
! Source source47.f90
use bigmod
...
call sub47
```

Submodules (F2008) – The Solution

```
module bigmod
...
interface
module subroutine sub1
...

module function func2
...

module subroutine sub47
...
end interface

end module bigmod
```

```
submodule (bigmod) bigmod_submod

contains

module subroutine sub1
... <implementation of sub1>

module function func2
... <implementation of func2>

module subroutine sub3
... <implementation of sub3>

end submodule bigmod_submod
```

Changes in the submodule do not force recompilation of uses of the module – as long as the interface does not change

IMPURE ELEMENTAL (F2008)

In Fortran 2003, ELEMENTAL procedures are PURE

- No I/O, no side-effects, can call only other PURE procedures

New IMPURE prefix allows non-PURE elemental procedures

- Can do I/O, call RANDOM NUMBER, etc.

Further C Interoperability (F2015)

TS29113 on “Further Interoperability of Fortran with C” to be part of Fortran 2015. Motivations include:

- Support needs of MPI3
- Provide Fortran equivalent of C’s “void*” – assumed type and rank
- Enable C code to manipulate array descriptors
- Extend interoperable interfaces to ALLOCATABLE, POINTER, OPTIONAL, assumed shape, character assumed length
- Extend ASYNCHRONOUS beyond I/O
- Relaxed restrictions

Uninitialized Variable Run-time Detection

Uninitialized variable checking using `[Q]init` option is extended to local, automatic, and allocated variables of intrinsic numeric type

Example:

```
4 real, allocatable, dimension(:) :: A
...
20 ALLOCATE (A(N))
...
49 do i = 1, N
50     Total = Total + A(I)
51 enddo
```

```
$ ifort -init=arrays,snan -g -traceback sample.F90 -o sample.exe
$ sample.exe
forrtl: error (182): floating invalid - possible uninitialized real/complex variable.
Image                PC                      Routine              Line                Source
...
sample.exe            0000000000402E12  MAIN__              50                  sample.F90
...
Aborted (core dumped)
```

Intel® Data Analytics Acceleration Library

2016

A C++ and Java API library of optimized analytics building blocks for all data analysis stages, from data acquisition to data mining and machine learning. Essential for engineering high performance Big Data applications.

New library targeting data analytics market

- **Customers:** analytics solution providers, system integrators, and application developers (FSI, Telco, Retail, Grid, etc.)
- **Key benefits:** improved time-to-value, forward-scaling performance and parallelism on IA, advanced analytics building blocks

Key features

- Building blocks highly optimized for IA to support all data analysis stages
- Support batch, streaming, and distributed processing with easy connectors to popular platforms (Hadoop, Spark) and tools (R, Python, Matlab)
- Flexible interfaces for handling different data sources (CSV, MySQL, HDFS, RDD (Spark))
- Rich set of operations to handle sparse and noisy data
- C++ and Java APIs

Data Management

Interfaces for data representation and access. Connectors to a variety of data sources and data formats, such as HDFS, SQL, CSV, ARFF, and user-defined data source/format.

Data Sources

Numeric Tables

Outliers Detection

Compression / Decompression

Serialization / Deserialization

Data Processing

Optimized analytics building blocks for all data analysis stages, from data acquisition to data mining and machine learning.

Data Modeling

Data structures for model representation, and operations to derive model-based predictions and conclusions.

Analysis

- PCA
- Variance-Covariance Matrix
- Distances
- Matrix decompositions (SVD, QR, Cholesky)
- EM for GMM
- Uni-/multi-variate outlier detection
- Statistical moments

Machine learning

- Linear regression
- Apriori
- K-Means clustering
- Naïve Bayes
- LogitBoost, BrownBoost, AdaBoost
- SVM

- Data layouts: AOS, SOA, homogeneous, CSR
- Data sources: csv, MySQL, HDFS/RDD
- Compression/decompression: ZLIB, LZO, RLE, BZIP2
- Serialization/deserialization

Optimization Notice

Copyright © 2015, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



Intel® Math Kernel Library (Intel® MKL) 11.3

Better performance with new two-stage API for Sparse BLAS routines

Additional Sparse Matrix Vector Multiplication API

- new two-stage API for Sparse BLAS level 2 and 3 routines

MKL MPI wrappers

- all MPI implementations are API-compatible but MPI implementations are **not ABI-compatible**
- MKL MPI wrapper solves this problem by providing an MPI-independent ABI to MKL

Support For Batched Small Matrix multiplication

- a single call executes multiple independent GEMM operation simultaneously

Support for Philox4x35 and ARS5 RNG

- two new pseudorandom number generators with a period of 2^{128} are highly optimized for multithreaded environment

Sparse Solver SMP improvements

- significantly improved overall scalability for Intel Xeon Phi coprocessors and Intel Xeon processors

Intel® Integrated Performance Primitives 9.0

Improved Threading Support

Additional optimization for Intel® Quark™, Intel® Atom™, and the processors with Intel® AVX2 instructions support

- Intel® Quark™: data compression, cryptography optimization
- Intel® Atom™: computation vision, image processing optimization
- Intel® AVX2: computer vision, image processing optimization

New APIs to support external threading



Improved CPU dispatcher

- Auto-initialization. No need for the CPU initialization call in static libraries.
- Code dispatching based on CPU features

Optimized cryptography functions to support SM2/SM3/SM4 algorithm

Custom dynamic library building tool

New APIs to support external memory allocation



Optimization Notice

Copyright © 2015, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Intel® Threading Building Blocks 4.3

Improved cross-platform support

Fully supported `tbb::task_arena`

- Task arenas provide improved control over workload isolation and the degree of concurrency

Dynamic replacement of standard memory allocation routines for OS X*

- Utilize the powerful TBB scalable allocator easily on OS X

Binary files for 64-bit Android* applications were added as part of the Linux* OS package

- Improvements to the Flow Graph features

Don't forget to check out [Flow Graph Designer](#)

- Several Improvements to examples and documentation

Intel® Advisor XE – New! Vectorization Advisor

Data Driven Vectorization Design

Have you:

- Recompiled with AVX2, but seen little benefit?
- Wondered where to start adding vectorization?
- Recoded intrinsics for each new architecture?
- Struggled with cryptic compiler vectorization messages?

Breakthrough for vectorization design

- What vectorization will pay off the most?
- What is blocking vectorization and why?
- Are my loops vector friendly?
- Will reorganizing data increase performance?
- Is it safe to just use pragma simd?

The screenshot shows the 'Threading and Vectorization Survey' window in Intel Advisor XE 2016. It displays a table of function call sites and loops, categorized by 'Vectorized' and 'Not Vectorized'. The table includes columns for 'Function Call Sites and Loops', 'Self Time', 'Total Time', 'Memory analysis', 'Compiler Vectorization', 'Loop Type', 'Gain Estimate', and 'Vectorized Loops'. A specific loop is highlighted with a gain estimate of 2.19727 and is vectorized using SSE2.

| Function Call Sites and Loops | Self Time | Total Time | Memory analysis | Compiler Vectorization | Loop Type | Gain Estimate | Vectorized Loops |
|-------------------------------|-----------|------------|-----------------|------------------------|-----------|---------------|------------------|
| [loop at mmult_serial.cp... | 0.000s | 10.100s | | Vectorized ... | Scalar | 2.19727 | SSE2 |
| [loop at mmult_serial.cp... | 0.000s | 10.100s | | | Scalar | | SSE2 |
| [loop in __libc_start_mai... | 0.000s | 10.100s | | | Scalar | | |

| Function Call Sites and Loops | Total Time % | Total Time | Self Time | Hot Loops | Vector... Loops | Location |
|-------------------------------|--------------|------------|-----------|-----------|-----------------|-----------------------------|
| Total | 100.0% | 10.100s | 0s | | | |
| lib_start_main | 100.0% | 10.100s | 0s | | | libc-2.12.so |
| [loop in __libc_st... | 100.0% | 10.100s | 0s | | | libc-2.12.so |
| main | 100.0% | 10.100s | 0s | | | mmult_ser..._1_mmult_serial |
| [loop at mm... | 100.0% | 10.100s | 0s | | | mmult_ser..._1_mmult_serial |
| [loop at m... | 100.0% | 10.100s | 0s | | SSE2 | mmult_ser..._1_mmult_serial |
| multiply_d | 100.0% | 10.100s | 0.0600s | | | mmult_ser..._1_mmult_serial |
| [loop ... | 99.4% | 10.040s | 10.0400s | | SSE2 | mmult_ser..._1_mmult_serial |

More Performance
Fewer Machine Dependencies

Intel® Advisor XE – Vectorization Advisor

Provides the data you need for high impact vectorization

Compiler diagnostics + Performance Data = All the data you need in one place

- Find “hot” un-vectorized or “under vectorized” loops.
- Trip counts

Recommendations – How do I fix it?

Correctness via dependency analysis

- Is it safe to vectorize?

Memory Access Patterns analysis

- Unit stride vs Non-unit stride access, Unaligned memory access, etc.

Intel® VTune™ Amplifier XE 2016

Enhanced GPU and Microarchitecture Profiling

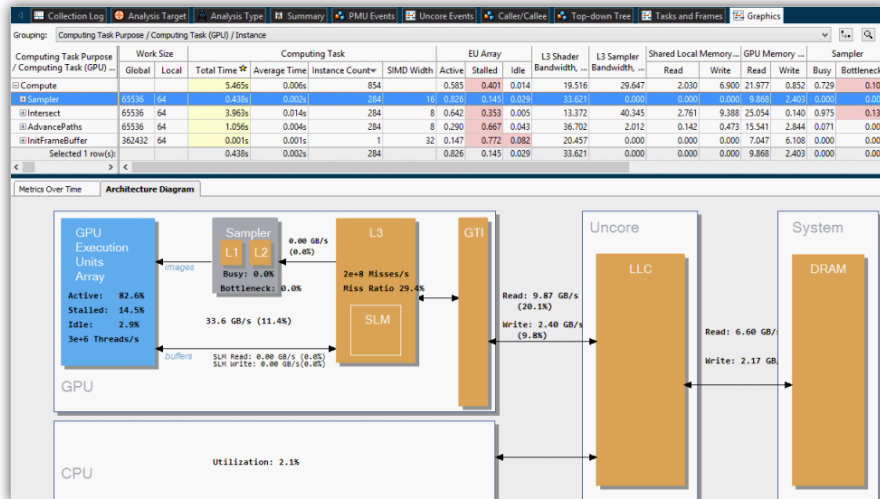
New OS and IDE support: Visual Studio* 2015 & Windows* 10 Threshold

GPU profiling

- GPU Architecture Annotation Diagram
- GPU profiling on Linux (OpenCL, Media SDK)

Microarchitecture tuning

- General Exploration analysis with confidence indication
- Driverless 'perf' EBS with stacks



Intel® VTune™ Amplifier XE 2016

Improved OpenMP and Hybrid Support

Intel OpenMP analysis enhancements

- Precise trace-based imbalance calculation that is especially useful for profiling small region instances
- Classification and issue highlighting of potential gains, e.g., imbalance, lock contention, creation overhead, etc.
- Detailed analysis of barrier-to-barrier region segments

| OpenMP Region / Function / Call Stack | OpenMP Potential Gain | | | | | | OpenMP Potential Gain (% of Collection Time) | | | | | | Elapsed Time | Number of OpenMP threads | Inst. Count | |
|---|-----------------------|-------------|----------|------------|----------|-------|--|-------------|------------|----------------|------------|-----------|--------------|--------------------------|-------------|--|
| | Imbalance | Lock Con... | Creation | Scheduling | Reduc... | Other | Imbalance (%) | Lock Con... | Cre... (%) | Scheduling (%) | Red... (%) | Other (%) | | | | |
| @@@_grad_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f314495 | 0.20% | 0s | 0.00% | 3.12% | 0s | 0s | 0.6% | 0.0% | 25.9% | 0.0% | 0.0% | 0.0% | 11.75s: 24 | 76 | 1 | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f315231 | 0.07% | 0.00% | 0s | 0s | 0s | 0s | 0.6% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.285s: 24 | 1 | 1 | |
| @@@Serial - outside any region | | | | | | | | | | | | | | | | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f339345 | 0.00% | 0s | 0s | 0s | 0s | 0s | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.001s: 24 | 75 | 1 | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f361365 | 0.00% | 0s | 0s | 0s | 0s | 0s | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.001s: 24 | 75 | 1 | |

MPI+OpenMP: multi-rank analysis on a compute node

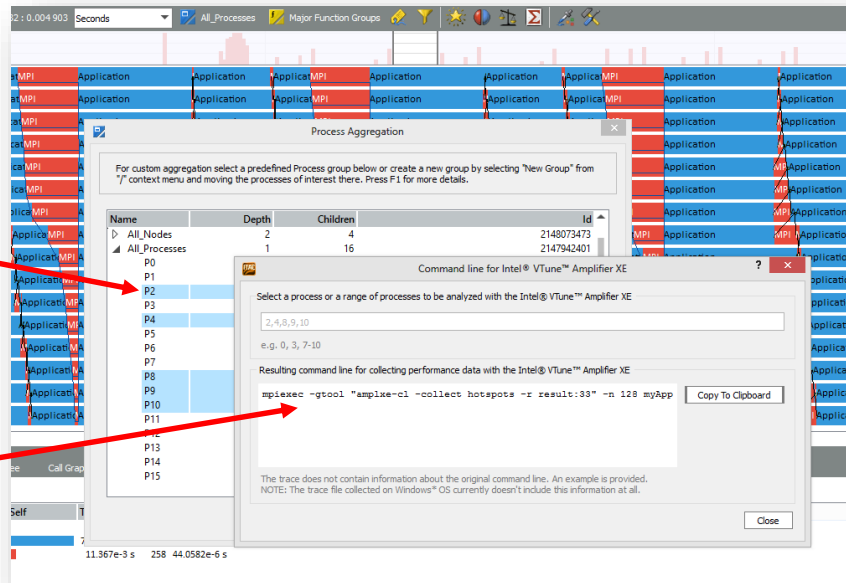
- Per-rank OpenMP potential gain and serial time metrics
- Per-rank Intel MPI communication busy wait time detection

| OpenMP Region / Function / Call Stack | OpenMP Potential Gain | | | | | | OpenMP Potential Gain (% of Collection Time) | | | | | | Elapsed Time | Number of OpenMP threads | Inst. Count | |
|--|-----------------------|-------------|--------|--------|--------|-------|--|-------------|------------|------------|------------|-----------|--------------|--------------------------|-------------|--|
| | Imbalance | Lock Con... | Cre... | Sch... | Red... | Other | Imbalance (%) | Lock Con... | Cre... (%) | Sch... (%) | Red... (%) | Other (%) | | | | |
| @@@_cong_grad_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f514495 | 3.04% | 0s | 0.00% | 0.00% | 0.00% | 0.00% | 34.6% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 11.095s: 24 | 76 | 1 | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f315231 | 0.56% | 0s | 0s | 0s | 0s | 0.00% | 0.9% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.286s: 24 | 1 | 1 | |
| @@@Serial - outside any region | | | | | | | | | | | | | | | | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f339345 | 0.00% | 0s | 0s | 0s | 0s | 0s | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.001s: 24 | 75 | 1 | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f361365 | 0.00% | 0s | 0s | 0s | 0s | 0s | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.001s: 24 | 75 | 1 | |
| @@@_MA2N_Somp\$parallel24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f363269 | 0.00% | 0s | 0s | 0s | 0s | 0s | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.0% | 0.00s: 24 | 1 | 1 | |

Improved Hybrid Profiling Support

Rank Selection/Multi-selection in ITAC for Profiling in VTune

- Automatically generate MPI command line script
- User will run command line script and automatically launch profiling of selected ranks by VTune




Run tool on specified ranks with gtool


```
$ mpirun -gtool "amplxe-cl -c advanced-hotspots \  
-r my_dir:all=node-wide" -n 4 -ppn 2 my_mpi_app
```

MPI Performance Snapshot

Scalable profiling for MPI and Hybrid



Lightweight – Low overhead profiling for 32K+ Ranks



Scalability- Performance variation at scale can be detected sooner



Identifying Key Metrics – Shows PAPI counters and MPI/OpenMP imbalances

MPI Performance Snapshot Summary

Application: ./hybrid
Ranks: 4
Used statistics: pcs_r4_f020.txt

Overview

| | |
|-----------------------------|--------|
| MPI Time: 14.70 sec | 48.99% |
| MPI Imbalance: 14.69 sec | 48.98% |
| Computation Time: 15.30 sec | 50.98% |
| OpenMP Time: 14.90 sec | 49.66% |
| OpenMP Imbalance: 7.37 sec | 24.57% |
| Serial Time: 0.40 sec | 1.32% |



Memory Usage

| | |
|-----------------------------------|---------|
| Peak memory consumption (rank 1): | 0.79 MB |
| Mean memory consumption: | 0.75 MB |

Per process memory usage affects the application scalability.

Performance by Metric

WallClock time: 30.00 sec
Total application lifetime. The time is elapsed time for the slowest process. This metric is the sum of the MPI Time and the Computation time below.

MPI Time: 14.70 sec 48.99%
Time spent inside the MPI library. High values are usually bad. This value is HIGH. The application is Communication-bound. [More details...](#)

MPI Imbalance: 14.69 sec 48.98%
Mean unproductive wait time per-process spent in the MPI library calls when a process is waiting for data. This time is part of the MPI time above. High values are usually bad. This value is HIGH. The application workload is NOT well balanced between MPI ranks. [More details...](#)

Computation Time: 15.30 sec 50.98%
Mean time per-process spent in the application code. This is the sum of the OpenMP Time and the Serial time. High values are usually good. This value is AVERAGE. The application is Computation-bound. [More details...](#)

OpenMP Time: 14.90 sec 49.66%
Mean time per process spent in the OpenMP parallel regions. High values are usually good and indicate that the application is well-threaded. This value is AVERAGE.

OpenMP Imbalance: 7.37 sec 24.57%
Mean unproductive wait time per-process spent in OpenMP parallel regions (normally at synchronization barriers). High values are usually bad. This value is HIGH. The application's OpenMP work sharing is NOT well load-balanced. [More details...](#)

Serial Time: 0.40 sec 1.32%
Mean application time per-process spent outside OpenMP parallel regions. High values may be good or bad depending on the application algorithm. This value is NEGLIGIBLE. This application is well parallelized via OpenMP directives.

Optimization Notice

Intel® Cluster Checker 3.0

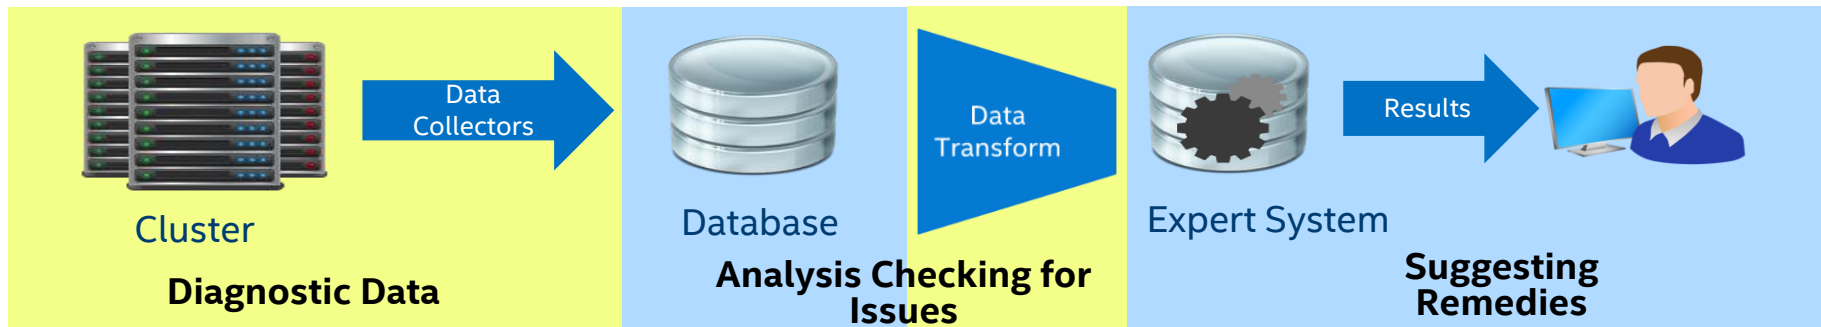
Cluster system expertise packaged into a utility

Provides Assistance

- Cluster Health Checks (on-demand, background)
- Diagnoses and remedies for common issues
- Compliance with Intel® Cluster Ready spec

Simplifies Cluster Computing Platforms

- Reduces need for specialized expertise
- Enables cluster health checks by applications
- Extensible and customizable, API



HPC Webinars

| Date | Title |
|-------------------------|--|
| Sep 1, 9:00 AM Pacific | What's New in Intel® Parallel Studio XE 2016? |
| Sep 15, 9:00 AM Pacific | Build Fast and Cross-Platform applications with Intel® Performance Libraries |
| Sep 22, 9:00 AM Pacific | Strategies for Tuning Multi-Level Parallelism |
| Sep 29, 9:00 AM Pacific | Faster Big data Analytics Using New Intel® Data Analytics Acceleration Library |
| Oct 6, 9:00 AM Pacific | Using Nested Parallelism in OpenMP* with James Reinders |
| Oct 13, 9:00 AM Pacific | From Serial to Awesome: Transforming Serial Code to Parallel Using Threading and Vectorization |
| Oct 20, 9:00 AM Pacific | From Serial to Awesome: Advanced Code Vectorization and Optimization |
| Oct 27, 9:00 AM Pacific | Performance Validation on Intel® Xeon® Processor and Xeon Phi™ Coprocessor Clusters |
| Nov 10, 9:00 AM Pacific | Nurture Happy Clusters With Intel® Cluster Checker |

<https://software.intel.com/en-us/events/hpc-webinars>

Optimization Notice

Copyright © 2015, Intel Corporation. All rights reserved.

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



Summary/Call to Action

Current Users:

- Go to <https://registrationcenter.intel.com> and get the latest version

New Users:

- Evaluation available at <https://software.intel.com/intel-parallel-studio-xe>

All Users:

- Product forums at <https://software.intel.com/forum>





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.

Copyright © 2015, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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