


# Agenda

Time	Title
<b>09:00-09:15</b>	Welcome, Introduction
09:15-09:45	Overview – Intel Processor Architecture Evolution
09:45-10:15	Intel's Software Development offerings at a glance
10:15-11:15	Strategies for Parallelism with Intel® Threading Methodologies
<b>11:15-12:30</b>	<b>Intel® Composer XE – the powerful compiler and performance libraries collection</b>
<b>12:30-13:30</b>	<i>Break</i>
<b>13:30-14:00</b>	Intel® Inspector XE - Detect memory and threading errors
<b>14:00-14:30</b>	Intel® Amplifier XE - Understand performance issues
<b>14:30-15:00</b>	Intel® Advisor XE – Get started with parallelization
<b>15:00-16:00</b>	Intel Development Tools for Multi-threading in a typical Development Cycle – Life-Demo – Wrap-up
<b>16:00- ...</b>	Q&A, opens, discussion



# Intel® Composer XE Overview

Hubert Haberstock  
Technical Consulting Engineer  
Intel GmbH

# What is the Composer XE?

- New name for the Intel compilers
  - Intel® Composer XE for Windows\*|Linux\*|Mac OS\* X (formerly Intel® Compiler Suite Professional Edition for Windows\*|Linux\*|Mac OS\* X)
  - Intel® C++ Composer XE for Windows\*|Linux\*|Mac OS\* X (formerly Intel® Compiler Professional Edition for Windows\*|Linux\*|Mac OS\* X )
  - Intel® Visual Fortran Composer XE for Windows\* (formerly Intel® Visual Fortran Compiler Professional Edition for Windows\*)
  - Intel® Fortran Composer XE for Linux\*|Mac OS\* X (formerly Intel® Visual Fortran Compiler Professional Edition for Linux\*|Mac OS\* X)
- Part of the new product bundle  
Intel® C++|Fortran|Parallel Studio XE
  - but also available separately
- Includes performance libraries, debugger, debugger extension, new threading methodologies



# Intel Composer XE 2013 SP1 - Overview

## Compilers, Performance Libraries, Debugging Tools

### Intel® C++

#### Composer XE 2013 SP1

- Intel® C++ Compiler XE 14.0
- Intel® Debugger + gdb add-ons
- Intel® Threading Building Blocks
- Intel® Cilk™ Plus
- Intel® Math Kernel Library
- Intel® Integrated Performance Primitives

### Intel® Fortran

#### Composer XE 2013 SP1

- Intel® Fortran Compiler XE 14.0
- Intel® Debugger + gdb add-ons
- Intel® Math Kernel Library
- Intel® Integrated Performance Primitives

Windows\*, Linux\*, Mac OS\*

- **Leading Performance Optimizing Compilers**

- Intel® C++ with Intel® Cilk™ Plus and Fortran Compiler
- Intel® Integrated Performance Primitives
- Intel® Math Kernel Library libraries

- **Standard Support**

- OpenMP\* 4.0
- Support for key parts of the latest Fortran and C++ standards, Visual Studio\* Shell for Visual Fortran\*

- **Compatibility – Mix and Match**

- Windows: Visual\* C++ and Visual Studio\* 2008, 2010, 2012
- Linux\*, OS X\*, gcc and, for C++ Eclipse & Xcode for Mac

- **Architecture Support**

- Intel compatible IA processors
- Intel® Xeon Phi™ product family,



# Leadership Application Performance



- More Performance for your C++ applications

- Just recompile
- Uses Intel® AVX and Intel® AVX2 instructions
- Intel® Xeon Phi™ product family support (Linux)
- Intel® Cilk™ Plus: Tasking and vectorization

NEW

NEW

- More Performance for your Fortran applications

- Just recompile
- Intel® Xeon Phi™ product family: Linux compiler, debugger support
- Access to Intel® AVX and Intel® AVX2 instructions (-xa or /Qxa)
- Auto-parallelizer & directives to access SIMD instructions
- Coarrays & synchronization constructs support parallel programming
- Loop optimization directives: VECTOR, PARALLEL, SIMD
- More control over array data alignment (align arrayNbytes)
- New in 2013 XE SP1 release: more Fortran 2008 support

NEW

NEW


NEW

NEW

NEW

# Up to 4x Faster Performance

with Intel® Advanced Vector Extensions 512 (Intel® AVX-512) Support



Intel® Compilers and Intel® Math Kernel Library will be updated in Q4 with AVX-512 support

- Significant leap to 512-bit SIMD support
- Increased compatibility with AVX
- One byte longer EVEX prefix, enabling additional functionality
- First implemented in the future Intel® Xeon Phi™ coprocessor, code named **Knights Landing**

Enables higher performance for the most demanding computational tasks

# Expanded C++ 11 support

New in SP1 (14.0 compilers)

- Unrestricted unions (Linux\*, OS X\*)
- Non-static data member initializers
- Explicit Virtual Overrides
- Allowing move constructors to throw
- Defining Move Special Member Functions
- Inline namespaces
- Rvalue references v2

Full list of actual C++ 11 support:

<http://software.intel.com/en-us/articles/c0x-features-supported-by-intel-c-compiler>

Full C++11 Support planned for next release

Excellent Support for C++ 11 on Windows\* and Linux\*

# Standards Conformance – C99

- Full support:
  - Restricted pointers (**restrict** keyword).
  - Variable-length Arrays
  - Flexible array members
  - Complex number support (**\_Complex** keyword)
  - Hexadecimal floating-point constants
  - Compound literals
  - Designated initializers
  - Mixed declarations and code
  - Macros with a variable number of arguments
  - Inline functions (**inline** keyword)
  - Boolean type (**\_Bool** keyword)
  - **long double** is 64 bit, not 128 bit (only Linux\*)
- Full list of actual C99 support:

<http://software.intel.com/en-us/articles/c99-support-in-intel-c-compiler>

# Excellent Fortran 2008 Support



- Maximum array rank has been raised to 31 dimensions (Fortran 2008 specifies 15)
- Recursive type may have ALLOCATABLE components
- Coarrays
  - CODIMENSION attribute
  - SYNC ALL statement
  - SYNC IMAGES statement
  - SYNC MEMORY statement
  - CRITICAL and END CRITICAL statements
  - LOCK and UNLOCK statements
  - ERROR STOP statement
  - ALLOCATE and DEALLOCATE may specify coarrays
  - Intrinsic procedures IMAGE\_INDEX, LCOBOUND, NUM\_IMAGES, THIS\_IMAGE, UCOBOUND
- CONTIGUOUS attribute
- MOLD keyword in ALLOCATE
- DO CONCURRENT
- NEWUNIT keyword in OPEN

G0 and G0.d format edit descriptor  
Unlimited format item repeat count specifier  
CONTAINS section may be empty  
Intrinsic procedures  
BESSEL\_J0, BESSEL\_J1, BESSEL\_JN,  
BESSEL\_YN, BGE, BGT, BLE, BLT, DSHIFTL,  
DSHIFTR, ERF, ERFC, ERFC\_SCALED, GAMMA,  
HYPOT, IALL, IANY, IPARITY, IS\_CONTIGUOUS,  
LEADZ, LOG\_GAMMA, MASKL, MASKR,  
MERGE\_BITS, NORM2, PARITY, POPCNT,  
POPPAR, SHIFTA, SHIFTL, SHIFTR,  
STORAGE\_SIZE, TRAILZ  
Additions to intrinsic module  
ISO\_FORTRAN\_ENV: ATOMIC\_INT\_KIND,  
ATOMIC\_LOGICAL\_KIND,  
CHARACTER\_KINDS, INTEGER\_KINDS, INT8,  
INT16, INT32, INT64, LOCK\_TYPE,  
LOGICAL\_KINDS, REAL\_KINDS, REAL32,  
REAL64, REAL128, STAT\_LOCKED,  
STAT\_LOCKED\_OTHER\_IMAGE,  
STAT\_UNLOCKED

## Support for Submodules highly requested

- Planned, but not committed yet

Full list of actual Fortran support: <http://software.intel.com/en-us/articles/intel-fortran-compiler-support-for-fortran-language-standards>

Leadership F2008 Support on Linux\*, Windows\* & OSX\*

# Expanded Fortran Capabilities in SP1

- Fortran 2008
  - ATOMIC\_DEFINE and ATOMIC\_REF
  - initialization of polymorphic INTENT(OUT) dummy arguments
  - standard handling of G format and of printing the value zero
  - polymorphic source allocation
- Co-array now supports Intel® Xeon Phi™ Coprocessor

Intel Fortran: Leadership Performance with a Leading Feature Set

# FORTRAN 2003 Support

Added in Intel® Fortran Composer XE 2013 SP1 (aka 14.0, released Q3/2013)

- User-defined derived type I/O
- Parts of the missing intrinsics

FORTRAN 2003 features still missing today:

- Parameterized derived types (soon)
- Transformational intrinsics, such as MERGE and SPREAD, in initialization expression (later)

Complete FORTRAN 2003 support


- Scheduled for next major release (15.0, Q4'2014)

Full list of actual Fortran support: <http://software.intel.com/en-us/articles/intel-fortran-compiler-support-for-fortran-language-standards>

# New Installer Features

- Online installer: A small installation program ~4 MB
  - Select the components to install
  - Components are downloaded and installed on demand from the Intel® Registration and Download Center
  - Requires an internet connection, may require proxy setting. For details, see <http://software.intel.com/en-us/articles/how-to-set-system-proxy>
  - Internet connection speed will impact installation time
- Offline: Full packages of Composer XE and Libraries available for offline installations
- New GUI installation for Linux\* available, --gui-mode

# Optional GUI installer on Linux\*



# Intel® XEON Phi Support

Intel® Many Integrated Core (Intel® MIC) support added since 2013 release ( compiler version 13.0 )

- Only Intel® Xeon Phi processor family (code name Knights Corner / KNC ), not KNF anymore
- One product only ! (no additional license required)
- Needs MPSS stack installed on host system to compile (not only for linking)
  - Very likely to be changed in future release

# Enhancement of Compiler Reporting

Intel Compilers offers two switches to get compiler reports:

**-vec-report** for vectorization

**-opt-report** for loop transformations, inlining, prefetching, ...

An initiative has started to enhance both reports. It will take some time to implement all enhancements requests we received and which we consider doable.

Some items improved or to be improved in future releases:

- More details on vectorization success or failure like
  - Unrolling
  - Remainder loop
  - Alignment
- Automated mapping of information to source code
- Optimization report will be enhanced too / will be made easier to read

# Vectorization Report

Provides details on vectorization success & failure:

Linux\*, Mac OS\* X: **-vec-report<n>**, Windows\*: **/Qvec-report<n>**

## **n Diagnostic Messages**

- 0 Tells the vectorizer to report no diagnostic information. Useful for turning off reporting in case it was enabled on command line earlier.**
- 1 Tells the vectorizer to report on vectorized loops. [default if **n** missing]**
- 2 Tells the vectorizer to report on vectorized and non-vectorized loops.**
- 3 Tells the vectorizer to report on vectorized and non-vectorized loops and any proven or assumed data dependences.**
- 4 Tells the vectorizer to report on non-vectorized loops.**
- 5 Tells the vectorizer to report on non-vectorized loops and the reason why they were not vectorized.**
- 6 Tells the vectorizer to emit greater detail when reporting on vectorized and non-vectorized loops and any proven or assumed data dependences: Introduced by compiler 13.0**
- 7 Very sophisticated vectorization report including support of source code annotation using Python helper scripts: Introduced by compiler 13.1**

# Level 6 Vectorization Report Sample

```
6: forall( i= 1:n)
7:     a(i)= a(i)-b(i)*d(i)
8:     c(i)= a(i)+c(i)
9: endforall
10:
11: do i=1,n
12:     a(i)= a(i)-b(i)*d(i)
13:     c(i)= a(i)+c(i)
14: enddo
```

## -vec-report3

```
12.f90(7): (col. 5) remark: LOOP WAS VECTORIZED.
12.f90(8): (col. 5) remark: LOOP WAS VECTORIZED.
12.f90(11): (col. 3) remark: LOOP WAS VECTORIZED.
```

## -vec-report6

```
12.f90(7): (col. 5) remark: vectorization support: unroll factor set to 2.
12.f90(7): (col. 5) remark: LOOP WAS VECTORIZED.
12.f90(8): (col. 5) remark: vectorization support: unroll factor set to 2.
12.f90(8): (col. 5) remark: LOOP WAS VECTORIZED.
12.f90(11): (col. 3) remark: vectorization support: unroll factor set to 2.
12.f90(11): (col. 3) remark: LOOP WAS VECTORIZED.
```

# Level 7 Vectorization Reporting

Provides very detailed and powerful vectorization reporting including summary statistic and source code annotation

- Details and required Python scripts available [here](#)
- Not documented in compiler manual yet !
- Python release must be 2.6.5 or younger

Sample usage:

```
ifort -c -vec-report7 loop.f90 2>&1 |  
./vecanalysis/vecanalysis.py -list
```

This results in statistics displayed and an annotated source file **loop\_ver.f90** to be created

# Some Intel-specific Extensions

- Available today / Composer XE 2013 SP1
  - **Btrace**: Get stack trace even in case stack corrupted
  - **PDBX**: Parallel Debug Extensions - Data Race Detection
  - TSX: Intel® Transactional Memory Extension support
  - Pointer Checking integration
- Being worked on for future release
  - Full (Intel-) FORTRAN support as provided by Intel® IDB
  - Full VLA (variable length array) debugging
    - In collaboration with community (“GBD Archer project” ) and Redhat
    - Needed for FORTRAN but also for C99

**Note:** No own, specific debugger GUI as for Intel® IDB

- But Eclipse-based GUI integration

# New Instruction Support

3<sup>rd</sup> Generation Intel® Core™ architecture code name Ivy Bridge:

- Extension **core-avx-i** for `-x` switches
- Enables support for all intrinsics introduced by Ivy Bridge (some 6) including RDRAND and 16 bit floating point instructions
- No implicit code generation for new instructions

4<sup>th</sup> Generation Intel® Core™ architecture code name Haswell:

- Extension **core-avx2** for `-x` switches
- Enables intrinsic support (like for Intel® TSX) and implicit code generation
- FMA instruction will be used by default but can be disabled using switch **-no-fma**
  - -valuable for FP consistence / rounding issues

# New Instruction Support : Broadwell

RDSEED	provide reliable seeds for pseudo-random number generator
ADCX, ADOX	large integer arithmetic addition
PREFETCHW	extending SW prefetch

- Supported via intrinsics :
  - Intel Compilers 13.0 Update 2
  - GNU GCC 4.8

# Manual Code Dispatch: A Note

- Intel ® TSX is an optional feature for Haswell
  - Not all models will have this feature
  - Thus two CPU “guard” names are/will be available
    - `core_4th_gen_avx`
    - `core_4th_gen_avx_tsx`

```
__declspec(cpu_specific(core_4th_gen_avx))
void dispatch_func()
{
    printf("\Code here can assume HSW NI but not TSX\n");
}

__declspec(cpu_specific(core_4th_gen_avx_tsx))
void dispatch_func()
{
    printf("\Code here can assume TSX to be available\n");
}
```

# Intel® Visual Fortran Composer XE 2013 with IMSL\* for Windows\*

- Same 13.1 compiler that is available separately
  - Highly optimizing Fortran compiler featuring scalable multi-threading with OpenMP, and introducing coarray Fortran (part of the Fortran 2008 standard), including new parallelism models
  - VAX Fortran and Compaq Visual Fortran compatibility
  - In addition to IMSL, it includes Intel® Math Kernel Library (Intel® MKL)
- Rogue Wave IMSL\* 6.0 Math Library highlights
  - Over 1,000 mathematical and statistical algorithms for developers of Fortran applications
  - Support for shared memory and distributed memory computing environments
  - High performance linear programming optimizer
  - ScaLAPACK integration for MPI, LAPACK integration for SMP
  - New probability density functions and inverses
  - Time series and forecasting additions
  - New Deployment Licensing terms and pricing

# Intel® Composer XE for Linux\*

Release 4<sup>th</sup> Sept 2013

Intel® Composer XE 2013 SP1 for Linux* (32/64 bit)	
• Intel® C++ Composer XE Intel® C++ Compiler XE for applications (icc / icpc / idb / gdb-ia)	2013 SP1 14.0.0.080 Build 20130728
• Intel® Fortran Composer XE Intel® C++ Compiler XE for applications (ifort / idbc)	2013 SP1 14.0.0.080 Build 20130728
• Intel® Integrated Performance Primitives (Intel® IPP)	8.0 Update 1
• Intel® Math Kernel Library (Intel® MKL)	11.1
• Intel® Threading Building Blocks (Intel® TBB)	4.1 Update 3
• MPSS package for linking to Intel® Xeon Phi™	2013 Update 3.3

- The 13.0 compilers (Intel® C++ and Fortran Compiler XE 13.0) introduced a large list of new features. All following updates improve stability & compatibility.
- Latest “minor” version is 14.0 (released 4<sup>th</sup> Sept. 2013)
- Note: Intel® Debugger is deprecated with the next major version.
  - But still needed for full Fortran debugging support

# Intel® Composer XE for Windows\* / OS X\*

Intel® Composer XE 2013 SP1 for Windows* (32/64 bit)	
<ul style="list-style-type: none"> <li>Intel® C++ Composer XE Intel® C++ Compiler XE for applications (icl)</li> </ul>	<b>2013 SP1</b> <b>14.0.0.103 Build 20130728</b>
<ul style="list-style-type: none"> <li>Intel® Visual Fortran Composer XE (incl. Microsoft Visual Studio Shell and Libraries*) Intel® Visual Fortran Compiler XE for applications (ifort)</li> </ul>	<b>2013 SP1</b> <b>14.0.0.013 Build 20130728</b>
<ul style="list-style-type: none"> <li>Intel® Integrated Performance Primitives (Intel® IPP)</li> </ul>	<b>8.0 Update 1</b>
<ul style="list-style-type: none"> <li>Intel® Math Kernel Library (Intel® MKL)</li> </ul>	<b>11.1</b>
<ul style="list-style-type: none"> <li>Intel® Threading Building Blocks (Intel® TBB)</li> </ul>	<b>4.2</b>
<ul style="list-style-type: none"> <li>MPSS package for linking to Intel® Xeon Phi™</li> </ul>	<b>2013 Update 3.1</b>

Intel® Composer XE 2013 SP1 for OS X* (32/64 bit)	
<ul style="list-style-type: none"> <li>Intel® C++ Composer XE (ICC/ICPC/IDB) Intel® C++ Compiler XE for applications (icc / icpc / idb)</li> </ul>	<b>2013 SP1</b> <b>14.0.0.074 Build 20130728</b>
<ul style="list-style-type: none"> <li>Intel® Fortran Composer XE (IFORT/IDB) Intel® C++ Compiler XE for applications (ifort / idb)</li> </ul>	<b>2013 SP1</b> <b>14.0.0.074 Build 20130728</b>

- OS X\* versions do not contain Intel® Integrated Performance Primitives!
- Development decoupled from Linux\*/Windows\* products (see change in version)

# Key Files Supplied with Compilers

## Windows\*

- Intel compiler
  - `icl.exe`, `ifort.exe`: C/C++ compiler, Fortran compiler drivers
  - `mcpcom.exe`, `fortcom.exe`: C/C++ or Fortran Compiler
  - `icl.cfg`, `ifort.cfg`: Default compiler options
  - `compilervars.bat`: Setup command window build environment (C/C++ and Fortran)
- Linker driver
  - `xilink.exe`: Invokes `link.exe`
- Intel include files, libraries
- Re-distributable files
  - `<install-dir>\ComposerXE-2011\redist\`



## Linux\*, Mac OS\* X

- Intel compiler
  - `icc`, `ifort`: C/C++ compiler, Fortran compiler
  - `mcpcom`, `fortcom`: C/C++ or Fortran Compiler
  - `compilervars.(c)sh`: Source scripts to setup the complete compiler/debugger/libraries environment (C/C++ and Fortran)
- Linker driver
  - `xild`: Invokes `ld`
- Intel include files, libraries
- Intel debugger
  - `idbc` (Command Line) Debugger, `idb` (GUI) Debugger (Linux\* only)
- GDB\* Enhancements
  - `gdb-ia` with `btrace` support and data race detection


# Command Line Build Environment - Windows

- Sets the environment variables for command line builds – common for C/C++ and Fortran

- Run C/C++/Fortran compilers from command line



# Visual Studio\* 2010/2012 Integration



- Switch between compilers (Intel / Visual C++) – **fully integrated**
- Additional Properties options setting when Intel Compiler is set

# Command Line Build Environment - Linux

- Sets the environment variables for command line builds – common for C/C++ and Fortran

```
bash-4.2$ source /opt/intel/composer_xe_2013_spl/bin/compilervars.sh intel64
bash-4.2$ icc -v
icc version 14.0.0 (gcc version 4.6.0 compatibility)
bash-4.2$ icc -V
Intel(R) C Intel(R) 64 Compiler XE for applications running on Intel(R) 64, Version 14.0.0.080 Build 20130728
Copyright (C) 1985-2013 Intel Corporation. All rights reserved.

bash-4.2$ ifort -v
ifort version 14.0.0
bash-4.2$ idbc
Intel(R) Debugger for applications running on Intel(R) 64, Version 13.0, Build [80.483.23]
(idb) quit
bash-4.2$ gdb-ia
No symbol table is loaded. Use the "file" command.
GNU gdb (GDB) 7.5-1.0.160
Copyright (C) 2012 Free Software Foundation, Inc; (C) 2013 Intel Corp.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law. Type "show copying"
and "show warranty" for details.
This GDB was configured as "x86_64-unknown-linux-gnu".
For information about how to find Technical Support, Product Updates,
User Forums, FAQs, tips and tricks, and other support information, please visit:
<http://www.intel.com/software/products/support/>.
(gdb) quit
bash-4.2$ gdb
GNU gdb (Ubuntu/Linaro 7.4-2012.04-0ubuntu2.1) 7.4-2012.04
Copyright (C) 2012 Free Software Foundation, Inc.
```

# Common Optimization Switches

	<b>Windows*</b>	<b>Linux*</b>
Disable optimization	/Od	-O0
Optimize for speed (no code size increase)	/O1	-O1
Optimize for speed (default)	/O2	-O2
High-level loop optimization	/O3	-O3
Create symbols for debugging	/Zi	-g
Multi-file inter-procedural optimization	/Qipo	-ipo
Profile guided optimization (multi-step build)	/Qprof-gen /Qprof-use	-prof-gen -prof-use
Optimize for speed across the entire program	/fast (same as: /O3 /Qipo /Qprec-div- /QxHost)	-fast (same as: -ipo -O3 -no- prec-div -static -xHost)
OpenMP 3.0 support	/Qopenmp	-openmp
Automatic parallelization	/Qparallel	-parallel

# Compiler Reports – Optimization Report

Compiler switch:

`/Qopt-report-phase[:phase]` (Windows\*)

`-opt-report-phase[=phase]` (Linux\*, Mac OS\* X)

`,phase`` can be:

- `ipo_inl` - Interprocedural Optimization Inlining Report
- `ilo` - Intermediate Language Scalar Optimization
- `hpo` - High Performance Optimization
- `hlo` - High-level Optimization
- `all` - All optimizations (not recommended, output too verbose)

Control the level of detail in the report:

`/Qopt-report[0|1|2|3]` (Windows\*)

`-opt-report[0|1|2|3]` (Linux\*, MacOS\* X)

If you do not specify the option, no optimization report is being generated; if you do not specify the level (i.e. `/Qopt-report`, `-opt-report`) level 2 is being used by the compiler.

# Optimization Report Example

```
icc -O3 -opt-report-phase=hlo -opt-report-phase=hpo  
icl /O3 /Qopt-report-phase:hlo /Qopt-report-phase:hpo
```

```
...  
LOOP INTERCHANGE in loops at line: 7 8 9  
Loopnest permutation ( 1 2 3 ) --> ( 2 3 1 )  
...  
Loop at line 8 blocked by 128  
Loop at line 9 blocked by 128  
Loop at line 10 blocked by 128  
...  
Loop at line 10 unrolled and jammed by 4  
Loop at line 8 unrolled and jammed by 4  
...  
... (10)... loop was not vectorized: not inner loop.  
... (8)... loop was not vectorized: not inner loop.  
... (9)... PERMUTED LOOP WAS VECTORIZED  
...
```

icc -vec-report2 (icl /Qvec-report2) for just the vectorization report

# Case study: Matrix calculation

```
#include <stdio.h>
#include <time.h>

#define NUM 1024
main()
{
    clock_t start, stop;
    int num;

    // initialize the arrays with data
    init_arr(3,-2,1,a);
    init_arr(-2,1,3,b);

    //start timing the matrix multiply code
    printf("NUM:%d\n",NUM);
    start = clock();
    multiply_d(a,b,c);
    stop = clock();

    // print simple test case of data to be sure multiplication is correct
    if (NUM < 5) {
        print_arr("a", a);
        print_arr("b", b);
        print_arr("c", c);
    }

    // print elapsed time
    printf("Elapsed time = %lf seconds\n",(((double)(stop - start)) / CLOCKS_PER_SEC));
}
```

# Case study: Matrix calculation

```
//routine to initialize an array with data
void init_arr(double row, double col, double off, double a[][NUM])
{
    int i,j;

    for (i=0; i< NUM;i++) {
        for (j=0; j<NUM;j++) {
            a[i][j] = row*i+col*j+off;
        }
    }
}
```

```
// routine to print out contents of small arrays
void print_arr(char * name, double array[][NUM])
{
    int i,j;

    printf("\n%s\n", name);
    for (i=0;i<NUM;i++){
        for (j=0;j<NUM;j++) {
            printf("%g\t",array[i][j]);
        }
        printf("\n");
    }
}
```

# Case study: Matrix calculation

```
static double a[NUM][NUM], b[NUM][NUM], c[NUM][NUM];  
void multiply_d(double a[][NUM], double b[][NUM], double c[][NUM]);
```

```
void multiply_d(double a[][NUM], double b[][NUM], double c[][NUM])  
{  
    int i,j,k;  
    double temp;  
    for(i=0;i<NUM;i++) {  
        for(j=0;j<NUM;j++) {  
            for(k=0;k<NUM;k++) {  
                c[i][j] = c[i][j] + a[i][k] * b[k][j];  
            }  
        }  
    }  
}
```

# Case study: Matrix calculation

- Compile the program matrix.c/matrix.f90 without optimization and with /O1 – compare results (e.g.):

```
icl /Od matrix.c /o noopt_matrix.exe  
icl /O1 matrix.c /o O1_matrix.exe
```

- Run and time the executable

```
noopt_matrix.exe      >>      30,25 sec.  
O1_matrix.exe        >>       9,6 sec.
```

[Intel® Core™ i7, 4-Core 64-bit, 1,6GHz]

# High-Level Optimizer (HLO)

- Compiler switches:  
  /**o2**, /**o3** (Windows\*), -**o2**, -**o3** (Linux\*)
- Loop level optimizations
  - loop unrolling, cache blocking, prefetching
- More aggressive dependency analysis
  - Determines whether or not it's safe to reorder or parallelize statements
- Scalar replacement
  - Goal is to reduce memory by replacing with register references

# Interprocedural Optimizations (IPO)

## Multi-pass Optimization

- Interprocedural optimizations performs a static, topological analysis of your application!
- ip: Enables inter-procedural optimizations for current source file compilation
- ipo: Enables inter-procedural optimizations across files
  - Can inline functions in separate files
  - Especially many small utility functions benefit from IPO


Windows*	Linux*
/Qip	-ip
/Qipo	-ipo

Enabled optimizations:

- Procedure inlining (reduced function call overhead)
- Interprocedural dead code elimination, constant propagation and procedure reordering
- Enhances optimization when used in combination with other compiler features

# Interprocedural Optimizations (IPO)

## Usage: Two-Step Process



# Inlining Functions

When the compiler inlines a function call, the function's code gets inserted into the caller's instruction stream

## Benefits:

- Reducing overhead of calling a function
  - writing the registers and parameters to/from stack
  - restore the registers when the function returns.
- Improving performance because the optimizer can procedurally integrate the called function and can do better optimizations
  - sub-expression elimination
  - copy propagation

## Drawbacks:

- Overuse of inlining can actually make programs slower. Depending on a function's size, inlining it can cause the code size to increase, resulting in more cache misses and more pressure on the instruction cache
- The speed benefits of inline functions tend to diminish as the function grows in size. At some point the overhead of the function call becomes small compared to the execution of the function body, and the benefit is lost.

# Techniques for Inlining Functions

- Compiler Switches

- Increase information provided to the compiler
  - `-ipo, -prof_use` (Linux), `/Qipo, /Qprof-use` (Windows)
- Change Compiler Heuristics
  - `-inline-factor=n` (default=100), `/Qinline-factor=n`
  - `-inline-level=0|1|2, /ob0|1|2`

- Inlining source code features

- Microsoft\* C/C++
  - Keywords: `inline, __inline, __forceinline`
- GCC C/C++
  - `__attribute__((always_inline))`
  - `__attribute__((noinline))`

# Case study: Include IPO

- Compile the program `matrix.c/matrix.f90` with optimization and with including `/Qipo` – compare results:

```
icl /O3 /Qipo matrix.c /Qopt-report-phase:ipo /o O3_ipo_matrix.exe
icl /O3 /Qipo /QxHost matrix.c /o O3_ipo_xHost_matrix.exe
[Don't expect much with IPO on this sample (why?)]
```

- Run and time the executable

```
O2_novec_matrix.exe    >>    2,12 sec.
O3_novec_matrix.exe    >>    1,6 sec.
O2_matrix.exe          >>    0,98 sec.
O3_matrix.exe          >>    0,83 sec.
O3_ipo_matrix.exe     >>    0,82 sec.
O3_ipo_xHost_matrix.exe >>    0,75 sec.
```

[Intel® Core™ i7,  
4-Core 64-bit,  
1,6GHz]

# Profile-Guided Optimizations (PGO)


- Static analysis leaves many questions open for the optimizer like:

- How often is  $x > y$
- What is the size of count
- Which code is touched how often

```
if (x > y)
    do_this();
else
    do_that();
```

```
for(i=0; i<count; ++i)
    do_work();
```

- Use execution-time feedback to guide (final) optimization
- Enhancements with PGO:
  - More accurate branch prediction
  - Basic block movement to improve instruction cache behavior
  - Better decision of functions to inline (help IPO)
  - Can optimize function ordering
  - Switch-statement optimization
  - Better vectorization decisions



# PGO Usage: Three Step Process

## Step 1

Compile + link to add instrumentation

```
icc -prof_gen prog.c
```

Instrumented executable:  
prog.exe

## Step 2

Execute instrumented program  
prog.exe (on a typical dataset)

Dynamic profile:  
12345678.dyn

## Step 3

Compile + link using feedback  
icc -prof\_use prog.c

Merged .dyn files:  
pgopti.dpi

Optimized executable:  
prog.exe

# Case study: Use PGO

- Compile the program matrix.c/matrix.f90 with /Qprof-gen /Qprof-use – compare results:

```
icl /Qprof-gen matrix.c /o pgen_matrix.exe
```

```
pgen_matrix.exe      [Why does this take so long?]
```

```
icl /Qprof-use /O3 /QxHost matrix.c /o puse_matrix.exe
```

- Run and time the executable

```
O2_novec_matrix.exe    >>      2,12 sec.
```

```
O3_novec_matrix.exe    >>      1,6 sec.
```

```
O2_matrix.exe         >>      0,98 sec.
```

```
O3_matrix.exe         >>      0,83 sec.
```

```
O3_ipo_matrix.exe     >>      0,82 sec.
```

```
O3_ipo_xHost_matrix.exe >>      0,75 sec.
```

```
puse_matrix.exe      >>      0,32 sec.
```

[Intel® Core™ i7,  
4-Core 64-bit,  
1,6GHz]

# Automatic Compiler Vectorization

## Processor Specific Optimizations

- Compiler vectorizer switch:  
`/Qvec` (Windows\*), `-vec` (Linux\*)
  - Implicitly included with optimization switch `/O2`, `-O2` or higher
  - Switch off autovectorizer explicitly with `/Qvec-`, `-vec-`
  - Automatically generates vector instructions using the multiple SIMD instruction set extensions SSE/SSE2/SSE3/SSSE3/SSE4.x/AVX
  - Operates at once with one instruction on, e.g.:  
4 float / 2 double values or 4 x 32-bit / 8 x 16-bit integers
- Processor specific extensions switches:  
Default switch is `-msse2` (Windows), `/arch:SSE2` (Linux)
  - Activated implicitly with optimization switch `/O2`, `-O2` or higher`/arch:<extension>` (Microsoft compatible), `-m<extension>` (Linux\*, Mac OS\* X)
  - No Intel processor check (but with sse4.x)
  - No Intel specific optimizations`/Qx<SIMDextension>`, `-x<SIMDextension>`
  - Code generation for specific target hardware
  - Enables additional extensions for Intel processors`/Qax<SIMDextension>`, `-ax<SIMDextension>`
  - Code generation for specific target hardware and extensions for Intel processors
  - Autodispatch switch `a` for generating additional default code pathSpecial switch `/QxHost`, `-xHost`
  - Automatically gives access to all the latest features of the processor you are working on
- `-mtune=<ARCH>` option on Linux\*/OS X\* to specify cpu targeting without generating instructions exclusive to that cpu

# Processor-specific Compiler Switches

(Intel® 64 and IA-32)


Windows	Linux
<b>/Qxsse2</b>	<b>-xsse2</b>
<b>/Qxsse3</b>	<b>-xsse3</b>
<b>/Qxssse3</b>	<b>-xssse3</b>
<b>/Qxsse4.1</b>	<b>-xsse4.1</b>
<b>/Qxsse4.2</b>	<b>-xsse4.2</b>
<b>/QxAVX</b>	<b>-xavx</b>
<b>/QxHOST</b>	<b>-xHost</b>
<b>/Qxsse3_atom</b>	<b>-xsse3_ATOM</b>

- Implies an Intel CPU ID check
- Runtime message if try to run on unsupported processor

Windows	Linux
<b>/arch:IA32</b>	<b>-mia32</b>
<b>/arch:sse2</b>	<b>-msse2 (default!)</b>
<b>/arch:sse3</b>	<b>-msse3</b>
<b>/arch:ssse3</b>	<b>-mssse3</b>

- No CPU ID check (but with sse4.x)
- Intel and non-Intel processors
- Illegal instruction error if run on unsupported processor
- At least Pentium4 required

# SIMD Types in Processors from Intel [1]




## MMX™

Vector size: 64bit

Data types: 8, 16 and 32 bit integers

VL: 2,4,8

For sample on the left:  $X_i$ ,  $Y_i$  16 bit integers



## Intel® SSE

Vector size: 128bit

Data types:


8,16,32,64 bit integers

32 and 64bit floats

VL: 2,4,8,16

Sample:  $X_i$ ,  $Y_i$  bit 32 int / float

# SIMD Types in Processors from Intel [2]




## Intel® AVX

Vector size: 256bit

Data types: 32 and 64 bit floats

VL: 4, 8, 16

Sample:  $X_i, Y_i$  32 bit int or float



## Intel® MIC

Vector size: 512bit

Data types:

32 and 64 bit integers

32 and 64bit floats

(some support for 16 bits floats)

VL: 8,16

Sample: 32 bit float

# Auto-Vectorization

## SIMD – Single Instruction Multiple Data

- Scalar mode

- one instruction produces one result

- SIMD processing

- with SSE or AVX instructions
- one instruction can produce multiple results

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




# Scalar and Packed SSE Instructions

The “vector” form of SSE instructions operating on multiple data elements simultaneously are called packed – thus vectorized SSE code means use of packed instructions

- Most of these instructions have a scalar version too operating only one element only


## `addss` Scalar Single-FP Add

↑↑ single precision FP data  
↑ scalar execution mode



## `addps` Packed Single-FP Add

↑↑ single precision FP data  
↑ packed execution mode



# References

[1] Aart Bik: "The Software Vectorization Handbook"



– [http://www.intel.com/intelpress/sum\\_vmmx.htm](http://www.intel.com/intelpress/sum_vmmx.htm)

[2] Randy Allen, Ken Kennedy: "Optimizing Compilers for Modern Architectures: A Dependence-based Approach"

[3] Steven S. Muchnik, "Advanced Compiler Design and Implementation"

[4] Intel Software Forums, Knowledge Base, White Papers, Tools Support - see <http://software.intel.com>  
Sample Articles :

- [software.intel.com/en-us/articles/a-guide-to-auto-vectorization-with-intel-c-compilers/](http://software.intel.com/en-us/articles/a-guide-to-auto-vectorization-with-intel-c-compilers/)
- [software.intel.com/en-us/articles/requirements-for-vectorizable-loops/](http://software.intel.com/en-us/articles/requirements-for-vectorizable-loops/)
- [software.intel.com/en-us/articles/performance-tools-for-software-developers-intel-compiler-options-for-sse-generation-and-processor-specific-optimizations/](http://software.intel.com/en-us/articles/performance-tools-for-software-developers-intel-compiler-options-for-sse-generation-and-processor-specific-optimizations/)



# Compiler Reports – Vectorization Report

Compiler switch:

`/Qvec-report<n>` (Windows)

`-vec-report<n>` (Linux)

Set diagnostic level dumped to stdout

n=0: No diagnostic information

n=1: (Default) Loops successfully vectorized

n=2: Loops not vectorized – and the reason why not

n=3: Adds dependency Information

n=4: Reports only non-vectorized loops

n=5: Reports only non-vectorized loops and adds dependency info

- Note:

For Linux, `-opt_report_phase hpo` provides additional diagnostic information for vectorization

# Case study: Use auto-vectorizer

- Compile the program matrix.c/matrix.f90 with optimization /O[2|3], vectorizer and create report:

```
icl /O2 matrix.c /Qvec-report2 /o O2_matrix.exe
```

```
icl /O3 matrix.c /Qvec-report2 /o O3_matrix.exe
```

```
icl /O3 /QxHost /Qvec-report2 matrix.c /o O3_xHost_matrix.exe
```

```
icl /fast matrix.c /o fast_matrix.exe
```

*[Do all loops vectorize? Is there room for improvement?]*

- Run and time the executable

```
noopt_matrix.exe >> 30,25 sec.
```

```
O1_matrix.exe >> 9,6 sec.
```

```
O2_novec_matrix.exe >> 2,12 sec.
```

```
O3_novec_matrix.exe >> 1,6 sec.
```

```
O2_matrix.exe >> 0,98 sec.
```

```
O3_matrix.exe >> 0,83 sec.
```

```
O3_xHost_matrix.exe >> 0,78 sec.
```

```
fast_matrix.exe >> 0,78 sec.
```

[Intel® Core™ i7,  
4-Core 64-bit,  
1,6GHz]

# Compiling for Intel® AVX


- Compile with `-xavx` (`/Qxavx` on Windows\*)
  - Vectorization works just as for SSE
- More FP loops can be vectorized than with SSE
  - Individually masked data elements
  - More powerful data rearrangement instructions
- Can create both SSE and AVX code paths in one binary
- Switches Linux: `-xavx`      Windows: `/Qxavx`
  - use additional `-x` (`/Qx`) switches to modify the default SSE code path; e.g `-xavx -xsse4.2` to target SSE4.2 (“Nehalem”) and AVX (“SandyBridge”)

# User-Mandated Vectorization

User-mandated vectorization is based on a new **SIMD Directive** (or “pragma”)


- The SIMD directive provides additional information to compiler to enable vectorization of loops ( at this time only inner loop )
- Supplements automatic vectorization but differently to what traditional directives like IVDEP, VECTOR ALWAYS do, the SIMD directive is more a command than a hint or an assertion: The compiler heuristics are completely overwritten as long as a clear logical fault is not being introduced

Relationship similar to OpenMP versus automatic parallelization:



# Positioning of SIMD Vectorization

## A Cilk Plus Feature



# SIMD Directive Notation

**C/C++:**        **#pragma simd [clause [,clause] ...]**

**Fortran:**        **!DIR\$ SIMD [clause [,clause] ...]**

Without any clause, the directive enforces vectorization of the (innermost) loop

Sample:

```
void add_fl(float *a, float *b, float *c, float *d, float *e, int n)
{
    #pragma simd
    for (int i=0; i<n; i++)
        a[i] = a[i] + b[i] + c[i] + d[i] + e[i];
}
```

Without SIMD directive, vectorization will fail since there are too many pointer references to do a run-time check for overlapping (compiler heuristic)

# Clauses of SIMD Directive

## **vectorlength**(n1 [,n2] ...)

- n1, n2, ... must be 2,4,8 or 16: The compiler can assume a vectorization for a vector length of n1, n2, ... to be save

## **private**(v1, v2, ...)

- variables private to each iteration; initial value is broadcast to all private instances, and the last value is copied out from the last iteration instance.

## **linear**(v1:step1, v2:step2, ...)

- for every iteration of original scalar loop, v1 is incremented by step1, ... etc. Therefore it is incremented by  $\text{step1} * (\text{vector length})$  for the vectorized loop.

## **reduction**(operator:v1, v2, ...)

- v1 etc are reduction variables for operation "operator"

## **[no]assert**

- reaction in case vectorization fails: Print a warning only ( noassert, the default) or treat failure as error and stop compilation

# Sample: SIMD Directive Vectorlength Clause

```
Void foo(float *a, float *b, float *c, int n)
{
    for (int k=0; k<n; k++)    c[k] = a[k] + b[k];
}
```

Due to the overlapping nature of array accesses from the different call sites, it might not be semantically correct to use restrict keyword or IVDEP directive ( there are dependencies between iterations for one call )

But it might be true for all calls, that e.g 4 consecutive iterations can be executed in parallel without violating any dependencies

```
void foo(float *a, float *b, float *c, int n)
{
    #pragma simd vectorlength(4)
    for (int k=0; k<n; k++ )    c[k] = a[k] + b[k];
}
```

# Auto-Parallelization

- Compiler automatically translates portions of serial code into equivalent multithreaded code with using these options:  
`/Qparallel, -parallel`
- The auto-parallelizer analyzes the dataflow of loops and generates multithreaded code for those loops which can safely and efficiently be executed in parallel.
- The auto-parallelizer report can provide information about program sections that could be parallelized by the compiler. Compiler switch:  
`/Qpar-report:0|1|2|3`  
`-par-report0|1|2|3`  
    ,0` is report disabled, ,3` maximum diagnostics level

# Case study: Auto-parallelization

- Compile the program `matrix.c/matrix.f90` with `/Qparallel /Qpar-report2` with and without additional vectorization/optimizations:

```
icl /Qparallel /Qpar-report2 /Qvec- /o par_novec_matrix.exe
```

```
icl /Qparallel /Qpar-report2 /o par_matrix.exe
```

```
icl /Qparallel /Qpar-report2 /O3 /Qvec-report2 /Qipo /QxHost /o  
par_vec_matrix.exe
```

```
icl /O3 /Qparallel /Qprof-use /O3 /Qipo /QxHost /o  
par_vec_puse_matrix.exe
```

*[Examine the parallel report. Is the program well parallelized? What are the possible pitfalls with thread-level parallelization?]*

- Run and time the executable

```
O3_ipo_xHost_matrix.exe    >>    0,75 sec.  
puse_matrix.exe           >>    0,32 sec.  
par_novec_matrix.exe      >>    0,90 sec.  
par_matrix.exe            >>    0,35 sec.  
par_vec_matrix.exe        >>    0,35 sec.  
par_vec_puse_matrix.exe   >>    0,12 sec.
```

[Intel® Core™ i7,  
4-Core 64-bit,  
1,6GHz]

# Auto-Parallelization vs. Auto-Vectorization

Explain difference

- **Auto-vectorization** which is DLP (Data Level Parallelization) uses wide registers (SSE) for operation of multiple instructions.
- **Auto-parallelization** which is TLP (Thread Level Parallelization) uses processor cores (and hardware threads) for parallelization of serial code.
- Possible pitfalls in auto-parallelization: Concurrent access by more threads to the same memory locations (data race!).

# Guided Automatic Parallelization (GAP)

- Feature of the C++/Fortran Compiler that offers advice and, when correctly applied, results in auto-vectorization or auto-parallelization of serial code.
  - Compiler switch `/Qguide` in parallel with `/O2` or higher gives advise for auto-vectorization
  - Compiler switch `/Qguide` in parallel with `/Qparallel` or higher gives advise for auto-parallelization
- Compiler runs in advisor mode
  - Compiler does not generate ANY code with `/Qguide`, `-guide` – vectorized or not.

# Intel® Guided Auto Parallelism (GAP)

## Let the Compiler Tell You What it Needs

- Motivation

- Effective, simplify way to add parallelism to your applications
- Use built-in compiler technology to speed parallelism development

- What is GAP?

- Compiler-based analyzer that provides guidance to developers to change code so it can be compiled to automatically optimize code through vectorization, parallelization, or data transformation
- Built upon existing auto-vectorization and auto-parallelization technology

- GAP does not

- Analyze code and find hotspots for threading (see Advisor)
- Verify threading correctness (use Inspector, Inspector XE)
- Do any performance/hotspot analysis (use Amplifier, VTune Amplifier XE)

**Developer Must Verify Semantics of GAP Recommendations**

# GAP – How it Works (Windows)


- Windows – right-click or Project Properties
  - GAP analysis appears in Output window



Right-click file or project,  
Fortran Composer pop-up

Choose single-file  
or whole project

OR you can use the  
Project Properties,  
Fortran, Diagnostics  
property page



# Vectorization Example [1]

```
void f(int n, float *x, float *y, float *z, float *d1, float *d2)
{
    for (int i = 0; i < n; i++)
        z[i] = x[i] + y[i] - (d1[i]*d2[i]);
}
```

GAP Message:

g.c(6): remark #30536: (LOOP) Add -Qno-alias-args option for better type-based disambiguation analysis by the compiler, if appropriate (the option will apply for the entire compilation). This will improve optimizations such as vectorization for the loop at line 6. [VERIFY] Make sure that the semantics of this option is obeyed for the entire compilation. [ALTERNATIVE] Another way to get the same effect is to add the "restrict" keyword to each pointer-typed formal parameter of the routine "f". This allows optimizations such as vectorization to be applied to the loop at line 6. [VERIFY] Make sure that semantics of the "restrict" pointer qualifier is satisfied: in the routine, all data accessed through the pointer must not be accessed through any other

***The compiler guides the user on source-change and on what pragma to insert and on how to determine whether that pragma is correct for this case***

# Vectorization Example [2]

```
void mul(NetEnv* ne, Vector* rslt
  Vector* den, Vector* flux1,
  Vector* flux2, Vector* num
{
  float *r, *d, *n, *s1, *s2;
  int i;
  r=rslt->data;
  d=den->data;
  n=num->data;
  s1=flux1->data;
  s2=flux2->data;

  for (i = 0; i < ne->len; ++i)
    r[i] = s1[i]*s2[i] +
    n[i]*d[i];
}
```

GAP Messages (simplified):

1. "Use a local variable to store the upper-bound of loop at line 29 (`variable:ne->len`) if the upper-bound does not change during execution of the loop"
  2. "Use "#pragma ivdep" to help vectorize the loop at line 29, if these arrays in the loop do not have cross-iteration dependencies: `r, s1, s2, n, d`"
- > Upon recompilation, the loop will be vectorized

# Parallelization Example

```
#define N 10000
double A[N], B[N];
int bar(int);
void foo(){
    int i;
    for (i=0;i<N;i++)
        A[i] = B[i] * bar(i);
}
```

Compiler invocation:

```
{L}: icc -c funcall.c
      -parallel
      -guide
{W}: icl /c funcall.c
      /Qparallel
      /Qguide
```

## GAP Message on Windows:

funcall.c(6): remark #30528: (PAR) Add "\_\_declspec(const)" to the declaration of routine "bar" in order to parallelize the loop at line 12. Alternatively, adding "\_\_declspec(concurrency\_safe(profitable))" achieves a similar effect.

[VERIFY] Make sure that the routine satisfies the semantics of this declaration.

[ALTERNATIVE] Yet another way to help the loop being parallelized is to inline the routine with "#pragma forceinline recursive". This method does not guarantee parallelization.

# Data Transformation Example

```
struct S3 {
    int a;
    int b; // hot
    double c[100];
    struct S2 *s2_ptr;
    int d; int e;
    struct S1 *s1_ptr;
    char *c_p;
    int f; // hot
};
```

```
...
for (ii = 0; ii < N; ii++){
    sp->b = ii;
    sp->f = ii + 1;
    sp++;
}
...
```

peel.c(22): remark #30756: (DTRANS) Splitting the structure 'S3' into two parts will improve data locality and is highly recommended. Frequently accessed fields are 'b, f'; performance may improve by putting these fields into one structure and the remaining fields into another structure. Alternatively, performance may also improve by reordering the fields of the structure. Suggested field order: 'b, f, s2\_ptr, s1\_ptr, a, c, d, e, c\_p'. [VERIFY] The suggestion is based on the field references in current compilation ...

# Compiler Floating Point Model

The Floating Point options allow to control the optimizations on floating-point data. These options can be used to tune the performance, level of accuracy or result consistency.

## Accuracy

Produce results that are “close” to the correct value

–Measured in relative error, possibly ulps (units in the last place)

## Reproducibility

Produce consistent results

- From one run to the next
- From one set of build options to another
- From one compiler to another
- From one platform to another

## Performance

Produce the most efficient code possible

–Default, primary goal of Intel® Compilers

These objectives usually conflict! Wise use of compiler options lets you control the tradeoffs.

# Compiler Floating-Point Model

## The Floating-Point Compiler Switch

`-fp-model keyword` (Linux\*, Mac OS\* X)

`/fp:keyword` (Windows\*)

Lets you choose the FP semantics at a coarse granularity and specify the compiler rules for

- Value safety
- FP expression evaluation
- FPU environment access
- Precise FP exceptions
- FP contractions
- Abrupt underflow (flush to zero)
  - Denormals are set to zero
  - May improve performance, esp. if HW doesn't support denormals

# Floating-Point Keywords

Controls consistency of floating point results by restricting certain optimizations. Values for *keywords* are

- `fast [=1|2]`; default is `fast=1`
  - Allows „value-unsafe“ optimizations (=default)
  - Allows aggressive optimizations at a slight cost in accuracy or consistency.
  - Some additional approximations allowed with `fast=2`
- `precise`
  - Enables only value-safe optimizations on floating point code.
- `source`
  - Implies `precise` and enables intermediates to be computed in source precision.
  - Source is the recommended form for the majority of situations on processors supporting Intel® 64 and IA-32 platforms when SSE are enabled with `/QxSSE2` or higher.

# Floating-Point Keywords (2)

- `double`
  - Implies `precise` and enables intermediates to be computed in double or extended precision.
  - Not available in Intel® Fortran Compilers
- `extended`
  - Rounds intermediate results to 64-bit (extended) precision
  - Enables value safe optimization
- `except`
  - Enables floating point exception semantics
- `strict`
  - Strictest mode of operation, enables both the `precise` and `except` options and disables contractions (i.e., `precise + strict + disable fma`)

# The `-fp-model<key>` Switch

Key	Value Safety	Expression Evaluation	FPU Environ. Access	Precise FP Exceptions	FP contract
precise source double extended	Safe	Varies Source Double Extended	No	No	Yes
strict	Safe	Varies	Yes	Yes	No
fast=1 (default)	Unsafe	Unknown	No	No	Yes
fast=2	Very Unsafe	Unknown	No	No	Yes
except	*/**	*	*	Yes	*
except-	*	*	*	No	*

\* These modes are unaffected. `-fp-model except[-]` only affects the precise FP exceptions mode.

\*\* It is illegal to specify `-fp-model except` in an unsafe value safety mode.

# Loop Profiler

## Identify Time Consuming Loops/Functions

- Compiler switch:

`/Qprofile-functions, -profile-functions`

- Insert instrumentation calls on function entry and exit points to collect the cycles spent within the function.

- Compiler switch:

`/Qprofile-loops=<inner|outer|all>,`

`-profile-loops= <inner|outer|all>`

- Insert instrumentation calls for function entry and exit points as well as the instrumentation before and after instrumentable loops of the type listed as the option's argument.

- Loop Profiler switches trigger generation of text (.dump) and XML (.xml) output files

- Invocation of XML viewer on command line:

```
java -jar loopprofviewer.jar <xml datafile>
```

# Loop Profiler Text Dump (.dump file)

time(abs)	time(%)	self(abs)	self(%)	call_count	exit_count	loop_ticks(%)	file:line
4378070322	99.83	1810826157	41.29	1	1	41.29	deflate.c:623
647499316	14.76	642829878	14.66	16768796	16768796	0.01	trees.c:961
1462208444	33.34	487754966	11.12	10546669	10546669	7.07	deflate.c:360
119744296	2.73	119686890	2.73	513	513	2.73	util.c:63
137053240	3.13	89563374	2.04	512	512	2.04	bits.c:185
198253943	4.52	66623288	1.52	512	512	1.46	deflate.c:478
47165828	1.08	47102278	1.07	1025	1025	0.00	util.c:153
69547871	1.59	32157819	0.73	344059	344059	0.66	trees.c:454
119928920	2.73	24484703	0.56	1536	1536	0.12	trees.c:611
132122614	3.01	12346758	0.28	513	513	0.00	zip.c:106
22904224	0.52	6738036	0.15	1537	1537	0.15	trees.c:571
6675927	0.15	6672487	0.15	1	1	0.00	gzip.c:1511
15997356	0.36	6029850	0.14	139288	139288	0.12	bits.c:151
2792164	0.06	2620132	0.06	1536	1536	0.06	trees.c:485
1290621	0.03	1175933	0.03	1024	1024	0.03	trees.c:699
495976	0.01	387220	0.01	513	513	0.01	trees.c:408
259246700	5.91	261552	0.01	512	512	0.00	trees.c:857
47392247	1.08	162869	0.00	1025	1025	0.00	util.c:121
5225406	0.12	113633	0.00	512	512	0.00	trees.c:791
4385684262	100.00	66840	0.00	1	1	0.00	gzip.c:424
630948	0.01	56083	0.00	1	1	0.00	deflate.c:289
4385610543	100.00	35086	0.00	1	1	0.00	gzip.c:704
6711753	0.15	32771	0.00	1	1	0.00	gzip.c:853
82503	0.00	23661	0.00	1	1	0.00	trees.c:335
53760	0.00	22140	0.00	513	513	0.00	bits.c:164
4378817661	99.84	19622	0.00	1	1	0.00	zip.c:35
26175	0.00	13585	0.00	1	1	0.00	gzip.c:1605
12528	0.00	12466	0.00	1	1	0.00	gzip.c:1583
30798	0.00	11268	0.00	512	512	0.00	bits.c:122
9294	0.00	9232	0.00	1	1	0.00	gzip.c:915
7575	0.00	7463	0.00	1	1	0.00	util.c:170
6237	0.00	6113	0.00	2	2	0.00	gzip.c:1421
4761	0.00	4699	0.00	1	1	0.00	util.c:283
2358	0.00	2234	0.00	2	2	0.00	util.c:183
9588	0.00	1153	0.00	1	1	0.00	gzip.c:1062
10218	0.00	862	0.00	1	1	0.00	gzip.c:989
8373	0.00	686	0.00	1	1	0.00	gzip.c:939
216	0.00	154	0.00	1	1	0.00	gzip.c:1703
87	0.00	25	0.00	1	1	0.00	bits.c:99
84	0.00	22	0.00	1	1	0.00	gzip.c:1398

# Loop Profiler Data Viewer GUI (copy from sl. 46)

The screenshot displays the Loop Profile Viewer GUI with two main sections: Function Profile and Loop Profile. Annotations highlight key features:

- Function Profile View:** A table showing function-level performance metrics.
- Column Headers:** Allow independent sorting for function and loop tables.
- Filtering Menu:** A dropdown menu for filtering data based on criteria like total time or self time.
- Loop Profile View:** A table showing detailed loop-level performance metrics.


Function	Function file:line	Time	% Time	Self time	% Self time	Call Count	% Time in Loops
_main	spec.c:286	33,737,882,427	99.80	52,724,829	0.16	1	0.09
_compressStream	bzin2.c:440	22,141,802,790	65.50	37,645,635	0.11	3	0.00
_h...		22,072,329,981	65.29	111,073,801	0.33		
_B2...		21,291,282,108	62.98	382,054	0.00		
_B2...		20,773,546,2...	61.45	805,260	0.00		
_ur...		11,543,357,...	34.15	1,509,243	0.00		
_B2...		11,519,777,937	34.08	2,278,698	0.01		
_B2...		11,400,811,637	33.74	69,950,540	0.21		
_mainSort	blocksort.c:805	11,377,841,172	33.53	1,090,262,703	3.23	25	3.20
_B22_decompress	decompress.c:147	10,916,277,582	33.05	10,916,277,582	32.29	1,177	13.81
_mainQSort3	blocksort.c:676	10,237,947,453	30.28	2,203,239,483	6.57	298,338	2.97
_mainSimpleSort	blocksort.c:540	3,978,755,360	22.99	3,978,755,360	11.94	1,177	3.62
_sendMTFValues	compress.c:165	3,149,709,581	16.62	3,149,709,581	9.61	1,177	0.67
_generateMTFValues	compress.c:243	3,565,312,664	12.62	3,565,312,664	10.85	1,177	5.96
_mainGtJ	blocksort.c:564	2,388,612,638	8.55	2,388,612,638	7.32	1,177	1.79
_bsW	blocksort.c:564	1,080,193,267	6.52	1,080,193,267	3.25	1,177	1.50
_B22_bzWriteClose64@28	bzlib.c:347	49,713	3.88	49,713	0.00	3	0.00
_copy_input_until_stop	bzlib.c:594	667,041,514	1.97	667,041,514	2.00	3,172	0.36
_unRLE_obuf_to_output_FAST	bzlib.c:594	336,452,554	1.00	336,452,554	1.00	3,169	0.00

Function	Function file:line	Loop file:line	Time	% Time	Self time	% Self time	Loop entries	Min iterations	Avg iterations	Max iterations
_B22_decompress	decompress.c:147	decompress.c:516	1,542,525,615	4.60	1,542,486,842	4.60	16,620,325	1	1	2
_generateMTFValues	compress.c:165	compress.c:243	2,897,058,042	8.60	2,200,189,958	6.50	5,573,353	1	59	254
_mainSimpleSort	blocksort.c:540	blocksort.c:564	983,191,878	2.90	389,571,523	1.20	1,919,830	1	1	15
_mainSimpleSort	blocksort.c:540	blocksort.c:564	1,072,097,649	3.20	363,132,708	1.10	1,781,679	1	1	16
_mainSimpleSort	blocksort.c:540	blocksort.c:564	1,109,670,579	3.30	388,082,294	1.10	1,622,744	1	1	15
_mainSort	blocksort.c:805	blocksort.c:805	10,359,855,054	30.60	120,685,201	0.40	6,400	256	256	256
_B22_bzWriteClose64@28	bzlib.c:347	bzlib.c:347	20,772,753,426	61.40	660,714	0.00	3,147	1	1	78
_uncompress	bzlib.c:594	bzlib.c:594	11,540,287,539	34.10	1,494,966	0.00	3	1,049	1,049	1,049
_B22_bzWriteClose64@28	bzlib.c:347	bzlib.c:347	1,307,063,997	3.90	34,869	0.00	3	6	23	50

# Static Security Analysis – (SSA)

## Enhanced Application Security Catch defects early in lifecycle



- Detects over 250 different kinds of errors and security risks, such as:
    - Buffer overruns and uninitialized variables
    - Heap corruption and bad pointer usage
    - Unchecked use of input data
    - Error prone usage of libraries and language features
    - Arithmetic overflow and divide by zero
  - True global analysis crossing subroutine and file boundaries
  - Easy set up and usability, delivers immediate results
  - Intuitive GUI & Command Line interface for Windows\* and Linux
  - Displays problems and associated source code
  - Tracks state associated with issues, even as source evolves and line numbers change
  - Provides filters and sorting to organize results and reduce clutter
- C, C++, Fortran for Windows\* and Linux support**

# Static Security Analysis (SSA)

- SSA uses the compiler technology to generate analysis information
  - Compiler switch:
    - `/Qdiag-enable:sc{[1|2|3]}`
    - `-diag-enable sc{[1|2|3]}`
- Viewing SSA analysis data requires the **Intel® Inspector XE** to manage the analysis
  - Inspector automatically started after compilation if installed on the same computer
  - SSA results can be viewed on different computer where the Inspector is available
    - Data collection using compiler front end
    - Data filtering
    - Display of results
    - Interactive source view

**NOTE: Using SSA requires a Parallel Studio XE license!**

# SSA Results - Evaluation with Intel Inspector XE

**Static Source Errors**

Intel® Parallel Inspector

Summaries/Subsets Relationships

ID	Problem	Sources	Mo.	State
P1	Invalid memory access	substringfinder.cpp	Co...	Not fixed
P2	Invalid memory access	output.c	So..	Not fixed
P3	Invalid memory access	output.c	So..	Not fixed
P4	Uninitialized memory access	substringfinder.cpp	So..	Not fixed
P5	Missing allocation	substringfinder.cpp	So..	Not fixed
P6	Invalid memory access	substringfinder.cpp	So..	Not fixed
P7	Invalid memory access	substringfinder.cpp	So..	Not fixed

**Invalid memory access: Observations in Problem Set**

ID	Description	Source	Function	Module	State
X4	Allocation site	substringfi...	main	SourceCheck...	Information
X5	Deallocation ...	substringfi...	process_part1	SourceCheck...	Information
X6	Read	substringfi...	process_part2	SourceCheck...	Not fixed

**Summaries/Subsets**

**Severity**

Error	7
-------	---

**Problem**

Invalid memory access	5
Missing allocation	1
Uninitialized memory...	1

**Source**

output.c	2
substringfinder.cpp	7

**Module**

MSVCR90D.dll	2
SourceCheckerProj, ...	7


**State**

Not fixed	7
-----------	---


Sort By Item Name

# Pointer Checker

Requires a Studio XE License!



- A key feature of Intel® Parallel Studio XE 2013
- Catches out-of-bounds memory accesses through pointers
  - Identifies and reports before memory corruption occurs!
  - Buffer overruns /overflows
  - Dangling Pointers
    - memory accesses through freed pointers
- Designed for use during application debugging and testing
- Security benefits from catching vulnerabilities prior to product release.



- Enabled via compile time switches.
- User API allows control over what happens when a violation is detected
- Implemented mostly in a runtime library which is automatically linked in by the compiler
- No change to structure layout or ABIs

*Key Benefit: Enable Incrementally*  
*Pointer Checker can be enabled on a single file, group of files or all files. Pointer Checker enabled code and non-enabled code can coexist!*

# Pointer Checker - Why?

C/C++ pointers have well defined semantics for memory range access:

```
// Pointer Bounds
p = malloc(size);
// Lower Bound(p) is (char *)p
// Upper Bound(p) is lower_bound(p) + size - 1
```

Buffer overflow/overrun anomaly:

- Violation of memory safety
- Data corruption
- Erratic program behavior
- Breach of system security
- Basis of many software vulnerabilities

```
char *buf = malloc(5);
for (int i=0; i<=5;i++)
{
    buf[i] = 'A' + i;
}
```

*Pointer Checker does bounds checking!*

# Implementation Details

- Each pointer variable has bounds associated with it.
  - Bounds for  $p = \text{malloc}(\text{size})$  are:  $p$  and  $p + \text{size} - 1$
  - Bounds for  $\&v$  are:  $\&v$  and  $\&v + \text{sizeof}(v) - 1$
  - Bounds for  $\&a[i]$  are: the bounds of  $\&a$ . A pointer to an array element is allowed to traverse the entire array.
  - Bounds of  $\&a.b$  are:  $\&a + \text{offset}(b)$  and  $\&a + \text{offset}(b) + \text{sizeof}(b) - 1$
- The pointer is allowed to go out of bounds, only memory accesses are checked to make sure the address is in bounds.
- Bounds for pointer variables in memory are kept in memory in a separate location mapped via the address of the pointer.
- Bounds for pointers in registers are in registers.

# Enabling Pointer Checker

Getting started is easy...

Meet requirements:

Supported Languages	Supported Architecture	Supported Platforms	Supported Processor
C, C++	IA-32, Intel® 64	Linux*, Windows*	Intel® Pentium® 4 processor or later, or compatible processor

Compile and build your application with:

`-check-pointers= [none | write | rw]` (Linux\* OS)

`/Qcheck-pointers: [none | write | rw]` (Windows\* OS)

- Pointer Checker is off by default
- Checks all indirect accesses through pointers and accesses to arrays

*One compiler switch enables Pointer Checker!*

# Enabling Pointer Checker [cont]

- Compile with Pointer Checker Enabling switch for "r" or "rw".
- Execute; Check for out-of-bounds violations (OOB) and if OOB is true or false positive.

## Sample Program

```
%cat main.c
1 #include<stdio.h>
2 #include<malloc.h>
3 #include "chkp.h"
4
5 int main () {
6 #ifdef REPORT
7
8   __chkp_report_control(__CHKP_REPORT_TRACE_LOG,
9 0);
10 #endif
11 char *buf = malloc(4);
12 int i;
13 for (i=0; i<=4; i++) {
14     printf(" %c",buf[i]);
15 }
16 for (i=0; i<=4; i++) {
17     buf[i] = 'A' + i;
18     printf(" %c",buf[i]);
19 }
20 printf ("\n");
21 return 0;
22 }
```

## Pointer Checker Enabling – Check Bounds

- **Compile without Pointer Checker enabling switch:**

```
% icc main.c -g;./a.out
A B C D E
```

- **Compile with Pointer Checker enabling option:**

```
% icc main.c -DREPORT -check-pointers=write -
rdynamic -g;./a.out
```

```
CHKP: Bounds check error
```

```
lb: 0xe67010
```

```
ub: 0xe67013
```

```
addr: 0xe67014
```

```
end: 0xe67014
```

```
size: 1
```

**Traceback is:**

```
./a.out(main+0x21f) [0x402de7] in file main1.c line
15
```

```
/lib64/libc.so.6(__libc_start_main+0xf4)
```

```
[0x35e8ald994] in file unknown line 0
```

```
...
```

```
A B C D E
```

```
CHKP Total number of bounds violations: 1
```

```
%
```

# Pointer Checker – Detailed Control

Model	Description	
Header File	Defines intrinsics and reporting functions (chkp.h)	
Compiler Options	-check-pointers (/Qcheck-pointers)	Enables pointer checker and adds associated libraries
	-check-pointers-dangling (/Q-check-pointers-dangling)	Enables checking for dangling pointer references
	-check-pointers-undimensioned (Qcheck-pointers-undimensioned)	Enables the checking of bounds for arrays without dimensions
Intrinsics	void * __chkp_lower_bound(void **)	Returns the lower bound associated with the pointer
	void * __chkp_upper_bound(void **)	Returns the upper bound associated with the pointer
	void * __chkp_kill_bounds(void *p)	Removes the bounds information to allow the pointer in the argument to access all memory.
	void * __chkp_make_bounds(void *p, size_t size)	Creates new bounds information within the allocated memory address for the pointer in the argument
Reporting API (Function/Enumeration)	void __chkp_report_control(__chkp_report_option_t option, __chkp_callback_t callback)	Determines how errors are reported
	__chkp_report_option_t {Enumerations: __CHKP_REPORT_LOG, __CHKP_REPORT_TRACEBACK, __CHKP_REPORT_CALLBACK, __CHKP_REPORT_BPT, __CHKP_REPORT_TERM}	Controls how out-of-bounds error are reported. Enumerations in header file
RTL Functions	Provides checking on C run-time library functions that manipulate memory through pointers	

# Pointer Checker – *Samples for Control*

## Check Arrays (with Bounds)/ Undimensioned Arrays

- **Check arrays without dimensions:**

```
-[no-]check-pointers-undimensioned  
  (Linux* OS)  
/Qcheck-pointers-undimensioned[-]  
  (Windows* OS)
```

## Check Runtime Library (RTL) functions / Intrinsics

### Wrapper libraries:

```
libchkpwrap.a> [Linux*]  
libchkpwrap.lib [Windows*]
```

### Intrinsics:

Write own wrappers for RTL functions  
Work with Enabled and Non-Enabled code  
Check and create correct bounds

## Check for Dangling Pointers

- **Compiler uses a wrapper for free()  
and  
delete operator.**

### Option:

```
-check-pointers-dangling=[none |  
heap | stack | all] (Linux* OS)  
/Qcheck-pointers-dangling:[none |  
heap | stack | all] (Windows* OS)
```

## Intrinsics - Example

```
extern void *wrap_malloc(size_t bytes) {  
    // code  
    ...  
    (void *) p = my_realloc(p, old_size +  
    100);  
    p = __chkp_kill_bounds(p);  
    p = (void*)__chkp_make_bounds(p,  
    old_size + 100);  
    return p;  
}
```

# Performance Overhead

- Runtime cost is relatively high, about 2x to 5x execution time effect.
- Code size increase from 20% to a very large increase (>100% plus), depending on the application.
- Pointer Checker is seen as a debug feature
- Deployed applications are expected to have Pointer Checker disabled
- Security benefits from catching vulnerabilities prior to product release is the trade-off.

# Coding Security Support from Intel

Feature	Required Tools
Switches of Intel® Compiler to control diagnostics: Remarks, Warnings and Errors	Compiler
Code Coverage and Test Prioritization	Compiler
Static Verifier ( Static Source Code Checking )	Compiler, Inspector XE – license of IPS-XE required
Switches of Intel® Compilers for run time checking ( like variables not initialized, interface in FORTRAN, ...)	Compiler
Intel® Inspector XE memory checking	Inspector XE
Intel® Inspector XE thread checking	Inspector XE
Pointer Checking	Compiler, license of IPS-XE required
Data race detection of Intel® Debugger and Intel-extended GDB	Compiler, Debugger

# Is Pointer Checking Redundant then ?

1. Investment and intrusiveness is different for all test methods
  - Inspector XE memory checking can extend run time by a factor far above 50 !
2. Pointer Checking can discover issues not found by other methods

```
void s(char *buf)
for (int i=0; i<=5;i++)
{
    buf[i] = 'A' + i;
}

int main()
{
    char buf[5];
    char *p=buf;
    s(p);
}
```

- Static Verifier will detect this only in case both modules are compiled in one step
- Inspector XE memory checking will miss the fault since memory is not from heap

# Intel® Fortran Composer XE

## Coarray Fortran (CAF)

- Simple extension to Fortran to make Fortran into a robust and efficient parallel programming language
- Single-Process-Multiple-Data programming model (SPMD).
  - Single program is replicated a fixed number of times
  - Each program instance has it's own set of data objects – called an "IMAGE"
  - Each image executes asynchronously and normal Fortran rules apply
  - Extensions to normal Fortran array syntax to allow images to reference data in other image(s)
- Part of the Fortran 2008 standard

# Compilation

- `ifort -coarray` **!Linux\***
- `ifort /Qcoarray` **!Windows\***

along with other options. Enables compiling for CAF. By default, executable will use as many cores (real and hyperthreaded) as are available.

```
ifort -coarray -coarray-num-procs=x
```

```
ifort /Qcoarray /Qcoarray-num-procs=x
```

along with other options. Sets number of images to "x". This option will soon change to `-coarray-num-images=x`, in our next update hopefully. See the RELEASE NOTES on each beta update this summer for the latest breaking news on option and env name changes.

- No Visual Studio property settings with initial beta products. Manually add to VS "Command Line" "Additional Options"

# Running (linux)

- Simple hello world:

```
program hello_image
  write(*,*) "Hello from image ", this_image(), &
    "out of ", num_images()," total images"
end program hello_image
```

```
ifort -coarray -o hello_image hello_image.f90
./hello_image
```

Hello from image	1 out of	4 total images
Hello from image	4 out of	4 total images
Hello from image	2 out of	4 total images
Hello from image	3 out of	4 total images

# Controlling the Number of Images, Command Line

- CAF is a SPMD model. “Images” similar to MPI “Processes”
- Environment variable can set number of images
- Environment variable overrides  
–coarray-num-procs compiler option

```
Linux host> export FOR_IMAGES_NUM=2
```

```
./hello_image
```

```
Window host> set FOR_IMAGES_NUM=2
```

```
hello_image.exe
```

```
Hello from image 1 out of 2 total images
```

```
Hello from image 2 out of 2 total images
```


# Further Reading

- Coarrays in the next Fortran Standard
  - <ftp://ftp.nag.co.uk/sc22wg5/N1801-N1850/N1824.pdf>
- The New Features of Fortran 2008
  - <ftp://ftp.nag.co.uk/sc22wg5/N1801-N1850/N1828.pdf>
- Fortran 2008 Standard (current draft)
  - <http://j3-fortran.org/doc/standing/links/007.pdf>

# Intel® Math Kernel Library

## Flagship math processing library

- Multi-core ready with excellent scaling
- Highly optimized, extensively threaded math routines for science, engineering and financial applications for maximum performance
- Automatic runtime processor detection ensures great performance on whatever processor your application is running on.
- Support for C and Fortran
- Optimizations for latest Intel processors including 4th-gen Core processors



Intel® Math Kernel Library  
Industry-leading high performance math library

Included in Intel® Parallel Studio XE & Intel® Cluster Studio

Linear Algebra	Fast Fourier Transforms	Vector Math	Random Number Generators	Summary Statistics	Data Fitting
<ul style="list-style-type: none"> <li>• BLAS</li> <li>• LAPACK</li> <li>• Sparse Solvers</li> </ul>	<ul style="list-style-type: none"> <li>• Multi-dimensional (up to 7D)</li> <li>• FFTW* Interfaces</li> </ul>	<ul style="list-style-type: none"> <li>• Trigonometric</li> <li>• Hyperbolic</li> <li>• Exponential, Logarithmic</li> <li>• Power / Root</li> <li>• Rounding</li> </ul>	<ul style="list-style-type: none"> <li>• Congruent</li> <li>• Recursive</li> <li>• Wichman-Hill</li> <li>• Mersenne Twister</li> <li>• Sobol</li> <li>• Neiderreiter</li> <li>• RDRAND-based</li> </ul>	<ul style="list-style-type: none"> <li>• Kurtosis</li> <li>• Variation coefficient</li> <li>• Quantiles, order statistics</li> <li>• Min/max</li> <li>• Variance-covariance</li> </ul>	<ul style="list-style-type: none"> <li>• Splines</li> <li>• Interpolation</li> <li>• Cell search</li> </ul>
<p>CLUSTER</p> <ul style="list-style-type: none"> <li>• ScaLAPACK</li> </ul>	<p>CLUSTER</p> <ul style="list-style-type: none"> <li>• Cluster FFT</li> </ul>				

Optimized Math Operations running on Windows\*, Linux\* and OS X\*



# Application Areas which could use MKL

- **Energy** - Reservoir simulation, Seismic, Electromagnetics, etc.
- **Finance** - Options pricing, Mortgage pricing, financial portfolio management etc.
- **Manufacturing** - CAD, FEA etc.
- **Applied mathematics**
  - Linear programming, Quadratic programming, Boundary value problems, Nonlinear parameter estimation, Homotopy calculations, Curve and surface fitting, Numerical integration, Fixed-point methods, Partial and ordinary differential equations, Statistics, Optimal control and system theory
- **Physics & Computer science**
  - Spectroscopy, Fluid dynamics, Optics, Geophysics, seismology, and hydrology, Electromagnetism, Neural network training, Computer vision, Motion estimation and robotics
- **Chemistry**
  - Physical chemistry, Chemical engineering, Study of transition states, Chemical kinetics, Molecular modeling, Crystallography, Mass transfer, Speciation
- **Engineering**
  - Structural engineering, Transportation analysis, Energy distribution networks, Radar applications, Modeling and mechanical design, Circuit design
- **Biology and medicine**
  - Magnetic resonance applications, Rheology, Pharmacokinetics, Computer-aided diagnostics, Optical tomography
- **Economics and sociology**
  - Random utility models, Game theory and international negotiations, Financial portfolio management

# Intel® MKL Contents

- **BLAS**
  - Basic vector-vector/matrix-vector/matrix-matrix computation routines.
- **Sparse BLAS**
  - BLAS for sparse vectors/matrices
- **LAPACK (Linear algebra package)**
  - Solvers and eigensolvers. Many hundreds of routines total!
  - C interface to LAPACK
- **ScaLAPACK**
  - Computational, driver and auxiliary routines for distributed-memory architectures
- **DFTs (General FFTs)**
  - Mixed radix, multi-dimensional transforms
- **Cluster DFT**
  - For Distributed Memory systems
- **Sparse Solvers (PARDISO, DSS and ISS)**
  - Direct and Iterative sparse solvers for symmetric, structurally symmetric or non-symmetric, positive definite, indefinite or Hermitian sparse linear system of equations
  - Out-Of-Core (OOC) version for huge problem sizes

# Intel® MKL Contents

- **VML (Vector Math Library)**
  - Set of vectorized transcendental functions, most of libm functions, but faster
- **VSL (Vector Statistical Library)**
  - Set of vectorized random number generators
  - SSL (Summary Statistical Library) : Computationally intensive core/building blocks for statistical analysis
- **PDEs (Partial Differential Equations)**
  - Trigonometric transform and Poisson solvers.
- **Optimization Solvers**
  - Solvers for nonlinear least square problems with/without boundary condition
- **Support Functions**

# References

- Intel® MKL product Information
  - <http://software.intel.com/en-us/intel-mkl/>
- User Discussion Forum
  - <http://software.intel.com/en-us/forums/intel-math-kernel-library/>
- What are the new software tools?
  - <http://whatif.intel.com>


# Intel® Integrated Performance Primitives (IPP)

## Multicore Power for Multimedia and Data Processing

- Features
  - Rapid Application Development
  - Cross-platform Compatibility & Code Re-Use
  - Highly optimized functions from 15 Domains
    - Images and Video
    - Communications and Signal Processing
    - Data Processing
  - Performance optimizations for latest 4<sup>th</sup> generation Intel Core processors and Atom processors

"I recently compared Intel IPP to the open source library zlib. I observed a *considerable performance improvement* (at least 10-15%) using IPP compared to zlib on the latest Intel processors. I've concluded that we should migrate our open source part of the project to utilize IPP."

Vadim Kavaleroy  
Investment Analyst  
Susquehanna International Group



# Intel® Integrated Performance Primitives

## Applications

Digital Media | Web/Enterprise Data | Embedded Communications | Scientific/Technical

*High level APIs and Codecs Interfaces and Code Samples*



## Intel® Integrated Performance Primitives 16 Function Domains

### Multimedia

- Image Processing
- Color Conversion
- JPEG/JPEG2000
- Video Coding
- Computer Vision
- Realistic Rendering

### Signal Processing

- Signal Processing
- Audio Coding
- Speech Coding
- Speech Recognition
- Vector Operations

### Data Processing


- Data Compression
- Data Integrity
- Cryptography
- String Processing
- Matrix Operations

*Cross-platform C/C++ API for Code Re-use*



Optimized 32-bit and 64-bit Multicore Performance

# Intel® Integrated Performance Primitives



# Intel® Integrated Performance Primitives

## Functions and Samples

Domains	Functions	Samples
1. Image Processing	<ul style="list-style-type: none"> <li>* Geometry transformations, such as resize/rotate</li> <li>* Linear and non-linear filtering operation on an image for edge detection, blurring, noise removal and etc for filter effect.</li> <li>* Linear transforms for 2D FFTs, DFTs, DCT.</li> <li>* image statistics and analysis</li> </ul>	<ul style="list-style-type: none"> <li>* Tiled Image Processing / 2D Wavelet Transform /C++ Image Processing Classes/Image Processing functions Demo</li> </ul>
2. Computer Vision	<ul style="list-style-type: none"> <li>* Background differencing, Feature Detection (Corner Detection, Canny Edge detection), Distance Transforms, Image Gradients, Flood fill, Motion analysis and Object Tracking, Pyramids, Pattern recognition, Camera Calibration</li> </ul>	<ul style="list-style-type: none"> <li>* Face Detection</li> </ul>
3. Color conversion	<ul style="list-style-type: none"> <li>* Converting image/video color space formats: RGB, HSV, YUV, YCbCr</li> <li>* Up/Down sampling</li> <li>* Brightness and contrast adjustments</li> </ul>	
4. JPEG Coding	<ul style="list-style-type: none"> <li>* High-level JPEG and JPEG2000 compression and decompression functions</li> <li>* JPEG/JPEG2000 support functions: DCT, Wavelet transforms, color conversion, downsampling</li> </ul>	<ul style="list-style-type: none"> <li>• UIC-Unified Image Codec/ Integration with the Independent JPEG Group (IJG) library</li> </ul>
5. Video Coding	<ul style="list-style-type: none"> <li>* VC-1, H.264, MPEG-2, MPEG-4, H.261, H.263 and DV codec support functions</li> </ul>	<ul style="list-style-type: none"> <li>* Simple Media Player/ Video Encoder / h.264 decoding</li> </ul>
6. Audio Coding	<ul style="list-style-type: none"> <li>* Echo cancellation and audio transcoding, BlockFiltering, Spectral Data prequantization.</li> </ul>	<ul style="list-style-type: none"> <li>* Audio Codec Console application</li> </ul>
7. Realistic Rendering	<ul style="list-style-type: none"> <li>* Acceleration Structures, Ray-Scene Intersection and Ray Tracing</li> <li>* Surface properties, shader support, tone mapping</li> </ul>	<ul style="list-style-type: none"> <li>* Ray Tracing</li> </ul>

# Intel® Integrated Performance Primitives

## Functions and Samples

Domains	Functions	Samples
8. Speech Coding	* Adaptive/Fixed Codebook functions, Autocorrelation, Convolution, Levinson-Durbin recursion, Linear Prediction Analysis & Quantization, Echo Cancellation, Companding	* G.168, G.167, G.711, G.722, G.722.1, G.722.2, AMRWB, Extended AMRWB (AMRWB+), G.723.1, G.726, G.728, G.729, RT-Audio, GSM AMR, GSM FR
9. Data Integrity	<ul style="list-style-type: none"> <li>▪ Error-Correcting Codes</li> <li>▪ Reed-Solomon</li> </ul>	
10. Data Compression	<ul style="list-style-type: none"> <li>* Entropy-coding compression: Huffman, VLC</li> <li>* Dictionary-based compression: LZSS, LZ77</li> <li>* Burrows-Wheeler Transform, MoveToFront, RLE, Generalized Interval Transformation</li> <li>* Compatible feature support for zlib and bzip2</li> </ul>	* zlib, bzip2, gzip-compatible /General data compression examples
11. Cryptography	* Big-Number Arithmetic / Rijndael, DES, TDES, SHA1, MD5, RSA, DSA, Montgomery, prime number generation and pseudo-random number generation (PRNG) functions	* Intel IPP crypto usage in Open SSL*
12. String Processing	* Compare, Insert, change case, Trim, Find, Regexp, Hash	* "ippgrep" - regular expression matching

# Intel® Integrated Performance Primitives

## Functions and Samples

Domains	Functions	Samples
13. Signal Processing	<ul style="list-style-type: none"> <li>* Transforms: DCT, DFT, MDCT, Wavelet (both Haar and user-defined filter banks), Hilbert</li> <li>* Convolution, Cross-Correlation, Auto-Correlation, Conjugate</li> <li>* Filtering: IIR/FIR/Median filtering, Single/Multi-Rate FIR LMS filters</li> <li>* Other: Windowing, Jaehne/Tone/Triangle signal generation, Thresholding</li> </ul>	* Signal Processing Function Demo
14. Vector Math	* Logical, Shift, Conversion, Power, Root, Exponential, Logarithmic, Trigonometric, Hyperbolic, Erf, Erfc	
15. Matrix Math	* Addition, Multiplication, Decomposition, Eigenvalues, Cross-product, transposition	
Other Common Functions	* CPUtypes, Thread Number Control, Memory Allocation	* Linkages/Different language support

*Intel IPP is suitable for a very wide range of applications*

*Video broadcasting, Video/Voice Conferencing*

*Consumer Multimedia*

*Medical Imaging, Document Imaging*

*Computer Vision /Object Tracking / Machine Learning*

*Databases and Enterprise Data Management*

*Information Security*

*Embedded Applications*


*Mathematical and Scientific*

# Intel® Integrated Performance Primitives

## Function naming convention and usage

### Function names

- ✓ are easy to understand
- ✓ directly indicate the purpose of the function via distinct elements
- ✓ each element has a fixed number of pre-defined values



Name Elements	Description	Examples
Prefix	Indicates the functional data type in 1D , 2D and Matrix	ipps, ippi, ippm
Base Name	Abbreviation for the core operation	Add, FFTFwd, LuDecomp
Data Type	Describes bit depth and sign	8u, 32f, 64f
Execution mode	Indicates data layout and scaling	ISfs, C1R, P


Each function performs a particular operation on a known type of data in a specific mode

# References

- Intel® IPP product Information
  - <http://software.intel.com/en-us/intel-ipp/>
- User Discussion Forum
  - <http://software.intel.com/en-us/forums/intel-integrated-performance-primitives>
- What are the new software tools?
  - <http://whatif.intel.com>

# Intel® Composer XE 3

## Debugging Tools Overview



- **For Linux\*/Mac OS\*:**
  - Intel® Debugger 13.0 (IDB)
    - deprecated, but still required for full **Fortran debugging**
  - Enhanced GDB\* Debugger
    - GNU gdb 7.5 + branch trace and data race detection

# Intel® Composer XE 2011

## Debugging Tools Overview (cont'd)

### Intel® Debugger (IDB)

Linux\*

Mac OS\*

- Intel® Parallel Studio XE
- Intel® C++ Studio XE
- Intel® Composer XE
- Intel® C++ Composer XE
- Intel® Fortran Composer XE

- Intel® Composer XE
- Intel® C++ Composer XE
- Intel® Fortran Composer XE

- Intel® Cluster Toolkit  
Compiler Edition

### Intel® Parallel Debugger Extension

Windows\*

- Intel® Parallel Studio XE
- Intel® C++ Studio XE
- Intel® Composer XE
- Intel® C++ Composer XE
- Intel® Visual Fortran Composer XE

# Intel® Debugger (IDB) Overview


Linux\*

Mac OS\*


- GUI + Command Line

- Command Line only

- 
- Support for TBB, Cilk Plus, OpenMP & Native Threads
  - GDB alike Command Line Syntax (also in GUI if available)
  - Languages: C/C++ & FORTRAN
  - Start GUI via `idb`, Command Line via `idbc`




# Intel® Debugger (IDB) GUI



- Eclipse based GUI with similar Look&Feel
  - No project/solution to set up – use debugger on-the-fly
  - Remembers windows and positions across debugging sessions
- Intuitive and always-ready to debug



# Intel® Debugger (IDB) Command Line



```
xterm
$ idbc
Intel(R) Debugger for applications running on Intel(R) 64, Version 12.1, Build [76,472,14]
(idb) █
```

- GDB alike syntax with similar Look&Feel
  - Useful for debugging where no GUI is available
  - Allows execution of scripts (e.g. regression testing)
- Low system requirements & familiar syntax



# Intel® Debugger (IDB)

## Information at a Glance

The screenshot shows the Intel Debugger (IDB) interface with several callouts pointing to specific features:

- Sources:** Points to the source code window showing a C function `foo` with a `malloc` call.
- Customizable Toolbar:** Points to the toolbar at the top of the debugger window.
- Assembly:** Points to the assembly window showing the machine code for the selected instruction.
- Customizable Views:** Points to the Locals window, which displays a table of local variables.
- Current State:** Points to the Breakpoints window, which shows the current state of the debugger.
- Customizable Hotkeys:** Points to the bottom right corner of the debugger window.


**Locals Window Data:**

Expression	Value	Type
tmp	0x0	long *
x	1	long



**Breakpoints Window Data:**

Location	File	Module	Address
foo (x=1)	memleak.c:15	memLeak	0x000000000400548
main (argc=1, argv=(char **) 0x7fff18c...	memleak.c:32	memLeak	0x0000000004005ad
__libc_start_main ()		libc-2.11.1.so	0x00007f189da99c4d

# Intel® Debugger (IDB) Customize I



Customize GUI





Customize Toolbars

# Intel® Debugger (IDB)

## Customize II

Windows showing Tables can be customized:

- Hide or Show Column Entries
- Change Column Width
- Rearrange Entries



# Intel® Debugger (IDB) Views

- **Console:**

Like Command Line; GDB Syntax

- **Threads:**

List of all Threads

- **Callstack:**

Show Backtrace

- **(Vector) Registers:**

CPU Registers

- **Memory:**

Manually inspect Memory

- **Assembler:**

Instruction level debugging

- **Breakpoints:**


Events to Stop (Code, Data, Synchronization)

- **Locals, (Vector) Evaluations:**

Show Variables; automatic or selected

- **Source Files:**

List of all Sources with Debug Information



# Intel® Debugger (IDB)

## Run Control

- **Run or Continue:**

Run or Continue Execution

- **Suspend:**

Pause Execution

- **Kill:**

Terminate Application

- **Interrupt:**

Abort Requests to IDB (e.g. complex evaluations)

- **(Instruction) Step Into/Over:**

Next Statement/Instruction following or not following calls/jumps

- **Run Until (Caller):**


Continue until Caller or until Location is reached

- **Restart (with Arguments):**

Terminate application & run again (providing different arguments)

- **Current Thread Set** (default: \$lasteventingthread):

Set of Threads that are affected by Run Control




Important Run Control

Full Run Control

# Intel® Debugger (IDB)




## Run Control Indicator




- The Run Control Indicator shows the current Application State
  - This Indicator is live Information from the Debugger
- ➔ In some cases (e.g. Inferior Calls) the Debugger might execute Code from the Application even if stopped!

# Intel® Debugger (IDB)

## Start a Debugging Session Easily



# Intel® Debugger (IDB) Signals




- Flexible and Easy Control of System Signals

# Intel® Debugger (IDB)

## Features: Vectors


- Vectors (SIMD) are composed from multiple elements of same Type
- IDB offers flexible Visualization and Access within **Vector Evaluations & Registers Views**



# Intel® Debugger (IDB)

## Features: Assembly View

- IDB supports Intel (default) and AT&T style Assembly
- Assembly is shown when no Source is known (can be customized)



# Intel® Debugger (IDB)


## Features: Breakpoints I

• IDB supports different kinds of Breakpoints:

-  Code Breakpoints
-  Thread Synchronization Breakpoints
-  Data Breakpoints

• Dynamic multiplicity\*:

- 1:1:  
one address
- 1:n (n > 1):  
multiple addresses
- 1:0:  
no address (yet)




\*: Not for Data Breakpoints

# Intel® Debugger (IDB)

## Features: Breakpoints II

### •Create or Modify Breakpoints:




### •Flexible Attributes:

- Condition
- Skip Count
- Action
- more later...

### •**Browse** buttons aid in selecting proper values:


- Symbol Selector
- Thread Set Selector



# Intel® Debugger (IDB)

## Features: C++ Template Support


- “Pretty print” C++ Templates in **Locals, (Vector) Evaluations** Views and also edit most of them:



“Pretty Print” & Edit

# Intel® Debugger (IDB)

## Features: Multithreading Support I




- Thread Types:
  - **Native (POSIX)**
  - **OpenMP**
  - **TBB**
  - **Cilk**

- Stop event (e.g. Breakpoint) interrupts all threads, except for **Uninterrupted** Threads
- Only Threads that are **Thawed** can continue execution
- **Frozen** Threads remain interrupted at all time

# Intel® Debugger (IDB)


## Features: Multithreading Support II



- The **Current Thread Set** is the set of Threads affected by Run Control
- Current Thread** defines what Data Views show, if related to Thread Context (e.g. Registers of a Current Thread)

# Intel® Debugger (IDB)


## Features: Multithreading Support III



- Breakpoints assist in Multithread Debugging:
  - **Thread Filter** of Code & Data Breakpoints
  - **Thread Filter** of Thread Synchronization Breakpoints create a Barrier – specified Threads of Set stop at Barrier and only continue if all Threads reached the Barrier
  - **Stopping Threads** will stop the specified Threads

# Intel® Debugger (IDB)

## Features: Support for Parallel Programming I



# Intel® Debugger (IDB)

## Features: Support for Parallel Programming I

- To use Parallel Programming Requirements enable **Compiler**

**Option:** `-debug parallel`

- For Thread Data Sharing Detection the Application Code will be instrumented:

- Calls to library routines added for memory accesses.
- These library routines will determine if read/write accesses to the same memory location are done by different threads (without synchronizations) during execution. Such are Race Conditions.
- Works for Native, TBB, Cilk Plus & OpenMP Threads

# Intel® Debugger (IDB)

## Features: Support for Parallel Programming II

- Thread Data Sharing Events:

The screenshot displays the Intel Debugger (IDB) interface. The 'Parallel' menu is open, showing 'Enable Detection' and 'Break on Event' options, both of which are checked. A yellow arrow labeled 'Enabled & Break' points to these options. Below the menu, the 'Thread Data Sharing Events' window is visible, showing a list of 'Shared Data Access' events. The events are listed as follows:

Event	Accesses	Threads
initial - 0x000000000060372c 4 Bytes, 2 accesses	2	2 threads
initial - 0x000000000060372c 4 Bytes, 5 accesses	5	1 thread
initial - 0x000000000060372c 4 Bytes, 4 accesses	4	4 threads
initial - 0x000000000060372c 4 Bytes, 4 accesses	4	4 threads

The 'Detected Races' window is also visible, showing a list of race events. The events are listed as follows:


Event	Thread
barrier.cpp(35) read, Thread=1 (OpenMP task=4)	Thread=1 (OpenMP task=4)
barrier.cpp(32) write, Thread=5 (OpenMP task=18)	Thread=5 (OpenMP task=18)
barrier.cpp(32) write, Thread=4 (OpenMP task=17)	Thread=4 (OpenMP task=17)
barrier.cpp(32) write, Thread=3 (OpenMP task=16)	Thread=3 (OpenMP task=16)

A yellow arrow labeled 'Detected Races' points to the race events. The status bar at the bottom indicates 'Analysis active, Filterset: Suppress, Active Filters: 0' and shows a red dot indicating a race at '0x0000000000401bd5 in phase1 () "/>

# Intel® Debugger (IDB)

## Features: Support for Parallel Programming III

- Thread Data Sharing Filters:



# Intel® Debugger (IDB)

## Features: Support for Parallel Programming IV

- Tasks, Task Spawn Tree & Teams Views:

The screenshot displays the Intel(R) Debugger interface with the following components:

- Source Code:** A C++ file named `barrier.cpp` is open, showing a parallel function `phasel(void)` using OpenMP pragmas. The code includes a `pragma omp parallel for` loop and a nested `for` loop.
- Tasks View:** A table listing active tasks with columns for ID, State, Type, Team, Parent, # Spawned, Thread, Created At, and Path.
- Task Spawn Tree:** A hierarchical view showing the spawning of tasks 19, 20, 21, and 22 from task 4.
- Teams View:** A table showing thread teams, including team 8 which has 4 threads.

ID	State	Type	Team	Parent	# Spawned	Thread	Created At	Path
4	Suspended	Implicit, ti	1	0	1	1	-	-
19	Running	Implicit, ti	8	4	0	1	unknown:43	-
20	Running	Implicit, ti	8	0	0	3	unknown:43	-
21	Running	Implicit, ti	8	0	0	4	unknown:43	-
22	Running	Implicit, ti	8	0	0	5	unknown:43	-

ID	Parent	# Threads	Created At	Path
1	0	1	-	-
8	1	4	unknown:43	-

0x00000000040204a in phase2 () "/nfs/iul/disks/iul\_team2/gzitzlsb/openmp/barrier.cpp":45

# Intel® Debugger (IDB)

## Features: Support for Parallel Programming V

- Barriers, Taskwaits & Locks Views:

The screenshot shows the Intel(R) Debugger interface. The main window displays the source code for `barrier.cpp` with the following lines highlighted:

```
20  
21# pragma omp barrier  
22   arr[i] = left + right;  
23  
24# pragma omp single  
25   computations++;  
26 }
```

Below the code editor, three views are visible: **Barriers**, **Taskwaits**, and **Locks**. The **Barriers** view is active and displays a table of barrier information:

ID	State	Type	Team	# Threads Reached	# Wait For Tasks	Created At	Path	
100100002	Wait	Threads	Workshare	10	1	0	unknown:24	-
120			share	12	1	0	unknown:24	-
140			share	14	1	0	unknown:24	-

Yellow callout boxes with arrows point to the code and the views:


- OpenMP Barriers** points to the `pragma omp barrier` line.
- OpenMP Locks** points to the `pragma omp single` line.
- OpenMP Waiting Tasks** points to the `computations++;` line.

# Intel® Debugger (IDB)

## Features: Support for Parallel Programming VI

- Serialize Execution:


- Execute **OpenMP** & **Cilk Plus** Code in one Thread
- Distinguish Common Errors & Parallelization Errors (one click)



# Intel® Debugger (IDB)

## Features: Threading Building Blocks (TBB) Support


- TBB Containers can be “pretty printed” and even partly modified in **Locals, (Vector) Evaluations Views**
- Step into** does not follow TBB internal Code (set `$usessteppingrules=0` to turn off default)



# Intel® Debugger (IDB)

## Features: Record & Execute Command File


- **Record** a session – all I/O or just commands/actions during a session
- Replay via **Execute Command File**



# Intel® Debugger (IDB)

## IDB + Inspector XE (memory analysis)

- IDB can be used with Inspector XE as debugger during memory analysis.
- For Inspector XE configure IDB as debugger via setting environment variable **\$INSPXE\_DEBUGGER** (if not set GDB is the default)




Start debugger when first error was found or with every application start

IDB supports all 3 types of memory analysis

# Intel GDB\* Enhancements for Linux\* and OS X\*

- Intel® Debugger (Intel® IDB) is deprecated:
  - No changes since Composer XE 2013 (13.0)
  - May be removed in a future release
- New debugger solution is based on GNU Project Debugger (GDB)\* and includes additional improvements & features:
  - Improved Fortran support
  - GNU GDB 7.5 based
  - Python\* 2.6 or greater required
- Linux only:
  - BTrace: Trace back branches taken before a crash
  - PDBX: Parallel Debug Extensions for data race detection
  - Pointer Checker: Assist in finding pointer issues
  - Intel® Transactional Synchronization Extensions

# Btrace – Based on Branch Trace Store



- Intel® Atom™ Processor and Intel® Core™ Family processors use *Branch Trace Store* (BTS) to keep track of branch ("return") targets for calls
- BTRACE uses this buffer to display stack trace "even in difficult cases" where program execution stack is corrupted completely

# Btrace: Example

## Program crashed and back trace not available

```
(gdb) btrace enable auto
(gdb) run
Starting program: ../gdb/trace/examples/function_pointer/stack64
Program received signal SIGSEGV, Segmentation fault. 0x00000002a in ?? ()
(gdb) bt
```

```
#0 0x0000000000000002a in ?? ()
#1 0x00000000000000017 in ?? ()
#2 0x000000000040050e in fun_B (arg=0x4005be) at src/stack.c:32
#3 0x0000000000000000 in ?? ()
```

Look at the branch trace.

List of blocks starts from the most recent block (ending at the current pc) and continues towards older blocks such that control flows from block n+1 to block n.

```
(gdb) btrace list 1-7
1 in ?? ()
2 in fun_B () at src/stack.c:36-37
3 in fun_B () at src/stack.c:32-34
4 in main () at src/stack.c:57
5 in fun_A () at src/stack.c:22-25
6 in fun_A () at src/stack.c:18-20
7 in main () at src/stack.c:51-56
```

# PDBX: Data Race Detection

- **Sample:** Assume two global variables  
a=1  
b=2
- And two threads T1, T2 executing  
T1: x = a + b  
T2: b = 42
- Value of x depends on execution order:  
If T1 runs before T2 → x = 3  
If T2 runs before T1 → x = 43
- Execution order is not guaranteed unless synchronization methods are used.

```
(gdb) pdbx enable
(gdb) continue
data race detected
1: write answer, 4 bytes from foo.c:36
3: read answer, 4 bytes from foo.c:40
Breakpoint -11, 0x401515 in foo () at foo.c:36
(gdb)
```

# PDBX: Requirements / Notes

- Intel® Compiler 12.1 or later
- Compilation by `-debug parallel` option
- Core, Atom, or Xeon Phi processor
- Pthreads and /or OpenMP
  - Works too for Intel® Cilk™ Plus and Intel® TBB tasking but currently many *false positives*
- Python 2.6+ (host)
- Rather intrusive execution mode
  - Both for time and memory consumption
  - But mode can be enabled/disabled any time during debug session
- Many filters available to fine tune PDBX debugging



# Backup

# Packages Overview

- The Intel® Composer XE is part of the respective Intel® Parallel Studio XE series and being installed within one unique setup
  - Please refer to the Introduction Module for the Intel® Parallel Studio XE setup
- The Intel® Composer XE is also available separately and can be installed independently from Parallel Studio XE.
  - for example when dedicated Composer updates are available

# Intel® Composer XE Standalone Packages

- Installation Fileset Combinations
  - Six filesets with C/C++ and Fortran for Windows and Linux
  - C/C++ and Fortran packages sync'd on same platform are always in sync and have the same package versions and features (except of language specific deviations)
  - Optional IPP packages additionally to the compiler packages due to licensing reasons
- Visual Studio\* Shell and Libraries included in the Intel® Visual Fortran Composer XE
  - Packages with VS Shell included don't require any Visual Studio installation on the system
  - This is the default package available from Intel's Registration Center product download.

# Composer XE Standalone Packages (2)

- Fileset Format

- [l|m|w]\_[c|f]compxe\_2013\_sp1.<updatenr>.<buildnr>.[exe|tgz|dmg]
  - *w* (Windows), *l* (Linux)
  - *c* (C/C++ package), *f* (Fortran package)
  - *exe* (Windows self-extracting executable), *tgz* (Linux tarball), *dmg* (OS X package file)

- Examples:

- `w_ccompxe_2013_sp1.0.103.exe`
  - Self-extracting Windows .exe file with C++ IA32 and Intel64 installers
- `l_fcompxe_2013_sp1.0.103.tgz`
  - Linux archive with Fortran IA32 and Intel64 insntallers

# Composer XE Optional Standalone Packages

- Optional Packages

- Due to licensing or packaging reasons some packages may be offered as additional standalone installations
- IPP Crypto and Fixed Size Function packages for C/C++ Compiler
  - I|m|w\_ccomp\_xe\_[gen|crypto]\_ipp\_7.x.x.xxx.exe|tgz|dmg
- Fortran packages without Visual Studio Shell and Libraries (requires an existing Visual Studio installation on the development machine)
  - w\_fcomp\_xe\_novsshell\_2013\_sp1.x.xxx.exe

# License File

- A FLEXlm\* license file is installed by the setup program at
  - <common files>\intel\licenses (Windows)
  - </opt/intel/licenses> or <\$HOME/intel/licenses> (Linux)
- Web users: a serial number is sent by email for web users
- CD users: a serial number is provided on the CD
  - Register to get one year free support and software updates
- Install-time license checking

# Product Activation

- 3 methods
  - Using license file (.lic)
    - License file is sent from Intel Registration Center upon serial number registration
  - Remote activation
    - Requires Internet access
    - Product S/N is registered at the Intel Registration Center and license file is sent via email
  - Evaluation installation
    - Does not require registration and license file

# License Types

- Single user
  - Product is licensed to be installed and used by a single user
- Concurrent users
  - Product may be installed and used on as many systems as desired, but no more than N users may compile concurrently
  - See FLEXIm user guide for more details
- Special licenses
  - Evaluation licenses: single-user license that expires (typically in 30 days)
  - CD license: single-user license, support service available upon registration
  - Beta license: expires after a certain time period and is used by beta compiler only

# Installation - Windows Setup

- GUI-driven and self-explaining installer
  - Self-extracting .exe file automatically starts the installer
  - Setup can be re-voked again to perform a repair/modify/remove operation of the toolsuite via
    - Windows Control Panel/Add or Remove Programs -> Intel Composer XE 2013 SP1 for Windows\*
    - <installidr>\setup\_x\_xxx\Setup.exe, e.g. C:\Program Files (x86)\Intel\Composer XE 2013 SP1\setup\_1\_139\Setup.exe
    - C:\Program Files (x86)\Intel\Download\<package>\Setup.exe

# Installation - Linux Setup

- Command line-driven installer/uninstaller

- Extract file to temp dir and run installer, e.g.:

```
$ cd
```

```
$ tar -xzvf l_ccomp_xe_2013_sp1.1.106.tgz
```

```
$ cd l_ccomp_xe_2013_sp1.1.106
```

```
$ ./install //choose root , sudo root or user installation
```


```
.....
```

```
$ <installdir> bin/uninstall.sh //root or user uninstall, for example:
```

```
$ sudo /opt/intel/l_ccomp_xe_2013_sp1.1.106/bin/uninstall.sh
```

# Directory Structure - Windows


C:\Program Files\Intel\Composer XE 2013 SP1\ (on 32-bit system)  
C:\Program Files (x86)\Intel\Composer XE 2013 SP1\ (on 64-bit system)



If not specified otherwise during the installation, the C/C++ Compiler and Fortran Compiler are being installed under the same directory structure

- **bin** - start scripts, C/C++ and/or Fortran compiler executables and DLLs (including source checker)
- **compiler:** Intel specific includes and libraries
- **Documentation:** Complete documentation of compiler, debugger and libraries
- **help:** Registration files for on-line help.
- **ipp:** IPP library files and demos
- **mkl:** MKL library files with examples, tests and benchmarks
- **redistrib:** Redistributable libraries, can be packed with application if required on runtimes the update number
- **Samples:** Visual Studio projects with compiler sample source code for C/C++ and/or Fortran and IPP
- **setup\_0:** Uninstall, repair or modify can be started with Setup.exe from here
- **tbb:** TBB library files with examples
- **VS Integration:** Plug-in files for Compiler integration into Visual Studio

# Directory Structure – Linux/OS X



- Based on Linux standard guidelines
- **bin** contains compiler source scripts `iccvars.(c)sh/ifortvars.(c)sh` and symbolic links to executables that can be invoked by user
- The **lib**, **include**, **ipp**, **mkl**, and **tbb** directories are symbolic links to directories with the same name in `composerxe` and contain the performance library headers and libraries
- **man** contains man pages for executable commands. It will be a symbolic link to the man structure under `composerxe`.
- **licenses** is the default licenses files (`.lic`) directory
- **composerxe** is a symbolic link pointing to the `composer_xe_2013_sp1` directory.
- **composer\_xe\_2013\_sp1** is a physical directory containing links to header files and libraries that are part of the latest Composer XE product configuration.
- **composer\_xe\_2013\_sp1.<n>.<bld>**, for example `composer_xe_2013_sp1.1.106` is the directory containing the files from a Composer release. `<n>` is the update number (starting at 0 for RTM) and `<bld>` is the build number.
- If `/opt` is a network drive, symbolic links are created if possible. They may or may not be visible on systems on which the product was not installed.

# Intel® C++ and Fortran Compiler XE 13.x Standards Conformance

- C (**icc**):
  - ISO/IEC 9899:1990 standard (**-std=c89**)
  - C90 plus GNU extensions (**-std=gnu89** [default])
  - ISO/IEC9899:1999 standard (**-std=c99** or **-std=c9x**)
  - C99 plus GNU extensions (**-std=gnu99**)
- C++ (**icpc**):
  - ISO/IEC 14882:1998 standard plus GNU\* extensions (**-std=gnu++98** [default])  
Note:  
Almost same as 2003 standard (ISO/IEC 14882:2003)
  - ISO/IEC 14882:2011 standard (**-std=c++0x**) or including GNU\* extensions (**-std=gnu++0x**)
- On Windows\* (**/Qstd=**), you can only specify values **c99** and **c++0x**!

# Intel® C++ and Fortran Compiler XE Standards Conformance – cont'd

- Fortran (**ifort**):
  - Fortran 90, FORTRAN 77 and FORTRAN IV (FORTRAN 66)
  - Fortran 95 (ISO/IEC 1539:1997) [baseline]
  - Fortran 2003 (ISO/IEC 1539:2004)
  - Fortran 2008 (ISO/IEC 1539:2010)
  
- More details here:  
<http://software.intel.com/en-us/articles/intel-fortran-compiler-support-for-fortran-language-standards/>

# Intel® C++ and Fortran Compiler XE 13.x Standards Conformance – C99

- Full support:

- Restricted pointers (**restrict** keyword).
- Variable-length Arrays
- Flexible array members
- Complex number support (**\_Complex** keyword)
- Hexadecimal floating-point constants
- Compound literals
- Designated initializers
- Mixed declarations and code
- Macros with a variable number of arguments
- Inline functions (**inline** keyword)
- Boolean type (**\_Bool** keyword)
- **long double** is 64 bit, not 128 bit (only Linux\*)

- More here:

<http://software.intel.com/en-us/articles/c99-support-in-intel-c-compiler>

# Intel® C++ and Fortran Compiler XE 13.x

## Standards Conformance – C++0x

- **Most important** already in Intel® C++ Composer XE 2011!
- New in Intel® C++ Composer XE 2013:
  - Additional type traits
  - Uniform initialization
  - Generalized constant expressions (partial support)
  - **noexcept**
  - Range based for loops
  - Conversions of lambdas to function pointers
  - Implicit move constructors and move assignment operators
  - Support for C++11 features in gcc 4.6 and 4.7 headers
- See full list here:  
<http://software.intel.com/en-us/articles/c0x-features-supported-by-intel-c-compiler/>

### Note:

Features marked as “partial” are not complete but at least implemented to the extent of the typical default compiler (GNU GCC\* or Microsoft Visual Studio\*)!

# Intel® C++ and Fortran Compiler XE 13.x Standards Conformance – Fortran 2003

- Almost complete!
- New:
  - Default initialization of polymorphic variables
  - The keyword **MODULE** may be omitted from **MODULE PROCEDURE** in a generic interface block when referring to an external procedure
- The following features will be added later:
  - User defined derived type I/O
  - Parameterized derived types
  - Transformational intrinsics, such as **MERGE** and **SPREAD**, in initialization expressions

# Intel® C++ and Fortran Compiler XE 13.x Standards Conformance – Fortran 2008

- Most important features already there
- New:
  - **ATOMIC\_DEFINE** and **ATOMIC\_REF**

Refer to the Release Notes for the full list!

## 3.7 Fortran 2003 and Fortran 2008 Feature Summary

The Intel Fortran Compiler supports many features that are new in Fortran 2003. Additional Fortran 2003 features will appear in future versions. Fortran 2003 features supported by the current compiler include:

- The Fortran character set has been extended to contain the 8-bit ASCII characters ~ \ [ ] ` ^ { } | # @
- Names of length up to 63 characters
- Statements of up to 256 lines
- Square brackets [ ] are permitted to delimit array constructors instead of (/ /)
- Structure constructors with component names and default initialization

And also here:

<http://software.intel.com/en-us/articles/intel-fortran-compiler-support-for-fortran-language-standards/>

# C++11 Support

## New C++11 features

enabled by switch

`/Qstd=c++11` (Windows), `-std=c++11` (Linux, OS X)  
(old `std=c++0x` switch still working)

- RVALUE references
- Variadic templates
- Extern templates
- Hexadecimal Floating Constants
- Atomic Types
- Right angle brackets
- Extended friend declarations
- Mixed string literal concatenations
- Support for long long
- Static assertions
- Universal character name literals
- Strongly-typed enums
- Lambda functions
- ...

# Pointer Checker



- Finds buffer overflows and dangling pointers before memory corruption occurs
- Powerful error reporting
- Integrates into standard debuggers (Microsoft, gdb, Intel)

## Dangling pointer

```
{  
  char *p, *q;  
  p = malloc(10);  
  q = p;  
  free(p);  
  *q = 0;  
}
```

## Buffer Overflow

```
{  
  char *my_chp = "abc";  
  char *an_chp = (char *) malloc (strlen((char *)my_chp));  
  memset (an_chp, '@', sizeof(my_chp));  
}
```

## CHKP: Bounds check error

### Traceback:

```
./a.out(main+0x1b2) [0x402d7a] in file mems.c at line 13
```

Pointer Checker Highlights Programming Errors For More Secure Applications

# Intel® Math Kernel Library 11.1

## New Features


- Conditional Numerical Reproducibility (CNR) for unaligned memory
  - Balances performance with reproducible results by allowing greater flexibility in code branch choice and by ensuring algorithms are deterministic. More information: [training site](#) or the Intel® MKL User's Guide).
  - This release extends the feature to remove memory alignment requirements
  - Memory alignment is still recommended for best performance

# Debugging on Linux\* and OS X\*

- Intel® Debugger (Intel® IDB) is deprecated:
  - No changes since Composer XE 2013 (13.0)
  - May be removed in a future release
- New debugger solution is based on GNU Project Debugger (GDB)\* and includes additional improvements & features:
  - Improved Fortran support
  - GNU GDB 7.5 based
  - Python\* 2.6 or greater required
- Linux only:
  - BTrace: Trace back branches taken before a crash
  - PDBX: Parallel Debug Extensions for data race detection
  - Pointer Checker: Assist in finding pointer issues
  - Intel® Transactional Synchronization Extensions

# Intel® MIC Architecture Debugging (Linux\*)


- Integration into Eclipse IDE\*:
  - Supports C/C++ & Fortran
  - Support for offload extensions
  - Multiple coprocessor cards
- Command line debugging of native code possible with GDB



# Intel® MIC Architecture Debugging (Windows\*)

- Integration into Microsoft Visual Studio\* 2012:

- Supports C/C++ & Fortran
- Support for offload extensions
- Multiple coprocessor cards



- Debugging native coprocessor applications also possible with remote attach

# Intel® Xeon Phi™ Coprocessors Offload Extensions

# Language Extensions for Offload New Features

- Three new clauses
  - `status` - affects handling of offload failures
  - `mandatory` - when coprocessor unavailable for offload and:
    - No `status` → code aborts
    - With `status` → user code directs action
  - `optional` - when coprocessor unavailable for offload and:
    - No `status` → code runs on CPU
    - With `status` → query value only – no user directed action available
- User-specified offload clauses override `-offload` compiler option settings

# Language Extensions for Offload New Features

- New `-offload (/Qoffload)` option w/keywords:
  - `none` - Offload directives are ignored and cause warnings at compile-time
  - `mandatory` (default) - Offload directives processed When coprocessor is unavailable for offload → code aborts
  - `optional` - Any offload directives are processed When coprocessor is unavailable → code runs on CPU
- These options are overridden by user-specified offload clauses

# Language Extensions for Offload New Features

- New environment variables
  - OFFLOAD\_DEVICES: Restrict process to using only the coprocessors specified
  - OFFLOAD\_INIT: Hint to the offload RTL when to initialize coprocessors
  - OFFLOAD\_REPORT: Enables different levels of tracing and statistical information from offload.
  - OFFLOAD\_ACTIVE\_WAIT: Controls keeping host CPU busy during DMA transfers

# OpenMP\* 4.0 RC2

# C++ Features

## Intel® Composer XE 2013 SP1

### • OpenMP\* 4.0 SIMD Constructs

- `#pragma omp simd [clause] // transform loops into SIMD loops`
  - *{for-loops}*
  - *clause is one of: safelen, linear, aligned, private, lastprivate, reduction, collapse(n)*
- `#pragma omp for simd [clause] // transform loops into SIMD loops and execute with a thread team`
  - *{for-loops}*
  - *clause is one of: safelen, linear, aligned, private, lastprivate, reduction, collapse(n)*
- `#pragma omp declare simd [clause] // create a SIMD version of a function`
  - *{function declaration or definition}*
  - *clause is one of: simdlen, linear, aligned, uniform, reduction, inbranch, notinbranch*
  - *Restriction: if the function is called in a SIMD loop, it cannot result in the execution of an OpenMP construct (but the function itself can be called from an OpenMP parallel region)*

# C++ Features

## Intel® Composer XE 2013 SP1

### • OpenMP\* 4.0 TARGET Constructs

- #pragma omp target data [clause] // create a device data environment for the extent of the target region
  - {structured block}
  - clause is one of: device(integer-expression), map(list), if(scalar-expression)
- #pragma omp target [clause] // create a device data environment and execute the construct on the device
  - {structured block}
  - clause is one of: device(integer-expression), map(list), if(scalar-expression)
- #pragma omp declare target // map functions and variables to a device; create device-specific versions of functions
  - {variable and function declarations}
- #pragma omp target update motion-clause [clause] // update variables between the device data environment and the host data environment
  - motion-clause is one of: to(list), from(list)
  - clause is one of: device(integer-expression), if(scalar-expression)

# Fortran Features

## Intel® Composer XE 2013 SP1

### • OpenMP\* 4.0 SIMD Constructs

- !\$omp simd [clause] ! transform loops into SIMD loops
  - {do-loops}
  - [!\$omp end simd] ! Optional end directive
  - clause is one of: safelen, linear, aligned, private, lastprivate, reduction, collapse(n)
- !\$omp do simd [clause] ! transform loops into SIMD loops and execute with a thread team
  - {do-loops}
  - [!\$omp end do simd] ! Optional end directive
  - clause is one of: safelen, linear, aligned, private, lastprivate, reduction, collapse(n)
- !\$omp declare simd [clause] ! create a SIMD version of a procedure
  - { function or subroutine definition }
  - clause is one of: simdlen, linear, aligned, uniform, reduction, inbranch, notinbranch
  - Restriction: if the procedure is called in a SIMD loop, it cannot result in the execution of an OpenMP construct (but the procedure itself can be called from an OpenMP parallel region)

# Fortran Features

## Intel® Composer XE 2013 SP1

### • OpenMP\* 4.0 TARGET Constructs

- !\$omp target data [clause] ! create a device data environment for the extent of the target region
  - {structured block}
  - clause is one of: device(integer-expression), map(list), if(scalar-expression)
- !\$omp target [clause] ! create a device data environment and execute the construct on the device
  - {structured block}
  - clause is one of: device(integer-expression), map(list), if(scalar-expression)
- !\$omp declare target (list)// map functions and variables to a device; create device-specific versions of functions
  - list is a comma-separated list of variables, procedures, and named common blocks
- !\$omp target update motion-clause [clause] // update variables between the device data environment and the host data environment
  - motion-clause is one of: to(list), from(list)
  - clause is one of: device(integer-expression), if(scalar-expression)

# Coarrays with Intel® Xeon Phi™ Coprocessors

# Coarrays with Intel® Xeon Phi™ Coprocessors (cont.)

- Restrictions on coarray w/offload regions
  - coindexing is **\*not\*** allowed
    - All accesses to coarrays within an offload region must be to the local copy of the coarray
  - SYNC ALL, SYNC MEMORY, SYNC IMAGES, or LOCK/UNLOCK use is **\*not\*** allowed in an offload region
  - Coarrays must **\*not\*** be allocated or deallocated within an offload region

# Coarrays with Intel® Xeon Phi™ Coprocessors (cont.)

- *Heterogeneous* coarray application
  - Some images run on the Intel® 64 host system and some run on an Intel® Xeon Phi™ coprocessor
- Requirements Overview (*heterogeneous*)
  - Refer to the RNs for specifics
  - Compile with:
    - `-coarray=coprocessor`
    - `-coarray-config-file=<file>` to specify coarray (CAF) configuration file

# Coarrays with Intel® Xeon Phi™ Coprocessors (cont.)

- Requirements Overview (*heterogeneous*) (cont.)
  - Creates host & Intel® MIC architecture native executable – must copy native exe to coprocessor
  - Coarray (CAF) configuration file - provides critical MPI config information
    - Path to native executable on the coprocessor
    - Number of images to run on the host and coprocessor
  - Run with:
    - Environment variable I\_MPI\_MIC set to ENABLE

# Coarrays with Intel® Xeon Phi™ Coprocessors (cont.)

- Native coarray application
  - Runs exclusively on the coprocessor
  - Compile with `-coarray` and `-mmic`
  - Coarray (CAF) configuration file is not required
  - Compilation creates single Intel® MIC architecture native executable image
  - Copy image along with referenced library shared objects to the card - execute as normal native executable

# Fortran 2003 Support: User Defined Derived Type Input and Output

# User Defined I/O

## Intel® Composer XE 2013 SP1

- Fortran 2003 Standard feature
- Custom subroutines can be used to handle input or output for objects of a derived type
- Motivation:
  - Allows user control of the way input and output is handled for derived type variables
  - The default I/O handlers can not be used for some derived types
    - Component variables printed in order of declaration
    - No pointers or allocatables

# User Defined I/O

## Intel® Composer XE 2013 SP1

- User defined I/O subroutines may be specified with an explicit interface or as a type bound generic procedure.
- Supports Formatted, Unformatted, Namelist, and List Directed I/O
  - read(formatted), write(formatted)
  - read(unformatted), write(unformatted)
- For Formatted I/O, format specified with dt edit descriptor
  - write (6, fmt=`(dt"my\_type"(3,4))`) foo
- Formatted user defined I/O subroutines may only use formatted I/O statements, and unformatted user defined I/O subroutines may only use unformatted I/O statements

# User Defined I/O

## Intel® Composer XE 2013 SP1

- All User Defined I/O subroutines must handle :
  - the derived type variable,
  - Unit – the I/O unit number
  - iostat
  - iomsg
- Formatted I/O routines also handle
  - Iotype - "Listdirected", "Namelist", or "DT"
  - v\_list – array containing values passed through the DT edit descriptor, defaults to a zero sized array
  - If I/O procedure could be called recursively, the routine must be labeled as recursive

# User Defined I/O Example

```
module ball_mod
  type ball
    integer, pointer :: x, y
    double precision :: r
    integer           :: id
  contains
    procedure :: write_ball
    generic   :: write(formatted) => write_ball
  end type
  interface read(formatted)
    module procedure read_ball
  end interface=
Contains
...
! subroutines write_ball & read_ball
! are on the following slide
end module
```

Type can not be  
output with  
default I/O

# User Defined I/O Example

```
subroutine write_ball (this, unit, iotype, vlist,
iostat, iomsg)
  class(ball), intent(in)           :: this
  integer, intent(in)               :: unit
  character (len=*), intent(in)     :: iotype
  integer, intent(in)               :: vlist(:)
  integer, intent(out)              :: iostat
  character (len=*), intent(inout)  :: iomsg
  !Local
  character (len=18) :: pfmt
  write (pfmt, '(a,i2,a,i2,a)' ) &
    '(a,i', vlist(1), ',i', vlist(2), ',a,f8.3)'
  write (unit, fmt=pfmt, iostat=iostat,&
    iomsg=iomsg) "ball position = &
    ", this%x, this%y, " area =", &
    3.14 * this%r**2
end subroutine
```

# User Defined I/O Example

```
subroutine read_ball (this, unit, iotype, vlist,
iostat, iomsg)
  class(ball), intent(inout)           :: this
  integer, intent(in)                  :: unit
  character (len=*), intent(in)        :: iotype
  integer, intent(in)                  :: vlist(:)
  integer, intent(out)                 :: iostat
  character (len=*), intent(inout)     :: iomsg

  if (iotype == "DTradius") then
    read (unit, fmt="(f8.4)", iostat=iostat, &
          iomsg=iomsg) this%r
  else if (iotype == "DTid") then
    read (unit, fmt="(i6)", iostat=iostat, &
          iomsg=iomsg) this%id
  endif
end subroutine
```

# User Defined I/O Example

```
program main
  use ball_mod
  integer, target :: x = 4, y = 5
  type(ball) :: ball_foo
  ball_foo%x => x
  ball_foo%y => y
  write (*, '(a)') "ball radius="
  read (*, '(dt"radius")') ball_foo
  write (*, '(dt(2,2))') ball_foo
end program main
```

```
:~> ifort ball.f90
:~> ./a.out
ball radius=
2.d0
ball position = 4 5 area = 12.560
```