




# Intel® Advisor XE Future Release

Threading Design & Prototyping  
Vectorization Assistant



# Parallel is the Path Forward

Intel® Xeon® and Intel® Xeon Phi™ Product Families are both going parallel



	Intel® Xeon® processor 64-bit	Intel® Xeon® processor 5100 series	Intel® Xeon® processor 5500 series	Intel® Xeon® processor 5600 series	Intel® Xeon® processor code-named Sandy Bridge EP	Intel® Xeon® processor code-named Ivy Bridge EP	Intel® Xeon® processor code-named Haswell EP
Core(s)	1	2	4	6	8	12	18
Threads	2	2	8	12	16	24	36
SIMD Width	128	128	128	128	256	256	256

Intel® Xeon Phi™ coprocessor Knights Corner	Intel® Xeon Phi™ processor & coprocessor Knights Landing <sup>1</sup>
61	72
244	288
512	512

\*Product specification for launched and shipped products available on [ark.intel.com](http://ark.intel.com).

1. Not launched or in planning.

More cores → More Threads → Wider vectors


# Don't use a single Vector lane!



To fully utilize the hardware you need to:

- Parallelize and
- Vectorize

# How much potential lies untapped today?



**Parallel + Vectorized is much faster than either one alone**

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

# Permission to design for all lanes

## Threading and Vectorization



Intel Advisor XE:	Threading	Vectorization
Today	✓	
Future	✓	✓

# Data Driven Threading Design


## Intel® Advisor XE – Thread Prototyping

Have you:

- Tried threading an app, but seen little performance benefit?
- Hit a “scalability barrier”? Performance gains level off as you add cores?
- Delayed a release that adds threading because of synchronization errors?

Breakthrough for threading design:

- Quickly prototype multiple options
- Project scaling on larger systems
- Find synchronization errors before implementing threading
- Separate design and implementation -Design without disrupting development



Part of Intel® Parallel Studio  
For Windows\* and Linux\* From \$1,599

**“Intel® Advisor XE has allowed us to quickly prototype ideas for parallelism, saving developer time and effort”**

*Simon Hammond*  
Senior Technical Staff  
**Sandia National Laboratories**

**Add Parallelism with Less Effort, Less Risk and More Impact**

<http://intel.ly/advisor-xe>

Optimization Notice



# Data Driven Vectorization Design

Intel® Advisor XE – Vectorization Advisor (future release)

In Beta  
soon.  
Sign-up!

Have you:

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

Breakthrough for vectorization design

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

Function Call Sites and Loops	Self Time	Total Time	Memory analysis	Compiler Vectorization	Loop Type	Gain Estimate	Vectorized Loops
[loop at mmult_serial.cp...	10.040s	10.040s		Vectorized ...	Scalar	2.19727	SSE2
[loop at mmult_serial.cp...	0.000s	10.100s		Scalar			SSE2
[loop in __libc_start_mai...	0.000s	10.100s		Scalar			

Function Call Sites and Loops	Total Time %	Total Time	Self Time	Hot Loops	Vector Loops	Location
Total	100.0%	10.100s	0s			
libc_start_mai...	100.0%	10.100s	0s			libc-2.12.so
[loop in __libc_start_mai...	100.0%	10.100s	0s			libc-2.12.so
main	100.0%	10.100s	0s			mmult_serial
[loop at mmult_serial.cp...	100.0%	10.100s	0s			mmult_serial
[loop at mmult_serial.cp...	100.0%	10.100s	0s		SSE2	mmult_serial
multiply_d	100.0%	10.100s	0.0600s			mmult_serial
[loop ...	99.4%	10.040s	10.0400s		SSE2	mmult_serial

**More Performance**  
**Fewer Machine Dependencies**

# Vectorization Advisor

Providing the data you need for high impact vectorization

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

- Find “hot” un-vectorized or “under vectorized” loops.
- Convince the compiler to vectorize

Recommendations – How do I fix it?

Correctness via dependency analysis

- Is it safe to vectorize?
- Data on specific variable causing the loop dependency

Memory Access Patterns analysis

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



# Vector Advisor Survey: all in one place

Diagram illustrating the Vector Advisor Survey interface, showing various data sources and analysis results.

**Vector indicator (filter)**

**Profile data (priority)**

**Loop type (filter)**

**Compiler data**

**Static analysis**

**Function Call Sites and Loops**

Function Call Sites and Loops	Self Time	Total Time	Memory Analysis	Compiler Vectorization	Why No Ve...	Gain Estimate	Vectorized Loops	Vectorization Traits
[loop at nbbody.cc:57 in main]	1,820s	1,820s		<Expand to see ...	<Expand t ...	<Expand to s ...	AVX	Square Roots; Inserts; Extracts; Masked Sto
i> [loop at nbbody.cc:57 in main]	1,810s	1,810s		Vectorized (Body, Peeled)		2,00	AVX	Square Roots; Inserts; Extracts; Masked Stores
i> [loop at nbbody.cc:57 in main]	0,010s	0,010s		Scalar				
i> [loop at nbbody.cc:54 in main]	0,000s	1,820s		Scalar	inner loop ...		AVX	Shuffles; Inserts; Extracts
i> [loop at nbbody.cc:54 in main]	0,000s	1,820s		Scalar	inner loop ...			

**Top Down** **Source** **Loop Assembly** **Assistance** **Recommendations** **Compiler Diagnostic Details**

File: nbbody.cc:57 main

Line	Source	Total Time	%	Loop
52	void Newton( size_t n, real dt ) {			
53	const real dtG = dt * G;			
54	for ( size_t i = 0; i < n; ++i ) {			3 640,
55	real dvx = 0, dvy = 0, dvz = 0;			
56	//#pragma vector always			
57	for ( size_t j = 0; j < n; ++j ) {	10,110ms		3 640,
	[loop at nbbody.cc:57 in main]			
	Scalar loop. Not vectorized			
	No loop transformations were applied			
	[loop at nbbody.cc:57 in main]			
	Vectorized AVX loop processing Float32; Float64; Int32; UInt32 data t			
	No loop transformations were applied			

**Top-down function-loops tree**

**Source tab**

**Suggestions and OpenMP4 snippets**

**Loop body/peel/reminder break-down and grouping**

# What Makes Vectorization Difficult?

- Non-contiguous memory access – Potential to vectorize but may be inefficient

- Non-unit strided access to arrays

```
for (i=0;i<N;i+=2) //Incrementing "i" by 2 is not unit-stride
```

- Indirect reference in a loop

```
for (i=0;i<N;i++)
```

```
    A[B[i]] = C[i]*D[i]; //We have to decode B[i] to find out  
                        //which element of A to reference
```

- Data dependencies

```
for (i=0;i<N;i++)
```

```
    A[i] = A[i-1]*C[i];
```

# Compiler diagnostics + Performance Data

## Find “hot” un-vectorized or “under vectorized” loops

All of the information you require to vectorize available on one screen!

The screenshot displays the Intel Advisor XE 2016 interface, specifically the 'Threading and Vectorization Survey' tool. The interface is divided into several sections:

- Filter bar:** Located at the top, it includes tabs for 'Summary', 'Survey Report', 'Correctness Report', 'Annotation Report', and 'Suitability Report'. Below these are filters for 'Filter by Loop Type' (Vectorized, Not Vectorized) and 'Filter by Source' (All).
- Flat loops list:** A table listing function call sites and loops. It includes columns for 'Self Time' and 'Total Time'. A callout points to the first row: '[loop at loopstl.cpp:1069 in s127\_]'. Below this, a 'Loop summary' section provides details for the selected loop, including 'Vectorized AVX loop processing Float32; Float64 data type(s) having Permute; Inserts; Extracts operations' and 'No loop transformations were applied'.
- Compiler Vectorization:** A section on the right showing 'Loop Type', 'Why No Vectorization?', 'Gain ...', and 'Vectorized Loops'. It includes expandable sections for 'Vectorized (Body)' and 'Remainder'.
- Source and assembly views:** A section at the bottom left showing a 'Top-down function-loops tree' with a tree view of function call sites. A callout points to the 'Source and assembly views' tab.
- Loop details:** A callout box on the right lists the information available for each loop: Self and total time, Compiler VEC/OPT report, Compiler gain estimate, Vectorization traits, and Vector instruction set.

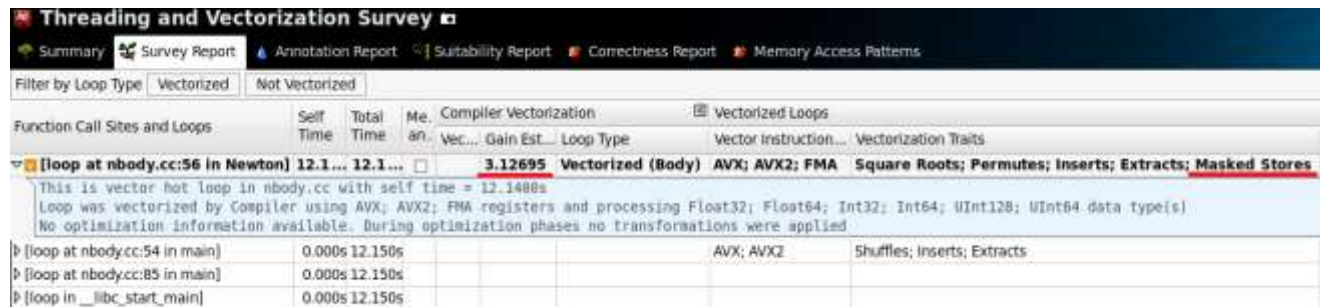
Function Call Sites and Loops	Self Time	Total Time
[loop at loopstl.cpp:1069 in s127_]	0.040s	0.040s
[loop at loopstl.cpp:1069 in s127_]	0.030s	0.030s
[loop at loopstl.cpp:1069 in s127_]	0.010s	0.010s
[loop at mains.F:594 in SET2D]	0.030s	0.030s
[loop at loopstl.cpp:275 in s112_]	0.030s	0.030s

Loop Type	Why No Vectorization?	Gain ...	Vectorized Loops
Vectorized (Body)	vectorization support ...	3.88	AVX
Remainder			
Vectorized (Body)		3.00	SSE; SSE2
Vectorized (Body)	vectorization support ...	1.51	AVX

Function Call Sites	Total Time %	Total Time	Self Time	Compiler VEC	Loop Type
Total	100.0%	1.951s	0s		
RtlUserThreadStart	100.0%	1.951s	0s		
BaseThreadInitThunk	100.0%	1.951s	0s		
_tmainCRTStartup	99.5%	1.941s	0s		
main	99.5%	1.941s	0s		
UNNAMED_MAIN	99.5%	1.941s	0s		
[loop at mains.F:157 in UNNAMED_MAIN]	99.5%	1.941s	0s		Scalar
	9.2%	0.180s	0s		
	0.7%	0.140s	0s		

# Gives estimated expected gain!

Gain estimates – Gives recommendations and the gain you can expect by using a different vector instruction or rewriting the control flow of your program.

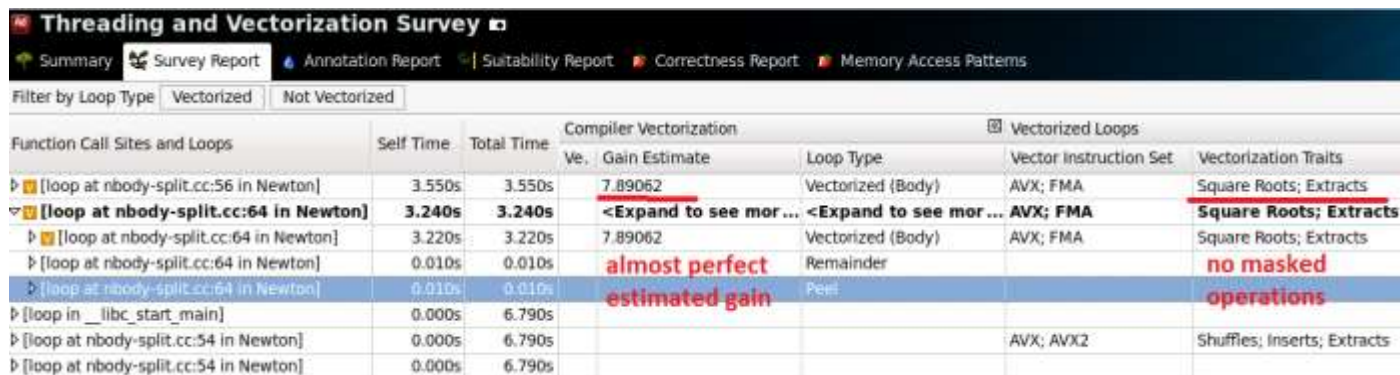


**Threading and Vectorization Survey**

Summary | **Survey Report** | Annotation Report | Suitability Report | Correctness Report | Memory Access Patterns

Filter by Loop Type: Vectorized | Not Vectorized

Function Call Sites and Loops	Self Time	Total Time	Me. an. Vec...	Compiler Vectorization	Vectorized Loops	Vector Instruction...	Vectorization Traits
<b>[loop at nbody.cc:56 in Newton]</b>	<b>12.1...</b>	<b>12.1...</b>		<b>3.12695</b>	<b>Vectorized (Body)</b>	<b>AVX; AVX2; FMA</b>	<b>Square Roots; Permutes; Inserts; Extracts; Masked Stores</b>
This is vector hot loop in nbody.cc with self time = 12.1480s Loop was vectorized by Compiler using AVX; AVX2; FMA registers and processing Float32; Float64; Int32; Int64; UInt32; UInt64 data type(s). No optimization information available. During optimization phases no transformations were applied.							
↳ [loop at nbody.cc:54 in main]	0.000s	12.150s				AVX; AVX2	Shuffles; Inserts; Extracts
↳ [loop at nbody.cc:85 in main]	0.000s	12.150s					
↳ [loop in __libc_start_main]	0.000s	12.150s					



**Threading and Vectorization Survey**

Summary | **Survey Report** | Annotation Report | Suitability Report | Correctness Report | Memory Access Patterns

Filter by Loop Type: Vectorized | Not Vectorized

Function Call Sites and Loops	Self Time	Total Time	Ve.	Gain Estimate	Loop Type	Vector Instruction Set	Vectorization Traits
↳ [loop at nbody-split.cc:56 in Newton]	3.550s	3.550s		7.89062	Vectorized (Body)	AVX; FMA	Square Roots; Extracts
<b>↳ [loop at nbody-split.cc:64 in Newton]</b>	<b>3.240s</b>	<b>3.240s</b>		<b>&lt;Expand to see mor ...</b>	<b>&lt;Expand to see mor ...</b>	<b>AVX; FMA</b>	<b>Square Roots; Extracts</b>
↳ [loop at nbody-split.cc:64 in Newton]	3.220s	3.220s		7.89062	Vectorized (Body)	AVX; FMA	Square Roots; Extracts
↳ [loop at nbody-split.cc:64 in Newton]	0.010s	0.010s		<b>almost perfect</b>	Remainder		<b>no masked</b>
↳ [loop at nbody-split.cc:64 in Newton]	0.010s	0.010s		<b>estimated gain</b>	Peel		<b>operations</b>
↳ [loop in __libc_start_main]	0.000s	6.790s					
↳ [loop at nbody-split.cc:54 in Newton]	0.000s	6.790s				AVX; AVX2	Shuffles; Inserts; Extracts
↳ [loop at nbody-split.cc:54 in Newton]	0.000s	6.790s					

# Vector Advisor:

- All the data in one place
  - Intel Compiler 15.x reports Integration
- Deep dive analysis

The screenshot displays the Intel Vector Advisor interface. At the top, a table lists function call sites and loops with their self and total times, and compiler vectorization details. Below this, a section titled 'Issue: Compile time constraints prevent loop optimization' explains the cause: internal time limits for the -O2 optimization level. A 'Recommendation' section provides further details. The bottom part of the image shows a detailed view of a specific loop (loop\_site\_139) with its memory access patterns, including stride information and access patterns.

Function Call Sites and Loops	Self Time	Total Time	Compiler Vectorization	Vectorized Loops
[loop at nbody.cc:22 in main]	1,864s	1,864s	Vectorized (Body)	Gain Estimate: 5.69, Vect.: AVX, Vectorization Tra...: Square Roots; Ins..., Vector Widths: 128/256, Vector ...: Float32; ...
[loop at nbody.cc:16 in main]	0,000s	1,864s	Scalar	inner loop was already vectorized, Vect.: AVX, Vectorization Tra...: Shuffles; Inserts; ..., Vector Widths: 128/256, Vector ...: Float32; ...
[loop at nbody.cc:97 in main]	0,000s	1,864s	Scalar	compile time constraints prevent ..., Vect.: AVX, Vectorization Tra...: 128/256, Vector ...: Float32; ...

**Issue: Compile time constraints prevent loop optimization**  
**Cause:** Internal time limits for the -O2 (Windows\* OS) or -O2 (Linux\* OS) optimization level prevented the compiler from determining a vectorization approach for this loop

**Recommendation**

Site Name	Site Function	Site Info	Loop-Carried Dependencies	Strides Distribution	Access Pattern
loop_site_203	runCRawLoops	runCRawLoops.cxx:1063	RAW:1	No information available	No information available
loop_site_139	runCRawLoops	runCRawLoops.cxx:622	No information available	39% / 36% / 25%	Mixed strides
loop_site_160	runCRawLoops	runCRawLoops.cxx:925	No information available	100% / 0% / 0%	All unit strides

**Memory Access Patterns**

ID	Stride	Type	Source	Modules	Alignment
P22	0; 0; 1	Unit stride	runCRawLoops.cxx:637	lcalc.exe	
P23	0; 0	Unit stride	runCRawLoops.cxx:638	lcalc.exe	
P24	0; 0	Unit stride	runCRawLoops.cxx:639	lcalc.exe	
P25	0; 0; 0	Unit stride	runCRawLoops.cxx:640	lcalc.exe	
P30	-1575; -63; -26; -25; -1; 0; 1; 25; 26; 63; 2164801	Variable stride	runCRawLoops.cxx:628	lcalc.exe	

# Convince the compiler to vectorize

Unvectorized loops / “under vectorized” loops

- Assumed dependencies
- Control structures preventing vectorization.
- Rewrite loops to vectorize – remove conditions, breaks and returns and many other techniques.

**Threading and Vectorization Survey** Intel Advisor XE 2016

Summary Survey Report Annotation Report Suitability Report Correctness Report Memory Access Patterns

Filter by Loop Type Vectorized Not Vectorized Filter by Source All Filter by Module All

Function Call Sites and Loops	Self Time	Total Time	Memory analysis	Compiler Vectorization	Compiler Optimization		
				Loop Type	Vectorization Message(s)	Gain Estimate	
▷ [loop at mmult_seri ...]		10.100s		Scalar			
▷ [loop at mmult_serial.cp ...]		10.100s		Scalar			
▷ [loop in __libc_start_mai ...]		10.100s		Scalar			

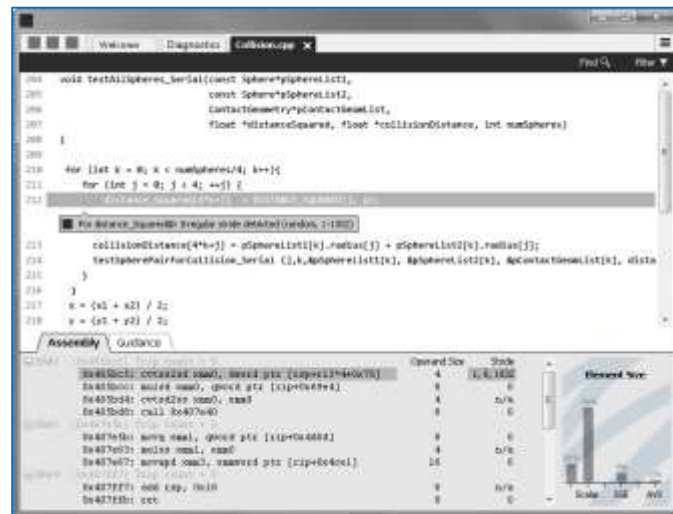
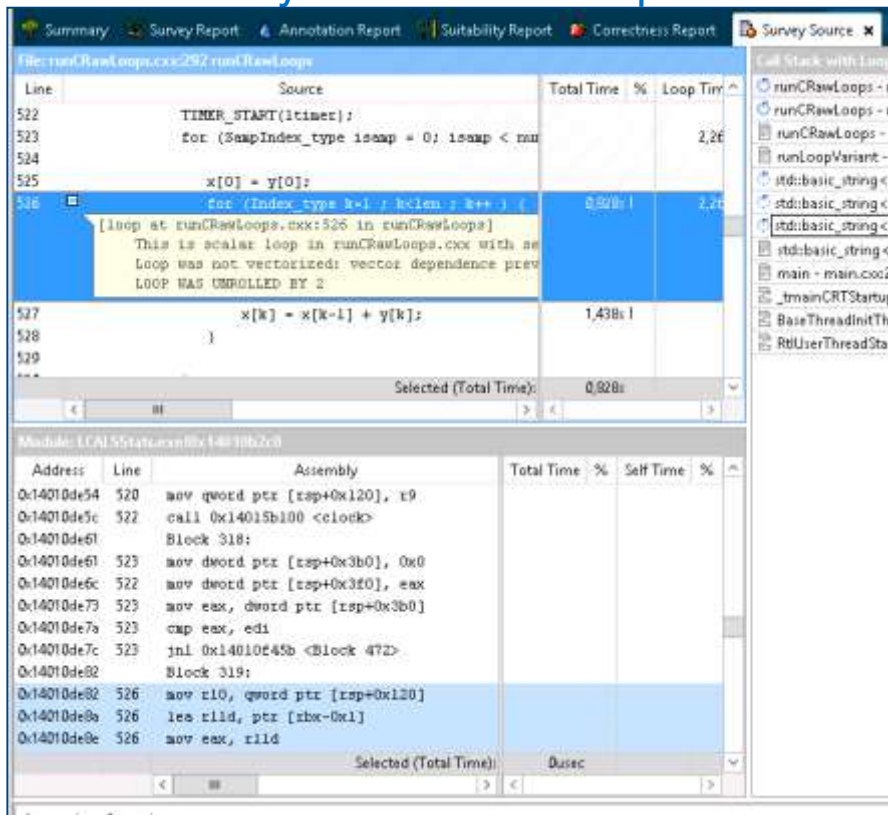
Top Down Source Assembly Assistance Recommendations

Function Call Sites and Loops	Total Time %	Total Time	Self Time	Hot Loops	Vector... Loops	Location
						Source Loc... Module
▽ Total	100.0%	10.100s	0s			
▽ __libc_start_main	100.0%	10.100s	0s			libc-2.1 ...
▽ [loop in __libc_start_main]	100.0%	10.100s	0s			libc-2.1 ...
▽ main	100.0%	10.100s	0s			mmult_seri ... 1_mmul ...
▽ [loop at mmult_serial.cpp:101 in main]	100.0%	10.100s	0s			mmult_seri ... 1_mmul ...
▽ [loop at mmult_serial.cpp:52 in mai ...]	100.0%	10.100s	0s		SSE2	mmult_seri ... 1_mmul ...
▽ multiply_d	100.0%	10.100s	0.0600s			mmult_seri ... 1_mmul ...
▷ [loop at mmult_serial.cpp:54 i ...]	99.4%	10.040s	10.0400s		Lo. SSE2	mmult_seri ... 1_mmul ...



# Deep source and assembly integration

## All the data you need in one place



# Recommendations – How do I fix it?




**SIZE:**

64B for Intel® Xeon Phi™,

32B for AVX1/2,

16B for SSE4.2 and below

## Alignment optimization



## Peel/remainder

- Typical vectorized loop consists of
  - Optional “peel” part
    - Needed to improve alignment
    - Scalar or slower vector
  - Main vector part
    - Fastest among the three.
  - “remainder” part
    - Due to  $\text{trip\_count} \% \text{VL} \neq 0$
    - Scalar or slower vector.
- Larger vector register means more iterations in peel/remainder
  - Align your data
  - Block to fight remainders



# Recommendations

Summary Survey Report Refinement Reports Annotation Report Suitability Report

Filter by Loop Type: Vectorized Not Vectorized Mark-up interesting loops to move on Filter by Source: All Filter by Module: All

Function Call Sites and Loops	Self Time	Total Time	Compiler Vectorization	Vectorized Loops	Vector Instruction Set	Vectorization Traits	Vector Widths	Vector Data Types	Compiler O
[loop in ru...	0.779s	1.380s	<input type="checkbox"/>	Loop was vectorized	AVX	Divisions; Square Roots; Inserts; Extracts	128/256	Float32; Float64	Reordered
[loop at runC...	0.770s	0.770s	<input type="checkbox"/>	vector dependence: assumed dependence between lines					Unrolled
[loop in ru...	0.769s	1.409s	<input type="checkbox"/>	loop was vectorized	AVX	Divisions; Square Roots; Inserts; Extracts	128/256	Float32; Float64	Reordered
[loop at ru...	0.738s	0.738s	<input type="checkbox"/>	<Expand to see more ...>	AVX	Type Conversions; Inserts; Extracts	128/256; 128	Float64	Unrolled
[loop in A...	0.660s	0.660s	<input type="checkbox"/>	loop was vectorized	AVX	Inserts	128/256	Float64	Unrolled
[loop at ru...	0.640s	0.640s	<input type="checkbox"/>	<Expand to see more ...>	AVX	Type Conversions; Inserts; Extracts	128/256; 128	Float64; Int32	Unrolled
[loop in runA...	0.609s	2.039s	<input type="checkbox"/>		AVX	Inserts	128/256	Float32; Float64	Reordered
[loop in runA...	0.589s	1.989s	<input type="checkbox"/>		AVX	Inserts	128/256	Float32; Float64	Reordered
[loop in A...	0.569s	0.569s	<input type="checkbox"/>	Loop was vectorized	AVX	Inserts	128/256	Float64	Unrolled
[loop at runC...	0.510s	3.260s	<input type="checkbox"/>	inner loop was already vectorized	AVX	Inserts; Extracts	128/256	Float64	

Top Down Source Loop Assembly Assembly Recommendations Compiler Generated Data

**Issues: 1**  
**Recommendations: 2**

**Issue: Ineffective Peeled/Remainder loop(s) present**

All or some source loop iterations are not executing in the loop body. Improve performance by moving source loop iterations from peeled/remainder loops to the loop body. Read more at [Glossary and Vector Essentials, Utilizing Full Vectors...](#)

**Use a smaller vector length**

Potential performance gain: Information not available until Beta  
Confidence this recommendation applies to your code: Information not available until Beta

The compiler chose a vector length, but the trip count might be smaller than that vector length. To fix: Identify a smaller vector length using a directive. For C/C++: `#pragma simd vectorlength(N)` with N less than actual trip count. For Fortran: `!DIRS SIMD VECTORLENGTH=N`.

**Specify the expected loop trip count**

Potential performance gain: Information not available until Beta  
Confidence this recommendation applies to your code: Information not available until Beta

The compiler cannot statically detect the trip count. To fix: Specify the expected number of iterations using a directive. For C/C++: `#pragma loop_count(N)`. For Fortran: `!DIRS LOOP_COUNT=N`. Read more at [User and Reference Guide for the Intel C++ Compiler 15.0 > Compiler Reference > Pragmas > Intel-specific Pragma Reference](#).

# End-user recommendations, performance penalties

Function Call Sites and Loops	Memory Analysis	
[loop at runBRowLoops.cxx:118 in runBRowLoops]	<input type="checkbox"/>	
[loop at runBRowLoops.cxx:82 in runBRowLoops]	<input type="checkbox"/>	
[loop at runCRowLoops.cxx:292 in runCRowLoops]	<input type="checkbox"/>	
[loop at runBRowLoops.cxx:55 in runBRowLoops]	<input type="checkbox"/>	
[loop at runOMPRawLoops.cxx:648 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	💡 2
[loop at runOMPRawLoops.cxx:679 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	
[loop at runOMPRawLoops.cxx:203 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	💡 1
[loop at runOMPRawLoops.cxx:187 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	
⚠️ [loop at runOMPRawLoops.cxx:173 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	💡 2
[loop at runOMPRawLoops.cxx:139 in runOMPRawLoops\$omp ...]	<input type="checkbox"/>	💡 1
[loop at runOMPRawLoops.cxx:86 in runOMPRawLoops\$omp\$...]	<input type="checkbox"/>	💡 2
[loop at runCRowLoops.cxx:713 in runCRowLoops]	<input type="checkbox"/>	💡 1
[loop at runOMPForAllLambdaLoops.cxx:126 in forall]	<input type="checkbox"/>	
[loop at runBForAllLambdaLoops.cxx:112 in forall]	<input type="checkbox"/>	
[loop at complex:617 in runOMPForAllLambdaLoops\$omp\$par ...]	<input type="checkbox"/>	💡 1
[loop at runCRowLoops.cxx:1019 in runCRowLoops]	<input type="checkbox"/>	

2  
8

## Issue: Peeled/Remainder loop(s) present

All or some source loop iterations are not executing in the kernel loop. Improve performance by moving source loop iterations from peeled/remainder loops to the kernel loop. Read more at [Vector Essentials](#), [Utilizing Full Vectors...](#)

### Recommendation: Align memory access

Projected maximum performance gain: High

Projection confidence: Medium

The compiler created a peeled loop because one of the memory accesses in the source loop does not start at a data boundary. Align the memory access and tell the compiler your memory access is aligned. This example aligns memory using a 32-byte boundary:

```
float *array;
array = (float *)_mm_malloc(ARRAY_SIZE*sizeof(float), 32);

// Somewhere else
__assume_aligned(array, 32);
// Use array in loop

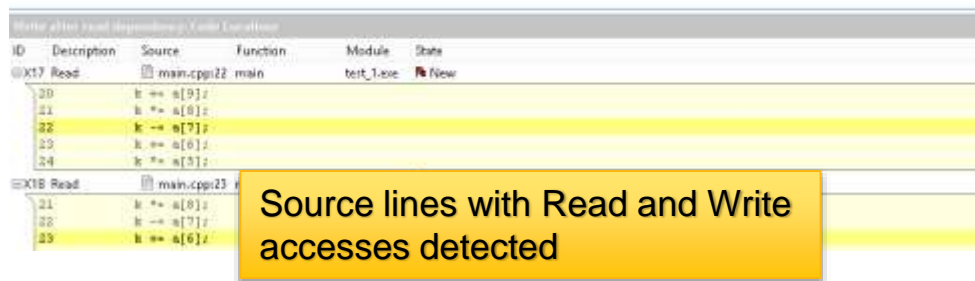
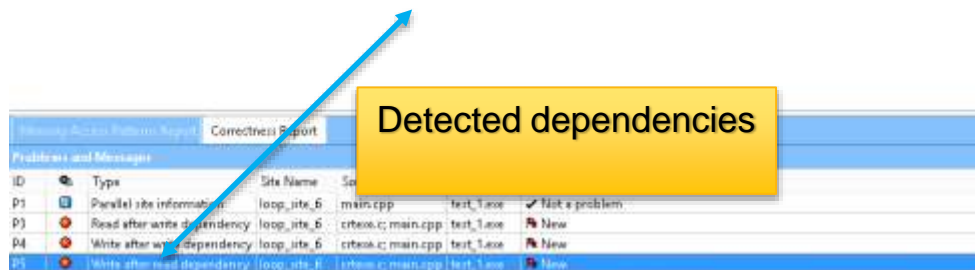
_mm_free(array);
```

Alternative: Declare a static aligned array using `__declspec(align(32)) float array[ARRAY_SIZE];` and use `__assume_aligned(array, 32);` before your loop.

### Recommendation: Use a smaller vector length

# Correctness – Is It Safe to Vectorize?

## Loop-carried dependencies analysis



Got recommendations to enforce vectorization of the loop:

1. Mark-up the loop and check for the presence of REAL dependencies
2. Explore dependencies in more details with code snippets

Are there dependencies in your loop preventing vectorization?

*(if you force the compiler to vectorize this could generate incorrect code)*


# Memory Access Patterns – Data Layout Is Key

## Vectorizing i-loop


Private:  $v$ ,  $c.q$ ,  $a[j]$ ,  $a[j][k]$ ,  $a[b[j]]$  ( $j \neq i$ ,  $k \neq i$ )



Unit-stride  $a[i]$ ,  $c.x[i]$ ,  $*(p++)$



Non-unit-stride:  $a[2*i]$ ,  $c[i].x$ ,  $a[i][j]$ ,  $a[i][0]$   
F90 Arrays in most cases  
if not “contiguous”



Gather/scatter:  $j = b[i]$ ;  $a[j]$ ,  $a[b.x[i]]$   
 $p = a[j]$ ;  $*p$  – Intel® IMCI: especially slow




- Private – good
  - Almost no alignment requirements
  - Any addressing if...
    - Not depend on vectorizable loop index
- Unit-Stride
  - Good – with one exception.
  - Subject to vector alignment
  - Out-of-order cores won't store-forward  
**masked (unit-stride) store.** ☹
  - On Intel® Xeon Phi™ correctness prevents efficient implementation of  
**masked (unit-stride) store**
- Strided, Gather/Scatter is less efficient
  - Perf varies on micro-arch and the actual index patterns.
  - Big latency is exposed if you have these on the critical path
  - Better if done at outer loop level if loop nest is vectorized

# Improve Vectorization

## Memory Access pattern analysis

### Mark-up loops for deeper analysis...



Site Name	Site Function	Site Info	Loop-Carried Dependencies	Strides Distribution	Access Pattern
loop_site_6	s114	loopstl.cpp:394	No information available	59% / 41%	Mixed strides
loop_site_7	s116	xfunctional:51	No information available	73% / 27%	Mixed strides
loop_site_8	s118	loopstl.cpp:574	No information available	79% / 21%	Mixed strides
loop_site_10	s123	loopstl.cpp:814	No information available	75% / 25%	Mixed strides
loop_site_12	s124	loopstl.cpp:870	No information available	75% / 25%	Mixed strides

**Unit stride vs non-unit stride accesses**

ID	Stride	Type	Source	Modules	Alignment
P59		Misaligned access	xfunctional:51	lcd_coi.exe	4; 4; 4
<pre> 49  _Ty operator()(const _Tys _Left, const _Tys _Right) const 50  { // apply operator* to operands 51      return (_Left * _Right); 52  } 53  };                     </pre> <p><b>Access to unaligned memory</b></p>					
P62	2	Non-unit stride	algorithm:5137	lcd_coi.exe	
<pre> 5135  if (_First != _Last) 5136      for (; ++_First != _Last; 5137          if (_DEBUG_LT(*_Found, *_First)) 5138              _Found = _First; 5139      return (_Found);                     </pre> <p><b>Stride is 2</b></p>					
P63	2	Non-unit stride	algorithm:5180	lcd_coi.exe	
P64	2	Non-unit stride	algorithm:5181	lcd_coi.exe	
P1	-6	Non-unit stride	loopstl.cpp:5001	lcd_coi.exe	
P5	0	Unit stride	loopstl.cpp:4950	lcd_coi.exe	
P5	0	Unit stride	loopstl.cpp:4964	lcd_coi.exe	
P5	0	Unit stride	loopstl.cpp:4992	lcd_coi.exe	
P5	0	Unit stride	loopstl.cpp:4996	lcd_coi.exe	
P5	0	Unit stride	loopstl.cpp:4997	lcd_coi.exe	
<pre> 4995  index = 1; 4996  max = ABS (a[1]); 4997  k += *inc; 4998  i_2 = *n; 4999  for (i_2 = 2; i_2 &lt;= i_2; ++i_2)                     </pre> <p><b>Stride 0 – writing to the same memory</b></p>					

ID	Stride	Type	Source	Modules
P3	2; 2	Non-unit stride	runARawLoops.coo427	LCALStats.exe
<pre> 425 426  Complex_type a0t = t0[it0+1]; 427  Complex_type a1t = t1[it0+1]; 428  Complex_type a2t = t2[it0+1] * fratio; 429                     </pre>				
P4	2; 2	Non-unit stride	runARawLoops.coo438	LCALStats.exe
<pre> 436  /* compute new A0 */ 437  Complex_type z3 = ( c1 * a1t + z3 * a2t ) * a1a1t ; 438  t0[it0+1] = a0t * c1a1t - ical * z3; 439 440  /* compute new A1 */                     </pre>				

# Vectorization Advisor

Providing the data you need for high impact vectorization

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

- Find “hot” un-vectorized or “under vectorized” loops.
- Convince the compiler to vectorize

Recommendations – How do I fix it?

Correctness via dependency analysis

- Is it safe to vectorize?
- Data on specific variable causing the loop dependency

Memory Access Patterns analysis

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



# Summary: Vector Advisor Alpha

## 4 Analysis Features for Efficient Vectorization

### 1. Compiler diagnostics with Performance Data

Function Call Site and Loops	Self Time	Total Time	Compiler Vectorization
[loop in runCFordLambdai.loops]	0.094s	0.094s	Scalar
[loop in runCFordLambdai.loops]	0.140s	0.140s	Scalar
[loop in std::complex_base::double_complex::...	0.031s	0.031s	Vectorized (Body)
Vectorized SSE2 SSE3 loop processing Float64 data type(s) having Divisions/ Square Roots operations Peel loop: loop state was disabled			
[loop in std::basic_string<char,traits,allocator>::...	0.000s	0.000s	Scalar
[loop in std::basic_string<char,traits,allocator>::...	0.000s	0.000s	Scalar
[loop in std::num_get<char,class std::istreambuf_iterator<char,traits>::...	0.000s	0.000s	Scalar

### 2. Recommendations on how to improve vectorization

**Issue: Peeled/Remainder loop(s) present**

All or some source loop iterations are not executing in the kernel loop. Improve performance by moving source loop iterations from peeled/remainder loops to the kernel loop. Read more at [Vector Essentials](#), [Utilizing Full Vectors](#).

**Recommendation: Align memory access**

Projected maximum performance gain: High  
Projection confidence: Medium

The compiler created a peeled loop because one of the memory accesses in the source loop does not start at a data boundary. Align the memory access and tell the compiler your memory access is aligned. This example aligns memory using a 32-byte boundary:

```
float *array;  
array = (float *)__builtin_malloc(ARRAY_SIZE*sizeof(float), 32);  
  
// Somewhere else  
__builtin_aligned(array, 32);  
// Use array in loop
```

### 3. Correctness Dependency Analysis

#### Problems and Messages

ID	Type	Site Name	Sources	Modules	State
P1	Parallel site information	site2	dqtest2.cpp	dqtest2	✓ Not a problem
P2	Read after write dependency	site2	dqtest2.cpp	dqtest2	New
P3	Read after write dependency	site2	dqtest2.cpp	dqtest2	New
P4	Write after write dependency	site2	dqtest2.cpp	dqtest2	New
P5	Write after write dependency	site2	dqtest2.cpp	dqtest2	New
P6	Write after read dependency	site2	dqtest2.cpp	dqtest2	New
P7	Write after read dependency	site2	dqtest2.cpp; idle.h	dqtest2	New




### 4. Memory Access Patterns Analysis

Site Name	Site Function	Site Info	Loop-Carried Dependencies	Strides Distribution	Access Pattern
loop_site_303	runCFordLoops	runCFordLoops.coc363	RAW:1	No information available	No information available
loop_site_339	runCFordLoops	runCFordLoops.coc622	No information available	88% / 36% / 2%	Mixed strides
loop_site_360	runCFordLoops	runCFordLoops.coc625	No information available	100% / 0% / 0%	All unit strides

ID	Stride	Type	Source	Modules	Alignment
P22	0: 0:1	Unit stride	runCFordLoops.coc637	icalls.exe	
P23	0: 0	Unit stride	runCFordLoops.coc638	icalls.exe	
P30	-1575: -63: -26: -25: -1: 0: 1: 25: 26: 63: 2164801	Variable stride	runCFordLoops.coc628	icalls.exe	

# Intel® Advisor XE is part of Intel® Parallel Studio XE

 <p>Intel® Parallel Studio XE 2015 Composer Edition</p>	 <p>Intel® Parallel Studio XE 2015 Professional Edition</p>	 <p>Intel® Parallel Studio XE 2015 Cluster Edition</p>
<p>Intel® C++ Compiler Intel® Fortran Compiler Intel® Threading Building Blocks Intel® Integrated Performance Primitives Intel® Math Kernel Library Intel® Cilk™ Plus Intel® OpenMP*</p>	<p>Intel® C++ Compiler Intel® Fortran Compiler Intel® Threading Building Blocks Intel® Integrated Performance Primitives Intel® Math Kernel Library Intel® Cilk™ Plus Intel® OpenMP*</p>	<p>Intel® C++ Compiler Intel® Fortran Compiler Intel® Threading Building Blocks Intel® Integrated Performance Primitives Intel® Math Kernel Library Intel® Cilk™ Plus Intel® OpenMP*</p>
	<p><b>Intel® Advisor XE</b> Intel® Inspector XE Intel® VTune™ Amplifier XE</p>	<p><b>Intel® Advisor XE</b> Intel® Inspector XE Intel® VTune™ Amplifier XE Intel® MPI Library Intel® Trace Analyzer and Collector</p>

For more information: <http://intel.ly/perf-tools>



# Join the beta!

## Intel® Advisor XE – Vectorization Advisor

Send e-mail to [vector\\_advisor@intel.com](mailto:vector_advisor@intel.com) to participate in the Vectorization Advisor beta.

Limited alpha access is available now under NDA

Public beta is coming late Q1 or Q2 2015

Sign-up now and we will contact you when we have more details.



# Legal Disclaimer & Optimization Notice

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

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Copyright © 2014, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

## Optimization Notice

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

Notice revision #20110804