

### Explicit Vectorization with OpenMP\* 4.0

This lab introduces explicit vectorization techniques, using the SIMD-focused features introduced in OpenMP 4.0. We follow a simple process:

1. Identify hotspots that can benefit from SIMD instructions.
2. Select a sufficiently wide data type, adjust floating-point accuracy, and consider mixed precision depending on algorithm phase.
3. Re-arrange data for SIMD efficiency.
4. Convert code to SIMD form.
5. Align data structures for SIMD efficiency.
6. Optimize memory access patterns.
7. Further (instruction-level) optimization.

This process can be applied to any code, and is not specific to OpenMP 4.0 – the conversion of the code to SIMD form could also make use of a compiler's auto-vectorization capabilities or array notation.

This lab focuses in particular on the SIMD features that target, for example Intel® Xeon Phi™ processors. This material uses the N-body problem to illustrate the effect of the SIMD features. For the matter of simplicity and to stay on topic, the presented code is not meant to be sophisticated as for example using an acceleration data structure i.e., the problem's complexity is straight forward  $O(N^2)$ .

Background: an **N-body simulation** is a simulation of a dynamical system of particles, usually under the influence of physical forces, such as gravity. In cosmology, they are used to study processes of non-linear structure formation such as the process of forming galaxy filaments and galaxy halos from dark matter in physical cosmology. Direct N-body simulations are used to study the dynamical evolution of star clusters.

## Building the Application

Decide whether you would like to work with GNU Fortran (`gfortran`) or Intel® Fortran compiler (`ifort`). The techniques covered in this lab are applicable to both, so you should choose whichever compiler you are most familiar with and have access to.

You can build the `nbody` application by using the Intel® toolchain with

```
make COMP=intel
```

or with the GNU toolchain with

```
make COMP=gnu
```

For both `gfortran` and `ifort` the optimization (and vectorization) report will be redirected to the file `nbody.optrpt`. Please notice that the structure of the report is dependent on the compiler and compiler version used!

## Configuring and Running the Application

Please be patient when doing the first runs. Because of the included warm-up calculations and the deactivated vectorization the total execution time will be in the range of 1.5-2.5 minutes.

Before running the application, be sure to set the OpenMP thread affinity by using the portable OpenMP 4.0 thread affinity syntax as below. The optional **verbose** setting available only with the Intel® OpenMP library will show the binding of threads to logical processors:

```
export OMP_PLACES=threads
export OMP_PROC_BIND=close
[export KMP_AFFINITY=verbose] # Intel OpenMP only
```

Then, to execute the application, run (for example):

```
./nbody 262144
```

Run the application now, and record the checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

### (Optional) Hotspots Analysis

If you have access to Intel® VTune™ Amplifier or any similar profiling tool such as **gprof**, you can try to profiling the application to see if you are able to find the application hotspots. To save time, we now jump right to the profiling results and conclude that the most time consuming routine of the application is **perform\_nbody()**. This is the routine where all the computation takes place and which will be chosen as an optimization target.

## Selecting Floating-Point Precision

Many applications use double precision by default, even though single precision may be adequate. Our application has been configured to allow the floating-point precision to be changed at compile time, allowing us to investigate the impact of this decision on performance and accuracy. Locate the statements within the `BodyData` -module at the top of the main file, and change the precision from double to single:

```
INTEGER, PARAMETER :: fptype = real32
```

Build and re-run the application, then record the new checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

Do the checksums match? How has the run time changed?

The impact of floating-point precision can be significant: double precision numbers are twice as large (in bytes) as single precision numbers, affecting both memory bandwidth requirements and the number of values that can be packed into SIMD units of a fixed bit-width. A speed-up of 2x is not unusual.

### **IMPORTANT:**

Before changing the floating-point precision of your own applications, you should analyze the numerical stability of your algorithm and assess the correctness of this optimization. In many cases, especially if the problem is somewhat ill-conditioned, double precision accuracy **is** needed to keep the computation numerically stable.

## Explicit Vectorization with OpenMP\* 4.0

### Enable compiler autovectorization

Edit the **Makefile** and comment the variable holding the **-no-vec** (for **ifort**) or **-no-ftree-vectorize** (for **gfortran**) flag and replace it with whatever is the highest performing vector ISA available on your system. This will enable the auto-vectorization by the compiler. Now re-build the application and check the vectorization report in **nbody.optrpt**.

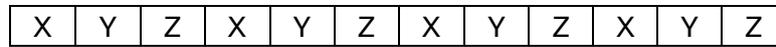
Re-run the application now, and record the checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

How much did the vectorization improve the performance; is this what you expected?

### Re-arrange Data for SIMD Efficiency

The code uses an Array-of-Structures (AoS) to store its position and acceleration data – that is, the Position and Acceleration arrays store the 3D quantities contiguously in memory, as shown below:



Do you think this layout is efficient for scalar execution? What about for SIMD execution?

To help with our investigation, our application produces copies of the input data arrays in both AoS and SoA formats (e.g. Position is split into `Position_X`, `Position_Y` and `Position_Z`). If the  $j$  loop in `perform_nbody()` is vectorized, which form of the arrays will be better?

Change the references to Position and/or Acceleration in `perform_nbody()`. If you do decide to make any changes to the Acceleration array, remember to update the `Checking()` function as well!

Build and re-run the application, then record the new checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

How has the run time changed?

## Explicit Vectorization with OpenMP\* 4.0

If the application runs slower, don't panic! The best data structure for scalar and SIMD execution are not always the same. As a rule of thumb:

### Data Accessed by Scalar Code

- Use AoS to improve cache locality.

### Data Accessed by SIMD Code with Linear Index

- SIMD lanes will access contiguous data elements.
- Use SoA (or AoSoA), to avoid gathers/scatters.

### Data Accessed by SIMD Code via Indirection

- SIMD lanes may access non-contiguous data elements.
- Use a new algorithm (to avoid indirection).

**OR**

- Use AoS to improve cache locality of gathers/scatters.

Applying this rule to our application (assuming that the  $j$  loop will be vectorized), only the layout of the Position array needs to be adjusted for SIMD efficiency – the Acceleration array is only accessed by scalar code, and can stay in AoS.

## Convert Code to SIMD Form

We are ready to convert our application to SIMD form. The OpenMP\* 4.0 standard supports two different ways to do this:

1. SIMD Loops  
Specify loops that should be vectorized by the compiler.
2. SIMD Functions  
Specify functions that will be called by vectorized loops.

The first of these methods is most applicable here, so go ahead and mark the *j* loop in `perform_nbody()` as being a SIMD loop. As a reminder, the syntax is as follows:

For C:

```
#pragma omp simd [clause[, clause ...]]
```

For Fortran:

```
!$omp simd [clause[, clause ...]]
```

where clause can be any of the following:

<code>safelen(length)</code>	Maximum distance between two iterations executed concurrently by a SIMD instruction.
<code>linear(list[:linear-step])</code>	List items are private and have a linear relationship with respect to the iteration space.
<code>aligned(list[:alignment])</code>	List items are aligned to a platform-dependent value (or the value of the optional parameter).

`private(list)`, `lastprivate(list)`, `reduction(reduction-identifier:list)` and `collapse(n)` are also supported, with functionality matching that of `omp for`.

Build and re-run the application, then record the new checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

Do the checksums match? How has the run time changed?

If your code produces different answers to what you expected, then you may have forgotten a clause. Repeat this exercise until you are happy with the checksum.

**Hint:** Make sure that your variables are declared as `private/reduction` where appropriate.

## Align Data Structures (Intel® Fortran only)

There are three steps to ensuring aligned data accesses in your SIMD applications:

1. Align memory.
2. Access memory in an aligned way.
3. Tell the compiler.

### Align Memory

Look at the *j* loop in `perform_nbody()`. Which arrays do you think need to be aligned?

When you have decided, update the definitions of these arrays to align them to a 64-byte boundary. We use a 64-byte boundary here because modern Intel® Xeon® processors have a 512-bit SIMD width ( $512 / 8 = 64$ ); 16- and 32-byte boundaries are sufficient when aligning arrays for the 128- and 256-bit instructions in SSE and AVX, respectively.

```
real(kind=fptype), allocatable :: array(:)
!dir$ attributes align:64 :: array
```

### Access Memory in an Aligned Way

An array will be “accessed in an aligned way” by a loop if:

- The base address of the array is appropriately aligned; and
- Every index used to access the array is a multiple of the SIMD width.

Is this true of the *j* loop in `perform_nbody()`?

### Tell the Compiler

Even when an array has been aligned using a function call or attribute, the compiler is still often unable to determine whether all accesses to that array will be aligned.

For best results, the programmer should tell the compiler all accesses are aligned when converting the code to SIMD form, as discussed in the previous section.

## Optimize Memory Access Patterns

So far, this lab has focused upon converting code to SIMD form (and doing so efficiently). The application speed-up we have achieved is (hopefully!) impressive, but we should not stop here – we can tune and optimize SIMD code in exactly the same way as we can tune and optimize scalar code.

Consider the memory-access pattern of our N-body application. Since we use a naïve all-pairs  $O(N^2)$  implementation, each body  $i$  is interacted with each body  $j$  and, for sufficiently large values of  $N$ , there are too many bodies to store in cache.

Cache-blocking is a well-known optimization that targets situations such as these, and remains applicable for SIMD codes. Block the  $j$  loop in `perform_nbody()`, following a similar code structure to that below:

```
#define CHUNK_SIZE 8192
!$omp parallel private(i,j,jj,...
  do jj = 1, number_of_bodies, CHUNK_SIZE
    !$omp do
      do i = 1, number_of_bodies
        ...

        !$omp simd private(...
          do j = jj, min(jj + CHUNK_SIZE, number_of_bodies)
            ...
          end do
        end do
      end do
    !$omp end do nowait
  !$omp end parallel
```

After applying this optimization, the code is able to keep a number of bodies (`CHUNK_SIZE`) in cache, and re-use this data for multiple iterations of the  $i$  loop.

Build and re-run the application, then record the new checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

How has the runtime changed? Play with the block size until you are seeing a performance improvement – the default size may not be optimal on all systems.

That cache blocking improves performance here highlights the importance of designing SIMD-friendly algorithms and converting code to SIMD form *before* spending time exploring other optimization opportunities. Codes that are compute-bound in scalar may be memory-bound in SIMD, and design decisions that are beneficial to scalar code may even be detrimental to SIMD code.

## Further Optimization (Intel® Fortran only)

In applications where high-latency math operations are a performance bottleneck, significant performance gains can be realised by reducing the accuracy of the Intel® math library.

To reduce the accuracy of the divide and square root operations in our N-body code, add the following two flags to the **Makefile**:

```
-fimf-domain-exclusion=15 -fimf-accuracy-bits=12
```

The first of these flags informs the compiler that we do not require math functions to return correct results for “uncommon” inputs (i.e., extremes, NaNs, infinities and denormals), and the second informs the compiler exactly how much precision we need for the remaining “common” inputs. The reader is referred to the Intel® Compiler documentation for an in-depth explanation of the operation of these flags.

Build and re-run the application, then record the new checksum and run time:

<b>Checksum:</b>	
<b>Run time:</b>	

How has the runtime changed?

People often fear that adjusting FP accuracy (as suggested above) is beyond what they want. Two examples: (1) processing de-normalized FP numbers (rather than flushing them to zero) is often not supported in “real hardware” or may be unsupported on some non-x86 architectures, and (2) consider calculating EXP(800) – did you know that the result is out of range with respect to double-precision? Therefore, it makes sense to eliminate such kind of input from the processing pipeline rather than processing it (normalized numbers, etc.).

## Homework: Outer-loop Vectorization, Array Notation

### Assignment 1

OpenMP 4.0 supports the application of its simd pragmas to outer- as well as inner-loops. Move the pragma from the *j* loop to the *i* loop, paying close attention to the clauses you supply.

How does this affect performance? Why?

### Assignment 2

Re-write the *j* loop using array notation.

## Legal Information

### 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

### Trademark Information

BlueMoon, BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Cilk, Core Inside, E-GOLD, Flexpipe, i960, Intel, the Intel logo, Intel AppUp, Intel Atom, Intel Atom Inside, Intel CoFluent, Intel Core, Intel Inside, Intel Insider, the Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel Sponsors of Tomorrow., the Intel Sponsors of Tomorrow. logo, Intel StrataFlash, Intel vPro, Intel Xeon Phi, Intel XScale, InTru, the InTru logo, the InTru Inside logo, InTru soundmark, Itanium, Itanium Inside, MCS, MMX, Pentium, Pentium Inside, Puma, skool, the skool logo, SMARTi, Sound Mark, Stay With It, The Creators Project, The Journey Inside, Thunderbolt, Ultrabook, vPro Inside, VTune, Xeon, Xeon Inside, X-GOLD, XMM, X-PMU and XPOSYS are trademarks of Intel Corporation in the U.S. and/or other countries.

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

### Technical Collateral Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

## Software Source Code Disclaimer

Any software source code reprinted in this document is furnished under a software license and may only be used or copied in accordance with the terms of that license.